



Supersymmetric Spectroscopy

Citation

Cordova, Clay Alexander. 2012. Supersymmetric Spectroscopy. Doctoral dissertation, Harvard University.

Permanent link

<http://nrs.harvard.edu/urn-3:HUL.InstRepos:9414575>

Terms of Use

This article was downloaded from Harvard University's DASH repository, and is made available under the terms and conditions applicable to Other Posted Material, as set forth at <http://nrs.harvard.edu/urn-3:HUL.InstRepos:dash.current.terms-of-use#LAA>

Share Your Story

The Harvard community has made this article openly available.
Please share how this access benefits you. [Submit a story](#).

[Accessibility](#)

©2012 - Clay Alexander Córdova

All rights reserved.

Thesis advisor
Cumrun Vafa

Author
Clay Alexander Córdova

Supersymmetric Spectroscopy

Abstract

We explore supersymmetric quantum field theories in three and four dimensions via an analysis of their BPS spectrum.

In four dimensions, we develop the theory of BPS quivers which provides a simple picture of BPS states in terms of a set of building block atomic particles, and basic quantum mechanical interactions. We develop efficient techniques, rooted in an understanding of quantum-mechanical dualities, for determining the spectrum of bound states, and apply these techniques to calculate the spectrum in a wide class of field theories including *ADE* gauge theories with matter, and Argyres-Douglas type theories.

Next, we explore the geometric content of quivers in the case when the four-dimensional field theory can be constructed from the six-dimensional $(2, 0)$ superconformal field theory compactified on a Riemann surface. We find that the quiver and its superpotential are determined by an ideal triangulation of the associated Riemann surface. The significance of this triangulation is that it encodes the data of geodesics on the surface which in turn are the geometric realization of supersymmetric particles.

Finally we describe a class of three-dimensional theories which are realized as supersymmetric domain walls in the previously studied four-dimensional theories. This leads to an understanding of quantum field theories constructed from the six-dimensional $(2, 0)$ superconformal field theory compactified on a three-manifold, and we develop the associated geometric dictionary. We find that the structure of the field theory is determined by a decomposition of the three-manifold

into tetrahedra and a braid which specifies the relationship between ultraviolet and infrared geometries. The phenomenon of BPS wall-crossing in four dimensions is then seen in these domain walls to be responsible for three-dimensional mirror symmetries.

Contents

Title Page	i
Abstract	iii
Table of Contents	v
Citations to Previously Published Work	vi
Acknowledgments	vii
Dedication	viii
1 Introduction and Summary	1
2 Quivers of $\mathcal{N} = 2$ QFTs	8
2.1 Introduction	8
2.2 BPS Quiver Quantum Mechanics	11
2.3 Quiver Mutation and Duality	29
2.4 $SU(2)$ Gauge Theories	51
2.5 $SU(N)$ Gauge Theories and Beyond	79
3 Quivers and Riemann Surfaces	102
3.1 Introduction	102
3.2 BPS Quivers of Complete Theories	103
3.3 Exceptional Complete Theories	128
4 Braids Walls and Mirrors	133
4.1 Introduction	133
4.2 Five-Branes on Three-Manifolds	138
4.3 R-flow, Domain Walls and a 4d-3d Link	170
4.4 4d BPS States of A_n Theories	184
4.5 Tetrahedra and Braids	191
4.6 Flows of General 4d $\mathcal{N} = 2$ Theories	244
A Self-Folded Triangles	253
Bibliography	256

Citations to Previously Published Work

Chapter 2 has appeared in the following paper:

“N=2 Quantum Field Theories and Their BPS Quivers,” M. Alim, S. Cecotti, C. Cordova, S. Espahbodi, A. Rastogi, C. Vafa. [hep-th/1112.3984](https://arxiv.org/abs/hep-th/1112.3984).

Chapter 3 has appeared in the following paper:

“BPS Quivers and Spectra of Complete N=2 Quantum Field Theories,” M. Alim, S. Cecotti, C. Cordova, S. Espahbodi, A. Rastogi, C. Vafa. [hep-th/1109.4941](https://arxiv.org/abs/hep-th/1109.4941).

Chapter 4 has appeared in the following paper:

“Braids, Walls, and Mirrors,” S. Cecotti, C. Cordova, C. Vafa. [hep-th/1110.2115](https://arxiv.org/abs/hep-th/1110.2115).

Citations of the form [hep-th/XXXX.XXXX](https://arxiv.org/abs/hep-th/XXXX.XXXX) can be found on the preprint server www.arxiv.org.

Acknowledgments

Completing this doctoral work has been a remarkable undertaking. Over the last five years I have experienced pessimism and discouragement in the face of challenging concepts and problems, and the optimism and excitement which comes through breakthroughs and understanding. Throughout, I have been sustained by a number of important colleagues, mentors, and family.

I have been lucky to have Cumrun Vafa as my advisor. Working closely with him over the last few years has been an incredible intellectual experience. Cumrun has a remarkable ability to marshal enthusiasm and to foster hard work. Working with him, one has the feeling that no problem is insurmountable. From him, I have begun to learn some of the more difficult and subtle aspects of science: how to give a good talk, how to be an effective group leader, how to ask the right questions, how to intuit the right answers.

I have also had the privilege of working on a number of projects with excellent collaborators. I would like to thank Murad Alim, Ashwin Rastogi, and Sam Espahbodi, and Jonathan Heckman for having energy when I had none. I would especially like to single out Sergio Cecotti as an outstanding coworker who kept me sharp and focused.

Throughout graduate school I have learned a tremendous amount from conversations with many people in the Harvard group. My peer group of graduate students have been unfailing comrades. And postdocs and professors including Andy Neitzke, Daniel Jafferis, and Xi Yin have been kind to share their time and insights. In particular I must single out Frederik Denef for guiding me through my adolescence of graduate school and always being there to listen to my ideas and confusions, and David Simmons-Duffin for helping me learn quantum field theory and for being a good friend.

Finally, I would like to thank my family for their constant support and encouragement. Chad, Jessy, Janet, Marsha, and Mitch have each in their own way encouraged me to persevere and to excel and for that I am eternally grateful. Above all, I thank my wife Allie for holding my hand every step of the way.

*Dedicated to my mother Janet,
who taught me that through determination
one may overcome any obstacle.*

Chapter 1

Introduction and Summary

One of the overarching and motivating questions for this thesis is to develop useful tools for understanding non-perturbative aspects of quantum field theories. The tools we will discuss apply to quantum field theories in three and four dimensions which have extended supersymmetries. In four dimensions, these extra symmetries make it possible to single out a subclass of distinguished particles in the theory, the supersymmetric states. The defining property of these states is that they are the lightest possible charged particles. More specifically, we will frequently be considering field theories whose low-energy description is that of an abelian gauge theory. The extended supersymmetry of the field theory allows us to prove that for each charge γ , there a lower bound on the mass of all particles carrying that charge

$$M_\gamma \geq |\mathcal{Z}(\gamma)|. \tag{1.0.1}$$

Where in the above, the quantity $\mathcal{Z}(\gamma)$ is a certain complex number, the central charge. The supersymmetric (BPS) states are those that saturate the bound. In a generic quantum field theory one may expect that massive charged particles which do not saturate the inequality (1.0.1) may decay to photons and lighter BPS states. When this is so, the BPS spectrum can be understood as simply the set of stable particles. As such, a detailed understanding of the BPS spectrum is first step

towards understanding the dynamics of any quantum field theory with extended supersymmetry.

Despite the obvious importance the BPS spectrum, until recently explicit computations of these states were not possible beyond the most elementary examples. In essence, part of the difficulty is that even for simple cases the BPS spectrum is often infinite and may be difficult to enumerate explicitly. What's more, the theories in question may often depend on additional parameters, moduli, and the spectrum is in general discontinuous as a function of these parameters, meaning simply that a particle which is supersymmetric and hence stable for a given value of the parameters may be unstable and decay for a different value of the parameters. A key breakthrough in this problem was a remarkable piece of technology, a so-called wall-crossing formula [1] discovered by mathematicians, which provides a general answer to the question of the discontinuities of the spectrum. This work directly inspired a number of results in physics interpreting the mathematics [2–5], and rekindled the hope that a systematic understanding of BPS states was possible.

The remaining difficulty in determining the supersymmetric spectrum of a given field theory is then to develop a useful framework where the states can be understood systematically. Such a framework is provided by *BPS quivers* [6–8] studied in detail in chapter 2 of this thesis. In physical terms, a BPS quiver is a description of the spectrum as quantum mechanical bound states of a small number of *atomic* BPS particles. The atoms and their interactions are conveniently encoded in graph where the atoms are indicated by nodes, and the interactions are indicated by arrows. Thus, a typical diagram which we will study takes the form given below:

$$\textcircled{1} \longrightarrow \textcircled{2} \tag{1.0.2}$$

In fact, this example is the BPS quiver which governs the spectrum of pure $SU(2)$ Yang-Mills [9,10] and as such is a paradigmatic example in the subject. The two nodes, indicate the two atoms, the so called magnetic monopole and dyon which are stable for all values of the parameters of the theory. The arrows specify interactions between these states, and the fact that there are two is an

indication of the strength of the the force. We will see that depending on the values of parameters of the theory, these same two atomic particles may support different numbers of bound states. As a result the wall-crossing formulas mentioned in the previous paragraph are automatically built in to the quiver description.

When a quantum field theory admits a quiver description of its spectrum, a great simplification is achieved. One has reduced the problem of calculating the states, for which previously there was no known method, to a definite problem in quantum mechanics. Even better, the entire spectrum which previously may have been quite intractable is reduced to primary constituents. This makes an understanding of the spectrum possible. Motivated by these ideas, in chapter 2 we accomplish the task of enumerating the BPS quivers for a huge variety of quantum field theories. We find that many basic operations in field theory such as adding matter, gauging flavor symmetries, and taking decoupling limits translate into simple graphical operations on the quiver. Thus, the quiver in many ways serves as an effective replacement for the Lagrangian description of the theory and indeed in many cases a BPS quiver description is available even when no Lagrangian is known.

One of the most significant features to emerge from our detailed analysis is the recognition that the set of atomic BPS particles governing the spectrum of a theory is in general not unique. This simple fact leads immediately to the idea that one may develop distinct inequivalent quantum mechanical descriptions of the BPS spectrum. These quantum mechanics problems are dual to one another. That is, they are mathematically distinct descriptions of the same physics. By understanding and exploiting these dualities we are able to determine an efficient algorithm for computing the states of the theory and calculate the full exact spectrum in a variety of physically interesting examples including all *ADE* gauge theories with fundamental matter.

Throughout chapter 2 our discussion of BPS quivers is rooted in four-dimensional physics, and we strive to emphasize the fact that the properties of a quiver are a simply a graphical way

of encoding a quantum field theory. While a tremendous amount of progress can be made in the subject from this perspective there remains a fundamental question which cannot be answered: *Why does the quiver exist at all?* This question may be tackled in the context of string theory and brings us to the subject of chapter 3. Here we study quantum field theories which arise from decoupling limits of type IIB string theory on non-compact Calabi-Yau manifolds, or equivalently, four-dimensional field theories that arise from the $(2,0)$ six-dimensional superconformal theory of M5-branes compactified on Riemann surfaces. As a general rule, whenever a supersymmetric quantum field theory is constructed from a higher dimensional theory via compactification, many of its properties observable at low-energies are encoded in the geometry of the compactification manifold. In our context, this principle is born out as follows: the BPS states of the theory are in one-to-one correspondence with a class of supersymmetric branes, which are in turn described by geodesics, i.e. curves of shortest length on the Riemann surface itself [11].

Given the above, in chapter 3 our primary aim is to determine how the geodesics on a Riemann surface are related to quivers and quantum mechanics. Building on previous work [12, 13] we are led to the observation that the combinatorial data of a geodesic can be encoded from its intersections with a triangulation of the surface. Such triangulations are naturally related to quivers. To give a concrete example, we consider the case of $SU(2)$ Yang-Mills theory whose quiver was given in equation (1.0.2). The surface in question is an annulus and the geometry is illustrated in Figure 1.1. This example illustrates the general features of the correspondence: interior lines of a triangulation, i.e. those not on the boundary of the surface, are nodes of the quiver, and when these lines share a triangle an arrow is produced.

The relationship between triangulations and quivers allows us to answer the basic question, of why in the context of these examples, a quiver description of the BPS spectrum is possible: BPS states are geodesics which are encoded in triangulations and hence quivers. In particular, this allows us to find a direct geometric analog of the quantum mechanical duality discussed in chapter




Figure 1.1: BPS trajectories and triangulations for $SU(2)$ Yang-Mills. In diagram (a), the associated triangulation of the annulus. Interior red dots denote points where geodesics may terminate. The two interior lines of the triangulation yield the two nodes of the quiver. The fact that these two interior lines share two triangles leads to the two arrows in the quiver. In (b) a sample BPS geodesic is shown in blue. This geodesic can be interpreted as a bound state of three particles of type 1 and two of type 2.

2: it is simply the fact that a given Riemann surface admits many triangulations each of which is capable of encoding the data of the BPS geodesics.

In the final chapter of this thesis we turn our attention to the study of three-dimensional quantum field theories with $\mathcal{N} = 2$ supersymmetry. These are naturally related to the four-dimensional theories studied in chapters 2 and 3 because the three-dimensional theories can often be thought of as trapped degrees of freedom living on a domain wall in a four-dimensional universe. In fact the examples we study are primarily of this type and this allows us to import directly the technology from our previous discussion.

We begin chapter 4 by focusing our attention on the class of three-dimensional field theories

that arise as domain walls in the models studied in chapter 3. This means that the field theories in question can be constructed by starting from the $(2, 0)$ six-dimensional superconformal theory of M5-branes and considering it on the geometry of a Riemann surface Σ times a real line \mathbb{R} where the parameters of the field theory, encoded by the geometry of Σ vary along the line \mathbb{R} . Thus the total geometry of the M5-brane is given by a three-dimensional Minkowski space, where the effective three-dimensional low-energy physics is observed, times a non-trivial three-manifold. We explain how properties of the theory are determined by the three-manifold geometry, for instance the effective gauge group and Chern-Simons levels describing the low-energy dynamics are encoded in homology. A primary role is again played by the BPS states, again described by a kind of geodesic on the three-manifold.

One of the most interesting results of this chapter concerns the phenomenon of three-dimensional mirror symmetry [14, 15]. In elementary terms this is a statement that two distinct quantum field theories may ultimately describe the same low-energy dynamics. A particularly useful example is supersymmetric quantum electrodynamics and the so-called XYZ model, a theory of three scalars with cubic interactions. At general energy scales these two field theories are distinct, however at extremely low energies they describe the same physics. In our context we make contact with this phenomenon by noting that the spectrum of particles trapped on the wall are inherited from the ambient four-dimensional universe whose spectrum is given by the BPS states. Thus, the BPS spectrum of the four-dimensional theory allows us to deduce the physics on the wall. Moreover the wall-crossing phenomenon, which related distinct spectra by variation of parameters in four-dimensions, shows that there are distinct three-dimensional descriptions of the same low-energy physics. In other words, wall-crossing in four-dimensions *implies* the existence of mirror symmetries in three dimensions.

We relate this mirror symmetry to the three-dimensional geometry of the M5-brane by introducing a combinatorial decomposition of the three-manifold into tetrahedra, analogous to

the decomposition of the Riemann surface into triangles used in chapter 3. In this case, each tetrahedron encodes the existence of a particle, and different ways of splitting the manifold into tetrahedra describe different mirror descriptions of the same low energy physics [16]. For example, the geometry relevant to the mirror pair described in the previous paragraph is illustrated in Figure 1.2.

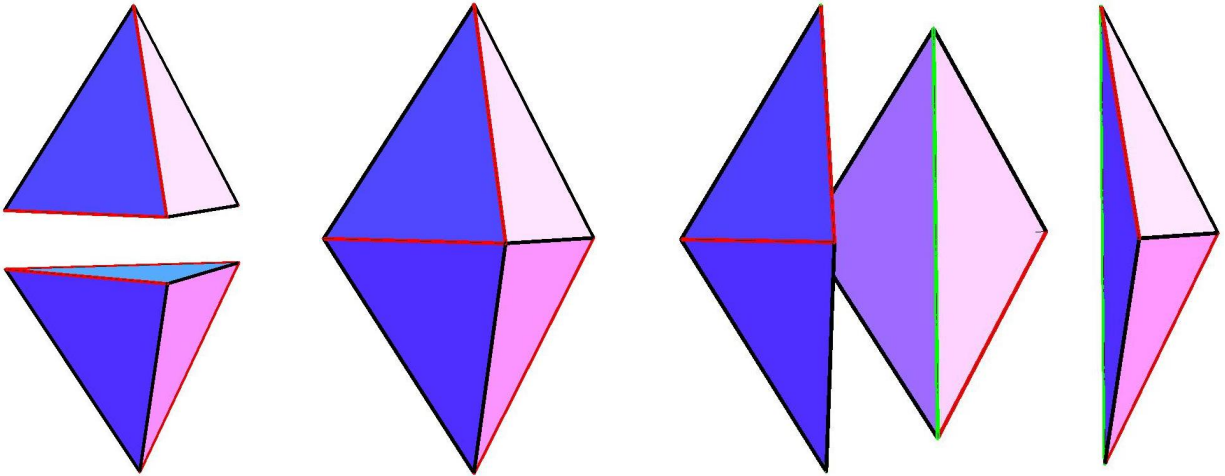


Figure 1.2: The geometry of the mirror symmetry between the XYZ model and $U(1)$ QED. In the center we see the bipiramid, the compactification manifold for the M5 brane. On the left, it is decomposed into two tetrahedra glued along a face. This describes a theory of electrodynamics with the two tetrahedra in correspondence with the electron and positron. On the right the same bipiramid, is decomposed into three tetrahedra glued along the green edge. The three tetrahedra are the X , Y , and Z particles of the XYZ model.

Finally, we conclude this dissertation by noting that the relationship between wall-crossing and mirror symmetry suggests the existence of fascinating dualities between exotic three-dimensional theories, potentially with infinitely many species of particles, yet to be discovered.

Chapter 2

Quivers of $\mathcal{N} = 2$ QFTs

2.1 Introduction

In the study of four-dimensional quantum field theories with extended supersymmetry, one of the most fruitful and enduring ideas has been the analysis of the spectrum of BPS particles. An understanding of this protected sector of the Hilbert space is often a key ingredient in testing field theory, and stringy dualities and played an important role in the foundational low-energy solution of $\mathcal{N} = 2$ gauge theories [9, 10]. More recently, significant progress has been made both in mathematics and in physics, in understanding the universal rules that govern potential decay processes of BPS particles [1, 2, 17–19], and continuing progress in this subject [3–5, 12, 13, 20–29] suggests that there are yet undiscovered structures lurking in the BPS spectra of field theories. However in spite of these dramatic developments, there exists no general method for calculating the BPS spectrum of a given field theory.

In this work we study a wide class of field theories where this difficulty is overcome. These are theories, whose spectra of BPS states can be calculated from the quantum mechanics of an associated BPS quiver. Such quivers originally arose in string theory constructions of quantum field theories [6–8, 30–33]. In that context, there is a natural class of BPS objects, namely D-branes,

and a quantum mechanical description of the BPS spectrum is provided by the worldvolume theory of the relevant branes. This string theory setup provides a simple way of organizing the spectrum into elementary BPS branes and their bound states and explains the non-abelian degrees of freedom needed in the quiver description.

While the geometric engineering perspective provides a useful source of examples, our focus in this paper is on analyzing the theory of BPS quiver directly from the point of view of quantum field theory. The class of BPS quiver theories is broad, and includes gauge theories coupled to massive hypermultiplets, Argyres-Douglas type field theories [34], and all theories defined by M5 branes on punctured Riemann surfaces [12, 28, 35–40].¹ For all of these theories the quiver appears to provide a simple and unique characterization of the theory, and one of the aims of this work is to illustrate in a variety of examples how simple graphical features and operations at the level of quivers translate into physical properties and constructions such as flavor symmetries, gauging, decoupling limits, and dualities.

To accomplish our task of exploring BPS quivers, we begin in section 2.2 with a detailed description of the way in which quiver quantum mechanics encodes the spectrum of BPS states. We then develop the theory of quiver representations, the holomorphic description of quiver quantum mechanics and explain how quivers yield a concrete method for studying wall-crossing phenomenon and review basic examples of these techniques. This material is known and is included for completeness and to provide context for subsequent extensions.

A significant feature of quiver description of the spectrum is that a fixed quiver typically describes the BPS particles only on a small patch of the moduli space of a given theory. A key role is then played by quantum mechanical dualities, encoded by quiver mutations, which relate distinct quivers valid in different regions of parameter space. These relationships between a priori distinct quantum mechanics are a one-dimensional version of Seiberg duality [41]. Their basic content is

¹At least one puncture is required. The punctureless case provides examples of theories without BPS quivers [12].

that the BPS spectrum can be decomposed into bound states of primitive particles in more than one way by suitable changes of the set of building block BPS states. In section 3 we discuss these dualities and analyze the constraints that they impose upon the BPS spectrum. Remarkably we find that these consistency conditions are so powerful that frequently they completely determine the BPS spectrum. This results in an algorithm, the *mutation method*, for calculating a spectrum that is far simpler than a direct investigation of the quantum mechanics.

In sections 4-5 we put the general theory to use by computing the BPS spectrum in a broad class of examples. We focus our attention on non-abelian gauge theories with ADE gauge group and fundamental matter. For all such theories we determine the quiver and frequently our mutation method is powerful enough to determine the BPS spectrum in a strongly coupled chamber where there are only finitely many BPS states. The spectrum in all chambers can then be deduced by the application of the wall crossing formula of Kontsevich and Soibelman [1]. Let \mathcal{B} denote the set of BPS particles at strong coupling, and $|\mathcal{B}|$ the number of such particles. Then a summary of the gauge theories whose strong coupling spectra we determine is:

- $SU(N_c)$ gauge theory coupled to N_f fundamentals.

$$|\mathcal{B}| = N_c(N_c - 1) + N_f(2N_c - 1)$$

- $SO(2N_c)$ gauge theory coupled to N_v vectors.

$$|\mathcal{B}| = 2N_c(N_c - 1) + N_v(4N_c + 1)$$

- E_6 gauge theory coupled to $N_{\mathbf{27}}$ fundamental $\mathbf{27}$'s.

$$|\mathcal{B}| = 72 + 73N_{\mathbf{27}}$$

- E_7 gauge theory.

$$|\mathcal{B}| = 126$$

- E_8 gauge theory.

$$|\mathcal{B}| = 240$$

One elegant feature of the above results can be seen in the limit where there is no matter whatsoever so that one is considering the strong coupling BPS spectrum of pure super-Yang-Mills with an arbitrary ADE gauge group. Then our results can be summarized by noting that the number of BPS particles is given simply by the number of roots in the associated Lie algebra.

2.2 BPS Quiver Quantum Mechanics

We begin with a four-dimensional $\mathcal{N} = 2$ field theory with Coulomb moduli space \mathcal{U} . Here by a point $u \in \mathcal{U}$ we will mean a specification of all supersymmetric parameters in the theory including Coulomb branch moduli, bare masses, and coupling constants. At a generic value of the moduli $u \in \mathcal{U}$, this field theory has a $U(1)^r$ gauge symmetry, and a low energy solution described by:

- A lattice Γ of electric, magnetic, and flavor charges of rank $2r + f$, where f is the rank of the flavor symmetry.
- A linear function $\mathcal{Z}_u : \Gamma \rightarrow \mathbb{C}$, the central charge function of the theory.² Central charges which couple to the electric and magnetic charges encode the effective coupling and theta angle of the infrared physics, while the central charges that couple to the flavor symmetries sample possible bare masses of matter in the theory.

The behavior of the central charge function as one varies the moduli fixes completely the effective action for the neutral massless fields. However, the description of the massive charged

²Here we explicitly indicate the u dependence by including a subscript on the central charge function. For notational simplicity, we will eventually drop the subscript and leave the u dependences implicit

particles is more subtle. According to the $\mathcal{N} = 2$ superalgebra, the central charge provides a lower bound on the masses of charged particles. The mass of a particle with charge $\gamma \in \Gamma$ satisfies

$$M \geq |\mathcal{Z}_u(\gamma)|. \quad (2.2.1)$$

The lightest charged particles are those that saturate the above bound - these are termed BPS. The spectrum of BPS states is a priori undetermined by the low energy solution of the theory alone, and it is precisely this question that we aim to address. We will describe a class of theories where the BPS spectrum can be computed and studied using the technology of quiver quantum mechanics.

2.2.1 Quivers and Spectra

In this section we lay the foundations for our ideas by describing the connection between quantum mechanical quiver theories and BPS spectra of four-dimensional quantum field theories. In the course of our analysis we will also discover various restrictions on the class of theories to which these quiver techniques apply. We first describe in section 2.2.1 how the BPS spectrum of the 4d theory at a fixed point in moduli space can frequently be used to define an associated quiver, and therefore to pose a supersymmetric quantum mechanics problem. We will then see in section 2.2.1 that the ground states of this supersymmetric quantum mechanics precisely reproduce the BPS spectrum. From this point of view, the quiver provides merely a clever way of organizing the BPS spectrum. However, the true power of the technique is that there exist many ways of producing a BPS quiver that *do not* assume a knowledge of the spectrum. These are briefly surveyed in section 2.2.1. It is through these methods that we can hope in turn to discover previously unknown spectra.

From BPS Spectra to BPS Quivers

Let us begin by fixing a point $u \in \mathcal{U}$ in moduli space. Suppose the occupancy of BPS states here is known. We will then explain how to construct a quiver that describes the theory at this point u .

To begin we split the BPS spectrum into two sets, the *particles* and the *antiparticles*. We define particles to be those BPS states whose central charges lie in the upper half of the complex \mathcal{Z} plane, and antiparticles those in the lower. CPT invariance ensures that for each BPS particle of charge γ , there is an antiparticle of charge $-\gamma$. Thus the full BPS spectrum consists of the set of BPS particles plus their associated CPT conjugate antiparticles. We will use the occupancy of the particles to construct a quiver.

Among the particles, we choose a minimal basis set of hypermultiplets. Since the lattice Γ has rank $2r + f$, our basis will consist of $2r + f$ BPS hypermultiplets. Let us label their charges γ_i . The particles in the basis set should be thought of as the elementary building blocks of the entire spectrum of BPS states. As such they are required to form a positive integral basis for all occupied BPS particles in the lattice Γ . This means that every charge γ which supports a BPS particle satisfies

$$\gamma = \sum_{i=1}^{2r+f} n_i \gamma_i. \quad n_i \in \mathbb{Z}^+ \quad (2.2.2)$$

We emphasize that the basis need not span Γ , but only the subset of occupied states in Γ . We will see in section 2.2.1 that this equation can be interpreted as saying that the BPS particle with charge γ can be viewed as a composite object built up from a set of elementary BPS states containing n_i particles of charge γ_i .

It is important to notice that the requirement that a set of states form a positive integral basis for the entire spectrum of BPS particles is quite strong, and in particular uniquely fixes a basis when it exists. To see this, we suppose that $\{\gamma_i\}$ and $\{\tilde{\gamma}_i\}$ are two distinct bases. Then there is a matrix n_{ij} relating them

$$\tilde{\gamma}_i = n_{ij} \gamma_j; \quad \gamma_i = (n^{-1})_{ij} \tilde{\gamma}_j. \quad (2.2.3)$$

However since both $\{\gamma_i\}$ and $\{\tilde{\gamma}_i\}$ form positive integral bases, the matrix n_{ij} and its inverse must have positive integral entries. It is easy to see that this forces both matrices to be permutations.

Thus the two bases can differ only by a trivial relabeling.

Now, given the basis of hypermultiplets $\{\gamma_i\}$ there is a natural diagram, a *quiver*, which encodes it. This quiver is constructed as follows:

- For each element γ_i in the basis, draw a node of the quiver.
- For each pair of charges in the basis compute the electric-magnetic inner product $\gamma_i \circ \gamma_j$. If $\gamma_i \circ \gamma_j > 0$, connect corresponding nodes γ_i and γ_j with $\gamma_i \circ \gamma_j$ arrows, each of which points from node j to node i .

To illustrate this construction, we consider the simple case of pure $SU(2)$ gauge theory at a large value of the Coulomb branch modulus, where the theory is governed by semiclassical physics. In terms of their associated electric and magnetic charges (e, m) , the occupied BPS states consist of:

$$\text{Vector multiplet } W - \text{boson} : (2, 0), \tag{2.2.4}$$

$$\text{Hypermultiplet dyons} : (2n, 1), (2n + 2, -1) \quad n \geq 0.$$

Choosing the particle half-plane represented in Fig. 2.1a, the unique basis is given by the monopole $(0, 1)$ and the dyon $(2, -1)$. The spectrum and the resulting quiver are then shown in Figure 2.1.

So, returning to the general story, we have given a map from BPS spectra to quivers. At this stage, we pause to point out important subtleties in this procedure. The first is that our identification of arrow being determined by the Dirac inner product glosses over the possibility of having arrows between nodes which point in opposite directions. In fact, what the Dirac product truly captures is the net number of arrows. It is a fortunate feature of all of the field theory examples discussed in this work, with the exception of section 6.2, the electric magnetic inner product accurately determines the arrows in the quiver. Further analysis of this issue occurs in our discussion of superpotentials in section 3.

A second important subtlety is that there exist field theories for which there is no BPS quiver whatsoever. To illustrate this, note that one assumption thus far was that we could find a

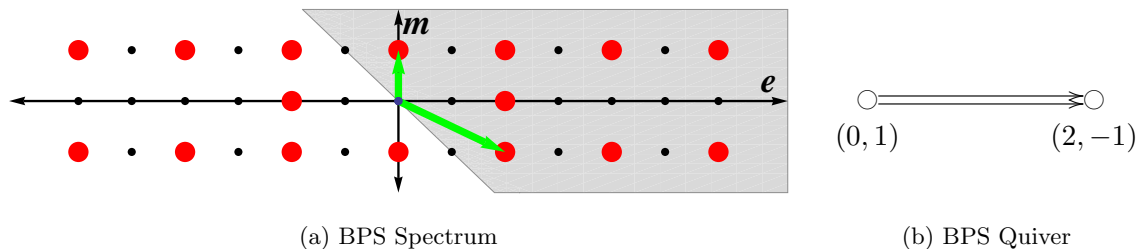


Figure 2.1: The spectrum and BPS quiver of $SU(2)$ Yang-Mills. In (a) the weak-coupling BPS spectrum, both particles and antiparticles, is plotted in the (e, m) plane. Red dots denote the lattice sites occupied by BPS states. The green arrows show the basis of particles given by the monopole and dyon. We have represented our choice of particle central charge half-plane by the grey region. In (b) the BPS quiver is extracted from this data. It has one node for each basis vector, and the double arrow encodes the symplectic product.

basis of hypermultiplets in the upper half of the central charge plane. By linearity of the central charge function, this gives a constraint on the occupied subset of Γ . In particular, since the set $\{\gamma_i\}$ forms a basis, we have for an arbitrary BPS particle of charge γ ,

$$\gamma = \sum_i n_i \gamma_i \implies \mathcal{Z}_u(\gamma) = \sum_i n_i \mathcal{Z}_u(\gamma_i). \quad n_i \geq 0 \quad (2.2.5)$$

Since $\mathcal{Z}(\gamma_i)$ all lie in the upper half-plane, (2.2.5) implies that the central charges of all BPS particles lie in a cone in the upper half of the central charge plane, bounded by the left-most and right-most $\mathcal{Z}(\gamma_i)$; we denote this the *cone of particles*.

One can see that many theories do not even have such a cone, and therefore don't have an associated BPS quiver. The simplest example is $\mathcal{N} = 4$ Yang-Mills with gauge group $SU(2)$. Because of S-duality, this theory has a spectrum of dyons with charges (p, q) , for p and q arbitrary coprime integers. It follows that the phases of the central charges of these dyons form a dense set in the unit circle in the central charge plane. In particular, there is no cone of particles and hence no quiver.

We can state the problem with $\mathcal{N} = 4$ Yang-Mills from the $\mathcal{N} = 2$ perspective: there is an adjoint hypermultiplet which is forced to be massless. The $\mathcal{N} = 2^*$ theory, where the adjoint is given a mass, *does* admit a BPS quiver, given in section 2.4. This situation is typical of gauge theories that become conformal when all mass deformations are turned off. A conformal field theory has no single particle states at all, let alone BPS states. A quiver description is therefore only possible when sufficiently many massive deformations of the theory exist and have been activated.

Alternative Constructions of BPS Quivers

Thus far we have explained how BPS quivers provide a way of describing certain properties of the basis for the BPS spectates at a fixed point in moduli. In the next section, we explain the reverse construction, that is, how to extract a BPS spectrum from a BPS quiver, and hence how a BPS quiver can be used as a convenient way for encoding the complete BPS spectrum. However, the most important application of BPS quivers is that they can be used to deduce an *unknown* BPS spectrum. One reason this is so, is that our construction of BPS quivers is completely local in the Coulomb branch moduli space \mathcal{U} . Given a point $u \in \mathcal{U}$ where the BPS spectrum is known, the quiver description of the spectrum is uniquely fixed if it exists. But, as will be clear by the conclusion of section 3, once a quiver is determined for a single modulus u , the quiver description of the entire moduli space \mathcal{U} is also fixed. Thus, we may determine the quiver in say a region of weak coupling where the physics is under control, and then use it to calculate the BPS spectrum at strong coupling.

Even more striking is the fact that BPS quivers can frequently be deduced by alternative geometric methods in various contexts in string theory, even when the BPS spectrum is unknown for any value of the moduli. The quiver methods described in the following sections can then be used to determine the spectrum from scratch.

The existing literature on the techniques used to extract BPS quivers is by now very vast,

in the following we outline some of the various interrelated approaches:³

- Building on the original orbifold construction of quiver gauge theories of [6] refs. [7, 8, 30, 43] provided the identification of the quiver nodes with a basis of BPS states obtained from fractional branes, these BPS quivers were further explored in [31, 32].
- The relation of the $4d$ quivers with the soliton spectrum in $2d$ [44] was studied in various places, see for example [45–47], more recently this $2d/4d$ correspondence and the associated construction of BPS quivers was discussed in [4].
- The toric methods of [48, 49] and the relation to dimer models [50] were used in [51] to construct a large class of quivers, their construction using mirror symmetry was studied in [52].
- Based on the geometric study of BPS states in SW theories pioneered in [11] and further studied in [13, 53], the BPS quivers can be obtained from triangulations of Riemann surfaces as described in [12, 28] using the relation of triangulations and quivers of [54]. Given a pair of M5-branes wrapping a Riemann surface \mathcal{C} , an ideal triangulation of \mathcal{C} can be used to determine the BPS quiver. We explore this idea in chapter 3 of this thesis.

Quiver Quantum Mechanics: From BPS Quivers to BPS Spectra

We now return to our general discussion of BPS quivers and explain how to deduce the full spectrum from the quiver. Thus far the BPS quiver we have introduced is merely a way of encoding a basis of BPS states $\{\gamma_i\}$ for a given $\mathcal{N} = 2$ theory. To construct a general BPS state, we must know, for a given charge

$$\gamma = \sum_i n_i \gamma_i \tag{2.2.6}$$

³See also [42] and references therein for an excellent recent exposition of the mathematical structures used to describe D-branes which includes in particular the associated quiver representation theory.

whether any particles of this charge exist in the theory, and if so, determine their degeneracy and spins. We attack this question by viewing the hypothetical state with charge γ as a quantum mechanical bound state of n_i copies of each basis particle γ_i . Since we seek a BPS particle, we introduce a four supercharge quantum mechanics problem and look for its supersymmetric ground states. The precise quantum mechanics theory is constructed from the BPS quiver and the charge γ in the following way: Let i index nodes of the quiver, and a index the arrows of the quiver. Then we introduce a gauge group for each node and bifundamental field B_{ij}^a for each arrow pointing $i \rightarrow j$,

$$\text{Gauge Group} = \prod_{\text{nodes } i} U(n_i), \quad \text{Matter} = \bigoplus_{\text{arrows } a} B_{ij}^a. \quad (2.2.7)$$

Thus, the BPS quiver, whose nodes and arrows were originally merely a presentation of a basis of hypermultiplets, now encodes the gauge groups and bifundamental matter of a quiver quantum mechanics.

This prescription can be motivated most easily when the four-dimensional field theory is engineered in string theory. In such a situation, BPS states are viewed as various supersymmetric bound states of D-branes. Then the nodes of our quiver correspond to a collection of basic supersymmetric branes and the arrows are bifundamental fields that arise at brane intersections. This also provides an elementary understanding of the appearance of *non-abelian* gauge fields in the quantum mechanics: they are the usual non-abelian degrees of freedom that arise when branes coincide. The quantum mechanics problem introduced above is then nothing but the worldvolume theory of a system of D-branes dimensionally reduced to 0+1 dimensions.

Returning to our general analysis, to assess the existence of a BPS particle with charge γ , we look for supersymmetric ground states on the Higgs branch of this quiver theory. These depend on two data which we must still specify:

- Fayet-Iliopoulos Terms

Since the gauge groups at each node are given by $U(n_i)$, the overall $U(1)$ at each node can couple to an independent FI-term θ_i . These parameters are fixed by the central charges $\mathcal{Z}_u(\gamma_i)$ of the constituent particles. We state this identification in the case that all the central charges point in nearly the same direction in the complex plane. Then let $\mathcal{Z}_u(\gamma)$ denote the central charge of a state with charge γ , and set

$$\theta_i = |\mathcal{Z}_u(\gamma_i)| \left(\arg(\mathcal{Z}_u(\gamma_i)) - \arg(\mathcal{Z}_u(\gamma)) \right). \quad (2.2.8)$$

For each node i in the quiver there is then a D-term equation of motion

$$\sum_{\substack{\text{arrows} \\ \text{starting at } i}} |B_{ij}^a|^2 - \sum_{\substack{\text{arrows} \\ \text{ending at } i}} |B_{ki}^a|^2 = \theta_i. \quad (2.2.9)$$

When the central charges are not nearly aligned, the identification of the FI parameters is more involved, and for now the reader should assume that the moduli are such that this approximation is valid.⁴ Later in section 2.2.2 we will see an elegant way of rephrasing our problem that completely avoids this issue.

- Superpotentials

Whenever there are non-trivial oriented cycles in the BPS quiver, the quantum mechanics theory admits a non-trivial gauge invariant superpotential \mathcal{W} which is a holomorphic function of the bifundamental fields. Our procedure for producing a quiver does not fix a superpotential; it is an independent datum of our construction which must be computed by alternative means. Later in section 2.3 we will see general constraints on \mathcal{W} . For now, we simply assume that \mathcal{W} is given. This superpotential yields F-term equations of motion

$$\frac{\partial \mathcal{W}}{\partial B_{ij}^a} = 0. \quad (2.2.10)$$

⁴Alternatively one may tune the central charges to near alignment. Since this involves no crossing of walls of marginal stability the spectrum is stable under this motion.

Having fully fixed the quantum mechanics, we now turn to the moduli space of supersymmetric ground states with charge γ , \mathcal{M}_γ .⁵ This space is simply the solution to the equations of motion described above, quotiented by the action of the unitary gauge groups.

$$\mathcal{M}_\gamma = \left\{ B_{ij}^a \left| \frac{\partial \mathcal{W}}{\partial B_{ij}^a} = 0, \sum_{\substack{\text{arrows} \\ \text{starting at } i}} |B_{ij}^a|^2 - \sum_{\substack{\text{arrows} \\ \text{ending at } i}} |B_{ki}^a|^2 = \theta_i \right. \right\} / \prod_i U(n_i). \quad (2.2.11)$$

If \mathcal{M}_γ is non-empty, then there exists a BPS particle in the spectrum with charge γ . To determine spins and degeneracy from \mathcal{M}_γ , we examine the structure of its cohomology. Specifically, since \mathcal{M}_γ is the moduli space of a theory with four supercharges, it is a Kähler manifold, and as such its cohomology automatically forms representations of Lefschetz $SU(2)$. For each such irreducible Lefschetz $SU(2)$ representation, we obtain a supersymmetric BPS multiplet. The spacetime spin of a multiplet is then determined by tensoring the Lefschetz spin with an overall $\mathcal{N} = 2$ hypermultiplet,

$$\text{Spin} = \text{Lefschetz} \otimes \left(\left[\frac{1}{2} \right] + 2 [0] \right). \quad (2.2.12)$$

Equation (2.2.12) can be intuitively understood by thinking about the worldvolume theory of a BPS particle. This worldvolume theory supports four supercharges and hence has an R-symmetry group of $SU(2)$ which is none other than the Lefschetz $SU(2)$ of the moduli space. On the other hand, the R-symmetry group of a brane, in this case our particle, can be identified with the group of rotations transverse to the worldvolume, which in turn controls the angular momentum of the state. Thus the Lefschetz $SU(2)$ computes the orbital angular momentum of the state, and the overall shift by $1/2$ in (2.2.12) simply takes into account the intrinsic spin contribution.

In practice the most important application of (2.2.12) is to distinguish vector multiplets from hypermultiplets. The latter are associated to Lefschetz multiplets of length zero, as would naturally occur if, say, \mathcal{M}_γ were a point. Meanwhile vector multiplets are associated to Lefschetz

⁵From now on, whenever we refer to supersymmetric ground states of the quiver quantum mechanics, we will always mean on the Higgs branch. The Coulomb branch can also be studied and gives rise to equivalent results for BPS spectra. [33]

multiplets of length two, the canonical example of which is $\mathcal{M}_\gamma \cong \mathbb{P}^1$. In complete generality the formula (2.2.12) tells us that if \mathcal{M}_γ has complex dimension d then there is guaranteed to be a BPS multiplet of spin $\frac{d+1}{2}$ with charge γ in the spectrum. Naive parameter counting gives the expected dimension of the \mathcal{M}_γ as

$$d = \sum_{B_{ij}^a} (n_i n_j) - \sum_{\text{nodes } i} n_i^2 - (\# \text{ F-term constraints}) + 1. \quad (2.2.13)$$

Here we have simply counted the degrees of freedom of the bifundamental fields, B_{ij}^a , and subtracted the gauge degrees of freedom and the F-term constraints. The addition of 1 is for the overall diagonal gauge group $U(1)_d \subset \prod_i U(1) \subset \prod_i U(n_i)$. Since all fields are bifundamental, no field is charged under the simultaneous $U(1)$ rotation of all gauge groups, so this gauge degree of freedom is actually redundant.

In summary, given a quiver we have defined a supersymmetric quantum mechanics problem, and the cohomology of the moduli spaces of ground states of this quantum mechanics determines the occupancy of BPS states.

2.2.2 Quiver Representations

While our supersymmetric quantum mechanics construction determines the BPS spectra as specified by a quiver, it is useful in practice to work in the language of quiver representation theory. Here the problem of determining the ground states of the supersymmetric quantum mechanics gets recast in a holomorphic framework. Our ability to rephrase the problem in terms of quiver representation theory arises from the fact that a supersymmetric moduli space of a theory with four supercharges, such as \mathcal{M}_γ , can be presented in two ways:

- As the solution to the F-term and D-term equations of motion modulo the action of the unitary gauge groups (this is what has been stated in (2.2.11)).

- As the solution to the F-term equations modulo the action of the complexified gauge group $\prod_i Gl(n_i, \mathbb{C})$, augmented by a stability condition.

It is the second notion of \mathcal{M}_γ that makes use of quiver representation theory.

To begin, we note that in a zero energy field configuration of supersymmetric quantum mechanics, the bifundamental fields are constants and hence their expectation values can be viewed as linear maps between vector spaces \mathbb{C}^{n_i} associated to each node. These expectation values are constrained by the condition that they must solve the F-term equations of motion $\partial\mathcal{W}/\partial B_{ij}^a = 0$. A quiver representation is by definition precisely a choice of complex vector spaces \mathbb{C}^{n_i} for each node, and linear maps $B_{ij}^a : \mathbb{C}^{n_i} \rightarrow \mathbb{C}^{n_j}$ for each arrow in a quiver subject to the F-term equations. So the data of a classical zero energy field configuration completely specifies a quiver representation (See [42] and references therein).

Given a quiver representation R , defined by vector spaces \mathbb{C}^{n_i} and maps B_{ij}^a an important notion in the following will be the subrepresentations $S \subset R$. A subrepresentation S is defined by a choice of vector subspaces $\mathbb{C}^{m_i} \subset \mathbb{C}^{n_i}$ for each node and maps $b_{ij}^a : \mathbb{C}^{m_i} \rightarrow \mathbb{C}^{m_j}$ for each arrow, such that all diagrams of the following form commute:

$$\begin{array}{ccc}
 \mathbb{C}^{n_i} & \xrightarrow{B_{ij}^a} & \mathbb{C}^{n_j} \\
 \uparrow & & \uparrow \\
 \mathbb{C}^{m_i} & \xrightarrow{b_{ij}^a} & \mathbb{C}^{m_j}
 \end{array} \tag{2.2.14}$$

To complete our holomorphic description of the moduli space we must still specify a stability condition that ensures that a given quiver representation R is related to a solution of the D-term equations in quiver quantum mechanics. To motivate this, note that a quiver rep R with vector spaces \mathbb{C}^{n_i} is related to the description of a particle with charge $\gamma_R = \sum n_i \gamma_i$. Then heuristically, a subrepresentation S of R can be thought of as a bound state of smaller charge which may, in principle, form one of the constituents of a decay of a particle of charge γ_R . To prohibit such a decay, we must restrict our attention to *stable* quiver representations. To define this notion

of stability we let $\mathcal{Z}_u(R)$ denote the central charge of a representation,⁶

$$\mathcal{Z}_u(R) \equiv \mathcal{Z}_u(\gamma_R) = \sum_i n_i \mathcal{Z}_u(\gamma_i). \quad (2.2.15)$$

By construction the central charge vector lies in the cone of particles in the upper half of the central charge plane. Then R is called stable if for all subrepresentations S other than R and zero, one has

$$\arg(\mathcal{Z}_u(S)) < \arg(\mathcal{Z}_u(R)). \quad (2.2.16)$$

We will refer to any subrepresentation S that violates this condition as a destabilizing subrepresentation. This condition is denoted Π -stability, and was studied in [7]. We take this to be the requisite notion of stability at general points in moduli space. One important consistency check on this choice is that when all the central charges are nearly aligned, the stability condition (2.2.16) reduces to the D-term equations of motion presented earlier [7, 55].

Given this notion of stability, we can now formulate the moduli space \mathcal{M}_γ as set of stable quiver representations modulo the action of the complexified gauge group.

$$\mathcal{M}_\gamma = \left\{ R = \{B_{ij}^a : \mathbb{C}^{n_i} \rightarrow \mathbb{C}^{n_j}\} \left| \frac{\partial \mathcal{W}}{\partial B_{ij}^a} = 0, R \text{ is } \Pi\text{-stable} \right. \right\} / \prod_i Gl(n_i, \mathbb{C}). \quad (2.2.17)$$

This is a completely holomorphic description of \mathcal{M}_γ , and in many examples is explicitly computable.

As a very elementary application, we note that the nodes of a quiver are always Π -stable reps. That is, consider γ_j as the representation given by choosing $n_i = \delta_{ij}$. This is always stable since it has no non-trivial subrepresentations, and thus in particular no destabilizing subreps. Furthermore, since there is only one non-zero vector space, all maps must be chosen zero; thus the moduli space \mathcal{M}_{γ_j} is given by a single point. We find that each node of a quiver gives a multiplicity one hypermultiplet BPS state.

⁶When we speak of the central charge of a representation, we are always referring to the central charge of the bound state associated to that representation.

2.2.3 Walls of Marginal Stability and Examples of Quiver Representations

The preceding discussion in this section has focused exclusively on utilizing BPS quivers to encode the spectrum of an $\mathcal{N} = 2$ quantum field theory at a specific point u on the Coulomb branch \mathcal{U} . BPS states are stable under infinitesimal variations of the modulus, and thus our description can be viewed as local theory of BPS particles adequate on a patch in \mathcal{U} . Of course we are interested in determining the spectrum across the entire moduli space, and this can also be achieved using the quiver.

In the quiver representation theory problem, the moduli u along with bare mass parameters and coupling constants enter the calculation through the central charge function \mathcal{Z}_u . From the perspective of quiver representation theory, these are changes in the stability conditions. For small deformations of the stability condition, the set of stable representations, and hence the BPS spectrum, is unchanged. However at certain real codimension one loci in moduli space we encounter walls of marginal stability where a supersymmetric particle decays. At the wall, the central charges of some representation R and its subrep S become aligned. On one side of the wall, $\arg \mathcal{Z}(S) < \arg \mathcal{Z}(R)$ so that R is stable, and hence some corresponding BPS particle exists. On the other side of the wall, the phases have crossed, and the stability condition has changed. We will have $\arg \mathcal{Z}(S) > \arg \mathcal{Z}(R)$, so the representation R is no longer stable, and the associated particle has disappeared from the BPS spectrum.

It is a virtue of the description of the spectrum in terms of stable quiver representations that these wall-crossing processes are completely explicit. Indeed the BPS quiver gives us a way to calculate directly the BPS spectrum on either side of a wall. One can then simply compare the answer on both sides, and see that properties such as the Kontsevich-Soibelman wall-crossing formula hold. In this section we study these wall crossing phenomena in the context of the Argyres-Douglas conformal theories.

A_2 Theory

Let's begin with a simplest possible example which demonstrates wall-crossing. We will consider the Argyres-Douglas A_2 theory, whose quiver is given by two nodes connected by a single arrow [4]. We will denote by \mathcal{Z}_i the central charges of the two basis particles,

$$\textcircled{1} \longrightarrow \textcircled{2} \quad (2.2.18)$$

No matter what the value of the central charges, the basis particles described by the nodes of the quiver are stable. Thus the spectrum always contains at least two hypermultiplets. Now let us search for a bound state involving n_1 particles of type γ_1 and n_2 particles of type γ_2 . According to the general theory developed in the previous sections we are to study a quiver representation of the following form

$$\mathbb{C}^{n_1} \xrightarrow{B} \mathbb{C}^{n_2} \quad (2.2.19)$$

To determine stability we investigate subrepresentations. Let's start with a subrepresentation of the form

$$\begin{array}{ccc} \mathbb{C}^{n_1} & \xrightarrow{B} & \mathbb{C}^{n_2} \\ \uparrow & & \uparrow \\ 0 & \xrightarrow{0} & \mathbb{C} \end{array} \quad (2.2.20)$$

There is no condition on the field B for this diagram to commute; it is always a subrepresentation.

Thus, stability of our bound state requires

$$\arg(\mathcal{Z}_2) < \arg(n_1 \mathcal{Z}_1 + n_2 \mathcal{Z}_2) \implies \arg(\mathcal{Z}_2) < \arg(\mathcal{Z}_1). \quad (2.2.21)$$

Next we consider a similar decay involving the first basis particle

$$\begin{array}{ccc} \mathbb{C}^{n_1} & \xrightarrow{B} & \mathbb{C}^{n_2} \\ \uparrow & & \uparrow \\ \mathbb{C} & \xrightarrow{0} & 0 \end{array} \quad (2.2.22)$$

If this is a subrepresentation, then stability demands that $\arg(\mathcal{Z}_1) < \arg(\mathcal{Z}_2)$, so (2.2.21) cannot be satisfied. Thus, to ensure the existence of a bound state we must forbid this subrepresentation, and hence we must choose B so that the diagram in (2.2.22) does not commute. Thus B should have no kernel, and in particular, we have $n_1 \leq n_2$.

Finally we consider a decay involving the subrepresentation

$$\begin{array}{ccc}
 \mathbb{C}^{n_1} & \xrightarrow{B} & \mathbb{C}^{n_2} \\
 \uparrow & & \uparrow \\
 \mathbb{C} & \xrightarrow{b} & \mathbb{C}
 \end{array} \tag{2.2.23}$$

It is clear that b can be chosen in such a way that this is always a subrepresentation. Then stability demands that the central charges satisfy

$$\arg(\mathcal{Z}_1 + \mathcal{Z}_2) < \arg(n_1\mathcal{Z}_1 + n_2\mathcal{Z}_2). \tag{2.2.24}$$

However, given that $n_1 \leq n_2$, and that \mathcal{Z}_2 has smaller phase than \mathcal{Z}_1 , it is not possible to satisfy the above inequality. It follows that the only possibility for a bound state is that (2.2.23) is not a subrepresentation, but an isomorphism of representations. So we only have the possibility of non-trivial moduli spaces for $n_1 = n_2 = 1$.

In summary, when $\arg(\mathcal{Z}_2) < \arg(\mathcal{Z}_1)$ this theory supports a bound state with charge $\gamma_1 + \gamma_2$. The moduli space of representations of this charge is given by the quotient of a single non-zero complex number B modulo the action of the complexified gauge group. Clearly this moduli space is just a point, and so this representation describes a single hypermultiplet. The complete spectrum for this example is depicted in Figure 2.2, and agrees with the known result for this theory [53]. This basic 2-3 decay process is known in various contexts as a primitive decay [18]. In formalism of Kontevich and Soibelman this wall-crossing gives rise to the pentagon identity of quantum dilogarithms.

