Acyclic Heaps of Pieces, I

R.M. GREEN   rmg@euclid.colorado.edu
Department of Mathematics, University of Colorado, Boulder, CO 80309-0395, USA

Abstract.   Heaps of pieces were introduced by Viennot and have applications to algebraic combinatorics, the-oretical computer science and statistical physics. In this paper, we show how certain combinatorial properties of heaps studied by Fan and by Stembridge are closely related to the properties of a certain linear map ∂_E associated to a heap E. We examine the relationship between ∂_E and ∂_F when F is a subheap of E. This approach allows neat statements and proofs of results on certain associative algebras (generalized Temperley–Lieb algebras) that are otherwise tricky to prove. The key to the proof is to interpret the structure constants of the aforementioned algebras in terms of the maps ∂.

Introduction

A heap is an isomorphism class of labelled posets satisfying certain axioms. Heaps have a wide variety of applications, notably to parallelism in computer science, but also to statistical physics and algebraic combinatorics. Many of these applications are discussed by Viennot in [17].

One of the oldest results in algebraic topology states that if k is a field and G is a finite, connected, directed graph with vertex set $V(G)$ and edge set $E(G)$, then the linear map $\partial : kE(G) \to kV(G)$ sending the edge $(v_i \to v_j)$ to $v_i - v_j$ has image of codimension 1 in $kV(G)$.

In this paper, we will introduce and study an analogue of the above situation for heaps. The definitions are completely general but we are particularly interested in heaps arising from fully commutative elements in Coxeter groups as studied by Stembridge in [15]. We will show how this theory may be applied to obtain neater proofs of results on certain associative algebras.

We now summarise the main results of the paper for easy reference. Section 1 sets up the basic definitions of heaps, including our definition of the map ∂ in Definition 1.2.1. A heap E will be called acyclic if ∂_E is injective, and E will be called strongly acyclic if it and all its maximal subheaps are acyclic. The main body of theory is developed in Section 2. We look at two combinatorial properties of heaps, properties P1 and P2, and show how they are related to the linear notions of being acyclic and strongly acyclic. In favourable circumstances, property P1 is equivalent to being acyclic (Theorem 2.4.4) and property P2 is equivalent to being strongly acyclic (Theorem 2.4.2).

We are interested in how $\ker \partial_E$ is related to $\ker \partial_F$ when F is a subheap of E. The Deletion Lemma (Theorem 2.1.1) shows that $\ker \partial$ changes in dimension by at most 1 when an element is removed from the heap. In Section 2.3, we look at certain specific kinds of subheaps for which $\ker \partial$ is exactly predictable from the original heap. These constructions exhibit a close relationship between the map ∂ and a certain quotient of the heap monoid algebra whose structure constants can be described in terms of ∂ (Theorem 3.2.3). We apply the theory here to give simpler proofs of certain results involving generalized Temperley–Lieb algebras; one such result is given in Proposition 3.4.2, and more are given in Section 4.1. In Section 4.2 we sketch the relationship between $\ker \partial$ and certain diagram calculi, and we conclude with some questions in Section 4.3.

Although we were led to this theory by questions about generalized Temperley–Lieb algebras, it is hoped that the theory in this paper will be of independent interest.

1. Heaps

In Section 1.1, we introduce the basic properties of heaps. We will tend to follow Viennot's notation [17]. We give our definition of ∂ in Section 1.2.

1.1. Basic definitions

Definition 1.1.1   Let P be a set equipped with a symmetric and reflexive binary relation $\mathcal{C}$. The elements of P are called (*basic*) *pieces*, and the relation $\mathcal{C}$ is called the *concurrency relation*.

A *labelled heap* with pieces in P is a triple $(E, \leq, \varepsilon)$ where $(E, \leq)$ is a finite (possibly empty) partially ordered set with order relation denoted by $\leq$ and ε is a map $\varepsilon : E \to P$ satisfying the following two axioms.

1. For every $a, b \in E$ such that $\varepsilon(a) \, \mathcal{C} \, \varepsilon(b)$, a and b are comparable in the order $\leq$.
2. The order relation $\leq$ is the transitive closure of the relation $\leq_{\mathcal{C}}$ such that for all $a, b \in E$, $a \leq_{\mathcal{C}} b$ if and only if both $a \leq b$ and $\varepsilon(a) \, \mathcal{C} \, \varepsilon(b)$.

We will sometimes express the relation $a \leq b$ by saying that "a is *above* b". The terms *minimal* and *maximal* applied to the elements of the labelled heap refer to minimality (respectively, maximality) with respect to $\leq$.

Parts 1 and 2 of Definition 1.1.1 correspond to axioms (i) and (ii'') respectively in Viennot's paper.

Example 1.1.2   Let $P = \{1, 2, 3\}$ and, for $x, y \in P$, define $x \, \mathcal{C} \, y$ if and only if $|x - y| \leq 1$. Let $E = \{a, b, c, d, e\}$ partially ordered by extension of the (covering) relations $a \leq c$, $b \leq c$, $c \leq d$, $c \leq e$. Define the map ε by the conditions $\varepsilon(a) = \varepsilon(d) = 1$, $\varepsilon(c) = 2$ and $\varepsilon(b) = \varepsilon(e) = 3$. Then $(E, \leq, \varepsilon)$ can easily be checked to satisfy the axioms of Definition 1.1.1 and it is a labelled heap. The minimal elements are a and b, and the maximal elements are d and e.

Definition 1.1.3 Let $(E, \leq, \varepsilon)$ and $(E', \leq', \varepsilon')$ be two labelled heaps with pieces in P and with the same concurrency relation, $\mathcal{C}$. An isomorphism $\phi : E \to E'$ of posets is said to be an *isomorphism of labelled posets* if $\varepsilon = \varepsilon' \circ \phi$.

A *heap* of pieces in P with concurrency relation $\mathcal{C}$ is a labelled heap (Definition 1.1.1) defined up to labelled poset isomorphism. The set of such heaps is denoted by $H(P, \mathcal{C})$. We denote the heap corresponding to the labelled heap $(E, \leq, \varepsilon)$ by $[E, \leq, \varepsilon]$.

We will sometimes abuse language and speak of the underlying set of a heap, when what is meant is the underlying set of one of its representatives.

Definition 1.1.4 Let $(E, \leq, \varepsilon)$ be a labelled heap with pieces in P and let F be a subset of E. Let ε' be the restriction of ε to F. Let $\mathcal{R}$ be the relation defined on F by $a \mathcal{R} b$ if and only if $a \leq b$ and $\varepsilon(a) \mathcal{C} \varepsilon(b)$. Let $\leq'$ be the transitive closure of $\mathcal{R}$. Then $(F, \leq', \varepsilon')$ is a labelled heap with pieces in P. The heap $[F, \leq', \varepsilon']$ is called a *subheap* of $[E, \leq, \varepsilon]$.

We will often implicitly use the fact that a subheap is determined by its set of vertices and the heap it comes from.

Definition 1.1.5 The *concurrency graph* associated to the class of heaps $H(P, \mathcal{C})$ is the graph whose vertices are the elements of P and for which there is an edge from $v \in P$ to $w \in P$ if and only if $v \neq w$ and $v \mathcal{C} w$.

Definition 1.1.6 Let $E = [E, \leq_E, \varepsilon]$ and $F = [F, \leq_F, \varepsilon']$ be two heaps in $H(P, \mathcal{C})$. We define the heap $G = [G, \leq_G, \varepsilon''] = E \circ F$ of $H(P, \mathcal{C})$ (which we call the superposition of E over F) as follows.

1. The underlying set G is the disjoint union of E and F.
2. The labelling map ε'' is the unique map $\varepsilon'' : G \to P$ whose restriction to E (respectively, F) is ε (respectively, ε').
3. The order relation $\leq_G$ is the transitive closure of the relation $\mathcal{R}$ on G, where $a \mathcal{R} b$ if and only if one of the following three conditions holds:

   (i) $a, b \in E$ and $a \leq_E b$;
   (ii) $a, b \in F$ and $a \leq_F b$;
   (iii) $a \in E$, $b \in F$ and $\varepsilon(a) \mathcal{C} \varepsilon'(b)$.

Remark 1.1.7 Definition 1.1.6 can easily be shown to be sound (see [17, Section 2]). It is immediate from the construction that E and F are subheaps of $E \circ F$. Note that Viennot calls $E \circ F$ "the superposition of F over E".

As in [17], we will write $a \circ E$ and $E \circ a$ for $\{a\} \circ E$ and $E \circ \{a\}$, respectively. Note that $a \circ E$ and $b \circ E$ are equal as heaps if $\varepsilon(a) = \varepsilon(b)$.

