Peeling phylogenetic 'oranges' *

Vincent Moulton,†
The Linnaeus Centre for Bioinformatics,
Uppsala University,
Box 598,
751 24 Uppsala,
Sweden.
vincent.moulton@lcb.uu.se
Tel: +46 18 471 6689
FAX: +46 18 471 6698
and
Mike Steel,
Biomathematics Research Centre,
Department of Mathematics and Statistics,
University of Canterbury,
Christchurch,
New Zealand.
m.steel@math.canterbury.ac.nz

22 Jan. 2003

Abstract

We investigate the combinatorics of a topological space that is generated by the set of edge-weighted finite trees. This space arises by multiplying the weights of edges on paths in trees and is closely connected to tree reconstruction problems involving finite state Markov processes. We show that this space is a contractible finite CW-complex whose face poset can be described via a partial order on semilabelled forests. We then describe some combinatorial properties of this poset, showing that, for example, it is pure, thin and contractible.

Keywords: Trees, forests, partitions, poset, contractibility

1 Introduction

Posets of trees and forests and associated spaces have been used as a tool in the representation theory of the symmetric group [7, 11]. However, recently such objects have also appeared in areas of classification such as evolutionary biology [1]. In this paper we introduce a poset on forests of semi-labelled trees that arises naturally from the set of edge-weighted trees. This space is closely connected to the reconstructability of trees under Markov random processes and has been called the *reconstruction quotient* in [12] and also been described by Junhyong Kim [6] as a space of "hyperdimensional oranges." We now define this topological space.

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 let $\mathcal{T}(X)$ denote the (finite) set of trees T that have X as their set of leaves (degree one vertices). 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

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

and let $\mathcal{E}(X)$ be the union of the subspaces $\mathcal{E}(X,T)$ of $[0,1]^{\binom{X}{2}}$ over all $T \in \mathcal{T}(X)$. We call $\mathcal{E}(X)$ the *edge-product space for trees on* X.

Apart from their intrinsic interest, a central motivation for investigating these spaces is that they are intimately connected to the problem of reconstructing trees from discrete leaf-colorations generated by a tree-indexed Markov process. Such processes are fundamental to modern molecular evolutionary biology, and other areas of classification in science. Briefly, this is because transition matrices are multiplied along paths, and if we are interested in reconstructing a discrete tree from the leaf-colorations induced by such random colorations, it has been shown ([5], [10], [12]) that we change nothing (topologically) in passing to the multiplication of real values (the determinants of the transition matrices). Consequently, the edge-product space is homeomorphic to the quotient space where edge-weighted trees are identified if they induce the same Markov process at the leaves.

So far little has been formally established about the topology or geometry of $\mathcal{E}(X)$ (or $\mathcal{E}(X,T)$) despite considerable interest in the properties of a related space where one adds rather than multiplies positive real numbers along paths in trees. This related 'additive' space has some attractive combinatorial properties (see for example, [1, 11]) and its metric properties are of interest in applications [1]. However it is the 'multiplicative' space that we study here which is the appropriate context for studying Markov process.

In this paper we will show that $\mathcal{E}(X)$ has a natural CW–complex structure for any finite set X, give a combinatorial description of the associated face poset, and use this description to determine some properties of this poset. We begin in Section 2 by providing some background terminology and results concerning X–trees and tree metrics. In Section 3 we describe a CW–complex structure on $\mathcal{E}(X)$ and show how it (and its face poset) can be naturally parameterized by a poset of X-forests, $\mathcal{S}(X)$. In Section 4, we determine some structural properties of $\mathcal{S}(X)$, in particular showing that it is pure and thin. In Section 5 we show that $\mathcal{E}(X)$ and the geometric realization of $\mathcal{S}(X)$ are both contractible. Finally in Section 6 we describe explicitly the fiber of the map Λ_T over any point in $\mathcal{E}(X,T)$ showing that it is a contractible regular cell complex, whose dimension can be readily computed.

2 Preliminaries on X–trees and tree metrics

In this section we present some material on trees that is important for the formulation of the results that follow later in the paper. 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 $V - \phi(X)$ have degree greater than two. We call the vertices in $V - \phi(X)$ *unlabelled.* Two X–trees $(T_1;\phi_1)$ and $(T_2;\phi_2)$ are isomorphic if there is a graph isomorphism $\alpha : V(T_1) \to V(T_2)$ such that $\phi_2 = \alpha \circ \phi_1$. For an X–tree $\mathcal{T} = (T;\phi)$ we let $E(\mathcal{T})$ denote $E(T)$, the set of edges of T.

A collection of bipartitions or *splits* of X is called a *split system* on X. We will write $A|B$ to denote the split $\{A,B\}$. Given a split system Σ on X and a subset Y of X, let

$$\Sigma|Y = \{B\cap Y|C\cap Y : B|C \in \Sigma, B\cap Y \neq \emptyset, C\cap Y \neq \emptyset\},$$

called the *restriction* of Σ to Y. If $\sigma = B|C \in \Sigma$, and $B\cap Y|C\cap Y$ is contained in $\Sigma|Y$ then we will denote $B\cap Y|C\cap Y$ by $\sigma|Y$. A split system Σ is said to be *pairwise compatible* if, for any two splits $A|B$ and $C|D$ in Σ, we have

$$\emptyset \in \{A\cap C, A\cap D, B\cap C, B\cap D\}.$$

Given an X–tree, $\mathcal{T} = (T;\phi)$, and an edge e of T, delete e from T and denote the vertices of the two connected components of the resulting graph by U and V. If we let $A = \phi^{-1}(U)$ and $B = \phi^{-1}(V)$ then it is easily checked that $A|B$ is a split of X, and that different edges of T induce different splits of X. We say that the split $A|B$ *corresponds to* edge e (and visa versa). Let $\Sigma(\mathcal{T})$ denote the set of all splits of X that are induced by this process of deleting one edge from $\mathcal{T}$. The following fundamental result is due to Buneman [2].

Proposition 2.1 *Let Σ be a split system on X. Then, there exists an X–tree $\mathcal{T}$ such that $\Sigma = \Sigma(\mathcal{T})$ if and only if Σ is pairwise compatible. Furthermore, in this case, $\mathcal{T}$ is unique up to isomorphism.*

Thus we may regard pairwise compatible split systems and (isomorphism classes of) X–trees as essentially equivalent. This allows us to make the following definitions that will be useful later.

- Given an X–tree, $\mathcal{T}$ and a non-empty subset Y of X let $\mathcal{T}|Y$ be the Y–tree for which $\Sigma(\mathcal{T}|Y) = \Sigma(\mathcal{T})|Y$.
- For an X–tree $\mathcal{T}$ and a Y–tree $\mathcal{T}'$, where $Y \subseteq X$, we say that $\mathcal{T}$ *displays* $\mathcal{T}'$ if $\Sigma(\mathcal{T}') \subseteq \Sigma(\mathcal{T}|Y)$ $(= \Sigma(\mathcal{T})|Y)$.