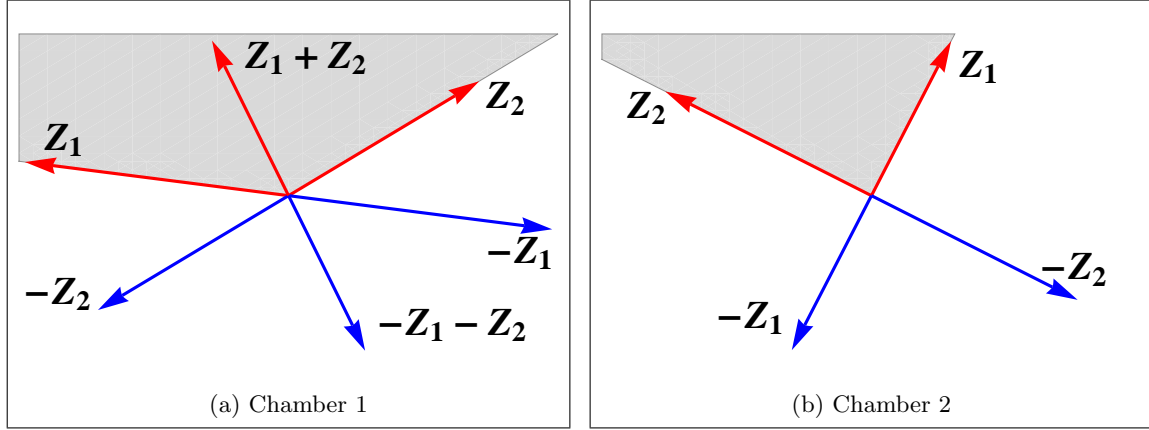


Figure 2.2: The chambers of the A_2 Argyres-Douglas theory. The BPS spectrum is plotted in the central charge plane. Particles are shown in red, antiparticles in blue. The cone of particles is the shaded grey region. In (a) the particles form a bound state. In (b) the bound state is unstable and decays.

A_3 Theory

As another example of quiver representation theory and wall-crossing we consider a quiver involving a non-trivial superpotential \mathcal{W} . The quiver, known to be related to the A_3 Argyres-Douglas theory is given by

$$\begin{array}{c}
 \textcircled{3} \\
 \alpha_3 \swarrow \quad \searrow \alpha_2 \\
 \textcircled{1} \quad \xrightarrow{\alpha_1} \quad \textcircled{2}
 \end{array}
 \tag{2.2.25}$$

We let α_i indicate the bifundamental field map exiting node i and \mathcal{Z}_i the central charge of node i . The quiver is equipped with a superpotential

$$\mathcal{W} = \alpha_3 \alpha_2 \alpha_1.
 \tag{2.2.26}$$

Minimization of \mathcal{W} implies that in any allowed field configuration all compositions of pairs of maps

vanish

$$\alpha_2 \circ \alpha_1 = 0, \quad \alpha_3 \circ \alpha_2 = 0, \quad \alpha_1 \circ \alpha_3 = 0. \quad (2.2.27)$$

We will show that this quiver has, up to relabeling the nodes, exactly two chambers with four or five BPS hypermultiplets respectively.

First, we note that as usual all of the node representation where the dimensions n_i of the associated vector space are given by $n_i = \delta_{ij}$ for $j = 1, 2, 3$ are stable and hence yield three hypermultiplets. Further, when one of the n_i vanishes, then two of the maps α must also vanish and the analysis reduces to the A_2 case considered in the previous section. This yields two or one bound states depending on whether the phases of the \mathcal{Z}_i are or are not cyclically ordered. To conclude the analysis of this quiver, we now wish to illustrate that there are no further bound states that arise from representations

$$\begin{array}{ccccc} & & \alpha_3 & & \\ & & \curvearrowright & & \\ \mathbb{C}^{n_1} & \xrightarrow{\alpha_1} & \mathbb{C}^{n_2} & \xrightarrow{\alpha_2} & \mathbb{C}^{n_3} \end{array} \quad (2.2.28)$$

with all n_i non-zero.

We begin by considering possible subrepresentations corresponding to node vectors, $(1, 0, 0)$, $(0, 1, 0)$, and $(0, 0, 1)$. These are only subrepresentations when α_i has a kernel for $i = 1, 2, 3$ respectively. Clearly not all of these can be subreps simultaneously or else the representation would already be destabilized. It follows that at least one of the α_i , say α_1 is injective and hence in particular $n_1 \leq n_2$.

Now we apply the F-term equations (2.2.27). From the fact that $\alpha_1 \circ \alpha_3 = \alpha_2 \circ \alpha_1 = 0$ and the fact that α_1 is injective we learn that both α_2 and α_3 have non-vanishing kernels. This means that both the node representations $(0, 1, 0)$ and $(0, 0, 1)$ are subreps so we deduce that \mathcal{Z}_1 must have largest phase for stability, and $\arg \mathcal{Z}_2, \arg \mathcal{Z}_3 < \arg(n_1 \mathcal{Z}_1 + n_2 \mathcal{Z}_2 + n_3 \mathcal{Z}_3)$.

However now we consider a subrepresentation with dimension vector $(1, 1, 0)$.

$$\begin{array}{ccccc}
 & & \alpha_3 & & \\
 & \curvearrowright & & \curvearrowleft & \\
 \mathbb{C}^{n_1} & \xrightarrow{\alpha_1} & \mathbb{C}^{n_2} & \xrightarrow{\alpha_2} & \mathbb{C}^{n_3} \\
 \uparrow i & & \uparrow j & & \uparrow 0 \\
 \mathbb{C} & \xrightarrow{\beta_1} & \mathbb{C} & \xrightarrow{\beta_2} & 0 \\
 & \curvearrowleft & & \curvearrowright & \\
 & & 0 & &
 \end{array} \tag{2.2.29}$$

This is a subrep exactly when the image of α_1 meets the kernel of α_2 non-trivially, which it does by the F-terms. Thus we learn that

$$\arg(\mathcal{Z}_1 + \mathcal{Z}_2) < \arg(n_1\mathcal{Z}_1 + n_2\mathcal{Z}_2 + n_3\mathcal{Z}_3). \tag{2.2.30}$$

Given the conditions on the \mathcal{Z}_i and the fact that $n_1 \leq n_2$, the above is impossible.

Thus we have arrived at a contradiction. It follows that for this quiver with the given superpotential there are no states with all n_i non-vanishing. Note that this conclusion is altered when the superpotential is turned off. In that case it is easy to check that the representation $(1, 1, 1)$ with all maps non-zero provides a stable hypermultiplet at all moduli. This completes our analysis of this quiver.

2.3 Quiver Mutation and Duality

We have seen how wall crossing is encoded into our quiver quantum mechanics picture. Walls of marginal stability correspond to hypersurfaces in which two central charges become aligned. The stability condition will differ on the two sides of this wall, and therefore there may be some representations which are stable on one side but not the other. There is in fact another type of hypersurface in moduli space that is strikingly relevant in our picture: hypersurfaces across which a fixed quiver quantum mechanics description of the BPS spectrum may break down entirely. Following [1] we will refer to these as walls of the second kind.

The situation is less dire than it may seem; we will be able to find another quiver description, valid on the other side of the wall. We will argue that the transformation of a quiver across a wall of the second kind is given by a canonical procedure, known as *quiver mutation* which describes a quantum mechanical duality relating the ground state spectra of two distinct quivers. Once the rule for transforming quivers at such walls is understood, we will be able to start with a quiver description at any point in moduli space and arrive at any other point by following an arbitrary path connecting them, doing the necessary mutations along the way. Further, in section 2.3.2 we will revisit this procedure and see that the same transformation can be made on quivers at a fixed point in moduli space, and in this case the transformation will take us between quivers that describe the same physics. We will then immediately exploit this duality to circumvent the computations involved in solving the representation theory problem.

Recalling that the nodes of a quiver all correspond to particles, and must therefore have central charges which lie in the upper half-plane, we see what can go wrong. As we tune moduli, our central charge function changes, and as we cross some real co-dimension 1 subspace in \mathcal{U} , the central charge of one of the nodes may exit the half-plane. This behavior defines the walls of the second kind. They are the loci in moduli space (including as usual masses and couplings) where the central charge of a basis particle becomes real

$$\mathcal{Z}_a(\gamma_i) \in \mathbb{R}. \tag{2.3.1}$$

Let us study the process of crossing a wall of the second kind in more detail. Consider the central charge configuration illustrated in Figure 2.3a where the BPS particles are described by the quiver Q . As moduli are varied, the central charge of one of the basis elements, \mathcal{Z}_1 rotates out of the upper half-plane and we arrive at the new configuration illustrated in Figure 2.3b.

The first thing to notice about this process is that, since no central charges align, no walls of marginal stability are crossed, and hence the total BPS spectrum (consisting of both particles

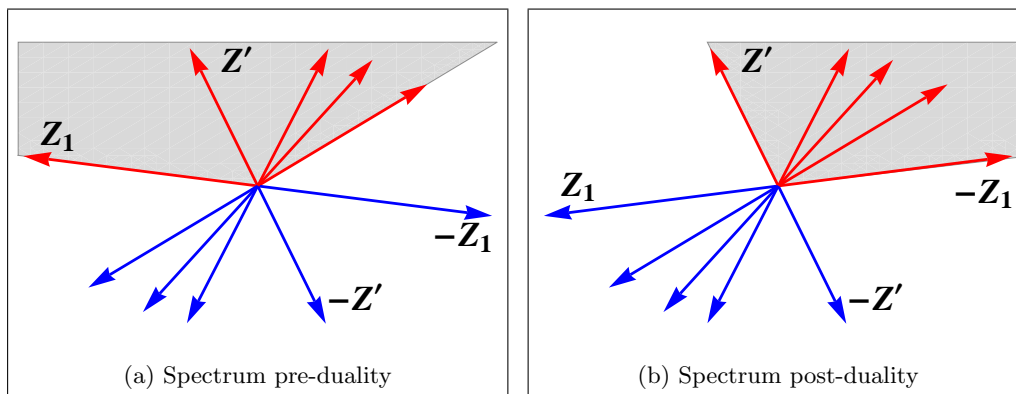


Figure 2.3: A discontinuity in the quiver description results in a quantum mechanical duality described by quiver mutation. In both diagrams the BPS spectrum is plotted in the central charge plane. Red lines denote particles while blue lines denote antiparticles. The gray shaded region indicates the cone of particles. In passing from (a) to (b) the particle with central charge Z_1 changes its identity to an antiparticle. The cone of particles jumps discontinuously and a new quiver description is required.

and antiparticles) is the same in Figures 2.3a and 2.3b. On the other hand, from the point of the quiver this process is discontinuous. After Z_1 has rotated out of the upper half of the central charge plane, it has changed its identity from a particle to an antiparticle. Then the original basis of particles encoded by the quiver Q is no longer an acceptable basis. Specifically, in passing from Figure 2.3a to Figure 2.3b, the cone of particles has jumped discontinuously and as a result the original quiver description of the BPS spectrum is no longer valid.

To remedy this deficiency we must introduce a new quiver \tilde{Q} that encodes the BPS spectrum in the region of moduli space described by Figure 2.3b. Since the total spectra of particles and antiparticles in Q and \tilde{Q} are identical, the physical relation between them is that of a duality: they are equivalent descriptions of the same total spectrum of BPS states. In the moduli space \mathcal{U} the regions of validity of Q and \tilde{Q} are sewn together smoothly along the loci where the central charge of an elementary basis particle is real. This sewing is illustrated in Figure 2.4

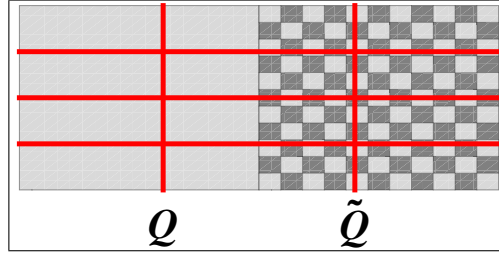


Figure 2.4: A cartoon of the moduli space and its relation to various BPS quiver descriptions. The red lines denote walls of marginal stability where the BPS spectrum jumps. The gray shaded region is the domain in moduli space where Q describes the BPS spectrum. The gray checkered region is the domain where \tilde{Q} describes the spectrum. The two descriptions are glued together smoothly away from the walls of marginal stability. Their interface is a wall of the second kind.

In section 2.3.1 we define the operation of mutation on a given quiver Q to produce the quiver \tilde{Q} , valid on the other side of the wall of the second kind. In section 2.3.2 we explain how the existence of the mutation operation, when interpreted as duality between different quiver descriptions, leads to a powerful and striking method for determining BPS spectra.

2.3.1 Quiver Mutation

As the preceding discussion indicates, a global description of the BPS spectrum across the entire Coulomb branch will require many quivers all glued together in the fashion described above. In this subsection we describe the algorithmic construction of this set of quivers by a graphical process known as quiver mutation. In the following subsection we justify these rules using arguments from quiver representation theory.

To define mutation, let us suppose that node γ_1 is the BPS particle in the quiver whose central charge \mathcal{Z}_1 is rotating out of the half-plane. We then seek to describe the dual quiver \tilde{Q} with corresponding nodes $\{\tilde{\gamma}_i\}$. Of course, since we have determined that a given spectrum of BPS particles admits at most one basis of BPS states, both \tilde{Q} and $\{\tilde{\gamma}_i\}$ are uniquely fixed. What's more,

the quiver \tilde{Q} can be described in a simple graphical way starting from Q . [45, 46, 56–60]. The new basis is given by

$$\tilde{\gamma}_1 = -\gamma_1 \quad (2.3.2)$$

$$\tilde{\gamma}_j = \begin{cases} \gamma_j + (\gamma_j \circ \gamma_1)\gamma_1 & \text{if } \gamma_j \circ \gamma_1 > 0 \\ \gamma_j & \text{if } \gamma_j \circ \gamma_1 \leq 0. \end{cases} \quad (2.3.3)$$

To construct \tilde{Q} graphically we follow the steps below:

1. The nodes of \tilde{Q} are in one-to-one correspondence with the nodes in Q .
2. The arrows of \tilde{Q} , denoted \tilde{B}_{ij}^a , are constructed from those of Q , denoted B_{ij}^a as follows:
 - (a) For each arrow B_{ij}^a in Q draw an arrow \tilde{B}_{ij}^a in \tilde{Q} .
 - (b) For each length two path of arrows passing through node 1 in Q , draw a new arrow in \tilde{Q} connecting the initial and final node of the length two path

$$B_{i1}^a B_{1j}^b \longrightarrow \tilde{B}_{ij}^c. \quad (2.3.4)$$

- (c) Reverse the direction of all arrows in \tilde{Q} which have node 1 as one of their endpoints.


$$\tilde{B}_{i1}^a \longrightarrow \tilde{B}_{1i}^a; \quad \tilde{B}_{1j}^a \longrightarrow \tilde{B}_{j1}^a. \quad (2.3.5)$$

3. The superpotential $\tilde{\mathcal{W}}$ of \tilde{Q} is constructed from the superpotential \mathcal{W} of Q as follows:
 - (a) Write the same superpotential \mathcal{W} .
 - (b) For each length two path considered in step 2(b) replace in \mathcal{W} all occurrences of the product $B_{i1}^a B_{1j}^b$ with the new arrow \tilde{B}_{ij}^c .
 - (c) For each length two path considered in step 2(b) $B_{i1}^a B_{1j}^b$ there is now a new length three cycle in the quiver \tilde{Q} formed by the new arrow created in step 2(b) and the reversed arrows in step 2(c)


$$\tilde{B}_{1i}^a \tilde{B}_{ij}^c \tilde{B}_{j1}^b. \quad (2.3.6)$$

Add to the superpotential all such three cycles.

As a simple example of this procedure we consider the A_3 quiver of section 2.2.3 shown on the left and its mutation at node 1 shown on the right.



$$\mathcal{W} = B_{12}B_{23}B_{31}$$



$$\mathcal{W} = \tilde{B}_{32}\tilde{B}_{23} + \tilde{B}_{32}\tilde{B}_{21}\tilde{B}_{13}$$

(2.3.7)

As the above example illustrates, the process of quiver mutation in general creates cycles of length two in our new quiver. From a physical perspective these are fields in the quiver quantum mechanics which admit a gauge invariant mass term. In the example above such mass terms are present in the quadratic piece of the potential $\tilde{B}_{32}\tilde{B}_{23}$. As is typical in physical theories, the massive fields decouple from the analysis of ground states and hence do not affect the BPS spectrum. We may therefore integrate them out. Thus to our list of quiver mutation rules we append the following final steps:

4. For each two-cycle in \tilde{Q} for which a quadratic term appears in \tilde{W} , delete the two associated arrows.
5. For each deleted arrow \tilde{B}_{ij}^a in step 4, solve the equation of motion

$$\frac{\partial \tilde{W}}{\partial \tilde{B}_{ij}^a} = 0. \quad (2.3.8)$$

Use the solution to eliminate \tilde{B}_{ij}^a from the potential.

In the example illustrated above, the only two cycle has quadratic terms in the superpotential and is therefore deleted from the quiver. This results in a vanishing superpotential and a

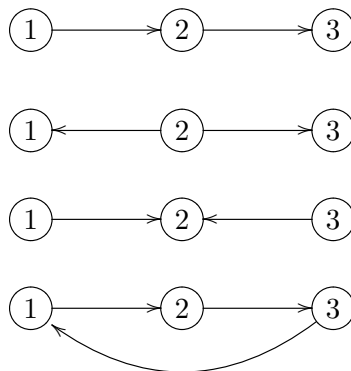
quiver of the following form.

$$\textcircled{2} \longrightarrow \textcircled{1} \longrightarrow \textcircled{3} \quad (2.3.9)$$

As a general rule, the study of BPS quivers is greatly complicated by the existence of pairs of opposite arrows whose associated fields cannot be integrated out from the superpotential. When this is never the case, that is when the potential \mathcal{W} is strong enough to integrate out to all opposite bifundamental fields after an arbitrary sequence of mutations, the potential is said to be *non-degenerate*. It is a fortunate simplification that for all of the BPS quivers related to quantum field theories that we discuss in this chapter the potential will turn out to be non-degenerate.

A_3 Revisited

To put the above theory of quiver mutation in perspective, it is useful to consider the simplest example where the phenomenon of wall of the second kind occur. This is the A_3 theory whose representation theory was investigated in section 2.2.3. There are in fact four distinct quivers for the A_3 theory related by mutation. These are given by



Let us name these four quivers respectively as L , O , I , and C . The representation theory of the C quiver was worked out in section 2.2.3. In particular we determined that C supports either 4 or 5 BPS states depending on moduli. The representation theory of the other quivers is also readily calculated. One finds that L has 6 distinct chambers, while both I and O have 4. If we denote

Chamber	Phase Conditions	Number of BPS States
L_1	$\theta_3 > \theta_2 > \theta_1$	3
L_2	θ_2 smallest, and $\theta_1, \theta_3 > \theta_{12}$	4
L_3	θ_2 largest, and $\theta_{23} > \theta_1, \theta_3$	4
L_4	$\theta_1 > \theta_{12} > \theta_3 > \theta_2$	5
L_5	$\theta_2 > \theta_1 > \theta_{23} > \theta_3$	5
L_6	$\theta_1 > \theta_2 > \theta_3$	6
O_1	θ_2 smallest	3
O_2	θ_2 intermediate	4
O_3	θ_2 largest, and $\theta_{12} < \theta_3$ or $\theta_{23} < \theta_1$	5
O_4	θ_2 largest, and $\theta_{12} > \theta_3$ and $\theta_{23} > \theta_1$	6
I_1	θ_2 largest	3
I_2	θ_2 intermediate	4
I_3	θ_2 smallest, and $\theta_3 < \theta_{12}$ or $\theta_1 < \theta_{23}$	5
I_4	θ_2 smallest, and $\theta_3 > \theta_{12}$ and $\theta_1 > \theta_{23}$	6
C_1	not cyclically ordered e.g. $\theta_2 > \theta_1 > \theta_3$	4
C_2	cyclically ordered e.g. $\theta_1 > \theta_2 > \theta_3$	5

Table 2.1: The chambers of the A_3 quivers before mutation equivalences are imposed. For each quiver labelled with node charges Z_i , θ_i denotes the argument of Z_i while θ_{ij} denotes the argument of $Z_i + Z_j$.

by θ_i the phase of \mathcal{Z}_i and θ_{ij} the phase of $\mathcal{Z}_i + \mathcal{Z}_j$, then the complete list of chambers is given in table 2.1.

In the global theory of A_3 these chambers are connected together across walls of the second kind where the quiver changes by a mutation. To understand mutations we then represent each chamber as a node in a graph and connect those mutation equivalent with directed arrows. For example we define the expression

$$Q_i \longrightarrow \tilde{Q}_j, \quad (2.3.10)$$

to mean that mutation in chamber i of quiver Q on the leftmost boundary ray leads to chamber j in the quiver \tilde{Q} . With these conventions the complete structure of walls of the second kind in the A_3 theory is encoded in the following diagrams.

$$\begin{array}{cccc}
 \begin{array}{c} L_1 \\ \downarrow \\ I_1 \end{array} \begin{array}{c} \nearrow \\ \searrow \\ O_1 \end{array} & & \begin{array}{c} I_2 \\ \nearrow \quad \searrow \\ L_2 \quad C_1 \quad L_3 \\ \searrow \quad \nearrow \\ O_2 \end{array} & & \begin{array}{c} L_5 \longrightarrow C_2 \longrightarrow L_4 \\ \uparrow \qquad \qquad \downarrow \\ I_3 \longleftarrow O_3 \end{array} & & \begin{array}{c} L_6 \\ \downarrow \\ O_4 \end{array} \begin{array}{c} \nearrow \\ \searrow \\ I_4 \end{array}
 \end{array} \quad (2.3.11)$$

Where in the above, some chambers have two arrows leaving them because one can change the leftmost ray without crossing a wall.

Justification of Mutation

The previous subsection gives a straightforward recipe for producing, from a given quiver Q , all of its related duals by considering mutations at various nodes. However we have not yet explained why this mutation rule is in fact correct.⁷ In this subsection we fill in this gap.⁷ Specifically our goal will be to derive the mutation rule, given the assumption that a quiver description \tilde{Q} exists after the transition illustrated by Figure 2.3.

⁷The arguments in this section are somewhat technical and could be skipped in a first reading.

The basic point is that the new elementary basis particles $\tilde{\gamma}_i$, are interpreted from the point of view of Q as certain bound states of the original basis particles γ_i . The key step is to identify which bound states.

Consider again the cone geometry illustrated in Figure 2.3. A special role is played by the two particles whose central charge rays form the boundary of the cone. Such particles must always be included in the basis because, as their central charges are on the boundary of the cone, there is no way to generate these states by positive linear combinations of other rays in the cone. Thus in Figure 2.3b the two states with central charges \mathcal{Z}' and $-\mathcal{Z}_1$ must appear as nodes of the quiver \tilde{Q} . Of these, the latter is easy to identify as the antiparticle of the mutated node, $-\gamma_1$, and hence this charge must be in the new basis. Meanwhile, in the following argument we will prove that the left-most ray, which we frequently refer to as the extremal ray, \mathcal{Z}' , is always a two particle bound state which may be identified explicitly.

To begin, we consider all connected length two subquivers of Q which involve the node γ_1 . For a given node γ_i there are k_i arrows pointing either from γ_i to γ_1 or from γ_1 to γ_i .

$$\begin{array}{ccc}
 \begin{array}{c} \textcircled{\gamma_1} \end{array} & \begin{array}{c} \xrightarrow{B_1} \\ \xrightarrow{B_2} \\ \vdots \\ \xrightarrow{B_{k_i}} \end{array} & \begin{array}{c} \textcircled{\gamma_i} \end{array} & \text{or} & \begin{array}{c} \textcircled{\gamma_i} \end{array} & \begin{array}{c} \xrightarrow{B_1} \\ \xrightarrow{B_2} \\ \vdots \\ \xrightarrow{B_{k_i}} \end{array} & \begin{array}{c} \textcircled{\gamma_1} \end{array} & (2.3.12)
 \end{array}$$

Let us describe the leftmost bound state supported by these two node quivers. In the case on the right of (2.3.12), γ_1 appears as a sink. Then, since $\mathcal{Z}(\gamma_1)$ has largest phase by hypothesis, γ_1 by itself is a destabilizing subrep of any possible bound state; thus no bound states can form.

On the other hand, in the case on the left of (2.3.12), where γ_1 appears as a source, bound states can exist. We consider a general representation of the form

$$\begin{array}{ccc}
 \mathbb{C}^n & \begin{array}{c} \xrightarrow{B_1} \\ \xrightarrow{B_2} \\ \vdots \\ \xrightarrow{B_{k_i}} \end{array} & \mathbb{C}^m & (2.3.13)
 \end{array}$$

To make a bound state with largest possible phase we wish to make a representation where n/m is as large as possible. However, it is not difficult to see that the ratio n/m is bounded. Indeed, since $\mathcal{Z}(\gamma_1)$ has largest phase, there is a potentially destabilizing subrepresentation involving only the particle γ_1 . Such a subrepresentation is described by k_i commutative diagrams of the form

$$\begin{array}{ccc} \mathbb{C}^n & \xrightarrow{B_j} & \mathbb{C}^m \\ \uparrow & & \uparrow \\ \mathbb{C} & \xrightarrow{0} & 0 \end{array} \quad (2.3.14)$$

In other words, the potential destabilizing subrepresentation is nothing but a non-zero vector which is simultaneously in the kernel of all of the maps B_j . But then a simple dimension count shows that

$$\text{dimension} \left(\bigcap_{j=1}^{k_i} \ker(B_j) \right) \geq n - k_i m. \quad (2.3.15)$$

And so in particular when the right-hand side of the above is positive, the subrepresentation (2.3.14) exists and hence the bound state is unstable. Thus we learn that stability requires

$$\frac{n}{m} \leq k_i. \quad (2.3.16)$$

Finally, it is not difficult to find a stable representation R which saturates the above bound. Indeed let us take $n = k_i$ and $m = 1$. Then the maps B_j are simply projections to a line. The stability constraint that the B_j have no common kernel implies that, up to gauge transformation, B_j can be taken to be the dual vector to the j th basis element in the vector space attached to γ_1 . So defined, the representation R is stable and has no moduli. Thus it gives rise to a hypermultiplet with charge

$$\gamma_i + k_i \gamma_1. \quad (2.3.17)$$

This completes the required analysis of quivers with two nodes. To summarize, in the region of parameter space where $\mathcal{Z}(\gamma_1)$ has largest phase, we have determined the extremal bound state of all two-node subquivers involving γ_1 . The charges of the extremal bound states are:

- If $\gamma_i \circ \gamma_1 < 0$ then the extremal bound state is simply γ_i .
- If $\gamma_i \circ \gamma_1 > 0$ then the extremal bound state is $\gamma_i + (\gamma_i \circ \gamma_1)\gamma_1$.

Now we claim that in the quiver Q with an arbitrary number of nodes, one of the two particle bound states we have identified above will still be the left-most extremal ray after $\mathcal{Z}(\gamma_1)$ exits the upper half-plane. To see this, we consider an arbitrary stable representation R of Q . We write the charge of R as

$$\gamma_R = n\gamma_1 + \sum_{\gamma_i \circ \gamma_1 > 0} m_i \gamma_i + \sum_{\gamma_j \circ \gamma_1 \leq 0} l_j \gamma_j \quad (2.3.18)$$

Let us focus in on the representation R near the node γ_1 . There are now many nodes connected to the node 1 by various non-zero maps. For those connections with $\gamma_i \circ \gamma_1 \leq 0$, the node γ_1 appears as a sink, for those with $\gamma_i \circ \gamma_1 > 0$, γ_1 appears as a source.

Our strategy is again to test whether R is stable with respect to decays involving the subrepresentation S with charge γ_1 . As in the two node case, in such a situation the connections where γ_1 is a sink are irrelevant. On the other hand, if S is really a subrepresentation then for each node link in the representation where node 1 is a source, we have commutative diagrams of the form (2.3.14).

Given that $\mathcal{Z}(\gamma_1)$ has largest phase, stability of R means that we must obstruct the existence of S . As in the analysis of the two node quivers we see that S will be a subrepresentation provided that the kernels of all maps exiting the node γ_1 have nonzero intersection. However, just as in (2.3.15) we can see that this leads to an a priori bound on n , the amount of γ_1 contained in the representation R . Explicitly we have

$$\text{dimension} \left(\bigcap_{\gamma_i \circ \gamma_1 > 0} \bigcap_{j=1}^{k_i} \ker(B_j) \right) \geq n - \sum_{\gamma_i \circ \gamma_1 > 0} k_i m_i. \quad (2.3.19)$$

Hence to obstruct the existence of the subrepresentation S we deduce the bound

$$n \leq \sum_{\gamma_i \circ \gamma_1 > 0} k_i m_i. \quad (2.3.20)$$

But now we can directly see that R cannot be extremal. We have

$$\begin{aligned} \arg(\mathcal{Z}(R)) &= \arg\left(nZ_1 + \sum_{\gamma_i \circ \gamma_1 > 0} m_i Z_i + \sum_{\gamma_j \circ \gamma_1 \leq 0} l_j Z_j\right) \\ &\leq \arg\left(\sum_{\gamma_i \circ \gamma_1 > 0} m_i (k_i Z_1 + Z_i) + \sum_{\gamma_j \circ \gamma_1 \leq 0} l_j Z_j\right). \end{aligned} \quad (2.3.21)$$

But the final expression in (2.3.21) is manifestly contained in the positive span of the two node extremal bound states, $k_i \gamma_1 + \gamma_i$, that we identified in our analysis of two node quivers. In particular, this means that R cannot be a boundary ray and hence is not extremal.

Thus we deduce that the left-most ray after mutation is one of the two particle bound states that we have identified in our analysis of two node quivers. Extremality then ensures that our new basis must include this two particle bound state. But finally we need only notice that the central charges of all the two node extremal bound states that we have discovered are independent parameters. Indeed letting the central charges vary in an arbitrary way, our conclusion is in fact that *all* the two node bound states which we have determined must in fact be in the new basis. In particular this means that the new basis of charges after mutation is completely fixed and we may write the transformation as follows:

$$\tilde{\gamma}_1 = -\gamma_1 \quad (2.3.22)$$

$$\tilde{\gamma}_j = \begin{cases} \gamma_j + (\gamma_j \circ \gamma_1) \gamma_1 & \text{if } \gamma_j \circ \gamma_1 > 0 \\ \gamma_j & \text{if } \gamma_j \circ \gamma_1 \leq 0 \end{cases} \quad (2.3.23)$$

As one can easily verify, the graphical quiver mutation rules described in the previous section are a direct consequence of computing the new BPS quiver \tilde{Q} from the symplectic products of the new basis of charges $\{\tilde{\gamma}_i\}$. This completes our argument justifying the mutation rules.

2.3.2 The Mutation Method: BPS Spectra from Quiver Dualities

We saw above that at walls of the second kind, we were forced to change our quiver description because the central charge of some state exited the upper half of the complex half-plane, thereby turning from a particle to an antiparticle. We might also consider what happens if we fix a modulus $u \in \mathcal{U}$ and then consider a different definition of the particle half-plane, \mathcal{H} . If we imagine continuously changing our choice from one \mathcal{H} to another, the situation is precisely the same as above; there is some parameter which we are tuning, and at some critical value the central charge of some state becomes such that it switches from particle to antiparticle.

In this case, however, we are remaining at a fixed point in moduli space, and so all of these quivers describe precisely the same physics. That is, they are dual descriptions of the BPS spectrum. In fact, there is a whole class of quivers related to each other by duality at each point in moduli space. We will now exploit this fact to produce for us, in many cases, the entire spectrum for free.

First, let us reiterate that a single form of the quiver already in principle determines exactly which BPS states in the theory are occupied, including their spin and multiplicity. To find the answer, one can solve the representation theory of the quiver with superpotential, which amounts to the linear algebra problem described in section 2.2.2. However, in practice this problem can become quite intractable. The mutation method we propose gets rid of all of the unsightly work required in solving the problem directly, and instead produces the spectrum using chains of dualities through different quiver descriptions of the theory.

Recall our first application of quiver rep theory in section 2.2.2, where we checked that nodes of the quiver always correspond to multiplicity one hypermultiplets. This fact, together with an examination of which states are forced to be nodes for various choices of half-plane \mathcal{H} , is at the heart of what we call the mutation method. Imagine that for our initial choice of \mathcal{H} , with BPS

basis $\{\gamma_i\}$, γ_1 is the node such that $\mathcal{Z}(\gamma_1)$ is left-most in \mathcal{H} .⁸ Say we then rotate our half-plane past it, and do the corresponding mutation to arrive at a new quiver description of the theory. This mutation includes an action on the charges of the quiver γ_i , as given in equation (3.2)-(3.3). Since this new quiver is a description of the BPS states of the same theory, its nodes are also multiplicity one hypermultiplets. Consequently, we have discovered some subset of states in the 4d theory which we can say must exist. In particular, we generate some new BPS states of the form $-\gamma_1, \gamma_i + (\gamma_i \circ \gamma_1)\gamma_1$. Of course, $-\gamma_1$ is just the antiparticle of the state γ_1 , so this is no additional information. However, the states $\gamma_i + m_i\gamma_1$ are completely new. To discover these same states from the original quiver would have involved solving the non-trivial representation theory problem studied in the previous subsection. We are able to avoid this headache by observing that, because of duality, these states must be in the spectrum for consistency.

So we have found that duality will trivially produce some subset of the spectrum as nodes of various dual quivers. But in fact it does much more: in many cases, mutation produces the full spectrum in this way. Imagine we're in a chamber with finitely many BPS states, and pick an arbitrary state γ which is a hypermultiplet of the 4d theory. Then we can rotate the half-plane \mathcal{H} so that γ is left-most. As usual, since the nodes of the quiver form a positive basis for states in \mathcal{H} , γ must itself be a node. Therefore, if we start with any quiver description, and start rotating $\mathcal{H} \rightarrow e^{-i\theta}\mathcal{H}$ until γ becomes left-most, we will go through a corresponding sequence of mutations, after which γ will simply be a node of the quiver.

It is then easy to see how to systematically generate the spectrum in any finite chamber. We start with any quiver description which is valid at our given point in moduli space, and start rotating the half-plane. Since there are only finitely many states, we will only pass through finitely many mutations before we return to the original half-plane $\mathcal{H} \rightarrow e^{2\pi i}\mathcal{H}$.⁹ The key point is every

⁸From now on we will abuse verbiage slightly and simply say that “ γ_1 is left-most.”

⁹Recall that for a given choice of \mathcal{H} , the quiver description is actually unique - there is a unique positive integral

state in the chamber is left-most at some point during this rotation, so every state will indeed show up as a node of one of the dual quivers. Since rotating past a state corresponds to mutating on the node corresponding to that state, if we do the entire sequence of mutations and record each state we've mutated on, we will have exhausted all states in the chamber.

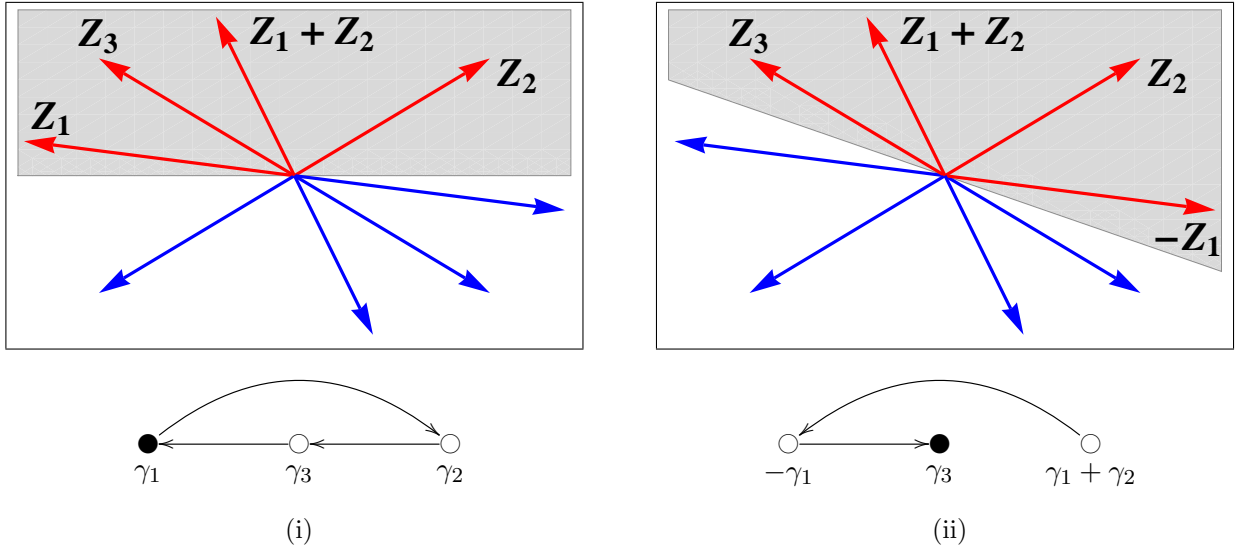
We can save a bit of work by making use of CPT: for any state γ in the spectrum, $-\gamma$ is also occupied. So instead of taking $\mathcal{H} \rightarrow e^{2\pi i}\mathcal{H}$, we can just rotate half-way, $\mathcal{H} \rightarrow e^{i\pi}\mathcal{H}$, ending up at the quiver which describes all the antiparticles.¹⁰ If we record every state γ we mutate on as \mathcal{H} is rotated, and then add all antiparticles $-\gamma$, we will have precisely the spectrum of the 4d theory. Note that we must repeat this procedure for each chamber, by doing mutations in some different order, as prescribed by the ordering of the phases of the central charges in that region of moduli space. As we discussed above any given quiver generally only covers some subset of moduli space; therefore, for different chambers, it will generally be necessary to apply this procedure to different mutation forms of the quiver.

Let's try an example. The representation theory for the Argyres-Douglas A3 theory was worked on in detail in section 2.3.2. We will see how to reproduce it with much less work in the present framework. We will assume that we are at a point in moduli space covered by the cyclic three node quiver. Imagine that γ_1 is leftmost. After the first mutation, the mutation that follows will depend on the ordering of γ_3 and $\gamma_1 + \gamma_2$. Suppose that γ_3 is to the left. Then the particle

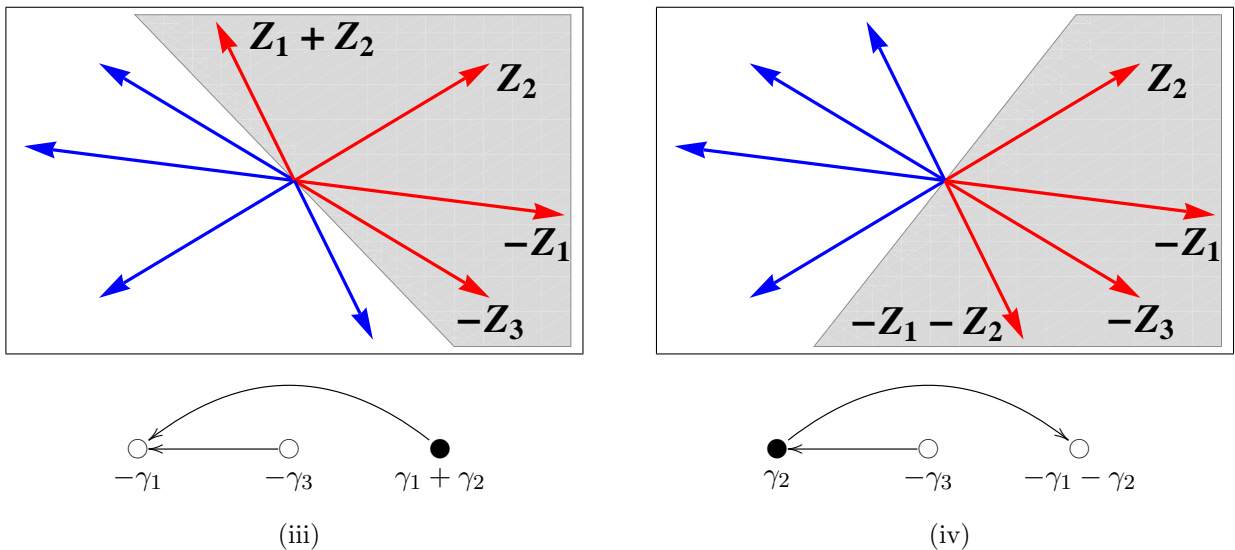
basis for the lattice of occupied BPS states, up to permutation. So we will also return to the original quiver up to permutation when \mathcal{H} undergoes a full rotation.

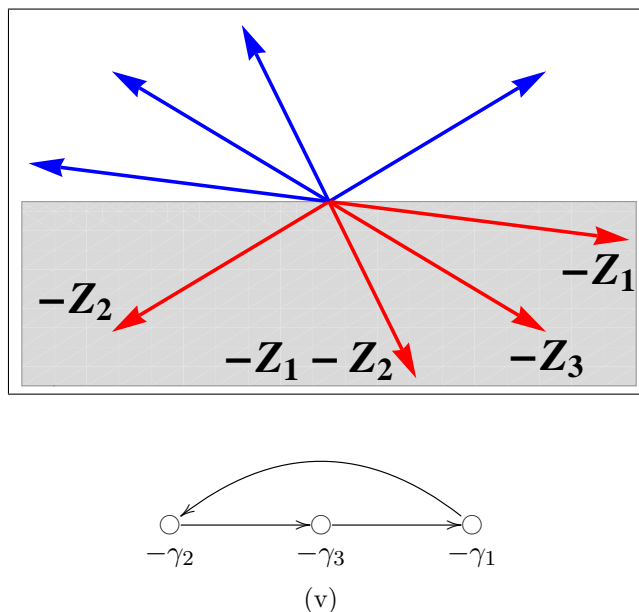
¹⁰By a similar argument as above, the final quiver will have nodes $-\gamma_i$.

half-plane, \mathcal{H} and associated quiver before (i) and after (ii) the first mutation at γ_1 are



In the above diagrams, we denote the left-most particle state in each quiver, which indicates the next node to be mutated, by drawing the corresponding node in black, \bullet . Now since the γ_i were in the original half-plane \mathcal{H} to begin with, it must be that $\gamma_1 + \gamma_2$ is to the left of $-\gamma_1$ and $-\gamma_3$ in the current half-plane. This is true in general: one never mutates on negative nodes in going through a π -rotation of \mathcal{H} from a quiver to its antiparticle quiver. The remaining mutations are completely fixed, and we find (iii,iv,v)





So we've arrived at the antiparticle quiver, which at the level of quiver without charges is the same, because the antisymmetric product is not affected by an overall sign on charges.¹¹ Therefore we've discovered a chamber with the states γ_1 , γ_2 , γ_3 and $\gamma_1 + \gamma_2$. This indeed agrees with one of the chambers found in 2.3.2. All of the chambers can similarly be mapped out, without ever doing the linear algebra analysis.

We pause here to emphasize two important points. The first is to recall that a quiver from the mutation class generically only covers a subset of moduli space. Therefore to map out all chambers, one must carry forth the above with the starting quiver being any one of the quivers in the mutation class. The second point is that, using the above method, one will not find any chamber covered by the cyclic quiver which contains the state $\gamma_1 + \gamma_2 + \gamma_3$. In the analysis of section 2.3.2, it was found that the $\gamma_1 + \gamma_2 + \gamma_3$ state was there in the quiver without superpotential, but killed when the (unique) non-degenerate superpotential was included. Thus we see that this mutation method knows about the associated non-degenerate superpotential indirectly. This is expected, because a

¹¹If you try to label nodes and keep track of them, which the drawings may subliminally suggest you do, in general you will return to $(-1) \times \text{permutation}$.

non-degenerate superpotential is required for the mutation rule written above to be sensible.

There are some simple non-trivial statements which we can immediately make based on this method. One is that any finite chamber can only contain hypermultiplets, with multiplicity one. The argument here is simply that any state in a finite chamber can be made into a node of some dual quiver, and nodes, as we've mentioned, can never correspond to higher spin objects or higher multiplicity hypers. Therefore, it would be inconsistent with duality to ever have a higher spin or higher multiplicity object in a finite chamber.

Now let's consider infinite chambers. An additional layer of complexity, as compared to the finite case, is that two dual quiver descriptions may be separated by an infinite sequence of mutations. This is because, as we rotate between two choices of \mathcal{H} , we will generically have infinitely many BPS states which rotate out to the left. Our method above depended on our ability to keep track of the sequence of mutations which happens as $\mathcal{H} \rightarrow e^{i\pi}\mathcal{H}$. Now the infinitude of states in some sense blocks us from competing this sequence of mutations. For example, if we start with a given quiver description, we can't explore beyond the closest accumulation ray in the \mathcal{Z} -plane. Because of this difficulty, we can't make a similarly definite statement about the method as it applies to infinite chambers. Indeed, for certain theories, such as $\mathcal{N} = 2^* SU(2)$ (the mass deformed $\mathcal{N} = 4$ theory), it appears that the method isn't sophisticated enough to exhaust the spectrum.¹²

However, as we will see in several examples, infinite chambers may also be understood by this method. Infinitude of the chamber is often due to higher spin objects, and we can often make progress by being just a bit clever. Note that any higher spin object must in fact be an accumulation ray of states in the central charge plane: If it weren't, we could rotate \mathcal{H} so that it was left-most, and as above, in this dual quiver description our higher spin state would be a node. Of course

¹²Of course we can always produce some arbitrarily large subset of states of the theory by mutating until exhaustion (of the mutator, that is).

this is a contradiction - nodes are always multiplicity one hypers. Higher multiplicity hypers must similarly be accumulation rays, a fact which may be less intuitive outside of this framework.

Before going on to examples and applications, we make some additional technical notes about the actual implementation of the mutation method. As we have described it here, we choose a point of the physical moduli space, compute central charges at that point, and mutate on the nodes in the order given by the ordering of phases of the central charges, as we tune $\mathcal{H} \rightarrow \mathcal{H}_\pi$. Instead, when exploring the possible BPS spectra, it is sometimes more practical to simply mutate on the nodes in any order, and then check two things: (1) that the ordering chosen is consistent, and (2) that the ordering chosen is realized somewhere in physical moduli space. By consistent, we mean that there exists some choice of central charges $\mathcal{Z}(\gamma_i)$ that correspond to the ordering chosen. As it turns out, there is no need to check the first point: as long as we mutate only on nodes whose charges are given by positive linear combinations of the original γ_i , then the ordering is consistent. Of course, we expect to only mutate on positive nodes since we are only rotating by π through the particle half-plane, and all particles should be given by positive integer linear combinations of the initial γ_i . Note that the only condition for consistency is that $\arg \mathcal{Z}(\gamma_1 + \gamma_2)$ lie between $\arg \mathcal{Z}(\gamma_1), \arg \mathcal{Z}(\gamma_2)$. In fact, the mutation method protects us from making inconsistent choices. Fix $\arg \mathcal{Z}(\gamma_1) > \arg \mathcal{Z}(\gamma_2)$, and suppose we have already mutated past γ_1 , but not yet γ_2 . Thus $-\gamma_1$ is in the positive integral span of the mutated quiver basis. Suppose both $\gamma_1 + \gamma_2$ and γ_2 to appear as nodes; this is an immediate contradiction with the fact that the nodes form a basis, since now γ_2 is both a basis element and a non-trivial linear combination of basis elements $(\gamma_1 + \gamma_2) + (-\gamma_1)$. So only one of these can appear as nodes and be mutated on next. If it is $\gamma_1 + \gamma_2$, there we are safe, and there is no inconsistency. If it is γ_2 , let's mutate past so that both $-\gamma_1, -\gamma_2$ are in the positive integral span of the mutated quiver basis; now it is impossible for $\gamma_1 + \gamma_2$ to appear as a node of the quiver, or else we can construct 0 as a non-trivial linear combination of basis elements $\gamma_1 + \gamma_2 + (-\gamma_1) + (-\gamma_2)$.

Therefore we can apply the mutation method by simply mutating on the positive nodes in any order we like, until we arrive at a quiver with all nodes labelled by negative charges, indicating that we have completed the rotation $\mathcal{H} \rightarrow \mathcal{H}_\pi$. It remains to be checked whether the ordering we have applied is actually physically realized in moduli space. We can dispense of this final check when the physical moduli space has complex dimension equal to the number of nodes. Then as we move in moduli space, it is possible to tune all central charges of nodes however we wish. These theories are known as *complete* theories, studied and classified in [12]. In a companion paper [28] we studied the application of these techniques to the class of complete theories. In the more general case of non-complete theories, existence of the desired changer in the physical moduli space must be checked by hand.

Quiver Mutation and Quantum Monodromy

The mutation method outlined in the previous section can be extended to compute not only the BPS spectrum, but also the full Kontsevich-Soibelman (KS) quantum monodromy operator itself [1, 12, 27]. In this section we briefly discuss these techniques.

To implement the KS formalism one first introduces the quantum torus algebra. Let i index the nodes of the quiver, as discussed in detail in previous sections, these nodes integrally generate the lattice of BPS charges. Then the quantum torus algebra is defined by:

- A generator Y_i for each node of the quiver.
- Commutation relations between the generators.

$$Y_i Y_j = q^{-\gamma_i \circ \gamma_j} Y_j Y_i, \quad (2.3.24)$$

where in the above, q is a parameter.

Given a general charge $\gamma = \sum_i n_i \gamma_i$ we introduce the operator Y_γ as a normal ordered product of

the corresponding generators:

$$Y_\gamma \equiv N[Y_1^{n_1} \cdot Y_2^{n_2} \cdots Y_m^{n_m}]. \quad (2.3.25)$$

The KS framework gives a characterization of the BPS spectrum in terms of a certain operator $M(q)$ which acts on the quantum torus algebra and is constructed as a product of certain quantum dilogarithm operators, $\Psi(Y_\gamma, q)$ built from the Y_γ . These operators act naturally on the quantum torus algebra by conjugation

$$Y_\alpha \rightarrow \Psi(Y_\gamma, q) Y_\alpha \Psi(Y_\gamma, q)^{-1}. \quad (2.3.26)$$

Meanwhile, the operation of quiver mutation studied in the previous sections also acts on the algebra through its action on the charges at various nodes. We let μ_k denote the operation on the charge lattice induced by quiver mutation at the k -th node. The induced action on the generators Y_i is then given in parallel to equations (3.2)-(3.3) as

$$\mu_k(Y_i) = \begin{cases} Y_k^{-1} & \text{if } i = k \\ Y_i & \text{if } \gamma_i \circ \gamma_k > 0 \\ Y_{\gamma_i + (\gamma_k \circ \gamma_k) \gamma_k} & \text{if } \gamma_i \circ \gamma_k < 0 \end{cases} \quad (2.3.27)$$

We can combine the action of conjugation by the quantum dilogarithm with quiver mutation to produce a quantum mutation operator which acts on the torus algebra

$$\mathcal{Q}_k = \text{Ad}(\Psi(Y_k, q)) \circ \mu_k. \quad (2.3.28)$$

The quantum mutation operator is the natural generalization of quiver mutation to the torus algebra. Furthermore, just as ordinary quiver mutations, like those studied in the previous section, allow us to easily determine the BPS spectrum, the quantum mutation operator allows us to write the full quantum monodromy operator $M(q)$. Specifically, in a chamber consisting of finitely many BPS states there exists a sequence of mutations which acts as the identity (up to a permutation of nodes) on the quiver Q

$$\mu_{k(s)} \cdots \mu_{k(2)} \mu_{k(1)} Q = Q. \quad (2.3.29)$$

A key feature of this sequence is that it is phase ordered; the state $k(1)$ is left-most, the state $k(2)$ is next to left-most and so on. Associated to this sequence is an ordered product of quantum mutation operators

$$\mathcal{Q}_{k(s)} \cdots \mathcal{Q}_{k(2)} \mathcal{Q}_{k(1)}. \quad (2.3.30)$$

The above operator can be expressed in terms of the adjoint action of a single operator which is none other than the desired operator $M(q)$. As a consequence of the fact that the original sequence of mutations in equation (2.3.29) is phase ordered, the operator $M(q)$ has the desired expression in terms of a phase ordered product over the BPS states of quantum dilogarithm operators [4, 61, 62]. In this way we recover the full KS monodromy operator from ordered mutation sequences.

2.4 $SU(2)$ Gauge Theories

We begin our study of examples with $SU(2)$ gauge theories. This is a natural starting point, as the BPS spectra of several of these theories have been worked out from different points of view [13, 63–65]. We will reproduce those results straight-forwardly from the mutation method. These examples serve as non-trivial confirmation of our framework, and also as a demonstration of the power of the mutation method.

2.4.1 Pure $SU(2)$

The quiver for pure $SU(2)$ gauge theory has been worked out in various papers [4, 12, 21, 32, 33]. Here we will content ourselves to fix it based on the known $SU(2)$ spectrum, as was done in section 2.2.1, and then check that the mutation method produces the correct spectrum.

