Shellability of a poset of polygonal subdivisions

Torsten Ekedahl

Stockholms universitet
SE-106 91 Stockholm
Sweden

teke@math.su.se

Submitted: Oct 31, 2008; Accepted: Dec 7, 2009; Published: Dec 15, 2009
Mathematics Subject Classification: 52B22, 18D50

Dedicated to Anders Björner on the occasion of his sixtieth birthday

Abstract

We introduce a sequence of posets closely related to the associahedra. In this we are motivated by reasons similar to those of Stasheff in the case of the associahedra. We make a study of this poset showing that it has an inductive structure with proper downwards intervals being products of smaller posets in the same series and associahedra. Using this we also show that they are thin dual CL-shellable and in particular that they are the face poset of a regular cell decomposition of the ball.

1 Introduction

The purpose of this article is to introduce a sequence of posets closely related to the face lattices of the associahedra and study their combinatorial properties, in particular it will be shown that they are shellable. The origin of these posets are in principle not relevant for such a study, nevertheless I shall start by briefly discussing it. The associahedra are relevant to the description of products which are associative only up to homotopy ("A_∞"-spaces). The prototypical such example is the path space of a topological space where the composition of paths is not associative but is associative up to homotopy and two maps from one space to another constructed out of such homotopies are homotopic and so on. Suppose now that the space is a manifold M and that we are really only interested in smooth paths. The problem is that the composition of smooth paths is usually not smooth. The solution would seem to be to smooth the composition but the problem then is that such a smoothing is not unique. Thus one is forced to speak about *a* composition and will have to contend with the ambiguities inherent in that. A direct way of expressing a composition is as a smooth map from the standard 2-simplex Δ_2 to

M, where the original two paths are the restriction of that map to the first two edges and the particular composite is the restriction to the third edge. The map itself is then the particular smoothing of the composition of the first two edges to the third edge. We write this as $a \cdot b \to ab$, where a and b are the original paths, $a \cdot b$ represents the composition, ab the smoothing and the arrow represents a particular smoothing with ab as end result.

Following the pattern of higher coherence conditions that made the associahedra appear in the first place, we assume that we have different choices of compositions $a \cdot b \to ab$, $b \cdot c \to bc$, $a \cdot bc \to abc$ and $ab \cdot c \to abc$ and we are looking for some supplementary coherence condition that would express that these choices of composites are coherently associative. In this particular case the composites fit together to give a smooth map from the boundary of Δ_3 to M and we could demand that this map extend smoothly to Δ_3 itself. However, we would rather have conditions that can be formulated in terms of composition of paths and the A_∞-structure. Hence using the homotopies given by the smoothings together with the associativity homotopy $a \cdot (b \cdot c) \xrightarrow{a \cdot b \cdot c} (a \cdot b) \cdot c$ gives us a mapping from the boundary of the pentagon into the path space of M:

$$
\begin{tikzcd}
a \cdot (b \cdot c) \arrow[rr, dash, "a\cdot b\cdot c"] \arrow[d, dash] & & (a \cdot b) \cdot c \arrow[d, dash] \\
a \cdot bc \arrow[dr, dash] & & ab \cdot c \arrow[dl, dash] \\
& abc &
\end{tikzcd}
$$

The coherence condition should then be that this map extend to the full pentagon. The pentagon of course is the second associahedron indicating that there is indeed a relation between this condition and the associahedron. If one analyses the next coherence condition, one arrives at something that is not an associahedron (see Fig. 2) but is visibly a polyhedron. In this note we shall give a general definition of these coherence conditions leading to posets that share the simplest combinatorial properties with the face posets of polyhedra such as being a lattice and being shellable. From the point of view of the original motivation shellability has the important consequence that the posets are face posets of a regular cell decomposition of a ball. However, its close relation with the associahedra seems to me an indication that they should be interesting from a purely combinatorial perspective.

Just as for the associahedra themselves everything is best phrased in terms of triangulations of n-gons. A new notion appears however. We shall need to consider not just triangulations of a fixed polygon but also triangulations of a subpolygon all of whose vertices are vertices of the larger polygon as well as a simple way of passing from a partial triangulation to a partial triangulation of a smaller polygon.

Concerning specific results our principal poset is $\mathcal{CC}_n$, the poset of compound collapses; by adding a smallest element we get $\mathcal{CC}_n^+$ which is proved to be a graded poset (Corollary 3.5) and a lattice (Proposition 3.7). Finally it is shown in Theorem 4.2 that $\mathcal{CC}_n^+$ is dual CL-shellable and thin. Analogously to the case of the associahedron we also have a local structure theorem in that every interval $]\hat{0}, x]$ is a product of smaller posets of compound collapses and face posets of associahedra (Proposition 3.3). As $\mathcal{CC}_n$ will be seen to be very analogous to the face poset of the associahedron the following problem comes very naturally.

Problem: Is $\mathcal{CC}_n$ the face poset of a convex polytope?

My own experience with polytopes is too meager to allow me to venture an opinion.

2 A low dimensional example

To motivate our subsequent deliberations we start by describing $\mathcal{CC}_5$, the first non-trivial example. (Note that in order to simplify we shall disregard some terminological distinctions that will be made later.)