Definition 1.1.8 A *trivial heap* is a heap $[E, \leq, \varepsilon]$ for which the order relation $\leq$ is trivial, meaning that no element of E is above any other.

1.2. The map ∂

We can now introduce our analogue of the graph theoretic phenomenon mentioned in the introduction; this is the central definition of this paper. Throughout Section 1.2, we let $[E, \leq, \varepsilon]$ be a heap in the set $H(P, \mathcal{C})$ with pieces in P and concurrency relation $\mathcal{C}$. We also fix a field, k.

Definition 1.2.1 Let V_0 be the set of elements of $[E, \leq, \varepsilon]$, i.e., the set of elements of (a representative of) the underlying poset, E. We call the elements of V_0 *vertices* and denote their k-span by C_0.

Let V_1 be the set of all pairs $(x, y) \in E \times E$ with $x < y$ and $\varepsilon(x) = \varepsilon(y)$ such that there is no element z for which we have both $\varepsilon(x) = \varepsilon(z) = \varepsilon(y)$ and $x < z < y$. We call the elements of V_1 *edges* and denote their k-span by C_1.

For all other integers $i \in \mathbb{Z} \backslash \{0, 1\}$, we define $C_i = 0$.

The k-linear map $\partial = \partial_E : C_1 \to C_0$ is defined by its effect on the edges as follows:

$$\partial : (x, y) \mapsto \sum_{\substack{x < w < y \\ \varepsilon(w) \; \mathcal{C} \; \varepsilon(x)}} w.$$

Remark 1.2.2 Note that, in the sum of Definition 1.2.1, we have $\varepsilon(x) = \varepsilon(y)$ but it is not possible for $\varepsilon(w) = \varepsilon(x)$ because of the conditions imposed on the edge (x, y).

Example 1.2.3 Consider the heap $[E, \leq, \varepsilon]$ arising from the labelled heap of Example 1.1.2. In this case, $C_1 = \{(a, d), (b, e)\}$ and $C_0 = \{a, b, c, d, e\}$. We have $\partial((a, d)) = c$ and $\partial((b, e)) = c$. It follows that $\ker \partial_E$ is of dimension 1 spanned by $[(a, d) - (b, e)]$ and that $\operatorname{coker} \partial_E$ is of dimension 4, spanned by the images of a, b, d, e.

The following simple Lemma explains the relationship between the dimensions of $\ker \partial$ and $\operatorname{coker} \partial$.

Lemma 1.2.4 *Let $[E, \leq, \varepsilon]$ be a heap in $H(P, \mathcal{C})$. Then*

$$\dim \operatorname{coker} \partial_E - \dim \ker \partial_E$$

is $|\varepsilon(E)|$, the number of different labels occurring on vertices of the heap.

Proof: It is clear from Definition 1.2.1 that

$$\dim \operatorname{coker} \partial_E - \dim \ker \partial_E = \dim C_0 - \dim C_1.$$

For each $p \in P$, let $n_p(E) = |\varepsilon^{-1}(p)|$; note that $\varepsilon^{-1}(p)$ is a chain in E. The definition of V_1 shows that the number of edges $(x, y) \in V_1$ with $\varepsilon(x) = \varepsilon(y) = p$ is $\max(0, \; n_p(E) - 1)$. This shows that $|V_1| = |E| - |\varepsilon(E)|$. The lemma follows from the fact that $|V_0| = |E|$. $\square$

It will be convenient to use the following definitions in the sequel.

Definition 1.2.5 Let $E = [E, \leq_E, \varepsilon]$ be a heap in $H(P, \mathcal{C})$ and let k be a field. If $v \in E$, we let $E(v) = [E(v), \leq_{E(v)}, \varepsilon']$ be the subheap of E obtained by defining $E(v) = E\backslash\{v\}$.

We say E is *acyclic* if $\ker \partial_E = 0$. We say E is *strongly acyclic* if E is acyclic and $E(v)$ is acyclic for all $v \in E$. We say v is an *image vertex* of E if $v \in \text{Im}(\partial_E)$.

A large portion of this paper will be concerned with the following two problems and their applications.

Problem 1.2.6 Can we characterize and classify the acyclic (respectively, strongly acyclic) heaps in $H(P, \mathcal{C})$?

Problem 1.2.7 If E is a heap in $H(P, \mathcal{C})$ and F is a subheap of E, how is $\ker \partial_F$ related to $\ker \partial_E$?

2. Properties of the map ∂

In Section 2 we prove some results about the map ∂ that hold in a general context. In Section 2.1, we will prove a deletion lemma that relates $\ker \partial_E$ to $\ker \partial_{E(v)}$, in the notation of Definition 1.2.5. In Section 2.2, we exhibit some general relationships between combinatorial and linear properties of heaps. In Section 2.3, we show how to "contract" a heap to a simpler one for which $\ker \partial$ is very similar; this turns out to be a very useful proof technique. We introduce the notion of a regular class of heaps $H(P, \mathcal{C})$ in Section 2.4; the heaps of these classes have particularly tractable properties.

Our motivation behind these results is to apply them to certain associative algebras in Section 3, but the results of Section 2 are related to each other in intriguing ways that shed light on Problems 1.2.6 and 1.2.7.

As before, we will fix a set $H(P, \mathcal{C})$ and a field, k.

2.1. The deletion lemma

The next theorem is the main result of Section 2.1. It is very useful for certain applications, including our proof of Proposition 3.4.2.

Theorem 2.1.1 (*Deletion lemma*) *Let $[E, \leq, \varepsilon]$ be a nonempty heap in $H(P, \mathcal{C})$, and let $v \in E$. Then $|\dim \ker \partial_E - \dim \ker \partial_{E(v)}| \leq 1$.*

To prove this, it is convenient to define the following vector spaces and maps.

Definition 2.1.2 Fix v as in Theorem 2.1.1.

(i) Let A_1 be the space C_1 associated with the heap E. Let A_0 be the quotient of the space C_0 associated to E by the 1-dimensional subspace $\langle v \rangle$. Let $\partial_A : A_1 \rightarrow A_0$ be the composition of the map ∂ associated to E with the natural epimorphism.

(ii) Let B_1 be the space C_1 associated to $E(v)$, let B_0 be the space C_0 associated to $E(v)$ and let ∂_B be the map ∂ associated to $E(v)$.

(iii) We define a k-linear map $f_1 : B_1 \to A_1$ by its effect on the edges of B as follows. If $(x, y) \in B_1$ is such that $\varepsilon(x) = \varepsilon(y) = \varepsilon(v)$ and $x < v < y$ in A, we define $f_1((x, y)) = (x, v) + (v, y)$. Otherwise, we define $f_1((x, y)) = (x, y)$. (This is well defined by the definition of subheaps and Definition 1.2.1.)

(iv) Let f_0 be the obvious k-isomorphism from B_0 to A_0.

Lemma 2.1.3 *In the notation of Definition* 2.1.2, *the following diagram commutes and has exact rows.*

$$
\begin{array}{ccccccccc}
0 & \to & B_1 & \overset{f_1}{\to} & A_1 & \to & \operatorname{coker} f_1 & \to & 0 \\
  &   & \downarrow \partial_B & & \downarrow \partial_A & & \downarrow & & \\
0 & \to & B_0 & \overset{f_0}{\to} & A_0 & \to & 0 & \to & 0
\end{array}
$$

Proof: Injectivity of f_1 follows from the fact that the images of basis elements in B_1 under f_1 have disjoint supports in A_1. (Recall that the fibres of ε are totally ordered.)

The rest of the claim follows from Definition 1.2.1, Definition 2.1.2 and a diagram chase. □

The Snake lemma immediately gives the following result.

Lemma 2.1.4 *Maintain the above notation. There is an exact sequence*

$$0 \to \ker \partial_{E(v)} \to \ker \partial_A \to \operatorname{coker} f_1 \to \operatorname{coker} \partial_{E(v)} \to \operatorname{coker} \partial_A \to 0.$$

In order to extract more information from Lemma 2.1.4, we need to know whether v is an image vertex or not (see Definition 1.2.5).

Lemma 2.1.5 *Maintain the above notation.*

 (i) *If v is an image vertex then*

$$\dim \ker \partial_A = \dim \ker \partial_E + 1$$

   and

$$\dim \operatorname{coker} \partial_A = \dim \operatorname{coker} \partial_E.$$

(ii) *If v is not an image vertex then*

$$\dim \ker \partial_A = \dim \ker \partial_E$$

   and

$$\dim \operatorname{coker} \partial_A = \dim \operatorname{coker} \partial_E - 1.$$

Proof: This is an exercise in linear algebra, using the definition of ∂_A. □

Lemma 2.1.6 *If v is the unique vertex x of E with $\varepsilon(v) = \varepsilon(x)$, then f_1 is surjective. Otherwise*, $\dim \operatorname{coker} f_1 = 1$.

