UC Berkeley
UC Berkeley Electronic Theses and Dissertations

Title
A Combinatorial Model of Lagrangian Skeleta

Permalink
https://escholarship.org/uc/item/882417wf

Author
Zorn, Alex

Publication Date
2018

Peer reviewed|Thesis/dissertation

eScholarship.org Powered by the California Digital Library
University of California

A Combinatorial Model of Lagrangian Skeleta

by

Alex Zorn

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Mathematics

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor David Nadler, Chair
Assistant Professor Vivek Shende
Professor Marjorie Shapiro

Spring 2018

A Combinatorial Model of Lagrangian Skeleta

Copyright 2018
by
Alex Zorn

Abstract

A Combinatorial Model of Lagrangian Skeleta

by

Alex Zorn

Doctor of Philosophy in Mathematics

University of California, Berkeley

Professor David Nadler, Chair

We investigate a collection of posets- combinatorial arboreal singularities- which are the strata posets of the arboreal singularities constructed by David Nadler. Nadler demonstrated that any Lagrangian skeleton admits a non-characteristic deformation into a skeleton with only arboreal singularities, suggesting that arboreal singularities form the basis for a combinatorial theory of Lagrangian skeleta.

In this document, we introduce a form of combinatorial data called a 'cyclic structure', which is essentially codimension-one data with compatibility conditions in codimensions two and three. We develop a comprehensive theory of isomorphisms of combinatorial arboreal singularities and cyclic arboreal singularities (singularities equipped with a cyclic structure). We show that a cyclic structure determines (up to quasi-equivalence) a sheaf of dg-categories on a combinatorial arboreal singularity, and investigate combinatorial properties of this sheaf. We describe a class of 'arboreal moves', which are local mutations of combinatorial arboreal spaces preserving this sheaf of categories. Finally, we discuss how this combinatorial picture is related to the geometric understanding of Lagrangian embeddings of arboreal singularities.

Throughout my life, my grandfather, Jens Zorn, has served as a constant source of inspiration and support, and has entertained and furthered my scientific curiosity. This dissertation is dedicated to him.

Contents

Contents **ii**

List of Figures **iv**

List of Tables **vii**

1 Introduction **1**
1.1 Main Results 2
1.2 Summary of the Paper 5

2 Combinatorial Arboreal Singularities and Arboreal Spaces **18**
2.1 Posets 18
2.2 Trees 21
2.3 Basic Properties 23
2.4 Some Examples: A_n and S_n singularities 25

3 Cyclic Structures **29**
3.1 Pre-Cyclic Structures 29
3.2 Directed Trees 30
3.3 Coherence 34
3.4 Cyclic Structures on $\mathcal{P}_{A_3}$ 34
3.5 Cyclic Structures on $\mathcal{P}_{S_4}$ 38
3.6 Cyclic Structures on A_n Singularities 42
3.7 Cyclic Structures on S_n Singularities 46

4 The Combinatorics of Cyclic Structures **50**
4.1 Isomorphisms of C-Arb Singularities 50
4.2 Reflection Isomorphisms 54
4.3 Isomorphisms of Cyclic C-Arb Singularities 59
4.4 Reflection Isomorphisms for Directed Trees 59
4.5 The Coherence Theorem 64

5 A Sheaf of Categories **68**

5.1 Dg-Categories . . . 68
5.2 Sheaves of dg-Categories . . . 71
5.3 Representations of Directed Trees . . . 73
5.4 The Trivalent Vertex . . . 75
5.5 Reflection Functors . . . 77
5.6 The Cyclic Structure Theorem . . . 80
5.7 Construction of the Reflection Functors . . . 81

6 Graded Structures and Graded Branes **99**
6.1 Graded Structures . . . 99
6.2 Branes . . . 111
6.3 The Representation Theory of Graded Branes . . . 118
6.4 Branes on A_n Singularities . . . 122
6.5 In Search of a 'Fukaya Category' . . . 125

7 Arboreal Moves **126**
7.1 The H-to-I Move . . . 126
7.2 Definition of Arboreal Moves . . . 129
7.3 The Topology of Branes . . . 130
7.4 Classification of Good Subsets . . . 132
7.5 Examples for A_n and S_n Trees . . . 144

8 The Geometry of Arboreal Spaces **148**
8.1 Stratified Spaces . . . 148
8.2 Arboreal Singularities and Arboreal Spaces . . . 150
8.3 Embeddings . . . 152
8.4 The Geometry of Branes . . . 164
8.5 The Geometry of Moves . . . 165

A Cyclic Orders on Sets **175**

B Infinity-Categories and Sheaves **177**
B.1 Quasi-Categories . . . 177
B.2 Sheaves . . . 180
B.3 Morphisms Between Sheaves . . . 182
B.4 Pullback Functors and Homotopy . . . 185

Bibliography **188**

List of Figures

1.1 The Geometry of the Neighborhood of a Codimension-One Cell 2
1.2 The A_3 (left) and S_4 (right) trees. 3
1.3 The H-to-I move 5
1.4 The A_n Tree 7
1.5 Illustrations of $\mathcal{P}_{A_2}$ (left) and $\mathcal{P}_{A_3}$ (right) 7
1.6 A diagram of the sheaf $\mathcal{Q}_\mu$ on $\mathcal{P}_{A_2}$ 10
1.7 The sheaf $\mathcal{Q}_\mu$ after applying a quasi-equivalence of global sections 10
1.8 Pre-Branes on $\mathcal{P}_{A_2}$ 12
1.9 The H-to-I move 13
1.10 Total Space of the H-to-I move 13
1.11 A Generalized Cyclic Interval $I(v, w, D)$. 14
1.12 The unique (up to equivalence) partitions of cyclic sets yielding moves of type A_3 and A_4 15

2.1 A Correspondence 22
2.2 A Correspondence with Edges/Terminal Vertices Labelled 24
2.3 The A_n Tree 25
2.4 Hasse diagram for $\mathcal{P}_{A_2}$ 26
2.5 Hasse diagram for $\mathcal{P}_{A_3}$ 26
2.6 Illustrations of $\mathcal{P}_{A_2}$ (left) and $\mathcal{P}_{A_3}$ (right) 27
2.7 Another Visualization of $\mathcal{P}_{A_3}$ 27
2.8 The S_n Tree 28

3.1 The Geometry of the Neighborhood of a Codimension-One Cell 29
3.2 A Directed Correspondence 31
3.3 A Codimension-One Cell, and the three Top Cells in the Induced Cyclic Order . 32
3.4 An Illustration of $\mathcal{P}_{A_3}\backslash\{\{e_0, e_1\}, \{e_2, e_3\}\}$. 36
3.5 A Completed Arrow Diagram 36
3.6 Translating Between Two Descriptions of Codimension-One Cells in $\mathcal{P}_{S_4}$. 39
3.7 An Illustration of $\mathcal{P}_{S_4}\backslash\{v_i, e_i\}^c$. 40
3.8 Constructing Pinwheel Orientations 41
3.9 Constructing Pinwheel Orientations 42

3.10 A Directed Tree (A_5, μ). 46
3.11 The Cyclic Order $\Delta(\mathcal{O}_\mu)$ on K_{A_5}. 46

4.1 A tree T and the chains in K_T 51
4.2 An illustration of the relation $\sim_C$: The subtree $\{v\}$ is a witness to the relation $C_1 \sim_C^0 C_3$, while $\{v, w\}$ is a witness to the relation $C_2 \sim_C^0 C_3$ 53
4.3 Cases B1 and B2 56
4.4 Case B3 56
4.5 (A_n, μ) 61
4.6 The Construction of $\mathfrak{q}$ and $\mathfrak{p}_1$ 66
4.7 The construction in Case 2b 67

5.1 The directed tree (A_2, μ) 75
5.2 Geometric Picture of the Indecomposable Representations 76

6.1 Illustration of Mondromy Axiom (M2) 100
6.2 Illustration of Monodromy Axiom (M3) 100
6.3 Arrow Diagram Setup With Labelled Cells 101
6.4 Arrow Diagrams Giving $m(\varphi_{t,x}^{\{e_0\}}\varphi_{z,t}^{\{e_2\}}\varphi_{y,z}^{\{e_1\}}\varphi_{x,y}^{\{e_3\}}) = -2$ (left) or 0 (right) 102
6.5 A Pinwheel Orientation 103
6.6 A Neighborhood of $\{v_1\}$ in the Pinwheel Orientation 103
6.7 Cells Near $\{v_3\}$ In $\mathcal{P}_{S_4}$ 104
6.8 Hasse diagram for $\mathcal{P}_{A_2}$ 104
6.9 Hasse diagram for $\mathcal{P}_{A_3}$ 105
6.10 Illustration of the Graded Structure G_μ 106
6.11 Pre-Branes on $\mathcal{P}_{A_2}$ 111
6.12 An Illustration of $L = \mathcal{P}_{A_3} \backslash \{\{e_0, e_1\}, \{e_2, e_3\}\}$ with two cyclic structures- On the left, L is *not* a Brane. 114

7.1 The H-to-I move 126
7.2 Total Space of the H-to-I move 127
7.3 The H-to-I move 127
7.4 An Invalid Move 128
7.5 A Second Invalid 'H-to-I-esque' Move 128
7.6 A Generalized Cyclic Interval $I(v, w, D)$. 133
7.7 Illustrating The Edge Expansion $\exp_e$ 135
7.8 Illustrating the Edge Cut cut_e. 136
7.9 Possible Partitions of K_{A_3}- Only the Fourth Gives a Move 145
7.10 The Partition Givng The Unique A_4 Move Up To Equivalence/Reversal 145
7.11 A Topological Picture of The A_4 Move 146
7.12 The Four Distinct Partitions of K_{S_4} Giving Moves 147

8.1 Embeddings Which Are Not Germ Equivalent 155

8.2 The embeddings i_+ (left) and i_- (right) of L_{A_2}. 157
8.3 A Lagrangian Embedding of L_{A_3} 158
8.4 Another Lagrangian Embedding of L_{A_3} 159
8.5 A Cross-Section of the Embedding at $\xi_y = t_0$ 159
8.6 The planes $L_T(v)$ and $L_T(w)$ 161
8.7 The planes $\mathcal{H}_{w_1}$,$\mathcal{H}_{w_2}$ in $L_v(T)$ 161
8.8 A Disallowed Embedding of $L_{A_2} \times \mathbb{R}$. 162
8.9 An Embedding of L_{A_3} Realizing a Non-Cyclic Structure 163
8.10 The non-brane in $(\mathcal{P}_{A_3}, \mathcal{O})$ relative to two Lagrangian embeddings 165
8.11 A Movie Realizing the 'H-to-I' Move 167
8.12 Realizing the $\mathcal{P}_{A_4}$ move with a movie 167
8.13 Labelled Cells In the $\mathcal{P}_{A_4}$ Movie 168
8.14 The Partition Corresponding to the $\mathcal{P}_{A_3}$ Move 169
8.15 One Movie Realizing a $\mathcal{P}_{S_4}$ Move 169
8.16 A Second $\mathcal{P}_{S_4}$ Movie 170
8.17 A movie realizing a new move? 170
8.18 Horizontal Cross-Sections of the $L_{A_3} \times \mathbb{R}$ Movie 171
8.19 Two New Moves? 172
8.20 Diagramming the Type I Noncompact Move of Type A_4 174
8.21 Diagramming the Type I Noncompact Move of type S_4 174

B.1 A Commutative Diagram Giving $\Gamma(Z; i^*\mathcal{Q}) \to \Gamma(\mathcal{P}; \mathcal{Q})$ 186

List of Tables

3.1 A Dictionary for Cyclic Structures on $\mathcal{P}_{A_3}$ 37

5.1 The Restrictions of Indecomposable Elements of Rep(A_2, μ). 75

6.1 Non-Branes In $(\mathcal{P}_{A_3}, \mathcal{O}_\mu)$ 114
6.2 Branes in $\mathcal{P}_{A_3}$, Four Ways 123

Acknowledgments

First and foremost, this dissertation would not have been possible without the incredible amount of aid and encouragement from my advisor, David Nadler. He has been with this project every step of the way- giving advice on directions to pursue, helping me understand technical details, and clarifying key concepts. Perhaps most importantly he has always believed in the value of my work, and his faith in my ideas and abilities has kept me motivated during the most stressful and frustrating parts of the writing process. I really cannot overstate his strength as a mentor during my six years at Berkeley.

I also owe thanks to useful conversations with Ben Gammage, Vivek Shende, and Laura Starkston.

I owe a great deal to Zvedelina Stankova, who I have worked with as a GSI on multiple occasions, and who gave me the opportunity to lead sessions with the Berkeley Math Circle. My growth as a math educator while at Berkeley has been in no small part due to her.

Finally, I want to thank the people at the UC Berkeley math department who have helped make the department a great community in which to study math, including Judie Filomeo, Barb Waller, and Vicky Lee.

Chapter 1

Introduction

This dissertation is part of an ongoing program of finding combinatorial models for symplectic topology, and combinatorial computations of associated symplectic invariants. In [15], Kontsevich conjectures roughly that the Fukaya category of a symplectic manifold has a combinatorial description in terms of a sheaf of differential graded categories along a Lagrangian skeleton. David Nadler has made substantial progress toward a resolution of this conjecture with the introduction of arboreal singularities in [20] and [24]. In this work we push Nadler's program further by undertaking a detailed combinatorial analysis of arboreal singularities.

The category of microlocal sheaves associated to a Lagrangian skeleton, developed by Kashiwara-Schapira [13], is intimately connected to the Fukaya category by the work of Nadler and Zaslow in [25] and [22]. In [20] and [24], Nadler introduces a collection of singular stratified spaces termed 'arboreal singularities', described using combinatorial data (namely, trees), and shows that these spaces form a fundamental set of combinatorial building blocks for singular Lagrangian skeleta, in the sense that arbitrary Lagrangian singularities can be non-characteristically deformed into a skeleton with only arboreal singularities. In addition, [20] includes a computation of the category of microlocal sheaves on the skeleton.

In recent years, the combinatorial description provided by arboreal singularities has seen applications to mirror symmetry in [10], [19], [23], and the combinatorics of Legendrian knots due to Shende, Treumann, Zaslow, and others in [31], [30], [29], [26]. There is also work in the direction of understanding the structure of Weinstein manifolds [9] via arboreal Lagrangian skeleta by Starkston in [33].

The main goals of this work are: (i) to describe a local combinatorial invariant (termed a 'cyclic structure') which is meant to capture the local symplectic geometry of an arboreal skeleton, (ii) to describe a sheaf of dg-categories on an arboreal poset associated to such an invariant, and (iii) to describe a class of combinatorial 'moves' which are meant to capture local modifications to an arboreal skeleta which do not affect the local symplectic geometry.

This work can be thought of a higher-dimensional version of the theory of ribbon graphs, which label cells in the moduli space of Riemann surfaces [11] [18]. The sheaf of dg-categories should be thought of as an extension of Sibilla, Treumann, and Zaslow's construction on ribbon graphs in [32]. The combinatorial moves generalize the classic H-to-I mutation of ribbon graphs, also known as the 2-2 Pachner move on the dual graph to a triangulation of a surface [27].

The combinatorial moves discussed here also appear in connection with cluster structures associated to Legendrian knots, for example in [31].

1.1 Main Results

For a tree T, we let L_T denote the arboreal singularity associated to T, and $\mathcal{P}_T$ denote the strata poset. Any codimension-one strata in L_T is in the closure of exactly three top dimensional strata, as pictured in 1.1. It follows that any codimension-one element of $\mathcal{P}_T$ has exactly three maximal elements greater than it.

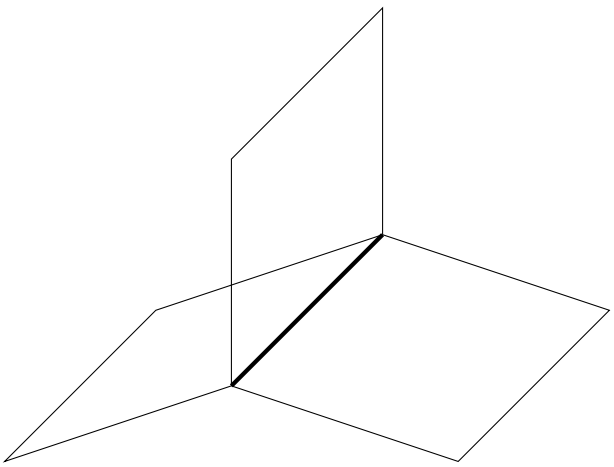


Figure 1.1: The Geometry of the Neighborhood of a Codimension-One Cell

More generally, we let a **C-Arb space** denote a poset $\mathcal{P}$ such that, for any $x \in \mathcal{P}$, the poset $N(x) = \{y \in \mathcal{P} \mid y \geq x\}$ is isomorphic to a C-Arb singularity. (We remark that C-Arb singularities are themselves C-Arb spaces). Then C-Arb spaces also have the property that every codimension-one element has exactly three elements greater than it.

Definition 1.1.1. *(Definition 3.1.1) A* ***pre-cyclic structure*** $\mathcal{O}$ *on a C-Arb space* $\mathcal{P}$ *is a choice, for every codimension-one element* x*, a cyclic ordering* $\mathcal{O}^x$ *on the three maximal elements greater than* x*.*

Among pre-cyclic structures, there is a collection of preferred structures which we term *cyclic structures*. These can be specified with a construction involving **directed trees**. A directed tree is a pair (T, μ), where μ is an orientation of the edges of T. To an orientation μ of T, one can define (Definition 3.2.3) a pre-cyclic structure $\mathcal{O}_\mu$ on $\mathcal{P}_T$. Then:

Definition 1.1.2. *(Definition 3.2.4) A pre-cyclic structure* $\mathcal{O}$ *on a C-Arb space* $\mathcal{P}$ *is a* ***cyclic structure*** *if, for every* $x \in \mathcal{P}$*, there exists a directed tree* (T, μ) *and an isomorphism* $\varphi : N(x) \xrightarrow{\sim} \mathcal{P}_T$ *such that* $\varphi^* \mathcal{O}_\mu = i^* \mathcal{O}$*, where* $i : N(x) \hookrightarrow \mathcal{P}$ *is the inclusion.*

A **cyclic C-Arb space** is a pair $(\mathcal{P}, \mathcal{O})$, where $\mathcal{O}$ is a cyclic structure on $\mathcal{P}$.

The first main result we prove shows that cyclic structures are distinguished among pre-cyclic structures by local coherence conditions. Namely:

Definition 1.1.3. *(Definition 3.3.1) Fix a tree* R*. A pre-cyclic structure* $\mathcal{O}$ *on a C-Arb space* $\mathcal{P}$ *is* $R-$***coherent*** *if for every embedding* $i : \mathcal{P}_R \hookrightarrow \mathcal{P}$*,* $i^* \mathcal{O}$ *is a cyclic structure.*

Theorem 1.1.1. ***The Coherence Theorem*** *(Theorem 3.3.1): A pre-cyclic structure is a cyclic structure iff it is* A_3-*coherent and* S_4-*coherent.*

A_3 is the tree with three vertices connected in a line, and S_4 is the tree with three vertices connected to a central vertex (Figure 1.2). One way to interpret the Coherence Theorem is that cyclic structure can be thought of as codimension-one data which satisfies codimension-two and codimension-three compatibility conditions.



Figure 1.2: The A_3 (left) and S_4 (right) trees.

Our next result concerns a sheaf of dg-categories $\mathcal{Q}_\mu$ on the poset $\mathcal{P}_T$ associated to a directed tree (T, μ). This is a sheaf meant to model a category of microlocal sheaves on the arboreal singularity L_T, as computed in [20]. In particular, the global sections $\Gamma(\mathcal{Q}_\mu; \mathcal{P}_T) \cong \mathrm{Rep}(T, \mu)$- the derived dg category of finite-dimensional representations of the directed tree (quiver) (T, μ).

We show that a cyclic structure is enough to determine the sheaf up to quasi-equivalence. Precisely:

Theorem 1.1.2. ***The Cyclic Structure Theorem*** *(Theorem 5.6.1): Let (T, μ) and (T', μ') be directed trees, and $\varphi : \mathcal{P}_T \xrightarrow{\sim} \mathcal{P}_{T'}$ be a poset isomorphism. Then $\varphi^* \mathcal{Q}_{\mu'}$ is quasi-equivalent to $\mathcal{Q}_\mu$ if and only if $\varphi^* \mathcal{O}_{\mu'} = \mathcal{O}_\mu$.*

Following this, we take the first steps toward constructing a canonical $A_\infty-$category associated to a cyclic arboreal space $(\mathcal{P}, \mathcal{O})$, which does not depend in isomorphisms with directed trees. As in [32], this category is meant to model the Fukaya category on a symplectic neighborhood of the space. In this document, we construct a set $\mathrm{Br}_{gr}(\mathcal{P}, G)$ of **graded branes** associated to an arboreal space $\mathcal{P}$ equipped with a **graded structure** G. Locally, graded structures up to isomorphism are equivalent to cyclic structures. We also specify a subcollection of objects $\mathrm{Rep}(T, \mu)$ termed **rank-one objects**, and prove the following:

Theorem 1.1.3. *(Propositions 6.3.2,6.3.3) Given a directed tree (T, μ) there is a canonical bijection:*

$$\left\{ \begin{array}{l} \textit{Rank one objects} \\ \textit{in Rep}(T, \mu) \end{array} \right\} \Big/ \textit{homotopy} \longleftrightarrow \mathrm{Br}_{gr}(\mathcal{P}_T, G_\mu)$$

Where G_μ is the graded structure on $\mathcal{P}_T$ associated to μ.

$\mathrm{Rep}(T, \mu)$ is the triangulated hull of the full subcategory of rank-one objects. It is therefore our hope that one could define a sheaf of $A_\infty-$categories whose objects are graded branes, such that locally one gets a category whose triangulated hull is quasi-equivalent to a category of representations of directed trees. An advantage of this perspective is that this category would manifestly respect all symmetries, for example those observed by Nadler in [21].

Our last main result concerns a generalization of the 'H-to-I move' (Figure 1.3) to higher-dimensional arboreal spaces. Namely, we construct a certain classification of combinatorial local 'mutations' of cyclic C-Arb spaces such that the total space of the mutation looks like an arboreal singularity $\mathcal{P}_T$. The 'H-to-I' move is such a move where $T = A_3$.

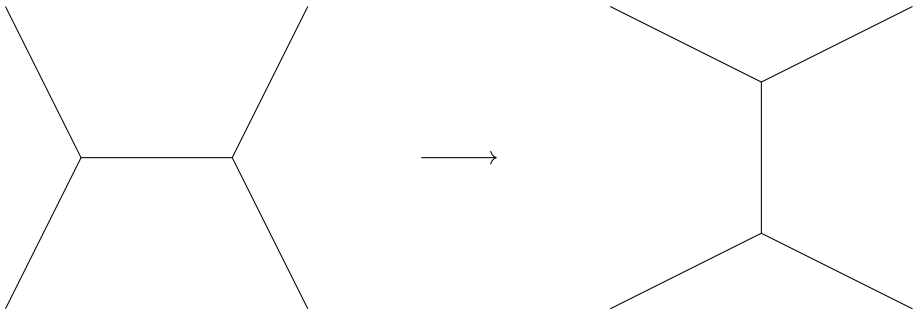


Figure 1.3: The H-to-I move

More concretely, for a tree T, let $K_T = E(T) \sqcup V_t(T)$, where $E(T)$ and $V_t(T)$ denote the set of edges of T and the set ofterminal vertices of T, respectively. We show that moves are classified by a directed tree (T, μ) along with a partition $K_T = K_1 \sqcup K_2$ of K_T into subsets which are *good* with respect to μ. We give an explicit combinatorial classification of good subsets:

Theorem 1.1.4. *(Theorem 7.4.1) A subset K of K_T is good if and only if it is not a* ***generalized cyclic interval*** *(definition 7.4.1).*

As a consequence, we are able to describe a move with total space $\mathcal{P}_{A_4}$, as well as four inequivalent moves with total space $\mathcal{P}_{S_4}$.

1.2 Summary of the Paper

We provide a brief summary of each chapter, outlining the main ideas and methods.

Summary of Chapter 2: Combinatorial Arboreal Singularities and Arboreal Spaces

In this chapter we introduce an abstract construction of posets arising from certain simply-behaved categories. We then apply this construction to the category of trees and tree correspondences in order to obtain our main objects of interest: Combinatorial Arboreal Singularities.

Briefly, a tree is a nonempty finite acyclic graph. A **correspondence of trees** (definition 2.2.2) is a diagram:

$$\mathfrak{p} = (R \xleftarrow{q} S \xhookrightarrow{i} T)$$

Where i is the inclusion of a subtree, and q is a quotient map (a surjective map with connected fibers). If one has another correspondence

$$\mathfrak{q} = (R' \overset{q'}{\twoheadleftarrow} S' \overset{i}{\hookrightarrow} R)$$

One can form the composition:

$$\mathfrak{p}' \circ \mathfrak{p} = (R' \overset{q' \circ q}{\twoheadleftarrow} q^{-1}(S') \overset{j}{\hookrightarrow} T)$$

The elements of a C-Arb singularity $\mathcal{P}_T$ are equivalence classes of correspondences $\mathfrak{p} = (R \overset{q}{\twoheadleftarrow} S \overset{i}{\hookrightarrow} T)$, denoted $[\mathfrak{p}]$. They are equipped with a partial order in which $[\mathfrak{p}] \geq [\mathfrak{p}']$ iff there exists $\mathfrak{q}$ such that $\mathfrak{p} = \mathfrak{p}' \circ \mathfrak{q}$. There is a unique minimal element $\hat{0}$ represented by:

$$\mathfrak{p}_0 = (T \overset{\sim}{\leftarrow} T \overset{\sim}{\rightarrow} T)$$

Maximal elements are in one-to-one correspondence with subtrees S of T via:

$$S \Leftrightarrow (\star \twoheadleftarrow S \hookrightarrow T)$$

Where $\star$ is the tree with one vertex (corollary 2.3.1).

This construction ensures that neighborhoods of points in C-Arb singularities are again C-Arb singularities. Explicitly, if we let:

$$\mathfrak{p} = (R \overset{q}{\twoheadleftarrow} S \overset{i}{\hookrightarrow} T)$$

Then $N([\mathfrak{p}]) = \{[\mathfrak{q}] \in \mathcal{P}_T \mid [\mathfrak{q}] \geq [\mathfrak{p}]\} \cong \mathcal{P}_R$.

An important property of $\mathcal{P}_T$ which we will exploit repeatedly is that it is **atomistic** (lemma 2.3.2), meaning that each element can be written as the least upper bound of a set of **atoms**$-$minimal elements of $\mathcal{P}_T \backslash \{\hat{0}\}$. Such elements are of the form $\mathfrak{p} = (R \overset{q}{\twoheadleftarrow} S \overset{i}{\hookrightarrow} T)$ with $|V(R)| = |V(T)| - 1$. They come in two forms: deleting a terminal vertex of T or contracting an edge of T. For this reason we will often make use of the notation $K_T = V_t(T) \sqcup E(T)$, where $V_t(T)$ denotes the set of terminal vertices of T and $E(T)$ the set of edges.

As a consequence of $\mathcal{P}_T$ being atomistic, we have an embedding $\mathcal{P}_T \hookrightarrow \mathscr{P}(K_T)$ (the power set of K_T).

Example 1.2.1. *Let A_n denote the tree with n vertices connected in a straight line (Figure 1.4). We label the vertices* 1 *to n from left to right.*

Then $|K_{A_n}| = n + 1$, and one can show (lemma 2.4.1) $\mathcal{P}_{A_n} \cong \{X \subseteq K_{A_n} \mid |K| \leq n - 1\}$. In other words, $\mathcal{P}_{A_n}$ is isomorphic to the cone over the $n - 2-$skeleton of the combinatorial $n-$simplex. We illustrate $\mathcal{P}_{A_2}$ and $\mathcal{P}_{A_3}$ in Figure 1.5.

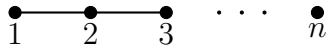


Figure 1.4: The A_n Tree

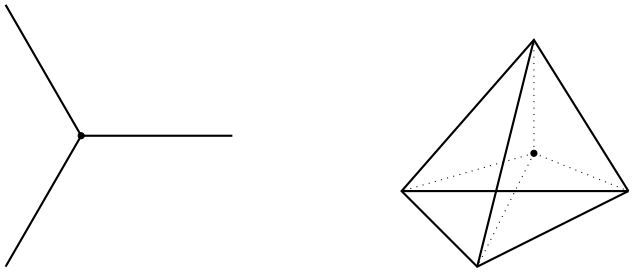


Figure 1.5: Illustrations of $\mathcal{P}_{A_2}$ (left) and $\mathcal{P}_{A_3}$ (right)

In Chapter 1 we give a similar description of the image of $\mathcal{P}_{S_n}$ in $\mathscr{P}(K_{S_n})$, where S_n is the tree with a central vertex connected to $n-1$ exterior vertices.

Summary of Chapter 3: Cyclic Structures

In this chapter we give the definition of a pre-cyclic structure and cyclic structure on a C-Arb space $\mathcal{P}$, as well as the construction of the cyclic structure $\mathcal{O}_\mu$ on $\mathcal{P}_T$ associated to an orientation μ of a tree T.

We provide an in-depth discussion of cyclic structures on $\mathcal{P}_{A_3}$ and $\mathcal{P}_{S_4}$, since according to the Coherence Theorem (theorem 1.1.1), these cyclic structures encode all the compatibility requirements distinguishing cyclic structures among pre-cyclic structures. As a part of this discussion, we construct a family pre-cyclic structures on $\mathcal{P}_{S_4}$ which are A_3-coherent, but not cyclic structures, showing that the S_4-coherence is a necessary condition.

We end the chapter by generalizing these observations to describe cyclic structures on $\mathcal{P}_{A_n}$ and $\mathcal{P}_{S_n}$. For instance, we show the following result:

Proposition 1.2.1. *(Proposition 4.4) Cyclic structures on $\mathcal{P}_{A_n}$ are in natural bijection with cyclic orders on K_{A_n}.*

Summary of Chapter 4: The Combinatorics of Cyclic Structures

In this section we undergo a detailed analysis of isomorphisms of C-Arb singularities, and of cyclic C-Arb singularities.

Note that we have an embedding $\mathcal{P}_T \hookrightarrow \mathscr{P}(K_T)$. So for any $\varphi \in \mathrm{Aut}(\mathcal{P}_T)$, we get $\overline{\varphi} \in \mathfrak{S}(K_T)$- the symmetric group on K_T. Recall $K_T = E(T) \sqcup V_t(T)$.

The key constructions of this chapter are the **reflection isomorphisms**: Let $v \in V(T)$ be a vertex of degree 2. If e, e' are the two edges in T incident to v, we construct a poset isomorphism:

$$\omega_v \in \mathrm{Aut}(\mathcal{P}_T)$$

Such that $\overline{\omega}_v$ is the transposition (ee'). Similarly, if $v \in V_t(T)$, and e is the unique edge incident to v, we construct $\omega_v \in \mathrm{Aut}(\mathcal{P}_T)$ such that $\overline{\omega}_v$ is the transposition (ve).

Our first main result in this section is that these reflection isomorphisms generate all nontrivial automorphisms of $\mathcal{P}_T$. In more detail: For any tree isomorphism $\sigma : T \xrightarrow{\sim} T'$, one obtains $i_\sigma : \mathcal{P}_T \xrightarrow{\sim} \mathcal{P}_{T'}$. Then:

Proposition 1.2.2. *(Corollary 4.2.2) If $\varphi : \mathcal{P}_T \xrightarrow{\sim} \mathcal{P}_{T'}$ is a poset isomorphism, there exists a tree isomorphism $\sigma : T \xrightarrow{\sim} T'$ such that $i_\sigma^{-1} \circ \varphi \in Aut(\mathcal{P}_T)$ is a product of reflection isomorphisms.*

We then extend these results to cyclic C-Arb singularities. If (T, μ) is a directed tree and $v \in V(T)$ is a source or sink, we define $r_v\mu$ to be the orientation μ with the edges incident to v reversed. We prove:

Lemma 1.2.1. *(Lemma 4.4.2) If v is a source or sink of degree ≤ 2, then $\omega_v : (\mathcal{P}_T, \mathcal{O}_\mu) \xrightarrow{\sim} (\mathcal{P}_T, \mathcal{O}_{r_v\mu})$ is an isomorphism of cyclic C-Arb singularities. In particular:*

$$\omega_v^* \mathcal{O}_{r_v\mu} = \mathcal{O}_\mu$$

We call such isomorphisms **directed reflection isomorphisms**. We also prove a similar classification result to proposition 1.2.2. As in that discussion, an isomorphism of σ directed trees induces an isomorphism i_σ of cyclic C-Arb singularities. Then:

Theorem 1.2.1. *(Theorem 4.4.1) If (T, μ) and (T', μ') are directed trees, and $\varphi : (\mathcal{P}_T, \mathcal{O}_\mu) \xrightarrow{\sim} (\mathcal{P}_{T'}, \mathcal{O}_{\mu'})$ is an isomorphism of cyclic C-Arb singularities, then there is an orientation λ on T and an isomorphism of directed trees $\sigma : (T, \lambda) \xrightarrow{\sim} (T', \mu')$ such that $i_\sigma^{-1} \circ \varphi : (\mathcal{P}_T, \mathcal{O}_\mu) \to (\mathcal{P}_T, \mathcal{O}_\lambda)$ is a product of directed reflection isomorphisms.*

As a first application of our understanding of isomorphisms of C-Arb singularities, we prove the Coherence Theorem (theorem 1.1.1).

Summary of Chapter 5: A Sheaf of Categories

Our main justification for our definition of cyclic structures is that a cyclic structure determines a sheaf of dg-categories on a C-Arb singularity up to quasi-equivalence. That is the content of the Cyclic Strucure Theorem (theorem 1.1.2). In this chapter we construct the sheaf of dg-categories $\mathcal{Q}_\mu$ associated to an orientation μ and prove this theorem.

To a directed tree (T, μ) one can associate the 'dg derived category' $\mathrm{Rep}(T, \mu)$. In this work, we consider $\mathrm{Rep}(T, \mu)$ to be the (formal) pretriangulated hull of the dg category $\mathrm{Rep}^0(T, \mu)$, which is a $k-$linear category with one object P_v for each vertex v.

If $\mathfrak{p} = (R \xleftarrow{q} S \xhookrightarrow{i} T)$, the stalk $\mathcal{Q}_{\mu,\mathfrak{p}}$ of the sheaf $\mathcal{Q}_\mu$ at the point $[\mathfrak{p}] \in \mathcal{P}_T$ is $\mathrm{Rep}(R, \lambda)$, where $\lambda = \mathfrak{p}^*\mu$ is the pullback of the orientation to R. The restriction map $c_\mathfrak{p} : \mathrm{Rep}(T, \mu) \to \mathrm{Rep}(R, \lambda)$ is characterized by its behavior on $\mathrm{Rep}^0(T, \mu)$:

$$c^0_\mathfrak{p} : \mathrm{Rep}^0(T, \mu) \to \mathrm{Rep}^0(R, \lambda)$$

$$c^0_\mathfrak{p}(P_v) = \begin{cases} P_{q(v)} & v \in V(S) \\ 0 & \text{Else} \end{cases}$$

(We must also characterize the behavior of $c^0_\mathfrak{p}$ on morphisms- see definition 5.3.2 for more details). This data defines a sheaf of dg-categories on the poset.

The proof of the Cyclic Structure Theorem (theorem 1.1.2), completed in this chapter, consists of two parts. The first assertion follows from a detailed analysis of the sheaf $\mathcal{Q}_\mu$ on $\mathcal{P}_{A_2}$. This sheaf can be viewed diagrammatically as in Figure 1.6. In that figure, V and W are finite-dimensional chain complexes and f is an injective morphism.

One sees that there is a cyclic rotation symmetry to this sheaf: Applying the quasi-equivalence $(V \xrightarrow{f} W) \mapsto (W \xrightarrow{i} \mathrm{Cone}(f))$ to the stalk at the vertex, one gets an equivalent presentation as shown in Figure 1.7:

And one notices that the restriction functors have been rotated clockwise. We prove:

Proposition 1.2.3. *(Proposition 5.4.1) Let (A_2, μ) be a directed tree. If $\varphi \in Aut(\mathcal{P}_{A_2})$ and $\varphi^*\mathcal{Q}_\mu$ is quasi-equivalent to $\mathcal{Q}_\mu$, then $\varphi^*\mathcal{O}_\mu = \mathcal{O}_\mu$.*

It follows easily that if (T, μ) and (T', μ') are directed trees, and $\varphi : \mathcal{P}_T \xrightarrow{\sim} \mathcal{P}_{T'}$ be a poset isomorphism such that $\varphi^*\mathcal{Q}_{\mu'}$ is quasi-equivalent to $\mathcal{Q}_\mu$. Then we must have $\varphi^*\mathcal{O}_{\mu'} = \mathcal{O}_\mu$.

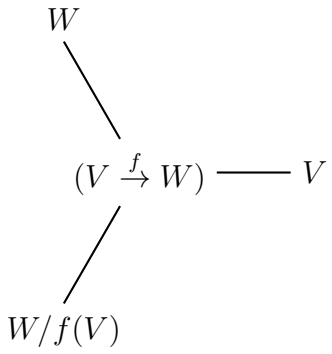


Figure 1.6: A diagram of the sheaf $\mathcal{Q}_\mu$ on $\mathcal{P}_{A_2}$

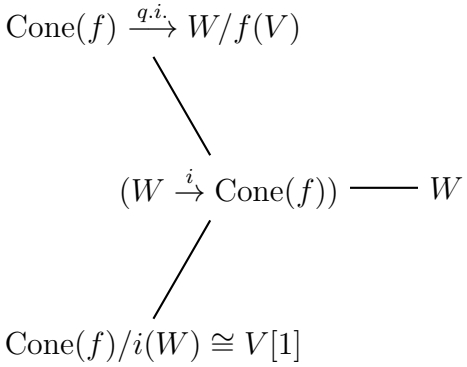


Figure 1.7: The sheaf $\mathcal{Q}_\mu$ after applying a quasi-equivalence of global sections

This is half of the Cyclic Structure Theorem. The bulk of the remainder of the chapter is devoted to showing the following:

Proposition 1.2.4. *(Proposition 5.5.1) Let (T, μ) be a directed tree, $v \in V(T)$ be a source or sink of degree ≤ 2, and $\omega_v : (\mathcal{P}_T, \mathcal{O}_\mu) \xrightarrow{\sim} (\mathcal{P}_T, \mathcal{O}_{r_v\mu})$ be the directed reflection isomorphism. Then there exists a quasi-equivalence:*

$$\mathcal{Q}_\mu \xrightarrow{\sim} \omega_v^* \mathcal{Q}_{r_v\mu}$$

Combining proposition 1.2.3 with proposition 1.2.4 yields the Cyclic Structure Theorem.

Proving proposition 1.2.4 requires the construction of a morphism:

$$\Gamma_v^+ : \mathcal{Q}_\mu \to \omega_v^* \mathcal{Q}_{r_v\mu}$$

When v is a source of degree ≤ 2, and likewise:

$$\Gamma_v^- : \mathcal{Q}_\mu \to \omega_v^* \mathcal{Q}_{r_v\mu}$$

When v is a sink. These morphisms are essentially the BGP reflection functors of [2], a classical construction which gives quasi-equivalences of the categories $\mathrm{Rep}(T, \mu)$ and $\mathrm{Rep}(T, r_v\mu)$ whenever v is a source or sink. Our proof extends these functors to give quasi-equivalences of *sheaves* of dg-catgories, meaning they must be compatible with the restriction functors.

Additionally, we must show these functors are inverses, which formally involves some higher coherence data. These constructions involve some tedious verifications of commutative diagrams, which we meticulously present in Chapter 5.

Summary of Chapter 6: Graded Structures and Graded Branes

While we have shown that one can construct a sheaf of dg-categories up to quasi-equivalence on a cyclic C-Arb singularity, we have not given a canonical construction- i.e. one that doesn't require the choice of an isomorphism with $(\mathcal{P}_T, \mathcal{O}_\mu)$. Such a construction is important if one wants to consider sheaves on general C-Arb spaces- since it gives an explicit recipe for gluing together local models.

We remark that such a canonical construction exists for A_n singularities, appearing in [21].

This chapter does not solve the general problem, but presents some ideas in this direction. We construct a set of objects, called **graded branes**, which are meant to serve as the objects in this theoretical canonical category. The use of the term brane is meant to be suggestive: We expect the correct answer to this problem is an $A_\infty-$category which models the Fukaya category of a symplectic neighborhood of an arboreal singularity.

To begin, we introduce the notion of a **graded structure** G on a C-Arb space $\mathcal{P}$ (definition 6.1.1). To a graded structure, we associate a pre-cyclic structure $\mathcal{O}_G$ (definition 6.1.2). We are able to show that $\mathcal{O}_G$ is in fact a cyclic structure (proposition 6.1.1) by using the Coherence Theorem (theorem 1.1.1). We also prove that, locally, each cyclic structure is induced by a unique (up to isomorphism) graded structure. In particular, for any orientation μ of T there is a canonical graded structure G_μ such that $\mathcal{O}_{G_\mu} = \mathcal{O}_\mu$.

We then define, for a C-Arb space, a set of *pre-branes* $\mathrm{Br}_{pre}(\mathcal{P})$, which are essentially subsets representing top-dimensional cycles over $\mathbb{Z}/2\mathbb{Z}$ (definition 6.2.1). As an example, pre-branes on A_2 are illustrated in Figure 1.8.

The set of pre-branes forms a vector space over $\mathbb{Z}/2\mathbb{Z}$. When $\mathcal{P} = \mathcal{P}_T$ we have a 'standard basis' of the form $\{L_v\}_{v\in V(T)}$ (though note- this basis is not invariant under automorphisms of $\mathcal{P}_T$).

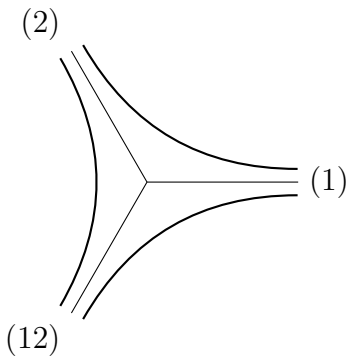


Figure 1.8: Pre-Branes on $\mathcal{P}_{A_2}$

For a cyclic C-Arb space $(\mathcal{P}, \mathcal{O})$, one has a subset of *branes* $\mathrm{Br}(\mathcal{P}, \mathcal{O}) \subset \mathrm{Br}_{pre}(\mathcal{P})$ given by a coherence condition in codimension-two (definition 6.2.3). We give a nice classification result:

Proposition 1.2.5. *(Proposition 6.2.1) Given a directed tree (T, μ), we can construct a bijection:*

$$Br(\mathcal{P}_T, \mathcal{O}_\mu) \setminus \{0\} \longleftrightarrow \left\{ \begin{array}{l} \textit{Subtrees of } T \textit{ isomorphic} \\ \textit{to } A_n \textit{ for some } n. \end{array} \right\}$$

Where 0 *is the 'zero brane' corresponding to the empty set.*

Furthermore, to a graded structure G one can construct the set of *graded branes* $\mathrm{Br}_{gr}(\mathcal{P}, G)$. There is a 'forgetful map'

$$\mathrm{Br}_{gr}(\mathcal{P}, G) \to \mathrm{Br}(\mathcal{P}, \mathcal{O}_G)$$

Whose nonempty fibers have a natural free $\mathbb{Z}-$action. When $\mathcal{P} = \mathcal{P}_T$, the forgetful map is surjective, and $\mathbb{Z}$ acts transitively on its fibers.

Our main result in this chapter is a representation-theoretic interpretation of branes. Namely, we say $X \in \mathrm{Rep}(T, \mu)$ is **rank-one** if, for any maximal correspondence $\mathfrak{p} = (\star \twoheadleftarrow S \overset{i}{\hookrightarrow} T)$, $c_{\mathfrak{p}} X \in C(k)$ has one-dimensional cohomology. ($C(k) \cong \mathrm{Rep}(\star, \cdot)$ is the dg category of finite-dimensional chain complexes over k). Then one has a bijection (propositions 6.3.2, 6.3.3):

$$\left\{ \begin{array}{l} \text{Rank one objects} \\ \text{in } \mathrm{Rep}(T, \mu) \end{array} \right\} \Big/ \text{homotopy} \longleftrightarrow \mathrm{Br}_{gr}(\mathcal{P}_T, G_\mu)$$

The rest of the chapter gives a more in-depth discussion of branes and graded structures in the case $T = A_n$. In particular, we provide a description of branes and graded branes that

doesn't explicitly reference the tree (like proposition 1.2.5 does).

Summary of Chapter 7: Arboreal Moves

In the theory of trivalent ribbon graphs there is a 'move' or 'mutation' that preserves the associated Riemann surface (Figure 7.1).

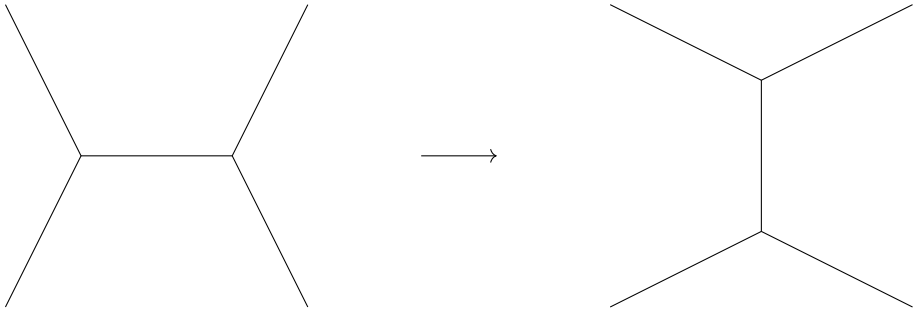


Figure 1.9: The H-to-I move

This move can be viewed as a transformation of cyclic C-Arb spaces. Viewing this transformation as a continuous deformation over time, the total space of this move is the A_3- singularity (Figure 7.2).

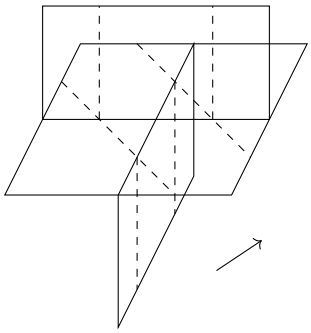


Figure 1.10: Total Space of the H-to-I move

In this section, our goal will be to generalize this observation to describe a wide class of moves of C-Arb spaces whose total space is a C-Arb singularity.

The process goes as follows: If T is a tree, for a subset K of K_T, we define:

$$U_K := \bigcup_{k \in K} N(\langle k \rangle)$$

And we say that K is **good** with respect to an orientation μ on T if the restriction $\Gamma(\mathcal{Q}_\mu; \mathcal{P}_T) \to \Gamma(\mathcal{Q}_\mu; U_K)$ is a quasi-equivalence of dg-categories. Then a partition: $K_T = K_1 \sqcup K_2$ into good sets determines an **arboreal move of type** T between cyclic C-Arb spaces:

$$(U_{K_1}, i_1^* \mathcal{O}_\mu) \to (U_{K_2}, i_2^* \mathcal{O}_\mu)$$

The bulk of this chapter is focused on a classification of good subsets. Our main tool will be the use of combinatorial homotopies (specifically, retracts) to conclude that certain maps are quasi-equivalences without explicitly computing a limit of dg-categories. We introduce the following notion:

Definition 1.2.1. *(Definition 7.4.1) Let (T, μ) be a directed tree. A **generalized cyclic interval** is a subset of K_T of the form $I(v, w, D)$, where:*

- *$v, w \in V(T)$*
- *D is a subset of the terminal vertices $V_t(T)$ of T.*
- *$I(v, w, D) \subset K_T$ consists of all elements of D, all edges between a point in D and the shortest path joining v and w, and all edges on the shortest path joining v and w which are pointing from v to w.*

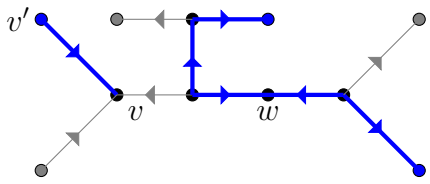


Figure 1.11: A Generalized Cyclic Interval $I(v, w, D)$.

An example is given in Figure 1.11. Elements of $I(v, w, D)$ are thick blue edges and blue terminal vertices, while elements not in $I(v, w, D)$ are thin gray lines and gray terminal vertices. D is the set of blue terminal vertices. The result is:

(Theorem 1.1.4) A subset K of K_T is good if and only if it is not a generalized cyclic interval.

We conclude this chapter with an explicit discussion on moves of type A_n and S_n. Explicitly, we have the following nice characterization. Recall (proposition 4.4) that giving a cyclic structure $\mathcal{O}$ on $\mathcal{P}_{A_n}$ is equivalent to giving a cyclic order on K_{A_n}. Then:

Proposition 1.2.6. *(Proposition 7.5.1) $K \subset K_{A_n}$ is a generalized cyclic interval with respect to an orientation μ iff it is a cyclic interval in the cyclic order determined by $\mathcal{O}_\mu$.*

As a result, moves of type A_n are equivalent to giving a partition of a cyclically ordered set of $n+1$ elements into two subsets, neither of which are cyclic intervals. For example, when $n = 3$, the unique valid partition of 4 cyclically ordered elements (two opposite pairs) gives the H-to-I move.



Figure 1.12: The unique (up to equivalence) partitions of cyclic sets yielding moves of type A_3 and A_4

Summary of Chapter 8: The Geometry of Arboreal Spaces

In this chapter, we bravely leave the combinatorial realm and attempt some honest symplectic geometry.

Of central importance is the notion of a stratified space. Our definition resembles that in [17].

Definition 1.2.2. *(Definition 8.1.1) A **stratified space** is a (Hausdorff, locally compact, second countable) topological space X which is can be written as a disjoint union: $X = \sqcup_{\alpha \in A} X_\alpha$, where the X_α are called **strata**. In addition we have:*

(i) Each X_α is equipped with the structure of a smooth manifold of some dimension, which is compatible with the subspace topology on X_α.

(ii) For $\alpha \neq \beta \in A$, if X_β intersects $\overline{X}_\alpha$, then $X_\beta \subset \overline{X}_\alpha$ and $dim(X_\beta) < dim(X_\alpha)$.

We then describe a construction given in [20] of the arboreal singularity L_T associated to a tree T. Each stratified space X has an associated poset $\mathscr{S}(X)$ whose elements are the

strata of X. We have $\mathscr{S}(L_T) \cong \mathcal{P}_T$- in other words, the strata of L_T are written $L_T(\mathfrak{p})$, for correspondences $\mathfrak{p} = (R \overset{q}{\twoheadleftarrow} S \overset{i}{\hookrightarrow} T)$.

Following this, we define an arboreal space as a stratified space which is locally isomorphic to a stratified space of the form $L_T \times \mathbb{R}^k$. As shown in (ref), L_T itself satisfies the definition: For any correspondence $\mathfrak{p} = (R \overset{q}{\twoheadleftarrow} S \overset{i}{\hookrightarrow} T)$ there is an open neighborhood of $L_T(\mathfrak{p})$ isomorphic to $L_R \times \mathbb{R}^{|T|-|R|}$.

We then discuss *Lagrangian embeddings* of arboreal spaces: Topological embeddings $i : L \hookrightarrow (M, \omega)$, where:

- (M, ω) is a symplectic manifold.
- i a smooth on strata and maps each strata to an isotropic submanifold.
- Certain regularity conditions on the relative positions of the strata (see definition 8.3.1).

We define the notion of a pre-cyclic structure on an arboreal space mimicking the definition for a C-Arb space, and show that a Lagrangian embedding induces a canonical pre-cyclic structure. In particular, for a tree T there is a natural bijection $\{\text{Pre-cyclic structures on } L_T\} \leftrightarrow \{\text{Pre-cyclic structures on } \mathcal{P}_T\}$.

What we term our 'main conjecture' for Lagrangian embeddings is a bijection of the form:

$$\left\{ \begin{array}{l} \text{'Regular' Lagrangian} \\ \text{embeddings of } L_T \end{array} \right\} \Big/ \text{equivalence} \longleftrightarrow \{\text{Cyclic Structures on } \mathcal{P}_T\}$$

While we give a plausible notion of 'equivalence', we are not able to give a satisfactory notion of 'regular'- though, we do explain why such a notion is necessary. Furthermore, we have a map:

$$\chi_T : \mathrm{Sh}(\mathcal{P}_T; \mathrm{dgCat}) \to \mathrm{Sh}(L_T; \mathrm{dgCat})$$

Yielding what we expect to be a commutative (up to quasi-equivalence) diagram:

$$\begin{array}{ccc} i : L_T \hookrightarrow (M, \omega) & \longrightarrow & (\mathcal{P}_T, \mathcal{O}) \\ \downarrow & & \downarrow \\ \mathrm{Sh}_{i(L_T)}(M) & \xleftarrow{\chi_T} & \mathcal{Q} \end{array}$$

Where i is a 'regular' Lagrangian embedding, $\mathrm{Sh}_{i(L_T)}(M)$ is the category of sheaves on M with microsupport along $i(L_T)$, and $\mathcal{Q} = \varphi^* \mathcal{Q}_\mu$ for some isomorphism $\varphi : (\mathcal{P}_T, \mathcal{O}) \xrightarrow{\sim} (\mathcal{P}_T, \mathcal{O}_\mu)$.

Lastly, we give a definition of 'arboreal moves' of cyclic arboreal spaces (arboreal spaces equipped with a cyclic structure). Here, the conjectures are even less precisely formulated- we are content to give some examples of geometric realizations of some of the combinatorial moves described in Chapter 7, as well as some geometric moves which are not as neatly captured by our formalism from that chapter.

Chapter 2

Combinatorial Arboreal Singularities and Arboreal Spaces

2.1 Posets

References for posets and the poset topology are [4], [35].

Definition 2.1.1. *A **poset** $\mathcal{P}$ is a set equipped with a partial order $\leq$. Throughout this paper we will view a poset as a T_0 topological space, where $U \subset \mathcal{P}$ is open if it is **upward closed**: for every $x \leq y \in \mathcal{P}$, if $x \in U$ then $y \in U$. A **open embedding** is a continuous map $\mathcal{P}_1 \hookrightarrow \mathcal{P}_2$, which is injective and has open image. For $x \in \mathcal{P}$, we have the **neighborhood of** x: $N(x) := \{y \in \mathcal{P} \mid x \leq y\}$, which is the intersection of all open sets containing x.*

We let **Pos** denote the category of posets whose morphisms are open embeddings.

Definition 2.1.2. *Let $\mathcal{P}$ be a poset. $\mathcal{P}$ is **graded** if there exists a* dimension function *$dim : \mathcal{P} \to \mathbb{Z}_{\geq 0}$ such that:*

(G1) Any minimal element x has $dim(x) = 0$.

(G2) If $x < y$, $dim(x) < dim(y)$.

(G3) If $x < y$ and there does not exist z with $x < z < y$, then $dim(y) = dim(x) + 1$

*$\mathcal{P}$ is **co-graded** if there exists a* codimension function *$codim : \mathcal{P} \to \mathbb{Z}_{\geq 0}$ which is a dimension function for $\mathcal{P}^{op}$: the poset $\mathcal{P}$ with the opposite order relation.*

Note that, for a graded (resp. co-graded) poset, the dimension function (resp. codimension function) is uniquely determined. If $\mathcal{P}$ is graded, we let $C_j(\mathcal{P}) = \{x \in \mathcal{P} \mid \dim(x) = j\}$.

We sometimes refer to these elements as $j-$*cells*. Similarly, if $\mathcal{P}$ is co-graded we let $C^j(\mathcal{P}) = \{x \in \mathcal{P} \mid \text{codim}(x) = j\}$.

Definition 2.1.3. *Let $\mathcal{P}$ be a poset with a least element $\hat{0}$ (that is, $\mathcal{P} = N(\hat{0})$). Then an* ***atom*** *of $\mathcal{P}$ is a minimal element of $\mathcal{P}\backslash\{\hat{0}\}$. In the case $\mathcal{P}$ is graded, atoms are the same as one-cells. $\mathcal{P}$ is* ***atomistic*** *if each element is the least upper bound of a set of atoms.*

Note that if $\mathcal{P}$ is atomistic, then there is an injection $\mathcal{P} \hookrightarrow \mathscr{P}(C_1(\mathcal{P}))$- where $\mathscr{P}(X)$ denotes the power set of X- which sends $x \in \mathcal{P}$ to $\{y \in C_1(\mathcal{P}) \mid y \leq x\}$.

We continue by reviewing a construction from [20], from an abstract perspective.

First, let $\mathcal{C}$ be any category[1]

Definition 2.1.4. *For an object $A \in Ob(\mathcal{C})$, let $\mathcal{C}_{/A}$ denote the slice category: Objects of $\mathcal{C}_{/A}$ are arrows $f : B \to A$, and a morphism $f \to g$ is a map h satisfying $f = gh$.*

Definition 2.1.5. *A category is* ***thin*** *if there is at most one morphism between any two objects. For any thin category $\mathcal{D}$, we can construct a poset $Pos(\mathcal{D})$ whose objects are isomorphism classes of objects in $\mathcal{D}$. If $X, Y \in Ob(\mathcal{D})$, and $[X]$ and $[Y]$ are the corresponding equivalence classes, we set $[X] \leq [Y]$ if and only if there exists a map from Y to X in $\mathcal{D}$.*

Now, let $\mathcal{C}$ be a category in which every morphism is monic. For $A \in \text{Ob}(\mathcal{C})$, $\mathcal{C}_{/A}$ is a thin category. Write $\mathcal{P}_A = \text{Pos}(\mathcal{C}_{/A})$. The following facts are readily verified: (We use $[f]$ to denote the equivalence class of a morphism $f \in \mathcal{C}_A$)

Lemma 2.1.1.

(i) *$\mathcal{P}_A$ has a unique minimal element, represented by the identity morphism $id_A : A \to A$ in $\mathcal{C}_A$.*

(ii) *If $f : A \to B$ is a morphism, the map $i_f : \mathcal{P}_A \to \mathcal{P}_B$ defined by $i_f([g]) = [f \circ g]$ is well-defined, and is an open embedding of posets. Furthermore, the image of $\mathcal{P}_A$ under i_f is the open neighborhood $N([f]) \subset \mathcal{P}_B$.*

[1]For this exposition our categories will be assumed to be essentially small- that is, the collection of isomorphism classes of objects forms a set, and morphisms between any two objects form a set.

□

Part (ii) of the lemma implies that the assignment $A \mapsto \mathcal{P}_A$, $f \mapsto i_f$ defines a functor $\mathcal{C} \to \mathbf{Pos}$.

Definition 2.1.6. *A **combinatorially graded category**[2] $\mathcal{C}$ is a category in which every morphism is monic, equipped with a functor $|\cdot| : \mathcal{C} \to \mathbb{N}$, where the category $\mathbb{N}$ has objects the positive integers, and a single morphism from m to n if and only if $m \leq n$. We require that $|\cdot|$ satisfies the following properties:*

(i) If $X, Y \in Ob(\mathcal{C})$ with $|X| = |Y|$, any morphism $X \to Y$ is an isomorphism.

(ii) If $X, Y \in Ob(\mathcal{C})$, if there exists a morphism $f : X \to Y$ and $|X| < |Y|$, there exists $Z \in Ob(\mathcal{C})$ with $|Z| = |Y| - 1$, and morphisms $g : X \to Z$, $h : Z \to Y$, with $f = hg$.

(iii) For any $X \in Ob(\mathcal{C})$, there exists $Y \in Ob(\mathcal{C})$ with $|Y| = 1$, and a morphism $f : Y \to X$.

Lemma 2.1.2. *If $\mathcal{C}$ is a combinatorially graded category, then for any $A \in Ob(\mathcal{C})$, $\mathcal{P}_A$ is graded and co-graded. In particular, for $f : B \to A$, $dim([f]) = |A| - |B|$ and $codim([f]) = |B| - 1$ are dimension/codimension functions.*

Proof. First we show dim is a dimension function.

For G1: we know $\mathcal{P}_A$ has a unique minimal element represented by id_A, and $\dim([\mathrm{id}_A]) = |A| - |A| = 0$, as desired.

For G2: If $[f] < [g]$, with $f : X \to A$ and $g : Y \to A$, there exists $h : Y \to X$ such that $g = f \circ h$. Since $|\cdot|$ is a functor, that implies $|Y| \leq |X|$. If $|Y| = |X|$, then h is an isomorphism by (i), which would make $[f] = [g]$, hence $|Y| < |X|$. So we have $\dim([f]) = |A| - |X| < |A| - |Y| = \dim([g])$.

For G3: Suppose $[f] < [g]$, with $f : X \to A$ and $g : Y \to A$, and there does not exist z with $[f] < [z] < [g]$. Then there exists $h : Y \to X$ such that $g = f \circ h$, and $|Y| < |X|$. By (ii) there exists Z with $|Z| = |X| - 1$, and morphisms $j : Y \to Z$, $k : Z \to X$, with $h = k \circ j$. So we have $[f] \leq [f \circ k] \leq [f \circ k \circ j] = [g]$. By assumption we must have $[g] = [f \circ k]$ or $[f] = [f \circ k]$. The latter is impossible, as $|Z| = |X| - 1$, so $[g] = [f \circ k]$, hence $|Y| = |Z| = |X| - 1$, and $\dim([g]) = \dim([f]) + 1$.

[2] Terminology is of my own creation, though most likely this notion exists elsewhere in the literature

To show that codim is a codimension function, the proofs of G2 and G3 are identical. To prove G1, note that by (iii) any minimal element in $\mathcal{P}_A^{op}$ must be of the form $f : B \to A$ with $|B| = 1$. Hence $\mathrm{codim}([f]) = |B| - 1 = 0$ as desired.

□

In this paper, we will use the above formalism for the categories of trees and directed trees to construct our central objects of interest: Combinatorial arboreal singularities.

2.2 Trees

Material in this section can also be found in [20].

To fix notation, for a set X, we let $\mathscr{P}(X)$ denote the power set of X. We review some basic definitions of graph theory (with conventions adopted to our purposes- we only consider finite nonempty graphs with no loops or multi-edges):

Definition 2.2.1. *A **graph** G is the data of a finite nonempty set $V(G)$ (whose elements are called **vertices**), along with a set of edges $E(G) \subseteq \mathscr{P}(V(G))$ consisting of subsets of $V(G)$ of size 2. A **path** in G is a sequence $(v_0, v_1, \ldots, v_n)$, where $v_i \in V(G)$, and $\{v_i, v_{i+1}\} \in E(G)$ for all $0 \leq i < n$. A graph is **connected** if every pair of vertices can be joined by a path. A **tree** is a connected graph with no cycles, where a cycle is a path $(v_0, v_1, \ldots, v_n = v_0)$, where $v_0, \ldots, v_{n-1}$ are all distinct vertices, and $n > 2$. If G_1 and G_2 are graphs, a map of graphs $f : G_1 \to G_2$ is the data of a function (which we also call f) $f : V(G_1) \to V(G_2)$ such that if $\{v, w\} \in E(G_1)$, either $f(v) = f(w)$ or $\{f(v), f(w)\} \in E(G_2)$.*

*If G and H are graphs, H is called a **subgraph** of G if $V(H) \subset V(G)$ and $E(H) \subset E(G)$. A **subtree** is a subgraph which is a tree. Note that if S is a subtree of T, S is determined by $V(S) \subseteq V(T)$- because of this we will occasionally refer to a subset of $V(T)$ as a subtree of T.*

*If T_1, T_2 are trees a map $q : T_1 \to T_2$ is a **quotient map** if it is surjective and, for $v \in V(T_2)$, $q^{-1}(\{v\})$ is a connected subtree of T_1.*

Definition 2.2.2. *The category **Tree** is the category whose objects are trees and whose morphisms are correspondences of trees. A **correspondence** of trees is a diagram:*

$$\mathfrak{p} = (R \overset{q}{\twoheadleftarrow} S \overset{i}{\hookrightarrow} T)$$

Where S is a subtree of T, and q is a quotient map. If we have a pair of correspondences:

$$\mathfrak{p} = (R \overset{q}{\twoheadleftarrow} S \overset{i}{\hookrightarrow} T) \qquad \mathfrak{p}' = (R' \overset{q'}{\twoheadleftarrow} S' \overset{i'}{\hookrightarrow} R)$$

The composition is given by:

$$\mathfrak{p}' \circ \mathfrak{p} = (R' \overset{q' \circ q}{\twoheadleftarrow} q^{-1}(S') \overset{j}{\hookrightarrow} T)$$

A correspondence is pictured in Figure 2.1. The subtree S consists of the fillled-in vertices and solid lines, while the edges contracted by q are the thick edges- we will use this convention to illustrate correspondences throughout this work.

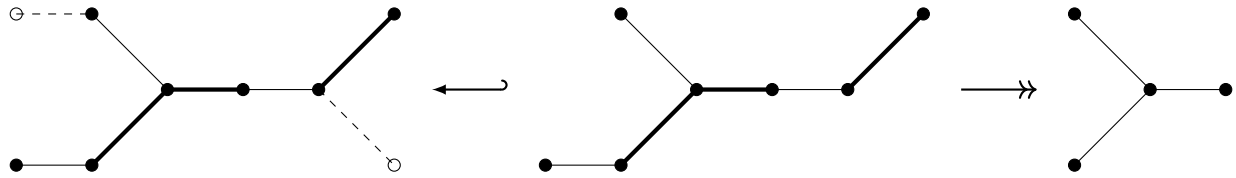


Figure 2.1: A Correspondence

It is easy to verify that the composition of correspondences is associative. In light of the above definition, we will sometimes use the notation $\mathfrak{p} : R \to T$ to denote a correspondence of the form $(R \twoheadleftarrow S \hookrightarrow T)$.

Lemma 2.2.1. *Each morphism in* **Tree** *is monic.*

Proof. Suppose we have morphisms:

$$\mathfrak{p} = (R \overset{q}{\twoheadleftarrow} S \overset{i}{\hookrightarrow} T) \qquad \mathfrak{q}_1 = (X \overset{w_1}{\twoheadleftarrow} Y_1 \overset{j_1}{\hookrightarrow} R) \qquad \mathfrak{q}_2 = (X \overset{w_2}{\twoheadleftarrow} Y_2 \overset{j_2}{\hookrightarrow} R)$$

With $\mathfrak{p} \circ \mathfrak{q}_1 = \mathfrak{p} \circ \mathfrak{q}_2$. Then $q^{-1}(Y_1) = q^{-1}(Y_2)$. Since q is surjective, this implies $Y_1 = Y_2$. Call this subtree Y. We have $w_1 \circ q = w_2 \circ q$ as maps $q^{-1}(Y) \to Y \to X$. Since q is surjective we have $w_1 = w_2$. □

Definition 2.2.3. *We can now apply the construction from the beginning of the section to a tree T to obtain a poset $\mathcal{P}_T$. A poset which is isomorphic to $\mathcal{P}_T$ for some tree T is a* ***combinatorial arboreal singularity****, or* ***C-Arb singularity****.*

Definition 2.2.4. *A* ***combinatorial arboreal space****, or* ***C-Arb space****, is a cograded poset $\mathcal{P}$ such that, for all $x \in \mathcal{P}$, $N(x)$ is a C-Arb singularity.*

2.3 Basic Properties

Let us examine some properties of $\mathcal{P}_T$. We use the notation from definition 2.1.5- Whenever $\mathfrak{p} : R \to T$ is a correspondence, $[\mathfrak{p}]$ will denote the element of $\mathcal{P}_T$.

First we introduce some terminology that will be crucial in what follows:

Definition 2.3.1. *For a tree T, let $|T| = \#(V(T))$. If $v \in V(T)$, the **degree** of v is the number of edges incident to v (equivalently, the number of vertices w with $\{v, w\} \in E(T)$). A vertex is called **terminal** if its degree is one. We let $V_t(T)$ denote the set of terminal vertices. Note that $T\backslash\{v\}$ is a subtree of T if and only if v is a terminal vertex of T.*

Lemma 2.3.1. *The category **Tree**, with $|\cdot|$ as defined above, is a combinatorially graded category. In particular:*

(i) If $\mathfrak{p} : R \to T$, then $|R| \leq |T|$ with equality if and only if $\mathfrak{p}$ is an isomorphism.

(ii) Let $\mathfrak{p} : R \to T$ be any correspondence with $|R| < |T|$. Then there exists a correspondence $\mathfrak{q} : R' \to T$ with $|R'| = |T| - 1$, and $[\mathfrak{p}] \geq [\mathfrak{q}]$.

(iii) Let $\star$ denote the tree with a single vertex. Then for any tree T there exists a correspondence $\mathfrak{p} : \star \to T$.

Proof.

(i) If $\mathfrak{p} = (R \overset{q}{\twoheadleftarrow} S \overset{i}{\hookrightarrow} T)$ is a correspondence, then it is clear $|R| \leq |S| \leq |T|$. If $|R| = |T|$, then $S = T$ and q must be an isomorphism, meaning $\mathfrak{p}$ is an isomorphism with inverse $(T \overset{q^{-1}}{\twoheadleftarrow} R \overset{=}{\hookrightarrow} R)$.