A further concept that will be useful to us is the notion of a tree metric, which we now describe. Suppose that $\mathcal{T} = (T;\phi)$ is an X–tree, and $w : E(T) \rightarrow \mathbb{R}^{>0}$. Let $d_{(\mathcal{T},w)} : \binom{X}{2} \rightarrow \mathbb{R}^{>0}$ be defined by

$$d_{(\mathcal{T},w)} = \sum_{e \in P(T;\phi(x),\phi(y))} w(e).$$

Any function $d : \binom{X}{2} \rightarrow \mathbb{R}^{>0}$ that can be written in this way is said to be a *tree metric (with representation $(\mathcal{T}, w)$)*. Recall that a *topological embedding* is a map between two topological spaces that is one-to-one and bicontinuous (i.e. a map that is a homeomorphism onto its image). Part (i) of the following lemma is a classic result - see for example Buneman [2]. For part (ii) the map described is injective by part (i), and it is bicontinuous by Theorem 2.1 of [8].

Lemma 2.2

(i) If d and d' are tree metrics on X with representations $(\mathcal{T}, w)$ and $(\mathcal{T}', w')$ respectively, then $d = d'$ if and only if $\mathcal{T}$ is isomorphic to $\mathcal{T}'$ and $w = w'$.

(ii) For each X–tree $\mathcal{T}$ the map from $(\mathbb{R}^{>0})^{E(\mathcal{T})}$ to $\mathbb{R}^{\binom{X}{2}}$ defined by $w \mapsto d_{(\mathcal{T},w)}$ is a topological embedding.

3 A cellular structure for the edge-product space

In this section we show that $\mathcal{E}(X)$ has a natural description as a CW-complex based on forests of trees that are vertex-labelled in a particular way. We begin with a definition.

An *X–forest* is a collection $f = \{(A, \mathcal{T}_A) : A \in \pi\}$ where

(i) π forms a partition of X, and

(ii) $\mathcal{T}_A$ is an A–tree for each $A \in \pi$.

Figure 1 illustrates an example of an X–forest.

We let $\mathcal{S}(X)$ denote the set of X–forests. A routine check (see also [12]) shows that $\mathcal{S}(X)$ is of size 3 and 15 when $|X| = 2$ and $|X| = 3$ respectively.

We now describe an order relationship on the set of X–forests which we show below gives a poset that is isomorphic to the face poset of $\mathcal{E}(X)$. Informally this order relation translates as follows - $f \leq g$ if the trees in f can be obtained from the trees in g by collapsing certain edges, and deleting certain other edges, with

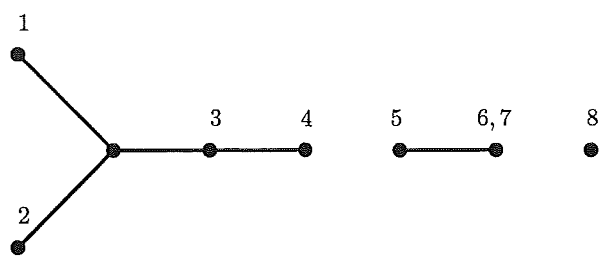

Figure 1: An X-forest for $X = \{1, 2, \ldots, 8\}$, with associated partition $\pi = \{\{1, 2, 3, 4\}, \{5, 6, 7\}, \{8\}\}$.

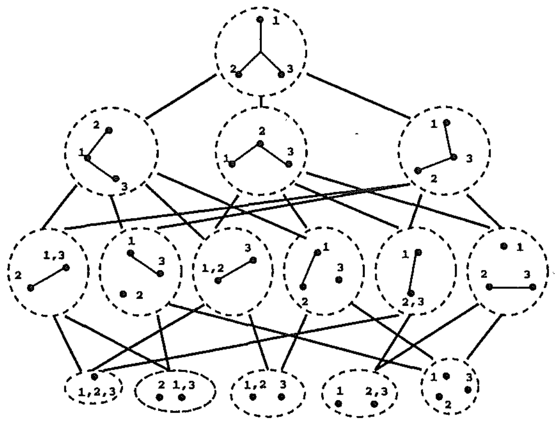

Figure 2: The Tuffley poset $\mathcal{S}(X)$ for $X = \{1, 2, 3\}$.

any resulting unlabelled vertices of degree 2 being suppressed. We now make this more formal using the terminology introduced in Section 2.

Let $f = \{(A, \mathcal{T}_A) : A \in \pi\}$ and $f' = \{(B, \mathcal{T}'_B) : B \in \pi'\}$ be two X–forests. We write $f' \leq f$ precisely if the following two conditions hold.

(O1) The partition π' is a refinement of the partition π.

(O2) If $A = \cup_{B \in J} B$ for some $A \in \pi$ and $J \subseteq \pi'$ then

(i) for all $B \in J$, $\mathcal{T}_A$ displays $\mathcal{T}'_B$, and

(ii) for all $B, C \in J$ with $B \neq C$ there exists $F|G \in \Sigma(\mathcal{T}_A)$ with $B \subseteq F$ and $C \subseteq G$.

The proof of the following lemma is routine.

Lemma 3.1 *$\leq$ is a partial order on $\mathcal{S}(X)$.*

The poset $\mathcal{S}(X)$ was first defined (slightly differently) by Christopher Tuffley [12], and accordingly we call it the *Tuffley poset* on X. In Figure 2 we picture $\mathcal{S}(X)$ for $X = \{1, 2, 3\}$. We now clarify its relationship to $\mathcal{E}(X)$.

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

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

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

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

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

Lemma 3.2 *For each X–forest $f = \{(A,\mathcal{T}_A) : A \in \pi\} \in \mathcal{S}(X)$ and map $\lambda = (\lambda_A : A \in \pi) \in \mathbf{B}(f)$, we have $\psi_f(\lambda) \in \mathcal{E}(X)$.*

Proof. Write $\mathcal{T}_A = (T_A;\phi_A)$. Let T'_A be the tree obtained from T_A by performing the following operation independently on each vertex: If v is a leaf, and $\phi_A^{-1}(v)$ has size $k \geq 2$, or if v is an interior vertex and $\phi_A^{-1}(v)$ has size $k \geq 1$ then make the elements in $\phi_A^{-1}(v)$ leaves by attaching each of them by a new edge to v (which is then regarded as an unlabelled vertex). In this way we obtain a tree T'_A that has leaf set A, and for which each edge of T_A has a corresponding edge of T'_A. Let λ_A be the edge weighting of T'_A that assigns the value $\lambda_A(e)$ to any edge e of T'_A that corresponds to an edge of T_A; otherwise $\lambda_A(e) = 1$. Finally, let T be any tree obtained by joining together the collection of trees $\{T'_A : A \in \pi\}$ by adding edges arbitrarily that have as their endpoints interior vertices in distinct trees from this set. Note that T has leaf set X. Let λ be the edge-weighting of T that agrees with λ_A for any edge in T'_A and that takes the value 0 for any edge that has its endpoint vertices in distinct trees from $\{T'_A : A \in \pi\}$. It is now easily seen that $\psi_f(\lambda)(x,y) = p_{(T,\lambda)}(x,y)$ for all $x,y \in X$ and so $\psi_f(\lambda) \in \mathcal{E}(X)$, as claimed. ∎

We now recall the definition of a finite CW-complex [3]. Suppose we have a Hausdorff topological space Y and a collection B_α of closed balls of various dimensions, together with associated maps $\psi_\alpha : B_\alpha \to Y$ where α ranges over a set A. The sets $o_\alpha = \psi_\alpha(\mathrm{Int}(B_\alpha))$ and $c_\alpha = \psi_\alpha(B_\alpha)$ are called the *open cells* and *closed cells* respectively, corresponding to α. In this setting, points may be regarded as 0-dimensional open cells.

