ROCKY MOUNTAIN JOURNAL
OF MATHEMATICS
Volume 54 (2024), No. 3, 631–674

A DUALITY FOR LABELED GRAPHS AND FACTORIZATIONS WITH APPLICATIONS TO GRAPH EMBEDDINGS AND HURWITZ ENUMERATION

NIKOS APOSTOLAKIS

The set of factorizations of permutations into m transpositions of some symmetric group $\mathcal{S}_n$ is naturally in bijection with the set of graphs of order n and size m with both edges and vertices labeled. We define a notion of duality (the *mind-body duality*) for factorizations and such labeled graphs and interpret it in terms of properly embedded graphs, a class of graphs embedded in a bounded compact oriented surface with all the vertices lying in the boundary, and show a close connection of this duality with the Hurwitz action of the braid group. Connections with the theory of cellularly embedded graphs are highlighted and hints of possible applications are given. In this paper we focus on developing the necessary theory, leaving specific applications and further developments for future projects.

1. Introduction

The research in this paper originated with the author reading [27] and attempting to understand the notion of duality implicit in the construction of the "*structural bijection*" defined there. That duality is closely related to the "duality" (the quotation marks are there because this "duality" is not involutory, a natural requirement in order to call any bijection a duality) for noncrossing trees defined in [29], and the structural properties of the bijection follow from the properties of that duality, namely from the fact that the neighborhood of a vertex is transformed to a path that, together with an arc of the boundary circle, forms the boundary of one of the region that the disk of the noncrossing tree is divided into by the tree. An exposition of the contents of [27] from this point of view is given in Section 5. A question posed in [27] is to find generalizations of the structural bijection, defined there for minimal factorizations of cycles, to more general classes of factorizations. We answer this question by making explicit, and clarifying, the implicit duality, which we call "mind-body duality" for reasons explained in Section 2C, and generalizing it so that it applies to factorizations of any permutation with any number of factors. We leave specific applications to bijective enumerations to future projects, see Section 6 for a sample of such projects planned for the near future.

Part of the work in Sections 2 and 5A was developed jointly with Kerry Ojakian. The computer algebra systems Sage [39] and GAP [26] were used extensively to confirm calculations and check conjectures at several stages of this project. It is my pleasure to thank Nichole McDaniel for help with drawing some of the figures. Finally I gratefully acknowledge the partial support of PSCCUNY Research Award TRADA-48-526.

2020 AMS *Mathematics subject classification:* primary 05C10; secondary 20F36, 57M12, 57M15.

Keywords and phrases: properly embedded graph, factorizations edge-labeled graphs, mind-body duality, Hurwitz action, Garside element of braid group, cellularly embedded graph, perfect trail double cover, medial digraph, perfect chain decomposition, self-dual embeddings, up-down numbers.

Received by the editors on December 6, 2020, and in revised form on February 6, 2023.

DOI: 10.1216/rmj.2024.54.631

By *factorization* we mean an expression of a permutation as a product of a sequence of transpositions. As observed by Dénes in [17], factorizations of permutations of a finite set V with m factors, are in bijection with graphs with vertex set V and m edges labeled by $[m] := \{1, \ldots, m\}$. This bijection assigns to a factorization a graph that has an edge connecting v and u labeled by i, if and only if, the i-th factor interchanges v and u. The mind-body duality is first defined in Section 2A for *edge-labeled graphs* using the fact that a labeling of the edges of a graph induces two dual structures on the graph: a *local edge order* (leo), and a *perfect trail double cover* (PTDC), see Definitions 2.3 and 2.4. The leo is simply the linear orders induced by the edge-labels on the set of edges incident at each vertex, while, in terms of factorizations, the PTDC is the set of trails formed by the trajectories of V under the successive application of the factors. The mind-body dual of an edge-labeled graph is then defined as the *intersection graph* of the family of these trajectories viewed as sets of edges, see Definition 2.8.

The study of factorizations, branched coverings, and their enumeration is a classical topic that goes back to Hurwitz (see [30]). The braid group, as a group of automorphisms of the free group seems to have appeared for the first time in that paper as well. It turns out that mind-body duality has a simple interpretation in terms of the Hurwitz action on the set of factorizations: it is simply the reverse of the action of the Garside element of the braid group (see Theorem 3.4), and this allows us to get explicit formulas for the dual of a factorization. Furthermore by exploiting a certain "operadic" property of the Garside element (see Theorem 3.6) we are able to calculate the dual of a factorization "locally" by relating the dual of a concatenation of factorizations to the concatenation of their duals. The mind-body duality is also closely related to the duality in the braid group defined by changing under-crossings to over-crossings and vice-versa (see Theorem 3.10).

If one interprets the mind-body duality in the context of topological graph theory, it turns out that it is a generalization of the duality of graphs embedded in closed oriented surfaces. In order to define that generalization we define the notion of a *proper embedding of a graph in a surface with boundary* (see Definition 4.1). A *properly embedded graph* (peg) is a graph embedded in a bounded surface in such a way that all the vertices lie in the boundary, with some technical conditions that ensure that the dual graph is also properly embedded. Namely we require every boundary component to contain at least one vertex, and each region to contain exactly one arc in its boundary. The prototype for this concept, that in this generality appears to be new, is a noncrossing tree. The class of *cellularly embedded graphs* (cegs) is contained in the class of pegs: a graph cellularly embedded in a closed oriented surface can be construed as a graph properly embedded in a surface with boundary in such a way that every boundary component of the surface contains exactly one vertex. In this context, leos are the analogue of *rotation systems*, and PTDCs the analogue of cycle double covers. A factorization (or a vertex-and-edge-labeled graph) gives an oriented surface with boundary with the graph pegged in it, via a straightforward generalization of the correspondence between rotation systems and cegs. The dual of a peg (defined in the obvious way, see Definition 4.7) is then isomorphic to the peg obtained by the mind-body dual of the factorization, and in the special case where the peg can also be considered a ceg, it coincides with the standard notion of duality for cegs (see Theorem 4.13).

We also give an alternative construction of the peg associated with a factorization, and of the mind-body duality, via the theory of branched coverings over the two dimensional disk. In this context, the peg corresponding to the factorization is essentially the lifting of a certain graph in the disk, a so-called *Hurwitz system*, and its dual is the lifting of a "dual" Hurwitz system, see Theorem 4.28. This interpretation of the

peg associated to a factorization is a generalization of a similar construction for cegs given by Arnold in [8], however our interpretation of duality via branched coverings appears to be new.

We hint at possible applications of this point of view to the theory of cegs by reproving some baby cases on the existence of self-dual complete graphs, and giving examples of self-dual embedding of the complete *digraph* on six vertices, where it is known that no self-dual embeddings of the corresponding complete (undirected) graph exist. We further remark that the theory of pegs is a refinement of the theory of cegs, that is more attuned to the graph theoretic properties of the graph: whether a graph can be pegged on to a given surface is not invariant under subdivisions; in fact (see Proposition 4.19) any ceg admits a subdivision that renders it the completion of a peg. We hope that this refinement will have future applications in explaining known, as well as discovering new, enumerative coincidences between various classes of Hurwitz numbers.

Even though in this work pegs are mainly used as a tool for a topological understanding of the mind-body duality and are not studied much as a topic on their own right, we do touch upon some natural questions that arise. One such natural question is the analogue of the genus question from the theory of cegs, namely: given a graph Γ, what can one say about the Euler characteristic and the number of boundary components of the pegs that arise from all possible edge-labelings of Γ? We consider that question in Section 4C and give a complete answer in the case of complete graphs, see Theorem 4.22. The general case is an interesting open question.

The main tool we develop is the *medial digraph* of an edge-labeled graph, or more generally a peg. This is a digraph that has vertices in bijection with the edges of the peg and there is an arc from a vertex a to a vertex b if and only if, the edge corresponding to b immediately follows the edge corresponding to a in the local edge order of some vertex of the peg, see Definition 2.10 and the paragraph following Definition 4.2. This notion is the analogue of the *medial graph* from the theory of cegs, which in the case of embeddings in oriented surfaces, also admits a natural digraph structure coming from the orientation of the surface. A peg can be encoded as a *perfect chain decomposition* (PCD) of its medial digraph (see Definition 2.16) and the mind-body duality corresponds to a natural duality on the set of PCDs (see Theorem 2.20).

We use the medial digraph to characterize the class of pegs obtained by factorizations in Proposition 2.13: a peg comes from a factorization if and only if its medial ditree is *directed acyclic graph* (dag). A peg can be "completed" into a ceg by gluing disks along the boundary components of its surface, and we use the medial digraph to characterize the class of cegs that are obtained as closures of pegs: a ceg is the closure of a peg if and only if its medial digraph admits a *feedback arc system* of cardinality the size of the ceg (see Proposition 4.12).

In the penultimate section, we examine the case of minimal transitive factorizations of a cycle, or equivalently edge-and-vertex-labeled trees, in the light of the developed theory. We first give an exposition of the results in [27], and we show that the mind-body duality at the level of factorizations can, in this case, be expressed via the mind-body duality at the level of rooted edge-labeled trees. In general, the mind-body dual of a factorization of a permutation π is a factorization of π^{-1}, but using rooted edge-labeled trees we show that one can define a duality between factorizations of the same cycle, enjoying the same structural properties as mind-body duality. This is a topic that will be further explored in [6]. As an application of the medial digraph (which in this case is a directed tree) we show that the set of self-dual edge-labeled trees is equinumerous with the set of alternating permutations (see Corollary 5.13).

In this work we develop the theory of pegs in enough detail to be able to treat the case of factorizations, postponing the fully developed theory for a future paper [3]. This and other forthcoming future directions of this project are outlined in the final section.

1A. Conventions and terminology. All graphs we consider are finite. We view graphs as one-dimensional complexes, with a set of 0-cells called *vertices* and a set of 1-cells called *edges*. Digraphs are graphs where every edge has been endowed with an orientation. We emphasize that every edge, *including loops*, admits two distinct orientations. In general we work with *loopless* graphs, i.e., all 1-cells are attached to two distinct 0-cells, except in Section 4, where digraphs are allowed to have loops. The edges of a digraph are sometimes called *darts* or *arcs* depending on the context. We use more or less standard graph theoretical terminology, with a few exceptions; in particular:

- We use Γ to denote a graph, sometimes endowed with extra structure. The set of vertices of a graph is denoted by V, and its set of edges by E. The *order* of Γ (i.e., $|V|$) is typically denoted by n, and its *size* (i.e., $|E|$) by m.
- We use χ to denote the Euler characteristic so that for a graph Γ, we have $\chi(\Gamma) = n - m$.
- The *neighborhood of a vertex v of* Γ, denoted $\nu(v)$ is the subgraph of Γ defined by all the edges that are incident to v. The *star of a vertex v of* Γ is the neighborhood of v in the first barycentric subdivision of Γ, the edges of the first barycentric subdivision are the *half-edges* of Γ.
- A *trail* is a walk without repeated edges. Every trail has a beginning and an ending vertex that may coincide.
- By a *cycle* we mean an equivalence class of closed trails, where two trails are equivalent if they have the same edges, i.e., the endpoint is not important. In the context of digraphs a cycle is always a *directed cycle*.

For a finite set V we denote the *symmetric group of V* by $\mathcal{S}_V$, if $V = [n] := \{1, \dots, n\}$ the symmetric group of V is denoted by $\mathcal{S}_n$. We use the usual terminology, e.g., *permutation*, *transposition*, *cycles*, etc even if V is not $[n]$. For a transposition $\tau = (s\,t)$ we call s, t the *moved points* of τ. We multiply permutations from left to right, so that $(1\,2)(2\,3) = (3\,2\,1)$. We use left and right exponential notation for conjugation in a group, i.e., $g^h := h^{-1}gh$ and ${}^h g := hgh^{-1}$.

2. Mind-body duality

In this section we define the *mind-body duality* first in a graph theoretical context and then in the more algebraic context of factorizations. We start by defining the main objects of study and their basic equivalence.

Definition 2.1. A *factorization in $\mathcal{S}_V$* is a sequence of transpositions $\rho = (\tau_i)_{1 \le i \le m}$, with $\tau_i \in \mathcal{S}_V$. The product of ρ is called its *total monodromy* or simply its *monodromy* and denoted by $\mu(\rho)$. We also say that ρ is a factorization *of* $\mu(\rho)$, and sometimes we'll call τ_i the i-th monodromy of ρ.[1]

The *reverse* of a factorization $\rho = (\tau_i)_{1 \le i \le m}$ is the factorization $\rho^\intercal := (\tau_{m+1-1})_{1 \le i \le m}$.

The *concatenation* of two factorization $\rho_1 = (\tau_i)_{1 \le i \le m_1}$ and $\rho_2 = (\tau'_i)_{1 \le i \le m_2}$ is the factorization $(\tau_1, \dots, \tau_{m_1}, \tau'_1, \dots, \tau'_{m_2})$. The concatenation of two factorizations ρ_1, ρ_2 will be denoted by $\rho_1 \rho_2$.

[1] This terminology comes from the theory of *branched coverings*, see Section 4D.

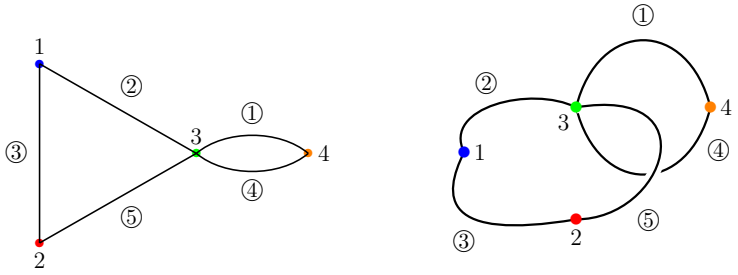

Figure 1. The graph associated with the factorization ρ of Example 2.2.

For a factorization $\rho = (\tau_i)$ and an element $\tau \in \mathcal{S}_V$, we use the notation ρ^τ (resp. ${}^\tau\rho$) for the factorization (τ_i^τ) (resp. $({}^\tau\tau_i)$).

An *edge-labeled graph* (e-graph for short) is a graph with edges labeled with elements of $[m]$ where m is the size of the graph, or equivalently a graph with a total order in the set of its edges. A *vertex-labeled graph* (or v-graph for short) is a graph whose vertices are labeled by $[n]$ where n is its order. An *edge-and-vertex-labeled graph* (or e-v-graph for short) has both vertices and edges labeled.

The *reverse* of an e-graph Γ, is the graph Γ^{T}, with the same vertices and edges as Γ and its edges relabeled according to $i \mapsto m + 1 - i$.

The *concatenation* of two e-graphs Γ_1 and Γ_2 of sizes m_1 and m_2, respectively, is the graph $\Gamma_1\Gamma_2$ with $V(\Gamma_1\Gamma_2) = V(\Gamma_1) \cup V(\Gamma_2)$ and $E(\Gamma_1\Gamma_2)$ consisting of the edges of Γ_1 with their labels unchanged, and the edges of Γ_2 with their labels increased by m_1.

There is an obvious bijection between the set of factorizations of $\mathcal{S}_V$ and the set of e-graphs with vertex set V: for a factorization ρ define the *associated e-graph of* ρ to be the graph $\Gamma(\rho)$ that has an edge with endpoints u, v, and labeled by i, if and only if the i-th monodromy of ρ is (u, v); conversely for an e-graph Γ with vertex set V, the *associated factorization of* Γ, $\rho(\Gamma)$ has the i-th monodromy interchanging the endpoints of the edge labeled i. This bijection specializes to a bijection between factorization of $\mathcal{S}_n$ of length m and e-v-graphs or order n and size m.

Clearly the bijection $\rho \mapsto \Gamma(\rho)$ preserves the notions of reverse and concatenation.

Example 2.2. The sequence $\rho = (3\,4), (1\,3), (1\,2), (3\,4), (2\,3)$ is a factorization of the cycle $(4\,3\,2\,1)$ in $\mathcal{S}_4$. The associated graph $\Gamma(\rho)$ is shown in Figure 1. We draw two versions of it, a standard planar drawing in the left, and one that the cyclic order of the edges at every vertex is consistent with the order induced by their labels, see Definition 2.3, on the right. The colors of the vertices are used later, see Section 2B.

2A. Mind-body duality for e-graphs. Let Γ be an e-graph. The edge labels induce two dual structures on Γ, a *local edge ordering* and a *perfect trail double cover*.

Definition 2.3. A *local edge ordering* (leo for short) of a graph Γ is an assignment of a linear order at the neighborhood of each vertex of Γ. We draw leos in such a way that the cyclic order induced by the standard (counterclockwise) orientation of the plane is consistent with the ordering of the edges at every vertex.

The edge labels of an e-graph induce a leo in the obvious way: the order of the edges around a vertex is given by the natural order of their labels.

A leo on a graph Γ induces a decomposition of the darts of Γ into chains. We define the relevant structure in more generality than is strictly needed for this paper in view of future planned work, see [3].

Definition 2.4. A *perfect trail double cover* (PTDC) of a graph Γ is a family $\mathcal{T}$ of positive length trails such that each edge of Γ belongs to exactly two trails of $\mathcal{T}$, and each vertex is the endpoint of two trails. We emphasize that trivial paths of length zero are not allowed, but *closed trails are allowed*; for a closed trail its endpoint is counted twice.

A PTDC is called *orientable* if each trail can be endowed with an orientation such that every edge is traversed once in each direction, in other words, $\mathcal{T}$ induces a decomposition of the darts of Γ into chains; a PTDC endowed with such a choice of orientations is called *oriented*. In that case every vertex is the start of exactly one trail (we denote that trail by $\overrightarrow{v}$) and the end of exactly one trail (we denote that trail by $\overleftarrow{v}$).

A PTDC is called *nonsingular* if for every interior (i.e., nonleaf) vertex the first edge of $\overrightarrow{v}$ and the last edge of $\overleftarrow{v}$ are distinct.

For the most part of this paper we will be dealing with oriented nonsingular PTDCs, and from now on, barring explicit mention to the contrary, we will use PTDC to mean an *oriented nonsingular* PTDC.

Given a leo on Γ define the *minimally increasing greedy trail* (migt) starting at v to be the trail $\overrightarrow{v}$ obtained as follows: the first edge of $\overrightarrow{v}$ is the smallest (with respect to the leo) edge in the neighborhood $\nu(v)$. We proceed inductively: once we have added an edge e from $\nu(u)$ to the trail, in the next step we add the smaller edge in the neighborhood of w, the other vertex of e, that is larger than e in the leo of w, provided that such edge exists. We stop when e is the last edge at the leo of w.[2]

As expected from the notation we have:

Lemma 2.5. *The migts of a leo give a* (*nonsingular*, *oriented*) *PTDC. Conversely*, *a* (*nonsingular*, *oriented*) *PTDC* $\mathcal{T}$ *gives a leo on its underlying graph*, *whose migts are the trails of* $\mathcal{T}$.

