11-2015

Unfolding Trees & Symmetrically-Associated Graphs

Kaelyn Flowerday
Portland State University

Unfolding Trees & Symmetrically-Associated Graphs

A Thesis
Presented to the
Department of Mathematics and Statistics
Portland State University

In Partial Fulfillment of the
Requirements for the Degree of
Bachelor of Science in Mathematics
with Departmental Honors

Kaelyn Flowerday
November 2015

Unfolding Trees & Symmetrically-Associated Graphs

Kaelyn Flowerday

ABSTRACT

An unfolding tree is an object reflecting the connectivity properties of a vector-labelled graph. First introduced in the context of theoretical computer science as a way of describing information flow in a neural net model of graph-structured data, unfolding trees have remained unexplored within graph theory. They give rise to an equivalence relation on the vertices of a graph, one which describes the connective environments of vertices but is not reducible to automorphism group orbits. This thesis formalizes unfolding trees and investigates their properties along with the implications of this vertex relation. This leads to the graph property of symmetric-association; graphs with this property have predictably-behaved unfolding trees. Symmetric-association is presented as a generalization of k-regularity, culminating in a Havel-Hakimi type result featuring a graph transformation that preserves unfolding trees.

Contents

List of Figures

Glossary of Notation

Notation	**Meaning**
$\mathbb{Z}$	Set of integers
$\mathbb{N}$	Set of non-negative integers
$\mathbb{Z}_+$	Set of positive integers
$[N]$	Interval of integers from 1 to N
$x := y$	Define x to be y
$f \restriction_S$	Restriction of function f to domain S
$\langle x_1, \ldots, x_m \rangle$	Ordered m-tuple
$u \sim_G v$	Vertices u, v are adjacent in G
$u \risingdotseq v$	Vertices u, v are unfolding equivalent
$u \doteq v$	Vertices u, v are similar
$\mathcal{N}_G(v)$	Neighborhood of vertex v in G
$d_G(v)$	Degree of vertex v in G
$\Delta(G)$	Maximum vertex degree in G
$\Lambda(G)$	Number of distinct labels on vertices of G
$G \cong H$	G is label isomorphic to H
$G \leq H$ $(G < H)$	G is a (proper) subgraph of H
$G[S]$	Subgraph of G induced by vertex subset S
$\ell(W)$	Length of walk W
$W[m]$	Length m initial segment of walk W
$W[\![m]\!]$	$(m+1)$th vertex in walk W
$T_v^d[\![m]\!]$	Set of vertices of level m in unfolding tree T_v^d

Chapter 1

Introduction

1.1 Origins & Motivation

UNFOLDING trees were first presented by Scarselli et al. in a paper on their Graph Neural Network (GNN) model [3]. The GNN model is a neural net architecture designed to predictively model properties of graph-structured data, such as molecular structures or a set of interlinked web pages. A distinguishing feature of the GNN model is that each vertex of the input graph is implemented as a computational unit, the output of which depends on the states of adjacent vertices. The other parameters determining the units' outputs are real vector labels on each of the vertices and edges of the input graph.

In order to prove that GNNs can approximate almost all vector-valued functions on graphs, the authors required a mathematical structure that could capture the dependence of one vertex's output on other vertices over time. To that end, they introduced unfolding trees, which encapsulate the routes between a given vertex and the rest of the graph and therefore the flow of information across these routes over multiple time-steps. The article then deployed unfolding trees as machinery in its proofs, showing that a vertex's unfolding tree could be encoded and then decoded by the GNN. From this they concluded that the architecture can approximate any vector-valued function on a graph domain, so long as the function preserves unfolding equivalence (a relation between vertices that we will cover at length in the present work).

Despite their potential interest for graph theory, unfolding trees do not appear outside of the theoretical computer science context of GNNs. Yet they seemed to us an interesting research subject for several reasons. As structures that reflect graph connectivity from the local viewpoint of a vertex—and that cannot be directly reduced to an extant graph-theoretic concept—unfolding trees are inherently intriguing. In addition to this, the original formulation in [3] lacked the rigor proper to mathematics. And, of course, the dearth of previous research promised fertile mathematical ground. For these reasons, unfolding trees served as the initial springboard for this work. From there, the process of investigating them led to our defining the class of symmetrically-associated graphs, which are significant not only for the properties of their unfolding trees, but possibly to graph theory more generally.

1.2 Thesis Structure

The foundations for discussing unfolding trees and symmetrically-associated graphs are laid down in Chapter 2, which contains the definitions of labelled graphs and of basic features related to these. Our definition is distinct from the variety of labelled graphs usually found in graph theory, wherein each of a graph's vertices receives a unique integer label [1]. Instead, we deal with labelled graphs whose vertices can be labelled with any vector of a predetermined length; in fact, most of our examples of labelled graphs involve vertices with non-unique labels.

Next, in Chapter 3 we define unfolding trees, as well as defining salient parts of their structure and proving some useful basic results about them. There we also discuss unfolding equivalence, a relation between vertices that holds when their unfolding trees are isomorphic. We then show in Chapter 4 that if two vertices are similar (*i.e.* are contained in the same orbit of their graph's automorphism group action) then they are unfolding equivalent. Following this, counterexamples of the converse are provided; the falsity of the converse is of interest because it means that unfolding equivalence is not merely a cryptomorphic re-formulation of similarity, but instead indicates a broader connectivity-related kinship between vertices.

Chapter 5 introduces symmetrically-associated graphs—which can be seen as a generalization of k-regular graphs—and sets out results that highlight the well-behaved patterns in their unfolding trees. Pursuing this, Chapter 6 considers an edge-swap graph transformation, called an n-switch, that preserves both symmetric-association and unfolding trees. The chapter concludes with a Havel-Hakimi type result, which states that two symmetrically-associated graphs are equivalent in terms of symmetric-association if and only if there exists a finite sequence of n-switches that transforms one graph into the other. Finally, Chapter 7 lists unexplored topics related to unfolding trees and symmetrically-associated graphs which we feel are promising.

Chapter 2

Labelled Graphs

2.1 Definitions & Basic Properties

To begin, we provide a definition of what shall be our primary object of study. This structure consists of a graph the vertices of which are assigned real vector labels.* For sake of simplicity, we consider only simple, finite, undirected graphs here; however, it should be noted that our formulation admits natural generalizations to more exotic graph structures. Conversely, labelled graphs can always be reduced to underlying unlabelled graphs, and so our results also hold in that context (unlabelled versions of important definitions and results can be found in Appendix B).

Definition 2.1.1: A **labelled graph** is an ordered triple $\langle V, E, L \rangle$, where the **vertex set** V is a collection of arbitrary elements ('vertices'), and the **edge set** E consists of unordered pairs of elements from V ('edges'). The **labeling** L is a function $L : V \to \mathbb{R}^n$ $(n \in \mathbb{Z}_+)$ which assigns a real vector ('label') to each of the vertices of G. ■

The rest of this section involves basic definitions and notation related to labelled graphs. To avoid clutter, only definitions and notations significantly different from those found in (unlabelled) graph theory are given here. If the reader is unfamiliar with elementary graph theory, we refer them to Appendix A.

* The labelled graphs presented in Scarselli et al. also included labels on the edges. Though the definitions and results herein could be extended to graphs with labelled edges, we found that their inclusion was notationally cumbersome and detracted from conceptual clarity. As such, we have omitted edge labels.

Notation 2.1.2: Unless otherwise specified, we use V, E, and L as our default symbols when discussing a generic graph G. When multiple graphs are in play, we use V_G, E_G, L_G or else $V(G)$, $E(G)$, $L(G)$.

The notation $u \sim_G v$ indicates that vertices u, v are adjacent in G. We use the symbol $\Lambda(G)$ for the total number of distinct labels on the vertex set of G (that is, $|L[V]|$). When there is no risk of ambiguity we abbreviate the notation for adjacency ($u \sim_G v$), neighborhood ($\mathcal{N}_G(v)$), and degree ($d_G(v)$) by eliding the subscript G. Similarly, we shorten $\Delta(G)$ and $\Lambda(G)$ to Δ and Λ, respectively. ■

Because any graph we treat has a finite vertex set, it immediately follows that there are only finitely many labels on the vertices, and so we simplify matters by arbitrarily enumerating them as l_i, where $1 \leq i \leq \Lambda$.

In order to define unfolding trees in Chapter 3, we must first introduce some machinery. A walk on a graph can be thought of as a trip across the vertices that starts at some vertex and travels along edges in order to visit vertices in a certain sequence, eventually stopping at some vertex. A path is a special type of walk in which no vertex is visited more than once.

Definition 2.1.3: Let $G = \langle V, E, L \rangle$ be a labelled graph. An ordered tuple $W = \langle v_0, v_1, \ldots, v_m \rangle \in V^m$ of vertices of G is called a **walk** if $v_i \sim_G v_{i+1}$ for all $1 \leq i < m$. A walk $P := \langle v_0, v_1, \ldots, v_m \rangle$ is a **path** if $i \neq j \Rightarrow v_i \neq v_j$ for every $i, j \in [0, m]$. We say that a walk is of **length** m when it consists of $m + 1$ vertices, *i.e.* when it traverses m edges. The length of W is written $\ell(W)$. ■

It is worth noting that this definition allows for walks of length 0, which are simply a single vertex. This should be kept in mind, for such single-vertex walks will arise later. The following is notation that will facilitate our work with labelled graphs and walks.

Notation 2.1.4: Let $W = \langle w_0, w_1, \ldots, w_m \rangle$ be a walk in some graph, and let j, k be integers where $1 \leq j \leq m$, $0 \leq k < m$. We write $W[j]$ to denote $\langle w_0, w_1, \ldots, w_j \rangle$ and $W[\![k]\!]$ to denote w_k. ■

2.2 Labelled Graph Isomorphisms

With the definition of a labelled graph established, we may define a stronger, label-preserving form of graph isomorphism.

Definition 2.2.1: Let $G = \langle V, E, L\rangle$ and $G' = \langle V', E', L'\rangle$ be labelled graphs. A bijection $\phi : V \rightarrowtail\!\!\!\!\twoheadrightarrow V'$ is a **labelled graph isomorphism** (or simply 'label isomorphism') if the following hold $\forall u, v \in V$:

(1) $u \sim_G v \;\Leftrightarrow\; \phi(u) \sim_{G'} \phi(v)$;

(2) $L(u) = L'(\phi(u))$.

When a labelled graph isomorphism exists, we say that G and G' are label isomorphic, writing $G \cong G'$. In the special case where $G = G'$ we call ϕ a **labelled graph automorphism**. ■

Figure 2.1 gives a pair of labelled graphs, $G_1 \coloneqq \langle V_1, E_1, L_1\rangle$ and $G_2 \coloneqq \langle V_2, E_2, L_2\rangle$, which are label isomorphic. To see why this is the case, consider the mapping

$$\begin{aligned} \phi : V_1 &\rightarrowtail\!\!\!\!\twoheadrightarrow V_2 \\ u_1 &\longmapsto v_1, \\ u_2 &\longmapsto v_3, \\ u_3 &\longmapsto v_2, \\ u_4 &\longmapsto v_5, \\ u_5 &\longmapsto v_4. \end{aligned}$$

A cursory examination confirms that ϕ is a label isomorphism. For example, checking part (1) of the definition for a pair of vertices, we see that $u_2 \sim_{G_1} u_5$ and thus

$$\phi(u_2) = v_3 \sim_{G_2} v_4 = \phi(u_5).$$

Similarly for (2), we have that

$$L_1(u_2) = b = L_2(v_3) = L_2(\phi(u_2))$$

as expected.

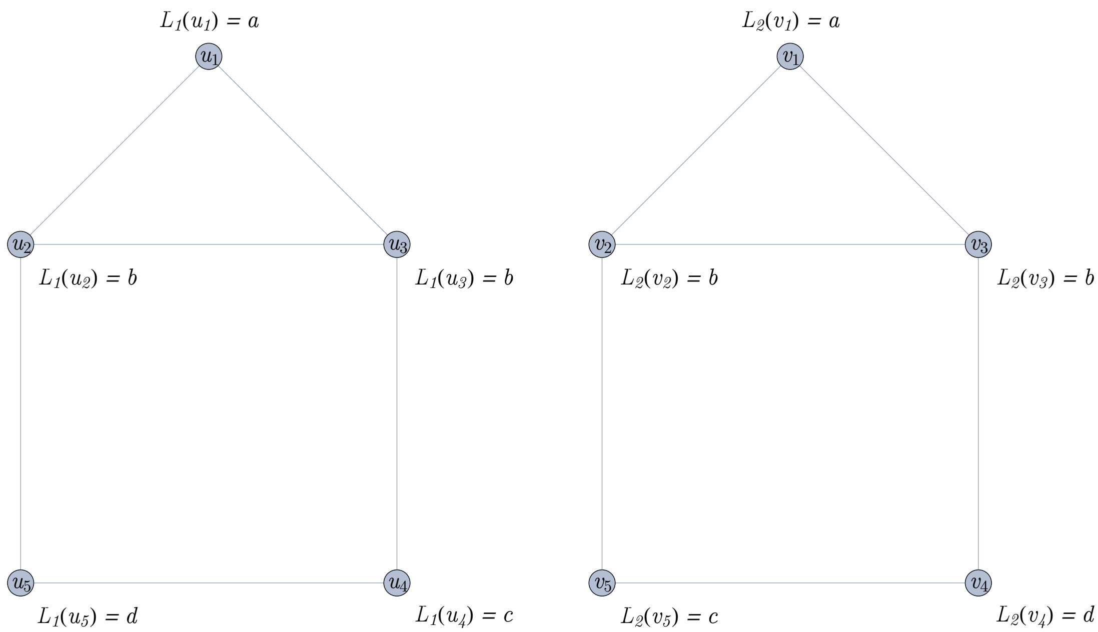

(a) A labelled graph, G_1, with vertex set $V_1 \coloneqq \{u_1, \dots, u_5\}$ and labeling L_1.

(b) Another labelled graph, G_2, with vertex set $V_2 \coloneqq \{v_1, \dots, v_5\}$ and labeling L_2.

Figure 2.1: Labelled versions of the house graph. The two graphs are label isomorphic.

As is the case with the various types of isomorphisms from other areas of mathematics, labelled graph isomorphisms always have inverse functions that are also label isomorphisms. This fact will be of use later, so we prove it here as a lemma.

Lemma 2.2.2: Let $G \coloneqq \langle V, E, L \rangle$, $G' \coloneqq \langle V', E', L' \rangle$ be labelled graphs for which there exists a labelled graph isomorphism $\phi : V \rightarrowtail\!\!\!\!\twoheadrightarrow V'$. Then the inverse function $\phi^{-1} : V' \rightarrowtail\!\!\!\!\twoheadrightarrow V$ exists and is also a label isomorphism.

Proof: To begin, recall that since ϕ is a bijection the inverse ϕ^{-1} exists and is also bijective. Thus, taking arbitrary $v' \in V'$ we know $\exists v \coloneqq \phi^{-1}(v') \in V$.

Let $u', v' \in V'$ be adjacent vertices and call their images under ϕ^{-1} by u, v, respectively. Since ϕ is a label isomorphism we immediately see that $\phi(u) = u' \sim_{G'} v' = \phi(v)$

implies $u \sim_G v$.

Now, for any $v' \in V$ and $v := \phi^{-1}(v')$ we have

$$L'(v') = L'\big(\phi(\phi^{-1}(v'))\big) = L(\phi^{-1}(v')) = L(v),$$

which is to say that ϕ^{-1} preserves labels. □

2.3 Vertex Similarity

THE language of labelled graph isomorphisms allows us to express a relationship on the vertex set of a labelled graph that, to put it roughly, holds between two vertices when they are interchangeable in terms of both graphic structure and labels.

Definition 2.3.1: Let G be a labelled graph, and let $u, v \in V$. We say that u and v are **similar** if $\exists \phi : V \rightarrowtail\!\!\!\!\twoheadrightarrow V$ a label automorphism such that $\phi(u) = v$. When this is the case, we write $u \doteq v$. ■

For example, in the graph G_1 from 2.1 (a), the vertices u_2, u_3 are similar—the label automorphism that makes this the case is the mapping switching u_2 and u_5 with u_3 and u_4, respectively, while leaving u_1 fixed. Of course, since the identity map on $V(G)$ is a label automorphism for any G, $v \doteq v$ always holds.