(ii) Let $\mathfrak{p} = \{R \overset{q}{\twoheadleftarrow} S \overset{i}{\hookrightarrow} T\}$ be any correspondence with $|R| < |T|$. If $|S| < |T|$, then there is a terminal vertex $v \in V_t(T)$ which is not in $V(S)$, and we can take $\mathfrak{q} = \{T\backslash\{v\} \overset{\sim}{\leftarrow} T\backslash\{v\} \hookrightarrow T\}$. If $|S| = |T|$, we have $|R| < |S|$, and so there are two neighboring vertices $v, w \in V(S)$ such that $q(v) = q(w)$. Then we can set $\mathfrak{q} = \{T/(v \sim w) \twoheadleftarrow T \overset{=}{\to} T\}$.

(iii) Given any vertex $v \in V(T)$ we can set $\mathfrak{p} = \{\star \overset{\sim}{\leftarrow} \{v\} \hookrightarrow T\}$.

□

Corollary 2.3.1. *For any tree T, $\mathcal{P}_T$ is a graded and co-graded poset. For $\mathfrak{p} : R \to T$, $dim([\mathfrak{p}]) = |T| - |R|$, and $codim([\mathfrak{p}]) = |R| - 1$. In particular:*

(i) $[\mathfrak{p}]$ *is maximal iff* $|R| = 1$*: Such correspondences are in bijection with subtrees* S *of* T*, via* $S \Leftrightarrow (\star \twoheadleftarrow S \hookrightarrow T)$*, where* $\star$ *is the tree with a single vertex.*

(ii) $[\mathfrak{p}]$ *is a one-cell iff* $|R| = |T| - 1$*. Such correspondences arise from contracting an edge or deleting a terminal vertex, so we have a bijection* $C_1(\mathcal{P}_T) \cong E(T) \sqcup V_t(T)$*.*

Corollary 2.3.2. *If* $\mathcal{P}_T \cong \mathcal{P}_{T'}$*, then* T *and* T' *have the same number of vertices.*

In fact, if $\mathcal{P}_T \cong \mathcal{P}_{T'}$, we must have $T \cong T'$. This is not trivial to prove, however, and is established in Chapter 4.

We have the following important lemma:

Lemma 2.3.2. $\mathcal{P}_T$ *is atomistic.*

Proof. We want to show that a correspondence $\mathfrak{p} = (R \overset{q}{\twoheadleftarrow} S \overset{i}{\hookrightarrow} T)$ is the least upper bound of the set $\{x \in C_1(\mathcal{P}_T) \mid x \leq [\mathfrak{p}]\}$. For the rest of this proof, denote this set by X.

Use the identification $C_1(\mathcal{P}_T) \cong E(T) \sqcup V_t(T)$. If $e \in E(T)$, we see that $e \in X$ if and only if e is disjoint from S (i.e. none of its endpoints are in S) or is in S and contracted by q. Similarly, if $v \in V_t(T)$, $v \in X$ if and only if v is not in S. As an example, consider the correspondince in Figure 2.2, (which uses the same convention of Figure 2.1) with the edges and terminal vertices labelled:

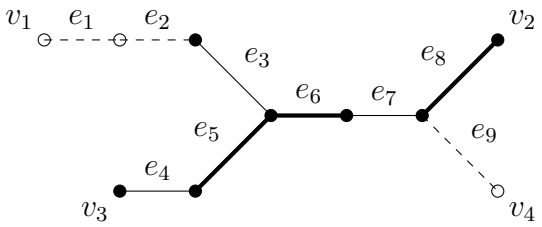


Figure 2.2: A Correspondence with Edges/Terminal Vertices Labelled

The set X is $\{v_1, v_4\} \cup \{e_1, e_5, e_6, e_8\}$.

We note that $v \in V(S)$ if and only if v *cannot* be connected by a path to a terminal vertex $w \in X$ such that each edge in the path is in X. So if $\mathfrak{q} = (R' \overset{q'}{\twoheadleftarrow} S' \overset{i}{\hookrightarrow} T)$ is another correspondence such that $x \leq [\mathfrak{q}]$ for each $x \in X$, we must have $S' \subset S$.

Secondly, suppose $v \in V(S')$, $w \in V(S)$, and v and w are connected by an edge $e \in X$, so e is contracted by q. Then $e \leq [\mathfrak{q}]$, meaning e is either disjoint from S' or contracted by q', and the assumption that $v \in V(S')$ means that the latter must hold.

These two statements imply that $[\mathfrak{q}] \geq [\mathfrak{p}]$. So $[\mathfrak{p}]$ is the least upper bound of X as claimed. □

Notation

Following the lemma, we introduce the following notation, which we use throughout this paper:

- For a tree T, let $K_T = E(T) \sqcup V_t(T)$. For $k \in K_T$, we let $\langle k \rangle \in C_1(\mathcal{P}_T)$ denote the corresponding one-cell.
- For a correspondence $\mathfrak{p}$, let $X_{\mathfrak{p}} = \{k \in K_T \mid \langle k \rangle \leq [\mathfrak{p}]\}$. By Lemma 2.5, this gives an injection from $\mathcal{P}_T$ into the power set $\mathscr{P}(K_T)$. (Though this is not an open embedding).
- It follows that there is an embedding $\mathrm{Aut}(\mathcal{P}_T) \hookrightarrow \mathfrak{S}(K_T)$, where $\mathfrak{S}(K_T)$ denotes the symmetric group on K_T. For $\varphi \in \mathrm{Aut}(\mathcal{P}_T)$, let $\overline{\varphi}$ denote its image in $\mathfrak{S}(K_T)$.
- More generally, if $\varphi : \mathcal{P}_T \xrightarrow{\sim} \mathcal{P}_{T'}$ is a poset isomorphism, let $\overline{\varphi} : K_T \xrightarrow{\sim} K_{T'}$ denote the corresponding bijection of sets.

2.4 Some Examples: A_n and S_n singularities

Description of $\mathcal{P}_{A_n}$

Let A_n denote the tree with n vertices connected in a straight line (Figure 2.3). We label the vertices 1 to n from left to right.

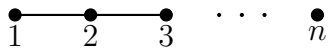


Figure 2.3: The A_n Tree

We'll examine the Hasse diagrams for $\mathcal{P}_{A_2}$ and $\mathcal{P}_{A_3}$. The notation used in this section to denote a correspondence is as follows: We list the vertices contained in S, grouping the fibers of q in parentheses. So $\mathcal{P}_{A_2}$ has four elements, represented by the following correspondences:

$$(1) = (\star \xleftarrow{\sim} \{1\} \hookrightarrow A_2)$$
$$(2) = (\star \xleftarrow{\sim} \{2\} \hookrightarrow A_2)$$
$$(12) = (\star \twoheadleftarrow A_2 \xrightarrow{\sim} A_2)$$
$$(1)(2) = (A_2 \xleftarrow{\sim} A_2 \xrightarrow{\sim} A_2)$$

And has Hasse diagram illustrated in Figure 2.4.

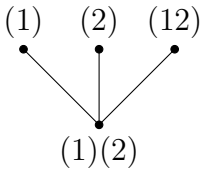


Figure 2.4: Hasse diagram for $\mathcal{P}_{A_2}$

$\mathcal{P}_{A_3}$ has eleven elements, including, for example:

$$(12)(3) = (A_2 \overset{q}{\twoheadleftarrow} A_3 \xrightarrow{\sim} A_3)$$
$$(23) = (\star \twoheadleftarrow \{2,3\} \hookrightarrow A_3)$$

Hasse diagram is illustrated in Figure 2.5.

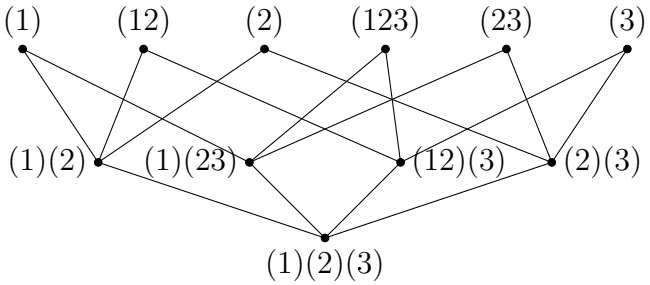


Figure 2.5: Hasse diagram for $\mathcal{P}_{A_3}$

Writing $[n] = \{1, 2, \ldots, n\}$, we observe that $\mathcal{P}_{A_2}$ is equivalent to the order poset of $\{S \subseteq [3] \mid |S| \leq 1\}$ and $\mathcal{P}_{A_3}$ is equivalent to the order poset of $\{S \subseteq [4] \mid |S| \leq 2\}$. This pattern generalizes. We have $K_{A_n} = E(A_n) \sqcup V_t(A_n)$ has $n+1$ elements. Then:

Lemma 2.4.1. *The image of $\mathcal{P}_{A_n}$ in $\mathscr{P}(K_{A_n})$ is $\{X \subseteq K_{A_n} \mid |X| \leq n-1\}$.*

Proof. Given a subset X of K_{A_n}, the proof of Lemma 2.3.2 suggests how to construct $\mathfrak{p}$: We have $v \in V(S)$ if and only v cannot be connected to a terminal vertex which is in X by a path, all of whose edges are in X. Then q contracts all the edges in $V(S)$ which are in X.

This will only fail if the subtree S determined by this procedure is empty. This occurs when X contains both the terminal vertices and all but one edge, or when X contains at least one one terminal vertex and all edges. These cases occur exactly when $|X| = n$ or $n+1$. □

Another way to phrase this result is that $\mathcal{P}_{A_n} \backslash \{\hat{0}\}$ is the face poset of the $n-1-$skeleton of a $n+1-$simplex (see [4]). This leads to the visualizations of $\mathcal{P}_{A_2}$ and $\mathcal{P}_{A_3}$ in Figure 2.6.

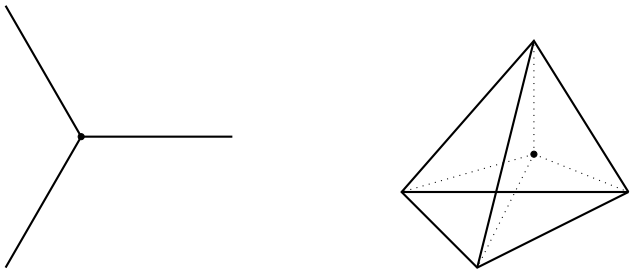


Figure 2.6: Illustrations of $\mathcal{P}_{A_2}$ (left) and $\mathcal{P}_{A_3}$ (right)

Another useful visualization of $\mathcal{P}_{A_3}$ is as a plane with two half-planes glued along the lines $x = 0$ and $y = 0$, see Figure 2.7.

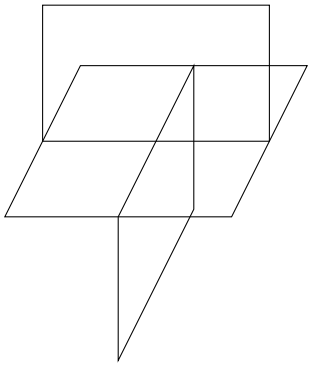


Figure 2.7: Another Visualization of $\mathcal{P}_{A_3}$

Remark 2.4.1. *Observe that $Aut(\mathcal{P}(A_n)) \cong \mathfrak{S}(K_{A_n})$. In general we will have more automorphisms of the poset $\mathcal{P}_T$ than of the tree T- this is investigated more thoroughly in sections 4-6.*

Description of $\mathcal{P}_{S_n}$

The S_n tree is a single vertex connected to $n-1$ other vertices (Figure 2.8).

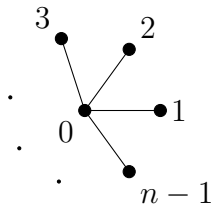


Figure 2.8: The S_n Tree

We take $n \geq 4$. The set K_{S_n} has size $|K| = 2(n-1)$: there are $n-1$ edges, and $n-1$ terminal vertices. For a terminal vertex v, let $e(v)$ denote the edge connected to v. We define an equivalence relation $\sim$ on K_{S_n} whose equivalence classes are $\{v, e(v)\}$.

Lemma 2.4.2. *The image of $\mathcal{P}_{S_n}$ in $\mathscr{P}(K_{S_n})$ is:*

$\{X \subseteq K_{S_n} \mid X^c$ *intersects each equivalence class of* $\sim$*, or is a single equivalence class of* $\sim\}$

Proof. It is straightforward to describe all possible correspondences $\mathfrak{p} = (R \overset{q}{\twoheadleftarrow} S \overset{i}{\hookrightarrow} T)$. There are two cases: if S does not contain the central vertex, then it must be a single terminal vertex v. Then we see $X = \{v, e(v)\}^c$, i.e. X^c is a single equivalence class of $\sim$.

If S does contain the central vertex, then S also contains some (possibly empty) subset $W \subset V_t(T)$. Then q is formed by contracting edges of the form $e(w)$ for $w \in W_0 \subset W$. Then we see $X = \{V_t(T) \backslash W\} \sqcup \{e(w) \mid w \in W_0\}$- in other words, X is arbitrary with the restriction that it intersects each equivalence class at most once, i.e. X^c insersects each equivalence class of $\sim$.

□

From the description, we see that $\mathrm{Aut}(\mathcal{P}(S_n))$ consists of $\varphi \in \mathfrak{S}(K)$ which descend to an automorphism of $K/\sim$. Explicitly, $\mathrm{Aut}(\mathcal{P}_{S_n}) \cong (\mathbb{Z}/2\mathbb{Z})^{n-1} \rtimes \mathfrak{S}_{n-1}$.

Analogously to the illustration in Figure 2.7, we can picture $\mathcal{P}_{S_n}$ as $\mathbb{R}^{n-1}$ with half-spaces glued along the hyperplanes $x_i = 0$.

Chapter 3

Cyclic Structures

3.1 Pre-Cyclic Structures

The notion of a cyclic structure is combinatorial data meant to capture the local geometry of a Lagrangian embedding of an arboreal singularity. This is discussed more thoroughly in Chapter 8, here we provide a brief non-rigorous explanation of the geometric motivation behind the definition.

The C-Arb singularity $\mathcal{P}_T \backslash \{\hat{0}\}$ is the face poset of a compact regular cell complex (see [4] for definitions), as proven in [20]. Therefore $\mathcal{P}_T$ is the strata poset of the cone over a compact regular cell complex- the resulting stratified space is the arboreal singularity L_T. sectionIf $x \in \mathcal{P}_T$ is a codimension-one cell, we have $N(x) \cong \mathcal{P}_{A_2}$. In L_T, this corresponds to the fact that codimension-one strata have neighborhoods homeomorphic to $\mathbb{R}^{n-1} \times \text{Cone}(\{a, b, c\})$ (refer to the Hasse diagram- Figure 2.4). This is illustrated in Figure 3.1.

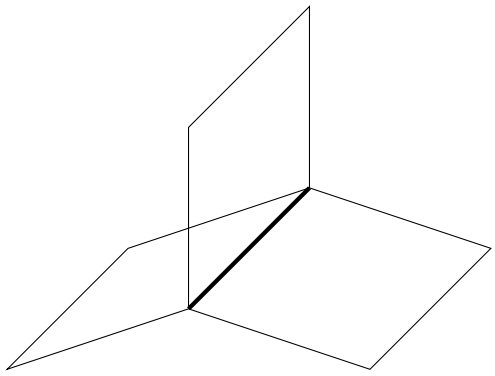


Figure 3.1: The Geometry of the Neighborhood of a Codimension-One Cell

If we have an embedding of $L = \mathbb{R}^{n-1} \times \text{Cone}(\{a, b, c\})$ into a $2n-$dimensional symplec-

tic manifold (M, ω) such that each strata is smoothly embedded and the top-dimensional strata are Lagrangian submanifolds, then under some extra assumptions, we can obtain an embedding of L into the symplectic normal bundle of the 'central' strata $\mathbb{R}^{n-1}$. This is a 2-dimensional vector bundle which has a symplectic form, i.e. an orientation, and so induces a cyclic order on the set $\{a, b, c\}$. (See appendix A for a discussion on cyclic orders on sets).

All this leads to the following definition:

Definition 3.1.1. *Let $\mathcal{P}$ be a C-Arb space. A **pre-cyclic structure** $\mathcal{O}$ on $\mathcal{P}$ is a choice, for each codimension-one cell $x \in \mathcal{P}$, a cyclic order $\mathcal{O}^x$ on the three top-cells containing x.*

There are no compatibility conditions in this definition, hence if a C-Arb space has k codimension-one cells, there are 2^k possible pre-cyclic structures.

Definition 3.1.2. *If $i : \mathcal{P} \hookrightarrow \mathcal{P}'$ is an open embedding of C-Arb spaces, and $\mathcal{O}$ is a pre-cyclic structure on $\mathcal{P}'$, let $i^*\mathcal{O}$ denote the **pullback** of $\mathcal{O}$, where if $x \in C^1(\mathcal{P})$ and $a, b, c > x$ we set:*

$$(i^*\mathcal{O})^x(a, b, c) \Leftrightarrow \mathcal{O}^{i(x)}(i(a), i(b), i(c))$$

A cyclic structure is a pre-cyclic structure satisfying certain compatibility conditions, which we will explain in the next section using directed tree.

3.2 Directed Trees

Definition 3.2.1. *A **directed tree** is a pair (T, μ), where T is a tree and μ is an **orientation** on T. Explicitly, μ is the data, for each edge $e \in E(T)$, a head $h(e)$ and a tail $t(e)$ in $V(T)$, such that $e = \{h(e), t(e)\}$. We say e "points from $t(e)$ to $h(e)$". The category **DirTree** has objects directed trees and morphisms correspondences of directed tree. A **correspondence of directed trees** is a correspondence*

$$\mathfrak{p} = ((R, \lambda) \overset{q}{\twoheadleftarrow} (S, \nu) \overset{i}{\hookrightarrow} (T, \mu))$$

Where each map preserves orientations. To be more explicit for q: whenever two vertices $v_1, v_2 \in V(R)$ are connected by an edge, there is a unique pair of vertices $w_1, w_2 \in V(S)$ connected by an edge such that $q(w_1) = v_1$ and $q(w_2) = v_2$. If that edge in S points from w_1 to w_2, then the edge in R points from v_1 to v_2.

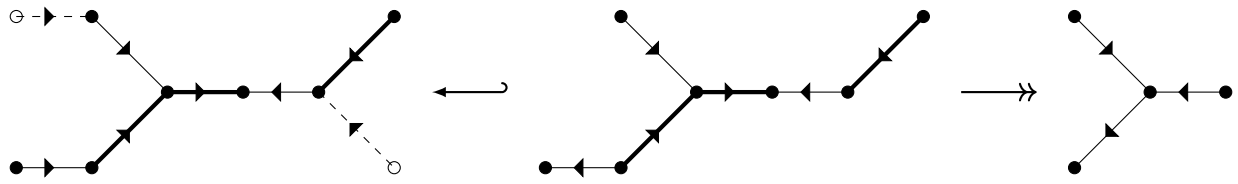


Figure 3.2: A Directed Correspondence

Figure 3.2 illustrates a directed version of the correspondence in Figure 2.1 in Chapter 2. We also let $(\star, \cdot)$ denote the directed tree with a single vertex and the trivial orientation.

The orientation μ of T determines the orientations ν, λ on S, R respectively. If $\mathfrak{p} : R \to T$ is a correspondence and μ is an orientation of T, we let $\mathfrak{p}^*\mu$ denote the induced orientation on R.

Another way of saying that is the forgetful functor $F : \mathbf{DirTree} \to \mathbf{Tree}$ is a discrete fibration:

Definition 3.2.2. *Let $F : \mathcal{D} \to \mathcal{C}$ be a functor. F is a **discrete fibration** if for every $c \in \mathcal{C}$, $d \in \mathcal{D}$, and $g : c \to F(d)$, there is a unique morphism f s.t. $F(f) = g$.*

This notion is useful for the following reason:

Lemma 3.2.1. *Suppose $F : \mathcal{D} \to \mathcal{C}$ is a discrete fibration. If $\mathcal{C}$ is a combinatorially graded category, then the functor $|\cdot| \circ F$ makes $\mathcal{D}$ a combinatorially graded category. Additionally, for any object $A \in Ob(\mathcal{D})$, F induces an isomorphism of slice categories $\mathcal{D}_A \to \mathcal{C}_{F(A)}$. Hence we get a canonical isomorphism $\mathcal{P}_A \xrightarrow{\sim} \mathcal{P}_{F(A)}$.*

Proof. The existance/uniqueness of lifts automatically implies that F induces an isomorphism of slice categories. $\mathcal{C}$ being a combinatorially graded category is a local property, in the sense that all relevant properties can be stated as properties of its slice categories, hence $\mathcal{D}$ is a combinatorially graded category as well. □

The poset $\mathcal{P}_{(T,\mu)}$ has objects which are directed correspondences $\mathfrak{p} : (R, \lambda) \to (T, \mu)$ up to equivalences. By Lemma 3.1, $\mathcal{P}_{(T,\mu)}$ is canonically isomorphic to $\mathcal{P}_T$.

Definition 3.2.3. *Given a directed tree (T, μ), there is a functorial pre-cyclic structure $\mathcal{O}_\mu$ on $\mathcal{P}_T \cong \mathcal{P}_{(T,\mu)}$. Namely, every codimension-one cell is represented by a unique correspondence of the form $\mathfrak{p} : (A_2, \lambda_0) \to (T, \mu)$, where λ_0 is the orientation on A_2 in which the edge*

points from vertex 1 *to vertex* 2. *The cyclic order* $\mathcal{O}_\mu^{[\mathfrak{p}]}$ *on the three top-cells containing* $[\mathfrak{p}]$ *is given by the successor function:*[1]

$$s([\mathfrak{p} \circ \langle 1 \rangle]) = [\mathfrak{p} \circ \langle e \rangle]$$
$$s([\mathfrak{p} \circ \langle e \rangle]) = [\mathfrak{p} \circ \langle 2 \rangle]$$
$$s([\mathfrak{p} \circ \langle 2 \rangle]) = [\mathfrak{p} \circ \langle 1 \rangle]$$

Recall (Corollary 2.3.1) that maximal cells in $\mathcal{P}_T$ are in bijection with subtrees S of T. Similarly, codimension-one cells determined by correspondences $\mathfrak{p} = (A_2 \overset{q}{\twoheadleftarrow} S \overset{i}{\hookrightarrow} T)$ are specified by the fibers $q^{-1}(1), q^{-1}(2)$- two subtrees of T separated by a single edge. The three top-cells greater than this correspondence are represented by the three subtrees $q^{-1}(1), q^{-1}(2), q^{-1}(A_2)$.

In determining the cyclic order on these top-cells, only the orientation of the edge separating the two subtrees matters. See Figure 3.3 illustrating a codimension-one cell for the directed tree in Figure 3.2.

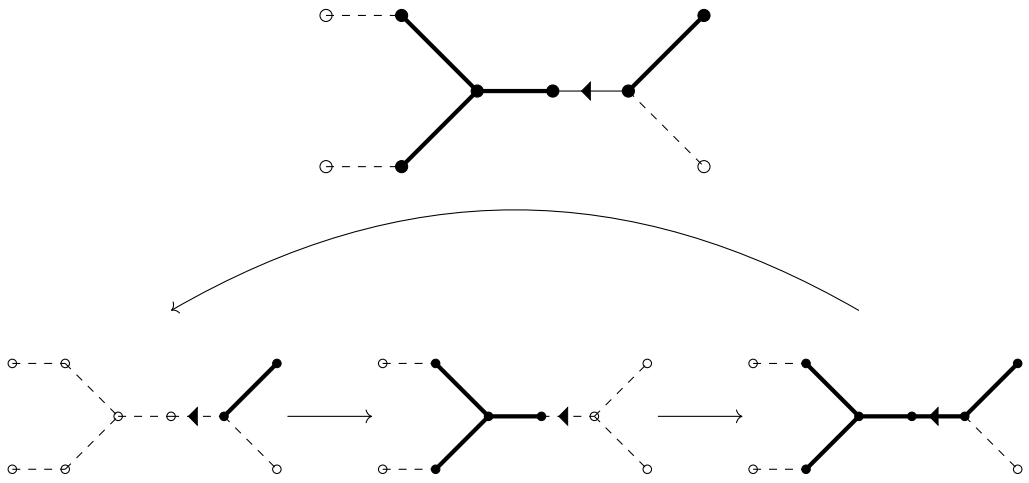


Figure 3.3: A Codimension-One Cell, and the three Top Cells in the Induced Cyclic Order

The meaning of 'functorial' in the above definition is as follows: let $i : \mathcal{P} \hookrightarrow \mathcal{P}'$ be an embedding of C-Arb singularities. Then given a pre-cyclic structure $\mathcal{O}$ on $\mathcal{P}'$, one can pull it back to obtain a pre-cyclic structure $i^*\mathcal{O}$ on $\mathcal{P}$. Then we have:

Lemma 3.2.2. *Suppose we have a correspondence of directed trees* $\mathfrak{p} : (R, \lambda) \to (T, \mu)$, *yielding an open embedding* $i_\mathfrak{p} : \mathcal{P}_R \hookrightarrow \mathcal{P}_T$. *Then* $i_\mathfrak{p}^* \mathcal{O}_\mu = \mathcal{O}_\lambda$.

[1]Refer to the notation introduced in Chapter 2

Proof. Consider a codimension-one cell represented by $\mathfrak{q} : (A_2, \lambda_0) \to (R, \lambda)$, where λ_0 is as discussed above. The cyclic order $\mathcal{O}_\mu^{i_\mathfrak{p}([\mathfrak{q}])} = \mathcal{O}_\mu^{[\mathfrak{p}\circ\mathfrak{q}]}$ is given by the successor function:

$$\begin{aligned} s([(\mathfrak{p}\circ\mathfrak{q})\circ\langle 1\rangle]) &= [(\mathfrak{p}\circ\mathfrak{q})\circ\langle e\rangle] \\ s([(\mathfrak{p}\circ\mathfrak{q})\circ\langle e\rangle]) &= [(\mathfrak{p}\circ\mathfrak{q})\circ\langle 2\rangle] \\ s([(\mathfrak{p}\circ\mathfrak{q})\circ\langle 2\rangle]) &= [(\mathfrak{p}\circ\mathfrak{q})\circ\langle 1\rangle] \end{aligned}$$

Since $[(\mathfrak{p}\circ\mathfrak{q})\circ\langle k\rangle] = i_\mathfrak{p}([\mathfrak{q}\circ\langle k\rangle])$ for $k \in K_{A_2}$, the cyclic order $i_\mathfrak{p}^*\mathcal{O}_\mu^{[\mathfrak{q}]}$ is given by the successor function:

$$\begin{aligned} s([\mathfrak{q}\circ\langle 1\rangle]) &= [\mathfrak{q}\circ\langle e\rangle] \\ s([\mathfrak{q}\circ\langle e\rangle]) &= [\mathfrak{q}\circ\langle 2\rangle] \\ s([\mathfrak{q}\circ\langle 2\rangle]) &= [\mathfrak{q}\circ\langle 1\rangle] \end{aligned}$$

Which is equal to $\mathcal{O}_\lambda^{[\mathfrak{q}]}$. □

This leads to the following definitions:

Definition 3.2.4. *Let $\mathcal{P}$ be a C-Arb space, and $\mathcal{O}$ a pre-cyclic structure on $\mathcal{P}$. $\mathcal{O}$ is a **cyclic structure** if, for all $x \in \mathcal{P}$, there exists a directed tree (T, μ) and an isomorphism $\varphi : N(x) \xrightarrow{\sim} \mathcal{P}_T$ such that $\varphi^*\mathcal{O}_\mu = i^*\mathcal{O}$, where $i : N(x) \hookrightarrow \mathcal{P}$ is the inclusion.*

Definition 3.2.5. *A **cyclic C-Arb singularity** (resp. cyclic C-Arb space) is a pair $(\mathcal{P}, \mathcal{O})$, where $\mathcal{P}$ is a C-Arb singularity (resp. space) and $\mathcal{O}$ is a cyclic structure on $\mathcal{P}$.*

Lemma 3.2.3.

(i) If $i : \mathcal{P} \hookrightarrow \mathcal{P}'$ is an open embedding of C-Arb spaces and $\mathcal{O}$ is a cyclic structure on $\mathcal{P}'$, $i^*\mathcal{O}$ is a cyclic structure on $\mathcal{P}$.

(ii) If (T, μ) is a directed tree, $\mathcal{O}_\mu$ is a cyclic structure on $\mathcal{P}_T$.

Proof. (i) Is clear from the local nature of the definition of cyclic structures. (ii) Follows from Lemma 3.2.2. □

Remark 3.2.1. *It is **not** the case (in general) that every cyclic structure on $\mathcal{P}_T$ is of the form $\mathcal{O}_\mu$, where μ is an orientation of T. Instead, every cyclic structure is of the form $\varphi^*\mathcal{O}_\mu$, where μ is an orientation of T and $\varphi \in Aut(\mathcal{P}_T)$. This follows trivially from the definition, once we show that there cannot exist an isomorphism $\mathcal{P}_T \xrightarrow{\sim} \mathcal{P}_{T'}$ when $T \ncong T'$, which we show in section 4.3.*

3.3 Coherence

Definition 3.2.4 is simple but somewhat unilluminating. In fact, cyclic structures can also be defined as pre-cyclic structures satisfying two local 'coherence conditions'. We provide a partial discussion of this point in the rest of this section, and a proof appears in Chapter 4.

Definition 3.3.1. *Fix a tree R. $\mathcal{O}$ is **R-coherent** if for every embedding $i : \mathcal{P}_R \hookrightarrow \mathcal{P}$, $i^*\mathcal{O}$ is a cyclic structure.*

It is clear from the definition that a cyclic structure $\mathcal{O}$ on a C-Arb space $\mathcal{P}$ is $R-$coherent for every tree R. Our main result regarding coherence is the following:

Theorem 3.3.1. *A pre-cyclic structure is a cyclic structure if and only if it is A_3-coherent and S_4-coherent.*

Recall that A_3 is the tree consisting of three vertices connected in a line, while S_4 is the tree with three vertices all connected to a central vertex.

In the following two sections, we will undergo a careful investigation of cyclic structures on $\mathcal{P}_{A_3}$ and $\mathcal{P}_{S_4}$, and we will also look at A_3-coherent pre-cyclic structures on $\mathcal{P}_{S_4}$ which are not cyclic structures. In the following sections we will examine cyclic structures on $\mathcal{P}_{A_n}$ and $\mathcal{P}_{S_n}$.

A complete proof of Theorem 3.3.1 will involve techniques developed in Chapter 4.

3.4 Cyclic Structures on $\mathcal{P}_{A_3}$

In this section we will carefully describe cyclic structures on $\mathcal{P}_{A_3}$, and give both a geometric and combinatorial way of understanding these objects.

To begin, write $K = K_{A_3} = V_t(A_3) \sqcup E(A_3) = \{1, e_1, e_2, 3\}$. Here $1, 3$ are terminal vertices, and e_1, e_2 the internal edges. For simplicity of notation we will write $1 = e_0$ and $3 = e_3$. Identify $\mathcal{P}_{A_3}$ with its image in $\mathscr{P}(K)$, which is $\{S \subseteq K \mid |S| \leq 2\}$.

If c is a cyclic order on K, there is a corresponding pre-cyclic structure $\mathcal{O}_c$ on $\mathcal{P}_{A_3}$ as follows:

Definition 3.4.1. *(Definition of $\mathcal{O}_c$) A codimension-one cell is given by $S \subseteq K$ with $|S| = 1$. If we write $K \backslash S = \{x, y, z\}$, the three top-cells containing S are $S \cup \{x\}, S \cup \{y\}, S \cup \{z\}$. We define:*

$$\mathcal{O}_c^S\{S \cup \{x\}, S \cup \{y\}, S \cup \{z\}\} \Leftrightarrow c(x, y, z)$$

.

We have:

Proposition 3.4.1. *The construction $c \mapsto \mathcal{O}_c$ gives a bijection between cyclic orders on K and cyclic structures on $\mathcal{P}_{A_3}$.*

Proof. If $\varphi : \mathcal{P}_{A_3} \xrightarrow{\sim} \mathcal{P}_T$ is a poset isomorphism, then $T \cong A_3$- this follows from Corollary 2.3.2. Hence any cyclic structure is of the form $\varphi^*\mathcal{O}_\mu$, where μ is an orientation of A_3 and $\varphi \in \mathrm{Aut}(\mathcal{P}_{A_3})$.

Observe that the association $c \mapsto \mathcal{O}_c$ is functorial, in the sense that if $\varphi \in \mathrm{Aut}(\mathcal{P}_{A_3})$ induces $\overline{\varphi} \in \mathfrak{S}(K_{A_3})$, then $\varphi^*\mathcal{O}_c = \mathcal{O}_{\overline{\varphi}^* c}$. (The pullback $\overline{\varphi}^* c$ is the cyclic order defined by $(\overline{\varphi}^* c)(x, y, z) \Leftrightarrow c(\varphi(x), \varphi(y), \varphi(z)))$.

An easy computation (the results of which are shown in Table 3.1 later in this section) verifies that for an orientation μ, there exists a cyclic order c on K with $\mathcal{O}_\mu = \mathcal{O}_c$. Then $\varphi^*\mathcal{O}_\mu = \mathcal{O}_{\overline{\varphi}^* c}$, so each cyclic structure corresponds to a cyclic order.

On the other hand, $\mathrm{Aut}(\mathcal{P}_{A_3}) \cong \mathfrak{S}(K)$ acts transitively on cyclic orders on K, so if one cyclic order yields a cyclic structure, all of them do. Hence each cyclic order corresponds to a cyclic structure. □

As a result, we see there are $6 = 3!$ distinct cyclic structures on $\mathcal{P}_{A_3}$.

We now present a visual way of diagramming pre-cyclic structures on $\mathcal{P}_{A_3}$. Recall from Chapter 1 that $\mathcal{P}_{A_3}$ can be visualized as a plane with two half-planes glued on- see Figure 2.7 Removing the top-cells $\{e_0, e_1\}$ and $\{e_2, e_3\}$ from $\mathcal{P}_{A_3}$, we can represent the remaining cells as a stratification of $\mathbb{R}^2$, illustrated in Figure 3.4:

In the figure, each of the four one-cells are pictured, along with two of the top-cells containing it. Suppose we have a pre-cyclic structure $\mathcal{O}$ on $\mathcal{P}_{A_3}$. For a one-cell $\{e\}$, let $\{e, x\}$, $\{e, y\}$ denote the top cells containing $\{e\}$ represented in the figure. One of these top-cells must be the successor of the other in the cyclic order $\mathcal{O}^{\{e\}}$. We can represent this by drawing an arrow from the predecessor to its successor, crossing $\{e\}$.

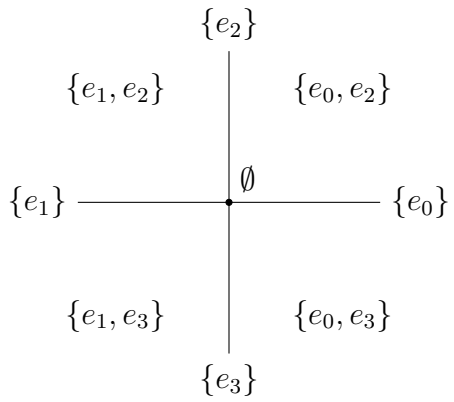


Figure 3.4: An Illustration of $\mathcal{P}_{A_3} \backslash \{\{e_0, e_1\}, \{e_2, e_3\}\}$.

For example, suppose the cyclic order $\mathcal{O}^{\{e_0\}}$ is induced by the total order $\{e_0, e_1\} < \{e_0, e_2\} < \{e_0, e_3\}$. Then, since $\{e_0, e_3\}$ is the successor of $\{e_0, e_2\}$, we draw an arrow pointing from the top-right quadrant to the bottom-right quadrant, crossing $\{e_0\}$. A completed diagram might look like Figure 3.5.

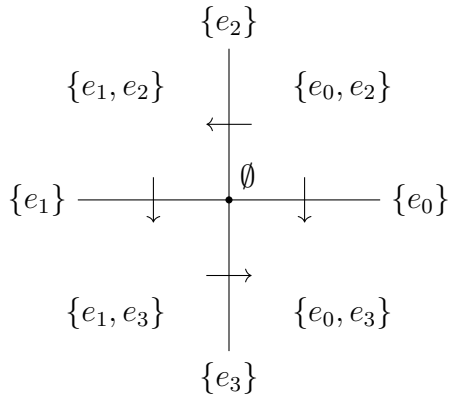


Figure 3.5: A Completed Arrow Diagram

One can check that the corresponding pre-cyclic structure is *not* a cyclic structure, by observing that it does not arise from a cyclic order.

Table 3.1 provides a translation between the three ways we have seen of describing cyclic structures on $\mathcal{P}_{A_3}$: Using an orientation on A_3, using a cyclic order on K, and using an 'arrow diagram'. Note that there are only four orientations of A_3 while there are six cyclic orders of K.

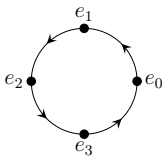
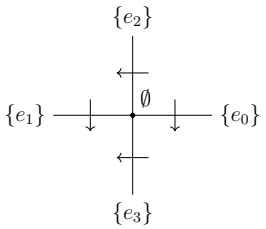
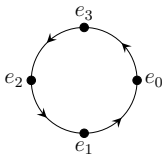
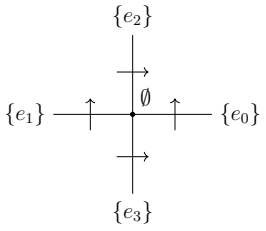
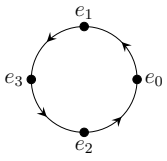
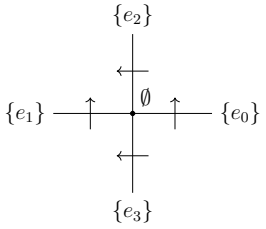
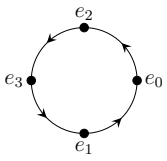
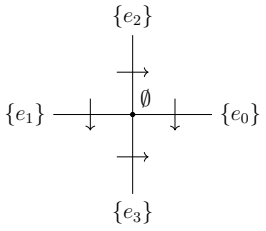
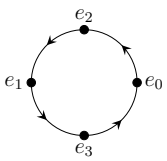
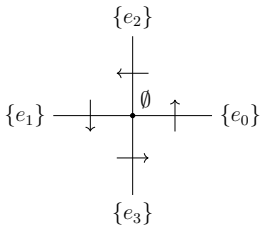
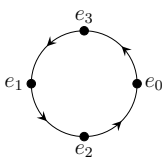
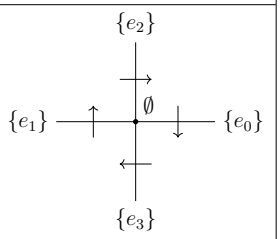
Directed Tree	Cyclic Order	Arrow Diagram
		
		
		
		
N/A		
N/A		

Table 3.1: A Dictionary for Cyclic Structures on $\mathcal{P}_{A_3}$

Definition 3.4.2. *An arrow diagram for $\mathcal{P}_{A_3}$ is **coherent** if arrows on opposite one-cells point in the same direction (the first four in the table), and **spiral** if the arrows form a spiral.*

We see an arrow diagram represents a cyclic structure if and only if it is coherent or spiral. However, as the automorphism group acts transitively on cyclic structures, the distinction between coherent and spiral structures involves some symmetry breaking. As a result, this distinction between is not really fundamental. In particular, while we drew this arrow diagram by removing cells $\{e_0, e_1\}$ and $\{e_2, e_3\}$, we could have also removed cells $\{e_0, e_2\}$ and $\{e_1, e_3\}$, or cells $\{e_0, e_3\}$ and $\{e_1, e_2\}$. For each of these presentations, four of the cyclic structures will yield coherent arrow diagrams, and the other two will yield spiral diagrams, but which ones yield which is altered.

To be more explicit: The spiral structures occur when the top-cells removed correspond to pairs of one-cells which are opposites in the cyclic order.

We will use understand this idea to cyclic structures on $\mathcal{P}_{S_4}$ via arrow diagrams, as well as A_3-coherent pre-cyclic structures on $\mathcal{P}_{S_4}$ which are not cyclic structures.

3.5 Cyclic Structures on $\mathcal{P}_{S_4}$

Now, consider $\mathcal{P}_{S_4}$. Write $K = K_{S_4} = E(S_4) \sqcup V_t(S_4) = \{e_1, e_2, e_3, v_1, v_2, v_3\}$. The image of $\mathcal{P}_{S_4}$ in $\mathscr{P}(K)$ is $\{S \subseteq K \mid \{v_i, e_i\} \nsubseteq S, 1 \leq i \leq 3\} \cup \{\{v_i, e_i\}^c \mid 1 \leq i \leq 3\}$. For the rest of this section we identify $\mathcal{P}_{S_4}$ with its image in $\mathscr{P}(K)$.

Definition 3.5.1. *A **polarization** is a function $\rho : K \to \{+, -\}$ such that $\rho(v_i) = -\rho(e_i)$ for $i = 1, 2, 3$.*

Definition 3.5.2. *Given a polarization ρ, the pre-cyclic structure $\mathcal{O}_\rho$ is defined as follows: A codimension-one cell is given by a set X with $|X| = 2$, such that there is a unique i where $\{e_i, v_i\} \cap X = \emptyset$. The three top-cells containing X are $X \cup \{e_i\}$, $X \cup \{v_i\}$, and $\{e_i, v_i\}^c$. If $\rho(e_i) = +$, the cyclic order is induced by the total order:*

$$X \cup \{v_i\} < X \cup \{e_i\} < \{e_i, v_i\}^c$$

Otherwise, it is the opposite.

Codimension-one cells can be also visualized as pairs of subtrees separated by a single edge. When $T = S_3$, one of these subtrees is a single terminal vertex $\{v\}$, while the other subtree, call it S contains the central vertex 0. The three top-cells greater than this codimension-one cell are indexed by the subtrees $\{v\}, S, S \cup \{v\}$. If the original correspondence is described by the set X, then the subtree S yields the set $X \cup \{v_i\}$, $S \cup \{v\}$ yields the set $X \cup \{e_i\}$, and $\{v\}$ the set $\{e_i, v_i\}^c$.

As an example, when $S = \{0, v_2\}$ and $v = v_1$, the translation between these two descriptions of codimension-one cells is illustrated in Figure 3.6. In the figure, $[\mathfrak{p}]$ is the codimension-one cell, and $[\mathfrak{p}_i]$, $1 \leq i \leq 3$ are the three top-cells greater than it.

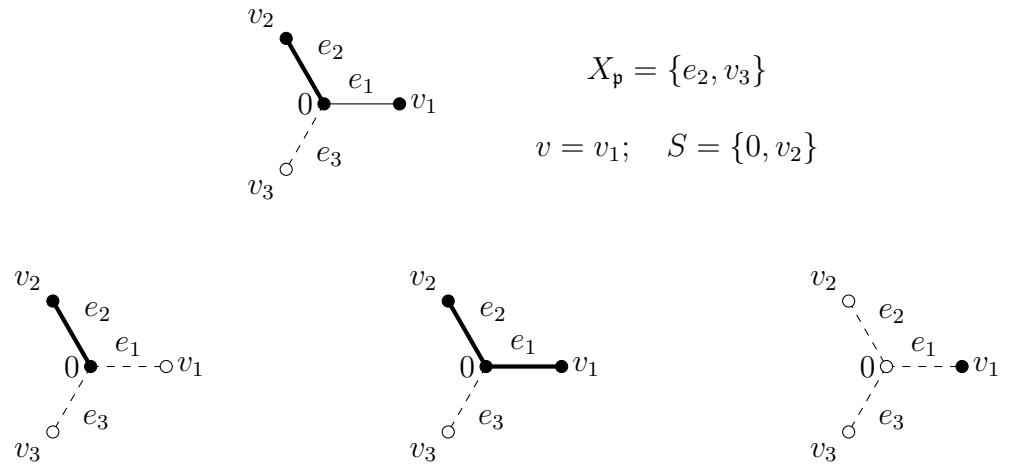


Figure 3.6: Translating Between Two Descriptions of Codimension-One Cells in $\mathcal{P}_{S_4}$.

Definition 3.5.3. *For an orientation μ of S_3, we can define a polarization $\rho = \rho(\mu)$ where $\rho(e_i) = +$ if edge e_i points toward vertex 0, and $-$ otherwise. One can easily check that $\mathcal{O}_\mu = \mathcal{O}_\rho$.*

Proposition 3.5.1. *The association $\rho \mapsto \mathcal{O}_\rho$ gives a bijection between polarizations and cyclic structures on $\mathcal{O}_{S_4}$. Therefore, does the association $\mu \mapsto \mathcal{O}_\mu$, where μ is an orientation of S_4.*

Proof. Suppose $\varphi : \mathcal{P}_{S_4} \xrightarrow{\sim} \mathcal{P}_T$ is an isomorphism. By Corollary 2.3.2, T must have 4 vertices. By observation, $\mathcal{P}_{S_4}$ and $\mathcal{P}_{A_4}$ are not isomorphic (one has five one-cells and the other has

six, for example), hence $T \cong S_4$. Hence any cyclic structure is of the form $\varphi^*\mathcal{O}_\mu$, where μ is an orientation of S_4 and $\varphi \in \mathrm{Aut}(\mathcal{P}_{S_4})$.

The association $\rho \mapsto \mathcal{O}_\rho$ is functorial, meaning that if $\varphi \in \mathrm{Aut}(\mathcal{P}_{S_4})$ induces $\overline{\varphi} \in \mathfrak{S}(K_{S_4})$, $\varphi^*\mathcal{O}_\rho = \mathcal{O}_{\overline{\varphi}^*\rho}$. Here $(\overline{\varphi}^*\rho)(x) = \rho(\varphi(x))$, for $x \in K$.

As explained in definition 3.5.3, for any orientation μ, there is a polarization $\rho = \rho(\mu)$ with $\mathcal{O}_\mu = \mathcal{O}_\rho$. Then $\varphi^*\mathcal{O}_\mu = \mathcal{O}_{\overline{\varphi}^*\rho}$, so each cyclic structure corresponds to a polarization.

On the other hand, for any polarization ρ there exists μ with $\mathcal{O}_\rho = \mathcal{O}_\mu$, so each polarization corresponds to a cyclic structure. □

We can present pre-cyclic structures on $\mathcal{P}_{S_4}$ using arrow diagrams as well. If we remove the three top-cells $\{v_i, e_i\}^c$, the remaining cells can be visualized as a stratification of $\mathbb{R}^3$ sliced by the three coordinate planes $\{x_i = 0\}, i = 1, 2, 3$. Identify the one-cell corresponding to the positive x_i-axis with $\{v_i\}$ and the negative x_i-axis with $\{e_i\}$. This is illustrated in Figure 3.7.

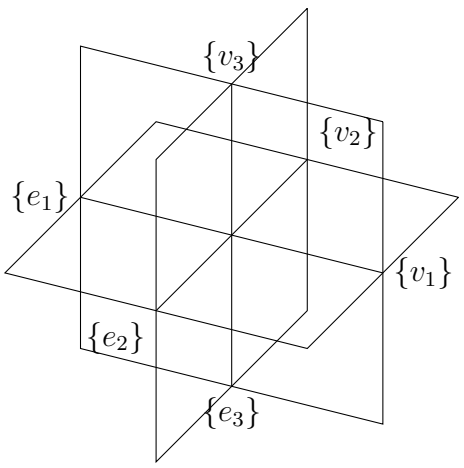


Figure 3.7: An Illustration of $\mathcal{P}_{S_4}\backslash\{v_i, e_i\}^c$.

As in $\mathcal{P}_{A_3}$, we can specify an arbitrary pre-cyclic structure on $\mathcal{P}_{S_4}$ by drawing an arrow crossing each of the twelve two-cells in the picture. Such a pre-cyclic structure is A_3-coherent if its restriction to each neighborhood isomorphic to $\mathcal{P}_{A_3}$ is a cyclic structure. There are six such neighborhoods, each corresponding to a one-cell, i.e. a positive or negative axis. Restricting to the neighborhood of a one-cell yields a geometric picture of a $\mathcal{P}_{A_3}$ as in the previous section. So we see that, in order for a pre-cyclic structure to be A_3-coherent, the arrows on the four two-cells attached to a one-cell must be oriented *coherently* or *spirally*

(see Definition 3.4.2). This is illustrated in Figure 3.8- the one-cell in question is bold, and two possible options for the arrows around that one-cell are shown:

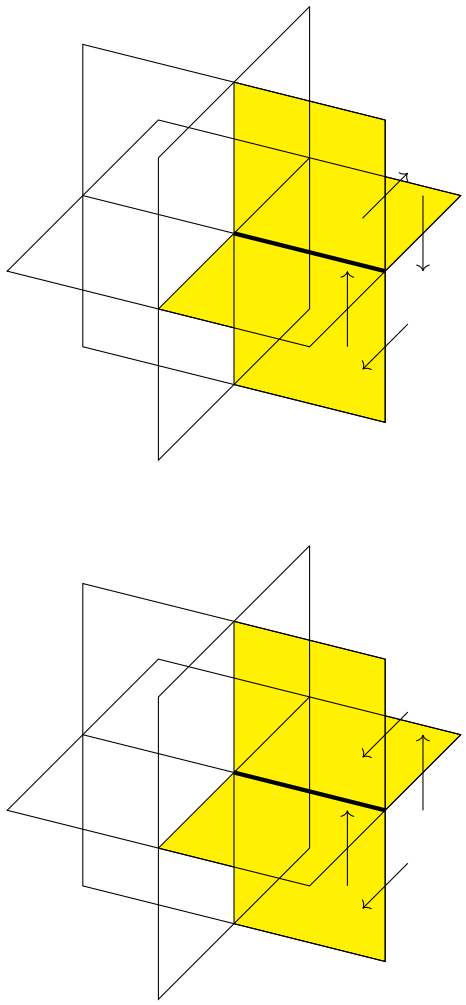


Figure 3.8: Constructing Pinwheel Orientations

One possibility is that the arrows on each hyperplane are oriented the same way (i.e. all pointing toward the same half-space). There are $8 = 2^3$ such pre-cyclic structures, and one can check that these are exactly the cyclic structures. Explicitly, we can write this structure as $\mathcal{O}_\mu$ or as $\mathcal{O}_\rho$, where the following are equivalent:

- The arrows on the hyperplane $\{x_i = 0\}$ all point toward the half-space $x_i > 0$ (i.e. toward $\{v_i\}$)
- The edge e_i is pointing away from the center in μ.

- $\rho(v_i) = +$.

However, this does not exhaust all possibilities for A_3-coherent pre-cyclic structures. To give an example: Fix an axis, and draw arrows on the eight two-cells attached to that axis so that they point cyclically about that axis (Figure 3.9).

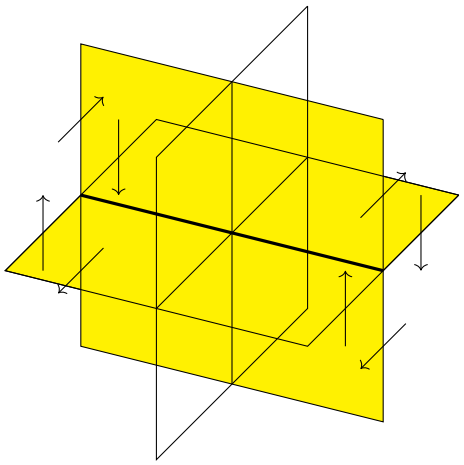


Figure 3.9: Constructing Pinwheel Orientations

For the remaining hyperplane, orient the arrows pointing toward the same half-space. There are $3 \cdot 2 \cdot 2 = 12$ of these *pinwheel orientations*. One can easily show that these are all the possible A_3-coherent pre-cyclic structures which are not coherent. This shows that the S_4-coherence condition is necessary.

3.6 Cyclic Structures on A_n Singularities

A_n singularities and cyclic structures on them have a beautiful description in terms of cyclic orders on finite sets. This cyclic symmetry also appears in the representation theory of A_n, as observed in [21]. We will investigate this connection more thoroughly later.

Recall that from 2.4.1, $\mathcal{P}_{A_n} \cong \{X \in \mathscr{P}(K_{A_n}) \mid |X| \leq n-1\}$. Throughout this section we identify these posets. Note that the codimension of a subset X is $|X^c| - 2$.

Definition 3.6.1. *Let Δ_{inj} denote the category of finite sets whose morphisms are injective maps, and $\boldsymbol{Arb}_A$ the full subcategory of **Pos** whose objects are $\mathcal{P}_{A_n}$, $n \geq 2$.*

Definition 3.6.2. *Define the functor* $\Delta : \boldsymbol{Arb}_A \to \Delta_{inj}$ *as follows:*

$$\Delta(A_n) = K_{A_n}$$

For an embedding $i : \mathcal{P}_{A_m} \hookrightarrow \mathcal{P}_{A_n}$*, since open embeddings preserve codimension, the image of the minimal point* $i(\emptyset) = X$ *is a subset of* K_{A_n} *of size* $n - m$*. We will define a functorial injection:*

$$\Delta(i) : K_{A_m} \hookrightarrow K_{A_n}$$

Whose image is X^c*, defined by:*

$$i(\{k\}) = X \cup \{\Delta(i)(k)\}$$

More generally, for $Y \subseteq K_{A_n}$, we have:

$$i(Y) = X \cup \Delta(i)(Y)$$

Definition 3.6.3. *Let* S *be a finite set. A* ***pre-cyclic order*** *on* S *is a ternary relation* c *satisfying all the properties of a cyclic order except transitivity.*

See appendix A for the axioms of cyclic orders. An equivalent way of thinking of pre-cyclic orders is that they are a choice of cyclic order on each 3-element subset of S- without the transitivity axiom there is no relationship among these choices.

Definition 3.6.4. *Let* $\boldsymbol{Arb}_A^{pre-cyc}$ *denote the category whose objects are pairs* $(\mathcal{P}_{A_n}, \mathcal{O})$*, where* $\mathcal{O}$ *is a pre-cyclic structure on* $\mathcal{P}_{A_n}$*, and whose morphisms* $(\mathcal{P}_{A_m}, \mathcal{O}) \to (\mathcal{P}_{A_n}, \mathcal{O}')$ *are embeddings* $i : \mathcal{P}_{A_m} \hookrightarrow \mathcal{P}_{A_n}$ *such that* $i^*\mathcal{O}' = \mathcal{O}$*. Let* $\boldsymbol{Arb}_A^{cyc}$ *denote the full subcategory of* $\boldsymbol{Arb}_A^{pre-cyc}$ *where* $\mathcal{O}$ *is a cyclic structure.*

Similarly, let $\Delta_{inj}^{pre-cyc}$ *denote the category whose objects are finite sets equipped with a pre-cyclic order and whose morphisms are injections preserving the ternary relation. Let* Δ_{inj}^{cyc} *denote the full subcategory for which the pre-cyclic order is a cyclic order.*

Given a pre-cyclic structure $\mathcal{O}$ on $\mathcal{P}_{A_n}$, one can construct a pre-cyclic order $c = \Delta(\mathcal{O})$ on K_{A_n} as follows: Each codimension-one cell in $\mathcal{P}_{A_n}$ is represented by $X \subseteq K_{A_n}$ with $|X| = n-2$. Write $K_{A_n} \backslash X = \{x, y, z\}$. The three top-cells containing X are $X \cup \{x\}, X \cup \{y\}$, and $X \cup \{z\}$. Then we write:

$$\mathcal{O}^X(X \cup \{x\}, X \cup \{y\}, X \cup \{z\}) \Leftrightarrow c(x, y, z)$$

It is clear that this construction yields a bijection between pre-cyclic structures on $\mathcal{P}_{A_n}$ and pre-cyclic orders in K_{A_n}.

Lemma 3.6.1. *The construction described above is functorial, i.e., it allows us to extend the functor* $\Delta : \boldsymbol{Arb}_A \to \Delta_{inj}$ *to a functor* $\Delta : \boldsymbol{Arb}_A^{pre-cyc} \to \Delta_{inj}^{pre-cyc}$.

Proof. Let $i : \mathcal{P}_{A_m} \hookrightarrow \mathcal{P}_{A_n}$, $\mathcal{O}$ be a pre-cyclic structure on $\mathcal{P}_{A_n}$, and $c = \Delta(\mathcal{O})$. We want to show $\Delta(i)^* c = \Delta(i^*\mathcal{O})$. So let $x, y, z \in K_{A_m}$, $X = K_{A_m} \backslash \{x, y, z\}$. Then $i(X) = K_{A_n} \backslash \{\Delta(i)(x), \Delta(i)(y), \Delta(i)(z)\}$, so:

$$\begin{aligned}(\Delta(i)^* c)(x, y, z) &\Leftrightarrow c(\Delta(i)(x), \Delta(i)(y), \Delta(i)(z)) \\ &\Leftrightarrow \mathcal{O}^{i(X)}(i(X) \cup \{\Delta(i)(x)\}, i(X) \cup \{\Delta(i)(y)\}, i(X) \cup \{\Delta(i)(z)\}) \\ &\Leftrightarrow (i^*\mathcal{O})^X(X \cup \{x\}, X \cup \{y\}, X \cup \{z\}) \\ &\Leftrightarrow \Delta(i^*\mathcal{O})(x, y, z)\end{aligned}$$

□

We are now able to prove the key proposition of this section:

Proposition 3.6.1. *Let* $\mathcal{O}$ *be a pre-cyclic structure on* $\mathcal{P}_{A_n}$. *The following are equivalent:*

1. $\mathcal{O}$ *is a cyclic structure.*
2. $\mathcal{O}$ *is* A_3-*coherent.*
3. $\Delta(\mathcal{O})$ *is a cyclic order on* K_{A_n}.

Proof. (1) ⇒ (2) is clear.

To show (2) ⇒ (3), we exploit the functor $\Delta : \mathbf{Arb}_A^{pre-cyc} \to \Delta_{inj}^{pre-cyc}$. Let $c = \Delta(\mathcal{O})$. If Y is a 4-element subset of K_{A_n}, we can find an embedding $i : \mathcal{P}_{A_3} \hookrightarrow \mathcal{P}_{A_n}$ such that $i(\emptyset) = Y^c$, meaning the image of $\Delta(i)$ is Y.

A_3 coherence implies the pullback $i^*\mathcal{O}$ is a cyclic structure. By functoriality, $\Delta(i^*\mathcal{O}) = \Delta(i)^* c$. Proposition 3.4.1 then implies $\Delta(i)^* c$ is a cyclic order. It follows that c is a cyclic order when restricted to any four-element subset of K_{A_n}, which is enough to ensure transitivity, so c is a cyclic order.

To show (3) ⇒ (1), suppose $\Delta(\mathcal{O}) = c$ is a cyclic order. If μ is any orientation on A_n, let $c_\mu = \Delta(\mathcal{O}_\mu)$. By the implication (1) ⇒ (3) already proved, c_μ is a cyclic order, hence there exists $\varphi \in \mathrm{Aut}(\mathcal{P}_{A_n})$ such that $\overline{\varphi}^* c_\mu = c$. (Recall $\overline{\varphi}$ is the permutation of K_T induced by an automorphism of $\mathcal{P}_T$-, which in this case equals $\Delta(\varphi)$). By the functoriality of Δ, $\mathcal{O} = \varphi^* \mathcal{O}_\mu$ is a cyclic structure.

□

To finish this section we will more explicitly describe the relationship between cyclic orders and orientations of A_n.

Proposition 3.6.2. *Let μ be an orientation on A_n. By proposition 4.4, $\Delta(\mathcal{O}_\mu) := c_\mu$ is a cyclic order on K_{A_n}. c_μ can be described using a successor function s as follows:*

- *If $v \in V_t(A_n)$, begin travelling from v toward the opposite endpoint. $s(v)$ is the first edge encountered pointing in the direction of travel. If no such edge is encountered, $s(v)$ is the opposite endpoint.*
- *If $e \in E(A_n)$, begin travelling toward an endpoint of A_n in the direction indicated by e. $s(e)$ is the next edge encountered pointing in the direction of travel. If no such edge is encountered before hitting an endpoint, $s(e)$ is the endpoint.*

In particular, write $K_{A_n} = \{e_0, e_1, \ldots, e_{n-1}, e_n\}$, where e_i denotes the edge joining vertex i to vertex $i+1$ for $1 \leq i < n$, e_0 is the vertex 1, and e_n is the vertex n. If μ is the orientation in which all edges point from vertex i to vertex $i+1$, $\Delta(\mathcal{O}_\mu)$ is the cyclic order induced by the total order $e_0 < e_1 \cdots e_{n-1} < e_n$.

Proof. In a cyclic order c, y is the successor of x if and only if $c(x, y, z)$ whenever $z \neq x, y$.

We use $\langle 1\rangle$, $\langle 2\rangle$, $\langle e\rangle$ to denote the one-cells of $\mathcal{P}_{A_2}$.

Suppose $e_i \in K_{A_n}$ with $0 \leq i < n$, and suppose either $i = 0$ or e_i is pointing to the right. (The other case can be handled by symmetry). Let j is the smallest index greater than i of an edge e_j pointing from vertex j to vertex $j+1$ (or $j = n$ if no such vertex does), we will show $s(e_i) = e_j$. Let $\mathfrak{p} = (A_2 \overset{q}{\twoheadleftarrow} S \overset{i}{\hookrightarrow} T)$ be a correspondence with $X_\mathfrak{p} = K_{A_n} \backslash \{e_i, e_j, e_k\}$. Write $\lambda = \mathfrak{p}^*\mu$. We can assume WLOG that the subtree $q^{-1}(\{1\})$ has vertices of smaller index than $q^{-1}(\{2\})$. There are three cases:

Case 1: If $k < i < j$, then $X_{\mathfrak{p}\circ\langle 1\rangle} = X \cup \{e_k\}$, $X_{\mathfrak{p}\circ\langle e\rangle} = X \cup \{e_i\}$, $X_{\mathfrak{p}\circ\langle 2\rangle} = X \cup \{e_j\}$, and the edge points from vertex 1 to vertex 2 in λ.

Case 2: If $i < k < j$, then $X_{\mathfrak{p}\circ\langle 1\rangle} = X \cup \{e_i\}$, $X_{\mathfrak{p}\circ\langle e\rangle} = X \cup \{e_k\}$, $X_{\mathfrak{p}\circ\langle 2\rangle} = X \cup \{e_j\}$, and the edge points from vertex 2 to vertex 1 in λ.

Case 3: If $i < j < k$, then $X_{\mathfrak{p}\circ\langle 1\rangle} = X \cup \{e_i\}$, $X_{\mathfrak{p}\circ\langle e\rangle} = X \cup \{e_j\}$, $X_{\mathfrak{p}\circ\langle 2\rangle} = X \cup \{e_k\}$, and the edge points from vertex 1 to vertex 2 in λ.

Verification that $c(e_i, e_j, e_k)$ in each case, using the definitions of $\mathcal{O}_\mu$ (definition 3.2.3) and $\Delta(\mathcal{O}_\mu)$, is straightforward and left to the reader. □

Example 3.6.1. *As an example, consider the directed A_5 pictured in Figure 3.10, with the edges/terminal vertices labelled as elements of K_{A_5}:*

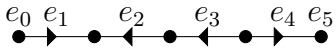


Figure 3.10: A Directed Tree (A_5, μ).

The successor of e_0 is e_1, since it points to the right. The successor of e_1 is e_4, which is the next edge that points to the right. The successor of e_4 is e_5, since there are no more edges pointing to the right. Continuing this, we get the cyclic order in Figure 3.11.

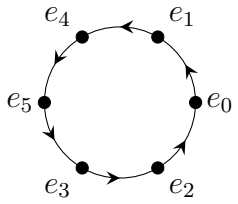


Figure 3.11: The Cyclic Order $\Delta(\mathcal{O}_\mu)$ on K_{A_5}.

3.7 Cyclic Structures on S_n Singularities

In this section we explain cyclic structures on $\mathcal{P}_{S_n}$. The exposition is analogous to the previous section.

We identify $\mathcal{P}_{S_n}$ with its image in $\mathscr{P}(K_{S_n})$, which is (lemma 2.4.2):

$\{X \subseteq K_{S_n} \mid X^c$ intersects each equivalence class of $\sim$, or is a single equivalence class of $\sim\}$

When X^c intersects each equivalence class of $\sim$, the codimension of X is the number of equivalence classes completely contained in X^c. Hence for an open embedding $i : S_m \hookrightarrow S_n$, the image $i(\emptyset) = X$ is a set such that X^c contains $m - 1$ equivalence classses in K_{S_n}. That inspires the following definition:

Definition 3.7.1. *Let $\square_{inj}$ denote the category whose objects are finite sets equipped with a relation $\sim$ whose equivalence classes are of size 2. A morphism $S \to T$ is a pair (i, X) where $i : S \hookrightarrow T$ is an injective map of sets satisfying $i(x) \sim i(y)$ when $x \sim y$, and $X \subseteq T \backslash S$ intersects each equivalence class of $T \backslash S$ exactly once. The composition $(i, X) \circ (j, Y) = (i \circ j, X \cup j(Y))$.*

Definition 3.7.2. *Let $\boldsymbol{Arb}_S$ denote the full subcategory of $\boldsymbol{Pos}$ whose objects are $\mathcal{P}_{S_n}$, $n \geq 4$.*

Definition 3.7.3. *Define a functor $\square : \boldsymbol{Arb}_S \to \square_{inj}$ as follows. $\square(\mathcal{P}_{S_n}) = K_{S_n}$. Let $i : \mathcal{P}_{S_m} \hookrightarrow \mathcal{P}_{S_n}$ be an embedding. If $i(\emptyset) = X$, then X^c contains $m-1$ equivalence classes in K_{S_n}, and intersects the rest once. We will construct:*

$$\square(i) = (j, Y)$$

Where the image of j is the $m-1$ equivalence classes contained in X^c, and $Y = X^c \backslash j(K_{S_m})$- it is defined by:

$$i(\{k\}) = X \cup \{j(k)\}$$

Definition 3.7.4. *Let $S \in Ob(\square_{inj})$. A **weak polarization** ρ is a function which assigns to every $x \in S$ and every set $Y \subset S$ which contains exactly one element of each equivalence class besides $[x]$, a sign $\rho_Y(x) \in \{+, -\}$. Furthermore if $x \sim y$ but $x \neq y$ we require $\rho_Y(x) = -\rho_Y(y)$. A **polarization** is a weak polarization which does not depend on the argument Y.*

We let $\square_{inj}^{w-pol}$ and $\square_{inj}^{pol}$ denote the categories consisting of pairs (S, ρ) where $S \in Ob(\square_{inj})$ and ρ is a weak polarization (resp. polarization). Note that when $(i, X) : S \to T$ is a morphism and ρ a weak polarization on T, the pullback $i^*\rho$ is defined by $i^*\rho_Y(v) = \rho_{Y \cup X}(i(v))$.

We let $\boldsymbol{Arb}_S^{pre-cyc}$ and $\boldsymbol{Arb}_S^{cyc}$ denote the category of pairs $(\mathcal{P}_{S_n}, \mathcal{O})$, where $\mathcal{O}$ is a pre-cyclic structure (resp. cyclic structure).

Given a pre-cyclic structure $\mathcal{O}$ on $\mathcal{P}_{S_n}$, we can construct a weak polarization $\square(\mathcal{O})$ as follows: Each codimension-one cell in $\mathcal{P}_{S_n}$ is given by a subset X of K_{S_n} which intersects each equivalence class of K_{S_n} once, except for one. Call the remaining class $\{x, y\}$. The three top-cells containing X are given by $X \cup \{x\}$, $X \cup \{y\}$, and $\{x, y\}^c$. (See the discussion of $\mathcal{P}_{S_4}$, including Figure 3.6). If:

$$\mathcal{O}^X(X \cup \{x\}, X \cup \{y\}, \{x, y\}^c)$$

We set $\square(\mathcal{O})_X(x) = -$, $\square(\mathcal{O})_X(y) = +$.

It is clear from this discussion that this establishes a bijection between weak polarizations and pre-cyclic structures. In fact, we have:

Lemma 3.7.1. *The above construction gives a functor* $\square : \boldsymbol{Arb}_S^{pre-cyc} \to \square_{inj}^{w-pol}$.

Proof. Let $\mathcal{O}$ be a pre-cyclic structure on $\mathcal{P}_{S_n}$ and $i : \mathcal{P}_{S_m} \hookrightarrow \mathcal{P}_{S_n}$ be an embedding, with $\square(i) = (j, Y)$. We want to show $\square(i^*\mathcal{O}) = j^*\square(\mathcal{O})$. Let X be a subset of K_{S_m} which intersects each equivalence class of K_{S_n} once, except for one: Call it $\{x, y\}$. Then:

$$\begin{aligned} j^*\square(\mathcal{O})_X(x) = - &\Leftrightarrow \square(\mathcal{O})_{X\cup Y}(j(x)) = - \\ &\Leftrightarrow \mathcal{O}^{j(X)\cup Y}(j(X)\cup Y\cup\{j(x)\}, j(X)\cup Y\cup\{j(y)\}, \{j(x), j(y)\}^c) \\ &\Leftrightarrow (i^*\mathcal{O})^X(X\cup\{x\}, X\cup\{y\}, \{x, y\}^c) \\ &\Leftrightarrow \square(i^*\mathcal{O})_X(x) = - \end{aligned}$$

□

Note that this convention agrees with our previous convention for polarizations on $\mathcal{P}_{S_4}$. In particular, we know from proposition 3.5.1 that a pre-cyclic structure $\mathcal{O}$ on $\mathcal{P}_{S_4}$ is a cyclic structure if and only if $\square(\mathcal{O})$ is a polarization. This idea generalizes to give a description of cyclic structures on $\mathcal{P}_{S_n}$, as the following proposition shows:

Proposition 3.7.1. *Let $\mathcal{O}$ be a pre-cyclic structure on $\mathcal{P}_{S_n}$, where $n \geq 4$. The following are equivalent:*

(1) $\mathcal{O}$ is a cyclic structure.