Proof. A vertex v of a graph Γ endowed with a leo is the endpoint of exactly two migts: $\overrightarrow{v}$ and $\overleftarrow{v}$. Now if the first edge e of $\overrightarrow{v}$ is the same as the last edge of $\overleftarrow{v}$, then e is both the first and the last edge in the leo of v and therefore is the only edge incident to v. So for an interior vertex v, the migts $\overleftarrow{v}$ and $\overrightarrow{v}$ are distinct.

Let e be an edge of Γ with endpoints v, u, and let e' (resp. e'') be the edge of Γ incident to v (resp. u) and immediately preceding e in the leo at v (resp. u), if such an edge exists. Then by the definition of migts, e belongs to two migts, $m_1 = \ldots, e', e, \ldots$ that transverses e from v to u, and $m_2 = \ldots, e'', e, \ldots$ that transverses e from u to v; of course if e' (resp. e'') does not exist then m_1 (resp. m_2) is simply $\overleftarrow{v}$ (resp. $\overleftarrow{u}$). If a migt does not pass through v or u it clearly can't contain e, and the other migts that pass through v or u do not contains e because migts are minimally increasing. So every edge belongs to exactly two migts that transverse it in opposite orientations.

The above can be summarized by saying that the local configuration of the migts in a neighborhood of a vertex v is as in Figure 4: we have a vertex of degree 4, there are four edges incident at a vertex with their order in the leo is indicated by the subscripts, and there are five migts that pass trough v.

[2]Recall that our graphs are loopless. Loops could be treated by considering half-edges but such generality is not needed in this paper.

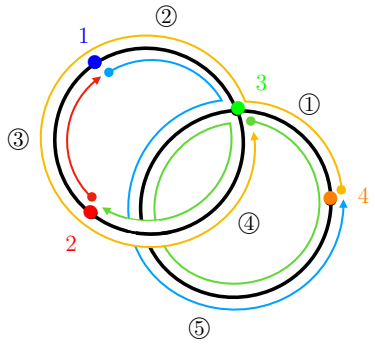

Figure 2. The graph from Figure 1 with its migts.

Conversely, the trails of a nonsingular oriented PTDC $\mathcal{T}$ that go through a vertex u are as in Figure 4. A leo at v can then be defined as the transitive closure of the relation defined by the rule that an edge e is less than an edge e' if e precedes e' in some trail. Clearly the migts of that leo are exactly the trails of $\mathcal{T}$. □

For example, the PTDC induced by the e-v-labeled graph in Figure 1 is shown in Figure 2.

If we start with a factorization ρ its monodromy can be recovered from the migts of the associated e-v-labeled graph. Recall that all PTDCs of graphs are assumed nonsingular and orientable.

Definition 2.6. The *monodromy digraph* of a PTDC is the digraph that has the same vertices as Γ and for each trail an arc from its beginning to its end, in other words there is an arc from v to u if and only if $\overleftarrow{v} = \overrightarrow{u}$.

Proposition 2.7. *The monodromy digraph of a PTDC is a functional digraph of a permutation in $\mathcal{S}_V$. If the PTDC comes from (the associated e-graph of) a factorization then that permutation is the monodromy of the factorization.*

Proof. The bidegree of every vertex of $\mu(\Gamma)$ is $(1, 1)$ because each vertex has exactly one incoming and one outgoing trail, so $\mu(\Gamma)$ is the functional digraph of a permutation.

To prove the second statement we need to prove that if $\overrightarrow{u} = \overleftarrow{v}$ then $\mu(u) = v$. In fact it's easy to see that the vertices of $\overrightarrow{u}$ form the trajectory of u under the successive applications of the elements of ρ. For, let e_1 be the first edge of $\overleftarrow{u}$ and v_1 its other endpoint, then e_1 is the first edge of $\nu(u)$ and therefore $(u\,v_1)$ is the first transposition in ρ that moves u. The next time u is moved, is when there is a monodromy $(v_1\,v_2)$ in ρ with index larger than the index of e_1, and this monodromy will correspond to the second edge e_2 of $\overrightarrow{u}$, and so on until will reach the last edge of $\overrightarrow{u} = \overleftarrow{v}$, which is the last edge of $\nu(v)$, and no further monodromies move v. It follows then that in the product μ of ρ we have $\mu(u) = v$. □

The above can be verified in Figure 2, $\mu(\rho) = (4\,3\,2\,1)$, and indeed $\overrightarrow{4} = \overleftarrow{3}$, $\overrightarrow{3} = \overleftarrow{2}$, $\overrightarrow{2} = \overleftarrow{1}$, and $\overrightarrow{1} = \overleftarrow{4}$.

The terminology PTDC and the monodromy digraph were inspired by [13].

Now we can define the *mind-body dual* of an e-labeled graph.

Definition 2.8. Let Γ be an e-graph. The *mind-body dual* e-graph is the graph Γ^* that has vertices the migts of Γ and an edge labeled i connecting two trails t_1 and t_2 if and only if the two trails share the edge of Γ labeled i.

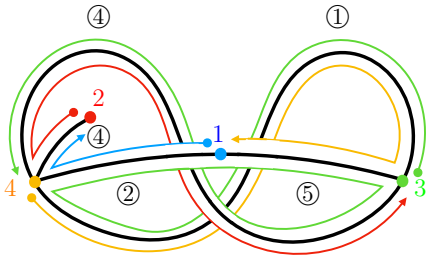

Figure 3. The mind-body dual of the graph in Figure 1.

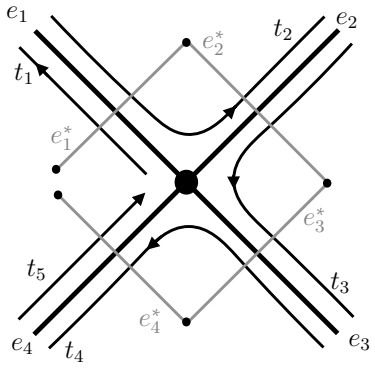

Figure 4. The local structure of migts.

There is a one-to-one correspondence between the edges of Γ and Γ^*, and corresponding edges have the same label, when needed we will denote the edge of Γ^* corresponding to the edge e of Γ by e^*. Since there are as many migts as vertices, Γ^* has as many vertices as Γ and there are two canonical ways to set up a correspondence between the vertices of Γ and the vertices of Γ^*: we can choose v^* to be $\overrightarrow{v}$ or $\overleftarrow{v}$. We choose the former, i.e., $v^* = \overrightarrow{v}$ but not much of what follows depends on that choice. We will comment when the choice makes a difference, see Remark 3.9 and Theorem 3.14.

Notice that if Γ is e-v-labeled, Γ^* is also e-v-labeled via the correspondence $v \mapsto v^*$. When no confusion is likely we will abuse language by talking as if Γ and Γ^* have the same vertices and edges.

For example, the mind-body dual of the e-v-graph of Figure 1 together with its migts is shown in Figure 3.

Theorem 2.9. (1) *If v is a vertex of Γ then the neighborhood of v^* in Γ^* consists of (the duals of) the edges of $\overrightarrow{v}$. The migt $\overrightarrow{v^*}$ consists of (the duals of) the edges of $\nu(v)$.*

(2) $(\Gamma^*)^* = \Gamma$.

(3) $\mu(\Gamma^*) = \mu(\Gamma)^{-1}$.

Proof. The arguments will be easier to follow if the reader refers to Figure 4.

The first statement of item (1) is obvious. To see the second let $e_1, \dots, e_d$ be the edges in $\nu(v)$ in the local ordering, then e_i and e_{i+1} are in some trail t_i and thus e_i^* connects t_i and t_{i+1}, and e_{i+1}^* is the smallest edge in t_{i+1} greater than e_i^*. So in the migt of v^*, e_{i+1}^* follows e_i^*.

Item (2) follows from item (1).

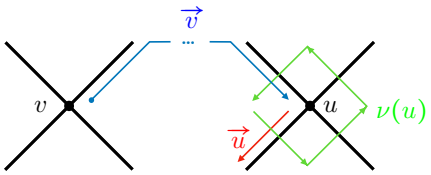

Figure 5. $\mu(\Gamma^*) = \mu(\Gamma)^{-1}$.

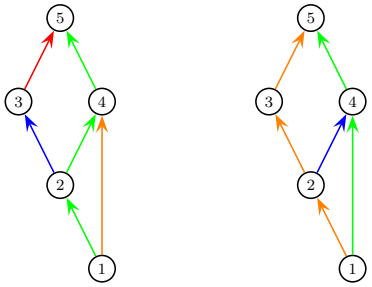

Figure 6. The medial digraph of the e-v-graph in Figure 1.

The proof of item (3) is illustrated in Figure 5: let $\mu = \mu(\Gamma)$ and $\mu^* = \mu(\Gamma^*)$, and assume that $m(v) = u$, we need to show that the migt of $\overrightarrow{u}$ in Γ^* ends in $\overrightarrow{v}$. The last edge of $\overrightarrow{v}$ is the last edge of $\nu(u)$. Now in Γ^* the migt of u^* consists of the edges dual to the edges in $\nu(u)$ so the last edge of this migt ends in $\overrightarrow{v}$. □

We end this subsection by noticing that one could more generally define the dual of a leo, and mutatis mutandis, almost all of the above would go through. This will indeed be done in [3].

2B. Medial digraphs. If we put together all the Hasse diagrams of the edge orders of a leo on Γ we obtain the *medial digraph* of the leo. For general definitions and terminology on digraphs we refer the reader to [9].

Definition 2.10. The *medial digraph* of a leo on Γ is the digraph $\mathcal{M}(\Gamma)$ with vertices the edges of Γ and an arc from edge a to edge b if and only if a immediately proceeds b in the local order around a vertex. A leo is called *e-realizable* if its medial digraph is a *dag*, in other words it has no (oriented) cycles.

For example, the medial digraph $\mathcal{M}$ of the e-v-graph in Figure 1 is shown in the left side of Figure 6, the edges of $\mathcal{M}$ are colored according to the vertex of Γ that they come from. The right side of Figure 6 shows the medial digraph of Γ^*, notice that, edge colors aside, the two digraphs coincide. In general, it follows from item (1) of Theorem 2.9 that:

Theorem 2.11. *For any factorization ρ we have $\mathcal{M}(\Gamma(\rho)^*) = \mathcal{M}(\Gamma(\rho))$.*

Notice also that the medial digraph in Figure 6 is a dag, and furthermore the edge labels endow it with a *topological sort*. We recall the definition:

Definition 2.12. A *topological sort* of a digraph is a total order of its vertices such that for two vertices u, v, we have that if there is an edge from u to v then $u < v$ in that order. Clearly a digraph admits a topological sort if and only if it is a dag.

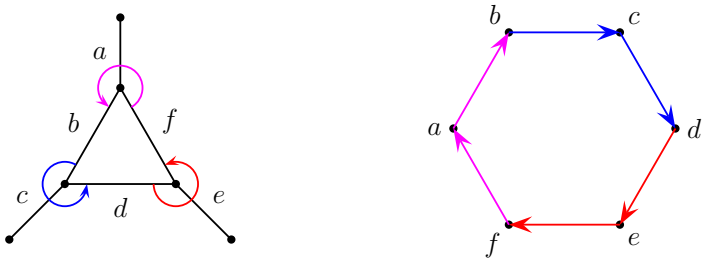

Figure 7. A non-e-realizable leo and its medial digraph.

A *topsorted dag* is a vertex-labeled dag such that the order of the vertices induced by their labeling is topological sort.

We note the somewhat subtle distinction of the two notions defined above. Clearly a topological sort of a dag gives a topsorted dag, but two different topological sorts of the same dag may give the same topsorted dag. For example, the updown ditree with three vertices (see Definition 5.11) admits two topological sorts, but because of the order two automorphism, there is only one topsorted dag whose underlying (unlabeled) dag is the updown ditree on three vertices.

Now we can prove:

Proposition 2.13. *A given leo is induced by an edge labeling of* Γ *if and only if it is e-realizable. Furthermore the edge labels induce a topological sort of the medial digraph.*

Proof. Clearly the medial digraph of an e-v-labeled graph contains no cycles since the local orders come from a global order. Conversely any dag admits a topological sort, that is a global order compatible with all the local orders thus giving a total order in the edges of Γ. □

Example 2.14. The left side of Figure 7 shows a non-e-realizable leo on a graph. Its (obviously nonacyclic) medial digraph is shown on the right.

These constructions were inspired by the ideas in Section 3 of [20]. The terminology *medial digraph* is meant to suggest an analogy with the medial graphs in the theory of graph embeddings, see, e.g., [7]. This analogy will be made precise in Section 4.

Before proceeding we prove a lemma:

Lemma 2.15. *The underlying graph of the medial digraph of a leo* Γ *has the same Euler characteristic as* Γ.

Proof. If Γ has m edges then $\mathcal{M}$ has m vertices. Furthermore a vertex v of degree d_v contributes $d_v - 1$ edges. So $\mathcal{M}$ has $\Sigma(d_v - 1) = 2m - n$ edges. So $\chi(\mathcal{M}) = m - (2m - n) = n - m$. □

Given a leo on Γ the local orders of every vertex induce a decomposition of the edges of $\mathcal{M}(\Gamma)$ into chains. For example, the graph in Figure 1 induce a chain decomposition of its medial digraph that is indicated by the coloring of the edges. We formalize this idea in the following definition.

Definition 2.16. A digraph M is a *binary digraph* if the in and out degree of every vertex is at most 2. A vertex of a binary digraph is *internal* if its in and out degree is at least 1.

A *perfect chain decomposition* (PCD for short) of a binary digraph M is a decomposition $\mathscr{C}$ of the edges of M into chains such that every vertex of M belongs to exactly two chains of $\mathscr{C}$, chains of length 0 are allowed.

One can now prove:

Lemma 2.17. (1) *The number of chains in any PCD of a binary digraph M is $2m-l$, where m is the number of vertices and l the number of edges.*

(2) *Given a binary digraph with ι internal vertices there are 2^ι PCDs on M. In particular, every medial digraph admits a PCD.*

(3) *Given any binary* ***dag*** *M, there is an e-graph whose medial digraph is M. In fact if ι stands for the number of internal vertices, τ for the number of topological sorts, and α for the number of automorphisms of M, then, up to isomorphism, there are*

$$\frac{2^\iota \tau}{\alpha}$$

e-graphs that have M as medial digraph.

Proof. Item (1) follows from the handshaking lemma: If there are k chains $c_1, \ldots, c_k$ with lengths $l_1, \ldots, l_k$, then c_i will have l_i+1 vertices, so $\sum_{i=1}^{k}(l_1+1)=2m$. On the other hand, $\sum_{i=1}^{k}(l_1+1)=l+k$.

For item (2) we note that there are 4 possible bidegrees for an internal vertex v: (1, 1), (1, 2), (2, 1), and (2, 2), and for each of these bidegrees there are two choices for joining the edges into 2 chains. If v has bidegree (1, 1), one choice is to join the two edges together for one of the chains and have the other chain to be the trivial chain v, while the other choice is to have one of the edges in the first chain and the other edge in the second. If v has bidegree (1, 2) with incoming edge (x, v) and outgoing edges (v, y) and (v, z) the fist choice is to join (x, v) and (v, y) and have (v, z) by itself, and the other choice is to join (x, v) and (v, z) and have (v, y) by itself. The case of bidegree (2, 1) is entirely similar. Finally for bidegree (2, 2) with incoming edges (x, v) and (y, v), and outgoing edges (v, z) and (v, w) the one choice is to join (x, v) with (v, z) and (y, v) with (v, w), and the other is to join (x, v) with (v, w) and (y, v) with (v, z). (See the first and third columns of Figure 8.)

For a noninternal vertex v there is only one choice: if the vertex is a leaf then one of the chains is the trivial chain v and the other contains the unique edge, while if v is a minimum (resp. maximum) each of the outgoing (resp. incoming) edges goes in to a different chain.

Making a choice in each of the vertices gives a PCD, and there are 2^ι such choices.

For item (3) we note that any PCD $\mathscr{C}$ of a binary digraph M defines a graph Γ with a leo, as follows: the vertices of Γ are the chains of $\mathscr{C}$, and for each vertex a of M there is an edge in Γ joining the vertices of Γ that correspond to the two chains that a belongs to. Clearly the neighborhood of a vertex c of G consists of the edges that correspond to the vertices of that chain in M, and the chain defines a total order on that neighborhood. By definition, the medial digraph of G is M and the PCD induced by the leo is $\mathscr{C}$.

If M is a dag, any topological sorting of M, gives an e-labeling for each of the 2^ι graphs Γ constructed above. Taking into account the action of the automorphism group of M gives the formula for the number of e-graphs that have M as medial digraph. □

The proof of item (3) identifies the set of (isomorphism classes of) e-graphs with medial dag M, with the set of (isomorphism classes of) PCDs of M, and (2) identifies PCDs of M with a set of binary choices,

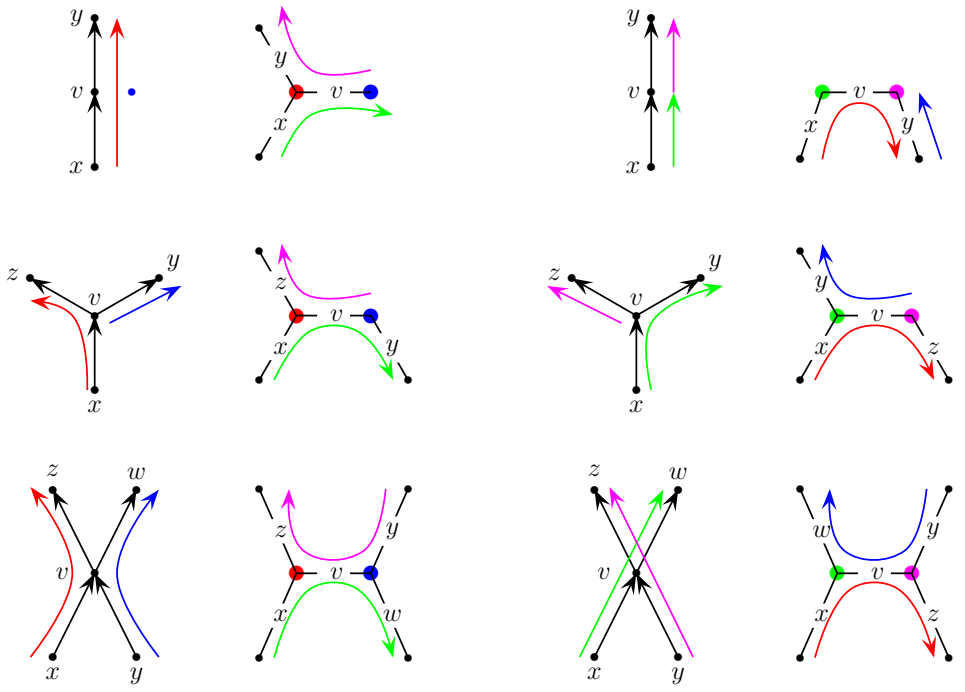

Figure 8. From PCDs to e-graphs and duality: choice 0 (left) and choice 1 (right).

one choice for each internal vertex. It turns out that under these identifications the mind-body dual of an e-graph is identified with the PCD obtained by making the opposite choice at every internal vertex. We make this precise below.