Then Y is a *finite CW-complex* and the collection $\{(B_\alpha,\psi_\alpha) : \alpha \in A\}$ is said to provide a *cell decomposition* of Y if A is finite, and the following three properties hold:

(cw1) $\psi_\alpha|_{\mathrm{Int}(B_\alpha)}$ maps $\mathrm{Int}(B_\alpha)$ homeomorphically onto o_α.

(cw2) Y is the disjoint union of all open cells.

(cw3) $c_\alpha - o_\alpha$ is a union of open cells of lower dimension.

The *face poset* of Y is the collection of closed cells c_α partially ordered by inclusion.

Theorem 3.3 *$\mathcal{E}(X)$ is a finite CW-complex, with cell decomposition $\{(\mathbf{B}(f), \psi_f) : f \in \mathcal{S}(X)\}$. Furthermore, the Tuffley poset $(\mathcal{S}(X), \leq)$ is isomorphic to the face poset of $\mathcal{E}(X)$ under the map that sends f to $\psi_f(\mathbf{B}(f))$.*

Proof. First we note that the number of X–forests is clearly finite. It thus suffices to establish the properties **(cw1)**, **(cw2)** and **(cw3)**.

To establish **(cw1)**, suppose that $f = \{(A, \mathcal{T}_A) : A \in \pi\} \in \mathcal{S}(X)$. For $x, y \in A$, and $\lambda_A \in \mathrm{Int}(B(\mathcal{T}_A))$, we have $p_{(\mathcal{T}_A, \lambda_A)}(x, y) \in (0, 1)$.

By Lemma 2.2(ii) the mapping

$$D_A : (\mathbb{R}^{>0})^{E(\mathcal{T}_A)} \to \mathbb{R}^{\binom{A}{2}}$$

$$w_A \mapsto d_{(\mathcal{T}_A, w_A)}$$

is a topological embedding.

Observe next that the map $(\exp -) : (\mathbb{R}^{>0})^{\binom{X}{2}} \to (0, 1)^{\binom{X}{2}}$ defined by

$$(t_{x,y}) \mapsto \exp(-t_{x,y}),$$

and the map $(-\log) : (0, 1)^{E(\mathcal{T})} \to (\mathbb{R}^{>0})^{E(\mathcal{T})}$ defined by

$$(t_e) \mapsto -\log(t_e),$$

are both homeomorphisms.

Now, if p_A denotes the restriction of p to $\binom{A}{2}$ for $A \in \pi$ then

$$p_A(x, y) = e^{-d_{(\mathcal{T}_A, -\log(\lambda_A))}(x,y)}$$

for all $x, y \in A$. Consequently, the map $\lambda_A \mapsto p_{(\mathcal{T}_A, \lambda_A)}$ is just the composition $(\exp -) \circ D_A \circ (-\log)$, which by the proceeding discussion is an embedding. It follows that the map ψ_f is bicontinuous and one-to-one on $\mathrm{Int}(\mathbf{B}(f))$ which establishes **(cw1)**.

For **(cw2)**, given $p \in \mathcal{E}(X)$, define an associated equivalence relation $\sim_p$ on X as follows: Write $x \sim_p y$ precisely if $p(x,y) \neq 0$. Let π_p denote the equivalence classes of $\sim_p$. Thus, for $x, y \in A \in \pi_p$ we may define $\delta_A : \binom{A}{2} \rightarrow \mathbb{R}^{\geq 0}$ by

$$\delta_A(x,y) = -\log(p(x,y)). \tag{1}$$

Notice that δ_A is a tree metric, and so, by Lemma 2.2(i), δ_A has a unique representation $(\mathcal{T}_A, w_A)$ where $\mathcal{T}_A = (T_A; \phi_A)$ is an A–tree and $w_A : E(\mathcal{T}_A) \rightarrow \mathbb{R}^{>0}$. Consequently, if we let $\lambda_A(e) = \exp(-w_A(e))$ for each edge e of $\mathcal{T}_A$ then $\lambda_A \in \mathrm{Int}(B(\mathcal{T}_A))$ and the restriction of p to $\binom{A}{2}$ is $p_{(\mathcal{T}_A, \lambda_A)}$. Let $f = \{(A, \mathcal{T}_A) : A \in \pi_p\}$ and $\lambda = (\lambda_A : A \in \pi_p)$. Then, $p = \psi_f(\lambda) \in \psi_f(\mathrm{Int}(\mathbf{B}(f))$ and since p determines f uniquely the disjointness property described in **(cw2)** also holds.

For **(cw3)** suppose that $p \in \psi_f(\mathbf{B}(f)) - \psi_f(\mathrm{Int}(\mathbf{B}(f)))$, where $f = \{(A, \mathcal{T}_A), A \in \pi\}$. Let p_A denote the restriction of p to $\binom{A}{2}$. Then, $p_A = p_{(\mathcal{T}_A, \lambda_A)}$ for some $\lambda_A : E(\mathcal{T}_A) \rightarrow [0,1]$. Consider

$$E_A^0 = \{e \in E(\mathcal{T}_A) : \lambda_A(e) = 0\} \text{ and } E_A^1 = \{e \in E(\mathcal{T}_A) : \lambda_A(e) = 1\}.$$

Contract all the edges of $\mathcal{T}_A$ in E_A^1. Also, delete each edge in E_A^0. Finally for any unlabelled vertex v of $\mathcal{T}_A$ that becomes, after this edge contraction and deletion process, incident with just two edges - say e_1, e_2 - we delete v and contract the path e_1, e_2 to obtain a single edge e, say, to which we assign the weight $\lambda_A(e_1)\lambda_A(e_2)$. In this way we obtain an X–forest $f' = \{(B, \mathcal{T}'_B) : B \in \pi'\}$ where $f' \leq f$ and an edge weighting $\lambda'_B : E(\mathcal{T}'_B) \rightarrow (0,1)$ for $B \in \pi'$, such that, for all $x, y \in X$,

$$p(x,y) = \begin{cases} p_{(\mathcal{T}'_B, \lambda'_B)}(x,y), & \text{if } \exists B \in \pi' \text{ with } x, y \in B; \\ 0, & \text{otherwise.} \end{cases}$$

Thus, there is an element in $\mathrm{Int}(\mathbf{B}(f))$ that maps to p under $\psi_{f'}$. Now, since $p \notin \psi_f(\mathrm{Int}(\mathbf{B}(f)))$ it follows that for at least one $A \in \pi$ we have $|E_A^0| + |E_A^1| \geq 1$ and so $f' < f$. This implies $\dim(f') < \dim(f)$ thereby establishing property **(cw2)**.