(2) $\mathcal{O}$ is S_4-coherent.

(3) $\square(\mathcal{O})$ is a polarization on K_{S_n}.

(4) $\mathcal{O} = \mathcal{O}_\mu$ for some orientation μ of S_n.

Proof. (1) $\Rightarrow$ (2) is clear.

For (2) $\Rightarrow$ (3), let $\rho = \square(\mathcal{O})$, where $\mathcal{O}$ is a S_4-coherent pre-cyclic structure on $\mathcal{P}_{S_n}$. Pick an equivalence class $\{a, b\} \subseteq K_{S_n}$. Let X and Y be two subsets of K_{S_n} which intersect each equivalence class of K_{S_n} once except $\{a, b\}$- we want to show $\rho_X(a) = \rho_Y(a)$. Assume that X and Y differ on a single equivalence class, that is, that there is a class $\{x, y\}$ such that $x \in X$ and $y \in Y$, but $X\backslash\{x\} = Y\backslash\{y\}$. Pick a third equivalence class $\{u, v\}$, with $u \in X, Y$.

We get a map $(j, Z) : K_{S_3} \to K_{S_n}$ in $\square_{inj}$, where j sends $v_1, e_1, v_2, e_2, v_3, e_3$ to a, b, x, y, u, v, respectively, and $Z = X\backslash\{x, u\} = Y\backslash\{y, u\}$. This map defines $i : \mathcal{P}_{S_4} \hookrightarrow \mathcal{P}_{S_n}$, with $\square(i) = (j, Z)$. Since $i^*\mathcal{O}$ is a cyclic structure, by functoriality $j^*\rho = \square(i^*\mathcal{O})$ is a polarization. So:

$$\rho_X(a) = j^*\rho_{\{v_2,v_3\}}(v_1) = j^*\rho_{\{e_2,v_3\}}(v_1) = \rho_Y(a)$$

We assumed X and Y differed only on a single equivalence class, but we can get between *any* two sets by repeatedly making changes on a single equivalence class.

For (3) $\Rightarrow$ (4), let ρ be a polarization. We define μ to be the orientation where the edge e_i points toward vertex 0 if $\rho(e_i) = +$, and toward vertex v_i if $\rho(e_i) = -$. It is easy to check that $\mathcal{O}_\rho = \mathcal{O}_\mu$.

Finally, (4) $\Rightarrow$ (1) follows immediately from the definition. $\square$

Remark 3.7.1. *An interesting result is that, in this case, all of the cyclic structures come from orientations. This is true in general when there are no vertices of degree* 2 *in the tree (we will not explicitly prove this, but it will follow in a straightforward way from general combinatorial descriptions of cyclic structures in the following section).*

Chapter 4

The Combinatorics of Cyclic Structures

4.1 Isomorphisms of C-Arb Singularities

In this section, we will prove a handful of combinatorial results involving arboreal singularities. An important component of this is the analysis of $\mathrm{Aut}(\mathcal{P}_T)$, the automorphism group of a C-Arb singularity, as well as isomorphisms $\mathcal{P}_T \xrightarrow{\sim} \mathcal{P}_{T'}$. In this second case it turns out we must have $T \cong T'$, but this is not immediately obvious.

Our first goal will be understanding how combinatorial properties of $\mathcal{P}_T$ relate to properties of the underlying tree T.

To begin, let T be a tree and $K_T = E(T) \sqcup V_t(T)$. As $\mathcal{P}_T$ is atomistic, for $\varphi \in \mathrm{Aut}(\mathcal{P}_T)$ we can associate $\overline{\varphi} \in \mathfrak{S}(K_T)$. We use notation from Chapter 2.

Definition 4.1.1. *Let $x, y \in K_T$. A **path** connecting x and y is a path $\{v_0, v_1, \ldots, v_k\}$ such that $v_0 \in x$ if $x \in E(T)$ or $v_0 = x$ if $x \in V_t(T)$, and likewise for v_k and y.*

Lemma 4.1.1. *Let T be a tree, and $x \neq y \in K_T$. Then there exists a correspondence $\mathfrak{p} : R \to T$ with $X_{\mathfrak{p}} = \{x, y\}^c$ if and only if x and y can be connected by a path where each vertex has degree ≤ 2.*

Proof. If there is such a path between x and y, let S be the subtree defined by that path. Then if $\mathfrak{p} = (\star \twoheadleftarrow S \hookrightarrow T)$, $X_{\mathfrak{p}} = \{x, y\}^c$. On the other hand, suppose there is a vertex v between x and y having degree > 2, and let $\mathfrak{p} = (R \twoheadleftarrow S \hookrightarrow T)$, with $x, y \notin X_{\mathfrak{p}}$. Then x and y are either edges having a vertex in common with S or vertices in S, hence $v \in V(S)$.

Since $\deg(v) > 2$, S has a boundary vertex which is not in (or equal to) x or y, and hence there is at least one more element in $X_{\mathfrak{p}}^c$. □

Definition 4.1.2. *Let $x, y \in K_T$. Define $x \sim y$ if $x = y$ or there exists $\mathfrak{p} : R \to T$ such that $X_{\mathfrak{p}} = \{x, y\}^c$. By the lemma, $\sim$ is an equivalence relation, we call the equivalence classes **chains**. An element $\varphi \in Aut(\mathcal{P}_T)$ **stabilizes chains** if $x \sim \overline{\varphi}(x)$ for all $x \in K_T$.*

Chains are essentially the connected components of T once the vertices of degree > 2 are removed. This is illustrated in Figure 4.1. We see that for each chain C there are at most two vertices with degree > 2 that are incident to an edge in C- such vertices are called **boundary vertices** of C.

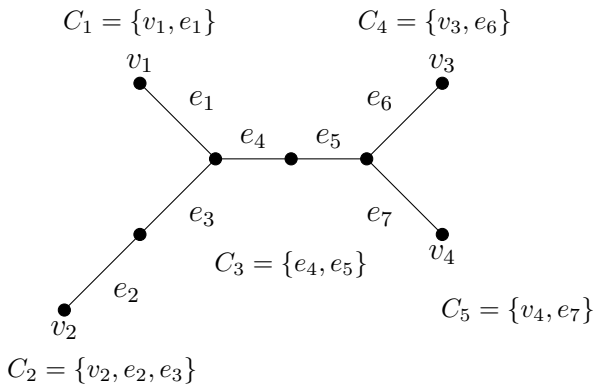


Figure 4.1: A tree T and the chains in K_T

The below definitions involve only the combinatorics of $\mathcal{P}_T$. We will see (lemma 4.1.2) that each of these notions has an interpretation in terms of the tree T (the names are meant to be suggestive).

Definition 4.1.3.

- *For a tree T, let $\mathcal{C}_T$ denote the set of chains.*
- *If $C \in \mathcal{C}_T$, we define $\sim_C^0$ to be the relation on $\mathcal{C}_T \backslash \{C\}$ where $C_1 \sim_C^0 C_2$ if and only if there exists a maximal correspondence $\mathfrak{p} : \star \to T$ such that $(X_{\mathfrak{p}})^c$ intersects each of C_1, C_2, C. We let $\sim_C$ denote the equivalence relation on $\mathcal{C}_T \backslash \{C\}$ generated by $\sim_C^0$.*
- *If $C, C_1, C_2 \in \mathcal{C}_T$, C **separates** C_1 and C_2 if C, C_1, C_2 are distinct and $C_1 \not\sim_C C_2$.*

- *Distinct chains $C_1, C_2 \in \mathcal{C}_T$ are **adjacent** if they are not separated by any chain C.*
- *A **node** is a maximal collection of pairwise adjacent chains. A **boundary node** of a chain is a node containing that chain.*
- *For a node $N \subset \mathcal{C}_T$, the **node relation** $\sim_N$ on $\mathcal{C}_T$ is defined by $C_1 \sim_N C_2$ if $C_1 = C_2$ or if C_1, C_2 are distinct, not both in N, and are not separated by any $C \in N$.*
- *A correspondence $\mathfrak{p}$ is **disjoint** from a node N if the set of chains intersecting $(X_p)^c$ is contained in a single equivalence class of $\sim_N$. $\mathfrak{p}$ **intersects** N if it is not disjoint from N.*

This terminology, which depends entirely on the combinatorics of the poset, is related to the combinatorics of the tree T as follows:

Lemma 4.1.2.

(i) C separates C_1 and C_2 iff any path in T connecting $x \in C_1$ to $y \in C_2$ contains both vertices of an edge in C. Equivalently, $C_1 \sim_C C_2$ iff C_1 and C_2 are on the same side of C.

(ii) C_1 and C_2 adjacent iff there is vertex of degree > 3 incident to an edge in each of them.

(iii) If v is a vertex of degree greater than 2, the subset N_v of chains containing an edge incident to v is a node, and this describes all nodes in $\mathscr{P}(\mathcal{C}_T)$.

(iv) For a node N_v, we can describe the node relation: $C_1 \sim_{N_v} C_2$ iff some $x \in C_1$ can be connected to some $y \in C_2$ by a path avoiding v.

(v) A correspondence $\mathfrak{p} = (R \overset{q}{\twoheadleftarrow} S \overset{i}{\hookrightarrow} T)$ intersects a node N_v iff $v \in V(S)$. In this case, $(X_\mathfrak{p})^c$ intersects at least one chain in each equivalence class of N_v.

Proof. (i)

Suppose $\mathfrak{p} = (\star \twoheadleftarrow S \hookrightarrow T)$ satisfies $(X_\mathfrak{p})^c \cap C \neq \emptyset$, and C has two boundary vertices. Then S can only contain one of the boundary vertices. This means that if $C_1 \sim_C^0 C_2$ then C_1, C_2 are on the same side of C. Since 'on the same side of' is an equivalence relation, we are done if we can show that $C_1 \sim_C C_2$ whenever C_1 and C_2 are on the same side of C.

For a subtree S of T, we say S is *a witness* to the relation $C_1 \sim_C C_2$ if, letting $\mathfrak{p} = (\star \twoheadleftarrow S \hookrightarrow T)$, $(X_\mathfrak{p})^c$ intersects each of C_1, C_2, C.

First assume that C_1, C_2 are on the same side of C, and C_1 is not in between C and C_2, nor C_2 in between C and C_1. Then one can find a subtree S such that each of C, C_1, C_2 contains an edge incident to only one vertex in S, so this S is a witness to the relation $C_1 \sim_C^0 C_2$.

Suppose on the other hand that C_1 is between C and C_2. Letting $v \in V(T)$ be the boundary vertex of C on the same side as C_1, C_2, we can find an edge e incident to v which is not between C and C_1. If we write C_3 for the chain containing e, by the argument in the previous paragraph we see $C_1 \sim_C^0 C_3$ and $C_2 \sim_C^0 C_3$, hence $C_1 \sim_C C_2$.

Figure 4.2 for an illustrates this argument.

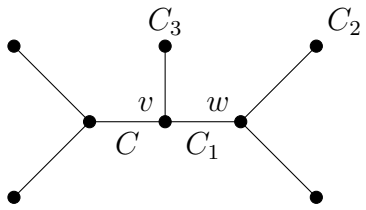


Figure 4.2: An illustration of the relation $\sim_C$: The subtree $\{v\}$ is a witness to the relation $C_1 \sim_C^0 C_3$, while $\{v, w\}$ is a witness to the relation $C_2 \sim_C^0 C_3$

$(ii), (iii)$

Both of these are clear from the definitions and part (i).

(iv)

If C_1 and C_2 are separated by a chain $C \in N_v$, then any path connecting C_1 to C_2 must pass through C, and hence through v. On the other hand, if neither $C_1, C_2 \in N_v$ and they are not separated by a chain in N_v, the shortest path connecting any $x \in C_1$ to $y \in C_2$ passes through at most one boundary vertex of a chain in N_v, and hence not through v.

If $C_1 \in N_v$ but $C_2 \notin N_v$, and they are not separated by a chain in N_v, then C_1 has a boundary vertex which is not equal to v which can be joined to C_2 by a path avoiding v.

(v)

If $v \in V(S)$, then each component of N_v has a chain which either has a boundary vertex of T which is also a boundary vertex of S, or has an edge in T only one of whose vertices are in S. So we see $(X_{\mathfrak{p}})^c$ intersects at least one chain in each equivalence class of N_v.

On the other hand, if $v \notin V(S)$, then S is completely disjoint from all the chains in all but one of the equivalence classes of N_v, hence $\mathfrak{p}$ is disjoint from N_v.

□

The conclusion from the above is that we can reconstruct the tree from the poset structure. To be more precise, let T, T' be trees, and $\sigma : T \xrightarrow{\sim} T'$ be a tree isomorphism. We can view σ as an isomorphism in **Tree**, namely $(T \overset{\sigma^{-1}}{\twoheadleftarrow} T' \overset{i}{\hookrightarrow} T')$. One has an induced isomorphism $i_\sigma : \mathcal{P}_T \xrightarrow{\sim} \mathcal{P}_{T'}$. Then:

Proposition 4.1.1. *Let $\varphi : \mathcal{P}_T \xrightarrow{\sim} \mathcal{P}_{T'}$ be an isomorphism of C-Arb singularities. Then there exists $\sigma : T \xrightarrow{\sim} T'$ such that $i_\sigma^{-1} \circ \varphi \in Aut(\mathcal{P}_T)$ stabilizes chains. If $T \ncong A_n$, then σ is unique.*

Proof. Since chains are defined purely in terms of the poset structure, it is clear that $\overline{\varphi}$ sends chains in K_T to chains in $K_{T'}$, i.e. determines a map $\mathcal{C}_T \to \mathcal{C}_{T'}$. If such a σ exists, we must have $\overline{i_\sigma} = \overline{\varphi}$ as maps $\mathcal{C}_T \to \mathcal{C}_{T'}$. When $T \ncong A_n$, there are no nontrivial automorphisms of T which act trivially on $\mathcal{C}_T$, meaning σ is unique. (When $T \cong A_n$, $n \geq 2$ there is only one chain, but there is a nontrivial automorphism of T).

To complete the proof, we need to show that there exists such a σ. However, by lemma 4.1.2, if $C_1, C_2 \in \mathcal{C}_T$ are chains sharing a boundary vertex, the same must be true of $\overline{\varphi}(C_1)$, $\overline{\varphi}(C_2)$, since this property can be discerned from the poset structure. Since this relationship completely determines the tree structure, the isomorphism σ must exist. □

4.2 Reflection Isomorphisms

We now tackle the converse: Suppose $\omega \in \mathfrak{S}(K_T)$ stabilizes chains. Does there necessarily exist $\varphi \in \mathrm{Aut}(\mathcal{P}_T)$ such that $\overline{\varphi} = \omega$? The answer, as we will see in this section, is yes. Our strategy will be to construct **reflection isomorphisms**, which are automorphisms of $\mathcal{P}_T$ which act as transpositions of adjacent elements in K_T.

Lemma 4.2.1. *Let $v \in V(T)$ be a vertex of degree 2, and e_1, e_2 the edges in $E(T)$ incident to v. Then there is a $\omega_v \in Aut(\mathcal{P}_T)$ such that $\overline{\omega}_v$ is the transposition $(e_1 e_2)$ switching e_1 and e_2. Additionally, if $v \in V(T)$ is a vertex of degree 1, and e is the edge incident to v, then there is a $\omega_v \in Aut(\mathcal{P}_T)$ such that $\overline{\omega}_v$ is the transposition (ve) switching v and e. These isomorphisms are called **reflection isomorphisms**.*

We will take the rest of this section to prove this lemma. It is enough to show that, for any correspondence $\mathfrak{p} : R \to T$, the set that results from taking $X_\mathfrak{p}$ and switching e_1, e_2 (or

v and e) is of the form $X_{\mathfrak{q}}$ for some other correspondence $\mathfrak{q} : R' \to T$. In our approach, we will explicitly construct the correspondence $\mathfrak{q}$. To begin, we have the following definition:

Definition 4.2.1. *A **pointed tree** is a pair (T, v), where either $(i) : v \in V(T)$ is a vertex of degree ≤ 2, or $(ii) : v = \hat{0}$, a symbol we introduce to mean the 'null vertex'. Let (T, v) be a pointed tree, and $\mathfrak{p} = (R \overset{q}{\twoheadleftarrow} S \overset{i}{\hookrightarrow} T)$ a correspondence. If $v \in V(S)$ and $q^{-1}(\{q(v)\}) = v$, write $\mathfrak{p}^* v := q(v)$, and in all other cases $\mathfrak{p}^* v = \hat{0}$. A **correspondence of pointed trees** $\mathfrak{p} : (R, w) \to (T, v)$ is a correspondence satisfying $w = \mathfrak{p}^* v$.*

For a correspondence of pointed trees $\mathfrak{p} : (R, w) \to (T, v)$, with $\mathfrak{p} = (R \overset{q}{\twoheadleftarrow} S \overset{i}{\hookrightarrow} T)$, we define the reflected correspondence $\mathcal{R}\mathfrak{p} : (R, w) \to (T, v)$ as follows. First, in the following cases, we define $\mathcal{R}\mathfrak{p} = \mathfrak{p}$:

(A1) If w, v are both in $V(T)$, or $w = v = \hat{0}$.

(A2) If $v \notin V(S)$ and none of the vertices adjacent to v are in $V(S)$[1] .

(A3) If $\deg(v) = 2$, $v \in V(S)$, and for both vertices v_1, v_2 adjacent to v, $v_i \in V(S)$ and $q(v_i) = q(v)$.

In the following cases, we modify $\mathfrak{p}$:

(B1) If $v \notin V(S)$, but there is $u \in V(S)$ adjacent to v, define $\mathcal{R}\mathfrak{p} = (R \overset{q'}{\twoheadleftarrow} S \cup \{v\} \overset{i}{\hookrightarrow} T)$, where $q' = q$ except $q'(v) = q(u)$.

(B2) If $v \in V(S)$, there is exactly one $u \in V(S)$ adjacent to v, and $q(v) = q(u)$. Then $\mathcal{R}\mathfrak{p} = (R \overset{q}{\twoheadleftarrow} S \backslash \{v\} \overset{i}{\hookrightarrow} T)$.

(B3) If $\deg(v) = 2$, suppose v is incident to $e_1 = \{v, v_1\}$ and $e_2 = \{v, v_2\}$. Suppose $v, v_1, v_2 \in V(S)$ and $q(v) = q(v_1) \neq q(v_2)$. If we write $q(v_1) = x$, $q(v_2) = y$, then $\mathcal{R}\mathfrak{p} = (R \overset{q'}{\twoheadleftarrow} S \overset{i}{\hookrightarrow} T)$, where $q' = q$ except $q'(v) = y$.

Cases B1 (Right Hand Side) and B2 (Left Hand Side) are illustrated in Figure 4.3. Case B3 is illustrated in Figure 4.4. These figures should be interpreted as in Chapter 2 (see, for example, Figure 2.1): Thick lines are those contracted by q, and dotted lines/empty circles are not contained in S.

A crucial fact is the following:

[1] $u \in V(T)$ is adjacent to v if $\{u, v\} \in E(T)$

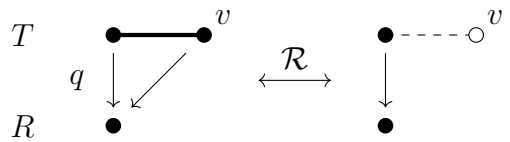


Figure 4.3: Cases B1 and B2

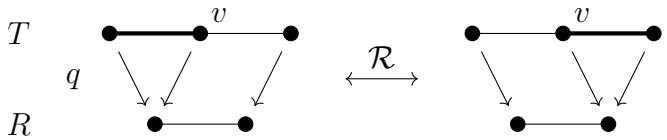
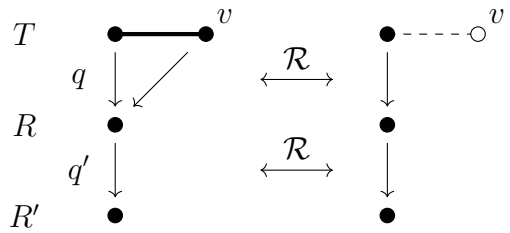


Figure 4.4: Case B3

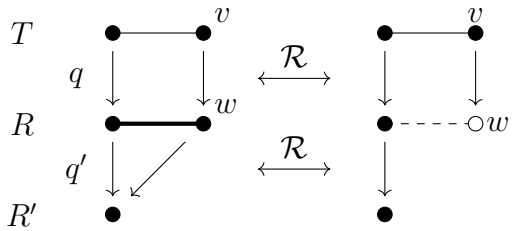
Lemma 4.2.2. *If we let* ***Tree***$_\bullet$ *denote the category of pointed trees,* $\mathcal{R}$ *defines a* functor ***Tree***$_\bullet \to$ ***Tree***$_\bullet$*, which acts at the identity on objects and satisfies* $\mathcal{R}^2 = Id$*. In particular,* $\mathcal{R}(\mathfrak{p} \circ \mathfrak{q}) = \mathcal{R}\mathfrak{p} \circ \mathcal{R}\mathfrak{q}$*.*

Proof. This proof amounts to checking a handful of cases. For each nontrivial case we illustrate $\mathfrak{p}$, $\mathfrak{q}$, $\mathcal{R}\mathfrak{p}$, and $\mathcal{R}\mathfrak{q}$ in a neighborhood of v, from which one can easily verify the identitites. Writing $\mathfrak{p} = (R \overset{q}{\twoheadleftarrow} S \overset{i}{\hookrightarrow} T)$, we can assume v, or at least one of its neighbors, is in S, otherwise $\mathcal{R}\mathfrak{p} = \mathfrak{p}$, $\mathcal{R}\mathfrak{q} = \mathfrak{q}$, and $\mathcal{R}(\mathfrak{p} \circ \mathfrak{q}) = \mathfrak{p} \circ \mathfrak{q}$. For similar reasons, if we write $\mathfrak{q} = (R' \overset{q'}{\twoheadleftarrow} S' \overset{i}{\hookrightarrow} R)$, we can assume v, or at least one of its neighbors, is in $q^{-1}(S')$.

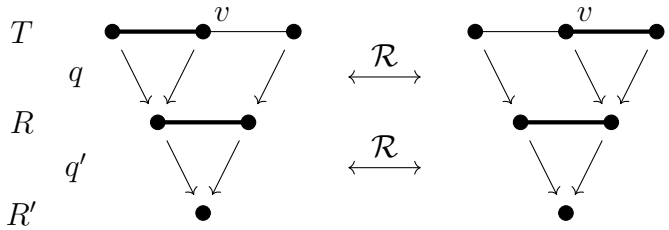
First assume v has one neighbor v_1 in S. If either $v \notin V(S)$ (RHS) or $v \in V(S)$ with $q(v) = q(v_1)$ (LHS), we have:



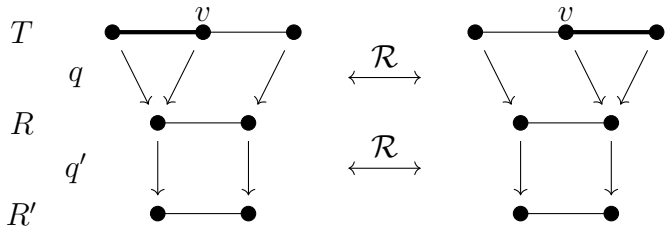
If $v \in V(S)$ with $q(v) \neq q(v_1)$, we have $q(v) = \mathfrak{p}^* v := w$. If either $w \notin V(S')$ (RHS) or $w \in V(S')$ with $q'(w) = q'(q(v_1))$ (LHS) we have:



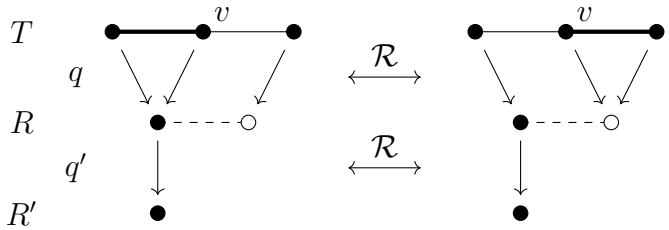
Now assume v has two neighbors, v_1 and v_2 in S. If $q(v) = q(v_1) = q(v_2)$ we see $\mathcal{R}\mathfrak{p} = \mathfrak{p}$, $\mathcal{R}\mathfrak{q} = \mathfrak{q}$, and $\mathcal{R}(\mathfrak{p} \circ \mathfrak{q}) = \mathfrak{p} \circ \mathfrak{q}$. Now suppose $q(v) = q(v_1) \neq q(v_2)$. Write $q(v_1) = x, q(v_2) = y$. If $x, y \in V(S')$, and $q'(x) = q'(y)$, we have:



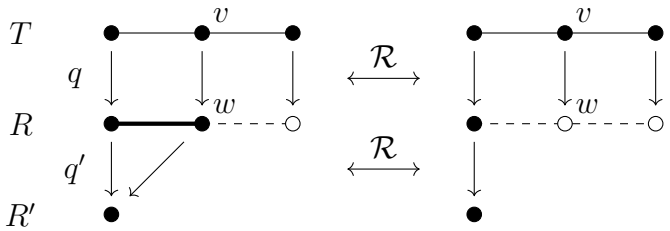
If $q'(x) \neq q'(y)$, we have:



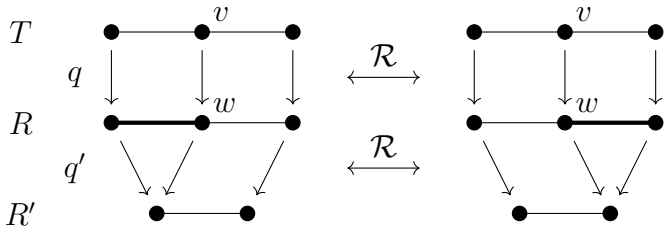
If $x \in V(S')$ but $y \notin V(S')$ (LHS) or the reverse (RHS) we have:



Finally, assume $q(v) \neq q(v_1) \neq q(v_2)$, so $\mathcal{R}\mathfrak{p} = \mathfrak{p}$. Write $w = q(v) = \mathfrak{p}^* v$, $w_1 = q(v_1)$, $w_2 = q(v_2)$. If $w, w_1 \in V(S')$ $w_2 \notin V(S')$, and $q'(w) \neq q'(w_1)$, $\mathcal{R}\mathfrak{q} = \mathfrak{q}$ and $\mathcal{R}(\mathfrak{p} \circ \mathfrak{q}) = \mathfrak{p} \circ \mathfrak{q}$. If $q'(w) = q'(w_1)$ (LHS), or only $w_1 \in V(S')$ (RHS), we have:



If $w, w_1, w_2 \in V(S')$ with $q(w) = q(w_1) \neq q(w_2)$ we have:



□

If $\mathfrak{p} : R \to T$, and $v \in V(T)$ has degree ≤ 2 or $v = \hat{0}$, $\mathfrak{p}$ lifts to a correspondence of pointed trees $(R, \mathfrak{p}^* v) \to (T, v)$. We use the notation $r_v\mathfrak{p}$ to denote the correspondence of trees $R \to T$ underlying the correspondence of pointed trees $\mathcal{R}\mathfrak{p}$. If we write $w = \mathfrak{p}^* v$, we can interpret the lemma as saying $r_v(\mathfrak{p} \circ \mathfrak{q}) = r_v\mathfrak{p} \circ r_w\mathfrak{q}$.

By the functoriality of $\mathcal{R}$, the map $\omega_v : [\mathfrak{p}] \mapsto [r_v\mathfrak{p}]$ is an automorphism of $\mathcal{P}_T$. Furthermore, it is easy to see from the description that if $\deg(v) = 2$ and e_1, e_2 are the two edges indicent to v, $\omega_v(\langle e_1 \rangle) = \langle e_2 \rangle$, $\omega_v(\langle e_2 \rangle) = \langle e_1 \rangle$, and for each other $k \in K_T$, $\omega_v(\langle k \rangle) = \langle k \rangle$, and a similar statement holds when $\deg(v) = 1$. This completes the proof of lemma 4.2.1.

Corollary 4.2.1. *Suppose $\omega \in \mathfrak{S}(K_T)$ stabilizes chains. Then there exists $\varphi \in Aut(\mathcal{P}_T)$ such that $\omega = \overline{\varphi}$, and φ is a product of reflection isomorphisms.*

Proof. Any such ω will act as a permutation on each chain. Any permutation can be generated by transpositions of adjacent elements, which can be realized by reflection isomorphisms. □

Corollary 4.2.2. *If $\varphi : \mathcal{P}_T \xrightarrow{\sim} \mathcal{P}_{T'}$ is a poset isomorphism, there exists a tree isomorphism $\sigma : T \xrightarrow{\sim} T'$ such that $i_\sigma^{-1} \circ \varphi \in Aut(\mathcal{P}_T)$ is a product of reflection isomorphisms.*

Proof. Follows immediately from proposition 4.1.1 and corollary 4.2.1. □

4.3 Isomorphisms of Cyclic C-Arb Singularities

In this section, we investigate isomorphisms of the form $\varphi : (\mathcal{P}_T, \mathcal{O}_\mu) \xrightarrow{\sim} (\mathcal{P}_{T'}, \mathcal{O}_{\mu'})$, where (T, μ) and (T', μ') are directed trees. That is, we require that $\varphi^* \mathcal{O}_{\mu'} = \mathcal{O}_\mu$.

By proposition 4.1.1, there exists an isomorphism $\sigma : T \to T'$ such that $i_\sigma^{-1} \circ \varphi$ stabilizes chains. Writing $\psi = i_\sigma^{-1} \circ \varphi$, we get $\psi : (\mathcal{P}_T, \mathcal{O}_\mu) \to (\mathcal{P}_T, \mathcal{O}_{i_\sigma^* \mu})$ is an isomorphism. Hence we can restrict our attention to the case when $T = T'$ and our isomorphism φ stabilizes chains.

As in the undirected case, we will see that any isomorphism of this kind can be written as a product of **directed reflection isomorphisms**. To explain this, we first need to extend the notion of reflection isomorphisms to the directed case.

4.4 Reflection Isomorphisms for Directed Trees

Definition 4.4.1. *Let (T, μ) be a directed tree. We let $\mathcal{V}^+(T)$, respectively $\mathcal{V}^-(T)$, denote the set consisting of all vertices $v \in V(T)$ of degree ≤ 2 which are* sources *(all incident edges point away from v), respectively* sinks *(all incident edges point toward v). Additionally, we let each set contain the null vertex $\hat{0}$.*

*A **+ pointed directed tree** is a triple (T, μ, v) with $v \in \mathcal{V}^+(T, \mu)$. A **- pointed directed tree** is a triple (T, μ, v) with $v \in \mathcal{V}^-(T, \mu)$. A **correspondence** of $\pm$ pointed directed trees is an orientation-preserving correspondence of $\pm$ pointed trees.*

We let $\mathbf{DirTree}_\bullet^+$, resp. $\mathbf{DirTree}_\bullet^-$, denote the category of + pointed directed trees (resp. - pointed directed trees).

Definition 4.4.2. *For a directed tree (T, μ), and $v \in V(T)$, define $r_v\mu$ to be the orientation on T with all arrows incident to v reversed. We also define $r_{\hat{0}}\mu = \mu$.*

Now we have:

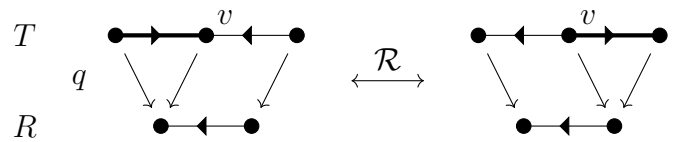
Lemma 4.4.1. *If $\mathfrak{p} : (R, \lambda, w) \to (T, \mu, v)$ is a correspondence of pointed directed trees, so is $\mathcal{R}\mathfrak{p} : (R, r_w\lambda, w) \to (T, r_v\mu, v)$. This gives functor $\mathcal{R}^\pm : \mathbf{DirTree}_\bullet^\pm \to \mathbf{DirTree}_\bullet^\mp$, which on objects acts as $\mathcal{R}^\pm(T, \mu, v) = (T, r_v\mu, v)$.*

Proof. We need to prove that $(r_v\mathfrak{p})^*(r_v\mu) = r_w\lambda$.

First assume $w \in V(R)$, $w = q(v)$ and $q^{-1}(\{w\}) = \{v\}$. Then any edge incident to w in R inherits its orientation from an edge incident to v in T, and $r_v\mathfrak{p} = \mathfrak{p}$, so the statement holds.

If not, then $w = \hat{0}$ so $r_w\lambda = \lambda$. Furthermore, if each edge incident to v in S is contracted by q, none of the edges in R inherits a direction from any of the edges whose direction we switched- explicitly $\mathfrak{p}^*\mu = \mathfrak{p}^*(r_v\mu) = (r_v\mathfrak{p})^*(r_v\mu)$, so the statement holds.

The last case left is illustrated in the figure below:



□

We can now prove:

Lemma 4.4.2. *$\omega_v \in Aut(\mathcal{P}_T)$ satisfies $\omega_v^*\mathcal{O}_{r_v\mu} = \mathcal{O}_\mu$.*

Proof. Consider a codimension-one cell in $\mathcal{P}_T$ represented by $\mathfrak{p} = (A_2 \overset{q}{\twoheadleftarrow} S \overset{i}{\hookrightarrow} T)$, and let $w = \mathfrak{p}^*v$. The functoriality of $\mathcal{R}$ gives the following commutative diagram:

$$\begin{array}{ccc} \mathcal{P}_T & \xrightarrow{\omega_v} & \mathcal{P}_T \\ {\scriptstyle i_{\mathfrak{p}}}\uparrow & & \uparrow{\scriptstyle i_{r_v\mathfrak{p}}} \\ \mathcal{P}_{A_2} & \xrightarrow{\omega_w} & \mathcal{P}_{A_2} \end{array}$$

Which reduces the proof to the case $T = A_2$. If $w = \hat{0}$, there is nothing to prove. If $w \in V(A_2)$, then $\overline{\omega}_w$ is a transposition on K_{A_2}, but for any orientation μ on A_2, $\mathcal{O}_\mu$ and $\mathcal{O}_{r_w\mu}$ give opposite cyclic orders, which shows $\omega_w^*\mathcal{O}_{r_w\mu} = \mathcal{O}_\mu$. □

The central result concerning directed reflection isomorphisms is as follows:

Theorem 4.4.1. *If (T, μ) and (T', μ') are directed trees, and $\varphi : (\mathcal{P}_T, \mathcal{O}_\mu) \xrightarrow{\sim} (\mathcal{P}_{T'}, \mathcal{O}_{\mu'})$ is an isomorphism of cyclic C-Arb singularities, then there is an orientation λ on T and an isomorphism of directed trees $\sigma : (T, \lambda) \xrightarrow{\sim} (T', \mu')$ such that $i_\sigma^{-1} \circ \varphi : (\mathcal{P}_T, \mathcal{O}_\mu) \to (\mathcal{P}_T, \mathcal{O}_\lambda)$ is a product of directed reflection isomorphisms.*

The argument is different depending on whether $T \cong A_n$ or not. In the next two parts we prove this theorem, first for A_n trees and then non-A_n trees.

Isomorphisms of Cyclic Singularities of A_n Type

In this part we prove theorem 4.4.1 in the case where $T \cong A_n$, which has some important differences from the general case. By the discussion at the beginning of the section, it is enough to consider $T = T' = A_n$. Furthermore, if μ and μ' are any two distinct orientations on T, there is a sequence of reflections which transform μ into μ'. Hence we can take $\mu = \mu'$ to be the orientation in which all edges point to the right (i.e. from i to $i+1$).

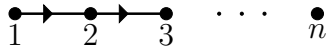


Figure 4.5: (A_n, μ)

Proposition 4.4.1. *If $\varphi \in Aut(\mathcal{P}_{A_n})$ satisfies $\varphi^* \mathcal{O}_\mu = \mathcal{O}_\mu$, then φ can be written as a product of reflection functors. Explicitly, the group $Aut(\mathcal{P}_{A_n}, \mathcal{O}_\mu)$ of such automorphisms is isomorphic to $\mathbb{Z}/(n+1)\mathbb{Z}$, and a generator can be written $\omega = \omega_1 \omega_2 \cdots \omega_n$, where ω_i is the reflection isomorphism at vertex i.*

Proof. By proposition , the automorphisms of $\mathcal{P}_{A_n}$ preserving $\mathcal{O}_\mu$ are the same as the automorphisms of K_{A_n} preserving the cyclic order induced by the total order $e_0 < e_1 < \cdots < e_n$. A generator for this group is $\theta(e_i) = e_{i+1}$, with index addition taken mod $n+1$. We will prove that the ω written in the statement of the proposition is well-defined, and $\overline{\omega} = \theta$.

To show ω is well-defined, note that n is a sink in A_n, so $\omega_n : (\mathcal{P}_{A_n}, \mathcal{O}_\mu) \to (\mathcal{P}_{A_n}, \mathcal{O}_{r_n \mu})$ is well-defined. Assume $\omega_k \cdots \omega_n$ is well-defined for some $1 < k \leq n$, and write $\mu^{(k)} := r_k \cdots r_n \mu$. Then the orientation of every edge has been flipped either twice or not at all, except the edge e_{k-1}, so vertex $k-1$ is a sink in $\mu^{(k)}$, meaning $\omega_{k-1} \omega_k \cdots \omega_n$ is well-defined. In $\mu^{(1)}$, every edge has been flipped twice so we have $\mu^{(1)} = \mu$, so $\omega \in \mathrm{Aut}(\mathcal{P}_{A_n}, \mathcal{O}_\mu)$.

We know $\overline{\omega}_i$ is the transposition $(e_{i-1} e_i)$, for $1 \leq i \leq n$. It follows that $\overline{\omega}(e_i) = \overline{\omega}_1 \cdots \overline{\omega}_n (e_i) = e_{i+1}$ as desired. □

Isomorphisms of Cyclic Singularities of Non-A_n Type

In this section we establish theorem 4.4.1 for trees not isomorphic to A_n. In some ways this situation is simpler; as we will see there is more rigidity imposed by the cyclic structure here.

Proposition 4.4.2. *Let T be a tree not isomorphic to A_n, and μ, μ' be two orientations on T. If $\varphi \in Aut(\mathcal{P}_T)$ satisfies $\varphi^*\mathcal{O}_\mu = \mathcal{O}_{\mu'}$, and φ stabilizes chains, then φ is a product of reflection isomorphisms.*

Proof. Note first that reflection isomorphisms stabilize all chains, since the induced bijection on K is the transposition of two edges which are connected by a vertex of degree 2, or an edge connected to a terminal vertex, and in each case both of these are contained in the same chain. The result follows from the following two observations:

(i) Let C be a chain with two boundary vertices v and w. The number of edges in C pointing from v to w in μ is the same as in μ'.

(ii) If $\mu = \mu'$, then φ is the identity.

(i) implies that we can obtain the orientation μ' from μ by a sequence of reflections, and (ii) implies that the product of the corresponding reflection isomorphisms must equal φ.

As a preliminary step, suppose $\varphi \in \mathrm{Aut}(\mathcal{P}_T)$ stabilizes chains, and suppose $\mathfrak{p} = (R \twoheadleftarrow S \hookrightarrow T)$ is a correspondence. Write $\varphi([\mathfrak{p}]) = [\mathfrak{p}']$, where $\mathfrak{p}' = (R' \overset{q'}{\twoheadleftarrow} S' \overset{i}{\hookrightarrow} T)$. Recall (lemma 4.1.2) that if $v \in V(T)$ is a vertex of degree > 2, $v \in V(S)$ iff $\mathfrak{p}$ intersects the node N_v. Since φ stabilizes chains, it also preserves the node relation $\sim_{N_v}$, so $\mathfrak{p}$ intersects the node N_v iff $\mathfrak{p}'$ does, in other words $v \in V(S)$ iff $v \in V(S')$.

Secondly, we use the notation $\langle 1\rangle$, $\langle 2\rangle$, $\langle e\rangle$ to denote the correspondences $\star \to A_2$. We remark that if $\mathfrak{p} = (A_2 \overset{q}{\twoheadleftarrow} S \overset{i}{\hookrightarrow} T)$ is a correspondence, and α is the edge in T joining the fibers $q^{-1}(1)$ and $q^{-1}(2)$, we have $X_{\mathfrak{p}\circ\langle e\rangle} = X_{\mathfrak{p}} \cup \{\alpha\}$.

We now proceed to the proofs of assertions (i) and (ii).

Proof of (i): Let α be in edge in a chain C with boundary vertices v and w. Write $\overline{\varphi}(\alpha) = \beta \in C$. We will show β must point the same direction as α. WLOG α points from v to w.

Let $\mathfrak{p} = (A_2 \overset{q}{\twoheadleftarrow} S \hookrightarrow T)$ be any correspondence with $v \in q^{-1}(1)$, $w \in q^{-1}(2)$, and α the edge joining the fibers $q^{-1}(1)$ and $q^{-1}(2)$. Our assumption is the edge in $\lambda = \mathfrak{p}^*\mu$ points from 1 to 2. Write $\mathfrak{p}' = (A_2 \overset{q'}{\twoheadleftarrow} S' \overset{i'}{\hookrightarrow} T)$, with $\varphi([\mathfrak{p}]) = [\mathfrak{p}']$, such that $\mathfrak{p}'^*\mu' = \lambda$. Since φ preserves the cyclic structure, there are three options:

$$(A)\begin{cases} \varphi([\mathfrak{p}\circ\langle 1\rangle]) = [\mathfrak{p}'\circ\langle 1\rangle] \\ \varphi([\mathfrak{p}\circ\langle e\rangle]) = [\mathfrak{p}'\circ\langle e\rangle] \\ \varphi([\mathfrak{p}\circ\langle 2\rangle]) = [\mathfrak{p}'\circ\langle 2\rangle] \end{cases};\qquad (B)\begin{cases} \varphi([\mathfrak{p}\circ\langle 1\rangle]) = [\mathfrak{p}'\circ\langle e\rangle] \\ \varphi([\mathfrak{p}\circ\langle e\rangle]) = [\mathfrak{p}'\circ\langle 2\rangle] \\ \varphi([\mathfrak{p}\circ\langle 2\rangle]) = [\mathfrak{p}'\circ\langle 1\rangle] \end{cases};$$

$$(C)\begin{cases} \varphi([\mathfrak{p}\circ\langle 1\rangle]) = [\mathfrak{p}'\circ\langle 2\rangle] \\ \varphi([\mathfrak{p}\circ\langle e\rangle]) = [\mathfrak{p}'\circ\langle 1\rangle] \\ \varphi([\mathfrak{p}\circ\langle 2\rangle]) = [\mathfrak{p}'\circ\langle e\rangle] \end{cases}$$

By our preliminary step, we know $\varphi([\mathfrak{p}\circ\langle 2\rangle])$ intersects N_v but not N_w, and $\varphi([\mathfrak{p}\circ\langle 1\rangle])$ intersects N_w but not N_v. Since $v, w \in V(S')$, $[\mathfrak{p}'\circ\langle e\rangle]$ intersects both N_v and N_w, which means option (A) is the correct one. That implies $q'(v) = 1$ and $q'(w) = 2$. Additionally, $X_{\mathfrak{p}'\circ\langle e\rangle} = \varphi(X_{\mathfrak{p}\circ\langle e\rangle}) = X_{\mathfrak{p}'} \cup \{\beta\}$, meaning β is the edge joining the fibers $q'^{-1}(1)$ and $q'^{-1}(2)$. Therefore, β also points from v to w.

Proof of (ii): First, let C be a chain with two boundary vertices v, w. We have shown $\overline{\varphi}$ must permute the edges between v and w pointing from v to w, we will show this permutation must be the identity. Pick two edges α, β between v and w, both pointing from v to w, with α closer to v than β. It is enough to prove that $\overline{\varphi}(\alpha) = \alpha'$ is also closer to v than $\overline{\varphi}(\beta) = \beta'$.

Construct the correspondence $\mathfrak{p} = (A_2 \overset{q}{\twoheadleftarrow} S \hookrightarrow T)$, where $q^{-1}(1)$ consists of the subtree connecting v to the closer vertex of α, and $q^{-1}(2)$ is the subtree connecting α to β. By our assumption, $\mathfrak{p}^*\mu = \lambda$ is the orientation where the edge points from 1 to 2. Let $\mathfrak{p}' = (A_2 \overset{q'}{\twoheadleftarrow} S' \hookrightarrow T)$ be a correspondence representing $\varphi([\mathfrak{p}])$, with $\mathfrak{p}'^*\mu' = \lambda$.

As before, we have options (A), (B), (C). The fact that $\mathfrak{p}\circ\langle 1\rangle$ does not intersect N_v eliminates option (B).

We have $X_{\mathfrak{p}\circ\langle e\rangle} = X_{\mathfrak{p}} \cup \{\alpha\}$, and $X_{\mathfrak{p}\circ\langle 2\rangle} = X_{\mathfrak{p}} \cup \{\beta\}$. If option (C) were true, then we would have $X_{\mathfrak{p}'\circ\langle e\rangle} = X_{\mathfrak{p}} \cup \{\beta'\}$, meaning β' is the edge joining the fibers $q'^{-1}(1)$ and $q'^{-1}(2)$. Furthermore, we would have $\mathfrak{p}'\circ\langle 1\rangle$ intersects N_v, meaning $v \in q'^{-1}(2)$. This would mean β' points from w toward v, contradicting (i).

This leaves option (A), which means α' is the edge joining the fibers $q'^{-1}(1)$ and $q'^{-1}(2)$, and $v \in q'^{-1}(1)$, so α' is closer to v than β'.

Lastly, suppose C is a chain with a single boundary vertex v. WLOG we can assume all edges point away from v in μ, since any situation can be turned into this one by a series of reflection isomorphisms. Then the same argument as before shows that if α is closer to v than β, the same must be true of $\overline{\varphi}(\alpha)$ and $\overline{\varphi}(\beta)$ (and this even works if β is a boundary vertex).

□

This establishes theorem 4.4.1 for $T \ncong A_n$.

4.5 The Coherence Theorem

We finish this section with a proof of theorem 3.3.1.

> **Theorem 3.3.1:** Let T be a tree, and $\mathcal{O}$ a pre-cyclic structure on $\mathcal{P}_T$. If $\mathcal{O}$ is A_3-coherent and S_4-coherent, then $\mathcal{O}$ is a cyclic structure.

Proof. **Step 1:** Use A_3-coherence to construct an orientation μ on T.

To begin, take vertices v and w of degree $\neq 2$, such that v and w can be connected by a path in which each intermediate vertex has degree 2. The path connecting these vertices is a subtree S isomorphic to A_n, and one has a correspondence:

$$\mathfrak{p} = (A_n \overset{q}{\twoheadleftarrow} S \overset{i}{\hookrightarrow} T)$$

Where q is an isomorphism, $q(v) = 1$ and $q(w) = n$.

To fix terminology, call such a subtree a **chain subtree**, and the associated correspondence a **chain correspondence**. Notice that for each edge e there is a unique (up to equivalence) chain correspondence $\mathfrak{p} = (A_n \overset{q}{\twoheadleftarrow} S \overset{i}{\hookrightarrow} T)$ with $e \in E(S)$.

Now, let u be a vertex between v and w, and write $q(u) = j$, where j is an integer between 1 and n. Then $\mathfrak{p}$ gives a correspondence of pointed trees $(A_n, j) \to (T, u)$, so by the functoriality of the reflection isomorphisms there is a commutative diagram

$$\begin{array}{ccc} \mathcal{P}_T & \xrightarrow{\omega_u} & \mathcal{P}_T \\ {\scriptstyle i_{\mathfrak{p}}}\uparrow & & \uparrow{\scriptstyle i_{\mathfrak{p}}} \\ \mathcal{P}_{A_n} & \xrightarrow{\omega_j} & \mathcal{P}_{A_n} \end{array}$$

Write $K_{A_n} = \{e_0, e_1, \ldots, e_n\}$, where e_i is the edge joining vertex i to $i+1$ for $1 \leq i < n$, and e_0, e_n are the terminal vertices. Taking transpositions centered on internal vertices allows us to obtain any permutation of K_{A_n} fixing $V_t(A_n) = \{e_0, e_n\}$. As a result, if $\overline{\eta} \in \mathfrak{S}(K_{A_n})$ is any automorphism fixing $\{e_0, e_n\}$, we can find $\eta \in \mathrm{Aut}(\mathcal{P}_{A_n})$ and $\omega \in \mathrm{Aut}(\mathcal{P}_T)$ such that there is a commutative diagram:

$$\begin{array}{ccc} \mathcal{P}_T & \xrightarrow{\omega} & \mathcal{P}_T \\ {\scriptstyle i_{\mathfrak{p}}}\uparrow & & \uparrow{\scriptstyle i_{\mathfrak{p}}} \\ \mathcal{P}_{A_n} & \xrightarrow{\eta} & \mathcal{P}_{A_n} \end{array}$$

We know $\mathcal{O}$ is A_3-coherent, so by proposition 4.4, $i_{\mathfrak{p}}^*\mathcal{O}$ is a cyclic structure, and $c = \Delta(i_{\mathfrak{p}^*}\mathcal{O})$ is a cyclic order on K_{A_n}. There exists a bijection $\overline{\eta} \in \mathfrak{S}(K_{A_n})$ fixing e_0 and e_n such that c is induced by the total order $e_0 < \overline{\eta}(e_1) < \cdots < \overline{\eta}(e_k) < e_n < \overline{\eta}(e_{n-1}) < \overline{\eta}(e_{n-2}) < \cdots < \overline{\eta}(e_{k+1})$.

For this $\overline{\eta}$, pick η and ω as in the above commutative diagram.

If we let λ denote the orientation of A_n where $e_1, \ldots, e_k$ are pointing to the right and $e_{k+1}, \ldots, e_{n-1}$ are pointing to the left, by proposition 3.6.2 the cyclic order $\Delta(\mathcal{O}_\lambda)$ is induced by the total order $e_0 < e_1 < \cdots < e_k < e_n < e_{n-1} < \cdots < e_{k+1}$. So we have $\mathcal{O}_\lambda = \eta^* i_{\mathfrak{p}}^* \mathcal{O} = i_{\mathfrak{p}}^* \omega^* \mathcal{O}$. Since ω is an automorphism of $\mathcal{P}_T$, $\mathcal{O}$ is a cyclic structure iff $\omega^*\mathcal{O}$ is, so we can replace $\mathcal{O}$ with $\omega^*\mathcal{O}$. So we can assume $i_{\mathfrak{p}}^*\mathcal{O} = \mathcal{O}_\lambda$ for some orientation λ.

Repeating this process for each chain correspondence, we can assume (after potentially applying an automorphism of $\mathcal{P}_T$) that for each chain correspondence $\mathfrak{p} : A_n \to T$, $\mathfrak{p}^*\mathcal{O} = \mathcal{O}_{\lambda(\mathfrak{p})}$ for some orientation $\lambda(\mathfrak{p})$. There is a unique orientation μ on T such that $\mathfrak{p} : (A_n, \lambda(\mathfrak{p})) \to (T, \mu)$ is an orientation-preserving correspondence. This tells us $i_{\mathfrak{p}}^*\mathcal{O} = i_{\mathfrak{p}}^*\mathcal{O}_\mu$ for any chain correspondence $\mathfrak{p}$.

Step 2: Our second step is to use S_4-coherence to show $\mathcal{O} = \mathcal{O}_\mu$.

Pick a correspondence $\mathfrak{p} = (A_2 \overset{q}{\twoheadleftarrow} S \overset{i}{\hookrightarrow} T)$. We are done if we can show $i_{\mathfrak{p}}^*\mathcal{O} = i_{\mathfrak{p}}^*\mathcal{O}_\mu$.

Step 2a: We consider a special case. If $\mathfrak{p}$ factors through a chain correspondence, then we are done. If not, assume there exists a chain subtree R such that $q^{-1}(2) \subseteq R$, and $q^{-1}(1)$ contains one of the endpoints v of R with the degree of v greater than 2.

We will construct a correspondence $\mathfrak{q} = (S_n \overset{q'}{\twoheadleftarrow} W \overset{i'}{\hookrightarrow} T)$ with $[\mathfrak{p}] \geq [\mathfrak{q}]$:

- $W = S \cup n(v)$, where $n(v)$ denotes the set of neighbors of the vertex $v \in V(T)$.
- We define the fiber of the central vertex, $q'^{-1}(0)$, to be $R \cap q^{-1}(1)$.
- Define the fiber $q'^{-1}(1)$ to be $q^{-1}(2)$.
- The rest of the fibers are the connected components of $W \backslash V$ (since v has degree > 2 there are at least two such connected components)

Let $\mathfrak{p}_0 : A_2 \to S_n$ denote the correspondence satisfying $\mathfrak{q} \circ \mathfrak{p}_0 = \mathfrak{p}$.

The pullback $i_{\mathfrak{q}}^*\mathcal{O}$ is an S_4-coherent cyclic structure on S_n. By proposition 3.7.1 it is a cyclic structure and equal to $\mathcal{O}_{\mu'}$ for some orientation μ' on S_n- so $i_{\mathfrak{p}}^*\mathcal{O} = i_{\mathfrak{p}_0}^*\mathcal{O}_{\mu'}$. We

are done, then (with this case), if we can show $\mathfrak{q}^*\mu$ agrees with μ' on the edge joining 0 and 1.

To accomplish this, let $\mathfrak{p}_1 : A_2 \to S_n$ denote the correspondence whose fiber over 1 is the central vertex 0, and whose fiber over 2 is the vertex 1.

Since $\mathfrak{q} \circ \mathfrak{p}_1$ factors though the chain correspondence $(A_n \twoheadleftarrow R \hookrightarrow T)$, $i^*_{\mathfrak{q}\circ\mathfrak{p}_1}\mathcal{O} = i^*_{\mathfrak{q}\circ\mathfrak{p}_1}\mathcal{O}_\mu = i^*_{\mathfrak{p}_1}\mathcal{O}_{\mathfrak{q}^*\mu}$. We also have $i^*_{\mathfrak{q}\circ\mathfrak{p}_1}\mathcal{O} = i^*_{\mathfrak{p}_1} i_{\mathfrak{q}^*}\mathcal{O} = i_{\mathfrak{p}_1^*}\mathcal{O}_{\mu'}$. That implies $\mathfrak{q}^*\mu$ agrees with μ' on the edge joining 0 and 1.

This strategy is illustrated in Figure 4.6. The figure shows the construction of $\mathfrak{q}$ and $\mathfrak{p}_1$ from $\mathfrak{p}$, and the edge e is the one whose direction is crucial.

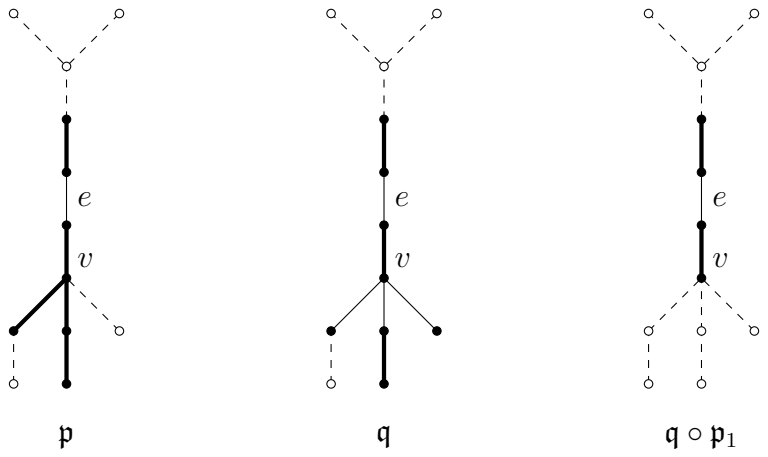


Figure 4.6: The Construction of $\mathfrak{q}$ and $\mathfrak{p}_1$

Step 2b: Finally, we consider the general case of a correspondence $\mathfrak{p} = (A_2 \overset{q}{\twoheadleftarrow} S \overset{i}{\hookrightarrow} T)$.

The edge connecting $q^{-1}(1)$ and $q^{-1}(2)$ is contained in a chain subtree R, call the endpoints of R v and w. We can assume $v \in q^{-1}(1)$ and $w \in q^{-1}(2)$, and both v and w have degree greater than 2 (if not, we are in one of the previous cases).

We can now repeat essentially the same argument as the previous case. Explicitly: we again construct a correspondence $\mathfrak{q} = (S_n \overset{q'}{\twoheadleftarrow} W \overset{i'}{\hookrightarrow} T)$:

- $W = S \cup n(v)$, where $n(v)$ denotes the set of neighbors of the vertex $v \in V(T)$.
- We define the fiber of the central vertex, $q'^{-1}(0)$, to be $R \cap q^{-1}(1)$.
- Define the fiber $q'^{-1}(1)$ to be $q^{-1}(2)$.

- The rest of the fibers are the connected components of $W \backslash V$ (since v has degree > 2 there are at least two such connected components)

Let $\mathfrak{p}_0 : A_2 \to S_n$ denote the correspondence satisfying $\mathfrak{q} \circ \mathfrak{p}_0 = \mathfrak{p}$.

$i_{\mathfrak{q}}^* \mathcal{O} = \mathcal{O}_{\mu'}$ for some orientation μ' on S_n, we are done if we can show $\mathfrak{q}^* \mu$ agrees with μ' on the edge joining 0 and 1. As before, let $\mathfrak{p}_1 : A_2 \to S_n$ denote the correspondence whose fiber over 1 is the central vertex 0, and whose fiber over 2 is the vertex 1:

The difference in this case is that $\mathfrak{q} \circ \mathfrak{p}_1$ does not necessarily factor through a chain correspondence- however, it is of the form covered in the previous case, so we again have $i_{\mathfrak{q} \circ \mathfrak{p}_1}^* \mathcal{O} = i_{\mathfrak{q} \circ \mathfrak{p}^*} \mathcal{O}_\mu$. See Figure 4.7 for an illustration of this case.

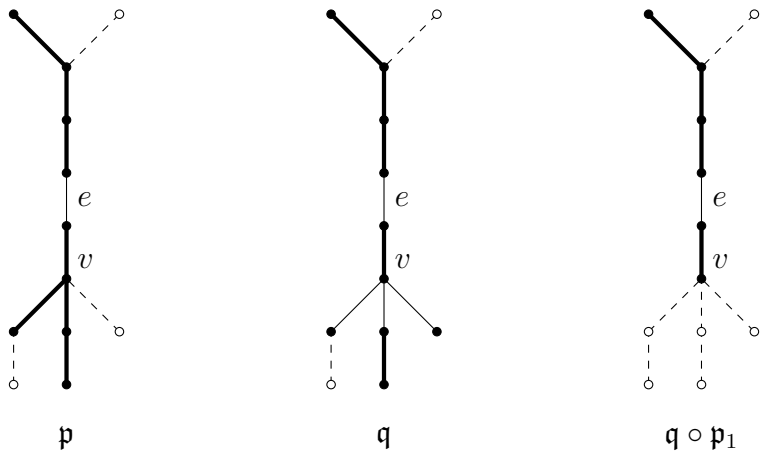


Figure 4.7: The construction in Case 2b

The rest of the argument is identical.

□

Chapter 5

A Sheaf of Categories

5.1 Dg-Categories

This section is centered around the representation theory of C-Arb singularities with cyclic structures. Of central importance is the notion of a dg-category. Standard references are [3], [14]. Quotients of dg-categories are explained in [7]. The model category structure and the underlying infinity category is discussed in [34], [8].

Throughtout this paper, we fix an algebraically closed field k of characteristic zero. All vector spaces will be assumed to be over k.

Definition 5.1.1. *A* ***dg-category*** $\mathcal{C}$ *is a category enriched over chain complexes. Explicitly, for objects* $X, Y \in \mathcal{C}$*,* $Hom(X, Y)$ *has the structure of a chain complex (where the differential increases degree). If* $f \in Hom^k(X, Y)$*, we say* f *is a homogeneous element, and write* $|f| = k$*. The following axioms are satisfied:*

(i) For objects X, Y, Z*: The composition morphism:*

$$Hom(Y, Z) \times Hom(X, Y) \to Hom(X, Z)$$

Is bilinear.

(ii) For homogeneous functions $f \in Hom(X, Y)$*,* $g \in Hom(Y, Z)$*,* $|gf| = |g| + |f|$*.*

(iii) With f *and* g *as in (ii), we have the* ***graded Leibniz rule****:*

$$d(fg) = (df)g + (-1)^{|f|} f(dg)$$

Definition 5.1.2. *If $\mathcal{C}$ and $\mathcal{D}$ are dg-categories, a **dg-functor** is a functor $F : \mathcal{C} \to \mathcal{D}$ such that, for any objects $X, Y \in \mathcal{C}$, the map:*

$$F : Hom_{\mathcal{C}}(X, Y) \to Hom_{\mathcal{D}}(FX, FY)$$

Is a map of chain complexes.

Definition 5.1.3. *The **graded homotopy category** $H^\bullet\mathcal{C}$ of a dg-category $\mathcal{C}$ is the category whose objects are the objects of $\mathcal{C}$, and for $X, Y \in Ob(\mathcal{C})$, we have $Hom_{H^\bullet\mathcal{C}}(X, Y) = H^\bullet(Hom_{\mathcal{C}}(X, Y))$. The **homotopy category** $H^0\mathcal{C}$ is the degree-zero portion of the graded homotopy category.*

$H^\bullet(\mathrm{Hom}_{\mathcal{C}}(X, Y))$ is a graded $k-$linear category. An **isomorphism** in such a category is a degree zero invertible map. A morphism in $\mathcal{C}$ whose image in $H^\bullet\mathcal{C}$ is an isomorphism is a **homotopy equivalence**. Two objects in $\mathcal{C}$ are **homotopic** if they are isomorphic in $H^\bullet\mathcal{C}$.

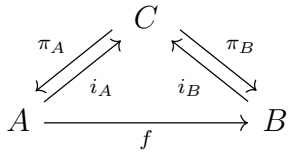
Note that a dg-functor $F : \mathcal{C} \to \mathcal{D}$ induces $H^\bullet\mathcal{F} : H^\bullet\mathcal{C} \to H^\bullet\mathcal{D}$.

Definition 5.1.4. *A dg-functor $F : \mathcal{C} \to \mathcal{D}$ is:*

- ***Quasi-fully faithful** if $H^\bullet F$ is fully faithful (it induces isomorphisms on hom sets).*
- ***Quasi-essentially surjective** if $H^\bullet F$ is essentially surjective (Any object in $H^\bullet\mathcal{D}$ is isomorphic to $H^\bullet F(X)$ for some $X \in Ob(H^\bullet\mathcal{C})$).*
- *A **quasi-equivalence** if $H^\bullet F$ is an equivalence of categories, that is, it is fully faithful and essentially surjective.*

Definition 5.1.5. *Let $\mathcal{C}$ be a dg-category.*

- *If $A \in Ob(\mathcal{C})$, an **k-tranlastion** of A is an object B such that there exist closed maps $f : A \to B$ and $g : B \to A$ with $fg = Id_B$, $gf = Id_A$, $|f| = -k$ and $|g| = k$.*
- *If $f : A \to B$ is a closed degree-zero morphism, a **cone** of f is an object C, with maps:*



With $|\pi_B| = |i_B| = 0$, $|\pi_A| = 1$, $|i_A| = -1$, $\pi_B i_B = Id_B$, $\pi_A i_A = Id_A$, $i_B\pi_B + i_A\pi_A = Id_C$, $\pi_A i_B = \pi_B i_A = di_B = d\pi_A = 0$, $di_A = i_B f$, $d\pi_B = -f\pi_A$.

- *A* ***zero object*** $Z \in Ob(\mathcal{C})$ *is an object with* $Id_Z = 0_Z$.
- *A dg category is* ***strongly pretriangulated*** *if it has a zero object, a k-translation of any object, and a cone of any closed degree zero morphism. Note that these objects are unique up to unique isomorphism, and are preserved by dg functors.*

Definition 5.1.6. *If* $\mathcal{C}$ *is strongly pretriangulated, a sequence of closed degree zero maps:*

$$A \xrightarrow{f} B \xrightarrow{g} C$$

is ***exact*** *if the map* $g \circ \pi_B : Cone(f) \to C$ *is a homotopy equivalence, equivalently if* $i_{B[1]} \circ f[1] : A[1] \to Cone(g)$ *is a homotopy equivalence. If* $F, G, H : \mathcal{C} \to \mathcal{D}$ *are dg-functors and* $\mathcal{D}$ *is strongly pretriangulated, a sequence of natural transformations:*

$$F \xrightarrow{\eta} G \xrightarrow{\phi} H$$

is called ***exact*** *for any object* $X \in Ob(\mathcal{C})$*, the sequence:*

$$FX \xrightarrow{\eta_X} GX \xrightarrow{\phi_X} HX$$

is exact.

If $\mathcal{C}$ is strongly pretriangulated, then the homotopy category $H^0\mathcal{C}$ is a triangulated category whose distinguished triangles are the image of exact sequences.

To any dg-category $\mathcal{C}$ there exists a strongly pretriangulated dg-category $\mathcal{C}^{pre-tr}$ called the **pretriangulated hull** of $\mathcal{C}$ and a canonical fully faithful functor $\mathcal{C} \hookrightarrow \mathcal{C}^{pre-tr}$ (see [5]). The pretriangulated hull satisfies a universal property:

Lemma 5.1.1. *If* $\mathcal{D}$ *is a strongly pretriangulated category and* $F : \mathcal{C} \to \mathcal{D}$ *is a functor, there is a unique (up to unique natural isomorphism) extension of* F *to a functor* $\mathcal{C}^{pre-tr} \to \mathcal{D}$.

We turn to the 2-categorical structure underlying dg-categories.

Definition 5.1.7. *Let* $F, G : \mathcal{C} \to \mathcal{D}$ *be dg-functors between dg-categories. A natural transformation* $\eta : F \to G$ *is:*

*(i) An **isomorphism** if $\eta_X : FX \to GX$ is an isomorphism for all $X \in Ob(\mathcal{C})$.*

*(ii) A **quasi-isomorphism** if η_X is a homotopy equivalence for all $X \in Ob(\mathcal{C})$.*

We will often make use of the following proposition in order to 'lift' properties of a dg-category $\mathcal{C}$ to properties of $\mathcal{C}^{pre-tr}$.

Lemma 5.1.2. *Let $\mathcal{C}$ be a dg-category, $\mathcal{D}$ be a strongly pre-triangulated dg-category, $F, G : \mathcal{C}^{pre-tr} \to \mathcal{D}$ be dg-functors. Let $F_{\mathcal{C}}, G_{\mathcal{C}}$ denote the restrictions of F and G to $\mathcal{C}$. Then if $\eta_{\mathcal{C}} : F_{\mathcal{C}} \to G_{\mathcal{C}}$ is a a natural transformation, there is a unique extension to $\eta : F \to G$. Furthermore:*

- *If $\eta_{\mathcal{C}}$ is an isomorphism, so is η.*
- *If $\eta_{\mathcal{C}}$ is a quasi-isomorphism, so is η.*
- *If $H : \mathcal{C}^{pre-tr} \to \mathcal{D}$ is another dg-functor, and:*

$$F_{\mathcal{C}} \xrightarrow{\eta_{\mathcal{C}}} G_{\mathcal{C}} \xrightarrow{\phi_{\mathcal{C}}} H_{\mathcal{C}}$$

Is an exact sequence of natural transformations, the unique extensions:

$$F \xrightarrow{\eta} G \xrightarrow{\phi} H$$

Also give an exact sequence.

5.2 Sheaves of dg-Categories

In this paper we will be dealing with sheaves of dg-categories on C-Arb singularities $\mathcal{P}_T$. By a sheaf, we mean a sheaf in the $\infty-$category sense- for details, see appendix B. The following definition will be important:

Definition 5.2.1. *Let $\mathcal{C}$ be an ordinary category, and let dg-Cat denote the $\infty-$category of dg-categories (over k). A **2-functor** $\mathcal{Q} : \mathcal{C} \to$ dg-Cat consists of the following data:*

(i) For each object $X \in Ob(\mathcal{C})$, a dg-category $\mathcal{Q}(X)$.

(ii) For each morphism $f : X \to Y$ in $\mathcal{C}$, a dg-functor $c_f : \mathcal{Q}(X) \to \mathcal{Q}(Y)$.

(iii) For each pair of morphisms f, g such that the composition fg exists, a quasi-isomorphism $\eta_{f,g} : c_f \circ c_g \to c_{fg}$.

(iv) Such that, for each triple f, g, h of morphisms such that the composition fgh exists, the following diagram commutes:

$$\begin{array}{ccc} c_f \circ c_g \circ c_h & \xrightarrow{\eta_{f,g} \diamond 1_{c_h}} & c_{fg} \circ c_h \\ \Big\downarrow{\scriptstyle 1_{c_f} \diamond \eta_{g,h}} & & \Big\downarrow{\scriptstyle \eta_{fg,h}} \\ c_f \circ c_{gh} & \xrightarrow{\eta_{f,gh}} & c_{fgh} \end{array}$$

Where $\diamond$ refers to the horizontal composition of functors.