The quiver we are studying is given by

$$\begin{array}{ccc} \circ & \xrightleftharpoons{\quad} & \circ \\ \gamma_1 = (0, 1) & & \gamma_2 = (2, -1) \end{array} \quad (2.4.1)$$

The strong coupling chamber is given by $\arg \mathcal{Z}(\gamma_2) > \arg \mathcal{Z}(\gamma_1)$. As we rotate \mathcal{H} , we have the following sequence of mutations

$$\begin{array}{ccc}
 \begin{array}{c} \circ \rightleftarrows \bullet \\ \gamma_1 \qquad \qquad \gamma_2 \end{array} & \begin{array}{c} \bullet \rightleftarrows \circ \\ \gamma_1 \qquad \qquad -\gamma_2 \end{array} & \begin{array}{c} \circ \rightleftarrows \circ \\ -\gamma_1 \qquad \qquad -\gamma_2 \end{array} \\
 \text{(i)} & \text{(ii)} & \text{(iii)}
 \end{array} \tag{2.4.2}$$

We see that we end with the antiparticle quiver, and that the only states in this chamber are γ_1 and γ_2 . This agrees with the well known result that only the monopole and dyon are stable at strong coupling.

We can move on to do the same analysis at weak coupling, where $\arg \mathcal{Z}(\gamma_1) > \arg \mathcal{Z}(\gamma_2)$.

$$\begin{array}{ccc}
 \begin{array}{c} \bullet \rightleftarrows \circ \\ \gamma_1 \qquad \qquad \gamma_2 \end{array} & \begin{array}{c} \circ \rightleftarrows \bullet \\ -\gamma_1 \qquad \qquad 2\gamma_1 + \gamma_2 \end{array} & \begin{array}{c} \bullet \rightleftarrows \circ \\ 3\gamma_1 + 2\gamma_2 \qquad \qquad -2\gamma_1 - \gamma_2 \end{array} \\
 \text{(i)} & \text{(ii)} & \text{(iii)} \\
 & & \tag{2.4.3} \\
 \dots & \begin{array}{c} (k+1)\gamma_1 + k\gamma_2 \\ \bullet \rightleftarrows \circ \\ -k\gamma_1 - (k-1)\gamma_2 \\ \text{(k+1)} \end{array} & \dots
 \end{array}$$

It is quite clear that we are in an infinite chamber. The entire sequence we'll find is obvious: we will have $(k+1)\gamma_1 + k\gamma_2$ for $k \geq 0$, with charge $(2k, 1)$. In the \mathcal{Z} plane these limit to the ray $\alpha\mathcal{Z}(\gamma_1 + \gamma_2)$. Notice that the (e, m) charge of $\gamma_1 + \gamma_2$ is $(2, 0)$. We're finding the expected accumulation ray associated with the vector, the W boson, in the weak coupling spectrum. In terms of rotating the half-plane, W is protected from being a node because it is an accumulation ray of hypermultiplet dyons. In terms of the mutations, the "quiver with W as a node" is infinitely many mutations away in the space of dualities, preventing a contradiction. As mentioned above,

this accumulation ray is blocking us from exploring the states lying in the rest of the central charge plane. We expect to only find one vector in the pure $SU(2)$ theory, but we have not yet found all the dyons. We would expect another set of dyons, $(2k, -1)$ which decompose as $k\gamma_1 + (k+1)\gamma_2$ for $k \geq 0$. These would all lie to the right of the W boson, $\gamma_1 + \gamma_2$; thus we need some method for exploring that region of the \mathcal{Z} -plane.

In this case, and in any case where there is only a single accumulation ray, we can get around this problem easily. To do so, we recall that our mutation rule came from rotating the half-plane clockwise, $\mathcal{H} \rightarrow e^{-i\theta}\mathcal{H}$. We'll refer to this as left-mutation, because it is associated with states rotating out of the left of \mathcal{H} . There should of course be a similar mutation rule corresponding to rotating the half-plane counter-clockwise instead, $\mathcal{H} \rightarrow e^{i\theta}\mathcal{H}$, which we will call right-mutation. Both of these rules can be expressed as an action of a linear operator on the set of charges γ_i which label the nodes of the quiver. If we call the usual left-mutation action on charges M_L , and the right-mutation action M_R , then we should have the obvious relations

$$M_L M_R = M_R M_L = Id_{\{\gamma\}} \quad (2.4.4)$$

One can check that the transformation which satisfies the above identities (for γ_1 rotating out of \mathcal{H}) is simply

$$\tilde{\gamma}_1 = -\gamma_1 \quad (2.4.5)$$

$$\tilde{\gamma}_j = \begin{cases} \gamma_j + (\gamma_1 \circ \gamma_j)\gamma_1 & \text{if } \gamma_1 \circ \gamma_j > 0 \\ \gamma_j & \text{if } \gamma_1 \circ \gamma_j \leq 0. \end{cases} \quad (2.4.6)$$

Pictorially, mutation to the left (on node 1) acts non-trivially on those nodes which 1 points to, while right mutation acts non-trivially on nodes which point to 1. With this new rule in hand, we can start with the quiver and begin mutating “to the right”. Then we'll explore the BPS

states starting from the right side of \mathcal{H} , as these are the ones leaving the half-plane. If there is only a single accumulation ray in \mathcal{H} , left and right-mutation together will allow us to explore both sides of it, filling out the entire \mathcal{Z} -plane except for the ray of the accumulation point.

Let's apply right mutation starting from the original $SU(2)$ quiver to find the remaining states. Now we use \otimes to indicate the right-most node which will be right-mutated next.

$$\begin{array}{ccc}
 \begin{array}{c} \circ \xrightarrow{\quad\quad\quad} \otimes \\ \gamma_1 \qquad\qquad\quad \gamma_2 \end{array} & \begin{array}{c} \otimes \xleftarrow{\quad\quad\quad} \circ \\ \gamma_1 + 2\gamma_2 \qquad\quad -\gamma_2 \end{array} & \begin{array}{c} \circ \xrightarrow{\quad\quad\quad} \otimes \\ -\gamma_1 - 2\gamma_2 \qquad\quad 2\gamma_1 + 3\gamma_2 \end{array} \\
 \text{(i)} & \text{(ii)} & \text{(iii)} \\
 & & (2.4.7) \\
 \dots & \begin{array}{c} k\gamma_1 + (k+1)\gamma_2 \\ \otimes \xleftarrow{\quad\quad\quad} \circ \\ -(k-1)\gamma_1 - k\gamma_2 \\ \text{(k+1)} \end{array} & \dots
 \end{array}$$

We have generated the states $k\gamma_1 + (k+1)\gamma_2 = (2k, -1)$. So mutation to the right obtains the dyons that we didn't see before, namely the ones lying on the other side of the vector. Since these states limit to the same ray in the \mathcal{Z} plane, at $\mathcal{Z}(\gamma_1 + \gamma_2)$, we have understood the stability of all states except those lying on this ray. To complete the analysis, in principle one should do the representation theory for states along the ray $\gamma_1 + \gamma_2$. At a generic choice of parameters, the only particles that may exist along this ray are of the form $n(\gamma_1 + \gamma_2)$.¹³ It turns out that there is indeed a single vector present with the expected charge. This seems slightly obnoxious, because we still have to do some representation theory, but keep in mind that the work has been drastically reduced in that we only have to check for representations along this ray.

To summarize, we have found the strong coupling $SU(2)$ spectrum by a completely trivial

¹³ This statement heavily relies on the fact that this theory is complete. If the central charges of nodes cannot be varied independently, and the theory is thus incomplete, then there are non-trivial relations satisfied by the central charges of nodes at all points of parameter space. For example, there may be a relation of the form $\gamma_k = \gamma_1 + \gamma_2$, satisfied for all parameter choices. Then the general particle at the ray $\mathcal{Z}(\gamma_1 + \gamma_2)$ is of the form $n(\gamma_1 + \gamma_2) + m\gamma_k$. We will see how this may come about in section 2.4.3.

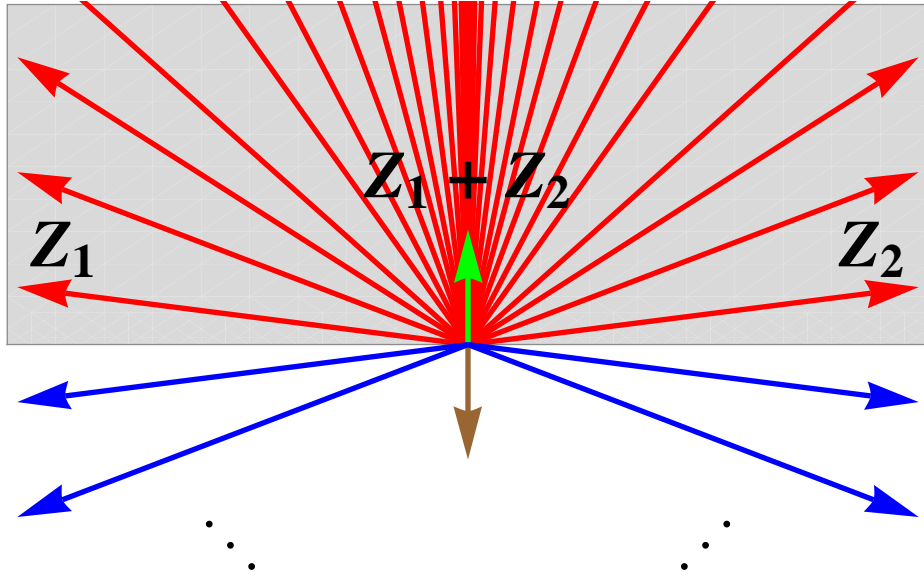


Figure 2.5: The BPS spectrum of pure $SU(2)$ gauge theory, plotted in the central charge \mathcal{Z} -plane. The spectrum contains a vector state with charge $\mathcal{Z}_1 + \mathcal{Z}_2$ (plotted in green), which is forced to occur in the \mathcal{Z} -plane at an accumulation ray of hypermultiplet states. On either side of the vector state, there is an infinite sequence of dyons whose central charges asymptotically approach the ray on which the vector lies. The mutation method is able to capture the full spectrum of the theory by rotating the half-plane to the left (yielding particles on the left of the vector particle) and the to right (yielding particles on the right of the vector particle).

application of the mutation method. For the weak coupling chamber, we introduced *right-mutation* to be able to explore the central charge \mathcal{Z} -plane on both sides of the accumulation ray at the W boson. Here we found, as expected, the W boson and the infinite tower of dyons. In Figure 2.5, we draw the spectrum in the \mathcal{Z} -plane to clarify how the mutation method is capable of obtaining all states of the theory. The well-known resulting spectra are summarized in the table below:

Strong coupling		Weak coupling
Monopole: $(0, 1)$		Positive dyons: $(2n, 1)$
Dyon: $(2, -1)$		Negative dyons: $(2n + 2, -1)$
		W boson: $(2, 0, 0)$

2.4.2 Adding matter

The quiver of $SU(2)$ $N_f = 1$ was deduced using general considerations in [12]. Here we simply recall their reasoning. We expect $2r + f = 3$ nodes of the quiver. First we note that we can tune the mass of the quark to infinity. Then the massive quark fields should decouple from the theory, leaving the BPS states of pure $SU(2)$. This suggest that the quiver should consist of the pure $SU(2)$ quiver (with the usual monopole and dyon charges) along with an additional node. In the decoupled limit, there should be additional states with (e, m) charges $(\pm 1, 0)$; the third node should correspond to one of these two charges. However, we need to make the correct choice for third node that allows *both* of these new states to be generated by positive linear combinations of the nodes. If we take $(1, 0)$, all nodes of the quiver have positive electric charge, and the state $(-1, 0)$ cannot be generated; the correct choice is then $(-1, 0)$, which can be combined with the W boson $(2, 0)$ of the $SU(2)$ subquiver to form $(1, 0)$. Computing electric-magnetic inner products, we find the following quiver:

$$\begin{array}{ccc}
 & \gamma_3 = (-1, 0) & \\
 & \circ & \\
 & \swarrow \quad \searrow & \\
 \gamma_1 = (0, 1) & \circ \rightleftarrows \circ & \gamma_2 = (2, -1)
 \end{array} \tag{2.4.8}$$

We can repeat this argument to add as many additional flavors as we like; the result is to produce N_f copies of the node γ_3 with different flavor charges.

$$(2.4.9)$$

Alternatively, we can add hypermultiplet matter charged under other representations of the gauge group. If instead of a fundamental $\mathbf{2}$ of $SU(2)$ we consider a \mathbf{j} rep of $SU(2)$, we find that γ_3 has charge $(-j, 0)$. Generalizing our analysis above, we conclude that if a quiver description of this theory exists, it is given by a similar quiver with j arrows $\gamma_3 \rightarrow \gamma_1, \gamma_2 \rightarrow \gamma_3$

$$(2.4.10)$$

Certainly this quiver can generate the full \mathbf{j} representation, raising the electric charge by adding the W boson. However, for $\mathbf{j} \neq \mathbf{2}$, it is possible that the quiver generates some additional representations of the gauge group. Indeed, it turns out that such a quiver will correspond to $SU(2)$ with a full $\otimes^j \mathbf{2}$ representation of the gauge group. We will see an explicit occurrence of this in section 2.4.8, where for $j = 2$ the quiver above (2.4.10) produces the matter representation $\mathbf{3} \oplus \mathbf{1}$.

2.4.3 Massless $N_f = 1$

Recall that a single quiver from the mutation class generally does not cover all of moduli space. If we start with a valid quiver description and move in moduli space, it may be that at some point the central charges $\mathcal{Z}(\gamma_i)$ no longer lie in a common half-plane. We deduced the $SU(2)$ with matter quivers in the decoupling limit of infinite quark mass, so there is no reason to expect it to cover the chamber with the bare mass of the quark set to zero. Actually, one can easily see that

the massless chamber should have $\mathcal{Z}(\gamma_1 + \gamma_2 - 2\gamma_3) = 0$ for the charges given in (2.4.3). Thus we have $\mathcal{Z}(\gamma_3) = -\frac{1}{2}\mathcal{Z}(\gamma_1) + \mathcal{Z}(\gamma_2)$. There is no way that the three central charges can lie in a single half-plane.

It is easy to remedy this situation by properly applying mutations. Imagine beginning at a point of parameter space where the quiver above is valid. Then we consider tuning parameters until we reach the desired point. As we do this, we should keep track of any states leaving the half-plane, and perform the appropriate mutations. This sounds as though it involves detailed knowledge of the moduli space geometry, but that turns out to be completely unnecessary. There is no need to restrict our path to the physical parameter space; instead we are free to move throughout full space of central charges for the theory. In other words, we are free to pretend that the theory is complete as we tune parameters.¹⁴ This drastically simplifies the procedure. Now we may start with a valid quiver at a certain choice of parameters, and then tune the central charges one-by-one to produce the arrangement at the desired endpoint in parameter space.

For the $N_f = 1$ quiver (2.4.3), let's keep γ_1, γ_2 fixed in the central charge plane, and tune γ_3 from its initial value within the half-plane by rotating it to the right. It will exit on the right, inducing a right-mutation on γ_3 . We should continue rotating γ_3 all the way to $\mathcal{Z}(\gamma_3) = -\mathcal{Z}(\gamma_1 + \gamma_2)$, and keep track of mutations of the charges of the mutated quiver. In this case, no additional mutations occur. This gives¹⁵

$$\begin{array}{ccc}
 & \gamma_3 = (1, 0, -1) & \\
 & \circ & \\
 & \nearrow \quad \searrow & \\
 \circ & \xrightarrow{\quad} & \circ \\
 \gamma_1 = (0, 1, 1/2) & & \gamma_2 = (1, -1, 1/2)
 \end{array} \tag{2.4.11}$$

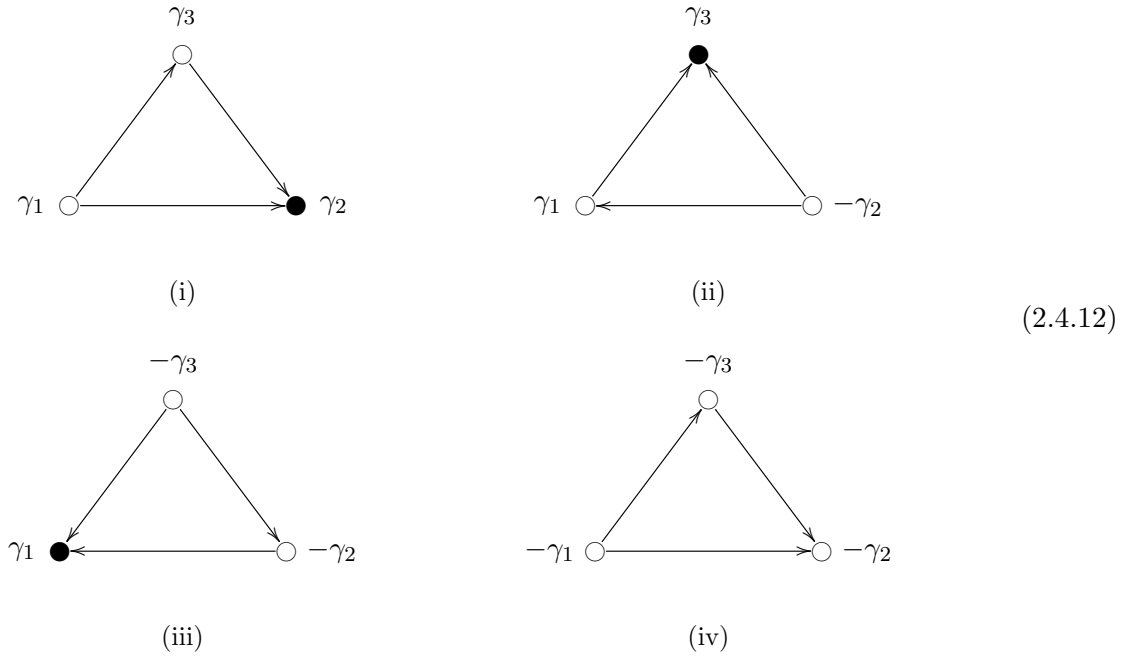
The flavor group for $N_f = 1$ is $SO(2) \cong U(1)$, so we label the charges of our states by their $U(1)$

¹⁴This theory actually *is* complete; however, in any other non-complete examples, the same approach is valid.

¹⁵The monopole and dyon acquire flavor charges [10], which we now include in the charge labels.

charge f ; the nodes then correspond to the electromagnetic and flavor charges (e, m, f) as given above. At zero bare mass, the central charge function only depends on the electric and magnetic charges of the states, so the third node is constrained as $\mathcal{Z}(\gamma_3) = \mathcal{Z}(\gamma_1) + \mathcal{Z}(\gamma_2)$. Thus, just as in the pure $SU(2)$ theory, there are only two distinct chambers, one with $\arg \mathcal{Z}(\gamma_1) > \arg \mathcal{Z}(\gamma_2)$, and the other with $\arg \mathcal{Z}(\gamma_2) > \arg \mathcal{Z}(\gamma_1)$. This will turn out to be a feature of all the massless examples we consider.

Let's start by exploring the chamber with $\mathcal{Z}(\gamma_2)$ ahead of $\mathcal{Z}(\gamma_1)$. We start by mutating on γ_2 , after which we have the nodes γ_1, γ_3 and $-\gamma_2$. γ_3 is now left most, so we must mutate on it next, and so on.



We see the only states in this chamber were the nodes of the original quiver and their antiparticles. We've discovered the strong coupling chamber of the $N_f = 1$ theory, whose spectrum indeed coincides with these hypermultiplets.

Now let's explore the other chamber. Here we take $\mathcal{Z}(\gamma_1)$ ahead of $\mathcal{Z}(\gamma_2)$. We have the

sequence

(i)

(ii)

(2.4.13)

(iii)

(iv)

We're clearly in an infinite chamber. Continuing in this way, we see our spectrum includes the states

$$(n+1)\gamma_1 + n(\gamma_2 + \gamma_3) = (2n, 1, 1/2)$$

$$(n+1)(\gamma_1 + \gamma_3) + n\gamma_2 = (2n+1, 1, -1/2)$$

As in the weak chamber of the pure $SU(2)$ theory, we are seeing the accumulation ray which should contain the W boson of the theory. Here we are actually getting twice as many hypermultiplets as in pure $SU(2)$ since we have states of both even and odd electric charge. We will identify the odd electric charge states as quark-dyon bound states.

As before, let's start with the original quiver and mutate to the right to study the BPS states on the other side of the accumulation ray. This generates the states

$$n(\gamma_1 + \gamma_3) + (n+1)\gamma_2 = (2n+1, -1, 1/2)$$

$$n\gamma_1 + (n+1)(\gamma_2 + \gamma_3) = (2n+2, -1, -1/2)$$

This sequence of states also accumulate at the same ray in the central charge plane; between these two sequences of infinities, the only central charges that can appear are proportional to $\mathcal{Z}(\gamma_1 + \gamma_2)$. We might again expect that these dyons limit to a single vector in the central charge plane. We could attempt to test this hypothesis by actually doing the representation theory along this ray, but instead let's appeal to some physical reasoning to see why this is indeed wrong. Namely, we're in the weak coupling chamber of the $N_f = 1$ theory. We would expect that this theory indeed contains BPS states corresponding to the fundamental quark hypermultiplet, and at zero bare mass the central charge of this hyper lies directly at the same BPS phase as the W boson. This is precisely the non-generic situation we hinted at in footnote 13.

Actually, given the non-genericity, something special has happened in this example. This quarks, given by γ_3 and by $\gamma_1 + \gamma_2$, appeared as nodes after a finite sequence of mutations. Note that we never mutated on these quark nodes, because the nodes we mutate on are left-most (or right-most) and being on an accumulation ray, the quark can never be made left-most (or right-most). Instead, they simply appeared as one of the other “interior” nodes in some of the dual quiver descriptions of the theory. This doesn't have to happen, and indeed won't happen in the undeformed $N_f = 2, 3$ cases below. We simply got lucky. If we hadn't seen the quark this way, we would have had to find it by hand. In either case, how can we be sure there are no other hypermultiplets lying on top of the vector, which aren't showing up as interior nodes elsewhere? One should consider a slightly deformed $N_f = 1$ theory with $m \neq 0$ and check that there are no additional hypermultiplets (aside from those predicted by wall-crossing formulae). In this way, one can check that there are no additional hypermultiplets coinciding with the vector when $m \rightarrow 0$. In principle, it is irrelevant whether or not the deformation we take is physically realized - thus, even in a non-complete theory, the same strategy works for understanding the particles along an accumulation ray. Alternatively, of course, one could always directly use quiver representation theory to rule out other states with that BPS phase.

Putting everything together, we find the following possible spectra for massless $N_f = 1$.

Strong coupling	Weak coupling
Quark: $(1, 0, -1)$	Quarks: $(1, 0, \pm 1)$
Monopole: $(0, 1, 1/2)$	Positive dyons: $(2n, 1, 1/2)$
Dyon: $(1, -1, 1/2)$	Negative dyons: $(2n + 2, -1, -1/2)$
	Quark-dyons: $(2n + 1, \mp 1, \pm 1/2)$
	W boson: $(2, 0, 0)$

where n ranges over integers $n \geq 0$. Along with their antiparticles, this collection agrees with the well known weak coupling spectrum of massless $SU(2)$ $N_f = 1$ ([64]).

2.4.4 Massive $N_f = 1$

For just one flavor, it is not too difficult to actually find all possible spectra of the theory with $m \neq 0$. It turns out that the acyclic quiver used in the previous subsection covers all chambers. Unfortunately, there is a great deal of redundancy in the full chamber spectrum - there are many distinct regions of moduli space that give the same spectrum due to dualities. By duality here, we mean the following: the spectrum depends only on the quiver and the central charges decorating the nodes, but not on the actual charge (e, m, f) labels themselves. Thus, there may be widely separated regions of moduli space that happen to have the same quiver and associated central charges, but different charge labels; consistency of this framework requires that such regions actually have spectra that are equivalent up to some appropriate $Sp(2r, \mathbb{Z})$ relabeling of charges. Here we will simply list the possible spectra, without choosing a particular point in moduli space or duality frame; the downside is that as a result we cannot give the charges of the states, since charge labels require a choice of duality frame.

- Minimal chamber: 3 nodes are the only BPS states.
- 4 state chamber: 3 nodes and 1 bound state hypermultiplet.
- 5 state chamber: 3 nodes and 2 bound state hypermultiplet.

- Weak coupling chambers, labelled by k . These consist of:
 - 2 quark hypermultiplets,
 - W boson vector multiplet,
 - Infinite tower of dyon hypermultiplets,
 - k additional quark-dyon bound state hypermultiplets, for $0 \leq k \leq \infty$.

This list exhausts the BPS spectra that can be supported by quivers in this mutation class. Embedded in this result are the two massless chambers, which correspond to the 3 state minimal chamber and the $k = \infty$ weak coupling chamber. It is a relatively straight-forward exercise to find all these chambers beginning with the minimal massless spectrum, by repeated application of the pentagon and $SU(2)$ wall-crossing identities.

2.4.5 Massless $N_f = 2$

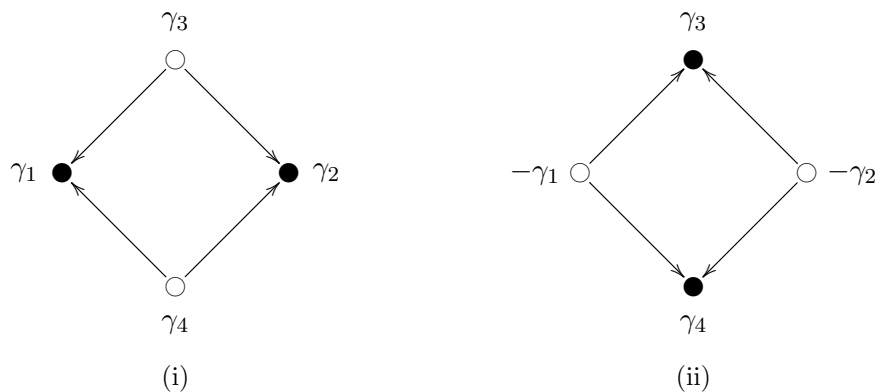
The relevant quiver for massless $N_f = 2$ follows from analogous mutations of the decoupling limit quiver (2.4.9) in section 2.4.2. Here we find

$$\begin{array}{c}
 \gamma_3 = (0, 1, 1/2, 0) \\
 \circ \\
 \swarrow \quad \searrow \\
 \gamma_1 = (1, -1, 0, -1/2) \quad \circ \quad \gamma_2 = (1, -1, 0, 1/2) \\
 \nwarrow \quad \nearrow \\
 \circ \\
 \gamma_4 = (0, 1, -1/2, 0)
 \end{array} \tag{2.4.14}$$

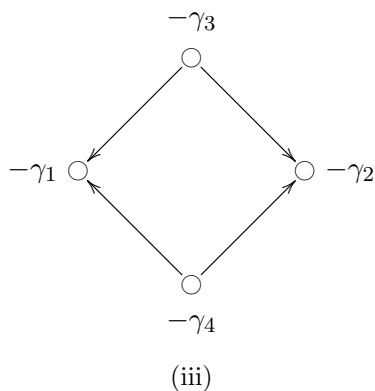
The flavor group is now $Spin(4) \cong SU(2) \times SU(2)$, and we will denote our states by (e, m, f_1, f_2) , where f_i are the charges under the $U(1)$ contained in the i th $SU(2)$ factor. We see that there are only two distinct values for the central charge between the four nodes when the bare masses vanish. This means that there will again only be two chambers, given by the relative ordering of $\mathcal{Z}(\gamma_1) = \mathcal{Z}(\gamma_2), \mathcal{Z}(\gamma_3) = \mathcal{Z}(\gamma_4)$.

There is a small added subtlety that was absent for $N_f = 1$. Namely, we technically can't rotate the central charge of a single node out of the half plane by itself. All mutations will happen for two nodes simultaneously. Also, as mentioned above, we don't get lucky in this example - the quarks don't show up as interior nodes of any of the quivers as we start mutating. If we mass deform the theory, however, the central charge of the quarks no longer coincides with the vector, and we will see them appear after a finite number of mutations. This tells us that there are the quark hypermultiplets lying on top of the vector when $m \rightarrow 0$, but no extra states. For simplicity, we will work out the $m = 0$ point and quote this result.

For strong coupling, we first mutate on nodes 1 and 2, and find



(2.4.15)



Thus we see that this chamber contains no bound states, and the only states are hypermultiplets contributed by the nodes. We have one hypermultiplet of electromagnetic charge $(1, -1)$

in the $(\mathbf{1}, \mathbf{2})$ rep of $SU(2) \times SU(2)$, and one of charge $(0, 1)$ in the $(\mathbf{2}, \mathbf{1})$.

The other chamber is of course more interesting. We have the following sequence of mutations:

$$\begin{array}{ccc}
 \begin{array}{c} \gamma_3 \\ \bullet \\ \swarrow \quad \searrow \\ \gamma_1 \circ \quad \gamma_2 \circ \\ \nwarrow \quad \nearrow \\ \bullet \\ \gamma_4 \end{array} & & \begin{array}{c} -\gamma_3 \\ \circ \\ \swarrow \quad \searrow \\ \gamma_1 + \gamma_3 + \gamma_4 \bullet \quad \gamma_2 + \gamma_3 + \gamma_4 \bullet \\ \nwarrow \quad \nearrow \\ \circ \\ -\gamma_4 \end{array} \\
 \text{(i)} & & \text{(ii)} \\
 & & \text{(2.4.16)} \\
 & & \begin{array}{c} \gamma_1 + \gamma_2 + \gamma_3 + 2\gamma_4 \\ \bullet \\ \swarrow \quad \searrow \\ -\gamma_1 - \gamma_3 - \gamma_4 \circ \quad -\gamma_2 - \gamma_3 - \gamma_4 \circ \\ \nwarrow \quad \nearrow \\ \bullet \\ \gamma_1 + \gamma_2 + 2\gamma_3 + \gamma_4 \end{array} \\
 & & \text{(iii)}
 \end{array}$$

Continuing in this way, we generate the states

$$n(\gamma_1 + \gamma_2 + \gamma_4) + (n + 1)\gamma_3 = (2n, 1, 1/2, 0)$$

$$(n + 1)\gamma_4 + n(\gamma_1 + \gamma_2 + \gamma_3) = (2n, 1, -1/2, 0)$$

$$(n + 1)(\gamma_1 + \gamma_3 + \gamma_4) + n\gamma_2 = (2n + 1, 1, 0, -1/2)$$

$$(n + 1)(\gamma_2 + \gamma_3 + \gamma_4) + n\gamma_1 = (2n + 1, 1, 0, 1/2).$$

On the other hand, mutating to the right gives the states

$$n(\gamma_1 + \gamma_3 + \gamma_4) + (n+1)\gamma_2 = (2n+1, -1, 0, 1/2)$$

$$n(\gamma_2 + \gamma_3 + \gamma_4) + (n+1)\gamma_1 = (2n+1, -1, 0, -1/2)$$

$$(n+1)(\gamma_1 + \gamma_2 + \gamma_4) + n\gamma_3 = (2n+2, -1, -1/2, 0)$$

$$n\gamma_4 + (n+1)(\gamma_1 + \gamma_2 + \gamma_3) = (2n+2, -1, 1/2, 0).$$

These fill out dyons $(2n, \pm 1)$ in the $(\mathbf{2}, \mathbf{1})$ and quark-dyons $(2n+1, \pm 1)$ in the $(\mathbf{1}, \mathbf{2})$. Trapped between the two infinite sequences we have the vector boson $\gamma_1 + \gamma_2 + \gamma_3 + \gamma_4 = (2, 0, 0, 0)$, which we identify as the W . The quarks also lie at the same BPS phase, and are given by $\gamma_2 + \gamma_4, \gamma_1 + \gamma_4, \gamma_2 + \gamma_3, \gamma_1 + \gamma_3$.

The two spectra are tabulated below, where we now assemble the states into representations of the full $SU(2) \times SU(2)$ with charges given as $(e, m)_{\mathbf{f}_1, \mathbf{f}_2}$:

Strong coupling	Weak coupling
Monopole: $(0, 1)_{\mathbf{2}, \mathbf{1}}$	Quarks: $(1, 0)_{\mathbf{2}, \mathbf{2}}$
Dyon: $(1, -1)_{\mathbf{1}, \mathbf{2}}$	Positive dyons: $(2n, 1)_{\mathbf{2}, \mathbf{1}}$
	Negative dyons: $(2n+2, -1)_{\mathbf{2}, \mathbf{1}}$
	Quark-dyons: $(2n+1, \pm 1)_{\mathbf{1}, \mathbf{2}}$
	W boson: $(2, 0)_{\mathbf{1}, \mathbf{1}}$

This agrees with the well known weak coupling spectrum of the $SU(2)$ $N_f = 2$ theory.

2.4.6 Massless $N_f = 3$

The $N_f = 3$ quiver is given, after mutations to reach the massless chamber, as

$$\begin{array}{c}
 \gamma_2 = (0, 1, -1, 1, 0) \\
 \circ \\
 \downarrow \\
 \gamma_5 = (0, 1, 1, 0, 0) \circ \longrightarrow \circ \longleftarrow \gamma_3 = (0, 1, 0, -1, 1) \\
 \uparrow \\
 \gamma_1 = (1, -2, 0, 0, 0) \\
 \downarrow \\
 \circ \\
 \gamma_4 = (0, 1, 0, 0, -1)
 \end{array} \tag{2.4.17}$$

The flavor group is $SO(6) \cong SU(4)$ and the nodes of the quiver make up a monopole of electric/magnetic charge $(0, 1)$ in the $\mathbf{4}$ of $SU(4)$, and a dyon of charge $(1, -2)$ in the $\mathbf{1}$. We have labelled the flavor charges as (e, m, q_1, q_2, q_3) , where q_i are the eigenvalues under the respective generators of the Cartan of $SU(4)$. The central charge degeneracy we experienced in the $N_f = 2$ case is again present, among γ_i for $2 \leq i \leq 5$. Half the spectrum will come as sets of 4 simultaneous mutations.

There are again two chambers, one with $\arg \mathcal{Z}(\gamma_5) > \arg \mathcal{Z}(\gamma_1)$, and the other with $\arg \mathcal{Z}(\gamma_1) > \arg \mathcal{Z}(\gamma_5)$. The second chamber is strong coupling, and just includes the particles that correspond to the original nodes of the quiver. In the other chamber, the mutations generate the spectrum

$$\begin{aligned}\gamma_i + n(\gamma_2 + \gamma_3 + \gamma_4 + \gamma_5) + 2n\gamma_1 &= (2n, 1, 1, 0, 0) \\ (n+1)(\gamma_2 + \gamma_3 + \gamma_4 + \gamma_5) + (2n+1)\gamma_1 &= (2n+1, 2, 0, 0, 0) \\ -\gamma_i + (n+1)(\gamma_2 + \gamma_3 + \gamma_4 + \gamma_5) + (2n+1)\gamma_1 &= (2n+1, 1, -1, 0, 0)\end{aligned}$$

The states in which γ_i appears are repeated for $1 \leq i \leq 4$. Thus we see that we have a magnetic charge 2 dyon that is a singlet under flavor $SU(4)$, as well as magnetic charge 1 dyons in the $\bar{\mathbf{4}}$ and quark-dyons in the $\mathbf{4}$.

As usual, the mutations to the right will fill out the dyons on the other side of the accumulation ray. Right mutation generates:

$$\begin{aligned}n(\gamma_2 + \gamma_3 + \gamma_4 + \gamma_5) + (2n+1)\gamma_1 &= (2n+1, -2, 0, 0, 0) \\ \gamma_i + n(\gamma_2\gamma_3 + \gamma_4 + \gamma_5) + (2n+1)\gamma_1 &= (2n+1, -1, 1, 0, 0) \\ -\gamma_i + (n+1)(\gamma_2 + \gamma_3 + \gamma_4 + \gamma_5) + (2n+2)\gamma_1 &= (2n+2, -1, -1, 0, 0)\end{aligned}$$

The vector W boson, is at an accumulation ray, and the subtlety about generating the quarks is the same as in the $N_f = 2$ case. Here the quarks are given by $\gamma_1 + \gamma_i + \gamma_j$, where $2 \leq i < j \leq 5$.

Strong coupling	Weak Coupling
Monopole: $(0, 1)_4$	Quarks: $(1, 0)_6$
Dyon: $(1, -2)_1$	Positive dyons: $(2n, 1)_4$
	Negative dyons: $(2n + 2, -1)_4$
	$m = 2$ dyons: $(2n + 1, \pm 2)_1$
	Quark-dyons: $(2n + 1, -1)_4$
	$(2n + 1, 1)_4$
	W boson: $(2, 0)_1$

2.4.7 $N_f = 4$

For $N_f = 4$ the massless theory is conformal; mass deformations break conformality. The quiver in the decoupling $m \rightarrow \infty$ limit is given as¹⁶

$$(2.4.18)$$

There are many additional subtleties in this BPS spectrum because it corresponds to a massive deformation of the conformal theory. In particular, there is no quiver that describes the $m \rightarrow 0$ limit; if we try to follow the strategy employed in the asymptotically free cases to trace the quiver from $m = \infty$ to $m = 0$, we find that any path goes through infinitely many mutations, preventing us from identifying a quiver for the $m = 0$ chamber.

Nonetheless, we may take a finite mass and find various chambers in which the mutation method can successfully compute BPS spectra. The following is an example of a finite chamber of this theory, with the BPS states listed in decreasing order of BPS phase:

$$\gamma_3, \gamma_4, \gamma_2, \gamma_1 + \gamma_3 + \gamma_4, \gamma_2 + \gamma_5, \gamma_2 + \gamma_6, \gamma_1 + \gamma_3, \gamma_1 + \gamma_4, \gamma_2 + \gamma_5 + \gamma_6, \gamma_1, \gamma_5, \gamma_6. \quad (2.4.19)$$

¹⁶Our analysis will break the $SO(8)$ flavor symmetry, so we suppress all flavor data.

This theory is complete, so, as previously discussed, this chamber must occur in physical moduli space.

In principle, the BPS spectrum can be worked out in all of moduli space by applying the KS wall crossing formula to this chamber. However, the spectrum in some regions of moduli space becomes extremely complicated. To give a general sense of this, we will describe some wall crossings in this theory, which were first studied in [13].

Focus on the first three states, $\gamma_3, \gamma_4, \gamma_2$. If we move γ_2 all the way to the left, we will produce $\gamma_2, \gamma_2 + \gamma_3, \gamma_2 + \gamma_4, \gamma_2 + \gamma_3 + \gamma_4, \gamma_3, \gamma_4$. Separating the rest of the spectrum into similar consecutive sets of three, analogous wall crossings will produce a spectrum of 24 states.

$$\begin{aligned}
&\gamma_2, \gamma_2 + \gamma_3, \gamma_2 + \gamma_4, \gamma_2 + \gamma_3 + \gamma_4, \gamma_3, \gamma_4, \\
&\gamma_2 + \gamma_5, \gamma_2 + \gamma_6, \gamma_1 + 2\gamma_2 + \gamma_3 + \gamma_4 + \gamma_5 + \gamma_6, \gamma_1 + \gamma_2 + \gamma_3 + \gamma_4 + \gamma_6, \\
&\quad \gamma_1 + \gamma_2 + \gamma_3 + \gamma_4 + \gamma_5, \gamma_1 + \gamma_3 + \gamma_4, \\
&\gamma_2 + \gamma_5 + \gamma_6, \gamma_1 + \gamma_2 + \gamma_3 + \gamma_5 + \gamma_6, \gamma_1 + \gamma_2 + \gamma_4 + \gamma_5 + \gamma_6, \\
&\quad 2\gamma_1 + \gamma_2 + \gamma_3 + \gamma_4 + \gamma_5 + \gamma_6, \gamma_1 + \gamma_3, \gamma_1 + \gamma_4, \\
&\gamma_5, \gamma_6, \gamma_1 + \gamma_5 + \gamma_6, \gamma_1 + \gamma_5, \gamma_1 + \gamma_6, \gamma_1.
\end{aligned} \tag{2.4.20}$$

Now we can produce various vectors by crossing states between the four sets of six; for example, $(\gamma_1 + \gamma_3 + \gamma_4) \circ (\gamma_2 + \gamma_5 + \gamma_6) = -2$, so exchanging them will produce a tower of dyons and a vector $\gamma_1 + \gamma_2 + \gamma_3 + \gamma_4 + \gamma_5 + \gamma_6 = (-2, 0)$, by the $SU(2)$ wall crossing identity. Similarly, exchanging γ_3, γ_4 with $\gamma_2 + \gamma_5, \gamma_2 + \gamma_6 =$ will produce a vector $2\gamma_2 + \gamma_3 + \gamma_4 + \gamma_5 + \gamma_6 = (-4, 2)$ along with two dyon towers and four additional hypers; this is just the wall crossing of massless $SU(2), N_f = 2$. Two more vectors will be generated by this procedure, $2\gamma_1 + \gamma_3 + \gamma_4 + \gamma_5 + \gamma_6 = (0, -2)$ and $\gamma_1 - \gamma_2 = (2, -2)$.¹⁷ These four vectors have non-trivial electric-magnetic inner products, and so additional wall crossing of the vectors will produce some highly complicated spectrum with infinitely

¹⁷To obtain this last vector, we must rotate the half-plane, allowing γ_2 to exit and mutating on γ_2 in the quiver.

many vectors.

One would expect such wild BPS behavior in the massless conformal limit, where conformal dualities produce some infinite set of vectors dual to the familiar W boson. It is interesting to observe that this complicated structure begins to emerge even with finite mass, in regions of moduli space where the quiver description is perfectly valid.

2.4.8 $\mathcal{N} = 2^*$

The $\mathcal{N} = 2^*$ theory is a massive deformation of conformal $\mathcal{N} = 4$, where we give the adjoint hypermultiplet some non-zero mass. Alternatively, it is simply a gauge theory with a massive hypermultiplet charged under the adjoint of the gauge group. For $SU(2)$ this is given, following the discussion in section 2.4.2, by the following quiver:

$$\begin{array}{c}
 \gamma_3 = (-2, 0) \\
 \circ \\
 \swarrow \quad \searrow \\
 \circ \quad \quad \quad \circ \\
 \gamma_1 = (0, 1) \quad \quad \quad \gamma_2 = (2, -1)
 \end{array}
 \tag{2.4.21}$$

As indicated in section 2.4.2, this quiver indeed turns out to generate matter content of the full $\mathbf{2} \otimes \mathbf{2} = \mathbf{3} \oplus \mathbf{1}$. Thus it gives the $\mathcal{N} = 2^*$ theory plus an uncharged singlet hypermultiplet. In [28], this quiver was obtained in studying the rank two Gaiotto theory on a torus with one puncture. We can understand this matter content from the point of view of [36]. We start with a pair of pants, corresponding to a half-hypermultiplet charged as a trifundamental under three $SU(2)$ flavor groups, represented by the three boundary components. Glueing together two boundary components of the pair of pants identifies the two $SU(2)$'s and gauges them. To form the punctured torus, we glue two legs together, producing an $SU(2)$ gauge group, and matter content $\mathbf{2} \otimes \mathbf{2} = \mathbf{3} \oplus \mathbf{1}$.

This fact can be checked from the BPS spectrum as follows. Consider the rep $\gamma_1 + \gamma_2 + \gamma_3$

of this quiver. This rep has charge $(0, 0)$ meaning that it is a pure flavor state. For $\mathcal{N} = 2^*$ we would expect such a hypermultiplet, corresponding to the state inside the $\mathbf{3}$ that is uncharged under the $U(1) \subset SU(2)$; if we add an uncoupled singlet, we would then expect this site of the charge lattice to be occupied by two BPS particles. Quiver representation theory finds the latter situation, as we now demonstrate.

The superpotential for this quiver was worked out in [28]. The result was

$$\mathcal{W} = X_{12}X_{23}X_{31} + Y_{12}Y_{23}Y_{31} + X_{12}Y_{23}X_{31}Y_{12}X_{23}Y_{31}. \quad (2.4.22)$$

Here, X_{ij}, Y_{ij} correspond to the two maps between nodes i, j in the representation. The resulting F-terms are of the form

$$X_{23}X_{31} + Y_{23}X_{31}Y_{12}X_{23}Y_{31} = 0, \quad (2.4.23)$$

$$X_{12}X_{23} + Y_{12}X_{23}Y_{31}X_{12}Y_{23} = 0, \quad (2.4.24)$$

$$X_{31}X_{12} + Y_{31}X_{12}Y_{23}X_{31}Y_{12} = 0, \quad (2.4.25)$$

$$Y_{23}Y_{31} + X_{23}Y_{31}X_{12}Y_{23}X_{31} = 0, \quad (2.4.26)$$

$$Y_{12}Y_{23} + X_{12}Y_{23}X_{31}Y_{12}X_{23} = 0, \quad (2.4.27)$$

$$Y_{31}Y_{12} + X_{31}Y_{12}X_{23}Y_{31}X_{12} = 0. \quad (2.4.28)$$

We are studying the rep $\gamma_1 + \gamma_2 + \gamma_3$, so all gauge groups are $U(1)$, and the bifundamental fields here are simply 1×1 matrices. In this example, we can solve the full equations by just truncating to the quadratic pieces and solving those, since setting the quadratic pieces to zero also sets the

quintic terms to zero.¹⁸

$$X_{23}X_{31} = 0 \tag{2.4.29}$$

$$X_{12}X_{23} = 0 \tag{2.4.30}$$

$$X_{31}X_{12} = 0 \tag{2.4.31}$$

$$Y_{23}Y_{31} = 0 \tag{2.4.32}$$

$$Y_{12}Y_{23} = 0 \tag{2.4.33}$$

$$Y_{31}Y_{12} = 0 \tag{2.4.34}$$

These will set two of the X 's and two of the Y 's equal to zero. We will focus on the two non-zero fields, X_i, Y_j , with $i, j \in \{(12), (23), (31)\}$. Before going on, we pause to consider what the possible moduli spaces may be. For any choice of i, j , there is enough gauge symmetry to set both X_i, Y_j to one; thus the moduli space is at most 9 points, one for each choice of (i, j) . Some of these points will be eliminated by the stability analysis. Note that II-stability does not distinguish between X, Y , so if $X_i, Y_j \neq 0$ is stable, then $X_j, Y_i \neq 0$ is also stable. We will show below that the stability analysis always yields a moduli space of 2 points.

The simplest way to proceed is a case-by-case analysis of the possible orderings of central charges. For each choice of orderings, we will consider the following cases of (i, j) : (a) (12, 23), (b) (23, 31), (c) (31, 12), (d) (12, 12), (e) (23, 23), (f) (31, 31). There are three more cases obtained by exchanging (i, j) . A simple study of commutative diagrams shows that, for (a) the subreps are $\gamma_3, \gamma_2 + \gamma_3$. By cyclic symmetry, (b) has subreps $\gamma_1, \gamma_3 + \gamma_1$, and for (c), $\gamma_2, \gamma_1 + \gamma_2$. For (d) we find subreps $\gamma_2, \gamma_3, \gamma_2 + \gamma_3, \gamma_1 + \gamma_2$; (e) and (f) have subreps given by cyclic symmetry. We can choose γ_1 to be the left-most node without loss of generality. Automatically, (e) and (f) are unstable due to the subrep γ_1 which has $\arg \mathcal{Z}(\gamma_1) > \arg \mathcal{Z}(\gamma_1 + \gamma_2 + \gamma_3)$. Suppose $\arg \mathcal{Z}(\gamma_1) > \arg \mathcal{Z}(\gamma_2) > \arg \mathcal{Z}(\gamma_3)$.

¹⁸There is also a solution given by nontrivial cancellation between the quadratic and quintic terms. However, the resulting moduli space is non-compact, so its cohomology contains no normalizable forms, and as such it does not contribute to the particle spectrum

Then rep (b) is destabilized by subrep γ_1 , and reps (c,d) are destabilized by subrep $\gamma_1 + \gamma_2$. Rep (a), on the other hand, is stable since its subreps have $\arg \mathcal{Z}(\gamma_1 + \gamma_2 + \gamma_3) > \arg \mathcal{Z}(\gamma_2 + \gamma_3) > \arg \mathcal{Z}(\gamma_3)$. So here the moduli space is 2 points, $X_{12}, Y_{23} \neq 0$ and $X_{23}, Y_{12} \neq 0$. Next, we consider $\arg \mathcal{Z}(\gamma_1) > \arg \mathcal{Z}(\gamma_1 + \gamma_2) > \arg \mathcal{Z}(\gamma_3) > \arg \mathcal{Z}(\gamma_2)$. Rep (a) is again stable, while rep (b) is destabilized by γ_1 and reps (c,d) are destabilized by $\gamma_1 + \gamma_2$. The final case we must study is $\arg \mathcal{Z}(\gamma_1) > \arg \mathcal{Z}(\gamma_3) > \arg \mathcal{Z}(\gamma_1 + \gamma_2) > \arg \mathcal{Z}(\gamma_2)$. Now we find that rep (c) is stable, while reps (a,d) are destabilized by γ_3 and rep (b) is destabilized by γ_1 . The conclusion is that the moduli space of the rep $(\gamma_1 + \gamma_2 + \gamma_3)$ is simply two points for any choice of parameters. Therefore, at all values in the parameter space of this theory, we find *two* hypermultiplets with no electric-magnetic charge. This confirms that the quiver is describing the Gaiotto construction, $\mathcal{N} = 2^*$ plus a single uncharged hypermultiplet.

The spectrum of this theory is extremely intricate for any chamber of the moduli space. We will demonstrate the existence of at least two vector particles for any choice of central charges. Without loss of generality, we take γ_1 to be leftmost. Then we should consider two cases. If $\arg \mathcal{Z}(\gamma_1) > \arg \mathcal{Z}(\gamma_2) > \arg \mathcal{Z}(\gamma_3)$, then the Π -stability analysis yields $\gamma_1 + \gamma_2 = (2, 0)$ and $\gamma_1 + \gamma_3 = (-2, 1)$ as stable vector particles. Alternatively, if $\arg \mathcal{Z}(\gamma_1) > \arg \mathcal{Z}(\gamma_3) > \arg \mathcal{Z}(\gamma_2)$, then $\gamma_1 + \gamma_2$ is a stable vector particle, along with either $(n + 1)\gamma_1 + n\gamma_2 + \gamma_3$ or $n\gamma_1 + (n + 1)\gamma_2 + \gamma_3$ for some choice of n . In any of the cases, the two vector particles identified have non-zero electric-magnetic inner product. Consequently, the stable vector states could form a highly complicated spectrum of bound states. The presence of multiple accumulation rays (one at each vector) obstructs the mutation method as defined from producing an unambiguous result for the spectrum. We can use left and right mutation to identify some set of dyons, along with the left-most and right-most vector states; however, the region of the \mathcal{Z} -plane between the two vectors could be arbitrarily wild. It would be interesting to try to develop an extension of the algorithm capable of computing the spectrum for this theory.

2.4.9 Flavor Symmetries and Gauging

The above $SU(2)$ examples involve a well-known $SO(2N_f)$ flavor symmetry at the massless point of parameter space. In fact, the quivers used in the analysis all display quite suggestive symmetries themselves. In this section we will study the relationship between global symmetries of the physical theory and discrete symmetries of the quiver. This will turn out to provide a powerful tool for constructing quivers for new theories by gauging global symmetries.

Suppose a physical theory has some known global symmetry. Generally speaking, turning on various deformations of the theory will break the global symmetry, so here we consider studying the theory at the precise point of parameter space that preserves the full global symmetry of interest. Of course, the BPS spectrum should reflect this symmetry. The first question we wish to explore is how this symmetry should be encoded in the BPS quiver.

It is possible that every state in the BPS spectrum might be singlet under the global symmetry; then it would be very difficult to find evidence for the symmetry in either the quiver or the full BPS spectrum. So we should refine the question a bit. Let us restrict to a global $SU(n)$ symmetry, and further, let us study the case in which there is some BPS hypermultiplet in the fundamental of $SU(n)$. In this case we can give a very straightforward answer to the question. The full fundamental multiplet of BPS states must have identical central charges. We simply choose our quiver half-plane so that this multiplet is left-most in the \mathcal{Z} -plane.¹⁹ Since they carry distinct flavor charges spanning the weight space, all n states of fundamental must occur in the quiver.²⁰ These

¹⁹This choice of half-plane will be impossible when the phase of central charge of the fundamental of hypermultiplets occurs at some accumulation ray of BPS states. In fact, this exact situation occurs in the case of $SU(2)$, $\mathcal{N} = 2^*$. This theory has an enhanced $SU(2)$ flavor symmetry at the massless point. However, we are never able to see the symmetry in the quiver (which has a single mutation form, given in section (2.4.21)). The massless theory is conformal, and the spectrum is dense; hence there is no half-plane that admits a positive integer basis. Barring this complication, there exists a half-plane that yields a mutation form of the quiver which explicitly presents the symmetry.