Finally it remains to show that the association $f \mapsto \psi_f(\mathbf{B}(f))$ preserves the poset structure - that is, $f \leq g$ implies $\psi_f(\mathbf{B}(f)) \subseteq \psi_g(\mathbf{B}(g))$.

Let $p \in \psi_f(\mathbf{B}(f))$. Since $f \leq g$, the trees in f are obtained from the trees in g by collapsing and deleting certain edges. Thus it is easily checked (following the types of arguments used earlier in this proof) that we may assign edge weight 1 to each edge of any tree in g that is collapsed, and edge weight 0 to each edge of any tree in g that is deleted, and assign the remaining edge weights to the trees in g appropriately to obtain an assignment $\lambda \in \mathbf{B}(g)$ such that $p = \psi_g(\lambda)$, as required. This completes the proof of Theorem 3.3. ∎

Notice that the cell decomposition given in Theorem 3.3 induces induces a corresponding cell decomposition of $\mathcal{E}(X,T)$.

4 Structural properties of the Tuffley Poset

In this section we provide an alternative description of the order $\leq$ on $\mathcal{S}(X)$ by explicitly describing the coatoms of any element $f \in \mathcal{S}(X)$. We use this description to show that the Tuffley poset has certain nice structural properties.

Let $f = \{(A, \mathcal{T}_A) : A \in \pi\} \in \mathcal{S}(X)$. Select one of the elements of f – say $(A, \mathcal{T}_A)$ – together with a split $B|C \in \Sigma(\mathcal{T}_A)$. Delete $(A, \mathcal{T}_A)$ from f and replace it by either one of the following:

- $(A, \mathcal{T}'_A)$ where $\Sigma(\mathcal{T}'_A) = \Sigma(\mathcal{T}_A) - \{B|C\}$, an operation that we call *edge contraction (on σ)*;
- the pair $(B, \mathcal{T}_A|B)$ and $(C, \mathcal{T}_A|C)$, an operation that we call *edge deletion (on σ)*.

Given an X–forest, $f = \{(A, \mathcal{T}_A) : A \in \pi\}$ let

$$\Sigma(f) = \cup_{A \in \pi} \Sigma(\mathcal{T}_A)$$

which, in view of Proposition 2.1, we may view as the set of edges in f.

Clearly, for any $\sigma \in \Sigma(f)$ the set obtained by contraction on σ, denoted f/σ, or by edge deletion on σ, denoted $f - \sigma$, results in an X–forest. Furthermore,

$$|\Sigma(f/\sigma)| = |\Sigma(f)| - 1, \tag{2}$$

and

$$|\Sigma(f)| - 3 \leq |\Sigma(f - \sigma)| \leq |\Sigma(f)| - 1. \tag{3}$$

We will say that the edge deletion $f \mapsto f - \sigma$ is *safe* if $|\Sigma(f - \sigma)| = |\Sigma(f)| - 1$.

The following easily checked lemma provides the graph theoretic interpretation of a safe edge deletion, where we say that a vertex in an X–tree is *unsupported* if it is unlabelled and of degree 3.

Lemma 4.1 *For an X-forest f, an edge deletion $f \mapsto f - \sigma$ is safe if and only if neither endpoint of the edge e that corresponds to σ in f 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 following result describes $\mathcal{E}(X)$ in terms of these operations, and establishes some further structural properties. To describe these we recall some further concepts concerning posets (see [3], [4]).

Let $(S, \leq)$ be an arbitrary poset.

- An element $f' \in S$ is a *coatom* of an element $f \in S$ if $f' < f$ and there is no element $g \in S$ satisfying $f' < g < f$.
- For $f, g \in S$ the *interval* between f and g, denoted $[f, g]$ is the set of all elements $h \in S$ satisfying $f \leq h \leq g$.
- $(S, \leq)$ is *pure* if all maximal chains have the same finite length, in which case there exists a *rank function* ρ on S that associates to each element $f \in S$ the length of a maximal chain that has f as its maximum element. The rank of an interval $[f, g]$ is defined as $\rho(g) - \rho(f)$.
- A poset is *thin* if any interval of rank 2 has cardinality four.

Theorem 4.2 *Suppose that X is a finite set and $f, f' \in \mathcal{S}(X)$. Then the following statements hold.*

(i) $f' \leq f$ if and only if f' can be obtained from f 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) f' is a coatom of f if and only if f' can be obtained from f by one elementary operation.

(iii) $\mathcal{S}(X)$ is a pure poset, and for an element $f = \{(A, \mathcal{T}_A) : A \in \pi\}$ of $\mathcal{S}(X)$, its rank, denoted $\rho(f)$ is given by

$$\rho(f) = |\Sigma(f)|.$$

(iv) $\mathcal{S}(X)$ is thin.

(v) The maximal elements of $\mathcal{S}(X)$ are precisely the elements f for which $f = \{(X, \mathcal{T})\}$ and with $|\Sigma(\mathcal{T})| = 2|X| - 3$.

(vi) The minimal elements of $\mathcal{S}(X)$ are precisely the X–forests of the form $f = \{(A, \mathcal{T}_A) : A \in \pi\}$, with $\Sigma(\mathcal{T}_A) = \emptyset$ for all $A \in \pi$.

Proof. (i) Suppose $f \in \mathcal{S}(X)$ and $\sigma \in \Sigma(f)$. Clearly $f/\sigma, f - \sigma \leq f$. It immediately follows that if $f' \in \mathcal{S}(X)$ and f' can be obtained from f by a sequence of contraction and deletion operations, then $f' \leq f$.

Conversely, suppose $f, f' \in \mathcal{S}(X)$ with $f' \leq f$. Let $f = \{(A, \mathcal{T}_A) : A \in \pi\}$ and $f' = \{(B, \mathcal{T}'_B) : B \in \pi'\}$ with $|\pi| \leq |\pi'|$. If $A = \cup_{B \in I_A} B$ for some $I_A \subseteq \pi'$, then $\Sigma(\mathcal{T}'_B) \subseteq \Sigma(\mathcal{T}_A|B)$, since $f' \leq f$. For $B \in I_A$, let

$$\Sigma_{AB} = \{E|F \in \Sigma(\mathcal{T}_A) : E \cap B|F \cap B \notin (\Sigma(\mathcal{T}_B) \cup \{\emptyset|B\})\}$$

and let $\Sigma_A = \bigcup_{B \in I_A} \Sigma_{AB}$, where $A \in \pi$ and $B \in \pi'$. For each $A \in \pi$ contract every split $\sigma \in \Sigma_A$ of $\mathcal{T}_A$ (in any order) to obtain a tree $\mathcal{T}^*_A$ with $\Sigma(\mathcal{T}^*_A) = \Sigma(\mathcal{T}_A) - \Sigma_A$. If $|\pi| = |\pi'|$, then this sequence of contractions yields f'.