We can as usual describe cell decompositions of the n-gon as partial parenthesisations of a product of $n-1$ symbols. As we shall deal also with k-gons for $k \leqslant n$ we shall do the following: We attach symbols, letters starting with a in our example, to the initial edges of the n-gon and then to an arbitrary edge we attach the concatenation of the symbols in order of the initial edges connecting the initial vertex of the edge to the final. (An initial vertex has been chosen and the polygon is then oriented counterclockwise.) Hence, if we attach a, b, and c to the initial edges of the 4-gon, the edge from the first to the third vertex will get label ab. The final edge will in general get the label obtained by concatenating all the labels of the initial edges. For instance if we start with a 6-gon with initial edge labels a, b, c, d, and e, then the 6-gon itself will correspond to the unparenthesised expression $a \cdot b \cdot c \cdot d \cdot e$ whereas the 4-gon consisting of the edge from the first to the third vertex, from the third to the fourth, from the fourth to the sixth, and the final edge will correspond to the unparenthesised expression $ab \cdot c \cdot de$ (cf., Fig. 1). The full concatenation $abcde$ will then correspond to the 1-gon consisting just of the

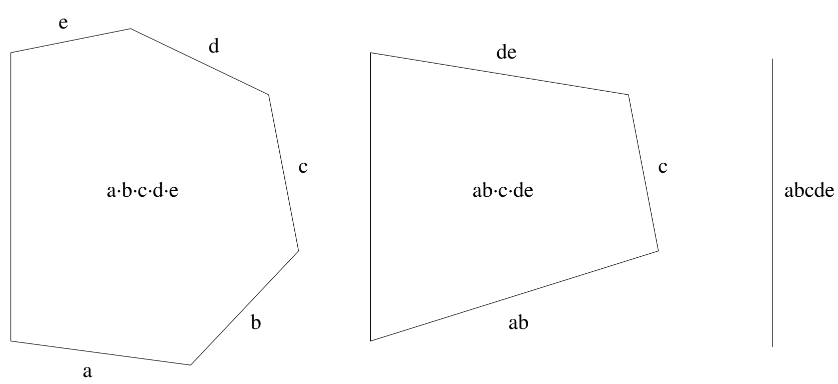

Figure 1: A 6-gon and a 4-gon and 1-gon contained in it.

final edge (*idem*).

Furthermore a cell decomposition of a k-gon, which is the convex hull of a set of vertices of the n-gon containing the initial and final vertex, will correspond to a partially parenthesised expression of concatenations of the labels of the initial edges such that the labels appear in increasing order. Thus $ab \cdot cd$ corresponds to the cell decomposition, using edges of the 5-gon, consisting of the edge connecting the first vertex to the third, the edge connecting the third to the fourth and (as always) the final edge. A compound collapse then corresponds to replacing possibly several but disjoint unparenthesised subexpressions with the corresponding concatenation. Continuing the last example $(a\cdot b)\cdot(c\cdot d)$ collapses to any of $ab \cdot (c \cdot d)$, $(a \cdot b) \cdot cd$, and $ab \cdot cd$.

Note further that according to our definitions any cell decomposition collapses to itself and we shall denote $C \to C \in \mathcal{CC}_n$ also by just C.

In Figure 2 we have assembled all the collapses (or as they shall be called later compound collapses) inside of a 5-gon except the ones corresponding to $a \cdot b \cdot c \cdot d$ and $a \cdot b \cdot c \cdot d \to abcd$. Edges represent incompletely parenthesised expressions with their (two as it were) complete parenthesisations at their ends. The arrows represent compound collapses, where the dotted arrows give those that correspond to 2-cells in what is obviously a regular cell decomposition of a 2-cell. Together with $a \cdot b \cdot c \cdot d$ and $a \cdot b \cdot c \cdot d \to abcd$ we get a regular cell decomposition of the 3-cell.

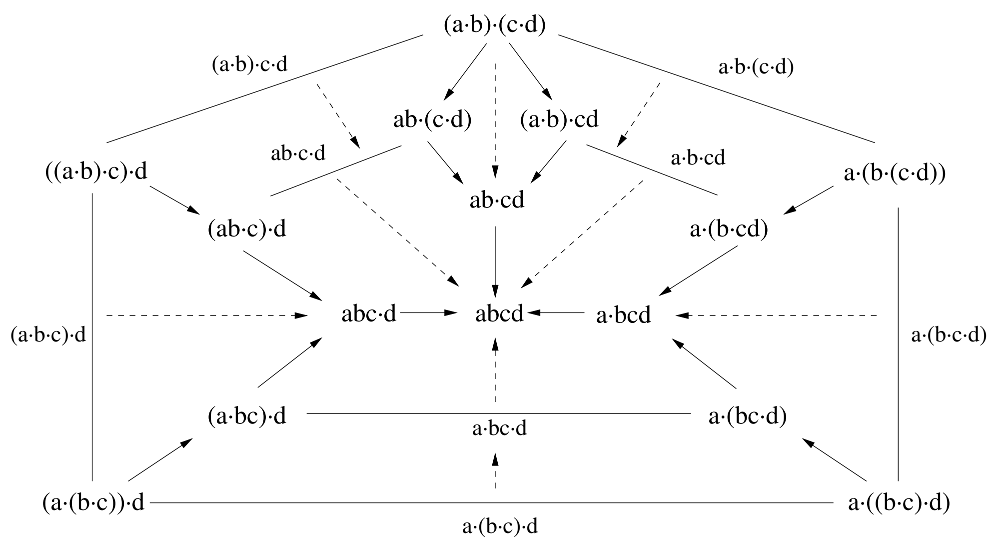

Figure 2: Convex cell decompositions of a 5-gon.

3 Combinatorics

We start with a convex planar n-gon, P_n. In inductive arguments we shall deal with many polygons and use P_n for the polygon currently considered, we shall also not particularly distinguish between a polygon as a convex set and as the set of its boundary edges or even vertices. We also pick one of its edges which we call the *final edge* and call the other edges its *initial edges*. We further orient the final edge so that it gets an *initial vertex* and a *final vertex*. The other edges are then oriented so that they point away from the initial vertex and towards the final one. (If P_n is explicitly embedded in the plane it seems reasonable to orient the edges so that one moves counterclockwise when moving along the initial edges from the initial to the final vertex.) In this way we also get a total order on the set of vertices of P_n; we go from a lower vertex to a higher one by moving along initial edges. A convex polygon whose vertices are a subset of the vertices of P_n will then have a unique edge connecting the first of its vertices to its last, called its *final edge*, the rest are its initial edges and come with natural orientations just as for P_n itself. We shall call such a polygon a *cap* or *cell*. For caps, but *not* for cells, we also allow the degenerate case consisting of just one edge.