Chapter 3

Unfolding Trees

3.1 The Unfolding Tree T_v^d

We now restate the definition of an unfolding tree—first given by Scarselli et al. [3]—which leads to a natural equivalence relation between vertices of a given graph. Informally, the unfolding tree of a vertex v in a labelled graph G can be thought of as a decision tree of walks starting at v in G. We start with a vertex corresponding to the length 0 walk $\langle v \rangle$, and then connect all of the length 1 walks starting at v to the vertex $\langle v \rangle$. From there we may grow the tree further by connecting length 2 walks to the length 1 walks that they extend. Each of vertex of the unfolding tree receives the same label as that of the last vertex in the corresponding walk, *e.g.* the label of an unfolding tree vertex $\langle v, x \rangle$ would be the same as the label of x in G.

The thrust of this definition is that the vertices and edges of the tree are not directly identified with those of G from which they were derived: so an unfolding vertex $\langle v, x, v \rangle$ is distinct from the root vertex $\langle v \rangle$. In this manner, the unfolding tree encodes the walk-structure (including labels) of G from the viewpoint of v.

Definition 3.1.1: Let $G = \langle V, E, L \rangle$ be a labelled graph, let u be any vertex of G, and take any $d \in \mathbb{N}$. The **unfolding tree** T_u^d of depth d at u is a labelled graph defined as follows. The vertex set of T_u^d is given by

$$V_T := \{w = \langle u, v_1, \ldots, v_m \rangle \mid w \text{ is a walk in } G \text{ and } \ell(w) \leq d\},$$

the set of walks starting at u of length at most d. The edge set E_T of the unfolding tree is induced by the following adjacency rule: for any $w, w' \in V_T$ where $\ell(w) \leq \ell(w')$,

$$w \sim_{T_u^d} w' \iff w = w'[\ell(w') - 1].$$

Finally, the labeling function L_T is defined for $w \coloneqq \langle u, v_1, \dots, v_m \rangle$ as

$$L_T(w) \coloneqq L(v_m).$$

■

It is easy to conflate vertex names with vertex labels, so it should be kept in mind that in the above definition we use walks in G to *name* vertices, and then we *label* the vertices according to labels from G. Figures 3.2 and 3.3 below show an example of an unfolding tree, specifically one created from the labelled version of the house graph G in Figure 3.1.

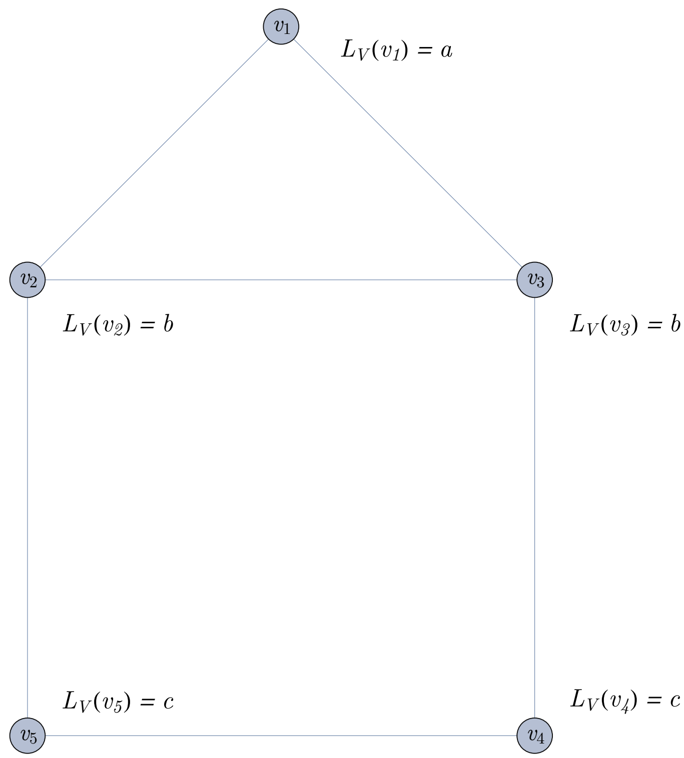

Figure 3.1: Labelled house graph G with vertex set $V \coloneqq \{v_1, \dots, v_5\}$.

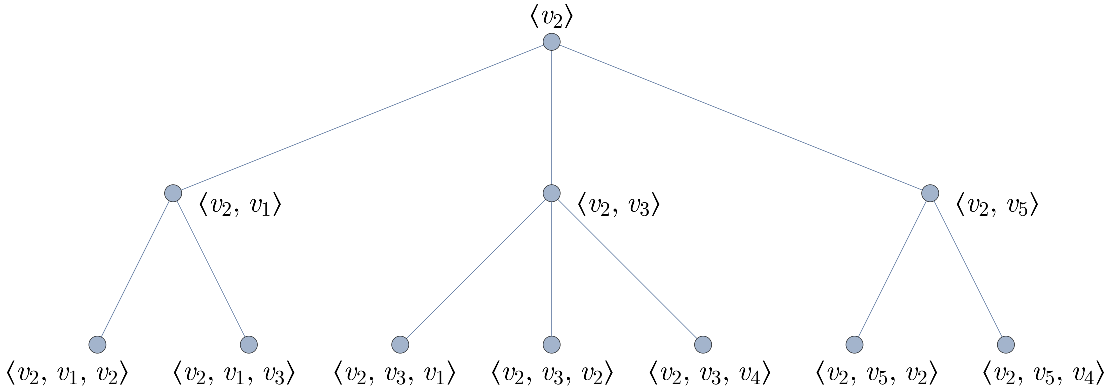

Figure 3.2: Unfolding tree $T^2_{v_2}$ of G with vertex names shown.

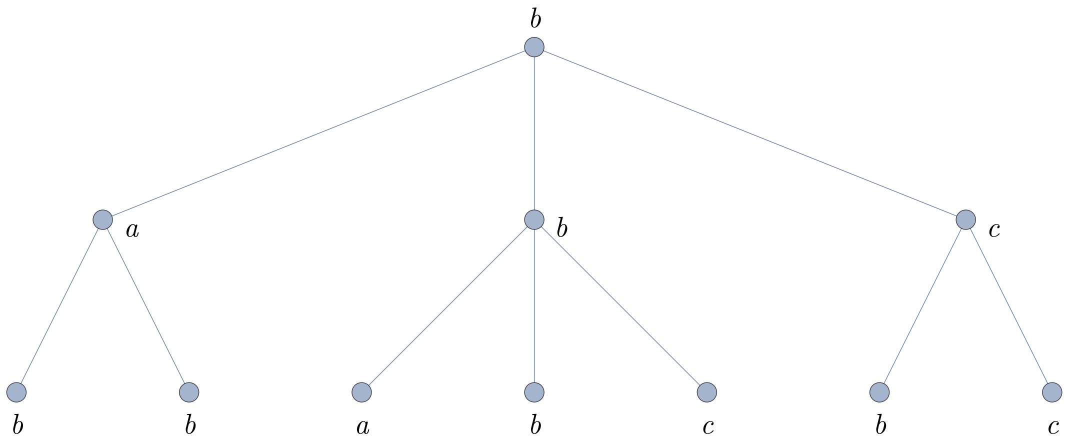

Figure 3.3: Vertex labeling of $T^2_{v_2}$.

As can be seen in Figure 3.2, the vertex set of $T^2_{v_2}$ consists of all walks in G of length at most 2 that start at v_2. Moreover, vertices are adjacent whenever the walk corresponding to one of them extends that of the other by a single vertex, *e.g.* the unfolding vertex $\langle v_2, v_5, v_4\rangle$ is adjacent to $\langle v_2, v_5\rangle$ since the walk corresponding to the former appends v_4 to that of the latter. In addition, Figure 3.2 demonstrates the way in which the labels of the unfolding tree are predicated on those of G. For example, $L_T\big(\langle v_2, v_1\rangle\big) = a = L(v_1)$ and $L_T\big(\langle v_2, v_1, v_2\rangle\big) = b = L(v_2)$.

3.2 Features of Unfolding Trees

HAVING just introduced unfolding trees, we now get acquainted with these graphs by setting out some vocabulary, notation and basic results related to them.

Definition 3.2.1: Let T_v^d be an unfolding tree. We call $\langle v\rangle \in V(T_v^d)$ the **root** of T_v^d. Given some $w \in V(T_v^d)$, the **level** of w in T_v^d is the length of w. The set of all vertices of level k in T_v^d is termed the **kth level of** T_v^d, written $T_v^d[\![k]\!]$. The vertices in $T_v^d[\![d]\!]$—the deepest level of the unfolding tree—are termed **leaves**. ■

Proposition 3.2.2: Let w be a vertex of level m in an unfolding tree T_v^d. Then $\mathcal{N}(w) \subseteq T_v^d[\![m-1]\!] \cup T_v^d[\![m+1]\!]$.

Proof: Let w be a vertex in the mth level of an unfolding tree T_v^d, and consider an arbitrary vertex $w' \in \mathcal{N}(w)$. Prima facie, there are three possibilities with regard to the level m' of w': either $m' < m$, or $m < m'$, or $m' = m$.

First, suppose that $m' < m$. Then by the definition of the edge set of an unfolding tree,

$$w \sim_{T_v^d} w' \quad\Rightarrow\quad w' = w[m-1] \quad\Rightarrow\quad m' = m-1.$$

Similarly, if instead we have $m' > m$, then

$$w \sim_{T_v^d} w' \quad\Rightarrow\quad w = w'[m'-1] \quad\Rightarrow\quad m' = m+1.$$

Finally, if $m' = m$ then neither $w = w'[m'-1]$ nor $w' = w[m-1]$ holds, and so $w \nsim_{T_v^d} w'$, a contradiction. Therefore w' is of level $m-1$ or $m+1$, and we are done. □

The above proposition means that a vertex of an unfolding tree only connects to vertices from the levels directly above and below it. This being the case, we now name the 'lower' and 'upper' subsets of an unfolding tree vertex's neighborhood.

Definition 3.2.3: Let w be a vertex of some unfolding tree T_v^d. The **daughter set** of w is $\mathcal{D}(w) := \{w' \in \mathcal{N}(w) \mid w' \in T_v^d[\![\ell(w)+1]\!]\}$, and its elements are referred to as **daughters** of w. If w is a non-root vertex, the **mother** $\mathcal{M}(w)$ of w is the vertex $w' \in \mathcal{N}(w)$ in level $\ell(w)-1$ of T_v^d. ■

Note that $\mathcal{M}(w)$ is unique: if x, y are neighbors of w with $\ell(x), \ell(y) < \ell(w)$ then it must be that $x = w[\ell(w) - 1] = y$. Also observe that $w \in \mathcal{D}(w') \Rightarrow \mathcal{M}(w) = w'$.

The astute reader may have realized that, humorously enough, we have not actually shown that every unfolding tree is a tree. Though we will not formally prove this, we take a moment to sketch the argument. The uniqueness of mothers together with Proposition 3.2.2 ensures that an unfolding tree contains no cycles: since there are no adjacent vertices of the same level, there must be a unique deepest vertex in the cycle, but then this vertex would have to be connected to two vertices above it, a contradiction. Moreover, every vertex w has a path to the root r, namely

$$P_w := \langle w,\, w[\ell(w) - 1], \ldots, w[1], r \rangle,$$

and so we can obtain a path between any two vertices w, w' by joining P_w and $P_{w'}$ into

$$P_{w,w'} := \langle w,\, w[\ell(w) - 1], \ldots, w[1],\, r,\, w'[1], \ldots, w'[\ell(w') - 1],\, w' \rangle.$$

So unfolding trees are connected as well as acyclic, and thus are trees.

3.3 Unfolding Equivalence

As mentioned, unfolding trees suggest a natural relationship between vertices, which we shall call unfolding equivalence.

Definition 3.3.1: Given vertices u, v of a labelled graph G, suppose that for all $d \geq 0$ we have $T_u^d \cong T_v^d$ and $\psi_d\big(\langle u \rangle\big) = \langle v \rangle$, where ψ_d is the isomorphism from T_u^d to T_v^d. When this is the case, we say that u and v are **unfolding equivalent**, writing $u \doteqdot v$. ■

Figure 3.4 shows the unfolding tree $T_{v_3}^2$, which is label isomorphic to $T_{v_2}^2$ from Figures 3.2 and 3.3 owing to the following function:

$$\begin{aligned}
\phi : V(T^2_{v_2}) &\rightarrowtail\!\!\!\!\twoheadrightarrow V(T^2_{v_3}) \\
\langle v_2\rangle &\longmapsto \langle v_3\rangle, \\
\langle v_2, v_1\rangle &\longmapsto \langle v_3, v_1\rangle, \\
\langle v_2, v_3\rangle &\longmapsto \langle v_3, v_2\rangle, \\
\langle v_2, v_5\rangle &\longmapsto \langle v_3, v_4\rangle, \\
\langle v_2, v_1, v_2\rangle &\longmapsto \langle v_3, v_1, v_3\rangle, \\
\langle v_2, v_1, v_3\rangle &\longmapsto \langle v_3, v_1, v_2\rangle, \\
\langle v_2, v_3, v_1\rangle &\longmapsto \langle v_3, v_2, v_1\rangle, \\
\langle v_2, v_3, v_2\rangle &\longmapsto \langle v_3, v_2, v_3\rangle, \\
\langle v_2, v_3, v_4\rangle &\longmapsto \langle v_3, v_2, v_1\rangle, \\
\langle v_2, v_5, v_2\rangle &\longmapsto \langle v_3, v_4, v_3\rangle, \\
\langle v_2, v_5, v_4\rangle &\longmapsto \langle v_3, v_4, v_5\rangle.
\end{aligned}$$

In fact, v_2 is unfolding equivalent to v_3.

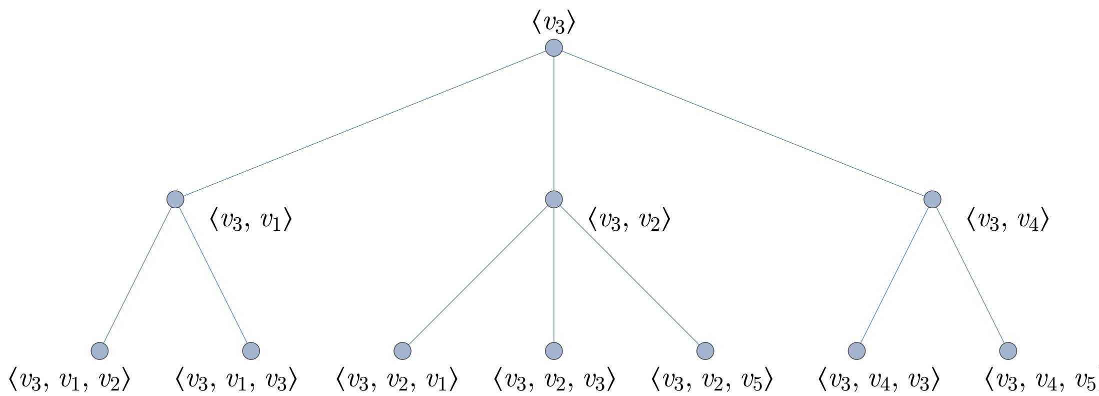

Figure 3.4: Unfolding tree $T^2_{v_3}$ of G (see Fig. 3.1) with vertex names shown.

Unfolding equivalence is, as the name suggests, an equivalence relation. This is easily verified, with all of the required properties inherited immediately from the symmetry, reflexivity and transitivity of isomorphism.

The following lemma will play a technical role in the proof of Theorem 5.1.3.

Proposition 3.3.2: Let T^d_u, T^d_v be arbitrary unfolding trees of G, H respectively, and suppose there exists a label isomorphism ϕ from T^d_u to T^d_v. Then for any vertex w of T^d_u, the level of $\phi(w)$ equals the level of w.

Proof: We prove the result by showing that $\ell(w)$ equals the distance from w to the root $\langle u \rangle$, and then showing that isomorphisms preserve distance.

Take any $w = \langle u, w_1, \ldots, w_m \rangle$ in T_u^d. Observe that the sequence

$$P \coloneqq \Big\langle w,\ w[m-1],\ w[m-2], \ldots,\ w[1],\ \langle u \rangle \Big\rangle$$

is a walk in T_u^d, since by Definition 3.1.1 we know that $w[i] \sim_{T_u^d} w[i-1]$ for any $i \in \mathbb{N}$. Furthermore, since P is constructed so that its vertices are all of different length, they must all be distinct; thus P is a path. Since T_u^d is a tree, this is the only path between the root and w, and so it is the shortest such path by default. Therefore the distance from $\langle u \rangle$ to w is $\ell(P) = \ell(w)$.

