Available online at www.sciencedirect.com

ScienceDirect

Advances in Applied Mathematics 41 (2008) 158–176

ADVANCES IN Applied Mathematics

www.elsevier.com/locate/yaama

A regular decomposition of the edge-product space of phylogenetic trees

Jonna Gill [a,1], Svante Linusson [b,*,2], Vincent Moulton [c], Mike Steel [d]

[a] *Matematiska Institutionen, Linköpings Universitet, 581 83 Linköping, Sweden*
[b] *Matematiska Institutionen, KTH, SE-100 44 Stockholm, Sweden*
[c] *School of Computing Sciences, University of East Anglia, Norwich, NR4 7TJ, UK*
[d] *Biomathematics Research Centre, Department of Mathematics and Statistics, University of Canterbury, Christchurch, New Zealand*

Received 14 August 2005; accepted 11 July 2006

Available online 19 November 2007

Abstract

We investigate the topology and combinatorics of a topological space called the edge-product space that is generated by the set of edge-weighted finite labelled trees. This space arises by multiplying the weights of edges on paths in trees, and is closely connected to tree-indexed Markov processes in molecular evolutionary biology. In particular, by considering combinatorial properties of the Tuffley poset of labelled forests, we show that the edge-product space has a regular cell decomposition with face poset equal to the Tuffley poset.

MSC: 05C05; 92D15

Keywords: Trees; Forests; Partitions; Poset; Regular cell complex; Recursive coatom ordering

* Corresponding author.
E-mail addresses: jogil@mai.liu.se (J. Gill), linusson@math.kth.se (S. Linusson), vincent.moulton@cmp.uea.ac.uk (V. Moulton), m.steel@math.canterbury.ac.nz (M. Steel).
[1] Supported by the Swedish Science Council.
[2] Partially supported by EC's IHRP program through grant HPRN-CT-2001-00272.

0196-8858/$ – see front matter
doi:10.1016/j.aam.2006.07.007

1. Introduction

For a tree T, we let $V(T)$ and $E(T)$ denote the sets of vertices and edges of T, respectively. For a fixed finite set X we are interested in the (finite) set of binary (i.e. trivalent) trees T that have X as their set of leaves (degree one vertices). They correspond to all the possible "phylogenetic trees" for X. See [12]. We study the space of all such trees with edge weights between 0 and 1. To each binary tree with edge weights we associate a point in a high-dimensional space in the following way. Given a map $\lambda : E(T) \to [0, 1]$ define

$$p = p_{(T,\lambda)} : \binom{X}{2} \to [0, 1]$$

by setting, for all $x, y \in X$,

$$p(x, y) = \prod_{e \in P(T;x,y)} \lambda(e),$$

where $P(T; x, y)$ is the set of edges in the path in T from x to y.

Let $\mathcal{E}(X, T) \subset [0, 1]^{\binom{X}{2}}$ denote the image of the map

$$\Lambda_T : [0, 1]^{E(T)} \to [0, 1]^{\binom{X}{2}}, \quad \lambda \mapsto p_{(T,\lambda)}$$

and let $\mathcal{E}(X)$ be the union of the subspaces $\mathcal{E}(X, T)$ of $[0, 1]^{\binom{X}{2}}$ over all binary trees T with X as its set of leaves. We call $\mathcal{E}(X)$ the *edge-product space for trees on X*. Note that a different notion of the space of trees was suggested by Billera, Holmes and Vogtmann in [2]. The main difference is that they add edge weights along paths to give a distance between leaves, rather than multiplying them. The edge-product space can be seen as a compactification of the space studied in [2], since we allow zero as an edge weight.

Apart from their intrinsic interest, a central motivation for investigating edge-product spaces is that they are intimately connected with tree-indexed Markov process in molecular evolutionary biology [5,8,12], as we now briefly outline. In these models there is a fixed matrix Q of transition rates between states of some set (e.g. nucleotide bases, amino acids), which forms a stationary and time-reversible Markov process. The process operates for some duration $d(e)$ on each edge e of T. Let $\lambda : E(T) \to [0, 1]$ be defined by $\lambda(e) = e^{-d(e)}$, and allow $\lambda(e)$ to equal 0 in order to model 'site saturation' (i.e. the limiting value as $d(e) \to \infty$). The Markov process, parameterised by the pair (T, λ), induces a (marginal) joint probability distribution on the set of state assignments to X. Furthermore it can be shown that two pairs (T, λ) and (T', λ') induce the same joint probability distribution precisely if $p_{(T,\lambda)} = p_{(T',\lambda')}$ (by extending the approach of [13] which established this result when Q is a symmetric 2×2 matrix). Consequently, the edge-product space defined above is homeomorphic to the quotient space where trees with λ-valued edge weights are identified if they induce the same Markov process at the leaves for a fixed rate matrix Q.

In [9] it was shown that $\mathcal{E}(X)$ has a natural CW-complex structure for any finite set X, and a combinatorial description of the associated face poset, called the Tuffley poset was given. It was also conjectured that $\mathcal{E}(X)$ is a regular cell complex. Here we prove that this conjecture holds.

Theorem 1.1. *The edge-product space $\mathcal{E}(X)$ has a regular cell decomposition with face poset given by the Tuffley poset.*

In particular, $\mathcal{E}(X)$ is homeomorphic to the geometric realization of the Tuffley poset (see [3, 12.4(ii)]). Note that as a consequence of this theorem we obtain an affirmative answer to the main conjecture in [1].

Our main technical tool is a proof that any interval of the Tuffley poset (with an artificial $\hat{0}$ added) has a recursive coatom ordering. This implies amongst other things that the order complex of every open interval is a sphere. Note that we do not describe an explicit recursive coatom ordering that is valid for any given interval. Instead, we have developed what we believe to be a new method for establishing the existence of recursive coatom orderings. In particular, we define a class of coatom orderings and show that for any interval we can always choose some ordering from that class that satisfies the conditions for a recursive coatom ordering.

We now describe the contents of the paper. In Section 2 we review some properties of "X-forests" and of the Tuffley poset, which can be regarded as an order relation on the collection of all X-forests. In Section 3 we prove that there exists a shelling order for the chains in every interval of the Tuffley poset which we require in order to prove Theorem 1.1 in Section 4. Finally, in Section 5 we present the proofs of some technicalities that we used in Section 3. As many of the cases are straightforward checks, where appropriate we will refer the reader to [6] where the full details are presented.

2. Trees, forests and the Tuffley poset

In this section we review some material concerning trees and related structures. Throughout this paper X will be a finite set.

An *X-tree* $\mathcal{T}$ is a pair $(T;\phi)$ where T is a tree, and $\phi: X \to V(T)$ is a map such that all vertices in $V(T)-\phi(X)$ have degree greater than two. Note that we do not require ϕ to be injective. We call X *labels* and the vertices in $V(T)-\phi(X)$ *unlabelled*. Two X-trees $(T_1;\phi_1)$ and $(T_2;\phi_2)$ are isomorphic if there is a graph isomorphism $\psi: V(T_1) \to V(T_2)$ such that $\phi_2 = \psi \circ \phi_1$. For an X-tree $\mathcal{T}=(T;\phi)$ we let $E(\mathcal{T})$ denote $E(T)$, the set of edges of T. By a *contraction* of an edge $e=\{v,u\}$ in an X-tree $\mathcal{T}=(T;\phi)$, we mean contraction of the edge in T, with identification of u and v and label $\phi^{-1}(u)\cup\phi^{-1}(v)$ for the new vertex.

An *X-forest* is a collection $\alpha = \{(A, \mathcal{T}_A)\colon\ A \in \pi\}$ where

(i) π forms a set partition of X, and
(ii) $\mathcal{T}_A$ is an A-tree for each $A \in \pi$.

We let $\mathcal{S}(X)$ denote the set of X-forests. We order the elements of $\mathcal{S}(X)$ by letting $\beta \leqslant \alpha$ if the trees in β can be obtained from the trees in α by contracting certain edges, and deleting certain other edges, with any resulting unlabelled vertices of degree 2 being suppressed.

The poset $\mathcal{S}(X)$ was first defined (slightly differently) by Christopher Tuffley [13], and it is thus called the *Tuffley poset* on X. In Fig. 1 we show the Hasse diagram of $\mathcal{S}(\{1,2,3\})$.

Given an X-forest α and an edge $e \in E(\alpha)$, we denote the X-forest obtained by contracting e with $e^c(\alpha)$, and the X-forest obtained by deleting e and suppressing any resulting vertices of degree 2 with $e^d(\alpha)$. From now on we will with an *edge deletion* always include the suppression

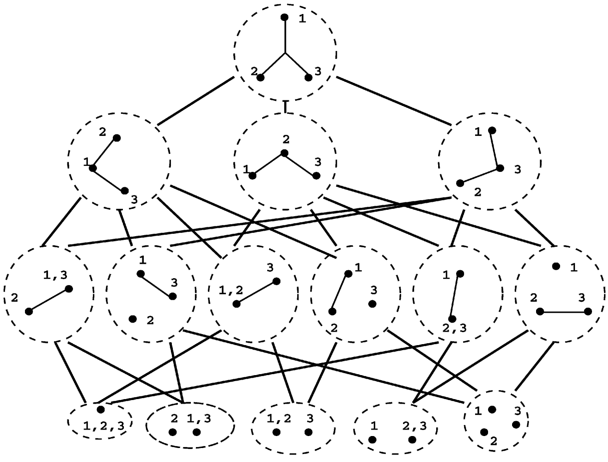

Fig. 1. The Tuffley poset $\mathcal{S}(X)$ for $X = \{1, 2, 3\}$.

of any degree two vertices. When α is clear from the context we will simply write e^c and e^d. Furthermore,

$$\bigl|E\bigl(e^c(\alpha)\bigr)\bigr| = \bigl|E(\alpha)\bigr| - 1, \tag{1}$$

and

$$\bigl|E(\alpha)\bigr| - 3 \leqslant \bigl|E\bigl(e^d(\alpha)\bigr)\bigr| \leqslant \bigl|E(\alpha)\bigr| - 1. \tag{2}$$

We will say that the edge deletion $\alpha \mapsto e^d(\alpha)$ is *safe* if $|E(e^d(\alpha))| = |E(\alpha)| - 1$. We say that a vertex in an X-tree is *unsupported* if it is unlabelled and of degree 3. We can easily conclude that for an X-forest α, an edge deletion $\alpha \mapsto e^d(\alpha)$ is safe if and only if neither endpoint of the edge e in α is unsupported.