We let $\mathcal{CD}_n$ be the set of *cell decompositions* of P_n, where a cell decomposition is a subset C of the set of edges connecting the vertices of P_n such that

- the final edge is a member of C,
- edges of C meet only at vertices of P_n and
- any edge which is a part of the boundary of the convex hull $\overline{C}$ of C is a member of C.

We shall also call the convex hull, $\overline{C}$, of C the *support* of C. If C is a convex cell decomposition, the closure of a component of $\overline{C} \setminus C$ (i.e., $\overline{C}$ minus the edges of C) is a cell and will be called a cell of the cell decomposition. It is clear that the boundary of $\overline{C}$ is a polygon, the *boundary polygon* of C.

When speaking of the *interior*, Z^o, of a subset of $\mathbf{R}^2$ we shall mean the *relative interior*, i.e., the interior of Z as a subset of the smallest affine subspace containing Z. If $C \in \mathcal{CD}_n$, then we put $C \cdot Z := \{\ell \in C \mid \ell^o \subseteq Z\}$.

If $C, D \in \mathcal{CD}_n$, then we say that D is a *compound collapse* of C, denoted $C \to D$, if $D \subseteq C$ and $C \setminus D$ is contained in the boundary polygon of C. We say then that C is the *collapse*r, denoted $(C \to D)_r$, and D the *collapse*e, denoted $(C \to D)_e$, of the compound collapse. As C is the union of D and the boundary polygon of C when specifying a compound collapse it is enough to specify D and the boundary polygon something we shall do without further comment. If $C, D \in \mathcal{CD}_n$, then D is said to be a *refinement* of C, denoted $C \geqslant D$, if C and D have the same support and $C \subseteq D$. If $C, D \in \mathcal{CD}_n$, then we say that D is a *collapse of* C, denoted $C \twoheadrightarrow D$, if there is a refinement C' of C and a compound collapse $C' \to D$, i.e., $C \geqslant C' \to D$. If $C, D \in \mathcal{CD}_n$, then we say that D is a *lapse*, denoted $C \rightarrowtail D$, of C if D is a refinement of some D' and there is a compound collapse $D' \to C$, i.e., $D \leqslant D' \to C$. We now let $\mathcal{CC}_n$ be the set of pairs $C, D \in \mathcal{CD}_n$ such that $C \to D$. We define a relation on $\mathcal{CC}_n$ by $(C, D) \geqslant (E, F)$ if $C \twoheadrightarrow E$ and $D \rightarrowtail F$. If P is a particular convex n-gon we shall also use $\mathcal{CC}(P)$ for the compound collapses of P.

Remark: These definitions could probably do with some further elaboration. An example of a compound collapse consists of the collapsing of a cap formed by some boundary edges of C and one interior edge of D onto the interior edge or more directly one just removes its boundary edges. A general compound collapse consists of several such collapses done simultaneously which means in particular that the caps involved meet at most in vertices. For a collapse one wants to collapse certain convex regions onto one of their sides but is stopped from doing it because some of the sides do not belong to C though the regions are contained in $\overline{C}$ and all but one of their faces lie in the boundary of C. One thus starts by adding those sides to C and has the liberty of adding other edges too as long as they don't lie in one of the regions. A lapse is formulated in a way so as to make it dual to a collapse but as we shall see in the next lemma one gets from C to D by adding any number of edges (in such a way so that we still get a cell decomposition). We shall later give a very precise description of the cover relation associated to $\geqslant$. It will indeed be seen to be a partial order.

If $C, D \in \mathcal{CD}_n$ and $\overline{C} \subseteq \overline{D}$, then a *residual boundary edge* of the pair $(\overline{D}, \overline{C})$ is an edge which lies in the boundary of $\overline{D}$ but is not contained in $\overline{C}$. The set of residual boundary edges will be denoted $R(\overline{D}, \overline{C})$.

Lemma 3.1 *Let $A, B \in \mathcal{CD}_n$.*

i) We have that $A \to B$ precisely when $\overline{B} \subseteq \overline{A}$, $A\cdot(\overline{A}\setminus\overline{B})^o = \emptyset$, and $A = B \cup R(\overline{A}, \overline{B})$.

ii) We have that $A \twoheadrightarrow B$ precisely when $\overline{B} \subseteq \overline{A}$, $A\cdot(\overline{A}\setminus\overline{B})^o = \emptyset$, and $A \subseteq B \cup R(\overline{A}, \overline{B})$.

iii) We have that $A \rightarrowtail B$ precisely when $A \subseteq B$.

Proof: Assume that $A \to B$ as $B \subseteq A$ we have $\overline{B} \subseteq \overline{A}$ and as all $\ell \in C \setminus D$ lie in the boundary of $\overline{C}$ we have $A \cdot (\overline{A} \setminus \overline{B})^o = \emptyset$ but also that $\ell \in R(\overline{A}, \overline{B})$ which gives $A = B \cup R(\overline{A}, \overline{B})$. Conversely, assume $\overline{B} \subseteq \overline{A}$, $A \cdot (\overline{A} \setminus \overline{B})^o = \emptyset$, and $A = B \cup R(\overline{A}, \overline{B})$ the last clearly giving $B \subseteq A$. An $\ell \in A \setminus B$ lies in $R(\overline{A}, \overline{B})$ and thus in the boundary of $\overline{A}$. This proves i).