Now consider the (coordinate-wise) image of P under ψ:

$$\psi[P] \coloneqq \Big\langle \psi(w),\ \psi(w[m-1]), \ldots,\ \psi(w[1]),\ \langle v \rangle \Big\rangle.$$

Since ψ preserves adjacency and P is a walk, it follows that $\psi[P]$ is a walk in T_v^d. Furthermore, because ψ is injective and the vertices of P are distinct, we know that the vertices of $\psi[P]$ are distinct. So $\psi[P]$ is a path in T_v^d from $\psi(w)$ to $\langle v \rangle$. For the same reasons as given for P, the length of $\psi[P]$ equals the distance from $\psi(w)$ to the root $\langle v \rangle$ with the same length as P. But we showed that this is exactly the level of $\psi(w)$, and so

$$\ell\big(w\big) = \ell\big(\psi[P]\big) = \ell\big(P\big) = \ell\big(\psi(w)\big).$$

□

3.4 Unfolding Trees of k-Regular Graphs

This section investigates properties of unfolding trees derived from graphs with a property known as k-regularity, with the aim of characterizing these graphs' unfolding trees and further elucidating the structure of unfolding trees.

A graph is said to be k**-regular** if all of its vertices are degree k. We show that the vertices of a k-regular graph have perfect k-ary unfolding trees (a **perfect** k**-ary tree** is one in which a vertex has k daughters unless it is in the last level of the tree.)

Proposition 3.4.1: Given any vertex v of a k-regular graph G, the unfolding tree T_v^d is a perfect k-ary tree for every $d \in \mathbb{Z}_+$.

Proof: Consider $w \in T_v^d$ with $\ell(w) < d$, and let $u := w[\![\ell(w)]\!] \in G$. Then $|\mathcal{D}(w)| = |\mathcal{N}_G(u)| = d_G(u)$, and since G is k-regular $d_G(u) = k$, so $|\mathcal{D}(w)| = k$. □

Corollary 3.4.2: If G is k-regular and all vertices of G have the same label, then $\forall u, v \in V$ we have $u \doteq v$.

Proof: Take an arbitrary $d \in \mathbb{N}$. We construct an isomorphism $\phi : T_u^d \rightarrowtail\!\!\!\!\twoheadrightarrow T_v^d$ by first sending one root to the other, and then proceeding at each successive level to map daughters of a vertex to daughters of its image. In this manner, if w is an unfolding vertex of level $m < d$ and $\phi(w) = w'$, then both w and w' have k daughters, which we pair with an arbitrary bijection from $\mathcal{D}(w)$ to $\mathcal{D}(w')$. This construction yields a bijection ϕ. Because daughters are by definition adjacent to their mothers, and since by Proposition 3.2.2 the unfolding tree contains only mother-daughter edges, condition (1) of Definition 2.1.1 is satisfied. Lastly, since the labeling of G is a constant function, part (2) of Definition 2.1.1 is vacuously true. □

For example, every cycle graph is 2-regular, and so the unfolding tree of any vertex of a cycle graph is a perfect binary tree.

Chapter 4

Similarity & Unfolding Equivalence

4.1 Similar Vertices Are Unfolding Equivalent

Having laid out the requisite preliminaries, we are now ready to state and prove our first substantial result. In it, we show that if two vertices of a labelled graph G are similar, then they are unfolding equivalent.

Theorem 4.1.1: Let G be a labelled graph, and let u, v be vertices of G. Then $u \doteq v \Rightarrow u \mathrel{\bar{\div}} v$.

Proof: Let $G \coloneqq \langle V, E, L \rangle$ be a labelled graph and let $u, v \in V$. We shall denote the vertex set of T_u^d by V_u and its edge set by E_u, and use a similar convention for T_v^d.

Suppose that $\phi : V \rightarrowtail\!\!\!\!\twoheadrightarrow V$ is a label automorphism such that $\phi(u) = v$. We claim that the mapping $\psi : V_u \longrightarrow V_v$, defined by

$$\psi\big(\langle u, x_1, x_2, \ldots, x_n \rangle\big) \coloneqq \langle v, \phi(x_1), \phi(x_2), \ldots, \phi(x_n) \rangle,$$

is a labelled graph isomorphism of T_u^d and T_v^d.

First we demonstrate that ψ satisfies part (1) of the definition of a label isomorphism. Suppose that $w \sim_{T_u^d} w'$ for some $w, w' \in V_u$; assume without loss of generality that the walk corresponding to w' is longer than that of w, and put $w' \coloneqq \langle u, w_1, w_2, \ldots, w_m \rangle$. Then by the definition of an unfolding tree it must be the case that $w = \langle u, w_1, \ldots, w_{m-1} \rangle$. Now consider the images of these under ψ, *i.e.* the vertices $\psi(w) = \langle v, \phi(w_1), \ldots, \phi(w_{m-1}) \rangle$ and $\psi(w') = \langle v, \phi(w_1), \ldots, \phi(w_m) \rangle$ of T_v^d; we must show that $\psi(w) \sim_{T_v^d} \psi(w')$.

Since $\psi(w')$ extends the tuple $\psi(w)$ by one vertex, we need only prove that both of $\psi(w), \psi(w')$ are walks in G of length at most d starting at v. Showing that they are walks is simple since—given that w' is by hypothesis a walk in G—we know that $w_i \sim_G w_{i+1}$ for all $0 \le i < m$, and therefore $\phi(w_i) \sim_G \phi(w_{i+1})$ because ϕ preserves adjacency. Moreover, since $\phi(u) = v$ we know that $\psi(w), \psi(w')$ are walks starting at v. Finally, because w' and $\psi(w')$ are of the same length $m \le d$, we know that $\psi(w')$ has length at most d, and similarly for $\psi(w)$. Thus $\psi(w), \psi(w')$ are adjacent vertices of T_v^d, and so $w \sim_{T_u^d} w' \Rightarrow \psi(w) \sim_{T_v^d} \psi(w')$. Because ϕ is a label isomorphism, by Lemma 2.2.2 it admits an inverse which is also a label isomorphism, and thus $\psi(w) \sim_{T_v^d} \psi(w') \Rightarrow w \sim_{T_u^d} w'$ also. So ψ satisfies part (1) of the definition of a label isomorphism.

Finally, we prove that ψ preserves labels. Keeping the above definition of w', observe that by virtue of the definition of an unfolding tree we know w' has the label $L(w_m)$. Now, since the last coordinate of $\psi(w')$ is $\phi(w_m)$, we get that

$$L_v(\psi(w')) = L(\phi(w_m)) = L(w_m) = L_u(w'),$$

meaning that ψ preserves labels. □

4.2 Counterexamples to $u \div v \Rightarrow u \doteq v$

BECAUSE similarity is defined in terms of label isomorphisms, which fully capture the structure of a labelled graph, Theorem 4.1.1 is perhaps unsurprising. What is not immediately clear is whether its converse is true. In fact, this author originally conjectured just that, and was surprised to find that it is not always the case. Below are some counterexamples.

4.2.1 Disconnected Cycles

THE vertices u, v of the graph in Figure 4.1 are unfolding equivalent, but clearly not similar. Their unfolding tree of depth 4 is shown in Figure 4.2. In fact, a given vertex of this graph is unfolding equivalent to every other vertex with the same label.

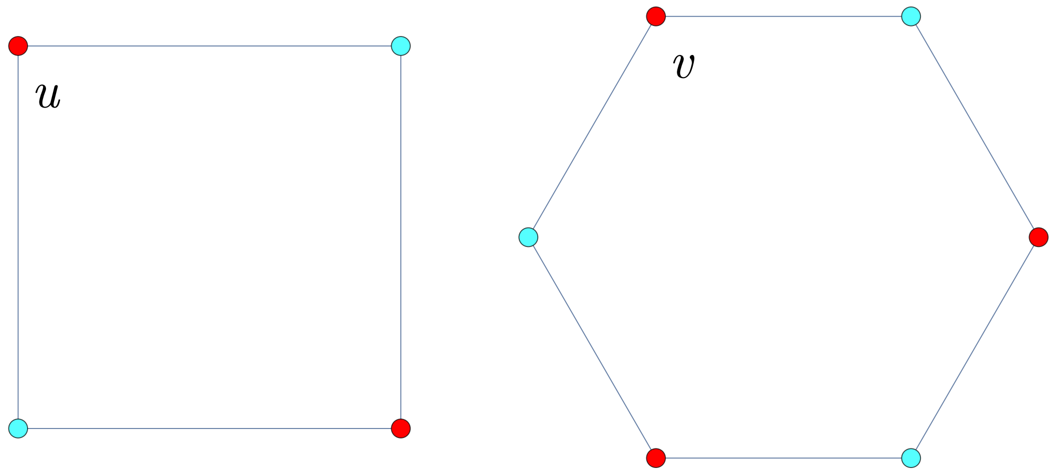

Figure 4.1: Disconnected graph with labeling represented by vertex color.

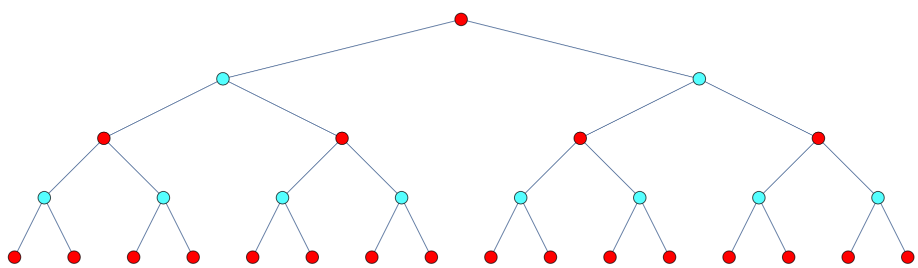

Figure 4.2: Unfolding tree $T_u^4 \cong T_v^4$ of the graph from Figure 4.1. Vertex labels are represented by the colors used in the original graph.

4.2.2 Edge-Switched Pentagonal Prisms

The previous example might hint that the result can be salvaged by requiring that the graph in question be connected. However, this is not the case. Figure 4.3 shows a connected graph obtained by swapping edges of a pentagonal prism with constant labeling. Similar vertices are presented with the same color; as can be seen, not all pairs of vertices are similar. However, recall from Corollary 3.4.2 that since this graph is 3-regular and has constant labeling, all vertices are unfolding equivalent.

Note that the graphs from Figures 4.1 and 4.3 are both regular. In light of

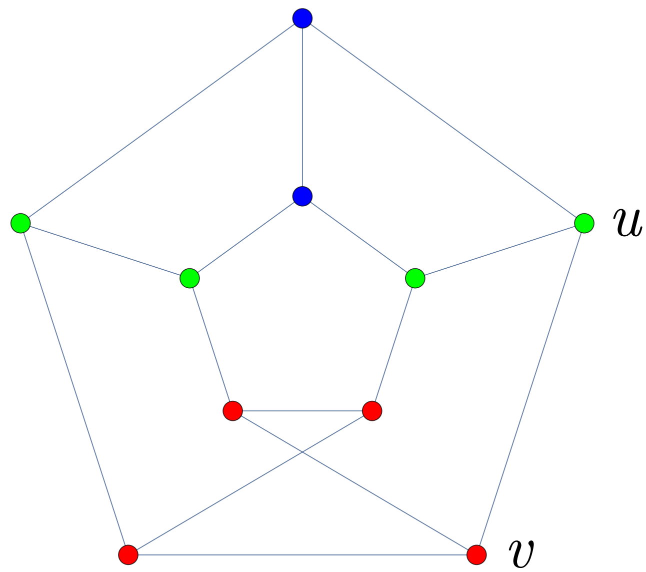

Figure 4.3: Edge-switched version of a pentagonal prism with constant labeling. Vertices of the same color are similar. Vertices u and v are unfolding equivalent but *not* similar.

this, we might try to further tighten the constraints by requiring non-regularity, in hopes of finding a class of graphs for which $u \fallingdotseq v \Rightarrow u \doteq v$. This would be to no avail: Figure 4.4 shows a graph obtained from the previous one by adding a vertex between each pair of vertices in the inner and outer pentagons. Once again, we take some constant function as our labelling. The graph is not regular, but each vertex is unfolding equivalent to every other vertex of the same degree. Since the sets of similar vertices shown do not coincide with the partition of vertices by degree, this graph also contains pairs of vertices which are unfolding equivalent but not similar, *e.g.* u and v.

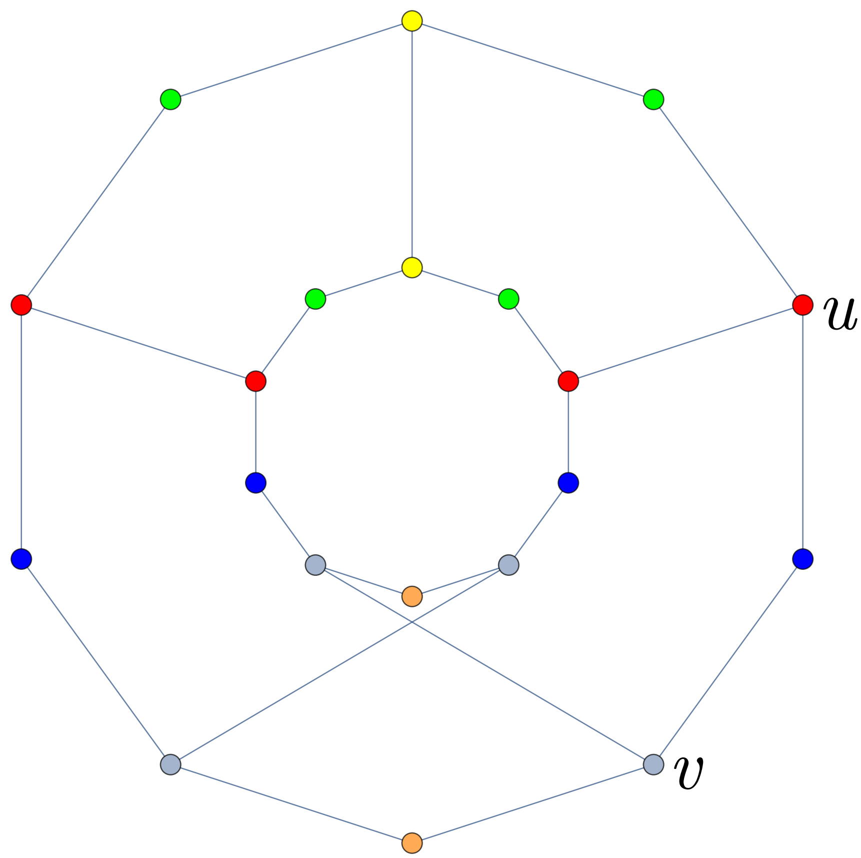

Figure 4.4: Modified version of the graph in Figure 4.3. As with that figure, the graph's labeling is constant, colors represent similarity classes, and u, v are unfolding equivalent but not similar.

Chapter 5

Symmetrically-Associated Graphs

THE previous chapter concluded with examples of graphs containing vertices that violate the converse of Theorem 4.1.1. The first couple of these were k-regular graphs and so, as we have discussed, had perfect k-ary unfolding trees at every vertex. However, the last example featured a graph that had vertices of varying degrees, and vertices were unfolding equivalent whenever they were of equal degree. This leads us to an intriguing question: is there a class of graphs wherein unfolding equivalence classes can be determined by properties of the vertices, without any direct reference to the unfolding trees? We answer in the affirmative by characterizing a class of graphs for which two vertices are unfolding equivalent exactly when they have the same degree and label. A graph of this type, which we term symmetrically-associated, is distinguished by the property that every vertex v has a fixed number m of neighbors with degree k and label l_λ (for any k, λ), where m is entirely determined by $d(v)$ and $L(v)$.

5.1 Definition & Example

IN this chapter and the subsequent one, we will make frequent use of the notation $[N]$, which is shorthand for the interval of positive integers from 1 to N.