Proof: If v satisfies the uniqueness property above then there can be no edges (x, y) in A_1 (or B_1) with either $x = v$ or $y = v$. It follows that f_1 is an isomorphism in this case.

On the other hand, if v does not satisfy the uniqueness property then the argument of Lemma 1.2.4 shows that $\dim A_1 = \dim B_1 + 1$, and the lemma follows.   □

Proof of Theorem 2.1.1:   Suppose that $v \in \mathrm{Im}(\partial_E)$; this is case (i) of Lemma 2.1.5. Since $\dim \mathrm{coker} f_1 \leq 1$ by Lemma 2.1.6, Lemma 2.1.4 shows that $\dim \ker \partial_{E(v)} - \dim \ker \partial_E$ is equal to 0 or 1.

Suppose now that $v \notin \mathrm{Im}(\partial_E)$. A similar argument based on case (ii) of Lemma 2.1.5 shows that $\dim \ker \partial_{E(v)} - \dim \ker \partial_E$ is equal to 0 or -1.   □

2.2.   *Heaps with additional properties*

The two properties P1 and P2 introduced in this section have combinatorial definitions, but as we shall see, they are related to properties of the map ∂ and they cast some light on Problem 1.2.6.

Definition 2.2.1 (Property P1)   Let $E = [E, \leq, \varepsilon] \in H(P, \mathcal{C})$ be a heap. We write $E(a) \prec^+ E$ (respectively, $E(a) \prec^- E$) if a is a maximal (respectively, minimal) vertex of E and there exists a maximal (respectively, minimal) vertex b of $E(a)$ with $\varepsilon(b) \neq \varepsilon(a)$ such that b is not maximal (respectively, minimal) in E. We write $E(a) \prec E$ if either $E(a) \prec^+ E$ or $E(a) \prec^- E$.

If there is a (possibly trivial) sequence $E_1 \prec E_2 \prec \cdots \prec E$ of heaps in $H(P, \mathcal{C})$ where E_1 is a trivial heap, we say that the heap E is *dismantlable* or that E has property P1.

Example 2.2.2   The heap E arising from Example 1.1.2 does not have property P1, but its subheap $E(a)$ does: consider the sequence of subheaps

$$\{d, e\} \prec \{c, d, e\} \prec \{b, c, d, e\}.$$

In Section 3.4, we shall exploit the relationship between property P1 and Fan's notion of left and right cancellability [5, Definition 4.2.4]. We avoid the term "cancellability" in this paper because of possible confusion with the use of this term in the theory of monoids.

Proposition 2.2.3   *A dismantlable heap is acyclic.*

Proof: Let E be a dismantlable heap and let

$$E_1 \prec E_2 \prec \cdots \prec E_l = E$$

be a chain of heaps with E_1 trivial. The proof is by induction on l. If $l = 1$, the claim is clear because a trivial heap has no edges.

For the inductive step, we treat the case where $E_{l-1} \prec^- E_l$; the other case is similar. Let a be a minimal element of E_l and let b be a minimal vertex of E_{l-1} that is not minimal in E_l with $\varepsilon(b) \neq \varepsilon(a)$. Suppose $\ker \partial_{E_l} \neq 0$ and let

$$\sum \lambda_i e_i$$

be a nontrivial element of $\ker \partial$. Since $\ker \partial_{E_{l-1}} = 0$, one of the edges e_i must involve the vertex a; let us write $e_i = (a, c)$ as a is minimal. (Note that $c \neq b$ as $\varepsilon(c) = \varepsilon(a) \neq \varepsilon(b)$.) Since b is minimal in E_{l-1}, the vertex b cannot occur with nonzero coefficient in $\partial(e_j)$ for any edge e_j with $j \neq i$. However, b occurs with coefficient 1 in $\partial(e_i)$, contrary to hypothesis. $\square$

Remark 2.2.4 The converse of Proposition 2.2.3 is false in general. Consider the class of heaps $H(P, \mathcal{C})$ for which the concurrency graph is a square whose corners (the elements of P) are consecutively labelled p_1, p_2, p_3, p_4. Let $E \in H(P, \mathcal{C})$ be a labelled heap whose underlying set is $\{a_1, a_2, a_3, a_4\}$ and $\varepsilon(a_i) = p_i$ for all i. Let $a_i < a_j$ whenever i is odd and j is even. The heap $[E, \leq, \varepsilon]$ is acyclic but not dismantlable.

The next property is modelled on Stembridge's characterization of full commutativity in [15, Proposition 2.3]. The term "convex chain" for a heap has its natural meaning: a chain

$$x_1 < x_2 < \cdots < x_t$$

of vertices in a heap is convex if and only if whenever $x_i < y < x_j$ for some y, the vertex y is an element of the chain.

Definition 2.2.5 (Property P2) We say a heap $E = [E, \leq, \varepsilon] \in H(P, \mathcal{C})$ has property P2 if it contains no convex chains of the form $x < y < z$ or $x < z$ with $\varepsilon(x) = \varepsilon(z)$ in either case.

Example 2.2.6 The heap arising from Example 1.1.2 does not have property P2. Although there are no chains of the form $x < z$ with $\varepsilon(x) = \varepsilon(z)$, the chains $a < c < d$ and $b < c < e$ each violate the other requirement.

Proposition 2.2.7 *A strongly acyclic heap has property P2.*

Proof: Let $E = [E, \leq, \varepsilon]$ be a heap that fails property P2. If E contains an convex chain of the form $x < z$ with $\varepsilon(x) = \varepsilon(z)$ then (x, z) is an edge in E and $\partial((x, z)) = 0$, meaning that E is not acyclic. The other possibility is that E contains a convex chain $x < y < z$ with $\varepsilon(x) = \varepsilon(z)$. In this case, the subheap $E(y)$ contains an edge (x, z) with $\partial((x, z)) = 0$, and $E(y)$ is not acyclic, meaning that E is not strongly acyclic. $\square$

Remark 2.2.8 The converse of Proposition 2.2.7 is false in general. Consider the set of heaps $H(P, \mathcal{C})$ defined in Remark 2.2.4. Let E be a labelled heap with elements $\{a_1, a_2, a_3,$

Figure 1.   Hasse diagram for the labelled heap $(E, \leq)$ of Remark 2.2.8.

$a_4, a_5, a_6\}$ labelled by the elements $p_1, p_3, p_2, p_4, p_1, p_3$ respectively. The Hasse diagram of $(E, \leq)$ is shown in figure 1. This gives a heap with edges $e_1 = (a_1, a_5)$ and $e_2 = (a_2, a_6)$. We have $\partial(e_1) = a_3 + a_4$ and $\partial(e_2) = a_3 + a_4$. This means that the heap arising from E is not acyclic, although the reader may easily check that the heap has property P2.

The following lemma is similar to [15, Lemma 3.1]. (See Remark 1.1.7 for the notation.)

Lemma 2.2.9   *Let $E = [E, \leq, \varepsilon]$ be a heap in $H(P, \mathcal{C})$. If E has property P2 but $a \circ E$ does not, then either there is a minimal element c of E with $\varepsilon(c) = \varepsilon(a)$ or there is a convex chain $a < b < d$ in $a \circ E$ with $\varepsilon(a) = \varepsilon(d) \neq \varepsilon(b)$.*

Proof: This is an easy consequence of the definition of property P2.   □

2.3.   *More on convex chains*

In Section 2.3, we give a precise answer to Problem 1.2.7 in certain special cases by showing that the subheaps F of a given heap E arising from a certain construction are such that $\ker \partial_F$ is predictable from $\ker \partial_E$. To define this construction, we need the concept of contraction along a convex chain.

Definition 2.3.1   Let $E = [E, \leq, \varepsilon]$ be a heap of $H(P, \mathcal{C})$, and let

$$\mathbf{c} = (x_1, x_2, \ldots, x_i) : x_1 < x_2 < \cdots < x_i$$

be a chain in E. We say $\mathbf{c}$ is *balanced* if $\varepsilon(x_1) = \varepsilon(x_i)$. If $\mathbf{c}$ is a balanced convex chain, we define the heap $E/\mathbf{c}$ to be the subheap of E obtained by omitting the vertices $x_2, x_3, \ldots, x_i$. We call the heap $E/\mathbf{c}$ the *contraction of E along* $\mathbf{c}$, and the number i is called the *length* of the chain.

Remark 2.3.2   If, in the above definition, we had omitted vertices $x_1, x_2, \ldots, x_{i-1}$ instead, we would have obtained the same heap: the fact that $\mathbf{c}$ is balanced makes the two corresponding labelled heaps isomorphic.