Assume now $A \twoheadrightarrow B$ and assume that $A \geqslant C \to B$. By i) we have $\overline{B} \subseteq \overline{C}$ and by definition $\overline{A} = \overline{C}$ so that $\overline{B} \subseteq \overline{A}$. As $A \subseteq C$ and $C\cdot(\overline{C}\setminus\overline{B}) = \emptyset$ by i) we get $A\cdot(\overline{A}\setminus\overline{B}) = \emptyset$ and also $A \subseteq C = B \cup R(\overline{C}, \overline{B})$ again by i) and $R(\overline{C}, \overline{B}) = R(\overline{A}, \overline{B})$ as $\overline{A} = \overline{C}$. Conversely, assume $\overline{B} \subseteq \overline{A}$, $A \cdot (\overline{A} \setminus \overline{B})^o = \emptyset$, and $A \subseteq B \cup R(\overline{A}, \overline{B})$. Putting $C := B \cup R(\overline{A}, \overline{B})$, C and A have the same support so that $C = B \cup R(\overline{C}, \overline{B})$ and we have $C \to B$ by i) and as $A \subseteq C$we also have $A \geqslant C$ which finishes the proof of ii).

Finally, if $A \rightarrowtail B$ there is a C with $B \leqslant C \to A$ we have $B \supseteq C$ and $C \supseteq A$ by i). Conversely, if $B \supseteq A$ we let C consist of the union of A and the elements of B that are not part of the boundary of $\overline{B}$. Then we have $B \leqslant C \to A$, which gives $A \rightarrowtail B$. □

Proposition 3.2 *The relation $\leqslant$ on $\mathcal{CC}_n$ is a partial order.*

Proof: The only property which is not clear is transitivity. Hence assume that we have $A \to B$, $C \to D$, and $E \to F$ as well as $A \twoheadrightarrow C$, $C \twoheadrightarrow E$, $B \rightarrowtail D$, and $D \rightarrowtail F$. It is obvious from Lemma 3.1 that $B \rightarrowtail F$. Now, again from Lemma 3.1, we have that $\overline{B} \subseteq \overline{D} \subseteq \overline{F} \subseteq \overline{E} \subseteq \overline{C} \subseteq \overline{A}$ and $A \cdot (\overline{A} \setminus \overline{B})^o = \emptyset$ and thus $A \cdot (\overline{A} \setminus \overline{E})^o = \emptyset$. What is left to show is then that $A \subseteq E \cup R(\overline{A}, \overline{E})$. Let $\ell \in A$ and assume first that $\ell \in A \cdot \overline{E}$. As we have that $A \subseteq C \cup R(\overline{A}, \overline{C})$ and $\overline{E} \subseteq \overline{C}$ we get that $\ell \in C$ and thus that $\ell \in C \cdot \overline{E}$. The fact that $C \subseteq E \cup R(\overline{C}, \overline{E})$ then implies that $\ell \in E$. Hence we may assume that $\ell \in A \cdot (\overline{A} \setminus \overline{E})$ but as we already know that $A \cdot (\overline{A} \setminus \overline{E})^o = \emptyset$ this implies that $\ell \in R(\overline{A}, \overline{E})$ which finishes the proof. □

If $C \to D$ is a compound collapse, a *cap* of it is the closure of a component of $\overline{C} \setminus \overline{D}$. It is clear that it is a cap in the sense introduced above.

We let $\mathcal{K}_n$ be the poset of partial triangulations of an $n+1$-gon.[1] With this indexing one of the consequences of the next proposition is that $\mathcal{K}_{n-1}$ is a coatom of $\mathcal{CC}_n$. In order to be able to deal with several copies of this poset at the same time, if P is a polygon then we shall also use $\mathcal{K}(P)$ for the partial triangulations of P.

Proposition 3.3 *Suppose $(C \to D) \geqslant (E \to F)$ in $\mathcal{CC}_n$. Let $E' := E \cdot \overline{D}$ and for each cap γ of $C \to D$ let $E_\gamma := E \cdot \gamma$ resp. $F_\gamma := F \cdot \gamma$. Then E' is a refinement of D, $E_\gamma \to F_\gamma$, $E = E' \cup \cup_\gamma E_\gamma$ and $F = E' \cup \cup_\gamma F_\gamma$. Conversely, given $C \to D$, a refinement E' of D and compound collapses $E_\gamma \to F_\gamma$ for each cap γ of $C \to D$ then we have a compound collapse $E \to F$ where $E = E' \cup \cup_\gamma E_\gamma$ and $F = E' \cup \cup_\gamma F_\gamma$ and $(C \to D) \geqslant (E \to F)$.*

In particular, the interval of elements below $C \to D$ is isomorphic to $\mathcal{K}_m \times \prod_\gamma \mathcal{CC}_{m_\gamma}$, where $\overline{D}$ is an $m+1$-gon and the cap γ is an m_γ-gon.

Proof: By the definition of compound collapse we have that $D \subseteq C$ and $F \subseteq E$ and by Lemma 3.1 $D \subseteq F$ so that D is contained in C, E and F. In particular, E' is a refinement of D. Now, γ lies in the closure of $\overline{C} \setminus \overline{D}$ which implies that all boundary edges of $\overline{E_\gamma}$ lies in E_γ. As $E \to F$, $F \cdot (\overline{E_\gamma} \setminus \overline{F_\gamma}) = \emptyset$ which gives $E_\gamma \to F_\gamma$. As $\overline{C}$ is the union of $\overline{D}$ and the (support of) the caps γ we get $E = E' \cup \cup_\gamma E_\gamma$ and $F = E' \cup \cup_\gamma F_\gamma$. The converse is similar.