So suppose $|\pi| < |\pi'|$. Since $f' \leq f$, for each $B \neq B' \in I_A$, there is some $E|F \in \Sigma(\mathcal{T}_A)$ with $B \subseteq E$ and $B' \subseteq F$. Let Σ^*_A denote the collection of all such splits $E|F$. Then $\Sigma^*_A \subseteq \Sigma(\mathcal{T}^*_A)$. Now, in case the edge of $\mathcal{T}^*_A$ corresponding to some $\sigma \in \Sigma^*_A$ contains an unsupported vertex, contract one of the other edges of $\mathcal{T}^*_A$ that is incident with this vertex. Perform all of these contractions (in any order) for each $A \in \pi$. The deletion of an edge corresponding to any $\sigma \in \Sigma^*_A$ in the resulting X-forest is safe. Delete all of these edges (in any order). The resulting X-forest equals f'. This completes the proof of (i).

(ii) This follows immediately from (i).

(iii) Suppose $f, g \in \mathcal{S}(X)$ with $g < f$. In view of (i), (2) and (3), we have $|\Sigma(f)| - |\Sigma(g)| \geq 1$ and if $|\Sigma(f)| - |\Sigma(g)| > 1$ then there must exist $h \in \mathcal{S}(X)$ with $g < h < f$. Now, suppose $g = h_1 < h_2 < \cdots < h_n = f$ is a maximal chain. Then it follows by our observations that $|\Sigma(h_{i+1})| - |\Sigma(h_i)| = 1$ for all $i = 1, \ldots, n-1$ and $|\Sigma(f)| - |\Sigma(g)| = n$. Thus (ii) holds.

(iv) Suppose $[f, g]$ is an interval in $\mathcal{S}(X)$ with rank 2, so that g can be obtained from f by two elementary operations. Then either both of these operations are contractions or both deletions, in which case it is easy to check that $|[f, g]| = 4$ holds, or one of these operations is a contraction and the other a deletion. For this latter situation it is also easy to check that $|[f, g]| = 4$ holds if the operations are performed on non-incident edges of f, whereas if the edges are incident a straight-forward case-by-case check yields the same conclusion.

(v) This follows as an easy consequence of the fact that a maximal compatible split system on X must have cardinality $2|X| - 3$ (see e.g. [2]).

(vi) If $f = \{(A, \mathcal{T}_A) : A \in \pi\}$ is a minimal element of $\mathcal{S}(X)$, then by (O2)(i) it follows that $\Sigma(\mathcal{T}_A) = \emptyset$ for all $A \in \pi$. Moreover, by (O2)(ii) it follows that any such element of $\mathcal{S}(X)$ is minimal. ∎

Note that part (v) of the previous theorem implies that the X–forests that correspond to the maximal elements of $\mathcal{S}(X)$ are precisely the X–trees $\mathcal{T} = (T;\phi)$ for which ϕ is a bijection from X to the set of leaves of T, and for which each interior vertex of T has degree 3. Moreover, in view of part (vi) there is an obvious bijection between the collection of partitions of X and the minimal elements of $\mathcal{S}(X)$, obtained by associating to the partition π the set $\{(A, \mathcal{T}_A) : A \in \pi\}$ where $\mathcal{T}_A$ is the A–tree consisting of a single vertex labelled by all the elements of A.

We end this section by making some general comments about the existence of upper and lower bounds for an arbitrary collection $\{f_1, f_2, \ldots, f_k\}$ of elements from $\mathcal{S}(X)$. First, even when $k = 2$ there may not exist an upper bound, or a lower bound, to this collection in $\mathcal{S}(X)$. Furthermore, even when upper bounds (respectively lower bounds) exist, there may not be a unique least upper bound (respectively greatest lower bound).

The existence question for upper bounds generalizes a well-known problem in computational biology called the *character compatibility problem* [9]. To understand this we require the following definitions.

- Suppose π is a partition of X, and $\mathcal{T} = (T;\phi)$ is an X–tree. Then π is said to be *convex* on $\mathcal{T}$ if and only if, for all $C, C' \in \pi$ with $C \neq C'$, there exists $A|B \in \Sigma(\mathcal{T})$ such that $C \subseteq A, C' \subseteq B$.
- A collection $\{\pi_1, \pi_2, \ldots, \pi_k\}$ of partitions of X is said to be *compatible* if and only if there exists an X–tree $\mathcal{T}$ so that π_i is convex on $\mathcal{T}$ for all $i \in \{1, 2, \ldots, k\}$.

The relevance of this condition to the Tuffley poset arises by associating each partition π of X to the rank 0 element $\{(A, \mathcal{T}_A) : A \in \pi\}$ of $\mathcal{S}(X)$, where $\Sigma(\mathcal{T}_A) = \emptyset$ for all $A \in \pi$. Furthermore, under this association we have the following result.

Proposition 4.3 *A collection $\{\pi_1, \pi_2, \ldots, \pi_k\}$ of partitions of X is compatible if and only if the set $\{f_1, f_2, \ldots, f_k\}$ of associated rank 0 elements of $\mathcal{S}(X)$ have an upper bound in $\mathcal{S}(X)$.*

Determining whether a collection of partitions is compatible is NP-complete [9], which suggests that it is unlikely that there is a good characterization for when an arbitrary subset of $\mathcal{S}(X)$ has an upper bound. It is not clear if there is a good characterization for when an arbitrary subset of $\mathcal{S}(X)$ has a lower bound, although it is possible to give reasonable characterizations for when a *pair* of elements in $\mathcal{S}(X)$ have an upper (or a lower) bound. We will describe these characterizations elsewhere when we consider further structural properties of the Tuffley poset including its Möbius function.

5 Topology of the edge-product space

In this section we consider topological properties of $\mathcal{E}(X)$ and $\mathcal{E}(X,T)$. Clearly, both of these spaces are compact since they are the continuous image of compact spaces. We now show that $\mathcal{E}(X)$ and $\mathcal{E}(X,T)$ are also contractible (formally, the identity map on each of these spaces is homotopic to a constant map) and so can be continuously 'shrunk' to a point. In particular, it follows that $\mathcal{E}(X)$ and $\mathcal{E}(X,T)$ are connected.