Definition 5.1.1: We say a labelled graph G is **symmetrically-associated** if $\forall k \in [\Delta]$, $\forall \lambda \in [\Lambda]$ there exists a sequence of non-negative integers

$$m^{k,\lambda} = \left\langle m^{k,\lambda}_{1,1}, m^{k,\lambda}_{1,2}, \ldots, m^{k,\lambda}_{1,\Lambda}, m^{k,\lambda}_{2,1}, \ldots, m^{k,\lambda}_{\Delta,\Lambda} \right\rangle$$

such that $\forall v \in V$ with $d(v) = k$, $L(v) = l_\lambda$ we have

$$m^{k,\lambda}_{i,j} = \left| \{u \in \mathcal{N}(v) \mid d(u) = i,\, L(u) = l_j\} \right|.$$

We impose the additional condition that $m^{k,\lambda} = \langle 0, 0, \ldots, 0\rangle$ whenever G has no vertices with degree k and label l_λ.

The sequence $m^{k,\lambda}$ is called the k,λ**-association** of G. When discussing a particular vertex v of degree k and label l_λ, we will occasionally refer to $m^{k,\lambda}$ simply as the **association** of v. The set $\{m^{k,\lambda} \mid k \in [\Delta], \lambda \in [\Lambda]\}$ of all k,λ-associations of G is called the **association profile** of G. ■

Notation 5.1.2: For convenience,* when working with a symmetrically-associated graph with vertex set V we follow the convention of writing $V^{k,\lambda}$ to represent the set of vertices with degree k and label l_λ. We also use the notation $n^{k,\lambda}$ to denote $|V^{k,\lambda}|$, the number of vertices in the graph with k,λ-association. ■

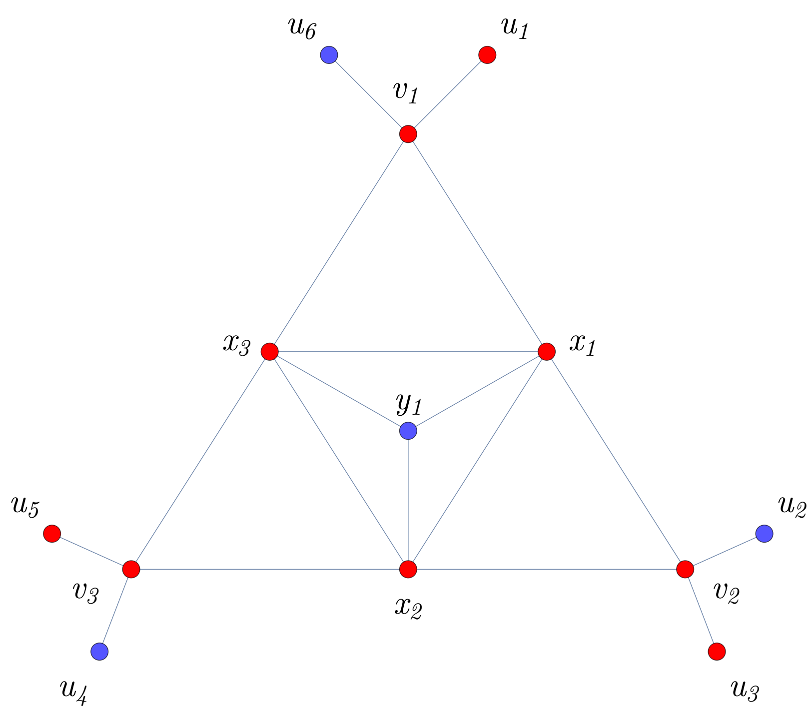

Figure 5.1: A symmetrically-associated graph with vertices colored by label.

* Relatively speaking. We are morally obligated to warn the reader that sub- and superscripts reach noxious levels in this chapter. For this we are sincerely sorry.

Figure 5.1 is an example of a symmetrically-associated graph. Symbolizing the red labels by l_1 and the blue labels by l_2, the association profile of the graph is

$$m^{1,1} = \langle \overbrace{\underbrace{0}_{l_1}, \underbrace{0}_{l_2}}^{\text{degree 1}}, \overbrace{\underbrace{0}_{l_1}, \underbrace{0}_{l_2}}^{\text{degree 2}}, \overbrace{\underbrace{0}_{l_1}, \underbrace{0}_{l_2}}^{\text{degree 3}}, \overbrace{\underbrace{1}_{l_1}, \underbrace{0}_{l_2}}^{\text{degree 4}}, \overbrace{\underbrace{0}_{l_1}, \underbrace{0}_{l_2}}^{\text{degree 5}} \rangle$$

$$m^{1,2} = \langle \overbrace{0,0}^{\text{deg 1}}, \overbrace{0,0}^{\text{deg 2}}, \overbrace{0,0}^{\text{deg 3}}, \overbrace{1,0}^{\text{deg 4}}, \overbrace{0,0}^{\text{deg 5}} \rangle,$$

$$m^{3,2} = \langle \overbrace{0,0}^{\text{deg 1}}, \overbrace{0,0}^{\text{deg 2}}, \overbrace{0,0}^{\text{deg 3}}, \overbrace{0,0}^{\text{deg 4}}, \overbrace{3,0}^{\text{deg 5}} \rangle,$$

$$m^{4,1} = \langle \overbrace{1,1}^{\text{deg 1}}, \overbrace{0,0}^{\text{deg 2}}, \overbrace{0,0}^{\text{deg 3}}, \overbrace{0,0}^{\text{deg 4}}, \overbrace{2,0}^{\text{deg 5}} \rangle,$$

$$m^{5,1} = \langle \overbrace{0,0}^{\text{deg 1}}, \overbrace{0,0}^{\text{deg 2}}, \overbrace{0,1}^{\text{deg 3}}, \overbrace{2,0}^{\text{deg 4}}, \overbrace{2,0}^{\text{deg 5}} \rangle.$$

The k, λ-associations with entries all equal to zero have been omitted. Examining one of the associations, the fact that $m^{4,1}_{1,1} = m^{4,1}_{1,2} = 1$ and $m^{4,1}_{5,1} = 2$ means that if a blue vertex has degree 4, then it has exactly one red degree 1 neighbor, one blue degree 1 neighbor and two red degree 5 neighbors. The vertices fall into the degree/label vertex subsets as follows: $V^{1,1} = \{u_1, u_3, u_5\}$, $V^{1,2} = \{u_2, u_4, u_6\}$, $V^{3,2} = \{y_1\}$, $V^{4,1} = \{v_1, v_2, v_3\}$ and $V^{5,1} = \{x_1, x_2, x_3\}$.

A moment's reflection reveals that all k-regular graphs with constant labelings are symmetrically-associated. Consider, for example, a 3-regular graph with all vertices labelled l_1; then $m^{k,\lambda}_{k',\lambda'} = 3$ exactly when $k = k' = 3$, $\lambda = \lambda' = 1$ and is zero otherwise. In fact, symmetric-association constitutes a generalization of k-regularity—one that, even in the unlabelled form (see Appendix B), is to the best of our knowledge absent from the literature. Furthermore, the uniformity the property imposes on vertices' 'degree/label environments' endows symmetrically-associated graphs with interesting unfolding properties, as demonstrated in the following theorem. The proof is rather technical.

Theorem 5.1.3: Let G be a symmetrically-associated graph. Then for all vertices u, v of G,

$$d(u) = d(v),\ L(u) = L(v) \iff u \doteqdot v.$$

Proof: **(Sufficiency)** Assuming $u \doteqdot v$, we know that there is a label isomorphism ψ from T_u^1 to T_v^1, and it must send $\langle u \rangle$ to $\langle v \rangle$ by Proposition 3.3.2; so since ψ preserves labels we see that the labels of $\langle u \rangle$ and $\langle v \rangle$ in their respective trees are equal. Then by the definition of an unfolding tree's labeling,

$$L_G(u) = L_{T_u^1}\big(\langle u \rangle\big) = L_{T_v^1}\big(\langle v \rangle\big) = L_G(v).$$

Next, because ψ is bijective, by Proposition 3.3.2 we know $|T_u^1[\![1]\!]| = |T_v^1[\![1]\!]|$. Furthermore, the size of $T_u^1[\![1]\!]$ equals the degree of u in G (and similarly for $T_v^1[\![1]\!]$ and $d_G(v)$), so combining this with the equality in the previous sentence gives $d_G(u) = d_G(v)$.

(Necessity) We use induction on the unfolding trees' depths to prove a stronger result: given the hypotheses, for any depth D there exists an isomorphism ψ_D from T_u^D to T_v^D with the property that, $\forall r \geq 0$,

$$\psi_D(w) = w' \implies d_G(w[\![r]\!]) = d_G(w'[\![r]\!]) \text{ and } L_G(w[\![r]\!]) = L_G(w'[\![r]\!]).$$

For the base case $D = 0$, we put $\psi_0\big(\langle u \rangle\big) \coloneqq \langle v \rangle$. This is clearly a bijection, and there are no edges to check, so ψ_0 is an isomorphism. Moreover, by hypothesis we know $d(u) = d(v)$ and $L(u) = L(v)$.

Now let $D \geq 1$ and suppose $\exists \psi_{D-1} : V\big(T_u^{D-1}\big) \rightarrowtail\!\!\!\!\twoheadrightarrow V\big(T_v^{D-1}\big)$ a label isomorphism satisfying our associaton-preservation requirement. Since $T_u^{D-1} \leq T_u^D$, we construct ψ_D by first letting $\psi_D \restriction_{T_u^{D-1}} \coloneqq \psi_{D-1}$. We now need only define ψ_D on the set of leaves $T_u^D[\![D]\!]$.

Call $T_u^D[\![D-1]\!] \eqqcolon \{w_1, \dots, w_m\}$, and let $w_i' \coloneqq \psi_{D-1}(w_i)$, noting that $T_v^D[\![D-1]\!] = \{w_i'\}_{i \in [m]}$ by Proposition 3.3.2. In addition, define

$$S_i \coloneqq \mathcal{D}(w_i), \qquad S_i' \coloneqq \mathcal{D}(w_i');$$
$$k_i \coloneqq d_G(w[\![D-1]\!]), \qquad \lambda_i \coloneqq L(w[\![D-1]\!]).$$

Since each vertex of $T_u^D[\![D]\!]$ has a unique mother in $T_u^D[\![D-1]\!]$, the S_i partition $T_u^D[\![D]\!]$. We proceed by defining ψ_D on each S_i.

We can write every S_i as the union of disjoint subsets

$$S_i = \bigcup_{\substack{k\in[\Delta]\\ \lambda\in[\Lambda]}} S_{i,k,\lambda},$$

where $S_{i,k,\lambda} \coloneqq \{x \in S_i \mid x[\![D]\!] \in V^{k,\lambda}\}$, the set of daughters of w_i whose last coordinate has k, λ-association in G. Additionally, we can express S'_i as a disjoint union of subsets $S'_{i,k,\lambda}$ defined analogously to the $S_{i,k,\lambda}$. Observe that

$$\mathcal{N}_G(w_i[\![D-1]\!]) = \{x[\![D]\!] \in V_G \mid x \in S_i\},$$

which is to say that each vertex in S_i corresponds to a neighbor of $w_i[\![D-1]\!]$ in G. Therefore $|S_{i,k,\lambda}| = m_{k,\lambda}^{k_i,\lambda_i}$ for all $k \in [\Delta]$, $\lambda \in [\Lambda]$. By the inductive hypothesis,

$$\psi_{D-1}(w_i) = w'_i \quad \Longrightarrow \quad d_G(w'_i[\![D-1]\!]) = k_i \text{ and } L(w'_i[\![D-1]\!]) = \lambda_i.$$

So since G is symmetrically-associated, we know $w'[\![D-1]\!]$ has the same association as $w[\![D-1]\!]$. Therefore $|S_{i,k,\lambda}| = |S'_{i,k,\lambda}|$, and so there exists a bijection $\theta_{i,k,\lambda}$ from $S_{i,k,\lambda}$ to $S'_{i,k,\lambda}$; fix such a $\theta_{i,k,\lambda}$ for each $S_{i,k,\lambda}$.

We define ψ_D on $T_u^D[\![D]\!]$ as follows. Given $x \in T_u^D[\![D]\!]$ such that $x \in S_{i,k,\lambda}$, we put

$$\psi_D(x) \coloneqq \theta_{i,k,\lambda}(x).$$

Before proceeding, we must show that this definition describes a function. Since

$$T_u^D[\![D]\!] = \bigcup_{i\in[m]} S_i = \bigcup_{i\in[m]} \bigcup_{\substack{k\in[\Delta]\\ \lambda\in[\Lambda]}} S_{i,k,\lambda},$$

and since the $S_{i,k,\lambda}$ are pairwise disjoint, we know that each leaf of T_u^D is contained in a unique $S_{i,k,\lambda}$. Moreover, the subset $S_{i,k,\lambda}$ has a unique function $\theta_{i,k,\lambda}$ associated with it. So the value $\psi_D(x)$ exists and is unique for each $x \in T_u^D[\![D]\!]$. This means that ψ_D is a function on $T_u^D[\![D]\!]$, and since $\psi_D \restriction_{T_u^{D-1}} = \psi_{D-1}$ is a function on the rest of T_u^D we conclude that the above definition yields a function $\psi_D : T_u^D \to T_v^D$.

Confirming that ψ_D is a label isomorphism, we first show that it is a bijection. By construction,

$$\begin{aligned}
\psi_D[T_u^D[\![D]\!]] &= \bigcup_{i\in[m]} \bigcup_{\substack{k\in[\Delta]\\ \lambda\in[\Lambda]}} \theta_{i,k,\lambda}[S_{i,k,\lambda}] \\
&\subseteq \bigcup_{i\in[m]} \bigcup_{\substack{k\in[\Delta]\\ \lambda\in[\Lambda]}} S'_{i,k,\lambda} \\
&= T_v^D[\![D]\!],
\end{aligned}$$

and therefore

$$\begin{aligned}\left(\psi_D[T_u^D[\![D]\!]]\right)\cap\left(\psi_D[T_u^{D-1}]\right) &= \left(\psi_D[T_u^D[\![D]\!]]\right)\cap\left(T_v^{D-1}\right)\\ &\subseteq \left(T_v^D[\![D]\!]\right)\cap\left(T_v^{D-1}\right)\\ &= \varnothing.\end{aligned}$$

That is to say, the image of T_v^{D-1} under ψ_D is disjoint from that of $T_v^D[\![D]\!]$. This, together with the inductive hypothesis that $\psi_D \restriction_{T_u^{D-1}} = \psi_{D-1}$ is bijective, implies that it suffices to show that $\psi_D \restriction_{T_u^D[\![D]\!]}$ is a bijection from $T_u^D[\![D]\!]$ to $T_v^D[\![D]\!]$. Now, suppose $\exists x, y \in T_u^D[\![D]\!]$ s.t. $\psi_D(x) = \psi_D(y)$. Letting $S_{i,k,\lambda}$ denote the unique subset containing x, we see that $y \in S_{i,k,\lambda}$ as well; for if, on the contrary, we had $y \in S_{j,h,\mu} \neq S_{i,k,\lambda}$ it would follow that both

$$\psi_D(y) = \theta_{j,h,\mu}(y) \in S'_{j,h,\mu},$$

and

$$\psi_D(y) = \psi_D(x) = \theta_{i,k,\lambda}(x) \in S'_{i,k,\lambda},$$

contradicting the fact that $S'_{j,h,\mu}$, $S'_{i,k,\lambda}$ are disjoint. Then $x, y \in S_{i,k,\lambda}$ implies that

$$\theta_{i,k,\lambda}(x) = \psi_D(x) = \psi_D(y) = \theta_{i,k,\lambda}(y),$$

and so since $\theta_{i,k,\lambda}$ is injective we conclude that $x = y$. Therefore $\psi_D \restriction_{T_u^D[\![D]\!]}$ is injective.

Next, recalling that $T_u^D[\![D]\!]$ is the union of the pairwise disjoint subsets S_i (and similarly for $T_v^D[\![D]\!]$ and the S'_i), and that

$$|S_i| = d_G(w_i[\![D-1]\!]) = d(w'_i[\![D-1]\!]) = |S'_i|,$$

we deduce that

$$\begin{aligned}|T_u^D[\![D]\!]| &= \sum_{i\in[m]} |S_i|\\ &= \sum_{i\in[m]} |S'_i|\\ &= |T_v^D[\![D]\!]|.\end{aligned}$$

So $\psi_D \restriction_{T_u^D[\![D]\!]}$ is an injection with domain and codomain of equal cardinality, and therefore is a bijection.