We define an *elementary operation* on an element of $\mathcal{S}(X)$ to be either an edge contraction, or a safe edge deletion. The covering relation in the poset will be denoted by the symbol $\lessdot$.

The following result which is a restatement of Theorem 4.2 of [9] describes $\mathcal{S}(X)$ in terms of these operations, and establishes some further structural properties of the Tuffley poset.

Theorem 2.1. *Suppose that X is a finite set and $\alpha, \beta \in \mathcal{S}(X)$. Then the following statements hold.*

(i) *$\beta \leqslant \alpha$ if and only if β can be obtained from α by any sequence of contraction and deletion operations, in which case we can insist that all contractions occur first, and that all the subsequent deletions are safe.*
(ii) *$\beta \lessdot \alpha$ if and only if β can be obtained from α by one elementary operation.*
(iii) *$\mathcal{S}(X)$ is a pure poset, and for an element $\alpha = \{(A, \mathcal{T}_A)\colon A \in \pi\}$ of $\mathcal{S}(X)$ its rank, denoted $\rho(\alpha)$, is given by*

$$\rho(\alpha) = \bigl|E(\alpha)\bigr|.$$

(iv) $\mathcal{S}(X) \cup \{\hat{0}\}$ *is thin* (*that is, all intervals of length* 2 *contain exactly four elements*).
(v) *The maximal elements of* $\mathcal{S}(X)$ *are precisely the binary* X*-trees with* $|X|$ *leaves. They have rank* $2|X| - 3$.
(vi) *The minimal elements of* $\mathcal{S}(X)$ *are precisely the* X*-forests with no edges. Hence, there is a minimal element for each set partition of* X.
(vii) *Suppose* α *is an* X*-forest, and that* α *has an interior vertex* v *labelled by* m. *Construct an* X'*-forest* β *by removing* v *from* α *and giving edge number* i *incident with* v, $1 \leqslant i \leqslant \deg(v)$, *a new vertex* v_i *that is labelled by* m_i, $m_i \notin X$ *and* $m_i \neq m_j$ *if* $i \neq j$. *Then* $[\hat{0}, \alpha]$ *is isomorphic to* $[\hat{0}, \beta]$.

3. Recursive coatom orderings

In this section we establish the following theorem.

Theorem 3.1. *There is a recursive coatom ordering for each interval* $[\hat{0}, \Gamma] \subset \mathcal{S}(X) \cup \{\hat{0}\}$. *In particular, every such interval is shellable.*

The proof of Theorem 3.1 is provided in Section 3.3.

Corollary 3.2. *The order complex of any interval* (α, β) *in* $\mathcal{S}(X) \cup \{\hat{0}\}$ *is homeomorphic to a sphere of dimension* $\rho(\beta) - \rho(\alpha)$. *In particular, the Möbius function* μ *of the Tuffley poset is given by* $\mu(\alpha, \beta) = (-1)^{\rho(\beta) - \rho(\alpha)}$.

Proof. We know that intervals in $\mathcal{S}(X) \cup \{\hat{0}\}$ are thin (see Theorem 2.1(iv)) and admit a recursive coatom ordering. The result now follows from [4, Theorem 4.7.24(i)] (see p. 169). The Möbius function is equal to the reduced Euler characteristic of the order complex of the open interval (α, β) (see [3, 9.14]). □

3.1. Preliminaries and definitions

Definition 3.3. A *recursive coatom ordering* for an interval $[\hat{0}, \Gamma]$ is an ordering $\alpha_1, \ldots, \alpha_t$ of its coatoms that satisfies the following two conditions:

(V1) For all $i < j$ and $\gamma < \alpha_i, \alpha_j$ there is a $k < j$ and an element β such that $\beta \lessdot \alpha_k, \alpha_j$ and $\gamma \leqslant \beta$.
(V2) For all $j = 1, \ldots, t$, $[\hat{0}, \alpha_j]$ admits a recursive coatom ordering in which the coatoms that come first in the ordering are those that are covered by some α_k where $k < j$.

A recursive coatom ordering of an interval has several implications, see e.g. [3].

The following notation is used. The edges incident with a vertex v in an X-forest will be denoted $e_1, \ldots, e_n$ (or $f_1, \ldots, f_n$). If the other vertex incident with an edge e_i has degree 3, the two other edges incident with that vertex will be denoted e_{i1} and e_{i2}, else e_{i1} and e_{i2} are not defined. See Fig. 2. Remember that the coatom obtained from an X-forest Γ by contraction of the edge e_i is $e_i^c(\Gamma)$, and the coatom obtained by safe deletion of the edge e_i is $e_i^d(\Gamma)$.

Convention 3.4. $e_1, \ldots, e_n$ (or $f_1, \ldots, f_n$) are not in general fixed labels for the edges incident with a vertex v. This convention is made to avoid having to write $e_{i_1}, \ldots, e_{i_n}$ where $i_1, \ldots, i_n$ is

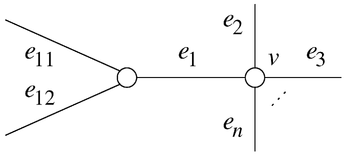

Fig. 2. Edge notation at the vertex v.

a permutation of $\{1,\dots,n\}$. If for example some condition uses e_1 and e_2, it will be true for all e_{i_1} and e_{i_2} with the same properties. The same applies to numbered components of an X-forest. This concerns all sections in this text.

Convention 3.5. By Theorem 2.1(vii) we may replace any X-forest α with labels on some internal vertices by a forest β of trees without labels on internal vertices and lower interval $[\hat{0},\beta]$ isomorphic to $[\hat{0},\alpha]$. We may therefore throughout the proof of the existence of a recursive coatom ordering assume that no internal vertices are labelled. We may also during the proof without loss of generality assume that each leaf has exactly one label.

To prove that there exists a recursive coatom ordering $\alpha_1,\dots,\alpha_t$ for $[\hat{0},\Gamma]$, a new condition for coatom orderings called (V3) is used. The idea is to replace the conditions for a recursive coatom ordering with a stronger but easier condition. In Theorem 3.16 it will be proven that all coatom orderings satisfying the condition (V3) are recursive coatom orderings. The following two definitions are made to simplify condition (V3) that will be given in Definition 3.8, and to facilitate dealing with (V3) later.

Definition 3.6. Let Γ be an X-forest, and v an interior vertex of Γ incident with the edges $e_1,\dots,e_n$. Let v_i be the other vertex incident with e_i, $1\leqslant i\leqslant n$. To simplify notation, we use the following symbols for coatoms of $[\hat{0},\Gamma]$:

$$\dot{e}_i^d(\Gamma)=\begin{cases} e_i^d(\Gamma) & \text{if } \deg(v_i)\neq 3,\ \deg(v)\geqslant 4,\\ e_{i1}^c(\Gamma) \text{ or } e_{i2}^c(\Gamma) & \text{if } \deg(v_i)=3,\end{cases}$$

$$\ddot{e}_i^d(\Gamma)=\begin{cases} e_i^d(\Gamma) & \text{if } \deg(v_i)\neq 3,\ \deg(v)\geqslant 4,\\ \{e_{i1}^c(\Gamma), e_{i2}^c(\Gamma)\} & \text{if } \deg(v_i)=3.\end{cases}$$

Note that $\dot{e}_i^d(\Gamma)$ and $\ddot{e}_i^d(\Gamma)$ are not defined if $\deg(v_i)\neq 3$, $\deg(v)=3$. When no confusion arises, Γ is often omitted.

Definition 3.7. Let $\mathcal{C}$ and $\mathcal{D}$ be disjoint sets of coatoms, and suppose there is a given coatom ordering. If all elements in $\mathcal{C}$ come before every element in $\mathcal{D}$ in the ordering, then $\mathcal{C}\lhd\mathcal{D}$. We will use $\mathcal{C}\ntriangleleft\mathcal{D}$ to denote the condition that at least one element in $\mathcal{C}$ comes after some element in $\mathcal{D}$.

Definition 3.8. Take an X-forest Γ, consisting of the non-trivial components $K_1,\dots,K_m$ (i.e. with at least one edge each). If a coatom ordering of $[\hat{0},\Gamma]$ satisfies the following conditions, then it is said to satisfy the condition (V3) (recall Convention 3.4):

(V3) (a) $\{\gamma \mid \gamma \lessdot \Gamma,\ \gamma = e^c \text{ or } \gamma = e^d \text{ where } e \in \bigcup_{i=1}^{\ell} K_i\} \ntriangleleft \{\gamma \mid \gamma \lessdot \Gamma,\ \gamma = e^c \text{ or } \gamma = e^d \text{ where } e \in \bigcup_{i=\ell+1}^{m} K_i\}$ for all $1 \leqslant \ell \leqslant m-1$.

(b) If v is an interior vertex in Γ, $n = \deg(v) = 3$ and $\ddot{e}_3^d$ is defined then $\{e_1^c, e_2^c\} \ntriangleleft \{\ddot{e}_3^d\}$ and $\{\ddot{e}_3^d\} \ntriangleleft \{e_1^c, e_2^c\}$.

(c) If v is an interior vertex in Γ and $n = \deg(v) \geqslant 4$ then $\{e_1^c, e_2^c\} \ntriangleleft \{\ddot{e}_3^d, \ddot{e}_4^d, \ldots, \ddot{e}_n^d\}$ and $\{\ddot{e}_3^d, \ddot{e}_4^d, \ldots, \ddot{e}_n^d\} \ntriangleleft \{e_1^c, e_2^c\}$.

(d) If v is an interior vertex in Γ and $n = \deg(v) \geqslant 4$ then $\{e_1^c, \ddot{e}_1^d, \ldots, e_k^c, \ddot{e}_k^d\} \ntriangleleft \{e_{k+1}^c, \ddot{e}_{k+1}^d, \ldots, e_n^c, \ddot{e}_n^d\}$, where $1 \leqslant k \leqslant n-1$.

The sub-conditions of (V3) are symmetric, hence the reversal of a coatom ordering satisfying (V3) also satisfies (V3).

Definition 3.9. Let Γ be an X-forest and v a vertex in Γ. The coatom α is said to be *near* v if α is obtained by (safe) deletion or contraction of an edge e_i incident with v or contraction of e_{i1} or e_{i2}.