This result then shows that an element below $C \to D$ is specified by a refinement of D and a compound collapse of each cap which gives the bijection of the final result and it is clear that partial order on such elements is the product order. □

Note that $\mathcal{CC}_n$ has a maximal element $C \to D$ where D is the final edge and C is the n-polygon itself. As is standard we shall denote it $\hat{1}$.

Proposition 3.4 *i) For $1 < i < n$ we let F_i be the compound collapse $C_i \to D_i$ where D_i consists of just the final edge and C_i consists of boundary edges of the convex hull of all the vertices but the i'th. For S a subset of $\{1, n\} \subset S \subseteq \{1, \dots, n\}$ we let F'_S be the compound collapse $C'_S \to D'_S$ where C'_S is the polygon P_n and D'_S is the convex hull of the vertices whose positions appear in S. Then the F_i and the F'_S are exactly the cocovers of the maximal element of $\mathcal{CC}_n$.*

ii) The covers of a compound collapse $A \to B$ are of one of the following forms and all such forms are covers:

- *An internal edge is removed from B (*"edge removal"*).*
- *A cap of B is collapsed (*"cap collapse"*).*
- *A small cap of A is added (*"small cap addition"*).*

Proof: If $(C' \to D') \geqslant (C \to D)$, then $D' \subseteq D$ and C' is a coarsening of a cell decomposition that has a compound collapse to C. In particular if D is a polygon then either $D' = D$ or D' consists only of the final edge. On the other hand if C contains the

[1] There are different choices of the indexing for the associahedra in the literature. We adopt here Stasheff's original, [St63a], indexing.

boundary of P_n then there is no non-trivial compound collapse to C and hence C would be a refinement of C'. Together this shows that if $C \to D = F'_S$ then $C' \to D' = \hat{1}$ so that F'_S is a cocover of $\hat{1}$. The argument for F_i is similar.

Conversely assume $C \to D$ is a cocover of $\hat{1}$. By our previous considerations we conclude that D is a polygon. Assume first that it is not equal to the final edge. Then if D is not equal to P_n, then there is a boundary edge ℓ of D that is not in the boundary of P_n. It is the final edge of a polygon γ whose edges other than ℓ lie in the boundary of P_n. Putting $C' := C \cup \gamma \setminus \ell$ we have that $C' \to D$ and $(C' \to D) > (C \to D)$ and $C' \to D$ is not equal to $\hat{1}$ which is a contradiction and the conclusion is that $C \to D = F'_S$ for some S. The case when D is the final edge leads in a similar fashion to F_i. This proves i).

As for ii) it follows immediately from Proposition 3.3 and i). $\square$

We now let $\mathcal{K}_n^+$ be the poset obtained by adding an artificial minimal element $\hat{0}$ to $\mathcal{K}_n$ and similarly $\mathcal{CC}_n^+$ is obtained from $\mathcal{CC}_n$ by adding $\hat{0}$.

Corollary 3.5 *$\mathcal{CC}_n^+$ is graded of length $n-2$.*

Proof: We argue by induction on n. It is enough, by Proposition 3.4, to show that the intervals $[\hat{0}, F_i]$ and $[\hat{0}, F'_S]$ are graded of length $n-3$. In the case of the F_i it follows from Proposition 3.3 that $[\hat{0}, F_i]$ is isomorphic to $\mathcal{CC}_{n-1}^+$ which by induction is indeed graded of length $n-3$. In the case of F_S, let $\gamma_1, \dots, \gamma_k$ be the caps of F'_S. We assume that D is an m-gon and γ_j an m_j-gon. Then we have that $\sum_j m_j - 2 + m = n$. On the other hand, by Proposition 3.3 we have that $[\hat{0}, F'_S]$ is isomorphic to $\left(\mathcal{K}_{m-1} \times \prod_j \mathcal{CC}_{m_j}\right)^+$. Now, $\mathcal{K}_{m-1}^+$ is graded of length $m-2$ and by induction $\mathcal{CC}_{m_j}^+$ is graded of length $m_j - 2$ and hence $[\hat{0}, F'_S]$ is graded of length $\sum_j m_j - 2 + m - 2 = n - 2$. $\square$

We say that two cell decompositions C and D are *compatible* if for any edge ℓ of C and any edge ℓ' of D distinct from ℓ, ℓ and ℓ' intersect at most in vertices of P_n.

Lemma 3.6 *If C and D are compatible then $C \cup D$ is a cell decomposition.*

Proof: The only condition that is not trivially fulfilled is that $C \cup D$ contains any edge of the boundary polygon of its convex hull. Assume therefore that E is such an edge which is not contained in $C \cup D$. Its two endpoints p and q must be endpoints of elements of $C \cup D$ but E cannot be an edge of C or D. Hence we may assume that p is larger than q (in the total order of the vertices of P_n) and is the final endpoint of $E' \in C$ and q is the initial endpoint of $E'' \in D$. Now, as E'' lies in the convex hull of $C \cup D$ its final endpoint must come after the final endpoint of E, i.e., p and similarly the initial endpoint of E' must come before q. That however implies that E' and E'' must meet in the interior of the convex hull $C \cup D$ which is a contradiction. $\square$

Proposition 3.7 *$\mathcal{CC}_n^+$ is a lattice. More precisely, the infimum of $A \to B$ and $A' \to B'$ is $\hat{0}$ unless B and B' are compatible and $\overline{B} \cup \overline{B'} \subseteq \overline{A} \cap \overline{A'}$. If these conditions are fulfilled the infimum $A'' \to B''$ is given as follows: $B'' = B \cup B'$ and the boundary of A'' is the largest polygon contained in $\overline{A} \cap \overline{A'}$.*