²⁰The weight space is only $n - 1$ -dimensional, so one may worry that only $n - 1$ of the states appear. However, the weights obey $\sum_i f_i = 0$ so that the last weight is given by a *negative* integer linear combination of the others. As long as the multiplet carries some non-zero electric-magnetic charge, the last state be linearly independent from the others. Then, to fill out the n states of the fundamental, all n states must appear in the quiver.

states of course have different global charges, but identical electric-magnetic charges. Since the quiver is only sensitive to electric-magnetic charges, we will find n identical nodes in the resulting quiver, and thus an S_n permutation symmetry that exchanges these identical nodes.

The above $SU(2)$ examples with massless matter illustrate this fact. For $N_f = 2$, we have an $SO(4) = SU(2) \times SU(2)$ flavor symmetry, which manifests as two S_2 discrete symmetries in the quiver, given by exchanging γ_1, γ_2 and γ_3, γ_4 . For $N_f = 3$, we have an $SO(6) = SU(4)$ symmetry, manifested as an S_4 on $\gamma_1, \gamma_2, \gamma_3, \gamma_4$. For $N_f = 4$, there should be a full $SO(8)$ flavor symmetry; however, it is only preserved at the massless conformal point, where we have no quiver description. For any mass deformation, the maximal symmetry is $SU(4)$, which corresponds to the obvious S_4 acting on $\gamma_3, \gamma_4, \gamma_5, \gamma_6$.

Alternatively, suppose we start with a quiver containing n identical nodes and an S_n symmetry. If we assign identical charges to these nodes, the resulting BPS spectra will be forced to organize into representations of $SU(n)$, because the quiver representation theory does not distinguish among these n identical nodes. The nodes themselves will form a multiplet in the fundamental representation, while bound states involving combinations of the identical nodes will form various tensor representations. Unfortunately, we cannot conclude from this that the full theory preserves this symmetry - perhaps it is preserved by the BPS states, but broken by some non-BPS states. Nonetheless, if we are expecting an $SU(n)$ global symmetry, it is quite natural to identify it with this discrete symmetry of the quiver.

From these observations, we can suggest a powerful rule for constructing quivers of new theories by gauging global symmetries of a theory with a known quiver. For now, let us focus on gauging a global $SU(2)$ symmetry that is manifested as an S_2 symmetry in the quiver acting on a pair of identical nodes. We will extend to general $SU(n)$ after we have discussed quivers of more general gauge theories. Physically, to gauge a symmetry, we add gauge degrees of freedom and couple them appropriately to the matter already present in the theory. At the level of the quiver,

the procedure is quite analogous. We should add two nodes of an $SU(2)$ subquiver to add the gauge degrees of freedom. Then we must couple to the existing pair of identical nodes to this subquiver to form a fundamental of the $SU(2)$. Recall that when we added a flavor to $SU(2)$, we added only one state of the doublet fundamental representation, because bound states would generate the second. Here we must do the same thing - we delete one of the nodes, and connect the other to the $SU(2)$ subquiver in an oriented triangle. The deleted state will now be generated by a bound state with the $SU(2)$ nodes.

To give an example, we can consider gauging one of the $SU(2)$ flavor symmetries of $SU(2)$, $N_f = 2$, which exchanges γ_1, γ_2 .

(2.4.35)

We have added an $SU(2)$ subquiver b, c and charged the flavor node γ_2 under it; now we have two $SU(2)$ gauge groups with a bifundamental matter field. In this case, we can actually see the weak coupling description of the resulting theory from the quiver, if we apply some mutations. Mutating on γ_1, γ_2, b, c in that order produces

(2.4.36)

in which there are two $SU(2)$ subquivers, each coupled to the same node as a fundamental matter state, producing a bifundamental.

This gauging procedure can be understood very nicely from the perspective of the Gaiotto curve [36]. That work studied the conformal $\mathcal{N} = 2$ theories that arise from wrapping stacks of n M5-branes on some punctured Riemann curve known as the Gaiotto curve; n is denoted the *rank* of the theory. The punctures correspond to mass deformations of the theory; an exactly conformal

theory would have all punctures turned off. In the case of two M5-branes, the resulting theories have gauge group $SU(2)^k$. We briefly recall the map between the Gaiotto curve and the weak coupling gauge theory description for the rank 2 case. Each puncture of the Gaiotto curve corresponds to an $SU(2)$ flavor symmetry. Such Riemann surfaces may be glued together at punctures by opening a hole at each puncture and glueing the two together with a tube. This results in gauging the diagonal subgroup of the $SU(2)$'s corresponding to the punctures. The sphere with three punctures corresponds to a half-hypermultiplet trifundamental under the three $SU(2)$'s associated to the three punctures. Then from the pair-of-pants decomposition of a Riemann surface, we can break any surface into some number of three-punctured spheres connected up in some way. From this, we may determine a weak coupling description of any such theory. Since the pair-of-pants decomposition is non-unique, there may be many different weak coupling descriptions; these are precisely the $\mathcal{N} = 2$ dualities studied in [36]. For our purposes, we simply want to note that this glueing procedure can be translated to the quiver gauging rule at the level of the quiver, if we can identify the appropriate S_2 symmetries in some mutation forms of each quiver. Then the quiver of the glued surface is precisely the quiver obtained by simultaneously gauging the S_2 symmetries in the two quivers. That is, we add an $SU(2)$ subquiver, remove one of each pair of identical nodes in the two quivers, and couple both of the remaining nodes to the same $SU(2)$ subquiver.

As another example, consider glueing the $SU(2)$, $N_f = 4$ quiver to itself other by gauging the diagonal subgroup $SU(2)_d \subset SU(2) \times SU(2) \subset SU(4)$. The original quiver presents $S_2 \times S_2 \subset S_4$ symmetries given by exchanging γ_3, γ_4 and γ_5, γ_6 respectively. The gauging procedure looks as follows

$$\begin{array}{ccc}
 \begin{array}{c} \gamma_5 \quad \gamma_6 \\ \swarrow \quad \searrow \\ \gamma_1 \quad \gamma_2 \\ \swarrow \quad \searrow \\ \gamma_3 \quad \gamma_4 \end{array} & \xrightarrow{\text{gauge}} & \begin{array}{c} \gamma_5 \\ \swarrow \quad \searrow \\ \gamma_1 \quad \gamma_2 \quad b \quad c \\ \swarrow \quad \searrow \\ \gamma_3 \end{array}
 \end{array} \tag{2.4.37}$$

For these rank 2 theories, there is actually a more systematic way to generate quivers for all surfaces via triangulations from special lagrangian flows, as developed in [12, 28]. The quiver gauging rule just described can in fact be understood from this triangulation view point, as explained in [12]. For example, the theory $SU(2)$, $N_f = 4$ corresponds to a sphere with 4 punctures; the gauged quiver shown above is known from that analysis to correspond to a torus with 2 punctures, which is precisely the surface produced after glueing two punctures from the 4-punctured sphere. Notice that, since the resulting surface contains 2 punctures, we would expect there to be two more $SU(2)$'s available for gauging. In fact, a mutation sequence can produce one S_2 in the quiver, but there is no way to produce two such symmetries. The analysis from the triangulation perspective shows that we can produce all but one S_2 in the quiver; that is, we can realize one fewer S_2 than the total number of punctures. Actually, there is a very good reason that we are unable to gauge the last $SU(2)$. If we did so, we would remove all punctures from the surface, and produce a quiver for a punctureless surface. However, a punctureless surface supports an exactly conformal theory - all mass deformations have been turned off. Hence the BPS spectrum would exhibit some duality, and in general be dense in the central charge plane, obstructing the existence of a quiver. Thus for consistency, it is necessary that we not be able to gauge the $SU(2)$ symmetry of a once-punctured surface. Nonetheless, we can be able to build up a quiver for any surface with at least one puncture, and these all agree with the quivers obtained from triangulations. For higher rank theories, the analog of the triangulation approach is not known; however, the gauging rules will allow us to construct quivers for a large class of theories whose quiver descriptions were previously unknown.

2.5 $SU(N)$ Gauge Theories and Beyond

In this section we apply the formalism discussed in the previous section to the examples of non-abelian ADE Yang-Mills theories with matter.

2.5.1 Construction of $SU(N)$ Quivers

Quivers for pure $SU(N)$ gauge theory were constructed in [4] via the $2d/4d$ correspondence studied there. These BPS quivers have also been studied previously in [32]. That work identified as nodes of the quiver a set of fractional branes in an orbifold phase of the geometries used in the type IIA geometric engineering [66, 67].²¹

Here we will provide a purely $4d$ motivation for that result, and use it to extend the proposal to $SU(N)$ gauge theory with arbitrary matter. First we fix some notation. We have been using (e, m) for electric and magnetic charges. Electric charges will naturally be associated to weights of the gauge group, and magnetic charges associated to roots. We denote simple roots α_i and fundamental weights ω_i ; the appropriate inner product is given by $\alpha_i \cdot \omega_j = \delta_{ij}$.

By the $2r + f$ counting, the quiver should consist of $2(N - 1)$ nodes. Let us consider the mutation form of the quiver that covers the decoupling limits in which each W boson associated to a simple root α_i separately becomes infinitely massive. In order to separately decouple these vectors, the $N - 1$ simple root W bosons must be disjointly supported as reps of the quiver. Since the reps supported on only one node cannot give vectors, and we only have $2(N - 1)$ nodes, each W boson must be supported on two distinct nodes. So we have two nodes b_i, c_i , forming an $SU(2)$ subquiver associated to each simple root. Then we simply need to choose charge assignments within

²¹Fractional branes as a basis of BPS quivers were studied in [7, 8, 30]. Their charges for $SU(N)$ were identified from a boundary CFT analysis in [68]. BPS particles with magnetic and electric charge in the IIA geometric engineering context correspond to even branes wrapped on cycles of the geometry. The fractional branes are identified with the monopoles and dyons which can become massless somewhere in moduli space, equivalently these states correspond to the vanishing cycles in the homology lattice of the Seiberg-Witten curves of these theories found in refs. [69–71]. See also [72] and references therein.

the $SU(2)$ subquivers. In order to obtain the associated W boson, the two nodes should have the charges of a consecutive pair of dyons, $((n_i + 1)\alpha_i, -\alpha_i), (-n_i\alpha_i, \alpha_i)$. The most obvious choice is just $n_i = 0$, the appropriate monopole and dyon for each simple root. If we make this choice, the result is precisely the quiver computed by [4] using the 2d/4d correspondence:

$$\begin{array}{ccc}
 \begin{array}{c} c_1 \\ \uparrow \\ \uparrow \\ b_1 \end{array} & \begin{array}{c} c_2 \\ \uparrow \\ \uparrow \\ b_2 \end{array} & \dots & \begin{array}{c} c_{N-1} \\ \uparrow \\ \uparrow \\ b_{N-1} \end{array} \\
 \swarrow & \searrow & & \swarrow \\
 & & & \searrow
 \end{array}
 \quad (2.5.1)$$

where $b_i = (0, \alpha_i)$ and $c_i = (\alpha_i, -\alpha_i)$.

The $SU(N)$ quivers we have deduced contain closed oriented cycles; thus the quiver requires a superpotential to be specified. The orbifold construction of [32] produces this superpotential by reducing the superpotential of the $\mathcal{N} = 4$ theory.²² Explicitly, the appropriate superpotential is given as,

$$\begin{array}{ccc}
 \begin{array}{c} c_1 \\ \uparrow \\ \uparrow \\ X_1 \quad Y_1 \\ \downarrow \\ b_1 \end{array} & \begin{array}{c} \xrightarrow{\phi_1} \\ b_2 \\ \leftarrow \phi'_2 \end{array} & \dots & \begin{array}{c} \leftarrow \phi'_{N-2} \\ c_{N-1} \\ \uparrow \\ X_{N-1} \quad Y_{N-1} \\ \downarrow \\ \xrightarrow{\phi_{N-2}} \\ b_{N-1} \end{array} \\
 & \begin{array}{c} \downarrow \\ X_2 \quad Y_2 \\ \uparrow \\ b_1 \leftarrow \phi'_1 \quad c_2 \rightarrow \phi_2 \end{array} & &
 \end{array}
 \quad (2.5.2)$$

with

$$\mathcal{W} = \sum_{i=1}^{N-2} X_i \phi'_i X_{i+1} \phi_i - Y_i \phi'_i Y_{i+1} \phi_i. \quad (2.5.3)$$

Before going on, we will demonstrate a weak-coupling check on this superpotential. The quivers given above explicitly display W bosons associated to the simple roots; the ordering $\arg \mathcal{Z}(b_i) > \arg \mathcal{Z}(c_i)$ ensures that there will be a W boson associated to the i th simple root. However, at weak coupling we would expect massive vector W bosons associated to all roots of the

²²The quiver (and superpotential) discussed on [32] is actually related by some mutations to the quiver we study here.

$SU(N)$ algebra, due to Higgsing of the gauge bosons. The set of massive vectors should fill out exactly one adjoint of the $SU(N)$, except for the Cartan elements, which remain massless.

Let us see how these additional vectors come about by first considering $SU(3)$. We seek a vector state corresponding to a representation with dimension vector $(1, 1, 1, 1)$. The superpotential is then

$$\mathcal{W} = X_1 \phi' X_2 \phi - Y_1 \phi' Y_2 \phi, \quad (2.5.4)$$

and the resulting F-terms are

$$\phi \phi' X_2 = \phi \phi' Y_2 = \phi \phi' X_1 = \phi \phi' Y_1 = 0, \quad (2.5.5)$$

$$\phi(X_1 X_2 - Y_1 Y_2) = \phi'(X_1 X_2 - Y_1 Y_2) = 0. \quad (2.5.6)$$

If both ϕ, ϕ' are zero, the rep is given by X_i, Y_i , and falls apart into the direct sum of two subreps, $b_1 + c_1, b_2 + c_2$. Such a situation is described as a decomposable representation; decomposable reps are never stable, since one of the two subreps must be to the left of decomposable rep in the \mathcal{Z} -plane. If ϕ, ϕ' are both nonzero, then X_i, Y_i are all zero by (2.5.5), and again the rep is decomposable. We are left with two cases, $\phi = 0, \phi' \neq 0$ and vice versa. Having set one of the ϕ 's to zero, there is one more equation in (2.5.6) that must be satisfied: $X_1 X_2 = Y_1 Y_2$. Naive dimension counting gives us $6 - 2 - 3 = 1$, so we have a vector. Gauge fixing sets ϕ (or ϕ') = $X_1 = Y_1 = 1$; then the actual moduli space is parameterized by $X_2 = 1/Y_2$, which forms \mathbb{P}_1 . Lefschetz $SU(2)$ gives exactly one vector of this charge, and no hypers. It remains to check the stability conditions. For $\phi = 0$, there are subreps $c_1, b_1 + c_1, b_1 + c_1 + c_2$; these are *not* destabilizing precisely when, in addition to the weak coupling conditions, we also have $\arg \mathcal{Z}(b_1 + c_1) < \arg \mathcal{Z}(b_2 + c_2)$. On the other hand, when $\arg \mathcal{Z}(b_1 + c_1) > \arg \mathcal{Z}(b_2 + c_2)$, then $c_1 + b_1$ is certainly a destabilizing subrep. Similarly, $\phi_2 = 0$ is stable precisely for $\arg \mathcal{Z}(b_1 + c_1) > \arg \mathcal{Z}(b_2 + c_2)$. Therefore, at any region in weak coupling, we find precisely one W boson of the desired charge.

Now we consider arbitrary $SU(N)$. By embedding the $SU(3)$ quiver as a subquiver of an

arbitrary $SU(N)$ quiver, we see that the specified superpotential (2.5.3) guarantees that exactly one W boson vector with charge $(\alpha_i + \alpha_{i+1}, 0)$ appears at weak coupling. It remains to check the W bosons associated to the rest of the roots, which have charges $(\sum_{i=j}^{j+k} \alpha_i, 0)$ for any $k > 1$. As representations, these are given by $\sum_{i=j}^{j+k} b_i + c_i$. It is clear that, for this analysis, we can simply focus on the subquiver formed by b_i, c_i for $j \leq i \leq j+k$; all other nodes (and maps involving them) are set to zero in this rep, and consequently, any superpotential terms from them are trivial. Thus we can simply study the rep $v = \sum_{i=1}^k b_i + c_i$ of the $SU(k+1)$ quiver and superpotential as shown above.

The F-terms are now a bit more subtle.

$$\phi_{i-1}\phi'_{i-1}X_{i-1} + \phi_i\phi'_iX_{i+1} = \phi_{i-1}\phi'_{i-1}Y_{i-1} + \phi_i\phi'_iY_{i+1} = 0, \quad (2.5.7)$$

$$\phi_i(X_iX_{i+1} - Y_iY_{i+1}) = \phi'_i(X_iX_{i+1} - Y_iY_{i+1}). \quad (2.5.8)$$

Again, not both ϕ_i, ϕ'_i can be zero, or else the rep is decomposable. However, it seems that perhaps both ϕ_i, ϕ'_i may be nonzero; since (2.5.7) now has two terms, this no longer forces the rep to become decomposable. Nonetheless, we can dispose of this possibility by stability. If both ϕ_i, ϕ'_i are nonzero, then either both ϕ_{i-1}, ϕ'_{i-1} are nonzero or X_{i+1}, Y_{i+1} are zero due to (2.5.7). By induction, we will find that X_j, Y_j are zero for some j . This situation cannot be Π -stable; because X_j, Y_j vanish, we have two subreps, b_j (which is now effectively a sink in the quiver), and $v - c_j$, the subrep where we set to zero c_j , (which is now an effective source in the quiver). It must be the case that one of these is destabilizing. If $\arg \mathcal{Z}(c_j) > \arg \mathcal{Z}(v)$, then we have $\arg \mathcal{Z}(b_j) > \arg \mathcal{Z}(c_j) > \arg \mathcal{Z}(v)$ so that b_j is destabilizing; otherwise $\arg \mathcal{Z}(v - c_j) > \arg \mathcal{Z}(v) > \arg \mathcal{Z}(c_j)$, so that $v - c_j$ is destabilizing.

Having dealt with this subtlety, we can continue with the analysis. The remaining case is that exactly one of ϕ_i, ϕ'_i is nonzero for each i ; this gives 2^k possibilities. First, we check the dimension of the parameter space: we start with $4k - 2$ maps and $2k - 1$ gauge symmetries; we have set $k - 1$ maps to zero, and we have $k - 1$ remaining constraints (2.5.8); thus $(4k - 2) - (2k -$

$1) - (k - 1) - (k - 1) = 1$. We may gauge fix ϕ_i (or ϕ'_i) = $X_i = Y_i = 1$ for $1 \leq i < N - 1$; then the moduli space is \mathbb{P}^1 parametrized by $X_{N-1} = 1/Y_{N-1}$. Thus we have 2^k vector states. Using stability, we will find that precisely one of these vectors is stable for any region of weak coupling. To see this, fix j and choose $\phi_j \neq 0$. Because of this choice, there is a subrep $\sum_{i=j+1}^k b_i + c_i$, which is destabilizing when $\arg \mathcal{Z} \left(\sum_{i=j+1}^k b_i + c_i \right) > \arg \mathcal{Z}(v) > \arg \mathcal{Z} \left(\sum_{i=1}^j b_i + c_i \right)$. If we had chosen $\phi'_j \neq 0$, we would have a subrep $\sum_{i=1}^j (b_i + c_i)$ which is destabilizing in exactly the opposite situation, $\arg \mathcal{Z} \left(\sum_{i=1}^j b_i + c_i \right) > \arg \mathcal{Z}(v) > \arg \mathcal{Z} \left(\sum_{i=j+1}^k b_i + c_i \right)$.²³ So we have arrived at the desired conclusion, namely, that we obtain precisely one vector for each root of $SU(N)$. With a bit more work it is possible to see that, up to field redefinitions, this is the unique superpotential at quartic order that properly produces exactly one set of W bosons. In principle this leaves the possibility of higher order terms in the superpotential, but the derivation of [32] shows that indeed no such terms arise.

2.5.2 General ADE-type Gauge Group

Some brief comments will allow us to extend the above analysis to arbitrary ADE-type (ie simply-laced) gauge group G . At weak coupling, we would again expect to be able to decouple the $\text{rank } G$ distinct $SU(2)$ subgroups, again with one corresponding to each simple root of the algebra. Then we would again find an $SU(2)$ subquiver for each simple root α_i . If we again make the ansatz of fixing charges $(0, \alpha_i), (\alpha_i, -\alpha_i)$, then we find that, for each line in the Dynkin diagram (ie

²³There are some additional subreps that should be considered, but ultimately play no role. For example, if $\phi_j \neq 0, \phi'_m \neq 0$ for $j < m$, then there is a subrep $\sum_{i=j+1}^m b_i + c_i$, which is destabilizing when $\arg \mathcal{Z} \left(\sum_{i=j+1}^m b_i + c_i \right) > \arg \mathcal{Z}(v)$. Suppose that neither subreps described above are destabilizing; then $\arg \mathcal{Z} \left(\sum_{i=j+1}^k b_i + c_i \right) < \arg \mathcal{Z}(v)$ and $\arg \mathcal{Z} \left(\sum_{i=1}^m b_i + c_i \right) < \arg \mathcal{Z}(v)$. Summing these inequalities, we find $\arg \mathcal{Z} \left(\sum_{i=j+1}^m b_i + c_i \right) < \arg \mathcal{Z}(v)$, so that this new subrep cannot be destabilizing. Further, if $c_i + b_i$ is a subrep, then so is c_i , but this again gives no additional destabilizing constraints since $\arg \mathcal{Z}(b_i) > \arg \mathcal{Z}(b_i + c_i) > \arg \mathcal{Z}(c_i)$.

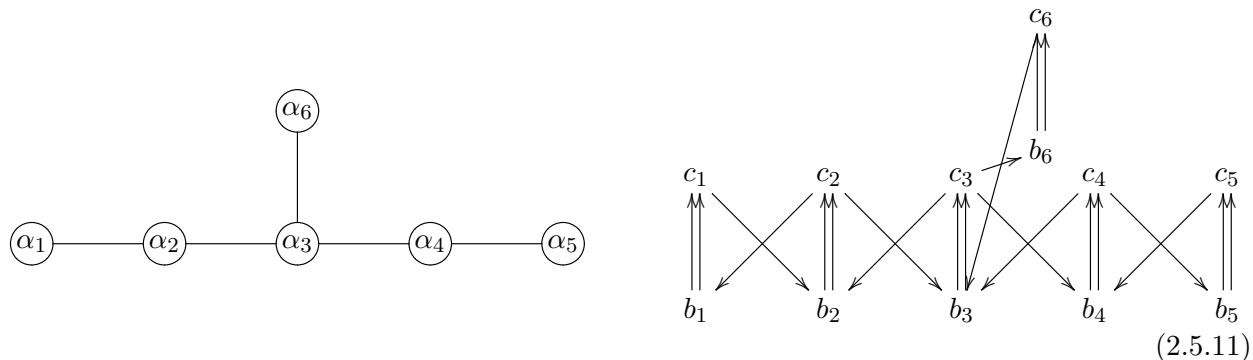
$\alpha_i \cdot \alpha_j = -1$), we must connect the respective $SU(2)$ subquivers as

$$\begin{array}{ccc}
 c_i & \xrightarrow{\phi_{ij}} & b_j \\
 \Uparrow X_i & & \Uparrow X_j \\
 b_i & & c_j \\
 \Downarrow Y_i & & \Downarrow Y_j \\
 & \xleftarrow{\phi'_{ij}} &
 \end{array}
 \tag{2.5.9}$$

with the quartic superpotential $\mathcal{W} = X_i \phi'_{ij} X_j \phi_{ij} - Y_i \phi'_{ij} Y_j \phi_{ij}$.

Thus there is a straightforward graphical prescription for constructing a quiver for pure SYM with simply-laced gauge group G , starting from the Dynkin diagram of G . For every node i of the Dynkin diagram, we draw an $SU(2)$ subquiver with nodes b_i, c_i ; for every line in the Dynkin diagram given $i - j$ we connect the $SU(2)$ subquivers as above, with the quartic superpotential. This is exactly the quiver $\widehat{A}_1 \boxtimes G$, which was found to describe these theories via $2d/4d$ in [12]. The superpotential guarantees the existence of some subset of the W bosons, namely those contained in any $SU(N)$ subquiver of the full G quiver; studying the full root system of W bosons becomes quite complicated, and we omit the analysis here. While the quartic terms must be present in the superpotential, there may or may not be some additional higher order terms. For clarity, we draw the Dynkin diagrams along with resulting quivers for D_4, E_6 .

$\tag{2.5.10}$



2.5.3 BPS Spectra of Pure $SU(N)$ SYM

In the following we will compute the BPS spectra of $SU(N)$ theories using the mutation method. We find a spectrum consisting of $N(N - 1)$ BPS particles and their antiparticles at strong coupling in agreement with the identification of the spectrum in this region with CFT states of [68].

For $N \geq 3$ these theories are not complete in the sense of ref. [12] since their charge lattice has rank $2(N - 1)$ while there are only N physical moduli that can be varied corresponding to half of the charges and the coupling of the theory. We will therefore not have the freedom to adjust all the central charges as we wish since some of them will be fixed by special geometry. To apply the mutation method we therefore need to compute the central charges in a chamber in moduli space and find a basis which has central charges lying in a half plane.

$SU(3)$

We begin with an analysis of the $SU(3)$ theory starting from the quiver discussed in section 2.5.1, which was obtained from a weak coupling analysis and which is verified by the $2d/4d$ correspondence [4]. We identify the nodes of the quiver with cycles in the SW geometry and compute their central charges to determine the ordering of the mutations. Furthermore, we track these cycles to the strong coupling region where we produce the full BPS spectrum consisting of 6

particles.

The central charge function is part of the IR data of the theory, and is thus specified by the SW solution. The $SU(N)$ SW curve can be written as [69–71]

$$y^2 = (P_{A_{N-1}}(x, u_i))^2 - \Lambda^{2N}, \quad P_{A_{N-1}}(x, u_i) = x^N - \sum_{i=2}^N u_i x^{N-i}, \quad (2.5.12)$$

where the u_i are the Casimirs parametrizing the Coulomb branch and Λ is the strong coupling scale. The SW differential is then given by [69–71]

$$\lambda(u_i) = \frac{1}{2\pi i} \frac{\partial P_{A_{N-1}}(x, u_i)}{\partial x} \frac{x dx}{y}, \quad (2.5.13)$$

and a BPS particle which is represented by a cycle γ on the SW curve has charge

$$Z_{u_i}(\gamma) = \int_{\gamma} \lambda(u_i). \quad (2.5.14)$$

Finally, the electric-magnetic inner product of two particles is computed by the intersection product of the associated cycles. We will use γ to refer to both the particle and associated cycle, and \circ to indicate both the electric-magnetic inner product and the intersection product.

We will calculate the central charge configuration for a weakly coupled point of the $SU(3)$ theory. For $SU(3)$ we set $u_2 = u$ and $u_3 = v$. The Casimirs u_i determine the vevs of the Cartan elements of $SU(N)$ semi-classically, and it can be checked that $u \rightarrow -\infty$ and $v = 0$ indeed corresponds to a weakly coupled point in $SU(3)$.

The $SU(N)$ theory has an $Sp(2N-2, \mathbb{Z})$ duality which is manifest in the different possible choices of symplectic homology basis that could be identified with electric and magnetic charges. We postpone the charge labeling and identify the nodes of the quiver directly with a choice of cycles in the geometry as shown in Fig. 2.6a.

The quiver obtained in this way at weak coupling should have a number of properties:

- The intersections of cycles must agree with the electric-magnetic inner product as defined by the quiver

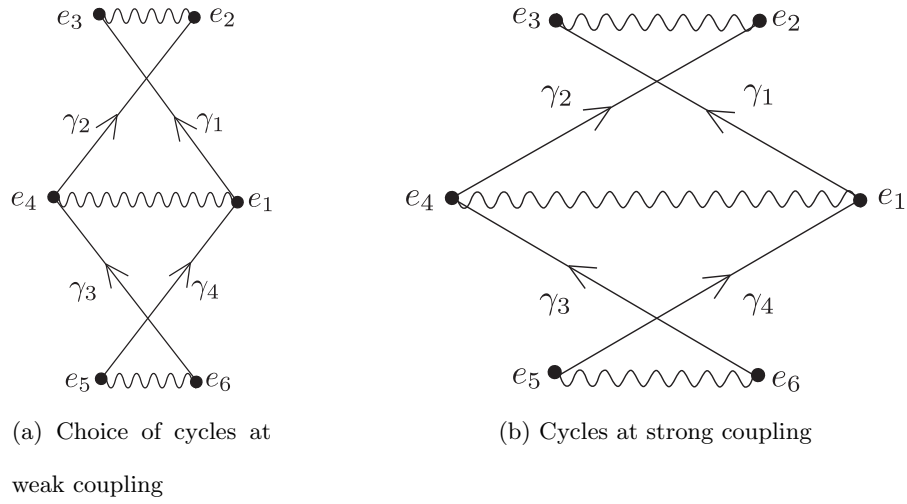


Figure 2.6: The choice of cycles in the x -plane at weak and strong coupling is shown in Figs. 2.6a, 2.6b respectively. $e_i, i = 1, \dots, 6$ denote the roots of $(x^3 - ux - v)^2 - \Lambda^6$ and become the sixth roots of unity as we tune the moduli to strong coupling and set $\Lambda = 1$.

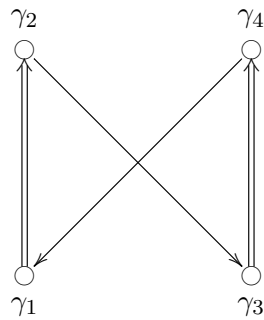


Figure 2.7: Quiver obtained from the intersections of the cycles in Figs. 2.6a, 2.6b.

- The central charges of all the nodes must lie in a common half-plane
- The apparent $SU(2)$ subquivers should be weakly coupled
- The central charges of the W bosons of the $SU(2)$ gauge groups should be vanishingly small compared to the central charges of the nodes in the $u \rightarrow -\infty$ limit

The last condition follows from the fact that the electrically charged objects should be parametrically light compared to the dyonic states of the theory at weak coupling, since here the electric particles are the fundamental degrees of freedom.

The choice of cycles in Fig. 2.6a meets these conditions. That the first is met is obvious, and the latter three can be explicitly checked by numerically computing the associated integrals of the SW differential along the given curves. This has been done, and the values of the central charges for large but finite $u < 0$ are as depicted in Fig. 2.8a. Since the $SU(2)$ subquivers are weakly coupled, we are in an infinite chamber, as expected at weak coupling. To apply the mutation method most efficiently we will tune the moduli to arrive in a chamber with a finite spectrum.

We can track the behavior of the quiver explicitly as we tune moduli. At walls of marginal stability nothing happens at the level of the quiver, while at walls of the second kind we must mutate to find a valid description on the other side. A generic path in the $SU(3)$ moduli space may pass through arbitrarily many - even infinitely many - walls of the second kind, thereby alluding an analysis. For $SU(3)$ there exists a path which takes us from weak coupling to the strongly coupled $u = 0$ point and passes through no walls of the second kind, thereby allowing a quite seamless transition between the understood weak coupling chamber and the strongly coupled chamber containing the $u = 0$ point.

We follow the straight line path with $v = \text{Im } u = 0$ from $u = -\infty$ to $u = 0$. The pairs of aligned central charges stay aligned along the entire path, and cross in tandem at a finite value of $u < 0$. All the while, all central charges remain in the upper half-plane. At $u = 0$, both $SU(2)$'s

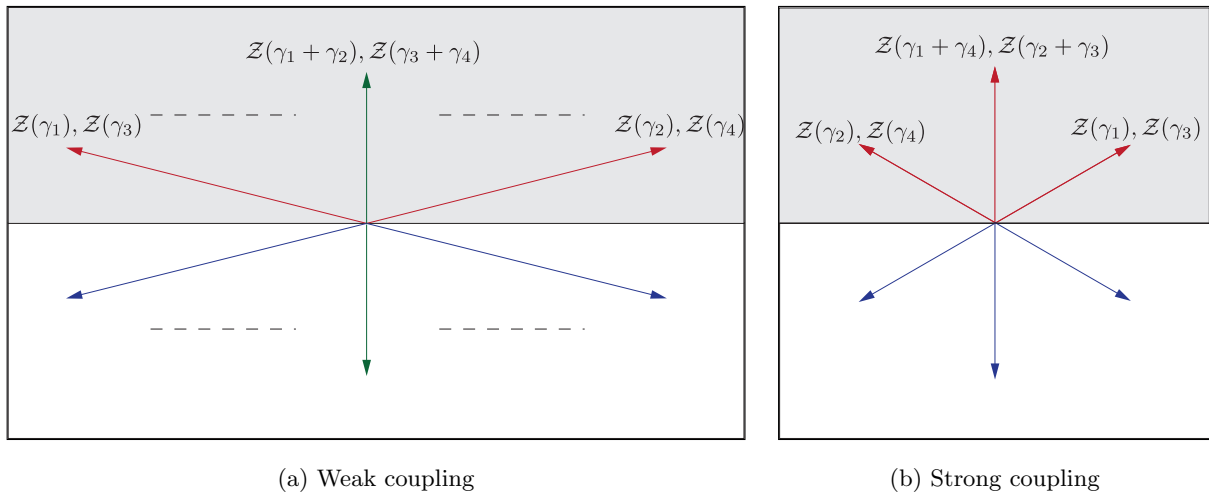


Figure 2.8: The central charges of BPS states of $SU(3)$ are depicted at weak (a) and strong (b) coupling respectively. At weak coupling, the left- and right-most nodes, along with the weak coupling W bosons are shown explicitly. The full spectrum at weak coupling is not known, but at least includes two infinite towers of dyons, which are not shown. In the limit of zero coupling, the left- and right-most nodes approach π separation and infinite length. As we tune towards strong coupling, the states γ_1, γ_3 and γ_2, γ_4 approach and cross each other. At strong coupling, the full finite spectrum of BPS states is depicted; the \mathbb{Z}_6 symmetry is manifest.

are strongly coupled, and the central charge configuration is as given in Fig. 2.8b. Now we simply apply the mutation algorithm with the central charges associated to this point in moduli space. What we find is a $N(N - 1) = 6$ state chamber with states

$$\gamma_2, \gamma_4, \gamma_2 + \gamma_3, \gamma_1 + \gamma_4, \gamma_1, \gamma_3. \quad (2.5.15)$$

Let us note some features of the strong coupling spectrum we have found. First of all, all states in the chamber correspond to vanishing cycles in the Seiberg-Witten geometry. That is, they all correspond to cycles which vanish somewhere on moduli space. This agrees with earlier intuition about the relation between the strong coupling $SU(N)$ spectrum and vanishing cycles of the SW geometry [32, 69, 71, 72].

The second feature, which will become quite important in our $SU(N)$ analysis below, is that the chamber we have found respects the $\mathbb{Z}_{2N} = \mathbb{Z}_6$ symmetry of the IR solution.

In principle one would hope that the same story carried over for the $SU(N)$ case. We would ideally start from weak coupling and tune moduli until we arrived at the strongly coupled $u_i = 0$ point, and then see that this point lied in a finite chamber with $N(N - 1)$ states. Unfortunately the situation becomes technically complicated, in a way we will briefly explain. Above, we chose a very particular path between the $u_i = 0$ point and weak coupling, along which the quiver passed through no walls of the second kind, where quiver mutation is necessary. This was a path which deformed the order 1 term in the defining polynomial of the Seiberg-Witten curve.

In the $SU(N)$ case it is always the x^{N-2} deformation which has this nice property. That is, if we deform the coefficient of the x^{N-2} term alone from the $u_i = 0$ point along certain directions in \mathbb{C} , the quiver will be extremely well behaved, just as above. The issue is that it is only in the $N = 3$ case that this deformation alone is sufficient to arrive at weak coupling. In all other cases there will be some unbroken subgroup which remains. Thus to get to weak coupling, we must deform lower order terms, but these are not nice in terms of the quiver description. In particular,

no simple choice seems to get from strong to weak coupling while only passing through a small number of walls of the second kind. Potentially such a path remains to be found, and the same method can then be generalized to the $SU(N)$ case. At present, we will proceed with a discussion of the $SU(N)$ case at $u = 0$ based on what we've learned in $SU(3)$.

$SU(N)$ at Strong Coupling

We now consider the general case of $SU(N)$ at strong coupling. Our objective is to determine the quiver, charge labels of nodes, and ordering of central charges at some point of strong coupling, and then compute the resulting spectrum via the mutation method. Of course, to honestly produce the quiver we would need to somehow find a basis of BPS states. However, the quiver has already been derived from other considerations, and motivated from a purely $4d$ perspective in 2.5.1. Here we will infer quiver along with charge labels at strong coupling by generalizing the results above for $SU(3)$.

Fix the moduli $u_i = 0$, so that the Seiberg-Witten curve is given as

$$y^2 = x^{2N} - \Lambda^{2N}, \quad (2.5.16)$$

with Seiberg-Witten differential

$$\lambda = \frac{1}{2\pi i} \frac{Nx^N dx}{y}. \quad (2.5.17)$$

We take a symplectic homology basis, a_i, b_i for $i = 1, \dots, N - 1$, with $a_i \circ a_j = b_i \circ b_j = 0$ and $a_i \circ b_j = \delta_{ij}$. The appropriate choice of cycles is shown in Figure 2.9. We have chosen the a_i 's to be the cycles that collapse as $u_N \rightarrow \infty$, since these are pure electric charges. There is still some ambiguity in choosing b cycles, which are pure magnetic monopoles with charges given by simple roots of $SU(N)$. We fix the ambiguity by choosing the b cycles to be ones that vanish somewhere in moduli space. This is a natural choice, since each of the simple roots has a full $SU(2)$ moduli space associated with it contained in the $SU(N)$ moduli space; by the original Seiberg-Witten

$SU(2)$ analysis, the monopole associated to each simple root becomes massless at some locus of the $SU(N)$ moduli space.

At the origin of moduli space, the curve has a \mathbb{Z}_{2N} discrete symmetry. If we denote ξ the generator of the symmetry, we have

$$\xi(x) = e^{-i\pi/N} x. \quad (2.5.18)$$

The action on the x -plane is simply a $-\pi/N$ rotation; on the central charge function \mathcal{Z} , this gives

$$\xi(\lambda) = -e^{-i\pi/N} \lambda \quad (2.5.19)$$

$$\xi(\mathcal{Z}(\gamma)) = -e^{-i\pi/N} \mathcal{Z}(\gamma). \quad (2.5.20)$$

This induces an exact symmetry of the quantum theory that will be quite useful. It indicates that BPS states will come in \mathbb{Z}_{2N} orbits; the magnitude of their central charges of cycles in an orbit are all identical, and their phases are distributed \mathbb{Z}_{2N} symmetrically in the complex plane. Again, by $SU(2)$ reasoning, each magnetic monopole with simple root charge will be a BPS state at the origin of moduli space. From Figure 2.9, it is clear that all the b_i 's are in distinct orbits. Thus we have obtained $(N - 1)$ distinct orbits, one for each simple root monopole with electric-magnetic charge $(0, \alpha_i)$; each orbit consists of $2N$ BPS states, N of which are particles, and N antiparticles.

To compute the periods, we integrate the Seiberg-Witten differential, to obtain

$$\int \lambda = \frac{1}{2\pi} \frac{N}{N+1} x^{N+1} {}_2F_1\left(\frac{1}{2}, \frac{N+1}{2N}, \frac{1}{2N} + \frac{3}{2}, 1\right) = \kappa(N)x^{N+1}, \quad (2.5.21)$$

where κ is some proportionality constant that depends on N but is independent of x . Evaluating the definite integral for the b_i 's shown in Figure 2.9, we find

$$\mathcal{Z}(b_j) = 2\kappa(N) i e^{i\pi/N} \sin \frac{j\pi}{N} \quad (2.5.22)$$

From the action of the ξ , we see that the full \mathbb{Z}_{2N} orbits of vanishing cycles will fill out all $2N$ -roots of unity (up to some overall phase $\arg(i e^{i\pi/N} \kappa(N))$) in the \mathcal{Z} -plane. This configuration of central charges is depicted in Figure 2.10

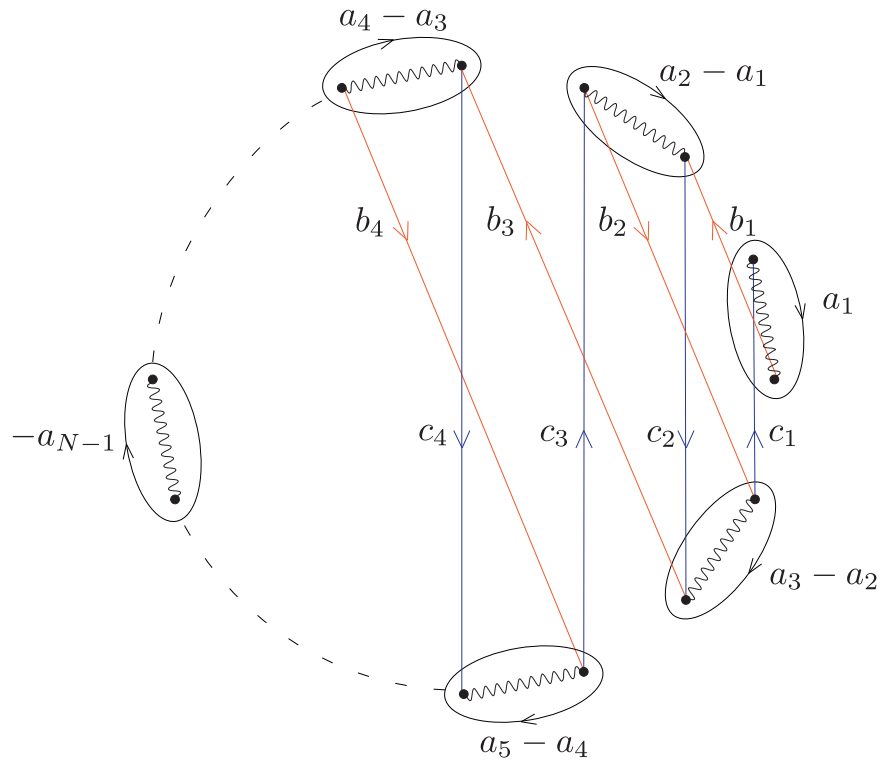


Figure 2.9: The Seiberg-Witten curve described by (2.5.16), shown as a double cover of the x -plane, with branch cuts as indicated. The labelled a_i, b_i cycles give a symplectic homology basis. The action of the \mathbb{Z}_{2N} symmetry rotates the plane by $e^{-i\pi/N}$, and thus rotates b_i into c_i . The b_i, c_i cycles constitute the positive integral basis of states that appear as nodes of the quiver. Note that we have taken a different convention for branch cuts than the one used in Fig. 2.6a. This choice is more convenient for the strong coupling analysis, and agrees with the conventions used in [71].

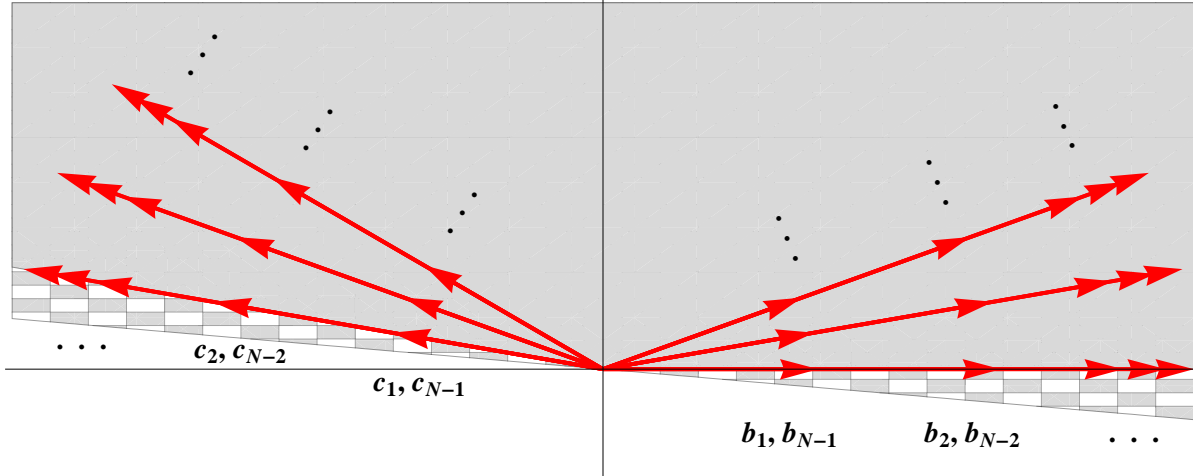


Figure 2.10: Central charges of vanishing cycles plotted in the \mathcal{Z} -plane (where we have rotated by some overall phase $\arg(i\kappa(N))$). The half-plane we use to construct the quiver is shown as the gray region. The b_j cycles have $\mathcal{Z}(b_j) \sim \sin \frac{j\pi}{N}$; note that $\mathcal{Z}(b_j) = \mathcal{Z}(b_{N-j})$. The b_j are therefore $N - 1$ distinct collinear states shown on the positive real axis. Each ray of collinear red arrows is a \mathbb{Z}_{2N} rotation of the b_j 's. There are N such rays in the half-plane, situated at $2N$ -roots of unity. In total we have $N(N - 1)$ states depicted in the diagram. The antiparticles in the opposite half-plane are not shown. The half-plane is chosen so that b_j are right-most BPS states, which forces c_j to be left-most BPS states. As explained in the analysis, for such a half-plane to exist, the region checkered in white and gray must be free of BPS states.

To continue, we now generalize from the $SU(2)$ and $SU(3)$ results. In those cases, the BPS spectra were precisely equivalent to the set of vanishing cycles of the Seiberg-Witten geometry. It is natural to imagine that for general N it is at least possible to choose a positive integral basis for BPS states that consists of vanishing cycles. The vanishing cycles do in fact span the homology lattice, so this is sensible assumption. As we will see, this allows us to obtain a quiver that agrees with (2.5.1), which was also proposed from other perspectives [4, 32]. Thus, we seek a positive integral basis of vanishing cycles; to do so, we must first choose a half-plane. Since the $N - 1$ b_i 's have the same phase, we may tune the half-plane to make them right-most vanishing cycles; then the b_i 's are forced to appear as $N - 1$ nodes of the quiver.²⁴ Having fixed this choice of half-plane, it is clear from Figure 2.10 that $c_i \equiv \xi(b_i)$ form $N - 1$ right-most vanishing cycles in the half-plane, and therefore must also appear in the quiver. These states are given as

$$c_i \equiv \xi(b_i) = \begin{cases} -a_{i-1} + 2a_i - a_{i+1} + b_i = (\alpha_i, \alpha_i) & \text{if } i \text{ is even} \\ -a_{i-1} + 2a_i - a_{i+1} - b_{i-1} - b_i - b_{i+1} = (\alpha_i, -\alpha_{i-1} - \alpha_i - \alpha_{i+1}) & \text{if } i \text{ is odd} \end{cases} \quad (2.5.23)$$

We now have specified $2(N - 1)$ nodes of the quiver; in fact, this is exactly the number of nodes in the quiver, by the counting $2r + f = 2(N - 1)$. At this point we have fully determined the quiver as follows:

The diagram shows a quiver with two rows of nodes. The bottom row has nodes b_1, b_2, \dots, b_{N-1} and the top row has nodes c_1, c_2, \dots, c_{N-1} . For each i , there are two vertical arrows between b_i and c_i : one pointing up and one pointing down. Additionally, there are diagonal arrows between adjacent nodes in the top row: $c_1 \rightarrow c_2$, $c_2 \rightarrow c_1$, $c_2 \rightarrow c_3$, $c_3 \rightarrow c_2$, and so on. Ellipses between the first and last parts of the diagram indicate the continuation of the pattern.

²⁴In principle, a bound state of multiple b_i 's would also have the same phase, and one might worry that some of these $N - 1$ states were actually bound states of the others. However, this is in fact impossible. The b_i are linearly independent cycles, so none can occur as a linear combination of the others; furthermore $b_i \circ b_j = 0$, so there exist no bound states of the form $b_i + b_j$. So *all* of the b_i cycles must appear as nodes of the quiver.

It is encouraging to note that mutation equivalences will allow us to make contact with the weak coupling discussion of section 2.5.1. The quiver we have obtained (2.5.24) is already of the same form as (2.5.1), but with different charge assignments. Mutating to the right on all b_{2i} and to the left on all b_{2i-1} will produce leave the quiver form unchanged, but transform the charges to $b_i = (0, -\alpha_i)$, $c_i = (\alpha_i, \alpha_i)$. These are precisely the weak coupling charges proposed in section 2.5.1, with some alternative choice of dyon pairs, $n_i = -1$. Note, however, that in order to realize these mutations, we must go through a large number of wall crossings, since we took left-mutations of some b_i , which, in our strong coupling calculation, are not left-most, but instead right-most.

We can use the quiver to compute the full BPS spectrum at this strong coupling chamber of moduli space. We begin by mutating on the left-most states, c_i . This produces a new set of charges, $c_i \rightarrow -c_i$, $b_i \rightarrow b_i + c_{i-1} + c_{i+1}$. The new states that replace the b_i are now left-most, again all at the same phase in the central charge plane. Focusing on the central charges of the nodes, we see that the charges of the new quiver are related to those of the original quiver by a rotation of $e^{-i\pi/N}$ (see Fig. 2.10). So as we continue mutating in phase order, this process of N coincident mutations simply repeats itself. Continuing in this way, a finite spectrum is exhibited by the mutation method with a mutation sequence of length $N(N-1)$,

$$c_1, c_2, \dots, c_{N-1}, b_1, b_2, \dots, b_{N-1}, c_1, c_2, \dots, c_{N-1}, b_1, b_2, \dots, b_{N-1}, \dots \quad (2.5.25)$$

The states produced in this way are,

$$\begin{array}{cccccc}
c_1, & c_2, & c_3, & \dots, & c_{N-1}, & \\
b_1 + c_2, & c_1 + b_2 + c_3, & c_2 + b_3 + c_4, & \dots, & c_{N-2} + b_{N-1}, & \\
b_2 + c_3, & b_1 + c_2 + b_3 + c_4, & c_1 + b_2 + c_3 + b_4 + c_5, & \dots, & c_{N-3} + b_{N-2} & \\
b_3 + c_4, & b_2 + c_3 + b_4 + c_5, & b_1 + c_2 + b_3 + c_4 + b_5 + c_6, & \dots, & c_{N-4} + b_{N-2} & \\
\vdots & \vdots & \vdots & \vdots & \vdots & \\
b_{N-1}, & b_{N-2}, & b_{N-1} & , \dots, & b_1 &
\end{array} \quad (2.5.26)$$

This array of states can be filled out iteratively after the first two rows are computed. The state μ_{ij} in position (i, j) with $i \geq 2$ is given by

$$\mu_{i-1, j-1} + \mu_{i-1, j+1} - \mu_{i-2, j}, \quad (2.5.27)$$

where we set $\mu_{ij} = 0$ for $j < 1$ and $j > N - 1$. It is slightly more economical to take as the base cases $i = 0, 1$ where we add $\mu_{0, j} = -b_j$, along with $\mu_{1, j} = c_j$ as already given. The resulting states precisely fill out the full set of $N(N - 1)$ vanishing cycles,

$$\boxed{|\mathcal{B}_{SU(N)}| = N(N - 1)}. \quad (2.5.28)$$

This result agrees with the computation of strong coupling BPS states via CFT methods [68] and is a strong confirmation of the techniques studied here.

2.5.4 Adding Matter

Adding arbitrary hypermultiplet matter to pure SYM with ADE-type gauge group is quite analogous to the procedure described in 2.4.2 for $SU(2)$. Consider adding hypermultiplet matter charged under the gauge group G in a representation R . Again, we tune the mass of the matter to infinity. Here, by similar decoupling reasoning we would expect to add as a node a an electrically charged lowest weight state of the matter representation R ; ie we should have electric-magnetic charge $(-d, 0)$ where $-d$ is the lowest weight of R . From this, positive linear combinations may generate the full representation R by adding various W bosons with charge $(\alpha_i, 0)$ to the new state $(-d, 0)$.

Having determined the charge of the new node $f = (-d, 0)$, it is straightforward to compute electric-magnetic inner products to fix the quiver. Explicitly, we may decompose the lowest weight $-d = -\sum_i d_i \omega_i$ where d_i are positive integers. Then $f \circ b_j = (-d, 0) \circ (0, \alpha_j) = -d_i(\omega_i \cdot \alpha_j) = -d_i$ and $f \circ c_j = (-d, 0) \circ (\alpha_j, -\alpha_j) = d_i$. Thus the new node has d_i arrows connected to each node of the i th $SU(2)$ subquiver, forming an oriented three-cycle. Again we run into the

subtlety seen in section 2.4.8: this quiver can certainly generate the matter rep R , but may in fact generate some additional matter representations. In fact, by adding such a node, we actually add the full tensor reducible representation $\otimes_i \mathbf{r}_i^{d_i}$, (where \mathbf{r}_i are the fundamental reps of the gauge group) instead of adding only the irreducible rep, R .