Example 2.3.3 Take the heap arising from Example 1.1.2. The chain $b < c < e$ is balanced and convex, and contraction along this chain gives the subheap with vertices $\{a, b, d\}$.

Lemma 2.3.4 *Let $E = [E, \leq, \varepsilon]$ be a heap and let $\mathbf{c}$ be a balanced convex chain of E of length 2. Then* $\dim(\ker \partial_E) = \dim(\ker \partial_{E/\mathbf{c}}) + 1$.

Proof: Let us write $B = E/\mathbf{c}$ for notational convenience. The chain $\mathbf{c}$ consists of (a covering pair of) elements $x < y$ with $\varepsilon(x) = \varepsilon(y)$.

We denote by B_1 (respectively, E_1) the span of the edges in B (respectively, E), and we denote by B_0 (respectively, E_0) the span of the vertices in B (respectively, E). Let g_1 be the map from B_1 to E_1 defined by its effect on edges as follows:

$$g_1((a, b)) = \begin{cases} (y, b) & \text{if } a = x, \\ (a, b) & \text{otherwise.} \end{cases}$$

Let g_0 be the map from B_0 to E_0 defined by its effect on vertices as follows:

$$g_0(v) = \begin{cases} x + y & \text{if } v = x, \\ v & \text{otherwise.} \end{cases}$$

Definition 1.2.1 implies that the left square in the diagram

$$
\begin{array}{ccccccccc}
0 & \to & B_1 & \overset{g_1}{\to} & E_1 & \to & \mathrm{coker} g_1 & \to & 0 \\
 & & \downarrow \partial_B & & \downarrow \partial_E & & \downarrow 0 & & \\
0 & \to & B_0 & \overset{g_0}{\to} & E_0 & \to & \mathrm{coker} g_0 & \to & 0
\end{array}
$$

commutes. The right square commutes because $\dim \mathrm{coker}\, g_1 = 1$ and the edge $(x, y) \in E_1$ lies in $\ker \partial_E \backslash \mathrm{Im}(g_1)$. The same argument shows that the connecting homomorphism z in the exact sequence

$$0 \to \ker \partial_B \to \ker \partial_E \to \mathrm{coker} g_1 \overset{z}{\to} \mathrm{coker}\, \partial_B \to \mathrm{coker}\, \partial_E \to \mathrm{coker} g_0 \to 0$$

provided by the Snake Lemma is zero, and the result follows. $\square$

Lemma 2.3.5 *Let $E = [E, \leq, \varepsilon]$ be a heap and let $\mathbf{c}$ be a balanced convex chain $x < y < z$ such that $\varepsilon(x) \neq \varepsilon(y)$. Then* $\dim \ker \partial_E = \dim \ker \partial_{E/\mathbf{c}}$.

Proof: Note that y is necessarily an image vertex as it is the image of the edge (x, z) under ∂. We will write $B = E/\mathbf{c}$ for notational convenience.

Suppose that y is the unique vertex c in E with $\varepsilon(y) = \varepsilon(c)$. Lemmas 2.1.4, 2.1.5 (i) and 2.1.6 then imply that

$$\dim \ker \partial_{E(y)} = \dim \ker \partial_E + 1,$$

and Lemma 2.3.4 shows that

$$\dim \ker \partial_{E(y)} = \dim \ker \partial_{E/\mathbf{c}} + 1,$$

because $E/\mathbf{c} = E(y)/\mathbf{c}'$ where $\mathbf{c}'$ is the balanced convex chain $x < z$ of $E(y)$. This proves the claim in this case, and we may now suppose that there exists a vertex $c \neq y$ with $\varepsilon(c) = \varepsilon(y)$. It is either possible to choose c so that (c, y) is an edge or to choose c so that (y, c) is an edge. We will assume the former case is possible; the other case follows mutatis mutandis.

We denote by B_1 (respectively, B_0) the span of the edges (respectively, vertices) of B. Let A_1 be the span of all edges of E except (x, z), and let A_0 be the quotient of the space C_0 (as in Definition 1.2.1) by the 1-dimensional subspace $\langle y \rangle$. The differential ∂_E induces a k-linear map ∂_A from A_1 to A_0. Let h_1 be the map from B_1 to A_1 defined by its effect on edges as follows:

$$h_1((a, b)) = \begin{cases} (a, y) + (y, b) & \text{if } (a, y) \text{ and } (y, b) \text{ are edges in } E, \\ (z, b) & \text{if } a = x, \\ (a, b) & \text{otherwise.} \end{cases}$$

Let h_0 be the map from B_0 to A_0 defined by its effect on vertices as follows:

$$h_0(v) = \begin{cases} x + z + \langle y \rangle & \text{if } v = x, \\ v + \langle y \rangle & \text{otherwise.} \end{cases}$$

A routine check using the definitions shows that the left square in the diagram

$$
\begin{array}{ccccccccc}
0 & \to & B_1 & \overset{h_1}{\to} & A_1 & \to & \operatorname{coker} h_1 & \to & 0 \\
 & & \downarrow \partial_B & & \downarrow \partial_A & & \downarrow s & & \\
0 & \to & B_0 & \overset{h_0}{\to} & A_0 & \to & \operatorname{coker} h_0 & \to & 0
\end{array}
$$

commutes. This induces a map s making the diagram commute. The map s is nonzero, since $\operatorname{coker} h_1$ is spanned by the image of (c, y), x occurs in the support of $\partial((c, y))$ and z does not. Since $\operatorname{coker} h_0$ is spanned by the image of x, s is an isomorphism and the Snake Lemma applied to this situation gives the exact sequence

$$0 \to \ker \partial_B \to \ker \partial_A \to 0 \to \operatorname{coker} \partial_B \to \operatorname{coker} \partial_A \to 0.$$

A linear algebra argument shows that $\dim \ker \partial_A = \dim \ker \partial_E$, because

$$\partial_E((x, z)) = y.$$

The conclusion follows. $\square$

Remark 2.3.6 It follows easily from the definition of property P2 that if E is a heap that does not have property P2 then it is possible to contract a balanced convex chain in E of one of the types given in Lemmas 2.3.4 and 2.3.5. This is the main point of the above two results.

2.4. Regular classes of heaps

We now introduce the notion of a regular class of heaps, and show that in a regular class of heaps, the converses to propositions 2.2.3 and 2.2.7 hold. We shall look at some examples of regular classes of heaps in Section 3.4.

Definition 2.4.1 A class of heaps $H(P, \mathcal{C})$ is said to be *regular* if any heap of $H(P, \mathcal{C})$ with property P2 also has property P1.

The counterexamples in Remarks 2.2.4 and 2.2.8 come from classes of heaps that are not regular.

Theorem 2.4.2 *Suppose that $H(P, \mathcal{C})$ is a regular class of heaps. Let $E = [E, \leq, \varepsilon']$ be a heap of $H(P, \mathcal{C})$ with property P2. Then*
 (i) *E is strongly acyclic*;
(ii) *if* $\ker \partial_{a \circ E} \neq 0$ *(respectively,* $\ker \partial_{E \circ a} \neq 0$*) then there is a minimal (respectively, maximal) vertex c of E such that $\varepsilon(c) = \varepsilon(a) \in P$, where ε is the map associated to the heap $a \circ E$ (respectively, $E \circ a$).*
In particular, property P2 and the property of being strongly acyclic coincide for heaps of $H(P, \mathcal{C})$.

Proof: The last claim is immediate from (i) and Proposition 2.2.7.

The proof of (i) and (ii) is by induction on n, the number of vertices in the heap E. If $n = 0$, there is nothing to prove. If E is nonempty but trivial, claims (i) and (ii) follow easily, and this deals with the case $n = 1$. Let $P(l)$ be the statement "claim (i) holds when E is a heap with l vertices and properties P1 and P2", and let $Q(l)$ be the statement "claim (ii) holds when E is a heap with $l - 1$ vertices and properties P1 and P2". We will be done if we can show that $P(l) \Rightarrow Q(l+1)$ and $(P(l) \wedge Q(l)) \Rightarrow P(l+1)$.

Suppose $P(l)$ holds, that E is a heap with l vertices and that $\ker \partial_{a \circ E} \neq 0$. (We omit consideration of the case $\ker \partial_{E \circ a} \neq 0$, which is similar.) It cannot be the case that $a \circ E$ has property P2, because if it did, it would have property P1 by assumption and would be acyclic by Proposition 2.2.3. By Lemma 2.2.9, we see that if the statement $Q(l+1)$ fails for the heap $a \circ E$, there must be a convex chain $\mathbf{c} = a < b < d$ in $a \circ E$ with $\varepsilon(a) = \varepsilon(d) \neq \varepsilon(b)$. Remark 2.3.2 shows that $(a \circ E)/\mathbf{c} = E(b)$, and by Lemma 2.3.5 we have

$$\dim \ker \partial_{a \circ E} = \dim \ker \partial_{E(b)}.$$

Since E is strongly acyclic, the right hand side is zero and we have a contradiction, proving $Q(l+1)$.