Proof: If we have $(A \to B), (A' \to B') \geqslant (C \to D)$, then $D \supseteq B, B'$ so that B and B' are compatible. Furthermore we have $\overline{B} \subseteq \overline{D} \subseteq \overline{CA'}$ and by symmetry $\overline{B} \cup \overline{B'} \subseteq \overline{A} \cap \overline{A'}$. Hence, if either B and B' are not compatible or $\overline{B} \cup \overline{B'} \subsetneq \overline{A} \cap \overline{A'}$, then the infimum is $\hat{0}$. Assume now that the conditions are fulfilled. For any $C \to D$ dominated as above by $A \to B$ and $A' \to B'$ we have $B'' \subseteq D$. By Proposition 3.6 B'' is a cell decomposition. It is also clear that $A \to B''$ and $A' \to B''$ and as we have $(A \to B) \geqslant (A \to B'')$, $(A' \to B') \geqslant (A' \to B'')$ as well as $C \to D$ being dominated by $A \to B''$ and $A' \to B''$ we may assume that $B = B'$. This means that the difference between B and A resp. A' are just edges that lie in the respective boundaries. Also, $\overline{C} \subseteq \overline{A} \cap \overline{A'}$ and it is clear that there is a largest polygon contained in $\overline{A} \cap \overline{A'}$ and adding its boundary edges to B thus gives an infimum. $\square$

4 Shellability

Our aim is now to show that $\mathcal{CC}_n^+$ is a thin shellable poset. More precisely we shall show that it is dual CL-shellable, cf., [BW83]. In our proof of the dual CL-shellability of $\mathcal{CC}_n$ we are going to use the equivalent condition of having a recursive coatom ordering. In the recursion necessary to verify that a particular coatom ordering is indeed a recursive one we are going to use the product structure of Proposition 3.3. It is however formulated in terms of $\mathcal{CC}_n$'s rather than $\mathcal{CC}_n^+$ so we start by formulating the condition of having a recursive coatom orderability in a way that uses a poset P (possibly) without a least element (and which is of course should be equivalent to the usual condition for P^+, the poset with a least element added). (We are otherwise using the notation and formulation of [BVSWZ, Def. 4.7.17] which appeared somewhat more convenient than the original one of [BW83, Def. 3.1].) Thus we assume that P^+ is graded and put, for $x \in P$, $[x] := \{y \in P \mid y \leqslant x\}$. A *recursive coatom ordering* of P consists of a total ordering $<^c$ of its coatoms and either $\mathrm{length}(P^+) \leqslant 2$ or $\mathrm{length}(P^+) > 2$ and for any coatom x there is a distinguished subset Q_x with $Q_x \neq \emptyset$ when x is not $<^c$-first such that

1. $[x] \cap (\cup_{x'<x}[x']) = \cup_{y \in Q_x}[y]$ and
2. $[x]$ has a recursive coatom ordering in which the elements of Q_x come first.

(The condition of non-emptiness of Q_j takes care of covering $\hat{0}$ of P^+.)

The fact we recursively are dealing with products of smaller posets will be used together with the fact that the condition of CL-shellability is inherited by products (noted in [BW83] with reference to [Bj08, Thm. 4.3]) with two modifications. First, as $(P \times Q)^+ \neq P^+ \times Q^+$ we are not talking about exactly the same product. Second, we are going to use this result recursively, i.e., apply it to the $[x_j]$, with the above notation, which forces us to verify an external condition on the recursive coatom ordering. Hence we need to be explicit on how the recursive coatom ordering on a product is obtained from orderings on the factors and during our arguments we shall in fact need to deal with several ways of obtaining such an ordering. (The fact that products of posets with recursive coatom orderings have some recursive coatom ordering is of course well-known, cf., [BW83, Comment before Thm. 8.3].) The following lemma takes care of the different versions that we shall need.

Lemma 4.1 *Let P and P' be posets with P^+ and P'^+ graded. Assume given recursive coatom orderings of P and P' together with initial sequences of them. Order the coatoms of $P \times P'$ as follows: Start with $(x, \hat{1})$ where x runs over the initial segment for P,in the order given, then take $(\hat{1}, x')$, where x' runs over the initial segment for P', then take $(x, \hat{1})$ with x running over the rest of the coatoms for P and finally end with $(\hat{1}, x')$ where x' runs over the rest of the coatoms of P'. This coatom ordering is a recursive coatom ordering.*

PROOF: We have to verify the recursive conditions and let us start with the case of the coatom being of the form $(x, \hat{1})$ with x in the initial segment of the coatom ordering of P. We then choose a recursive coatom ordering of $[x]$ and an initial segment Q_x for which $[x] \cap \cup_{z<x}[z] = \cup_{y \in Q_x}[y]$. We then order the coatoms of $[(x, \hat{1})] = [x] \times P'$ by putting the coatoms $(y, \hat{1})$ first and in the order given and the coatoms (x, y) last in the given recursive coatom ordering of P'. If we let the Q for $(x, \hat{1})$ be $\{(y, \hat{1}) \mid y \in Q_x\}$ we easily see that the conditions needed are fulfilled. The case of a coatom of the form $(\hat{1}, x')$ with x' in the initial segment of P' is completely analogous. Consider next the case when the coatom is of the form $(x, \hat{1})$ with x *not* in the initial segment. Choose as before a recursive coatom ordering of $[x]$ with a Q_x as above. We define a coatom ordering on $[(x, \hat{1})] = [x] \times P'$ by taking the $(Q_x, \hat{1})$ first, then the given initial segment of P' times x, then the rest of the coatoms of P (times $\hat{1} \in P'$) and last the rest of the coatoms of P' (times x). By induction this gives a recursive coatom ordering of $[(x, \hat{1})]$ and letting the Q for this element be the union of $(Q_x, \hat{1})$ and the initial segment of P' the necessary conditions are fulfilled. The case of coatoms of the form $(\hat{1}, x')$ with x' not in the initial segment of P' is analogous. □

With this lemma we are now ready to prove dual CL-shellability after having introduced one more notion. For a cover $(A \to B) \prec (C \to D)$ we define its *residual* as follows. If the cover is a small cap addition the residual is the small cap added, if it is a cap collapse it is the cap that is collapsed and if it is an edge removal the residual is the removed edge.