For any given binary digraph one could identify the two choices of connecting arcs at each internal vertex with 0 and 1. This identification can be done canonically for vertices of bidegree $(1, 1)$, say the choice that connects the two arcs is 0, and the choice that doesn't is 1. For the other types of internal vertices an identification has to be made arbitrarily at each vertex. One way to accomplish a uniform encoding is to draw M in the plane and then use the orientation of the plane to say, for example, that for vertices of bidegree $(1, 2)$ (resp. $(2, 1)$) choice 0 is to connect the single incoming (resp. outgoing) edge with the left outgoing (resp. incoming) edge, while for vertices of bidegree $(2, 2)$ choice 0 means to connect the left incoming to the left outgoing edge. After such an identification has been made, the proof of item (2) constructs a bijection from the set of all PCDs on M to the set of all function $s : I \to \{0, 1\}$, where I is the set of internal vertices of M. (See the first and third column of Figure 8)

Definition 2.18. The *dual* of a function $s : I \to \{0, 1\}$ is the function $s^* : I \to \{0, 1\}$ defined by $s^*(v) = 1 - s(v)$.

The *dual* $\mathcal{C}^*$ of a PCD on a binary digraph M constructed using the function s, is the PCD $\mathcal{C}$ constructed using $s*$.

Lemma 2.19. *The PTDC of an e-graph* Γ *also induces a PCD on* $\mathcal{M}(\Gamma)$. *The PCD induced from the PTDC of* Γ *is the dual of the PCD induced by the leo of* Γ.

Proof. The first statement follows from Theorem 2.11. The proof of the second is in Figure 8. The first column shows for each type of internal vertex with the PCD imposed by choice 0, the second shows the e-graph constructed from that PCD, the third column the PCD induced by the migts. Notice that in each case the PCD in the third column is exactly the PCD that corresponds to choice 1. The fourth column shows the graph constructed from choice 1 with it's migts. One can readily verify that the PCD imposed by those migts is exactly the PCD imposed by choice 0. □

An immediate corollary is:

Theorem 2.20. *For an e-graph* Γ, *the leo of* Γ^* *induces on* $\mathcal{M}(\mu)$ *the PCD dual to the PCD induced by the leo of* Γ.

For example, see Figure 6 that shows the medial digraph of the graph of Figure 1 in the left, and of its dual (the graph in Figure 3) in the right. The different chains of the PCDs are indicated by the different colors of the edges, and chains of length 0 are not shown since their presence can be deduced.

2C. Mind-body dual of a factorization. Now we can transfer this notion of duality to factorizations.

Definition 2.21. Let $\rho = (\tau_1, \ldots, \tau_m)$ be a factorization in $\mathcal{S}_n$. Then its mind-body dual ρ^* is defined to be the factorization associated to the mind-dual e-v-graph associated with ρ.

The term mind-body duality comes from the following amusing interpretation of a transposition introduced on the episode *The prisoner of Benda* (sixth season, episode 10) of the animated sitcom *Futurama* and elaborated on, for example, in [23]. In this scenario there is a machine that interchanges the minds of any two bodies that enter its two booths and each such exchange can be encoded by a transposition. A sequence of transpositions can then be interpreted as a series of mind exchanges from the point of view of the bodies. The dual sequence is then the series of body exchanges that the corresponding minds experience. This follows from the fact that $\nu(i)$, the neighborhood of a vertex i with its leo, stands for the sequence of mind exchanges that the body i experiences, while the trail $\vec{i}$ is the trajectory of the mind i, and so it describes the sequence of body exchanges that the mind i experiences.

Pursuing this interpretation a bit, we have a set of minds M and a set of bodies B, of the same cardinality, and an initial assignment of minds to bodies $\alpha_0 : M \to B$, say each mind is assigned to the body it's born in. Choosing an identification of M with $[n]$, and pushing it forward via α_0 to an identification of B with $[n]$ we can consider α_0 to be the identity permutation in $\mathcal{S}_n$ and any other mind-body assignment as a permutation $\alpha \in \mathcal{S}_n$. That way $\mathcal{S}_n$ acts on the set of mind-body assignments on the left by permuting the minds and on the right by permuting the bodies. The dual of a permutation of minds π, with respect to a mind-body assignment α is the permutation of bodies π^{*_α} that has the same effect in α as π. In other words we want $\pi\,\alpha = \alpha\,\pi^{*_\alpha}$, and it follows that

$$\pi^{*_\alpha} = \pi^{\alpha}. \tag{1}$$

Now given a factorization $\rho = \tau_1, \ldots, \tau_n$, and considering it as a sequence of mind exchanges, it's mind-body dual is the factorization ρ^* which when considered as a sequence of body exchanges, has the same effect in the mind-body assignment as ρ, at *every* step. The mind-body assignments we obtain by applying ρ to α_0 are, $\alpha_1 = \tau_1\alpha_0 = \tau_1,\ \alpha_2 = \tau_2\alpha_1 = \tau_2\tau_1,\ \ldots,\ \alpha_n = \tau_n \cdots \tau_2\tau_1$.

Taking into account equation (1) we have the following explicit formula for the mind-body dual of a factorization:

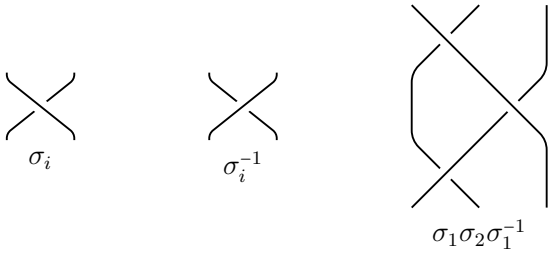

Figure 9. Generators and multiplication in the braid group.

Theorem 2.22. *For a factorization $\rho = \tau_1, \ldots, \tau_n$ we have*

$$\rho^* = \tau_1, \tau_2^{\tau_1}, \ldots, \tau_n^{\tau_{n-1}\cdots\tau_1}. \tag{2}$$

We note that this formula is an expanded version of Theorem 3.4 in the next section.

It is amusing to explain the properties of mind-body duality described in Theorem 2.9 in terms of the mind-exchange interpretation. For example, the monodromy of the dual is the inverse of the monodromy of the original, because from the point of view of the minds, there are body-mind assignments and the initial body-mind assignment is of course α_0^{-1}.

The author would like to stress that despite the use of this terminology, he does not subscribe to the philosophically untenable position of Cartesian dualism that is implicitly assumed.[3]

3. The Hurwitz action

The braid group on m strands B_m is the group generated by $m-1$ generators σ_i for $i \in [m-1]$ and relations: $\sigma_i\sigma_j = \sigma_j\sigma_i$ if $|i-j| > 1$, and $\sigma_i\sigma_{i+1}\sigma_i = \sigma_{i+1}\sigma_i\sigma_{i+1}$. For details about the braid groups we refer the reader to [12] and [31]. We view the braid group as the *mapping class group* of a 2-dimensional disc $\mathbb{D}_m^2$, with m distinguished points called the punctures; that is, B_m is the group of isotopy classes of orientation preserving self homeomorphisms of $\mathbb{D}_m^2$ that fix the boundary circle pointwise and permute the m punctures. For details about mapping class groups of surfaces and this interpretation of the braid groups we refer the reader to [12] and [24]. We will represent braids graphically and our convention is that the positive generator σ_i is represented diagrammatically by the i-th strand going *over* the $(i+1)$-th and that multiplication in the braid group happens from top to bottom, see Figure 9.

Since the fundamental group of a disc with m punctures is F_m, the free group on m-generators, the interpretation of elements of B_m as self homeomorphisms of $\mathbb{D}_m^2$, induces a left action of B_m by automorphisms on F_m. If $x_1, \ldots, x_m$ are the free generators of F_m then the action of the generator σ_i is given, on the generators of F_m by, $\sigma_i x_j = x_j$ for $j \neq i, i+1$, while $\sigma_i x_i = x_i x_{i+1} x_i^{-1}$ and $\sigma_i x_{i+1} = x_i$. It follows that B_m acts on the right on the set of homomorphisms $\mathsf{F}_m \to G$, for any group G and in particular for G a symmetric group. A factorization ρ is a sequence of elements in a symmetric group, and therefore can be construed as a representation of F_m to that group. So we have a right action of B_m on the set of all factorizations in any symmetric group, this action is called the *Hurwitz action*. The action

[3]The author, after long deliberations, decided to not use the term *husband-wife duality* alluding to a more risquè interpretation, and to leave such an interpretation to the reader if (s)he is so inclined.

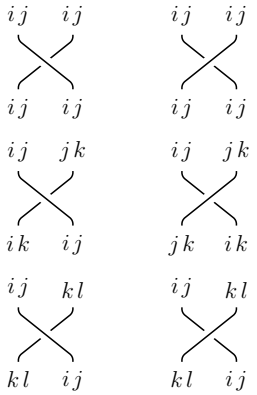

Figure 10. The Hurwitz action on factorizations.

of a generator σ_i on the factorization $\rho = \tau_1, \ldots, \tau_m$ is given by[4]

$$(\rho\,\sigma_i)_k = \begin{cases} \tau_k & \text{if } k \neq i, i+1, \\ {}^{\tau_i}\tau_{i+1} & \text{if } k = i, \\ \tau_i & \text{if } k = i+1. \end{cases} \tag{3}$$

The Hurwitz action can be described diagrammatically as in Figure 10, where i, j, k, l are distinct elements of $[n]$ and $i\ j$ stands for the transposition $(i\ j)$. For more details about the Hurwitz action, see [2].

Using the bijection between factorizations of $\mathcal{S}_n$ and e-graphs on $[n]$ (see Definition 2.1), we can transfer this to a B_m action on the set of e-labeled graphs on $[n]$ with m edges. It is easily seen that if Γ is such an e-v-graph then $\Gamma\sigma_i$ is obtained from Γ by interchanging the labels of the i-th and $(i+1)$-th edge and then "sliding" the $(i+1)$-th edge along the i-th, while $\Gamma\sigma_i^{-1}$ is obtained by interchanging the i-th and $(i+1)$-th labels and then sliding the i-th edge along the $(i+1)$-th. We interpret a slide of an edge along a nonadjacent edge to have no effect. This action on e-v-labeled graphs, which we'll also call the *Hurwitz action*, is shown in Figure 11, only the edges labeled i and $i+1$ are shown since the other edges are not affected.

Notice that this action descends at the level of e-labeled graphs (just forget the v-labels in Figure 11). We will still call it the Hurwitz action since no confusion is likely to arise. We remark that this B_m action on the set of e-graphs, was also noted in [15], for the case of e-trees.

3A. The duality in terms of the Hurwitz action. We first define the notions that we need in order to provide the promised characterization.

Definition 3.1. For $i < j \leq m$ define the braid $\delta_{i,j}$ to be the braid that takes the j-th point, and moves it to the i-place going under all in between strands, leaving all the other strands unchanged, i.e., $\delta_{i,j} := \sigma_{j-1}\sigma_{j-2}\ldots\sigma_i$. We will write simply δ_m for $\delta_{1,m}$.

The braid $\lambda_{i,j}$ is defined to be the braid that takes the i-th point, and moves it to the j-th place going over all the in between strands, i.e., $\lambda_{i,j} := \sigma_i \ldots \sigma_{j-1}$. We simply write λ_m for $\lambda_{1,m}$.

[4]Recall that ${}^h g$ stands for hgh^{-1}. Since transpositions are involutions, we could have used $\tau_{i+1}^{\tau_i}$ in the formula, but we choose to write the formula in a way that applies for any elements of any group.

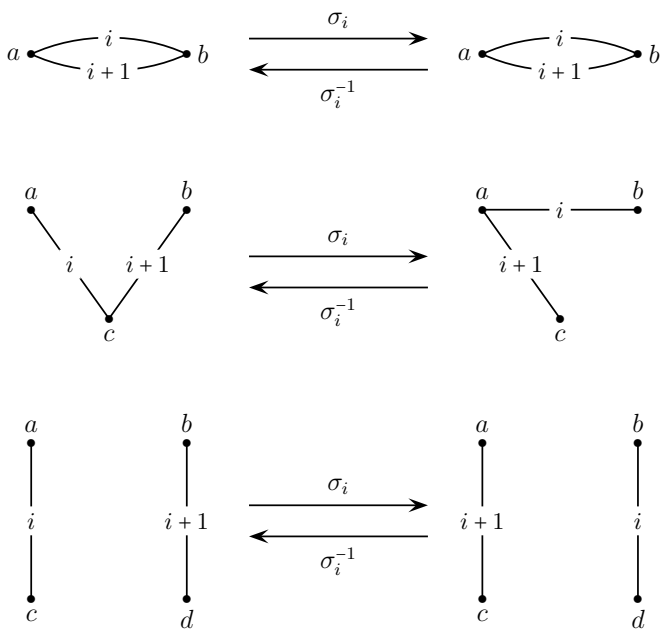

Figure 11. The Hurwitz action on e-v-graphs.

We also define $\Delta_{i,j} := \delta_{i,j}\,\delta_{i+1,j}\ldots\delta_{j-1,j}$. We simply write Δ_m for $\Delta_{1,m}$ and call it the *Garside element* of B_m.

We summarize some of the properties of the Garside element in the following proposition. All of these properties are either well known or follow easily from the definitions.

Proposition 3.2. (1) *For all $i \in [m-1]$ we have $\sigma_i^{\Delta_m} = \sigma_{m-i}$.*

(2) *Δ_m^2 is central in B_m. In fact it generates the center of B_m.*

(3) *As an element of the mapping class group of $\mathbb{D}_m^2$, Δ_m is represented by a homeomorphism that, leaving the boundary circle fixed, twists a smaller disk that contains all the punctures by π.*

(4) $\Delta_m = \delta_m\,\Delta_{2,m}$.

(5) $\Delta_m = \Delta_{2,m}\,\lambda_m$.

If e is an edge of a graph Γ and t a trail in Γ ending in a vertex incident to e then we refer to the operation of detaching e from the end vertex of t and attaching it to the beginning vertex as *sliding the edge e along the trail t*.

Lemma 3.3. *Let Γ be an e-v-graph with m edges. If the m-th edge of Γ has endpoints (v_1, v_2) then $\Gamma\delta_m$ is obtained from Γ by sliding edge m along the migts $\overleftarrow{v_1}$ and $\overleftarrow{v_2}$ and relabeling its edges according to $i \mapsto (i+1) \bmod m$.*

Proof. Let $i_1 < i_2 < \ldots < i_k < i_{k+1} = m$ be the edges of the union of the trails $\overleftarrow{v_1}$ and $\overleftarrow{v_2}$. Then by the definition of migts as minimally increasing, the edges with labels l with $i_k < l < m$ are not adjacent to edge m, the edges with labels l in the range $i_{k-1} < l < i_k$ are not adjacent to the edge i_k, and so on. So if we write

$$\delta_m = (\sigma_{m-1}\ldots\sigma_{i_k+1})(\sigma_{i_k}\ldots\sigma_{i_{k-1}+1})\ldots(\sigma_{i_1}\ldots\sigma_1),$$

then the action of the first factor has the effect of increasing the labels of the edges in the range $i_k < l < m$ by one and relabeling edge m as edge $i_k + 1$ without changing the underlying graph. The action of the second factor on the resulting e-v-labeled graph is then to slide the edge i_{k+1} along edge i_k, increase the labels in the range $i_{k-1} < l < i_{k+1}$ by one, and relabel i_{k+1} as $i_{k-1} + 1$. This pattern continues so that the overall effect of the action by δ_m is to increase all the labels in the range $1 \leq l < m$ by one, relabel the edge originally labeled m as 1 and slide it along all the edges in the trails leading to v_1 or v_2. □

We can now give the characterization of mind-body duality in terms of the Hurwitz action.

Theorem 3.4. *Let Γ be an e-v-graph of size m. Then*

$$\Gamma^* = (\Gamma \Delta_m)^{\intercal}. \tag{4}$$

Proof. We proceed by induction on the number of edges m. For $m = 1$ the theorem is obvious. Let the edge labeled m be incident to vertices v_1 and v_2, v_1' be the starting vertex of $\overleftarrow{v_1}$, and v_s' be the starting vertex of $\overleftarrow{v_2}$. Consider the e-v-labeled graph $\Gamma_1 = \Gamma \setminus m$ obtained by deleting edge m. After we attach the edge m to Γ_1 the trail that ended in v_1 gets augmented by m and ends in v_2 while the trail that ended in v_2 gets augmented by m and ends in v_1, and all the other trails are the same. It follows that Γ^* is obtained from Γ_1^* by attaching a new edge labeled m to the vertices v_1' and v_2'.

On the other hand since $\Delta_m = \delta_m \Delta_{2,n}$, by Lemma 3.3, $\Gamma \Delta_m$ is obtained by $\Gamma_1 \Delta_{m-1}$ by increasing all edge labels by one and attaching an edge labeled 1 to v_1' and v_2'. Now by induction we have that $\Gamma_1 \Delta_{m-1} = (\Gamma_1^*)^{\intercal}$ and so it follows that $\Gamma \Delta_m$ is obtained from $(\Gamma_1^*)^{\intercal}$ by increasing all edge labels by one and adding an edge labeled 1 attached to the vertices v_1' and v_2'. Taking the reverse we conclude that $(\Gamma \Delta_m)^{\intercal}$ is obtained by Γ_1^* by attaching a new edge labeled m to the vertices v_1' and v_2'. □

Using this we can get the following formula for the mind-body dual of a factorization:

Corollary 3.5. *For a factorization $\rho = \tau_1, \ldots, \tau_m$ we have*

$$\rho^* = \tau_1, {}^{\tau_1}\tau_2, \ldots, {}^{\tau_1 \ldots \tau_{m-1}}\tau_m. \tag{5}$$

Proof. One can prove by induction that $\rho\, \Delta_m = {}^{\tau_1 \ldots \tau_{m-1}}\tau_m, \ldots, {}^{\tau_1}\tau_2, \tau_1$ using Proposition 3.2(4). □

Notice that since all transpositions are involutions, this formula is the same as (2) in Theorem 2.22.

The Garside element plays a central role in the theory of braid groups and can be written in terms of the generators in many interesting ways, and each of these ways gives some information for the mind-body dual of a factorization. We give a very general description of many of these properties, using the language of operads.