The above definition is made since the sub-conditions (V3)(b), (V3)(c), and (V3)(d) only deal with the coatoms near an interior vertex v in an X-forest.

3.2. Outline of proof of Theorem 3.1

To prove Theorem 3.1 the following method is used. A class of recursive coatom orderings is created. This class has the property that if $\alpha_1, \ldots, \alpha_j, \ldots, \alpha_t$ is a coatom ordering in the class, then for each $1 \leqslant j \leqslant t$ the class contains a coatom ordering for $[\hat{0}, \alpha_j]$ in which the coatoms that come first in the ordering are those that are covered by some α_k where $k < j$. In particular, the class is defined to be all coatom orderings satisfying the condition (V3).

That property (V1) of Definition 3.3 follows from (V3) is shown in Lemma 3.10. The property which implies (V2) of Definition 3.3 is shown in Lemma 3.12 and Theorem 3.15 with the help of Lemmas 3.13 and 3.14. This is done by first finding certain orderings of the coatoms near each interior vertex in α_j, and then combining them to a coatom ordering for $[\hat{0}, \alpha_j]$.

The results are put together in Theorem 3.16 to show that a coatom ordering satisfying (V3) is a recursive coatom ordering. Finally Theorem 3.1 follows from Theorem 3.15 which implies that there is always a coatom ordering of $[\hat{0}, \Gamma]$ satisfying (V3), and Theorem 3.16.

Since the proofs of Lemmas 3.10–3.14 are very technical, they will be presented separately in Section 5.

3.3. There is a recursive coatom ordering for $[\hat{0}, \Gamma]$

The following lemma is obviously necessary, and will be proven in Section 5.1.

Lemma 3.10. *Let Γ be an X-forest, and let $\alpha_1, \ldots, \alpha_t$ be the coatoms of* $[\hat{0}, \Gamma]$. *If $\alpha_1, \ldots, \alpha_t$ satisfies* (V3)*, then it also satisfies part* (V1) *of Definition* 3.3.

To prove that a coatom ordering satisfying (V3) also satisfies part (V2) of Definition 3.3, Lemma 3.12 and Theorem 3.15 are needed.

The following definition is useful since $\mathcal{A} \diamond \mathcal{B}$ is a necessary condition for the possibility of ordering $\mathcal{A} \cup \mathcal{B}$ so that the ordering satisfies $\mathcal{A} \triangleleft \mathcal{B}$ and (V3).

Definition 3.11. Let Γ be an X-forest, and let $\mathcal{A}$ and $\mathcal{B}$ be disjoint sets of coatoms of $[\hat{0}, \Gamma]$. Then $\mathcal{A}$ and $\mathcal{B}$ are said to be *compatible* with the condition (V3) if $\mathcal{A} \lhd \mathcal{B}$ is not forbidden by any single sub-condition of (V3). If $\mathcal{A}$ and $\mathcal{B}$ are compatible with (V3), we write $\mathcal{A} \diamond \mathcal{B}$. If $\mathcal{A}$ and $\mathcal{B}$ are not compatible with (V3), we write $\mathcal{A} \not\diamond \mathcal{B}$. Since the condition (V3) is symmetric, $\mathcal{A} \diamond \mathcal{B} \Leftrightarrow \mathcal{B} \diamond \mathcal{A}$.

Lemma 3.12. *Let $\alpha_1, \dots, \alpha_j, \dots, \alpha_t$ be a coatom ordering of $[\hat{0}, \Gamma]$ satisfying (V3). Fix j, and consider the interval $[\hat{0}, \alpha_j]$. Let $\mathcal{A} = \{\gamma \mid \gamma \lessdot \alpha_k, \alpha_j$ and $k < j\}$ and let $\mathcal{B} = \{\gamma \mid \gamma \lessdot \alpha_k, \alpha_j$ and $k > j\}$. (The sets $\mathcal{A}$ and $\mathcal{B}$ are disjoint since $\mathcal{S}(X) \cup \{\hat{0}\}$ is thin.) Then $\mathcal{A} \diamond \mathcal{B}$.*

The above lemma is shown in Section 5.3. To prove Theorem 3.15 the following method is used. For each interior vertex a coatom ordering for the coatoms near that vertex is found, an ordering that satisfies $\mathcal{A} \lhd \mathcal{B}$ and (V3). Then these coatom orderings are combined to a coatom ordering of all coatoms. Hence these orderings must agree on the order of coatoms that are near more than one vertex, which we will now make sure by the following investigation and Lemmas 3.13, 3.14.

Let v_1 and v_2 be adjacent interior vertices, and denote the edge between them e_1. Let the edges incident with v_1 be denoted $e_1, \dots, e_n$. The coatoms near v_1 which are also near v_2 are the following:

$$\begin{array}{ll} e_1^c, e_2^c, \text{ and } e_3^c & \text{if } \deg(v_1) = 3,\ \deg(v_2) \geqslant 4, \\ e_1^c, e_2^c, e_3^c, e_{11}^c, \text{ and } e_{12}^c & \text{if } \deg(v_1) = 3,\ \deg(v_2) = 3, \\ e_1^c \text{ and } e_1^d & \text{if } \deg(v_1) \geqslant 4,\ \deg(v_2) \geqslant 4, \\ e_1^c, e_{11}^c, \text{ and } e_{12}^c & \text{if } \deg(v_1) \geqslant 4,\ \deg(v_2) = 3. \end{array}$$

If v_2 is adjacent to an interior vertex v_3 then some coatoms can be near both v_1 and v_3, but in that case these coatoms are near both v_1 and v_2 too.

Lemma 3.13. *Let α be an X-forest, and let v be an interior vertex of α with degree 3. Suppose the coatoms near v are partitioned into two sets $\mathcal{C}$ and $\mathcal{D}$ where $\mathcal{C} \diamond \mathcal{D}$. Then there is always an ordering of $\mathcal{C} \cup \mathcal{D}$ that satisfies (V3) and $\mathcal{C} \lhd \mathcal{D}$, and has a given order of e_1^c, e_2^c, e_3^c, and e_{11}^c, e_{12}^c (if they are defined) that is not forbidden by (V3)(b) or $\mathcal{C} \lhd \mathcal{D}$.*

Lemma 3.14. *Let α be an X-forest, and let v be an interior vertex of α with $\deg(v) \geqslant 4$. Suppose the coatoms near v are partitioned into two sets $\mathcal{C}$ and $\mathcal{D}$ where $\mathcal{C} \diamond \mathcal{D}$. Then there is always an ordering of $\mathcal{C} \cup \mathcal{D}$ that satisfies (V3) and $\mathcal{C} \lhd \mathcal{D}$, and has a given order of those of e_1^c and $\ddot{e}_1^d$ that are in $\mathcal{C}$, and the same for $\mathcal{D}$.*

Lemmas 3.13 and 3.14 are proven in Section 5.4.

Theorem 3.15. *Let α be an X-forest, and let $\mathcal{A}$ and $\mathcal{B}$ be a disjoint bipartition of the coatoms of $[\hat{0}, \alpha]$. If $\mathcal{A} \diamond \mathcal{B}$ then there is a coatom ordering for $[\hat{0}, \alpha]$ satisfying (V3) and $\mathcal{A} \lhd \mathcal{B}$.*

Proof. Let $\mathcal{A}_v = \{\gamma \in \mathcal{A} \mid \gamma$ is near $v\}$ and let $\mathcal{B}_v = \{\gamma \in \mathcal{B} \mid \gamma$ is near $v\}$ for an interior vertex v in α. Then $\mathcal{A} \diamond \mathcal{B}$ implies that $\mathcal{A}_v \diamond \mathcal{B}_v$.

It is now possible to find a coatom ordering that satisfies (V3) in the following way. Choose a component of α, and then choose an interior vertex v_p in that component. Take an ordering of the coatoms near that vertex that satisfies (V3) and $\mathcal{A}_{v_p} \lhd \mathcal{B}_{v_p}$. This is possible by Lemmas 3.13 and 3.14 with $\mathcal{C} = \mathcal{A}_{v_p}$ and $\mathcal{D} = \mathcal{B}_{v_p}$. Then for each of the vertices v_i, $i \in I$, adjacent to the first vertex, choose an ordering of the coatoms near v_i that does not contradict the earlier chosen ordering and satisfies (V3) and $\mathcal{A}_{v_i} \lhd \mathcal{B}_{v_i}$. This is possible since the coatoms near two adjacent vertices are exactly those possible to choose order of in Lemmas 3.13 and 3.14.

Then continue recursively in the same way with the vertices adjacent to them, until all interior vertices in the component are dealt with. It is now easy to combine the orderings for all interior vertices in the component to produce a coatom ordering that satisfies (V3) and does not contradict $\mathcal{A} \lhd \mathcal{B}$.

Do the same for each component. In the end, combine the orderings of the coatoms in the components by choosing one component that has at least two coatoms in $\mathcal{A}$ (if such a component exists) and put the first coatom of its ordering first in the ordering of $\mathcal{A}$ and the last coatom in $\mathcal{A}$ of its ordering last in the ordering of $\mathcal{A}$. Do the same with $\mathcal{B}$. Since $\mathcal{A} \diamond \mathcal{B}$ this implies that the result is a coatom ordering of $[\hat{0}, \alpha]$ satisfying (V3) and $\mathcal{A} \lhd \mathcal{B}$. □

Theorem 3.16. *Let Γ be an X-forest, and let $\alpha_1, \ldots, \alpha_t$ be a coatom ordering for* $[\hat{0}, \Gamma]$. *If $\alpha_1, \ldots, \alpha_t$ satisfies* (V3)*, then it is a recursive coatom ordering.*

Proof. Assume $\alpha_1, \ldots, \alpha_t$ is a coatom ordering for $[\hat{0}, \Gamma]$ satisfying (V3). It will be shown by induction over $|E(\Gamma)|$ that $\alpha_1, \ldots, \alpha_t$ is then a recursive coatom ordering. Remember that if α is a coatom of $[\hat{0}, \Gamma]$ then $|E(\alpha)| = |E(\Gamma)| - 1$.

(I) If $|E(\Gamma)| \leqslant 1$ there is only one or two coatoms. In both these cases all possible coatom orderings satisfy (V3) and are recursive coatom orderings.

(II) Assume that if $0 \leqslant |E(\Gamma)| \leqslant q$ $(q \geqslant 1)$, then every coatom ordering of $[\hat{0}, \Gamma]$ satisfying (V3) is a recursive coatom ordering.