Remark 5.2.1. *By lemma 5.1.2, a 2-functor $\mathcal{Q}^0 : \mathcal{C} \to$ dg-Cat can be extended to a 2-functor $\mathcal{Q} : \mathcal{C} \to$ dg-Cat whose image lies in the subcategory of strongly pretriangulated dg-categories, by setting $\mathcal{Q}(X) = \mathcal{Q}^0(X)^{pre-tr}$, and extending the rest of the data.*

Now, let $\mathcal{C}$ be a category in which every morphism is monic. For $A \in \mathrm{Ob}(\mathcal{C})$, recall the construction of the poset $\mathcal{P}_A = \mathrm{Pos}(\mathcal{C}_{/A})$ from Chapter 2. Then a 2-functor $\mathcal{Q} : \mathcal{C}^{op} \to$ dg-Cat determines a sheaf of dg-categories $\mathcal{Q}_A$ on $\mathcal{P}_A$ for $A \in \mathrm{Ob}(\mathcal{C})$, where:

(i) The stalk $\mathcal{Q}_{A,[f]}$ of $\mathcal{Q}_A$ at $f : B \to A$ is $\mathcal{Q}(B)$.

(ii) When $[f] \geq [g]$, meaning $f = gh$, the restriction $\mathcal{Q}_{A,[g]} \to \mathcal{Q}_{A,[f]}$ is given by c_h.

Definition 5.2.2. *Let $\mathcal{C}$ be a category, and $\mathcal{Q}, \mathcal{Q}'$ 2-functors $\mathcal{C} \to$ dg-Cat, with associated data c, η, resp. c', η'. A **left morphism** $\mathcal{F} : \mathcal{Q} \to \mathcal{Q}'$ consists of:*

(i) For each $X \in Ob(\mathcal{C})$, a dg-functor $\mathcal{F}(X) : \mathcal{Q}(X) \to \mathcal{Q}'(X)$.

(ii) For each morphism $f : Y \to X$, a quasi-isomorphism $\gamma_f : c'_f \circ \mathcal{F}(Y) \to \mathcal{F}(X) \circ c_f$.

(iii) Such that, for morphisms $f : Y \to X$ and $g : Z \to Y$, the following diagram commutes:

$$\begin{array}{ccccc} c'_f \circ c'_g \circ \mathcal{F}(X) & \xrightarrow{1_{c'_f} \diamond \gamma_g} & c'_f \circ \mathcal{F}(Y) \circ c_g & \xrightarrow{\gamma_f \diamond 1_g} & \mathcal{F}(Z) \circ c_f \circ c_g \\ \Big\downarrow{\scriptstyle \eta'_{f,g} \diamond 1_{\mathcal{F}(X)}} & & & & \Big\downarrow{\scriptstyle 1_{\mathcal{F}(Z)} \diamond \eta_{f,g}} \\ c'_{fg} \circ \mathcal{F}(X) & & \xrightarrow{\gamma_{fg}} & & \mathcal{F}(Z) \circ c_{fg} \end{array}$$

*A **right morphism** is defined similarly, except $\gamma_f : \mathcal{F}(X) \circ c_f \to c'_f \circ \mathcal{F}(Y)$, and in the commutative diagram in (iii), the horizontal arrows point in the other direction.*

A left or right morphism between 2-functors $\mathcal{Q} \to \mathcal{Q}'$ determines a morphism $\mathcal{Q}_A \to \mathcal{Q}'_A$ of sheaves of dg-categories on $\mathcal{P}_A$ for any $A \in \mathrm{Ob}(\mathcal{C})$- see appendix B for a detailed demonstration of this fact.

There is also a notion of when left/right morphisms are quasi-inverses, and thus induce quasi-equivalences on sheaves. In a later section, when we need this, we will outline explicitly what data we need. For more details see appendix B.

5.3 Representations of Directed Trees

We describe the construction of a sheaf of dg-categories on the cyclic C-Arb singularity $(\mathcal{P}_T, \mathcal{O}_\mu)$, where μ is an orientation on a tree T, by defining a 2-functor $\mathbf{DirTree}^{op} \to$ dg-Cat. For other perspectives on this construction, see [20] or [28][1].

To begin, we define a dg category associated to the tree (T, μ).

Definition 5.3.1. *Let (T, μ) be a directed tree. The category $Rep^0(T, \mu)$ is a $k-$linear category generated where:*

- *For each $v \in V(T)$, there is an object P_v.*
- *If there exists a directed path from v to w in (T, μ), there is a distinguished nonzero morphism $e_{w,v} : P_w \to P_v$, such that $Hom(P_w, P_v) = Span(e_{w,v})$.*
- *$e_{v,w} \circ e_{u,v} = e_{u,w}$.*
- *If there does not exist a directed path from v to w in (T, μ), then $Hom(P_w, P_v) = 0$.*

This can be thought of as a dg-category where all $e_{w,v}$ are in degree zero and all differentials are zero. The category $Rep(T, \mu) := (Rep^0(T, \mu))^{pre-tr}$.

We make some remarks regarding the theory of quiver representations- see [6] for proofs and a detailed exposition. First, one has the classical notion of the abelian category $\mathrm{Mod}(T, \mu)$ of finite-dimensional representations of (T, μ). Objects of this category assign a finite-dimensional $k-$vector space X_v to each $v \in V(T)$, and a morphism $\alpha_e^X : X_{t(e)} \to X_{h(e)}$

[1] While both of those papers deal explicitly with *rooted* trees, there are no essential differences

to each edge $e \in E(T)$. Morphisms $X \to Y$ of this category are a collection of maps $f_v : X_v \to Y_v$ satisfying $\alpha_e^Y f_{t(e)} = f_{h(e)} \alpha_e^X$ for all $e \in E(T)$.

For each $v \in V(T)$, there is a corresponding projective object $P_v \in \mathrm{Mod}(T, \mu)$, which assigns:

$$(P_v)_u = \begin{cases} k & \text{There exists a directed path from } v \text{ to } u \\ 0 & \text{Else} \end{cases}$$

And all arrows α_e are the identity when both endpoints are assigned k, and zero otherwise.

One can see easily that the category $\mathrm{Rep}^0(T, \mu)$ is isomorphic to the subcategory of $\mathrm{Mod}(T, \mu)$ on the objects $\{P_v\}_{v \in V(T)}$. Furthermore ([6]), any object in $\mathrm{Mod}(T, \mu)$ has a finite projective resolution whose objects are direct sums of the P_v- hence $\mathrm{Rep}(T, \mu)$ is a dg-enhancement of the bounded derived category of $\mathrm{Mod}(T, \mu)$. This category is sometimes called the 'bounded derived dg category', and can also be presented as the dg-quotient ([7]) of the dg-category of finite-dimensional chain complexes of $(T, \mu)-$modules by the subcategory of acyclic complexes.

Returning to the goal at hand: We want to define a 2-functor $\mathcal{Q} : \mathbf{DirTree}^{op} \to$ dg-Cat.

Definition 5.3.2. *We define $\mathcal{Q}^0 : \boldsymbol{DirTree}^{op} \to$ dg-Cat as follows:*

(i) For a directed tree (T, μ), we set $\mathcal{Q}^0(T, \mu) = Rep^0(T, \mu)$.

(ii) For a correspondence $\mathfrak{p} = (R \xleftarrow{q} S \xhookrightarrow{i} T)$, define the functor $c^0_{\mathfrak{p}} : Rep^0(T, \mu) \to Rep^0(R, \lambda)$, where $\lambda = \mathfrak{p}^* \mu$, by:

- $c^0_{\mathfrak{p}}(P_v) = \begin{cases} P_{q(v)} & v \in V(S) \\ 0 & Else \end{cases}$
- $c^0_{\mathfrak{p}}(e_{v,w}) = \begin{cases} e_{q(v),q(w)} & v, w \in V(S) \\ 0 & Else \end{cases}$

(iii) Note that $c^0_{\mathfrak{p}'} \circ c^0_{\mathfrak{p}} = c^0_{\mathfrak{p} \circ \mathfrak{p}'}$, hence we can set $\eta^0_{\mathfrak{p},\mathfrak{p}'}$ to be the identity transformation.

We now define $\mathcal{Q}$ to be the pretriangulated hull of $\mathcal{Q}^0$, as in remark 5.2.1.

To verify that $c^0_{\mathfrak{p}}$ is well-defined, first notice that by the definition of $\lambda = \mathfrak{p}^* \mu$, if there is a path from w to v in (T, μ) there will be a path from $q(w)$ to $q(v)$ in (R, λ). It also bears mentioning that if $e_{u,w} = e_{v,w} \circ e_{u,v}$ and $u, w \in V(S)$, then we also must have $v \in V(S)$, since no directed path can leave S and reenter it. (It would be a problem if it were not, since then we would have $c^0_{\mathfrak{p}}(e_{u,v}) = c^0_{\mathfrak{p}}(e_{v,w}) = 0$ but $c^0_{\mathfrak{p}}(e_{u,w}) = e_{q(u),q(w)} \neq 0$).

In the following sections we will analyze some properties of this sheaf.

5.4 The Trivalent Vertex

The simplest case to investigate is the trivalent vertex: $\mathcal{P}_{A_2}$. A_2 denotes the tree with two vertices, labelled 1 and 2, and a single edge connecting them. We let μ denote the orientation with the arrow pointing from 1 to 2 (Figure 5.1).

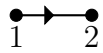


Figure 5.1: The directed tree (A_2, μ)

We use the notation $(1), (2), (12)$ for the three correspondences $\star \to \mathcal{P}_{A_2}$ (see, for example, Figure 2.4). We have the following classical fact about $\mathrm{Rep}(A_2, \mu)$:

Lemma 5.4.1. *The indecomposable objects in Rep(A_2, μ) are (up to shifts and homotopy):*

- P_1
- P_2
- $P_2[1] \xrightarrow{e_{2,1}} P_1$

These correspond to the indecomposables $k \to k$, $k \to 0$, and $0 \to k$, respectively, in $\mathrm{Mod}(A_2, \mu)$. We are using the notation of twisted complexes to describe $P_2[1] \xrightarrow{e_{2,1}} P_1 \cong \mathrm{Cone}(e_{2,1}) \in \mathrm{Rep}(A_2, \mu)$. See definition 5.7.1 later in this chapter.

We can compute the restrictions $c_{\mathfrak{p}}X$, for $\mathfrak{p} = (1), (2), (12)$, this is given in Table 5.1, and illustrated geometrically (with the cyclic order $\mathcal{O}_\mu$ agreeing with the counterclockwise order) in Figure 5.2. The figure should be interpreted as showing the 'support' of the indecomposables, i.e. indicating where the restrictions are nonzero.

$\mathfrak{p}$	P_1	P_2	$P_2[1] \to P_1$
(1)	k	0	k
(2)	0	k	$k[1]$
(12)	k	k	0

Table 5.1: The Restrictions of Indecomposable Elements of $\mathrm{Rep}(A_2, \mu)$.

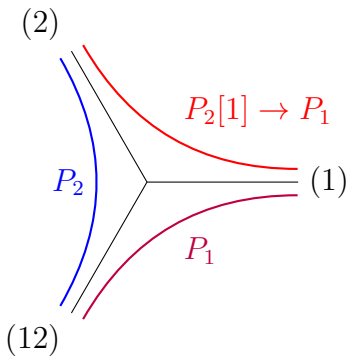


Figure 5.2: Geometric Picture of the Indecomposable Representations

Note that:

$$H^\bullet(\mathrm{Hom}(P_2, P_1)) \cong H^\bullet(\mathrm{Hom}(P_1, P_2[1] \to P_1)) \cong H^\bullet(P_2[1] \to P_1, P_2[1]) \cong k$$

While:

$$H^\bullet(\mathrm{Hom}(P_1, P_2)) \cong H^\bullet(\mathrm{Hom}(P_2, P_2[1] \to P_1)) \cong H^\bullet(P_2[1] \to P_1, P_1) \cong 0$$

As suggested by this computation, the sheaf $\mathcal{Q}_\mu$ encodes the data of the cyclic structure $\mathcal{O}_\mu$. That's the content of the following proposition:

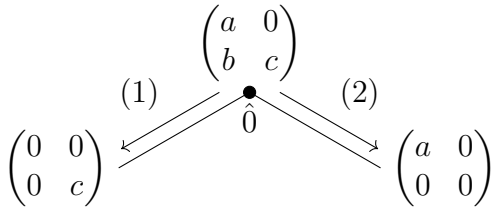
Proposition 5.4.1. *Let $\varphi \in Aut(\mathcal{P}_{A_2})$. Then $\varphi^*\mathcal{Q}_\mu \cong \mathcal{Q}_\mu$ if and only if $\overline{\varphi} \in \mathfrak{S}(K_{A_2})$ is an even permutation (that is, if and only if φ preserves the cyclic order $\mathcal{O}_\mu$).*

Proof. We prove 'only if' completely. The proof of 'if' is a consequence of the formalism of the remainder of the section, but we will sketch the argument here as well.

For 'Only If':, assume $\overline{\varphi}$ is the transposition (e_0e_2). The other cases are similar (or follow from 'if').

For any object $X \in \mathrm{Ob}(\mathrm{Rep}(A_2, \mu))$, we get a sheaf $\mathscr{H}om_{\mathcal{Q}_\mu}(X, X)$ of chain complexes over k (see appendix B for details). This gives a presheaf of graded algebras $\mathcal{E}^\bullet(X) = H^\bullet(\mathscr{H}om_{\mathcal{Q}_\mu}(X, X))$.

Let $X = P_1 \oplus P_2$. If $\varphi^*\mathcal{Q}_\mu$ were quasi-equivalent to $\mathcal{Q}_\mu$, there would be an object $Y \in \mathrm{Rep}(A_2, \mu)$ such that $\mathcal{E}^\bullet(Y) \cong \varphi^*\mathcal{E}^\bullet(X)$. Looking at the stalks of $\mathcal{E}^\bullet(X)$ we get:



(We do not need to consider the stalk at (12)). We know Y, if it exists is homotopic to a direct sum of indecomposables. By the stalks $\mathcal{E}^\bullet(Y)_{(1)} \cong \mathcal{E}^\bullet(Y)_{(2)} \cong k$, using Table 5.1 we see Y is either a shift of $P_2[1] \to P_1$, or a shift of $P_1 \oplus P_2[n]$. The stalk at $\{\hat{0}\}$ eliminates all but a shift of $P_1 \oplus P_2$,i.e. Y is homotopic to X.

But we do not have $\varphi^*\mathcal{E}^\bullet(X) \cong \mathcal{E}^\bullet(X)$: If we take products of ideals in $\mathcal{E}^\bullet(X)_{\hat{0}}$ we get $\ker(c_{(1)})\ker(c_{(2)}) = 0$ but $\ker(c_{(2)})\ker(c_{(1)}) \neq 0$, implying an asymmetry between $c_{(1)}$ and $c_{(2)}$.

For 'If' (Sketch): Assume φ is the permutation $(1) \mapsto (2) \mapsto (12) \mapsto (1)$. A quasi-equivalence $\mathcal{Q} \to \varphi^*\mathcal{Q}$ can be specified by a morphism of 2-functors, that is, a quasi-equivalence $F_{\hat{0}} : \mathrm{Rep}(A_2, \mu) \to \mathrm{Rep}(A_2, \mu)$, and for each correspondence $\mathfrak{p} : \star \to A_2$ a quasi-equivalence $F_{\mathfrak{p}} : \mathrm{Rep}(\star, \cdot) \to \mathrm{Rep}(\star, \cdot)$, such that the following diagram commutes up to homotopy:

$$\begin{array}{ccc} \mathrm{Rep}(A_2, \mu) & \xrightarrow{F_{\hat{0}}} & \mathrm{Rep}(A_2, \mu) \\ {\scriptstyle c_{\mathfrak{p}}}\downarrow & & \downarrow{\scriptstyle c_{\varphi(\mathfrak{p})}} \\ \mathrm{Rep}(\star, \cdot) & \xrightarrow{F_{\mathfrak{p}}} & \mathrm{Rep}(\star, \cdot) \end{array}$$

We define $F^0_{\hat{0}} : \mathrm{Rep}^0(A_2, \mu) \to \mathrm{Rep}(A_2, \mu)$ by $F^0_{\hat{0}}(P_1) = P_2[1] \longrightarrow P_1$, and $F^0_{\hat{0}}(P_2) = P_1$, and let $F_{\hat{0}} : \mathrm{Rep}(A_2, \mu) \to \mathrm{Rep}(A_2, \mu)$ denote the extension guaranteed by lemma 5.1.2. We also have: $F_{(12)} = F_{(2)}$ is the identity, while $F_{(1)}$ is the shift operator $X \mapsto X[1]$.

The reader can check that these maps do indeed satisfy the required properties. □

5.5 Reflection Functors

Now, recall the definition of the category of pointed trees $\mathbf{DirTree}^{\pm}$, whose objects are triples (T, μ, v), with $v \in \mathcal{V}^{\pm}(T, \mu)$.

Our goal of this section will be to 'categorify' the reflection functor $\mathcal{R}^{\pm} : \mathbf{DirTree}^{\pm} \to \mathbf{DirTree}^{\mp}$ described in Chapter 4. The machinery here is effectively an extention of the classical BGP reflection functors- see [2]- to this sheaf of categories.

Because of a small sign issue, we need to slightly modify the category:

Definition 5.5.1. *The category* ***DirTree***$^{\pm,s}$ *of* ***signed*** $\pm$ ***pointed directed trees*** *is the category whose objects are quadruples* (T, μ, v, s)*, where* s *is a* ***sign function*** $s : n(v) \to \{-1, +1\}$*, where* $n(v)$ *denotes the neighbors of* v*. When* $v = \hat{0}$ *is the null vertex, we set* $n(v) = \emptyset$*. When* $deg(v) = 2$*, we require that* s *assigns opposite signs to the two neighbors of* v*.*

Suppose $\mathfrak{p} = (R \xleftarrow{q} S \xhookrightarrow{i} T)$ *is a correspondence and* $v \in \mathcal{V}^+(T)$*, then there is a functorial injection* $j : n(w) \to n(v)$*. Specifically, if* $v \in V(S)$ *with* $q(v) = w$*, and* $q^{-1}(w) = \{v\}$*, then* j *is the unique function with* $q \circ j = id$*, and in all other cases* $n(w) = \emptyset$*. So a sign function* $s : n(v) \to \{-, +\}$ *pulls back to* $\mathfrak{p}^* s = s \circ j : n(w) \to \{+, -\}$*- this defines morphisms in* ***DirTree***$^{+,s}$ *as those which pullback sign functions correctly.*

Finally, we extend the functor $\mathcal{R}$ *by declaring that it reverses the sign function, i.e.* $\mathcal{R}(T, \mu, v, s) = (T, r_v\mu, v, \overline{s})$*, where* $\overline{s}(u) = -s(u)$*.*

By a 'categorification' of $\mathcal{R}^{\pm}$ involves the following data:

(i) For an object $(T, \mu, v, s) \in \mathbf{DirTree}^{\pm,s}$, we will construct a functor $\Gamma_v^{\pm} : \mathrm{Rep}(T, \mu) \to \mathrm{Rep}(T, r_v\mu)$. Note that $\Gamma_v^{\pm}$ will depend on μ and s, but we omit these from the notation to avoid clutter.

(ii) For a correspondence $\mathfrak{p} : (R, \lambda, w, s') \to (T, \mu, v, s)$, with $v \in \mathcal{V}^+(T, \mu)$, we will construct a quasi-isomorphism:

$$\gamma_{\mathfrak{p}}^+ : c_{r_v\mathfrak{p}} \circ \Gamma_v^+ \to \Gamma_w^+ \circ c_{\mathfrak{p}}$$

Satisfying, for $\mathfrak{p} : (R, \lambda, w, s') \to (T, \mu, v, s)$ and $\mathfrak{p}' : (R', \lambda', w', s'') \to (R, \lambda, w, s')$, the commutativity of:

$$\begin{array}{ccccc}
c_{r_w\mathfrak{p}'} \circ c_{r_v\mathfrak{p}} \circ \Gamma_v^+ & \xrightarrow{1_{c_{r_v\mathfrak{p}'}} \diamond \gamma_{\mathfrak{p}}^+} & c_{r_w\mathfrak{p}'} \circ \Gamma_w^+ \circ c_{\mathfrak{p}} & \xrightarrow{\gamma_{\mathfrak{p}'}^+ \diamond 1_{c_{\mathfrak{p}}}} & \Gamma_{w'}^+ \circ c_{\mathfrak{p}'} \circ c_{\mathfrak{p}} \\
\downarrow{\scriptstyle \eta_{r_v\mathfrak{p}, r_w\mathfrak{p}'} \diamond 1_{\Gamma_v^+}} & & & & \downarrow{\scriptstyle 1_{\Gamma_{w'}^+} \diamond \eta_{\mathfrak{p},\mathfrak{p}'}} \\
c_{r_v(\mathfrak{p}\circ\mathfrak{p}')} \circ \Gamma_v^+ & & \xrightarrow{\qquad\qquad\gamma_{\mathfrak{p}\circ\mathfrak{p}'}^+\qquad\qquad} & & \Gamma_{w'}^+ \circ c_{\mathfrak{p}\circ\mathfrak{p}'}
\end{array}$$

(iii) When $v \in \mathcal{V}^-(T,\mu)$, a quasi-isomorphism:

$$\gamma^-_{\mathfrak{p}} : \Gamma^-_w \circ c_{\mathfrak{p}} \to c_{r_v\mathfrak{p}} \circ \Gamma^-_v$$

Satisfying, for $\mathfrak{p} : (R,\lambda,w,s') \to (T,\mu,v,s)$ and $\mathfrak{p}' : (R',\lambda',w',s'') \to (R,\lambda,w,s')$, the commutativity of:

$$\begin{array}{ccccc}
c_{r_w\mathfrak{p}'} \circ c_{r_v\mathfrak{p}} \circ \Gamma^-_v & \xleftarrow{1_{c_{r_v\mathfrak{p}'}} \diamond \gamma^-_{\mathfrak{p}}} & c_{r_w\mathfrak{p}'} \circ \Gamma^-_w \circ c_{\mathfrak{p}} & \xleftarrow{\gamma^-_{\mathfrak{p}'} \diamond 1_{c_{\mathfrak{p}}}} & \Gamma^-_{w'} \circ c_{\mathfrak{p}'} \circ c_{\mathfrak{p}} \\
\downarrow{\scriptstyle \eta_{r_v\mathfrak{p},r_w\mathfrak{p}'} \diamond 1_{\Gamma^-_v}} & & & & \downarrow{\scriptstyle 1_{\Gamma^-_{w'}} \diamond \eta_{\mathfrak{p},\mathfrak{p}'}} \\
c_{r_v(\mathfrak{p}\circ\mathfrak{p}')} \circ \Gamma^-_v & & \xleftarrow{\qquad\qquad\gamma^-_{\mathfrak{p}\circ\mathfrak{p}'}\qquad\qquad} & & \Gamma^-_{w'} \circ c_{\mathfrak{p}\circ\mathfrak{p}'}
\end{array}$$

(iv) A quasi-isomorphism:

$$\varepsilon^+_v : \Gamma^-_v \circ \Gamma^+_v \to \mathrm{Id}$$

Satisfying, for $\mathfrak{p} : (R,\lambda,w,s') \to (T,\mu,v,s)$:

$$\begin{array}{ccc}
\Gamma^-_w \circ c_{r_v\mathfrak{p}} \circ \Gamma^+_v & \xrightarrow{\gamma^-_{\mathfrak{p}} \diamond 1_{\Gamma^+_v}} & c_{\mathfrak{p}} \circ \Gamma^-_v \circ \Gamma^+_v \\
\downarrow{\scriptstyle 1_{\Gamma^-_w} \diamond \gamma^+_{\mathfrak{p}}} & & \downarrow{\scriptstyle 1_{c_{\mathfrak{p}}} \diamond \varepsilon^+_v} \\
\Gamma^-_w \circ \Gamma^+_w \circ c_{\mathfrak{p}} & \xrightarrow{\varepsilon^+_w \diamond 1_{c_{\mathfrak{p}}}} & c_{\mathfrak{p}}
\end{array}$$

(v) A quasi-isomorphism:

$$\varepsilon^-_v : \mathrm{Id} \to \Gamma^+_v \circ \Gamma^-_v$$

Satisfying, for $\mathfrak{p} : (R,\lambda,w,s') \to (T,\mu,v,s)$:

$$\begin{array}{ccc}
c_{\mathfrak{p}} & \xrightarrow{1_{c_{\mathfrak{p}}} \diamond \varepsilon^-_v} & c_{\mathfrak{p}} \circ \Gamma^+_v \circ \Gamma^-_v \\
\downarrow{\scriptstyle \varepsilon^-_w \circ 1_{c_{\mathfrak{p}}}} & & \downarrow{\scriptstyle \gamma_{\mathfrak{p}^+} \diamond 1_{\Gamma^-_v}} \\
\Gamma^+_w \circ \Gamma^-_w \circ c_{\mathfrak{p}} & \xrightarrow{1_{\Gamma^+_w} \diamond \gamma^-_{\mathfrak{p}}} & \Gamma^+_w \circ c_{\mathfrak{p}} \circ \Gamma^-_v
\end{array}$$

If we temporarily let $\pi^{\pm} : \mathbf{DirTree}^{\pm,s} \to \mathbf{DirTree}$ denote the forgetful functors, the bullet points (i), (ii) define a left 2-functor from $\mathcal{Q} \circ \pi^+$ to $\mathcal{Q} \circ \mathcal{R}^+ \circ \pi^+$, and hence a morphism of sheaves $\mathcal{Q}_\mu \to \omega_v^* \mathcal{Q}_{r_v \mu}$ on the poset $\mathcal{P}_T$, whenever $(T, \mu, v, s) \in \mathrm{Ob}(\mathbf{DirTree}^{+,s})$.

Similarly, bullet point (iii) gives a morphism of sheaves $\mathcal{Q}_\mu \to \omega_v^* \mathcal{Q}_{r_v \mu}$ when $(T, \mu, v, s) \in \mathrm{Ob}(\mathbf{DirTree}^{-,s})$.

The bullet points (iv) and (v) show that these morphisms are quasi-inverses, hence are each quasi-equivalences. Hence as a result of this construction we will have:

Proposition 5.5.1. *Let (T, μ) be a directed tree, $v \in V(T)$ be a source or sink of degree ≤ 2, and $\omega_v : (\mathcal{P}_T, \mathcal{O}_\mu) \xrightarrow{\sim} (\mathcal{P}_T, \mathcal{O}_{r_v \mu})$ be the directed reflection isomorphism. Then there exists a quasi-equivalence:*

$$\mathcal{Q}_\mu \xrightarrow{\sim} \omega_v^* \mathcal{Q}_{r_v \mu}$$

We will delay this construction for a bit so that we can emphasize the most important result of this chapter.

5.6 The Cyclic Structure Theorem

Assuming proposition 5.5.1, we can show:

Theorem 5.6.1. *(**Cyclic Structure Theorem**) Let (T, μ) and (T', μ') be directed trees, and suppose $\varphi : \mathcal{P}_T \xrightarrow{\sim} \mathcal{P}_{T'}$ is an isomorphism. Then $\varphi^* \mathcal{Q}_{\mu'}$ is quasi-equivalent to $\mathcal{Q}_\mu$ if and only if $\varphi^* \mathcal{O}_{\mu'} = \mathcal{O}_\mu$.*

Proof. By 4.4.1, if $\varphi^* \mathcal{O}_{\mu'} = \mathcal{O}_\mu$, then φ can be written as a product of reflection isomorphisms and isomorphisms of the form i_σ, where σ is an isomorphism of directed trees. Then proposition 5.5.1 implies $\varphi^* \mathcal{Q}_{\mu'}$ is quasi-equivalent to $\mathcal{Q}_\mu$.

On the other hand, assume $\varphi^* \mathcal{Q}_{\mu'}$ is quasi-equivalent to $\mathcal{Q}_\mu$. Pick a codimension-one cell $[\mathfrak{p}] \in \mathcal{P}_T$, where $\mathfrak{p} : A_2 \to T$ is a correspondence, and let $[\mathfrak{p}'] = \varphi([\mathfrak{p}])$. We can arrange it so $\mathfrak{p}^* \mu = \mathfrak{p}'^* \mu' = \lambda$. Then we have a commutative diagram:

$$\begin{array}{ccc} \mathcal{P}_T & \xrightarrow{\varphi} & \mathcal{P}_{T'} \\ {\scriptstyle i_{\mathfrak{p}}}\uparrow & & \uparrow{\scriptstyle i_{\mathfrak{p}'}} \\ \mathcal{P}_{A_2} & \xrightarrow{\varphi_0} & \mathcal{P}_{A_2} \end{array}$$

By assumption, $\varphi_0^*\mathcal{Q}_\lambda \cong \varphi_0 i_{\mathfrak{p}'}^*\mathcal{Q}_{\mu'} \cong i_{\mathfrak{p}}^*\varphi^*\mathcal{Q}_{\mu'} \cong i_{\mathfrak{p}}^*\mathcal{Q}_\mu \cong \mathcal{Q}_\lambda$. So proposition 5.4.1 implies $\varphi_0^*\mathcal{O}_\lambda = \mathcal{O}_\lambda$. Since the codimension-one cell was arbitrary, $\varphi^*\mathcal{O}_{\mu'} = \mathcal{O}_\mu$. □

In particuar, if $(\mathcal{P}, \mathcal{O})$ is a cyclic C-Arb singularity, then we can define a sheaf of dg-categories $\mathcal{Q}$ on $\mathcal{P}$ by picking a directed tree (T, μ), an isomorphism $\varphi : \mathcal{P} \xrightarrow{\sim} \mathcal{P}_T$ with $\varphi^*\mathcal{O}_\mu = \mathcal{O}$, and setting $\mathcal{Q} = \varphi^*\mathcal{Q}_\mu$. By the Cyclic Structure Theorem, this sheaf of categories doesn't depend on the choice of (T, μ) and φ- it is determined up to quasi-equivalence by $\mathcal{O}$.

Remark 5.6.1. *While a cyclic structure determines a sheaf of categories up to quasi-equivalence, it does not determine a sheaf up to* natural *quasi-equivalence. The quasi-equivalence can be sensitive to the order in which the reflection functors are applied.*

As a concrete example of this phenomenon, consider the tree (A_2, μ), where μ points from vertex 1 to vertex 2. The isomorphism $r = \omega_1\omega_2 \in Aut(\mathcal{P}_{A_2})$ satisfies $r^*\mathcal{O}_\mu = \mathcal{O}_\mu$, and the above construction gives a quasi-equivalence $F : \mathcal{Q}_\mu \xrightarrow{\sim} r^*\mathcal{Q}_\mu$, which coincides with the functor described in the proof of 5.4.1. Then r^3 is the identity on $\mathcal{P}_{A_2}$, but F^3 is the shift functor $X \mapsto X[1]$. This grading ambiguity thwarts naive attempts to define a canonical sheaf associated to a cyclic structure, and is part of the motivation for the definition of a graded structure in Chapter 6.

The remainder of this chapter will be a construction of the reflection functors and related data, as explained in section 5.5. It can be skipped without consequence for the rest of the paper.

5.7 Construction of the Reflection Functors

We must proceed through steps $(i) - (v)$ described in section 5.5. As a preliminary step, we introduce the notion of twisted complexes as a model for the pretriangulated hull, which we will use for computational purposes.

Twisted Complexes

We discuss twisted complexes as a model for $\mathcal{C}^{pre-tr}$ when $\mathcal{C}$ is a $k-$linear category- that is, a dg-category in which all morphisms are in degree zero and all differentials are zero. For a more general notion, see (ref). Bondal-Kapranov.

Definition 5.7.1. *Let $\mathcal{C}$ be a $k-$linear category. $\mathcal{C}^\oplus$ is the graded $k-$linear category which contains formal shifts and finite direct sums of objects in $\mathcal{C}$ (including the zero object). A*

twisted complex *over $\mathcal{C}$ is a pair (C, d), where $C \in Ob(\mathcal{C}^{\oplus})$ and $d : C \to C$ is a degree-one morphism satisfying $d^2 = 0$. It can also be written as a sequence:*

$$\cdots \xrightarrow{d_{i-2}} C_{i-1} \xrightarrow{d_{i-1}} C_i \xrightarrow{d_i} C_{i+1} \xrightarrow{d_{i+1}} \cdots$$

Where each $C_i \in C^{\oplus}$ is an object concentrated in degree i, each d_i is a degree-one morphism, $d_{i+1}d_i = 0$, and only finitely many C_i are nonzero.

The category of twisted complexes is a dg-category in a straightforward way, and it admits a natural embedding of $\mathcal{C}$. We have:

Lemma 5.7.1. *The category of twisted complexes is a model for the dg-category $\mathcal{C}^{pre-tr}$. In particular, shifts and cones of twisted complexes can be computed in the usual way.*

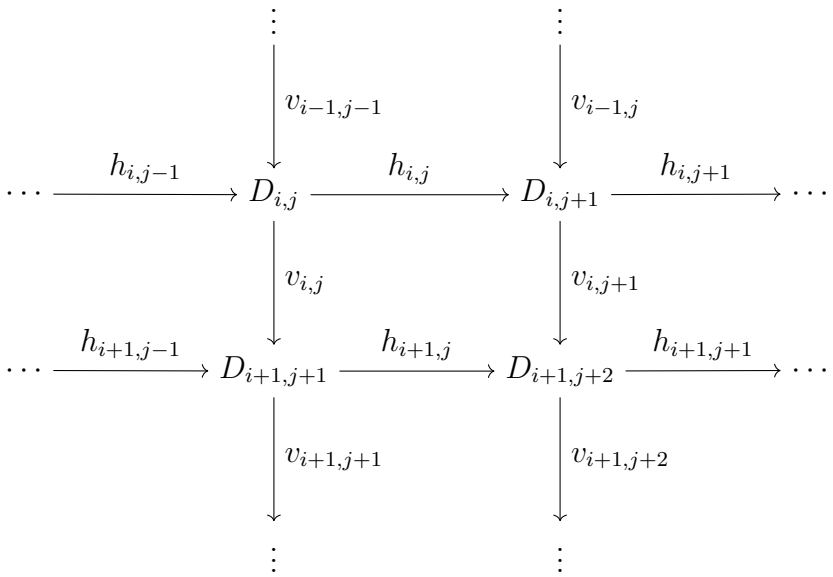
Suppose $\mathcal{C}$ and $\mathcal{D}$ are $k-$linear categories, and we have a functor $F : \mathcal{C} \to \mathcal{D}^{pre-tr}$. If we take a complex in $\mathcal{C}^{pre-tr}$:

$$C = \cdots \xrightarrow{f_{i-2}} C_{i-1} \xrightarrow{f_{i-1}} C_i \xrightarrow{f_i} C_{i+1} \xrightarrow{f_{i+1}} \cdots$$

The naive extension of F to C is a **double complex**. Write:

$$F(C_i) = \cdots \xrightarrow{h_{i,j-2}} D_{i,j-1} \xrightarrow{h_{i,j-1}} D_{i,j} \xrightarrow{h_{i,j}} D_{i,j+1} \xrightarrow{h_{i,j+1}} \cdots$$

Where the $D_{i,j} \in \mathcal{D}^{\oplus}$ are concentrated in degree j. Then applying F to the maps f_i yields a family of degree-one maps $v_{i,j} : D_{i,j} \to D_{i+1,j+1}$. We get an object:



Where we have $v^2 = h^2 = 0$ and $vh = hv$. Such an object is a double complex.

Definition 5.7.2. *The **totalization** of a double complex as in the above diagram is an object $T \in \mathcal{D}^{pre-tr}$:*

$$T = \cdots \xrightarrow{d_{j-2}} T_{j-1} \xrightarrow{d_{j-1}} T_j \xrightarrow{d_j} T_{j+1} \xrightarrow{h_{j+1}} \cdots$$

Where:

$$T_j = \oplus_i D_{i,j}$$

$$d_j = \sum_i v_{i,j} - (-1)^i h_{i,j}$$

In this way, we can compute the extension of F to a functor $\mathcal{C}^{pre-tr} \to \mathcal{D}^{pre-tr}$.

The Functors $\Gamma^{\pm}$

We now return to our goal at hand- defining the objects and verifying the relations outlined in steps $(i) - (v)$ earlier in this chapter.

Definition 5.7.3. *Let $(T, \mu, v, s) \in$ **DirTree**$^{\pm,s}$ We define a functor $\Gamma_v^{\pm} : Rep(T, \mu) \to Rep(T, r_v\mu)$ by first defining $\Gamma_v^{0,\pm} : Rep^0(T, \mu) \to Rep(T, r_v\mu)$, and extending it as in lemma 5.1.2. Let $\{P_w\}_{w \in V(T)}$ denote the objects in $Rep^0(T, \mu)$, and $\{\hat{P}_w\}_{w \in V(T)}$ the corresponding objects in $Rep^0(T, r_v\mu)$. Also, let $e_{w,u}$, $\hat{e}_{w,u}$ denote the morphisms.*

- *Set $\Gamma_v^{0,\pm}(P_w) = \hat{P}_w$ if $w \neq v$, and $\Gamma_v^{0,\pm}(e_{u,w}) = \hat{e}_{u,w}$ for $u, w \neq v$.*
- *If $v = \hat{0}$, we've completely defined $\Gamma_v^{0,\pm}$.*
- *If v is a **source**, then:*

$$\Gamma_v^{0,+}(P_v) = \hat{P}_v[1] \xrightarrow{d} \bigoplus_{w \in n(v)} \hat{P}_w$$

$$d = \sum_{w \in n(v)} s(w) \cdot i_w \hat{e}_{v,w}$$

Where i_w is the inclusion of the summand $\hat{P}_w$, and s is the sign function. (Note that v is a sink with respect to $r_v\mu$). If $n(v) = \emptyset$ (i.e. T has only one vertex) we let the empty direct sum equal zero. On morphisms, first $\Gamma_v^{0,+}(e_{v,v})$ is the identity on $\Gamma^{0,+}(P_v)$. Now suppose there exists a directed path from v to u in (T, μ), and $u \neq v$. There is a unique $w \in n(v)$ such that there is a directed path from w to u. Then:

$$\Gamma_v^{0,+}(e_{u,v}) = i_w \hat{e}_{u,w}$$

- *If v is a **sink**:*

$$\Gamma_v^{0,-}(P_v) = \bigoplus_{w \in n(v)} \hat{P}_w \xrightarrow{d} P_v[-1]$$

$$d = \sum_{w \in n(v)} s(w) \cdot \hat{e}_{w,v} \pi_w$$

Where π_w is the projection onto the $\hat{P}_w$ summand. Again, the empty sum is zero. On morphisms, again $\Gamma_v^{0,-}(e_{v,v})$ is the identity on $\Gamma_v^{0,-}(P_v)$. Then suppose there is a directed path from u to v in (T, μ), and $u \neq v$. There is a unique $w \in n(v)$ such that there is a directed path from u to w. Then:

$$\Gamma^{0,-}(e_{v,u}) = \hat{e}_{w,u} \pi_w$$

It is clear that these are well-defined and respect composition.

Definitions of $\gamma_{\mathfrak{p}}^{\pm}$

We've finished part (i) of the program outlined in the beginning of the section. We now explain the constructions of $\gamma_{\mathfrak{p}}^{\pm}$.

First assume $v \in \mathcal{V}^+(T)$. We will construct, for $\mathfrak{p} : (R, \lambda, w, s') \to (T, \mu, v, s)$:

$$\gamma_{\mathfrak{p}}^{0,+} : c_{r_v \mathfrak{p}}^0 \circ \Gamma_v^{0,+} \xrightarrow{\sim} \Gamma_w^{0,+} \circ c_{\mathfrak{p}}^0$$

Which is a quasi-isomorphism of functors. Then, by lemma 5.1.2, there is a unique extension of $\gamma_{\mathfrak{p}}^{0,+}$ to the desired quasi-isomorphism $\gamma_{\mathfrak{p}}^{+}$. Let $\mathfrak{p} = (R \xleftarrow{q} S \xhookrightarrow{i} T)$. There are a few easy cases, where the definition of γ^0 is clear- i.e., both sides are the same object (up to, perhaps, a re-indexing of a direct sum). These include:

- If, $u \neq v$, $u \in V(S)$ we have $c_{r_v \mathfrak{p}}^0 \circ \Gamma_v^{0,+}(P_u) = \Gamma_w^{0,+} \circ c_{\mathfrak{p}}^0(P_u) = \hat{P}_{q(u)}$.
- If, on the other hand, $u \neq v$ and $u \notin V(S)$, $c_{r_v \mathfrak{p}}^0 \circ \Gamma_v^{0,+}(P_u) = \Gamma_w^{0,+} \circ c_{\mathfrak{p}}^0(P_u) = 0$

We remark that at this point all that's left is to define:

$$\gamma_{\mathfrak{p}}^{0,+}(P_v) : c_{r_v \mathfrak{p}}^0 \circ \Gamma_v^{0,+}(P_v) \to \Gamma_w^{0,+} \circ c_{\mathfrak{p}}^0(P_v) \tag{5.1}$$

Continuing with the easy cases:

- If v nor any of its neighbors are in $V(S)$, both sides of (5.1) are zero.

- When $q(v) = w$ and $q^{-1}(w) = \{v\}$, both sides of (5.1) are:

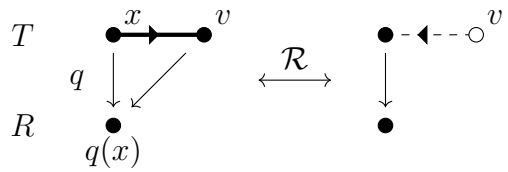
$$\hat{P}_w[1] \xrightarrow{d} \bigoplus_{u \in n(w)} \hat{P}_u$$

$$d = \Sigma_{u \in n(w)} s'(u) \cdot i_u \hat{e}_{w,u}$$

We now examine the remaining three nontrivial cases in detail. We remark that once we define $\gamma_{\mathfrak{p}}^{0,+}$ on objects, we have to verify that it is, in fact, a natural transformation.

Case I: v has one neighbor in $V(S)$, and either $v \notin V(S)$ or $v \in V(S)$ has the same image as its neighbor under q.

Near v, the correspondence looks as follows:



First suppose the left-hand-side is illustrating $\mathfrak{p}$, so $v \in V(S)$. Then:

$$\Gamma_{\hat{0}}^{0,+} \circ c_{\mathfrak{p}}^0(P_v) = c_{r_v\mathfrak{p}}^0 \circ \Gamma_v^{0,+}(P_v) = \hat{P}_{q(x)}$$

This is again an easy case. However, if the right-hand-side is illustrating $\mathfrak{p}$, so $v \notin V(S)$:

$$\Gamma_{\hat{0}}^{0,+} \circ c_{\mathfrak{p}}^0(P_v) = 0$$

$$c_{r_v\mathfrak{p}}^0 \circ \Gamma_v^{0,+}(P_v) = \hat{P}_{q(x)}[1] \xrightarrow{d} \hat{P}_{q(x)}$$

$$d = \pm \hat{e}_{q(x),q(x)}$$

The sign depending on the sign function s. We define $\gamma^{0,+}(P_v) = (\hat{P}_{q(x)}[1] \xrightarrow{d} \hat{P}_{q(x)}) \xrightarrow{0} 0$, which is an isomorphism in the homotopy category $H^\bullet(\mathrm{Rep}(R, \lambda))$.

Now, suppose $y \in V(S)$, and there exists a directed path from x to y in S (and hence also from v to y). $\gamma_{\mathfrak{p}}^{0,+}$ being a natural transformation requires the commutativity of the diagram:

$$\begin{array}{ccc} c_{r_v\mathfrak{p}}^0 \circ \Gamma_v^{0,+}(P_y) & \xrightarrow{c_{r_v\mathfrak{p}}^0 \circ \Gamma_v^{0,+}(e_{y,v})} & c_{r_v\mathfrak{p}}^0 \circ \Gamma_{(T,v,\mu)}^{0,+}(P_v) \\ \downarrow{\scriptstyle \gamma_{\mathfrak{p}}^{0,+}(P_y)} & & \downarrow{\scriptstyle \gamma_{\mathfrak{p}}^{0,+}(P_v)} \\ \Gamma_{\hat{0}}^{0,+} \circ c_{\mathfrak{p}}^0(P_y) & \xrightarrow{\Gamma_{\hat{0}}^{0,+} \circ c_{\mathfrak{p}}^0(e_{y,v})} & \Gamma_{\hat{0}}^{0,+} \circ c_{\mathfrak{p}}^0(P_v) \end{array}$$

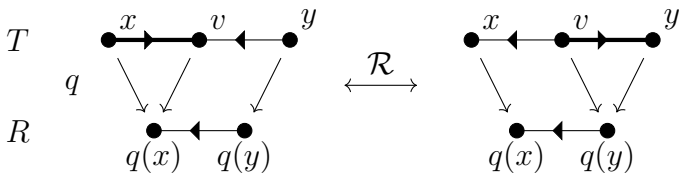
Which reduces to:

$$\begin{array}{ccc} \hat{P}_{q(y)} & \xrightarrow{\hat{e}_{q(y),q(x)}} & \hat{P}_{q(x)}[1] \to \hat{P}_{q(x)} \\ \Big\downarrow{\scriptstyle \hat{e}_{q(y),q(y)}} & & \Big\downarrow{\scriptstyle 0} \\ \hat{P}_{q(y)} & \xrightarrow{0} & 0 \end{array}$$

Note that this could not be commutative if we had instead tried to define $\gamma^+_{\mathfrak{p}} : \Gamma^+_w \circ c_{r_v\mathfrak{p}} \to c_{\mathfrak{p}} \circ \Gamma^+_v$- this is why the data Γ^+_v, γ^+ gives a *left* 2-functor. When v is a sink this same phenomenon forces Γ^-_v to be a *right* 2-functor.

Case II: v has two neighbors in $V(S)$, one of which has the same image as v under q.

Near v, the correspondence looks as follows:



We are assuming v is a source, so the right-hand-side is illustrating $\mathfrak{p}$ and the left-hand-side $r_v\mathfrak{p}$. Computing both sides of 5.1, we have:

$$\Gamma^{0,+}_{\hat{0}} \circ c^0_{\mathfrak{p}}(P_v) = \hat{P}_{q(y)}$$

$$c^0_{r_v\mathfrak{p}} \circ \Gamma^{0,+}_v(P_v) = \hat{P}_{q(x)}[1] \xrightarrow{d} (\hat{P}_{q(x)} \oplus \hat{P}_{q(y)})$$
$$d = s(x) \cdot i_{q(x)}\hat{e}_{q(x),q(x)} + s(y) \cdot i_{q(y)}\hat{e}_{q(x),q(y)}$$

In this case we set $\gamma^{0,+}_{\mathfrak{p}}(P_v) = \pi_{q(y)} + \hat{e}_{q(x),q(y)}\pi_{q(x)}$, which is a homotopy equivalence. This is the crucial point requiring the sign function: If we had $s(x) = s(y)$ then $\gamma^{0,+}_{\mathfrak{p}}(P_v)$ would not be a closed morphism (much less a homotopy equivalence).

As in case I, to ensure this is a natural transformation, we need to check commutativity of the diagrams, whenever $z \in V(S)$ $z \neq v$, and there exists a directed path from v to z in (T, μ):

$$\begin{array}{ccc} c^0_{r_v\mathfrak{p}} \circ \Gamma^{0,+}_v(P_z) & \xrightarrow{c^0_{r_v\mathfrak{p}} \circ \Gamma^{0,+}_v(e_{y,v})} & c^0_{r_v\mathfrak{p}} \circ \Gamma^{0,+}_v(P_v) \\ \Big\downarrow{\scriptstyle \gamma^{0,+}_{\mathfrak{p}}(P_z)} & & \Big\downarrow{\scriptstyle \gamma^{0,+}_{\mathfrak{p}}(P_v)} \\ \Gamma^0_{\hat{0}} \circ c^0_{\mathfrak{p}}(P_z) & \xrightarrow{\Gamma^{0,+}_{\hat{0}} \circ c^0_{\mathfrak{p}}(e_{y,v})} & \Gamma^{0,+}_{\hat{0}} \circ c^0_{\mathfrak{p}}(P_v) \end{array}$$

When z is to the left of v, the diagram becomes:

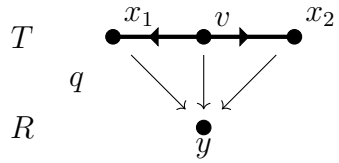
$$\begin{array}{ccc} \hat{P}_{q(z)} & \xrightarrow{i_{q(x)}\hat{e}_{q(z),q(x)}} & \hat{P}_{q(x)}[1] \to (\hat{P}_{q(x)} \oplus \hat{P}_{q(y)}) \\ \Big\downarrow{\scriptstyle \hat{e}_{q(z),q(z)}} & & \Big\downarrow{\scriptstyle \pi_{q(y)} + \hat{e}_{q(x),q(y)}\pi_{q(x)}} \\ \hat{P}_{q(z)} & \xrightarrow{\hat{e}_{q(z),q(y)}} & \hat{P}_{q(y)} \end{array}$$

And both compositions are $\hat{e}_{q(z),q(y)}$. Similarly, when z is to the right:

$$\begin{array}{ccc} \hat{P}_{q(z)} & \xrightarrow{i_{q(y)}\hat{e}_{q(z),q(y)}} & \hat{P}_{q(x)}[1] \to (\hat{P}_{q(x)} \oplus \hat{P}_{q(y)}) \\ \Big\downarrow{\scriptstyle \hat{e}_{q(z),q(z)}} & & \Big\downarrow{\scriptstyle \pi_{q(y)} + \hat{e}_{q(x),q(y)}\pi_{q(x)}} \\ \hat{P}_{q(z)} & \xrightarrow{\hat{e}_{q(z),q(y)}} & \hat{P}_{q(y)} \end{array}$$

Again, both compositions are $\hat{e}_{q(z),q(y)}$.

Case III: v has two neighbors $x_1, x_2 \in V(S)$, and $q(v) = q(x_1) = q(x_2)$. For brevity, write $q(x_1) = q(v) = q(x_2) = y$:



Here $r_v\mathfrak{p} = \mathfrak{p}$. Then:

$$\Gamma^{0,+}_{\hat{0}} \circ c_{\mathfrak{p}}(P_v) = \hat{P}_y$$

And:

$$c_{\mathfrak{p}} \circ \Gamma^{0,+}_{\hat{0}}(P_v) = \hat{P}_y \xrightarrow{d} \hat{P}^{(1)}_y \oplus \hat{P}^{(2)}_y$$

$$d = s(x_1) \cdot i^{(1)}_y \hat{e}_{y,y} + s(x_2) \cdot i^{(2)}_y \hat{e}_{y,y}$$

Where we use the superscripts to distinguish coordinates. Similarly to case II, we let $\gamma^{0,+}_{\mathfrak{p}}(P_v) = \pi^{(1)}_y + \pi^{(2)}_y$- the sign function ensures $\gamma^{0,+}_{\mathfrak{p}}(P_v)$ is a homotopy equivalence. For commutativity, if there exists a directed path from z to x_α in S, $\alpha = 1, 2$, the commutative diagram reduces to:

$$\begin{array}{ccc} \hat{P}_{q(z)} & \xrightarrow{i^{(\alpha)}_y \hat{e}_{q(z),y}} & \hat{P}_y[1] \to (\hat{P}^{(1)}_y \oplus \hat{P}^{(2)}_y) \\ \Big\downarrow{\scriptstyle \hat{e}_{q(z),q(z)}} & & \Big\downarrow{\scriptstyle \pi^{(1)}_y + \pi^{(2)}_y} \\ \hat{P}_{q(z)} & \xrightarrow{\hat{e}_{q(z),y}} & \hat{P}_y \end{array}$$

Regardless of α, both compositions are $\hat{e}_{q(z),y}$.

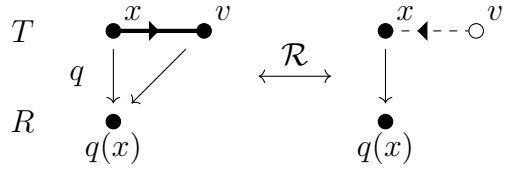
When v is a **sink**, much is the same. All of the cases which were 'easy' when v was a source are still that way. We do have to consider the two special cases. The analysis is mostly identical, though for completeness we include it (omitting some details). One crucial difference is the direction of γ^-, we are looking to construct:

$$\gamma^{0,-}_{\mathfrak{p}}(P_v) : \Gamma^{0,-}_w \circ c_{\mathfrak{p}}(P_v) \to c_{r_v \mathfrak{p}} \circ \Gamma^{0,-}_v(P_v)$$

The necessity for this can be seen by analyzing Case I:

Case I: v has one neighbor in $V(S)$, and either $v \notin V(S)$ or $v \in V(S)$ has the same image as its neighbor under q.

Near v, the correspondence looks as follows:



Suppose the left-hand-side is illustrating $\mathfrak{p}$, so $v \in V(S)$. Then:

$$\Gamma^{0,-}_{\hat{0}} \circ c^0_{\mathfrak{p}}(P_v) = c^0_{r_v \mathfrak{p}} \circ \Gamma^{0,-}_v(P_v) = \hat{P}_{q(x)}$$

If the right-hand-side is illustrating $\mathfrak{p}$, so $v \notin V(S)$:

$$\Gamma^{0,-}_{\hat{0}} \circ c^0_{\mathfrak{p}}(P_v) = 0$$

$$c^0_{r_v\mathfrak{p}} \circ \Gamma^0_v(P_v) = \hat{P}_{q(x)} \xrightarrow{d} \hat{P}_{q(x)}[-1]$$

$$d = \pm \hat{e}_{q(x),q(x)}$$

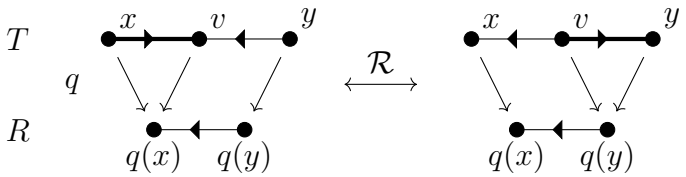
Again, the sign depends on the sign function $s(x)$. $\gamma^{0,-}_{\mathfrak{p}}(P_v) = 0 \xrightarrow{0} (\hat{P}_{q(x)} \xrightarrow{d} \hat{P}_{q(x)}[-1])$, which is a homotopy equivalence. Now, suppose $y \in V(S)$, and there exists a directed path from y to x, hence also from y to v. We need commutativity of the diagram:

$$\begin{array}{ccc} c^0_{r_v\mathfrak{p}} \circ \Gamma^{0,-}_v(P_v) & \xrightarrow{c^0_{r_v\mathfrak{p}} \circ \Gamma^{0,-}_v(e_{v,y})} & c^0_{r_v\mathfrak{p}} \circ \Gamma^{0,-}_v(P_y) \\ \uparrow{\scriptstyle \gamma^{0,-}_{\mathfrak{p}}(P_y)} & & \uparrow{\scriptstyle \gamma^{0,-}_{\mathfrak{p}}(P_v)} \\ \Gamma^{0,-}_{\hat{0}} \circ c^0_{\mathfrak{p}}(P_v) & \xrightarrow{\Gamma^{0,-}_{\hat{0}} \circ c^0_{\mathfrak{p}}(e_{v,y})} & \Gamma^{0,-}_{\hat{0}} \circ c^0_{\mathfrak{p}}(P_y) \end{array}$$

Which reduces to:

$$\begin{array}{ccc} \hat{P}_{q(x)} \to \hat{P}_{q(x)}[-1] & \xrightarrow{\hat{e}_{q(x),q(y)}} & \hat{P}_{q(y)} \\ \uparrow{\scriptstyle 0} & & \uparrow{\scriptstyle \hat{e}_{q(y),q(y)}} \\ 0 & \xrightarrow{0} & \hat{P}_{q(y)} \end{array}$$

Case II: Near v, the correspondence looks as follows:



Now we assume the left-hand-side illustrates $\mathfrak{p}$. We have:

$$\Gamma^{0,-}_{\hat{0}} \circ c^0_{\mathfrak{p}}(P_v) = \hat{P}_{q(x)}$$

$$c^0_{r_v\mathfrak{p}} \circ \Gamma^{0,-}_v(P_v) = (\hat{P}_{q(x)} \oplus \hat{P}_{q(y)}) \xrightarrow{d} \hat{P}_{q(y)}[-1]$$

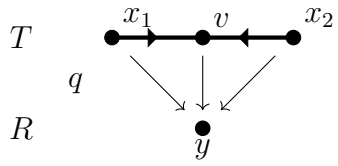
$$d = s(x) \cdot \hat{e}_{q(x),q(y)} \pi_{q(x)} + s(y) \cdot \hat{e}_{q(y),q(y)} \pi_{q(y)}$$

We set:

$$\gamma^{0,-}_{\mathfrak{p}}(P_v) = i_{q(x)} + i_{q(y)} \hat{e}_{q(x),q(y)}$$

Again the sign function ensures $\gamma^{0,-}$ is closed. Verifying the associated commutative diagrams is just as before and is left to the reader.

Case III: Near v the correspondence looks like:



We have:

$$\Gamma^{0,-}_{\hat{0}} \circ c^0_{\mathfrak{p}}(P_v) = \hat{P}_y$$

$$c^0_{\mathfrak{p}} \circ \Gamma^{0,-}_v(P_v) = \hat{P}^{(1)}_y \oplus \hat{P}^{(2)}_y \xrightarrow{d} \hat{P}_y[-1]$$

$$d = s(x_1) \cdot e_{y,y} \pi^{(1)}_y + s(x_2) \cdot e_{y,y} \pi^{(2)}_y$$

We set:

$$\gamma^{-}_{\mathfrak{p}}(P_v) = i^{(1)}_y + i^{(2)}_y$$

And again we leave the commutative diagram to the reader.

Compatibility across Compositions of Restrictions

We now verify the commutative diagrams in (ii) and (iii).

First, let $\mathfrak{p} : (R, \lambda, w, s') \to (T, \mu, v, s)$ and $\mathfrak{p}' : (R', \lambda', w', s'') \to (R, \lambda, w, s')$ be correspondences in $\mathbf{DirTree}^{+,s}$. We require:

$$\begin{array}{ccccc}
c_{r_w \mathfrak{p}'} \circ c_{r_v \mathfrak{p}} \circ \Gamma^+_v & \xrightarrow{1_{c_{r_v \mathfrak{p}'}} \diamond \gamma^+_{\mathfrak{p}}} & c_{r_w \mathfrak{p}'} \circ \Gamma^+_w \circ c_{\mathfrak{p}} & \xrightarrow{\gamma^+_{\mathfrak{p}'} \diamond 1_{c_{\mathfrak{p}}}} & \Gamma^+_{w'} \circ c_{\mathfrak{p}'} \circ c_{\mathfrak{p}} \\
\downarrow{\scriptstyle \eta_{r_v \mathfrak{p}, r_w \mathfrak{p}'} \diamond 1_{\Gamma^+_v}} & & & & \downarrow{\scriptstyle 1_{\Gamma^+_{w'}} \diamond \eta_{\mathfrak{p},\mathfrak{p}'}} \\
c_{r_v(\mathfrak{p} \circ \mathfrak{p}')} \circ \Gamma^+_v & & \xrightarrow{\gamma^+_{\mathfrak{p} \circ \mathfrak{p}'}} & & \Gamma^+_{w'} \circ c_{\mathfrak{p} \circ \mathfrak{p}'}
\end{array}$$

If we are in $\mathbf{DirTree}^{-,s}$, the diagram is:

$$\begin{array}{ccccc}
c_{r_w\mathfrak{p}'} \circ c_{r_v\mathfrak{p}} \circ \Gamma_v^- & \xleftarrow{1_{c_{r_v\mathfrak{p}'}} \diamond \gamma_{\mathfrak{p}}^-} & c_{r_w\mathfrak{p}'} \circ \Gamma_w^- \circ c_{\mathfrak{p}} & \xleftarrow{\gamma_{\mathfrak{p}'}^- \diamond 1_{c_{\mathfrak{p}}}} & \Gamma_{w'}^- \circ c_{\mathfrak{p}'} \circ c_{\mathfrak{p}} \\
\Big\downarrow{\scriptstyle \eta_{r_v\mathfrak{p}, r_w\mathfrak{p}'} \diamond 1_{\Gamma_v^-}} & & & & \Big\downarrow{\scriptstyle 1_{\Gamma_{w'}^-} \diamond \eta_{\mathfrak{p},\mathfrak{p}'}} \\
c_{r_v(\mathfrak{p}\circ\mathfrak{p}')} \circ \Gamma_v^- & & \xleftarrow{\gamma_{\mathfrak{p}\circ\mathfrak{p}'}^-} & & \Gamma_{w'}^- \circ c_{\mathfrak{p}\circ\mathfrak{p}'}
\end{array}$$

It is enough to verify this diagram commutes on each generator $\{P_u\}_{u \in V(T)}$. As before, when $u \neq v$ both directions yield the identity or the zero map trivially.

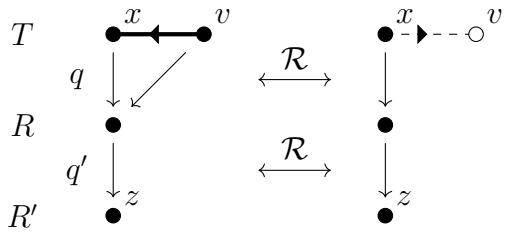
Write $\mathfrak{p} = (R \overset{q}{\twoheadleftarrow} S \overset{i}{\hookrightarrow} T)$ and $\mathfrak{p}' = (R' \overset{q'}{\twoheadleftarrow} S' \overset{i'}{\hookrightarrow} R)$. Note that if $v \notin V(q^{-1}(S'))$, then the bottom-right corner of the + diagram is zero, meaning both compositions are zero. The same is true of the top-right corner of the − diagram.

If $v \in V(q^{-1}(S'))$, and $(q' \circ q)^{-1}(v) = \{v\}$, then all objects in the diagram are the same (up to perhaps a reindexing of a direct sum) and all the maps are the canonical isomorphism realizing this.

We now check a few nontrivial cases. We will only check the cases when v is a source, the cases where v is a sink are similar. We assume $\mathfrak{p}$ is represented on the left of each diagram.

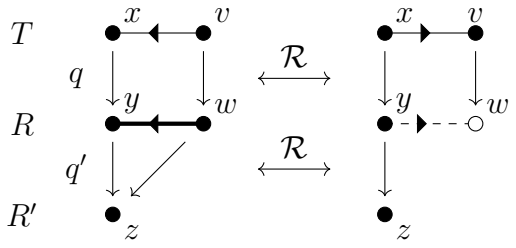
Cases I: v has one neighbor x in $V(q^{-1}(S'))$, and $q'(q(v)) = q'(q(x))$.

Case I.a:



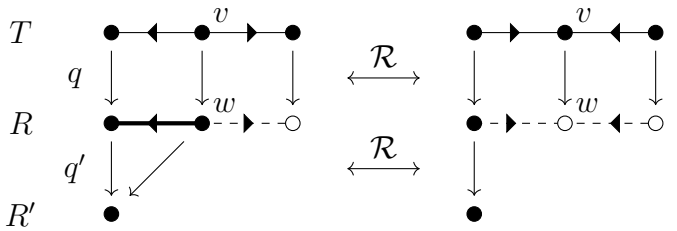
It is easy to see all objects are $\hat{P}_z$ and all maps are $\hat{e}_{z,z}$ in the commutative diagram.

Case I.b:



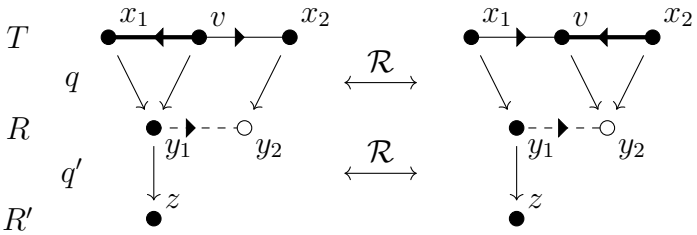
All objects are $\hat{P}_y$. Looking at the top-left horizontal arrow of the commutative diagram: $\gamma^+_{\mathfrak{p}}(P_v) : c_{r_v\mathfrak{p}} \circ \Gamma^+_v(P_v) \to \Gamma^+_w \circ c_{\mathfrak{p}}(P_v)$ is the identity map on $\hat{P}_w[1] \to \hat{P}_y$, and applying $c_{r_w\mathfrak{p}'}$ yields the identity on $\hat{P}_y$. All other maps are easily $\hat{e}_{y,y}$.

Case I.c:



This works identically to case I.b.

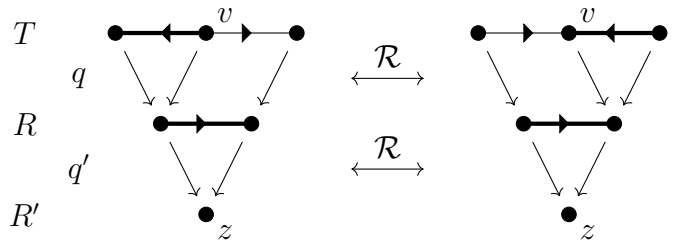
Case I.d:



All objects are $\hat{P}_z$. Looking at the top-left horizontal arrow: $\gamma^+_{\mathfrak{p}}(P_v) : c_{r_v\mathfrak{p}} \circ \Gamma^+_v(P_v) \to \Gamma^+_w \circ c_{\mathfrak{p}}(P_v)$ is the map $(\hat{P}_{y_2} \to \hat{P}_{y_2} \oplus \hat{P}_{y_1}) \to \hat{P}_{y_1}$ given by $\pi_{y_1} + e_{y_2,y_1}\pi_{y_2}$. Applying $c_{r_w\mathfrak{p}'}$ gives the identity $\hat{e}_{z,z}$. All other maps are clearly $\hat{e}_{z,z}$.

Cases II: v has two neighbors x_1, x_2 in $V(q^{-1}(S'))$, and $q'(q(x_1)) = q'(q(v))$, which can either equal $q'(q(x_2))$ or not.

Case II.a:

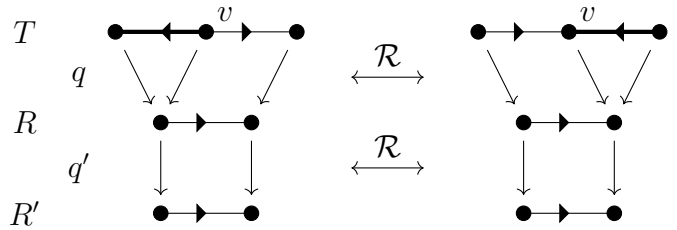


The diagram reduces to:

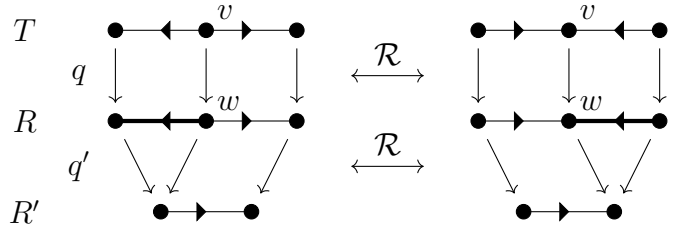
$$\begin{array}{ccccc} \hat{P}_z[1] \to \hat{P}_z^{(1)} \oplus \hat{P}_z^{(2)} & \xrightarrow{\pi_z^{(1)} + \pi_z^{(2)}} & \hat{P}_z & \longrightarrow & \hat{P}_z \\ \downarrow & & & & \downarrow \\ \hat{P}_z[1] \to \hat{P}_z^{(1)} \oplus \hat{P}_z^{(2)} & & \xrightarrow{\pi_z^{(1)} + \pi_z^{(2)}} & & \hat{P}_z \end{array}$$

Where all unlabelled arrows are the identity map, which commutes.

Case II.b:



Case II.c:



Verification of Case II.b and II.c is straightforward and left to the reader.

Case II.d: Either $q(x_1) = q(v) = q(x_2)$, or $q(x_1) \neq q(v) \neq q(x_2)$ and $q'(q(x_1)) = q'(q(v)) = q'(q(x_2))$. In each of these cases, each correspondence is fixed by $\mathcal{R}$.

Defining $\varepsilon_v^{\pm}$

We now come to the final parts of our program: (iv) and (v). We will construct two natural quasi-isomorphisms:

$$\varepsilon_v^+ : \Gamma_v^- \circ \Gamma_v^+ \to \mathrm{Id}$$
$$\varepsilon_v^- : \mathrm{Id} \to \Gamma_v^+ \circ \Gamma_v^-$$

As before, we define $\varepsilon_v^{\pm,0}$ on $\mathrm{Rep}^0(T,\mu)$. We see that the only nontrivial part of the definition is $\varepsilon_v^{\pm,0}(P_v)$.

As a shorthand, write:

$$P_{n(v)} := \bigoplus_{w \in n(v)} P_w$$

First, let (T,μ,v,s) be an object in $\mathbf{DirTree}^+$, with $v \neq \hat{0}$. We have:

$$\Gamma_v^- \circ \Gamma_v^+(P_v) = \Gamma_v^-(\hat{P}_v[1] \xrightarrow{d} \hat{P}_{n(v)})$$

This computation can be carried out using the *totalization* of a complex (definition 5.7.2):

$$\Gamma_v^- \begin{pmatrix} \hat{P}_v[1] \\ \downarrow d \\ \hat{P}_{n(v)} \end{pmatrix} = \mathrm{Tot} \begin{pmatrix} P_{n(v)}[1] & \xrightarrow{d'} & P_v \\ & \searrow^{\Gamma_v^-(d)} & \\ & P_{n(v)} & \end{pmatrix} = \begin{matrix} P_{n(v)}[1] \\ \downarrow d''=d'+\Gamma_v^-(d) \\ P_v \oplus P_{n(v)} \end{matrix}$$

We have:

$$\Gamma_v^-(d) = \sum_{w \in n(v)} s(w) \cdot i_w \Gamma_v^-(\hat{e}_{v,w}) = \sum_{w \in n(v)} s(w) \cdot i_w \pi_w$$

$$d' = \sum_{w \in n(v)} \overline{s}(w) \cdot i_v e_{w,v} \pi_w$$

And so:

$$d'' = \sum_{w \in n(v)} s(w) \cdot i_w \pi_w + \sum_{w \in n(v)} \overline{s}(w) \cdot i_v e_{w,v} \pi_w$$

Define:

$$\varepsilon_v^{+,0}(P_v) = \pi_v + \sum_{w \in n(v)} e_{w,v} \pi_w$$

$\varepsilon^{+,0}(P_v)$ is a homotopy equivalence- we remark that $\varepsilon^{+,0}(P_v)$ being closed depends on our convention that $\overline{s}(w) = -s(w)$.