All the braid groups can be "put together" into an algebraic structure called the *braid operad*. We refer the reader to [33] for (some) details. The composition in the braid operad is given by *cabling*, that is, the composition

$$B_m \times B_{i_1} \times \ldots \times B_{i_m} \to B_{i_1 + \cdots + i_m}$$

sends $(\beta, \beta_1, \ldots, \beta_m)$ to $\beta[\beta_1, \ldots, \beta_m]$ defined informally[5] as follows: think of the strands of β as "cables" where several strands are weaved together, the cable that corresponds to the k-th strand is weaved according to the braid $\beta_k \in B_{i_k}$. The braid $\beta[\beta_1, \ldots, \beta_m]$ is then the braid that results if we forget the

[5]We won't give the formal definition since it would take us far afield. We hope this informal description is enough for the reader to fill the details if (s)he wishes.

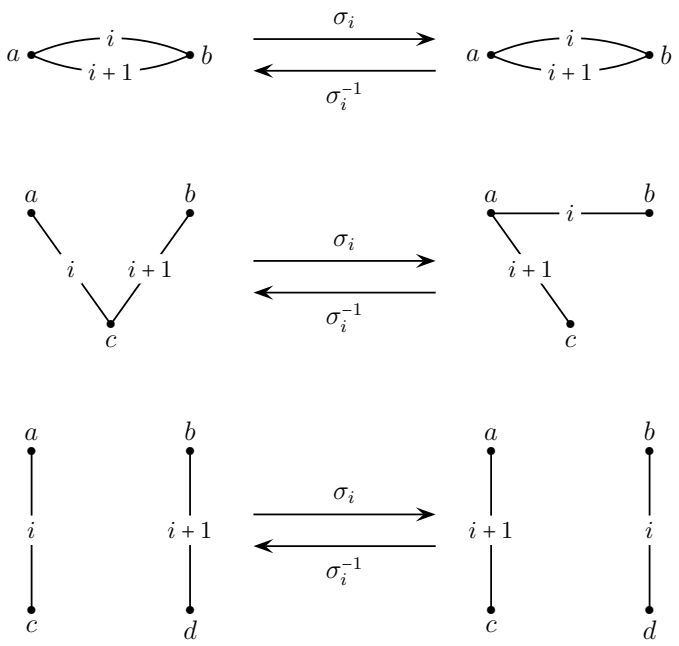

Figure 12. First part of proof of Theorem 3.6.

"cable structure" and view all the strands of all the cables as strands of new bigger braid in $B_{i_1+\cdots+i_m}$. See, for example, Figure 14, for cabling using Garside elements. With that notation in place we can now state the following property of the Garside element, which the author feels it should be well known but wasn't able to find a reference in the literature. For the statement of the following, we take $\Delta_1 = 1$, the unique one strand braid; we also remark that $\Delta_2 = \sigma_1$.

Theorem 3.6. *The family $(\Delta_k)_{k\geq 1}$ is a suboperad of the braid operad isomorphic to the associative operad.*[6] *Indeed, for all positive integers $i_1, \ldots, i_m$ we have*

$$\Delta_m[\Delta_{i_1}, \ldots, \Delta_{i_m}] = \Delta_{i_1+\cdots+i_m}.$$

Proof. We proceed by induction on m. For $m = 1$ the result is obvious. For $m = 2$, i.e., proving $\Delta_2[\Delta_{k_1}, \Delta_{k_2}] = \Delta_{k_1+k_2}$ we first observe that if $k_2 = 1$, this is simply a restatement of the definition of $\Delta_{k_2+1} := \delta_{k_2+1}\Delta_{2,k_2+1}$ (see Definition 3.1). This can be seen in the top of Figure 12. In the diagrammatic calculations we use the convention that a cable weaved according to Δ_k is denoted by a thick strand carrying a box labeled k. If we assume that the result has been proved for k_2, then the bottom of Figure 12, and the definition of $\Delta_{k_1+k_2+1}$ as $\delta_{k_1+k_2+1}\Delta_{2,k_1+k_2+1}$ proves it for $k_2 + 1$.

Assuming now that the result has been proved for m we use induction on k_m to prove it for $m + 1$. For $k_m = 1$, it is again a restatement of the definition of the Garside element. Assuming that it has been proved for k_{m+1} Figure 13, proves it for $k_{m+1} + 1$, using the case $m = 2$ that was proved above. This concludes the induction and the proof. □

[6]Thanks to Najib Idrissi for observing this in this comment in MathOverflow.

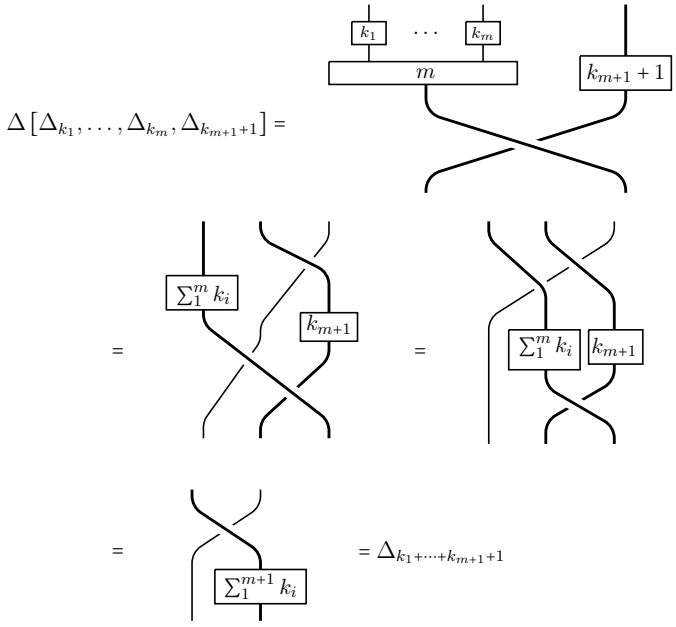

Figure 13. Second part of proof of Theorem 3.6.

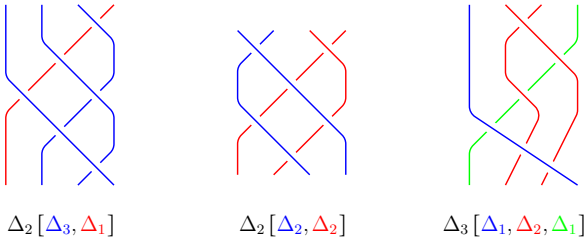

Figure 14. Δ_4.

For example, Figure 14 shows the Garside element Δ_4 as $\Delta_2[\Delta_1, \Delta_3]$, (its definition) in the left, as $\Delta_2[\Delta_2, \Delta_2]$ in the center, and as $\Delta_3[\Delta_1, \Delta_2, \Delta_1]$ in the right.

Using the operadic property of the Garside elements we can prove the following generalization of Corollary 3.5, that allow one to compute the mind-body dual of a concatenation of factorizations "piecewise".

Theorem 3.7. *For a factorization that is a concatenation of k factorizations* $\rho = \rho_1\,\rho_2\ldots\,\rho_k$, *we have*:

$$\rho^* = \rho_1^{*\,\mu(\rho_1)}(\rho_2^*)^{\,\mu(\rho_1\ldots\rho_{k_1})}(\rho_k^*).$$

Corollary 3.8. *If* $\rho = \tau_1, \tau_1, \tau_2, \tau_2, \ldots, \tau_k, \tau_k$ *is a factorization of the identity permutation into k pairs of identical transpositions, then* $\rho^* = \rho$.

Proof. This follows form the fact that $\Delta_{2k} = \Delta_k[\sigma_1, \ldots, \sigma_1]$. Each σ_1 fixes the corresponding pair and since the product of each pair is the identity permutation, the conjugations resulting from action of Δ_k have no effect. □

We conclude this subsection with a remark.

Remark 3.9. Recall that in Section 2 we made the convention $v^* = \overrightarrow{v}$ and we remarked that not much would change if we had made the convention $v^* = \overleftarrow{v}$ instead. In our context, if we had made that convention then formula (4) would read

$$\Gamma^* = (\Gamma \Delta_m^{-1})^{\mathsf{T}}$$

instead. Indeed, there is a straightforward analogue of Lemma 3.3 and the proof of Theorem 3.4 would go through almost verbatim. For a different proof see Theorem 3.14 in the next subsection.

3B. A closer look at the relation of the Hurwitz action and duality. So far we have seen two involutions on the set of our objects of study: mind-body duality $x \mapsto x^*$ and reversion $x \mapsto x^{\mathsf{T}}$. These involutions are related to analogous involutions on the braid group. In this subsection we explore that relation.

It's easy to check that the assignment $\sigma_i \mapsto \sigma_i^{-1}$ defines an (outer) automorphism $* : B_m \to B_m$.[7] If $\beta \in B_m$ we will denote its image under this automorphism by β^*. Diagrammatically a diagram for β^* is obtained from a diagram of β by reversing all the crossings, i.e., turning over-crossings to under-crossings and vice versa.

We have the following relation of this automorphism of B_n and mind-body duality:

Theorem 3.10. *Let $\beta \in B_m$ and Γ an e-graph with m edges. Then the Hurwitz action has the property*

$$(\Gamma\beta)^* = \Gamma^*\beta^*.$$

Proof. Let Γ be an e-v-labeled graph, it suffices to prove that for all $i \leq m-1$ we have $(\Gamma\sigma_i)^{-1} = \Gamma^*\sigma_i^{-1}$. There are three cases, the edges i and $i+1$ have two, one, or no vertices in common.

In the first case $\Gamma\sigma_i = \Gamma$ so we need to prove that $\Gamma^*\sigma_i^{-1} = \Gamma^*$, i.e., the edges i and $i+1$ have two vertices in common in Γ^* as well. This is the case, because both trails, say t_1 and t_2, that contain the edge i have to continue with $i+1$, thus in Γ^* both edges i and $i+1$ connect t_1 to t_2.

In the second case, in Γ there are exactly three migts that contain edge i or $i+1$, t_0 that contains both i and $i+1$, t_1 that contains only i, and t_2 that contains only $i+1$. Then in Γ^* edge i connects t_0 to t_1, and edge $i+1$ connects t_0 to t_2; in particular i and $i+1$ are adjacent at t_0. After the action of σ_i, all migts except these three remain the same, while t_0 changes by loosing the edge $i+1$, t_1 remains the same except that the edge that was labeled i is now labeled $i+1$, and t_2 changes by replacing the edge $i+1$ with two edges i and $i+1$. So in the dual of $\Gamma\sigma_i$ the vertex t_2 is connected to t_0 by an edge labeled i and to t_1 by an edge labeled $i+1$. See Figure 15.

Finally in the third case Γ and $\Gamma\sigma_i$ have the same underlying graph and their labeling differs only in that the edges i and $i+1$ have exchanged labels. We need to prove that the same is true for Γ^* and $\Gamma^*\sigma_i^{-1}$. This is the case because a migt that contains edge i cannot contain edge $i+1$, since it starts with edges less or equal to i and either ends in an endpoint of i or continues with the largest edge greater than i incident at an endpoint of i, which is greater than $i+1$. So the edges i and $i+1$ cannot belong to the same migt. □

The reverse of a factorization is also related to an involution of the braid group. Define the *reverse* of a braid $\beta \in B_m$ to be the dual of β conjugated by the Garside element, i.e., $\beta^{\mathsf{T}} := (\beta^*)^{\Delta_m}$. Then we have:

[7] Actually (see [21]) it is the only nontrivial outer automorphism of B_m.

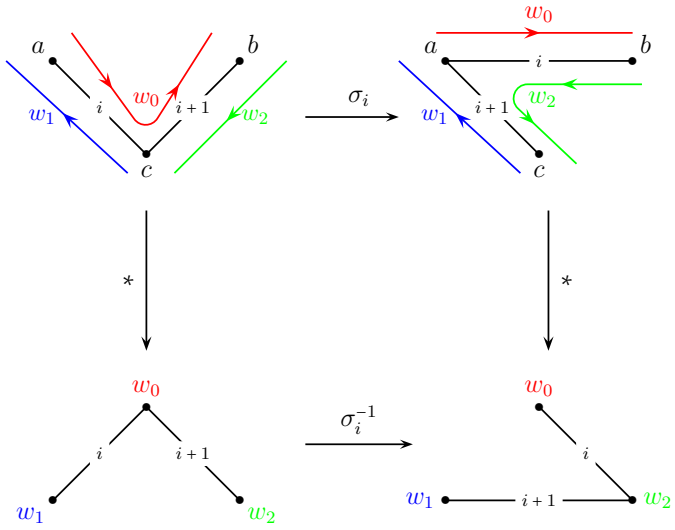

Figure 15. The second case in the proof of Theorem 3.10.

Theorem 3.11. *For any factorization with m monodromies we and any braid $\beta \in B_m$ we have*

$$(\rho\beta)^{\mathsf{T}} = \rho^{\mathsf{T}}\beta^{\mathsf{T}}.$$

Proof. It suffices to prove this for the standard generators σ_k of the braid group B_m. We have that for $1 \le i \le n$ then $\rho^{\mathsf{T}}(i) = \rho(m+1-i)$. Now

$$(\rho\sigma_k)(i) = \begin{cases} \rho(i) & \text{if } i \neq k, k+1, \\ \rho(k+1)^{\rho(k)} & \text{if } i = k, \\ \rho(k) & \text{if } i = k+1, \end{cases}$$

so that

$$(\rho\sigma_k)^{\mathsf{T}}(i) = (\rho\sigma_k)(m+1-i) = \begin{cases} \rho(m+1-i) & \text{if } m+1-i \neq k, k+1, \\ \rho(k+1)^{\rho(k)} & \text{if } m+1-i = k, \\ \rho(k) & \text{if } m+1-i = k+1 \end{cases}$$

$$= \begin{cases} \rho(m+1-i) & \text{if } i \neq m+1-k, m-k, \\ \rho(k+1)^{\rho(k)} & \text{if } i = m+1-k, \\ \rho(k) & \text{if } i = m-k. \end{cases}$$

Note that for the standard generators $\sigma_i \in B_m$ we have $\sigma_k^{\mathsf{T}} = \sigma_{n-k}^{-1}$. So

$$(\rho\sigma_k)^{\mathsf{T}}(i) = \rho^{\mathsf{T}}\sigma_{n-k}^{-1}(i) = \begin{cases} \rho^{\mathsf{T}}(i) & \text{if } i \neq m-k, m-k+1, \\ \rho^{\mathsf{T}}(m-k+1) & \text{if } i = m-k, \\ \rho^{\mathsf{T}}(m-k)^{\rho^{\mathsf{T}}(m-k+1)} & \text{if } i = m-k+1 \end{cases}$$

$$= \begin{cases} \rho(m+1-i) & \text{if } i \neq m-k, m-k+1, \\ \rho(k) & \text{if } i = n-k, \\ \rho(k+1)^{\rho(k)} & \text{if } i = m-k+1. \end{cases}$$

□

We conclude this section by having a closer look at the action of the Garside element Δ_m.

Proposition 3.12. *Let ρ be a factorization, then*

$$\rho\Delta_m^2 = {}^{\mu(\rho)}\rho.$$

Proof. We will use the well known fact that $\Delta_m^2 = \lambda_m^m$, where $\lambda_m = \sigma_1 \dots \sigma_{m-1}$ (see Definition 3.1). It is easy to see that for an e-v-graph Γ, $\Gamma\lambda_m$ is obtained from Γ by sliding all the edges along the edge labeled 1 and then relabel all the edges according to $i \to i-1 \bmod m$, or equivalently, interchanging the two vertex labels of the edge labeled 1 and then relabeling the edges according to $i \mapsto i-1 \bmod m$. After m iterations all the edges have their original labels, and each edge has interchanged its labels, in the order of the original e-labeling. This means that the vertex labels have been relabeled according to $\mu(\Gamma)$. □

At the level of e-graphs the vertex labels are not important so it follows:

Corollary 3.13. *Action by Δ_m^2 fixes all e-graphs, and therefore the action of Δ_m is an involution on the set of all e-graphs.*

Recall that in Section 2A we made the convention that in the case of an e-v-graph Γ the vertex labeling of Γ^* is that the vertex that corresponds to the migt $\overrightarrow{v}$ gets the same label as v. With this definition of duality we have that $(\Gamma^*)^\intercal = \Gamma\Delta$. Let $\bar{*}$ be the dual defined by the convention that the vertex of the dual that corresponds to the migt $\overleftarrow{v}$ gets the same label as v, in other words, $v^{\bar{*}} = \overleftarrow{v}$. Then we can prove that:

Theorem 3.14. *For an e-v-labeled graph with m edges, or factorization with m monodromies, x, we have*

$$x^{\bar{*}} = (x\Delta_m^{-1})^\intercal.$$

In particular for a factorization $\rho = \tau_1 \dots \tau_m$ we have

$$\rho^{\bar{*}} = \tau_1^{\tau_2\tau_3\dots\tau_m}, \tau_2^{\tau_3\dots\tau_m}, \dots, \tau_{m-1}^{\tau_m}, \tau_m.$$

We will need the following lemma in the proof:

Lemma 3.15.
$$\Delta_m^\intercal = \Delta_m^* = \Delta_m^{-1}.$$

Proof. If $\Delta_m^* = \Delta_m^{-1}$ then $\Delta_m^\intercal = (\Delta_m^{-1})^{\Delta_m} = \Delta_m^{-1}$. To prove the former we proceed by induction. For $m = 1, 2$ it is clear. Assuming that it has been proved for m we have

$$\Delta_{m+1}^* = \delta_{m+1}^* \Delta_{2,m+1}^* = \lambda_{m+1}^{-1} \Delta_{2,m+1}^{-1} = (\Delta_{2,m+1}\lambda_{m+1})^{-1} = \Delta_{m+1}^{-1},$$

where we used items (4) and (5) of Proposition 3.2, and the easily checked fact that $\delta_m^* = \lambda_m^{-1}$. □

Proof of Theorem 3.14. Clearly the migts of $\Gamma^\intercal$, for an e-v-graph Γ are the inverses of the migts of Γ, and therefore $\Gamma^{\bar{*}}$ is the reverse of the mind-body dual of the reverse of Γ. So

$$\Gamma^{\bar{*}} = ((\Gamma^\intercal)^*)^\intercal = \big(((\Gamma^\intercal)\Delta_m)^\intercal\big)^\intercal = \Gamma^\intercal\Delta_m = (\Gamma\Delta^\intercal)^\intercal = (\Gamma\Delta_m^{-1})^\intercal.$$

The formula for the $\rho^{\bar{*}}$ is just the expanded form of this. □

3C. Loop braid group action. The *loop braid group* LB_m in m strands is an extension of the braid group and has been defined several times in the literature under a variety of different names and points of view, we refer the reader to [16] for a survey of the different manifestations of these groups. We will follow the spirit of [25]: LB_m is generated by $2(m-1)$ generators $\sigma_1, \dots, \sigma_{m-1}$ and $s_1, \dots, s_{m-1}$, the σ_i generate a subgroup isomorphic to B_m and the s_i a subgroup isomorphic to $\mathcal{S}_m$ (s_i stands for the transposition $(i\ i+1)$) and there are three types of additional relations involving generators of both types: $\sigma_i\, s_j = s_j\, \sigma_i$ for $|i-j|>1$, $s_i\, s_{i+1}\sigma_i = \sigma_{i+1}\, s_i\, s_{i+1}$, and $\sigma_i\, \sigma_{i+1}\, s_i = s_{i+1}\, \sigma_i\, \sigma_{i+1}$.