(III) Take an X-forest Γ such that $|E(\Gamma)| = q + 1$. Let $\alpha_1, \ldots, \alpha_t$ be a coatom ordering for $[\hat{0}, \Gamma]$ satisfying (V3). Then Lemma 3.10 implies that (V3) $\Rightarrow$ (V1).

Fix j, $1 \leqslant j \leqslant t$. Let $\mathcal{A} = \{\gamma \mid \gamma \lessdot \alpha_k, \alpha_j \text{ and } k < j\}$. Since $\mathcal{S}(X) \cup \{\hat{0}\}$ is thin, the remaining coatoms in $[\hat{0}, \alpha_j]$ are $\mathcal{B} = \{\gamma \mid \gamma \lessdot \alpha_k, \alpha_j \text{ and } k > j\}$. From Lemma 3.12 and Theorem 3.15 it now follows that there is a coatom ordering for $[\hat{0}, \alpha_j]$ satisfying (V3) and $\mathcal{A} \lhd \mathcal{B}$. By the induction assumption it is a recursive coatom ordering. Hence (V2) is also satisfied. Thus $\alpha_1, \ldots, \alpha_t$ is a recursive coatom ordering for $[\hat{0}, \Gamma]$. □

Proof of Theorem 3.1. Let Γ be an X-forest. The existence of a coatom ordering satisfying (V3) for $[\hat{0}, \Gamma] \subset \mathcal{S}(X) \cup \{\hat{0}\}$ is implied by Theorem 3.15 with $\alpha = \Gamma$, $\mathcal{A} = \emptyset$, and $\mathcal{B} = \{\gamma \mid \gamma \lessdot \Gamma\}$. That this coatom ordering is a recursive coatom ordering for $[\hat{0}, \Gamma]$ follows from Theorem 3.16. □

3.4. *An example of a coatom ordering satisfying* (V3)

Figure 3 shows an X-forest Γ. The coatom ordering of $[\hat{0}, \Gamma]$ given by $g_1^c\ e_2^d\ e_3^c\ e_4^d\ g_4^c\ g_2^c\ f_1^d\ g_5^c\ e_2^c\ e_3^d\ e_4^c\ f_2^c\ g_3^c\ f_3^c\ f_1^c\ f_2^d\ f_3^d$ satisfies (V3), and thus it is a recursive coatom ordering.

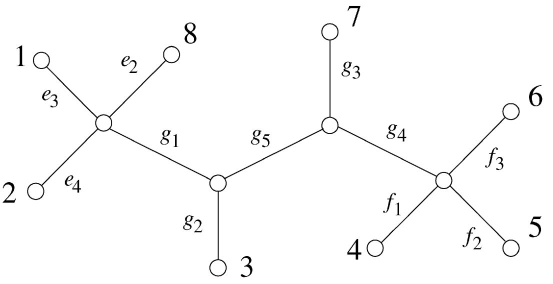

Fig. 3. The X-forest Γ.

4. The edge-product space is a regular cell complex

In this section we will assume that the reader is familiar with basic concepts of point-set topology. We will also make use of some purely topological results that for convenience we state in Appendix A.

To an X-tree $\mathcal{T}$, we associate the closed ball $\mathbf{B}(\mathcal{T}) = [0,1]^{E(\mathcal{T})}$ and open ball $\mathrm{Int}(\mathbf{B}(\mathcal{T})) = (0,1)^{E(\mathcal{T})}$. More generally, for an X-forest $\alpha = \{(A, \mathcal{T}_A)\colon A \in \pi\}$, we let $\mathbf{B}(\alpha) = \prod_{A\in\pi} \mathbf{B}(\mathcal{T}_A)$ and let $\mathrm{Int}(\mathbf{B}(\alpha)) = \prod_{A\in\pi} \mathrm{Int}(\mathbf{B}(\mathcal{T}_A))$. Note that $\mathbf{B}(\alpha)$ (respectively $\mathrm{Int}(\mathbf{B}(\alpha))$) is homeomorphic to a closed (respectively open) ball of dimension $\sum_{A\in\pi} |E(\mathcal{T}_A)|$ and accordingly we will refer to this quantity as the *dimension of* α, denoted $\dim(\alpha)$.

Given an X-tree $\mathcal{T} = (T;\phi)$ and map $\lambda : E(T) \to [0,1]$ define $p_{(\mathcal{T},\lambda)} : \binom{X}{2} \to [0,1]$ by setting

$$p_{(\mathcal{T},\lambda)}(x,y) = \prod_{e\in P(T;\phi(x),\phi(y))} \lambda(e),$$

where the empty product is taken as 1.

We can extend the correspondence $\lambda \mapsto p_{(\mathcal{T},\lambda)}$ to X-forests as follows. Given an X-forest $\alpha = \{(A, \mathcal{T}_A)\colon A \in \pi\}$ let $\psi_\alpha : \mathbf{B}(\alpha) \to [0,1]^{\binom{X}{2}}$ be defined by setting, for $\lambda = (\lambda_A\colon A \in \pi)$,

$$\psi_\alpha(\lambda)(x,y) = \begin{cases} p_{(\mathcal{T}_A,\lambda_A)}(x,y), & \text{if } \exists A \in \pi \text{ with } x,y \in A,\\ 0, & \text{otherwise.}\end{cases}$$

We begin by proving two useful lemmas. Let $\delta(\mathbf{B}(\alpha))$ denote the boundary of the ball $\mathbf{B}(\alpha)$. The following lemma describes a useful property of the map ψ_α.

Lemma 4.1. *Let $\alpha = \{(A, \mathcal{T}_A)\colon A \in \pi\}$ be an X-forest. Then,*

$$\psi_\alpha\big(\mathrm{Int}\big(\mathbf{B}(\alpha)\big)\big) \cap \psi_\alpha\big(\delta\big(\mathbf{B}(\alpha)\big)\big) = \emptyset.$$

Proof. Suppose $\psi_\alpha(\mathrm{Int}(\mathbf{B}(\alpha))) \cap \psi_\alpha(\delta(\mathbf{B}(\alpha))) \neq \emptyset$—we will show that this leads to contradictions. This assumption implies that for some $\lambda_1 \in \mathrm{Int}(\mathbf{B}(\alpha))$, and $\lambda_2 \in \delta(\mathbf{B}(\alpha))$ we have $\psi_\alpha(\lambda_1)(x,y) = \psi_\alpha(\lambda_2)(x,y)$ for all $x, y \in X$. Now, if there exists an edge e of α with $\lambda_2(e) = 0$ then select a pair $x, y \in X$ that are separated by e but contained in the same component of α. Then, $\psi_\alpha(\lambda_1)(x,y) = \psi_\alpha(\lambda_2)(x,y) = 0$, and this implies that $\lambda_1 \in \delta(\mathbf{B}(\alpha))$, a contradiction. Thus we may suppose that for every edge e of α we have $\lambda_2(e) > 0$ and so therefore also

$\psi_\alpha(\lambda_1)(x,y) = \psi_\alpha(\lambda_2)(x,y) > 0$ for all x, y that belong to any component tree $\mathcal{T}_A$ of α. Now if we let $d_i(x,y) := -\log(\psi_\alpha(\lambda_i))(x,y)$ for all x, y in $\mathcal{T}_A$, then d_i describes, for each pair $x, y \in X$, the sum of the real-valued weights $-\log(\lambda_i)(e)$ over all edges e of $\mathcal{T}_A$ that separate x and y (i.e. d_i is a distance function on X induced by this edge weighting). Now, it is a well-known and easily established result that two edge weightings of an X-tree induce the same distance function on X if and only if the two edge weightings are identical (see e.g. Lemma 2.2(i) of [9]). Consequently, since $d_1 = d_2$ it follows that λ_1 and λ_2 agree on each edge of $\mathcal{T}_A$. Since this applies for each component $A \in \pi$ it follows that $\lambda_1 = \lambda_2$. But this is impossible since $\lambda_1 \in \mathrm{Int}(\mathbf{B}(\alpha))$ and $\lambda_2 \in \delta(\mathbf{B}(\alpha))$. □

Using the following lemma we will later be able to restrict our attention to trees as opposed to forests.

Lemma 4.2. *If* $\alpha = \{(A, \mathcal{T}_A)\colon\ A \in \pi\}$ *then*

$$\psi_\alpha\big(\mathbf{B}(\alpha)\big) \cong \prod_{A\in\pi} \big(\psi\big(\mathbf{B}(\mathcal{T}_A)\big)\big).$$

Proof. This follows from Lemma A.3 of Appendix A, taking $I = \pi$, and for $A \in I$, $Z_A = \mathbf{B}(\mathcal{T}_A)$, and for $\lambda, \lambda' \in \mathbf{B}(\mathcal{T}_A)$, writing $\lambda R_A \lambda'$ if and only if $p_{(\mathcal{T},\lambda|_A)} = p_{(\mathcal{T}_A,\lambda'|_A)}$. □

We now recall the definition of a regular cell complex. In [3, Section 12.4] it states that a family of balls (homeomorphs of B^d, $d \geqslant 0$) in a Hausdorff space Y is a set of closed balls of a *regular cell complex* if and only if the interiors of the balls partition Y and the boundary of each ball is a union of other balls.

Consider the set

$$\mathcal{C} := \big\{\psi_\alpha\big(\mathbf{B}(\alpha)\big)\colon\ \alpha \in \mathcal{S}(X)\big\}.$$

We claim that this forms a set of closed balls of a regular cell complex (decomposition of $\mathcal{E}(X)$) where the boundary of each ball $\psi_\alpha(\mathbf{B}(\alpha))$, denoted $\delta(\psi_\alpha(\mathbf{B}(\alpha)))$, is defined by $\delta(\psi_\alpha(\mathbf{B}(\alpha))) := \psi_\alpha(\delta(\mathbf{B}(\alpha)))$ (so that, by Lemma 4.1, the interior of each ball $\psi_\alpha(\mathbf{B}(\alpha))$ is given by $\mathrm{Int}(\psi_\alpha(\mathbf{B}(\alpha))) = \psi_\alpha(\mathbf{B}(\alpha)) - \delta(\psi_\alpha(\mathbf{B}(\alpha))) = \psi_\alpha(\mathrm{Int}(\mathbf{B}(\alpha))))$.