We need to show that $\varepsilon_v^{+,0}$ is, in fact, a natural transformation. So take $x \neq v \in V(T)$ such that there exists a directed path from v to x in T. We need to check the commutativity of:

$$\begin{array}{ccc} \Gamma_v^- \circ \Gamma_v^+(P_x) & \xrightarrow{\Gamma_v^- \circ \Gamma_v^+(e_{x,v})} & \Gamma_v^- \circ \Gamma_v^+(P_v) \\ \Big\downarrow \varepsilon^{+,0}(P_x) & & \Big\downarrow \varepsilon^{+,0}(P_v) \\ P_x & \xrightarrow{e_{x,v}} & P_v \end{array}$$

There is a unique $u \in n(v)$ such that there is a path from u to x. The diagram becomes:

$$\begin{array}{ccc} P_x & \xrightarrow{i_u e_{x,u}} & P_{n(v)}[1] \to P_v \oplus P_{n(v)} \\ \Big\downarrow e_{x,x} & & \Big\downarrow \pi_v + \Sigma e_{w,v}\pi_w \\ P_x & \xrightarrow{e_{x,v}} & P_v \end{array}$$

And we see both compositions equal $e_{x,v}$.

Next, let (T, μ, v, s) be an object in $\mathbf{DirTree}^-$, with $v \neq \hat{0}$. A similar computation gives us:

$$\Gamma_v^+ \circ \Gamma_v^-(P_v) = P_v \oplus P_{n(v)} \xrightarrow{d} P_{n(v)}[-1]$$

$$d = \sum_{w \in n(v)} s(w) \cdot i_w e_{w,w} \pi_w + \sum_{w \in n(v)} \overline{s}(w) i_w e_{v,w} \pi_v$$

Define:

$$\varepsilon_v^-(P_v) = i_v + \sum_{w \in n(v)} i_w e_{v,w}$$

For the requisite commutative diagram: Assume there is a directed path from x to v in T. We check the commutativity of:

$$\begin{array}{ccc} P_v & \xrightarrow{e_{v,x}} & P_x \\ \Big\downarrow \varepsilon_v^{-,0}(P_v) & & \Big\downarrow \varepsilon_v^{-,0}(P_x) \\ \Gamma_v^+ \circ \Gamma_v^-(P_v) & \xrightarrow{\Gamma_v^+ \circ \Gamma_v^-(e_{v,x})} & \Gamma_v^+ \circ \Gamma_v^-(P_x) \end{array}$$

Which becomes:

$$\begin{array}{ccc} P_v & \xrightarrow{\quad e_{v,x} \quad} & P_x \\ \Big\downarrow{\scriptstyle i_v + \sum i_w e_{v,w}} & & \Big\downarrow{\scriptstyle e_{x,x}} \\ P_v \oplus P_{n(v)} \to P_{n(v)}[-1] & \xrightarrow{\quad e_{u,v}\pi_u \quad} & P_x \end{array}$$

And we see both compositions equal $e_{v,x}$.

Hence these definitions give natural quasi-isomorphisms $\varepsilon_v^{\pm}$.

Compatibility of $\varepsilon_v^{\pm}$

The last step is to check the commutativity diagrams in parts (iv) and (v) of our program.

Let $\mathfrak{p} : (R, \lambda, w, s') \to (T, \mu, v, s)$ be a correspondence. We need to check the commutativity of:

$$\begin{array}{ccc} \Gamma_w^- \circ c_{r_v\mathfrak{p}} \circ \Gamma_v^+ & \xrightarrow{\quad \gamma_{\mathfrak{p}}^- \diamond 1_{\Gamma_v^+} \quad} & c_{\mathfrak{p}} \circ \Gamma_v^- \circ \Gamma_v^+ \\ \Big\downarrow{\scriptstyle 1_{\Gamma_w^-} \diamond \gamma_{\mathfrak{p}}^+} & & \Big\downarrow{\scriptstyle 1_{c_{\mathfrak{p}}} \diamond \varepsilon_v^+} \\ \Gamma_w^- \circ \Gamma_w^+ \circ c_{\mathfrak{p}} & \xrightarrow{\quad \varepsilon_w^+ \diamond 1_{c_{\mathfrak{p}}} \quad} & c_{\mathfrak{p}} \end{array}$$

As in all cases so far, we only need to check this when applied to P_v. Writing $\mathfrak{p} = (R \xleftarrow{q} S \xhookrightarrow{i} T)$, if $v \notin V(S)$ the bottom-right corner of the diagram is 0, hence we get commutativity automatically. The other cases will take some care.

Case I: $q^{-1}(q(v)) = \{v\}$. Then each object is $\Gamma_w^- \circ \Gamma_w^+(P_w)$, with the right and bottom arrows being ε_w^+. Furthermore, $\gamma_\rho^{\pm}$ are canonical isomorphisms on each object, which implies that the top and left arrows are the identity.

Case II: v has one neighbor x in $V(S)$, and $q(v) = q(x) = z$.

Then $w = \hat{0}$ and the bottom arrow is just the identity $e_{z,z}$. Also, the map $\gamma_{\mathfrak{p}}^+(P_v)$ is $\hat{e}_{z,z}$, hence the left arrow is $e_{z,z}$. The right arrow is:

$$c_{\mathfrak{p}}\left[\ (P_x[1] \to P_x \oplus P_v) \xrightarrow{\ e_{x,v}\pi_x + \pi_v\ } P_v\ \right]$$

$$= \ (P_z[1] \to P_z^{(1)} \oplus P_z^{(2)}) \xrightarrow{\pi_z^{(1)}+\pi_z^{(2)}} P_z$$

The top arrow is the map $\gamma_{\mathfrak{p}}^-(\hat{P}_v[1] \to \hat{P}_x)$. Note that we have:

$$\gamma_{\mathfrak{p}}^-(\hat{P}_v) = 0 \xrightarrow{0} (P_z[1] \to P_z)$$

And

$$\gamma_{\mathfrak{p}}^-(\hat{P}_x) = P_z \xrightarrow{e_{z,z}} P_z$$

So the top arrow is:

$$P_z \xrightarrow{i_z^{(1)}} (P_z[1] \to P_z^{(1)} \oplus P_z^{(2)})$$

Putting this together yields:

$$\begin{array}{ccc} P_z & \xrightarrow{\quad i_z^{(1)} \quad} & P_z[1] \to P_z^{(1)} \oplus P_z^{(2)} \\ \Big\downarrow{\scriptstyle e_{z,z}} & & \Big\downarrow{\scriptstyle \pi_z^{(1)} + \pi_z^{(2)}} \\ P_z & \xrightarrow{\quad e_{z,z} \quad} & P_z \end{array}$$

Which certainly commutes.

Case III: v has two neighbors x, y in $V(S)$, and $q(x) = q(v) = z$.

Write $u = q(y)$, this demonstration will not depend on whether $u = z$ or not.

We have $w = \hat{0}$, and the bottom arrow is $e_{z,z}$. Also, the left arrow is effectively $\gamma_{\mathfrak{p}}^+$, i.e. it is the map: $(P_u[1] \to P_z \oplus P_u) \to P_z$ given by $\pi_z + e_{u,z}\pi_u$.

The right arrow is:

$$c_{\mathfrak{p}}\left[\ (P_{n(v)}[1] \to P_v \oplus P_{n(v)}) \xrightarrow{e_{x,v}\pi_x + e_{y,v}\pi_y + \pi_v} P_v \ \right]$$

$$= \ (P_{zu}[1] \to P_z^{(1)} \oplus P_z^{(2)} \oplus P_u) \xrightarrow{\pi_z^{(1)}+\pi_z^{(2)}+e_{u,z}\pi_u} P_z$$

And the top arrow is again obtained using:

$$\gamma_{\mathfrak{p}}^-(\hat{P}_v) = \ P_u \xrightarrow{e_{u,u}+i_z e_{u,z}} (P_u \oplus P_z \to P_z[-1])$$

$$\gamma_{\mathfrak{p}}^-(\hat{P}_x \oplus \hat{P}_y) = P_z \oplus P_u \xrightarrow{Id} P_z \oplus P_u$$

Hence the top arrow is:

$$\begin{array}{ccc} P_u[1] & \xrightarrow{e_{u,u} + i_z e_{u,z}} & P_{uz}[1] \\ \downarrow & & \downarrow \\ P_z \oplus P_u & \xrightarrow{i_z^{(1)} e_{z,z} \pi_z + i_u e_{u,u} \pi_u} & P_z^{(1)} \oplus P_z^{(2)} \oplus P_u \end{array}$$

For each case, we need to verify two things: First, that these are actually natural transformations (i.e. that they commute with morphisms), and second that they are compatible with the restriction functors.

In the case where v is a source, suppose $x \in V(T)$ with $x \neq v$ and suppose there is a directed path from v to x in μ. Let $w \in n(v)$ be the vertex such that there exists a directed path from w to x. We need to check commutativity of the diagram, but it is enough to check commutativity in degree zero, since the bottom-right entry only has degree zero components. We get:

$$\begin{array}{ccc} P_z \oplus P_u & \xrightarrow{i_z^{(1)} \pi_z + i_u \pi_u} & P_z^{(1)} \oplus P_z^{(2)} \oplus P_u \\ \pi_z + e_{u,z} \pi_u \downarrow & & \downarrow \pi_z^{(1)} + \pi_z^{(2)} + e_{u,z} \pi_u \\ P_z & \xrightarrow{e_{z,z}} & P_z \end{array}$$

Which, again, commutes.

Chapter 6

Graded Structures and Graded Branes

6.1 Graded Structures

Definition of Graded Structures

Definition 6.1.1. *Let $\mathcal{P}$ be an arboreal poset. A **graded structure** G on $\mathcal{P}$ consists of the following data:*

(i) For each top cell x, a $\mathbb{Z}-$torsor G_x.

*(ii) For each triple x, y, z where x, y are top-cells and z is a codimension-one cell with $z \leq x, y$, an isomorphism $\varphi^z_{x,y} : G_x \xrightarrow{\sim} G_y$. These are called the **structure maps**.*

Satisfying the Monodromy Relations. *Any automorphism φ of G_x is of the form $a \mapsto a + n$, for some $n \in \mathbb{Z}$. Define the **monodromy** of such an automorphism to be $m(\varphi) = n$. Then:*

(M1) $m(\varphi^z_{x,y}\varphi^z_{y,x}) = 0$.

(M2) If x, y, z are the three top-cells containing a codimension-one cell t, $m(\varphi^t_{z,x}\varphi^t_{y,z}\varphi^t_{x,y}) = \pm 1$ (Figure 6.1).

(M3) Suppose w is a codimension-two cell, and x_1, x_2, x_3, z_1, z_2, z_3 are as in Figure 6.2. Then: $m(\varphi^{z_1}_{x_3,x_1}\varphi^{z_3}_{x_2,x_3}\varphi^{z_2}_{x_1,x_2}) = 0$

For the remainder of this section we will need some simple facts about the monodromy function:

Lemma 6.1.1.

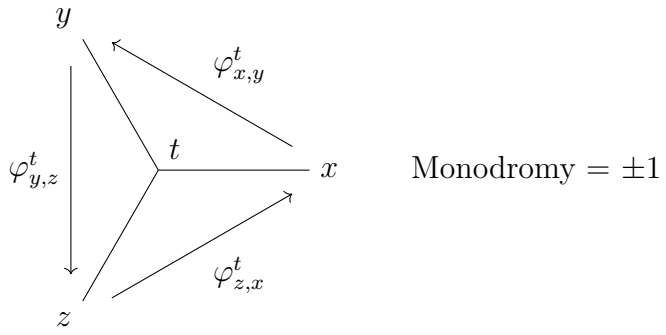


Figure 6.1: Illustration of Mondromy Axiom (M2)

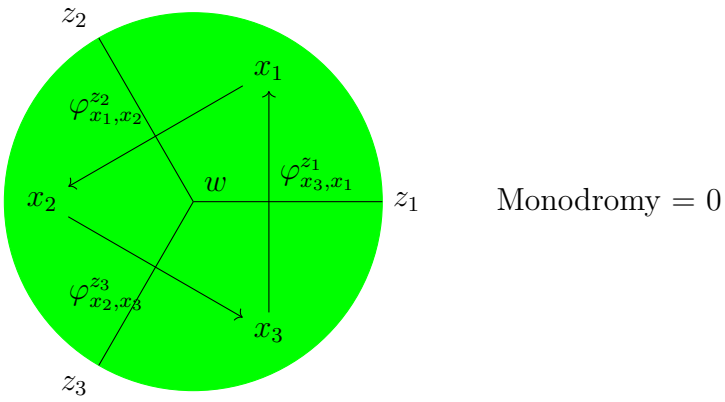


Figure 6.2: Illustration of Monodromy Axiom (M3)

- *If G_0 is a $\mathbb{Z}-$torsor, the map $m : Aut(G_0) \to \mathbb{Z}$ is a group isomorphism.*
- *If G_0, G_1 are $\mathbb{Z}-$torsors, and $a : G_0 \to G_1$ and $b : G_1 \to G_0$ are isomorphisms, then $m(ab) = m(ba)$.*

Proof. Straightforward □

We note some simple consequences of the definition. First, $(M1)$ implies that $\varphi^z_{x,y} = (\varphi^z_{y,x})^{-1}$. Hence if t, x, y, z are as in $(M2)$, and $m(\varphi^t_{z,x}\varphi^t_{y,z}\varphi^t_{x,y}) = 1$, then we have $m(\varphi^t_{y,x}\varphi^t_{z,y}\varphi^t_{x,z}) = -1$. Also, if we have $x_1, x_2, x_3, z_1, z_2, z_3$ as in $(M3)$, then $\varphi^{z_3}_{x_2,x_3}\varphi^{z_2}_{x_1,x_2} = \varphi^{z_1}_{x_1,x_3}$.

Definition 6.1.2. *Let $\mathcal{P}$ be an arboreal poset with graded structure G. The associated pre-cyclic structure $\mathcal{O}_G$ is defined as follows: Let t be a codimension-one cell, and x, y, z the three top-cells containing it. Then the cyclic order is induced by $x < y < z$ if and only if $m(\varphi^t_{z,x}\varphi^t_{y,z}\varphi^t_{x,y}) = -1$.*

Referring to Figure 6.1: If the cyclic order is the counterclockwise order, the monodromy of the illustrated map is -1.

Proposition 6.1.1. *For any graded structure G, $\mathcal{O}_G$ is a cyclic structure.*

Proof. By theorem 3.3.1, it is enough to show that $\mathcal{O}_G$ satisfies the A_3 and S_4 coherence conditions. First we show A_3- so let G be a graded structure on $\mathcal{P}_{A_3}$. Recall the setup for the arrow diagram in Chapter 3 (see Figures 3.4 and 3.5), where we use letters to abbreviate the two-cells for convenience (Figure 6.3).

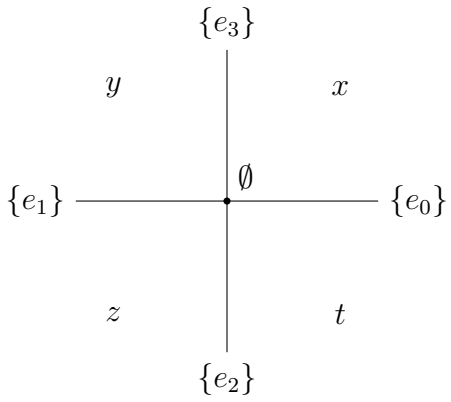


Figure 6.3: Arrow Diagram Setup With Labelled Cells

We will compute $m(\varphi_{t,x}^{\{e_0\}}\varphi_{z,t}^{\{e_2\}}\varphi_{y,z}^{\{e_1\}}\varphi_{x,y}^{\{e_3\}})$ in two ways. First, let $a = \{e_2, e_3\}$. By M3 we can factor $\varphi_{y,z}^{\{e_1\}} = \varphi_{a,z}^{\{e_2\}}\varphi_{y,a}^{\{e_3\}}$ and $\varphi_{t,x}^{\{e_0\}} = \varphi_{a,x}^{\{e_3\}}\varphi_{t,a}^{\{e_2\}}$. We get:

$$\begin{aligned} m(\varphi_{t,x}^{\{e_0\}}\varphi_{z,t}^{\{e_2\}}\varphi_{y,z}^{\{e_1\}}\varphi_{x,y}^{\{e_3\}}) &= m((\varphi_{a,x}^{\{e_3\}}\varphi_{t,a}^{\{e_2\}})\varphi_{z,t}^{\{e_2\}}(\varphi_{a,z}^{\{e_2\}}\varphi_{y,a}^{\{e_3\}})\varphi_{x,y}^{\{e_3\}}) \\ &= m(\varphi_{t,a}^{\{e_2\}}\varphi_{z,t}^{\{e_2\}}\varphi_{a,z}^{\{e_2\}}\varphi_{y,a}^{\{e_3\}}\varphi_{x,y}^{\{e_3\}}\varphi_{a,x}^{\{e_3\}}) \\ &= m(\varphi_{t,a}^{\{e_2\}}\varphi_{z,t}^{\{e_2\}}\varphi_{a,z}^{\{e_2\}}) + m(\varphi_{y,a}^{\{e_3\}}\varphi_{x,y}^{\{e_3\}}\varphi_{a,x}^{\{e_3\}}) \end{aligned}$$

Similarly, let $b = \{e_0, e_1\}$. Factoring $\varphi_{x,y}^{\{e_3\}} = \varphi_{b,y}^{\{e_1\}}\varphi_{x,b}^{\{e_0\}}$ and $\varphi_{z,t}^{\{e_2\}} = \varphi_{b,t}^{\{e_0\}}\varphi_{z,b}^{\{e_1\}}$, we get:

$$\begin{aligned} m(\varphi_{t,x}^{\{e_0\}}\varphi_{z,t}^{\{e_2\}}\varphi_{y,z}^{\{e_1\}}\varphi_{x,y}^{\{e_3\}}) &= m(\varphi_{t,x}^{\{e_0\}}(\varphi_{b,t}^{\{e_0\}}\varphi_{z,b}^{\{e_1\}})\varphi_{y,z}^{\{e_1\}}(\varphi_{b,y}^{\{e_1\}}\varphi_{x,b}^{\{e_0\}})) \\ &= m(\varphi_{x,b}^{\{e_0\}}\varphi_{t,x}^{\{e_0\}}\varphi_{b,t}^{\{e_0\}}\varphi_{z,b}^{\{e_1\}}\varphi_{y,z}^{\{e_1\}}\varphi_{b,y}^{\{e_1\}}) \\ &= m(\varphi_{x,b}^{\{e_0\}}\varphi_{t,x}^{\{e_0\}}\varphi_{b,t}^{\{e_0\}}) + m(\varphi_{z,b}^{\{e_1\}}\varphi_{y,z}^{\{e_1\}}\varphi_{b,y}^{\{e_1\}}) \end{aligned}$$

Now we note that these monodromies only depend on the pre-cyclic structure $\mathcal{O}_G$, which can be encoded using an arrow diagram. For example, $m(\varphi_{t,a}^{\{e_2\}}\varphi_{z,t}^{\{e_2\}}\varphi_{a,z}^{\{e_2\}}) = -1$ if and only

if the arrow on edge $\{e_2\}$ points from z to t.

Given an arrow diagram, we can define c_i to be -1 if the arrow on edge $\{e_i\}$ points counterclockwise around the origin, and 1 otherwise. Then the preceeding algebra gives:

$$m(\varphi_{t,x}^{\{e_0\}}\varphi_{z,t}^{\{e_2\}}\varphi_{y,z}^{\{e_1\}}\varphi_{x,y}^{\{e_3\}}) = c_2 + c_3 = c_0 + c_1$$

Since each $c_i = \pm 1$, the only options are if all $c_i = 1$, all $c_i = -1$, or c_0, c_1 have opposite signs and c_2, c_3 have opposite signs.

It is now straightforward to see that the arrow diagram is either (see definition 3.4.2):

- Coherent, if $m(\varphi_{t,x}^{\{e_0\}}\varphi_{z,t}^{\{e_2\}}\varphi_{y,z}^{\{e_1\}}\varphi_{x,y}^{\{e_3\}}) = 0$
- Spiral, if $m(\varphi_{t,x}^{\{e_0\}}\varphi_{z,t}^{\{e_2\}}\varphi_{y,z}^{\{e_1\}}\varphi_{x,y}^{\{e_3\}}) = \pm 2$

(See Figure 6.4). Meaning $\mathcal{O}_G$ is a cyclic structure.

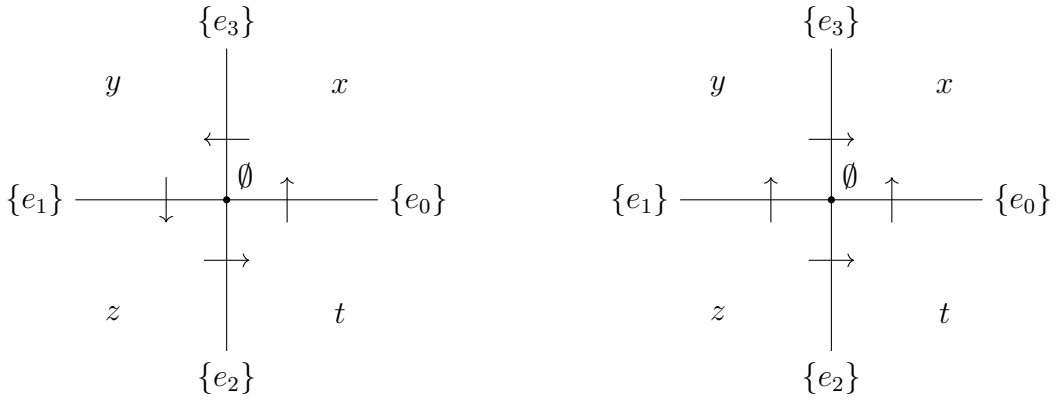


Figure 6.4: Arrow Diagrams Giving $m(\varphi_{t,x}^{\{e_0\}}\varphi_{z,t}^{\{e_2\}}\varphi_{y,z}^{\{e_1\}}\varphi_{x,y}^{\{e_3\}}) = -2$ (left) or 0 (right)

Now let G be a graded structure on $\mathcal{P}_{S_4}$. We know $\mathcal{O}_G$ is A_3-coherent, so it suffices to rule out the possibility of a 'pinwheel orientation' discussed in Chapter 3. WLOG we can consider the arrow diagram pictured in Figure 6.5, in which arrows on the $x_2 = 0$ and $x_3 = 0$ hyperplanes point counterclockwise about the x_1-axis (as viewed from the positive side), and arrows on the $x_1 = 0$ hyperplane point toward $x_1 > 0$.

Focussing on a neighborhood of $\{v_1\}$, we get Figure 6.6.

Let u denote the top-cell $\{v_2, e_2, v_3, e_3\}$. By the reasoning in the previous section, we have:

$$m(\varphi_{t,x}^{\{v_1,v_2\}}\varphi_{z,t}^{\{v_1,e_3\}}\varphi_{y,z}^{\{v_1,e_2\}}\varphi_{x,y}^{\{v_1,v_3\}}) = -2$$

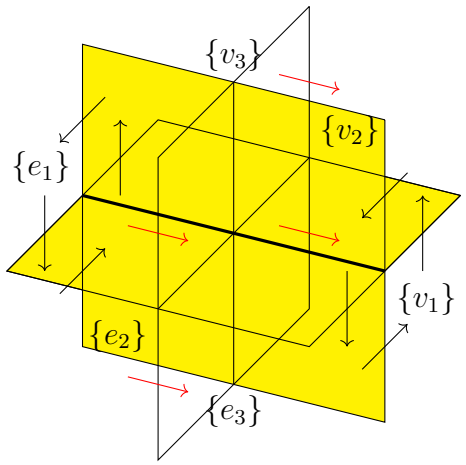


Figure 6.5: A Pinwheel Orientation

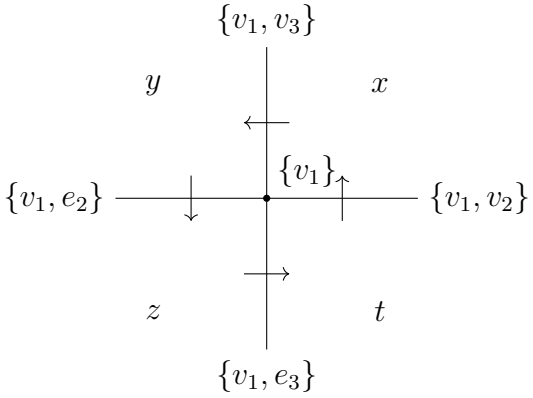


Figure 6.6: A Neighborhood of $\{v_1\}$ in the Pinwheel Orientation

On the other hand, we have the closed subset in a neighborhood of $\{v_3\}$ pictured in Figure 6.7 which allows us to apply (M3).

$(M3)$ gives us: $\varphi_{x,y}^{\{v_1,v_3\}} = \varphi_{u,y}^{\{v_3,e_2\}}\varphi_{x,u}^{\{v_2,v_3\}}$. A similar identity holds for $\varphi_{y,z}^{\{v_1,e_2\}}$, etc. We get:

$$\begin{aligned}
&m(\varphi_{t,x}^{\{v_1,v_2\}}\varphi_{z,t}^{\{v_1,e_3\}}\varphi_{y,z}^{\{v_1,e_2\}}\varphi_{x,y}^{\{v_1,v_3\}}) \\
&\quad = m((\varphi_{u,x}^{\{v_2,v_3\}}\varphi_{t,u}^{\{e_3,v_2\}}) \circ (\varphi_{u,t}^{\{e_3,v_2\}}\varphi_{z,u}^{\{e_2,e_3\}}) \circ (\varphi_{u,z}^{\{e_2,e_3\}}\varphi_{y,u}^{\{v_3,e_2\}}) \circ (\varphi_{u,y}^{\{v_3,e_2\}}\varphi_{x,u}^{\{v_2,v_3\}})) \\
&\quad = m(\varphi_{u,x}^{\{v_2,v_3\}}(\varphi_{t,u}^{\{e_3,v_2\}}\varphi_{u,t}^{\{e_3,v_2\}})(\varphi_{z,u}^{\{e_2,e_3\}}\varphi_{u,z}^{\{e_2,e_3\}})(\varphi_{y,u}^{\{v_3,e_2\}}\varphi_{u,y}^{\{v_3,e_2\}})\varphi_{x,u}^{\{v_2,v_3\}}) \\
&\quad = m(\varphi_{u,x}^{\{v_2,v_3\}}\varphi_{x,u}^{\{v_2,v_3\}}) = 0
\end{aligned}$$

With the last deduction using $(M1)$. This contradicts our previous calculation. □

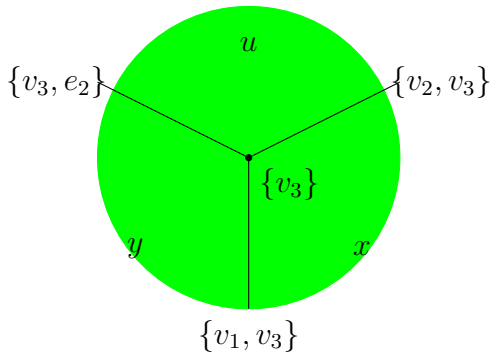


Figure 6.7: Cells Near $\{v_3\}$ In $\mathcal{P}_{S_4}$

Note on Notation

Due to the nature of graded structures, we will frequently be referencing top-cells in $\mathcal{P}_T$. Recall from Chapter 2 (corollary 2.3.1) that these are in bijection with subtrees of T. For brevity, in the remainder of this section, we will use the notation $[S]$ to designate the top-cell in $\mathcal{P}_T$ determined by S. Additionally, if S is an explicit list of vertices, say $S = \{v_1, v_2\}$, we will write $[v_1, v_2]$ to denote $[S]$ instead of the lengthier $[\{v_1, v_2\}]$.

Additionally, in this section we will be doing many computations on A_2 and A_3. We will be using notation from Chapter 2 to express correspondences $\mathfrak{p} : R \to A_2$ and $\mathfrak{p} : R \to A_3$ (The Hasse diagrams are shown in Figures 6.8, 6.9 for convenience).

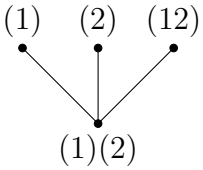


Figure 6.8: Hasse diagram for $\mathcal{P}_{A_2}$

As a final remark: We have a general notation for one-cells: For $k \in K_T$ we have $\langle k \rangle \in C_1(\mathcal{P}_T)$- and this introduces a somwhat confusing conflict in the case when $T = A_2$. Namely, when $v \in V_t(T)$, the notation $\langle v \rangle$ refers to *deleting* the vertex v, hence comparing those notations when $T = A_2$ we have:

$$[(1)] = \langle 2 \rangle$$
$$[(2)] = \langle 1 \rangle$$
$$[(12)] = \langle e \rangle$$

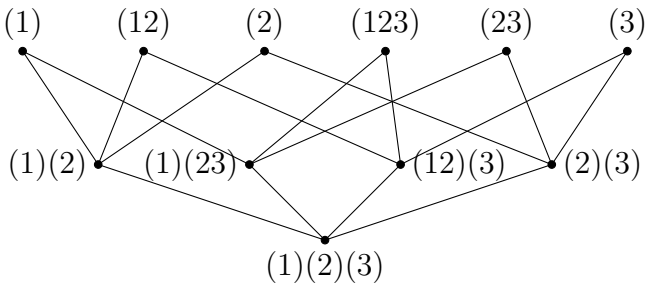


Figure 6.9: Hasse diagram for $\mathcal{P}_{A_3}$

Construction of G_μ

We have seen that any graded structure determines a cyclic structure. In this section we will see that, locally, graded structures are effectively equivalent to cyclic structures.

First, we will construct a graded structure which induces $\mathcal{O}_\mu$ when (T, μ) is a directed tree:

Definition 6.1.3. *Let (T, μ) be a directed tree. We define the graded structure G_μ on $\mathcal{P}_T$ as follows: For all top-cells x, we take $G_{\mu,x} = \mathbb{Z}$. Then an isomorphism $G_{\mu,x} \xrightarrow{\sim} G_{\mu,y}$ can be identified with an integer. To emphasize that this notation is* additive*, if $\varphi : \mathbb{Z} \to \mathbb{Z}$ is defined by $\varphi(x) = x + n$, we will write $\varphi = +n$. (Hence the notation $\varphi = +0$ denotes the* identity*, not the zero map).*

Let z be a codimension-one cell represented by $\mathfrak{p} = (A_2 \xleftarrow{q} S \hookrightarrow T)$, where the edge in $\mathfrak{p}^*\mu$ points from 1 to 2, and write $z_1 = [\mathfrak{p} \circ \langle 1 \rangle] = [\mathfrak{p} \circ (2)]$, $z_2 = [\mathfrak{p} \circ \langle 2 \rangle] = [\mathfrak{p} \circ (1)]$, $z_e = [\mathfrak{p} \circ \langle e \rangle] = [\mathfrak{p} \circ (12)]$. Then the map $\varphi^z_{z_1,z_2} = +1 = -\varphi^z_{z_2,z_1}$, and all other maps are $+0$.

Another way of interpreting this definition: Suppose S, S' are subtrees with $[S], [S']$ sharing a codimension-one cell z in their closure. We have $\varphi^z_{[S],[S']} = -1$ if S and S' are disjoint and the edge that connects them points from $[S]$ to $[S']$, $+1$ if the edge points in the opposite direction, and $+0$ in all other cases. This is illustrated in Figure 6.10.

Proposition 6.1.2. *G_μ as defined above is a graded structure, and $\mathcal{O}_{G_\mu} = \mathcal{O}_\mu$.*

Proof. The cyclic order on z_1, z_e, z_3 in $\mathcal{O}_\mu^{[\mathfrak{p}]}$ is the total order induced by $z_1 < z_e < z_2$, and we see this agrees with the cyclic order in $\mathcal{O}_{G_\mu}^{[\mathfrak{p}]}$. So we need to show that G_μ is actually a graded structure. Monodromy axioms $(M1)$ and $(M2)$ are clear from the definition, so we need to show $(M3)$. The maps involved $(M3)$ are all in the neighborhood of a codimension-two cell

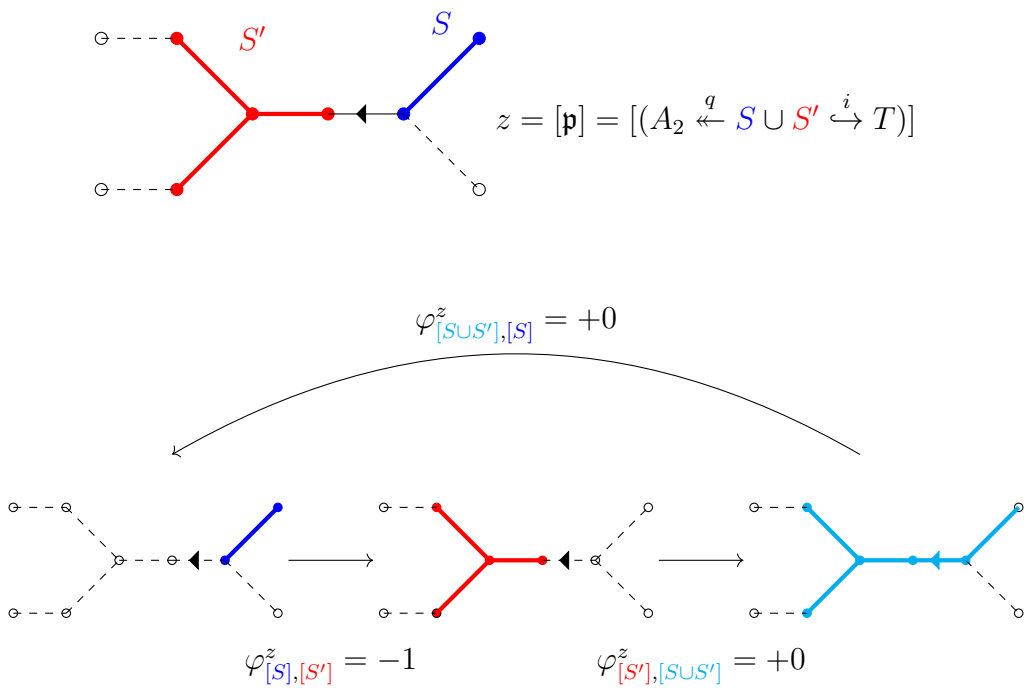


Figure 6.10: Illustration of the Graded Structure G_μ

w, meaning that it is enough to verify this axiom when $T = A_3$.

To verify the axiom, we need to select three one-cells z_1, z_2, z_3. This uniquely determines the three top-cells x_1, x_2, x_3 referenced in the axiom. Note that applying some permutation to $\{z_1, z_2, z_3\}$ can only impact the monodromy $m(\varphi^{z_1}_{x_3,x_1}\varphi^{z_3}_{x_2,x_3}\varphi^{z_2}_{x_1,x_2})$ by changing its sign, so we only need to show the monodromy is zero once for each of the four subsets of three one-cells of $\mathcal{P}_{A_3}$.

Suppose the three one-cells are $z_1 = (1)(2)$, $z_2 = (1)(23)$, and $z_3 = (2)(3)$. Then $x_1 = (1)$, $x_2 = (23)$, and $x_3 = (2)$.

Then $\varphi^{z_2}_{x_1,x_2} = \pm 1$, depending on the orientation of the edge $\{1, 2\}$- it is -1 when the edge points from 1 to 2 and $+1$ otherwise. $\varphi^{z_3}_{x_2,x_3} = +0$. $\varphi^{z_1}_{x_3,x_1} = \pm 1$, also depending on the orientation of the edge $\{1, 2\}$, but this time it is $+1$ when the edge points from 1 to 2 and -1 otherwise. So the total monodromy is zero as desired.

The case when $z_1 = (2)(3)$, $z_2 = (12)(3)$, $z_3 = (1)(2)$, is the same as this by applying an automorphism of A_3 switching vertices 1 and 3.

Suppose we have $z_1 = (1)(2)$, $z_2 = (1)(23)$, $z_3 = (12)(3)$. Then $x_1 = (1)$, $x_2 = (123)$,

$x_3 = (12)$.

In this case, $\varphi^{z_2}_{x_1,x_2} = \varphi^{z_3}_{x_2,x_3} = \varphi^{z_1}_{x_3,x_1} = 0$, so the total monodromy is zero. Applying the automorphism of A_3 switching vertices 1 and 3 gives the last case.

□

Isomorphisms of Graded Structures

In this section, we show graded structures are locally determined up to isomorphism by cyclic structures:

Definition 6.1.4. *Let $\mathcal{P}$ be a combinatorial arboreal space, and G, G' be graded structures with structure maps φ, φ'. An **isomorphism** $\tau : G \to G'$ consists of an isomorphism $\tau_x : G_x \xrightarrow{\sim} G'_x$ for each top-cell x, such that $\tau_y \varphi^z_{x,y} = \varphi'^z_{x,y} \tau_x$ for all top-cells x, y and codimension-one cells $z \leq x, y$.*

We say τ **intertwines** $\varphi^z_{x,y}$ and $\varphi'^z_{x,y}$ when the above relation holds.

Proposition 6.1.3. *If G and G' are graded structures on $\mathcal{P}_T$ with $\mathcal{O}_G = \mathcal{O}_{G'}$, then G and G' are isomorphic. Furthermore, the group of automorphisms of any graded structure G is naturally isomorphic to $\mathbb{Z}$.*

Proof. We prove both statements simultaneously by induction on the number of vertices of T. When T has a single vertex, a graded structure is the same as a single $G-$torsor, so both statements are true.

For the inductive step, we can assume after applying an automorphism to $\mathcal{P}_T$ (and using the fact that 'isomorphic' is a transitive relation) that $G = G_\mu$ for some orientation μ of T. Let G' be a graded structure with $\mathcal{O}_{G'} = \mathcal{O}_G = \mathcal{O}_\mu$. If v is a terminal vertex of T, let $\mathfrak{p} = (T\backslash\{v\} \xleftarrow{\sim} T\backslash\{v\} \hookrightarrow T)$. By induction, there exists an isomorphism $\tau : i^*_{\mathfrak{p}} G' \to i^*_{\mathfrak{p}} G$. Then we are done if we can prove:

Claim: There is a unique extension of τ to an isomorphism $G' \to G$.

To prove the claim- we first use the commutativity property to show there is only one choice for $\tau_{[S]}$ when $[S] \in \mathcal{P}_T\backslash N([\mathfrak{p}])$, i.e. when $v \in S$. We then use (M2) and (M3) to show this extension is an isomorphism.

Let S be a subtree containing v which is not equal to $\{v\}$, write $S_0 = S\backslash\{v\}$. Let z be the class of the correspondence $\{A_2 \overset{q}{\twoheadleftarrow} S \overset{i}{\hookrightarrow} T\}$, with $q(v) = 1$ and $q(S\backslash\{v\}) = 2$. The following diagram needs to commute:

$$\begin{array}{ccc} G'_{[S]} & \xrightarrow{\varphi'^{z}_{[S],[S_0]}} & G'_{[S_0]} \\ {\scriptstyle \tau_{[S]}}\downarrow & & \downarrow{\scriptstyle \tau_{[S_0]}} \\ \mathbb{Z} & \xrightarrow{\varphi^{z}_{[S],[S_0]}} & \mathbb{Z} \end{array}$$

Since all other maps in this diagram are defined, this defines $\tau_{[S]}$. If $S = \{v\}$, let $\hat{v}$ denote the unique vertex connected to v. Write z for the class of $(A_2 \overset{q}{\twoheadleftarrow} \{v, \hat{v}\} \overset{i}{\hookrightarrow} T)$. We've defined $\tau_{[v,\hat{v}]}$ in the previous step, hence the diagram:

$$\begin{array}{ccc} G'_{[v]} & \xrightarrow{\varphi'^{z}_{[v],[v,\hat{v}]}} & G'_{[v,\hat{v}]} \\ {\scriptstyle \tau_{[v]}}\downarrow & & \downarrow{\scriptstyle \tau_{[v,\hat{v}]}} \\ \mathbb{Z} & \xrightarrow{\varphi^{z}_{[v],[v,\hat{v}]}} & \mathbb{Z} \end{array}$$

Defines $\tau_{[v]}$ uniquely. We are done if we can show τ intertwines all structure maps $\varphi^z_{[S],[S']}$ and $\varphi'^z_{[S],[S']}$. When $v \notin S, S'$, this is guaranteed by our assumption on τ. Aside from this, there are a handful of cases, each of which is solved by a suitable application of (M2) or (M3).

Case I: Assume $S = \{v\}$, and $v \notin S'$. This case is a little bit tricky: We know $\hat{v} \in S'$, where $\hat{v}$ is the unique vertex connected to v, but $S'\backslash\{\hat{v}\}$ may have multiple connected components. We induct on the number of connected components of $S'\backslash\{\hat{v}\}$, with the base case being *zero* connected components, i.e. $S' = \{\hat{v}\}$.

For the base case, let $\mathfrak{p} = (A_2 \overset{q}{\twoheadleftarrow} \{v, \hat{v}\} \overset{i}{\hookrightarrow} T)$, so $z = [\mathfrak{p}]$. Then we have a diagram:

$$\begin{array}{ccccc} G'_{[v]} & \xrightarrow{\varphi'^{z}_{[v],[v,\hat{v}]}} & G'_{[v,\hat{v}]} & \xrightarrow{\varphi'^{z}_{[v,\hat{v}],[\hat{v}]}} & G'_{[\hat{v}]} \\ {\scriptstyle \tau_{[v]}}\downarrow & & {\scriptstyle \tau_{[v,\hat{v}]}}\downarrow & & \downarrow{\scriptstyle \tau_{[\hat{v}]}} \\ \mathbb{Z} & \xrightarrow{\varphi^{z}_{[v],[v,\hat{v}]}} & \mathbb{Z} & \xrightarrow{\varphi^{z}_{[v,\hat{v}],[\hat{v}]}} & \mathbb{Z} \end{array}$$

The left square commutes by the definition of $\tau_{[v]}$, while the right square commutes by the definition of $\tau_{[v,\hat{v}]}$. Since $\mathcal{O}_{G'} = \mathcal{O}_G$, we have an equality of monodromies:

$$m(\varphi'^{z}_{[\hat{v}],[v]}\varphi'^{z}_{[v,\hat{v}],[\hat{v}]}\varphi'^{z}_{[v],[v,\hat{v}]}) = m(\varphi^{z}_{[\hat{v}],[v]}\varphi^{z}_{[v,\hat{v}],[\hat{v}]}\varphi^{z}_{[v],[v,\hat{v}]})$$

Hence τ intertwines $\varphi'^{z}_{[v],[\hat{v}]}$ and $\varphi^{z}_{[v],[\hat{v}]}$.

For the inductive step, we can take a subtree S_0 with $S_0 \subseteq S'$ and $\hat{v} \in S_0$, such that $S_0 \backslash \{\hat{v}\}$ has one fewer connected component than $S' \backslash \{\hat{v}\}$, and $S' \backslash S_0$ has a single connected component. Let $\mathfrak{p} = (A_3 \overset{q}{\twoheadleftarrow} S' \overset{i}{\hookrightarrow} T)$, where $q(v) = 1$, $q(S_0) = 2$, and $q(S' \backslash S_0) = 3$.

Then $z = [\mathfrak{p} \circ (1)(23)]$. Let $z_1 = [\mathfrak{p} \circ (1)(2)]$ and $z_2 = [\mathfrak{p} \circ (2)(3)]$. Then consider the diagram:

$$\begin{array}{ccccc} G'_{[v]} & \xrightarrow{\varphi'^{z_1}_{[v],[S_0]}} & G'_{[S_0]} & \xrightarrow{\varphi'^{z_2}_{[S_0],[S']}} & G'_{[S']} \\ \downarrow{\scriptstyle \tau_{[v]}} & & \downarrow{\scriptstyle \tau_{[S_0]}} & & \downarrow{\scriptstyle \tau_{[S']}} \\ \mathbb{Z} & \xrightarrow{\varphi^{z_1}_{[v],[S_0]}} & \mathbb{Z} & \xrightarrow{\varphi^{z_2}_{[S_0],[S']}} & \mathbb{Z} \end{array}$$

The left square commutes by the inductive assumption. The right square commutes because $v \notin S_0, S'$. By (M3) we have $\varphi^{z}_{[v],[S']} = \varphi^{z_2}_{[S_0],[S']}\varphi^{z_1}_{[v],[S_0]}$, and likewise for φ', so τ intertwines $\varphi^{z}_{[v],[S']}$ and $\varphi'^{z}_{[v],[S']}$.

Case II: Assume $v \in S$, but $S \neq \{v\}$, and $S' = \{v\}$. Letting $S_0 = S \backslash \{v\}$ and $\mathfrak{p} = (A_2 \overset{q}{\twoheadleftarrow} S \overset{i}{\hookrightarrow} T)$ with $q(v) = 1$ and $q(S_0) = 2$, we have $z = [\mathfrak{p}]$. Then in the following diagram:

$$\begin{array}{ccccc} G'_{[S]} & \xrightarrow{\varphi'^{z}_{[S],[S_0]}} & G'_{[S_0]} & \xrightarrow{\varphi'^{z}_{[S_0],[v]}} & G'_{[v]} \\ \downarrow{\scriptstyle \tau_{[S]}} & & \downarrow{\scriptstyle \tau_{[S_0]}} & & \downarrow{\scriptstyle \tau_{[v]}} \\ \mathbb{Z} & \xrightarrow{\varphi^{z}_{[S],[S_0]}} & \mathbb{Z} & \xrightarrow{\varphi^{z}_{[S_0],[v]}} & \mathbb{Z} \end{array}$$

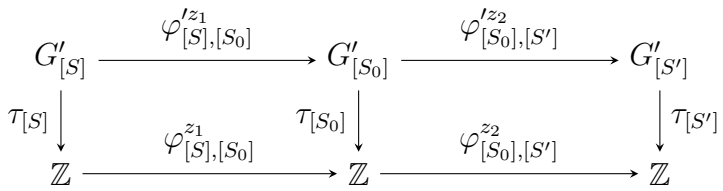
The left square commutes by definition of $\tau_{[S]}$, and the right square commutes by Case I. Then, since $\mathcal{O}_G = \mathcal{O}_{G'}$, we have equality of monodromies:

$$m(\varphi'^{z}_{[v],[S]}\varphi'^{z}_{[S_0],[v]}\varphi'^{z}_{[S],[S_0]}) = m(\varphi^{z}_{[v],[S]}\varphi^{z}_{[S_0],[v]}\varphi^{z}_{[S],[S_0]})$$

τ also intertwines $\varphi^{z}_{[S],[v]}$ and $\varphi'^{z}_{[S],[v]}$.

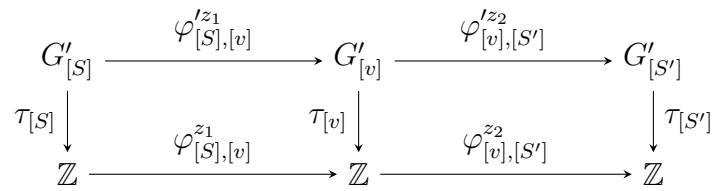
Case III: Assume $v \in S$, $S \neq \{v\}$, and S, S' are disjoint. If there is a codimension-one cell z less than S, S', then there is one edge separating S and S'. Let $S_0 = S \backslash \{v\}$. If we let $\mathfrak{p} = (A_3 \overset{q}{\twoheadleftarrow} S \cup S' \overset{i}{\hookrightarrow} T)$, where $q(v) = 1$, $q(S_0) = 2$, $q(S') = 3$, then $z = [\mathfrak{p} \circ (12)(3)]$.

We set $z_1 = [\mathfrak{p} \circ (1)(2)]$, and $z_2 = [\mathfrak{p} \circ (2)(3)]$. Consider the diagram:



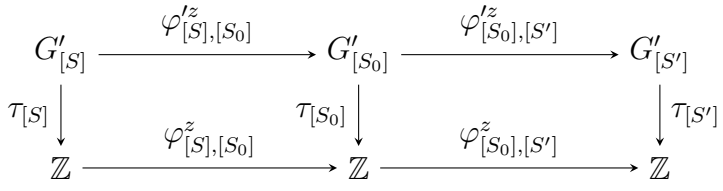
The left square commutes by the definition of $\tau_{[S]}$. The right square commutes, since neither S_0 nor S' contain v. So the whole rectangle commutes, and by (M3) we have $\varphi^z_{[S],[S']} = \varphi^{z_2}_{[S_0],[S']}\varphi^{z_1}_{[S],[S_0]}$, and likewise for φ'. So τ intertwines $\varphi^z_{[S],[S']}$ and $\varphi'^z_{[S],[S']}$.

Case IV: Assume $\{v\} \in S$, $S \neq \{v\}$, and $S \subseteq S'$. Let $\mathfrak{p} = (A_3 \overset{q}{\twoheadleftarrow} S' \overset{i}{\hookrightarrow} T)$, where $q(v) = 1$, $q(S\backslash\{v\}) = 2$, and $q(S'\backslash S) = 3$. Then $z = [\mathfrak{p} \circ (12)(3)]$. Let $z_1 = [\mathfrak{p} \circ (1)(2)]$, and $z_2 = [\mathfrak{p} \circ (1)(23)]$. Consider the diagram:



Both squares commute by Case II (the left side using the inverse of Case II), hence the full rectanges commute, which implies the desired result using (M3).

Case V: In the last case: Assume $\{v\} \in S$ but $S \neq \{v\}$, $v \notin S'$, and $S' \subseteq S$. Write $S_0 = S\backslash S'$. Use the correspondence $\mathfrak{p} = (A_2 \overset{q}{\twoheadleftarrow} S \overset{i}{\hookrightarrow} T)$, with $q(S_0) = 1$, and $q(S') = 2$. Then $z = [\mathfrak{p}]$. Consider the diagram:



The left square commutes by (the inverse of) Case IV, and the right square commutes by either Case I or Case III, depending on whether $S_0 = \{v\}$ or not. Then, since $\mathcal{O}_G = \mathcal{O}_{G'}$, we have equality of monodromies:

$$m(\varphi'^z_{[S'],[S]}\varphi'^z_{[S_0],[S']}\varphi'^z_{[S],[S_0]}) = m(\varphi^z_{[S'],[S]}\varphi^z_{[S_0],[S']}\varphi^z_{[S],[S_0]})$$

τ also intertwines $\varphi^z_{[S],[S']}$ and $\varphi'^z_{[S],[S']}$.

The above options (and their inverses) cover all cases. □

6.2 Branes

Pre-Branes

Let $\mathcal{P}$ be a combinatorial arboreal space.

Definition 6.2.1. *A subset $L \subseteq \mathcal{P}$ is a **pre-brane** if it satisfies the following criteria:*

(i) L is equal to the closure of its interior. Equivalently, L is closed and, for $x \in L$, there exists a $y \in L$ with $x \leq y$ and y maximal in $\mathcal{P}$.

(ii) If $x \in L$ is a codimension one cell, so $N(x) \cong \mathcal{P}_{A_2}$, L contains exactly two out of the three maximal cells greater than x (See Figure 6.11).

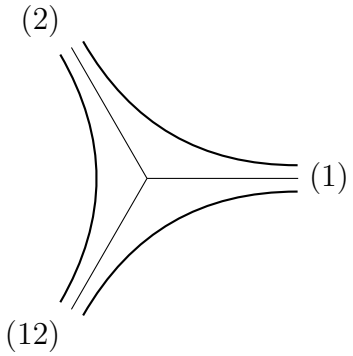


Figure 6.11: Pre-Branes on $\mathcal{P}_{A_2}$

Note the similarity of Figure 6.11 to Figure 5.2. This is no accident- We will see there is a close relationship between branes and the representation theory of (T, μ).

Let $\mathrm{Br}_{pre}(\mathcal{P})$ denote the set of pre-branes in $\mathcal{P}$. Note that a pre-brane is determined by the set of top-cells it contains, by property 1. If we let $C^0(\mathcal{P})$ denote the set of top-cells, then we can think of $\mathrm{Br}_{pre}(\mathcal{P})$ as a subset of $\mathrm{Fun}(C^0(\mathcal{P}), \mathbb{Z}/2\mathbb{Z})$. In fact, we can see that it is a *subspace* of this $\mathbb{Z}/2\mathbb{Z}-$vector space: If we let:

$$\partial : \mathrm{Fun}(C^0(\mathcal{P}), \mathbb{Z}/2\mathbb{Z}) \to \mathrm{Fun}(C^1(\mathcal{P}), \mathbb{Z}/2\mathbb{Z})$$

$$\partial(f)(y) = \sum_{y<x} f(x)$$

Then condition (ii) implies $\mathrm{Br}_{pre}(\mathcal{P}) = \mathrm{Ker}(\partial)$.

Whenever a sum of pre-branes is referenced it will be with respect to this structure. The identity element is the empty brane $\emptyset \subset \mathcal{P}$- we will typically refer to it as 0.

We remark that for an arbitrary C-Arb space $\mathcal{P}$, the assignment $U \mapsto \mathrm{Br}_{pre}(U)$ defines a sheaf of vector spaces on $\mathcal{P}$.

For an arboreal singularity $\mathcal{P}_T$, we have a nice way to find a basis for $\mathrm{Br}_{pre}(\mathcal{P}_T)$.

Definition 6.2.2. *Let v be a vertex in T. We define $L_v \subset \mathcal{P}_T$ by:*

$$L_v = \{[\mathfrak{p}] \mid \mathfrak{p} = (R \overset{q}{\twoheadleftarrow} S \overset{i}{\hookrightarrow} T), v \in V(S)\}$$

For a vertex $v \in V(T)$, we have a correspondence $\mathfrak{p}_v = (\star \overset{\sim}{\leftarrow} \{v\} \hookrightarrow T)$. We can define the linear functional $\chi_v : Br_{pre}(\mathcal{P}_T) \to \mathbb{Z}/2\mathbb{Z}$ by:

$$\chi_v(L) = \begin{cases} 1 & [\mathfrak{p}_v] \in L \\ 0 & [\mathfrak{p}_v] \notin L \end{cases}$$

Lemma 6.2.1. *L_v is a pre-brane. Furthermore, $\{L_v\}_{v \in V(T)}$ and $\{\chi_v\}_{v \in V(T)}$ are dual bases.*

Proof. First we show that L_v satisfies the local A_2 condition. Let $\mathfrak{p} = \{A_2 \overset{q}{\twoheadleftarrow} S \overset{i}{\hookrightarrow} T\}$ be a correspondence with $[\mathfrak{p}] \in L_v$ (so, $v \in V(S)$). The three top-cells containing $[\mathfrak{p}]$ are:

$$[\mathfrak{p} \circ \langle 2 \rangle] = (\star \twoheadleftarrow q^{-1}(1) \hookrightarrow T) \qquad [\mathfrak{p} \circ \langle 1 \rangle] = (\star \twoheadleftarrow q^{-1}(2) \hookrightarrow T) \qquad [\mathfrak{p} \circ \langle e \rangle] = (\star \twoheadleftarrow S \hookrightarrow T)$$

Then $[\mathfrak{p} \circ (12)] \in L_v$, and exactly one of $[\mathfrak{p} \circ \langle 1 \rangle], [\mathfrak{p} \circ \langle 2 \rangle]$ are in L_v depending on whether $q(v) = 1$ or $q(v) = 2$.

For the second assertion, we have $\chi_v(L_w) = \delta_{v,w}$, so it remains to show that if $\chi_v(L) = 0$ for all $v \in V(T)$, then $L = 0$.

To show this, assume $L \neq 0$. Pick a cell $[S] \in L$, with S having a minimal number of vertices. If S has at least two vertices, we can construct a correspondence $\mathfrak{p} = (A_2 \twoheadleftarrow S \hookrightarrow T)$. Since L is closed, $[\mathfrak{p}] \in L$, and by the A_2- condition we must have exactly one of $[q^{-1}(1)] \in L$ or $[q^{-1}(2)]$ in L. But $|q^{-1}|$ and $|q^{-1}(2)|$ are both less than $|S|$, which contradicts the minimality assumption. So $|S| = 1$, meaning $S = \{v\}$ for some $v \in V(T)$, and so $\chi_v(L) \neq 0$. □

We end this section with a brief discussion about pullbacks:

Let $\mathfrak{p} = (R \overset{q}{\twoheadleftarrow} S \overset{i}{\hookrightarrow} T)$ denote a correspondence, which induces the open embedding $i_{\mathfrak{p}} : \mathcal{P}_R \to \mathcal{P}_T$. Then there is a linear map $i_{\mathfrak{p}}^* : \mathrm{Br}_{pre}(\mathcal{P}_T) \to \mathrm{Br}_{pre}(\mathcal{P}_R)$ which sends L to $i_{\mathfrak{p}}^{-1}(L)$.

There is a nice description of this linear map in the basis discussed above:

Lemma 6.2.2.

$$i_{\mathfrak{p}}^*(L_v) = \begin{cases} L_{q(v)} & v \in V(S) \\ 0 & v \notin V(S) \end{cases}$$

Proof. For a subtree W of R, write $\mathfrak{q} = (\star \twoheadleftarrow W \hookrightarrow R)$. Then $i_{\mathfrak{p}}^*(L_v)$ contains the top-cell $[\mathfrak{q}]$ if and only if L_v contains the top-cell $[\mathfrak{p} \circ \mathfrak{q}] = [(\star \twoheadleftarrow q^{-1}(W) \hookrightarrow T)]$. This is true if and only if $q^{-1}(W)$ contains v, i.e. if and only if $v \in V(S)$ and $q(v) \in W$. □

The similarity of Lemma 6.2.2 and the restriction functors $c_{\mathfrak{p}}$ reinforce an analogy between branes and representations: $L_v \leftrightarrow P_v$. This is a special case of a more general correspondence.

Branes

Let $(\mathcal{P}_{A_3}, \mathcal{O})$ denote a cyclic A_3 singularity. By proposition 4.4, $c = \Delta(\mathcal{O})$ is a cyclic order on $K_{A_3} = \{e_0, e_1, e_2, e_3\}$. Fix an isomorphism $K_{A_3} \cong [4] := \{1, 2, 3, 4\}$ such that c is induced by the total order $1 < 2 < 3 < 4$, and identify $\mathcal{P}_{A_3}$ with its image $\{S \subseteq [4] \mid |S| \leq 2\} \subset \mathscr{P}([4])$.

Now, write $L_0 = \overline{\{\{1,2\},\{2,3\},\{3,4\},\{4,1\}\}} = \mathcal{P}_{A_3} \backslash \{\{1,3\},\{2,4\}\} \in \mathrm{Br}_{pre}(\mathcal{P}_{A_3})$. Then we define the set of **branes** on $\mathcal{P}_{A_3}$. to be $\mathrm{Br}(\mathcal{P}_{A_3}, \mathcal{O}) := \mathrm{Br}_{pre}(\mathcal{P}_{A_3}) \backslash \{L_0\}$.

Note that L_0 does not depend on our choice of isomorphism $K_{A_3} \cong [4]$, since it is fixed under cyclic-order-preserving automorphisms of [4].

Definition 6.2.3. *Let $(\mathcal{P}, \mathcal{O})$ be a cyclic C-Arb space. A **brane** is a pre-brane L satisfying an A_3-coherence condition: If $i : \mathcal{P}_{A_3} \hookrightarrow \mathcal{P}$ is an open embedding, then $i^*(L) \in Br(\mathcal{P}_{A_3}, i^*\mathcal{O})$. Denote the set of branes by $Br(\mathcal{P}, \mathcal{O})$.*

Then a brane is a pre-brane satsifying a coherence condition in codimension two. There are geometric reasons for imposing this coherence condition- see Chapter 8 for more on this. As we will see in this section, this definition of brane can be justified by representation theory of $\mathrm{Rep}(T, \mu)$.

One can understand the above restriction using the arrow diagram from Chapter 3. Removing $\{e_0, e_1\}$ and $\{e_2, e_3\}$ from $\mathcal{P}_{A_3}$, the remaining cells form a pre-brane. The cyclic

structure on $\mathcal{P}_{A_3}$ can be illustrated by drawing arrows crossing the one-cells, forming a diagram which is either *coherent* or *spiral* (definition 3.4.2). Then the pre-brane $L = \mathcal{P}_{A_3}\backslash\{\{e_0, e_1\}, \{e_2, e_3\}\}$ is a brane if and only if the diagram is coherent. (Figure 6.12).

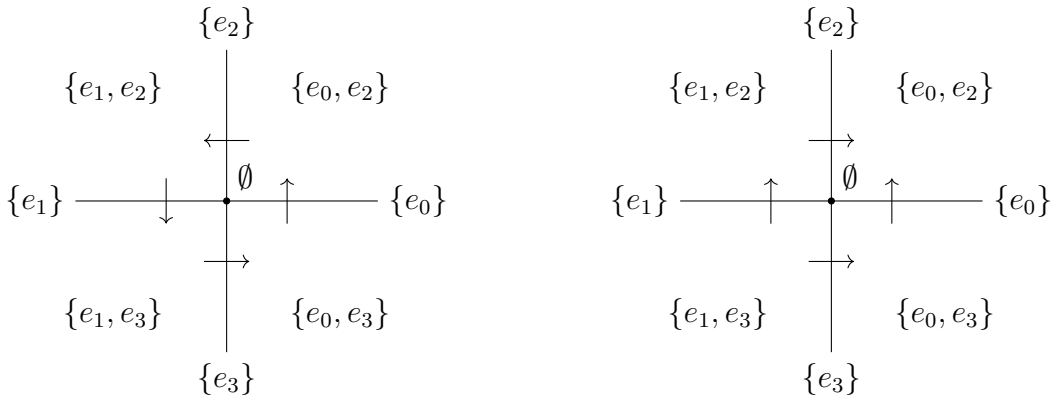


Figure 6.12: An Illustration of $L = \mathcal{P}_{A_3}\backslash\{\{e_0, e_1\}, \{e_2, e_3\}\}$ with two cyclic structures- On the left, L is *not* a Brane.

We now turn to a concrete description of the set of branes on the cyclic C-Arb singularity $(\mathcal{P}_T, \mathcal{O}_\mu)$ in terms of the basis $\{L_v\}_{v \in V(\vec{T})}$ for $\mathrm{Br}_{pre}(\mathcal{P}_T)$.

First, we analyze the case $T = A_3$. We introduce the term 'non-brane' for a pre-brane which is not a brane. For each orientation of A_3, we can use Table 3.1 in Chapter 3 to identify the unique non-brane. The results are listed in Table 6.1.

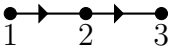
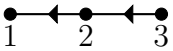
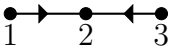
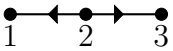
Tree	Non-Brane
	$L_1 + L_2 + L_3$
	$L_1 + L_2 + L_3$
	$L_1 + L_3$
	$L_1 + L_3$

Table 6.1: Non-Branes In $(\mathcal{P}_{A_3}, \mathcal{O}_\mu)$

Now, let T be a directed tree, and A a subtree of T isomorphic to A_n (we call A an A_n-subtree). We define $V_+(A) \subset V(A)$ to be the set of sources in A, and $V_-(A)$ to be the set of sinks- we only consider whether they are a source/sink *in* A (not in T- so, for example, the endpoints of A always qualify), and by convension if A consists of a single vertex we will call it a source.

Proposition 6.2.1. *Given an A_n-subtree A of T, define:*

$$L_A = \sum_{v \in V_+(A) \cup V_-(A)} L_v$$

Then the set of nonzero branes $Br(\mathcal{P}_T, \mathcal{O}_\mu)\backslash\{0\}$ is in bijection with the set of A_n subtrees via $A \leftrightarrow L_A$.

Proof. Given A, we know $L_A \in \mathrm{Br}_{pre}(\mathcal{P}_T)$, so we first show that it is a brane. Let $\mathfrak{p} = (A_3 \xleftarrow{q} S \xhookrightarrow{i} T)$. Write $\lambda = \mathfrak{p}^*\mu$.

Referring to Table 6.1, we want to show $i_\mathfrak{p}^* L_A$ is not $L_1 + L_2 + L_3$ or is not $L_1 + L_3$, depending on λ. Assume we have $i_\mathfrak{p}^* L_A = L_1 + aL_2 + L_3$, where $a \in \mathbb{Z}/2\mathbb{Z}$. Then A must intersect $q^{-1}(1)$ and $q^{-1}(3)$, which means a is the number of sources or sinks of A contained in $q^{-1}(2)$, counted mod 2. If both arrows in λ are pointing the same direction (i.e. both to the left or both to the right), there are an even number of sources or sinks in A contained in $q^{-1}(2)$, meaning $a = 0$. Similarly, if both arrows in λ are pointing in opposite directions, $a = 1$. So by Table 6.1, $i^* L_A \in \mathrm{Br}(\mathcal{P}_{A_3}, \mathcal{O}_\lambda)$.

Now, let $L \in \mathrm{Br}(\mathcal{P}_T, \mathcal{O}_\mu)$, and write:

$$L = \sum_{v \in W} L_v$$

Where $W \subset V(T)$. We will use Table 6.1 along with certain correspondences $\mathfrak{p} : A_3 \to T$ to show that $W = V_+(A) \cup V_-(A)$ for some A_n-subtree A.

Step 1: If $v \in V(T)$ is a vertex of degree > 2, we show W cannot contain a vertex in more than two connected components of $T\backslash\{v\}$.

Suppose for the sake of contradiction that it did. Then we could find $v_1, v_2, v_3 \in W$ such that no vertices between v_i and v are in W for each i. Let e_1, e_2, e_3 be the edges adjacent to v in the connected components of v_1, v_2, v_3 respectively. Then at least two of the three, say

e_1 and e_2, must either both be pointing toward v or both away from v.

Let $\mathfrak{p} = (A_3 \overset{q}{\twoheadleftarrow} S \overset{i}{\hookrightarrow} T)$, where $q^{-1}(1)$ is the path connecting v_1 to e_1 and $q^{-1}(3)$ is the path connecting v_2 to e_2. If $v \notin W$, let $q^{-1}(2) = \{v\}$, and if $v \in W$, let $q^{-1}(2)$ be the path connecting v to v_3. Then $i_{\mathfrak{p}}^* L = L_1 + L_3$, and by assumption on the orientations of e_1 and e_2 and table 6.1, this is not a brane.

Step 2: By step 1, we know all the vertices in W lie in some A_n-subtree. Let A be the smallest A_n-subtree containing all the vertices in W (so, in particular, the endpoints of A are in W). Then for a vertex $v \in V(A)$ which is not an endpoint, we show $v \in W$ iff v is a source or sink in A.

To see why, since A is an A_n-tree and both endpoints of A are in W, there exist two vertices $v_1, v_2 \in W$ such that no vertices between v_i and v are in W for each i. Let e_1, e_2 be the edges adjacent to v in the connected components of v_1, v_2 respectively. Let $\mathfrak{p} = (A_3 \overset{q}{\twoheadleftarrow} S \overset{i}{\hookrightarrow} T)$, where $q^{-1}(1)$ is the path connecting v_1 to e_1, $q^{-1}(2) = \{v\}$, and $q^{-1}(3)$ is the path connecting v_2 to e_2.

So we have $i_{\mathfrak{p}}^* L = L_1 + L_3 + aL_2$, where $a \in \mathbb{Z}/2\mathbb{Z}$ depends on whether $v \in W$ or not. By Table 6.1, if v is a source or sink, since L is a brane we must have $i_{\mathfrak{p}}^* L = P_1 + P_2 + P_3$, i.e. $v \in W$. Similarly, if v is neither a source or sink, $i_{\mathfrak{p}}^* L = P_1 + P_3$, so $v \notin W$.

□

Graded Branes

Let $(\mathcal{P}, G)$ be a graded C-Arb space.

Definition 6.2.4. *A **graded brane** is a pair (L, g), where $L \in Br_{pre}(\mathcal{P})$, and g is a choice $g_x \in G_x$ for all top-cells $x \in L$, such that whenever z is a codimension-one cell in L and $x, y \in L$ are top-cells containing z in their closure, $\varphi^z_{x,y}(g_x) = g_y$. Denote the set of graded branes by $Br_{gr}(\mathcal{P}, G)$.*

The use of the term graded brane (rather than graded pre-brane) is explained in the following lemma:

Lemma 6.2.3. *If $(L, g) \in Br_{gr}(\mathcal{P}, G)$, then $L \in Br(\mathcal{P}, \mathcal{O}_G)$.*

Proof. It is enough to prove this when $\mathcal{P} = \mathcal{P}_{A_3}$. In our discussion following definition 6.2.3, we see that if $L \in \mathrm{Br}_{pre}(\mathcal{P}_{A_3})$ is a non-brane (with respect to the cyclic structure $\mathcal{O}_G$), we can represent the cyclic structure as an arrow diagram on L, with the arrows representing a

cyclic orientation. But by our argument in the proof of proposition 6.1.1, for such an arrow diagram, the monodromy going counterclockwise around the four top-cells of L equals ± 2, hence there is no possible grading g on L. $\square$

In fact, the above is an if-and-only-if locally:

Proposition 6.2.2. *If T is a tree, the map $Br_{gr}(\mathcal{P}_T, G) \to Br(\mathcal{P}_T, \mathcal{O}_G)$ of $(L, g) \mapsto L$ realizes the former as a $\mathbb{Z}-$bundle over the latter (where the $\mathbb{Z}-$action is the obvious shifting).*

Proof. **Step 1:** First, let g^1, g^2 be two gradings of a brane L. We want to show they differ by the same constant on any top-cell in L- that is, if S, S' are subtrees of T in L, $g^1_{[S]} - g^2_{[S]} = g^1_{[S']} - g^2_{[S']}$.