Now we check that $x \sim_{T_u^D} y \Leftrightarrow \psi_D(x) \sim_{T_v^D} \psi_D(y)$. Again, since ψ_{D-1} is a label isomorphism, we need only check edges incident to vertices in $T_u^D[\![D]\!]$. To that end, suppose we have $x \in T_u^D$, $y \in T_u^D[\![D]\!]$ s.t. $x \sim_{T_u^D} y$. Since y belongs to the deepest level, applying Proposition 3.2.2 gives

$$\mathcal{N}(y) \subseteq T_u^D[\![D-1]\!] \cup T_u^D[\![D+1]\!] = T_u^D[\![D-1]\!]$$

and thus x must be in level $D-1$. So $x = w_i$ for some $i \in [m]$, implying that $y \in S_i$. It follows that $\psi_D(y) \in S_i'$, which by definition of S_i' means that $\psi_D(y) \sim_{T_v^D} \psi_D(x)$. Turning to the other direction of the biconditional, assume instead that $x' \sim_{T_v^D} y'$ for some $x' \in T_v^D$, $y' \in T_v^D[\![D]\!]$. By the same logic as above we get that $x' = w_i'$ for some $i \in [m]$ and so $y' \in S_i'$. Then since $\psi_D^{-1}[S_i'] = S_i$, we have

$$\psi_D^{-1}(y') \in S_i \;\Rightarrow\; \psi_D^{-1}(y') \sim_{T_u^D} w_i = \psi_D^{-1}(x').$$

Finally, we show that ψ_D satisfies our additional stipulation that

$$\psi_D(w) = w' \quad \Longrightarrow \quad d_G(w[\![r]\!]) = d_G(w'[\![r]\!]) \text{ and } L_G(w[\![r]\!]) = L_G(w'[\![r]\!]).$$

As before, since ψ_{D-1} has this property by inductive hypothesis, we check only leaves. Recall that any unfolding tree vertex x in $S_{i,k,\lambda}$ maps to one in $S'_{i,k,\lambda}$. By construction of those subsets, this implies that the last coordinates of x and $\psi_D(x)$ both have k, λ-association in G, and we are done. □

5.2 Association Profiles

Just as two non-isomorphic graphs may have the same degree sequence, it is possible for two non-isomorphic symmetrically-associated graphs to have identical association profiles. The definition below allows us to describe such pairs of symmetrically-associated graphs.

Definition 5.2.1: We say that the symmetrically-associated graphs G and H are **association equivalent** if they have the same number of vertices, the same set of vertex labels, and the same association profile (when their labels are enumerated in the same order). ■

Given a pair of association equivalent graphs, we might wonder whether it is possible for them to have different values for $n^{k,\lambda}$. As it turns out, if G is connected then the $n^{k,\lambda}$ are uniquely determined by the association profile and total number of vertices. We now prove this.

Lemma 5.2.2: If G is a symmetrically-associated graph with association profile $\{m^{k,\lambda}\}$, then $\forall k, k' \in [\Delta],\ \lambda, \lambda' \in [\Lambda]$ we have $n^{k,\lambda} m^{k,\lambda}_{k',\lambda'} = n^{k',\lambda'} m^{k',\lambda'}_{k,\lambda}$.

Proof: We shall count the set of edges between $V^{k,\lambda}$ and $V^{k',\lambda'}$ in two ways. Since the $n^{k,\lambda}$ vertices of $V^{k,\lambda}$ each have $m^{k,\lambda}_{k',\lambda'}$ edges to neighbors in $V^{k',\lambda'}$, there are $n^{k,\lambda} m^{k,\lambda}_{k',\lambda'}$ edges in total. Counting by the same method, this time in terms of vertices of $V^{k',\lambda'}$, we can also express the number of edges as $n^{k',\lambda'} m^{k',\lambda'}_{k,\lambda}$. Therefore $n^{k,\lambda} m^{k,\lambda}_{k',\lambda'} = n^{k',\lambda'} m^{k',\lambda'}_{k,\lambda}$. □

Theorem 5.2.3: If G, H are a pair of connected association equivalent graphs, then the number of vertices in each with k, λ-association is equal $\forall k \in [\Delta]$, $\forall \lambda \in [\Lambda]$.

Proof: We let $n^{k,\lambda}_G$, $n^{k,\lambda}_H$ denote the number of vertices with k, λ-association in G and H, respectively; because they have the same association profile, we need not make this distinction for the $m^{k,\lambda}$. Before beginning the proof proper, let us outline it. Our goal is to show that $n^{k,\lambda}_G = n^{k,\lambda}_H$ for any valid k and λ. We will first assume that for all k, λ pairs either $n^{k,\lambda}_G \neq n^{k,\lambda}_H$ or else $n^{k,\lambda}_G = n^{k,\lambda}_H = 0$, and will show that this assumption leads to a contradiction. Then, negating this assumption, we know there exists some pair k, λ such that $n^{k,\lambda}_G = n^{k,\lambda}_H \neq 0$. Finally, we will show that the existence of such a pair implies that $n^{k',\lambda'}_G = n^{k',\lambda'}_H$ holds for *any* choice of $k' \in [\Delta]$, $\lambda' \in [\Lambda]$.

Assume, for the sake of contradiction, that $\forall k \in [\Delta]$, $\forall \lambda \in [\Lambda]$ we have either $n^{k,\lambda}_G \neq n^{k,\lambda}_H$ or $n^{k,\lambda}_G = n^{k,\lambda}_H = 0$. Since $V(G)$ is nonempty and G is connected there exists a pair $k \in [\Delta]$, $\lambda \in [\Lambda]$ s.t. $n^{k,\lambda}_G \neq 0$; fix some such k, λ pair. Furthermore, since G is connected there is some pair $i \in [\Delta]$, $j \in [\Lambda]$ s.t. $m^{k,\lambda}_{i,j}$ is nonzero, and therefore $m^{i,j}_{k,\lambda} \neq 0$ also. So then by Lemma 5.2.2 we have

$$n^{k,\lambda}_G m^{k,\lambda}_{i,j} = n^{i,j}_G m^{i,j}_{k,\lambda},$$

$$n^{k,\lambda}_H m^{k,\lambda}_{i,j} = n^{i,j}_H m^{i,j}_{k,\lambda},$$

and taking the difference of these gives

$$m^{k,\lambda}_{i,j}(n^{k,\lambda}_G - n^{k,\lambda}_H) = m^{i,j}_{k,\lambda}(n^{i,j}_G - n^{i,j}_H).$$

By hypothesis we know $n_G^{k,\lambda} - n_H^{k,\lambda} \neq 0$. Supposing without loss of generality that $n_G^{k,\lambda} > n_H^{k,\lambda}$ and recalling that $m_{i,j}^{k,\lambda}$ is positive, we see that the LHS of the above must be positive, and thus the RHS is as well. Because $m_{k,\lambda}^{i,j} > 0$, this implies $n_G^{i,j} > n_H^{i,j}$.

Taking any $k' \in [\Delta]$, $\lambda' \in [\Lambda]$ such that $n_G^{k',\lambda'} \neq 0$, we now repeatedly apply this line of argument to show that $n_G^{k',\lambda'} > n_H^{k',\lambda'}$ always holds. Keeping the k, λ pair from above and taking an arbitrary k', λ' pair, since G is connected it contains a path from a vertex with k, λ-association to one with k', λ'-association. Given such a path $\langle u_1, \ldots, u_r \rangle$ put $d_i := d_G(u_i)$, $l_i := L_G(u_i)$ $(1 \leq i \leq r)$. We immediately get that $n^{d_i,l_i} \neq 0$. And since $u_i \sim_G u_{i+1}$ for any $1 \leq i < r$, we also know that $m_{d_{i+1},l_{i+1}}^{d_i,l_i} > 0$ for each $1 \leq i < r$. So then, employing the same argument as before,

$$n_G^{k,\lambda} > n_H^{k,\lambda}$$

$$\Rightarrow 0 > m_{d_2,l_2}^{k,\lambda}(n_G^{k,\lambda} - n_H^{k,\lambda}) = m_{k,\lambda}^{d_2,l_2}(n_G^{d_2,l_2} - n_H^{d_2,l_2})$$
$$\Rightarrow n_G^{d_2,l_2} > n_H^{d_2,l_2}$$

$$\Rightarrow 0 > m_{d_3,l_3}^{d_2,l_2}(n_G^{d_2,l_2} - n_H^{d_2,l_2}) = m_{d_2,l_2}^{d_3,l_3}(n_G^{d_3,l_3} - n_H^{d_3,l_3})$$
$$\Rightarrow n_G^{d_3,l_3} > n_H^{d_3,l_3}$$

$$\vdots$$

$$\Rightarrow 0 > m_{k',\lambda'}^{d_{r-1},l_{r-1}}(n_G^{d_{r-1},l_{r-1}} - n_H^{d_{r-1},l_{r-1}}) = m_{d_{r-1},l_{r-1}}^{k',\lambda'}(n_G^{k',\lambda'} - n_H^{k',\lambda'})$$
$$\Rightarrow n_G^{k',\lambda'} > n_H^{k',\lambda'}.$$

With this, we have shown that $n_G^{k',\lambda'} \geq n_H^{k',\lambda'}$ for any $k' \in [\Delta], \lambda' \in [\Lambda]$—with equality only if both are zero—and that there exist k, λ such that $n_G^{k,\lambda} > n_H^{k,\lambda}$. From this we get our contradiction, for now we have

$$\sum_{\substack{k \in [\Delta] \\ \lambda \in [\Lambda]}} n_G^{k,\lambda} > \sum_{\substack{k \in [\Delta] \\ \lambda \in [\Lambda]}} n_H^{k,\lambda},$$

but the LHS is exactly $|V(G)|$ and the RHS $|V(H)|$. By hypothesis these are both n, leaving us with the contradiction $n > n$, and so we conclude that $\exists k \in [\Delta], \lambda \in [\Lambda]$

such that $n_G^{k,\lambda} = n_H^{k,\lambda} \neq 0$.

Having shown this, we now proceed to prove that it implies that $n_G^{k,\lambda} = n_H^{k,\lambda}$ for all k, λ. Accordingly, let $k \in [\Delta]$, $\lambda \in [\Lambda]$ be such that $n_G^{k,\lambda} = n_H^{k,\lambda} \neq 0$, and take arbitrary $k' \in [\Delta]$, $\lambda' \in [\Lambda]$. As before, we use a path $\langle u_1, \ldots, u_r \rangle$ in G from a vertex with k, λ-association to one with k', λ'-association, and all of the observations made above about such a path still hold. However, this time we prove the pairwise equality of successive $n_G^{d_i,l_i}, n_H^{d_i,l_i}$ using Lemma 5.2.2:

$$
\begin{aligned}
& & n_G^{k,\lambda} &= n_H^{k,\lambda} \\
\implies & & m_{k,\lambda}^{d_2,l_2} n_G^{d_2,l_2} = m_{d_2,l_2}^{k,\lambda} n_G^{k,\lambda} &= m_{d_2,l_2}^{k,\lambda} n_H^{k,\lambda} = m_{k,\lambda}^{d_2,l_2} n_H^{d_2,l_2} \\
& & m_{k,\lambda}^{d_2,l_2} n_G^{d_2,l_2} &= m_{k,\lambda}^{d_2,l_2} n_H^{d_2,l_2} \\
& & n_G^{d_2,l_2} &= n_H^{d_2,l_2} \\
& & &\vdots \\
\implies & & m_{d_{r-1},l_{r-1}}^{k',\lambda'} n_G^{k',\lambda'} = m_{k',\lambda'}^{d_{r-1},l_{r-1}} n_G^{d_{r-1},l_{r-1}} &= m_{k',\lambda'}^{d_{r-1},l_{r-1}} n_H^{d_{r-1},l_{r-1}} = m_{d_{r-1},l_{r-1}}^{k',\lambda'} n_H^{k',\lambda'} \\
& & m_{d_{r-1},l_{r-1}}^{k',\lambda'} n_G^{k',\lambda'} &= m_{d_{r-1},l_{r-1}}^{k',\lambda'} n_H^{k',\lambda'} \\
& & n_G^{k',\lambda'} &= n_H^{k',\lambda'}.
\end{aligned}
$$

In summary, we have shown that $\exists k \in [\Delta], \lambda \in [\Lambda]$ s.t. $n_G^{k,\lambda} = n_G^{k,\lambda} \neq 0$, and then showed that this implies $n_G^{k',\lambda'} = n_H^{k',\lambda'}$ $\forall k' \in [\Delta], \lambda' \in [\Lambda]$, our desired result. □

Note that in the case where G, H have constant labelings, the consequent of Theorem 5.2.3 reduces to the statement that the graphs have identical degree sequences. The reader should be cautioned that the theorem does *not* state that $G \cong H$.

Chapter 6

n-Switch Sequences

6.1 Definition

An n-switch is a graph transformation that changes the edge set of a graph without affecting the vertices' associations. This 'swap' transformation generalizes a 2-switch, which exchanges edges of a graph while leaving degrees fixed [4, p. 46].

Definition 6.1.1: For $2 \leq n \in \mathbb{Z}$, let $\langle v_0, v_1, \ldots, v_{2n-1}, v_0 \rangle$ be a sequence of vertices in a graph G. Suppose there exist $k, k' \in [\Delta]$ and $\lambda, \lambda' \in [\Lambda]$ such that:

1. $v_i \neq v_{i+1}$ for all i;
2. $v_i \sim_G v_{i+1}$ if and only if i is odd;
3. $v_i \in V^{k,\lambda}$ if i is odd;
4. $v_i \in V^{k',\lambda'}$ if i is even.

Then the graph transformation

$$\sigma : G \longmapsto \langle V, E', L \rangle,$$

where

$$E' := \left(E \smallsetminus \left\{\{v_i, v_{i+1}\} \mid 1 \leq i < 2n,\ i \text{ is odd}\right\}\right) \cup \left\{\{v_j, v_{j+1}\} \mid 1 \leq j < 2n,\ j \text{ is even}\right\}$$

is called an n-**switch** of G. ∎

When the labeling of G is constant and $n = 2$, Definition 6.1.1 coincides with a traditional 2-switch with the additional property that the endpoints of the two edges switched have the same pair of degrees.

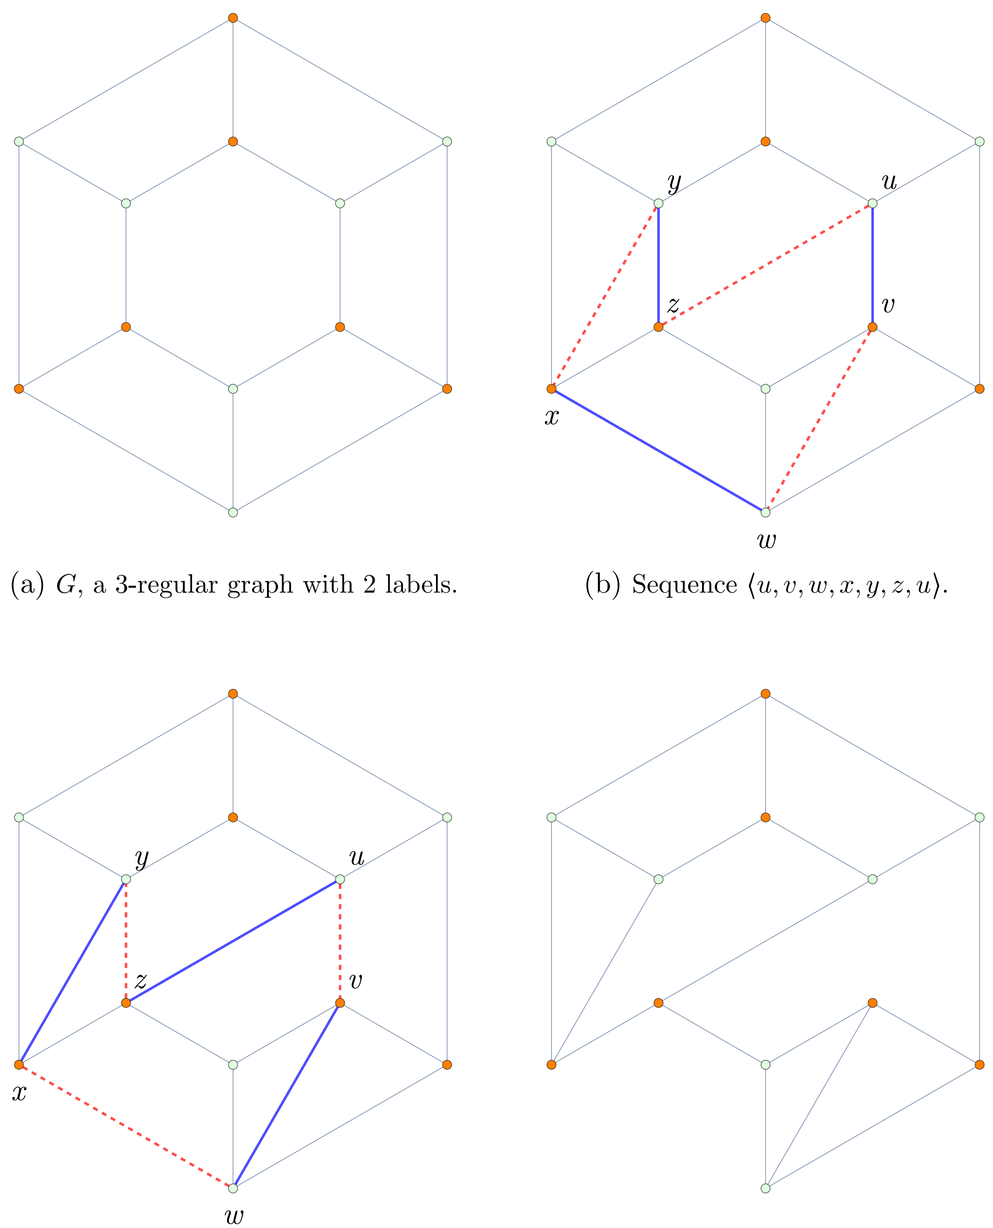

(a) G, a 3-regular graph with 2 labels.