To help prove our claim we first present a proposition that is a reformulation of some results appearing in [9]. Let

$$\mathcal{S}(X)_{<\alpha} := \big\{\beta \in \mathcal{S}(X)\colon\ \beta < \alpha\big\}.$$

Proposition 4.3. *The following statements hold*:

(i) $\mathcal{E}(X)$ *is the* disjoint *union of the elements of*

$$\big\{\psi_\alpha\big(\mathrm{Int}\big(\mathbf{B}(\alpha)\big)\big)\colon\ \alpha \in \mathcal{S}(X)\big\}.$$

(ii) *For* $\alpha \in \mathcal{S}(X)$, $\delta(\psi_\alpha(\mathbf{B}(\alpha)))$ *is the union of the elements of*

$$\big\{\psi_\beta\big(\mathbf{B}(\beta)\big)\colon\ \beta \in \mathcal{S}(X)_{<\alpha}\big\},$$

and the disjoint *union of the elements of*

$$\{\psi_\beta(\mathrm{Int}(\mathbf{B}(\beta))):\ \beta \in \mathcal{S}(X)_{<\alpha}\}.$$

(iii) *If α is an X-tree, then for each $y \in \psi_\alpha(\mathbf{B}(\alpha))$, $\psi_\alpha^{-1}(y)$ is a contractible regular cell complex.*

Proof. Parts (i) and (ii) follow from [9, Theorem 3.3] and the definition of $\delta(\psi_\alpha(\mathbf{B}(\alpha)))$. Part (iii) is [9, Proposition 6.5]. □

By Proposition 4.3(i), the interiors of the elements of $\mathcal{C}$ partition $\mathcal{E}(X)$, and by Proposition 4.3(ii) the boundary of each element of $\mathcal{C}$ is equal to the union of other elements in $\mathcal{C}$. Hence to show that $\mathcal{C}$ is the set of closed balls of a regular cell complex it suffices to prove the following.

Theorem 4.4. *For all $\alpha \in \mathcal{S}(X)$, the set $\psi_\alpha(\mathbf{B}(\alpha))$ is homeomorphic to* $[0,1]^{\dim(\alpha)}$.

Proof. By Lemma 4.2 it suffices to prove the theorem for $\alpha \in \mathcal{S}(X)$ an X-tree.

To prove the theorem we use induction on $\dim(\alpha)$. It can easily be checked that the result holds for $\dim(\alpha) = 0, 1, 2, 3$.

Now suppose that $d := \dim(\alpha) > 3$, and that $\psi_\beta(B(\beta))$ is homeomorphic to $[0,1]^{\dim(\beta)}$ for all $\beta \in S(X)$ such that $\dim(\beta) < d$.

By Proposition 4.3(ii) and the inductive hypothesis, $\delta(\psi_\alpha(\mathbf{B}(\alpha)))$ is a regular cell complex, with set of closed balls equal to

$$\{\psi_\beta(\mathbf{B}(\beta)):\ \beta \in \mathcal{S}(X)_{<\alpha}\}.$$

Moreover, this complex has face poset isomorphic to $(\mathcal{S}(X)_{<\alpha}, \leqslant)$ (cf. [9, Theorem 3.3]). By Theorem 2.1 the poset $[\hat{0}, \alpha]$ obtained by adding a minimal and a maximal element to $(\mathcal{S}(X)_{<\alpha}, \leqslant)$ is thin and graded (graded means pure with a unique minimal and maximal element) with length $d+1$, and by Theorem 3.1 $[\hat{0}, \alpha]$ has a recursive coatom ordering. It follows by [4, Theorem 4.7.24(i)][3] that $\psi_\alpha(\delta(\mathbf{B}(\alpha)))$ is homeomorphic to $\delta([0,1]^d)$, the $(d-1)$-dimensional sphere.

It now follows that the set $\psi_\alpha(\mathbf{B}(\alpha))$ is homeomorphic to $[0,1]^d$ by applying Proposition 4.3(iii) together with Corollary A.2 of Appendix A with $g = \psi_\alpha$, $B = B(\alpha)$, and $Z = \psi_\alpha(B(\alpha))$. □

5. Proof of some combinatorial lemmas

In this section we give the proofs of Lemmas 3.10–3.14. Since many of the proofs require cases that are straightforward to check (but quite detailed to write out), when appropriate we will refer the reader to [6] where the full details are presented.

[3] Denoting the face poset of a cell complex Δ by $\mathcal{F}(\Delta)$, this theorem states that if P is a graded poset of length $d+2$, then $P \cong \mathcal{F}(\Delta) \cup \{\hat{0}, \hat{1}\}$ for some shellable regular cell decomposition Δ of the d-sphere if and only if P is thin and admits a recursive coatom ordering.

5.1. Reformulation of (V1) *with implications*

In this section we prove Lemma 3.10.

Recall part (V1) of Definition 3.3: *For all* $i < j$ *and* $\gamma < \alpha_i, \alpha_j$ *there is a* $k < j$ *and an element* β *such that* $\gamma \leqslant \beta \lessdot \alpha_k, \alpha_j$.

An equivalent formulation is:

For each interior vertex v of Γ, where $\deg(v) = n$, the following conditions apply (recall Convention 3.4):

(V1)(a) $\{e_1^c, e_2^c\} \ntriangleleft \{e_{31}^c, e_{32}^c\}$ when $n = 3$, e_{31}, e_{32} are defined;
(V1)(b) $\{e_1^c, e_2^c\} \ntriangleleft \{\ddot{e}_3^d, \ldots, \ddot{e}_n^d\}$ when $n \geqslant 4$;
(V1)(c) $\{e_1^d, e_2^d\} \ntriangleleft \{e_3^c, e_4^c\}$ when $n = 4$, $e_1^d, e_2^d \lessdot \Gamma$;
(V1)(d) $e_1^d e_1^c \ntriangleleft \{\ddot{e}_2^d, \ldots, \ddot{e}_n^d\}$ when $n \geqslant 4$, $e_1^d \lessdot \Gamma$;
(V1)(e) $e_1^c e_1^d \ntriangleleft \{e_2^c, \ddot{e}_2^d, \ldots, e_n^c, \ddot{e}_n^d\}$ when $n \geqslant 4$, $e_1^d \lessdot \Gamma$;
(V1)(f) furthermore, if there is a component K in Γ with only one edge e, then the coatoms e^c and e^d are not the two first coatoms in the ordering.

That the new formulation of (V1) is equivalent to the original one, follows from the investigation of common elements of $[\hat{0}, \alpha_i]$ and $[\hat{0}, \alpha_j]$ which is made in Section 5.2.

Proof of Lemma 3.10. It is easy to show that (V3)(b) $\Rightarrow$ (V1)(a), (V3(c) $\Rightarrow$ (V1)(b), (V3)(c) $\Rightarrow$ (V1)(c), (V3)(d) $\Rightarrow$ (V1)(e), and (V3)(a) $\Rightarrow$ (V1)(f). Thus it only remains to show that (V1)(d) holds. Suppose the coatom ordering satisfies (V3) but not (V1)(d). Then some e_i^c, $2 \leqslant i \leqslant n$, has to come before e_1^c in the ordering because of (V3)(d). But then $\{e_1^c, e_i^c\} \triangleleft \{\ddot{e}_2^d, \ldots, \ddot{e}_n^d\}$, which contradicts (V3)(c). Hence (V3)(c) and (V3)(d) imply (V1)(d). □

5.2. Common elements of $[\hat{0}, \alpha_1]$ *and* $[\hat{0}, \alpha_2]$

Let Γ be an X-forest, and let α_1 and α_2 be different coatoms of $[\hat{0}, \Gamma]$. To reformulate the condition (V1) it is important to find the common elements of $[\hat{0}, \alpha_1]$ and $[\hat{0}, \alpha_2]$. For every pair α_1, α_2 there is either some δ such that $[\hat{0}, \alpha_1] \cap [\hat{0}, \alpha_2] = [\hat{0}, \delta_1]$ or δ_1 and δ_2 such that $[\hat{0}, \alpha_1] \cap [\hat{0}, \alpha_2] = [\hat{0}, \delta_1] \cup [\hat{0}, \delta_2]$ (see Lemma 5.1).

To reformulate (V1) we also need to find all $\beta \in \mathcal{S}(X)$ such that $\delta_i < \beta \lessdot \alpha_j$ when δ_i is not covered by α_j, $i = 1, 2$, $j = 1, 2$.

From Lemma 5.1 it now follows that the first and second formulation of (V1) are equivalent. This lemma is also needed in the proofs of Lemmas 5.3 and 3.12.

Lemma 5.1. *Let* Γ *be an* X*-forest, and let* α_1 *and* α_2 *be distinct coatoms of* $[\hat{0}, \Gamma]$. *For every pair of* α_1 *and* α_2, *Table* 1 *gives* δ *such that* $[\hat{0}, \alpha_1] \cap [\hat{0}, \alpha_2] = [\hat{0}, \delta_1]$ *or* δ_1 *and* δ_2 *such that* $[\hat{0}, \alpha_1] \cap [\hat{0}, \alpha_2] = [\hat{0}, \delta_1] \cup [\hat{0}, \delta_2]$. *Furthermore, Table* 1 (*column* 3) *states whether or not* δ_i *is covered by* α_1 *and* α_2, *and if not, column* 4 *gives* β *such that* $\delta_i < \beta \lessdot \alpha_j$ *for* $j = 1, 2$.

If nothing else is specified, all operations are made on Γ. This means that for example $f_1^c f_2^d = f_1^c(f_2^d) = f_1^c(f_2^d(\Gamma))$. Let $e_1, \ldots, e_n$ be the edges incident with a vertex v_0 in Γ. Let the other vertex incident with e_i be v_i, and let $m_i = \deg(v_i)$. If $m_1 \geqslant 3$, let $f_2, \ldots, f_m$ be the other edges incident with v_1. Let g be an edge of Γ that is not adjacent to e_1.