If there is a codimension-one cell z with $z \leq [S], [S']$, then this conclusion is clear. More generally, we can construct a graph H_L whose vertices are labelled by top-cells $[S] \in L$, and whose edges connect pairs of subtrees sharing a codimension-one cell in their closure. It is enough to show that H_L is connected. In fact, H_L is connected more generally for pre-branes, as we show.

We use strong induction on the number of vertices of T, with the base case $T = \star$ being obvious. For the inductive step, let $L \in \mathrm{Br}_{pre}(\mathcal{P}_T)$, and $[S], [S'] \in V(H_L)$. If S, S' are disjoint and there is a single edge between them in T, there is a codimension-one cell z with $z \leq [S], [S']$, so there is an edge between $[S]$ and $[S']$ in H_L. If S, S' are disjoint and there are multiple edges of T between them, there is a correspondence $\mathfrak{p} = (A_3 \overset{q}{\twoheadleftarrow} \tilde{S} \overset{i}{\hookrightarrow} T)$ such that $q^{-1}(1) = S$ and $q^{-1}(3) = S'$. Letting $\mathfrak{q} = \mathfrak{p} \circ (1)(2)$, the A_2 condition at $N([\mathfrak{q}])$ implies L contains either $[q^{-1}(2)]$ or $[q^{-1}(\{1, 2\})]$, each of which are connected to $[S]$ and $[S']$ in H_L.

Now suppose $S \cap S'$ is nonempty. WLOG $S \backslash S' \neq \emptyset$, let S_1 be a connected component of $S \backslash S'$. There is a correspondence $\mathfrak{p} = (A_2 \overset{q}{\twoheadleftarrow} S \overset{i}{\hookrightarrow} T)$ with $q(S_1) = 1$ and $q(S \backslash S_1) = 2$. The A_2 condition at $N([\mathfrak{p}])$ implies L contains either $[S_1]$ or $[S \backslash S_1]$. In the former case, $[S_1]$ is connected to $[S]$ and $[S']$ in H_L. In the latter case, $[S]$ is connected to $[S \backslash S_1]$, so it remains so show $[S]$ and $[S \backslash S_1]$ lie in the same connected component of H_L.

Letting $\tilde{S} = (S \cup S') \backslash S_1$, and setting $\mathfrak{q} = (\tilde{S} \overset{q}{\twoheadleftarrow} \tilde{S} \overset{i}{\hookrightarrow} T)$, we have $i^*_{\mathfrak{q}} L \in \mathrm{Br}_{pre}(\tilde{S})$, and there is a map of graphs $H_{i_{\mathfrak{q}^* L}} \to H_L$ which sends $[R]$ to $[R]$ when $R \subseteq \tilde{S}$. By induction, $H_{i_{\mathfrak{q}^* L}}$ is connected, so $[S \backslash S_1]$ and $[S']$ are in the same connected component in $H_{i_{\mathfrak{q}^* L}}$, and so also in H_L.

Step 2: Now we show that any brane has a grading.

By proposition 6.1.3, we can assume WLOG that $G = G_\mu$, so $\mathcal{O}_G = \mathcal{O}_\mu$. By proposition 6.2.1, we can write $L = L_A$ for some A_n-subtree A.

If $A = \{v\}$ has a single vertex, then $L_A = L_v$. For any subtree S with $v \in S$, define $g_{[S]} = 0$. Then for any pair of top-cells $[S], [S'] \in L$, if there exists a codimension-one cell z in the closure of both, we must have $S \subseteq S'$ or $S' \subseteq S$. In either case, the definition of G_μ says $\varphi^z_{[S],[S']} = +0$, so this satisfies the condition of being a grading.

If A has multiple vertices, let $V_+(A)$ denote the sources in A and $V_-(A)$ the sinks. For any two elements in $V_+(A)$ there is an element of $V_-(A)$ between them, and likewise with the signs reversed. So for any subtree S, the quantity $\chi_A(S) = \#(S \cap V_+(A)) - \#(S \cap V_-(A))$ is either 1, 0, or -1. $[S] \in L_A$ if and only if S has an odd number of elements of A, meaning $\chi_A(S) \pm 1$. Define $g_{[S]} = 0$ if $\chi_A(S) = 1$, $g_{[S]} = -1$ if $\chi_A(S) = -1$.

Now, let $[S], [S'] \in L_A$, and assume there is a codimension-one cell z in the closure of both. If $[S]$ and $[S']$ are disjoint, suppose the edge e in between is pointing from $[S]$ to $[S']$ (so $\varphi^z_{[S],[S]'} = -1$). Then the closest element in $S \cap (V_+(A) \cup V_-(A))$ to e is a source, which means $\#(S \cap V_+(A)) \geq \#(S \cap V_-(A))$, and so $\chi_A(S) = 1 \Rightarrow g_{[S]} = 0$. For the same reason $\chi_A(S') = -1$, so $g_{[S']} = -1$, we get $\varphi^z_{[S],[S']}(g_{[S]}) = 0 - 1 = -1 = g_{[S']}$ as desired.

Lastly, if $S \subseteq S'$, then $A \cap S \subset A \cap S'$ are two A_n trees in T. Also, one of the endpoints of $A \cap S$ must coincide with one of the endpoints of $A \cap S'$. If the closest vertex in $V_+(A) \cap V_-(A)$ to that endpoint is a source in A, $\chi_A(S) = \chi_A(S') = 1$, and both are -1 if that closest vertex is a sink, so $g_{[S]} = g_{[S']}$. We also have $\varphi^z_{[S],[S']} = +0$, so $\varphi^z_{[S],[S']}(g_{[S]}) = g_{[S']}$.

□

6.3 The Representation Theory of Graded Branes

We now turn to a representation-theoretic interpretation of graded branes. To begin, fix a directed tree (T, μ), and an object $X \in \mathrm{Rep}(T, \mu) = \mathcal{Q}_\mu(\mathcal{P}_T)$. We define the **support** of X to be the set $\mathrm{supp}(X) = \{[\mathfrak{p}] \in \mathcal{P}_T \mid c_\mathfrak{p}(X) \ncong 0\}$, which is a closed subset of $\mathcal{P}_T$.

Definition 6.3.1. *Let $X \in Rep(T, \mu)$. For any maximal element $[S] \in \mathcal{P}_T$, we have $c_{[S]}(X) \in Rep(\star, \cdot) \cong C(k)$, the dg-category of finite-dimensional chain complexes over k. X is **rank-one** if for all maximal elements $[S] \in supp(X)$, we have $H^\bullet(c_{[S]}(X))$ is one-dimensional. As a consequence, $c_{[S]}(X)$ is homotopy equivalent to a shift of k. If X is rank one and $[S] \in supp(S)$, define $g^X_{[S]} \in \mathbb{Z}$ by $c_{[S]}(X) \sim k[-g^X_{[S]}]$.*

Proposition 6.3.1. *For any rank-one object, $(supp(X), g^X) \in Br_{gr}(\mathcal{P}_\mu, G_\mu)$.*

Proof. By definition, is it clear that if $\mathfrak{p} : R \to T$ is a correspondence and $X \in \mathrm{Rep}(T, \mu)$ is rank one, then $c_{\mathfrak{p}}(X)$ is also rank-one.

So assume X is rank-one, we show that $(\mathrm{supp}(X), g^X)$ is a graded brane. Note that the condition of being a graded brane is determined by restriction to open neighborhoods of codimension-one cells, so by the above comment it is enough to prove this for $T = A_2$.

WLOG we can assume in μ the single edge points from vertex 1 to vertex 2. X is a direct sum of indecomposable objects. First assume X is indecomposable. The indecomposable objects in $\mathrm{Rep}(A_2, \mu)$ are, up to shifts and homotopy, P_1, P_2, and $P_2[1] \to P_1$ (see lemma 5.4.1). Since homotopy does not change the support, and shifts only change g up to a constant, neither will affect $(\mathrm{supp}(X), g^X)$ being a graded brane. Thus we only need to check these three objects, which is routine. We can also see from Table 5.1 that if X is a direct sum of more than one indecomposable object then X cannot be rank-one. □

We remark briefly that the object $P_2[1] \to P_1$ has $g^X_{[1]} = 0$ and $g^X_{[2]} = -1$, which fixes our convention on G_μ that the monodromy of structure maps in the cyclic order is -1.

We also have the following:

Proposition 6.3.2. *The map $X \mapsto (supp(X), g^X)$ gives a surjection:*

$$\left\{ \begin{array}{c} \textit{Rank-One Objects} \\ \textit{in Rep}(T, \mu) \end{array} \right\} \Big/ \textit{Homotpy Equivalence} \to \left\{ \begin{array}{c} \textit{Graded Branes} \\ \textit{in } (\mathcal{P}_\mu, G_\mu) \end{array} \right\}$$

Proof. For any graded brane $(L, g) \in (\mathcal{P}_\mu, G_\mu)$, we construct a rank one object X with $(\mathrm{supp}(X), g^X) = (L, g)$. We have $L = L_A$ for some A_n-subtree A, and by proposition 6.2.2 up to shifts we can assume g is the grading given in that proof. We repeat the explanation of g here: If $[S] \in L_A$, then $\chi_A(S) := \#(S \cap V_+(A)) - \#(S \cap V_-(A)) = 1$ or -1, we get $g_{[S]} = 0$ in the first case and $g_{[S]} = -1$ in the second case.

The object we claim works is given by the twisted complex (see Chapter 4):

$$X_A = \bigoplus_{v \in V_-(A)} P_v[1] \xrightarrow{d} \bigoplus_{w \in V_+(A)} P_w$$

Where $d = \Sigma e_{v,w}$, and the sum is over all $v \in V_-(A)$ and $w \in V_+(A)$ such that there is a directed path from w to v in (T, μ).

We check that $(\mathrm{supp}(X_A), g^{X_A}) = (L_A, g)$, as claimed. Let $P_\star$ denote the generator of $\mathrm{Rep}(\star)$, then for a subtree S we have:

$$c_{[S]}(X_A) = \bigoplus_{v \in S \cap V_-(A)} P_\star[1] \xrightarrow{d} \bigoplus_{w \in S \cap V_+(A)} P_\star$$

$$\cong k^{S \cap V_-(A)}[1] \xrightarrow{d} k^{S \cap V_+(A)}$$

We can write d in matrix notation as follows: First suppose $[S] \in L_A$. If $g_{[S]} = 0$, then moving along $S \cap A$ we have an alternating sequence $w_1, v_1, w_2, \ldots, w_n, v_n, w_{n+1}$ of sources/sinks, starting and ending with a source. Then the matrix entries are $(d)_{v_i,w_i} = (d)_{v_i,w_{i+1}} = 1$, and all everything else zero- from this we easily see $c_{[S]}(X_A) \cong k$, hence $[S] \in \mathrm{supp}(A)$ and $g_{[S]}^{X_A} = 0 = g_{[S]}$. Similarly, when $g_{[S]} = -1$, $S \cap A$ alternates between sinks and sources starting and ending with a source, and $c_{[S]}(X_A) \cong k[1]$, meaning $g_{[S]}^{X_A} = -1 = g_{[S]}$.

Lastly, suppose $[S] \notin L_A$. If S does not intersect $V_-(A) \cup V_+(A)$, $c_{[S]}(X_A) = 0$. If it does intersect, then since $[S] \notin L_A$, $\#(S \cap V_-(A)) = \#(S \cap V_+(A))$ and the map d in $c_{[S]}(X_A)$ is an isomorphism, so $c_{[S]}(X_A) \cong 0$. Either way, $[S] \notin \mathrm{supp}(X_A)$.

We also see from this calculation that X is rank-one. □

In fact, the above correspondence is actually a bijection. In order to show this, we prove a representation-theoretic lemma:

Lemma 6.3.1. *Let (T, μ) be a directed tree with more than one vertex, and let v be a terminal vertex of T. Write $\mathfrak{p}_v = (\star \xleftarrow{\sim} \{v\} \hookrightarrow T)$, and $\mathfrak{p}_v^\circ = (T\backslash\{v\} \xleftarrow{\sim} T\backslash\{v\} \hookrightarrow T)$. If we let $Rep_{\backslash v}(T, \mu)$ denote the subcategory of $Rep(T, \mu)$ consisting of objects X whose support does not contain $[\mathfrak{p}_v]$, the functor $c_{\mathfrak{p}_v^\circ} : Rep_{\backslash v}(T, \mu) \to Rep(T\backslash\{v\}, \mu)$ is a quasi-equivalence.*

Proof. Define the functors $j_v^0 : \mathrm{Rep}^0(\star, \cdot) \to \mathrm{Rep}^0(T, \mu)$ and $j_{\backslash v} : \mathrm{Rep}^0(T\backslash\{v\}, \mu) \to \mathrm{Rep}^0(T, \mu)$ by:

$$j_v^0(P_\star) = P_v$$
$$j_{\backslash v}^0(P_w) = P_w$$

Furthermore, if v is a source, we have natural transformations fitting into an exact sequence:

$$j_v^0 c_{\mathfrak{p}_v}^0 \to \mathrm{Id} \to j_{\backslash v}^0 c_{\mathfrak{p}_v^\circ}^0$$

Which, when applied to P_v yields:

$$P_v \xrightarrow{\sim} P_v \to 0$$

When applied to P_w for $w \neq v$ yields:

$$0 \to P_w \xrightarrow{\sim} P_w$$

And, when applied to $e_{w,v}$ for $w \neq v$, yields the map of complexes:

$$\begin{array}{ccccc} 0 & \xrightarrow{0} & P_w & \xrightarrow{e_{w,w}} & P_w \\ \downarrow{\scriptstyle 0} & & \downarrow{\scriptstyle e_{w,v}} & & \downarrow{\scriptstyle 0} \\ P_v & \xrightarrow{e_{v,v}} & P_v & \xrightarrow{0} & 0 \end{array}$$

By lemma 5.1.2 this data can be extended into an exact sequence of natural transformations in $\mathrm{Rep}(T,\mu)$. Since the image of $j_{\backslash v}$ lies in $\mathrm{Rep}_{\backslash v}(T,\mu)$, and $c_{\mathfrak{p}_v^\circ} j_{\backslash v}$ is the identity, $j_{\backslash v}$ is quasi-fully faithful. Additionally, if $X \in \mathrm{Ob}(\mathrm{Rep}_{\backslash v}(T,\mu))$, the exact sequence gives a homotopy equivalence $X \sim j_{\backslash v} c_{\mathfrak{p}_v^\circ} X$, so $j_{\backslash v}$ is quasi-essentially surjective. It follows that $j_{\backslash v}$ is a quasi-equivalence, and so $c_{\mathfrak{p}_v^\circ}$ is as well.

When v is a sink, the direction of the triangle is reversed but the argument is unchanged. □

Proposition 6.3.3. *The correspondence in proposition 6.3.2 is a bijection. In particular, any two rank-one objects in Rep(T,μ) with the same support are homotopy-equivalent up to a shift.*

Proof. Let $L \in \mathrm{Br}(\mathcal{P}_T, \mathcal{O}_\mu)$, and suppose $\mathrm{supp}(X) = \mathrm{supp}(Y) = L$, with X, Y rank-one objects. By the above lemma, if there is a terminal vertex v such that $[v] \notin L$, we can reduce to the tree $T\backslash\{v\}$. The classification result says there is always such a terminal vertex unless $T = A_n$.

In the case of A_n, for $n \geq 2$, by applying reflection isomorphisms we can assume that all edges point from vertex i to vertex $i+1$. If L contains both terminal vertices, then we must have $L = L_1 + L_n$. Recall (proposition 4.4.1) there is an automorphism ω of $(\mathcal{P}_{A_n}, \mathcal{O}_\mu)$ which acts as a cyclic rotation on K_{A_n}: Writing $K_{A_n} = \{e_0, e_1, \ldots, e_n\}$ we have $\overline{\omega}(e_i) = e_{i+1}$. We claim that $[1] \notin \omega(L)$.

Indeed, considered as a subset of K_{A_n}, we have $[1] = \{e_0, e_1\}^c$, and $[A_n] = \{e_0, e_n\}^c$, and note that $[A_n] \notin L$, and $\omega([A_n]) = [1]$, so $[1] \notin \omega(L)$. Hence we can apply this rotation, and then reduce to the tree $A_n\backslash\{1\}$.

This reduces us, inductively, to the base case $T = A_1$ where the claim is obvious. □

The above proposition suggest we should think of graded branes as objects in our sheaf of categories. This idea is discussed briefly at the end of this chapter.

6.4 Branes on A_n Singularities

In this section we give a brief discussion of $\mathrm{Br}(\mathcal{P}_{A_n}, \mathcal{O})$, where $\mathcal{O}$ is a cyclic structure.

Proposition 6.4.1. *Let μ be the orientation on A_n in which all vertices point from vertex i to $i+1$. Then:*

$$Br(\mathcal{P}_{A_n}, \mathcal{O}_\mu) = \{\emptyset\} \cup \{L_v \mid v \in V(A_n)\} \cup \{L_v + L_w \mid v \neq w \in V(A_n)\}$$

Proof. This is an easy corollary of proposition 6.2.1. □

Our first goal will be to describe branes in a more invariant way. Writing $\mathcal{P}_{A_n} \cong \{X \subseteq K_{A_n} \mid |X| \leq n-1\}$, with $K_{A_n} = \{e_0, e_1, \ldots, e_n\}$, one can show:

- For $1 \leq i \leq n$, $X \in L_i$ iff X does not contain the cyclic interval $[e_0, e_{i-1}]$, nor the cyclic interval $[e_i, e_{j-1}]$.
- For $1 \leq i < j \leq n$, $X \in L_i + L_j$ iff X does not contain the cyclic interval $[e_i, e_{j-1}]$, nor the cyclic interval $[e_j, e_{i-1}]$.

See appendix A for the notion of cyclic intervals. We see that we each nonzero brane corresponds with a way of dividing K_{A_n} into two disjoint nonempty cyclic intervals.

Notice that such a division can be seen as drawing a line connecting two 'gaps'. We can also think of a gap as an adjacent pair $g_i = \{e_{i-1}, e_i\}$ of elements in K_{A_n}. Then branes are lines connecting two such pairs. We diagram this identification explicitly for A_3 in Table 6.2:

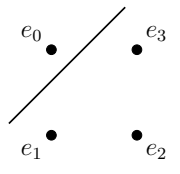
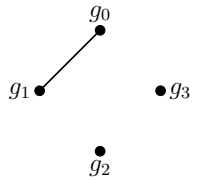
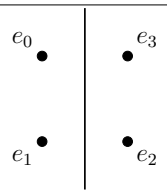
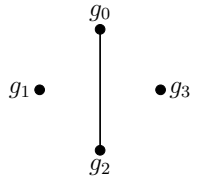
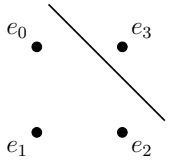
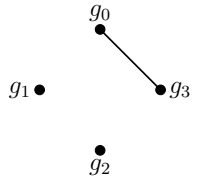
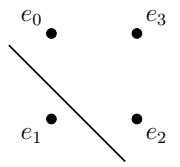
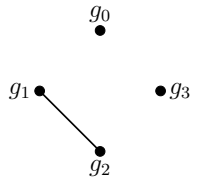
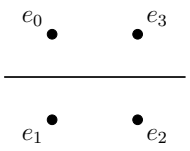
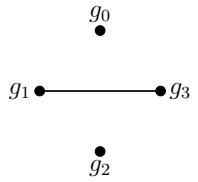
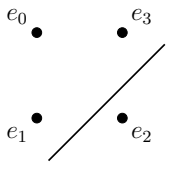
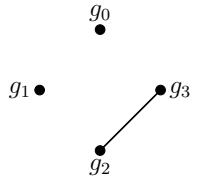
Brane	'Division' Diagram	'Gap' Diagram	Set-Theoretic Description
L_1			$\{X \subseteq K_{A_3} \mid$ $[e_0, e_0] \nsubseteq X$ $[e_1, e_3] \nsubseteq X\}$
L_2			$\{X \subseteq K_{A_3} \mid$ $[e_0, e_1] \nsubseteq X$ $[e_2, e_3] \nsubseteq X\}$
L_3			$\{X \subseteq K_{A_3} \mid$ $[e_0, e_2] \nsubseteq X$ $[e_3, e_0] \nsubseteq X\}$
$L_1 + L_2$			$\{X \subseteq K_{A_3} \mid$ $[e_1, e_1] \nsubseteq X$ $[e_2, e_0] \nsubseteq X\}$
$L_1 + L_3$			$\{X \subseteq K_{A_3} \mid$ $[e_1, e_2] \nsubseteq X$ $[e_3, e_0] \nsubseteq X\}$
$L_2 + L_3$			$\{X \subseteq K_{A_3} \mid$ $[e_2, e_2] \nsubseteq X$ $[e_3, e_1] \nsubseteq X\}$

Table 6.2: Branes in $\mathcal{P}_{A_3}$, Four Ways

In particular, there is a full diagrammatic calculus for branes:

Definition 6.4.1. *Let Δ^{cyc}_{surj} denote the category of finite sets with a cyclic order whose morphisms are order-preserving injections. We define a functor $Br : \Delta^{cyc}_{surj} \to \boldsymbol{Set}$ by:*

- *For a set C with a cyclic order, $Br(C)$ consists of the zero brane 0, and the line segments $L(x, y)$ joining pairs of elements in C.*

- *For a surjection $f : C \to D$, the map $B_f : Br(C) \to Br(D)$ is defined by $B_f(0) = 0$, $B_f(L(x,y)) = L(f(x), f(y))$ if $f(x) \neq f(y)$, and $B_f(L(x,y)) = 0$ if $f(x) = f(y)$.*

The main idea of this section is:

Proposition 6.4.2. *Let $\boldsymbol{Arb}_A^{cyc}$ denote the category whose objects are pairs $(\mathcal{P}_{A_n}, \mathcal{O})$, where $\mathcal{O}$ is a cyclic structure on $\mathcal{P}_{A_n}$ (see proposition 4.4). Then there is a contravariant functor $\Delta^\circ : \boldsymbol{Arb}_A^{cyc} \to \Delta_{surj}^{cyc}$ and a natural identification $Br(\mathcal{P}_{A_n}, \mathcal{O}) \cong Br(\Delta^\circ(\mathcal{P}_{A_n}, \mathcal{O}))$, such that for an open embedding $i : \mathcal{P}_{A_m} \hookrightarrow \mathcal{P}_{A_n}$ we have the commutative diagram:*

$$\begin{array}{ccc} Br(\mathcal{P}_{A_n}, \mathcal{O}) & \xrightarrow{\quad i^* \quad} & Br(\mathcal{P}_{A_m}, i^*\mathcal{O}) \\ \sim\downarrow & & \downarrow\sim \\ Br(\Delta^\circ(\mathcal{P}_{A_n}, \mathcal{O})) & \xrightarrow{B_{\Delta^\circ i}} & Br(\Delta^\circ(\mathcal{P}_{A_m}, i^*\mathcal{O})) \end{array}$$

Spelling it out:

- *$\Delta^\circ(\mathcal{P}_{A_n}, \mathcal{O})$ consists of sets of pairs $\{x, s(x)\} \in K_{A_n}$ which are adjacent in the cyclic order determined by $\mathcal{O}$. Denote this set $K^\circ_{A_n}$.*
- *An embedding $i : \mathcal{P}_{A_m} \hookrightarrow \mathcal{P}_{A_n}$ determines an embedding $\Delta(i) : K_{A_m} \hookrightarrow K_{A_n}$ (see definition 3.6.2). Then $\Delta^\circ(i) : K^\circ_{A_n} \twoheadrightarrow K^\circ_{A_m}$ sends the pair $\{x, y\}$ to the unique adjacent pair $\{a, b\}$ with $[x, y] \subseteq [\Delta(i)(a), \Delta(i)(b)]$.*
- *The identification $Br(\mathcal{P}_{A_n}, \mathcal{O}) \cong Br(K^\circ_{A_n})$ is as described at the beginning of this section and illustrated for $\mathcal{P}_{A_3}$ in Table 6.2.*

The proof is straightforward, and we omit it. This can be extended to describe graded Lagrangian branes as well:

Definition 6.4.2. *A **graded cyclic set** is a finite set C equipped with a cyclic order, along with:*

(i) For each $x \in C$, a $\mathbb{Z}-$torsor I_x.

(ii) For each $x \in C$, an isomorphism $\varphi_x : I_x \xrightarrow{\sim} I_{s(x)}$, where $s(x)$ denotes the successor function.

Note that if $y = s^k(x)$ we get an isomorphism $\varphi_x^k : I_x \xrightarrow{\sim} I_y$. Then we require:

(M1) If $|S| = n$, $m(\varphi_x^n) = 2 - n$.

The set of ***graded branes*** $Br_{gr}(C)$ *consists of the empty brane* 0 *as well as pairs* $(L(x,y),(\iota_x,\iota_y))$*, where* $x \neq y \in S$*, and if* $y = s^k(x)$ *we have* $\iota_y = \varphi^k(\iota_x) + k - 1$*. Note that the monodromy condition implies that this condition is symmetric in* x *and* y*.*

Then, just as before, we can functorially assign a graded cyclic set to a pair $(\mathcal{P}_{A_n}, G)$, where G is a graded structure on $\mathcal{P}_{A_n}$, such that the graded Lagrangian branes are identified . Explicitly, for a pair $\{x,y\} \in K^{\circ}_{A_n}$, we define the $\mathbb{Z}-$torsor $I_{\{x,y\}}$ to equal $G_{\{x,y\}^c}$, and the map $\varphi_{\{x,y\}}$ to equal $\varphi^{\{x,y,z\}^c}_{\{x,y\}^c,\{y,z\}^c}$. We leave further details to the reader.

6.5 In Search of a 'Fukaya Category'

The discussion in this section is largely informal, and refers to a 'Fukaya category' and 'A_∞ categories' without defining them. For references on Fukaya categories see [1], and for $A_\infty-$categories see [16].

To any cyclic C-Arb singularity $(\mathcal{P}_T, \mathcal{O})$, we explained in Chapter 5 how to construct a sheaf of dg-categories on $\mathcal{P}_T$ up to quasi-equivalence. This sheaf of categories is meant to model the Fukaya category of an arboreal singularity embedded in a symplectic manifold, or (following [25]) a category of microlocal sheaves. We discuss this perspective further in Chapter 8.

The objects of the Fukaya category $\mathrm{Fuk}(M,\omega)$ of a symplectic manifold are (graded) Lagrangian submanifolds of M. As shown in propositions 6.3.2 and 6.3.3, graded branes in $\mathcal{P}_T$ with respect to the graded structure g_μ have a natural interpretation in terms of objects in the sheaf $\mathcal{Q}_\mu$.

This raises a hope that, given a graded C-Arb singularity $(\mathcal{P}_T, G)$ one could define an $A_\infty-$category whose objects are the graded Lagrangian branes, which only depends on G (and not on, say, any isomorphism with a directed tree graded structure).

Such a goal has been partially realized in [20]. In that paper, Nadler shows how to (functorially) construct a ($\mathbb{Z}/2\mathbb{Z}-$graded) $A_\infty-$ category $\mathcal{C}_S$ associated to any cyclic set S, with $\mathrm{Ob}(\mathcal{C}_S) = S$. The triangulated hull of this category is quasi-equivalent to $\mathrm{Rep}(A_n, \mu)$, where $n+1 = |S|$. This construction can easily be modified using the notion of 'graded cyclic set' to construct a $\mathbb{Z}-$graded category.

A potential future goal, then, would be to extend Nadler's construction to arbitrary trees.

Chapter 7

Arboreal Moves

7.1 The H-to-I Move

In the theory of trivalent ribbon graphs there is a 'move' or 'mutation' that preserves the associated Riemann surface (Figure 7.1).

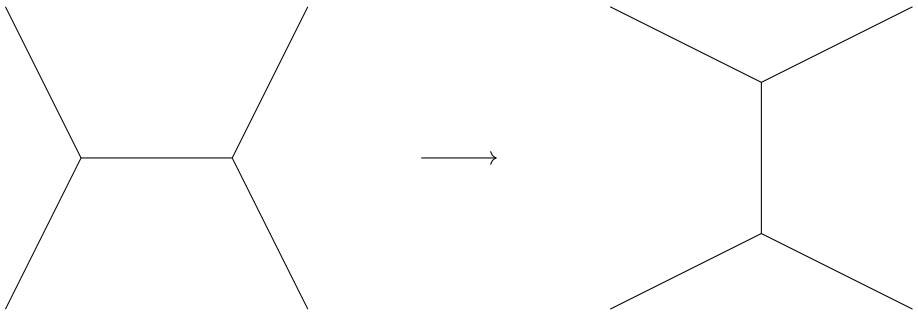


Figure 7.1: The H-to-I move

This move can be viewed as a transformation of cyclic C-Arb spaces. Viewing this transformation as a continuous deformation over time, the total space of this move is the A_3- singularity (Figure 7.2).

In this section, our goal will be to generalize this observation to describe a wide class of moves of C-Arb spaces whose total space is a C-Arb singularity. Before introducing the formal definition, we make a few comments about combinatorial aspects of the H-to-I move.

Combinatorially, we can describe the H-to-I move as follows: Begin with the cyclic C-Arb singularity $(\mathcal{P}_{A_3}, \mathcal{O}_\mu)$, with all arrows in μ pointing to the right. We identify elements of $\mathcal{P}_{A_3}$ with their image in $\mathscr{P}(K_{A_3})$, and recall (section) that the cyclic structure $\mathcal{O}_\mu$ corresponds

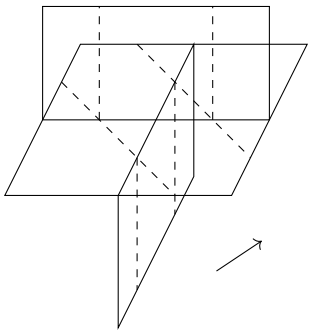


Figure 7.2: Total Space of the H-to-I move

to the cyclic order $e_0 < e_1 < e_2 < e_3$.

In the H-to-I move, the 'before' picture can be thought of as the poset $N(\{e_0\}) \cup N(\{e_2\})$, and the 'after' as $N(\{e_1\}) \cup N(\{e_3\})$. We redraw the move in Figure 7.3, this time labelling cells.

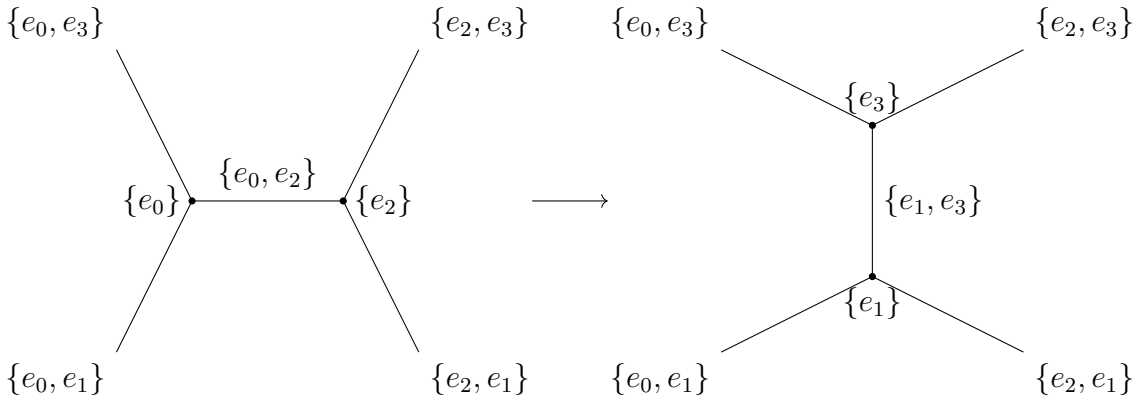


Figure 7.3: The H-to-I move

Note that in each case the induced cyclic structure is faithfully represented by the figure. For example, the cyclic order on the top-cells containing $\{e_0\}$ is induced by the total order $\{e_0, e_1\} < \{e_0, e_2\} < \{e_0, e_3\}$, and in the figure those cells appear in that order going counterclockwise about $\{e_0\}$.

It is also worth examining some non-examples. If we were to set 'before' to $N(\{e_0\}) \cup N(\{e_1\}) \cup N(\{e_2\})$ and 'after' to $N(\{e_3\})$, we would have the situation pictured in 7.4.

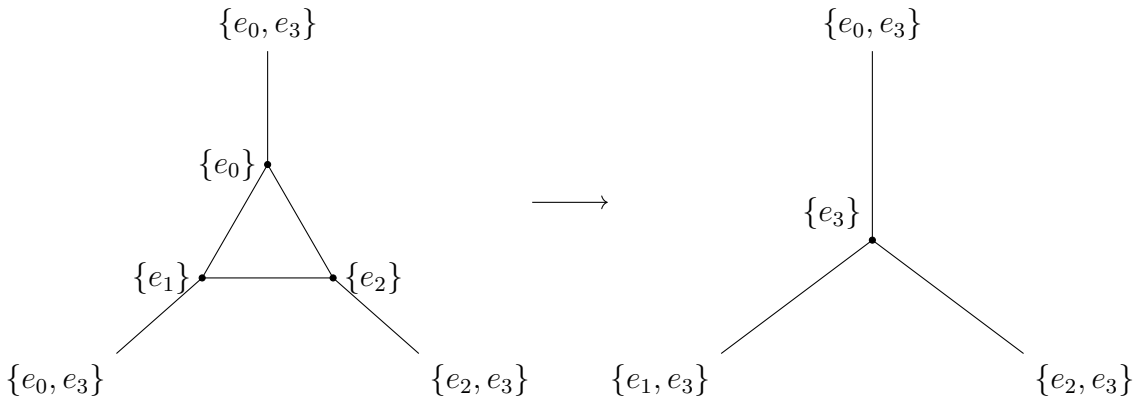


Figure 7.4: An Invalid Move

More subtlely, a move $N(\{e_0\}) \cup N(\{e_1\})$ to $N(\{e_2\}) \cup N(\{e_3\})$ would look like 7.5. The crossing paths in the figure is necessary to preserve the boundary and have the correct cyclic orders at each vertex.

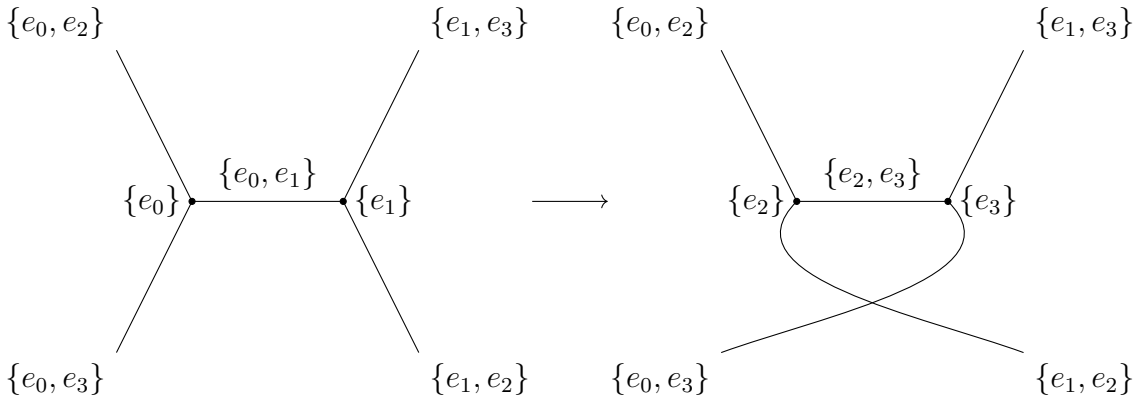


Figure 7.5: A Second Invalid 'H-to-I-esque' Move

We see that there are topological restrictions to valid moves- the move in Figure 7.4 is not a homotopy equivalence- and also restrictions imposed by the cyclic structure- the move in Figure 7.5 introduces unwanted 'crossings' if we want to maintain the cyclic orders about each vertex.

To try to capture these qualities in our combinatorial definition, we will use the sheaf of dg categories constructed in Chapter 5.

7.2 Definition of Arboreal Moves

Definition 7.2.1. *Given a tree T, a **combinatorial arboreal move of type** T between cyclic C-Arb spaces $(\mathcal{P}_1, \mathcal{O}_1)$ and $(\mathcal{P}_2, \mathcal{O}_2)$ consists of the following:*

(i) *An orientation μ on T*

(ii) *Open embeddings $i_\alpha : \mathcal{P}_\alpha \hookrightarrow \mathcal{P}_T$, with $i_\alpha^* \mathcal{O}_\mu = \mathcal{O}_\alpha$, for $\alpha = 1, 2$. ($\mathcal{P}_T$ is called the **total space** of the move).*

(iii) *Let U_α denote the image of i_α, and $\hat{0}$ the minimal element of $\mathcal{P}_T$. Then $U_1 \cup U_2 = \mathcal{P}_T \backslash \{\hat{0}\}$.*

(iv) *U_1 is the smallest open set containing $U_1 \backslash U_2$, and likewise with the indices switched.*

(v) ***Preservation of Categories:*** *The restrictions $\Gamma(\mathcal{P}_T; \mathcal{Q}_\mu) \to \Gamma(U_\alpha; \mathcal{Q}_\mu)$ are quasi-equivalences of dg-categories.*

We will frequently abbreviate "combinatorial arboreal move of type T" to simply "move of type T".

A brief remark on condition (iv): Note that $i_2 i_1^{-1}$ gives an isomorphism $i_1^{-1}(U_2) \xrightarrow{\sim} i_2^{-1}(U_1)$. Letting $Z_1 = (i_1^{-1}(U_2))^c$, and $Z_2 = (i_2^{-1}(U_1))^c$, we can think of the move as mutating Z_1 into Z_2. For example, In the case of the H-to-I move, Z_1, Z_2 are the internal closed line segments. Condition (iv) then says the mutation is 'local', in the sense that $\mathcal{P}_\alpha = N(Z_\alpha)$, for $\alpha = 1, 2$.

Definition 7.2.2. *Two moves of type T, given by the data $i_\alpha^1 : (P_\alpha, \mathcal{O}_\alpha) \hookrightarrow (\mathcal{P}_T, \mathcal{O}_\mu)$ and $i_\alpha^2 : (P_\alpha, \mathcal{O}_\alpha) \hookrightarrow (\mathcal{P}_T, \mathcal{O}_{\mu'})$ are **equivalent** if there is an isomorphism of total spaces $\varphi : (\mathcal{P}_T, \mathcal{O}_\mu) \xrightarrow{\sim} (\mathcal{P}_T, \mathcal{O}_{\mu'})$ such that $i_\alpha^2 = \varphi \circ i_\alpha^1$ for $\alpha = 1, 2$.*

Then a move is determined up to equivalence by the open sets U_α.

We introduce the following notation: For a subset $K \subset K_T$ we write:

$$U_K := \bigcup_{k \in K} N(\langle k \rangle)$$

Lemma 7.2.1. *If we have an move of type T then we can find a partition $K_T = K_1 \sqcup K_2$ such that $U_\alpha = U_{K_\alpha}$.*

Proof. Define $K_\alpha = \{k \in K_T \mid \langle k \rangle \in U_\alpha\}$. We will show that K_1, K_2 give a partition of K_T, and $U_\alpha = U_{K_\alpha}$.

Since $U_1 \cup U_2 = \mathcal{P}_{\vec{T}} \backslash \{\hat{0}\}$, we must have $\langle k \rangle \in U_1 \cup U_2$ for all $k \in K_T$. Hence $K_T = K_1 \cup K_2$. If $k \in K_1 \cap K_2$, the set $\{\langle k \rangle, \hat{0}\}$ is closed and violates property (iii). So $K_T = K_1 \sqcup K_2$.

Since U_α is open, we have $U_{K_\alpha} \subseteq U_\alpha$. To show equality, WLOG set $\alpha = 1$ and let $x \in U_1$. By property (iv), there exists $y \in U_1 \backslash U_2$ with $y \leq x$. Then there is $k \in K$ with $\langle k \rangle \leq y$. Since $y \notin U_2$ we must have $k \notin K_2$, so $k \in K_1$, hence $x \in U_{K_1}$. □

Definition 7.2.3. *A subset $K \subseteq K_T$ is **good** with respect to an orientation μ on T if the restriction $\Gamma(\mathcal{P}_T; \mathcal{Q}_\mu) \to \Gamma(U_K; \mathcal{Q}_\mu)$ is a quasi-equivalence.*

We can also say $K \subseteq K_T$ is good with respect to a cyclic structure $\mathcal{O}$ on $\mathcal{P}_T$ if for any orientation μ on T and $\varphi \in Aut(\mathcal{P}_T)$ with $\varphi^* \mathcal{O}_\mu = \mathcal{O}$, $\overline{\varphi}(K)$ is good with respect to μ. The Cyclic Structure Theorem (theorem 5.6.1) implies this definition makes sense.

Hence, up to equivalence, moves of type T are given by orientations μ on T and partitions $K_T = K_1 \sqcup K_2$ into sets which are good with respect to μ.

The following sections will be dedicated to a classification of good subsets of T with respect to μ.

7.3 The Topology of Branes

In this section we show a negative result: That $K = K_T$ is *not* a good subset of K_T. To accomplish this we need some topological information about branes of the form L_v, for $v \in V(T)$.

A **regular cell complex** is a cell complex such that the closure of any open cell is a closed cell. To a regular cell complex we can associate a poset, called the **face poset**, whose elements are cells and whose relation is given by $x \leq y$ iff $x \subset \overline{y}$. From Nadler ([20]), (alternatively, see Chapter 8), we have:

Proposition 7.3.1. *Let L_v denote the closed set of all $[\mathfrak{p}]$, where $\mathfrak{p} = (R \overset{q}{\twoheadleftarrow} S \overset{i}{\hookrightarrow} T)$ is a correspondence with $v \in V(S)$ (see definiton 6.2.2). Let $n = |T| - 1$ denote the top dimension of a cell in $\mathcal{P}_T$. Then:*

(i) $L_v \setminus \{\hat{0}\}$ *is the face poset of a* $n-1-$*sphere (with the convention that* $S^{-1} = \emptyset$*).*

(ii) For $v \neq w$, $(L_v \cap L_w) \setminus \{\hat{0}\}$ *is the face poset of a closed* $n-1-$*ball, whose interior is given by correspondences* $\mathfrak{q} = (R \overset{q}{\twoheadleftarrow} S \overset{i}{\hookrightarrow} T)$ *with* $q(v) = q(w)$.

Proof. Nadler's construction explicitly realizes $L_v \setminus \{\hat{0}\}$ as a face poset of the standard unit $n-1-$sphere $S^{n-1} \subset \mathbb{R}^n$. Similarly, $(L_v \cap L_w) \setminus \{\hat{0}\}$ is realized as the face poset of the closed subset of $S^{n-1} \subset \mathbb{R}^n$ cut out by a finite (nonzero) number of equations of the form $x_i \geq 0$, where x_i is a coordinate on $\mathbb{R}^n$. This space is homeomorphic to a closed $n-1-$ball, and the statement about the interior is also implied by the construction. □

For any $P_v, P_w \in \mathrm{Rep}^0(T, \mu)$, the assignment:

$$U \mapsto \mathrm{Hom}_{\mathcal{Q}_\mu}(P_v \mid_U, P_w \mid_U)$$

Is a sheaf on $\mathcal{P}_T$ valued in $\mathrm{Ch}(k)$ (the underlying $\infty-$category of bounded chain complexes over k) which we call $\mathscr{H}om_{\mathcal{Q}_\mu}(P_v, P_w)$.

Lemma 7.3.1. *We can describe the sheaf* $\mathscr{H}om_{\mathcal{Q}_\mu}(P_v, P_w)$ *explicitly: Let* $\mathfrak{p} = (R \overset{q}{\twoheadleftarrow} R \overset{i}{\hookrightarrow} T)$*. Then the stalk:*

$$\mathscr{H}om_{\mathcal{Q}_\mu}(P_v, P_w)_{[\mathfrak{p}]} = \begin{cases} k & v, w \in V(S) \text{ and all edges between } v \text{ and } w \\ & \text{pointing from } v \text{ to } w \text{ are contracted by } q \\ 0 & \text{Else.} \end{cases}$$

And the maps between stalks are the identity when both stalks are k*, and* 0 *otherwise.*

Proof. This follows immediately from the definitions of $\mathrm{Rep}^0(R, \lambda)$ for a tree R and the restriction functors $c_{\mathfrak{p}} : \mathrm{Rep}^0(T, \mu) \to \mathrm{Rep}^0(R, \lambda)$. □

Proposition 7.3.2. *Let* (T, μ) *be a directed tree, and let* $K = K_T$*. For* $v, w \in V(T)$*, write* $\mathcal{F}_{v,w} = \mathscr{H}om_{\mathcal{Q}_\mu}(P_v, P_w)$*. Then:*

(i) If $v \in V(T)$*, the restriction* $\Gamma(\mathcal{P}_T; \mathcal{F}_{v,v}) \to \Gamma(U_K; \mathcal{F}_{v,v})$ *is not a quasi-isomorphism.*

(ii) If $v \neq w \in V(T)$ *and there exists a directed path from* v *to* w*, then the restriction* $\Gamma(\mathcal{P}_T; \mathcal{F}_{v,w}) \to \Gamma(U_K; \mathcal{F}_{v,w})$ *is not a quasi-isomorphism.*

Proof. It will suffice to show that the cohomology $\Gamma^\bullet(\mathcal{P}_T; \mathcal{F}_{v,w})$ is not isomorphic to $\Gamma^\bullet(U_K; \mathcal{F}_{v,w})$.

Let $n = |T| - 1$ be the dimension of top-cells.

In case (i), $\Gamma^\bullet(\mathcal{P}_T; \mathcal{F}_{v,v}) \cong k$. By lemma 7.3.1, we have $\mathcal{F}_{v,v}$ is the constant sheaf on the closed subset $L_v \backslash \{\hat{0}\}$. Hence by proposition 7.3.1, when $n > 0$, $\Gamma^\bullet(U_K; \mathcal{F}_{v,v}) \cong H^\bullet(S^{n-1}) \cong k \oplus k[1-n]$. When $n = 0$, $\Gamma^\bullet(U_K; \mathcal{F}_{v,v}) \cong H^\bullet(S^{-1}) = 0$. In either case, $\Gamma^\bullet(U_K; \mathcal{F}_{v,v})$ is not isomorphic to k.

In case (ii), $\Gamma^\bullet(\mathcal{P}_T; \mathcal{F}_{v,w}) = 0$. Note $n \geq 1$, since $v \neq w$. $\mathcal{F}_{v,w}$ is zero outside the closed subset (of U_K), $Z_{v,w} = (L_v \cap L_w) \backslash \{\hat{0}\}$. Furthermore, since there is a directed path from v to w, the stalk $(\mathcal{F}_{v,w})_{[\mathfrak{p}]}$ is nonzero for $[\mathfrak{p}] \in Z_{v,w}$ only when $q(v) = q(w)$. By proposition 7.3.1 and lemma 7.3.1, we have:

$$\Gamma^\bullet(U_K; \mathcal{F}_{v,w}) \cong H^\bullet(B^{n-1}, \partial B^{n-1}) \cong k[1-n]$$

□

As a corollary, the restriction $\Gamma(\mathcal{P}_T; \mathcal{Q}_\mu) \to \Gamma(U_K; \mathcal{Q}_\mu)$ is not quasi-fully faithful when $K = K_T$. This particular result does not help us classify moves (As the decomposition $K_T = K_T \sqcup \emptyset$ obviously does not give a move), however the above proposition will serve as a crucial 'base case' in the proof of our main result later in this chapter.

7.4 Classification of Good Subsets

The Result

Let us state the main result first, then we will give a detailed proof.

Definition 7.4.1. *Let (T, μ) be a directed tree. A **generalized cyclic interval** is a subset of K_T of the form $I(v, w, D)$, where:*

- *$v, w \in V(T)$*
- *D is a subset of the terminal vertices $V_t(T)$ of T.*
- *$I(v, w, D) \subset K_T$ consists of all elements of D, all edges between a point in D and the shortest path joining v and w, and all edges on the shortest path joining v and w which are pointing from v to w.*

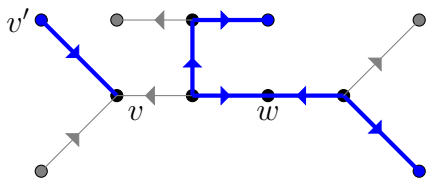


Figure 7.6: A Generalized Cyclic Interval $I(v, w, D)$.

An example is given in Figure 7.6. Elements of $I(v, w, D)$ are thick blue edges and blue terminal vertices, while elements not in $I(v, w, D)$ are thin gray lines and gray terminal vertices. D is the set of blue terminal vertices. We remark briefly that the generalized cyclic interval does not necessarily uniquely determine the data v, w, D: In the figure, we have $I(v, w, D) = I(v', w, D)$.

The reason for the term 'generalized cyclic interval' will be explained in the next section. The main result of this section is:

Theorem 7.4.1.

(i) $U_K \hookrightarrow K_T$ *is a weak equivalence with respect to* $\mathscr{H}om_{\mathcal{Q}_\mu}(P_v, P_w)$ *if and only if* K *is not a generalized cyclic interval of the form* $I(v, w, D)$.

(ii) $K \subset K_T$ *is good if and only if it is not a generalized cyclic interval.*

A proof of this result will involve some careful combinatorial arguments. The reader may skip directly to the examples for A_n and S_n trees with no crucial gaps.

Combinatorial Retracts

We introduce two simple combinatorial lemmas which will do most of the heavy lifting with respect to quasi-equivalences of sheaves of dg-categories. These lemmas can be stated more generally using the language of sheaves valued in an arbitrary $\infty-$category with small limits. We leave the proofs and definitions to appendix B. It is enough to understand the statements of these facts in order to follow the arguments in the following section.

Let $\mathcal{P}$ be a poset. We let $\mathrm{Sh}(\mathcal{P}; \mathcal{C})$ denote the category of sheaves on $\mathcal{P}$ valued in an $\infty-$category $\mathcal{C}$ with small limits.

Definition 7.4.2. *Let* $\mathcal{Q} \in Sh(\mathcal{P}; \mathcal{C})$. *The inclusion of a subset* $i : A \hookrightarrow \mathcal{P}$ *is a* ***weak equivalence*** *with respect to* $\mathcal{Q}$ *if the map of global sections* $\Gamma(\mathcal{Q}; \mathcal{P}) \to \Gamma(i^*\mathcal{Q}; A)$ *is an equivalence in* $\mathcal{C}$.

Definition 7.4.3. *Let $\mathcal{P}$ be a poset, and A any subset. A **combinatorial retract** from $\mathcal{P}$ to A is a map $r : \mathcal{P} \to A$ satisfying:*

(i) $r(x) = x$ *for* $x \in A$.

(ii) If $x \leq y$, $r(x) \leq r(y)$.

(iii) Either (a) $r(x) \leq x$ *for all* $x \in \mathcal{P}$, *or (b):* $r(x) \geq x$ *for all* $x \in \mathcal{P}$

*If (iii)(a) holds, r is a **downward combinatorial retract**, and if (b) holds it is an **upward combinatorial retract**.*

Lemma 7.4.1. *Let A be a subset of a poset $\mathcal{P}$, and $\mathcal{Q} \in Sh(\mathcal{P}; \mathcal{C})$.*

(i) If there exists a downward combinatorial retract $r : \mathcal{P} \to A$, $i : A \hookrightarrow \mathcal{P}$ is a weak equivalence with respect to $\mathcal{Q}$.

(ii) If there exists an upward combinatorial retract $r : \mathcal{P} \to A$ and a morphism $F : r^*i^*\mathcal{Q} \to \mathcal{Q}$ such that $i^*F : i^*\mathcal{Q} \to i^*\mathcal{Q}$ is the identity morphism, then $i : A \hookrightarrow \mathcal{P}$ is a weak equivalence with respect to $\mathcal{Q}$.

Proof. See appendix B (lemma B.4.1).

□

In practice, our sheaves will be defined using 2-functors $(\mathcal{Q}, c, \eta) : \mathcal{P} \to$ dg-Cat, see definition 5.2.1. When this is the case, when we have an upward combinatorial retract r, the required morphism in (ii) can be specified by a left or right morphism of 2-functors $i^*r^*\mathcal{Q} \to \mathcal{Q}$ (definition B.3.1) given by the data F, γ, such that, when $x \in A$, $\mathcal{F}_x = \mathrm{Id}_{\mathcal{Q}_x}$, and when $x \leq y \in A$, $\gamma_{xy} = \mathrm{Id}_{c_{xy}}$.

We will also deal with the case $\mathcal{C} = \mathrm{Ch}(k)$, the underlying $\infty-$category of bounded chain complexes over k with the standard (Quillen) model structure. In this case, sheaves of chain complexes will all specified by functors $\mathcal{P} \to \mathrm{Ch}(k)$, and morphisms of these specified by natural transformations.

We also have the following lemma:

Lemma 7.4.2. *Let $\mathcal{Q} \in Sh(\mathcal{P}; \mathcal{C})$. Suppose $\mathcal{C}$ has a zero object 0 (i.e. an object which is initial and terminal). If $j : U \hookrightarrow \mathcal{P}$ is the inclusion of an open subset of $\mathcal{P}$ such that $j^*\mathcal{Q} \cong 0$, then $i : \mathcal{P} \backslash U \hookrightarrow \mathcal{P}$ is a weak equivalence with respect to $\mathcal{Q}$.*

Proof. See appendix B (lemma B.4.2) □

Edge Expansions, Cuts, and Contractions

Let (T, μ) be a directed tree. We describe three combinatorial retracts which will be instrumental in all our arguments in this section.

If $e \in E(T)$ is an edge, write $A_e \subset \mathcal{P}_T$ for the subset consisting of all $[\mathfrak{p}]$, where $\mathfrak{p} = (R \overset{q}{\twoheadleftarrow} S \overset{i}{\hookrightarrow} T)$, such that either (i) $e \notin E(S)$ or (ii) $e \in E(S)$ and is not contracted by q.

Definition 7.4.4. *If $e \in E(T)$ is an edge, we have the **edge expansion** map $exp_e : \mathcal{P}_T \to A_e$ defined as follows: Let $\mathfrak{p} = (R \overset{q}{\twoheadleftarrow} S \overset{i}{\hookrightarrow} T)$ be a correspondence. Letting $e = \{v, w\}$, if $v, w \in S$ and $q(v) = q(w)$, we define $exp_e([\mathfrak{p}]) = [\mathfrak{p}']$, where $\mathfrak{p}' = \mathfrak{p}$ except the edge e is not contracted by q'. For all other $\mathfrak{p}$, $exp_e([\mathfrak{p}]) = [\mathfrak{p}]$. (See Figure 7.7)*

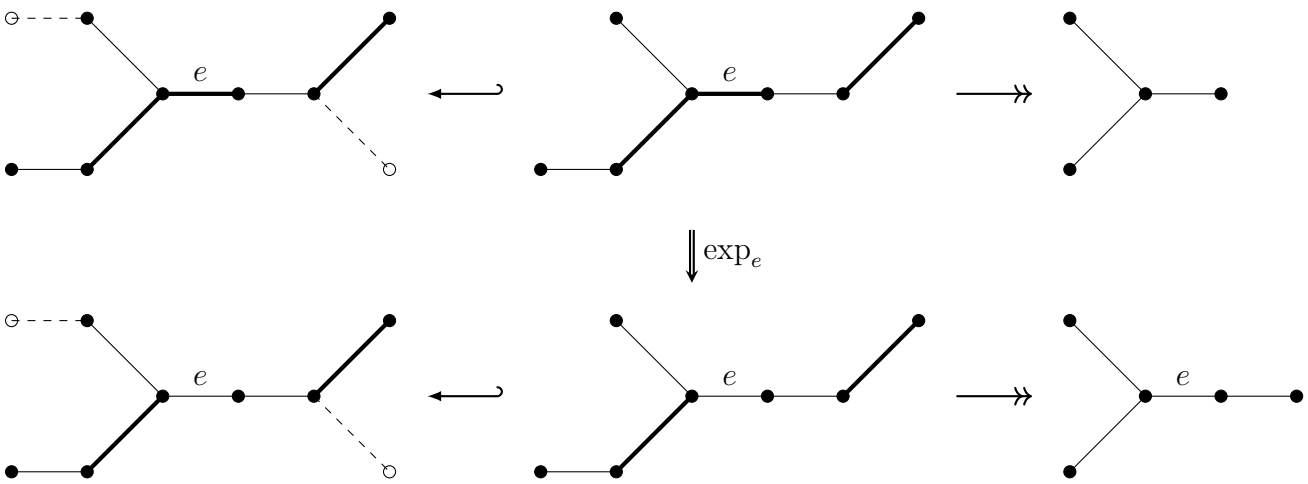


Figure 7.7: Illustrating The Edge Expansion $\exp_e$

Then $\exp_e$ is a downward combinatorial retract. Now, any edge e divides T into two disjoint subtrees Σ_1, Σ_2. For a subtree Σ of T, let $\mathfrak{p}_\Sigma := (\Sigma \overset{\sim}{\leftarrow} \Sigma \overset{i}{\hookrightarrow} T)$.

Definition 7.4.5. *For e, Σ_1, Σ_2 as above, define the **edge cut**:*

$$cut_e : A_e \backslash N([\mathfrak{p}_{\Sigma_2}]) \to N([\mathfrak{p}_{\Sigma_1}])$$

As follows: If $\mathfrak{p} = (R \overset{q}{\twoheadleftarrow} S \overset{i}{\hookrightarrow} T)$ and $[\mathfrak{p}] \in A_e \backslash N([\mathfrak{p}])$, then we know S intersects Σ_1, and $q(\Sigma_1)$ and $q(\Sigma_2)$ are disjoint in R (by virtue of e not being contracted). We define:

$$cut_e([\mathfrak{p}]) = [(R \cap q(\Sigma_1) \overset{q}{\twoheadleftarrow} S \cap \Sigma_1 \overset{i}{\hookrightarrow} T)]$$

(See Figure 7.8).

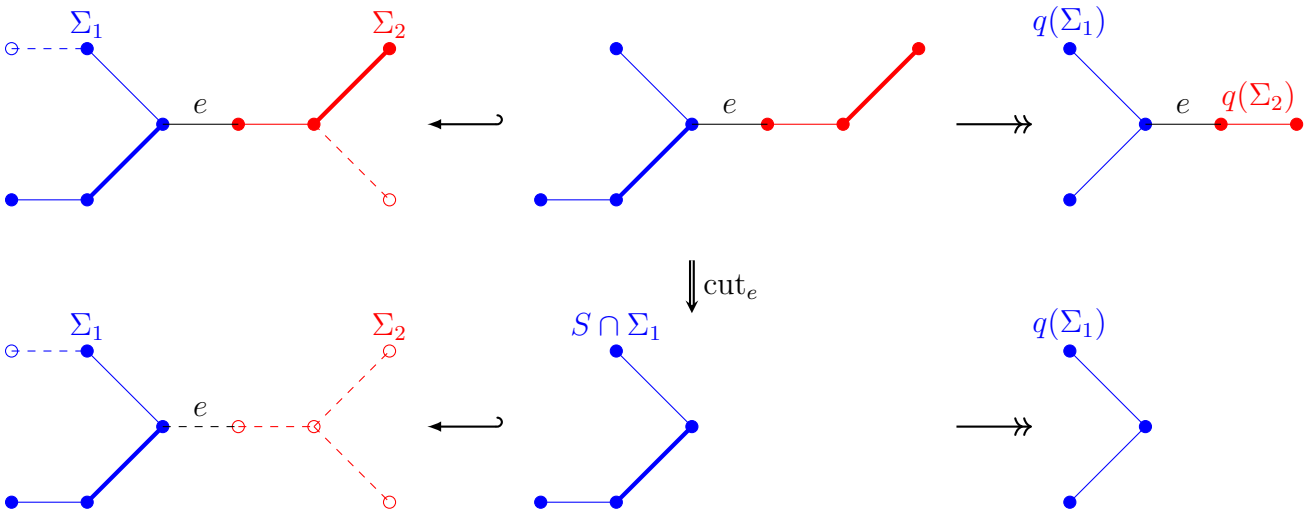


Figure 7.8: Illustrating the Edge Cut cut_e.

Then cut_e is an upward combinatorial retract. Finally, we define edge contraction- which is just as it sounds. If $e = \{v, w\}$, let $Z_e \subset \mathcal{P}_T$ be the closed subset $L_v \cap L_w$- that is, classes of correspondences $[(R \overset{q}{\twoheadleftarrow} S \overset{i}{\hookrightarrow} T)]$ with $v, w \in V(S)$ (so, $e \in E(S)$).

Definition 7.4.6. *Define the* ***edge contraction map*** $con_e : Z_e \to Z_e \cap N(\langle e \rangle)$, *where* $con_e([\mathfrak{p}])$ *is the same as* $[\mathfrak{p}]$ *except the edge* e *is contracted.*

Then con_e is an upward combinatorial retract.

Proof of Theorem 7.4.1

We will need some preliminary results:

Definition 7.4.7. *If* Σ *is a subtree of* T, *we define the subsheaf* $\mathcal{Q}_{\Sigma,\mu}$ *of* $\mathcal{Q}_\mu$ *as follows: Let* $Rep^0_\Sigma(T, \mu)$ *denote the full subcategory of* $Rep^0(T, \mu)$ *on the objects* $\{P_v\}_{v \in V(\Sigma)}$. *If* $\mathfrak{p} = (R \overset{q}{\twoheadleftarrow} S \overset{i}{\hookrightarrow} T)$ *is a correspondence, write* $\mathfrak{p}^*\Sigma = \{q(w), w \in V(\Sigma \cap S)\}$, *and* $\lambda = \mathfrak{p}^*\mu$. *Then the restriction functor* $c^0_\mathfrak{p}$ *maps* $Rep^0_\Sigma(T, \mu)$ *to* $Rep^0_{\mathfrak{p}^*\Sigma}(R, \lambda)$.

Then, this data assembles to define the subsheaf $\mathcal{Q}_{\Sigma,\mu}$.

Before we state our first result, let us introduce the following notation: Let Σ be a subtree of T, such that each terminal vertex of Σ is either a terminal vertex of T or is incident to a unique edge in $E(T)\backslash E(\Sigma)$. Then we have an inclusion $\kappa : K_\Sigma \hookrightarrow K_T$, which sends each edge in $E(\Sigma)$ to its image in $E(T)$, and every terminal vertex $v \in V_t(\Sigma)$ to either its image in $V_t(T)$ or the unique edge incident to V in $E(T)\backslash E(\Sigma)$.

Our first main result is the following:

Proposition 7.4.1. *Let $K \subset K_T$, and suppose $e \in K_T\backslash K$ is an edge in T. Removing e divides T into two subtrees Σ_1, Σ_2, where we assume e points from Σ_1 to Σ_2. Write $\kappa_\alpha : K_{\Sigma_\alpha} \hookrightarrow K_T$ for the inclusions. Then:*

(i) The inclusion $N([\mathfrak{p}_{\Sigma_\alpha}]) \cap U_K \hookrightarrow U_K$ is a weak equivalence with respect to $\mathcal{Q}_{\Sigma_\alpha,\mu}$ for $\alpha = 1, 2$.

(ii) the inclusion $U_K \hookrightarrow \mathcal{P}_T$ is a weak equivalence with respect to $\mathcal{Q}_{\Sigma_\alpha,\mu}$ if and only if $[\mathfrak{p}_{\Sigma_\alpha}] \in U_K$ or $\kappa_i^{-1}(\Sigma_i)$ is good with respect to (Σ_α, μ).

(iii) There is a functorial exact sequence of functors on $\Gamma(U_K; \mathcal{Q}_\mu)$:

$$\mathcal{F}_1 \to Id \to \mathcal{F}_2$$

Where the image of $\mathcal{F}_1$ lies in $\Gamma(U_K; \mathcal{Q}_{\Sigma_1,\mu})$, and the image of $\mathcal{F}_2$ lies in the image of $\Gamma(U_K; \mathcal{Q}_{\Sigma_2,\mu})$.

Proof. (i) WLOG $\alpha = 1$ (the direction of the edge e will not impact this step).

Write A_e for the set of correspondences in which e is either avoided or contracted, and let $\exp_e : \mathcal{P}_T \to A_e$ denote the edge expansion, which is a downward combinatorial retract. Since $e \notin K$, U_K is invariant under $\exp_e$. Abbreviate $A_{e,K} = A_e \cap U_K$- lemma 7.4.1 then implies that the inclusion $A_{e,K} \hookrightarrow U_K$ is a weak equivalence with respect to any sheaf on U_K.

Note that $\mathcal{Q}_{\Sigma_1,\mu}$ is zero when restricted to $N([\mathfrak{p}_{\Sigma_2}])$, hence (ref) the inclusion:

$$A_{e,K} \setminus N([\mathfrak{p}_{\Sigma_2}]) \hookrightarrow A_{e,K}$$

is a weak equivalence with respect to $\mathcal{Q}_{\Sigma_1,\mu}$. Then:

$$\mathrm{cut}_e : A_{e,K} \setminus N([\mathfrak{p}_{\Sigma_2}]) \to A_{e,K} \cap N([\mathfrak{p}_{\Sigma_1}])$$

Is an upward combinatorial retract. Writing $r = \mathrm{cut}_e$, we need to construct a map:

$$i^* r^* \mathcal{Q}_{\Sigma_1} \to \mathcal{Q}_{\Sigma_1}$$

Notice that, for $\mathfrak{p} = (R \xleftarrow{q} S \xhookrightarrow{i} T)$ with $[\mathfrak{p}] \in A_{e,K} \setminus N([\mathfrak{p}_{\Sigma_2}])$, we have:

$$\mathcal{Q}_{\Sigma_1,[\mathfrak{p}]} = \mathrm{Rep}_{q(\Sigma_1)}(R, \lambda) \qquad \mathcal{Q}_{\Sigma_1,\mathrm{cut}_e([\mathfrak{p}])} = \mathrm{Rep}_{q(\Sigma_1)}(R \cap q(\Sigma_1), \lambda)$$

There is a natural equivalence between these categories, which establishes that the inclusion:

$$A_{e,K} \cap N([\mathfrak{p}_{\Sigma_1}]) \hookrightarrow A_{e,K} \setminus N([\mathfrak{p}_{\Sigma_2}])$$

Is a weak equivalence with respect to $\mathcal{Q}_{\Sigma_1}$. Then since $N([\mathfrak{p}_{\Sigma_1}]) \subset A_e$:

$$A_{e,K} \cap N([\mathfrak{p}_{\Sigma_1}]) = U_K \cap N([\mathfrak{p}_{\Sigma_1}])$$

And we are done.

(ii) Consider the diagram of subsets of $\mathcal{P}_T$:

$$\begin{array}{ccc} N([\mathfrak{p}_{\Sigma_1}]) \cap U_K & \longrightarrow & N(\mathfrak{p}_{\Sigma_1}) \\ \downarrow & & \downarrow \\ U_K & \longrightarrow & \mathcal{P}_T \end{array}$$

The left arrow is a weak equivalence with respect to $\mathcal{Q}_{\Sigma_1}$, as was established in part (i), and the right arrow is clearly a weak equivalence. Therefore, the bottom arrow is a weak equivalence with respect to $\mathcal{Q}_{\Sigma_1}$ if and only if the top arrow is.

If $[\mathfrak{p}_{\Sigma_1}] \in U_K$, the top arrow is the identity map. If not, then all elements of K are either edges in $E(\Sigma_1)$ or terminal vertices of T which lie in $V(\Sigma_1)$. So for $\mathfrak{q} : R \to \Sigma_1$, $[\mathfrak{p} \circ \mathfrak{q}] \in N(\mathfrak{p}_{\Sigma_1}) \cap U_K$ if and only if $[\mathfrak{q}] \in U_{\kappa_1^{-1}(K)}$:

$$\begin{array}{ccc} N([\mathfrak{p}_{\Sigma_1}]) \cap U_K & \xrightarrow{\sim} & U_{\kappa_1^{-1}(K)} \\ \downarrow & & \downarrow \\ N([\mathfrak{p}_{\Sigma_1}]) & \xrightarrow{\sim} & \mathcal{P}_{\Sigma_1} \end{array}$$

So in this case, $N(\mathfrak{p}_{\Sigma_1}) \cap U_K \hookrightarrow N([\mathfrak{p}_{\Sigma_1}])$ is a weak equivalence with respect to $\mathcal{Q}_{\Sigma_1}$ if and only if $\kappa_1^{-1}(K)$ is good with respect to (Σ_1, μ).

The proof for $\alpha = 2$ is identical.

(iii) Since $A_{e,K}$ is a downward combinatorial retract of U_K, we can instead construct a triangle on $\Gamma(A_{e,K}; i^*\mathcal{Q}_\mu)$, where $i : A_{e,K} \hookrightarrow U_K$. We will do this by constructing functors of *sheaves*: $i^*\mathcal{Q}_\mu \to i^*\mathcal{Q}_\mu$, and then taking global sections.