Now suppose $P(l)$ and $Q(l)$ hold and that E has $l + 1$ vertices. We may assume that E is not trivial. Because E has property P1, we have $E' \prec E$ for some heap E. We deal with the case where $E' \prec^- E$, the other case being similar. In this case we have $E = a \circ E'$ for some vertex a. Since E has property P1 by assumption, it is acyclic by Proposition 2.2.3 and there is a minimal element $b \in E'$ (with $\varepsilon(b) \neq \varepsilon(a)$) that is not minimal in E.

Suppose also that E is not strongly acyclic. Then then there is an element $v \in E$ such that $\ker \partial_{E(v)} \neq 0$. We cannot have $v = a$ because $E' = E(a)$ inherits property P2 from E and is therefore acyclic. Suppose $v \neq b$ and let $[\sum \lambda_i e_i]$ be a nontrivial element of $\ker \partial_{E(v)}$, where the e_i are edges in $E(v)$. It is not possible for any of the edges e_i to involve the vertex a, because $b \neq v$ would occur with coefficient 1 in the image of any edge (a, c) but would not occur in the image of any other edge as b is minimal in E'. This means that $[\sum \lambda_i e_i]$ would be a nontrivial element of $\ker \partial_{E'}$, contradicting $P(l)$ applied to the heap E'. We conclude that $v = b$. The heap $E'(b)$ inherits property P2 from E' and is therefore acyclic. Since $E(b) = a \circ E'(b)$ and $\ker \partial_{E(b)} \neq 0$, we apply statement $Q(l)$ to $E(b)$ and conclude that there is a minimal vertex c of $E'(b)$ with $\varepsilon(c) = \varepsilon(a)$. Now $a < b < c$ is a convex chain in E with $\varepsilon(a) = \varepsilon(c)$, which contradicts the fact that E has property P2 and completes the proof.   □

Remark 2.4.3   It is possible to find classes of heaps $H(P, \mathcal{C})$ for which property P2 does not imply property P1. For example, the heap in Remark 2.2.8 has property P2 but not property P1. In Section 3.4, we shall look at some examples of classes of heaps $H(P, \mathcal{C})$ that do satisfy the hypotheses of Theorem 2.4.2.

Theorem 2.4.4   *In a regular class of heaps, every acyclic heap is dismantlable, so property P1 and the property of being acyclic coincide in this case.*

Proof: The second assertion is immediate from the first and Proposition 2.2.3.

Let $E = [E, \leq, \varepsilon]$ be an acyclic heap of the regular class $H(P, \mathcal{C})$. If E has property P2, we are done. If not, Remark 2.3.6 shows that there is a balanced convex chain in E of one of the types mentioned in Lemmas 2.3.4 and 2.3.5. Since E is finite, there is a sequence

$$E = E_0, E_1, \ldots, E_l$$

where E_l has property P2 and for each i, E_{i+1} is obtained from E_i by contraction of a balanced convex chain as above. Since E is acyclic, and E_l is acyclic by Proposition 2.2.3, we see from lemmas 2.3.4 and 2.3.5 that there will never be an opportunity to apply Lemma 2.3.4 in this sequence. We will be done if we can show that if $\mathbf{c}$ is a balanced convex chain $x < y < z$ for which $\varepsilon(x) \neq \varepsilon(y)$, then E is dismantlable if $E/\mathbf{c}$ is dismantlable.

We proceed by induction on $n = |E/\mathbf{c}|$. If $n = 1$, the heap E consists solely of the chain $x < y < z$, which is dismantlable by inspection. Suppose the statement is true for $n = l$, and that $|E/\mathbf{c}| = l + 1$. Choose a such that $(E/\mathbf{c})(a) \prec E/\mathbf{c}$ as in Definition 2.2.1. We will deal with the case $(E/\mathbf{c})(a) \prec^- E/\mathbf{c}$, the other case being similar.

If $a \neq x$ then the heap $E(a)$ can be contracted by Lemma 2.3.5 to the heap $(E/\mathbf{c})(a)$; the latter heap has property P1 by construction and $E(a)$ is dismantlable by the inductive

hypothesis. Let b be minimal in $(E/\mathbf{c})(a)$ but not in $E/\mathbf{c}$, with $\varepsilon(b) \neq \varepsilon(a)$. Then b lies in $E(a)$ and is minimal in $E(a)$ but not in E. This shows that E is dismantlable.

The other possibility is that $a = x$, in which case we have a sequence

$$E/\mathbf{c} \prec E(x) \prec E,$$

and the claim follows from Definition 2.2.1. $\square$

3.   Quotients of heap monoid algebras

We show in Section 3 that the dimensions of $\ker \partial_E$ of a heap E have a nice interpretation as the structure constants of a certain algebra associated to the heap E. This fact is our main motivation in this paper.

3.1.   The heap monoid and some related structures

We now introduce the heap monoid associated to an arbitrary heap. This is naturally isomorphic to the commutation monoid (or "free partially abelian monoid") appearing in the work of Cartier and Foata [2]. The Mazurkiewicz traces [13] used to study concurrency in computer science are another variant of the same idea, and there is a large body of literature about them.

Definition 3.1.1   A class of heaps $H(P, \mathcal{C})$ has a natural monoid structure with composition given by the superposition map $\circ$ of Definition 1.1.6. We call this monoid the *heap monoid*.

Definition 3.1.2   Let A be a set and let A^* be the free monoid generated by A. Let C be a symmetric and antireflexive relation on A. The *commutation monoid* $\mathrm{Co}(A, C)$ is the quotient of the free monoid A^* by the congruence $\equiv_C$ generated by the commutation relations:

$$ab \equiv_C ba \quad \text{for all } a, b \in A \quad \text{with} \quad a \, C \, b.$$

To explain the relationship between heap monoids and commutation monoids, it is convenient to consider linear extensions of heaps. These are also discussed in [17, Section 3] and [15, Section 1.2].

Definition 3.1.3   Let $(E, \leq)$ be a poset with n elements. A *natural labelling* of $(E, \leq)$ is a bijection $\pi : E \to [n] = \{1, 2, \ldots, n\}$ such that $a < b$ implies that $\pi(a) < \pi(b)$. If $[E, \leq, \varepsilon]$ is a heap of $H(P, \mathcal{C})$ and $p_i = \varepsilon(\pi^{-1}(i))$, we call the word

$$p_1 p_2, \ldots, p_n$$

of P^* a *linear extension* of E.

The following is a standard result about heaps, and a proof may be found in [17, Proposition 3.4].

Proposition 3.1.4   *Let $E = [E, \leq, \varepsilon]$ be a heap of $H(P, \mathcal{C})$, let C be the complementary relation of $\mathcal{C}$, and let π be a natural labelling of E. If we regard the words P^* as representing elements of $\mathrm{Co}(P, C)$, then the map sending E to its linear extension in P^* under π is independent of the choice of π, and induces an isomorphism of monoids $H(P, \mathcal{C}) \to \mathrm{Co}(P, C)$.*

The following quotient of the monoid algebra will be of interest in our applications.

Definition 3.1.5   Maintain the above notation. Let $\mathcal{A}$ be the ring of Laurent polynomials $\mathbb{Z}[v, v^{-1}]$, let $\delta := v + v^{-1}$, and let $\mathcal{A}\mathrm{Co}(P, C)$ be the monoid algebra of $\mathrm{Co}(P, C)$ over $\mathcal{A}$. We define the *generalized Temperley–Lieb algebra $TL(P, \mathcal{C})$* to be the $\mathcal{A}$-algebra obtained by quotienting $\mathcal{A}\mathrm{Co}(P, C)$ by the relations

$$ss = \delta s,$$
$$sts = s \quad \text{if } s \neq t \quad \text{and} \quad s\,\mathcal{C}\,t,$$

where $s, t \in P$.