Table 1
Description of elements less than both α_1 and α_2

α_1, α_2	δ or δ_1, δ_2	Is $\delta \lessdot \alpha_1, \alpha_2$?	$\delta < \beta \lessdot \alpha_j$
e_1^c, g^c	$\delta = g^c e_1^c$	Yes	
e_1^c, g^d	$\delta = g^d e_1^c$	Yes	
e_1^d, g^d	$\delta = g^d e_1^d$	Yes	
e_1^c, e_2^c	$\begin{cases} \delta_1 = e_2^c e_1^c \\ \delta_2 = e_3^d \dots e_n^d \end{cases}$	Yes If $n = 3,\ m_3 \neq 3$	 $\dot{e}_3^d e_1^c \lessdot e_1^c, \dot{e}_3^d$
e_1^c, e_2^d	$\delta = e_2^d e_1^c$	Yes	
e_1^d, e_2^d	$\delta = e_2^d e_1^d$	If $n \geqslant 5$	$e_3^c e_1^d \lessdot e_1^d, e_3^c$
e_1^c, e_1^d	$\delta = \hat{0}$ if $n = m_1 = 1$	If $\lvert E(\Gamma)\rvert = 1$	$\begin{cases} \text{All } \gamma \lessdot e_1^c \\ \text{All } \gamma \lessdot e_1^d \end{cases}$
	$\delta = e_2^d \dots e_n^d$ if $n \geqslant 4, m_1 = 1$	No	$\begin{cases} \dot{e}_2^d e_1^c \lessdot e_1^c, \dot{e}_2^d \\ \dot{e}_2^d e_1^d \lessdot e_1^d, \dot{e}_2^d \\ e_2^c e_1^d \lessdot e_1^d, e_2^c \end{cases}$
	$\begin{cases} \delta_1 = e_2^d \dots e_n^d \\ \delta_2 = f_2^d \dots f_m^d \end{cases}$ if $n, m_1 \geqslant 4$	No No	As δ above Corresponding

Proof. With the help of the definition of $\leqslant$ it is rather straightforward to obtain the results in Table 1. $\square$

5.3. $\mathcal{A}$ and $\mathcal{B}$ are compatible with (V3)

The following definition is used in Sections 5.4 and 5.3.

Definition 5.2. Let Γ be an X-forest and e an edge in Γ, and let α be $e^c(\Gamma)$, $e^d(\Gamma)$, $\dot{e}^d(\Gamma)$ or $\ddot{e}^d(\Gamma)$. The symbol $\langle\alpha\rangle$ denotes α if α is defined and $\alpha \lessdot \Gamma$.

In this section we prove Lemma 3.12. To do this, we require some preliminary results.

Let $\alpha_1, \dots, \alpha_j, \dots, \alpha_t$ be a coatom ordering of $[\hat{0}, \Gamma]$ satisfying (V3). Fix j, and consider the interval $[\hat{0}, \alpha_j]$. Let $\mathcal{A} = \{\gamma \mid \gamma \lessdot \alpha_k, \alpha_j \text{ and } k < j\}$ and let $\mathcal{B} = \{\gamma \mid \gamma \lessdot \alpha_k, \alpha_j \text{ and } k > j\}$. The sets $\mathcal{A}$ and $\mathcal{B}$ are disjoint since $\mathcal{S}(X) \cup \{\hat{0}\}$ is thin.

Lemma 5.3. *If v is an interior vertex of α_j, $\mathcal{A}_v = \{\gamma \in \mathcal{A} \mid \gamma \text{ is near } v\}$, and $\mathcal{B}_v = \{\gamma \in \mathcal{B} \mid \gamma \text{ is near } v\}$, then $\mathcal{A}_v \diamond \mathcal{B}_v$.*

Proof. Let $\mathcal{A}' = \{\beta \lessdot \Gamma \mid \gamma \lessdot \beta \text{ and } \gamma \in \mathcal{A}_v\}$ and $\mathcal{B}' = \{\beta \lessdot \Gamma \mid \gamma \lessdot \beta \text{ and } \gamma \in \mathcal{B}_v\}$. Then $\mathcal{A}' \subseteq \{\alpha_1, \dots, \alpha_{j-1}\}$ and $\mathcal{B}' \subseteq \{\alpha_{j+1}, \dots, \alpha_t\}$. This lemma will be proven by assuming $\mathcal{A}_v \not\diamond \mathcal{B}_v$ for some interior vertex v in α_j, and then deducing that $\alpha_1, \dots, \alpha_t$ does not satisfy (V3) (which is a contradiction) by showing that $\mathcal{A}' \lhd \{\alpha_j\} \lhd \mathcal{B}'$ is forbidden by one of the sub-conditions (V3)(b), (V3)(c), and (V3)(d).

First, note that since all coatoms near v are obtained by contracting or deleting edges in the same component of α_j and (V3) is symmetric, $\mathcal{A}_v \not\diamond \mathcal{B}_v$ implies that $\mathcal{A}_v \lhd \mathcal{B}_v$ is forbidden by one of the sub-conditions (V3)(b), (V3)(c), and (V3)(d).

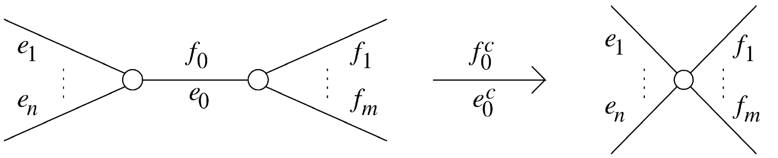

Fig. 4. Case 1.

To simplify the proof, we use the following notation. Let $\mathcal{C} \subseteq \mathcal{A}_v$ and $\mathcal{D} \subseteq \mathcal{B}_v$. Then $\mathcal{C} \lhd \mathcal{D} \longleftarrow \mathcal{C}' \lhd \{\alpha_j\} \lhd \mathcal{D}'\}$ means that $\mathcal{C}' = \{\beta \lessdot \Gamma \mid \gamma \lessdot \alpha_j, \beta$ where $\gamma \in \mathcal{C}\}$ and $\mathcal{D}' = \{\beta \lessdot \Gamma \mid \gamma \lessdot \alpha_j, \beta$ where $\gamma \in \mathcal{D}\}$. Since $\mathcal{S}(X) \cup \{\hat{0}\}$ is thin, $\mathcal{C}'$ and $\mathcal{D}'$ are well defined and are always disjoint. Observe that if $\mathcal{C} \lhd \mathcal{D} \longleftarrow \mathcal{C}' \lhd \{\alpha\} \lhd \mathcal{D}'$, then $\mathcal{D} \lhd \mathcal{C} \longleftarrow \mathcal{D}' \lhd \{\alpha\} \lhd \mathcal{C}'$. Since (V3) is symmetric it is not necessary to check the reverse of any condition.

Based on the result in Lemma 5.1, we can divide the rest of the proof into cases as below. For each case, (V3)(b), (V3)(c), and (V3)(d) are checked.

The interior vertex v in α_j corresponds to one interior vertex v' or two interior vertices v_1' and v_2' in Γ. Let the edges incident with v' or v_1' be denoted $e_0, \ldots, e_n$, and let the edges incident with v_2' be denoted $f_0 = e_0, f_1, \ldots, f_m$.

Case 1. The coatom α_j is obtained from Γ by contracting the edge e_0 incident with v_1' and v_2', where $\deg(v_1') \geqslant 3$ and $\deg(v_2') \geqslant 3$. This case is illustrated in Fig. 4, where $n, m \geqslant 2$.

We will deduce the desired contradiction for the following case. Suppose $m \geqslant 3$ and that $\mathcal{A}_v \lhd \mathcal{B}_v$ is forbidden by (V3)(c). If $e_1^c, e_2^c \in \mathcal{A}_v$ and $\ddot{e}_3^d, \ldots, \ddot{e}_n^d, \ddot{f}_1^d, \ldots, \ddot{f}_m^d \in \mathcal{B}_v$, then by Lemma 5.1 $\{e_1^c, e_2^c\} \lhd \{\ddot{e}_3^d, \ldots, \ddot{e}_n^d, \ddot{f}_1^d, \ldots, \ddot{f}_m^d\} \longleftarrow \{e_1^c, e_2^c\} \lhd \{f_0^c\} \lhd \{\ddot{e}_3^d, \ldots, \ddot{e}_n^d, \ddot{f}_1^d, \ldots, \ddot{f}_m^d\}$. If $f_0^d \lessdot \Gamma$, then (V3)(c) implies that $\{e_1^c, e_2^c, f_0^d\} \lhd \{f_0^c\} \lhd \{\ddot{e}_3^d, \ldots, \ddot{e}_n^d, \ddot{f}_1^d, \ldots, \ddot{f}_m^d, f_1^c, \ldots, f_m^c\}$. If f_0^d is not covered by Γ, then $n = 2$ and $\ddot{f}_0^d = \{e_1^c, e_2^c\}$. This means that $\{\ddot{f}_0^d, f_0^c\} \lhd \{\ddot{f}_1^d, f_1^c, \ldots, \ddot{f}_m^d, f_m^c\}$ which is forbidden by (V3)(c). This is a contradiction since $\alpha_1, \ldots, \alpha_t$ satisfies (V3).

The other cases are dealt with similarly (see [6, Section 5.2] for details).

Case 2. The coatom α_j is obtained from Γ by deleting the edge e_0 incident with v', where $\deg(v') \geqslant 4$.

We deduce the desired contradiction for the following case. Suppose $n \geqslant 4$ and that $\mathcal{A}_v \lhd \mathcal{B}_v$ is forbidden by (V3)(d). Then for some $1 \leqslant k \leqslant n-1$ $e_1^c, \ddot{e}_1^d, \ldots, e_k^c, \ddot{e}_k^d \in \mathcal{A}_v$ and $e_{k+1}^c, \ddot{e}_{k+1}^d, \ldots, e_n^c, \ddot{e}_n^d \in \mathcal{B}_v$. By Lemma 5.1 we have $\{e_1^c, \ddot{e}_1^d, \ldots, e_k^c, \ddot{e}_k^d\} \lhd \{e_{k+1}^c, \ddot{e}_{k+1}^d, \ldots, e_n^c, \ddot{e}_n^d\} \longleftarrow \{e_1^c, \ddot{e}_1^d, \ldots, e_k^c, \ddot{e}_k^d\} \lhd \{e_0^d\} \lhd \{e_{k+1}^c, \ddot{e}_{k+1}^d, \ldots, e_n^c, \ddot{e}_n^d\}$. Since $e_0^c \lessdot \Gamma$, e_0^c has to come before or after e_0^d in the coatom ordering. This implies that $\{e_0^c, \ddot{e}_0^d, \ldots, e_k^c, \ddot{e}_k^d\} \lhd \{e_{k+1}^c, \ddot{e}_{k+1}^d, \ldots, e_n^c, \ddot{e}_n^d\}$ or $\{e_1^c, \ddot{e}_1^d, \ldots, e_k^c, \ddot{e}_k^d\} \lhd \{e_0^c, \ddot{e}_0^d, e_{k+1}^c, \ddot{e}_{k+1}^d, \ldots, e_n^c, \ddot{e}_n^d\}$ which are both forbidden by (V3)(d). This is a contradiction since $\alpha_1, \ldots, \alpha_t$ satisfies (V3).