The definition is straightforward. Let $\mathfrak{p} = (R \xleftarrow{q} S \xhookrightarrow{i} T)$ be a correspondence with $[\mathfrak{p}] \in A_{e,K}$, and let $P_v \in \mathrm{Rep}^0(R, \lambda)$. Since e is not contracted, the subsets $q(\Sigma_1), q(\Sigma_2)$ partition R (though one of these may be empty). Then if $v \in q(\Sigma_1)$, the exact sequence is:

$$P_v \to P_v \to 0$$

And if $v \in q(\Sigma_2)$, it is:

$$0 \to P_w \to P_w$$

We also need to know the behavior on morphisms: $e_{v,w}$ should be taken to a map of sequences. For $e_{v,w}$, if $v, w \in q(\Sigma_1)$ or $v, w \in q(\Sigma_2)$, the definition is clear. If not, recall that the edge e points from Σ_1 to Σ_2, so if $e_{v,w}$ exists we must have $v \in q(\Sigma_2)$ and $w \in q(\Sigma_1)$, and to this the functor assigns:

$$\begin{array}{ccccc} 0 & \xrightarrow{0} & P_v & \xrightarrow{e_{v,v}} & P_v \\ \downarrow{\scriptstyle 0} & & \downarrow{\scriptstyle e_{v,w}} & & \downarrow{\scriptstyle 0} \\ P_w & \xrightarrow{e_{w,w}} & P_w & \xrightarrow{0} & 0 \end{array}$$

That this is compatible with restrictions is evident from the definition.

□

Recall that for $v, w \in V(T)$, we have the sheaf $\mathscr{H}om_{\mathcal{Q}_\mu}(P_v, P_w)$. This is a sheaf valued in the $\infty-$category $\mathrm{Ch}(k)$, and is therefore amenable to the methods in the previous section. Namely:

Lemma 7.4.3. *Let $K \subseteq K(T)$, and $v, w \in V(T)$ with $v \neq w$. Then if either:*

(i) There exists an edge $e \in K$ between v and w pointing from w to v, or:

(ii) There exists an edge $e \notin K$ between v and w pointing from v to w

Then the inclusion $i : U_K \hookrightarrow \mathcal{P}_T$ is a weak equivalence with respect to $\mathscr{H}om_{\mathcal{Q}_\mu}(P_v, P_w)$.

Proof. Abbreviate $\mathcal{F}_{v,w} = \mathscr{H}om_{\mathcal{Q}_\mu}(P_v, P_w)$.

(i) Let Z_e be the closed subset of $\mathcal{P}_T$ of correspondences containing the edge e, and $\mathrm{con}_e : Z_e \to Z_e \cap N(\langle e \rangle)$ the edge contraction map. Any correspondence $\mathfrak{p} = (R \xleftarrow{q} S \xhookrightarrow{i} T)$ with $v, w \in V(S)$ must be in Z_e, so $\mathcal{F}_{v,w} \cong 0$ on $U_K \setminus Z_e$, meaning $Z_e \hookrightarrow U_K$ is a weak

equivalence with respect to $\mathcal{F}_{v,w}$.

Since e points from w to v, contracting e cannot change whether there is a directed path from w to v. This means, for $[\mathfrak{p}] \in Z_e$,

$$(\mathcal{F}_{v,w})_{[\mathfrak{p}]} \to (\mathcal{F}_{v,w})_{r[\mathfrak{p}]}$$

Is an isomorphism of chain complexes. This implies by lemma 7.4.1 that the inclusion:

$$Z_e \cap N(\langle e \rangle) \hookrightarrow Z_e$$

Is a weak equivalence with respect to $\mathcal{F}_{v,w}$. Now, consider the diagram:

$$\begin{array}{ccc} Z_e \cap N(\langle e \rangle) & \longrightarrow & U_K \\ \downarrow & & \downarrow \\ N(\langle e \rangle) & \longrightarrow & \mathcal{P}_T \end{array}$$

We've shown the top arrow is a weak equivalence with respect to $\mathcal{F}_{v,w}$. The left arrow is one since both subsets admit a downward combinatorial retract to $\langle e \rangle$. The bottom one is a weak equivalence since the map of stalks:

$$(\mathcal{F}_{v,w})_{\hat{0}} \to (\mathcal{F}_{v,w})_{\langle e \rangle}$$

Is an isomorphism. Hence the right arrow is as well.

(ii) Since there is an edge between v and w pointing from v to w, $\Gamma(\mathcal{P}_T; \mathcal{F}_{v,w}) \cong 0$.

Let $i : U_K \hookrightarrow \mathcal{P}_T$ denote the inclusion. Writing $A_{e,K} = A_e \cap U_K$ as before, we know the inclusion $j : A_{e,K} \hookrightarrow U_K$ is a weak equivalence with respect to any sheaf of $\infty-$categories on U_K. Since e points from v to w, $\mathcal{F}_{v,w}$ is zero on $A_{e,K}$, meaning $\Gamma(A_{e,K}; j^* i^* \mathcal{F}_{v,w}) \cong \Gamma(U_K; i^* \mathcal{F}_{v,w}) \cong 0$.

□

We are now ready for a proof of theorem 7.4.1:

(i) $U_K \hookrightarrow K_T$ is a weak equivalence with respect to $\mathscr{H}om_{\mathcal{Q}_\mu}(P_v, P_w)$ if and only if K is not a generalized cyclic interval of the form $I(v, w, D)$.

(ii) $K \subset K_T$ is good if and only if it is not a generalized cyclic interval.

Proof. (i) Let $v, w \in V(T)$, $K \subseteq K_T$. If $K = I(v, w, D)$ for some D, then D must equal the set of terminal vertices of T which are in K. Unless explicitly stated otherwise, for the rest of this proof D will refer to this set. To abbreviate notation, we let $\mathcal{F}_{v,w} = \mathscr{H}om_{\mathcal{Q}_\mu}(P_v, P_w)$. We divide the analysis into cases.

Case 1: There exists an edge in K between v and w pointing from w to v, or an edge not in K between v and w pointing from v to w.

By lemma 7.4.1 the inclusion $U_K \hookrightarrow \mathcal{P}_T$ is a weak equivalence with respect to $\mathcal{F}_{v,w}$. If either of these hold, $K \neq I(v, w, D)$.

Case 2: There is a vertex $u \in D$, and an edge $e \notin K$ which is between u and the shortest path joining v and w.

Then $K \neq I(v, w, D)$. Additionally, e separates T into two disjoint subtrees, Σ_1 and Σ_2, WLOG $v, w \in V(\Sigma_1)$. Since $\langle u \rangle \leq [\mathfrak{p}_{\Sigma_1}]$ we have $[\mathfrak{p}_{\Sigma_1}] \in U_K$, and so proposition 7.4.1 implies the inclusion $U_K \hookrightarrow \mathcal{P}_T$ is a weak equivalence with respect to $\mathcal{Q}_{\Sigma_1,\mu}$. P_v and P_w are both objects of the subcategory $\mathcal{Q}_{\Sigma_1,\mu}(\mathcal{P}_T)$, and so $U_K \hookrightarrow \mathcal{P}_T$ is a weak equivalence with respect to $\mathcal{F}_{v,w}$.

Case 3: Cases 1,2 do not hold. This gives $I(v, w, D) \subseteq K$. Suppose there is an edge e not between v and w which is not in K.

Then e separates T into two subtrees Σ_1, Σ_2, WLOG $v, w \in \Sigma_1$- by our assumption, then Σ_2 cannot contain any elements of D. If $[\mathfrak{p}_{\Sigma_1}] \in U_K$, then by the same reasoning as the previous paragraph, $U_K \hookrightarrow \mathcal{P}_T$ is a weak equivalence with respect to $\mathcal{F}_{v,w}$, and K contains some edge in Σ_2, hence is not equal to $I(v, w, D)$. If $[\mathfrak{p}_{\Sigma_1}] \notin U_K$, we can draw the same diagram of subsets as in the proof of 7.4.1:

$$\begin{array}{ccc} N(\mathfrak{p}_{\Sigma_1}) \cap U_K & \longrightarrow & N(\mathfrak{p}_{\Sigma_1}) \\ \downarrow & & \downarrow \\ U_K & \longrightarrow & \mathcal{P}_T \end{array}$$

By the same reasoning as in that proof, the bottom arrow of this diagram is a weak equivalence with respect to $\mathcal{F}_{v,w}$ if and only if the top arrow is. Furthermore, $K =$

$I(v, w, D) \subset K_T$ if and only if $\kappa_1^{-1}(K) = I(v, w, D) \subset K_{\Sigma_1}$ (recall $D \subset \Sigma_1$). So we can inductively assume K contains all edges except those which are between v and w pointing from w to v.

Case 4: K contains all edges except those which are between v and w pointing from w to v.

If this is the case, K is a generalized cyclic interval if and only if D contains every terminal vertex in $V_t(T)$, except for v or w (if those happen to be terminal vertices). To see why, if there were a terminal vertex $u \notin D$ which is not equal to v or w, then we have $u \notin D$, but the unique edge incident to u is in K, which cannot be the case if $K = I(v, w, D)$.

Case 4a: K is not a generalized cyclic interval.

That is, there is a terminal vertex of u not in D which is not equal to v or w. Let ω_u denote the reflection isomorphism at u. $U_K \hookrightarrow \mathcal{P}_T$ is a weak equivalence with respect to $\mathcal{F}_{v,w}$ if and only if $U_{\overline{\omega}_u}(K) \hookrightarrow \mathcal{P}_T$ is a weak equivalence with respect to $\mathscr{H}om_{\mathcal{Q}_{r_u\mu}}(\hat{P}_v, \hat{P}_w)$. Since $\overline{\omega}_u$ acts as a transposition of u and the unique edge e incident to u, and $e \in K$, this puts us in Case 2.

Case 4b: K is a generalized cyclic interval.

Finally, we assume K contains all edges except for the ones between v and w pointing from v to w, and all terminal vertices not equal to v or w, so K is a generalized cyclic interval. Let $Z_{v,w}$ denote the set of correspondences $[\mathfrak{p}] = [(R \overset{q}{\twoheadleftarrow} S \overset{i}{\hookrightarrow} T)]$ with $v, w \in V(S)$. Let $A_{v,w}$ be the open subset of $Z_{v,w}$ consisting of correspondences where all edges pointing from w to v (those which are not in K) are contracted. Lastly, let $\mathfrak{p}_{v,w}$ denote the correspondence $(R \overset{q}{\twoheadleftarrow} T \overset{\sim}{\to} T)$, where q contracts all the edges not in K. Note that $A_{v,w} = Z_{v,w} \cap N([\mathfrak{p}_{v,w}])$. We have the following diagram of subsets:

$$\begin{array}{ccccc} A_{v,w} \cap U_K & \longrightarrow & Z_{v,w} \cap U_K & \longrightarrow & U_K \\ \downarrow & & & & \downarrow \\ N([\mathfrak{p}]) \setminus \{[\mathfrak{p}_{\mathfrak{v},\mathfrak{w}}]\} & \longrightarrow & N([\mathfrak{p}_{v,w}]) & \longrightarrow & \mathcal{P}_T \end{array}$$

Now, 7.3.2 implies the bottom-left horizontal arrow is *not* a weak equivalence with respect to $\mathcal{F}_{v,w}$. It will follow that the right vertical arrow also is not, as long as the

other four arrows are.

The top-right horizontal arrow is a weak equivalence with respect to $\mathcal{F}_{v,w}$, since $\mathcal{F}_{v,w}$ is zero on the complement of $Z_{v,w}$. The left vertical arrow is a weak equivalence for the same reason: K contains all edges not contracted by $\mathfrak{p}_{v,w}$, and all terminal vertices not equal to v or w (and no correspondence in $A_{v,w}$ is in $N(\langle v \rangle)$ or $N(\langle w \rangle)$ if either is a terminal vertex) we get:

$$A_{v,w} \cap U_K = (N([\mathfrak{p}_{v,w}]) \backslash \{[\mathfrak{p}_{v,w}]\} \cap Z_{v,w}$$

Now consider the top-left horizontal arrow. The map $r : Z_{v,w} \to A_{v,w}$ given by applying all edge contractions con_e for e between v and w pointing from w to v is an upward combinatorial retract. Also, contracting edges pointing from w to v will not affect whether there is a directed path from w to v, hence the restriction $\mathcal{F}_{v,w}(N(x)) \to \mathcal{F}_{v,w}(N(r(x)))$ is an isomorphism for any x. This implies $A_{v,w} \hookrightarrow Z_{v,w}$ is a weak equivalence with respect to $\mathcal{F}_{v,w}$. Furthermore, since none of the contracted edges are in K, r maps $Z_{v,w} \cap U_K$ to $A_{v,w} \cap U_K$.

Lastly, the map of stalks:

$$(\mathcal{F}_{v,w})_{\hat{0}} \to (\mathcal{F}_{v,w})_{[\mathfrak{p}_{v,w}]}$$

Is an isomorphism, which implies the bottom-right horizontal arrow is also a weak equivalence with respect to $\mathcal{F}_{v,w}$. This completes the proof of Case 4b, and thus of (i).

(ii) By part (i), we know that the functor $\Gamma(\mathcal{P}_T; \mathcal{Q}_\mu) \to \Gamma(U_K; \mathcal{Q}_\mu)$ is quasi fully faithful if and only if K is not a generalized cyclic interval, so it suffices to show that, if K is not a generalized cyclic interval, the restriction is quasi essentially surjective.

If $K = K_T$, then K is a generalized cyclic interval. If K contains all edges but omits at least one terminal vertex, we can apply a reflection isomorphism at that terminal vertex. So we can assume there is an edge e which is not in K. Then e divides T into two subtrees Σ_1 and Σ_2.

Since the restriction functor is quasi fully faithful, its essential image is a strongly pretriangulated subcategory of $\Gamma(U_K; \mathcal{Q}_\mu)$. By proposition 7.4.1, if its essential image contains $\Gamma(U_K; \mathcal{Q}_{\Sigma_1,\mu})$ and $\Gamma(U_K; \mathcal{Q}_{\Sigma_2,\mu})$, it is quasi essentially surjective. Thus it suffices to show the inclusion $U_K \hookrightarrow \mathcal{P}_T$ is a weak equivalence with respect to $\mathcal{Q}_{\Sigma_1,\mu}$ and $\mathcal{Q}_{\Sigma_2,\mu}$.

By proposition 7.4.1, if $[\mathfrak{p}_{\Sigma_1}], [\mathfrak{p}_{\Sigma_2}] \in U_K$, this requirement is satisfied. If not, say $[\mathfrak{p}_{\Sigma_1}] \notin U_K$, all edges and terminal vertices in K are internal to Σ_1. Since K is

nonempty, we get $[\mathfrak{p}_{\Sigma_2}] \in U_K$. If $\Sigma_1 = \{u\}$ were a single (necessarily terminal) vertex, this would force $K = \emptyset$ or $K = \{u\}$, each of which are generalized cyclic intervals, so this cannot occur. So assume Σ_1 has more than one vertex. Using the notation from the beginning of this section: If $\kappa_1^{-1}(K)$ were a generalized cyclic interval $I(v, w, D) \subset K_{\Sigma_1}$, we would have $K = I(v, w, D)$ is also a generalized cyclic interval. Hence we can assume $\kappa^{-1}(K)$ is not a generalized cyclic interval, and therefore by induction is good with respect to (Σ_1, μ). Then by proposition 7.4.1, $U_K \hookrightarrow \mathcal{P}_T$ is a weak equivalence with respect to $\mathcal{Q}_{\Sigma_1,\mu}$.

□

7.5 Examples for A_n and S_n Trees

In this section, we discuss moves with total space A_n and S_n.

Moves of type A_n

Recall (proposition 4.4) that cyclic structures on $\mathcal{P}_{A_n}$ are in bijection with cyclic orders on K_{A_n}. The key result is that, in this case, **generalized cyclic intervals are cyclic intervals**. (See appendix A for the notion of a cyclic interval). This explains the terminology. This is the content of the following proposition:

Proposition 7.5.1. *A subset $K \subseteq K_{A_n}$ is good with respect to a cyclic structure if and only if K is not a cyclic interval in the associated cyclic order on K_{A_n}. (See appendix A for the definition of cyclic interval).*

Proof. It suffices to consider the case where the cyclic structure is $\mathcal{O}_\mu$, and μ is the orientation where all edges are pointing to the right. If we let e_i denote the edge $\{i, i+1\}$, e_0 denote the vertex 1 and e_n denote the vertex n, proposition 3.6.2 implies that the cyclic order on K_{A_n} is induced by the total order $e_0 < e_1 < \cdots < e_n$.

Consider the generalized cyclic interval $I(i, j, D)$. We consider possible cases for i, j, D.

If $i < j$ and $D = \emptyset$, then $I(i, j, D)$ is the closed interval $[e_i, e_{j-1}]$. If $D = \{1\}$, then $I(i, j, D)$ is $[e_0, e_{j-1}]$. If $D = \{2\}$, $I(i, j, D) = [e_i, e_n]$. If $D = \{1, 2\}$, then $I(i, j, D) = K_{A_n}$. Each of these are cyclic intervals.

If $i \geq j$ and $D = \emptyset$, then $I(i, j, D) = \emptyset$. If $D = \{1\}$, $I(i, j, D) = [e_0, e_{j-1}]$. If $D = \{2\}$, $I(i, j, D) = [e_i, e_n]$. Finally, if $D = \{1, 2\}$, $I(i, j, D) = [e_i, e_{j-1}]$ (where this interval wraps

around). Each of these are cyclic intervals.

Then we see any generalized cyclic interval is a cyclic interval, and every cyclic interval is a generalized cyclic interval. □

We can illustrate a partition $K_{A_n} = K_1 \sqcup K_1$ by representing elements of K_{A_n} as points on a circle (respecting the cyclic order), and filling in the elements of K_1 while keeping the elements of K_2 empty circles. Let us examine some small cases.

For K_{A_3} there are four possible partitions up to equivalence and switching 1 and 2. See Figure 7.9. Only the fourth gives a move-the 'H-to-I' move. The second and third are the non-moves illustrated at the beginning of this chapter (Figures 7.4 and 7.5). The first with $K_2 = \emptyset$ clearly does not give a move.

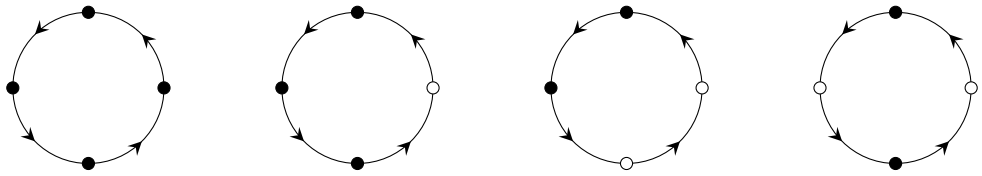


Figure 7.9: Possible Partitions of K_{A_3}- Only the Fourth Gives a Move

When $T = A_4$, up to equivalence and switching 1 and 2 there is only one possible move with total space $\mathcal{P}_{A_4}$, pictured in Figure 7.10.

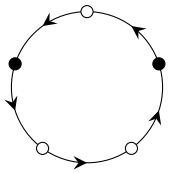


Figure 7.10: The Partition Givng The Unique A_4 Move Up To Equivalence/Reversal

We can give a topological picture of this move, as a continuous move of cell complexes (Figure 7.11). It looks a bit like a tetrahedral H-to-I move.

The above picture does not really capture the symplectic geometry of the move. We'll investigate that more in the final section.

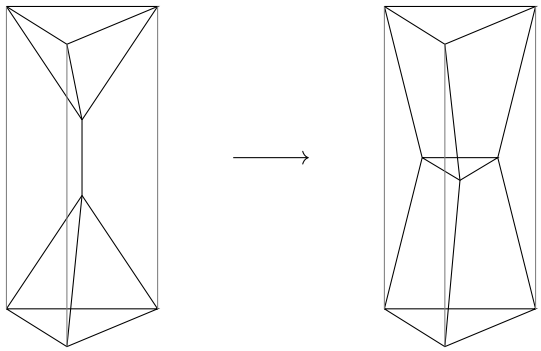


Figure 7.11: A Topological Picture of The A_4 Move

Moves of type S_n

By proposition 3.7.1, we know cyclic structures on $\mathcal{P}_{S_n}$ are in bijection with polarizations on K_{S_n}.

Proposition 7.5.2. *A subset $K \subseteq K_{S_n}$ is a good with respect to a cyclic structure if and only if either:*

(i) There are at least three equivalence classes of K_{S_n} such that K contains exactly one element of that equivalence class.

(ii) There are two equivalence classes of K_{S_n} such that K contains exactly one element of that equivalence class, and those two elements have the same sign with respect to the polarization.

Proof. As in the previous part, we can prove this by looking at all possible generalized cyclic intervals. Denote the elements of K_{S_n} by e_i, v_i, with $1 \leq i \leq n$. Let us use the term 'prohibited set' to denote the subsets of K_{S_n} which are not of the form described in the proposition. Then our goal is to show that generalized cyclic intervals are exactly the prohibited sets.

WLOG we can assume the cyclic structure is $\mathcal{O}_\mu$, where all edges are pointing outward. The $\rho(e_i) = -$ and $\rho(v_i) = +$ (see the proof of proposition 3.7.1).

Letting 0 denote the central vertex, the generalized cyclic interval $I(0, 0, D)$ either contains each equivalence class $\{e_i, v_i\}$ or is disjoint from it, depending on whether $v_i \in D$, hence is a prohibited set. Furthermore, each prohibited set which either contains or is disjoint from each equivalence class is of this kind.

The generalized cyclic intervals $I(0, v_i, D)$ and $I(v_i, 0, D)$ are disjoint from each equivalence class $\{e_j, v_j\}$ with $j \neq i$ and $v_j \notin D$, and contains each equivalence class $\{e_j, v_j\}$ with

$j \neq i$ and $v_j \in D$, so are prohibited sets. Furthremore, each prohibited set which either contains or is disjoint from all but one equivalence class if of this kind.

$I(v_i, v_i, D)$ contains each equivalence class $\{e_j, v_j\}$ with $j \neq i$ and $v_j \in D$, and is disjoint from each equivalence class $\{e_j, v_j\}$ with $j \neq i$ and $v_j \notin D$, so is a prohibited set.

$I(v_i, v_j, D)$ contains each equivalence class $\{e_k, v_k\}$ with $k \neq i, j$ and $v_k \in D$, and is disjoint from each equivalence class $\{e_k, v_k\}$ with $k \neq i, j$ and $v_k \notin D$. Furthermore, $I(v_i, v_j, D)$ contains e_j, but not e_i, by the direction of the arrows. Therefore if $I(v_i, v_j, D)$ contains exactly one element of $\{e_i, v_j\}$ and of $\{e_j, v_j\}$, it contains e_j and v_i, which have opposite signs. (This occurs when $v_i \in D$ but $v_j \notin D$). Hence $I(v_i, v_j, D)$ is a prohibited set, and any prohibited set which either contains or is disjoint from all but two equivalence classes is of this kind.

□

As a corollary, we can enumerate moves with total space $\mathcal{P}_{S_4}$, up to equivalence. We can diagram K_{S_4} as three pairs of objects, each pair consisting of a + object and a − object with respect to the polarization. As before, we fill in the elements of K_1 while keeping elements of K_2 empty. Then the four possible moves, up to equivalence and switching 1 and 2, are pictured in Figure 7.12.

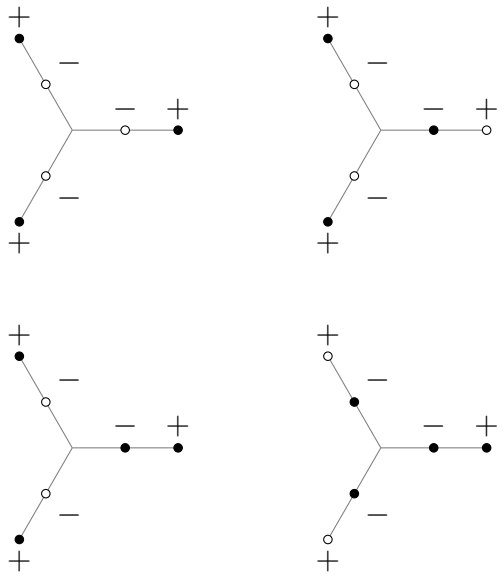


Figure 7.12: The Four Distinct Partitions of K_{S_4} Giving Moves

Chapter 8

The Geometry of Arboreal Spaces

This chapter is an overview of the geometric theory of arboreal singularities. We make some comments relating the geometric situation to the combinatorial situation, and present some conjectures at various levels of precision.

8.1 Stratified Spaces

Before we introduce arboreal singularities, we make some comments on stratified spaces. The definition below (and much of the following discussion) is inspired by [17].

Definition 8.1.1. *A **stratified space** is a (Hausdorff, locally compact, second countable) topological space X which is can be written as a disjoint union: $X = \sqcup_{\alpha \in A} X_\alpha$, where the X_α are called **strata**. In addition we have:*

(i) Each X_α is equipped with the structure of a smooth manifold of some dimension, which is compatible with the subspace topology on X_α.section

(ii) For $\alpha \neq \beta \in A$, if X_β intersects $\overline{X}_\alpha$, then $X_\beta \subset \overline{X}_\alpha$ and $dim(X_\beta) < dim(X_\alpha)$.

A stratified space is called **n-dimensional** if any strata X_α is contained in the closure of some X_β with $\dim(X_\beta) = n$.

Definition 8.1.2. *Given a stratified space $X = \sqcup_{\alpha \in A} X_\alpha$, the **strata poset** $\mathscr{S}(X)$ of X is the poset whose elements are A, and whose relation is given by $\beta \leq \alpha$ iff $X_\beta \subset \overline{X}_\alpha$.*

We have some basic constructions on stratified spaces: Any open subset U of a stratified space X is also a stratified space, whose strata U_α are the nonempty intersections $X_\alpha \cap U$.

Additionally, for any stratified spaces X and Y, the cross product $X \times Y$ has the structure of a stratified space whose strata are cross products of strata of X and Y.

Definition 8.1.3. *Let X and Y be stratified spaces. An **open embedding** is a map $i : X \hookrightarrow Y$ which is topologically an open embedding, such that for any $x \in X_\alpha$, there exists an open subset U of X_α containing x and a strata Y_β of Y such that $\varphi(U) \subset Y_\beta$ and $\varphi\,|_U$ is a diffeomorphism.*

We close with two lemmas which will be useful in the next section:

Lemma 8.1.1. *Let $i : X \hookrightarrow Y$ be an open embedding of stratified spaces. If the strata of X are connected, then i determines an map $s_i : \mathscr{S}(X) \to \mathscr{S}(Y)$ of posets, such that whenever $x \in X_\alpha$, $i(x) \in Y_{s_i(\alpha)}$. Furthermore s_i is order-preserving and has open image.*

Proof. Let X_α and Y_β be strata of X and Y respectively. The definition of open embedding implies that $i^{-1}(Y_\beta) \cap X_\alpha$ is an open subset of X_α. Since X_α is connected, the image $i(X_\alpha)$ is contained in a unique strata of Y, this defines s_i.

To show s_i is order-preserving: we need to show $X_\alpha \subset \overline{X}_\beta$ implies $Y_{s_i(\alpha)} \subset \overline{Y}_{s_i(\beta)}$. If we take $x \in X_\alpha$, then any open set containing X intersects X_β, hence any open set containing $i(x)$ intersects $i(X_\beta)$, which establishes this fact.

To show s_i has open image: we need to show $Y_{s_i(\alpha)} \subset \overline{Y}_\beta$, implies β is in the image of s_i. This follows from i being an open embedding, hence the image of X is open and so must intersect Y_β. □

Due to our definition of the category **Pos** (definition 2.1.2), we would like to identify topological conditions which imply s_i is an open embedding of posets- the only piece the above lemma doesn't give us is the injectivity of s_i.

Definition 8.1.4. *We say a stratified space X is a **local model** if the poset $\mathscr{S}(X)$ has a unique minimal element and the strata of X are connected. X is **locally connected** if, for every strata X_α and every $x \in \overline{X}_\alpha$, there is an open set U containing x such that $X_\alpha \cap U$ is connected.*

Lemma 8.1.2. *Let $i : X \hookrightarrow Y$ be an open embedding of stratified spaces. If X is a local model and Y is locally connected, the map s_i is injective (and therefore an open embedding of posets).*

Proof. Let α_0 denote the minimal element of $\mathscr{S}(X)$, and pick $x \in X_{\alpha_0}$. Fix a strata Y_β with β in the image of s_i- it follows from the continuity of s_i that $i(x) \in \overline{Y}_\beta$, so we can find an open set U containing $i(x)$ such that $Y_\beta \cap U$ is connected.

Now suppose $\alpha \in \mathscr{S}(X)$ with $s_i(\alpha) = \beta$, so $i(X_\alpha) \subseteq Y_\beta$. From the continuity of s_i, we have $i(x) \in \overline{i(X_\alpha)}$, so $i(X_\alpha) \cap U$ is nonempty. Furthermore, the definition of open embedding implies $i(X_\alpha)$ is open in Y_β. This implies:

$$Y_\beta \cap U = \coprod_{s_i(\alpha)=\beta} i(X_\alpha) \cap U$$

Is a covering of $Y_\beta \cap U$ by disjoint nonempty open sets, so the connectedness of $U \cap Y_\beta$ implies s_i is injective. □

8.2 Arboreal Singularities and Arboreal Spaces

We now explain the geometric construction of arboreal singularities, as presented in [20]:

Definition 8.2.1. *Given a tree T, for $v \in V(T)$, we define the Euclidean space $L_T(v) = \mathbb{R}^{V(T)\backslash\{v\}}$, whose coordinates are written $x_w(v)$, for $w \in V(T)\backslash\{v\}$. The **arboreal singularity** L_T is the quotient space:*

$$L_T = \left(\coprod_{v \in V(T)} L_T(v) \right) / \sim$$

*Where $\sim$ is the equivalence relation generated by the **edge relations**: If $e \in E(T)$, with $e = \{v, w\}$, we set $(x_u(v))_{u \in V(T)\backslash\{v\}} \sim_e (x_u(w))_{u \in V(T)\backslash\{w\}}$ if:*

$$x_w(v) = x_v(w) \geq 0$$

$$x_u(v) = x_u(w) \qquad \text{for all } u \in V(T)\backslash\{v, w\}$$

It is easy to see that the quotient map is a closed embedding restricted to $L_T(v)$, hence we can consider $L_T(v)$ to be a subset of L_T. For a fixed $x \in L_T$. We have a subset S_x of $V(T)$ given by:

$$S_x = \{v \in V(T) \mid x \in L_T(v)\}$$

In [20] it is proven that S_x is a subtree. Now, for $e = \{v, w\} \in E(S_x)$, by assumption we have $x_w(v) = x_v(w) \geq 0$. Define the subset $E_x \subset E(S_x)$ to be those edges $e = \{v, w\}$ such that $x_w(v) = x_v(w) > 0$. Then we get:

$$q : S_x \twoheadrightarrow R_x$$

The quotient map obtained by contracting the edges in E_x. This gives a correspondence:

$$\mathfrak{p}_x = (R_x \overset{q}{\twoheadleftarrow} S_x \overset{i}{\hookrightarrow} T)$$

Definition 8.2.2. *Let T be a tree, and $\mathfrak{p} : R \to T$ a correspondence. Define:*

$$L_T(\mathfrak{p}) = \{x \in L_T \mid [\mathfrak{p}_x] = [\mathfrak{p}]\}$$

The following statement is proven in [20]:

Proposition 8.2.1.

(i) $L_T(\mathfrak{p})$ is an open cell of dimension $|V(T)| - |V(R)|$. In particular, for $v \in V(S)$, $L_T(\mathfrak{p}) \subset L_T(v)$ is cut out by equations of the form $x_w(v) = 0$, $x_w(v) > 0$, $x_w(v) < 0$.

(ii) For correspondences $\mathfrak{p} : R \to T$, $\mathfrak{q} : R' \to T$, $L_T(\mathfrak{q}) \subset \overline{L_T(\mathfrak{p})}$ if and only if $\mathfrak{q} \leq \mathfrak{p}$.

Corollary 8.2.1. *L_T is a stratified space whose strata poset is naturally identified with $\mathcal{P}_T$. (The smooth structure on $L_T(\mathfrak{p})$ is given by its description as an open subset of Euclidean space).*

From this definition we can establish the following lemma, which implies lemma 7.3.2:

Lemma 8.2.1.

(i) $L_T(\mathfrak{p}) \subseteq L_T(v)$ if and only if $\mathfrak{p} = (R \overset{q}{\twoheadleftarrow} S \overset{i}{\hookrightarrow} T)$ satisfies $v \in V(S)$- i.e. if and only if $[\mathfrak{p}] \in L_v$ (definition 6.2.2). Hence $L_v \setminus \{\hat{0}\}$ is the face poset of $S^{n-1} \subseteq L_T(v)$.

(ii) For $v \neq 0$, $L_T(v) \cap L_T(w)$ is the closed subset of $L_T(v)$ given by $x_u(v) \geq 0$ for $u \in V(T)$ either equal to w or between v and w. $L_T(\mathfrak{p})$ is in the interior of $L_T(v) \cap L_T(w)$ iff all edges between v and w are contracted by q. It follows that $(L_v \cap L_w) \setminus \{\hat{0}\}$ is the face poset of $S^{n-1} \cap \{x_u(v) \geq 0\}$, with the interior given by $[\mathfrak{p}]$ for which all edges between v and w are contracted.

Proof. Both are straightforward from the definition of $\mathfrak{p}_x$ above. □

Note that L_T has a unique zero-dimensional strata- call it $\{\hat{0}\}$.

Definition 8.2.3. *A **arboreal space** is an $n-$dimensional stratified space L such that, for any $x \in L$ there is a tree T and an open embedding $i : L_T \times \mathbb{R}^k \hookrightarrow L$ such that $i(\hat{0}, \mathbf{0}) = x$.*

It is worth noting that L_T itself is an arboreal space according to the above definition. Actually, we can say more: Note that L_T is a local model and locally connected, hence so is $L_T \times \mathbb{R}^k$ for any k. The following is shown in [20]:

Lemma 8.2.2. *Let $\mathfrak{p} : R \to T$ be a correspondence. Then there exists an open embedding $i : L_R \times \mathbb{R}^k \hookrightarrow L_T$ such that the image of $\{\hat{0}\} \times \mathbb{R}^k$ is the strata $L_T(\mathfrak{p})$, and furthermore the open embedding $s_i : \mathscr{S}(L_R \times \mathbb{R}^k) \to \mathscr{S}(L_T)$ is naturally identified with $i_{\mathfrak{p}} : \mathcal{P}_R \hookrightarrow \mathcal{P}_T$.*

Definition 8.2.4. *If L is an arboreal space, an **arboreal chart** is a triple (U, T, φ), where U is an open subset of L, T is a tree, and $\varphi : L_T \times \mathbb{R}^k \to L$ is an open embedding whose image is U.*

For arboreal charts (U, T, φ) and (V, R, φ'), with $V \subseteq U$, $\varphi^{-1} \circ \varphi'$ is an open embedding $L_R \times \mathbb{R}^m \to L_T \times \mathbb{R}^k$, which determines an open embedding of posets $\mathcal{P}_R \hookrightarrow \mathcal{P}_T$.

Our discussion in this section will focus on the relationship between the combinatorial understanding of C-Arb singularities and a geometric understanding of arboreal singularities. There are not many concrete results here, but there will be a few conjectures (stated at varying levels of precision) as well as a handful of specific constructions and pictures in low-dimensional cases.

8.3 Embeddings

Let L be an $n-$dimensional arboreal space.

For a symplectic vector space (V, ω), if W is a subspace, write $W^\perp = \{v \in V \mid \omega(v, w) = 0 \forall w \in W\}$. Recall that W is **isotropic** if $W \subset W^\perp$, **coisotropic** if $W^\perp \subset W$, and **Lagrangian** if it is both isotropic and coisotropic. A submanifold X of a symplectic manifold (M, ω) is isotropic/coisotropic/Lagrangian iff $T_x X \subseteq T_x M$ is the appropriate type of subspace, for all $x \in X$.

Definition 8.3.1. *A **Lagrangian embedding** of L consists of a symplectic manifold M of dimension $2n$ and a topological embedding $i : L \hookrightarrow M$ such that:*

(i) For each strata $L_\alpha \subset L$, $i : L_\alpha \hookrightarrow M$ is a smooth embedding whose image is an isotropic submanifold of M. (By an abuse of notation we will sometimes identify L_α and $i(L_\alpha)$ when it will not cause confusion).

(ii) ***Whitney's Condition B:*** *([17]) Suppose $L_\alpha \subset \overline{L}_\beta$ are two strata, and $\{x_i\}_{i\in\mathbb{N}}$ and $\{y_i\}_{i\in\mathbb{N}}$ are subsequences of L_α, L_β respectively, converging to $x \in L_\alpha$. Then if the tangent planes $T_{y_i}L_\beta$ converge to $V \subset T_xM$ and the secant lines x_iy_i converge to a line $\ell \subset T_xM$[1], we have $\ell \subset V$.*

(iii) Suppose L_α is a strata, and $\{x_i\}_{i\in\mathbb{N}}$ and $\{y_i\}_{i\in\mathbb{N}}$ are subsequences of L converging to $x \in L_\alpha$. If the secant line x_iy_i converges to a line $\ell \subset T_xM$, then $\ell \subset (T_xL_\alpha)^\perp$.

Definition 8.3.2. *A **pre-cyclic structure** on L is the data of, for every arboreal chart (U, T, φ), a pre-cyclic structure $\mathcal{O}_{U,\varphi}$ on $\mathcal{P}_T$, such that if (V, R, φ') is an arboreal chart with $V \subseteq U$ and $i : \mathcal{P}_R \hookrightarrow \mathcal{P}_T$ denotes the associated open embeding, $i^*\mathcal{O}_{U,\varphi} = \mathcal{O}_{V,\varphi'}$. A **cyclic structure** is a pre-cyclic structure in which $\mathcal{O}_{U,\varphi}$ is a cyclic structure for every arboreal chart.*

Note that by our definition of pre-cyclic structures on posets, a pre-cyclic structure on L is determined by its behavior on arboreal charts of the form (U, A_2, φ). Additionally, we see pre-cyclic structures on L_T are in bijection with pre-cyclic structures on $\mathcal{P}_T$.

The following proposition connects symplectic geometry and combinatorics:

Proposition 8.3.1. *Lagrangian embeddings of L induce pre-cyclic structures on L.*

Proof. It is enough to show that a Lagrangian embedding $i : L_{A_2} \times \mathbb{R}^{n-1} \to (M, \omega)$, where M is a $2n-$dimensional symplectic manifold, determines a canonical cyclic order on $\mathcal{P}_{A_2}$.

Denote the strata of L_{A_2} by $\{\hat{0}\}, a, b, c$. Write $X = i(\hat{0} \times \mathbb{R}^{n-1})$, $X_\alpha = i(\alpha \times \mathbb{R}^{n-1})$ for $\alpha = 1, 2, 3$, and $L = i(L_{A_3} \times \mathbb{R}^{n-1})$.

The *symplectic normal bundle* E, given by $E_x = (T_xX)^\perp/(T_xX)$, is a 2-dimensional symplectic vector bundle over X. The isotropic embedding theorem (ref) states that we can find an open neighborhood U of X which is symplectomorphic to an open neighborhood of the zero section in $T^*X \oplus E$. Let $\pi_X : U \to X$ denote the associated projection.

Fixing a Riemannian metric on the vector bundle $T^*X \oplus E$ determines a smooth distance-squared function $\rho : U \to \mathbb{R}$ such that $X = \rho^{-1}(0)$. Whitney's condition B implies (reference: Mather) that we can pick U small enough such that $(\pi_X, \rho) : U \to X \times \mathbb{R}$ is a submersion when restricted to any of the strata X_α, $\alpha = a, b, c$.

[1] With respect to a coordinate chart around x

Fix $x \in X$, then $\pi^{-1}(x)$ is a submanifold diffeomorphic to an open neighborhod of the origin in $T^*_X \oplus E_x$, and $X_\alpha \cap \pi^{-1}(X)$ is a one-dimensional submanifold. Letting $p : T^*_x X \oplus E_x \to E_x$, we claim that p is injective on $L \cap \pi^{-1}(x)$ in some neighborhood of x. If not, we could find $x_i, y_i \in L \cap \pi^{-1}(X)$ be sequences in $L \cap \pi^{-1}(x)$ converging to x such that $x_i - y_i \in T^*_x X$. By the compactness of projective space, we can take some subsequence such that $\mathrm{Span}(x_i - y_i)$ converges to $\ell \subset T^*_x X$, contradicting requirement (iii).

So, for small enough ε, $q(L \cap \pi^{-1}(x) \cap \rho^{-1}(\varepsilon)) \subset E_x$ is three distinct points on an oriented circle, with each point contained in one X_α giving a cyclic order on the strata $\{X_a, X_b, X_c\}$. Furthermore, these vary continuously as x, ε change. Lastly, the isotropic embedding theorem implies that different choices for the symplectomorphism $T^*X \oplus E \cong U$ are related by a continuous family, hence this gives a canonical well-defined cyclic order on $\mathcal{P}_{A_2}$.

□

As a consequence, we have a map:

$$\{\text{Lagrangian embeddings of } L_T\} \longrightarrow \{\text{Pre-cyclic structures on } \mathcal{P}_T\}$$

Our **main conjecture** is a bijection of the following form:

$$\left\{ \begin{array}{l} \text{'Regular' Lagrangian} \\ \text{embeddings of } L_T \end{array} \right\} \Big/ \text{equivalence} \longleftrightarrow \{\text{Cyclic Structures on } \mathcal{P}_T\} \tag{8.1}$$

We discuss various aspects of the main conjecture in the subsequenct sections, including the notions of 'equivalence' and 'regularity' which have not been mentioned yet.

Equivalence

The first part of the main conjecture we discuss is the idea of equivalence of Lagrangian embeddings. We begin with the following 'too strict' definition:

Definition 8.3.3. *Let L be an arboreal space. Two Lagrangian embeddings $i_1 : L \hookrightarrow (M_1, \omega_1)$ and $i_2 : L \to (M_2, \omega_2)$ are **germ equivalent** if there exist open neighborhoods U_α of $i_\alpha(L)$ in M_α, for $\alpha = 1, 2$, and a symplectomorphism $\varphi : U_1 \to U_2$ such that $\varphi \circ i_1 = i_2$.*

The main issue with germ equivalence as a notion of equivalence is that symplectomorphisms act linearly on tangent spaces and are consequently far too rigid. For instance, germ equivalence fails to identify the three embeddings of L_{A_2} into $\mathbb{R}^2$ pictured in Figure 8.1.

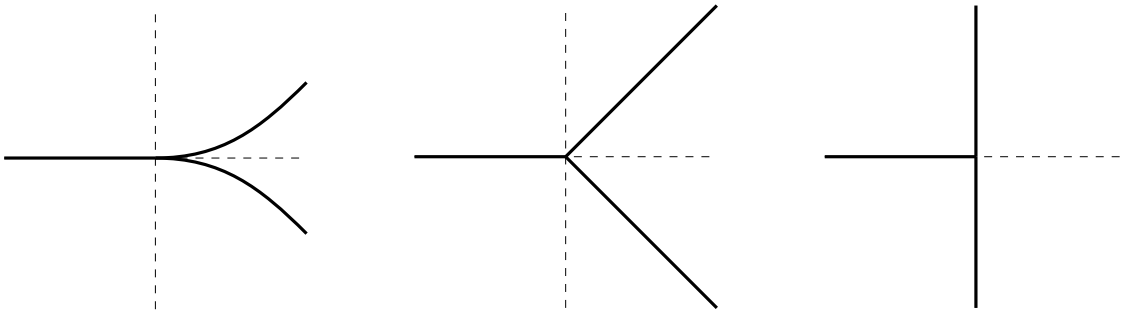


Figure 8.1: Embeddings Which Are Not Germ Equivalent

To state our proposed definition, we first extend the notion of germ equivalence:

Definition 8.3.4. *If $i_\alpha : L \hookrightarrow (M_\alpha, \omega_\alpha)$, $\alpha = 1, 2$ are Lagrangain embeddings, and $\pi_\alpha : M_\alpha \to P$ are submersions onto some smooth manifold P, i_1, i_2 are **germ equivalent relative to** P if there exists U_1, U_2, φ as in definition 8.3.3 with $\pi_2 \circ \varphi = \pi_1$.*

Definition 8.3.5. *An **equivalence** between two Lagrangian embeddings $i_1 : L \to (M_1, \omega_1)$ and $i_2 : L \to (M_2, \omega_2)$ consists of a Lagrangian embedding $H : L \times \mathbb{R} \to (\tilde{M}, \tilde{\omega})$ and a submersion $t : \tilde{M} \to \mathbb{R}$ satisfying:*

(i) Writing $\pi : L \times \mathbb{R} \to \mathbb{R} \to \mathbb{R}$ for the projection onto $\mathbb{R}$, $\pi = t \circ H$.

(ii) For $\alpha = 1, 2$ and $\varepsilon > 0$, write $B_\alpha(\varepsilon) = (\alpha - \varepsilon, \alpha + \varepsilon) \subset \mathbb{R}$. Write $j_\alpha : B_\alpha(\varepsilon) \hookrightarrow T^* B_\alpha(\varepsilon)$ for the Lagrangian embedding as the zero section. Then there exists ε such that the embeddings

$$i_\alpha \times j_\alpha : L \times B_\alpha(\varepsilon) \hookrightarrow M_\alpha \times T^* B_\alpha(\varepsilon)$$

and:

$$H : L \times B_\alpha(\varepsilon) \hookrightarrow t^{-1}(B_\alpha(\varepsilon))$$

Are germ equivalent relative to $B_\alpha(\varepsilon)$. (Both right hand sides have a natural projection to $B_\alpha(\varepsilon)$).

Then we have:

Lemma 8.3.1. *If there exists an equivalence between to Lagrangian embeddings, they induce the same pre-cyclic structure on L.*

Proof. Let (U, A_2, φ) be an arboreal chart on L. Then $(U \times \mathbb{R}, A_2, \tilde{\varphi})$ is an arboreal chart on $L \times \mathbb{R}$. The Lagrangian embedding H then induces a cyclic order on $\mathcal{P}_{A_2}$. For $\varepsilon > 0$, the subchart $(U \times B_\alpha(\varepsilon), A_2, \tilde{\varphi})$ inherits the same cyclic order from H, which by germ equivalence is the same cyclic order as induced by $i_\alpha \times j_0$. This is the same as the cyclic order induced by i_α. □

Thus a rightward-pointing map in the main conjecture (8.1) is well-defined. Showing the injectivity of this map is a central goal of this program going forward- unfortunately, we were unable to make much progress on this.

Realizing Cyclic Structures using Cooriented Hypersurfaces

We now turn to the question of the *surjectivity* of the rightward-pointing map in the main conjecture (8.1).

Definition 8.3.6. *A **realization** of a cyclic structure $\mathcal{O}$ on $\mathcal{P}_T$ is a Lagrangian embedding $i : L_T \hookrightarrow (M, \omega)$ inducing $\mathcal{O}$.*

Let us not worry about the adjective 'regular' for now. Then the surjectivity of the map is the same as the assertion that every cyclic structure has a realization. We will begin this discussion by analyzing certain types of Lagrangian embeddings of $L_{A_2} \times \mathbb{R}^k$.

L_{A_2} is the union of two 1-dimensional Euclidean spaces: $L_{A_2}(1)$ with coordinate $x_2(1)$, and $L_{A_2}(2)$ with coordinate $x_1(2)$. It has four strata: We denote the correspondences by $\{\hat{0}\}$, $\mathfrak{p}_{(1)}$, $\mathfrak{p}_{(2)}$ and $\mathfrak{p}_{(12)}$ as in chapter 2. (The notation has been slightly modified to emphasize that $\mathfrak{p}_{(1)}$, etc. are correspondences). Then:

$$L_{A_2}(\{\hat{0}\}) = \{\hat{0}\} = \{\mathbf{0} \in L_{A_2}(1)\} = \{\mathbf{0} \in L_{A_2}\}$$

$$L_{A_2}(\mathfrak{p}_{(1)}) = \{x \in L_{A_2}(1) \mid x_2(1) < 0\}$$

$$L_{A_2}(\mathfrak{p}_{(2)}) = \{x \in L_{A_2}(2) \mid x_1(2) < 0\}$$

$$L_{A_2}(\mathfrak{p}_{(12)}) = \{x \in L_{A_2}(1) \mid x_2(1) > 0\} = \{x \in L_{A_2}(2) \mid x_1(2) > 0\}$$

Let's define a pair of Lagrangian embeddings $i_\pm$ of L_{A_2} as a *conical Lagrangian* in $T^*\mathbb{R}$. Write coordinates on $T^*\mathbb{R}$ as (x, η_x). Then the embeddings are given by:

$$i_\pm(x) = \begin{cases} (x_2(1), 0) & x \in L_{A_2}(1) \\ (0, \pm x_1(2)) & x \in L_{A_2}((2)) \end{cases}$$

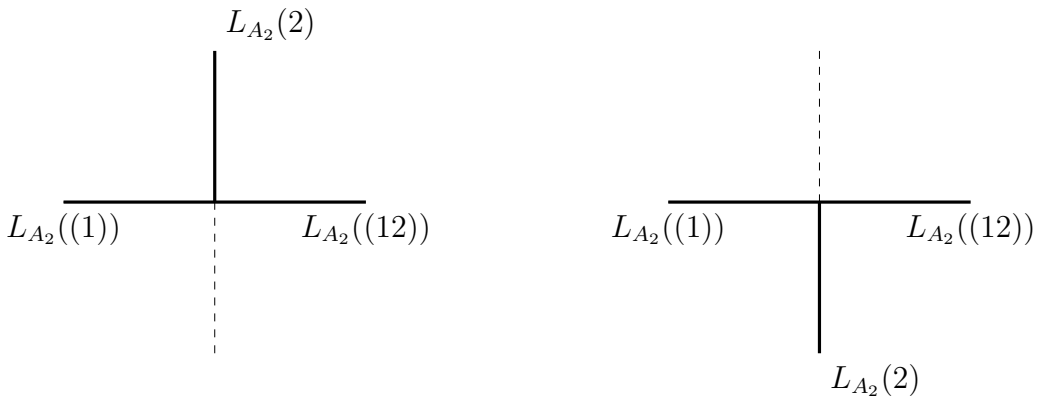


Figure 8.2: The embeddings i_+ (left) and i_- (right) of L_{A_2}.

They are drawn in Figure 8.2.

From the figure and our conventions in Chapter 2, we see that i_+ induces the cyclic structure $\mathcal{O}_\mu$ where μ points from vertex 2 to vertex 1.

More generally, let H be a smooth hypersurface inside an $n-$dimensional manifold M. Let $T^*_H M \subset T^*M$ denote the conormal bundle to H. We denote by S^*M the spherical projectivization of T^*M, whose points consist of pairs $(x, [v])$, where $x \in M$, $v \neq 0 \in T^*_x M$, and $[v] = [w]$ if $w = \lambda v$ for some $\lambda > 0$. $S^*_H M \subset S^*M$ is the spherical projectivization of $T^*_H M$.

Definition 8.3.7. *A **co-orientation** of a hypersurface H is a section σ of $S^*_H M$.*

If H is a co-oriented hypersurface, we have a conical[2] Lagrangian submanifold $\mathcal{L}_H \subset T^*M = \{(x, v) \mid x \in H, [v] = \sigma(x)\}$. Furthermore, the singular conical Lagrangian $\underline{\mathcal{L}_H} = \mathcal{L}_H \sqcup M \times \{\mathbf{0}\}$ is an embedded arboreal space.

Locally near H we can find a coordinate neighborhood $U \cong \mathbb{R}^n$ with coordinates $\{x_1, \ldots, x_n\}$ such that $H \cap U = \{x_1 = 0\}$ and $[dx_1] = \sigma$. Then we have an arboreal chart:

$$(T^*U \cap \underline{\mathcal{L}_H}, L_{A_2} \times \mathbb{R}^{n-1}, \varphi)$$

Where, writing $T^*(U) = T^*\mathbb{R} \times T^*\mathbb{R}^{n-1}$ $\varphi = i_+ \times j$, j being the inclusion of the zero section.

As a result, if we let ξ_1 denote the dual coordinate to x_1, locally near U we see the cyclic structure is induced by the total order:

$$\{x_1 < 0\} < \{x_1 > 0\} < \{x_1 = 0, \xi_1 > 0\}$$

[2]A subset of T^*M is conical if it is invariant under scaling by a positive scalar.

Where each of these denote subsets of $T^*U \cap \underline{\mathcal{L}_H}$- the first two subsets of the zero section and the last a subset of $\mathcal{L}_H$.

Now, let's consider two hypersurfaces: Let $H_1 = \{y = 0\}$ and $H_2 = \{x = 0\}$ be curves in $\mathbb{R}^2$, cooriented toward $y > 0$ and $x > 0$ respectively. Then $\underline{\mathcal{L}_{H_1}} \cup \underline{\mathcal{L}_{H_2}}$ is a conical Lagrangian in $T^*\mathbb{R}^2$ which is isomorphic to L_{A_3} (Figure 8.3).

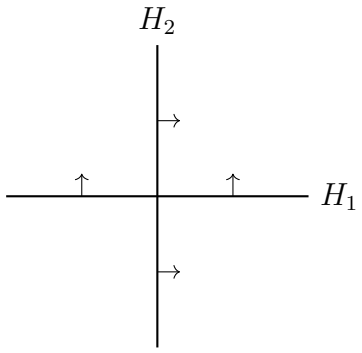
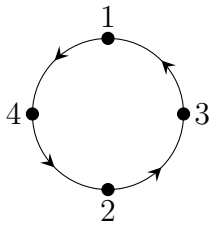


Figure 8.3: A Lagrangian Embedding of L_{A_3}

Our discussion allows us to determine the induced pre-cyclic structure. One can check that it is, in fact, a cyclic structure. Explicitly, labelling the one-cells $1, 2, 3, 4$ going in a counterclockwise order starting from $\{x > 0, y = 0\}$, we can determine the cyclic order. In the neighborhood of 1, the top-cells containing it are $\{1, 2\}$ (the top-right quadrant), $\{1, 3\}$ (the conormal direction $\mathcal{L}_{H_1}$), and $\{1, 4\}$ (the bottom right quadrant). Our discussion shows that the induced cyclic order is induced by the total order:

$$\{1, 4\} < \{1, 2\} < \{1, 3\}$$

Using similar considerations, we can deduce that the cyclic structure on L_{A_3} corresponds to the cyclic order on the one-cells:



Another important figure for us is constructed as follows: Let $f : \mathbb{R} \to \mathbb{R}$ be a smooth function such that $f(x) = 0$ for $x \leq 0$ and $f(x) > 0$, $f'(x) > 0$ for $x > 0$. Let $H_1 = \{x, f(x)\}$

and $H_2 = \{y = 0\}$ be hypersurfaces in $\mathbb{R}^2$ both cooriented toward $y > 0$. Then $\mathcal{L} = \underline{\mathcal{L}_{H_1}} \cup \underline{\mathcal{L}_{H_2}}$ is a conical Lagrangian in $T^*\mathbb{R}^2$ which is isomorphic to L_{A_3}. This is pictured in Figure 8.4, with the one-cells labelled.

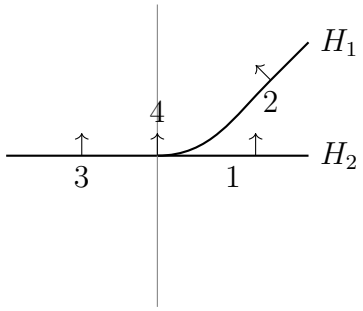


Figure 8.4: Another Lagrangian Embedding of L_{A_3}

The one-cell labelled 4 is, writing (x, y, ξ_x, ξ_y) for coordinates on $T^*\mathbb{R}^2$, the subset $\{(0, 0, 0, t), t > 0\}$. Looking at a cross-section $\{\xi_y = t_0\}$ for some fixed $t_0 > 0$, we get the 3-d picture in Figure 8.5.

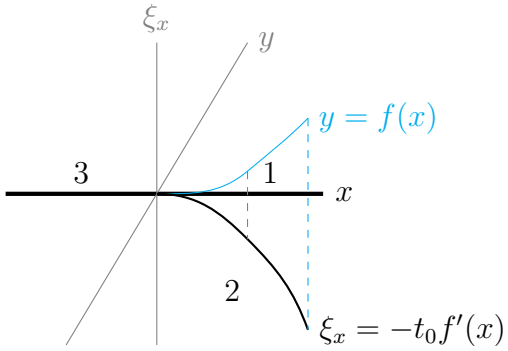


Figure 8.5: A Cross-Section of the Embedding at $\xi_y = t_0$

We see by projecting onto the $(x, \xi_x)-$plane that the cyclic order on the top-cells is induced by the total order:

$$\{4, 1\} < \{4, 3\} < \{4, 2\}$$

These constructions are generalized in [20][3]. Explicitly, let T be a tree, and $v \in V(T)$ a vertex. Nadler constructs an embedding $i : L_T \hookrightarrow T^*\mathbb{R}^{V(T)\setminus\{v\}}$ with the following properties:

[3]The Constructions in [20] refer to *Legendrian embeddings*- but the theory is essentially the same

- For $w \neq v$, there is a smooth function $h_w : \mathbb{R}^{V(T)\setminus\{v\}} \to \mathbb{R}$ with nonvanishing differential.
- h_w yields a coordinted hypersurface $\mathcal{H}_w = \{h_w = 0\}$ cooriented toward $\{h_w > 0\}$.
- The image of L_T is:
$$\bigcup_{w\neq v} \underline{\mathcal{L}_{\mathcal{H}_w}}$$
- More specifically, the image of $L_T(v)$ is the zero section, and the image of $L_T(w)$ is $\mathcal{L}_{\mathcal{H}_w}$ and the subset $\{h_w > 0\}$ of the zero section.

See that paper for the details of the construction. The vertex $v \in V(T)$ is called a **root** and the pair (T, v) a rooted tree. There is an associated orientation μ_v of the rooted tree in which all edges are pointing toward the root. We can prove:

Proposition 8.3.2. *The cyclic structure given by the embedding is $\mathcal{O}_{\mu_v}$.*

Proof. To simplify notation, we identify L_T with its image in $T^*\mathbb{R}^{V(T)\setminus\{v\}}$. Let $\mathfrak{p} = (A_2 \overset{q}{\twoheadleftarrow} S \overset{i}{\hookrightarrow} T)$ be a correspondence. First assume $v \in V(S)$, WLOG $v \in q^{-1}(1)$. Pick any vertex $w \in q^{-1}(2)$. Then $[\mathfrak{p}]$ lies in the hyperplane $\mathcal{H}_w$, and the three top-cells containing $L_T(\mathfrak{p})$ are:
$$L_T(\mathfrak{p} \circ \mathfrak{p}_{(1)}) \subset L_T(v) \setminus L_T(w)$$
$$L_T(\mathfrak{p} \circ \mathfrak{p}_{(2)}) \subset L_T(w)\setminus L_T(v)$$
$$L_T(\mathfrak{p} \circ \mathfrak{p}_{(12)}) \subset L_T(w) \cap L_T(v)$$
The situation is pictured in Figure 8.6.

Therefore the cyclic order is given by:
$$[\mathfrak{p} \circ (1)] < [\mathfrak{p} \circ (12)] < [\mathfrak{p} \circ (2)]$$
So the arrow between the trees is pointing toward v- which is consistent with μ.

Now suppose $v \notin V(S)$. Let w_1 be the closest vertex in S to v, WLOG $w_1 \in q^{-1}(1)$. Pick $w_2 \in q^{-1}(2)$. Then $L_T(w_2) \cap L_T(v) \subset L_T(w_1) \cap L_T(v)$, and we are in the situation pictured in Figure 8.7.

Then $L_T(\mathfrak{p})$ lies in the fiber of the cotangent bundle over the origin in the image (though note that the origin represents a codimension-two subspace), $L_T(\mathfrak{p} \circ (1)) \subset L_T(w_1)\setminus L_T(w_2)$ lies in the fiber of the cotangent bundle over $\{x > 0\}$, and $L_T(\mathfrak{p} \circ (12)) \subset L_T(w_1) \cap L_T(w_2)$

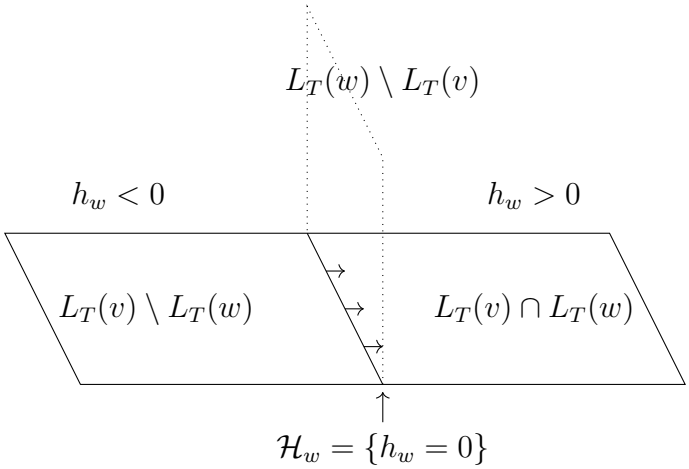


Figure 8.6: The planes $L_T(v)$ and $L_T(w)$

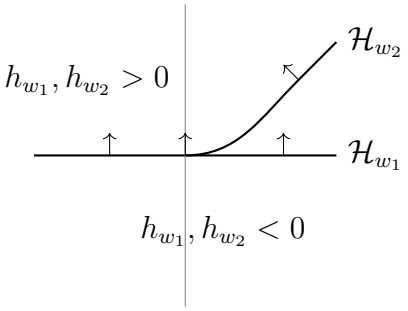


Figure 8.7: The planes $\mathcal{H}_{w_1}$,$\mathcal{H}_{w_2}$ in $L_v(T)$

lies in the fiber of the cotangent bundle over $\{x < 0\}$. From our previous computation this is enough to determine the cyclic order:

$$[\mathfrak{p} \circ (2)] < [\mathfrak{p} \circ (1)] < [\mathfrak{p} \circ (12)]$$

This is again consistent with $\mathcal{O}_\mu$. □

Conjecture 8.3.1. *Any cyclic structure has a realization.*

It is enough to show that any cyclic structure of the form $\mathcal{O}_\mu$ has a realization. We expect that there is a similar strategy involving cooriented hypersurfaces, perhaps of the following form: Fix a source $v \in V(T)$. Then:

- For $w \neq v$, there is a smooth function $h_w : \mathbb{R}^{V(T)\setminus\{v\}} \to \mathbb{R}$, and a cooriented hypersurface $\mathcal{H}_w$ as before.

- The image of $L_T(w)$ is the zero section, and the union of $\mathcal{L}_{\mathcal{H}_w}$ and the subset $\{h_w > 0\}$ of the zero section is the image of the strata contained in L_A, the brane associated to the subtree connecting v to w (see proposition 6.2.1).

Regularity and Bad Embeddings

In this section we give two 'bad' examples exhibiting behavior we would like to prohibit. The first is already prohibited by requirement (iii) of a Lagrangian embedding, and this examples justifies that condition. The second is not prohibited by the given definition, which leads us to believe that the definition given is lacking- this explains our use of the mysterious word 'regular' in the main conjecture (8.1).

First: Let $H_1 = \{x, f(x)\}$ and $H_2 = \{y = 0\}$ be curves in $\mathbb{R}^2$ as constructed in the previous section, and define $\mathcal{L} = \{(x, v) \mid x \in H_1, v \in T^*_{H_1,x}\mathbb{R}^2\} \cup \{(x, v) \mid x \in H_2, v \in T^*_{H_2,x}\mathbb{R}^2\}$- so we incude *both* directions in the conormal bundle and *exclude* the zero section (Figure 8.8).

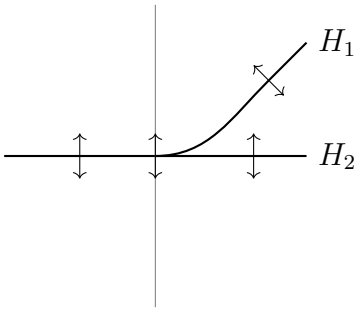


Figure 8.8: A Disallowed Embedding of $L_{A_2} \times \mathbb{R}$.

This gives a seemingly well-behaved embedding of $L_{A_2} \times \mathbb{R}$- but it violates requirement (iii): The secant line connecting $(x, 0, 0, 0)$ and $(x, f(x), 0, 0)$ approaches the line $(0, t, 0, 0)$ (with coordinates (x, y, ξ_x, ξ_y)), but the tangent space to the strata at the origin is $(0, 0, 0, s)$, and these are not orthogonal. Indeed, this embedding is disallowed for a reason: The cyclic order on the three Lagrangian half-planes surrounding the central strata flips as we cross the zero section!

For another example which is *not* disallowed by our definition, Figure 8.9 shows a Lagrangian embedding of L_{A_3} using cooriented half-hypersurfaces.

One can check using the methods of the previous section that the pre-cyclic structure induced by the above embedding is *not* a cyclic structure. This suggests that the notion of 'Lagrangian Embedding' is missing some extra condition- this is the meaning of the myste-

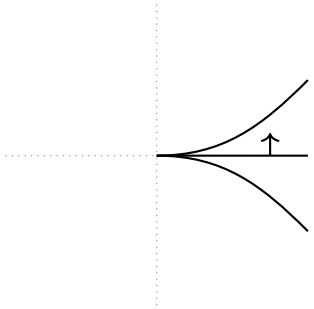


Figure 8.9: An Embedding of L_{A_3} Realizing a Non-Cyclic Structure

rious term 'regular' in our main conjecture.

The 'regularity' condition would be something which prohibits embeddings of the above form. At this point, I have no concrete definition for regularity. One (somewhat unsatisfying) idea is:

> **Proposed Definition:** A Lagrangian embedding is regular if the induced pre-cyclic structure is a cyclic structure.

However, the hope is that there's a natural geometric condition which yields the above as a theorem. By the coherence theorem, it is enough to ensure that the cyclic structures induced by the embeddings are A_3-coherent and S_4-coherent- meaning the desired geometric condition is a requirement in codimensions 2 and 3.

As a final comment in this section: In Chapter 3 we construct pre-cyclic structures on $\mathcal{P}_{S_4}$ which are A_3-coherent but not S_4-coherent. It is an interesting question whether there are 'bad' embeddings of $\mathcal{P}_{S_4}$ realizing these pre-cyclic structures, analogous to Figure 8.9.

A Sheaf of Categories

We conclude our discussion of the 'main conjecture' by mentioning the representation theory of Lagrangian embeddings. The canonical references for the sheaf of categories discussed here is [13].

Let M be a smooth manifold. We let $\mathrm{Sh}(M)$ denote the dg category of cohomologically constructible complexes of sheaves on M over a fixed field k. To an object $\mathcal{F}$ of $\mathrm{Sh}(M)$, one can define the singular support $SS(\mathcal{F})$, which is a conical Lagrangian subspace of $\mathrm{Sh}(M)$.

To an open set U of T^*M, one can form the dg category $\mathrm{Sh}(M,U)$, which is the dg quotient of $\mathrm{Sh}(M)$ by the subcategory of sheaves $\mathcal{F}$ with $SS(\mathcal{F}) \cap U = \emptyset$. This data assembles

to define a presheaf of dg-categories on T^*M. We let $\mathrm{MSh}(M)$ denote the sheafification of this presheaf.

If $\Lambda \subset T^*M$ is a conical Lagrangian subspace, one can construct a category $\mathrm{Sh}_\Lambda(M)$ of constructible sheaves with singular support contained in Λ. Similarly, $\mathrm{Sh}_\Lambda(M, U)$ is the dg-quotient of $\mathrm{Sh}_\Lambda(M)$ by the subcategory of sheaves whose singular support is disjoint from U. We let $\mathrm{MSh}_\Lambda(M)$ denote the associated sheaf of dg-categories on T^*M.

In [20], Nadler essentially shows that, using the embedding $L_T \hookrightarrow \mathbb{R}^n$ coming from a choice of root, $\mathrm{Sh}_{L_T}(\mathbb{R}^n) \cong \mathrm{Rep}(T, \mu)$. where μ is the orientation induced by the root. Furthermore, for a correspondence $\mathfrak{p} : R \to T$ one has an open neighborhood U of $L_T(\mathfrak{p})$ in T^*M, and the restriction:

$$\mathrm{Sh}_{L_T}(\mathbb{R}^n) \to \mathrm{Sh}_{L_T(\mathfrak{p})}(U)$$

Coincides with the restriction functor $c_\mathfrak{p} : \mathrm{Rep}(T, \mu) \to \mathrm{Rep}(R, \lambda)$. It follows that the sheaf $\mathrm{MSh}_{L_T}(\mathbb{R}^n)$ is quasi-equivalent to the sheaf $\mathcal{Q}_\mu$- more precisely, one can construct a 'geometric realization' functor:

$$\chi_T : \mathrm{Sh}(\mathcal{P}_T; \mathrm{dgCat}) \to \mathrm{Sh}(L_T; \mathrm{dgCat})$$

And we have $\chi_T(\mathcal{Q}_\mu)$ is quasi-equivalent to the pull-back of $\mathrm{MSh}_{L_T}(\mathbb{R}^n)$ to L_T.