Proposition 5.1

(i) *For every* $0 \leq \beta \leq 1$, *and* $p \in \mathcal{E}(X)$,

$$\beta \cdot p \in \mathcal{E}(X).$$

(ii) $\mathcal{E}(X)$ *and* $\mathcal{E}(X,T)$ *are contractible.*

Proof. For part (i) suppose the $p \in \mathcal{E}(X)$, i.e. $p = p_{(T,\lambda)}$ for some tree $T = (V, E)$ with leaf set X and $\lambda : E \to [0,1]$. For $\beta \in [0,1]$ let $\lambda_\beta : E \to [0,1]$ be defined as follows:

$$\lambda_\beta(e) = \begin{cases} \lambda(e), & \text{if } e \text{ is an interior edge;} \\ \lambda(e)\sqrt{\beta}, & \text{otherwise.} \end{cases}$$

Then it is easily checked that

$$\beta \cdot p = p_{(T,\lambda_\beta)}$$

and so $\beta \cdot p \in \mathcal{E}(X)$, which establishes (i).

For part (ii) the map $H : \mathcal{E}(X) \times [0,1] \to \mathcal{E}(X)$

$$p \mapsto \beta \cdot p$$

is a homotopy from a constant map to the identity map on $\mathcal{E}(X)$ and so $\mathcal{E}(X)$ is contractible. Furthermore, H restricts to $\mathcal{E}(X,T)$ to provide a homotopy from the constant map to the identity map on $\mathcal{E}(X,T)$. This completes the proof. ∎

It is worth noting for any $x \in X$, there is a natural embedding $e : \mathcal{E}(X - \{x\}) \to \mathcal{E}(X)$ and a natural surjection $f : \mathcal{E}(X) \to \mathcal{E}(X - \{x\})$ such that $f \circ e$ is the identity map on $\mathcal{E}(X - \{x\})$; thus $\mathcal{E}(X - \{x\})$ is a *retract* of $\mathcal{E}(X)$.

The map e is defined as follows: For any $p \in \mathcal{E}(X - \{x\})$, let $e(p) : \binom{X}{2} \to [0,1]$ satisfy

$$e(p)(y,y') = \begin{cases} p(y,y'), & \text{for } \{y,y'\} \subseteq X - \{x\}; \\ 0, & \text{otherwise.} \end{cases}$$

Let $f(p)$ be the restriction of p to $\binom{X-\{x\}}{2}$. Then it is straight-forward to verify that $e(p) \in \mathcal{E}(X)$, that $f(p) \in \mathcal{E}(X - \{x\})$ and that the maps e and f are continuous with $f \circ e$ the identity map on $\mathcal{E}(X - \{x\})$.

Describing $\mathcal{E}(X)$ up to homeomorphism appears (as might be expected) to be a somewhat harder problem. To understand this problem it would be useful to relate topological properties of $\mathcal{E}(X)$ with those of the geometric realization $\|\mathcal{S}(X)\|$ of the order (simplicial) complex of $\mathcal{S}(X)$ [3, Section 9.3]. The following proposition, which can be regarded as a discrete analogue of Proposition 5.1, implies that $\mathcal{E}(X)$ and $\|\mathcal{S}(X)\|$ are at least homotopy equivalent.

Proposition 5.2 $||\mathcal{S}(X)||$ *is contractible.*

Proof. Consider the following two maps

$$m_1, m_2 : \mathcal{S}(X) \to \mathcal{S}(X)$$

defined as follows. For $f = \{(A, \mathcal{T}_A) : A \in \pi\} \in \mathcal{S}(X)$, $m_1(f)$ replaces $(A, \mathcal{T}_A)$ by $(A, \mathcal{T}_A^0)$, where $\Sigma(\mathcal{T}_A^0) = \Sigma(\mathcal{T}_A) \cup \{\{a\}|A - \{a\} : a \in A\}$.

The map $m_2(f)$ replaces $(A, \mathcal{T}_A)$ by $(A, \mathcal{T}_A')$, where $\Sigma(\mathcal{T}_A') = \emptyset$.

Then, for all $f \in \mathcal{S}(X)$ we have

$$f \leq m_1(f), \text{ and } m_2(f) \leq f.$$

Furthermore, $m_2 \circ m_1(f) = f_0$ where $f_0 = \{(\{x\}, \mathcal{T}_x) : x \in X\}$, where $\Sigma(\mathcal{T}_x) = \emptyset$, which is a minimal element of $\mathcal{S}(X)$. From [3, Corollary 10.12], it follows that $||\mathcal{S}(X)||$ is homotopic to $||m_1(\mathcal{S}(X))||$, which in turn is homotopic to $||m_2 \circ m_1(\mathcal{S}(X))||$. However this last space consists of a single element, and so $||\mathcal{S}(X)||$ is contractible as claimed. ∎

If the cell decomposition $\mathcal{C}_X = \{(\mathbf{B}(f), \psi_f) : f \in \mathcal{S}(X)\}$ of $\mathcal{E}(X)$ given in Theorem 3.3 were *regular*, that is, for each $f \in \mathcal{S}(X)$ ψ_f maps $\mathbf{B}(f)$ homeomorphically onto its image, then it would follow that $\mathcal{E}(X)$ is homeomorphic to $\|\mathcal{S}(X)\|$ (cf.[3, 12.4 (ii)]). It is straight-forward to check that $\mathcal{C}_X$ is regular when $|X| \leq 3$, and, using topological arguments, that the cell decomposition induced by $\mathcal{C}_X$ on $\mathcal{E}(X,T)$ is regular when T is a tree having exactly one interior vertex (Bill Baritompa, personal communication). Moreover, it can be shown that $\mathcal{C}_X$ is regular if $\mathcal{S}(X)$ has a *recursive coatom ordering*, and that such orderings exist for $\mathcal{S}(X)$ when $|X| \leq 4$.

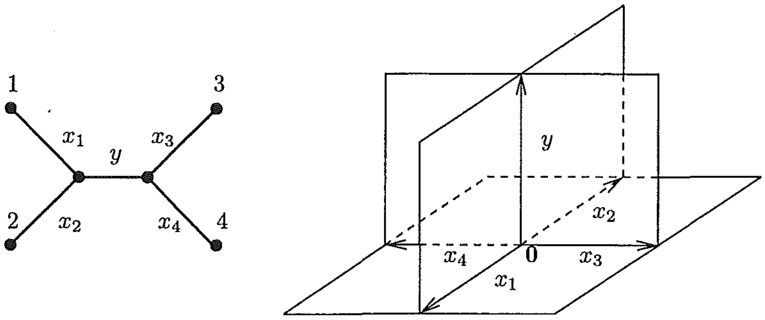

(a) (b)

Figure 3: (a) A tree $T \in \mathcal{T}(X)$ for $X = \{1, 2, 3, 4\}$; (b) the fiber of the map Λ_T over $\mathbf{0}$.

6 Structure of the fibers of Λ_T

We conclude with a description of the topological and combinatorial structure of the fibers of the map Λ_T over points in $\mathcal{E}(X,T)$ and show that they have attractive topological and combinatorial properties. Part of our motivation for investigating these fibers is to obtain a better understanding of the topology of $\mathcal{E}(X)$.