The other cases are dealt with similarly (see [6, Section 5.2] for details).

The remaining cases concern coatoms α_j obtained from Γ by contracting or deleting the edge g not incident with v'. In this situation $\deg(v) = \deg(v')$.

Let e_0 be incident with v' and v_0' in Γ. If g is not adjacent to any edge e_i, then the coatoms near v are obtained from α_j by the same operations as the coatoms near v' in Γ. Furthermore, if a coatom γ near v is obtained from α_j by a certain operation, then γ is covered by the coatom near v' which is obtained from Γ by the same operation. Hence, if $\mathcal{A}_v \lhd \mathcal{B}_v$ is forbidden by one of (V3)(b), (V3)(c), and (V3)(d), then $\mathcal{A}' \lhd \mathcal{B}'$ obviously is forbidden by the same condition. This also applies when g is adjacent to e_0, if α_j is obtained from Γ by contracting g and $\deg(v_0') \geqslant 4$, or if α_j is obtained from Γ by deleting g and $\deg(v_0') \geqslant 5$. There are now two cases left.

Case 3. The coatom α_j is obtained from Γ by contracting the edge g adjacent to e_0, where $\deg(v_0') = 3$.

The conditions (V3)(b), (V3)(c) and (V3)(d) can be checked in a similar way to Cases 1 and 2. Details can be found in [6, Section 5.2].

Case 4. The coatom α_j is obtained from Γ by deleting the edge g adjacent to e_0, where $\deg(v_0') = 4$.

The conditions (V3)(b), (V3)(c) and (V3)(d) can be checked in a similar fashion to Cases 1 and 2. Details can be found in [6, Section 5.2].

By checking the conditions for all cases it is found that $\mathcal{A}_v \lhd \mathcal{B}_v$ is not forbidden by any of the conditions (V3)(b), (V3)(c), and (V3)(d). Thus $\mathcal{A}_v \diamond \mathcal{B}_v$ for every interior vertex v in α_j. $\square$

We will now prove that $\mathcal{A} \lhd \mathcal{B}$ is not forbidden by the condition (V3)(a) in Lemmas 5.5 and 5.6. The following notation will be helpful.

Definition 5.4. If γ is a coatom of $[\hat{0}, \beta]$ which is obtained by contraction or deletion of an edge of the component K in β, then we write $\gamma \in K$.

Recall Convention 3.4. The X-forest α_j is obtained from Γ by contracting or deleting an edge e_0. Sometimes contraction or deletion of an edge results in splitting of a component in two or more parts (recall Convention 3.5). Let $K_1, \ldots, K_{\ell+1}$ be the non-trivial components in Γ, and suppose $e_0 \in K_{\ell+1}$. Thus α_j has the non-trivial components $K_1, \ldots, K_\ell, K_1', \ldots, K_s'$ where $s \geqslant 0$. Let v_1 and v_2 be the vertices incident with e_0, let $e_0, \ldots, e_n$ be the edges incident with v_1, and let $e_0, f_1, \ldots, f_m$ be the edges incident with v_2. Then $\deg(v_1) = n + 1$ and $\deg(v_2) = m + 1$.

Lemma 5.5. *If $s \geqslant 2$, then $\{\gamma \lessdot \alpha_j \mid \gamma \in \bigcup_{i=1}^{k} K_i'\} \ntriangleleft \{\gamma \lessdot \alpha_j \mid \gamma \in \bigcup_{i=k+1}^{s} K_i'\}$ for all $1 \leqslant k \leqslant s - 1$.*

Proof. Suppose $\{\gamma \mid \gamma \in \bigcup_{i=1}^{k} K_i'\} \lhd \{\gamma \mid \gamma \in \bigcup_{i=k+1}^{s} K_i'\}$ for some $1 \leqslant k \leqslant s - 1$. We show that $\alpha_1, \ldots, \alpha_t$ does not satisfy (V3), a contradiction from which the lemma follows. The following two cases arise. The details are straightforward and similar to those in Lemma 5.3. They will be omitted here, but can be found in [6, Section 5.3].

Case 1. v_1 is unlabelled, v_2 is labelled, and $\alpha_j = e_0^c(\Gamma)$. By Convention 3.5 it follows that $s = n$, and in α_j that $e_i \in K_i'$ and e_i is incident with a leaf for $1 \leqslant i \leqslant n$.

It is easy to show that $\{e_1^c, \ddot{e}_1^d, \ldots, e_k^c, \ddot{e}_k^d\} \lhd \{e_{k+1}^c, \ddot{e}_{k+1}^d, \ldots, e_n^c, \ddot{e}_n^d\}$ implies that $\alpha_1, \ldots, \alpha_t$ does not satisfy (V3).

Case 2. The vertices v_1 and v_2 are unlabelled and $\alpha_j = e_0^d(\Gamma)$. Hence $s = 2$.

Using $\{e_1^c, \langle \ddot{e}_1^d \rangle, \ldots, e_n^c, \langle \ddot{e}_n^d \rangle\} \lhd \{f_1^c, \langle \ddot{f}_1^d \rangle, \ldots, f_m^c, \langle \ddot{f}_m^d \rangle\}$, it is easy to show that $\alpha_1, \ldots, \alpha_t$ does not satisfy (V3). $\square$

Lemma 5.6. *$\{\gamma \lessdot \alpha_j \mid \gamma \in \bigcup_{i=1}^{k} K_i\} \not\lhd \{\gamma \lessdot \alpha_j \mid \gamma \in \bigcup_{i=k+1}^{\ell} K_i \cup \bigcup_{i=1}^{s} K_i'\}$ for all $1 \leqslant k \leqslant \ell$ if $s > 0$, and for all $1 \leqslant k \leqslant \ell - 1$ if $s = 0$.*

Proof. Let $\mathcal{C}_1 = \{\gamma \lessdot \alpha_j \mid \gamma \in \bigcup_{i=1}^{k} K_i\}$, $\mathcal{C}_2 = \{\gamma \lessdot \alpha_j \mid \gamma \in \bigcup_{i=k+1}^{\ell} K_i\}$, and $\mathcal{C}_3 = \{\gamma \lessdot \alpha_j \mid \gamma \in \bigcup_{i=1}^{s} K_i'\}$. Also, let $\mathcal{D}_1 = \{\beta \lessdot \Gamma \mid \beta \in \bigcup_{i=1}^{k} K_i\}$, $\mathcal{D}_2 = \{\beta \lessdot \Gamma \mid \beta \in \bigcup_{i=k+1}^{\ell} K_i\}$, and $\mathcal{D}_3 = \{\beta \lessdot \Gamma \mid \beta \in K_{\ell+1}\}$. The lemma now states that $\mathcal{C}_1 \not\lhd (\mathcal{C}_2 \cup \mathcal{C}_3)$.

Note that $\mathcal{D}_i = \{\beta \lessdot \Gamma \mid \gamma \lessdot \alpha_j, \beta \text{ where } \gamma \in \mathcal{C}_i\}$ for $i = 1, 2$. The set $\mathcal{D}_3$ is the disjoint union of $\{\beta \lessdot \Gamma \mid \gamma \lessdot \alpha_j, \beta \text{ where } \gamma \in \mathcal{C}_3\}$ and the set $\mathcal{D}'$ defined by

$$\mathcal{D}' = \begin{cases} \{e_0^c, \langle e_0^d \rangle\} & \text{if } \alpha_j = e_0^d, n \neq 3, m \neq 3 \text{ or if } \alpha_j = e_0^c; \\ \{e_0^c, e_0^d, \langle e_1^d \rangle, \langle e_2^d \rangle, \langle e_3^d \rangle\} & \text{if } \alpha_j = e_0^d,\ n = 3,\ m \neq 3; \\ \{e_0^c, e_0^d, \langle e_1^d \rangle, \langle e_2^d \rangle, \langle e_3^d \rangle, \langle f_1^d \rangle, \langle f_2^d \rangle, \langle f_3^d \rangle\} & \text{if } \alpha_j = e_0^d,\ n = m = 3. \end{cases}$$

Now, suppose $\mathcal{C}_1 \lhd (\mathcal{C}_2 \cup \mathcal{C}_3)$. Then it is easy to deduce that $\mathcal{D}_1 \lhd (\mathcal{D}_2 \cup \mathcal{D}_3)$ if $s \geqslant 1$ and that $(\mathcal{D}_1 \cup \mathcal{D}_3) \lhd \mathcal{D}_2$ or $\mathcal{D}_1 \lhd (\mathcal{D}_2 \cup \mathcal{D}_3)$ if $s = 0$, since $\mathcal{C}_1 \lhd (\mathcal{C}_2 \cup \mathcal{C}_3) \longleftarrow \mathcal{D}_1 \lhd \{\alpha_j\} \lhd (\mathcal{D}_2 \cup (\mathcal{D}_3 \setminus \mathcal{D}'))$. The details are straightforward and therefore omitted (see [6, Section 5.3]). This is forbidden by the condition (V3)(a), hence $\alpha_1, \ldots, \alpha_t$ does not satisfy (V3). Since this is a contradiction, $\mathcal{C}_1 \not\lhd (\mathcal{C}_2 \cup \mathcal{C}_3)$. $\square$

Proof of Lemma 3.12. From Lemma 5.3 follows that $\mathcal{A} \lhd \mathcal{B}$ is not forbidden by (V3)(b), (V3)(c), or (V3)(d). Lemmas 5.5 and 5.6 imply that $\mathcal{A} \lhd \mathcal{B}$ is not forbidden by (V3)(a). Thus $\mathcal{A} \diamond \mathcal{B}$. $\square$

5.4. *The order of the coatoms near two vertices*

In this section the following notation will be used. Let α be an X-forest and $\mathcal{C}$ a set of coatoms of α. In an ordering of $\mathcal{C}$, $[\beta]$ denotes β if $\beta \in \mathcal{C}$.

Proof of Lemma 3.13. It is easy to show that there always exists an order of e_1^c, e_2^c, e_3^c, $\langle e_{11}^c \rangle$, and $\langle e_{12}^c \rangle$ that satisfies (V3)(b) and $\mathcal{C} \lhd \mathcal{D}$.