Theorem 4.2 *$\mathcal{CC}_n^+$ is a thin dual CL-shellable lattice.*

PROOF: To begin let us record that we shall use without further mention that as $\mathcal{CC}_n^+$ and $\mathcal{K}_n^+$ are lattices, the intersection $[x] \cap [y]$ in $\mathcal{CC}_n$ resp. $\mathcal{K}_n$ is, when non-empty, equal to $[x \wedge y]$.

In order to verify the recursive conditions of a (proposed) recursive coatom ordering we shall use Proposition 3.3. This forces us to first discuss recursive coatom orderings on $\mathcal{K}_n$. We can define a partial pseudo-order, the *chord order*, on the chords (i.e., non-boundary line segments connecting two vertices) of P_{n+1} by saying that $\ell_1 \leqslant^c \ell_2$ if they intersect in at most vertices of P_{n+1} and if the interior of ℓ_1 lies in the component of $\overline{P_{n+1}} \setminus \ell_2$ that contains the final edge. An extension of the chord order to a total order on the chords will be a called an *admissible order*. Each chord ℓ defines a coatom, x_ℓ, of $\mathcal{K}_n$ given by the union of ℓ and P_{n+1} and all coatoms are obtained in that way. Hence an admissible order gives a total order on the coatoms of P_{n+1} and the contention is that any admissible order gives a recursive coatom order. A chord ℓ of P_{n+1} decomposes P_{n+1} into two polygons, one, P', which has the final edge of P_{n+1} as its final edge and one, P'', which has ℓ as its final edge. It is clear that the chord orders on P' and P'' are the orders induced from the chord order on P_{n+1} and in particular an admissible order on the chords of P_{n+1} induces an admissible order on the chords of P' and P''. Hence by induction we may assume that these latter orders give recursive coatom orderings. Furthermore, the interval $[x_\ell]$ is isomorphic to the product $\mathcal{K}(P') \times \mathcal{K}(P'')$ and thus has a recursive coatom ordering by Lemma 4.1 and induction. Now, every coatom of $[x_\ell]$ whose second component in the product decomposition is $\hat{1}$ is of the form $x_{\ell'} \wedge x_\ell$, in the lattice structure of $\mathcal{K}_n$, for some $\ell' \leqslant^c \ell$. This means that we may choose the Q for x_ℓ to be equal to the coatoms of the form $x_{\ell'} \wedge x_\ell$ for ℓ' smaller than ℓ in the admissible order.

Turning to $\mathcal{CC}_n$ we shall use the same technique as for $\mathcal{K}_n$ using Proposition 3.3. This means that to begin with we choose once and for all an admissible order on the chords of P_n which induces a recursive coatom ordering on each $\mathcal{K}$-subposet arising from a subpolygon of P_n. We are going to use the following convention to describe the coatoms and more generally elements of an interval $[x]$. By Proposition 3.3 this interval is the product of the $\mathcal{CC}$'s of the caps of x and the $\mathcal{K}$ of its collapsee. Hence, each coatom corresponds to a coatom of either the $\mathcal{CC}$ of a cap or the $\mathcal{K}$ of the collapsee and we shall say that the coatom *belongs* to that cap or collapsee. General elements correspond to sequences of elements in the $\mathcal{CC}$'s of the caps and one element of the $\mathcal{K}$ of the collapsee and the components of that sequence will be called the components of the elements *in* the respective caps or collapsee. We now have to choose an ordering of the coatoms of $\mathcal{CC}_n$ in a uniform enough way so that it applies also to all subpolygons but we also have to tell in which order we are going to write the factors of an interval as it is needed to get a coatom ordering on the product. Starting with the coatom ordering we have by Proposition 3.4 two types of coatoms, the F_i, which we shall say are of *small cap type*, for $1 < i < n$ and F'_S, of *cap type*, for $\{1, n\} \subset S \subseteq \{1, \ldots, n\}$. We now put the F_i first, ordered by the reverse of the natural order of the i (even though as we shall see the order among the F_i doesn't matter). After that we take the F'_S ordered as follows: We put S before T if there is an i such that $j > i \Rightarrow (j \in S \iff j \in T)$ and $i \in T$ but $i \notin S$. We shall call this order the *lex-order*.

Next we need to decide in which order to put the factors when representing an interval $[x]$. We do this by going through the caps of x starting with the one containing the largest vertices, in their polygon order, and continuing with the next largest and so on. Finally, for an arbitrary $x \in \mathcal{CC}_n$ we order its coatoms as follows: We first take all the coatoms of small cap type of the caps of x ordered within each cap as above and between caps by the order of the caps that has just been given. We then take all coatoms of cap type of the caps ordered within caps as per above and between caps by the cap order. Last, we take all the coatoms corresponding to the $\mathcal{K}$ of the collapsee of x ordered by the given admissible total order of chords of P_n. Combining Lemma 4.1 and Proposition 3.3, the result we have already obtained about recursive coatom orderings of $\mathcal{K}$'s and using induction we may assume that these coatom orderings give recursive coatom orderings on all $[x]$ but $x = \hat{1}$.