(b) Sequence $\langle u, v, w, x, y, z, u \rangle$.

(c) The sequence after switching.

(d) $\sigma(G)$, graph obtained by switching.

Figure 6.1: An example showing a vertex sequence inducing a 3-switch, and the results of applying that switch. Vertices are colored by label.

Although the above definition might seem technical, its meaning can be intuitively grasped. Basically, the sequence described is like a walk with every other edge missing, and with the vertices alternating between one of two associations. Then the switch is performed by 'rotating' this sequence so that the missing edges are added and vice versa. An example is shown in Figure 6.1. In it, we can see that every other pair of contiguous vertices in the sequence is adjacent, *e.g.* $u \sim v$ but $v \nsim w$. Moreover, letting l_1 represent orange labels and l_2 green labels, the sequence alternates between vertices with $3,1$-association and those with $3,2$-association.

6.2 Association Profiles & Unfolding Trees under n-Switches

In our extended analogy, association equivalence generalizes equality of degree sequences and symmetric-association generalizes k-regularity. So if, as we have hinted, n-switches play the part of 2-switches in that same analogy, we would expect them to preserve association profiles, just as 2-switches preserve degree sequences. We now prove exactly that.

Proposition 6.2.1: If G is a symmetrically-associated graph and σ is an n-switch of G, then $\sigma(G)$ is association equivalent to G.

Proof: Let $\sigma(G) = \langle V, E_\sigma, L\rangle$, and take arbitrary $v \in V$ with k, λ-association. Fix arbitrary $k' \in [\Delta]$, $\lambda' \in [\Lambda]$, and let

$$S_G := \{v' \in \mathcal{N}_G(v) \mid d_G(v') = k', L(v') = l_{\lambda'}\},$$

and define $S_{\sigma(G)}$ analogously. Because an n-switch by definition leaves the vertex set and labeling unaffected, we need only confirm that it also preserves the association profile. In other words, it suffices to show that $|S_G| = |S_{\sigma(G)}|$.

If C is the length n sequence of vertices that induces σ, then if v does not appear in C its neighborhood is unaltered and we are done. Suppose instead that v appears m times in C. By the definition of an n-switch, C contains vertices with two types of associations, not necessarily distinct. Since v is in C, we know one of these types must be the k, λ-association of G. If the other is not the k', λ'-association, then $S_G = S_{\sigma(G)}$ and we are done; assume otherwise.

Due to the construction of E' in Definition 6.1.1, σ causes v to lose m neighbors from $V^{k',\lambda'}$, while at the same time gaining m neighbors from $V^{k',\lambda'}$. Therefore

$|S_{\sigma(G)}| = |S_G| - m + m = |S_G|$. It can now be easily verified that the association profile of G is a valid association profile of $\sigma(G)$. □

Corollary 6.2.2: If v is a vertex of a symmetrically-associated graph G and σ is an n-switch of G, then T_v^d in G is isomorphic to T_v^d in $\sigma(G)$ for any depth d.

Proof: By the proof of Proposition 6.2.1, G is association equivalent to $\sigma(G)$, and v has the same association in both. Since the proof of Theorem 5.1.3 demonstrates the isomorphism of unfolding trees, while only having recourse to the association profile of the graph in question, the current result follows by symmetry. □

Picking up our extended analogy yet again, since any two graphs with the same degree sequence can be changed into each other by iteratively applying 2-switches [2], it should also be the case that two association equivalent graphs can be made identical through repeated n-switches. This can indeed be proved, and in fact the technique we use to do so was inspired by a proof of Hakimi [2, p. 136-37].

Theorem 6.2.3: If G, H are connected symmetrically-associated graphs, then G and H are association equivalent if and only if there exists a finite sequence of n-switches transforming one into the other.

Proof: **(Sufficiency)** By Proposition 6.2.1 the image of a graph under an n-switch is association equivalent to the original graph. The image under multiple n-switches must then also be association equivalent, and so we get the sufficiency direction for free.

(Necessity) Now suppose that $G = \langle V_G, E_G, L_G \rangle$, $H = \langle V_H, E_H, L_H \rangle$ are association equivalent. By Theorem 5.2.3 this implies that $|V_G^{k,\lambda}| = |V_H^{k,\lambda}|$ $\forall k \in [\Delta]$, $\lambda \in [\Lambda]$, and so we may identify each $u \in V_G$ with a unique $v \in V_H$ such that $d_G(u) = d_H(v)$, $L_G(u) = L_H(v)$. Having done so, we can safely commit a helpful abuse of notation by putting $G = \langle V, E_G, L \rangle$, $H = \langle V, E_H, L \rangle$.

Let

$$K := \langle V, E_G \oplus E_H, L \rangle,$$
$$B := E_K \cap E_G,$$
$$R := E_K \cap E_H,$$

where $\oplus$ denotes the symmetric difference. Then K is a 'difference graph' with all edges in either E_G or E_H but not both; E_K is composed of disjoint subsets B, R

of 'blue' edges from G and 'red' edges from H, respectively. In this way, $|E_K|$ can be thought of as an indicator of how different G and H are from each other, with $|E_K| = 0$ only if $G \cong H$. In the following, we partition B and R into subsets of edges connecting vertices with particular associations:

$$B^{k,\lambda}_{k',\lambda'} \coloneqq \{\{v,v'\} \in B \mid v \in V^{k,\lambda}, v' \in V^{k',\lambda'}\} = B^{k',\lambda'}_{k,\lambda},$$

$$R^{k,\lambda}_{k',\lambda'} \coloneqq \{\{v,v'\} \in R \mid v \in V^{k,\lambda}, v' \in V^{k',\lambda'}\} = R^{k',\lambda'}_{k,\lambda}.$$

The lever of our proof is the fact that any vertex v in K is incident to m edges from $B^{k',\lambda'}_{k,\lambda}$ iff it is also incident to m edges from $R^{k',\lambda'}_{k,\lambda}$; otherwise v would have an association in G that differs from its association in H, a contradiction. We refer to this phenomenon as blue-red parity.

We will show that if $|E_K| > 0$, then there exists an n-switch of G that results in a new difference graph K' for which $|E_{K'}| < |E_K|$. Since $|E_K|$ is finite, this suffices to prove our result.

To that end, let $e_0 \coloneqq \{v_0, v_0'\} \in R^{k,\lambda}_{k',\lambda'}$. To simplify notation, hereafter we will write

$$\tilde{B} \coloneqq B^{k,\lambda}_{k',\lambda'}, \qquad \tilde{R} \coloneqq R^{k,\lambda}_{k',\lambda'}.$$

By blue-red parity, we know there exist distinct blue edges $e_1 = \{v_0, v_1\} \in \tilde{B}$ and $e_1' = \{v_0', v_1'\} \in \tilde{B}$. We know v_1, v_1' are distinct from v_0, v_0' since otherwise there would be multiple edges between v_0 and v_0'. However, it is still possible that $v_1 = v_1'$. If they are distinct, blue-red parity again dictates that each is incident to an edge from $\tilde{R}$. If instead $v_1 = v_1'$, then by blue-red parity there must be two edges from $\tilde{R}$ incident to v_1. In either case, call these edges $e_2 = \{v_1, v_2\}$, $e_2' = \{v_1', v_2'\}$.

Our strategy is to continue in steps, forcing the existence of 2 new, distinct edges

$$e_i \coloneqq \{v_{i-1}, v_i\}, \qquad e_i' \coloneqq \{v_{i-1}', v_i'\}$$

at each step. The color alternates from step to step, with $e_i, e_i' \in \tilde{R}$ if i is even and $e_i, e_i' \in \tilde{B}$ otherwise. As mentioned, if $v_i \neq v_i'$ then it is always possible to get $v_{i+1} = v_{i+1}'$. What is more, it is possible that v_i or v_i' has already appeared at an earlier step. However, for our purposes it will not matter whether or not the vertices

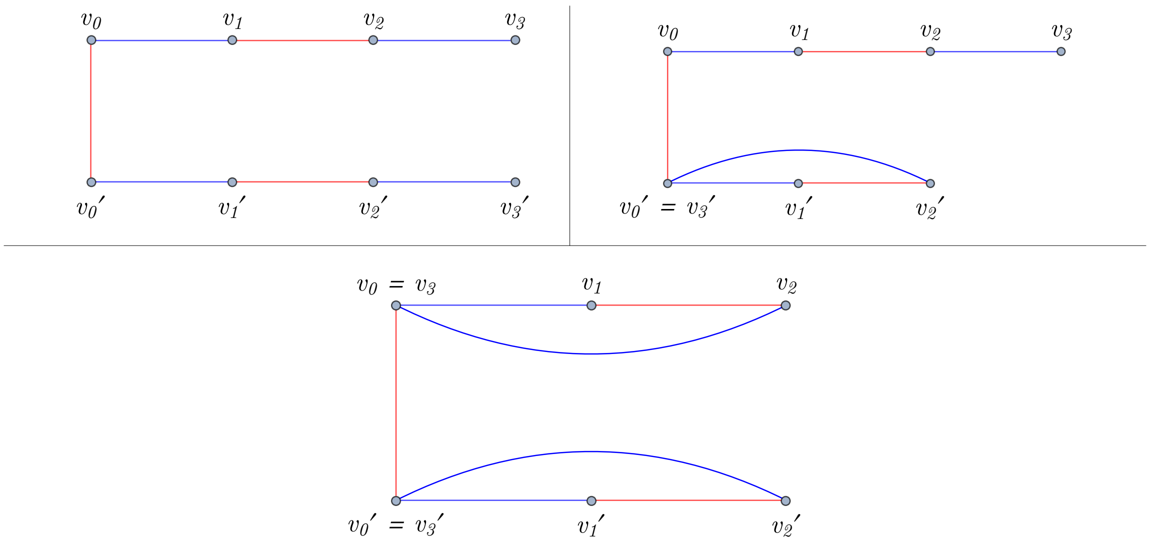

Figure 6.2: Possibilities for v_3, v_3' in the case where $v_3 \neq v_3'$ and $v_3 \neq v_2'$, $v_3' \neq v_2$.

at each step are distinct, and so we will use the names v_i and v_i' always.

Suppose that $v_i \neq v_{i-1}'$ and $v_i' \neq v_{i-1}$. Regardless of the identity of v_i, v_i', we know by blue-red parity that there must be two additional edges from $\tilde{B}$ if i is even, or two from $\tilde{R}$ if i is odd. The only time this is not the case is when we reach some step $n \geq 2$ for which $v_n = v_{n-1}'$, $v_n' = v_{n-1}$, since in this case v_{n-1} and v_{n-1}' fulfill each other's blue-red parity requirement (see Figures 6.2, 6.3 and 6.4 for an example with v_3, v_3'). But K only has finitely many edges, and so eventually we must arrive at such a step n.

Now, form the sequence

$$S := \langle v_0, v_1, \ldots, v_{n-1}, v_{n-1}', v_{n-2}', \ldots, v_0', v_0 \rangle.$$

We claim that this sequence induces an n-switch of G. To that end, rename the vertices of S as $\langle u_0, u_1, \ldots, u_{2n-1}, u_0 \rangle$. First note that, by construction of B and R, whenever a blue edge connects two vertices in S they must be adjacent in G, and whenever a red edge connects them they must be non-adjacent in G. So $u_i \sim_G u_{i+1}$ iff i is odd. Furthermore, all of the edges join a vertex of k, λ-association to one of k', λ'-association, so $u_0 = v_0 \in S^{k,\lambda}$ necessitates that $u_i \in V^{k,\lambda}$ if i is even and $u_i \in V^{k',\lambda'}$ if i is odd.

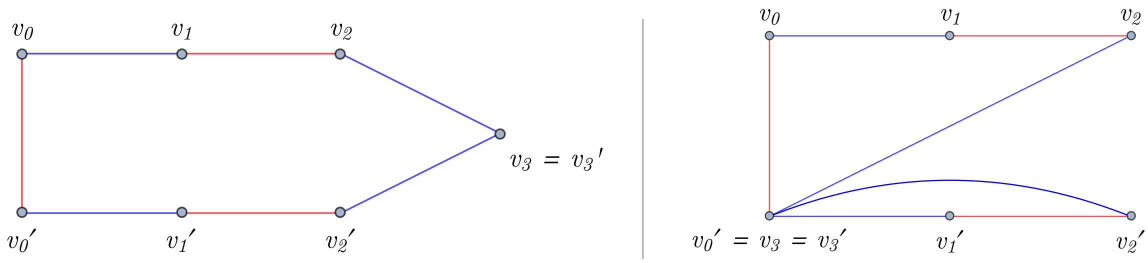

Figure 6.3: Possibilities for v_3, v_3' in the case where $v_3 = v_3'$.

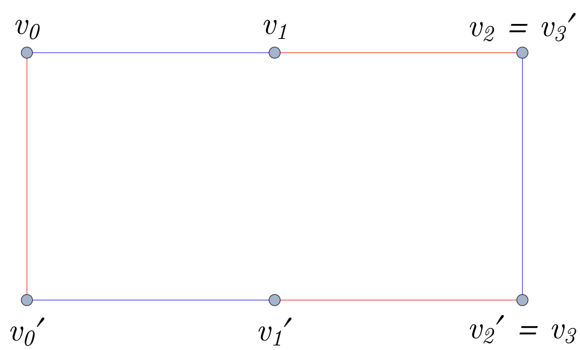

Figure 6.4: The sole possibility for v_3, v_3' in the case where $v_3 = v_2'$, $v_3' = v_2$.

Having established these properties, we know that S yields an n-switch σ of G. Furthermore, the n edges removed from G by σ are not in H, whereas the n edges added were. Thus $|E_{K'}| = |E_K| - 2n$, where $K' := \langle V, E_{\sigma(G)} \oplus E_H, L \rangle$, and we are done. $\square$

Because n-switches preserve associations, the $G_j := \sigma(G_{j-1})$ obtained at the jth step in the switch sequence from G to H will have the same association profile as G.

An example will help illustrate Theorem 6.2.3. Figure 6.5 below displays two non-isomorphic association equivalent graphs G and H, and Figure 6.6 shows their difference graph. As can be seen from the latter figure, there are two blue-red alternating cycles S_1 and S_2. So these cycles induce a sequence of two n-switches (one 3-switch and one 2-switch, induced by S_1 and S_2, respectively) taking G to H. The subfigures of Figure 6.7 show the intermediary graphs obtained by applying this sequence, with subfigure (b) isomorphic to H.