We can propose one very clear consistency check on this procedure. Due to the structure of $\mathcal{N} = 2$ hypermultiplets, adding a hypermultiplet in rep R adds a multiplet of states in $R \oplus \bar{R}$. Thus, in principle, adding matter in rep R is equivalent to adding matter in rep \bar{R} . For the fundamental \mathbf{N} of $SU(N)$, the lowest weight of \mathbf{N} is $-\omega_{N-1}$, while the lowest weight of $\bar{\mathbf{N}}$ is $-\omega_1$. This creates some ambiguity in defining the quiver of $SU(N)$ $N_f > 1$.

$$(2.5.29)$$

By the above discussion, any choice of $0 \leq k \leq N_f$ seems to give a possible quiver for this theory. For consistency, the representation theory of all of these quivers must be equivalent. One can easily check that the quivers are in fact mutation equivalent. To move node f_i from the left to the right, apply the following sequence of mutations: $f_i, b_1, c_1, b_2, c_2, \dots, b_{N-1}, c_{N-1}$; a similar reversed sequence $f_j, b_{N-1}, c_{N-1}, b_{N-2}, c_{N-2}, \dots, b_1, c_1$ moves node f_j from right to left. We can move the f_i one by one across the quiver, and any two choices of k will be connected via these mutation sequences. Thus by the general reasoning of section 2.3, these quivers do in fact correspond to identical physical theories.

2.5.5 BPS States of SQCD

We now wish to extend our analysis of strong-coupling SYM to include arbitrary fundamental quark hypermultiplets coupled to the gauge group. Recall that our rule for coupling matter was valid with all masses tuned parametrically large. With a suitable definition of charges, only the N_f flavor nodes will carry flavor charge,²⁵ and decouple from the pure gauge theory when masses are scaled up. We again study the origin of the Coulomb branch, and expect the light pure gauge degrees of freedom to reproduce the finite spectrum given above. Finally, we must fix the central charge phases of the flavor nodes; we choose all of them to be to the left of the c_i ; for definiteness, let $\arg \mathcal{Z}(f_1) > \arg \mathcal{Z}(f_2) > \dots > \arg \mathcal{Z}(f_{N_f})$. Having fixed all parameters of the theory, we may use the mutation method to compute a finite spectrum. For each flavor f_k , we find, in phase order

$$f_k, f_k + b_1, f_k + b_1 + c_1, f_k + b_1 + c_1 + b_2, \dots, f_k \sum_{i=1}^{N-1} b_i + c_i, \quad (2.5.30)$$

given by mutation sequence

$$f_k, b_1, c_1, b_2 \dots c_{N-1}. \quad (2.5.31)$$

As discussed in section 2.4.4, the charges assigned to nodes are dependent on some choice of ‘duality frame.’ If we take the charge assignments found at weak coupling, $b_i = (0, \alpha_i)$, $c_i = (\alpha_i, -\alpha_i)$, we can see a nice consistency check on this result. With these charges, the flavor states found above contain N pure electric (ie, zero magnetic charge) states with charges forming a fundamental \mathbf{N} of the $SU(N)$, given by $f_k + \sum_{i=1}^k b_i + c_i$, $0 \leq k \leq N - 1$. The remaining states are then some additional $N - 1$ additional flavor dyon states.

Since the flavor nodes are to the left with parametrically large masses, any state with flavor occurs before any of the light pure gauge degrees of freedom; by our choice of central charges,

²⁵Recall that in our analysis of $SU(2)$ with flavor, the natural assignment of charges gave flavor charge to the nodes of the $SU(2)$ subquiver, along with the additional flavor node. This was simply a familiar choice of convention; by redefining electric and magnetic charges, we can arrange a configuration in which only the additional matter node carries flavor charge.

the flavor states occur in order. All states with flavor charge f_1 occur first, and then all states with charge f_2 and so on. Continuing with the mutation method, the set of $N(N-1)$ gauge dyons will be found after all the flavor states described above. The full spectrum is given by $N_f(2N-1)+N(N-1)$ BPS hypermultiplets, consisting of $2N-1$ flavor states for each fundamental, and $N(N-1)$ pure gauge strong coupling dyons,

$$\boxed{|\mathcal{B}_{\text{SQCD}}| = N_f(2N-1) + N(N-1)}. \quad (2.5.32)$$

2.5.6 Further ADE examples

Here, we briefly review some additional finite chambers of ADE-type gauge theories that may be obtained by the mutation method. For these examples, the period computation done in section 2.5.3 becomes much more complicated. We will skip that calculation, and instead simply identify a finite mutation sequence that generalizes the one found there for $SU(N)$.

For pure SYM with DE-type gauge group, the quiver was given in section 2.5.2. There exists a finite mutation sequence for any of the ADE-type quivers whose number of states is exactly the total number of roots of G ,

$$\boxed{|\mathcal{B}_{ADE}| = \dim(\text{adjoint}) - \text{rank}(G)}. \quad (2.5.33)$$

This spectrum can be interpreted as a monopole-dyon pair for every positive root. The mutation sequence is given as before

$$c_1, c_2, \dots, c_n, b_1, b_2, \dots, b_n, c_1, c_2, \dots, c_n, \dots \quad (2.5.34)$$

We can also study ADE-type groups with additional matter representations, by following the same strategy as 3.2.4. We fix the pure gauge degrees of freedom at the strong coupling, finite chamber point discussed above, and take large mass limit for the matter. By choosing the phase of the matter nodes to be left-most, we force all states with flavor charge to be further

left than the pure gauge states. For an A-type group (ie $SU(N)$), in addition to quarks, we may couple antisymmetric tensor representations, and find a finite chamber. Generalizing from the SQCD result, there is some duality frame for which the flavor states organize into $\frac{1}{2}N(N-1)$ pure electric states whose charges fill out the antisymmetric tensor of $SU(N)$, along with some number of additional dyon states. Note that by contrast, an $SU(N)$ theory with matter in the symmetric tensor rep can never have a finite chamber. The symmetric tensor is given as a the highest weight representation of the tensor $\mathbf{N} \otimes \mathbf{N}$. By the prescription of section 2.5.4, the resulting quiver would contain a subquiver of the form studied for the $SU(2)$, $\mathcal{N} = 2^*$ theory. In section 2.4.8, we showed that this any chamber of this quiver contains at least two vector particles, and thus cannot have finitely many states. Furthermore, the presences of at least two accumulation rays obstructs the mutation method. The larger quiver for $SU(N)$ with a symmetric tensor will produce at least all the states obtained from its subquiver, and thus it will suffer from the same complications.

For a D-type group, $SO(2n)$ with matter in vector representation of $SO(2n)$, we find a finite chamber of $4(n+1)$ flavor states, along with the $2n(n-1)$ gauge states. Here the flavor states contain $2n$ pure electric states whose charges fill out a $\mathbf{2n}$ -vector of $SO(2n)$, along with $2n+1$ additional flavor dyon states. With N_v vector representations, we find

$$\boxed{|\mathcal{B}_{SO(2n)}| = N_v(4n+1) + 2n(n-1).} \quad (2.5.35)$$

We also find a finite chamber for E_6 with matter in the smallest fundamental representation, $\mathbf{27}$; the flavor states contain pure electric charges filling out the fundamental representation, along with 46 additional flavor dyon states; a theory with N_f $\mathbf{27}$'s yields

$$\boxed{|\mathcal{B}_{E_6}| = 73N_f + 72.} \quad (2.5.36)$$

For E_8 , one may not expect any finite chamber, since the smallest fundamental is the adjoint, and the resulting theory is $\mathcal{N} = 2^*$, that is, a massive deformation of a conformal $\mathcal{N} = 4$ theory.

Chapter 3

Quivers and Riemann Surfaces

3.1 Introduction

In this chapter we aim to broaden our understanding of BPS quivers by studying a simple, and largely geometric set of examples. We investigate BPS quivers in a class of $\mathcal{N} = 2$ quantum field theories known as *complete theories* [12]. These theories are defined by the property that as one varies all parameters (including moduli, couplings and bare masses), the number of independent central charges is equal to the rank of the charge lattice. Completeness is a strong assumption about a field theory and is typically not satisfied. However, a rich class of examples of such theories includes all the four dimensional $\mathcal{N} = 2$ models that can be obtained by wrapping a pair of M5 branes on a punctured Riemann surface. These are the so-called rank two Gaiotto theories. [35, 36, 39, 73].

Because of their simplicity, the class of complete theories defined by pairs of M5 branes on Riemann surfaces will be the focus of our investigation in this work. Broadly speaking, our aim is to determine and understand the BPS quiver in such examples. To accomplish the task of determining the BPS quiver, in section 3.2 we reconstruct these complete theories via geometric engineering in type IIB string theory on a local Calabi-Yau threefold. [11, 66, 67]. Such an approach

has the advantage that the BPS states can be explicitly identified as D3-branes wrapping special lagrangian cycles in the Calabi-Yau. This makes the appearance of a quiver in the BPS state counting problem manifest: the quiver simply encodes the world volume quantum mechanics of the D3-branes. [6] However, we can go further and pass from this implicit description of the quantum mechanics of D3-branes to an explicit algorithm for constructing the BPS quiver. As we review there, the structure of the quiver is completely encoded by a certain triangulation of the Gaiotto curve, the Riemann surface where the pair of M5 branes lives. Further, we explain how the same triangulation allows one to compute the superpotential for the quiver, and in this way makes the task of determining the full BPS quiver data for any given example an algorithmic procedure.

Finally, in section 3.3 we undertake a brief investigation of complete theories with BPS quivers which do not come from Gaiotto type constructions. In [12] such theories were classified. They consist of eleven exceptional theories which are not of the Riemann surface type. For all these examples except one, we determine an associated superpotential and a finite chamber of BPS states.

3.2 BPS Quivers of Complete Theories

We now turn to our primary interest of determining the BPS quivers, superpotentials, and spectra for complete theories. In this section we focus on determining the BPS quiver for those complete theories that coincide with the rank two Gaiotto theories.¹ By construction, all such theories are intrinsically determined by a Riemann surface \mathcal{C} decorated by a number of marked points defined by the punctures. By the conclusion of this analysis, we will see that the BPS quiver, together with its superpotential, is encoded combinatorially in a triangulation of this decorated surface.

¹In fact, among such theories, BPS quivers exist only for theories given by a Riemann surface *with* some punctures. The case with no punctures describes an exactly conformal theory and its BPS states do not admit a simple description.

We will construct these models using geometric engineering [11, 66, 67] in type IIB string theory on a non-compact Calabi-Yau threefold. The threefolds in question can be built up starting from a Riemann surface \mathcal{C} . We start with a four complex-dimensional space described by a rank three complex vector bundle over \mathcal{C} . Explicitly

$$K_{\mathcal{C}} \oplus K_{\mathcal{C}} \oplus K_{\mathcal{C}} \rightarrow \mathcal{C}, \quad (3.2.1)$$

where in the above $K_{\mathcal{C}}$ denotes the canonical line bundle of holomorphic one-forms on the Riemann surface \mathcal{C} . In general the surface \mathcal{C} is punctured at a finite number of points $p_i \in \mathcal{C}$ and thus is non-compact.

Next we select a particular holomorphic quadratic differential ϕ on \mathcal{C} . As a quadratic differential, ϕ transforms under holomorphic changes of coordinates on \mathcal{C} as follows

$$\phi'(x') = \phi(x) \left(\frac{dx}{dx'} \right)^2. \quad (3.2.2)$$

To completely specify the problem, we must also fix the limiting behavior of ϕ at the ideal boundaries of \mathcal{C} , namely the punctures p_i . Near each such puncture the quadratic differential is permitted to have a pole of finite order. We fix the non-normalizable behavior of ϕ as a boundary condition and therefore impose that near p_i

$$\phi(x) \sim \frac{1}{x^{k_i+2}} dx^2 + \text{less singular terms}. \quad (3.2.3)$$

The integer $k_i \geq 0$ associated to each puncture is invariant under changes of coordinates. It is an important aspect of the construction, which we return to in section 3.2.1.²

Given this data our Calabi-Yau threefold is then defined by introducing local coordinates (u, v, y) on the fiber of the vector bundle (3.2.1) and solving the following equation

$$uv = y^2 - \phi(x). \quad (3.2.4)$$

²The reason for the exclusion of the case $k_i = -1$ is that such fluctuations in ϕ are normalizable, and hence are not fixed as part of the boundary conditions.

The associated holomorphic three-form Ω is given by

$$\Omega = \frac{du}{u} \wedge dy \wedge dx. \quad (3.2.5)$$

It is then known that finite mass strings probing the singularity of this geometry engineer a 4d field theory with $\mathcal{N} = 2$ supersymmetry. The Seiberg-Witten curve Σ of such a theory is given by a double cover of \mathcal{C} , and we obtain the Seiberg-Witten differential by integrating Ω over a non-trivial 2-cycle in the fiber.

$$\Sigma = \{(x, y) | y^2 = \phi(x)\}; \quad \lambda = \int_{S^2(x)} \Omega = y dx = \sqrt{\phi}. \quad (3.2.6)$$

By varying the quadratic differential we obtain a family of Seiberg-Witten curves, and in this way the Coulomb branch \mathcal{U} of the theory is naturally identified with the space of quadratic differentials obeying the boundary conditions (3.2.3).

It is also known that many of the simplest interesting gauge theories can be geometrically engineered in this fashion. For example taking \mathcal{C} to be a sphere with two punctures p_i both with $k_i = 1$ constructs the pure $SU(2)$ theory. In general the class of field theories constructed in this way yields asymptotically free or conformal theories with gauge groups given by a product of $SU(2)$'s, together with various scaling and decoupling limits of such field theories. They are exactly the type IIB version of the rank two Gaiotto theories constructed using M-theory in [36], and, as we have mentioned above, in that context \mathcal{C} is referred to as the Gaiotto curve.

For our present purposes, the primary advantage of building an $\mathcal{N} = 2$ quantum field theory in string theory is that the set of supersymmetric objects in string theory, the BPS branes, is known. In our case we seek a brane whose physical interpretation in four-dimensions is a charged supersymmetric particle of finite mass. Thus the worldvolume of the brane should be an extended timelike worldline in Minkowski space times a volume minimizing compact cycle in the Calabi-Yau (3.2.4). Since type IIB has only odd dimensional branes, the only possibility is that BPS states are described geometrically by Dirichlet three-branes wrapping special lagrangian three-cycles.

Thus we are reduced to a classical, if difficult, geometric problem of counting special lagrangians [53, 74]. These are compact lagrangian three-manifolds N on which the holomorphic three-form has a constant phase

$$\Omega|_N = e^{i\theta}|\Omega|. \quad (3.2.7)$$

The central charge of such a brane is given by

$$\mathcal{Z}_u(N) = \int_N \Omega, \quad (3.2.8)$$

and the phase θ in the above is identified with the argument of the central charge of the 4d particle defined by N

$$\theta = \arg \mathcal{Z}(N). \quad (3.2.9)$$

Now one of the key observations of [11] is that, in the geometries described by (3.2.4), the counting of special lagrangians can in fact be phrased entirely as a problem in \mathcal{C} . To exhibit this feature we use the fact that all of our special lagrangians are embedded inside the vector bundle (3.2.1) and hence admit a natural projection to \mathcal{C} . The image of this projection is a certain one cycle η in \mathcal{C} whose topology depends on the topology of N . Each special lagrangian also wraps a non-trivial S^2 in the fiber, which shrinks to zero at the zeros of ϕ . The possibilities in our examples are as follows, and are illustrated in Figure 3.1:

- $N \cong S^3$. Such special-lagrangians are discrete. Their quantization yields hypermultiplets in 4d. When this three-sphere is projected to \mathcal{C} we obtain an interval η stretching between two zeros of the quadratic differential ϕ .
- $N \cong S^1 \times S^2$. This class of special-lagrangians always come in one-parameter families. Their quantization yields a vector multiplet in 4d. The projection of any such $S^1 \times S^2$ to \mathcal{C} is a closed loop η .

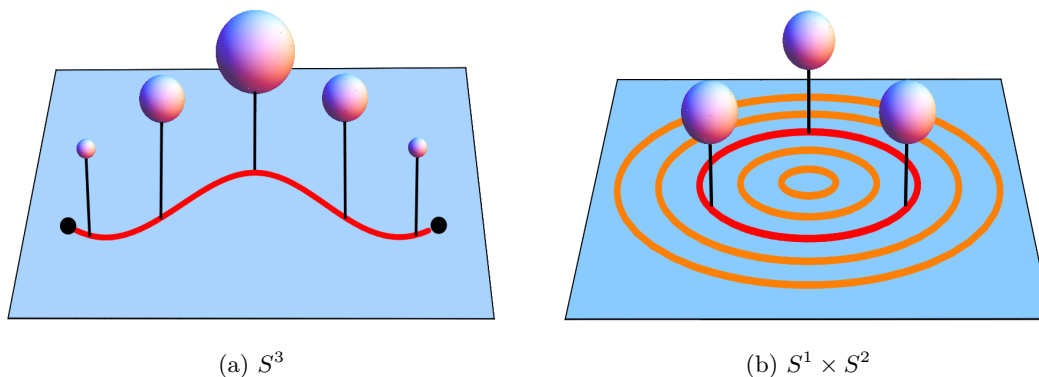


Figure 3.1: Special-Lagrangian geometry in the Calabi-Yau. The blue denotes a patch of the surface \mathcal{C} . The red trajectory denotes the cycle η and the S^2 fibers are indicated schematically above \mathcal{C} . In (a) the topology of the cycle η is an interval which terminates at two zeros of ϕ . The S^2 fibers shrink at these end points yielding a total space of an S^3 . In (b), the cycle η has the topology of a circle, and the total space is $S^1 \times S^2$. Such special-lagrangians always come in one parameter families indicated in orange.

The shape of η in \mathcal{C} is constrained by the special Lagrangian condition (3.2.7) on N . Explicitly if we let $t \in \mathbb{R}$ parametrize η then the condition of constant phase Ω reduces to

$$\sqrt{\phi}|_{\eta} = e^{i\theta} dt. \quad (3.2.10)$$

The ambiguity in choosing the square root appearing in the above reflects the physical fact that for every BPS particle there is also an associated BPS antiparticle of opposite charge. Choosing the opposite sign for the square root then sends $\theta \rightarrow \theta + \pi$, i.e. it replaces a BPS particle by its antiparticle.

We have now arrived at an elegant statement of the problem of calculating BPS states in this class of quantum field theories. Our goal, however, is not directly to use this structure to compute the BPS states, but rather to extract the BPS quiver of this theory. In the following we will explain a natural way to extract such a quiver from a global analysis of the flow equations

(3.2.10).

3.2.1 Triangulations from Special-Lagrangian Flows

Our goal in this section will be to encode certain topological and combinatorial data about the special lagrangian flow in terms of a triangulation of the surface \mathcal{C} . Our basic strategy will be to analyze the local and asymptotic properties of the flow on \mathcal{C} defined by (3.2.10). This is a problem which is well-studied in mathematics [75] and has received much attention in the present physical context [2, 13, 22, 26, 53]. We will confine ourselves to a brief self-contained review. Since a quiver is constructed from hypermultiplets, our focus will be on the trajectories of this flow which interpolate between the zeros of ϕ . Thus a special role will be played by these trajectories.

To begin, we investigate the local nature of the flow near each zero. We assume that this is a simple zero so that, in some holomorphic coordinate $w(x)$ centered at the zero of ϕ , the flow equation (3.2.10) takes the local form

$$\sqrt{w}dw = e^{i\theta}dt \implies w(t) = \left(\frac{3}{2}e^{i\theta}t + w_0^{3/2} \right)^{2/3}. \quad (3.2.11)$$

Because of the three roots of the right-hand-side of the above, each zero has three trajectories emanating from it. These trajectories make angles of $2\pi/3$ with each other and separate a local neighborhood centered on them into three distinct families of flow lines, as illustrated in Figure 3.2.

Aside from the zeros, which can serve as endpoints for BPS trajectories, the other distinguished points for the flow are the punctures of \mathcal{C} . Since the punctures form ideal boundaries of \mathcal{C} , they should be thought of as lying at strictly infinite distance. Thus the behavior of the flow equation near these points governs the asymptotic properties of trajectories at very late and early times. In a local neighborhood centered on the puncture $p_i \in \mathcal{C}$, the flow equation is asymptotically

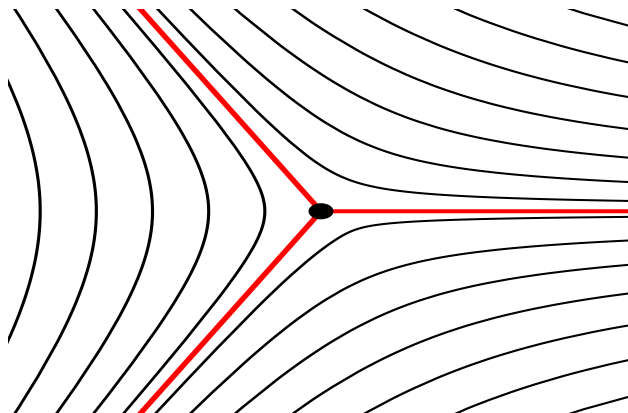


Figure 3.2: The local structure of the flow near a zero of ϕ shown as a black dot at the center of the diagram. The red trajectories are the three flow lines which pass through the zero. The black trajectories denote other generic flow lines.

given by

$$\frac{dw}{w^{1+k_i/2}} = e^{i\theta} dt. \quad (3.2.12)$$

We split our analysis of the solutions into two cases depending on the order $k_i + 2$ of the pole in ϕ at the puncture:

- *Regular Punctures:* $k_i = 0$

The regular punctures in \mathcal{C} are naturally associated to flavor symmetries and hence mass parameters of the engineered field theory [36]. In our analysis this manifests itself in the following way: the residue of the pole in the flow equation is a coordinate invariant complex parameter that is part of the boundary data of the geometry. Restoring this parameter to the asymptotic flow equation we then have.

$$m \frac{dw}{w} = e^{i\theta} dt. \quad (3.2.13)$$

The parameter m is the residue of a first order pole in the Seiberg-Witten differential and can be interpreted as a bare mass parameter.

We deduce the behavior of the late time trajectories by integrating (3.2.13). The solution with initial condition w_o takes the form

$$w(t) = w_o \exp\left(m^{-1}e^{i\theta}t\right). \quad (3.2.14)$$

Assume that the BPS angle θ has been chosen so that $m^{-1}e^{i\theta}$ is not purely imaginary. Then the solution (3.2.14) is a logarithmic spiral. Asymptotically all trajectories spiral in towards the puncture as illustrated in Figure 3.3.

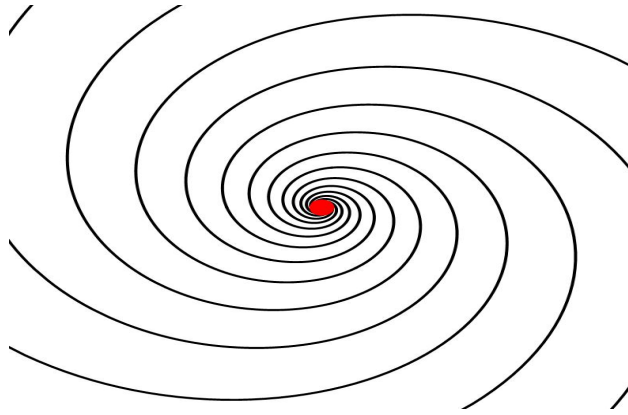


Figure 3.3: The local flow near a regular puncture indicated in red. The flow lines are spirals terminating at the puncture.

- *Irregular Punctures:* $k_i > 0$

In the case of irregular punctures, we find power law behavior for the asymptotic trajectories upon integrating (3.2.12):

$$w(t) = \left(\frac{-2e^{i\theta}}{k_i}t + \frac{1}{w_o^{k_i/2}}\right)^{-2/k_i}. \quad (3.2.15)$$

A key feature of this solution is that it exhibits Stokes phenomena. For large $|t|$ the trajectories converge to the origin $w = 0$ along k_i distinct trajectories. We account for this behavior of the flows by cutting out a small disk in the surface \mathcal{C} centered on the origin in the w plane.

In terms of the metric structure of \mathcal{C} this hole is to be considered of strictly infinitesimal size. The modified surface now has a new ideal boundary S^1 , and the k_i limiting rays of the flows are replaced by k_i marked points on this boundary. This procedure is illustrated in Figure 3.4.

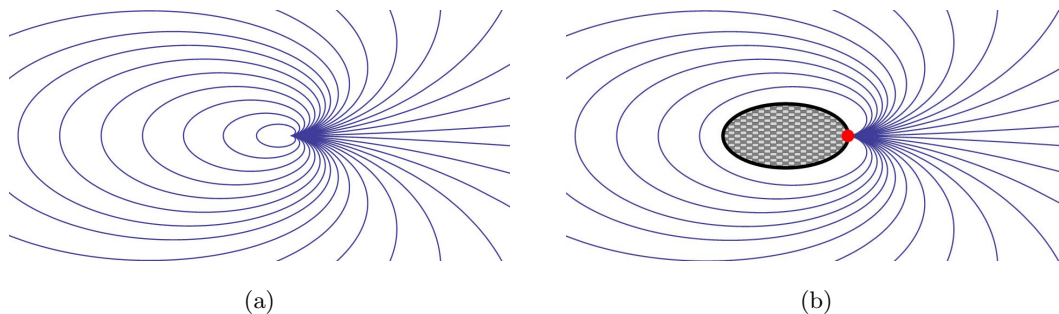


Figure 3.4: Asymptotic flows near an irregular puncture with $k = 1$. In (a) the flow lines converge along a single ray, the rightward horizontal direction. In (b), the surface \mathcal{C} is modified by cutting out the small gray checkered region. This surface now has a boundary, depicted by the black curve. On the modified surface with boundary, generic flows terminate at a point, indicated in red, on the boundary.

For each puncture p_i with $k_i > 0$ we perform the operation described above. At the conclusion of this procedure our modified surface \mathcal{C} now has an ideal boundary component S_i^1 for each irregular puncture p_i and further each S_i^1 is decorated with k_i marked points. From now on, when discussing flows with irregular punctures, the symbol \mathcal{C} shall mean this modified surface, equipped with boundary components containing marked points for each irregular puncture.

Armed with the above, it is easy to deduce the global structure of the flow diagram on \mathcal{C} , that is, the global picture of the solutions to

$$\sqrt{\phi} = e^{i\theta} dt. \quad (3.2.16)$$

We first choose the BPS angle θ *generically*. This means that there are no BPS trajectories in

the flow, and hence no finite length trajectories connecting zeros of ϕ as well as no closed circular trajectories. There are then two types of flow lines:

- *Separating Trajectories*

These are flow lines which have one endpoint at a zero of ϕ and one endpoint at a regular puncture or marked point on the boundary of \mathcal{C} . Separating trajectories are discrete and finite in number.

- *Generic Trajectories*

These are flow lines which have both endpoints at either regular punctures or marked points on the boundary. Generic trajectories always come in one parameter families.

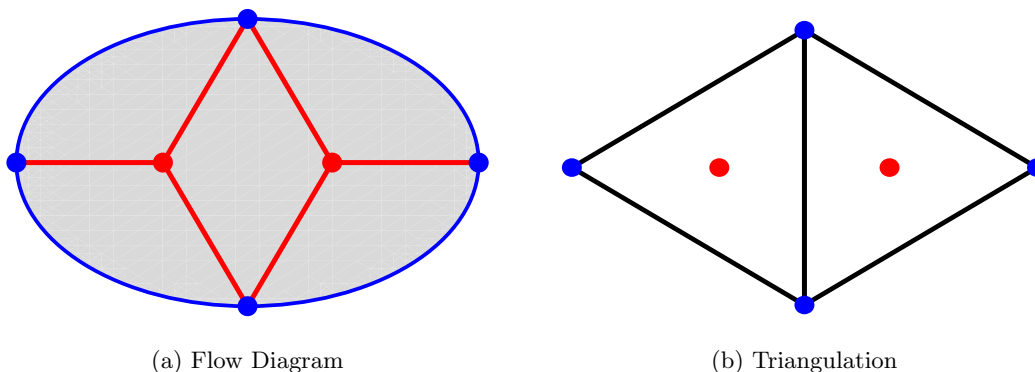


Figure 3.5: An example flow diagram and its associated triangulation. In (a) we have a global flow diagram on a disc with four marked points on the boundary. The red dots are the zeros of ϕ and the associated separating trajectories are the red lines. The gray cells denote one parameter families of generic flows. All flow lines end on the four marked blue dots on the boundary. In (b) we have extracted the associated triangulation. Each black line is a generic flow line selected from each one parameter family. The resulting triangles each contain one zero of ϕ by construction.

A useful way to encode the topological structure of these flow diagrams is the following. We consider our surface \mathcal{C} with boundary. It has marked points in the interior for each regular

puncture, and marked points on the boundary given by the order of the pole of ϕ at the associated irregular puncture. Then, for each one parameter family of generic trajectories, we choose exactly one representative trajectory and draw an arc on \mathcal{C} connecting the indicated marked points. An example is indicated in Figure 3.5b. This procedure produces an *ideal triangulation* of \mathcal{C} where each diagonal of the triangulation terminates at two marked points. Further, by construction, each triangle contains exactly one zero of ϕ . Generally it is possible for the flow to produce an ideal triangulation with *self-folded* triangles; these result in some technical complications which we address in appendix A.

In summary, for a fixed quadratic differential ϕ and generic angle θ , we have produced an ideal triangulation of \mathcal{C} by studying trajectories of

$$\sqrt{\phi} = e^{i\theta} dt. \tag{3.2.17}$$

The combinatorial structure of this triangulation encodes properties of the flow, and we will see in the remainder of this section how to directly extract a BPS quiver and superpotential from this triangulation. Throughout the discussion it will be important to inquire how the triangulation varies as the data (ϕ, θ) varies. The quadratic differential ϕ labels a point in the Coulomb branch of the gauge theories in question, and thus it is natural to fix this data and study the BPS spectrum at fixed point in moduli space. By contrast, the angle θ is completely arbitrary. Any generic angle θ can be used, and different angles will produce distinct triangulations. Demanding that ultimately our results are independent of θ will give a powerful constraint in the upcoming analysis.

3.2.2 BPS Quivers from Ideal Triangulations

We have now arrived at the structure of an ideal triangulation on the surface \mathcal{C} . From this data there is a simple algorithmic way to extract a quiver [54]. As a preliminary definition, we refer to an edge in the triangulation as a diagonal, δ , if the edge does not lie on a boundary of \mathcal{C} . Then proceed as follows:

- For each diagonal δ in the triangulation, draw exactly one node of the quiver.
- For each pair of diagonals δ_1, δ_2 find all triangles for which the specified diagonals are both edges. For each such triangle, draw one arrow connecting the nodes defined by δ_1 and δ_2 . Determine the direction of the arrow by looking at the triangle shared by δ_1 and δ_2 . If δ_1 immediately precedes δ_2 going counter-clockwise around the triangle, the arrow points from δ_1 to δ_2 .

In [12] many aspects of these quivers were explored and it was argued that these are exactly the BPS quivers of the associated quantum field theories. We now provide a full explanation of this proposal.

We first address the identification of the diagonals of the triangulation with the nodes of the quiver. As we have previously explained, our triangulation is constructed at a fixed value of the central charge angle θ appearing in (3.2.10). This angle has been chosen such that no BPS states have a central charge occupying this angle. Now let us imagine rotating θ . Eventually we will reach a critical value θ_c where a BPS hypermultiplet occurs and the structure of the flow lines will jump discontinuously. The key observation is that each triangle in the triangulation contains exactly one zero of ϕ . Then, since BPS hypermultiplets are trajectories which connect zeros of ϕ , a BPS hypermultiplet trajectory must cross some number of diagonals in the triangulation to traverse from one zero to another. A simple example of this is illustrated in Figure 3.6(b).

What the above example illustrates is that each diagonal δ labels an obvious candidate BPS hypermultiplet trajectory, connecting the two zeros in the two triangles which have δ as a common boundary. Further any hypermultiplet trajectory which crosses multiple diagonals can be viewed homologically as a sum of the elementary BPS trajectories which cross only one diagonal. Therefore, diagonals should be nodes of the BPS quiver.

Next let us justify why arrows in the quiver should be described by triangles in the triangulation. Each elementary hypermultiplet, corresponding to a diagonal in the triangulation, lifts to a three-sphere in the Calabi-Yau. Since these three spheres form nodes of the quiver, the lattice generated by their homology classes is naturally identified with the charge lattice Γ of the theory. Further the symplectic pairing given by the electric magnetic inner-product is precisely the intersection pairing on these homology classes. Thus for each intersection point of the three-spheres, we should put an arrow connecting the associated nodes. On the other hand it is clear that this intersection number can be calculated by projecting the three-spheres to \mathcal{C} and then simply counting the signed number of endpoints that the associated trajectories share. Each shared endpoint is naturally associated to the triangle containing it; so the triangles correspond to arrows between nodes.

The result of this section is that, given a Riemann surface \mathcal{C} defining a 4d, $\mathcal{N} = 2$ quantum field theory, we have produced a natural candidate BPS quiver. It is quite interesting to note that as a result of recent mathematical work [54], these quivers are all of *finite mutation type*. In other words, repeated mutations of vertices produce only a finite number of distinct quiver topologies. In fact this property is equivalent to the more physically understandable property of completeness [12]. The set of finite mutation type quivers (or equivalently, the set of complete theories) consists precisely of the quivers associated to triangulated surfaces, as described above, along with a finite number of exceptional cases, discussed in section 3.3 [76].

We can give one strong consistency check on our proposal for the BPS quivers as follows. Observe that, to a given Riemann surface theory \mathcal{C} we have in fact produced not one quiver but many. Indeed our quivers are constructed from the triangulation produced from a fixed value θ of the BPS angle where there are no BPS states. So in fact our assignment is

$$(\mathcal{C}, \theta) \longrightarrow Q_\theta = \text{BPS Quiver.} \quad (3.2.18)$$

As the central charge phase θ varies over a small region, the flow evolves continuously and the incidence data of the triangulation encoded in Q_θ remains fixed. However, as θ varies past a BPS state, the flow lines and triangulation will jump discontinuously, as illustrated in the basic example of Figure 3.6. This results in a new quiver $Q_{\theta'}$, distinct from Q_θ . Both of these quivers Q_θ and $Q_{\theta'}$ are natural candidates for the BPS quiver of theory defined by \mathcal{C} , and hence we should expect that the quantum mechanics theories they define are equivalent. In other words consistency of our proposal demands that all quivers of the form Q_θ for any given θ are mutation equivalent. Happily, a simple theorem [54] shows that this is indeed the case: the set of quivers obtained from triangulations of a given surface precisely forms a mutation class of quivers.

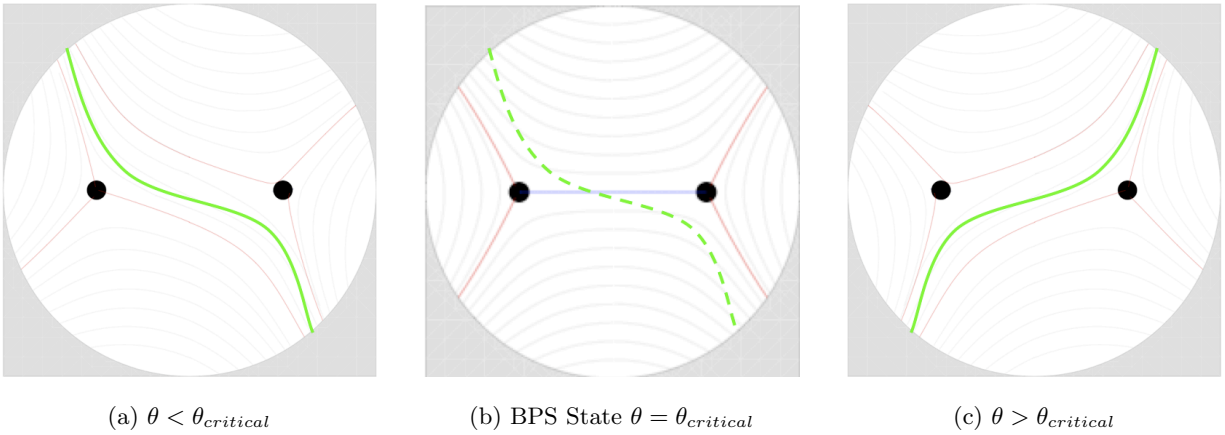


Figure 3.6: Evolution of the special lagrangian flows with the BPS angle θ . In each picture the black dots indicate the branch points of the cover where flows emerge. Red trajectories are flows that emerge from the branch points and terminate on the boundary at $|x| = \infty$, while gray trajectories indicate generic flow lines. The green trajectory denotes a representative of a generic flow line which can serve as an edge in the triangulation. In (b) the BPS angle of the flow aligns with the phase of the central charge and a new kind of trajectory, shown in blue, traverses between branch points. Afterwards in (c) the green line has flipped.

Actually, we can say more. If we tune θ from 0 to 2π , we will see that every BPS hypermultiplet corresponds to a jump of the triangulation, and gives a new choice of quiver. This approach to computing BPS spectra was studied in [13]. As was described there, the discontinuous jump of triangulation, or *flip*, at each BPS state γ is given by simply removing the diagonal crossed by γ , and replacing it with the unique *other* diagonal that gives an ideal triangulation.³ As argued in [54], at the level of the quiver, this flip corresponds precisely to a mutation at the associated node. Thus, if we forget about the surface \mathcal{C} and triangulation, and instead focus on the quiver itself, we see that we are simply applying the mutation method to compute Π -stable representations! This seems to be a deep insight into how the naively unrelated problems of finding special lagrangians and computing Π -stable quiver representations are in fact equivalent. Recall, however, that the mutation method made no reference to completeness of the theory. While the triangulations and flips exist for some set of complete theories, the mutation method is more general, and can be applied any BPS quiver. In [77], and the previous chapter we explored applications of the mutation method to non-complete theories.

In later sections of this paper we will see further evidence for this proposal by recovering the BPS quivers of well-known quantum field theories. However, before reaching this point let us illustrate one important subtlety which we have glossed over in the above. Consider the possible structure in an ideal triangulation of some Riemann surface \mathcal{C} , as illustrated in Figure 3.7. According to the rules of this section, for each bivalent puncture in the triangulation we will obtain, as indicated, a cycle of length two in the quiver. These are fields in the quiver theory which could, in principle, admit a gauge invariant mass term in the superpotential. The quantum mechanics described by the quiver will be rather complicated, if no such mass term is generated. In the next section we will argue that the natural potential for these theories does indeed generate all possible

³To clarify, once we remove the diagonal of the appropriate BPS state, we are left with some quadrilateral in our ‘triangulation.’ To produce a true triangulation, we may add one of the two possible diagonals that would cut the quadrilateral into a triangle. A flip is simply given by taking the choice that differs from the original triangulation.

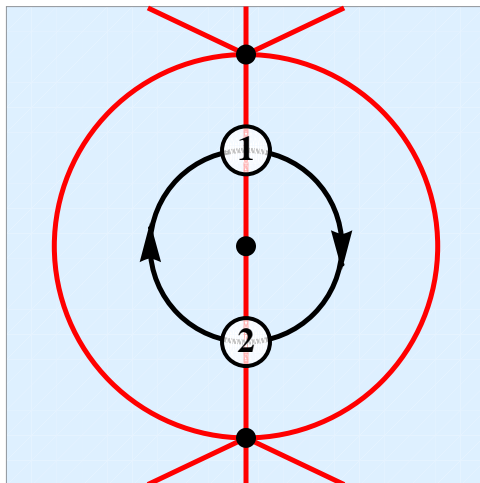


Figure 3.7: A bivalent puncture in the triangulation gives rise to a two-cycle in Q . The blue denotes a patch of \mathcal{C} . Red lines indicate diagonals and marked points are punctures. The nodes of the quiver for the two indicated diagonals are drawn. The bivalent puncture implies that there is a two cycle in the quiver indicated by the black arrows.

gauge invariant mass terms and therefore simplifies the resulting quivers considerably.

3.2.3 The Superpotential

The previous subsection identified a quiver associated to any ideal triangulation, and further suggested that this quiver is naturally the BPS quiver of the associated gauge theory. In this subsection we will complete this picture by describing a natural superpotential for such a quiver, recently developed in the mathematics literature [78–80]. We will then argue on general grounds, essentially as a consequence of completeness, that this superpotential yields the necessary F-flatness conditions for the quiver quantum mechanics theory.

We will build up the superpotential starting from the elementary case of an acyclic quiver. Since such a quiver has no cycles, there are simply no gauge invariant terms to be written and $\mathcal{W} = 0$.

Next we consider an arbitrary quiver Q which, by a sequence of mutations, is connected to an acyclic quiver. Since Q is the quiver of a complete theory, all of its central charges are free parameters that can be varied arbitrarily as one scans over parameter space. It follows that the sequence of mutations connecting Q to its dual acyclic form is in fact realizable by physical variation of parameters. Hence, following the mutation rules of section 2.3, the superpotential for the quiver Q is completely fixed by the acyclic quiver with trivial potential.

The argument of the previous paragraph shows that the \mathcal{W} assigned to any such quiver Q is completely fixed, however complicated the sequence of mutations leading from the acyclic form to Q may be. Surprisingly, there exists an elementary description of this superpotential in terms of the local incidence data of the triangulation of \mathcal{C} which gives rise to Q . This description has been developed in [78]. For any quiver Q mutation equivalent to an acyclic quiver, the superpotential \mathcal{W} is computed as follows:

- Let T denote a triangle in \mathcal{C} . We say T is *internal* if all of its edges are formed by diagonals, that is none of the sides of T are boundary edges in \mathcal{C} . Then each edge of T represents a node of the quiver and the presence of the internal triangle T implies that these nodes are connected in the quiver in the shape of a three-cycle. For each such triangle T we add the associated three-cycle to \mathcal{W} . This situation is illustrated in Figure 3.8a.
- Next let p be an internal, regular puncture in \mathcal{C} . Then some number n of edges in the triangulation end at p . Further since p is an internal puncture which does not lie on the boundary of \mathcal{C} it follows that each such edge terminating at p is in fact a diagonal and hence a node of the quiver. The n distinct nodes are connected in an n -cycle in the quiver and we add this cycle to \mathcal{W} . This situation is illustrated in Figure 3.8b.

For quivers with multiple arrows between two given nodes, it is important to keep track of which triangle the arrow arises from when writing down the superpotential. The superpotential

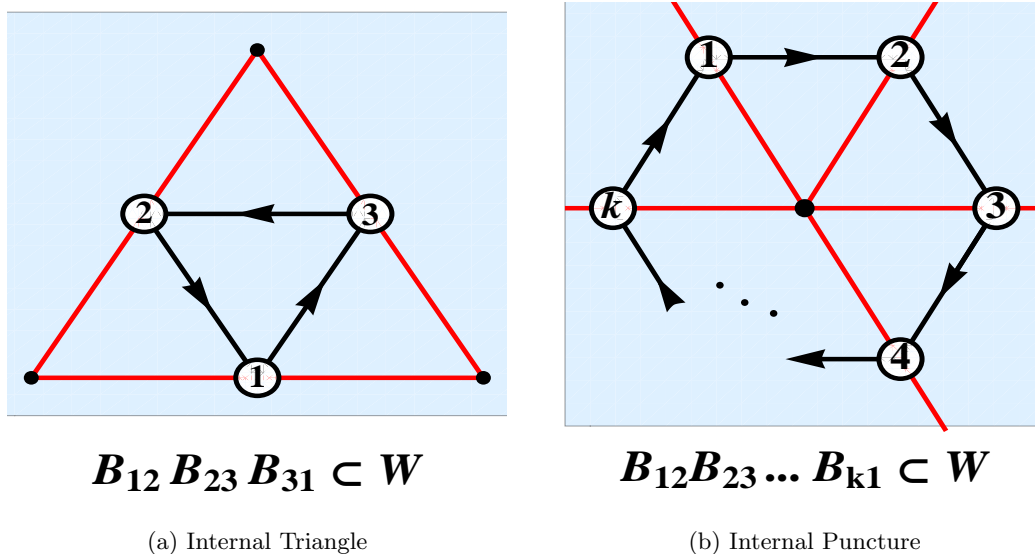


Figure 3.8: The two distinct structures in the triangulation which contribute to the potential. The blue region denotes a patch of \mathcal{C} , the red edges are diagonals in the triangulation. These correspond to nodes of the quiver which we have indicated on the triangulation. The black arrows connecting the nodes are the arrows in the quiver induced by the shared triangles shown in the diagram. In (a) an internal triangle gives rise to a three-cycle in \mathcal{W} in (b) an internal puncture of valence k gives rise to a k -cycle in \mathcal{W} .

must be written with a fixed, consistent assignment of arrows to triangles; inconsistent choices are *not* equivalent, and will generally give the wrong answer.

The observation that the superpotential can be determined in such an elementary way from the incidence data of the triangulation is striking. It strongly suggests that \mathcal{W} is a local object that can be determined patch by patch on \mathcal{C} . Granting for the moment that this is so allows us to immediately generalize to any theory determined by an arbitrary Riemann surface \mathcal{C} . We can simply extend the simple rules given above to all quivers.

One important consequence of this extension is that it automatically ensures that all of our superpotentials will be compatible with mutation. That is, just as in equation (3.2.18),

we have now constructed a map from a Riemann surface \mathcal{C} and an angle θ to a quiver Q and superpotential \mathcal{W} . However the angle θ is arbitrary. As θ rotates, in general the triangulation \mathcal{T} of \mathcal{C} will undergo a series of flips and arrive at a new triangulation $\tilde{\mathcal{T}}$. From this new triangulation we can determine the quiver $(\tilde{Q}, \tilde{\mathcal{W}})$. On the other hand we have previously noted that flips in the triangulation are the geometric manifestation of quiver mutation. Thus we have two independent ways of determining the dual quiver and superpotential:

- Compute $(\tilde{Q}, \tilde{\mathcal{W}})$ from (Q, \mathcal{W}) by performing a sequence of mutations.
- Compute $(\tilde{Q}, \tilde{\mathcal{W}})$ from the new triangulation $\tilde{\mathcal{T}}$

A necessary condition for a consistent superpotential is that the two computations yield the same answer. In [78] it was proved that this is the case.

The above argument shows that our proposal for the superpotential is consistent with the quiver dualities described by mutation. However, it depends fundamentally on our locality hypothesis for the superpotential. As we will now argue, using the completeness property of the field theories in question, we can give a strong consistency check on this assumption.

All of our arguments thus far involve constraints on \mathcal{W} that arise from mutation. As we mentioned in section 2.3 mutations may be forced when, as we move around in moduli space, the central charges rotate out of the chosen half-plane. Most importantly, all these rotations are physically realized, since in a complete theory all central charges are free parameters.

Of course the central charges of the theory come not just with phases but also with magnitudes. In a complete theory we are also free to adjust these magnitudes arbitrarily. Let us then consider the limit in parameter space where the magnitude of the central charge associated to a node δ becomes parametrically large compared to all other central charges

$$|\mathcal{Z}(\delta)| \longrightarrow \infty. \tag{3.2.19}$$

In this limit, the BPS inequality implies that all particles carrying the charge δ become enormously massive and decouple from the rest of the spectrum. At the level of the quiver Q this decoupling operation is described as follows: simply delete from the quiver the node δ and all arrows which start or end at δ . This produces a new quiver \tilde{Q} with one node fewer than Q . The superpotential for the resulting quiver theory \tilde{Q} is then determined simply by setting to zero all fields transforming under the gauge group indicated by δ .

Following our interpretation of nodes of the quiver as diagonals in a triangulation, it is possible to describe this decoupling operation at the level of the Riemann surface \mathcal{C} itself. Consider the diagram of Figure 3.9a which depicts the local region in \mathcal{C} containing a diagonal δ traversing between two punctures or marked points p_i . The decoupling operation to destroy the node δ is then realized by excising a small disc containing δ as a diameter and no other diagonals. The result of this procedure is shown in Figure 3.9b. It is clear from our construction of BPS quivers from triangulations that this decoupling operation produces a new surface $\tilde{\mathcal{C}}$, whose BPS quiver is exactly \tilde{Q} , the quiver with the node δ decoupled. We may therefore determine the superpotential \mathcal{W} for \tilde{Q} by applying the incidence rules described in this section to the new surface $\tilde{\mathcal{C}}$.

In summary, we see that there are two distinct ways for computing the superpotential for the quiver \tilde{Q} :

- Determine from \mathcal{C} the superpotential for the quiver Q . Then reduce to \tilde{Q} by deleting the node δ .
- Determine directly from the surface $\tilde{\mathcal{C}}$ the superpotential for the quiver \tilde{Q} .

Consistency of our proposal demands that the two methods give rise to the same superpotential. It is easy to see directly that this is the case. Indeed the effect of the surgery operation illustrated in Figure 3.9 is to change the two triangles T_i to external ones, and to change the points p_i to marked points on the boundary. Clearly this eliminates from the superpotential exactly those terms in

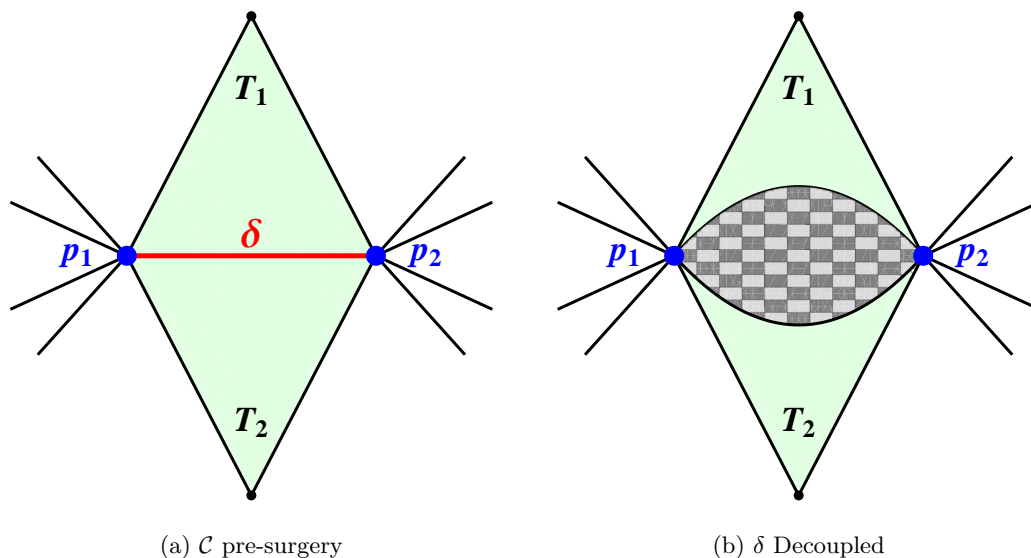


Figure 3.9: The node decoupling surgery for a typical diagonal δ . In (a) we see a patch of \mathcal{C} focused on the region involving a typical diagonal δ . In (b) δ has decoupled leaving a new a new Riemann surface $\tilde{\mathcal{C}}$ which differs from \mathcal{C} by the addition of a new boundary component which encloses the checkered region and has two marked points p_i .

which fields charged under the node δ appear.

By completeness, the decoupling limit argument can be applied to an arbitrary node in a BPS quiver and yields a strong consistency check on the locality hypothesis and thus our proposal for the superpotential.

Let us remark that the superpotential we have constructed naturally resolves the headache proposed at the end of section 3.2.2. By construction, every two-cycle in a quiver arises from a bivalent puncture of the corresponding triangulation. For each bivalent puncture there is now a quadratic term in the superpotential that lifts the fields involved in the associated two-cycle. Thus we may integrate out and cancel all possible two-cycles to produce a two-acyclic quiver.

Finally, before turning to examples, we point out that it would be interesting to calculate this superpotential directly from a string theory construction. While several plausibility and

consistency arguments have been given, a direct calculation may certainly lead to further insight.

3.2.4 Examples from $SU(2)$ Gauge Theory

In this section we illustrate the rules developed above by cataloguing the BPS quivers, with their required superpotential, for simple theories given by a single $SU(2)$ gauge group with matter and asymptotically free or conformal coupling. Of course each theory comes with a number of quivers related by mutations and we need only derive one. Consistent with our previous discussion, for those examples involving irregular punctures, we will present triangulations of surfaces with boundary. In [77] and the previous chapter, the representation theory of these quivers was studied, and found to agree with the well known BPS spectra of the associated theories.