Since $\deg(v) = 3$, it follows that (V3) is satisfied if (V3)(b) is satisfied. Order the elements of $\mathcal{C} \cup \mathcal{D}$ so that $\mathcal{C} \lhd \mathcal{D}$ and e_1^c, e_2^c, e_3^c, $\langle e_{11}^c \rangle$ and $\langle e_{12}^c \rangle$ have the given order. Then put $\langle e_{21}^c \rangle$ and $\langle e_{22}^c \rangle$ as far from each other as possible in the ordering without violating $\mathcal{C} \lhd \mathcal{D}$, and do the same with $\langle e_{31}^c \rangle$ and $\langle e_{32}^c \rangle$. It is now easy to see that this ordering satisfies (V3) and $\mathcal{C} \lhd \mathcal{D}$, and has the given order of e_1^c, e_2^c, e_3^c, $\langle e_{11}^c \rangle$ and $\langle e_{12}^c \rangle$. $\square$

Proof of Lemma 3.14. If $\gamma_1, \ldots, \gamma_k$ is an ordering of $\mathcal{C}$ such that $\{\gamma_1, \ldots, \gamma_i\} \diamond \{\gamma_{i+1}, \ldots, \gamma_k\} \cup \mathcal{D}$ for all $1 \leqslant i \leqslant k-1$, and $\delta_1, \ldots, \delta_\ell$ is an ordering of $\mathcal{D}$ such that $\mathcal{C} \cup \{\delta_1, \ldots, \delta_i\} \diamond \{\delta_{i+1}, \ldots, \delta_\ell\}$ for all $1 \leqslant i \leqslant \ell - 1$, then the ordering $\gamma_1, \ldots, \gamma_k, \delta_1, \ldots, \delta_\ell$ of $\mathcal{C} \cup \mathcal{D}$ satisfies (V3) and $\mathcal{C} \lhd \mathcal{D}$.

The symmetry of (V3) implies that if it is possible to find orderings as above for $\mathcal{C}$, then it is also possible to find the desired orderings for $\mathcal{D}$. Hence it suffices to find such orderings of $\mathcal{C}$. In this case the sub-conditions of (V3) that need to be checked are (V3)(c) and (V3)(d).

If $\mathcal{D} = \emptyset$, we may give $\mathcal{C}$ the ordering $e_2^c\ \ddot{e}_3^d\ \ldots\ \ddot{e}_2^d\ e_3^c$. If $|\mathcal{C}| \leqslant 2$, then give $\mathcal{C}$ any ordering.

Now suppose $\mathcal{D} \neq \emptyset$ and $|\mathcal{C}| \geqslant 3$. If there exists an edge f_1 such that $f_1^c \in \mathcal{C}$ and $\dot{f}_1^d \in \mathcal{D}$, then $|\mathcal{C}| \geqslant 3$ and (V3)(c) imply that there is some $\dot{f}_2^d \in \mathcal{C}$. If also $f_2^c \in \mathcal{C}$, then (V3)(c) implies that there is some $\dot{f}_3^d \in \mathcal{C}$. Hence either $\dot{f}_2^d \in \mathcal{C}$, $f_2^c \in \mathcal{D}$, or there is some $\dot{f}_2^d \in \mathcal{C}$ such that $f_2 \neq e_1$. Now it is easy to show that the orderings $f_1^c\ \dot{f}_2^d\ \ldots\ [f_2^c]$ or $f_{11}^c\ f_1^c\ \dot{f}_2^d\ \ldots\ [f_2^c]$ of $\mathcal{C}$ will do.

If $f_i^c \in \mathcal{C} \Rightarrow \ddot{f}_i^d \in \mathcal{C}$ for all edges f_i incident with v, then (V3)(d) implies that there exists an edge f_1 such that $\ddot{f}_1^d \in \mathcal{C}$ and $f_1^c \in \mathcal{D}$. If f_2^c and/or f_3^c are in $\mathcal{C}$, then $\ddot{f}_2^d$ and/or $\ddot{f}_3^d$ are also in $\mathcal{C}$. Then the ordering $\ddot{f}_1^d\ [f_2^c]\ [\ddot{f}_3^d]\ [f_3^c]\ \ldots\ [\ddot{f}_2^d]$ of $\mathcal{C}$ will do.

Thus in all cases above it is possible to choose an ordering that has the given order of those of e_1^c and $\ddot{e}_1^d$ that are in $\mathcal{C}$. □

Acknowledgments

We thank Anders Björner for valuable comments on an earlier version of this manuscript. We also thank Richard Hain and especially Robert Daverman for their help with the topological results stated in Appendix A. Steel thanks the Marsden Fund.

Appendix A

In this appendix we present some topological results that we use in Section 4. For basic terminology and results concerning point-set topology see, for example, [10].

Proposition A.1. *Suppose that B is a ball, with boundary S, a sphere. If $f : S \to S$ is a continuous map such that $f^{-1}(p)$ is contractible for each $p \in S$, then f extends to a continuous map $F : B \to B$ such that F restricted to* $\mathrm{Int}(B)$ *is injective and* $F(\mathrm{Int}(B)) = \mathrm{Int}(B)$.

Proof. The conditions on f ensure that f is a cellular map, and therefore it can be approximated by homeomorphisms—for the definition of these terms, and the justification of the last statement, see [7] and (for the dimension 4 case) [11]. The existence of an extension F with the properties promised by Proposition A.1 now follows by [7, Lemma 1]. □

Now, given topological spaces Y, Z, and a map $f : Y \to Z$, let R_f be the equivalence relation on Y induced by f, i.e. the equivalence relation that identifies elements of Y that are mapped by f to the same element of Z. Recall that if Y is compact, and f is a continuous surjection with Z Hausdorff, then Z is homeomorphic to the quotient space Y/R_f. With this fact and the previous proposition in hand, we now present a result that is used in the proof of Theorem 4.4.

Corollary A.2. *Let Z be a Hausdorff topological space, and let B be a d-dimensional ball with boundary S, a sphere. Suppose that $g : B \to Z$ is a continuous surjection whose restriction to* $\mathrm{Int}(B)$ *is bicontinuous and injective. Suppose furthermore that $g(S)$ is homeomorphic to S, $g(S) \cap g(\mathrm{Int}(B)) = \emptyset$, and that for each $q \in g(S)$, $g^{-1}(q)$ is contractible. Then Z is homeomorphic to B.*

Proof. By assumption the space Z is homeomorphic to the quotient space B/R_g. Now, by Proposition A.1, the map $g|_S$ extends to a continuous map $F : B \to B$ that maps $\mathrm{Int}(B)$ injectively onto $\mathrm{Int}(B)$. In particular, it follows that B is homeomorphic to B/R_F. Now, $R_F = R_g$ since F and g are both injective on $\mathrm{Int}(B)$, $F(\mathrm{Int}(B)) \cap F(S) = g(\mathrm{Int}(B)) \cap g(S) = \emptyset$, and $F|_S = g$. Consequently, $B/R_F = B/R_g$, from which it immediately follows that Z is homeomorphic to B. □

We conclude this appendix with a result that is used in the proof of Lemma 4.2.

Lemma A.3. *Let I be a finite index set. For each $i \in I$, let Z_i be compact topological space, and R_i be an equivalence relation on Z_i such that $Y_i := Z_i/R_i$ is Hausdorff. Let R be the (product) equivalence relation on $Z := \prod_{i\in I} Z_i$ defined by $(z_i)_{i\in I} R (z_i')_{i\in I}$ if and only if $z_i R_i z_i'$ for each $i \in I$. Then Z/R is homeomorphic to $\prod_{i\in I} Y_i$.*

Proof. Set $Y := \prod_{i\in I} Y_i$, and let $f : Z \to Y$ be the map that takes $(z_i)_{i\in I}$ to $([z_i])_{i\in I}$, where, for $z_i \in Z_i$, $[z_i]$ denotes the equivalence class of z_i under relation R_i.

Set $W = Z/R$, and let $g : Z \to W$ be the associated quotient mapping. By the universal property of quotient mappings, there is a continuous mapping $h : W \to Y$ with $h \circ g = f$. By standard properties of quotient mappings h is a bijection. But since Z is compact, so is W, and since each Y_i is Hausdorff, so is Y. Hence h is a continuous bijection from a compact space to a Hausdorff space, and is therefore a homeomorphism. □

References

[1] W.P. Baritompa, The space of edge-weighted trees is a Euclidean cell for trees with exactly one interior vertex, Report No. UCDMS2003/3, Department of Mathematics and Statistics, University of Canterbury, Christchurch, New Zealand, 2003.
[2] L.J. Billera, S.P. Holmes, K. Vogtmann, Geometry of the space of phylogenetic trees, Adv. in Appl. Math. 27 (4) (2001) 733–767.
[3] A. Björner, Topological methods, in: Handbook of Combinatorics, MIT Press, Cambridge, MA, 1995, pp. 1819–1872.
[4] A. Björner, M. Las Vergnas, B. Sturmfels, N. White, G. Ziegler, Oriented Matroids, Cambridge Univ. Press, Cambridge, 1993.
[5] J. Felsenstein, Inferring Phylogenies, Sinauer Press, Sunderland, MA, 2004.
[6] J. Gill, The k-assignment polytope and the space of evolutionary trees, Licentiate thesis, Linköping Studies in Science and Technology, Thesis No. 1117, available at http://www.ep.liu.se/lic/science_technology/11/17/, 2004.
[7] W. Haver, A characterization theorem for cellular maps, Bull. Amer. Math. Soc. 76 (1970) 1277–1280.
[8] J. Kim, Slicing hyperdimensional oranges: The geometry of phylogenetic estimation, Molecular Phyl. Evol. 17 (2000) 58–75.
[9] V. Moulton, M. Steel, Peeling phylogenetic ‘oranges’, Adv. in Appl. Math. 33 (2004) 710–727.
[10] J. Munkres, Topology, 2nd ed., Pearson Education, 1999.
[11] F. Quinn, Ends of maps III: Dimensions 4 and 5, J. Differential Geom. 17 (1982) 503–521.
[12] C. Semple, M. Steel, Phylogenetics, Oxford Univ. Press, Oxford, 2003.
[13] C. Tuffley, Trees and Ps and things that sneeze: Markov process models of site substitution, MSc thesis, University of Canterbury, Christchurch, New Zealand, 1997.