Lemma 3.1.6   *The isomorphism of Proposition 3.1.4 induces an isomorphism between the algebra $TL(P, \mathcal{C})$ and the quotient of $\mathcal{A}H(P, \mathcal{C})$ by the relations*

$$E = \delta E/\mathbf{c} \text{ if } \mathbf{c} \text{ is a balanced convex chain of length } 2,$$
$$E = E/\mathbf{c} \text{ if } \mathbf{c} \text{ is a balanced convex chain } x < y < z \text{ with } \varepsilon(x) \neq \varepsilon(y).$$

Proof: If $p_1 p_2, \ldots, p_n$ is a word in P^* corresponding to a heap $E = [E, \leq, \varepsilon]$ in $H(P, \mathcal{C})$ with natural labelling π, then equation (7) of [17, Section 3] shows that $\pi^{-1}(i) < \pi^{-1}(j)$ if and only if there is a sequence

$$1 \leq i = i_1 < \cdots < i_t = j \leq n$$

such that $p_{i_m} \mathcal{C} p_{i_{m+1}}$ for $1 \leq m < t$. It follows from this observation that subwords of $p_1 p_2, \ldots, p_n$ of the form ss (respectively, sts) as in Definition 3.1.5 correspond to chains in E of length 2 (respectively, 3) as described in the statement.

The converse implication follows from the standard fact that for any convex chain in a poset, there exists a linear extension of the poset in which the members of the chain appear consecutively.   □

3.2.   *Structure constants for $TL(P, \mathcal{C})$*

In Section 3.2, we exhibit a free $\mathcal{A}$-basis for $TL(P, \mathcal{C})$ using Bergman's diamond lemma [1], and show that in favourable circumstances the structure constants of the algebra with

respect to this basis are closely related to dimensions of $\ker \partial$ for certain heaps. This allows us to explore our main application. We remark that the idea for finding this basis essentially comes from Graham's thesis [7, Theorem 6.2].

In order to use Bergman's diamond lemma we need to recall some terminology from [1]. Let R be a commutative ring and let X be a nonempty set. Let $\leq_X$ be a semigroup partial order on X^*: that is, if λ, μ, ν are (possibly empty) words in X^* and $\mu \leq \nu$, then we have $\lambda\mu \leq \lambda\nu$ and $\mu\lambda \leq \nu\lambda$. We say $\leq_X$ satisfies the descending chain condition if any sequence $\lambda_1 >_X \lambda_2 >_X \cdots$ terminates. A reduction system S for RX^* is a set of rules of the form $s : \mu_s \mapsto a_s$, where μ_s lies in X^* and a_s lies in RX^*. The R-module maps $RX^* \to RX^*$ used to apply rules are known as reductions; these may consist of several rules performed sequentially. The two-sided ideal $\mathcal{I}(S)$ of RX^* is that generated by all elements $\mu_s - a_s$ for all rules $s \in S$. We say s is *compatible* with $\leq_X$ if a_s can be written as a linear combination of monomials strictly less than μ_s in $\leq_X$, and we say S is compatible with $\leq_X$ if each of its rules is.

An *overlap ambiguity* occurs when there are two rules s_1 and s_2 such that there exist monomials ν_2 and ν_1 with $\mu_{s_1}\nu_2 = \nu_1\mu_{s_2}$; it is said to be *resolvable* if there are reductions t_1 and t_2 with $t_1(a_{s_1}\nu_2) = t_2(\nu_1 a_{s_2})$. An *inclusion ambiguity* occurs when there are two rules s_1 and s_2 such that there exist monomials λ and ν with $\mu_{s_2} = \lambda\mu_{s_1}\nu$; it is said to be *resolvable* if there are reductions t_1 and t_2 with $t_1(\lambda a_{s_1}\nu) = t_2(a_{s_2})$.

A reduction t is said to act trivially on $a \in RX^*$ if $t(a) = a$, and if all reductions act trivially on a, we say a is irreducible. The set of irreducible elements arising from S is denoted $\mathrm{Irr}(S)$. A normal form of $a \in RX^*$ is an element $b \in \mathrm{Irr}(S)$ to which a can be reduced; it is not immediate that normal forms always exist or that they are unique.

The following theorem is part of Bergman's diamond lemma, which is proved in [1, Theorem 1.2].

Theorem 3.2.1 (Bergman) *Let R be an associative, commutative ring with 1. Let X be a nonempty set, let $\leq_X$ be a semigroup partial order on X^* and let S be a reduction system for RX^*. If S is compatible with $\leq_X$ and $\leq_X$ satisfies the descending chain condition then the following are equivalent*:

 (i) *All ambiguities in S are resolvable.*
 (ii) *Every element $a \in RX^*$ has a unique normal form which equals $t(a)$ for some reduction t.*
(iii) $RX^* = \mathrm{Irr}(S) \oplus \mathcal{I}(S)$.

Proposition 3.2.2 *The quotient of $\mathcal{A}H(P, \mathcal{C})$ described in Lemma* 3.1.6 *has as a free $\mathcal{A}$-basis the images of those heaps in $H(P, \mathcal{C})$ with property P2.*

Proof: We use Theorem 3.2.1. Let $R = \mathcal{A} = \mathbb{Z}[v, v^{-1}]$, let $X = P$ and let $\leq_X$ be the partial order such that $b_1 b_2, \ldots, b_l \leq_X c_1 c_2, \ldots, c_m$ if and only if $b_1 b_2, \ldots, b_l$ is a subsequence of $c_1 c_2, \ldots, c_m$; this is a semigroup partial order. There are two kinds of reduction rules.

For the first kind of reduction rule, we take

$$\mu_1 = pp_{i_1} \cdots p_{i_l} p$$

where the p_{i_n} (if there are any) are distinct from p and commute with p. This is precisely the condition for the occurrences of p to correspond by Proposition 3.1.4 to heap elements a and b for which $a < b$ is a balanced convex chain. The element a_1 in this case is given by

$$\delta p p_{i_1} \cdots p_{i_l} = \delta p_{i_1} \cdots p_{i_l} p.$$

For the second kind of reduction rule, we take

$$\mu_2 = p p_{i_1} \cdots p_{i_l} p' p_{j_1} \cdots p_{j_m} p$$

where all the p_{i_n} and p_{j_n} (if there are any) are distinct from p and commute with p, and p' does not commute with p. This is precisely the condition for the letters p, p', p to correspond respectively to heap elements a, b, c for which $a < b < c$ is a balanced convex chain with $\varepsilon(a) \neq \varepsilon(b)$. The element a_2 in this case is given by

$$a_2 = p p_{i_1} \cdots p_{i_l} p_{j_1} \cdots p_{j_m} = p_{i_1} \cdots p_{i_l} p_{j_1} \cdots p_{j_m} p.$$

In each case, a_i is a multiple of a strictly shorter monomial than μ_i. Since every heap is a finite set, $\leq_X$ has the descending chain condition.

We now show that all ambiguities are resolvable. Most of the possible inclusion ambiguities are easily seen to be resolvable. The only difficult case arises from words of the form

$$p \cdots (p'' \cdots p' \cdots p'') \cdots p,$$

where p and p'' are distinct and commute with each other but neither commutes with p', and where the both the whole word and the parenthetic expression shown are of the form μ_2 as above. This type of ambiguity resolves to

$$\delta p \cdots (p'' \cdots \widehat{p'} \cdots \widehat{p''}) \cdots \widehat{p}$$

(where the hats denote omission) whether the outermost or the innermost reduction rule is applied first.

The overlap ambiguities correspond to chains in the heap of the following kinds:

1. Convex chains $b < c < d$ with $\varepsilon(b) = \varepsilon(c) = \varepsilon(d)$.
2. Convex chains $b < c < d < e$ with either $\varepsilon(b) = \varepsilon(c) = \varepsilon(e) \neq \varepsilon(d)$ or $\varepsilon(b) = \varepsilon(d) = \varepsilon(e) \neq \varepsilon(c)$.
3. Chains $b < c < d < e < f$ with $b < c < d$ convex and $d < e < f$ convex and $\varepsilon(c) \neq \varepsilon(b) = \varepsilon(d) = \varepsilon(f) \neq \varepsilon(e)$ (but not necessarily with $\varepsilon(c) = \varepsilon(e)$).

These ambiguities are easily seen to be resolvable. The definition of property P2 shows that a word in P^* will be irreducible in this reduction system if and only if it has property

P2, so by Theorem 3.2.1 the heaps with property P2 represent a basis for the quotient algebra. □

The following theorem gives a nice interpretation of the dimensions of $\ker \partial$ of heaps in terms of the algebra $TL(P, \mathcal{C})$ and its structure constants.

Theorem 3.2.3

(i) *Let D be an arbitrary heap of $H(P, \mathcal{C})$. Then there exists a unique heap $G \in H(P, \mathcal{C})$ with property P2 such that $D = \delta^m G$ in the quotient algebra $TL(P, \mathcal{C})$. We have*

$$\dim \ker \partial_D = m + \dim \ker \partial_G.$$

In particular, if $H(P, \mathcal{C})$ is a regular class of heaps then $m = \dim \ker \partial_D$.