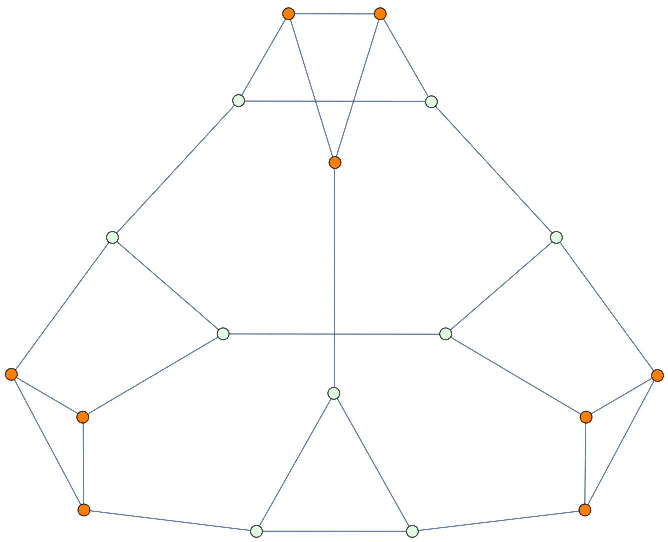

(a) Drawing of graph G.

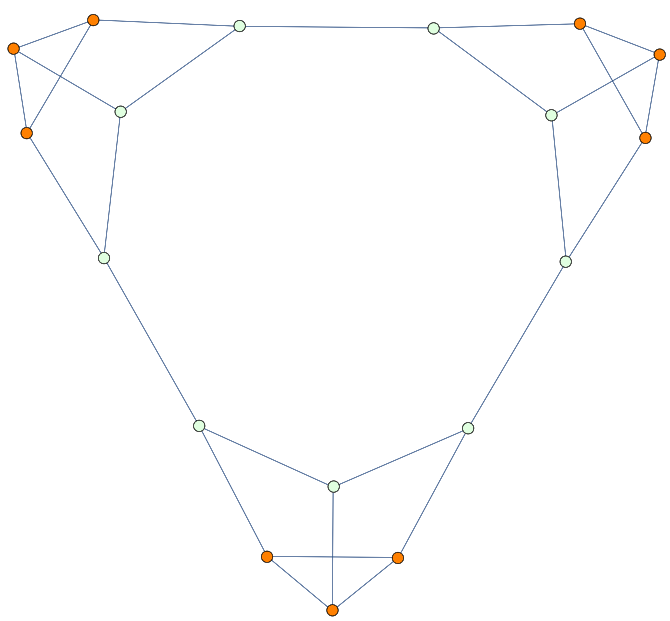

(b) Drawing of graph H.

Figure 6.5: A pair of non-isomorphic association equivalent graphs. Vertex colors represent labels.

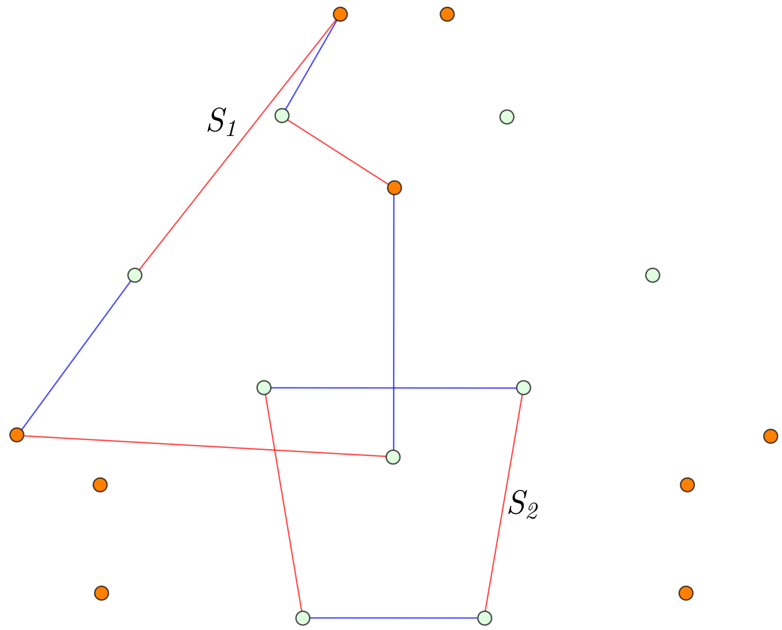

Figure 6.6: Difference graph K with edge set $E_G \oplus E_H$; the blue edges are from G and the red edges from H. Vertices are arranged as in the drawing of G.

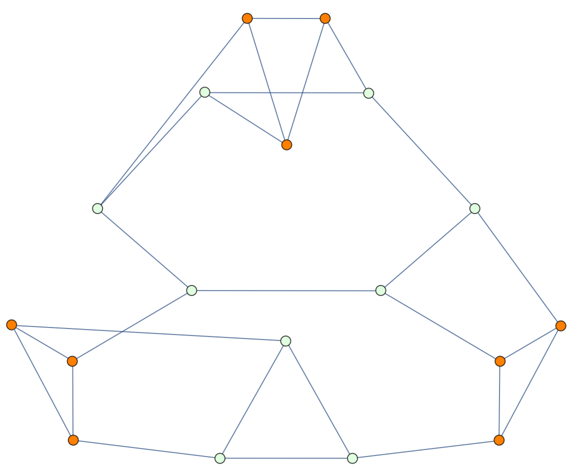

(a) Result of applying the 3-switch induced by S_1 to G.

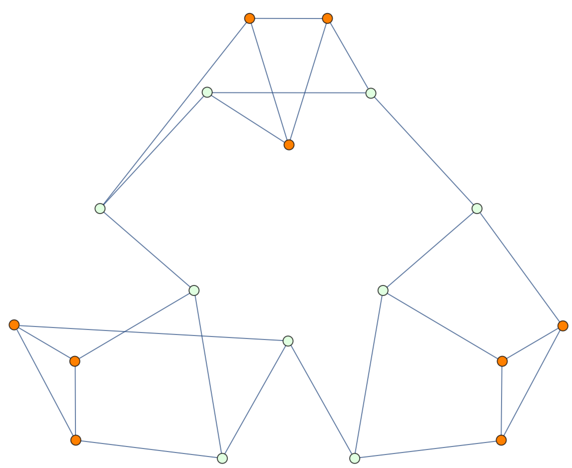

(b) Graph isomorphic to H, obtained by applying the 2-switch induced by S_2 to the graph in (a).

Figure 6.7: Results of applying the n-switch sequence to G.

Chapter 7

Directions for Further Research

In conclusion, we summarize what we view as the most promising areas for future research suggested by this work. Of course, since unfolding trees and symmetrically-associated graphs have not received attention in pure mathematics until this point, the range of possible topics is too great to exhaust here. Nonetheless, what follows are those aspects that most intrigued (or vexed) us during the course of our research.

As mentioned in Section 4.2, there are numerous examples of graphs with pairs of vertices that are unfolding equivalent but not similar. Much of the earlier work on this thesis was dedicated to finding a meaningful constraint that would make unfolding equivalence coincide with similarity, but this proved to be a much more difficult endeavor than it initially seemed. Of course, there are trivial constraints that make this the case (such as requiring that all vertices have different labels and thus making all the unfolding equivalence classes singletons), but these fail to impart insight into the patterns at work in unfolding trees.

Another possible avenue is that of topology. Because unfolding trees derive from the connective structure of a graph, one might reasonably suppose that they encode topological information of some sort. Because the tools of topology required to test such suppositions fell outside the scope of the present work, we did not pursue such a course. However, we passed by a number of interesting topological questions. For example, it appears to be the case that every graph is a quotient space of each of its unfolding trees. Additionally, it might be the case that symmetric-association can be characterized in topological terms.

Appendix A

Elementary Graph Theory

For elegance of presentation, the definitions and notation found in Chapter 2 are those for which there is no obvious extension of definitions from elementary graph theory to labelled graphs. Here we go over some fundamentals of graph theory, adapted to the context of labelled graphs.

First are definitions concerning the vertices of a labelled graph.

Definition A.1: In a graph $G = \langle V, E, L \rangle$ two vertices u, v are **adjacent** if there is an edge $\{u, v\} \in E$ between them. The **neighborhood** $\mathcal{N}_G(v)$ of v is the set of vertices adjacent to v. The **degree** of v is $d_G(v) := |\mathcal{N}_G(v)|$, the number of vertices adjacent to v in G. ■

Notation A.2: We denote $\max_{v \in V} d_G(v)$, the maximum degree of any vertex in G, by the symbol $\Delta(G)$. ■

Definition A.3: The **distance** between two vertices u, v of a graph G is the length of the shortest path (see Definition 2.1.3) between them. ■

Related to paths (Definition 2.1.3) is the notion of connectedness, a concept essential to graph theory.

Definition A.4: A graph G is said to be **connected** if for every pair of vertices u, v in G there exists a path from u to v. ■

In other words, a connected graph is one in which each vertex can be 'reached' from every other vertex by moving along edges.

Next we have subgraphs, which are the natural substructure for graphs (akin to subgroups in group theory or subspaces in topology).

Definition A.5: Given $G = \langle V_G, E_G, L_G \rangle$, we say that a labelled graph $H = \langle V_H, E_H, L_H \rangle$ is a **subgraph** of G if

$$V_H \subseteq V_G,\ E_H \subseteq E_G,\ L_H \subseteq L_G.$$

The binary relation $H \leq G$ denotes that H is a subgraph of G. We sometimes wish to communicate that $H \neq G$, and in this case we call H a **proper** subgraph of G, writing $H < G$. ∎

Occasionally we wish to single out a subgraph that only contains a subset of vertices from the larger graph. This intuitively natural object is formalized by the definition of an induced subgraph.

Definition A.6: For some subset S of the vertex set of a graph G, the **subgraph of G induced by S** is

$$G[S] \coloneqq \langle V', E', L' \rangle,$$

where

$$\begin{aligned} V' &\coloneqq S, \\ E' &\coloneqq \{e \in E_G \mid e \subseteq S\}, \\ L' &\coloneqq L \restriction_S . \end{aligned}$$

We call $G[S]$ an **induced subgraph**. ∎

Appendix B

Alternate Definitions & Results for Unlabelled Graphs

HERE we restate the important definitions and proofs from this thesis, adapted to apply to unlabelled graphs. In this appendix, the terms 'graph,' 'isomorphism,' and 'automorphism' all refer to their unlabelled versions, and the symbol $\cong$ denotes unlabelled graph isomorphism.

Definition 3.1.1*: Let $G = \langle V, E \rangle$ be a graph, let u be any vertex of G, and take any $d \in \mathbb{N}$. The **unfolding tree** T_u^d of depth d at u is a graph defined as follows. The vertex set of T_u^d is given by

$$V_T := \{w = \langle u, v_2, \ldots, v_m \rangle \in V^m \mid w \text{ is a walk in } G \text{ and } \ell(w) \leq d\},$$

the set of walks starting at u of length at most d. The edge set E_T of the unfolding tree is induced by the following adjacency rule: for any $w, w' \in V_T$ where $\ell(w) < \ell(w')$,

$$w \sim_{T_u^d} w' \quad \Longleftrightarrow \quad w = w'[\ell(w') - 1].$$

■

Definition 3.3.1*: For any vertices u, v of a graph G, suppose that for all $d \geq 0$ we have $T_u^d \cong T_v^d$ and $\psi_d(\langle u \rangle) = \langle v \rangle$, where ψ_d is the isomorphism from T_u^d to T_v^d. When this is the case, we say that u and v are **unfolding equivalent**, writing $u \risingdotseq v$. ■

Theorem 4.1.1*: Let u, v be vertices of a graph G. Then $u \doteq v \Rightarrow u \risingdotseq v$.

Proof: Let $G = \langle V, E\rangle$ be a graph and let $u, v \in V$. We denote the vertex set of T_u^d by V_u and its edge set by E_u, and use a similar convention for T_v^d.

Suppose that $\phi : V \rightarrowtail\!\!\!\!\twoheadrightarrow V$ is an automorphism such that $\phi(u) = v$. We claim that the mapping $\psi : V_u \to V_v$, defined by

$$\psi\big(\langle u, x_2, x_3, \ldots, x_k\rangle\big) \coloneqq \langle v, \phi(x_2), \phi(x_3), \ldots, \phi(x_k)\rangle,$$

is an isomorphism of T_u^d and T_v^d.

Suppose that $w \sim_{T_u^d} w'$ for some $w, w' \in V_u$; assume without loss of generality that the walk corresponding to w' is longer than that of w, and put $w' \coloneqq \langle u, w_1, w_2, \ldots, w_m\rangle$ where $m \leq d$. Then by the definition of an unfolding tree it must be the case that $w = \langle u, w_1, \ldots, w_{m-1}\rangle$. Now consider the images of these under ψ, *i.e.* the vertices $\psi(w) = \langle v, \phi(w_1), \ldots, \phi(w_{m-1})\rangle$ and $\psi(w') = \langle v, \phi(w_1), \ldots, \phi(w_m)\rangle$ of T_v^d; we must show that $\psi(w) \sim_{T_v^d} \psi(w')$. Since $\psi(w')$ extends the tuple $\psi(w)$ by one vertex, we need only prove that both of $\psi(w), \psi(w')$ are walks in G of length at most d starting at v. Showing that they are walks is simple since—given that w' is by hypothesis a walk in G—we know that $w_i \sim_G w_{i+1}$ for all $1 \leq i < m$, and therefore $\phi(w_i) \sim_G \phi(w_{i+1})$ because ϕ preserves adjacency. Moreover, since $\phi(u) = v$ we know that $\psi(w), \psi(w')$ are walks starting at v. Finally, because w' and $\psi(w')$ are of the same length $m < d$, we know that $\psi(w')$ has length at most d, and similarly for $\psi(w)$. Thus $\psi(w), \psi(w')$ are adjacent vertices of T_v^d, and so $w \sim_{T_u^d} w' \Rightarrow \psi(w) \sim_{T_v^d} \psi(w')$. Because ϕ is an isomorphism, it admits an inverse which is also an isomorphism, and thus $\psi(w) \sim_{T_v^d} \psi(w') \Rightarrow w \sim_{T_u^d} w'$ also. So ψ preserves adjacency, and thus is an isomorphism. □

Definition 5.1.1*: We say a graph G is **symmetrically-associated** if $\forall k \in [\Delta]$ there exists a sequence of non-negative integers

$$m^k = \langle m_1^k,\ m_2^k, \ldots,\ m_\Delta^k\rangle$$

such that $\forall v \in V$ with $d(v) = k$ we have

$$m_j^k = \big|\,\{u \in \mathcal{N}(v) \mid d(u) = j\}\,\big|.$$

We impose the additional condition that $m^k = \langle 0, 0, \ldots, 0\rangle$ whenever G has no vertices of degree k.

The sequence m^k is called the **k-association** of G. When discussing a particular vertex v of degree k, we will occasionally refer to m^k simply as the **association** of v. The set $\{m^k \mid k \in [\Delta]\}$ of all k-associations of G is called the **association profile** of G. ■

Theorem 5.1.3*: Let G be a symmetrically-associated graph. Then for all vertices u, v of G,

$$d(u) = d(v) \iff u \doteqdot v.$$

Proof: **(Sufficiency)** Assuming $u \doteqdot v$, we know that there is an isomorphism ψ from T_u^1 to T_v^1. Because ψ is bijective, by Proposition 3.3.2 we know $|T_u^1[\![1]\!]| = |T_v^1[\![1]\!]|$. Furthermore, the size of $T_u^1[\![1]\!]$ equals the degree of u in G (and similarly for $T_v^1[\![1]\!]$ and $d_G(v)$), so combining this with the equality in the previous sentence gives $d_G(u) = d_G(v)$.

(Necessity) We use induction on the unfolding trees' depths to prove a stronger result: given the hypotheses, for any depth D there exists an isomorphism ψ_D from T_u^D to T_v^D with the property that, $\forall r \geq 0$,

$$\psi_D(w) = w' \implies d_G(w[\![r]\!]) = d_G(w'[\![r]\!]).$$

For the base case $D = 0$, we put $\psi_0\big(\langle u \rangle\big) \coloneqq \langle v \rangle$. This is clearly a bijection, and there are no edges to check, so ψ_0 is an isomorphism. Moreover, by hypothesis we know $d(u) = d(v)$.