LB_m is isomorphic to a subgroup of the automorphism group of F_m, where the σ_i acts like a braid, while s_i interchanges the i-th and $(i+1)$-th generators. It follows that the Hurwitz action extends to an action of LB_m, where s_i just interchanges the i-th and $(i+1)$-th monodromy (or the labels of the corresponding edges for an e-graph).

Definition 3.16. The element of LB_n corresponding to the permutation $\prod_{i=1}^{\lfloor m/2 \rfloor}(i, m+1-i)$ is denoted by D_m. Notice that D_m is the image of Δ_m under the standard surjection $B_m \to \mathcal{S}_m$.

The *dualizer* is the element $d_m := \Delta_m D_m \in LB_m$.

It is clear that $\rho D_m = \rho^{\intercal}$, so we have the following theorem justifying the name dualizer:

Theorem 3.17. *If x is an e-graph with m edges, or a factorization with m monodromies, we have*

$$x^* = x d_m.$$

4. Properly embedded graphs

The reader may have noticed the close analogy of the mind-body dual with the dual of a graph embedded in a (closed oriented) surface. In this section we elaborate on that analogy. We refer the reader to [28] and [32] for the rich theory of graph embeddings and maps. Most of the constructions in this section are entirely analogous to the usual case of embeddings in a closed surface, so some of the details are skipped trusting the reader to supply them.

Definition 4.1. Let F be an oriented surface with boundary and Γ a graph. A *proper embedding* of Γ into F is an embedding $i : \Gamma \to F$ such that:

(1) The vertices of Γ are mapped in the boundary of F, i.e., $i(V) \subset \partial F$, and the interior of each edge of Γ is mapped into the interior of F.

(2) $F \setminus i(\Gamma)$ is a disjoint union of simply connected subsets, called the *regions* of the embedding, and the interior of each region is homeomorphic to an open disc.

(3) $\partial F \setminus i(V)$ is a disjoint union of open intervals, called the *arcs* of the embedding, and the closure of each region contains exactly one arc.

A *properly embedded graph* (peg for short) is a graph endowed with a proper embedding into a surface. We will abuse notation by not distinguishing a peg from its image, and we will use the same symbol to denote a peg, its underlying graph, or even the surface that the graph is embedded.

An *isomorphism of pegs* is an orientation preserving homeomorphism of their surfaces that restricts to a graph isomorphism on the images of the graphs. Two pegs are called *isomorphic* if there is an isomorphism between them.[8]

[8]More nuanced notions of maps and equivalences between pegs will be considered in the planned work [3].

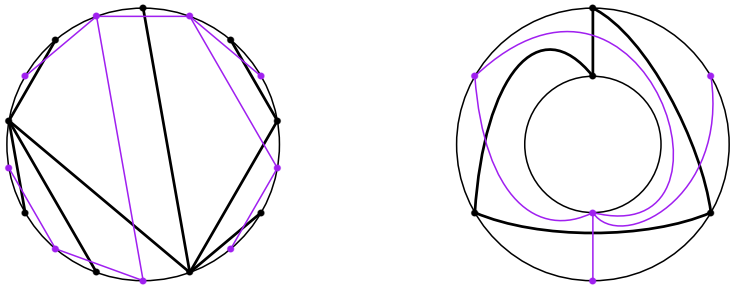

Figure 16. Properly embedded tree (left) and unicycle (right).

We emphasize that, contrary to the usual convention in the theory of graph embeddings, neither Γ nor the surface F is assumed connected. We also emphasize that the regions are not open or closed and that their closure is not necessarily a closed disk. For example, the closure of one of the regions of the peg in the right side of Figure 16 is a "pinched annulus".

The prototype of a peg is a *noncrossing tree*. Noncrossing trees are well studied in the literature, see, for example, [20] and [36], from our perspective a noncrossing tree is a tree properly embedded in a disk. The left side of Figure 16 shows a noncrossing tree, and the right side shows a unicycle[9] properly embedded in an annulus (drawn with thick black lines). See also the right side of Figure 18 that shows the graph of Figure 1 properly embedded in a torus with a disk removed.

If Γ is properly embedded in F then the vertices and arcs endow the boundary with the structure of a 1-complex, i.e., a graph. The orientation of ∂F further endows it with a digraph structure, and clearly that digraph is the functional digraph of a permutation.

Definition 4.2. The *monodromy digraph of a peg* Γ is the inverse of the digraph ∂F described above. The *monodromy* of Γ is the permutation of V defined by this digraph. As usual we will use the same symbol $\mu(\Gamma)$ to denote the monodromy of a peg or its functional digraph.

Given a graph Γ pegged in the oriented surface F, the orientation of F induces a leo structure on Γ: for any vertex v order the edges in $\nu(v)$ starting from the rightmost one and proceeding according to the orientation.

The leo structure of the definition above means that a peg Γ has a *medial digraph*. That digraph can be considered actually embedded (not properly) in F: simply put the vertex of M in the middle of the edge of Γ that it corresponds to, and draw the edge that connects two consecutive edges inside the region of Γ in whose boundary they lie.

Given an e-graph (or more generally a leo) there is a natural way to construct a peg, analogous to the way one defines a cellularly embedded graph from a rotation scheme, see [28], and the discussion in Section 4B. One obtains a sort of *ribbon graph* with the vertices in the boundary instead of the interior. We describe this procedure with more details below.

Definition 4.3. The *peg associated with a leo* Γ is obtained as follows: Thicken the star of each vertex, into a ribbon surface as shown in Figure 17 (for a vertex of degree 4), so that the vertex is at the boundary and each edge is thickened into a ribbon with the ribbons arranged according to the local order around

[9]That is a connected graph with exactly one cycle.

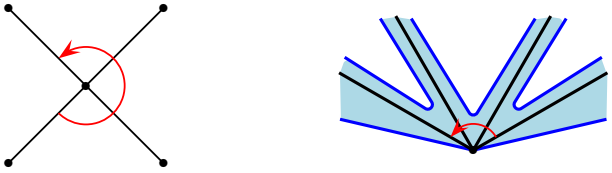

Figure 17. Thickening the star of a vertex.

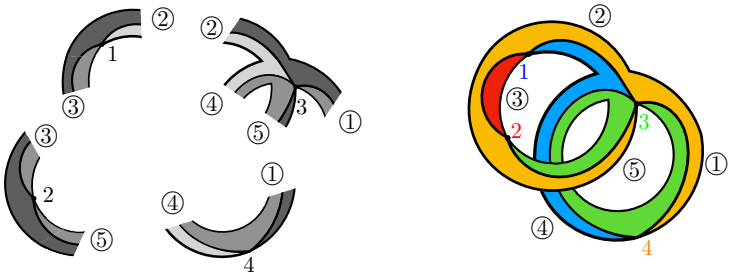

Figure 18. The peg of the e-v-graph of Figure 1.

that vertex. Then glue, *via an orientation reversing homeomorphism of the interval* the free ends of any pair of ribbons that correspond to the two half edges of the same edge of Γ together to get a surface F with Γ embedded in it in such a way that the vertices are on the boundary. We use $P(\Gamma)$ to denote the peg of a leo Γ.

For example, the left side of Figure 18 shows the thickening of the stars of the vertices of the e-graph of Figure 1 while the right side shows them assembled into a peg.

Notice that the boundary of each region (colored with different colors) is the union of a migt of Γ and an arc of $\partial P(\Gamma)$. In fact one can easily see that this is always the case:

Lemma 4.4. *$P(\Gamma)$ is indeed a peg and is isomorphic to the peg obtained by attaching a half disk, that is, a domain homeomorphic to $\{(x, y) \in \mathbb{D}^2 : y \geq 0\}$, along each migt of Γ, in such a way that the part of the boundary of the half disk that lies on the real axis is identified with the migt.*

Proof. It's clear that the resulting complex is an orientable ribbon surface, with the vertices of the graph in the boundary, and the interior of the edges in the interior. To see that the regions are indeed open discs with exactly one arc in their boundary, notice that each thickened edge separates its ribbon in to two simply connected regions whose interior is an open disc, and when we glue the half edges together the four regions are glued together according to the migts of Γ because the gluing is happening via an orientation reversing homeomorphism. □

We note that an alternative construction of $P(\Gamma)$ as the total space of a branched covering over the disc is given in Section 4D below, see Theorem 4.28.

We have the following general properties:

Lemma 4.5. *Let Γ be an e-graph. Then*:

(1) *$P(\Gamma)$ has the same number of connected components as Γ.*

(2) $\chi(P(\Gamma)) = \chi(\Gamma)$.

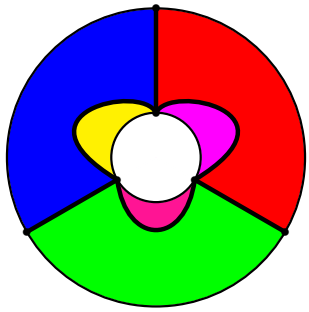

Figure 19. A peg that doesn't come from an e-graph.

(3) *The monodromy digraph of* $P(\Gamma)$ *is isomorphic to the monodromy digraph of* Γ.

(4) *If* Γ *is connected and* $\mu(\Gamma) \in \mathcal{S}_V$ *has* b *disjoint cycles then genus of* $P(\Gamma)$ *is*

$$g = \tfrac{1}{2}(2 + m - n - b).$$

Proof. Clearly the surface of $P(\Gamma)$ can be homotopically collapsed to Γ so items (1) and (2) are immediate. Item (3) follows from the fact that the boundary of a region of $P(\Gamma)$ consists of a migt and an arc, and since $P(G)$ is oriented, the migt and the arc have opposite orientations. Item (4) is also clear, since the resulting surface has b boundary components. □

We remark that the notion of peg is more general than that of e-graph. Indeed given any leo Γ, e-realizable or not, one can define the peg $P(\Gamma)$. For example, the peg of Figure 19 does not come from an e-graph, in fact it is $P(\Gamma)$ for the non-e-realizable leo of Example 2.14 (see Figure 7).

Definition 4.6. An e-peg is a peg that is $P(\Gamma)$ for some e-graph Γ.

4A. The dual of a peg. We can now interpret mind-body duality in terms of pegs. Given a peg there is a natural way to define its dual analogous to the way one defines the dual of a cellularly embedded graph.

Definition 4.7. The *dual peg* of a peg Γ is the peg Γ^* obtained as follows:

- If Γ is embedded in the surface F then Γ^* is embedded is $F^\intercal$, that is, F endowed with the opposite orientation.
- The vertices of Γ^* are in one-to-one correspondence with the regions of Γ; when we draw Γ^* we place its vertices on the arcs of the corresponding regions.
- The edges of Γ^* are in one-to-one correspondence with the edges of Γ, the edge e^* that corresponds to the edge e connects the vertices of Γ^* that correspond to the two regions of Γ that e lies in the boundary of.

For example, the duals of the pegs in Figure 16 are shown with lighter purple lines. The reason we consider Γ^* to be embedded in the oppositely oriented surface than that of Γ will become clear later in this section.

This notion of duality for noncrossing trees is implicit in [27], although they do not explicitly consider the dual as a noncrossing tree. An explicit notion of dual for noncrossing trees was defined in [29], however properly speaking the notion defined there, though closely related to ours, is not a real duality

since it fails to be involutory; rather it reflects the action of the Garside element (see Section 3). See also Section 6B.

Since the boundary of each region of Γ contains exactly one arc and exactly one migt of Γ, and two regions share an edge if and only if the corresponding migts do, the following proposition is clear.

Proposition 4.8. *If Γ is an e-graph then, the peg that corresponds to Γ^* is the dual of the peg that corresponds to Γ, i.e., $P(\Gamma^*) = P(\Gamma)^*$.*

4B. Relation with cellularly embedded graphs. There is a direct connection with the usual cellular embeddings to a closed surface. We will use the abbreviation *ceg* to refer to a cellularly embedded graph in a closed surface.

As noted the construction of a peg from a leo, given in the previous subsection is directly analogous to the construction of a ceg from a rotation system. Using that we can see that any peg can be "completed" to a ceg.

Definition 4.9. The *closure* of a compact oriented surface with boundary F is the surface $\overline{F}$ obtained by F by attaching 2-cells along its boundary components.

If $i : \Gamma \to F$ is a peg then its *closure* is the ceg $\bar{i} : \Gamma \to \overline{F}$ obtained by composing i with the natural inclusion $F \hookrightarrow \overline{F}$. When we abuse notation and refer to a peg by its underlying graph Γ we will further abuse the notation by denoting its closure by $\overline{\Gamma}$.

As we have already remarked, the construction of the peg of an e-graph is very similar to the construction of a ceg given a rotation system. In fact the following is clear.

Lemma 4.10. *The rotation system of $\overline{\Gamma}$ is obtained by the leo of Γ by "completing" each local edge order to a cyclic order.*

A natural question that arises at this point is what cegs are completions of an e-peg? An obvious necessary condition is that there should be at least as many vertices as regions, but as we will see shortly this is not sufficient, for example, we will see that the ceg in Figure 22 does not come from an e-graph. In order to answer this question we introduce the concept of the *medial digraph* of a ceg.

The medial graph $\mathcal{M}(\Gamma)$ of a ceg Γ is defined (see, for example, [7]) as the graph that has vertices the edges of Γ, and an edge connecting any two vertices that correspond to a pair of consecutive edges in the rotation system given by the embedding. Since we consider embeddings into *oriented* surfaces we observe that $\mathcal{M}(\Gamma)$ has a natural orientation given by the cyclic order, so we will refer to $\mathcal{M}(\Gamma)$ as the *medial digraph* of the ceg Γ.

The medial digraph M of a ceg Γ is embedded in the same surface that the original graph is. Indeed one can place the vertex of the medial digraph in the middle of the corresponding edge of the graph, and draw the edge connecting the vertices that correspond to two consecutive edges inside the region that has the corresponding edges in its boundary. If Γ has n vertices then the edges of $\mathcal{M}(\Gamma)$ can be colored with n colors corresponding to the vertices of Γ, where all the edges corresponding to a rotation around a vertex are colored by the color of that vertex. Notice that if e is an edge then at the corresponding vertex of $\mathcal{M}$ there are edges of two different colors, say blue and red and the cyclic order induced from the embedding is red-in, blue-out, blue-in, red-out.

The regions of the medial digraph correspond to either vertices, or regions of the original graph, and their boundaries consist of (oriented) cycles. The boundary cycle of a region that corresponds to a vertex

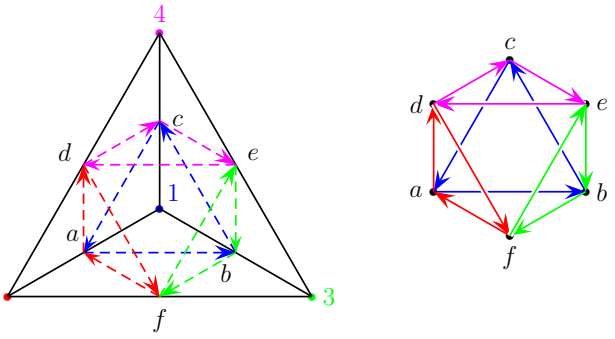

Figure 20. An embedded K_4 and its medial digraph.

of the original graph is *monochromatic*, and the boundary cycle of a region that corresponds to a region of the original graph is *properly edge colored*, that is any two adjacent edges are colored with different colors.[10] These concepts are illustrated in Figure 20 for the standard genus 0 embedding of the complete graph on four vertices.

The medial digraph of the dual of a ceg is isomorphic to the medial digraph of the original graph, or its inverse depending on whether we consider the dual ceg embedded in the oppositely oriented, or the same surface as the original. We make the convention that *for a graph* Γ *embedded in the closed surface* F, *it's dual* Γ^* *is embedded in* $F^\intercal$.

If Γ is a peg then the medial digraph of the ceg $\overline{\Gamma}$ is obtained from $\mathcal{M}(\Gamma)$ by adding an edge (of the appropriate color) from the last to the first vertex of every chain in the PCD of Γ. Removing those edges from $\mathcal{M}(\overline{\Gamma})$, makes it into a dag.

Definition 4.11. A *feedback arc set* (FAS) in a digraph D is a set of edges whose removal converts D into a dag (see, for example, [9, Chapter 10].)

A set of edges in an edge-colored digraph is called *diverse* if it contains a representative of each color.

Clearly a ceg Γ that comes from a factorization, i.e., it is the completion of an e-peg, has a diverse FAS, since each vertex of Γ gives a monochromatic cycle in $\mathcal{M}(\Gamma)$. With these concepts in place we can now state that:

Proposition 4.12. *The ceg* Γ, *with* n *vertices*, *comes from a factorization exactly when its medial digraph has a FAS of cardinality* n.

Proof. If Γ is an e-peg, then $\mathcal{M}(\Gamma)$ is a dag contained in, and having exactly n edges less than, $\mathcal{M}(\overline{\Gamma})$.

Conversely, since the medial digraph $\mathcal{M}(\Gamma)$ of a ceg has n monochromatic cycles, a FAS S with n edges will contain an edge in the cycle of each vertex. Removing S therefore will give a leo structure on Γ, whose medial digraph is a dag, and whose completion will be the rotation system of Γ. □

As an example consider the genus 0 embedding of K_4 in Figure 20. The set of edges $\{de, fa, bd, eb\}$ is a diverse FAS for its medial digraph. Taking a topological sort of the resulting dag give us the factorization $(1, 2), (1, 3), (2, 4), (1, 4), (2, 3), (3, 4)$ of id, see Figure 21.

As an example of an nonpeggable ceg consider the plane graph in Figure 22. It's clear that its colored digraph does not admit a diverse FAS so it does not come from an e-graph.

[10] Recall that we consider loopless graphs. For graphs with loops this observation is not necessarily true.

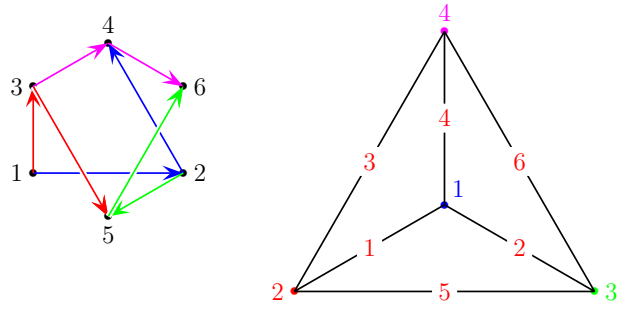

Figure 21. An e-labeling for the embedded K_4 in Figure 20.