(ii) *Suppose $H(P, \mathcal{C})$ is a regular class of heaps and let E and F be heaps with property P2 regarded as basis elements of $TL(P, \mathcal{C})$ via Proposition 3.2.2. Then there is a basis element G of $TL(P, \mathcal{C})$ for which $E \circ F = \delta^{\dim \ker \partial_{E \circ F}} G$.*

Proof: Part (ii) follows from part (i), so we prove the former.

The existence of G comes from Theorem 3.2.1 and Proposition 3.2.2: the normal form of D will be a multiple of only one basis element because the reduction rules of Proposition 3.2.2 take monomials to multiples of monomials. The equation

$$\dim \ker \partial_D = m + \dim \ker \partial_G$$

comes from comparing Lemma 3.1.6 with lemmas 2.3.4 and 2.3.5. We apply the latter two lemmas repeatedly to the heap D until no further reductions are possible and we are left with the heap G. Along the way, the dimension of $\ker \partial$ will decrease by 1 precisely when an extra factor of δ appears in Lemma 3.1.6. For the last claim, observe that if $H(P, \mathcal{C})$ is regular then G has property P1 because it has property P2, and the conclusion follows from Proposition 2.2.3. □

3.3.   Heaps arising from Coxeter groups

In Section 3.4, we will apply the theory developed here to results involving Coxeter groups, and for this we need to explain the connection between Coxeter groups and heaps. This is a major theme of the paper [15].

We only consider the case of simply laced Coxeter groups, as this is all we need for our purposes. This simplifies the definitions somewhat.

Definition 3.3.1   Let Γ be a graph. The (simply laced) Coxeter group $W(\Gamma)$ associated to Γ is given by generators $\{s_i : i \in S\}$ where S is the set of vertices of Γ, and defining relations

$$\begin{aligned} s_i^2 &= 1, \\ s_i s_j s_i &= s_j s_i s_j \quad && \text{if } i, j \text{ are adjacent in } \Gamma, \\ s_i s_j &= s_j s_i \quad && \text{otherwise.} \end{aligned}$$

We now define the fully commutative elements of a Coxeter group. These were studied in Fan's thesis [4] in the simply laced case under the name "commutative elements", and a general definition was given by Stembridge [15, Section 1.1].

Definition 3.3.2 Let $W = W(\Gamma)$ be a Coxeter group as above. Clearly every element w can be written as

$$w = s_{i_1} s_{i_2}, \ldots, s_{i_r}$$

for some generators s_i. If r is minimal for a given w, we define the *length* of w to be r, and we call the associated product of generators *reduced*. If any two reduced expressions for w may be transformed into each other by repeated applications of relations of the form $s_i s_j = s_j s_i$ as in Definition 3.3.1, we say w is *fully commutative*. The set of fully commutative elements of W is denoted by W_c.

Elements of Coxeter groups give rise to heaps as follows; see also [15, Section 1.2]. The next procedure is inverse to that given by Proposition 3.1.4.

Definition 3.3.3 Let Γ be a graph, let S be the set of vertices of Γ and let $\mathcal{C}$ be the relation on S defined by $s_i \, \mathcal{C} \, s_j$ if and only if either $s_i = s_j$ or s_i and s_j are adjacent vertices. Let C be the complementary relation of $\mathcal{C}$. Let $w = s_{i_1}, \ldots, s_{i_l} \in S^*$ be an arbitrary word in the generators S of a simply laced Coxeter group $W = W(\Gamma)$. The word w gives a labelled heap $(E, \leq_E, \varepsilon)$ where $E = \{1, 2, \ldots, l\}$, $\varepsilon(j) = s_{i_j}$, and the relation $\leq_{\mathcal{C}}$ of condition 2 of Definition 1.1.1 can be defined by

$$a \leq_{\mathcal{C}} b \Leftrightarrow a \leq b \quad \text{and} \quad \varepsilon(a) \, \mathcal{C} \, \varepsilon(b),$$

where $\leq$ is the usual ordering on integers. The partial order $\leq_E$ is the transitive extension of $\leq_{\mathcal{C}}$, and the heap of $H(S, \mathcal{C})$ corresponding to the given labelled heap is by definition the heap of the word w. Moreover, the heap of w may be recovered from the element of $\mathrm{Co}(S, C)$ represented by w.

The following result is a special case of Stembridge's [15, Proposition 2.3].

Proposition 3.3.4 *Let Γ be a graph with set of vertices S. A word in S^* is a reduced expression for an element of $W_c \subset W = W(\Gamma)$ if and only if the associated heap (as in Definition 3.3.3) has property P2.*

We can use the graph Γ to define an analogue of the algebra $TL(P, \mathcal{C})$ of Section 3.1.

Definition 3.3.5 Let Γ be a graph with set of vertices S and let $\mathcal{A} = \mathbb{Z}[v, v^{-1}]$ with $\delta = v + v^{-1}$. We define the $\mathcal{A}$-algebra $TL(\Gamma)$ by generators $\{e_i : i \in S\}$ and relations

$$\begin{aligned}
e_i^2 &= \delta e_i, \\
e_i e_j e_i &= e_i && \text{if } i, j \text{ are adjacent in } \Gamma, \\
e_i e_j &= e_j e_i && \text{otherwise.}
\end{aligned}$$

Remark 3.3.6   Definition 3.3.5 is compatible with Definition 3.1.5: if $H(P, \mathcal{C})$ is a class of heaps with concurrency graph Γ, then $TL(\Gamma)$ is canonically isomorphic to $TL(P, \mathcal{C})$.

The following result is well known.

Proposition 3.3.7   *Let Γ be a graph with vertex set S and let w be a fully commutative element of the corresponding Coxeter group with reduced expression $s_{i_1} s_{i_2}, \ldots, s_{i_r}$. Then the element $e_w := e_{i_1} e_{i_2}, \ldots, e_{i_r}$ is independent of the choice of reduced expression and the set $\{e_w : w \in W_c\}$ is a free $\mathcal{A}$-basis of $TL(\Gamma)$ (called the monomial basis).*

Proof: The first assertion is an easy consequence of the definition of full commutativity. The second assertion follows from Proposition 3.2.2, Lemma 3.1.6 and Proposition 3.3.4.
□

3.4.   An application to algebra

We will be particularly interested in heaps whose concurrency graph is of type E_n, namely the graph shown in figure 2.

The numbering above is chosen to agree with that in Fan's paper [5, Section 6.3]. The elements of P will be the integers $0, 1, 2, \ldots, n - 1$. We emphasize that n is arbitrary; in particular, it can be greater than 8. Graphs of type A_{n-1} (respectively, D_{n-1}) may be obtained from the graph in figure 2 by omitting the vertices numbered 0 (respectively, 1). Graphs that are isomorphic to graphs of type A_n, D_n or E_n will be called *graphs of type ADE*. We will also be concerned with graphs of type $\tilde{A}_{n-1}$ for $n \geq 3$; these are n-gons with vertices numbered consecutively 1 up to n.

Interesting examples of our main results arise as follows.

Theorem 3.4.1 (*Fan*)   *Let $H(P, \mathcal{C})$ be a class of heaps whose concurrency graph is of type ADE or of type $\tilde{A}_{n-1}$ with n odd. Then $H(P, \mathcal{C})$ is a regular class of heaps.*

Proof: The *ADE* case is a restatement of [5, Lemma 4.3.1], and the case of type $\tilde{A}_{n-1}$ is a restatement of [6, Proposition 3.1.2].
□

The following subtle property of the algebras $TL(\Gamma)$ (see Definition 3.3.5) is originally due in the *ADE* case to Graham [7, Lemma 9.10, Lemma 9.13], who proved it using a


Figure 2.   Coxeter graph of type E_n.

complicated combinatorial argument. It is of key importance in certain applications as we explain in Section 4.

Proposition 3.4.2 *Let* Γ *be the concurrency graph of a regular class of heaps,* $H(P, \mathcal{C})$. *Consider an arbitrary word*

$$e = e_{i_1} e_{i_2}, \ldots, e_{i_r}$$

in the generators for TL(Γ), and define $w \in W_c$ *such that* $e = \delta^m e_w$ *for* $w \in W_c$.
 (i) *If we omit one generator from the word* e *to form*

$$e' = e_{i_1} \cdots \widehat{e_{i_l}} \cdots e_{i_r}$$

(where the hat denotes omission) and write $e' = \delta^{m'} e_{w'}$ *for some* $w' \in W_c$, *then* $m' \in \{m-1, m, m+1\}$.
(ii) *If* e *is of the form* e_x *for some* $x \in W_c$ *then* $m' = 0$.