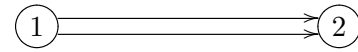
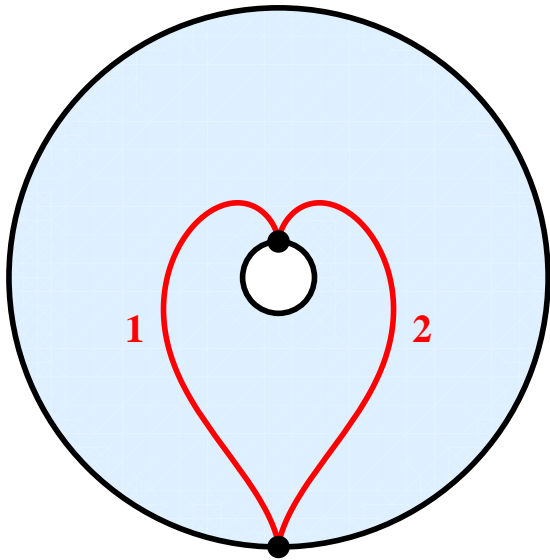
Before enumerating the examples, we take a moment to fix conventions. Throughout, in all triangulations, red labeled lines denote diagonals, which appear as nodes of the quiver, while black lines denote boundary components. Both regular punctures and marked points on the boundary are indicated by black dots. Bifundamental fields corresponding to arrows in the quiver will be denoted by X_{ij} and Y_{ij} where i and j label the initial and final vertex of the arrow respectively.

Asymptotically Free Theories

We first study quivers for $SU(2)$ theories with asymptotically free gauge coupling.

- $SU(2)$

This theory is constructed on an annulus with one marked point at each boundary.

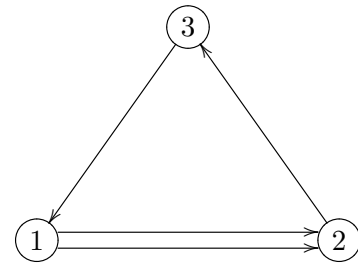
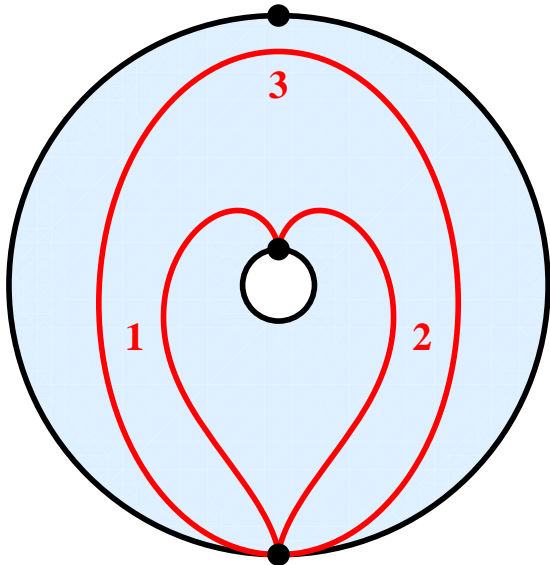


$$\mathcal{W} = 0.$$

Of course this is exactly the quiver for $SU(2)$ Yang-Mills.

- $SU(2)$ $N_f = 1$

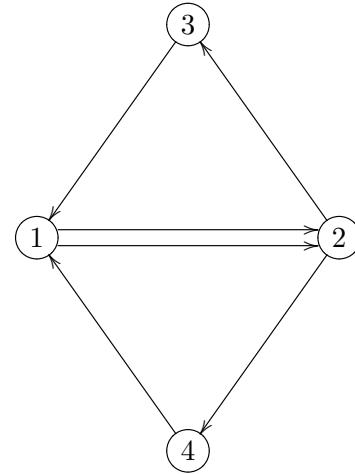
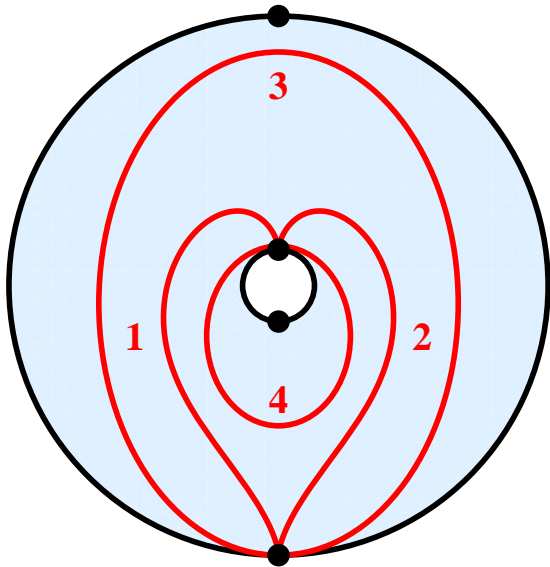
This theory is constructed on an annulus with one marked point on one boundary component, and two marked points on the remaining boundary component.



$$\mathcal{W} = X_{12}X_{23}X_{31}.$$

- $SU(2)$ $N_f = 2$

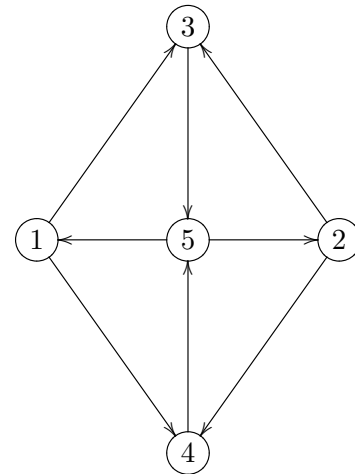
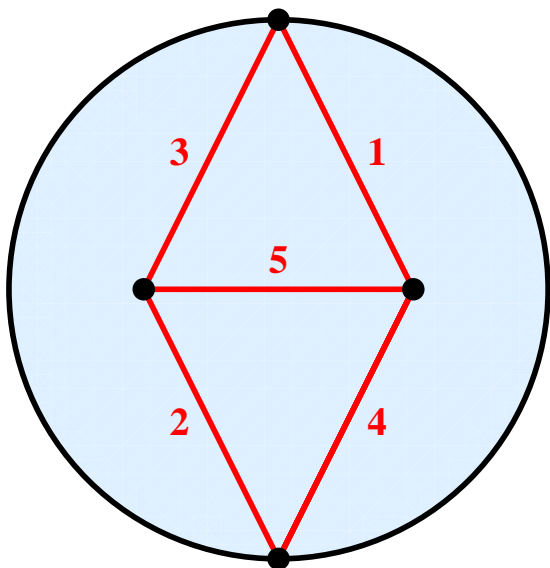
This theory is constructed on an annulus with two marked points on each boundary component.



$$\mathcal{W} = X_{12}X_{23}X_{31} + Y_{12}X_{24}X_{41}.$$

- $SU(2)$ $N_f = 3$

This theory is constructed on a disc with two marked points on the boundary and two punctures.





$$\begin{aligned} \mathcal{W} = & X_{13}X_{35}X_{51} + X_{23}X_{35}X_{52} \\ & + X_{14}X_{45}X_{51} + X_{24}X_{45}X_{52}. \end{aligned}$$

Conformal Theories

While the previous examples illustrate many general features, all the quivers given there are mutation equivalent to quivers without oriented cycles. Thus for those cases the potential is completely fixed by the mutation rules of section 2.3. Now we will consider the case of $SU(2)$ Yang-Mills theories with vanishing beta functions where the conformal invariance is broken only by mass terms. Such quivers arise from triangulations of closed Riemann surfaces and never have acyclic quivers. As such, our proposal for the superpotential is the only known way of constructing \mathcal{W} .

- $SU(2)$ $N_f = 4$

This theory is constructed on a sphere with four punctures. We draw the associated triangulation on a plane omitting the point at infinity.

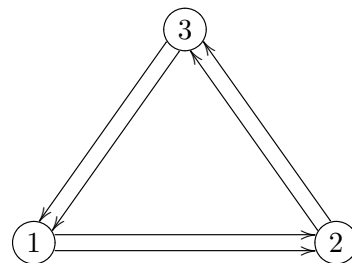
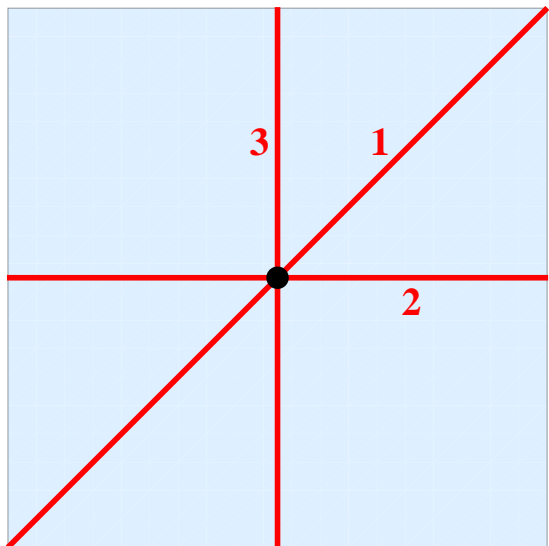


$$\begin{aligned} \mathcal{W} = & X_{15}X_{52}X_{24}X_{41} + X_{13}X_{32}X_{26}X_{61} \\ & + X_{15}X_{52}X_{26}X_{61} + X_{13}X_{32}X_{24}X_{41}. \end{aligned}$$

Notice that this triangulation contains two bivalent punctures; the quiver and superpotential above are obtained after integrating out the corresponding two-cycles.

- $SU(2)$ $\mathcal{N} = 2^*$.

This theory is constructed on a torus with one puncture. We draw the triangulation on a quadrilateral where opposite sides are identified.



$$\begin{aligned} \mathcal{W} &= X_{12}X_{23}X_{31} + Y_{12}Y_{23}Y_{31} \\ &+ X_{12}Y_{23}X_{31}Y_{12}X_{23}Y_{31}. \end{aligned}$$

It is amusing to note that the this quiver for the $\mathcal{N} = 2^*$ theory is in fact invariant under mutation and, consistent with our general discussion, our potential is also mutation invariant.

Building from the examples in this section the reader can easily construct the BPS quiver for a complete theory associated to any arbitrary Riemann surface.

3.3 Exceptional Complete Theories

Thus far in our analysis in this paper we have studied complete gauge theories that are canonically related to Riemann surfaces. These Riemann surface examples constitute all but finitely many of the complete theories with BPS quivers. More generally, the full classification of complete theories consists of [12, 76]:

- All quivers associated to triangulated surfaces, as described in subsection 3.2.2.
- 9 quivers corresponding to $E_n, \widehat{E}_n, \widehat{\widehat{E}}_n$ type Dynkin diagrams, for $n = 6, 7, 8$. E_n and \widehat{E}_n correspond to the usual finite and affine Dynkin diagrams; $\widehat{\widehat{E}}_n$ is given in Figure 3.10.

- Derksen-Owen quivers, X_6, X_7 , given in Figure 3.10 [81].

Having thoroughly investigated the BPS quivers and spectra for complete theories associated to Riemann surfaces, we now take our investigation to its logical conclusion and investigate the BPS spectra of the 11 exceptional cases. By construction, the examples of quivers described here have no interpretation in terms of triangulated surfaces. Thus a priori we have no independent method for fixing the superpotential, and we simply proceed with an ad hoc case by case investigation.⁴

3.3.1 $E_n, \widehat{E}_n, \widehat{\widehat{E}}_n$

The E_n quivers correspond to physical theories that are generalizations of the Argyres-Douglas superconformal theories, and were studied with the affine \widehat{E}_n quivers in [4]. These quivers are acyclic, and thus have no superpotential. Acyclic quivers always contain a chamber in which the only stable states are those given by the nodes themselves. Thus these theories have finite chambers, where the BPS spectra consists of only the nodes themselves.

The $\widehat{\widehat{E}}_n$ quivers were also explored in [12]. They are given by glueing linear acyclic quivers to the quiver of $SU(2)$, $N_f = 3$, (see Figure 3.11). The only cycles available in these quivers are those of the $SU(2)$, $N_f = 3$ quiver; thus we can decouple the acyclic linear pieces as described in subsection 3.2.3. The linear subquivers do not participate in the superpotential, since they are not involved in any cycles of the full quiver; therefore this decoupling does not change the superpotential at all. The superpotential for these quivers is simply the one given by $SU(2)$, $N_f = 3$, shown in Figure 3.11. Since the quivers involved in the glueing (i.e. A_n linear quivers and $SU(2)$, $N_f = 3$) have finite chambers⁵ we conclude that the $\widehat{\widehat{E}}_n$ quivers also have finite chambers.

⁴After completing the manuscript, we were informed that these potentials (excluding X_7) were independently obtained in [82] from slightly different considerations.

⁵We have not described an explicit finite chamber for the $SU(2)$, $N_f = 3$ quiver. However, since it corresponds to a Riemann surface with boundary, namely the disc with two marked points on the boundary and two punctures, we know that a finite chamber exists.

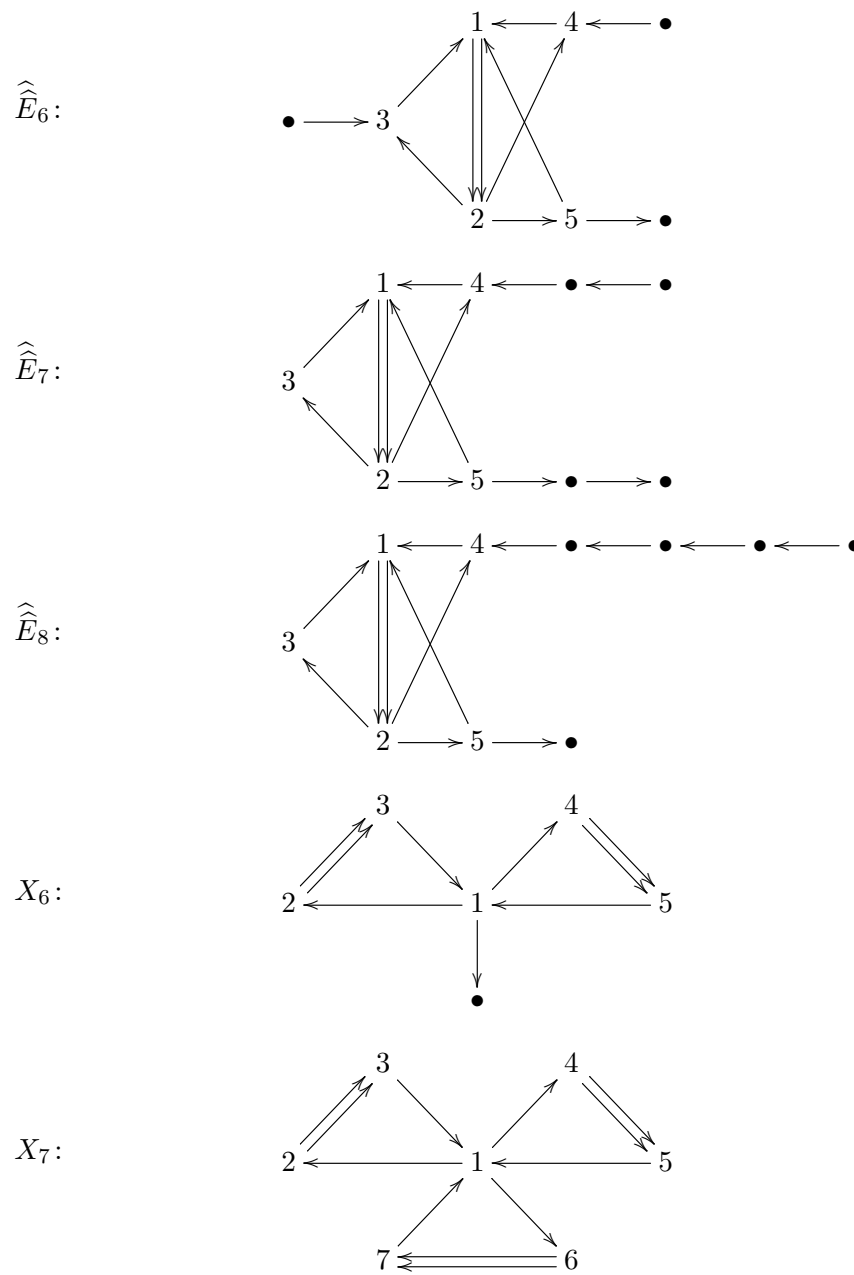


Figure 3.10: The three elliptic E -type Dynkin diagrams oriented as to give finite mutation quivers, and the two Derksen–Owen quivers.

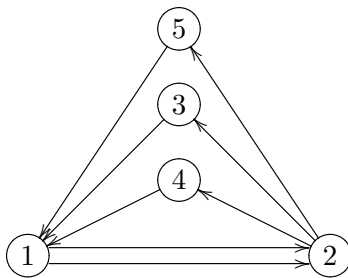


Figure 3.11: Quiver of $SU(2)$, $N_f = 3$. The superpotential is given by $\mathcal{W} = X_{12}X_{23}X_{31} + Y_{12}X_{24}X_{41} + (X_{12} + Y_{12})X_{25}X_{51}$. Notice that this quiver is embedded as a subquiver of the \widehat{E}_n quivers, as shown in Fig. 3.10. A decoupling argument indicates that this gives the correct superpotential for studying the \widehat{E}_n quivers.

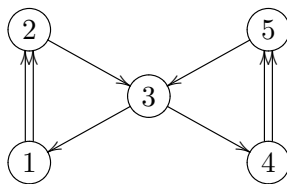


Figure 3.12: Quiver of the annulus with one marked point on each boundary and one puncture, $(0, 1, 2, \{1, 1\})$. The superpotential is given by $\mathcal{W} = X_{12}X_{23}X_{31} + X_{34}X_{45}X_{53} + Y_{12}X_{23}X_{34}Y_{45}X_{53}X_{31}$. Note that this quiver is embedded as a subquiver in X_6, X_7 .

3.3.2 X_6, X_7

The corresponding theories to the Derksen-Owen quivers were also studied in [12]. The X_7 theory is an $SU(2)^3$ gauge theory with a massive hypermultiplet trifundamental. The X_6 theory is a certain decoupling limit of the X_7 .

The X_6 theory can be decoupled to the quiver corresponding to a punctured annulus, with one marked point on each boundary $(0, 1, 2, \{1, 1\})$ without losing any cycles. Thus its superpotential is simply given by the triangulation construction for that theory, as shown in Figure 3.12. Since X_6 can be obtained from a quiver glueing of the punctured annulus quiver to a one-node quiver, this theory also has a finite chamber.

Finally, we consider X_7 . No node of this quiver can be decoupled without removing an oriented cycle, so the approaches used for the other exceptional quivers will not apply. However, the mutation class consists of only two quivers [81]; thus it is easy to check by hand that a proposed superpotential provides a quadratic mass term for all two-cycles generated under mutation. Furthermore, decoupling node 7 should yield the quiver X_6 , with the superpotential given there. From this we are able to guess the superpotential, $\mathcal{W} = X_{12}X_{23}X_{31} + X_{14}X_{45}X_{51} + X_{16}X_{67}X_{71} + Y_{12}X_{23}X_{34}Y_{45}X_{51} + Y_{45}X_{53}X_{36}Y_{67}X_{73}X_{34} + Y_{67}X_{73}X_{31}Y_{12}X_{23}$, which has the desired properties. In principle there are infinitely many higher order terms that could be added to this potential and preserve these properties; this is simply the minimal guess. Exhaustive computational searches via the mutation method have failed to yield a finite chamber for this quiver. Although we have no proof of this statement, it appears that this quiver does not admit any finite chamber.

Chapter 4

Braids Walls and Mirrors

4.1 Introduction

The previous chapters of this dissertation have served to illustrate the intricate nature of the BPS spectrum of a general four-dimensional theory. One outcome of these calculations appears to be the fact that the BPS states are a powerful invariant of the ultraviolet behavior of the quantum field theory. In particular it is natural to conjecture that for each set of BPS charges and degeneracies, there is at most one consistent $\mathcal{N} = 2$ theory. Thus, the BPS structure, which can roughly be viewed as an IR data, appears powerful enough to reconstruct the full UV description of the theory. However, despite these developments, there is no simple explicit map from the four-dimensional BPS data to the UV description of the field theory. In particular, as the BPS states carry both electric and magnetic charges, there is in general no local Lagrangian description of their interactions.

The story may be simpler for $\mathcal{N} = 2$ theories in three dimensions. These theories are close cousins of $\mathcal{N} = 2$ theories in four dimensions, but they have half as much supersymmetry. They enjoy a real central charge in the BPS algebra. Moreover they have the advantage, compared to their 4d cousin, that even abelian theories are UV-complete. Thus, one can in principle hope that

given an effective Lagrangian description of all the BPS states, the same Lagrangian may describe the theory in the UV.

In this chapter, we provide a link between the study of 4d BPS states and 3d BPS states by constructing 3d $\mathcal{N} = 2$ theories from parent 4d $\mathcal{N} = 2$ theories. One natural way to carry out such a reduction, is to consider a one-parameter family of 4d models, parameterized by an extra circle where as one goes around the circle one identifies the two sides up to some symmetry transformation. In [83,84], the corresponding symmetry was an element of the S-duality group of the $\mathcal{N} = 2$ theory. Another choice, studied in [3,4] was to use the R-symmetry of conformal $\mathcal{N} = 2$ theories to reduce the theory. Our construction of 3d theories from 4d theories is close in spirit to [3,4], except that the circle is replaced by real line with suitable boundary conditions at infinity.

In our description, the 3d theory will appear as a 3d domain wall inside a 4d theory. This wall is characterized by a one-parameter flow of the 4d BPS central charges Z_i . As we traverse the thickness of the 3d wall from one side to the other, the Z_i vary along parallel lines while preserving their phase order, and the boundary conditions of the flow are such that asymptotically all central charges become infinitely large. As a consequence of these boundary conditions, all degrees of freedom of the 4d bulk theory, except the massless $U(1)$ gauge multiples, become infinitely heavy and decouple from the 3d wall theory. However, the BPS particles of the 4d theory have finite mass on the wall and are trapped there. Thus, the result of this construction is a 3d theory with $\mathcal{N} = 2$ supersymmetry whose BPS states are inherited from the parent 4d and are potentially gauged under the $U(1)$ symmetries of the bulk. We call this flow of the 4d theory the ‘R-flow’ due to the fact that at the two boundaries the central charges have flipped sign and hence have undergone an R-symmetry rotation, $Z_i \rightarrow e^{i\theta} Z_i$, by $\theta = \pi$. The most important feature of this wall is that, because the R-flow respects the phase order of the 4d central charges, each BPS state of a given 4d chamber will give rise to a trapped particle on the 3d domain wall.

Our reduction of a given 4d theory to a 3d theory is not unique, as one could in principle

start the R-flow from different chambers in 4d related by BPS wall-crossing. This results in a different set of trapped modes on the 3d wall and hence determines a correspondence between 3d theories constructed by R-flow from a parent 4d model, and the BPS chambers of the parent. We thus have the analog of induced wall-crossing in three dimensions, and as we will see, this 3d wall-crossing phenomenon can be interpreted as mirror symmetry. As a result, we find a set of 3d dual theories which are labeled by chambers of the parent 4d theory. And further, the 4d Kontsevich-Soibelman wall-crossing formula enforces partition function equality of these 3d dual theories. For example, the simplest non-trivial $\mathcal{N} = 2$ superconformal theory is the A_2 Argyres-Douglas theory [34]. In this case we have two or three 4d BPS states depending on the choice of chamber, and under R-flow, these lead to two dual theories in 3d, known as $N_f = 1$ SQED and the XYZ model [15].

In the process of reduction of 4d theories to 3d, supersymmetry demands that we vary all the central charges along parallel lines. However, this is not generally possible for arbitrary 4d theories, as the space of allowed central charges is a subspace of all allowed complex numbers. An exception is the case of ‘complete’ $\mathcal{N} = 2$ theories, which are characterized by the property that all their central charges can be varied arbitrarily [12]. Thus, the reduction of complete $\mathcal{N} = 2$ theories from 4d to 3d will constitute the main example of this paper.

As shown in [12], and discussed in chapter 3, with the exception of eleven cases, all complete theories (which have BPS quivers) come from two M5-branes wrapping a punctured Riemann surface as studied for example in [13, 22, 26, 35, 36]. Thus, their reduction to 3d will correspond to two M5-branes wrapping a one-parameter family of Riemann surfaces. In other words, it is a 3d theory determined to two M5-branes wrapping a 3d geometry M . For this class of theories we make contact with the recent work [16]. The BPS data of the 4d parent theory is governed by a triangulation of the associated Riemann surface. During the R-flow, this triangulation evolves by a sequence of flips each of which corresponds to a 4d BPS state. Remarkably, exactly the sequence

of flips prescribed by R-flow determines a decomposition of the 3d geometry M into tetrahedra, where each tetrahedron is in direct correspondence with a 3d BPS particle. This picture leads to an explanation and extension of the rules proposed in [16].¹ Each 4d chamber, which corresponds to a dual 3d description, maps to a particular tetrahedral decomposition M , and 4d wall-crossings, reinterpreted as 3d mirror symmetries, manifest themselves as changes in the number of tetrahedra.

The 3d geometry described by M , together with its decomposition into tetrahedra, encodes the physics of two M5-branes and hence can be viewed as a non-abelian UV data of the theory. However, in the IR, this non-abelian structure is higgsed to an abelian one. As a result, the physics is captured by the geometry of a *single* recombined M5-brane \widetilde{M} which is a double cover of M branched along a knot. The IR geometry \widetilde{M} is the direct 3d analog of the Seiberg-Witten curve for 4d $\mathcal{N} = 2$ theories and we develop its properties in detail. We find that the R-flow of the parent 4d theory naturally determines a braid diagram presentation of the branching knot. The geometry of this branching braid is the key to deciphering the 3d theory. Each intersection of the branching braids describes a massless 3d particle. Giving the particle mass resolves the intersection and, in simple cases, determines a correspondence between general 3d particles and braid moves. Further, the geometry of the braid also encodes the existence of certain superpotentials. These superpotentials are generated by M2-brane instantons ending on the M5-brane, and are seen as primitive polygons in the braid diagram.

We illustrate these ideas in the context of the *ADE* Argyres-Douglas theories. For example, for A_n theories, there are various chambers ranging from n particles to $n(n+1)/2$ particles. This in turn translates to a UV 3d geometry with a minimum number of n tetrahedra and a maximum of $n(n+1)/2$ tetrahedra. In the IR this same theory is described by a branching braid on $n+1$ strands, with particles described by braid moves, and generically cubic and quartic superpotentials. For the R-flow of the *E*-case, as we will demonstrate, we can still obtain the resulting 3d theory.

¹ For instance we find that some, but not all superpotential terms arise from tetrahedra sharing an edge.

However, since these do not correspond to multiple M5-branes, the corresponding 3d theories are not captured by a 3d geometry.

Perhaps the most exciting new 3d theories correspond to the case where the bulk 4d theory has infinitely many BPS particles. This for example, happens for the weak coupling phase of pure $SU(2)$. In this case there are infinitely many dyonic BPS states. However, unlike the 4d case where the dyons have unbounded masses, their reduction to 3d can lead to nearly equal and finite mass for the trapped dyons. Moreover the 4d vector W-bosons can also be trapped on the 3d wall. In this way it appears that all of the infinitely many BPS states of the 4d theory assemble themselves into a representation of $SL(2, \mathbb{R})$, and it is natural to conjecture that the trapped W-bosons mean that the $SL(2, \mathbb{R})$ symmetry is gauged. What is remarkable, is that this theory also has a strongly coupled phase with only two particles which should describe its 3d dual. In terms of the 3d geometry, this phenomenon corresponds to situations where, as the hyperbolic structure is varied, the manifold M goes from having a finite ideal tetrahedralization, to a decomposition into infinitely many accumulating tetrahedra. It appears that similar phenomena have been studied in the math literature [85, 86].

The organization of this chapter is as follows: In section 2 we study the geometry of M5-branes wrapping special Lagrangian cycles of Calabi-Yau threefolds, leading to $\mathcal{N} = 2$ theories in their three uncompactified directions. We explore the emergence of the 3d recombined M5-brane geometry which encodes the 3d gauge theory, as a direct 3d analog of Seiberg-Witten geometry for $\mathcal{N} = 2$ theories in 4d. We focus on the main example of the paper which involves two M5-branes. In this case we explain how the geometry of the branching knot encodes the 3d physics. In particular we show how the Seifert surface of the knot encodes the description of the $U(1)^k$ Chern-Simons gauge theories, with the Seifert matrix giving the matrix of the Chern-Simons levels. In section 3 we introduce the main notion of R-flow and explain our reduction of $\mathcal{N} = 2$ theories in 4d to $\mathcal{N} = 2$ theories in 3d. In section 4 we recall some basic facts about 4d BPS states described in previous

chapters and isolate the key features relevant to our constructions. In section 5 we provide some concrete examples, and study the R-flow of 4d A_n Argyres-Douglas theories and their resulting braids. In particular, we explain how the A_1 theory (a 4d free hypermultiplet) maps to two M5-branes wrapping the tetrahedron geometry and show how the double cover of the tetrahedron is a special Lagrangian lens space in \mathbb{C}^3 . We also show how the chambers of the A_2 theory map to the XYZ model and $N_f = 1$ SQED, and explain how the 4d wall-crossing leads to 3d mirror symmetry. We also discuss some aspects of other A_n theories and show that they correspond to UV 3d geometries comprised of n -pyramids (in the minimal chamber). In section 6 we discuss the case where we have infinitely many particles corresponding to weak coupling limit of $SU(2)$.

4.2 Five-Branes on Three-Manifolds

One purpose of this paper is to describe a class of three-dimensional $\mathcal{N} = 2$ quantum field theories which can be engineered by wrapping M5-branes on three-manifolds which we generically denote by M . In later sections of the paper our primary applications will be to the case involving two five-branes though the geometry described in this section applies more generally.

4.2.1 Three-Dimensional $\mathcal{N} = 2$ Gauge Theories

Let us begin by recalling the basic parameters and properties of the field theories in question [15]. We will be focused on describing the degrees of freedom in the infrared on the Coulomb branch where all non-abelian gauge symmetries have been higgsed to a product of $U(1)$ factors. The data of such a field theory is then:

- A gauge group $U(1)^N$.
- A flavor group $U(1)^F$ with an associated real mass parameter m_i for each $U(1)$ factor.
- A symmetric matrix of k_{ij} of Chern-Simons terms.

- A spectrum of charged chiral matter multiplets X_i .
- A superpotential, \mathcal{W} , a holomorphic function of chiral fields.

An important fact is that in three dimensions, abelian gauge fields with field strength F are dual to scalars γ via the relation

$$*F = d\gamma. \quad (4.2.1)$$

Charge quantization means that γ is periodic, and in simple cases the resulting theory after duality enjoys a flavor $U(1)$ which acts on the dual photon γ as a shift. Under this duality, the real mass parameter m of the dual flavor symmetry can be interpreted as the real FI parameter ζ of the original gauge group. However, in general it is not always true that shifts of the dual photon appear as flavor symmetries of the theory. If σ denotes expectation value of the real scalar in the $U(1)$ gauge multiplet, then after duality the monopole operator

$$\mathcal{M} = \exp(\sigma + i\gamma) \quad (4.2.2)$$

is a chiral field which carries charge under the candidate $U(1)$ flavor symmetry. In particular, if say \mathcal{M} appears in the superpotential, then the flavor symmetry will be broken and correspondingly there is no real mass parameter, or equivalently no FI-term for the original gauge theory.

Next we consider the central charge of 3d $\mathcal{N} = 2$ theories. Just as in four-dimensional $\mathcal{N} = 2$ theories, the superalgebra admits the appearance of a central term Z which sets the BPS bound for the masses of particles carrying $U(1)$ charges. However, unlike the situation in four dimensions where Z is complex, in three dimensions the central charge is real. If q_j and f_i denote gauge and flavor charges respectively, then the total central charge of a particle is

$$Z(q, f) = \sum_j q_j \sigma_j + \sum_i f_i m_i. \quad (4.2.3)$$

Where in (4.2.3) we have implicitly included FI terms as real masses to dual flavor groups. Then, as stated above, charged particles satisfy a bound on their mass

$$m \geq |Z|. \quad (4.2.4)$$

Charged BPS states saturate the above and, in the simplest case of minimal spin, form chiral multiplets.

Finally, we take a moment to discuss Chern-Simons terms. In general, we study theories involving fermions and thus $\mathbb{R}^{1,2}$ (or any other manifold on which we study a three-dimensional field theory) is equipped with a choice of spin structure. In this situation, the correct quantization condition for the level matrix is half-integral units, $k_{ij} \in \frac{1}{2}\mathbb{Z}$ [87]. For convenience, we therefore introduce the notation $\hat{k}_{ij} \equiv 2k_{ij}$. Then, \hat{k}_{ij} is integrally quantized. Concretely, given a collection of $U(1)$ gauge fields A_i with canonically normalized kinetic terms, \hat{k}_{ij} appears in the action as

$$\sum_{ij} \frac{\hat{k}_{ij}}{4\pi} \int A_i \wedge dA_j. \quad (4.2.5)$$

From now on we will always work with the quantity \hat{k}_{ij} , and refer to it as the level.

We also note that in three dimensions, CS levels receive anomalous contributions from integrating massive fermions at one loop. Specifically, if $(q_F)_i$ denotes the vector of $U(1)$ gauge charges of a chiral fermion F with mass m_F then the effective levels are related to the bare ones as

$$(\hat{k}_{ij})_{\text{eff}} = (\hat{k}_{ij})_{\text{bare}} + \sum_F (q_F)_i (q_F)_j \text{sign}(m_F). \quad (4.2.6)$$

For answering questions about the physics in the extreme IR it is the effective levels which are the relevant ones. Indeed, assuming that all matter fields are massive, they may be integrated out leaving a pure Yang-Mills-CS theory with level matrix $(\hat{k}_{ij})_{\text{eff}}$. However, from the right-hand-side of (4.2.6) we can see that the effective levels depend on the masses of fields which in turn depend on the parameters and moduli (m_i, ζ_j, σ_k) . By contrast the bare CS terms are a globally well-defined property of a theory. Thus, in the following, when we compute CS terms we will always have in

mind the bare contribution. The effective levels can then be determined from a knowledge of the spectrum and an application of (4.2.6).

4.2.2 One Five-Brane

Now we study a class of three-dimensional $\mathcal{N} = 2$ field theories that can be constructed from M-theory. We let Q denote a Calabi-Yau threefold and consider M-theory on the spacetime

$$\mathbb{R}^{1,4} \times Q. \tag{4.2.7}$$

We pick a linear subspace $\mathbb{R}^{1,2} \subset \mathbb{R}^{1,4}$ and consider a three-manifold M embedded inside Q as a special Lagrangian. We then consider the effective three-dimensional field theory determined by a single M5 brane on

$$\mathbb{R}^{1,2} \times M. \tag{4.2.8}$$

In the field theory limit, which is all that is relevant for this paper, we are interested only in the local dynamics near M inside the Calabi-Yau Q . Then, we may consider a scaling limit where Q is taken to be non-compact and gravity is decoupled from the degrees of freedom determined by the five-brane. By construction, the resulting field theory admits four supercharges and hence has $\mathcal{N} = 2$ supersymmetry in the three-dimensional sense. We will see that the structure of this field theory is intimately related to the geometry of M .

Geometry of the Coulomb Branch

A basic observation is that there are scalar degrees of freedom describing the small fluctuations of the special Lagrangian M inside the local Calabi-Yau threefold Q . To characterize these, we first note that near M , Q can be modeled by the cotangent bundle T^*M , and hence to describe deformations it suffices to think of M embedded inside its own cotangent bundle as the special Lagrangian zero section. To be explicit, we may introduce a system of local coordinates x_i

on M . Then, any one-form can be expressed locally as $y_i dx_i$, and hence the y_i provide a natural set of coordinates on the cotangent directions to M . In terms of these our starting point for studying deformations is therefore the special Lagrangian

$$M = \{(x, 0) \in T^*M\}. \quad (4.2.9)$$

Consider a deformation of M . Since M deforms in its cotangent bundle its local motion is described by activating a certain one-form λ . In other words, M has deformed to the locus of points M' of the form

$$M' = \{(x, \lambda(x))\} \subset T^*M. \quad (4.2.10)$$

To minimize the energy, the deformation M' must also be special Lagrangian. In the linear approximation, such deformations are canonically identified with the space of harmonic one-forms on M . To see this we note that in terms of the local coordinates (x, y) on T^*M the symplectic form ω has the canonical expression

$$\omega = dy_1 \wedge dx_1 + dy_2 \wedge dx_2 + dy_3 \wedge dx_3 = d(y_i dx_i). \quad (4.2.11)$$

Therefore on the deformed locus M' , the symplectic form restricts to

$$\omega|_{M'} = d(\lambda). \quad (4.2.12)$$

Since we wish M' to be Lagrangian, the restriction of ω to M' must vanish and hence λ must be closed.

We can perform a similar calculation with the local holomorphic three-form Ω on T^*M . Restricted to the deformation locus M' the imaginary part of Ω appears to first order in λ as

$$\Im(\Omega|_{M'}) = d(*\lambda). \quad (4.2.13)$$

To ensure that the deformation is special, the imaginary part of Ω must vanish when restricted to M' . This implies that the one-form λ is co-closed and hence harmonic on the original three-manifold M .

Thus, in the linear approximation, the classical moduli space of special Lagrangian deformations of the three-manifold M can be identified with the vector space of harmonic one-forms which can in turn be identified with the cohomology group $H^1(M, \mathbb{R})$ via Hodge theory. To generalize beyond the linear approximation we now invoke a theorem of Mclean [88] which ensures that every first order supersymmetric fluctuation of M can in fact be integrated to a supersymmetric deformation of finite size. Hence, the full non-linear classical moduli space of deformations of the special Lagrangian M inside Q can be identified with a manifold whose tangent space at M is canonically the space of harmonic one-forms.

Now, supersymmetry dictates that all fields must appear in representations of the $\mathcal{N} = 2$ superalgebra. In particular, this means that the real scalars we have found must in fact be paired with other bosons. To find the remaining half of the bosonic fields, we recall that the five-brane theory supports a two-form field B propagating on its worldvolume. This field can be activated for zero cost in energy provided that the field strength is vanishing $dB = 0$. On the other hand, B itself is only defined up to gauge transformations which shift its value by an exact two-form. Hence, flat B fields on M yield a space of deformations of dimension $h^2(M, \mathbb{R})$. To be completely precise we should also note that as a gauge field, B is naturally a periodic variable and hence the correct cohomology measuring B is valued in \mathbb{R}/\mathbb{Z} . If we combine these scalars with those arising from fluctuations of M we find that locally, the classical five-brane moduli space can be parameterized by

$$H^1(M, \mathbb{R}) \times H^2(M, \mathbb{R}/\mathbb{Z}). \quad (4.2.14)$$

Three-dimensional Poincaré duality ensures that the two vector spaces introduced above are of equal dimension and implies that these scalars fill out $\mathcal{N} = 2$ chiral multiplets.

To a low-energy three-dimensional observer in $\mathbb{R}^{1,2}$, the scalar degrees of freedom that we have identified have a natural interpretation in terms of the classical coordinates on the Coulomb branch of an effective $U(1)^{b_1(M)}$ gauge theory. Indeed the one-form λ is characterized by its periods

and describes the expectation values of the real adjoint scalars σ appearing in $\mathcal{N} = 2$ vector multiplets. Meanwhile the circle valued variables described by the periods of B are the expectation values of dual photons γ . If we introduce a basis of one-cycles α_i and a Poincaré dual basis of two-cycles β_j then an explicit set of local coordinates along the moduli space is given by

$$\int_{\alpha_i} \lambda = \sigma_i, \quad \int_{\beta_j} B = \gamma_j. \quad (4.2.15)$$

The fact that the moduli space can be coordinatized in terms of periods is the starting point for a kind of real special geometry which governs the classical effective action.² This real special geometry is the three-dimensional counterpart to the holomorphic special geometry of four-dimensional $\mathcal{N} = 2$ systems. However unlike the situation there where non-renormalization theorems protect the form of the metric from quantum corrections, a three-dimensional $\mathcal{N} = 2$ system has only four supercharges and hence the metric is subject to quantum corrections. Nevertheless, the observation that the central charges of particles can be characterized in terms of the periods of a one-form λ will play a crucial role in the remainder of this paper. These central charges are protected from quantum corrections involving chiral multiplets [15, 90], and hence the periods of λ will remain meaningful when we study the quantum behavior of the theory.

There is an important subtlety in the above description of the Coulomb branch which arises due to the fact that three-dimensional Yang-Mills theories admit Chern-Simons terms. In the presence of a non-vanishing level \hat{k} the equation of motion for a three-dimensional $U(1)$ gauge field with field strength F and Yang-Mills coupling e is modified to

$$\Delta F \sim (\hat{k}e^2)^2 F. \quad (4.2.16)$$

² Indeed, in generalizing beyond the linear approximation, one finds a real prepotential $F(\sigma)$ characterized by the condition that $\frac{\partial F}{\partial \sigma_i} = \int_{\beta_i} * \lambda$. In terms of F the full non-linear metric on the classical moduli space is then [89]

$$ds^2 = \frac{\partial^2 F}{\partial \sigma_i \partial \sigma_j} d\sigma_i \otimes d\sigma_j + \left(\frac{\partial^2 F}{\partial \sigma_i \partial \sigma_j} \right)^{-1} d\gamma_i \otimes d\gamma_j.$$

This equation means that the propagating photon has been given a non-zero mass $\hat{k}e^2$. In particular this implies that the expectation value of the dual photon γ is frozen to zero. By supersymmetry the same is in fact true for the adjoint scalar, σ . In our geometric context this has the following significance. The quantities γ and σ are measured by periods of the two form B and the one-form λ over a certain two-cycle β and dual one-cycle α . If these periods are frozen to zero then at the level of cohomology valued in \mathbb{R} the associated cycles cannot be detected, and hence the Betti numbers $b_1(M)$ and $b_2(M)$ have each been decreased by one unit.

However, the cycle can still be detected by the more refined data of the integer valued homology. In the presence of a non-zero level \hat{k} for the $U(1)$, there are observables given by the holonomy of the gauge field along cycles C in $\mathbb{R}^{1,2}$

$$\exp\left(iq \oint_C A\right). \quad (4.2.17)$$

For such an operator, the charge q is naturally valued in $\mathbb{Z}_{\hat{k}}$. Indeed, given two such observables, the correlation function is [91]

$$\left\langle \exp\left(iq_1 \oint_{C_1} A\right) \exp\left(iq_2 \oint_{C_2} A\right) \right\rangle \sim \exp\left(\frac{2\pi i q_1 q_2}{\hat{k}} [C_1, C_2]\right). \quad (4.2.18)$$

Here the quantity $[C_1, C_2]$ denotes the integer valued linking number between the curves C_i . From the form of this correlation function, we see that if q vanishes mod \hat{k} then the Wilson line (4.2.17) has trivial correlation functions thus illustrating that q is valued in $\mathbb{Z}_{\hat{k}}$.

In our context, gauge charges for the theory are captured by $H_1(M, \mathbb{Z})$. Then, if the CS level is \hat{k} we see from the above discussion that we only expect mod \hat{k} charges. In other words the CS level is \hat{k} if and only if $H_1(M, \mathbb{Z}) = \mathbb{Z}_{\hat{k}}$. This is our desired result: CS levels for the $U(1)$ gauge theory are encoded in geometry by torsion classes in $H_1(M, \mathbb{Z})$. We can extend this observation to the case where we have many $U(1)$'s. In full generality, the relationship between the homology of M and the gauge theory in $\mathbb{R}^{1,2}$ is as follows. Let the gauge theory be that of n $U(1)$ gauge fields with a level matrix \hat{k}_{ij} . Then $H_1(M, \mathbb{Z})$ is generated by n elements Γ_i modulo relations defined by

the image of \hat{k}_{ij}

$$H_1(M, \mathbb{Z}) \cong \bigoplus_{i=1}^n \mathbb{Z}[\Gamma_i] / \left(\hat{k}_{ij} \Gamma_j = 0 \right). \quad (4.2.19)$$

This equation has the key feature that $b_1(M)$ counts the number of zero-eigenvalues of \hat{k}_{ij} and hence captures the number of propagating massless gauge fields. The remaining non-degenerate part of \hat{k}_{ij} encodes the torsion structure of the homology. The fact that the charges of $U(1)^n$ Chern-Simons theories are captured by (4.2.19) is well known (see *e.g.* [92, 93]), and we return to concrete applications of this formula in our study of examples in section 4.2.4.

Finally, to complete our geometric description of the massless sector of the Coulomb branch, we will now describe how to include FI terms and real masses into the description. Both of these deformations are naturally associated with activating bulk moduli of the ambient Calabi-Yau Q . In fact both arise from a variation in the Kähler class of Q . To see this let us suppose that the symplectic form ω is varied to a new class

$$\omega \rightarrow \omega + \delta\omega. \quad (4.2.20)$$

Then, since ω enters in determining the Lagrangian condition on submanifolds, the deformation above enters our description as a modification in the behavior of the one-form λ as in (4.2.12)

$$d\lambda = \delta\omega|_M \neq 0. \quad (4.2.21)$$

The interpretation of the above modification depends on the behavior of $\delta\omega$ restricted to M . Specifically, since $\delta\omega$ is closed its restriction to M can be in general a sum of terms which are cohomologically trivial or non-trivial. We examine the effects of each of these:

- $\delta\omega$ restricts to M to an exact form $d\eta$. Then, the one-form λ is modified to include a contribution from η

$$\lambda \rightarrow \lambda + \eta. \quad (4.2.22)$$

Such a modification permits λ to develop periods over contractible one-cycles in M and it is these periods which are interpreted as real mass parameters. They are well defined as a consequence of the fact that the symplectic form is closed.

To see the connection of the modulus $\delta\omega$ to a flavor symmetry, we note that this modulus is in the same $\mathcal{N} = 2$ multiplet as the bulk $U(1)$ gauge field A which descends from the reduction of the M-theory three-form C as

$$C = \delta\omega \wedge A. \tag{4.2.23}$$

From the point of view of the five-brane, the field A is non-dynamical, and therefore fields charged under A carry a flavor charge. The expectation value of the scalar modulus $\delta\omega$ then determines the associated mass.

To be precise, one should view the non-vanishing contribution to $d\lambda$ as being supported at infinity in the Calabi-Yau Q , and the real mass as a kind of residue. This is analogous to how mass parameters appear in 4d $\mathcal{N} = 2$ theories described by wrapping an M5-brane on a Riemann surface. There is a one-form ϕ on the Riemann surface, the Seiberg-Witten differential, which characterizes the normal motion of the brane. The embedding of the Riemann surface in Q is non-compact and has ends which appear asymptotically as $\mathbb{R} \times S^1$. The periods of ϕ over these asymptotic circles are then the mass parameters of the theory [13, 16, 22, 26, 35, 36]. We can equivalently describe this feature by compactifying the Riemann surface, and allowing ϕ to have residues. This means that ϕ is no longer closed as $d\phi \sim \delta(x)$. Similarly in our three-dimensional context, the embedding of M in Q can have ends which appear asymptotically as $\mathbb{R} \times \mathcal{C}$ for some Riemann surface \mathcal{C} . Then, the one-form λ can have periods over cycles in \mathcal{C} which encode the real masses. Compactifying M , simply means that λ is no longer closed as above.

- $\delta\omega$ restricts to M to a non-exact form. In that case we make use of a basis β_i of cohomologically

non-trivial two-forms and expand $\delta\omega$. Then, the modification of the equation defining λ is

$$d\lambda = \sum_i \zeta_i \beta_i. \quad (4.2.24)$$

Here, the real coefficients ζ_i appearing in the above expansion are naturally interpreted as FI parameters. Observe that, there is one such constant for each two-cycle, dual to β_i which are non-trivial not only inside M but also in Q . Later, when we describe M2-branes we will see that it is exactly these cycles which give rise to dual flavor symmetries.

The fact that these parameters are indeed the FI terms can be understood by noting that in the presence of non-vanishing ζ_i there is no solution to the above equation. As in our description of Chern-Simons levels this is interpreted as a destruction of the two-cycle dual to β_i . As a consequence of this we see that the parameter ζ_i has the correct physical effect of higgsing the associated $U(1)$ gauge group. Again, as in the case of real masses, one can make λ closed at the expense of deleting certain loci.

BPS M2-Branes and Instantons

The massless $U(1)$ gauge multiplets we have identified constitute an important subset of the information defining the Coulomb branch of the three-dimensional field theory determined by a five-brane on a three-manifold M . To complete the description, we now incorporate charged chiral multiplets and superpotentials. As we will see, all such objects arise from the possibility of M2-branes ending on the M5-brane and altering the physics.

First, let us discuss the inclusion of chiral multiplets in the field theory. We recall that because an M5 supports the two-form field B an M2 may end on an M5 in two spacetime dimensions while remaining consistent with charge conservation. Thus, to make a particle in three dimensions we may consider a two-brane whose worldvolume meets the five-brane along a timelike direction in $\mathbb{R}^{1,2}$ and a one-cycle Γ in M , as illustrated in Figure 4.1. Let D denote the two-cycle in the Calabi-

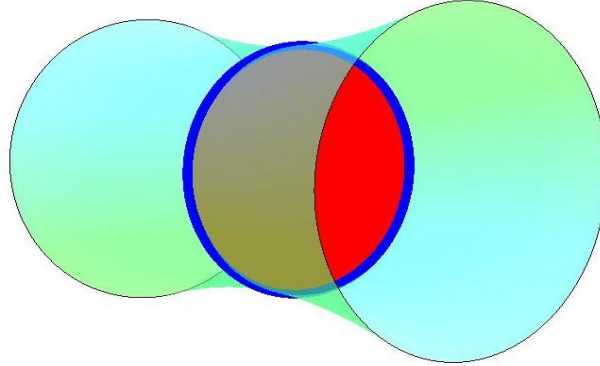


Figure 4.1: A BPS M2 brane. A three-manifold M , shown in green, sits inside the ambient Calabi-Yau Q , shown in white. The three-manifold supports a non-trivial one-cycle Γ shown in blue. A minimal M2 disc, shown in red, can end on this cycle and describes a BPS particle in $\mathbb{R}^{1,2}$.

Yau Q defined by the spatial directions of the M2. Then the mass m of the associated particle in $\mathbb{R}^{1,2}$ is determined by the volume of D . However, since we are interested in chiral multiplets we are interested in short representations of the supersymmetry algebra and hence in BPS M2 branes. Thus the cycle D must be minimal in its homology class and is therefore holomorphic. As a result the volume of D is fixed by the Kähler form

$$m = \int_D \omega. \quad (4.2.25)$$

However, locally near M we may use (4.2.12) to write $\omega = d\lambda$. Then since $\partial D = \Gamma$ we use Stokes' theorem in the above to obtain

$$m = \int_{\Gamma} \lambda = Z \quad (4.2.26)$$

Where Z is the central charge of the particle as measured by the periods of λ . This fact clarifies why it is the periods of the one-form λ which measures the central charges of charged particles. In the far infrared, all matter particles in $\mathbb{R}^{1,2}$ can be described by two-branes, and hence the gauge charge lattice of the theory is naturally identified with the set of one-cycles $H_1(M, \mathbb{Z})$. The one-form λ pairs with these charges and hence its periods can encode the central charges of the

field-theory.

Geometrically, a non-trivial chiral multiplet of charge Γ is described by a two-brane with topology of a disc. The existence of this disc means that while the one-cycle Γ may be non-trivial in M , when considered as a cycle in the ambient Calabi-Yau Q it is homologically trivial. It is exactly these cycles which become contractible in the ambient space that give rise to charged matter. Those one-cycles in M which remain non-contractible in Q describe gauge groups which have no associated charged chiral particles.

The above discussion of chiral multiplets sets the stage for other ways in which two-branes can influence the three-dimensional physics. Indeed, because an M2 can end on an M5 in two spacetime dimensions, its interpretation to a low-energy observer depends upon how many of the macroscopic dimensions the two-brane occupies. If a two-brane ends along a compact two-cycle in M then it occupies zero macroscopic dimensions and hence exists at a point in $\mathbb{R}^{1,2}$. Such an object is naturally interpreted as an instanton. One way to understand this is to examine the contributions to the action of this instanton. Since the two-brane carries B field charge, if it ends on the cycle β_j in M then its action will receive a contribution of the form

$$\exp\left(i\int_{\beta_j} B\right) = \exp\left(i\gamma_j\right). \quad (4.2.27)$$

Thus, the instanton action is weighted by a phase determined by the expectation value of the dual photon. This is familiar from the general structure three-dimensional field theories. It also serves to illustrate why it is the periods of B which measure the expectation values of the dual photons. The charges of possible instantons are naturally labeled by two-cycles, and it is with these objects that B can naturally pair.

Similar to the case of M2-brane particles, the presence of the instanton ending on the cycle β implies that while β is a non-trivial cycle in M it is homologically trivial in Q . This in turn implies that the associated monopole operator \mathcal{M} is present in the Lagrangian of the theory,

and hence the dual flavor symmetry is broken. Note that this further clarifies why the parameters appearing in (4.2.24) are indeed the FI parameters. Dual flavor symmetries appear only for those non-trivial two-cycles in M which remain non-trivial in Q . For those two-cycles in M which are contractible in Q there are M2-brane instantons, and the dual flavor symmetry is broken.