Consider the map Λ_T from $[0,1]^{E(T)}$ to $\mathcal{E}(X,T)$ defined by $\lambda \mapsto p_{(T,\lambda)}$. Figure 3 illustrates the 2-dimensional fiber $\Lambda_T^{-1}(\mathbf{0})$ of Λ_T over $\mathbf{0}$ for a tree $T \in \mathcal{T}(X)$ where $X = \{1, 2, 3, 4\}$.

In order to describe the structure of the fibers of the map Λ_T in general we need to introduce a number of further concepts and results.

Suppose $T = (V, E)$ is a tree with X a subset of its leaf set. A set $I \subseteq E$ of edges is said to be an *isolating set for X in T* if the graph $(V, E - I)$ has no two elements of X in the same component. An isolating set I of X in T is said to be *minimal* if no proper subset of I is an isolating set for X in T.

Proposition 6.1 *Let $T = (V, E)$ be a tree with leaf set containing X. Then any minimal isolating set for X in T has cardinality $|X| - 1$.*

Proof. We use induction on $|X|$. The result clearly holds for $|X| \leq 2$. Suppose I is a minimal isolating set for X in T, where $|X| > 2$. Select a leaf vertex $l \in X$, and let $e = \{v, l\}$ be the edge of T incident with l. Let T' denote the tree obtained from T by deleting leaf l and edge e.

There are two possibilities: (i) $e \in I$, and (ii) $e \notin I$. In case (i) let $I' = I - \{e\}$. Then I' is an isolating set for $X - \{l\}$ in T'. Furthermore, I' is a minimal isolating set for $X - \{l\}$ in T', for if a proper subset I'' of I' were an isolating subset for $X - \{l\}$ in T' then $I'' \cup \{e\}$ would be an isolating subset for X in T, which contradicts the minimality assumption on I. Thus, by the inductive hypothesis, $|I'| = |X - \{l\}| - 1$ and so $|I| = |X| - 1$ which establishes the inductive step in case (i).

Now consider case (ii). Then for each element $x \in X - \{l\}$ the path from v to x includes at least one edge in I. Select any element $x_0 \in X - \{l\}$ and let e_0 denote the first edge on the path from v to x_0 that lies in I. Then $I - \{e_0\}$ is an isolating set for $X - \{l\}$ in T', and as in case (i) it is also easily verified that $I - \{e_0\}$ is a minimal isolating set for $X - \{l\}$ in T'. Thus, by the inductive hypothesis we have $|I - \{e_0\}| = |X - \{l\}| - 1$, and so $|I| = |X| - 1$, thereby establishing the inductive step in case (ii). This completes the proof. ∎

We now describe the structure of the fibre of Λ_T over the element $\mathbf{0} \in \mathcal{E}(X,T)$.

Proposition 6.2 *Let T be a tree with leaf set X. Then $\Lambda_T^{-1}(\mathbf{0})$ is a contractible regular cell complex whose dimension is equal to the number of interior vertices of T. Furthermore, this space is homeomorphic to the (geometric realization of the) poset of isolating sets for X in T ordered by reverse inclusion.*

Proof. For each isolating set I for X in $T = (V,E)$ let

$$\Lambda_I := \{\lambda : E \to [0,1] : \lambda(e) = 0 \text{ for all } e \in I\}.$$

Note that Λ_I is a closed cell of dimension $|E| - |I|$, which takes the maximum value $|E| - |X| + 1$ by Proposition 6.1. Furthermore in any tree with leaf set X, $|E| - |X| + 1$ is the number of interior vertices of T.

Let $\Lambda = \Lambda_T^{-1}(\mathbf{0})$. Now, $\lambda \in \Lambda$ if and only if $\{e \in E : \lambda(e) = 0\}$ is an isolating set for X in T. Consequently, $\Lambda = \cup_I \Lambda_I$. Furthermore, for isolating sets I, I' we have $\Lambda_I \cap \Lambda_I' = \Lambda_{I \cup I'}$, and for $I \subseteq I'$ we have $\Lambda_I' \subseteq \Lambda_I$. It follows that Λ is a regular cell complex. The last statement in the theorem follows immediately from [3, 12.4 (ii)].

To show that Λ is contractible it suffices to note that the map $H : \Lambda \times [0,1] \to \Lambda$ defined by $H(\lambda, t)(e) = (1-t)\lambda(e)$ is a homotopy from the identity map to a constant map on Λ. This completes the proof. ∎

We now extend Proposition 6.2 to describe the topology of the fibre of Λ_T over an arbitrary point in $\mathcal{E}(X,T)$. We begin with a useful lemma.

Lemma 6.3 *For any fixed value $\theta \in (0,1)$ consider the following subset Λ_θ, of $[0,1]^k$ defined by:*

$$\Lambda_\theta = \{(\lambda_1, \ldots, \lambda_k) \in [0,1]^k : \prod_{i=1}^{k} \lambda_i = \theta\}.$$

Then, for each $k \geq 1$, Λ_θ is homeomorphic to a closed $(k-1)$-dimensional ball.

Proof. First note that $\theta > 0$ implies that $\lambda_i > 0$ for all i, for any vector $\lambda \in \Lambda_\theta$. We may therefore apply the map $t \mapsto \frac{\log(t)}{\log(\theta)}$ to each component of each element of Λ_θ to obtain a homeomorphism from Λ_θ onto

$$\{(x_1, x_2, \ldots, x_k) \subseteq (\mathbb{R}^{\geq 0})^k : \sum_{i=1}^{k} x_i = 1\},$$

which is the $(k-1)$–dimensional simplex. ∎

Now, let T be a tree with leaf set X, let e be an edge of T, and suppose that $p \in \mathcal{E}(X,T)$. We say that e is *isolated relative to* p if $p(x,y) = 0$ for all pairs $x, y \in X$ for which the path in T connecting x and y contains e. Let $I(p)$ (respectively $NI(p)$) denote the sets of edges of T that are isolated (respectively not isolated) relative to p.

We now define relations on $I(p)$ and $NI(p)$. For two edges e, e' with either $\{e, e'\} \subseteq I(p)$ or $\{e, e'\} \subseteq NI(p)$ write $e \sim e'$ if either $e = e'$ or e and e' are adjacent and all the edges that are incident with both e and e' are isolated relative to p.

Let us now take the transitive closure of $\sim_p$ restricted to pairs of edges from $I(p)$ to form an equivalence relation on $I(p)$. Similarly, take the transitive closure of $\sim_p$ restricted to pairs of edges from $NI(p)$ to form an equivalence relation on $NI(p)$. We will let $C(p)$ denote the equivalence classes of $NI(p)$. We illustrate these concepts with an example in Figure 4.