Proof: Part (i) is a consequence of the Deletion Lemma (Theorem 2.1.1) and Theorem 3.2.3 (i), first with $D = e$, $G = e_w$, and then with $D = e'$, $G = e_{w'}$.

For part (ii), Theorem 2.4.2 shows that the heap of $x \in W_c$ is strongly acyclic, and the conclusion follows from Theorem 3.2.3(i). $\square$

4.   Applications and questions

We conclude with a survey of how the results of this paper are related to results in the literature, and some possible directions for future research.

4.1.   Canonical bases for Hecke algebra quotients

In [9], the author and J. Losonczy introduced canonical bases (IC bases) for the generalized Temperley–Lieb algebras of [7] and showed that in the case of the algebras $TL(\Gamma)$ where Γ is of type *ADE*, the basis of Proposition 3.2.2 is the canonical basis. This relies heavily on Proposition 3.4.2. An argument similar to our proof of Proposition 3.4.2 shows that any strongly acyclic heap corresponds to a canonical basis element for $TL(\Gamma)$ that is given by a monomial in the generators e_i. (Not all canonical basis elements are of this form.)

Although Theorem 2.4.2(ii) has so far only appeared as a by-product, it has some nice applications of its own, one of which is to allow a recurrence formula for canonical basis elements for $TL(\Gamma)$ similar to that given by Kazhdan and Lusztig [11, Section 2.2] in the case of Hecke algebras. It also gives recurrence formulae for analogues of inverse Kazhdan–Lusztig polynomials (denoted by $\tilde{Q}_{x,w}$ in [9, Lemma 3.5]).

Our results also have applications to generalized Temperley–Lieb algebras arising from non-simply-laced Coxeter graphs. These algebras have bases indexed by heaps satisfying conditions similar to, but weaker than property P2; see [15, Proposition 2.3] for the exact condition. The canonical bases for these algebras are described in [8] for Coxeter types B and H. In these two special cases, the heaps indexing the basis are acyclic, which is

essentially a neater restatement of [8, Lemma 3.1.1]; the Deletion Lemma (Theorem 2.1.1) then gives [8, Lemma 3.1.3]. Another result of [8] that is more easily phrased in terms of the map ∂ is [8, Proposition 3.1.9], which classifies the possible image vertices in a basis heap.

In some cases, such as Coxeter type B, it is possible to find for every element w in the Coxeter group a reduced expression whose heap (see Definition 3.3.3) is acyclic. The results of [10, Section 2] are essentially consequences of this observation. The analogous claim in type D is false, which makes Losonczy's argument in [12] much more difficult.

We remark that if the basis heaps mentioned above are acyclic, it is possible to state a (more complicated) version of Theorem 3.2.3 for the non-simply-laced case.

We hope to give details of these applications separately.

4.2. Computing $\dim \ker \partial_E$

Certain of the algebras $TL(\Gamma)$ may be understood by a calculus of diagrams rather than using generators and relations. In these cases, a word in the generators e_i of $TL(\Gamma)$ may be represented by a diagram, and Proposition 3.1.4 then shows how to represent a heap as a diagram. We outline here how these diagrams may be used to calculate $\dim \ker \partial_E$ essentially by inspection. The proofs are too long to present here, particularly in the case of type E_n below, but we hope to give details separately.

Type A_n. Suppose $H(P, \mathcal{C})$ is a class of heaps with concurrency graph Γ of type A_n, in other words, Γ is a line. In this case, $TL(\Gamma)$ is the Temperley–Lieb algebra of [16] which has a well known diagram calculus whose origins can be traced back to [14]. If E is a heap of $H(P, \mathcal{C})$ then $\dim \ker \partial_E$ is the number of loops occurring in the corresponding diagram.

Type E_n. Suppose $H(P, \mathcal{C})$ is a class of heaps with concurrency graph Γ of type E_n, so that it is isomorphic to a graph such as that shown in figure 2. Let E be a heap in a class $H(P, \mathcal{C})$ that has concurrency graph Γ. A diagram calculus for $TL(\Gamma)$ was described by tom Dieck in [3], and in fact this gives a faithful representation of the algebra, although this is not proved in [3]. In this case the diagrams are non-intersecting curves drawn inside a rectangle whose endpoints lie on the boundary, and there are certain discs ("pillars") lying in the connected components of the complement of these curves. If E is a heap of $H(P, \mathcal{C})$ then

$$\dim \ker \partial_E = a + \sum_{i \in I} \max(0,\ b_i - 1),$$

where a is the number of loops containing no pillars, I is the set of connected components in the diagram (including the insides of loops) and b_i is the number of pillars in component i.

Type $\tilde{A}_{n-1}$. Suppose $H(P, \mathcal{C})$ is a class of heaps with concurrency graph Γ of type $\tilde{A}_{n-1}$ ($n \geq 3$); in other words, let Γ be an n-gon. A diagram calculus for $TL(\Gamma)$ is given in [6, Section 4]. The diagrams consist of certain curves inscribed on the surface of a cylinder

whose endpoints (if any) lie on the boundary. A diagram contains a number a of loops contractible on the cylinder, and a number b of loops that are not contractible on the cylinder. (These numbers may be zero.) If E is a heap of $H(P, \mathcal{C})$ then

$$\dim \ker \partial_E = a + c(\max(0,\ b-1)),$$

where $c = 1$ if $4 \mid n$ or k has characteristic 2, and $c = 0$ otherwise.

4.3. Concluding questions

In light of the results of this paper, it would be interesting to have an answer to the following graph theoretic problem.

Problem 4.3.1   Find necessary and sufficient conditions on a graph Γ for it to be the concurrency graph of a regular class of heaps.

   In a future paper, we will present a complete solution to Problem 4.3.1 under the assumption that the set of pieces, P, is finite. That paper will also show how, with considerably more work, Theorem 2.4.2 may be sharpened.
   Another intriguing direction for future research is suggested by the following

Question 4.3.2   Do the results of this paper have applications to concurrency in computer science?

Acknowledgments

References

 1. G.M. Bergman, "The diamond lemma for ring theory," *Adv. Math.* **29** (1978), 178–218.
 2. P. Cartier and D. Foata, "Problèmes combinatoires de commutation et réarrangements," *Lecture Notes in Mathematics* **85** (1969), Springer-Verlag New York/Berlin.
 3. T. tom Dieck, "Bridges with pillars: A graphical calculus of knot algebra," *Topology Appl.* **78** (1997), 21–38.
 4. C.K. Fan, "A Hecke algebra quotient and properties of commutative elements of a Weyl group," Ph.D. thesis M.I.T. 1995.
 5. C.K. Fan, "Structure of a Hecke algebra quotient," *J. Amer. Math. Soc.* **10** (1997) 139–167.
 6. C.K. Fan and R.M. Green, "On the affine Temperley–Lieb algebras," *Jour. L.M.S.* **60** (1999), 366–380.
 7. J.J. Graham, "Modular representations of Hecke algebras and related algebras," Ph.D. thesis University of Sydney, 1995.
 8. R.M. Green, "Decorated tangles and canonical bases," *J. Algebra* **246** (2001), 594–628.
 9. R.M. Green and J. Losonczy, "Canonical bases for Hecke algebra quotients," *Math. Res. Lett.* **6** (1999), 213–222.
10. R.M. Green and J. Losonczy, "A projection property for Kazhdan–Lusztig bases," *Int. Math. Res. Not.* **1** (2000), 23–34.

11. D. Kazhdan and G. Lusztig, "Representations of Coxeter groups and Hecke algebras," *Invent. Math.* **53** (1979), 165–184.

12. J. Losonczy, "The Kazhdan–Lusztig basis and the Temperley–Lieb quotient in type D," *J. Algebra* **233** (2000), 1–15.

13. A. Mazurkiewicz, "Trace theory," Petri nets, applications and relationship to other models of concurrency Lecture Notes in Computer Science 255, Springer Berlin–Heidelberg–New York, 1987, pp. 279–324.

14. R. Penrose, "Angular momentum: An approach to combinatorial space-time," *Quantum Theory and Beyond* E. Bastin, (ed.), Cambridge University Press, Cambridge, 1971, pp. 151–180.

15. J.R. Stembridge, "On the fully commutative elements of Coxeter groups," *J. Algebraic Combin.* **5** (1996), 353–385.

16. H.N.V. Temperley and E.H. Lieb, "Relations between percolation and colouring problems and other graph theoretical problems associated with regular planar lattices: Some exact results for the percolation problem," *Proc. Roy. Soc. London Ser. A* **322** (1971), 251–280

17. G.X. Viennot, "Heaps of pieces, I: basic definitions and combinatorial lemmas," *Combinatoire Énumérative*, G. Labelle and P. Leroux (eds.), Springer-Verlag, Berlin, 1986, pp. 321–350.