Finally, to construct a superpotential for the chiral fields we may consider an two-brane geometry which is a hybrid of the two elementary geometries described above. We fix background chiral particles X_i described geometrically by M2-brane discs ending on a collection of one-cycles $\Gamma_i \subset M$. Then, we find an M2 world-membrane that mediates an interaction between these objects. Topologically the worldvolume of this membrane is a three-manifold with boundary. This three-manifold lies entirely in the internal geometry Q and has boundary along the M2 discs describing the particles and along a two dimensional surface in M whose boundary is the union of the Γ_i . An example of this geometry is illustrated in Figure 4.2a. When the world-volume of this membrane is minimal, it describes a supersymmetric interaction and hence can give a contribution to the superpotential for the chiral fields. In practice, the most relevant case of this phenomenon occurs in the limit where the masses of the chiral particles become small and the superpotential is important. In that limit, the one-cycles labeling the charge of the chiral fields collapse to points and the two-brane we are describing is a handlebody whose boundary Riemann surface lies on a two-cycle in M and has a number of marked points corresponding to the insertion of massless chiral fields as shown in Figure 4.2b. That such instanton-like brane geometries make contributions to the superpotential is familiar from a variety of similar situations.

4.2.3 Many Five-Branes

When multiple five-branes wrap a three-manifold M , the resulting non-abelian dynamics gives rise to a strongly interacting field theory in three dimensions. Nevertheless in the IR on the Coulomb branch, we can still make use of the geometry described in the previous section to encode

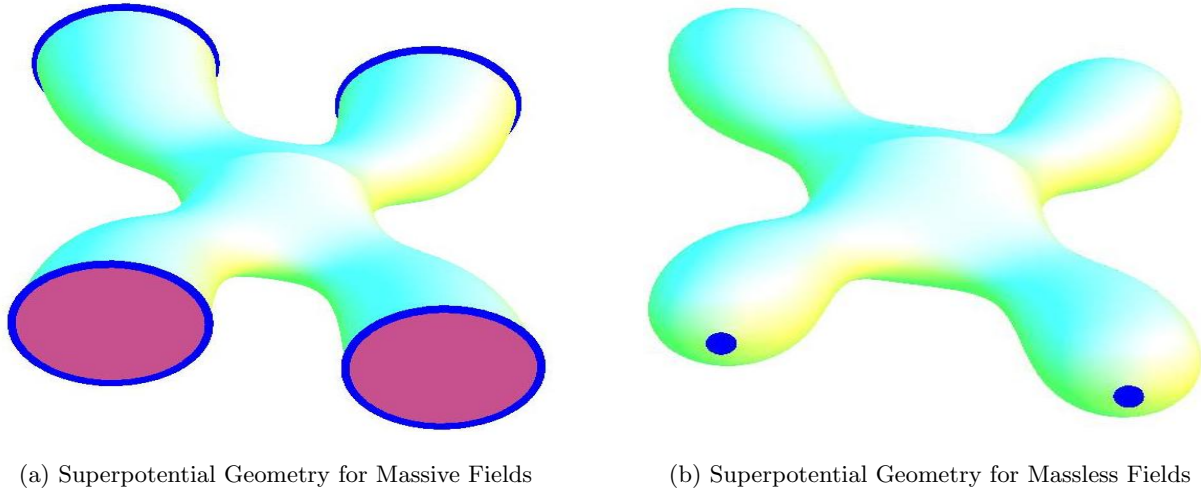


Figure 4.2: M2 brane contributions to the superpotential. In (a), we have four massive BPS states described by the pink M2 discs ending on blue one-cycles in the three-manifold. A three-dimensional closed M2 has boundary on these discs and along a two-dimensional locus in M and mediates a quartic interaction between the BPS particles. In (b), the BPS states become massless and the membrane geometry degenerates to a solid ball with four marked points, whose boundary lies entirely in M .

the physics. The key observation is again to recognize the effective scalar degrees of freedom. Just as for the case of a single five-brane, the transverse motion of the branes can be viewed as taking place in the cotangent bundle T^*M . If there are a total of n five-branes wrapping M then there are naively n independent one-forms λ_i on M which describe the motion of each individual five-brane. The reason that this assertion is naive is that it fails to account for the possibility that, after activating fields, the n distinct branes will recombine into a single connected object. In fact, a generic point on the Coulomb branch of the field theory is described by a geometry of this sort, and thus this possibility must be taken into account.

Fortunately, there is an elementary way to allow for brane recombination. We simply permit the possibility that the n objects λ_i are not individually globally well-defined but instead

permute amongst themselves as we traverse the manifold M [94]. Said differently, the one-forms λ_i are permitted to have one-dimensional branch loci and, under circling the branch locus, they are acted on by S_n , the permutation group on n letters. Such a structure naturally encodes brane recombination and gives rise to a three-manifold \widetilde{M} which is a n -sheeted cover of M . By definition, \widetilde{M} is exactly the three-manifold where the n branched one-forms λ_i glue together to yield a single, globally well-defined, harmonic one-form λ . We can encode this condition in equations by noting that λ defines completely the locus of the three-manifold cover \widetilde{M} inside the cotangent bundle of the base T^*M . Thus, knowledge of λ is equivalent to knowledge of the induced metric on the special Lagrangian \widetilde{M} and hence defines a hodge star operation $*_\lambda$. Then, the supersymmetric equations defining the IR geometry are

$$d\lambda = d*_\lambda \lambda = 0. \quad (4.2.28)$$

These are a set of non-linear relations on λ or equivalently the λ_i . They state that λ is harmonic in the induced metric which it determines.

Conceptually, the advantage of passing to the cover \widetilde{M} is that in the infrared all of the physics that is described by n five-branes wrapping M is completely encoded by the recombined brane \widetilde{M} . The virtue of this description is that the effective description is that of a *single* five-brane on \widetilde{M} . It is therefore naturally abelian and described by the geometry of the previous section. For example, it is the periods of the harmonic one-form λ on \widetilde{M} which determine the real central charges of the three dimensional field theory. Thus all of the non-abelian dynamics of multiple five-branes is encoded in the geometry of the covering manifold.

It is natural to interpret the existence of the cover \widetilde{M} , and its central role in the field theory, as a parallel to a similar structure which occurs in four-dimensional $\mathcal{N} = 2$ gauge theories which arise from placing n five-branes on a Riemann surface Σ . Just as above, the infrared dynamics of that theory are determined by brane recombination. The transverse motion of a single five brane is again identified with fluctuations in the contangent directions to the compactification manifold,

namely Σ . Thus, for each brane we expect a holomorphic one-form ϕ_i which parametrizes the position of the i -th brane. Brane recombination implies that the ϕ_i are not well-defined and instead have branch cuts where they mix. On passing to an n -fold cover $\tilde{\Sigma}$ these one-forms glue together to a single globally well-defined object ϕ . This cover $\tilde{\Sigma}$ is the Seiberg-Witten curve and ϕ is the Seiberg-Witten differential [35]. Their geometry and periods completely encode the low-energy action [9, 10]. The manifold \tilde{M} , whose abstract existence we have eluded to in this section plays a similar role in the three-dimensional physics, and in later sections where we study explicit examples, our primary task will be to determine \tilde{M} .

4.2.4 Two Five-Branes

For most of our explicit examples in later sections, we will be interested in the specialization to the case where the number of five-branes, before recombination, is two. Then, the IR five-brane geometry is that of a branched double cover $\tilde{M} \rightarrow M$. In this section we discuss in more detail the resulting three-dimensional topology and its relation to the physics.³ In practice our primary examples will be to the case where M is a three-sphere and for the remainder of this section we make that restriction.⁴ Our specific goal will be to determine the homology group $H_1(\tilde{M}, \mathbb{Z})$. As we have argued in equation (4.2.19), complete knowledge of this homology is equivalent to a description of the gauge boson sector of the field theory on $\mathbb{R}^{1,2}$, with propagating fields captured by the Betti number $b_1(\tilde{M})$ and non-trivial levels \hat{k}_{ij} encoded in the torsion classes of $H_1(\tilde{M}, \mathbb{Z})$. Further, in section 4.2.4, we also illustrate how the M2 brane geometries discussed in previous sections can be visualized more directly in the case of a double cover.

³ Much of this geometry is classical. For an introduction see [95].

⁴ In terms of the topology of the cover \tilde{M} this is no restriction. Indeed, every orientable compact three-manifold can be presented as a double branched cover of S^3 [95].

Seifert Surfaces

We begin with the elementary observation that S^3 has trivial topology. From this it follows that all the resulting topology of the cover is encoded in its branch structure over S^3 . Since \widetilde{M} is smooth, the branch locus is required to be a smooth embedded closed submanifold of M dimension one. Topologically, the branch locus is therefore a union of circles. However, the circles may be embedded in S^3 in a complicated fashion and hence form a non-trivial knot \mathcal{K} .⁵ The topology of the cover \widetilde{M} is completely fixed by \mathcal{K} . To construct the cover we first proceed by drawing a branch sheet. This is a smooth two-dimensional surface F whose boundary is the given knot \mathcal{K} . A classical theorem of Seifert, Frankl, and Pontrajgin ensures that such a surface always exists, and that further one may assume F to be orientable. When this is so, F is referred to as a *Seifert surface* for the knot \mathcal{K} . Some examples are illustrated in Figure 4.3.

Once a Seifert surface has been specified, the double cover \widetilde{M} can be constructed explicitly. We take two copies of S^3 and cut them both along F . In each three-sphere this creates a boundary which is topologically two copies of the Seifert surface, F_+ and F_- , intersecting along \mathcal{K} . Then we glue F_+ in one S^3 to F_- in the other S^3 and vice versa to create \widetilde{M} which is a two-to-one cover of S^3 except over the knot \mathcal{K} . Of course, as is the case with branched covers of Riemann surfaces, there are in general many choices of branch sheets, and so given a knot \mathcal{K} its Seifert surface is not unique. However, the topology of the cover \widetilde{M} is independent of this choice and thus any convenient Seifert surface can be used for the branch sheet. In practice this construction is useful since many properties of the cover can be deduced directly from any given Seifert surface.

One useful quantity that we may extract from the Seifert surface F is the homology $H_1(\widetilde{M})$. Indeed, since S^3 has no non-trivial one-cycles, all cycles in \widetilde{M} can be localized to a neighborhood of \mathcal{K} and must involve the knot if they are to be non-trivial. This is quite similar

⁵ In this paper by the term *knots* we will refer to both knots and links, and whenever we really mean “knot” we shall emphasize it.

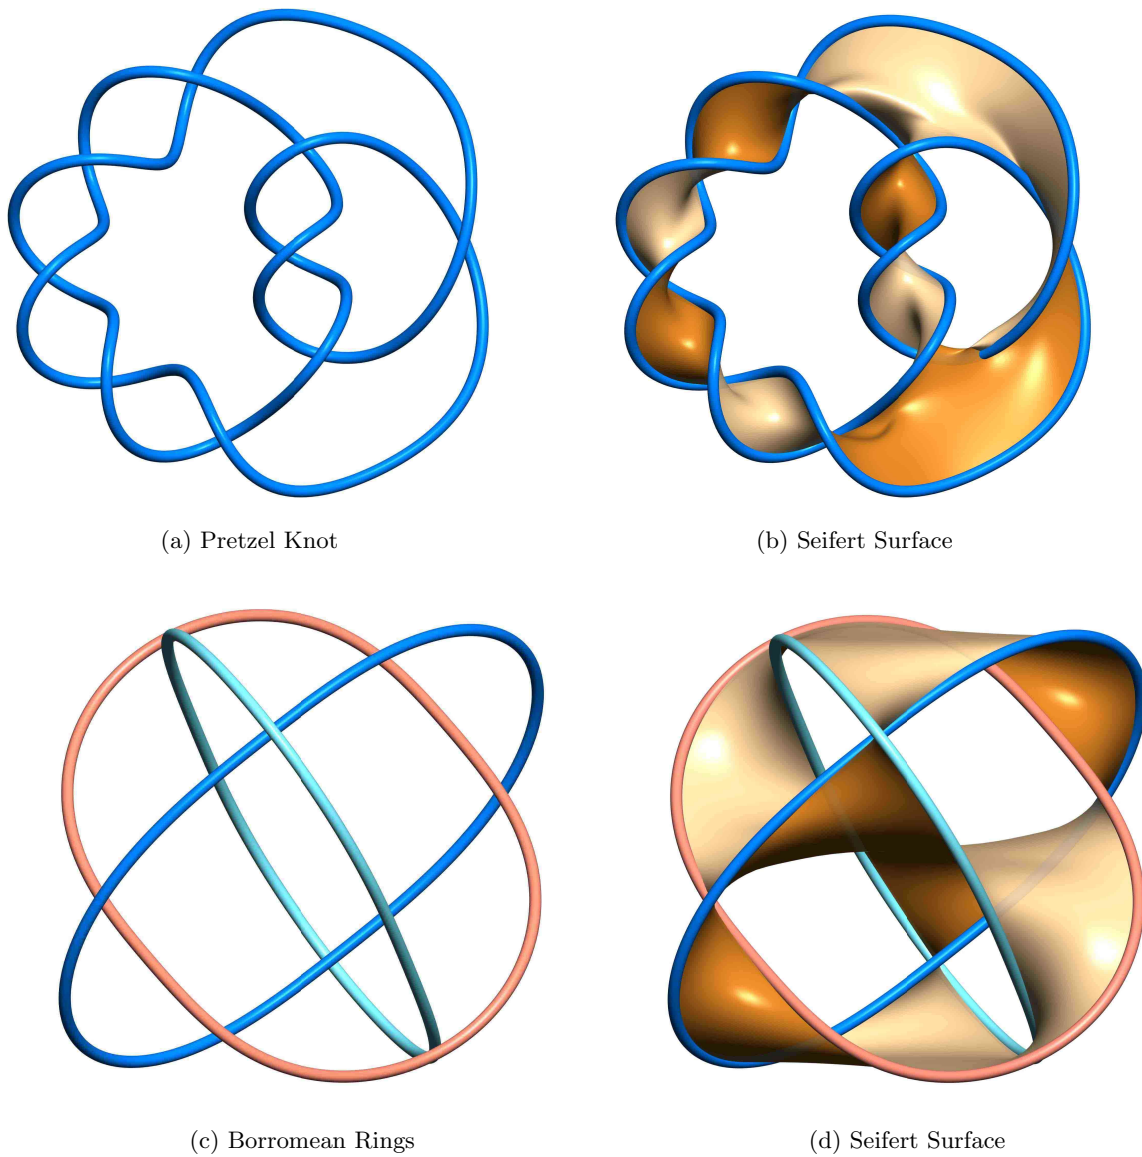


Figure 4.3: Knots and associated Seifert surfaces. In (a) and (c) we see non-trivial knots. In (b) and (d), we see Seifert surfaces whose boundaries are the given knots. The knots are an example of possible branch loci for a double cover of S^3 . The Seifert surfaces are then the branch sheets defining the cover geometry.

to the case of the homology of a branched cover of the two-sphere. There, the branch locus is a number of points, and the branch lines are segments I connecting pairs of these points. The homology of the double cover is then generated by the classes in the complement of the branch lines, $H_1(S^2 - \cup I)$. Our situation is exactly parallel, save for the fact that it takes place in one higher dimension. The homology of \widetilde{M} , a branched double cover of S^3 , is generated by the cycles in the complement of the branch sheet $H_1(S^3 - F)$.

We can make a further simplification by observing that cycles in $S^3 - F$ and cycles in F are naturally dual. Indeed, given $\alpha \in H_1(S^3 - F)$ and $\beta \in H_1(F)$ we represent α and β by simple, oriented, closed curves and compute the linking number $[\alpha, \beta] \in \mathbb{Z}$. This determines a pairing

$$[\cdot, \cdot] : H_1(S^3 - F) \times H_1(F) \longrightarrow \mathbb{Z}. \quad (4.2.29)$$

And in fact the pairing is an isomorphism. This implies that a natural spanning set of one-cycles generating $H_1(\widetilde{M})$ is given by the homology classes on the Seifert surface itself, $H_1(F)$.

Now, although the homology $H_1(F)$ generates the homology of the cover, typically when considered in \widetilde{M} , many of these cycles are in fact homologically trivial. Thus, $H_1(F)$ is generally an overcomplete set of cycles, and our task is to determine which classes on the Seifert surface become trivial in \widetilde{M} . To do so, we introduce the concept of a Seifert matrix. This is a $b_1(F) \times b_1(F)$ integral matrix defined as follows.

- Definition: Let Γ_i be a one-cycle in F . Since F is oriented we can define Γ_i^+ as a small pushoff of Γ_i out of F in the positive direction. Then the Seifert matrix A is the matrix of linking numbers A_{ij} between the one-cycle Γ_i and the pushoff Γ_j^+ .

Armed with this matrix we can then say that the first homology of the cover \widetilde{M} is generated by $H_1(F)$ modulo the relations defined by the symmetrized Seifert matrix

$$H_1(\widetilde{M}) \cong H_1(F) / \text{Image}(A + A^T). \quad (4.2.30)$$

Thus for example, the rank of the kernel of the map $A + A^T$ computes the first Betti number of \widetilde{M} . Meanwhile if the homology of \widetilde{M} is a finite abelian group, then the order of this group is computed by $|\det(A + A^T)|$.⁶

The fact (4.2.30) should be directly interpreted in terms of our general discussion of $U(1)$ gauge fields and CS terms in (4.2.19). In general, the gauge multiplet sector of the theory is described by a collection of abelian gauge fields and a matrix \hat{k}_{ij} of levels. This data translates into a description of the homology $H_1(\widetilde{M})$. The $U(1)$ gauge fields are a generating set of classes in the homology, and the matrix \hat{k}_{ij} describes the relations among these generators. In (4.2.30) we see exactly this description and hence we propose that:

- $U(1)$ gauge fields are given by generators Γ_i of $H_1(F)$.
- CS levels \hat{k}_{ij} are given by the symmetrized Seifert matrix $A + A^T$.

In this description, the fact that the Seifert surface is non-unique translates to a statement about equivalence of various CS theories. Any Seifert surface may be used to describe the physics, and distinct surfaces give distinct sets of $U(1)$'s and level matrices \hat{k}_{ij} all of which determine the same theory.

To derive the above proposal, we first phrase the computation of linking numbers in a more familiar language of differential forms as follows. Each homology class Γ_i in $H_1(F)$ can be represented by a cycle that is topologically an unknot embedded in S^3 . Thus, the pushoff Γ_i^+ bounds an embedded disc $D_i \subset S^3$. The symmetrized Seifert matrix of linking numbers is then given by computing the intersection number of Γ_i with D_j

$$(A + A^T)_{ij} = \Gamma_i \cap D_j + \Gamma_j \cap D_i. \quad (4.2.31)$$

⁶ Incidentally, the Seifert matrix can also be used to define the Alexander polynomial of the knot by the definition $\mathcal{A}_K(t) = \det(A - tA^T)$. Then the determinant above, which computes the order of $H_1(\widetilde{M})$ when finite, is the Alexander polynomial evaluated at $t = -1$.

However, the intersection number on the right of the above can alternatively be computed in terms of the Poincaré dual form to the disc D_i [96]. Specifically, since the disc D_i are bounded by the cycles Γ_i^+ the above is

$$(A + A^T)_{ij} = \int_M \alpha_i \wedge d\alpha_j. \quad (4.2.32)$$

Equation (4.2.32) gives a direct way of seeing that the symmetrized Seifert matrix computes the levels \hat{k}_{ij} . For each of the one-forms α_i introduced above we may consider an associated $U(1)$ gauge field A_i by decomposing the two-form field B propagating on the fivebrane \widetilde{M} . This gauge field may be massless, or massive depending on the resulting equation of motion. To examine this issue, we consider the self-dual 3-form field strength $T = dB$ written as

$$T = \alpha_i \wedge F_i + *\alpha_i \wedge *F_i. \quad (4.2.33)$$

Then, the equation of motion $dT = 0$ implies in particular

$$d\alpha_i \wedge F_i + *\alpha_i \wedge d*F_i = 0. \quad (4.2.34)$$

Wedge the above equation with α_j and integrate over \widetilde{M} to find

$$\left(\int_{\widetilde{M}} \alpha_j \wedge *\alpha_i \right) d*F_i + (A + A^T)_{ji} F_i = 0. \quad (4.2.35)$$

The normalization matrix $\int_{\widetilde{M}} \alpha_j \wedge *\alpha_i$ determines the metric on the space of $U(1)$ gauge fields, and with this identification, (4.2.35) is exactly the equation for a collection of $U(1)$ vectors with a level matrix \hat{k}_{ij} given by the symmetrized Seifert matrix $A + A^T$.

As a sample application of these ideas, we consider the case of a cover branched over the unknot. Of course, because the unknot bounds a disc we may choose this as the Seifert surface. Then glueing together two S^3 's with a branching sheet given by a disc is simply taking the connected sum of the two S^3 's. This means that the cover \widetilde{M} is again an S^3 and hence has trivial homology. However for a more interesting choice we can take as the Seifert surface F a torus minus a disc

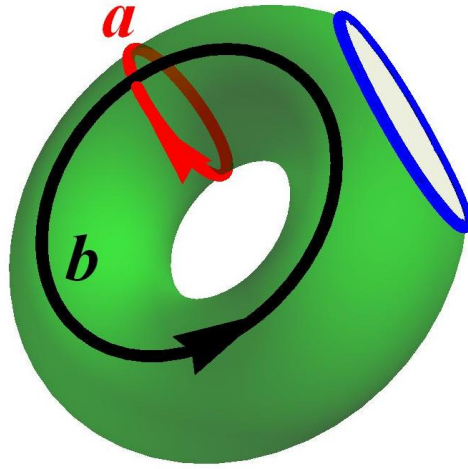


Figure 4.4: A Seifert surface F for the unknot. The green torus minus a disc bounds the unknot shown in blue. The red and black cycles, a and b , are a basis for the homology of F . The pushoff a^+ is linked with b .

as shown in Figure 4.4. A basis of $H_1(F)$ is then the two one-cycles a and b shown in the Figure. Taking the outward direction of the torus to be the positive orientation we then conclude that there is a non-vanishing linking number $+1$ between b and the pushoff a^+ . Hence in this case we find that the symmetrized Seifert matrix is

$$A + A^T = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}. \quad (4.2.36)$$

From this we deduce that both a and b are in the image $A + A^T$ and hence trivial in the homology of \widetilde{M} . Thus we recover the correct result that $H_1(\widetilde{M})$ is vanishing.

Physically, the example given above describes a known duality [97]. A $U(1) \times U(1)$ gauge theory together with level matrix

$$A + A^T = \hat{k}_{ij} = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} \quad (4.2.37)$$

is equivalent to a trivial theory of no gauge group whatsoever. In fact, this example suffices to

understand the general result that the gauge multiplet theory is independent of the choice of F . Indeed, given a fixed knot \mathcal{K} any two topologically distinct Seifert surfaces F and F' with the genus of F' larger than F , differ by excising some number n of discs from F and attaching n handles like those shown in Figure 4.4 to obtain F' . At the level of the Seifert matrices this has the effect

$$(A + A^T)|_{F'} = (A + A^T)|_F \bigoplus_{i=1}^n \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}. \quad (4.2.38)$$

In other words, adding an irrelevant handle simply adds a trivial theory in the form of a $U(1) \times U(1)$ with off-diagonal level matrix (4.2.36) and hence does not modify the physics.

Checkerboards and Tait Graphs

In the forthcoming applications of this paper, it will be important for us to have a more explicit method for determining a Seifert surface for a given knot and computing Chern-Simons levels. One way to produce such a surface is to use a so-called *checkerboard coloring* of the knot. This is an assignment of black versus white to each region enclosed by the planar projection of the knot. It has the property that regions which share an edge have opposite colors. Given any knot, there is no obstruction to constructing a checkerboard coloring. Indeed, we simply consider the local structure of the knot near a given crossing. If we forget the data of which component passes over versus under, the crossing appears as the intersection of two lines, and separates the plane into four regions as shown in Figure 4.5a. Then, we choose a pair of non-adjacent regions and shade them as shown in Figure 4.5b. We continue this shading procedure consistently to all the remaining crossing in the knot. At the conclusion we have constructed a checkerboard coloring.

The relation of checkerboard colorings to Seifert surfaces is simply that the shaded regions of the coloring define a two-dimensional surface F whose boundary is the knot \mathcal{K} . The interior of each shaded region is a disc and at the crossings, these discs are glued together by twisted bands.

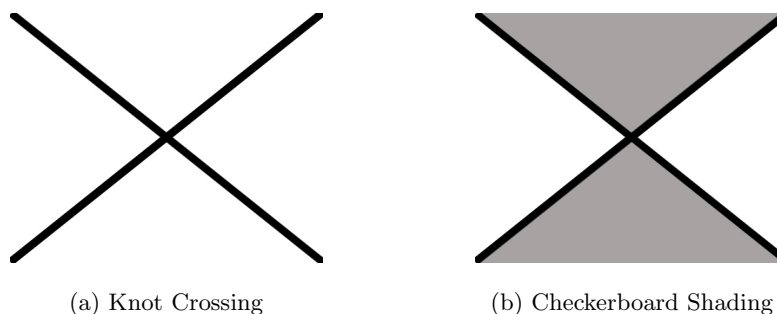


Figure 4.5: Local definition of checkerboard coloring. In (a) a planar projection of a crossing in a knot diagram. In (b) a checkerboard coloring at the crossing.

Thus the shaded regions of a checkerboard diagram determine a surface F whose boundary is the knot \mathcal{K} .⁷

Notice that as a consequence of the construction, there is a natural notion of duality among checkerboard colorings of the knot. Given such a coloring, we may exchange the black and white regions to produce a new coloring. A specific example of this is given for the trefoil knot in Figure 4.6. In later sections we will see that this basic black-white duality of checkerboard colorings has an interesting physical interpretation in terms of duality of 3d field theories.

However, for our present purposes our main interest in checkerboard colorings is simply that they provide a convenient way of determining the homology of the cover $\widetilde{\mathcal{M}}$, and therefore a method for determining a set of $U(1)$ gauge fields and a level matrix \hat{k}_{ij} .

Let us first fix the basis of cycles. These are manifest in the checkerboard coloring. Each white region of the diagram is, by construction, a hole in the surface defined by the shaded regions of the coloring. Therefore there is a non-trivial cycle defined by simply moving the boundary of the given white region slightly into the shaded region. Thus, if R_1, \dots, R_{n+1} denote the white regions

⁷ In many situations this surface is orientable and hence meets the requirements to be called a Seifert surface for \mathcal{K} . Sometimes, however the surface is non-orientable and is technically therefore not a Seifert surface. This causes no problem from the point of view of using such a surface as a branch sheet to construct a cover. Further as we describe below, all of the pertinent results of the previous section, go through unmodified.

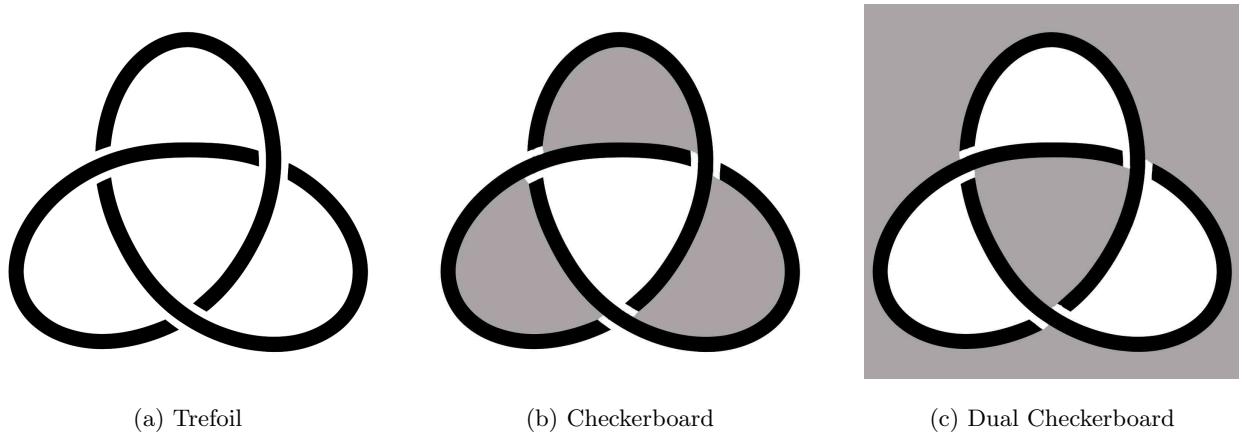


Figure 4.6: Checkerboard colorings for the trefoil knot. In (a) the trefoil knot. In (b) and (c) its two dual colorings. The shaded regions can be interpreted as surfaces with boundary the trefoil.

of the checkerboard there are associated cycles Γ_i encircling R_i . Examples are illustrated in Figure 4.7. One of these cycles, say Γ_{n+1} can be generated in homology of the surface F by the remaining n . The remaining cycles $\Gamma_1, \dots, \Gamma_n$ are an explicit basis for the homology of the surface. In physical language these are the defining generators for a $U(1)$ gauge theory.

Next we extract the CS levels. As in our general discussion, these levels are determined by a linking number computation. However in the case of the checkerboard coloring there is a simple more algorithmic way of determining the levels. First we associate to each crossing c in the diagram a sign function $\zeta(c) = \pm 1$ determined by whether the overstrand or understrand is to the left or the right of the surface F as illustrated in Figure 4.8. The CS matrix is then most easily determined by keeping track of all the $n + 1$ cycles associated to all the white regions, as opposed to just the n in a spanning set. Specifically, we construct an $(n + 1) \times (n + 1)$ matrix whose rows and columns index the white regions. Then:

- The off-diagonal elements \hat{k}_{ij} for $i \neq j$ are determined by summing over all crossings where

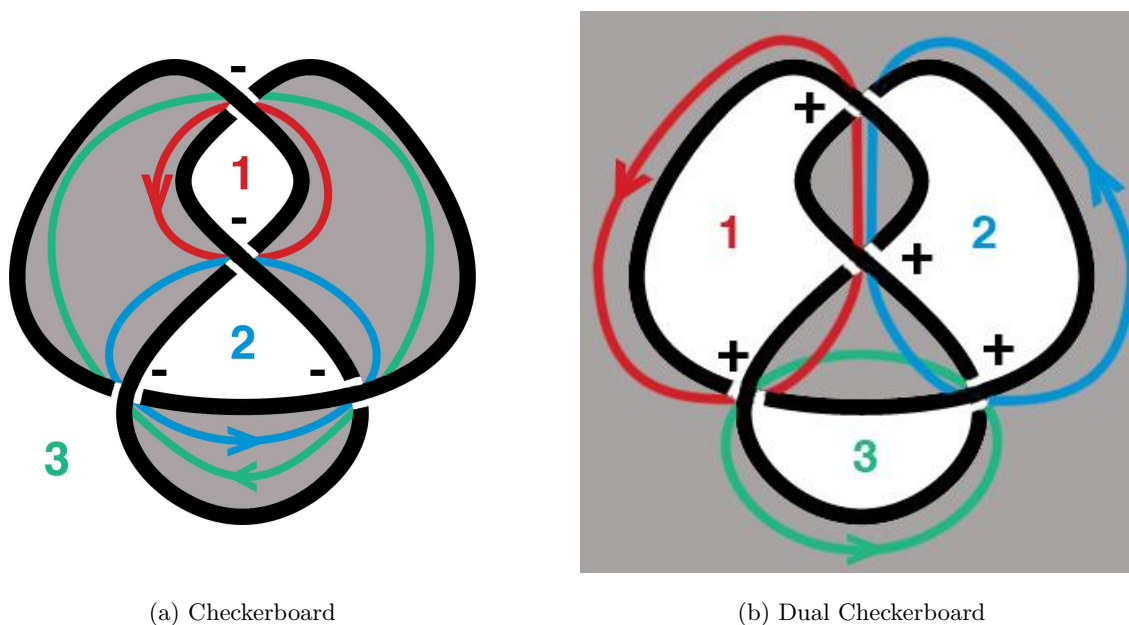


Figure 4.7: Dual checkerboard colorings for the figure eight knot. The gray regions denote the Seifert surface of the black knot. The colored cycles denote the spanning set of the homology of the surface given by the enclosed white regions of the checkerboard. The signs at each crossing indicate the contribution to the Chern-Simons levels.

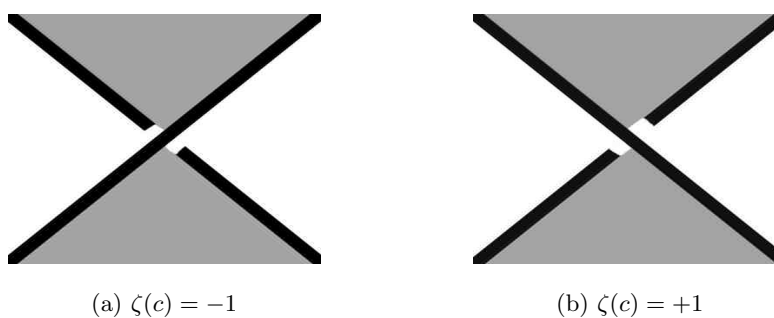


Figure 4.8: Positive and negative crossings in a knot.

regions i and j meet taken with sign.

$$\hat{k}_{ij} = \sum_{i,j \text{ crossings}} \zeta(c) \tag{4.2.39}$$

- The diagonal elements \hat{k}_{ii} are determined by the condition that the sum of all entries in any row vanishes

$$\hat{k}_{ii} = - \sum_{j \neq i} \hat{k}_{ji}. \quad (4.2.40)$$

- At the conclusion of the computation, eliminate the $(n + 1)$ -st row and column to obtain an $n \times n$ matrix of levels \hat{k}_{ij} involving the $U(1)$'s related to the white regions R_1, \dots, R_n .

In this context of checkerboard colorings, the matrix \hat{k}_{ij} is known as the *Goerizt form*, and the equations above provide us with an algorithmic recipe for determining CS levels.

For concreteness, let us now apply this construction to the case of the knot shown in Figure 4.7. As illustrated, there are two dual checkerboards each of which has three white regions. Then the the 2×2 level matrix for the cycles defined by regions 1 and 2 are given respectively by

$$\hat{k}_{ij} = \begin{pmatrix} 2 & -1 \\ -1 & 3 \end{pmatrix}, \quad \text{and} \quad \hat{k}_{ij} = \begin{pmatrix} -3 & 2 \\ 2 & -3 \end{pmatrix}. \quad (4.2.41)$$

As a consistency check on this computation, note that the two level matrices determined by the checkerboard and its dual have identical determinants. Indeed as we have previously described, when finite, $|\det(\hat{k}_{ij})|$ computes the order of the first homology group of the cover \widetilde{M} , and hence is invariant to the choice of surface.

The combinatorics of checkerboard colorings can also be conveniently encoded in a so-called *Tait graph*. Given a coloring we extract the graph as follows:

- For each white region R_i draw a node of the diagram.
- For each crossing connecting white regions i and j connect the corresponding nodes with a link.
- Attach a sign ± 1 to each link by evaluating ζ of the corresponding crossing.

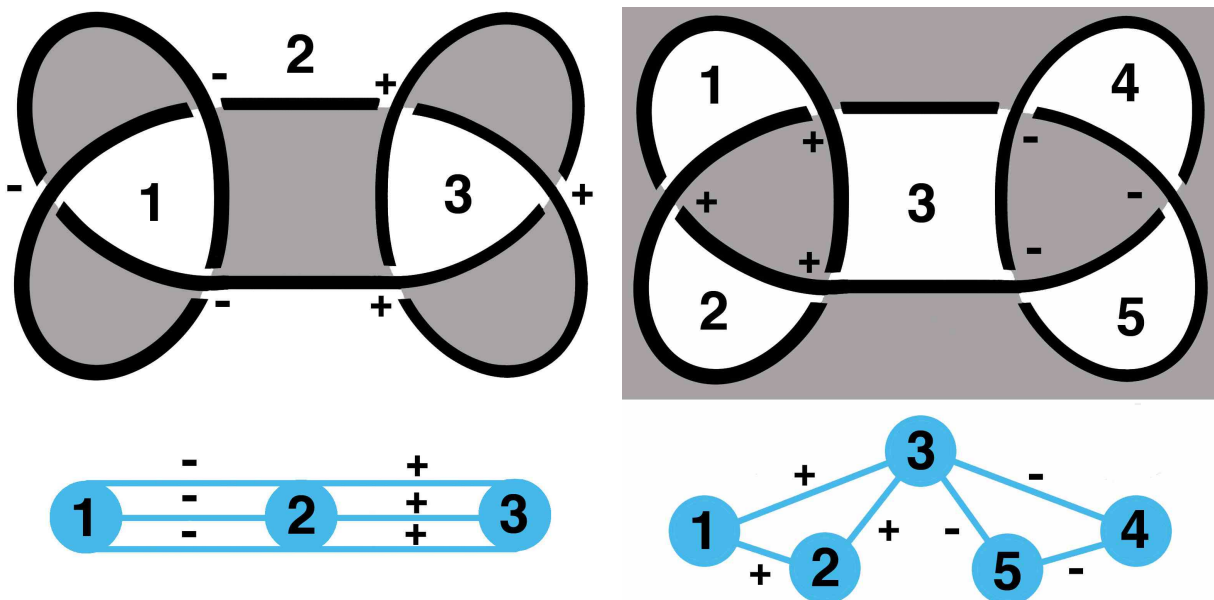


Figure 4.9: Checkerboards and associated Tait graphs for the square knot. The two checkerboards are related by a black white exchange. And correspondingly, the two Tait graphs are dual.

As a sample application of this notion, we consider the two checkerboards of the knot illustrated in Figure 4.9.

Notice from this example that the black-white duality between checkerboard colorings of the knot maps to duality of the corresponding Tait graphs. Specifically, given a graph G , to construct the dual \hat{G} we simply:

- Replace each cell of G with a dual vertex of \hat{G} .
- Replace each link in G with a dual link in \hat{G} .
- Change the sign of each link relative to its dual.

The Tait graph encodes completely the gauge content of the theory. Each node describes a white region, and hence corresponds to a cycle in the surface F . Up to removing one such cycle or equivalently one node in the graph, these are exactly the $U(1)$ gauge fields. Similarly, the CS

level matrix \hat{k}_{ij} is determined by summing over the links connecting nodes i and j weighted as in (4.2.39) by the sign of the link. This structure is completely universal and depends only on the topology of the double cover \widetilde{M} . In subsequent sections however, we will see examples where the Tait diagram encodes more than merely the gauge group and levels. Indeed, in Section 4.5, after determining the matter content of the theory, we will see that the Tait diagram plays the role of the quiver diagram for the gauge theory in $\mathbb{R}^{1,2}$.

BPS M2-Branes, Instantons and Double Covers

In addition to the homology of the cover, which encodes the gauge sector of the theory, there are other physical quantities of interest that can be read directly from the knot. Of particular importance to us in later sections will be the possible BPS M2 brane geometries. Let us begin with the case of an M2 brane disc describing a particle in three dimensions. The boundary of this disc is a circle Γ inside the double cover \widetilde{M} . Now project Γ to the base S^3 . If the projection is a circle then, since S^3 has vanishing homology, the cycle is contractible and the particle carries no gauge charges. We will therefore ignore this case. The remaining possibility is that the projection is an interval connecting two pieces of the branching knot \mathcal{K} as shown in Figure 4.10a. Such a particle can in principle carry gauge charges depending on whether or not the cycle Γ is non-trivial in \widetilde{M} .

Geometrically, Γ is partitioned into two pieces, one on each sheet of the cover \widetilde{M}_{\pm} , both of which project to the given segment in S^3 . The two segments are then glued together at their intersection with \mathcal{K} which occupies both sheets. This description also makes elementary why the mass of such a particle is determined by λ as

$$m = \int_{\Gamma} |\lambda|. \quad (4.2.42)$$

In this case, $|\lambda|$ is the physical height separating the two sheets \widetilde{M}_{\pm} of the cover, and the above integral, by definition, computes the area of the disc illustrated in Figure 4.10b.

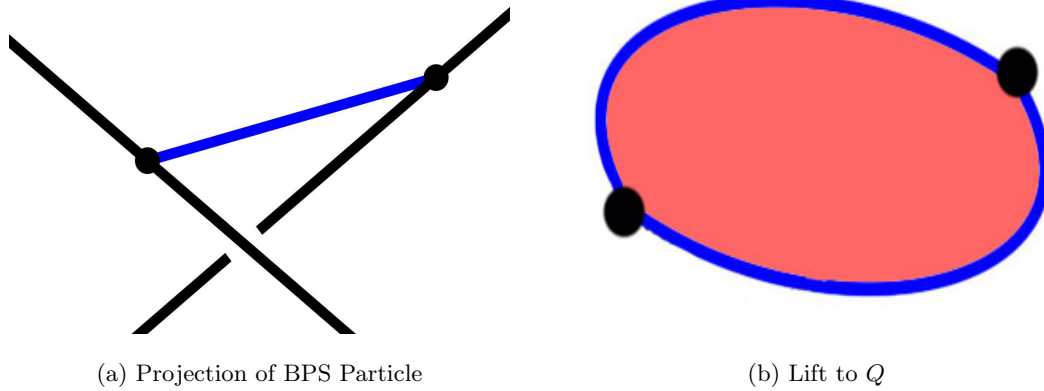


Figure 4.10: Projections of BPS M2-brane particles to the base. A portion of the branching knot \mathcal{K} is shown in black. In (a), the boundary of an M2 disc appears as a blue segment connecting two pieces of the knot. In (b), the blue segment is doubled in \widetilde{M} to make a closed cycle S^1 . The red disc, which lies in $Q - \widetilde{M}$, shows how Γ is filled in to make the full M2-brane geometry.

M2 brane contributions to the superpotential can also be seen from the knot diagram. Suppose first that we consider the case of interactions among *massless* particles. According to the geometry described in the previous paragraph, this means that the segment projections shown in blue in Figure 4.10a have all collapsed to points, and as a result the knot \mathcal{K} has developed self-intersections. Let us further assume that these self-crossings form the vertices of a closed polygon whose boundary lies entirely in \mathcal{K} . Then, there is an M2 brane instanton in Q whose boundary projects to the polygon and which can give rise to interactions among the massless fields at the vertices as illustrated in Figure 4.11.

To be explicit, in the Calabi-Yau Q , the polygon lifts to a three-ball B^3 whose boundary S^2 is P_{\pm} , the interior of the polygons on each of the sheets \widetilde{M}_{\pm} of the cover. These two polygons have been glued together along their common boundary in the branching knot \mathcal{K} whose boundary

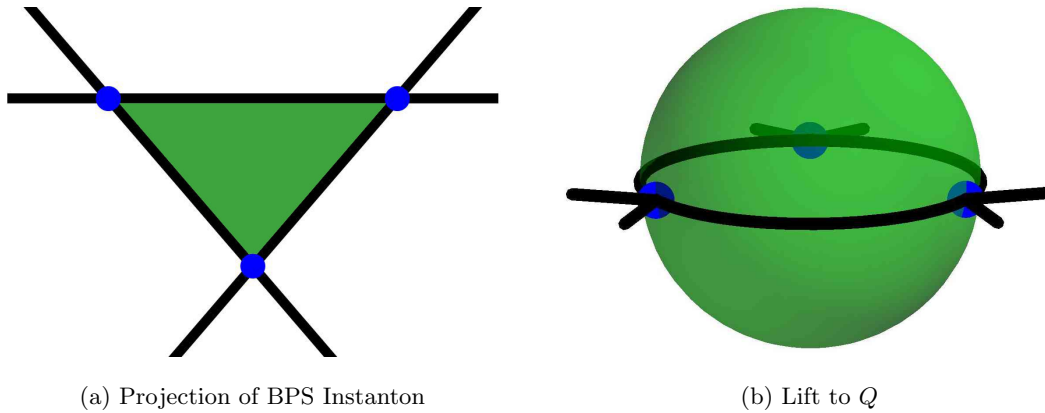


Figure 4.11: Projections of BPS M2-brane instanton to the base. A portion of the branching knot \mathcal{K} is shown in black. The knot has self-intersections supporting massless particles shown in blue. In (a), the interior of the polygon, shown in green, is the projection to \widetilde{M} , of the boundary of an M2 instanton. In (b), we see the lift of the M2 instanton to the Calabi-Yau. Its boundary is doubled to an S^2 presented as two hemispheres glued along the knot. In the interior, this S^2 is filled in to make a three-ball.

is

$$\partial B^3 = S^2 = P_+ \bigcup_{\mathcal{K}} P_- \subset \widetilde{M}. \quad (4.2.43)$$

And a wrapped M2 brane over B^3 leads to a superpotential term involving the product of the massless chiral fields, one for each of the vertices of the polygon. Deforming the geometry and making the chiral fields massive, resolves the self-crossings of the knot diagram. This can be described by an M2 brane instanton, whose boundary will also include a disc ‘plug’ for each massive chiral field as in the general description of Figure 4.2.

4.3 R-flow, Domain Walls and a 4d-3d Link

The geometry described in the previous section gives an abstract prescription for extracting the effective three-dimensional $\mathcal{N} = 2$ system that arises when multiple five-branes wrap a three-manifold. We first determine the IR covering geometry \widetilde{M} , then we compute the spectrum of BPS M2-brane particles which give rise to chiral multiplets, and finally determine their interactions from the various M2-brane contributions to the superpotential. However, in practice it may be difficult to carry out this procedure. The first difficulty is that we have no general method for determining \widetilde{M} , or equivalently the branching knot $\mathcal{K} \subset M$. Moreover, even if the topology of \widetilde{M} were deduced, this only suffices to describe the gauge groups and flavor symmetries but not the BPS states. In the infrared, it is the BPS states which describe the charged chiral multiplets of the theory, and thus extracting the spectrum of these particles is a crucial step in determining the behavior of the quantum field theory.

In principle, the BPS states are completely encoded by the one-form λ on the cover. Indeed, given that λ defines the local central charge density, it follows that the boundary of a BPS M2 brane is an integral curve of the flow defined by λ . In equations, if γ denotes this boundary one-cycle and s is a local coordinate on γ , then the BPS condition requires

$$\lambda|_{\gamma} = ds. \tag{4.3.1}$$

For example, a chiral field arises as a solution to the above equation whose endpoints are on the branch knot as described in the previous section. This is the three-dimensional analog of the flow equation defining BPS states in 4d $\mathcal{N} = 2$ theories [11]. However, without explicit knowledge of λ , there is no method for determining the BPS particles in the theory and hence no way of directly fixing the IR behavior of the model.

In this section, we introduce our main technique for circumventing these difficulties. We consider the special case where the M5-branes are related to a flow of a parent 4d theory, and

use the knowledge of the spectrum of the 4d BPS states and Seiberg-Witten geometry to find the answer for the resulting 3d theory. Let us recall that in the 4d case, instead of (4.3.1), the BPS spectrum is determined by integral curves of the Seiberg-Witten differential

$$\phi|_{\gamma} = e^{i\theta} ds, \quad (4.3.2)$$

where there is a solution only if θ is chosen to equal the phase of the central charge Z of the BPS state

$$Z = |Z|e^{i\theta}. \quad (4.3.3)$$

Now let us consider the 3d case. Our ansatz, of viewing the 3d theory as a one-parameter variation of the 4d theory, implies that we are studying a domain wall, where roughly

$$\lambda = e^{i\theta(t)} \phi + c.c. + \dots \quad (4.3.4)$$

In this way, we will have solutions to the 3d BPS equation (4.3.1) at specific ‘times’ t_i during the one-parameter variation, such that $\theta(t_i) = \theta(\text{BPS})$ for some BPS particle. This is exactly the characteristic feature of the time evolution defined by R-twisting [3, 4]. However, there is one important difference in our context: unlike the case in R-twisting where the central charges rotate in phase as we evolve in time, in order to preserve a standard 3d supersymmetry, we need to make the central charges *flow along parallel lines*. This can be achieved by the suitable choice of the ... terms in equation (4.3.4), as will be discussed later in this section. The result, which we shall call the ‘R-flow,’ is a one-parameter variation of the parent 4d theory, characterized in terms of the variation of the 4d central charges Z_i by two simple features:

- The real part of the Z_i is constant along the flow.
- Along the flow, the Z_i retain their phase order.

These features are shown in Figure 4.12. As we will see, these two properties mean that the time

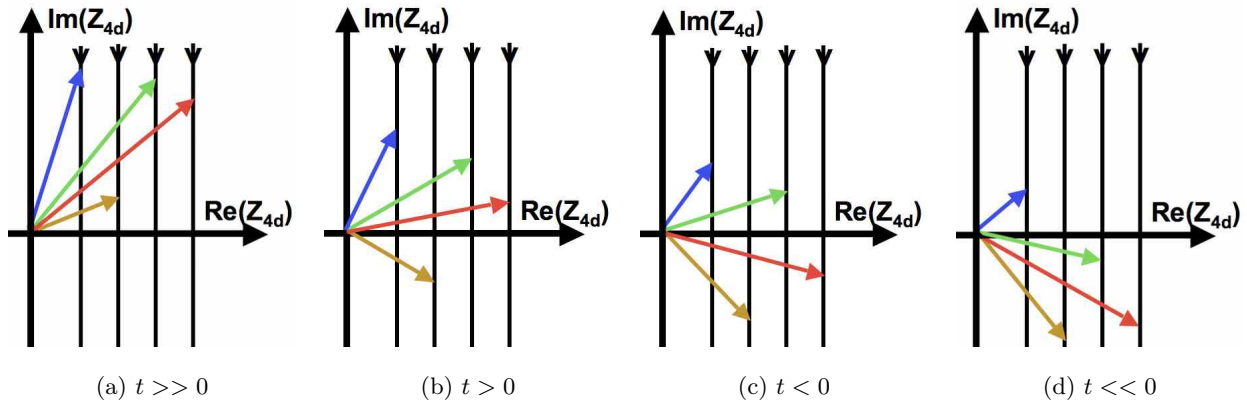


Figure 4.12: Evolution of the central charges along an R-flow. The colored rays denote the central charges of the 4d theory. As we move along the flow, these rays flow along parallel lines and maintain their phase order.

evolution in the R-flow respects the BPS spectrum of the 4d theory, and ultimately implies that the 3d BPS spectrum of chiral fields living on the domain wall is in one-to-one correspondence with the BPS spectrum of the ambient 4d theory.

Finally, a crucial aspect of the domain wall construction is the observation that, in general, such a domain wall field theory couples non-trivially to the bulk four-dimensional physics. However, our interest is in constructing honest three-dimensional theories which have an independent existence. Thus, in addition to the construction of the domain walls, we must also take a decoupling limit in which the four-dimensional theory decouples and the three-dimensional theory on the wall remains. This decoupling limit amounts to a specification of boundary conditions for the R-flow, where as $|t| \rightarrow \infty$ we also have $|Z| \rightarrow \infty$. Then, the full trajectories of the 4d central charges during an R-flow are infinite parallel lines.

4.3.1 Domain Wall Geometry

Consider any number five-branes which wrap a Riemann surface Σ and give rise to a four-dimensional $\mathcal{N} = 2$ gauge theory. The local geometry is then

$$\mathbb{R}^{1,4} \times T^*\Sigma \times \mathbb{C}. \quad (4.3.5)$$

To engineer a macroscopically four-dimensional system we choose linear subspaces $\mathbb{R}^{1,2} \subset \mathbb{R}^{1,4}$ and $\mathbb{R} \subset \mathbb{C}$ and place the five-branes on

$$\mathbb{R}^{1,2} \times \Sigma \times \mathbb{R}. \quad (4.3.6)$$

Such a geometry supports a natural class of defects which describe domain walls. The linear subspace $\mathbb{R} \subset \mathbb{C}$ is replaced by a non-trivial path $\gamma(t)$. Asymptotically for $t \rightarrow \pm\infty$ this path approaches horizontal asymptotes, and combines with the fixed $\mathbb{R}^{1,2}$ dimensions to make a macroscopically four dimensional theory described by five-branes on $\Sigma(\pm\infty)$. However in the interior of the path there is a non-trivial kink along which we allow the parameters describing the Riemann surface, and hence the parameters of the four-dimensional field theory, to vary, $\Sigma \rightarrow \Sigma(t)$. Since this defect is codimension one in space it describes a domain wall in the four-dimensional $\mathcal{N} = 2$ system [83]. The total geometry is illustrated in Figure 4.13.

In terms of the geometry of the previous section, we can phrase the domain wall construction as follows. The three-manifold, which supports the five-branes is $\Sigma \times \mathbb{R}$. We coordinatize \mathbb{R} by t and loosely refer to as “time,” and we allow the parameters describing the Riemann surface Σ to vary with t . The asymptotic boundaries of the three-manifold M , namely $\Sigma \times \{-\infty\}$ and $\Sigma \times \{+\infty\}$ encode the fact that this domain wall theory does not have an independent existence but couples to the bulk four-dimensional theory.