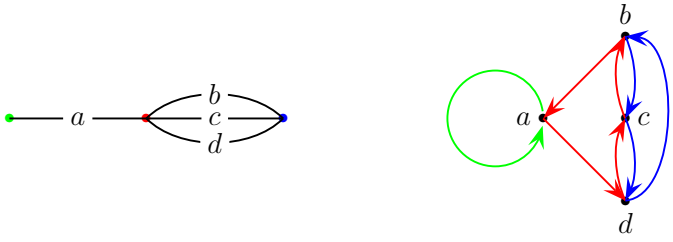

Figure 22. A nonpeggable ceg and it's medial digraph.

In general, if Γ is peg then $\overline{\Gamma^*}$ and $(\overline{\Gamma})^*$ are not isomorphic as graphs. If $\rho = (1\,3), (1\,2), (1\,3)$ then $\rho^* = (1\,3), (2\,3), (1\,2)$, and the reader can easily check that the duals of the completions of the corresponding pegs, which are planar, are different.

However, one can easily see that:

Theorem 4.13. *If for a peg* Γ *we have* $\mu(\Gamma) = \mathrm{id}$ *then* $(\overline{\Gamma})^* = \overline{\Gamma^*}$.

Proof. If $\mu(\Gamma) = \mathrm{id}$ then each boundary component of $P(\Gamma)$ contains exactly one point, and the closure of each region is a "pinched" annulus, and so the regions of $\overline{\Gamma}$ are in one-to-one correspondence with the regions of Γ, and this correspondence obviously preserves incidence relations between regions and edges. Since the vertices, edges, and regions of Γ and $\overline{\Gamma}$ are in one-to-one correspondence and those correspondences preserve incidence relations, it follows that $(\,\overline{\Gamma}\,)^* = \overline{\Gamma^*}$. □

Example 4.14. It is known that the complete graph K_n admits a self-dual[11] embedding into a closed oriented surface, if and only if, $n \equiv 0$ or $1 \pmod 4$, (see [37]). These are exactly the degrees for which the complete graph has an even number of edges, and so it's possible to give K_n an e-labeling with monodromy equal to the identity. We can ask then, whether the theory of e-pegs we developed can be used to prove this result. In this work we will give a simple proof for the cases $n = 4$ and $n = 5$ to illustrate the basic ideas.

For $n = 4$, start with the self-dual factorization (see Corollary 3.8)

$$\rho_{0;4} := (1\,2), (1\,2), (2\,3), (2\,3), (3\,4), (3\,4)$$

[11] In the sense that the underlying graph of the dual ceg is also complete.

and notice that if $\beta = \sigma_2\sigma_4\sigma_3^{-1}$ then $\rho_0\,\beta = (1\,2), (1\,3), (2\,4), (1\,4), (2\,3), (3\,4)$, whose associated graph is complete is the complete e-v-graph on the right side of Figure 21. On the other hand, we also have that $\rho_0\,\beta^* = (1\,2), (2\,3), (1\,4), (1\,3), (2\,4), (3\,4)$, whose associated graph is also complete. Using Theorems 4.13 and 3.10 we conclude that the e-labeling of K_4 in Figure 21 gives a self-dual embedding of K_4 into the sphere.

For $n = 5$, one can start with the self-dual factorization

$$\rho_{0;5} := (1\,2), (1\,2), (2\,3), (2\,3), (3\,4), (3\,4), (4\,5), (4\,5), (1\,5), (1\,5)$$

and observe that if $\beta = \delta_{2,10}\,\sigma_3^{-1}\sigma_5^{-1}\sigma_7^{-1}\sigma_9^{-1}$ then

$$\rho_{0;5}\,\beta = (1\,2), (2\,5), (2\,3), (1\,3), (3\,4), (2\,4), (4\,5), (3\,5), (1\,5), (1\,4)$$

and

$$\rho_{0;5}\,\beta^* = (1\,2), (1\,5), (3\,5), (2\,5), (2\,4), (2\,3), (1\,3), (3\,4), (4\,5), (1\,4)$$

both factorizations with complete associated e-v-graphs, thus giving self-dual embeddings of K_5 into a torus.

Remark 4.15. Even though K_6 does not admit self-dual embeddings into a closed surface, it can be self-dually pegged in surfaces with boundary. Indeed the factorization

$$\begin{aligned}\rho_1 := {} & (1,2), (3,5), (1,3), (4,6), (2,4), (1,4), (5,6), (1,6),\\ & (2,3), (2,5), (1,5), (3,4), (4,5), (2,6), (3,6)\end{aligned}$$

pegs K_6 into a torus with three boundary components, and the factorization

$$\begin{aligned}\rho_2 := {} & (1,2), (3,6), (1,3), (4,5), (4,6), (2,4), (2,3), (1,5),\\ & (1,4), (5,6), (3,4), (2,5), (3,5), (1,6), (2,6)\end{aligned}$$

pegs it, self-dually, into a genus 2 surface with one boundary component.

In general we can use mind-body duality to get self-dual embeddings of graphs into closed surfaces by gluing a pair of dual pegs along their common boundary.

Definition 4.16. Let P_1, P_2 be two pegs, with $\mu(P_2) = \mu(P_1)^{-1}$ and $f : \partial P_1 \to \partial P_2$ an orientation reversing homeomorphism, that maps the vertices of P_1 to the vertices of P_2. The *boundary connected sum* of P_1 and P_2, with respect to f is the ceg $P_1 \#_f P_2$ defined as follows:

- The surface of $P_1 \#_f P_2$ is the boundary connected sum of the surfaces of P_1 and P_2 with respect to f.
- $P_1 \#_f P_2$ has the same vertices as P_1 and P_2, and edges the union of the edges of P_1 and the edges of P_2.

Clearly $P_1 \#_f P_2$ is a ceg and each of its regions is obtained by gluing a region of P_1 with the region of P_2 that has the same boundary arc, along their common boundary. Now we can prove:

Theorem 4.17. *For every peg P, the boundary connected sum $P \#_{\mathrm{id}} P^*$ is self-dually embedded.*

Proof. Let C be the boundary connected sum. First since the orientations of P and P^* are opposite id is orientation reversing so C is defined. Now observe that each region of C is obtained by gluing a region of P and a region of P^* along their boundary arc. If we choose the vertex dual to a region to lie along that common arc, we see that $C^* = P^* \#_{\mathrm{id}} P$. □

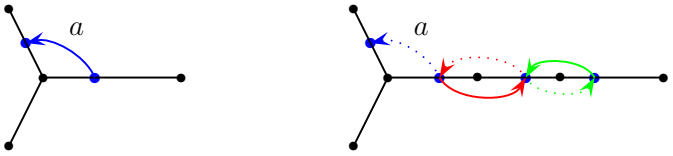

Figure 23. Subdividing to get a diverse FAS.

Example 4.18. Consider $\overrightarrow{K}_6$ the complete digraph with six vertices, any two of which are connected by a pair of opposite edges. Theorem 4.17 and Remark 4.15 imply that $\overrightarrow{K}_6$ admits self-dual embeddings into a surface of genus 4. Furthermore this is a digraph embedding in the sense of [14], i.e., the boundary of every region is a directed cycle.

We end this subsection by remarking that the theory of pegs is a refinement of the theory of cegs, more attuned to the graph theoretical properties of the graph. For example, if a graph is cellularly embeddable into a closed surface then so is any graph homeomorphic to it. This is not the case with pegs. Indeed we have:

Proposition 4.19. *Any ceg has a subdivision that is the closure of an e-peg.*

Proof. Let Γ be a ceg. Subdividing an edge of Γ, adds a pair of opposite arcs with a new color to the medial digraph of Γ. Chose a FAS S for $\mathcal{M}(\Gamma)$, and let S' be a minimal diverse subset of S (it's clear that any FAS contains a diverse set). Now for every arc a in $S \setminus S'$ subdivide the edge of Γ that corresponds to the beginning vertex of a, twice, and then replace, a with a pair of the new edges, each going in opposite direction. The resulting set S'' is a diverse FAS for the medial digraph of the subdivided ceg. See Figure 23. □

4C. On the genus and number of boundary components of e-pegs. For an e-graph Γ, the Euler characteristic of $P(\Gamma)$ is the Euler characteristic of Γ, and the number of boundary components b equals the number of cycles in the disjoint cycle decomposition of its monodromy $\mu(\Gamma)$. A natural question that arises is: given a graph Γ what can we say about the values of b that arise from the different edge-labelings of Γ? The following proposition provides two obvious necessary, but not in general sufficient, conditions.

Proposition 4.20. *For every edge-labeling of* Γ, *the number* b *of disjoint cycles of* $\mu(\Gamma)$, *or equivalently the number of boundary components of* $P(\Gamma)$, *satisfies*

$$b \leq n \quad \textit{and} \quad b \equiv \chi(\Gamma) \pmod 2.$$

Proof. Every boundary component of $P(\Gamma)$ contains at least one vertex of Γ, so $b \leq n$. Also, since the genus of a closed surface is an integer, by item (4) of Lemma 4.5 we infer that b has the same parity as $\chi(\Gamma)$. □

To see that these conditions are not sufficient consider the graph Γ shown on the left side of Figure 22: we have $n = 3$, and $\chi(\Gamma) = -1$ so the value $b = 3$ satisfies both conditions, but, as one can easily see, for every edge-labeling of Γ the monodromy is a 3-cycle, so that $b = 1$.

Question 4.21. For what class of graphs are the conditions of Proposition 4.20 sufficient?

We don't know the complete answer but we can prove that complete graphs belong in that class.

Theorem 4.22. *For every complete graph K_n all values of b allowed by Proposition 4.20 occur. Actually all possible conjugacy classes of μ consistent with Proposition 4.20 occur.*

In the proof of Theorem 4.22 we will use the following lemma, the proofs of which are straightforward.

Lemma 4.23. *Let Γ be an e-graph of size m with $\mu(\Gamma)$ a d-cycle, and let Γ' be the graph obtained from Γ by adding one new vertex v and connecting it by an edge labeled $m+1$ to a vertex u that is moved by the cycle $\mu(\Gamma)$. Then $\mu(\Gamma')$ is a $(d+1)$-cycle; namely if $\mu(\Gamma) = (\ldots, u', u, u'', \ldots)$ then $\mu(\Gamma') = (\ldots, u', v, u, u'', \ldots)$.*

Definition 4.24. Let v be a vertex, and e an edge of the e-graph Γ not incident to v. Then a *T-operation from the vertex v on the edge e* is the following modification: If i is the label of e, increase all edge-labels greater than i by 2, and change i to $i+1$. Then add two new edges labeled i and $i+2$ connecting v to the endpoints of e. See Figure 24.

Lemma 4.25. *The operation T doesn't change the monodromy of the graph.*

Proof of Theorem 4.22. We proceed by induction on n. The theorem is obvious for $n = 1, 2, 3$ and proven in Figure 25 for $n = 4$. Assume then the theorem proven for all values less than n.

If $n \equiv 0 \pmod 4$, then μ has to induce a partition into an even number of parts b. If $b = n$ then $\mu = \mathrm{id}$. To construct an e-labeling of K_n with $\mu(K_n) = \mathrm{id}$, start with an e-labeling of K_{n-4} with $\mu(K_{n-4}) = \mathrm{id}$ and an e-labeling of K_4 using labels, $\binom{n}{2} - 6, \ldots, \binom{n}{2}$ with $\mu(K_4) = \mathrm{id}$. Then chose a partition of the vertices of K_{n-4} into $\frac{1}{2}\binom{n-4}{2}$ pairs, and apply a T operation from each vertex of K_4 on each of the edges of K_{n-4} determined by those pairs. The result is a K_n with monodromy equal to id.

If $b < n$, and the corresponding partition is $k_1 + \cdots + k_b = n$, with $k_b > 1$, chose an e-labeled K_{n-1} with monodromy of type $(k_1, \ldots, k_b - 1)$ and add a new vertex v. Choose a vertex u of K_{n-1} that belongs in the $(k_b - 1)$-cycle and partition the rest of the vertices of K_{n-1} into pairs. Now add an edge connecting v and u and label it $\binom{b-1}{2} + 1$, resulting in a graph with monodromy of type $(k_1, \ldots, k_b)$. Then apply T

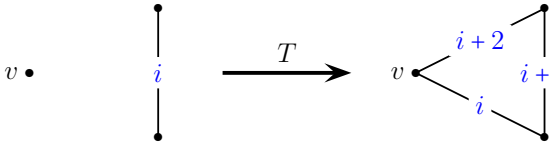

Figure 24. The T-operation.

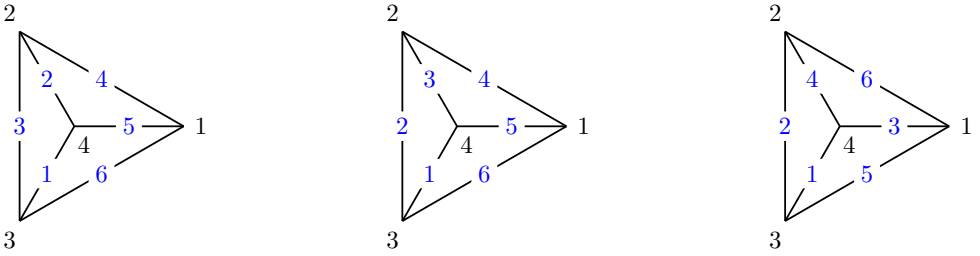

Figure 25. Proof of Theorem 4.22 for $n = 4$: $\mu = (1, 2, 3)$ (left), $\mu = (1, 2)(3, 4)$ (center) and $\mu = \mathrm{id}$ (right).

operations from v to the edges determined by the partition into pairs of the remaining vertices of K_{n-1}. At the end we get an e-labeled K_d whose monodromy has type $(k_1, \ldots, k_b)$.

If $n \equiv 1 \pmod 4$, then the possible values of b are $1, 3, \ldots, n$. Let $n = k_1 + \cdots + k_b$ be the partition induced by μ. If k_1 is odd, by the inductive hypothesis we can find an e-labeled K_{n-1} with monodromy of type $(k_2, \ldots, k_b, 1, \ldots, 1)$, where there are $k_1 - 1$ ones (if $k_1 = 1$ there are no ones). Add a new vertex and connect it to the $k_1 - 1$ fixed points of $\mu(K_{n-1})$ with edges labeled $\binom{n-1}{2} + 1, \ldots, \binom{n-1}{2} + k_1 - 1$. The resulting graph has monodromy of type $(k_1, \ldots, k_b)$. At this stage, there are an even number of vertices of K_{n-1} not connected to the new vertex, so we can partition them into pairs and apply T moves from the new vertex to get an e-labeled K_n with the desired monodromy.

If k_1 is even, then we start with an e-labeled K_{n-1} with monodromy of type $(k_2, \ldots, k_b - 1, 1, \ldots, 1)$, where there are k_1 ones. Add a new vertex v and connect it to one of the vertices of the $(k_b - 1)$-cycle with an edge labeled $\binom{n-1}{2} + 1$, and to $k_1 - 1$ of the fixed points of $\mu(K_{n-1})$ by edges labeled $\binom{n-1}{2} + 2, \ldots, \binom{n-1}{2} + k$. The resulting graph has monodromy of type $(k_1, \ldots, k_b)$, and there are an even number of vertices not connected to the new vertex. Then we can use T moves to get an e-labeled K_n with the desired monodromy.

If $n \equiv 2 \pmod 4$, then $b = 1, \ldots, n-1$. If $n = k_1 + \ldots + k_b$ is a partition of n then $b_K > 1$ and we can choose a K_{n-1} with monodromy of type $(k_1, \ldots, k_b - 1)$. Add a new vertex and connect it to a vertex of the $(k_b - 1)$-cycle by a edge labeled $\binom{n-1}{2} + 1$. The result is a graph with monodromy of type $(k_1, \ldots, k_b)$ and an even number of edges unconnected to the new vertex. So we can apply T-moves from the new vertex to get an e-labeled K_n with the desired monodromy.

If $n \equiv 3 \pmod 4$, then $b = 2, \ldots, n-1$. Let $n = k_1 + \ldots + k_b$ be a partition of n. If k_1 is odd, choose a K_{n-1} with monodromy of type $(k_2, \ldots, k_b, 1, \ldots, 1)$, where there are $k_1 - 1$ ones. Add a new vertex and connect it to the fixed points of $\mu(K_{n-1})$ by edges labeled $\binom{n-1}{2} + 1, \ldots, \binom{n-1}{2} + k_1 - 1$. This gives a graph with monodromy of the right type and an even number of vertices not connected to the new vertex. So we can use T-moves to complete the proof.

If k_1 is even, choose a K_{n-1} with monodromy of type $(k_2, \ldots, k_b - 1, 1, \ldots, 1)$, where there are k_1 ones. Add a new vertex and connect it to a vertex on the $k_b - 1$-cycle by an edge labeled $\binom{n-1}{2} + 1$, and to $k_1 - 1$ fixed points of $\mu(K_{n-1})$ with edges labeled, $\binom{n-1}{2} + 2, \ldots, \binom{n-1}{2} + k$. This gives a graph with the right monodromy and an even number of edges not connected to the new vertex. So again we can use T-moves to complete the proof. This completes the inductive step and the proof of the theorem. □

Item (4) of Lemma 4.5 implies that in general peggable embeddings have genera in the upper part of the genus range of a graph. These and related topics will be addressed in more detail in [3].

4D. Branched coverings interpretation. In this subsection we provide an alternative construction of the peg associated with a factorization via the theory of branched coverings of the two-dimensional disk $\mathbb{D}^2$.

Recall that the B_m-action on the free group comes from the fact that B_m is the *mapping class group relative to the boundary* of a $\mathbb{D}^2_m$, a disk with m punctures, while the fundamental group $\pi_1(\mathbb{D}^2_m)$ is a free group with m generators, see, for example, [12], [31], and [24]. A factorization ρ can be thought of as a representation $\pi_1(\mathbb{D}^2_m) \to \mathcal{S}_n$ and therefore gives a covering of $\mathbb{D}^2$ branched over m points, and the Hurwitz action can be thought of as an action of B_m to the set of (equivalence classes of) branched coverings over the disk $\mathbb{D}^2$. For details on branched coverings see [11] and [2], the later describes the Hurwitz action in detail.

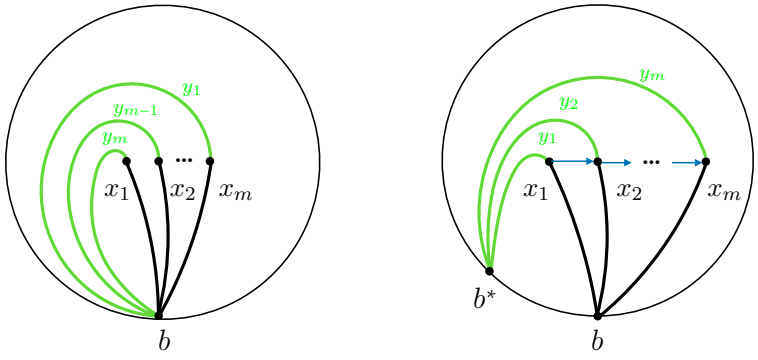

Figure 26. The standard Hurwitz system and its image under Δ (left), and its dual (right).