Now let $D \geq 1$ and suppose $\exists \psi_{D-1} : V(T_u^{D-1}) \rightarrowtail\!\!\!\!\twoheadrightarrow V(T_v^{D-1})$ a isomorphism satisfying our degree-preservation requirement. Since $T_u^{D-1} \leq T_v^{D-1}$, we construct ψ_D by first letting $\psi_D \restriction_{T_u^{D-1}} \coloneqq \psi_{D-1}$. We now need only define ψ_D on the set of leaves $T_u^D[\![D]\!]$.

Call $T_u^D[\![D-1]\!] \eqqcolon \{w_1, \ldots, w_m\}$, and let $w_i' \coloneqq \psi_{D-1}(w_i)$, noting that $T_v^D[\![D-1]\!] = \{w_i'\}_{i \in [m]}$ by Proposition 3.3.2. In addition, define

$$\begin{aligned} S_i &\coloneqq \mathcal{C}(w_i), \\ S_i' &\coloneqq \mathcal{C}(w_i'), \\ k_i &\coloneqq d_G(w[\![D-1]\!]). \end{aligned}$$

Since each vertex of $T_u^D[\![D]\!]$ has a unique mother in $T_u^D[\![D-1]\!]$, the S_i partition $T_u^D[\![D]\!]$. We proceed by defining ψ_D on each S_i.

We can write every S_i as the union of disjoint subsets

$$S_i = \bigcup_{k \in [\Delta]} S_{i,k},$$

where $S_{i,k} := \{x \in S_i \mid d_G(x[\![D]\!]) = k\}$. Additionally, we can express S_i' as a disjoint union of subsets $S_{i,k}'$ defined analogously to the $S_{i,k}$. Observe that

$$\mathcal{N}_G(w_i[\![D-1]\!]) = \{x[\![D]\!] \in V_G \mid x \in S_i\},$$

which is to say that each vertex in S_i corresponds to a neighbor of $w_i[\![D-1]\!]$ in G. Therefore $|S_{i,k}| = m_k^{k_i}$ for all $k \in [\Delta]$. By the inductive hypothesis,

$$\psi_{D-1}(w_i) = w_i' \quad \Longrightarrow \quad d_G(w_i'[\![D-1]\!]) = k_i.$$

So since G is symmetrically-associated, we know $w'[\![D-1]\!]$ has the same association as $w[\![D-1]\!]$. Therefore $|S_{i,k}| = |S_{i,k}'|$, and so there exists a bijection $\theta_{i,k}$ from $S_{i,k}$ to $S_{i,k}'$; fix such a $\theta_{i,k}$ for each $S_{i,k}$.

We define ψ_D on $T_u^D[\![D]\!]$ as follows. Given $x \in T_u^D[\![D]\!]$ such that $x \in S_{i,k}$, we put

$$\psi_D(x) := \theta_{i,k}(x).$$

Before proceeding, we must show that this definition describes a function. Since

$$T_u^D[\![D]\!] = \bigcup_{i \in [m]} S_i = \bigcup_{i \in [m]} \bigcup_{k \in [\Delta]} S_{i,k},$$

and since the $S_{i,k}$ are pairwise disjoint, we know that each leaf of T_u^D is contained in a unique $S_{i,k}$. Moreover, the subset $S_{i,k}$ has a unique function $\theta_{i,k}$ associated with it. So the value $\psi_D(x)$ exists and is unique for each $x \in T_u^D[\![D]\!]$. This means that ψ_D is a function on $T_u^D[\![D]\!]$, and since $\psi_D \restriction_{T_u^{D-1}} = \psi_{D-1}$ is a function on the rest of T_u^D we conclude that the above definition yields a function $\psi_D : T_u^D \to T_v^D$.

Confirming that ψ_D is an isomorphism, we first show that it is a bijection. By construction,

$$\begin{aligned}\psi_D[T_u^D[\![D]\!]] &= \bigcup_{i\in[m]}\bigcup_{k\in[\Delta]}\theta_{i,k}[S_{i,k}] \\ &\subseteq \bigcup_{i\in[m]}\bigcup_{k\in[\Delta]} S'_{i,k} \\ &= T_v^D[\![D]\!],\end{aligned}$$

and therefore

$$\begin{aligned}\left(\psi_D[T_u^D[\![D]\!]]\right)\cap\left(\psi_D[T_u^{D-1}]\right) &= \left(\psi_D[T_u^D[\![D]\!]]\right)\cap\left(T_v^{D-1}\right) \\ &\subseteq \left(T_v^D[\![D]\!]\right)\cap\left(T_v^{D-1}\right) \\ &= \varnothing.\end{aligned}$$

That is to say, the image of T_v^{D-1} under ψ_D is disjoint from that of $T_v^D[\![D]\!]$. This, together with the inductive hypothesis that $\psi_D\restriction_{T_u^{D-1}} = \psi_{D-1}$ is bijective, implies that it suffices to show that $\psi_D\restriction_{T_u^D[\![D]\!]}$ is a bijection from $T_u^D[\![D]\!]$ to $T_v^D[\![D]\!]$. Now, suppose $\exists x, y \in T_u^D[\![D]\!]$ s.t. $\psi_D(x) = \psi_D(y)$. Letting $S_{i,k}$ denote the unique subset containing x, we see that $y \in S_{i,k}$ as well; for if, on the contrary, we had $y \in S_{j,l} \neq S_{i,k}$ it would follow that both

$$\psi_D(y) = \theta_{j,l}(y) \in S'_{j,l},$$

and

$$\psi_D(y) = \psi_D(x) = \theta_{i,k}(x) \in S'_{i,k},$$

contradicting the fact that $S'_{j,l}$, $S'_{i,k}$ are disjoint. So $x, y \in S_{i,k}$ implies that

$$\theta_{i,k}(x) = \psi_D(x) = \psi_D(y) = \theta_{i,k}(y),$$

and so since $\theta_{i,k}$ is injective we conclude that $x = y$. Therefore $\psi_D\restriction_{T_u^D[\![D]\!]}$ is injective.

Next, recalling that $T_u^D[\![D]\!]$ is the union of the pairwise disjoint subsets S_i (and similarly for $T_v^D[\![D]\!]$ and the S'_i), and that

$$|S_i| = d_G(w_i[\![D-1]\!]) = d(w'_i[\![D-1]\!]) = |S'_i|,$$

we deduce that

$$\begin{aligned} |T_u^D[\![D]\!]| &= \sum_{i\in[m]} |S_i| \\ &= \sum_{i\in[m]} |S_i'| \\ &= |T_v^D[\![D]\!]|. \end{aligned}$$

So $\psi_D \restriction_{T_u^D[\![D]\!]}$ is an injection with domain and codomain of equal cardinality, and therefore is a bijection.

Next, we check that $x \sim_{T_u^D} y \Leftrightarrow \psi_D(x) \sim_{T_v^D} \psi_D(y)$. Again, since ψ_{D-1} is an isomorphism, we need only check edges incident to vertices in $T_u^D[\![D]\!]$. To that end, suppose we have $x \in T_u^D$, $y \in T_u^D[\![D]\!]$ s.t. $x \sim_{T_u^D} y$. Since y belongs to the deepest level, applying Proposition 3.2.2 gives

$$\mathcal{N}(y) \subseteq T_u^D[\![D-1]\!] \cup T_u^D[\![D+1]\!] = T_u^D[\![D-1]\!]$$

and thus x must be in level $D-1$. So $x = w_i$ for some $i \in [m]$, implying that $y \in S_i$. It follows that $\psi_D(y) \in S_i'$, which by definition of S_i' means that $\psi_D(y) \sim_{T_v^D} \psi_D(x)$. Turning to the other direction of the biconditional, assume instead that $x' \sim_{T_v^D} y'$ for some $x' \in T_v^D$, $y' \in T_v^D[\![D]\!]$. By the same logic as above we get that $x' = w_i'$ for some $i \in [m]$ and so $y' \in S_i'$. Then since $\psi_D^{-1}[S_i'] = S_i$, we have $\psi_D^{-1}(y') \in S_i \Rightarrow \psi_D^{-1}(y') \sim_{T_u^D} w_i = \psi_D^{-1}(x')$.

Finally, we show that ψ_D satisfies our additional stipulation

$$\psi_D(w) = w' \quad \Longrightarrow \quad d_G(w[\![r]\!]) = d_G(w'[\![r]\!]).$$

As before, since ψ_{D-1} has this property by inductive hypothesis, we check only leaves. Recall that any unfolding tree vertex x in $S_{i,k}$ maps to one in $S_{i,k}'$. By construction of those subsets, this implies that the last coordinates of x and $\psi_D(x)$ both have degree k in G, and we are done. □

Definition 6.1.1*: For $2 \le n \in \mathbb{Z}$, let $\langle v_0, v_1, \ldots, v_{2n-1}, v_0 \rangle$ be a sequence of vertices in a graph G. Suppose there exist $j, k \in [\Delta]$ such that:

1. $v_i \neq v_{i+1}$ for all i;
2. $v_i \sim_G v_{i+1}$ if and only if i is odd;

3. $d(v_i) = j$ if i is odd;

4. $d(v_i) = k$ if i is even.

Then the graph transformation

$$\sigma : G \longmapsto \langle V, E', L \rangle,$$

where

$$E' := \left(E \setminus \left\{\{v_i, v_{i+1}\} \mid 1 \leq i < 2n,\ i \text{ is odd}\right\}\right) \cup \left\{\{v_j, v_{j+1}\} \mid 1 \leq j < 2n,\ j \text{ is even}\right\}$$

is called an n-**switch** of G. ■

Proposition 6.2.1*: If G is a symmetrically-associated graph and σ is an n-switch of G, then $\sigma(G)$ is association equivalent to G.

Proof: The proof is nearly identical to that of the original result, the only difference being that we replace instances of 'k, λ-association' with 'k-association.' □

Proposition 6.2.2*: If v is a vertex of a symmetrically-associated graph G and σ is an n-switch of G, then T_v^d in G is isomorphic to T_v^d in $\sigma(G)$ for any depth d.

Proof: By the proof of Proposition 6.2.1*, G is association equivalent to $\sigma(G)$, and v has the same association in both. Since the proof of Theorem 5.1.3* demonstrates the isomorphism of unfolding trees, while only having recourse to the association profile of the graph in question, the current result follows by symmetry. □

Theorem 6.2.3*: If G, H are symmetrically-associated connected graphs, then G and H are association equivalent if and only if there exists a finite sequence of n-switches transforming one into the other.

Proof: By Proposition 6.2.1*, n-switches respect association equivalence, and so we get the sufficiency direction for free.

Now suppose that $G = \langle V_G, E_G \rangle$, $H = \langle V_H, E_H \rangle$ have the same association profile. By Theorem 5.2.3* this implies that $|V_G^k| = |V_H^k|\ \forall k \in [\Delta]$, and so we may identify each $u \in V_G$ with a unique $v \in V_H$ such that $d_G(u) = d_H(v)$. Having done so, we can safely commit a helpful abuse of notation by putting $G = \langle V, E_G \rangle$, $H = \langle V, E_H \rangle$.

Let

$$K := \langle V, E_G \oplus E_H \rangle,$$
$$B := E_K \cap E_G,$$
$$R := E_K \cap E_H,$$

where $\oplus$ denotes the symmetric difference. Then K is a 'difference graph' with all edges in either E_G or E_H but not both; E_K is composed of disjoint subsets B, R of 'blue' edges from G and 'red' edges from H, respectively. In this way, $|E_K|$ can be thought of as an indicator of how different G and H are from each other, with $|E_K| = 0$ only if $G \cong H$. In the following, we partition B and R into subsets of edges connecting vertices with particular associations:

$$B^k_{k'} := \{\{v, v'\} \in B \mid v \in V^k, v' \in V^{k'}\} = B^{k'}_k,$$

$$R^k_{k'} := \{\{v, v'\} \in R \mid v \in V^k, v' \in V^{k'}\} = R^{k'}_k.$$

The lever of our proof is the fact that any vertex v in K is incident to m edges from $B^{k'}_k$ iff it is also incident to m edges from $R^{k'}_k$; otherwise v would have an association in G that differs from its association in H, a contradiction. We refer to this phenomenon as blue-red parity.

We will show that if $|E_K| > 0$, then there exists an n-switch of G that results in a new difference graph K' for which $|E_{K'}| < |E_K|$. Since $|E_K|$ is finite, this suffices to prove our result.

To that end, let $e_0 := \{v_0, v'_0\} \in R^k_{k'}$. To simplify notation, hereafter we will write

$$\tilde{B} := B^k_{k'}, \qquad \tilde{R} := R^k_{k'}.$$

By blue-red parity, we know there exist distinct blue edges $e_1 = \{v_0, v_1\} \in \tilde{B}$ and $e'_1 = \{v'_0, v'_1\} \in \tilde{B}$. We know v_1, v'_1 are distinct from v_0, v'_0 since otherwise there would be multiple edges between v_0 and v'_0. However, it is still possible that $v_1 = v'_1$. If they are distinct, blue-red parity again dictates that each is incident to an edge from $\tilde{R}$. If instead $v_1 = v'_1$, then by blue-red parity there must be two edges from $\tilde{R}$ incident to v_1. In either case, call these edges $e_2 = \{v_1, v_2\}$, $e'_2 = \{v'_1, v'_2\}$.

Our strategy is to continue in steps, forcing the existence of 2 new, distinct edges

$$e_i := \{v_{i-1}, v_i\}, \qquad e'_i := \{v'_{i-1}, v'_i\}$$

at each step. The color alternates from step to step, with $e_i, e_i' \in \tilde{R}$ if i is even and $e_i, e_i' \in \tilde{B}$ otherwise. As mentioned, if $v_i \neq v_i'$ then it is always possible to get $v_{i+1} = v_{i+1}'$. What is more, it is possible that v_i or v_i' has already appeared at an earlier step. However, for our purposes it will not matter whether or not the vertices at each step are distinct, and so we will use the names v_i and v_i' always.

Suppose that $v_i \neq v_{i-1}'$ and $v_i' \neq v_{i-1}$. Regardless of the identity of v_i, v_i', we know by blue-red parity that there must be two additional edges from $\tilde{B}$ if i is even, or two from $\tilde{R}$ if i is odd. The only time this is not the case is when we reach some step $n \geq 2$ for which $v_n = v_{n-1}'$, $v_n' = v_{n-1}$, since in this case v_{n-1} and v_{n-1}' fulfill each other's blue-red parity requirement. But K only has finitely many edges, and so eventually we must arrive at such a step n.

Now, form the sequence

$$S := \langle v_0,\, v_1, \ldots,\, v_{n-1},\, v_{n-1}',\, v_{n-2}', \ldots,\, v_0',\, v_0 \rangle.$$

We claim that this sequence induces an n-switch of G. To that end, rename the vertices of S as $\langle u_0, u_1, \ldots, u_{2n-1}, u_0 \rangle$. First note that, by construction of B, R, whenever a blue edge connects two vertices in S they must be adjacent in G, and whenever a red edge connects them they must be non-adjacent in G. So $u_i \sim_G u_{i+1}$ iff i is odd. Furthermore, all of the edges join a vertex of k-association to one of k'-association, so $u_0 = v_0 \in S^k$ necessitates that $u_i \in V^k$ if i is even and $u_i \in V^{k'}$ if i is odd.

Having established these properties, we know that S yields an n-switch σ of G. Furthermore, the n edges removed from G by σ are not in H, whereas the n edges added are. Thus $|E_{K'}| = |E_K| - 2n$, where $K' := \langle V, E_{\sigma(G)} \oplus E_H, L \rangle$, and we are done. $\square$

Bibliography

[1] Gallian, J.A. A dynamic survey of graph labeling. *The Electronic Journal of Combinatorics* 16, no. 6 (2009): 1-219.

[2] Hakimi, S.L. On realizability of a set of integers as degrees of the vertices of a linear graph. *Journal of the Society for Industrial & Applied Mathematics* 10, no. 3 (1962): 496-506.

[3] Scarselli, F., M. Gori, A.C. Tsoi, M. Hagenbuchner, and G. Monfardini. Computational capabilities of graph neural networks. *IEEE Transactions on Neural Networks* 20, no. 1 (2009): 81-102.

[4] West, D.B. *Introduction to graph theory.* Vol. 2. Upper Saddle River: Prentice hall, 2001.