It thus remains to prove that this ordering is a recursive coatom ordering of $\mathcal{CC}_n$ and for that we need to identify the Q_x and show that they fulfil the required conditions. We start with the case when the coatom x is equal to F_i. The interval $[F_i]$ is just $\mathcal{CC}_{n-1} = \mathcal{CC}((F_i)_r)$. Furthermore, the coatoms of $[F_i]$ of small cap type are exactly the compound collapses of the form $F_i \wedge F_j$ for $i \neq j$ and the order on the F_j induces the coatom ordering we have recursively defined on $[F_i]$. Hence we can let the Q for F_i consist of its coatoms of small cap type. This verifies the recursive condition for the case when x is of small cap type.

Assume now instead that we are dealing with a coatom F'_S of cap type. We let its Q consist of all the coatoms of its caps. This is an initial segment of its coatoms and it remains to verify the first part of the recursive coatom ordering conditions.

It is clear that $[F_i] \cap [F'_S]$ is empty unless $i \notin S$ and is otherwise the coatom of $[F'_S]$ belonging to the cap of F'_S which contains i and there it is the small cap type coatom whose collapser does not contain i. This means that the interval of any small cap type coatom belonging to the caps of F'_S are of the form $[F_i] \cap [F'_S]$ and thus all intersections $[F_i] \cap [F'_S]$ are contained in intersections which are coatoms in $[F_i] \cap [F'_S]$.

Assume now that T comes before S in the lex-order. Hence there is an i such that $j > i \Rightarrow (j \in S \iff j \in T)$ and $i \in T$ but $i \notin S$. Let γ be the cap of F'_S containing i. In the product decomposition of $[F'_S]$ we have that the component of $F'_T \wedge F'_S$ in a cap that comes before γ is $\hat{1}$ while the component in γ is *not* $\hat{1}$ but rather a coatom c of cap type. There then is a T' such that $F'_{T'} \wedge F'_S$ is a coatom belonging to γ and its component in γ is c. (In fact T' is given by the conditions that $S \subset T'$ and $T' \setminus S$ are the vertices of c not in S.) In particular we have that $F'_T \wedge F'_S$ is below $F'_{T'} \wedge F'_S$ in the partial order of $\mathcal{CC}_n$. As also every coatom of F'_S belonging to a cap where it is of cap type is of the form $F'_{T'} \wedge F'_S$ (with T' constructed as above) we have verified the required conditions for a recursive coatom ordering.

For thinness let us assume that $x \prec y \prec z$.

If $x \prec y$ and $y \prec z$ have different forms then it is clear that there is a $x \prec y' \prec z$ such that the residual of $x \prec y'$ is equal to that of $y \prec z$ and the residual for $y' \prec z$ to that of $x \prec y$ and that y and y' are the only elements in the interval (x, z). If $x \prec y$ and $y \prec z$ both are edge removals, then x and z differ by two internal edges and they can be removed in any order giving again two elements in (x, z). For the case when both are cap collapses, then the caps can be collapsed in any order unless the final edge of one of them is an initial edge of the other. In that case we get an $x \prec y' \prec z$ by letting $x \prec y'$ be the edge removal where the edge removed is the common final edge of one cap and initial edge of the other and $y' \prec z$ is the cap collapse collapsing the union of the two caps (see Fig. 3).

The remaining case is that both covers are small cap additions. Again if the two small cap residuals do not have an edge in common they can be added in any order. If they do have a common edge their union is a quadrangle with the common edge as a diagonal. One may then make two small cap additions by using the other diagonal (see Fig. 3). □

We have an immediate corollary.

Corollary 4.3 *$\mathcal{CC}_n$ is the face poset of a regular PL-cell decomposition of a ball.*

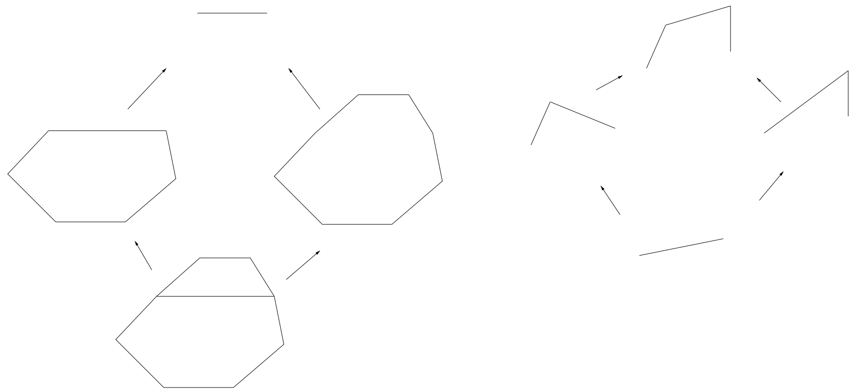

Figure 3: Interfering cap collapses and small cap additions.

Proof: This follows from [Bj84]. □

References

[BVSWZ] A. Björner, M. Las Vergnas, B. Sturmfels, N. White, and G. M. Ziegler, *Oriented matroids*, second ed., Encyclopedia of Mathematics and its Applications, vol. 46, Cambridge University Press, Cambridge, 1999.

[BW83] A. Björner and M. Wachs, *On lexicographically shellable posets*, Trans. Amer. Math. Soc. **277** (1983), no. 1, 323–341.

[Bj08] A. Björner, *Shellable and Cohen-Macaulay partially ordered sets*, Trans. Amer. Math. Soc. **260** (1980), no. 1, 159–183.

[Bj84] A. Björner, *Posets, regular CW complexes and Bruhat order*, European J. Combin. **5** (1984), no. 1, 7–16.

[St63a] J. D. Stasheff, *Homotopy associativity of H-spaces. I*, Trans. Amer. Math. Soc. **108** (1963), 275–292.