A free generating set of π_1 can be represented by a *Hurwitz system* i.e., an ordered system of arcs connecting each branching point to the basepoint of the disk, which we take to be on the boundary circle, and whose interiors are pairwise disjoint. An arc in the system represents the loop that starts at the basepoint, follows the arc to a small neighborhood of the puncture, goes once around the puncture counterclockwise and returns to the basepoint along the arc. To be concrete we consider the unit disk in $\mathbb{R}^2$, the basepoint b to be $(0, -1)$, while the m branching points to be equally spaced along the interval $[-1, 1]$, and we take the *standard Hurwitz system* h_0, to be the m straight line segments connecting the basepoint to the branching points, this is shown in the left side of Figure 26 in black. The B_m action on $\mathbb{F}_m$ determines a left action on the set of (isotopy classes of) Hurwitz systems, for example the image of h_0 under the action of the Garside element Δ_m is shown in the left side of Figure 26 in green.

A factorization in $\mathcal{S}_n$ gives a *simple* branched covering, that is a branched covering where the preimage of each branching point has only one singular point and $n-2$ regular points called *pseudosingular*. There is an explicit model for the simple branched covering corresponding to a factorization ρ (see, for example, [11] and [2]) and we briefly recall that construction. First choose a *cut system* for $\mathbb{D}^2_m$, consisting of m segments connecting each branch point to the boundary, to be explicit we take the m vertical segments in the upper half disk, and "cut" the disk open along this cut system. Then take n labeled copies of the disk (the sheets of the covering) and glue them together along the cuts and according to the monodromy sequence, that is, for $i = 1, \ldots, m$ if the i-th monodromy of ρ is $(k\, l)$ we glue the i-th cuts of the k-th and l-th sheet together, and "sew" back together the i-th cut of any other sheet. The surface resulting from all these gluings is the total space of the covering. This construction is illustrated in Figure 27 for the factorization $(1\,2), (2\,3)$.

Definition 4.26. The *essential preimage* of a Hurwitz system is defined to be the union of all the preimages of the arcs that contain a singular point.

The *reverse* of a Hurwitz system h is the Hurwitz system $h^\intercal$ that has the same arcs as h but in reverse order: the i-th arc of $h^\intercal$ is the $(m+1-i)$-th arc of h.

The *dual* of a Hurwitz system h is the Hurwitz system $h^* := (\Delta_m h)^\intercal$ with basepoint slightly to the left of the basepoint of h.

For example, the dual of the standard Hurwitz system is shown in green in the right side of Figure 26.

We can now prove the following theorem.

Theorem 4.27. *Let $p : F \to D^2$ be a simple branched covering. Then for any Hurwitz system h, the essential preimage of h is a graph pegged in the total space of the covering. Moreover the dual peg is the essential preimage of the dual Hurwitz system h^*.*

Proof. Let x be an arc of h, then $p^{-1}(x)$ consists of n arcs, one in each sheet of the covering, Of these preimages, only two are essential, and they meet at the same singular point. So x contributes to the essential preimage an arc connecting two points in the boundary of the total space, namely the preimages of b in the two sheets that are glued together in the cut that corresponds to x. It follows that the essential preimage of h is a graph Γ embedded in F with all the vertices in the boundary. To see that Γ is indeed properly embedded we first observe that the complement of the full preimage of h consists of n disjoint domains with one arc in their boundary, and interior homeomorphic to an open disc. Indeed, if we remove h from D^2 we are left with a contractible set, with exactly one arc on the boundary and its interior homeomorphic to a disc. That set has n homeomorphic preimages, that constitute the complement of the whole preimage of h in F. A component of the complement of Γ is obtained from a component of the complement of the full preimage, by possibly inserting some semiopen arcs, and it is easily seen that this still results in a contractible set with exactly one arc in its boundary and interior homeomorphic to an open disc.

The fact that the essential preimage of h^* is Γ^* follows from Theorem 3.4. □

From the explicit construction of the branched covering $p : F \to \mathbb{D}^2$, associated with a factorization ρ described above one can easily see that the essential preimage of the standard Hurwitz system h_0 is a graph isomorphic to $\Gamma(\rho)$. Indeed if the k-th monodromy of ρ is $(i\ j)$, then the whole preimage of the k-th arc of h consists of $n-2$ "short" arcs connecting the $n-2$ preimages of the basepoint x_0 in the sheets with labels different than i, j to the $n-2$ points in the pseudosingular locus above x_k, and one "long" arc that consists of the two preimages of the arc that start at the sheets labeled i and j and are glued together at the singular point above x_k, see Figure 27, for an example when $m = 2$ and $n = 3$. As a corollary then we have an alternative construction of the peg associated with an e-graph Γ, and its mind-body dual:

Theorem 4.28. *For a factorization ρ, the peg $P(\Gamma(\rho))$ is the essential preimage of the standard Hurwitz system in the branched covering determined by ρ. Furthermore it's dual peg is the essential preimage of h_0^*.*

An example of the theorem is shown in Figure 27 for the factorization $(1, 2)$, $(2, 3)$ and its mind-body dual $(1, 2)$, $(1, 3)$.

Remark 4.29. The medial digraph of the peg $P(\Gamma)$, for an e-graph Γ can also be interpreted via the associated branched covering. Consider the $n-1$ oriented intervals connecting x_i to x_{i+1}, for $i = 1, \ldots, n-1$, shown in blue on the right side of Figure 26. Clearly the subset of the preimage of those intervals, consisting only of those arcs that connect two singular points is a digraph isomorphic to the medial digraph of Γ.

The relation of e-labeled graphs with branched covering was observed in [8] (see also [32]). However, they consider branched coverings over the sphere S^2, by adding an additional branched point with monodromy $\mu(\Gamma)^{-1}$, so that in effect they obtain the ceg $\overline{\Gamma}$.

If ρ is a factorization of the identity permutation, it determines not only a branched covering p of the 2-disk but also a branched cover $\bar{p}$ of the sphere 2-sphere S^2. This is so because the fundamental

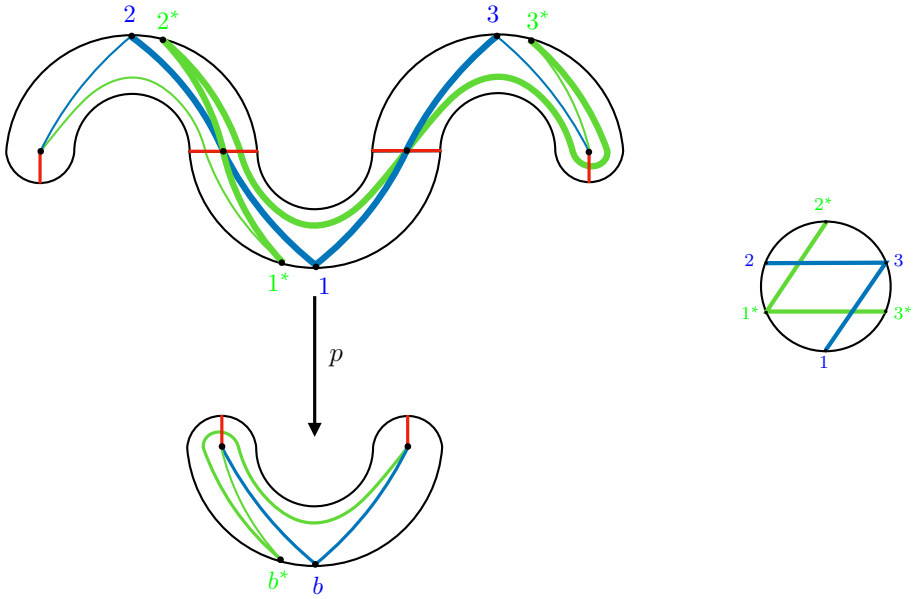

Figure 27. The branched covering (1 2), (2 3) and its dual (1 2), (1 3).

group of S^2 with m punctures has a presentation with m generators $x_1, \ldots, x_m$, corresponding to loops going around each puncture, and a single relation $x_1 \ldots x_m = \mathrm{id}$. In that case the essential preimage of the Hurwitz system under $\bar{p}$ is a ceg, and it is easy to see that it is the completion of the peg obtained as the essential preimage under p. One can also easily see the following:

Theorem 4.30. *Let ρ_1 and ρ_2 be factorizations with $\mu(\rho_2) = \mu(\rho_1)^{-1}$, and P_1, P_2 the corresponding pegs. Then $P_1 \#_{\mathrm{id}} P_2$ is the ceg obtained from the concatenation of the factorizations ρ_1 ρ_2 interpreted as a branched cover of the 2-sphere.*

5. On cycles and trees

The research on this paper started by the author reading [27]. That paper provides a "structural" bijection ϕ from the set of minimal factorizations of an n-cycle to the set of trees on $[n]$. In this section we give an exposition of that and related topics in light of the present work.

The set of all cyclic permutations on $[n]$ is denoted by $\mathscr{C}_n$ and we choose the *standard cyclic permutation* to be $\zeta_0 := (n, n-1, \ldots, 1)$. $\mathscr{E}_n$ stands for the set of edge-labeled trees (*e-trees* for short), with n vertices, $\mathscr{E}_n^*$ for the set of rooted e-trees with n-vertices, $\mathscr{V}_n$ for the set of vertex labeled trees (*v-trees* for short) with n vertices, and $\mathscr{L}_n$ for the set of edge and vertex labeled trees (*e-v-trees* for short).

We denote the set of all factorizations in $\mathscr{S}_n$ of length m by $\mathscr{F}_m(n)$. For a permutation $\pi \in \mathscr{S}_n$ and $m \in \mathbb{N}$, the set of all factorizations of π as a product of m transpositions is denoted by $\mathscr{F}_m^\pi$.

5A. Bijections between $\mathscr{F}_n^{\zeta_0}$ and $\mathscr{V}_n$.

In [17], Dénes proved that:

Theorem 5.1. *The graph of a factorization $\rho \in \mathscr{F}_{n-1}(n)$ is a tree if and only if $\mu(\rho) \in \mathscr{C}_n$.*

Proof. If the graph $\Gamma(\rho)$ is a tree, then by item (2) of Lemma 4.5 we have that $P(\Gamma)$ is an orientable surface with Euler characteristic 1, and therefore a disk. It follows that all vertices of Γ lie in a circle and therefore $\mu(\rho)$ is an n-cycle.

Conversely, if $\mu(\rho)$ is an n-cycle, $P(\rho)$ is connected and therefore $\Gamma(\rho)$ is a connected graph with Euler characteristic $n-(n-1)=1$. Therefore $\Gamma(\rho)$ is a tree. □

Using this result one can then establish that:

Theorem 5.2 [17]. *For any $\zeta \in \mathscr{C}_n$ there is a bijection $f_\zeta : \mathscr{F}^\zeta \to \mathscr{E}_n^*$.*

Proof. Given $\rho \in \mathscr{F}^\zeta$ we obtain $f(\rho)$ by taking the corresponding graph in $\mathscr{L}_n$, declaring, say, 1 to be the root and forgetting the v-labels. To go back, starting from a rooted e-tree t, by Theorem 5.1, $\mu(t)$ is an n-cycle in $\tau \in \mathscr{S}_V$ and we can label the vertices of t so that the root is labeled 1 and arranging so that the label of $\tau(v)$ is the image under ζ of the label of v for all vertices v. □

Putting all these $(n-1)!$ bijections together one obtains:

Theorem 5.3 [17]. *There is a bijection $D : \mathscr{C}_n \times \mathscr{E}_n^* \to \mathscr{L}_n$.*

Corollary 5.4. $$|\mathscr{F}_{n-1}^{\zeta_0}| = n^{n-2}.$$

Proof. By Cayley's result, see, for example, [34], L_n has cardinality $(n-1)!\,n^{n-2}$ and since $\mathscr{C}_n$ has cardinality $(n-1)!$ it follows by Theorem 5.3 that $\mathscr{E}_n^*$ has cardinality n^{n-2} and therefore so does $\mathscr{F}_{n-1}^{\zeta_0}$ by Theorem 5.2. □

The fibers of D are rather complicated, given a tree $t \in \mathscr{V}_n$ there is an e-labeling of t that makes it being in the image of $\{\zeta\} \times \mathscr{E}_n$ if and only if t is noncrossing with respect to ζ. If the degree sequence of t is $d_1, \ldots, d_n$ then there are $(d_1)! \ldots (d_n)!$ such cycles ζ, see [22] and [20]. It follows that one cannot extract a bijection $\mathscr{F}_{n-1}^{\zeta_0} \to \mathscr{V}_n$ from D, and Dénes in [17] posed the problem of finding such an explicit bijection.

Based on the observations above, Moszkowski, in [35], realized that in order to solve the problem one has to *delabel* the vertices of the trees, and provided the following solution:

Theorem 5.5. *There is a bijection $S : \mathscr{F}_{n-1}^{\zeta_0} \to \mathscr{V}_n$.*

Proof. S is the composition of f_{ζ_0} with a bijection $\mathscr{E}_n \to \mathscr{V}_n$ defined by labeling the root of the tree 1, increasing all e-labels by 1 and then sliding each e-label to the vertex of that edge that is further away from the root. Starting from a v-tree, we can recover the rooted e-tree, by declaring the vertex 1 to be the root, and decreasing the vertex labels by 1 and then sliding each v-label to its incident edge that is closest to the root; therefore S is a bijection. □

We remark that the description above comes from [27], and is also contained in [38].

Goulden and Yong in [27] constructed a new bijection $\phi : \mathscr{F}_n^{\zeta_0} \to \mathscr{V}_n$, which with our notation is defined by the following diagram:

$$\begin{array}{ccc} \mathscr{F}_n^{\zeta_0} & \xrightarrow{f_{\zeta_0}} & \mathscr{E}_n^* \\ \phi \,\vdots\, \downarrow & & \downarrow * \\ \mathscr{V}_n & \xleftarrow{S} & \mathscr{E}_n^* \end{array}$$

where $* : \mathscr{C}_n^* \to \mathscr{C}_n^*$ is the mind-body dual. This bijection enjoys two "structural" properties which we now explain.

For a transposition τ define it's *difference index* to be the cyclic distance of its moved points, in other words, if $\tau = (s, t)$, $s < t$ then $\delta(\tau) = \min\{t - s, n - t + s\}$, and for a factorization $\rho = \tau_1, \ldots, \tau_n$ define its *difference distribution* to be $(d_1, \ldots, d_n)$, where d_i is the number of elements in ρ with difference index i.

For an edge e of a tree $t \in \mathscr{V}_n$ define its *edge-deletion index* $\varepsilon(e)$ to be the minimum of the orders of the two trees that result from t after we delete e, and the *edge-deletion distribution* of t to be $(a_1, \ldots, a_n)$ where a_i is the number of edges of t with edge-deletion index i.

For a factorization $\rho \in \mathscr{F}_{n-1}^{\zeta_0}$ define its *degree distribution*, $d(\rho)$to be the degree distribution of the associated e-v-tree. For a vertex i of a v-tree $t \in \mathscr{V}_n$ define its *maximal minimally increasing path* to be the path obtained by starting at i follow the edge that leads to the smallest of its neighbors and then keep going to the smaller neighbor that is larger than the vertex we are in, for as long as such neighbors exist. The *path-length distribution* of t is the sequence $l(t) = (l_1, \ldots, l_n)$ where l_i is the number of vertices of t that have maximal minimal increasing path of length i.

Theorem 5.6. *The bijection $\phi : \mathscr{F}_n^{\zeta_0} \to \mathscr{V}_n$ satisfies*

(1) $\delta(\rho) = \varepsilon(\phi(\rho))$,

(2) $d(\rho) = l(\phi(\rho))$.

Proof. Both properties follow from the properties of mind-body duality. The first one from the fact that if τ has difference index k then e^*, the dual of the corresponding edge e of the associated tree, will have deletion index k. The second one follows from the fact that the maximal minimally increasing paths are the image of migts under S. □

Item (1) was observed in [27], item (2) is not explicitly stated there, although it is implicit in the discussion.

5B. Mind-body duality for rooted e-graphs and flagged PCDs. It turns out that for e-v-trees, mind-body duality at the level of factorizations can be described via the functions f_ζ (see Theorem 5.2) and mind-body duality at the level of rooted e-trees. Mind-body duality can be extended to *rooted* e-graphs in an obvious way: if Γ is a rooted e-graph with root v_0 then we can take the root of Γ^* to be v_0^*, i.e., the trail that starts at v_0. Alternatively, the Hurwitz action of the (loop) braid group extends to rooted e-graphs, the root just stays the same, and we can use Theorem 3.17 to define the dual of a rooted e-graph. It's straightforward to check that these two approaches define the same notion. It's also easy to see that the following diagram commutes, where $*$ stands for mind-body duality of the relevant sets:

$$\begin{array}{ccc} \mathscr{F}_{n-1}^{\zeta_0} & \xrightarrow{f_{\zeta_0}} & \mathscr{C}_n^* \\ {\scriptstyle *}\downarrow & & \downarrow{\scriptstyle *} \\ \mathscr{F}_{n-1}^{\zeta_0^{-1}} & \xleftarrow{f_{\zeta_0^{-1}}^{-1}} & \mathscr{C}_n^* \end{array}$$

Remark 5.7. We note that one could use $f_{\zeta_0}^{-1}$ for the bottom arrow to define duality between factorizations of the standard cycle, i.e., one could define a duality $\mathscr{F}_{n-1}^{\zeta_0} \to \mathscr{F}_{n-1}^{\zeta_0}$, to be the conjugate $f_{\zeta_0}^{-1} \circ * \circ f_{\zeta_0}$. This

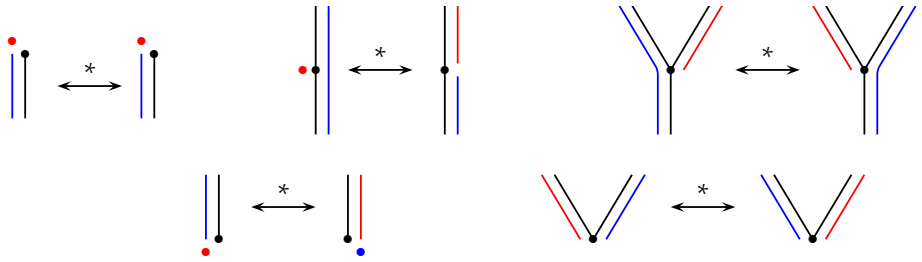

Figure 28. The flag of the dual of a flagged PCD.

observation will be used in [6] to define a "true" duality for noncrossing trees, and study it's properties. See also Remark 5.15 below.