Definition 8.3.8. *Two embeddings $i_\alpha : L \hookrightarrow T^*M_\alpha$, $\alpha = 1, 2$, of L as a conical Lagrangian are **weakly equivalent** if the sheaves $MSh_{i_\alpha(L)}(M_\alpha)$ are quasi-equivalent sheaves of dg-categories on L.*

Conjecture 8.3.2. *If $i : L \hookrightarrow T^*M$ induces the cyclic structure $\mathcal{O}_\mu$ on $\mathcal{P}_T$, the sheaf $MSh_L(M)$ produces (under the functor ref) a sheaf quasi-equivalent to $\mathcal{Q}_\mu$.*

As a corollary of the conjecture and the Cyclic Structure Theorem (theorem 5.6.1), one would be able to conclude that two embeddings are weakly equivalent if and only if they induce the same cyclic structure. This can be seen as a weak form of the main conjecture.

8.4 The Geometry of Branes

We make a brief comment on the geometry of branes. In Chapter 6, we introduce the notion of a *pre-brane* (definition 6.2.1), and specify a subcollection of such objects as *branes* (definition 6.2.3). While the definition of brane is partially justified algebraically by the connection to representation theory- namely, the bijection between graded branes and rank-one objects

in $\mathrm{Rep}(T, \mu)$ (propositions 6.3.2 and 6.3.3)- we also hope that the definition will turn out to have geometric content. More specifically, we would hope that, given an embedding, the branes can be 'smoothed' into Lagrangian submanifolds, while non-branes cannot.

We present here an illustration of the non-brane in $\mathcal{P}_{A_3}$, relative to the embeddings in Figures 8.3 and 8.4. This is pictured in Figure 8.10.

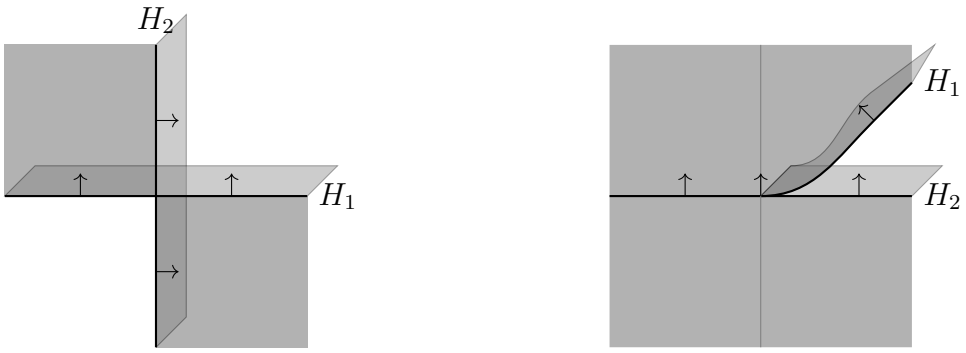


Figure 8.10: The non-brane in $(\mathcal{P}_{A_3}, \mathcal{O})$ relative to two Lagrangian embeddings

We remark that in each case there does not exist a constructible sheaf on $\mathbb{R}^2$ whose singular support is the non-brane.

8.5 The Geometry of Moves

We move on from a discussion of embeddings to a discussion of arboreal moves:

Definition 8.5.1. *Let L_T be an arboreal singularity. An **abstract elementary arboreal move** with total space L_T is a function $\pi : L_T \to \mathbb{R}$ satisfying the following:*

(i) $\pi(\hat{0}) = 0$.

(ii) Away from 0, π is a locally trivial fibration. Explicitly: If $t \neq 0$ there exists $\varepsilon > 0$ such that $\pi^{-1}(t - \varepsilon, t + \varepsilon)$ is isomorphic, as a stratified space, to $L \times (t - \varepsilon, t + \varepsilon)$, where L some arboreal space and the isomorphism preserves projection onto $\mathbb{R}$.

(iii) Away from $\hat{0}$, π is a locally trivial fibration. Explicitly: For every open neighborhood U of $\hat{0}$, there exists an open set V and $\varepsilon > 0$ such that $\pi^{-1}(-\varepsilon, \varepsilon) \subseteq U \cup V$, and $\pi\,|_V$ is a locally trivial fibration at 0.

This notion will be less important for us than the following:

Definition 8.5.2. *Let L_T be an arboreal singularity with cyclic structure $\mathcal{O}$. An **embedded elementary arboreal move** with total space $(L_T, \mathcal{O})$ is a Lagrangian embedding $i : L_T \hookrightarrow T^*(M \times \mathbb{R})$, where M is a smooth manifold, inducing $\mathcal{O}$, such that:*

(i) L_T is a conical Lagrangian subspace of $T^*(M \times \mathbb{R})$.

(ii) If $\pi : T^*(M \times \mathbb{R}) \to \mathbb{R}$ is the projection, $\pi : L_T \to \mathbb{R}$ is an abstract elementary arboreal move.

*(iii) **Preservation of Categories:** The restriction $Sh_{L_T}(M) \to Sh_{L_T}(\pi^{-1}(-\infty, 0))$ is a quasi-equivalence, as is $Sh_{L_T}(M) \to Sh_{L_T}(\pi^{-1}(0, \infty))$.*

An ideal definition would share the quality of the 'abstract' notion that doesn't rely on an explicit embedding, while somehow encoding the preservation of categories requirement of the embedded moves. Such a definition escapes us at the time of writing.

Definition 8.5.3. *If $(\mathcal{P}_T, \mathcal{O})$ is a cyclic C-Arb singularity, and $K_T = K_1 \sqcup K_2$ is a partition of K_T into good sets with respect to $\mathcal{O}$, a **realization** of the combinatorial move $U_{E_1} \to U_{E_2}$ is an embedded elementary arboreal move with total space $(L_T, \mathcal{O})$ such that $L_T(\mathfrak{p})$ intersects $\pi^{-1}(-\infty, 0)$ iff $[\mathfrak{p}] \in U_{E_1}$, and likewise with $\pi^{-1}(0, \infty)$ and U_{E_2}.*

We have very little to say when it comes to theorems or precise conjectures about arboreal moves. We have what might be called "dreams":

> **Dream 1:** The combinatorial situation faithfully represents the geometric situation, in that the combinatorial classification of moves gives a geometric classification of moves up to equivalence.

> **Dream 2:** Any two arboreal Lagrangian skeleta of a symplectic manifold can be related by a sequence of arboreal moves.

We will not make any claims regarding dream 2. In the following sections, we will explore some embedded elementary arboreal moves where the dimension of M is one or two, meaning L_T has dimension two or three. These examples will shed light on dream 1 (and add some complications!)

Realization of Combinatorial Moves

Just as in a previous section, where we constructed Lagrangian embeddings using cooriented hypersurfaces, here we will construct Lagrangian embeddings into $T^*(M \times \mathbb{R})$ using *moving* cooriented hypersurfaces. We will call these constructions *movies*.

As a first example, consider the H-to-I move. We can represent it as two cooriented points moving past each other in one dimension, pictured in 8.11.

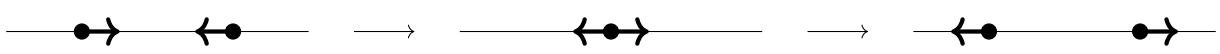


Figure 8.11: A Movie Realizing the 'H-to-I' Move

Let us look one dimension higher. Movies of cooriented curves in $\mathbb{R}^2$ can yield elementary arboreal moves whose total space is three-dimensional, so $T = A_4$ or S_4. By proposition 7.5.1, we can see there is effectively one combinatorial move with total space $\mathcal{P}_{A_4}$. It can be realized as the movie in Figure 8.12.

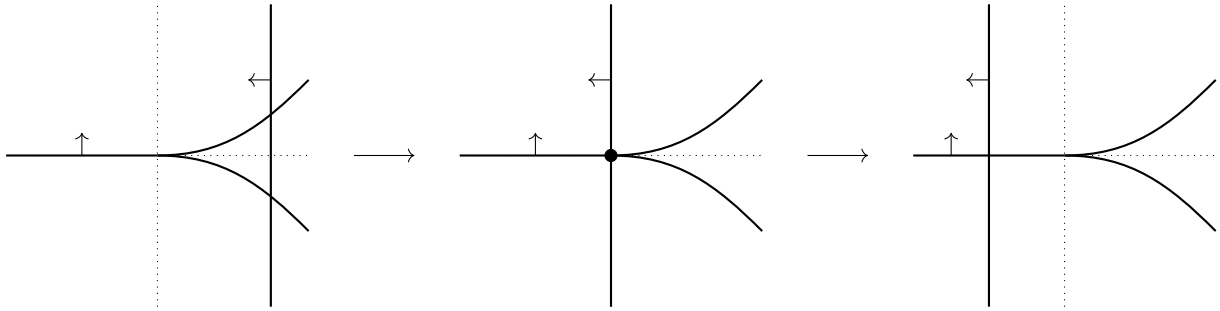


Figure 8.12: Realizing the $\mathcal{P}_{A_4}$ move with a movie

Preservation of categories holds for such a movie. Recall that *points* in the movie are *one-cells* of L_{A_4}, i.e. elements of K_{A_4}- ignoring the $t = 0$ slice, where the origin is the minimal strata $\hat{0}$. We can label all the cells appearing in the zero-section of the $t < 0$ state, this is shown in Figure 8.13.

The labelling of the one-cells is arbitrary. We can restrict to neighborhoods of codimension-one strata to deduce pieces of the cyclic order:

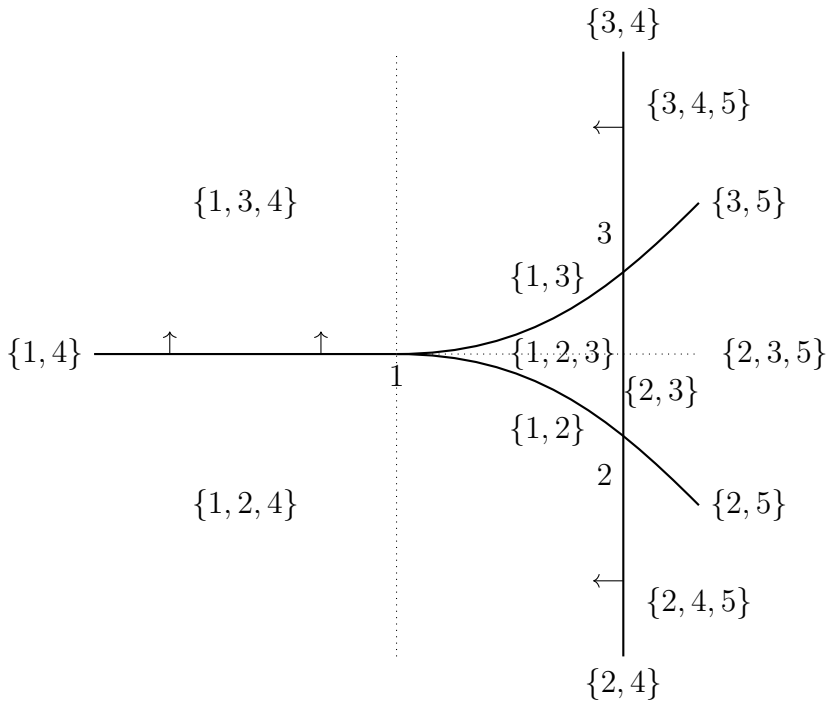
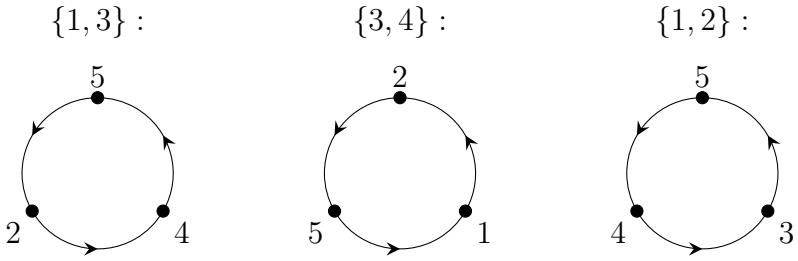


Figure 8.13: Labelled Cells In the $\mathcal{P}_{A_4}$ Movie



The pieces above are enough to deduce the full cyclic order, illustrated in Figure 8.14, with (as in Chapter 7) filled in dots correspond to K_1, or $t < 0$.

There are four combinatorial moves with total space $\mathcal{P}_{S_4}$. A first example is given in Figure 8.15, with the polarization diagram (we leave the detailed calculation to the reader).

Here is another, different-looking example, in Figure 8.16.

This second move raises some interesting questions. In particular, we notice that if we just shifted the horizontal line slightly vertically, the second move would look like a sequence of two moves: The first S_4 move, and another simpler one we haven't diagrammed yet shown

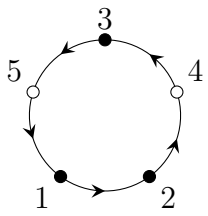


Figure 8.14: The Partition Corresponding to the $\mathcal{P}_{A_3}$ Move

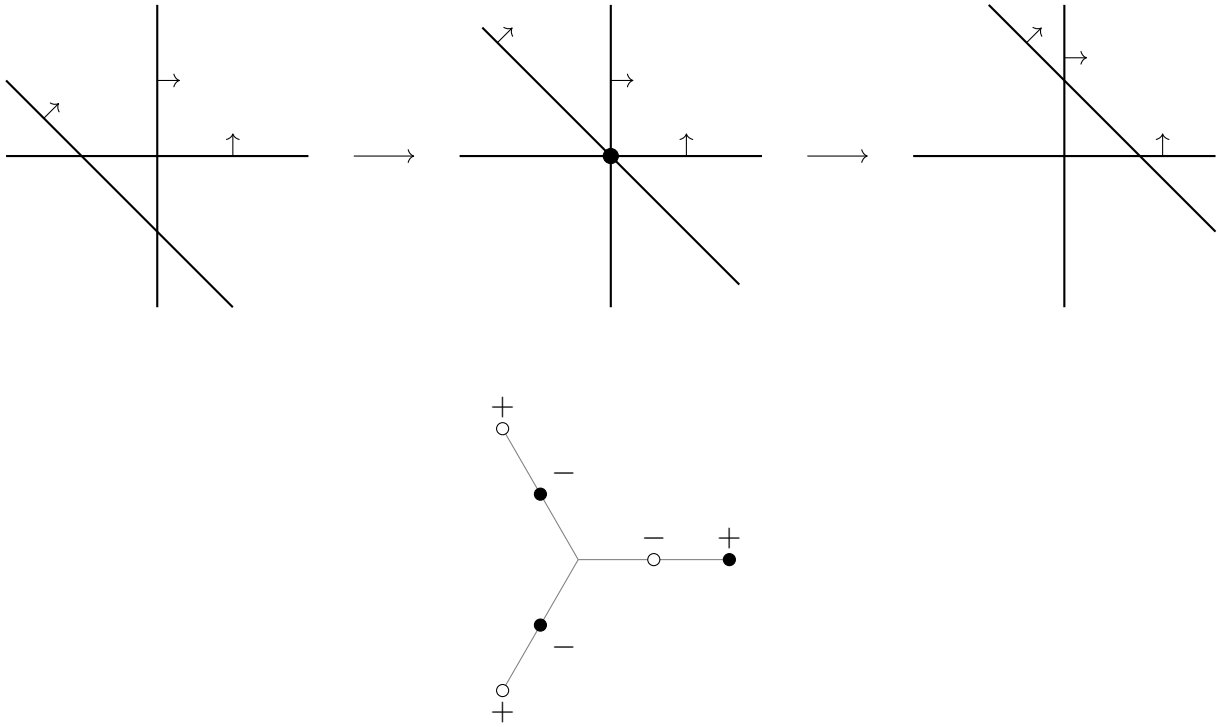


Figure 8.15: One Movie Realizing a $\mathcal{P}_{S_4}$ Move

in Figure 8.17.

Toward a Classification of Moves

Pursuing 'Dream '1', it is sensible to wonder what role the movie in Figure 8.17 plays. One can show the move in Figure 8.17 has total space $L_{A_3} \times \mathbb{R}$, so it is *not*, strictly speaking, an elementary arboreal move as defined in the previous section. However, it seems to deserve the status of being an elementary move, which suggests we extend the definition of elementary move to allow total spaces of the form $L_T \times \mathbb{R}^k$:

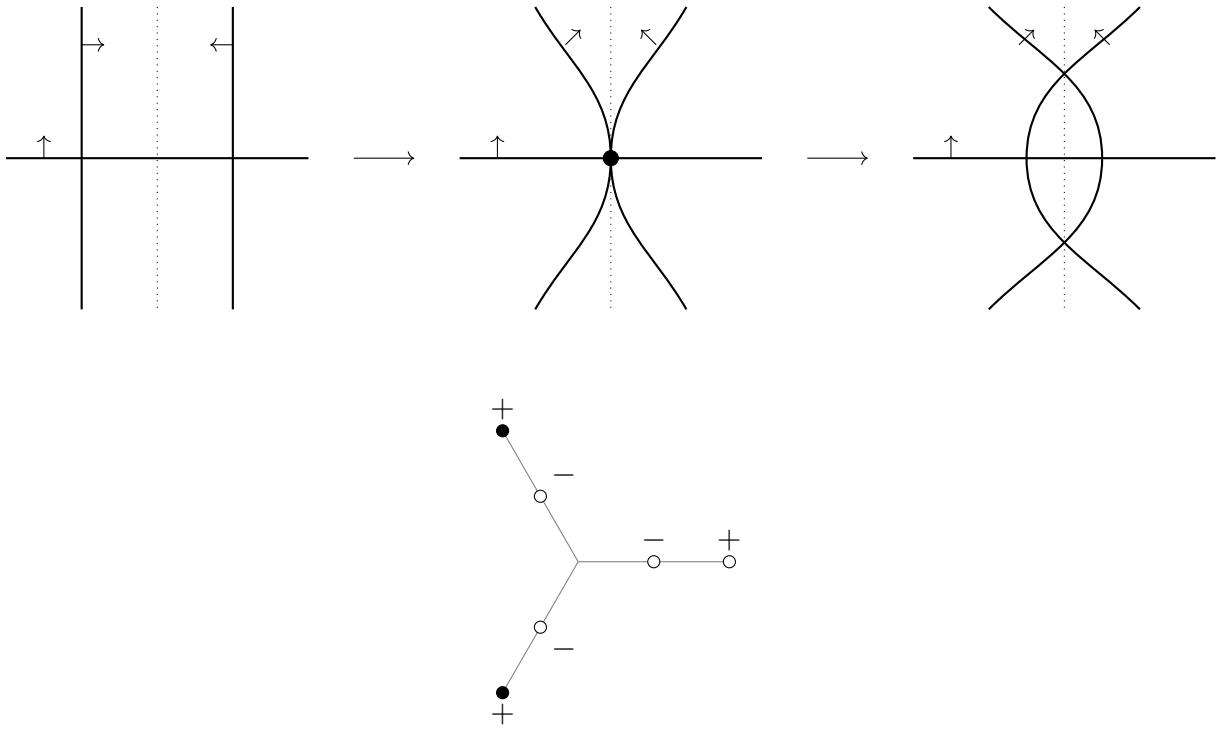


Figure 8.16: A Second $\mathcal{P}_{S_4}$ Movie

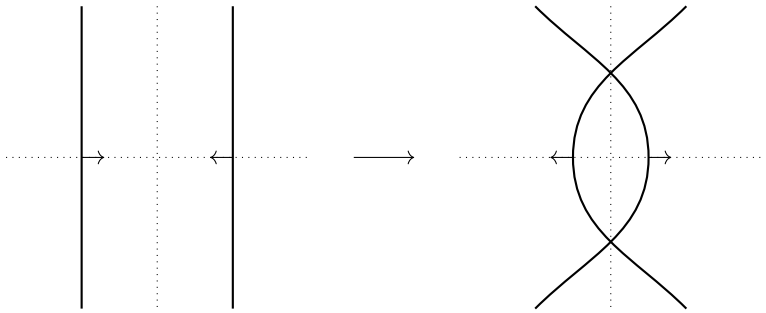


Figure 8.17: A movie realizing a new move?

Definition 8.5.4. *A **generalized elementary move** is an elementary move with total space $L_T \times \mathbb{R}^k$ for some nonnegative integer k*

To justify our assertion that this move has total space $L_{A_3} \times \mathbb{R}$, we can interpret it as a *move of a move*. To explain, look at the ending position in the figure. If we make horizontal

slices and read the image from the top-down, it looks like the H-to-I move (two points moving past each other in one dimension, Figure 8.11), then the H-to-I move run in the opposite direction. For an illustration of these cross-sections, see Figure 8.18.

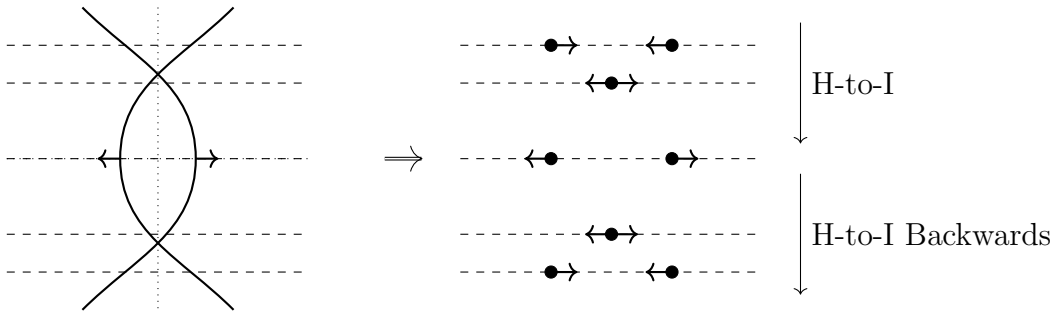


Figure 8.18: Horizontal Cross-Sections of the $L_{A_3} \times \mathbb{R}$ Movie

We can generalize this phenomenon: Any move with total space $L_T \times \mathbb{R}^k$ can be transformed into a move with total space $L_T \times \mathbb{R}^{k+1}$ where $t < 0$ corresponds to 'running the move from the end to the start to the end', while $t > 0$ corresponds to 'stay at the end', or the opposite. Formally:

Definition 8.5.5. *Let $\pi : L_T \times \mathbb{R}^k \to \mathbb{R}$ be a generalized elementary move. The **minus-suspension** is the generalized elementary move $\overline{\pi}^- : L_T \times \mathbb{R}^{k+1} \to \mathbb{R}$ defined by $\overline{\pi}^-(x, s) = \pi(x) - s^2$. The **plus-suspension** is defined by $\overline{\pi}^+(x, s) = \pi(x) + s^2$.*

Then the move illustrated in Figure 8.17 is the plus-suspension of the H-to-I move. We can iterate suspensions as well. Applying two minus-suspensions of the H-to-I move results in a move which looks like two intersecting cooriented spheres being pulled apart, while applying a minus-suspension and a plus suspension gives two 'saddles' moving past each other.

A more sophisticated conjecture might read:

Conjecture 8.5.1. *Any generalized elementary move is, up to equivalence, the iterated suspension of an elementary move.*

Going beyond this, there is another idea raised by Figure 8.16- namely, that some of the combinatorial arboreal moves described in Chapter 7 have realizations which are not 'generic', in the sense that a small perturbation reveals them to be compositions of other (possibly generalized) elementary moves. My conjecture is that any combinatorial move

$K_T = E_+ \sqcup E_-$ in which any boundary chain is entirely contained within either E_+ or E_- is of this form. Formulation of a precise conjecture would require a well defined notion of generic move and of composition of moves, both of which are beyond the scope of this discussion.

To close this section, I want to illustrate two movies which also appear in the context of Weinstein isotopies of Lagrangian submanifolds- see Figure 8.19.

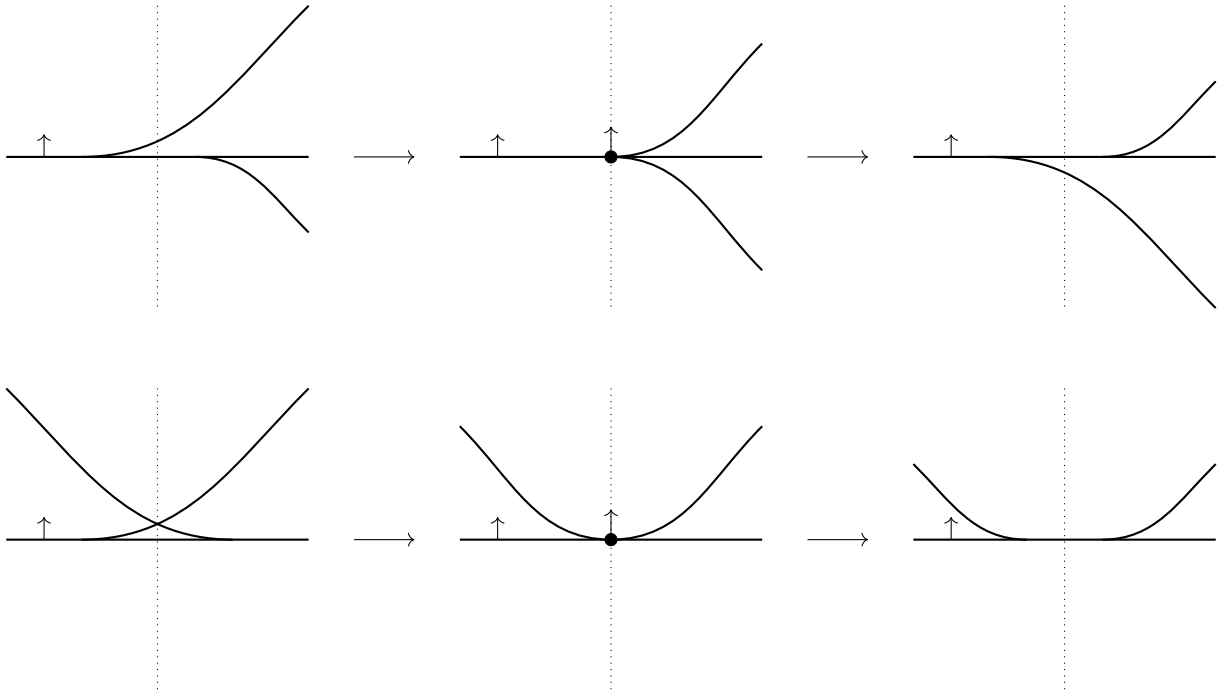


Figure 8.19: Two New Moves?

These moves don't appear in any of the pictures I've drawn so far and, in fact, are not elementary moves according to the definition. One can check that the total spaces of these moves are arboreal- the first is L_{A_4} and the second is L_{S_4}. However, both of them 'change the topology at infinity' (in fact, they look like an H-to-I move at infinity)- we might call them **non-compact moves**, while what we have been discussing are **compact moves**.

Specifically, it is property (iii) in the definition which ensures this idea of compactness, and which both of these moves violate. Looking at the $t = 0$ slice for both moves, we see that there is a one-dimensional cell which is entirely contained within that slice (the conormal vector at the origin).

If we want to include such moves, this suggests we should define non-compact elementary moves by eliminating property (iii)- though the observation that each of these non-compact moves is an H-to-I move at infinity perhaps suggests something more subtle. We propose the following notion:

Definition 8.5.6. *A **Type I Noncompact Elementary Move** with total space L_T is a function $\pi : L_T \to \mathbb{R}$ satisfying all the properties in definition 8.5.1 except property (iii), and additionally satisfying:*

(i) There is a unique non-minimal strata $L_T(\mathfrak{p})$ which is entirely contained within $\pi^{-1}(0)$. (It follows that $L_T(\mathfrak{p})$ is one-dimensional)

(ii) Letting $U \cong L_R \times \mathbb{R}$ denote an open neighborhood of $L_T(\mathfrak{p})$, the composition $L_R \times \{0\} \hookrightarrow U \xrightarrow{\pi} \mathbb{R}$ is an elementary move with total space L_R.

We can similarly modify the combinatorial definition:

Definition 8.5.7. *A **Type I noncompact combinatorial move** with total space $\mathcal{P}_T$ equipped with a cyclic structure $\mathcal{O}$ is a decomposition $K_T = K_1 \sqcup \{e_0\} \sqcup E_2$, such that K_1, K_2 are good, and the triple of data $(U_{E_1} \cap N(\langle e_0 \rangle), U_{E_2} \cap N(\langle e_0 \rangle), N(\langle e_0 \rangle))$ determines a combinatorial arboreal move with total space $N(\langle e_0 \rangle)$.*

With this notion, we can view the Type I noncompact moves in Figure 8.19 as realizations of noncompact combinatorial moves diagrammed in Figures 8.20 and 8.21, with the double nested circle denoting $\{e_0\}$, and the one-cells labelled:

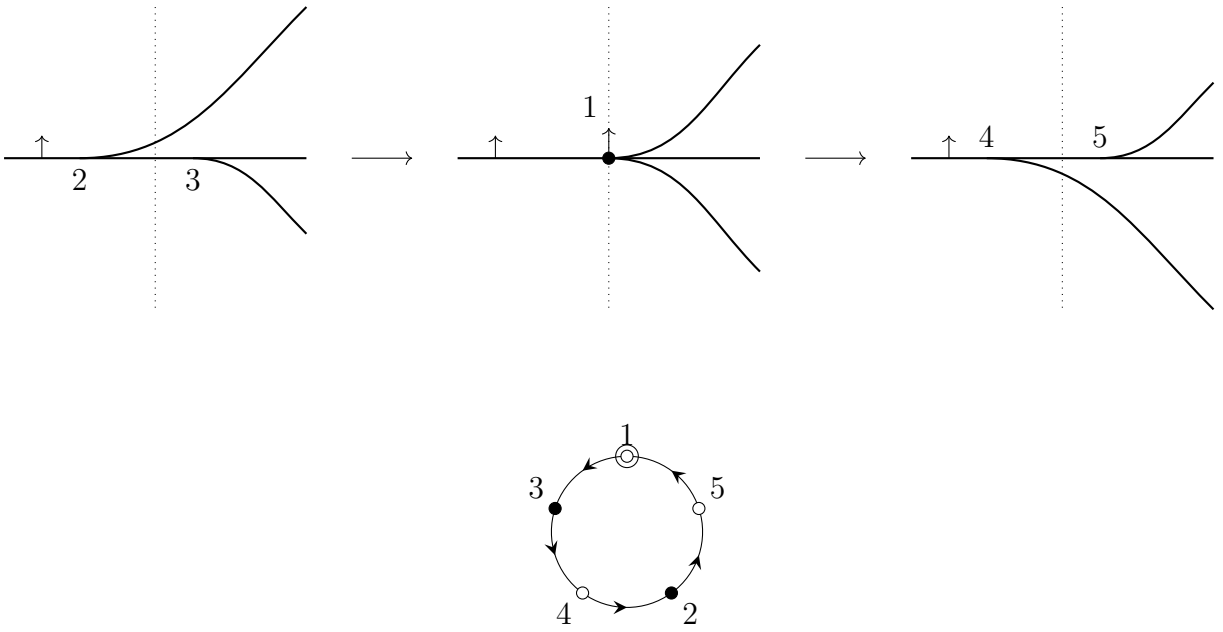


Figure 8.20: Diagramming the Type I Noncompact Move of Type A_4

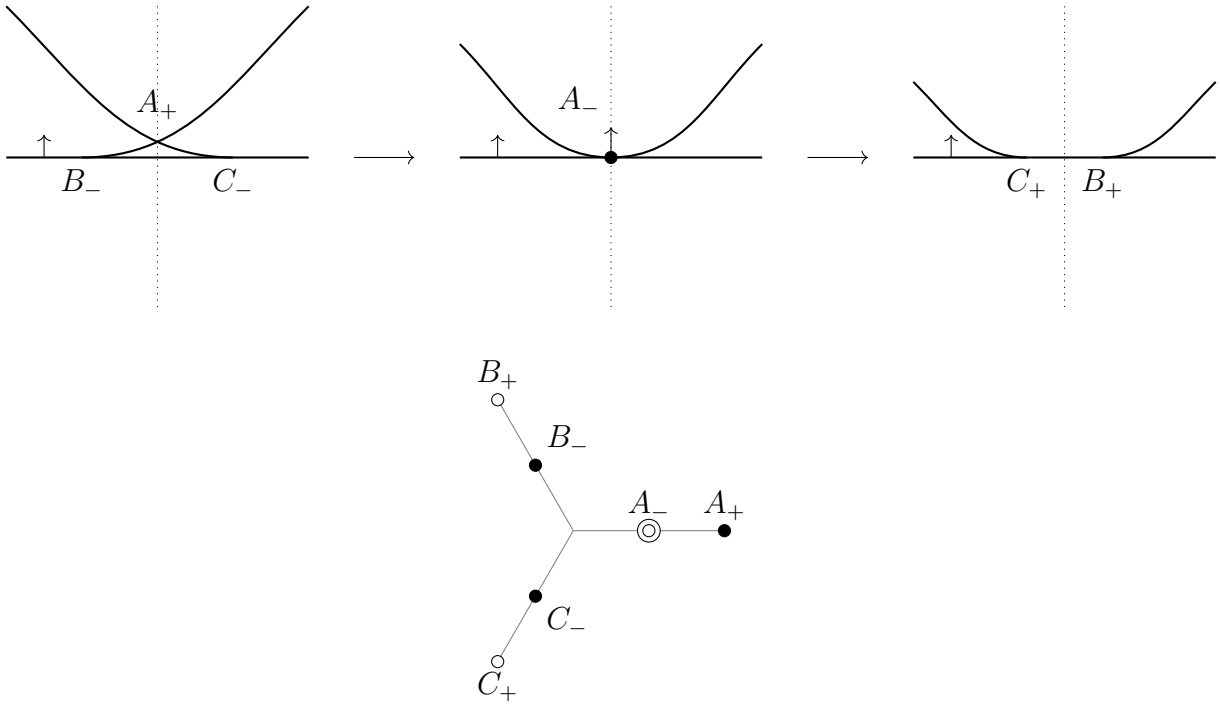


Figure 8.21: Diagramming the Type I Noncompact Move of type S_4

Appendix A

Cyclic Orders on Sets

There are multiple notions of a cyclic order, which we discuss here. Many proofs are straightforward and have been omitted. A reference is [12].

Definition A.0.1. *Let S be a finite set. A **cyclic order** on S is a ternary relation c on S satsifying:*

(i) ***Cyclicity:*** *If $c(x, y, z)$, then $c(y, z, x)$.*

(ii) ***Antisymmetry:*** *If $c(x, y, z)$, then not $c(x, z, y)$.*

(iii) ***Transitivity:*** *If $c(x, y, z)$ and $c(x, z, t)$, then $c(x, y, t)$.*

(iv) ***Totality:*** *If x, y, z are distinct, then either $c(x, y, z)$ or $c(x, z, y)$.*

Note that, by antisymmetry, if $c(x, y, z)$, then x, y, z are all distinct. Cyclic orders have a standard geometric interpretation:

Lemma A.0.1. *Giving a cyclic order on a set is equivalent to choosing an embedding $i : S \hookrightarrow S^1$ up to isotopy (Where S^1 is the standard oriented unit circle). An embedding i induces a cyclic order c on S, where $c(x, y, z)$ if and only if x, y, z are distinct, and in travelling counterclockwise from $i(x)$ to $i(z)$, we pass $i(y)$.*

Furthermore, we can understand cyclic orders in terms of successor functions:

Definition A.0.2. *A **successor function** $s : S \to S$ is a function such that, for any $x \in S$, $S = \{x, s(x), \ldots, s^{n-1}(x)\}$.*

Lemma A.0.2. *Giving a cyclic order on a set S is equivalent to fixing a successor function $s : S \to S$. The successor function determined by a cyclic order c is uniqely determined by: $c(x, s(x), y)$ whenever $y \neq x, s(x)$.*

A typical way of constructing a cyclic order is by giving a successor function. Another is using a total order:

Definition A.0.3. *Given a total order $<$ on S, the **induced cyclic order** is the one where $c(x, y, z)$ if and only if $x < y < z$, or $y < z < x$, or $z < x < y$.*

In particular, given the set $S = \{0, 1, \ldots, n-1\}$, the following cyclic orders are all equivalent:

- The cyclic order induced by the total order $0 < 1 < \ldots < n - 1$.
- The cyclic order described by the successor function $s(k) = k + 1$, with addition taken mod n.
- The cyclic order described by the embedding $S \hookrightarrow S^1$ given by $k \mapsto e^{\frac{2k\pi i}{n}}$.

Lemma A.0.3. *If S has n elements, there are $(n-1)!$ cyclic orders on S.*

Proof. Fix $x \in S$. Then successor functions on S are in bijection with permutations of $S \backslash \{x\}$ via $s \leftrightarrow (s(x), s^2(x), \ldots, s^{n-1}(x))$. □

In Chapters 6 and 7, the notion of a cyclic interval is referenced.

Definition A.0.4. *If S is a set with cyclic order c, and $x, y \in S$, then the **open cyclic interval** (x, y) is defined to be $\{z \in S \mid c(x, z, y)\}$. The **closed cyclic interval** $[x, y]$ is defined to be $(x, y) \cup \{x, y\}$. A **cyclic interval** is an open or closed cyclic interval in S.*

Appendix B

Infinity-Categories and Sheaves

B.1 Quasi-Categories

We use quasi-categories as a model for $\infty-$categories (or, $(\infty,1)-$categories). A canonical reference is [16]. Information on the $\infty-$category underlying model categories and the model category of dg-categories can be found in [8], [3], [34], see also the review in [32].

Definition B.1.1. *An $\infty-$category is a simplicial set $\mathcal{C}$ in which all inner horns can be filled- meaning, every map $\Lambda^n_k \to \mathcal{C}$ with $0 < k < n$ can be extended to a map $\Delta^n \to \mathcal{C}$.*

Recall that a simplicial set is a functor $\mathcal{C} : \Delta^{op} \to \mathbf{Set}$, where Δ is the *simplex* category whose objects are the totally ordered sets $[n] = \{0, 1, \ldots, n\}$, and whose morphisms are order-preserving functions. [1] The set $\mathcal{C}([n])$ is called the set of $n-$simplices of $\mathcal{C}$. Δ^n is the simplicial set $\mathrm{Hom}_\Delta(\cdot, [n])$, and the horn Λ^n_k for $0 \leq k \leq n$ is the sub-simplicial set consisting of homs whose image avoids $k \in [n]$.

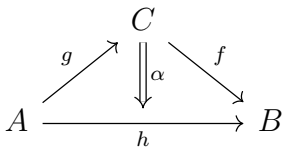
In a simplicial set X, for $n \geq 0$ and $0 \leq k \leq n$ we have the *face maps* $d^n_k : \mathcal{C}([n]) \to \mathcal{C}([n-1])$ given by the inclusions $[n-1] \hookrightarrow [n]$ whose image avoids k, and the *degeneracy maps* $s^n_k : \mathcal{C}([n]) \to \mathcal{C}([n+1])$ given by the surjections $[n+1] \to [n]$ whose preimage of k consists of two elements.

By *objects* of $\mathcal{C}$ we will mean the set of zero-simplices, and by *morphisms* $A \to B$ we will mean one-simplices f such that $d^1_1(f) = A$ and $d^1_0(f) = B$.

When $n = 2$ and $k = 1$, the condition that an inner horn $\Lambda^2_1 \to \mathcal{C}$ can be filled can be interpreted as saying that for any pair of morphisms $f : A \to B$ and $g : B \to C$, there is a third morphism (the 'composition') $h : A \to C$ and a 2-simplex $\alpha \in \mathcal{C}([2])$ such that

[1] In practice, we will usually have *class* of objects and not a set- this technicality is easily incorporated into the theory (for example by replacing 'set' with 'discrete category') and we will not worry about it here.

$d_0^2(\alpha) = g$, $d_1^2(\alpha) = h$, $d_2^2(\alpha) = f$:



h and α do not have to be unique, but one can show that there is a contractible space of choices for h.

A common source of ∞-categories is as the nerve of small categories:

Definition B.1.2. *If C is a small category, the **nerve** $\mathcal{N}C$ is the $\infty-$category where:*

- *An element of the set $\mathcal{N}C([n])$ is an $n-$tuple of composable morphisms*

$$X_0 \to X_1 \to \cdots \to X_n$$

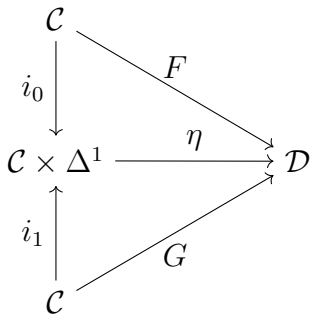
- *The face maps $d_k^n : \mathcal{N}C([n]) \to \mathcal{N}C([n-1])$ are given by removing the kth object X_k, and composing the two maps meeting at X_k if $0 < k < n$.*
- *The degeneracy maps $s_k^n : \mathcal{N}C([k]) \to \mathcal{N}C([k+1])$ are given by duplicating X_k and inserting an identity morphism.*

If $\mathcal{P}$ is a poset, we can form the $\infty-$category $\mathcal{N}\mathcal{P}$ by viewing $\mathcal{P}$ as a category in which there is a unique morphism $x \to y$ when $x \leq y$.

We introduce some basic definitions:

Definition B.1.3.

- *Given two $\infty-$categories $\mathcal{C}, \mathcal{D} : \Delta^{op} \to$ **Set**, a **functor** $F : \mathcal{C} \to \mathcal{D}$ is a natural transformation.*
- *Let J be the category with two objects, 0 and 1, and two non-identity morphisms $x : 0 \to 1$ and $y : 1 \to 0$ which are inverses. A morphism $f : A \to B$ in an $\infty-$category $\mathcal{C}$ is an **equivalence** if there is a functor $F : \mathcal{N}J \to \mathcal{C}$ such that $F(x) = f$.*
- *If $\mathcal{C}$ and $\mathcal{D}$ are $\infty-$categories, the product $\mathcal{C} \times \mathcal{D}$ is defined by $(\mathcal{C} \times \mathcal{D})([k]) = \mathcal{C}([k]) \times \mathcal{D}([k])$, with maps defined component-wise.*
- *If $F, G : \mathcal{C} \to \mathcal{D}$ are functors, a **homotopy** $F \to G$ is a functor $\eta : \mathcal{C} \times \Delta^1 \to \mathcal{D}$ making the following diagram commute:*



Where, for $\alpha = 0, 1$, $i_\alpha : \mathcal{C}[k] \to (\mathcal{C} \times \Delta^1)[k] = \mathcal{C}[k] \times Hom([k], [1])$ *is the map* $S \mapsto (S, \pi_\alpha)$, π_α *being the constant function* $j \mapsto \alpha$.

When $\mathcal{C}$ and $\mathcal{D}$ are $\infty-$categories, $\mathrm{Hom}(\mathcal{C}, \mathcal{D})$ also has the structure of an $\infty-$category, where the objects are functors and the morphisms are homotopies- higher simplices can be similarly defined, see [16].

Another source of $\infty-$categories is the localization of simplicial model categories- that is, model categories endowed with a compatible simplicial enrichment. The $\infty-$category constructed- which we call the 'underlying $\infty-$category', has objects which are the fibrant and cofibrant objects of the model category, and morphisms are the usual morphisms of these objects. Model categories are equipped with a subcollection of morphisms called 'weak equivalences'. The key property of the underlying $\infty-$category is that the equivalences (in the sense of definition B.1.3) are exactly the weak equivalences.

For us, the two key examples of simplicial model categories are:

- dg-Cat, the category of dg-categories, equipped with the Dwyer-Kan model structure. The weak equivalences are quasi-equivalences of dg-categories (see [8]).
- Ch, the category of bounded chain complexes [2], equipped with the standard (Quillen) model structure. The weak equivalences are quasi-isomorphisms.

We restate an important definition here:

Definition B.1.4. *Let* $\mathcal{C}$ *be an ordinary category, and let dg-Cat denote the* $\infty-$*category of dg-categories (over* k*). A* ***2-functor*** $\mathcal{Q} : \mathcal{C} \to$ *dg-Cat consists of the following data:*

[2] In this paper, the differential in chain complexes always increases the degree. These are sometimes called 'cochain complexes'

(i) For each object $X \in Ob(\mathcal{C})$, a dg-category $\mathcal{Q}_X$.

(ii) For each morphism $f : X \to Y$ in $\mathcal{C}$, a dg-functor $c_f : \mathcal{Q}_X \to \mathcal{Q}_Y$.

(iii) For each pair of morphisms f, g such that the composition fg exists, a quasi-isomorphism $\eta_{g,f} : c_f \circ c_g \to c_{fg}$

(iv) Such that, for each triple f, g, h of morphisms such that the composition fgh exists, the following diagram commutes:

$$\begin{array}{ccc} c_f \circ c_g \circ c_h & \xrightarrow{\eta_{g,f} \diamond Id_{c_h}} & c_{fg} \circ c_h \\ {\scriptstyle Id_{c_f} \diamond \eta_{h,g}} \downarrow & & \downarrow {\scriptstyle \eta_{h,fg}} \\ c_f \circ c_{gh} & \xrightarrow{\eta_{gh,f}} & c_{fgh} \end{array}$$

Where $\diamond$ refers to the horizontal composition of functors.

The above notion will allow for the construction of $\infty-$functors:

Lemma B.1.1. *If $\mathcal{I}_n$ denotes the poset with $n+1$ objects $0 < 1 < \ldots < n$ (which can also be viewed as a category), then:*

(i) A 2-functor $\mathcal{I}_n \to$ dg-Cat describes an n-simplex in dg-Cat.

(ii) For any ordinary category $\mathcal{C}$, a 2-functor $\mathcal{C} \to$ dg-Cat determines an $\infty-$functor $\mathcal{N}\mathcal{C} \to$ dg-Cat.

B.2 Sheaves

Now, let X be a topological space. The collection of open sets τ is a poset, where we let $U \leq V$ if and only if $V \subseteq U$.

Definition B.2.1. *If X is a topological space and $\mathcal{C}$ is an $\infty-$category, a **presheaf** on X valued in $\mathcal{C}$ is a functor $\mathcal{F} : \mathcal{N}\tau \to \mathcal{C}$.*

There is a notion of $\infty-$categorical limits, which are well-defined up to a contractible space of choices. Then we have:

Definition B.2.2. *If X is a topological space and $\mathcal{C}$ is an $\infty-$category with small limits, a **sheaf** on X valued in $\mathcal{C}$ is a presheaf $\mathcal{F}$ satisfying the following: Let U be any open set, and $\mathscr{U}$ any open cover of U for which, if $V_1 \in \mathscr{U}$ and $V_2 \subseteq V_1$ then $V_2 \in \mathscr{U}$. Then the map:*

$$\mathcal{F}(U) \to \varprojlim_{V \in \mathscr{U}} \mathcal{F}(V) \tag{B.1}$$

Is an isomorphism.

We let $\mathrm{Sh}(X;\mathcal{C})$ denote the category of sheaves on X. Morphisms in $\mathrm{Sh}(X;\mathcal{C})$ are homotopies of functors. A sheaf on X is determined by its values on a base for the topology of X. Explicitly:

Proposition B.2.1. *Let $\mathcal{B}$ be a base for a topology τ on X. Define $Sh_{\mathcal{B}}(X;\mathcal{C})$ to be the category of functors $\mathcal{N}\mathcal{B} \to \mathcal{C}$ satisfying condition B.1 restricted to $\mathcal{B}$. Then the inclusion $\mathcal{N}\mathcal{B} \hookrightarrow \mathcal{N}\tau$ gives an equivalence of categories $Sh_{\mathcal{B}}(X;\mathcal{C}) \cong Sh(X;\mathcal{C})$.*

As a result, we can define:

Definition B.2.3. *If $\mathcal{P}$ is a poset and $\mathcal{C}$ is an $\infty-$category with small limits, a **sheaf** on $\mathcal{P}$ valued in $\mathcal{C}$ is an $\infty-$functor $\mathcal{F} : \mathcal{N}\mathcal{P} \to \mathcal{C}$. For $x \in \mathcal{P}$, $\mathcal{F}_x \in \mathcal{C}$ is called the **stalk** of $\mathcal{F}$ at x. The morphisms $\mathcal{F}_x \to \mathcal{F}_y$ when $x \leq y$ are the **generization maps**.*

This is equivalent to the previous definition. To see why, let U be an open subset of $\mathcal{P}$, considered as a topological space. If $\mathcal{F}$ is a sheaf on $\mathcal{P}$, we define:

$$\Gamma(\mathcal{F};U) = \varprojlim_{x \in U} \mathcal{F}_x$$

Then we have $\Gamma(\mathcal{F};N(x)) \cong \mathcal{F}_x$, and the $N(x)$ form a base of the topology for $\mathcal{P}$. Furthermore, any open cover of $N(x)$ must include $N(x)$, meaning there are no extra compatibility conditions among the stalks.

We can then show:

Lemma B.2.1. *Let $\mathcal{C}$ be a category in which every morphism is monic, and $\mathcal{Q} : \mathcal{C}^{op} \to dg\text{-}Cat$ be a 2-functor. Then for each object $A \in Ob(\mathcal{C})$, $\mathcal{Q}$ defines a sheaf $\mathcal{Q}_A$ on the poset $\mathcal{P}_A$.*

Proof. Let $(\mathcal{Q}, c, \eta)$ denote the triple of data associated to a 2-functor. For $A \in \mathrm{Ob}(\mathcal{C})$ we set:

- For $f : B \to A$, the stalk of $\mathcal{Q}_A$ at $[f] \in \mathcal{P}_A$ is:

$$\mathcal{Q}_{A,[f]} = \mathcal{Q}(B)$$

- For $h : C \to B$, let $g = fh$. Then $[g] \geq [f]$ and the generization map:

$$c_{[f][g]} : \mathcal{Q}_{A,[f]} \to \mathcal{Q}_{A,[g]} := c_h : \mathcal{Q}(B) \to \mathcal{Q}(C)$$

- For $j : D \to C$, let $k = gj$. Then:

$$\eta_{[f][g],[g][k]} := \eta_{h,j}$$

This defines a 2-functor, and hence a sheaf. □

B.3 Morphisms Between Sheaves

We have the following notion:

Definition B.3.1. *Let $\mathcal{C}$ be a category, and $\mathcal{Q}, \mathcal{Q}'$ 2-functors $\mathcal{C} \to$ dg-Cat, with associated data c, η, resp. c', η'. A **left morphism** $\mathcal{F} : \mathcal{Q} \to \mathcal{Q}'$ consists of:*

(i) For each $X \in Ob(\mathcal{C})$, a dg-functor $\mathcal{F}(X) : \mathcal{Q}(X) \to \mathcal{Q}'(X)$.

(ii) For each morphism $f : Y \to X$, a quasi-isomorphism $\gamma_f : c'_f \circ \mathcal{F}(Y) \to \mathcal{F}(X) \circ c_f$.

(iii) Such that, for morphisms $f : Y \to X$ and $g : Z \to Y$, the following diagram commutes:

$$\begin{array}{ccccc} c'_f \circ c'_g \circ \mathcal{F}(X) & \xrightarrow{1_{c'_f} \diamond \gamma_g} & c'_f \circ \mathcal{F}(Y) \circ c_g & \xrightarrow{\gamma_f \diamond 1_g} & \mathcal{F}(Z) \circ c_f \circ c_g \\ \downarrow{\scriptstyle \eta'_{f,g} \diamond 1_{\mathcal{F}(X)}} & & & & \downarrow{\scriptstyle 1_{\mathcal{F}(Z)} \diamond \eta_{f,g}} \\ c'_{fg} \circ \mathcal{F}(X) & & \xrightarrow{\gamma_{fg}} & & \mathcal{F}(Z) \circ c_{fg} \end{array}$$

*A **right morphism** is defined similarly, except $\gamma_f : \mathcal{F}(X) \circ c_f \to c'_f \circ \mathcal{F}(Y)$, and in the commutative diagram in (iii), the horizontal arrows point in the other direction.*

Importantly for us, left/right morphisms define morphisms of sheaves:

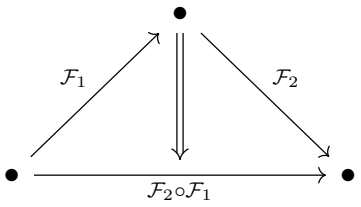
Lemma B.3.1. *Let $\mathcal{C}$ be a category in which every morphism is monic, and $\mathcal{Q}, \mathcal{Q}' : \mathcal{C}^{op} \to$ dg-Cat are 2-functors. If $F : \mathcal{Q} \to \mathcal{Q}'$ is a left or right morphism, F defines a morphism of sheaves.*

Proof. In light of definition B.1.3, we need to construct an $\infty-$functor $\mathcal{N}\mathcal{P}_A \times \Delta^1 \to$ dg-Cat fitting into the appropriate commutative diagram. Note that $\mathcal{N}\mathcal{P} \times \Delta^1 \cong \mathcal{N}(\mathcal{P}_A \times \mathcal{I}_1)$, where $\mathcal{I}_1$ is the poset $\{0 < 1\}$. We construct a 2-functor (also called $\mathcal{F}$), $\mathcal{F} : \mathcal{P}_A \times \mathcal{I}_1 \to$ dg-Cat, with C, N the compatibility data:

- Let $f : B \to A$. We let $[f]_0, [f]_1$ denote elements of $\mathcal{P}_A \times \mathcal{I}_1$. The 2-functor sends:
  (i) $\mathcal{F}_{[f]_0} = \mathcal{Q}_{A,[f]} = \mathcal{Q}(B)$
  (ii) $\mathcal{F}_{[f]_1} = \mathcal{Q}'_{A,[f]} = \mathcal{Q}(B)$
- Now suppose $h : C \to B$, and write $g = fh$. The generization maps:
$$C_{[f]_0[g]_0} : \mathcal{F}_{[f]_0} \to \mathcal{F}_{[g]_0}$$
$$C_{[f]_1[g]_1} : \mathcal{F}_{[f]_1} \to \mathcal{F}_{[g]_1}$$
are given by c_h and c'_h, respectively. We set:
$$C_{[f]_0[g]_1} = \mathcal{F}(C) \circ c_h : \mathcal{Q}(B) \to \mathcal{Q}'(C)$$
- Let $j : D \to C$, and $k = gj$. Then we have:
$$C_{[g]_0[k]_1} \circ C_{[f]_0)[g]_0} = \mathcal{F}(D) \circ c_j \circ c_h$$
$$C_{[f]_0[k]_1} = \mathcal{F}(D) \circ c_{jh}$$
And we can set $N_{[g]_0[k]_1,[f]_0[g]_0} = \eta_{j,h}$. We also have:
$$C_{[g]_1[k]_1} \circ C_{[f]_0[g]_1} = c'_j \circ \mathcal{F}(C) \circ c_h$$
And we can set $N_{[g]_1[k]_1,[f]_0[g]_0}$ to be the composition:
$$c'_j \circ \mathcal{F}(C) \circ c_h \xrightarrow{\gamma_j \diamond 1_{c_h}} \mathcal{F}(D) \circ c_j \circ c_h \xrightarrow{\eta_{j,h}} \mathcal{F}(D) \circ c_{jh}$$
We also have $N_{[g]_0[k]_0,[f]_0[g]_0} = \eta_{j,h}$, and $N_{[g]_1[k]_1,[f]_1[g]_1} = \eta'_{j,h}$.

Then the commutative diagram in the definition of left morphism ensures that all the properties of being a 2-functor are satisfied. The construction is analogous for right morphisms. □

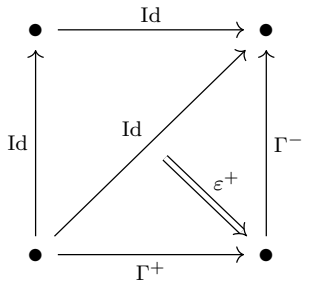
Given 2-functors $\mathcal{F}_1, \mathcal{F}_2 : \mathcal{P}_A \times \mathcal{I}_1 \to$ dg-Cat such that $\mathcal{F}_1 \circ i_1 = \mathcal{F}_2 \circ i_0 : \mathcal{P}_A \to$ dg-Cat, there is a natural way to compose them to get $\mathcal{F}_2 \circ \mathcal{F}_1 : \mathcal{P}_A \times \mathcal{I}_1 \to$ dg-Cat, and furthermore, these morphisms fit into a 2-simplex in the $\infty-$category $\mathrm{Hom}(\mathcal{N}\mathcal{P}, \text{dg-Cat})$:

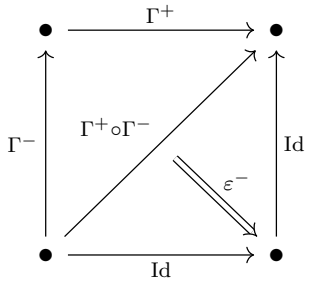


We conclude this section with a discussion on our construction in Chapter 5: We construct quasi-isomorphisms:

$$\varepsilon_v^+ : \Gamma_v^- \circ \Gamma_v^+ \to \mathrm{Id}$$
$$\varepsilon_v^- : \mathrm{Id} \to \Gamma_v^+ \circ \Gamma_v^-$$

Satisfying certain commutativity properties. The commutativity properties are exactly what is needed to create diagrams of 2-simplices in $\mathrm{Hom}(\mathcal{NP}, \text{dg-Cat})$:





The construction of the required 2-simplex and verification of that claim are just as in B.3.1. The existence of these diagrams imply the data (Γ_v^+, γ^+) and (Γ_v^-, γ^-) give quasi-inverse morphisms of sheaves of dg-categories.

B.4 Pullback Functors and Homotopy

In this section we discuss the proof of the lemmas in Chapter 7. Fix an $\infty-$category $\mathcal{C}$ with small limits.

If $f : \mathcal{P}_1 \to \mathcal{P}_2$ is a morphism of posets, it induces a functor of $\infty-$categories $f : \mathcal{N}\mathcal{P}_1 \to \mathcal{N}\mathcal{P}_2$[3].

Definition B.4.1. *If $\mathcal{Q} : \mathcal{N}\mathcal{P}_2 \to \mathcal{C}$ is an object in $Sh(\mathcal{P}_2; \mathcal{C})$, we define the* ***pullback*** $f^*\mathcal{Q} = \mathcal{Q} \circ f \in Sh(\mathcal{P}_1; \mathcal{C})$.

Now, let $\mathcal{Q} : \mathcal{N}\mathcal{P} \to \mathcal{C}$ be a sheaf. There is a poset $\mathrm{C}(\mathcal{P})$, the *cone* over $\mathcal{P}$, which is obtained by adjoining a minimal element $\hat{0}$ to $\mathcal{P}$. The limit:

$$\varprojlim_{x \in \mathcal{P}} \mathcal{Q}_x$$

Is given by the universal (in the $\infty-$categorical sense) functor $\overline{\mathcal{Q}}$ making the following diagram commute:

$$\begin{array}{ccc} N\mathcal{P} & & \\ {\scriptstyle j_{\mathcal{P}}}\downarrow & \searrow^{\mathcal{Q}} & \\ \mathcal{N}\mathrm{C}(\mathcal{P}) & \xrightarrow[\overline{\mathcal{Q}}]{} & \mathcal{C} \end{array}$$

Now, let $\mathcal{P}$ be a poset, and $i : Z \hookrightarrow \mathcal{P}$ be the inclusion of a subset. This gives a commutative diagram:

$$\begin{array}{ccc} \mathcal{N}Z & \xrightarrow{\quad i \quad} & \mathcal{N}\mathcal{P} \\ \downarrow{\scriptstyle j_Z} & & \downarrow{\scriptstyle j_{\mathcal{P}}} \\ \mathcal{N}\mathrm{C}(Z) & \xrightarrow{\bar{i}} & \mathcal{N}\mathrm{C}(\mathcal{P}) \end{array}$$

Which leads to:

$$\begin{array}{ccc} N\mathcal{Z} & & \\ {\scriptstyle j_Z}\downarrow & \searrow^{i^*\mathcal{Q}} & \\ \mathcal{N}\mathrm{C}(Z) & \xrightarrow[\bar{i}^*\overline{\mathcal{Q}}]{} & \mathcal{C} \end{array}$$

[3]It should not cause confusion to also call this f

By universality, this gives a morphism $\Gamma(\mathcal{P};\mathcal{Q}) \to \Gamma(Z; i^*\mathcal{Q})$. We say i is a **weak equivalence** with respect to $\mathcal{Q}$ when this morphism is an equivalence. It is enough to construct a morphism $H : \mathcal{N}\mathrm{C}(\mathcal{P}) \to \text{dg-Cat}$ making the following diagram commute:

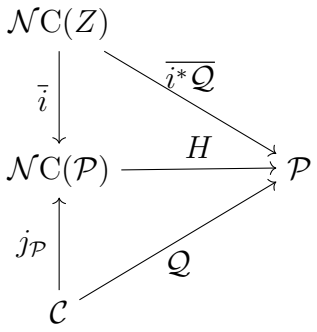


Figure B.1: A Commutative Diagram Giving $\Gamma(Z; i^*\mathcal{Q}) \to \Gamma(\mathcal{P};\mathcal{Q})$

Where $\overline{i^*\mathcal{Q}}$ is the universal morphism extending $i^*\mathcal{Q} : \mathcal{N}Z \to \text{dg-Cat}$.

We can now prove the two lemmas from Chapter 7:

Lemma B.4.1. *(Lemma 7.4.1 in the text) Let A be a subset of a poset $\mathcal{P}$, and $\mathcal{Q} \in Sh(\mathcal{P};\mathcal{C})$.*

(i) If there exists a downward combinatorial retract $r : \mathcal{P} \to A$, $i : A \hookrightarrow \mathcal{P}$ is a weak equivalence with respect to $\mathcal{Q}$.

(ii) If there exists an upward combinatorial retract $r : \mathcal{P} \to A$ and a morphism $F : r^*i^*\mathcal{Q} \to \mathcal{Q}$ such that $i^*F : i^*\mathcal{Q} \to i^*\mathcal{Q}$ is the identity morphism, then $i : A \hookrightarrow \mathcal{P}$ is a weak equivalence with respect to $\mathcal{Q}$.

Proof. We construct a diagram as in Figure B.1. H is determined on all objects and on all simplices within $\mathcal{P}$. So let

$$\Delta = p < x_1 < \cdots < x_n$$

, we construct the value $H(\Delta)$.

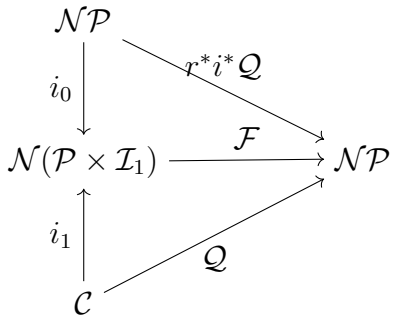
First, consider the case where $n = 1$, i.e. we have a simplex:

$$p < x$$

We have a one-simplex $f = \overline{i^*\mathcal{Q}}(p < r(x))$. If r is a downward combinatorial retract, we have another one-simplex $g = \mathcal{Q}(r(x) < x)$. Furthermore we have $\partial^0 f = \overline{i^*\mathcal{Q}}(r(x)) = \mathcal{Q}(r(x)) =$

$\partial^1 g$, hence this data gives a map $\Lambda^2_1 \to \mathcal{C}$, which by assumption is fillable. Other simplices can be defined inductively in a similar way.

If r is an upward combinatorial retract, we are given the data of a morphism $\mathcal{F} : r^* i^* \mathcal{Q} \to \mathcal{Q}$, i.e. (see definition B.1.3)



Then to the one-simplex $p < x$, let $f = \overline{i^* \mathcal{Q}}(p < r(x))$, and $g = \mathcal{F}((x, 0) < (x, 1))$. We have $\partial^1 g = \mathcal{F}((x, 0)) = r^* i^* \mathcal{Q}(x) = \mathcal{Q}(r(x))$, so again we get a map $\Lambda^2_1 \to \mathcal{C}$ which by assumption is fillable. The other simplices can similarly be defined. □

We also provide a proof of:

Lemma B.4.2. *(Lemma 7.4.2 in the text) Let U be open in $\mathcal{P}$, and $Z = \mathcal{P} \backslash U$. Let $i : Z \hookrightarrow \mathcal{P}$ and $j : U \hookrightarrow \mathcal{P}$ denote the inclusions. If $\mathcal{Q} : \mathcal{P} \to \mathcal{C}$ is a sheaf, and $j^* \mathcal{Q} = 0$, then i is a weak equivalence with respect to $\mathcal{Q}$.*

Proof. Since U is open, we can construct H as in Figure B.1 by extending $\overline{i^* \mathcal{Q}}$ to be zero on simplices which aren't contained in the image of $\mathcal{N}\mathrm{C}(\mathcal{P})$. □

Bibliography

[1] Denis Auroux. "A beginner's introduction to Fukaya categories". In: *arXiv:1301.7056* (2013).

[2] I.N. Bernstein, I.M. Gelfand, and V.A. Ponomarev. "Coxeter Functors and Gabriel's Theorem". In: *Russian Math. Surv.* 28 (1973), pp. 17–32.

[3] Toën Bertrand. "Lectures on DG-categories". In: *Topics in Algebraic and Topological K-Theory* (2011).

[4] A. Bjorner. "Posets, regular CW complexes, and Bruhat order". In: *European Journal of Combinatorics* 5 (1984).

[5] A. I. Bondal and M. M. Kapranov. "Enhanced triangulated categories". In: *Mathematics of the URSS Sbornik* 70 (1991).

[6] Michel Brion. *Representations of Quivers*. URL: `https://www-fourier.ujf-grenoble.fr/~mbrion/notes_quivers_rev.pdf`.

[7] Vladimir Drinfeld. "Quotient of dg-categories". In: *Journal of Algebra* 272.2 (2004).

[8] W. Dwyer and D. Kan. "Simplicial localization of categories". In: *Journal of Pure and Applied Algebra* 17 (1980).

[9] *From Stein to Weinstein and Back*. Vol. 59. Amer. Math. Soc., 2010.

[10] Benjamin Gammage and David Nadler. "Mirror symmetry for honeycombs". In: *arXiv:1702.03255* (2017).

[11] J.L. Harer. "The virtual cohomological dimension of the mapping class group of an orientable surface". In: *Inventiones mathematicae* 84 (1986), pp. 157–176.

[12] Edward V. Huntington. "A Set of Independent Postulates for Cyclic Order". In: *Proceedings of the National Academy of Sciences of the United States of America* 2 (1916).

[13] Masaki Kashiwara and Pierre Schapira. *Sheaves on Manifolds*. Springer-Verlag, 1990.

[14] International Congress of Mathematicians, ed. *On differential graded categories*. Vol. 2. 2006.

[15] Maxim Kontsevich. *Symplectic Geometry of Homological Algebra*. URL: `http://www.ihes.fr/~maxim/TEXTS/Symplectic_AT2009.pdf`.

[16] Jacob Lurie. *Higher Topos Theory*. Princeton University Press, 2009.

[17] John Mather. *Notes on Topological Stability*. July 1970.

[18] Motohico Mulase and Michael Penkava. "Ribbon Graphs, Quadratic Differentials on Riemann Surfaces, and Algebraic Curves Defined over $\overline{\mathbb{Q}}$". In: *1998* 2.4 (Asian Journal of Mathematics).

[19] David Nadler. "A combinatorial calculation of the Landau-Ginzburg model $M = \mathbb{C}^3$, $W = z_1z_2z_3$". In: *arXiv:1507.08735* (2015).

[20] David Nadler. "Arboreal Singularities". In: *Geometry & Topology* 21 (2017).

[21] David Nadler. "Cyclic Symmetries of A_n-Quiver Representations". In: *Advances in Mathematics* 269 (2015).

[22] David Nadler. "Microlocal branes are constructible sheaves". In: *Selecta Mathematica* 15.4 (2009).

[23] David Nadler. "Mirror symmetry for the Landau Ginzburg A-model $M = \mathbb{C}^n$, $W = z_1 \cdots z_n$". In: *arXiv:1601.02977* (2016).

[24] David Nadler. "Non-Characteristic Expansions of Legendrian Singularities". In: *arXiv 1507.01513* ().

[25] David Nadler and Eric Zaslow. "Constructible Sheaves and the Fukaya Category". In: *J. Amer. Math Soc.* 22 (2009).

[26] Lenhard Ng et al. "Augmentations are Sheaves". In: *arxiv:1502.04939* (2015).

[27] Udo Pachner. "P.L. Homeomorphic Manifolds are Equivalent by Elementary Shellings". In: *European Journal of Combinatorics* 12.2 (1991).

[28] Vivek Shende and Alex Takeda. "Symplectic structures from topological Fukaya categories". In: *arXiv:1605.02721* (2016).

[29] Vivek Shende, David Treumann, and Harold Williams. "On the combinatorics of exact Lagrangian surfaces". In: *arXiv:1603.07449* (2016).

[30] Vivek Shende, David Treumann, and Eric Zaslow. "Legendrian knots and constructible sheaves". In: *Inventiones mathematicae* 207.3 (2014).

[31] Vivek Shende et al. "Cluster Varieties from Legendrian Knots". In: *arxiv:1512.08942* (2015).

[32] Nicolá Sibilla, David Treumann, and Eric Zaslow. "Ribbon Graphs and Mirror Symmetry". In: *Selecta Mathematica, New Series* (2014).

[33] Laura Starkston. "Arboreal Singularities in Weinstein Skeleta". In: *arXiv:1707.03446* (2017).

[34] Bertrand Toën. "The homotopy theory of dg-categories and derived Morita theory". In: *arXiv:math/0408337* (2004).

[35] M.L. Wachs. "Poset topology: tools and applications". In: *Geometric Combinatorics, IAS/PCMI lecture notes series* 13 (2007).