Lemma 6.4 *Let $T = (V, E)$ be a tree with leaf set X, and suppose $p \in \mathcal{E}(X,T)$.*

(i) The edges in any equivalence class of $I(p)$ form a connected subgraph of T.

(ii) The edges in any equivalence class C of $NI(p)$ form a path in T and $\alpha_p = \prod_{e \in C} \lambda(e)$ is uniquely determined by p.

(iii) For any $\lambda' : E(T) \to [0,1]$ let $p' = p_{(T,\lambda')}$. If $I(p') = I(p)$ and $NI(p') = NI(p)$ and $\alpha_p = \prod_{e \in C} \lambda'(e)$ for all equivalence classes C of $NI(p)$ we have $p = p'$.

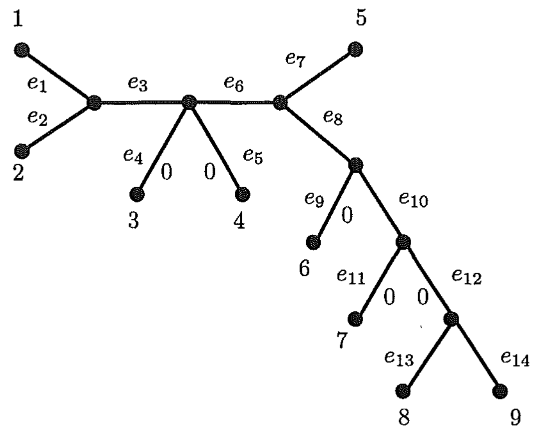

Figure 4: For the tree T pictured, associate a map $\lambda : E(T) \longrightarrow [0,1]$ which maps each edge weighted 0 to 0 and any other edge to an element of $(0,1)$. For the associated map $p = p_{(T,\lambda)}$ in $\mathcal{E}(X,T)$, we have $C(p) = \{\{e_1\},\{e_2\},\{e_3,e_6,e_7\},\{e_{13},e_{14}\}\}$.

Proof. Part (i) and the first part of part (ii) are clear from the definition of the equivalence relations. For an equivalence class C of $NI(p)$ let u and v denote the endpoints of the corresponding path. Then we may select leaves x,y and w,z such that there are edge-disjoint paths from u to x and u to y, and edge-disjoint paths from v to w and v to z, and such that $\lambda(e) \neq 0$ for each edge e on each of these four paths, and for each edge e on the path between u and v (if u is a leaf of T we take $x = y = u$, while if v is a leaf of T we take $w = z = v$). Then $C_p = \sqrt{\frac{p(x,w)p(y,z)}{p(x,y)p(w,z)}}$. The proof of part (iii) is straight-forward, and we leave the details to the reader. ∎

Note that the edges of T are now partitioned by p into two types - isolated edges, which form subtrees, and non-isolated edges which form paths. Let $n_1(p), n_2(p), \ldots,$ denote the number of interior vertices of the subtrees of T induced by the equivalence classes of isolated edges.

Proposition 6.5 *For any point $p \in \mathcal{E}(X,T)$, $\Lambda_T^{-1}(p)$ is a contractible regular cell complex of dimension $\sum_{i\geq 1} n_i + \sum_{A\in C(p)}(|A|-1)$.*

Proof. By Lemma 6.4 $\Lambda_T^{-1}(p)$ is precisely the collection of those $\lambda : E \longrightarrow [0,1]$ for which

- $\lambda(e_1)\dots\lambda(e_r) = \alpha_p$ for any equivalence class C of $NI(p)$, where α_p is the value described by Lemma 6.4(ii).
- For each equivalence class E' of $I(p)$, if we regard the resulting subtree $T' = (V, E')$ of T as having leaf set U, and let λ' denote the restriction of λ to E' then $p_{(T',\lambda')}(x, y) = 0$ for all $x, y \in U$.

It follows that $\Lambda_T^{-1}(p)$ is homeomorphic to the Cartesian product of cells of dimension $|A| - 1$ for each equivalence class A from $C(p)$ (by Lemma 6.3), and fibres over zero of subtrees of T. These latter spaces are regular cell complexes whose dimension is precisely the number of interior vertices of the subtree by Proposition 6.2. Consequently, $\Lambda_T^{-1}(p)$ is a regular cell complex whose dimension is as claimed.

To establish the contractability claim we construct a homotopy by considering the two types of edges, as follows. For each isolated edge e let

$$H(\lambda(e), t) = (1 - t)\lambda(e).$$

For an equivalence class $\{e_1, \dots, e_r\} \in C(p)$ let

$$H(\lambda(e), t) = \lambda(e)^{1-t} \left(\prod_{i=1}^{r} \lambda(e) \right)^{t/r}.$$

Then, as t varies from 0 to 1, H provides a homotopy from the identity map to a constant map on $\Lambda_T^{-1}(p)$ and so this space is contractible, as claimed. This completes the proof. ∎

To illustrate this last proposition, the space $\Lambda_T^{-1}(p)$ for the element p described in Figure 4 is homeomorphic to the Cartesian product $[0, 1] \times [0, 1]^2 \times F$, where F is the space picture in Figure 3(b).

References

[1] L. J. Billera, S. P. Holmes and K. Vogtmann, Geometry of the space of phylogenetic trees, *Advances in Applied Math.* **27**(4) (2001), 733-767.

[2] P. Buneman, The recovery of trees from measures of dissimilarity, *in* "Mathematics in the Archaeological and Historical Sciences", pp. 387-395, Edinburgh University Press, 1971.

[3] A. Björner, Topological methods, *in* "Handbook of Combinatorics", pp. 1819-1872, Elsevier Science, 1995.

[4] A. Björner, M. Las Vergnas, B. Sturmfels, N. White, and G. Ziegler, "Oriented Matroids", Cambridge University Press, 1993.

[5] J. T. Chang, Full reconstruction of Markov models on evolutionary trees: identifiability and consistency, *Mathematical Biosciences* **137** (1996*a*), 51-73.

[6] J. Kim, Slicing hyperdimensional oranges: The geometry of phylogenetic estimation *Molecular Phylogenetics and Evolution* **17** (2000), 58-75.

[7] D. Kozlov, Complexes of directed trees, *J. Comb. Theory Ser. A* **88** (1999), 112-122.

[8] V. Moulton and M. Steel, Retractions of finite distance functions onto tree metrics *Discrete Applied Mathematics* **91** (1999), 215-233.

[9] M. Steel, The complexity of reconstructing trees from qualitative characters and subtrees *Journal of Classification* **9** (1992), 91-116.

[10] M. Steel, Recovering a tree from the leaf colourations it generates under Markov model *Applied Mathematical Letters* **7** (1994), 19-23.

[11] H. Trappmann and G. Ziegler, Shellability of complexes of trees, *J. Comb. Theory, Ser. A* **82** (1998), 179-185.

[12] C. Tuffley, " Trees and Ps and things that sneeze. Markov process models of site substitution" MSc thesis, University of Canterbury, Christchurch, New Zealand, 1997.