Clearly choosing a root for an e-graph, or more generally a leo, is equivalent to choosing one of the chains of the PCD of its medial digraph, and we can translate mind-body duality of rooted e-graphs (or leos) in terms of PCDs with a distinguished chain. We formalize this in the following definition.

Definition 5.8. A *flagged PCD* on a binary digraph M is a PCD $\mathcal{C}$ on M endowed with a distinguished chain $f \in \mathcal{C}$ called its *flag*.

For a chain c in a PCD we use the notation $\alpha(f)$ (respectively $\omega(f)$) to denote the first (respectively last) vertex of c.

The *mind-body dual* of a flagged PCD $(\mathcal{C}, f)$ is the flagged PCD $(\mathcal{C}^*, f^*)$ where f^* is defined as follows: $\alpha(f^*) = \alpha(f)$ and if f is the only chain that starts at $\alpha(f)$ then f^* is the only chain of $\mathcal{C}^*$ that starts at $\alpha(f)$, otherwise the first edge of f^* is the outgoing edge incident at $\alpha(f)$ that does not belong to f, if no such edge exist then f^* is a trivial chain, see Figure 28 where the flags of the relevant PCDs are shown in red.

By Lemma 2.15 the underlying graph of the medial digraph of an e-tree is a tree.

Definition 5.9. A *ditree* is a digraph whose underlying graph is a tree.

We finish this subsection by mentioning that the enumeration of factorizations, or equivalently simple branched coverings, is a well established area of research with connections to geometric topology, algebraic geometry, and mathematical physics, and there are many results with bijective proofs for various "Hurwitz numbers". See, for example, [19]. The author hopes that the notion of mind-body duality introduced in this paper will help provide explicit bijections explaining known enumerative coincidences, as well discovering new ones.

5C. Self-dual e-trees. In every context where an interesting concept of duality is defined, a natural question that arises is whether there are any self-dual objects. The question for general graphs will be studied in further projects, in this subsection we concentrate on trees. In the context of mind-body duality it is obvious that there are no self-dual factorizations or e-v-trees[12] since the monodromy of the dual is the inverse of the monodromy of the original object.[13] For e-trees the question is meaningful and has an interesting answer.

[12] with the trivial exceptions of $n = 1, 2$ where the unique objects are obviously self-dual

[13] See, however, Remarks 5.7 and 5.15 at the end of this section.

Definition 5.10. An e-tree t is called *self-dual* if $t^* = t$.

Definition 5.11. For an integer n, the *updown ditree with n vertices* is the ditree with vertices $x_1, \ldots, x_n$ and an edge from x_{2i-1} to x_{2i}, and an edge from x_{2i} to x_{2i+1} for each $i = 1, \ldots, \lfloor \frac{n}{2} \rfloor$.

The *downup ditree with n vertices* is the ditree with vertices $x_1, \ldots, x_n$ and an edge from x_{2i} to x_{2i-1}, and an edge from x_{2i+1} to x_{2i} for each $i = 1, \ldots, \lfloor \frac{n}{2} \rfloor$.

A *zigzag ditree* is an updown or downup ditree.

For even n the updown and downup ditrees are isomorphic, while for odd n there are two (inverse to each other) zigzag ditrees. The zigzag ditrees are the Hasse diagrams of the zigzag (or *fence*) posets. See [41], page 157, Exercise 23 in Chapter 3.

With this definitions in place, we can now prove:

Theorem 5.12. *An e-tree t is self-dual if and only if its medial ditree $\mathcal{M}(t)$ is a zigzag ditree.*

Proof. The medial ditree of an e-tree is a topsorted binary ditree with a PCD. It is clear that the set of zigzag ditrees coincides with the set of binary ditrees with no internal vertices. So we need to prove that an e-tree is self-dual if and only if its medial ditree has no internal vertices.

By Definition 2.18, it follows that if $\mathcal{M}(t)$ has no internal vertices then $\mathcal{M}(t)$ is self-dual and hence, by Theorem 2.20, t is self dual.

Conversely, since a PCD and its dual, differ at every internal vertex, and $\mathcal{M}(t)$, being topsorted, has labeled vertices, it follows that if $\mathcal{M}(t)$ has internal vertices then t is not self-dual. □

It is well known that the number of topological sorts of a zigzag ditree with n vertices is given by the n-th Euler up/down number. This is sequence A000111 in [40]. This sequence enumerates (among other things) the set of *alternating permutations*, see [1].

The fact that A000111 enumerates the set of topological sorts of a given zigzag ditree is not enough to conclude that it also enumerates self-dual e-trees, because of the presence of automorphisms. Indeed, while for even n the zigzag ditree has no nontrivial automorphisms, for odd n there is a nontrivial automorphism of order 2. However, for odd n there are two, inverse to each other, zigzag ditrees and that introduces a factor of 2 that compensates. So we have:

Corollary 5.13. *The number of self-dual e-trees with n vertices is equal to the $(n-1)$-th Euler up/down number.*

The self-dual e-trees for $n = 3, 4, 5$ are shown in Figure 29, on the left side we have the zigzag ditree(s), on the center all possible topsorted zigzag ditrees, and on the right the corresponding e-graphs.

One can also ask if there are any self-dual rooted e-trees and if so, how many. The answer turns out to be again the Euler up/down numbers. To see this notice that it follows from the discussion in Section 5B (see Definition 5.8) that in order for a flagged PCD to be self-dual it is necessary that it's flag is a trivial chain whose only vertex is a maximal leaf. Therefore an rooted e-tree is self-dual if and only if it's medial ditree is a zigzag ditree, and the flag of its PCD is a maximal leaf. For even m each of the updown and downup ditrees with m vertices has exactly one maximal leaf, while for odd n the updown ditree has no maximal leaf, while the downup ditree has exactly two. So we also have the following corollary:

Corollary 5.14. *The number of rooted self-dual e-trees with n vertices is also equal to the $(n-1)$-th Euler up/down number.*

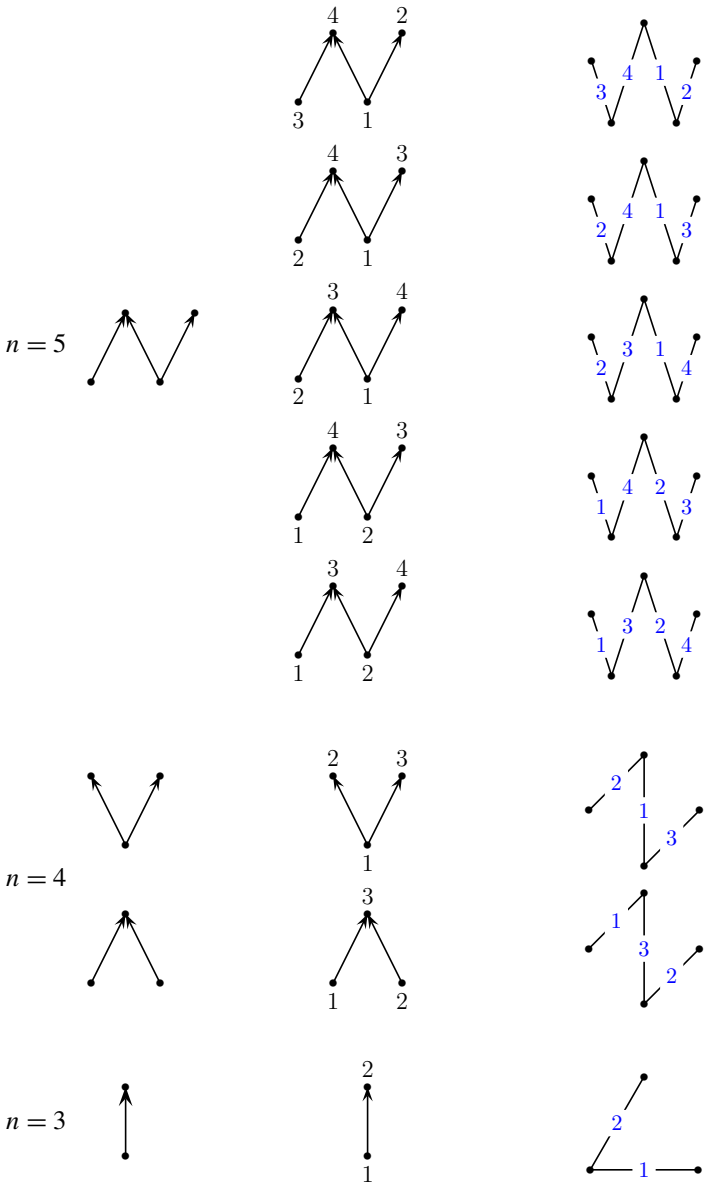

Figure 29. Self dual e-trees.

Remark 5.15. This means that if we use the alternative duality of Remark 5.7 there are self-dual e-v-trees, since that duality is simply a conjugate of the duality of rooted e-trees.

6. Future directions

As we mentioned in the introduction, we expect that the main application of this work will be in finding new, as well as explaining already known results about Hurwitz numbers. Pegs are more attuned to the graph theoretical properties of the graph than cegs.

We conclude by listing a few further works that will use the theory developed in this paper.

6A. Almost minimal factorizations of cycles. The next class of e-graphs after trees is the class of e-unicycles, i.e., e-graphs with a unique cycle. Using item (2) of Lemma 4.5, and the classification of surfaces, we see that the surface of the peg of such an e-unicycle is an annulus. It follows that the monodromy of an e-unicycle is a product of two disjoint cycles, this was also observed in [8] using a branched covering argument. The mind-body dual of an e-unicycle is also an e-unicycle and so one obtains interesting "structural" bijections between different classes of e-unicycles.

A particularly simple case is the case of e-unicycles whose monodromy has a fixed point. The set of these unicycles, with $n+1$ vertices has cardinality n^n and a question posed in [27], is (or rather can be interpreted to be) whether mind-body duality can be used to "explain" this simple counting. It turns out that the migt of the fixed vertex of such a unicycle is (the closed trail corresponding to) the unique cycle, and so one of the "structural" properties of mind-body duality is that it takes the neighborhood of the fixed point to the unique cycle, and that fact can be used to provide bijections between various subsets of the set of those e-unicycles.

The factorization that corresponds to an e-unicycle with a fixed vertex expresses an n-cycle as the product of $n+1$ transpositions, while the minimum number of transpositions needed is $n-1$, so we call such factorizations *almost minimal*. Those and related topics will be studied in [4].

6B. Duality for noncrossing trees. As we've mentioned a notion of duality for noncrossing trees has been defined in [29], however in the context that it was defined (vertex-labeled noncrossing trees), that duality is not involutory: the dual of the dual is not the original, it becomes involutory, and coincides with the peg duality as we defined it in Definition 4.7, only if we descend to the level of unlabeled noncrossing trees. That "duality" is closely related with the action of the Garside element (see Section 3A) and it's periodic with period a multiple of n.

In [6] we use the idea of Remark 5.7 to define a "true" duality for (labeled) noncrossing trees, ask the question "how many self-dual noncrossing trees are there?" and get an interesting answer.

6C. Duality for increasing trees. The class of increasing trees is well studied in the literature, see, for example, [10], these are rooted v-trees in which the children of every vertex have labels greater than the vertex. It follows that the root is labeled 1, and we can apply the inverse of the sliding operation $\mathscr{E}_n^* \to \mathscr{V}_n$ defined in the proof of Theorem 5.5 to convert the class of increasing trees to a class of rooted e-trees that turns out to be closed under the mind-body duality. So one can define a duality in the set of increasing trees and study its properties. This will be done in [5] where interesting bijection are obtained for several classes of increasing trees.

We mention that the set of (topsorted) medial ditrees of increasing trees consists of those binary ditrees that have exactly one minimum, and that set is obviously in bijection with binary increasing trees, which in turn are in bijection (see, e.g., [18]) with the set of alternating permutations. So the Euler up/down numbers appear again!

6D. General theory of properly embedded graphs. The focus of this paper is on e-graphs and factorizations, and we developed enough of the theory of pegs to be able to treat this case. However, there is a more general theory of pegs, that treats the case of pegs whose medial digraph is not a dag, as well as the case of graphs properly embedded in nonorientable surfaces. One can even consider *semipegs*,

where some of the vertices lie in the boundary, and some in the interior of the surface. While most of the ingredients for such a theory are already contained, or have been hinted on, in this work, there are a few new ingredients needed for such an extension. We plan to pursue this in a future work [3].

References

[1] D. André, "Sur les permutations alternées", *J. Mathé. Pures Appliquées* **7** (1881), 167–184.

[2] N. Apostolakis, "On 4-fold covering moves", *Algebr. Geom. Topol.* **3** (2003), 117–145.

[3] N. Apostolakis, "On graphs properly embedded on bounded surfaces", preprint, 2018. In preparation.

[4] N. Apostolakis, "Almost minimal factorizations of permutations and unicyclic graphs", preprint, 2019. In preparation.

[5] N. Apostolakis, "A duality for increasing trees", preprint, 2019. In preparation.

[6] N. Apostolakis, "Non-crossing trees, quadrangular dissections, ternary trees, and duality-preserving bijections", *Ann. Comb.* **25**:2 (2021), 345–392.

[7] D. Archdeacon, "The medial graph and voltage-current duality", *Discrete Math.* **104**:2 (1992), 111–141.

[8] V. I. Arnold, "Topological classification of complex trigonometric polynomials and the combinatorics of graphs with an identical number of vertices and edges", *Funktsional. Anal. i Prilozhen.* **30**:1 (1996), 1–17, 96.

[9] J. Bang-Jensen and G. Gutin, *Digraphs*: *theory*, *algorithms and applications*, Springer, 2001.

[10] F. Bergeron, P. Flajolet, and B. Salvy, "Varieties of increasing trees", pp. 24–48 in *CAAP '92* (Rennes, 1992), Lecture Notes in Comput. Sci. **581**, Springer, Berlin, 1992.

[11] I. Berstein and A. L. Edmonds, "On the construction of branched coverings of low-dimensional manifolds", *Trans. Amer. Math. Soc.* **247** (1979), 87–124.

[12] J. S. Birman, *Braids*, *links*, *and mapping class groups*, Annals of Mathematics Studies **82**, Princeton University Press, Princeton, NJ, 1974.

[13] J. A. Bondy, "Perfect path double covers of graphs", *J. Graph Theory* **14**:2 (1990), 259–272.

[14] C. P. Bonnington, M. Conder, M. Morton, and P. McKenna, "Embedding digraphs on orientable surfaces", *J. Combin. Theory Ser. B* **85**:1 (2002), 1–20.

[15] F. Catanese and B. Wajnryb, "The fundamental group of generic polynomials", *Topology* **30**:4 (1991), 641–651.

[16] C. Damiani, "A journey through loop braid groups", *Expo. Math.* **35**:3 (2017), 252–285.

[17] J. Dénes, "The representation of a permutation as the product of a minimal number of transpositions, and its connection with the theory of graphs", *Magyar Tud. Akad. Mat. Kutató Int. Közl.* **4** (1959), 63–71.

[18] R. Donaghey, "Alternating permutations and binary increasing trees", *J. Combinatorial Theory Ser. A* **18** (1975), 141–148.

[19] E. Duchi, D. Poulalhon, and G. Schaeffer, "Bijections for simple and double Hurwitz numbers", preprint, 2014. arXiv

[20] S. Dulucq and J.-G. Penaud, "Cordes, arbres et permutations", *Discrete Math.* **117**:1-3 (1993), 89–105.

[21] J. L. Dyer and E. K. Grossman, "The automorphism groups of the braid groups", *Amer. J. Math.* **103**:6 (1981), 1151–1169.

[22] M. Eden and M. P. Schützenberger, "Remark on a theorem of Dénes", *Magyar Tud. Akad. Mat. Kutató Int. Közl.* **7** (1962), 353–355.

[23] R. Evans and L. Huang, "Mind switches in *Futurama* and *Stargate*", *Math. Mag.* **87**:4 (2014), 252–262.

[24] B. Farb and D. Margalit, *A primer on mapping class groups*, Princeton Mathematical Series **49**, Princeton University Press, Princeton, NJ, 2012.

[25] R. Fenn, R. Rimányi, and C. Rourke, "The braid-permutation group", *Topology* **36**:1 (1997), 123–135.

[26] The GAP Group, *GAP (version 4.7.8)*, 2015, available at https://www.gap-system.org/.

[27] I. Goulden and A. Yong, "Tree-like properties of cycle factorizations", *J. Combin. Theory Ser. A* **98**:1 (2002), 106–117.

[28] J. L. Gross and T. W. Tucker, *Topological graph theory*, Dover Publications, Mineola, NY, 2001.

[29] M. C. Herando, *Complejidad de estructuras geométricas y combinatorias*, Ph.D. thesis, Universitat Politècnica de Catalunya, 1999, available at http://hdl.handle.net/10803/6720.

[30] A. Hurwitz, "Ueber die Anzahl der Riemann'schen Flächen mit gegebenen Verzweigungspunkten", *Math. Ann.* **55**:1 (1901), 53–66.

[31] C. Kassel and V. Turaev, *Braid groups*, Graduate Texts in Mathematics **247**, Springer, New York, 2008.

[32] S. K. Lando and A. K. Zvonkin, *Graphs on surfaces and their applications*, Encyclopaedia of Mathematical Sciences **141**, Springer, Berlin, 2004.

[33] M. Markl, S. Shnider, and J. Stasheff, *Operads in algebra, topology and physics*, Mathematical Surveys and Monographs **96**, American Mathematical Society, Providence, RI, 2002.

[34] J. W. Moon, *Counting labelled trees*, Canadian Mathematical Monographs **1**, Canadian Mathematical Congress, Montreal, QC, 1970.

[35] P. Moszkowski, "A solution to a problem of Dénes: a bijection between trees and factorizations of cyclic permutations", *European J. Combin.* **10**:1 (1989), 13–16.

[36] M. Noy, "Enumeration of noncrossing trees on a circle", *Discrete Math.* **180**:1-3 (1998), 301–313.

[37] D. Pengelley, "Self-dual orientable embeddings of K_n", *J. Combinatorial Theory Ser. B* **18** (1975), 46–52.

[38] D. Poulalhon, *Graphes et décompositions de permutations*, Mémoire de DEA **59**, École Polytechnique, 1997.

[39] The Sage Developers, *SageMath (version 8.0)*, 2017, available at https://www.sagemath.org.

[40] N. J. A. Sloane, "Euler or up/down numbers", pp. A000111 in *The on-line encyclopedia of integer sequences*, 1991.

[41] R. P. Stanley, *Enumerative combinatorics*, vol. 1, Cambridge Studies in Advanced Mathematics **49**, Cambridge University Press, Cambridge, 1997.

NIKOS APOSTOLAKIS: nikolaos.apostolakis@bcc.cuny.edu
Department of Mathematics and Computer Science, Bronx Community College of The City University of New York, Bronx, NY, United States