Now, if the variation of parameters of the Riemann surface Σ is done in an arbitrary way, then the domain wall will break all the supersymmetry of the problem. If we wish to preserve 3d $\mathcal{N} = 2$ supersymmetry, then the domain wall should be half-BPS, and the supersymmetries

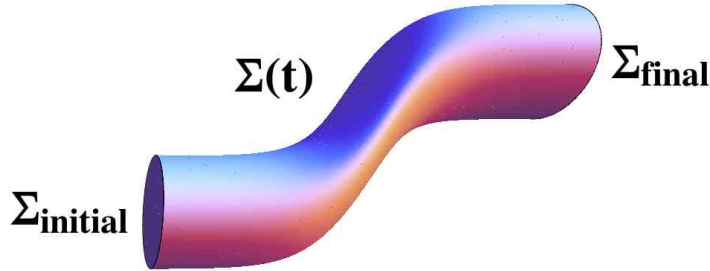


Figure 4.13: The domain wall geometry. Asymptotically, the five-branes wrap the Riemann surfaces Σ_{initial} and Σ_{final} and give rise to theories with four macroscopic dimensions. In the interior, the parameters of Σ vary and describe a domain wall.

preserved in 3d are embedded inside the 4d $\mathcal{N} = 2$ superalgebra as a subalgebra. Such 3d $\mathcal{N} = 2$ subalgebras are labeled by a choice of angle. Note that this also matches the central charge structure. A 3d theory with $\mathcal{N} = 2$ supersymmetry has a real central charge. To get a reduction from $\mathcal{N} = 2$ theory in 4d, which has a complex central charge, to the one in 3d, with a real central charge, we must choose a ‘real’ subspace in the 4d complex central charge plane. Let us choose this direction to correspond to the real axis in the complex plane of the 4d central charges. Then, the four-dimensional and three-dimensional central charges obey by the relation

$$Z_{3d}^i = \text{Re}(Z_{4d}^i). \quad (4.3.7)$$

In terms of the IR Coulomb branch geometry, the domain wall construction means that there is a relationship between the SW curve $\widetilde{\Sigma}$ of the parent 4d $\mathcal{N} = 2$ model and the IR three-manifold \widetilde{M} . Specifically, \widetilde{M} is a one parameter thickening of the SW curve

$$\widetilde{M} = \widetilde{\Sigma}(t) \times \mathbb{R}_t. \quad (4.3.8)$$

This means that every non-trivial one-cycle Γ in \widetilde{M} is inherited from $\widetilde{\Sigma}$. As a result, (4.3.7) yields a simple relationship between the periods of the Seiberg-Witten differential ϕ on the Seiberg-Witten

curve $\widetilde{\Sigma}$, and the periods of the harmonic one-form λ on \widetilde{M}

$$\int_{\Gamma} \lambda = \Re \left(\int_{\Gamma} \phi \right). \quad (4.3.9)$$

In the primary case of interest in this paper, the parent 4d theory can be described by two M5-branes on Σ . Hence, the SW curve $\widetilde{\Sigma}$ is a branched double cover of Σ , where the branch points of the cover are exactly the zeros of the SW differential. In this case, the presentation (4.3.8) of the IR three-manifold implies that \widetilde{M} is a branched double cover of $\Sigma \times \mathbb{R}$, where the branch locus is exactly the one-dimensional *strands* swept out by the zeros of the SW differential during the time evolution. This fact will be of crucial importance to us throughout the remainder of this work.

An Elementary Example

Let us now turn to the most basic example of this construction. We consider an $\mathcal{N} = 2$ theory in 4d which is the theory of a free massive hypermultiplet. We can model this geometrically as above by taking Σ to be simply the complex plane with coordinate x and placing a pair of five-branes there with suitable boundary conditions at infinity. Then the Seiberg-Witten geometry is given by the following curve and differential

$$y^2 = x^2 - m, \quad \phi = y dx. \quad (4.3.10)$$

In the above, the function $y(x) = \pm\sqrt{x^2 - m}$ describes the separation between the two branes. The sign ambiguity in $y(x)$ is consistent with the fact that the two five-branes are indistinguishable. At $y(x) = 0$ the two branches of the function $y(x)$ exchange and hence the two M5 branes connect up into a single smooth object. This is consistent with the general geometry described in section 2: the IR Coulomb branch physics is governed by the geometry of a single smooth five-brane related to the UV description by brane recombination. In this case, the recombined brane is an infinite cylinder which is a branched double cover of the complex plane.

The BPS hypermultiplet of the theory can also be seen from the general analysis in section 2. The non-trivial one-cycle Γ in the cylinder, describes a charge in the 4d IR physics. However in the ambient Calabi-Yau geometry, this cycle is contractible. Physically this means that there is an M2 brane disc which ends on the cycle Γ . The boundary of the disc is a circle made up of two halves, each half corresponds to an interval on each of the two M5 branes stretched between the two branch points. It gives rise to a BPS particle with central charge m in four-dimensions.

Now we would like to construct a domain wall in this theory by considering a one-parameter family of these SW geometries. Thus, we let m vary as a function of a parameter t as

$$m(t) = m_0 + it \quad (4.3.11)$$

Further, we take m_0 to be real, and this will be the resulting three-dimensional central charge. The flow of Z_{4d} is illustrated in Figure 4.14.




Figure 4.14: The flow of the 4d central charge for the A_1 theory. The blue line indicates the trajectory of $Z_{4d}(t)$ shown in blue. The 3d central charge m_0 , is the real part of Z_{4d} .

As a result of this one-parameter variation, the UV description of the theory is two five-branes which wrap the three dimensional space (x, t) . In the IR, the theory is described by a single five-brane described as a branched double cover over (x, t) and given by the equation

$$y^2 = x^2 - (m_0 + it). \quad (4.3.12)$$

Further, the one-form λ and the SW differential are related as

$$\lambda = y(t)dx + \bar{y}(t)d\bar{x} + f dt. \quad (4.3.13)$$

Where f denotes separation of the M5 branes in the cotangent direction to t and is chosen so that $d\lambda = 0$. Notice that this satisfies the key requirement (4.3.7) for preserving three-dimensional $\mathcal{N} = 2$ supersymmetry, namely the periods of λ over one-cycles at constant time t are simply the real parts of the periods of ϕ , and hence the three-dimensional central charge is simply m_0 .

An important fact is that already in this simple example we can see topology changing transitions occurring in the IR geometry as 3d parameters are varying. Specifically, consider the branch loci of the cover. These are given by the two curves

$$x_{\pm} = \pm\sqrt{m_0 + it} \quad (4.3.14)$$

Note that when $m_0 = 0$ the two branch lines meet at $(x, t) = (0, 0)$. Also note that the branch lines have reconnected as m_0 goes from positive values to negative. This reconnection is illustrated in Figure 4.15. In terms of the IR geometry the topology of the cover is jumping as m_0 passes through




Figure 4.15: The reconnection process. The strands are illustrated in black and the blue line indicates the projection of the boundary of the BPS M2 brane. In (c), when the mass m_0 of the particle vanishes, the two strands touch and their individual identity is ambiguous. As m_0 becomes negative the strands reconnect.

zero. In later sections we will interpret these topology changes in terms of mirror symmetries.

General Flows

Let us return to the general discussion of domain walls. We start with an $\mathcal{N} = 2$ theory in four dimensions given by a SW geometry. This can be an arbitrary $\mathcal{N} = 2$ theory, and not necessarily one arising from M5 branes. However, for simplicity of exposition here we describe it for the case of M5-branes. We start with an IR M5-brane geometry in 4d given by the Seiberg-Witten curve $\widetilde{\Sigma}$. We then consider a one real parameter family of these theories to yield the IR three-manifold \widetilde{M} . The real one-form λ and the SW differential are related as

$$\lambda = \phi(t) + \overline{\phi}(t) + gdt \quad (4.3.15)$$

We require that λ is closed, which in particular requires

$$d\lambda = 0 \rightarrow \left(\frac{d\phi}{dt} + c.c. \right) dt + d(g dt) = 0. \quad (4.3.16)$$

This means that $d(\phi + \overline{\phi})/dt$ is cohomologically trivial, and hence has no periods on the SW curve, which in turn implies that the periods of $\text{Re}(\phi)$ do not change with t . Thus to preserve supersymmetry, all the central charges will have to move along straight vertical lines as we move through the flow parameter t . Can this be arranged?

In general the answer is simply “no.” The various central charges of the 4d theory will be related in an intricate way determined by special geometry. In particular they are not independent parameters and hence there is no reason to believe that they can be varied in any particular prescribed way. However, there does exist a class of $\mathcal{N} = 2$ theories, the so-called *complete theories* which have exactly enough moduli and coupling constants to treat the central charges as independent parameters. Further, except for a finite number of exceptional cases, all complete $\mathcal{N} = 2$ theories can be described by two M5-branes wrapping a punctured Riemann surface with suitable boundary conditions at the punctures [12]. Our primary examples will be in the case involving *two* five-branes where we can preserve 3d $\mathcal{N} = 2$ supersymmetry via a flow of 4d central charges in

straight vertical lines. However, in section 8 we also discuss exceptional complete theories which are not M5-brane theories.

For non-complete theories, we cannot vary the central charges arbitrarily if we wish to maintain having a UV-complete theory. However, even for non-complete theories the central charges can vary arbitrarily if one does not insist on a UV completion and views them as effective theories which are UV incomplete, but can in principle be embedded in a UV complete theory. An example of this is pure $SU(N)$ gauge theory, where the non-renormalizable terms $\text{tr } \Phi^k$ with $k > 2$ can in principle be generated if $SU(N)$ is embedded in a bigger UV complete theory.

In any case, our primary examples in this paper will all be associated to domain walls in theories described by pairs of five-branes wrapping punctured Riemann surfaces. In fact our main focus discussed in section 5 will all be generalizations of the free hypermultiplet theory, where the UV five-brane Riemann surface is again the complex plane \mathbb{C} . In that case, the total internal Calabi-Yau threefold is simply \mathbb{C}^3 with its standard symplectic, and holomorphic structure. The abstract flow of Riemann surfaces studied in this section is then a specific instance of a Joyce-Harvey flow construction [98–100] of special Lagrangians.⁸ These equations turn out to be difficult to solve. Luckily, many features of what we need are independent of the detailed solution.

For another perspective on the domain wall geometry we can ask for the dual description for these theories in type IIB. The dual to an M5 brane is a local ALE fibration of the form

$$uv = P(x, y, t). \tag{4.3.17}$$

Where in the above the equation $P(x, y, t) = 0$ defines the locus of five-branes in the original geometry and as t varies describes a one-parameter family of SW geometries. Abstractly, the

⁸ The general structure of this flow equation is as follows. We consider a Riemann surface Σ and a one parameter family of real analytic embeddings $\psi_t : \Sigma \rightarrow \mathbb{C}^3$. Then given any positive bivector X on Σ one studies the following flow equation

$$\frac{\partial \psi_t^i}{\partial t} = g^{ij} (\psi_{t*} X)^{kl} \Re(\Omega)_{jkl}.$$

It is then a fact that if the symplectic form vanishes on the initial surface $\psi_0(\Sigma)$ then the three-manifold swept out by Σ as one varies through time is special Lagrangian.

equation (4.3.17) defines a one-parameter family of Calabi-Yau threefolds, and supersymmetry demands that the resulting seven-dimensional total space have G_2 holonomy. Then, the one-form λ we have discussed is promoted to the three-form ρ which determines the G_2 structure. This three-form fixes the metric completely and hence satisfies the equation

$$d\rho = d *_{\rho} \rho = 0. \quad (4.3.18)$$

These are the analogs of the harmonicity of the one-form λ . If we fix the boundary conditions for the flow, then the G_2 version of Yau's theorem implies that the metric is characterized completely by the three-dimensional real central charges.

4.3.2 Decoupling Limits and R-Twisting

In the previous section we have described a class of domain walls which exist in four-dimensional $\mathcal{N} = 2$ systems described by five-branes on Riemann surfaces. Such domain walls are characterized by the fact that the flow of the 4d central charges is on vertical straight lines. In general such walls will have complicated interactions with the ambient 4d field theories. In this section, we take the key step of decoupling the bulk physics leaving only the remaining 3d $\mathcal{N} = 2$ system. In the process, we see how the domain wall geometries described in this section can be interpreted in terms of R-twisting.

The most important observation is simply the BPS bound in the bulk 4d theory. This states that all charged particles have a mass m which satisfies

$$m \geq |Z_{4d}|. \quad (4.3.19)$$

Consider this bound applied to the bulk 4d theories living at the endpoints of the flow defining the domain wall. To decouple the 4d charged particles from the low-energy physics, it suffices to make them infinitely massive. On the other hand according to the BPS bound (4.3.19) this will be

achieved, provided that the initial and final central charges of the flow have infinite length. Thus, decoupling demands that for both the initial and final condition

$$|Z_{4d}^i| \rightarrow \infty. \quad (4.3.20)$$

Since the flow demands that the 4d central charges evolve along straight lines, the above equation implies that the trajectories in the complex Z_{4d} plane swept out by the central charges during the flow are infinite vertical lines which cross the real axis at the values dictated by the three dimensional real central charges.

If these boundary conditions for the flow are satisfied then all 4d charged particles have infinite mass in the bulk and decouple from the domain wall. However the massless 4d neutral gauge multiplets remain unaffected by this limit. For these fields, which have independent $U(1)$ coupling constants, we are free to choose their three-dimensional physics. We can take these coupling constants to be finite in which case we are left with dynamical gauge field in three dimensions, or we may dial these constants to zero in which case the resulting $U(1)$ appears as a flavor symmetry in three dimensions. Thus, what the decoupling limit (4.3.20) naturally produces is in fact a class of, in general distinct, 3d theories labeled by a choice of whether the $U(1)$'s are gauged or ungauged. We will examine this freedom in detail in section 5. It turns out, that there are some additional $U(1)$'s coupled to the chiral fields which are massed up by Chern-Simons terms, but are nevertheless necessary for describing the full content of the theory.

Finally, we come to a crucial ansatz of our theory of domain walls. We have succeeded in producing a decoupled 3d $\mathcal{N} = 2$ theory, but so far there is no simple relationship between the spectrum of BPS chiral multiplets on the wall and the spectrum of the bulk four-dimensional theory. As the central charges of the 4d theory flow in general they flow at different speeds and cross each other at various times. Such crossings lead to the wall-crossing phenomenon. If they occur they imply that the effective spectrum of the 4d theory is changing, and hence during the flow the 4d

theory is crossing walls of marginal stability in its moduli space. As a key simplifying assumption, we will now assume that such crossings do not happen at any time during the evolution. Thus our assumption for the rates of flow of central charge is:

- During the flow the central charges retain whatever phase order they started with.

This is a natural assumption for solutions to the Joyce flow equations. For example in the context of Janus domain walls, such BPS walls do indeed exist [101,102]. One can see that a simple ansatz satisfying the above assumption is given by taking the 4d central charges to flow linearly in some coordinate with a speed controlled by their real part

$$Z_i(t) = Z_i^0 - i\Re(Z_i^0)t. \quad (4.3.21)$$

Given such an ansatz for the flow, one can see that the central charges retain their phase order and have constant real part.

Let us now take stock of the resulting properties of the domain walls we have described.

- They are characterized by a phase ordered flow of the central charges along vertical lines.
- In the decoupling limit, the central charges begin at $i\infty$ and terminate at $-i\infty$.

Notice that if we ignore the length of the central charges, which varies during the flow in time, the first property is identical to the evolution of the central charges generated by an R-symmetry rotation $Z \rightarrow e^{i\theta}Z$. Further, the decoupling limit boundary conditions can also naturally be interpreted as saying that as time evolves, the central charges rotate through angle of π .

Thus, we have in a sense succeeded in making the R-twisting compactification physical. To preserve the standard supersymmetry, the central charges flow along straight vertical lines and hence their lengths during the evolution are not constant. In this sense, the time evolution we have constructed is not merely a phase rotation on Z . However, when our ansatz for the rates of flow is satisfied, the central charges of 4d flow in way which respects their phase order and in this sense

the time coordinate we have constructed can be interpreted as, essentially, the phase of Z . Further, our decoupling limit boundary conditions mean that under the complete time evolution the phase rotates by π . As in the general story of R-twisting this leads to a simple relationship between the 4d BPS spectrum and the 3d BPS spectrum, and for this reason we refer to the flow as the ‘R-flow.’

3d BPS Spectrum from Trapped 4d BPS States

Now we come to the central consequence of the decoupling limit developed in the previous section:

- The 3d BPS chiral spectrum is in one-to-one correspondence with the 4d BPS spectrum.

To see this fundamental fact, we observe that each 4d central charge, Z_{4d} , is the central charge of a certain BPS particle in 4d. However for a typical point along the flow such a particle is not BPS in the three-dimensional sense. Indeed, to be BPS in three dimensions $Z_{4d}(t)$ must align with real direction defining the 3d central charges. This means that at the time $t = t_i$ when $Z_{4d}^i(t_i) = \text{real}$ the corresponding BPS state will be a 3d BPS state with the same central charge. In other words, the 4d BPS state with central charge Z_i is trapped in the wall at $t = t_i$. Note that this is physically sensible, in the sense that the 4d mass $m(t) = |Z_i(t)|$ is minimized at t_i , where the length of $Z_i(t)$ is minimized. Thus, the particle is trapped at $t = t_i$ simply by energy considerations. If the boundary condition of the flow were such that the asymptotic central charges had finite length, then the difference in length between the 3d central charge and the 4d central charge at either side of the wall would be finite and the 3d BPS chiral particle, while trapped on the wall, could escape out to the bulk for a finite cost in energy. However, in the decoupling limit where the 4d central charges become parametrically large as $|t| \rightarrow \infty$, the potential energy well trapping the 3d particle becomes infinitely deep, and the 4d bulk physics decouples.

Finally if we now invoke our ansatz where the order of the phases of the central charges do not change, it follows from our discussion above that *for each chamber of the 4d theory, we*

get a chamber of a 3d theory, where the corresponding 4d BPS states are trapped and become BPS states of the 3d theory. Given that by changes of parameters in the 4d theory we can go from one chamber to another (passing through walls of marginal stability), it suggests that the same should be true for the 3d theory, at least as far as the IR behaviour is concerned. In particular the initial conditions for the R-flow which can vary continuously, should not affect the IR dynamics. In other words we should get *a family of dual 3d theories labeled by chambers of the parent 4d theory*. In the remainder of this paper we provide evidence for this claim through a study of explicit examples. We aim to illustrate that, via this correspondence, the three-dimensional version of wall-crossing is mirror symmetry.

4.4 4d BPS States of A_n Theories

At the conclusion of the previous section, we have arrived at a class of domain wall theories whose 3d BPS particles are in one-to-one correspondence with the ambient 4d BPS particles. In order to apply this useful fact to the study of 3d field theories, we will need to make use of various methods for counting 4d BPS states. These have been discussed in detail in the previous chapters of this dissertation. In this section we present a brief review of recap of these techniques: ideal triangulations, and BPS quivers. Our goal is to isolate the essentially algorithmic features of each method in the simplest possible examples, the so-called A_n Argyres-Douglas theories [34], relevant for subsequent analysis in this work. We refer the reader to the previous chapters and to the original papers [11, 13, 22, 26, 28, 53, 77] for a complete treatment.

4.4.1 Ideal Triangulations

The first method we recall is that of ideal triangulations developed in detail in [13]. We consider a pair of M5 branes wrapping a Riemann surface Σ . The Seiberg-Witten geometry is described by a quadratic differential ϕ^2 on Σ . As in previous sections, ϕ^2 defines a double cover $\tilde{\Sigma}$

of Σ . This double cover is the Seiberg-Witten curve and on $\tilde{\Sigma}$, the one-form ϕ is the Seiberg-Witten differential.⁹ The key idea in this method is to recognize that a BPS M2-brane, describing a BPS particle in four-dimensions, must have minimal area. In particular the boundary of this two-brane defines a one-cycle in the Seiberg-Witten curve and its length must be minimal.

It is straightforward to translate this idea into concrete equations formulated on the original curve Σ itself. The boundary one-cycle of the M2-brane projects to Σ and defines a curve γ parametrized by a variable s . Since the central charge is measured by ϕ it is this quantity which characterizes the notion of minimal length and hence $\gamma(s)$ solves the differential equation (4.3.2) [11].

$$\phi|_{\gamma} = e^{i\theta} ds. \quad (4.4.1)$$

As discussed in the previous section, the angle θ entering the equation above is the angle of the central charge of the particle defined by the two-brane.

For most of the remainder of this paper, we will be focused on a simple class of examples involving theories which have a finite number of BPS hypermultiplets and no BPS vector multiplets. These examples are the so-called Argyres-Douglas *ADE* theories [34]. The A_n case, which will be our main focus in this section, is characterized by a particularly simple geometry. The curve Σ which supports a pair of five-branes is just the complex plane \mathbb{C} . Giving this plane the complex coordinate x , the Seiberg-Witten differential ϕ defining the central charge density is given by a polynomial in x of degree $n + 1$

$$\phi = \sqrt{P_{n+1}(x)} dx. \quad (4.4.2)$$

In these cases the BPS counting problem is particularly simple as we will see below.

For the vast majority of angles θ , there is no BPS state whose central charge has that given angle, and hence no finite length solution to the flow equation (4.4.1). In this case, we can

⁹ We are being a bit loose with notation here. On the base Σ the quantity ϕ^2 is not the square of anything. Only on passing to $\tilde{\Sigma}$ does it have a globally well defined square root.

draw a simple combinatorial picture which characterizes the global asymptotic properties of the flow. Then, if we perturb the angle θ by a small amount, we will not encounter any BPS states and hence the combinatorial diagram will be stable. On the other hand, if we tune θ a large amount past a critical angle which supports a BPS hypermultiplet, the global flow diagram will jump in a definite way. As θ varies from 0 to π we encounter all BPS particles in the theory and thus the BPS spectrum is realized geometrically as a sequence of moves in the flow diagram.

In the context of our simple A_n theories, let us now be more specific and introduce the asymptotic flow diagram, an *ideal triangulation*, and the operation on it a *flip*, determined by a BPS state. We first draw a large circle the x plane which defines the asymptotic boundary of \mathbb{C} . Then, on this circle we mark the $(n + 3)$ -rd roots of unity, which makes the boundary circle into an $(n + 3)$ -sided polygon. The complex plane is represented by the interior of this polygon. We then triangulate this space using only lines that end at the vertices of the polygon. So defined, we have constructed an ideal triangulation of the $(n + 3)$ -gon. This triangulation has the important property that each triangle contains exactly one zero of the differential ϕ . An example for the case of A_1 is shown in Figure 4.16a.

Now that we have introduced ideal triangulations, it remains to explain how BPS states are visualized in this setup. As we have described above, the BPS states are sudden changes in the flow as we rotate θ . This means that they are described by operations, known as flips, which change the triangulation. The flip operation can be performed on any internal (*i.e.* non-boundary) edge E in the triangulation. We first delete E making a quadrilateral, and then replace E with E' , the unique other edge in the quadrilateral which forms a triangulation as shown in Figure 4.16b. The name of the operation, a flip, is justified by the fact that the new triangulation is related to the old one by rotating the edge E .

In terms of the flow equation (4.4.1), the significance of a flip is easy to explain. Each triangle in the triangulation contains exactly one zero of ϕ and the trajectories of the flow equation

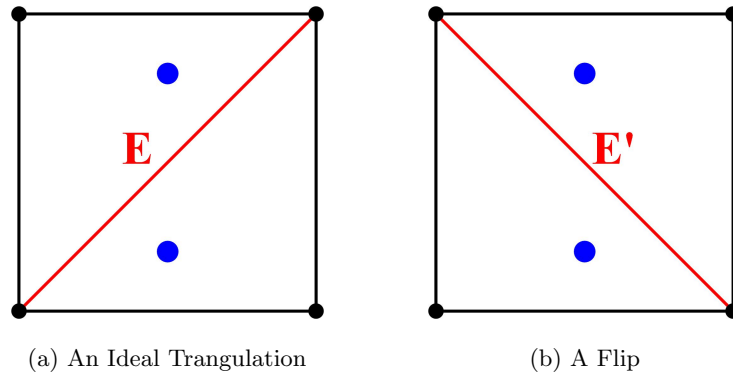


Figure 4.16: A sample ideal triangulation. In (a), we have an ideal triangulation of the square used for counting BPS states in the A_1 model. The red line denotes an interior edge. The blue points the zeros of the differential ϕ^2 . In (b), the edge E has flipped to E' .

(4.4.1) emanating from ϕ asymptotically terminate on the vertices of the triangle containing the zero. Meanwhile, an internal edge E is an object at the interface of two triangles, and hence equivalently two zeros. As the BPS angle is rotated towards a critical value, a pair of trajectories, one from each zero, become near to each other. Exactly at the critical angle the trajectories connect, leading to a BPS hypermultiplet described by a segment which crosses the edge E . Just after the critical angle the trajectories again separate and the edge E is replaced with E' .

Given that an individual BPS state appears as a flip, the complete BPS spectrum is then characterized by a sequence of flips. To describe the sequence, we note that as the BPS angle rotates from 0 to π , all BPS particles will be seen by the flow and hence all flips will occur. On the other hand, as θ rotates through π the quadratic differential returns to itself, except that the asymptotic vertices rotate counterclockwise by an $(n + 3)$ -rd root of unity. In other words, the vertices of the polygon rotate one unit to the left. These facts determine how a BPS spectrum is encoded in a sequence of flips:

- A BPS spectrum of A_n is a sequence of flips on the internal edges of a triangulation such that,

after all flips have occurred, the ideal triangulation has returned to itself up to a rotation by an $(n + 3)$ -rd root of unity.

The allowed sequences of flips also satisfy several minimality properties. Namely, if the edge E flips to E' then the edge E' is not the next edge to flip, and if the sequence at any point reaches the initial triangulation rotated by an $(n + 3)$ -rd root of unity, then it must terminate.

As developed in detail in [13], the most fascinating and useful aspect of this description of BPS spectra is the ease with which one can describe wall-crossing. In this context, the fact that there exists more than one chamber of BPS states is simply reflected in the fact that there exists more than one sequence of flips satisfying the above criteria. Indeed, in the simplest example, the A_1 model, there is exactly one possible sequence of flips shown in Figure 4.16, and hence the BPS spectrum consists of exactly one BPS hypermultiplet as described in section 3. However the A_2 model, corresponding to the pentagon, already exhibits two such sequences and hence two chambers of BPS spectra, as illustrated in Figure 4.17. This geometric fact will be significant for us in our study of 3d field theories in section 5 and beyond.

4.4.2 Quivers and Mutation

A second, equivalent method for studying BPS states of the A_n models is to make use of BPS quivers and mutations [28, 77]. In this method, BPS states are described by studying a quiver quantum mechanics on the worldvolume of a BPS particle. In the A_n examples the quiver is given by an oriented version of the A_n Dynkin diagram.

$$\begin{array}{ccccccc}
 \bigcirc & \longrightarrow & \bigcirc & \longrightarrow & \cdots & \longrightarrow & \bigcirc \\
 \gamma_1 & & \gamma_2 & & \cdots & & \gamma_n
 \end{array} \tag{4.4.3}$$

In (4.4.3) γ_i denote the charges of an elementary basis of hypermultiplet which are always stable states. Since charges are defined by one-cycles on the SW-curve each γ_i is associated to an element of $H_1(\tilde{\Sigma})$. The number of arrows between the nodes of the quiver is then fixed by computing the




Figure 4.17: The two BPS chambers of the A_2 model realized as a sequence of flips. The upper-left pentagon, and upper-right pentagons, are the initial and final triangulations respectively. Each arrow denotes a flip on one of the internal edges shown in red. Following the solid arrows we find a chamber with two BPS states. Following the dashed arrows we find a chamber with three BPS states.

intersection product of cycles, $\gamma_i \circ \gamma_j$, or equivalently the four-dimensional electric-magnetic inner product of the associated 4d particles. All remaining BPS particles in the spectrum can be viewed as supersymmetric bounds states of these, which exist in the quiver quantum mechanics theory defined by (4.4.3) [6–8].

In comparing the method of ideal triangulations to that of BPS quivers, the quiver diagram plays the role of the ideal triangulation of a polygon. It is a basic combinatorial diagram which encodes information about the spectrum. The analogous operation to a flip is then a quiver *mutation* which acts on the quiver to produce a new quiver. This operation can be defined on any node i of

the quiver, and acts on the charges as follows:

$$\begin{aligned} \gamma_i &\longrightarrow -\gamma_i \\ \gamma_j &\longrightarrow \begin{cases} \gamma_j + (\gamma_i \circ \gamma_j)\gamma_1 & \text{if } \gamma_i \circ \gamma_j > 0 \\ \gamma_j & \text{if } \gamma_i \circ \gamma_j \leq 0 \end{cases} \end{aligned} \quad (4.4.4)$$

Thus, after mutation we can form a new quiver by computing the intersection of the charges on the right-hand-side of (4.4.4).

Now, in the method of ideal triangulations, each BPS state is associated to a flip. Similarly in the method of BPS quivers each BPS state is associated to a mutation. It then follows that the complete BPS spectrum is captured by certain sequences of mutations. These sequences are defined by the following properties [77]:

- The initial quiver appears as in (4.4.3) with node charges γ_i .
- The final quiver has charges $-\gamma_i$.
- At each step one may mutate on any node whose charge γ can be expressed as

$$\gamma = \sum_i n_i \gamma_i, \quad (4.4.5)$$

where in the above the n_i are non-negative integers.

Let us see how the two examples considered in the previous section, the A_1 and A_2 theories, are described using this method. In the case of A_1 , the quiver consists of one node and there is trivially one possible sequence of mutations.

$$\begin{array}{ccc} \bigcirc & \longrightarrow & \bigcirc \\ \gamma_1 & & -\gamma_1 \end{array} \quad (4.4.6)$$

This agrees with our identification of this theory as a single free hypermultiplet. There are no interactions and hence no wall-crossing. Meanwhile, in the case of the A_2 theory things are more

interesting. The two spectra described in Figure 4.17 map to two possible sequences of mutations.

The first sequence, with two BPS particles, is:

$$\begin{array}{ccccccc}
 \bigcirc & \longrightarrow & \bigcirc & \longrightarrow & \bigcirc & \longleftarrow & \bigcirc & \longrightarrow & \bigcirc & \longrightarrow & \bigcirc & & (4.4.7) \\
 \gamma_1 & & \gamma_2 & & \gamma_1 & & -\gamma_2 & & -\gamma_1 & & -\gamma_2 & &
 \end{array}$$

While the second sequence describing the second chamber with three BPS particles is:

$$\begin{array}{ccccccccccc}
 \bigcirc & \longrightarrow & \bigcirc & \longrightarrow & \bigcirc & \longleftarrow & \bigcirc & \longrightarrow & \bigcirc & \longrightarrow & \bigcirc & \longrightarrow & \bigcirc & \longrightarrow & \bigcirc & & (4.4.8) \\
 \gamma_1 & & \gamma_2 & & -\gamma_1 & & \gamma_1 + \gamma_2 & & \gamma_2 & & -\gamma_1 - \gamma_2 & & -\gamma_2 & & -\gamma_1 & &
 \end{array}$$

One can easily generalize from these examples to determine the spectrum in the various chambers of A_n theories for larger n . In our applications of this method to 3d $\mathcal{N} = 2$ theories in later sections, one detail of these calculations will be important to us:

- At the conclusion of a sequence of mutations the original quiver charges $\{\gamma_i\}$, as a set, have been changed to $\{-\gamma_i\}$. However, they may have also undergone a non-trivial permutation by an element $\chi \in S_n$. Indeed, in the case of the first chamber of A_2 described by (4.4.7) χ is the identity element, while in the case of the second sequence, described by (4.4.8) χ is the non-trivial element in S_2 . This permutation proves important for our considerations later in this paper.

4.5 Tetrahedra and Braids

Armed with the technology of the previous section, we now return to our general discussion of 3d $\mathcal{N} = 2$ theories constructed as domain walls in 4d $\mathcal{N} = 2$ theories. Our aim will be to apply the techniques of ideal triangulations and quiver mutations to develop a detailed geometrical toolkit for extracting the physics of the domain wall.

Throughout all of the examples discussed in this section, the 4d theory will be one of the A_n models whose BPS spectra we have now described in some detail. In the UV these theories are

determined by a pair of five-branes wrapping the complex plane \mathbb{C} and this leads to a particularly simple geometry of the associated three-manifold M defining the domain wall theory. To be specific, M is simply a thickening of the complex plane to $\mathbb{C} \times \mathbb{R}$, where \mathbb{R} describes the time parameter of the R-flow in section 3. Along this flow all the central charges move in vertical straight lines, and central charges cross the real axis in phase order. As we have previously noted this means that each 4d BPS state will appear as a 3d BPS chiral particle trapped along the wall. Further, if we ignore the length of the 4d central charges along the flow and concentrate only on their angles, then we may interpret the fact that the particles cross in phase order as an identification of the time coordinate with the BPS angle θ of the 4d central charges. In this section our aim will be to make use of this fact to determine a concrete Lagrangian description of the field theory on the domain wall.

First, we study the structure of the three-manifold M . As we have described above, $M = \mathbb{C} \times \mathbb{R}$, however the boundary conditions on the circle at infinity in the complex plane are fixed for all time. Thus, we will in fact work in a quotient three-manifold defined by identifying these asymptotic regions for all time. It then follows that our three-manifold M can be viewed as an infinite solid ball with an asymptotic S^2 boundary. This boundary two sphere is naturally partitioned into two components, the northern hemisphere corresponding to the initial boundary condition, and the southern hemisphere corresponding to the final boundary condition. We will refer to these hemispheres as the “front” and “back” face of our three-manifold respectively. The equatorial circle of the boundary S^2 is where the front and back faces are glued together and is the boundary circle inside \mathbb{C}^2 that is identified for all time. Further, both the front and back face of our three-manifold describe an A_n theory, and as such these faces are naturally equipped with ideal triangulations of $(n + 3)$ -gons governing their BPS spectra. Since the complete flow through time corresponds to a rotation of the BPS angle by π , the final triangulation differs from the initial triangulation by a rotation by $\frac{2\pi}{n+3}$. An example of the geometry for the case of A_2 is shown in

Figure 4.18.

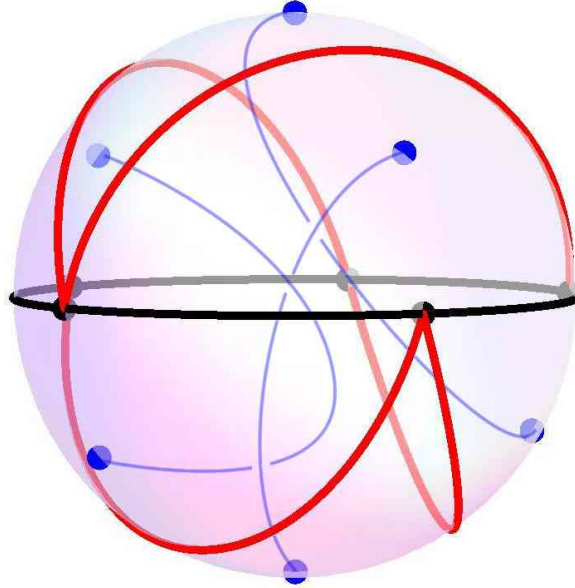


Figure 4.18: The manifold M and its boundary triangulation for the case of A_2 . M is an infinite solid ball. Its boundary two sphere has two faces given by the northern and southern hemispheres. Each face is a triangulated pentagon. The vertices of the pentagon are shown in black while the arcs in the triangulation are shown in red. The blue dots are the zeros of the SW differential. One such zero occurs in every triangle. As we flow through time, the zeros on the front face interpolate to those on the back face.

As we flow through time, the initial triangulation will evolve by a sequence of flips as described in section 4. We will see that this sequence of flips will naturally endow the three-manifold M with a decomposition into tetrahedra. Since the 4d BPS states correspond to both tetrahedra and trapped 3d BPS particles, we then learn that each tetrahedron in the manifold M will encode the existence of a 3d BPS particle. In this way we will make contact with the work of [16]: the tetrahedron is a kind of basic BPS building block of these 3d theories. Further as we will see, the fact that the ambient 4d theories can undergo wall crossing, and hence have

different numbers of flips, becomes the statement that a given three-manifold admits many distinct decompositions into tetrahedra. In our context, these distinct tetrahedral decompositions of M will encode different dual descriptions of the same IR field theory.

Next in our analysis, we describe the IR geometry which is given by a branched double cover $\widetilde{M} \rightarrow M$. Since M is an infinite solid ball its topology is trivial. Thus, up to data at infinity, the situation is exactly the same as that of double covers of S^3 described in section 2. In particular, \widetilde{M} is completely fixed by the associated branch locus knot in M . In our context, this knot is exactly the set of zeros of the one-form λ , or equivalently the zeros of the evolving Seiberg-Witten differential ϕ . On the front face of M the differential ϕ for the A_n model has exactly $n+1$ zeros and each zero resides in a triangle in the ideal triangulation. As we flow through time the zeros evolve continuously and sweep out a *braid* composed of $n+1$ strands. As we will argue, the structure of this braid completely determines the 3d physics with BPS particles in direct correspondence with the crossings in the braid diagram. An example is shown in Figure 4.19.

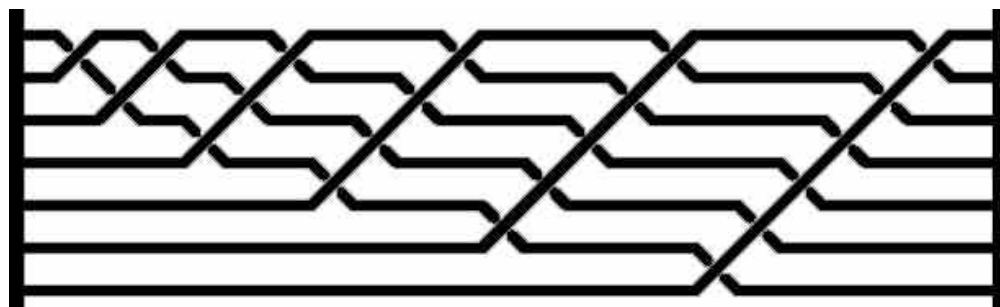


Figure 4.19: A braid with seven strands describing a chamber in the A_6 theory. Each strand follows the evolution in time of a zero of ϕ . 3d BPS particles are described by braid moves. The endpoints are the zeros of the initial and final SW differential.

To complete the description of the 3d theory from its braid diagram, there is one final step: we must turn the braid into a knot; that is we choose a way of identifying the endpoints of the strands of the braid in pairs to turn all components of the braid into closed loops. This step is physically

natural from a number of perspectives. First, our three-manifold M is non-compact, and hence we must impose boundary conditions. These boundary conditions involve specifying a choice of which cycles in \widetilde{M} are contractible at infinity and which remain non-trivial. Since all cycles in the cover \widetilde{M} can be localized to a neighborhood of the branching link, this choice is equivalent to a specification of how the braid is capped off to form a closed knot. Alternatively, from the perspective of the domain wall theory we can see the need for boundary conditions as follows. At the conclusion of the decoupling limit described in section 3 all the massive BPS states of the ambient 4d theory have decoupled. However there remains the coupling to the $U(1)$ gauge and flavor symmetries. To completely specify the theory on the wall we must specify how we couple our 3d field theory to these vectors. Since the coupling constants of these $U(1)$'s are arbitrary parameters, we can choose whether in three dimensions a given $U(1)$ appears as a gauge or global symmetry. In fact such coupling choices for the A_n are in direct correspondence elements of $Sp(2n, \mathbb{Z})$, where the various S transformations act by changing the set of gauged versus global $U(1)$'s and the T transformations appear as changes in the Chern-Simons levels. We will see how these facts are made geometrically manifest in the course of our analysis.

4.5.1 The Tetrahedron theory

We begin with the simplest example of domain walls in the A_1 theory. In 4d, this is the theory of a free hypermultiplet, and the R-flow of central charges for this example was studied in section 3. In this case, the boundary triangulations of the front and back face are squares, and as we flow through time the triangulation evolves by a single flip to produce a single tetrahedron shown in Figure 4.20. We know that this flip is naturally associated to a 3d BPS chiral particle which has become trapped on the wall, and thus this theory of two five-branes on a tetrahedron supports exactly one BPS chiral particle. The mass of this particle, m_0 is the real part of the 4d mass of the parent 4d hypermultiplet.




Figure 4.20: The tetrahedron associated to the A_1 domain wall. The tetrahedron is viewed with its front face pointing out of the page. The square of the A_1 theory is given by the black edges. The red diagonal flips as one flows from the front to the back face. The black dots denote the two zeros of the SW differential on the front face. As we flow through time, these zeros evolve to the two zeros on the back face shown in white. In the process they sweep out two strands.

To study the geometry in more detail, we track the evolving zeros of the SW differential as we move through the geometry of the tetrahedron. In each triangle in both the front and back face there is one zero, and as time flows they determine a braid composed of two strands. At exactly one critical time the strands of the braid become closest to each other and the BPS chiral particle in 3d appears. We encode this fact in the braid diagram by drawing exactly one braid move as shown in Figure 4.21a.

In terms of the geometry of section 2, the single BPS particle appears as a segment connecting the two strands of the braid. Since the BPS particle is also associated to the one braid move we can view this BPS segment as being localized at the crossing in the braid diagram.

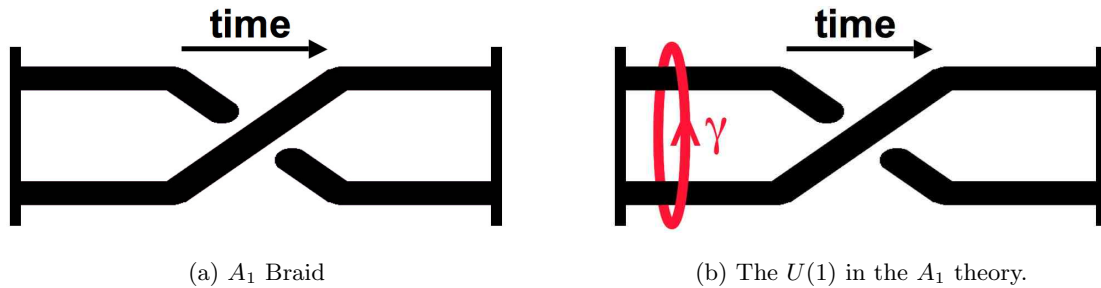


Figure 4.21: In (a), the braid for the tetrahedron theory. The single particle is encoded by the single braid move. In (b), the red cycle γ circles the two strands of the braid. The particle at the crossing is charged under this cycle.

As explained in section 2 such a particle carries a $U(1)$ charge under the cycle γ which wraps around the two strands of the braid illustrated in Figure 4.21. Depending on boundary conditions to be specified, the cycle γ may be non-contractible, in which case it is gauged, or it may be contractible at infinity, in which case the $U(1)$ will survive as a flavor symmetry of the theory. Thus in either case, the BPS particle carries a unit charge under this $U(1)$. Note that in the limit where the particle is massless, the two strands of the braid intersect. Thus, we can view the separation between the strands as proportional to the mass of the particle and the effect of going from overcross to undercross corresponds to changing the sign of the mass for the chiral field. Finally, we will always make the convention that time flows from left to right in the braid diagram. So defined the configuration of Figure 4.21a encodes a charged BPS particle with charge $+1$ under the cycle γ .

Thus far, the braid we have introduced is simply a diagrammatic notation for the rather trivial particle content of the tetrahedron theory. However, the reason that the braid is useful is that operations on the braid diagram have a natural physical interpretation. We will illustrate this feature throughout the course of our analysis. To begin, the first and most basic point we address is the proof that the field theory we have defined is canonically associated to the braid.

What this means is the following. The braid group on two strands is an infinite cyclic group which is generated by a single element b which acts on the two strands, as in Figure 4.21a, by braiding the lower strand over the upper strand in time order. Then the tautological relationship $b^{-1}b = 1$ translates to the clear geometrical fact that an insertion of an overcross followed by and undercross at any point in the braid is trivial as illustrated in Figure 4.22.




Figure 4.22: The braid group relation $b^{-1}b = 1$. In the gray region, the strands braid with each other in an arbitrary manner.

Now, in our physical context we may ask whether the relation in Figure 4.22 is satisfied. To address this we follow the tentative dictionary set in the previous paragraphs. For each crossing in the diagram we add a single chiral particle to the theory. Thus, in the left of Figure 4.22, the relevant region where $b^{-1}b$ has been inserted corresponds to two particles X and Y . However, as we will argue later in this section, these particles have opposite $U(1)$ charges. This means that there is an invariant superpotential term

$$W = \mu XY. \tag{4.5.1}$$

Furthermore, we know from our general discussion of M2-brane contributions to the superpotential in section 2 that exactly in this situation we expect to find such a quadratic contribution to W . Indeed, the region of the overcross followed by undercross bounds a disc which is precisely the projection of an M2 describing a quadratic interaction between the inserted particles as illustrated in Figure 4.23.

Now, in equation (4.5.1) the parameter μ is a complex (as opposed to real) mass for the fields X and Y . Such a mass term means that the fields X and Y are irrelevant in the infrared and



Figure 4.23: The superpotential coupling the fields corresponding to the insertion of bb^{-1} . The green region lifts to the boundary of an M2-brane instanton which gives rise to a quadratic interaction between the particles.

may be safely removed from the spectrum. This should be contrasted with the case of particles with non-zero real masses. In the latter case, even though such fields are massive, their real mass is detected by the partition function of the theory as we will discuss in sections 7 and 8. By contrast, the partition function is independent of complex masses such as μ and thus we may freely take these to be parametrically large. Doing so, we find that the insertion of $b^{-1}b$ in the braid diagram is physically equivalent to inserting the identity, i.e. no particles whatsoever. In this way, we have verified the braid group relation described by Figure 4.22.

Boundary Conditions and $SL(2, \mathbb{Z})$

Next in our analysis, we turn to the discussion of boundary conditions for the theory of two M5-branes on the tetrahedron. As we have previously discussed, what the domain wall and decoupling limit constructs for us is a 3d theory, together with an arbitrary choice of coupling to the background $U(1)$ multiplet. On such field theories, there is a natural action of $SL(2, \mathbb{Z})$ [97, 103] defined by the action of its S and T generators as:

- T acts to increase the Chern-Simons level of the background $U(1)$ by $\hat{k} \rightarrow \hat{k} + 1$.
- S acts to gauge the $U(1)$ in three dimensions, and introduces a new background $U(1)$ which is the dual flavor group.

Thus, $SL(2, \mathbb{Z})$ does not act as a duality group, but simply acts on such a theory to produce a new one. As we will see, in our context, this $SL(2, \mathbb{Z})$ is realized as acting on our choice of boundary conditions.

The simplest way to study the boundary conditions is to consider the IR geometry $\widetilde{M} \rightarrow M$. This is a double cover of M branched over the braid described in the previous section. In particular, the boundary of M as an S^2 which contains the four endpoints of the braid, two from the front face and another two from the back face. It follows that the boundary of \widetilde{M} is a double cover of S^2 branched over four points, and therefore $\partial\widetilde{M}$ is a torus. The three-manifold \widetilde{M} fills in this boundary smoothly, and is thus a solid torus.

Alternatively, one can also see the fact that \widetilde{M} is a solid torus by recalling that the tetrahedron theory is determined by a one-parameter thickening of the A_1 theory in 4d. The Seiberg-Witten curve for the latter is a cylinder. Then, \widetilde{M} is a thickening of this cylinder. It has boundary the SW cylinders associated to the front and back face A_1 theories which are connected at their respective ends to make the surface ∂M into a torus as illustrated in Figure 4.24.




Figure 4.24: The IR geometry for the tetrahedron theory. In (a), we see the SW curve, in this case a cylinder, for the 4d A_1 theory. In (b), the three-manifold \widetilde{M} obtained as a thickening of the SW curve. Topologically this thickened cylinder has an asymptotic boundary of a torus.

Now we are equipped to specify boundary conditions. We will modify the manifold \widetilde{M} by adding data at infinity which turns it into a closed manifold without boundary \widetilde{M}_c . Then, all fields are required to be well-behaved on \widetilde{M}_c . Since \widetilde{M} has boundary given by a torus, to close \widetilde{M} means to glue it to another three-manifold whose boundary is a torus, in other words we simply glue \widetilde{M} to another solid torus. From this description, we see that our choices of boundary conditions are labeled by the gluing map $g : T^2 \rightarrow T^2$ that specifies how the boundary tori are glued. Up to isotopy, such gluings g are specified by their $SL(2, \mathbb{Z})$ action on the homology of the boundary of the torus. The manifolds \widetilde{M}_c that we obtain from such gluing are exactly the *lens spaces*. For example, gluing two solid tori with the identity map makes, $S^2 \times S^1$, while gluing with the S transformation produces S^3 . More generally, given p and q relatively prime, we consider the following element of $SL(2, \mathbb{Z})$:

$$g = \begin{pmatrix} m & n \\ p & q \end{pmatrix}. \quad (4.5.2)$$

Where in the above m, n are chosen such that g has determinant one. Then, the three-manifold obtained by gluing two solid tori with the map g is the Lens space $L(p, q)$.¹⁰

One can see from this description that the S and T generators have the desired physical effect of gauging, and shifting the level \hat{k} respectively. Indeed, for example consider as a starting point the theory on $S^3 = L(1, 0)$. This manifold has no homology and hence no gauge group. Acting with S changes the gluing to produce $S^2 \times S^1$. Since this has first Betti number one, the $U(1)$ has been gauged, which is indeed the appropriate action for the generator S . Similarly, we can act on the S^3 theory with the transformation T^p . This means that we are gluing two solid tori with the map

$$g = T^p S = \begin{pmatrix} p & -1 \\ 1 & 0 \end{pmatrix}. \quad (4.5.3)$$

¹⁰ Recall that, for any choice of signs $L(\pm p, \pm q)$ are all identical. Thus we can be somewhat lax about signs in the following.

This again produces the $S^3 = L(1, 0)$. However, the integer p in the above is physical as it encodes the CS level of the coupling of the theory to the background $U(1)$ flavor symmetry now given by $\hat{k} = p$. Indeed, to make this manifest we can now further act by S . This gauges the $U(1)$ which is now at level p . It is specified by the gluing map

$$g = ST^pS = \begin{pmatrix} 1 & 0 \\ p & 1 \end{pmatrix}, \quad (4.5.4)$$

and hence results in the Lens space $L(p, 1)$. This space has first homology group that is pure torsion $H_1(L(p, 1)) \cong \mathbb{Z}_p$, and thus, as explained in section 2 describes a gauged $U(1)$ CS theory at level $\hat{k} = p$ as desired.

We can further illuminate this $SL(2, \mathbb{Z})$ structure by alternatively studying it from the point of view of the branching braid which encodes the structure of the cover $\widetilde{M} \rightarrow M$. The $SL(2, \mathbb{Z})$ action on the homology of the boundary T^2 of \widetilde{M} is obtained by motions involving the four branch points in the cover $T^2 \rightarrow S^2$. Since these four endpoints are precisely the endpoints of the braid, this means that the $SL(2, \mathbb{Z})$ action can be seen as acting on the braid. To describe this action, we must first state how we specify boundary conditions at the level of a braid. Our infinite tetrahedron can be compactified to S^3 by adding a point at infinity. As discussed in section 2, a double branched cover of S^3 is completely specified by its branching knot \mathcal{K} . Thus, to specify the boundary conditions we must close our braid into a knot. We do this by identifying the endpoints of the braid in pairs. Specifically, we glue the initial points at $t = -\infty$ together, and the final points at $t = +\infty$ together. In this way make a closed knot as illustrated in Figure 4.25.

In general, for the A_1 theory, we will always specify boundary conditions by gluing initial and final points as above. We illustrate this diagrammatically, with the dashed string shown in Figure 4.25b to emphasize that this gluing is boundary data at infinity. With this prescription, we can now specify completely the geometry of the compactified double cover \widetilde{M}_c . For example, in the case of the braid of Figure 4.21a describing the basic tetrahedron theory, this procedure produces

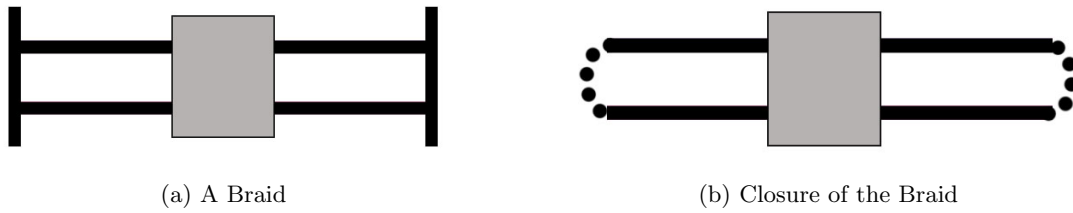


Figure 4.25: Imposing boundary conditions. In (a) we see a braid. In the gray region the strands braid with each other in an arbitrary manner. In (b), the strands are connected by the dashed string to make a closed knot.

an unknot. Then \widetilde{M}_c is double cover of S^3 branched along the unknot and hence is also an S^3 .

Now we are equipped to discuss the action of $SL(2, \mathbb{Z})$ on closed braids. Let us first consider the T generator. This is to act by increasing the CS level for the background $U(1)$ by $\hat{k} \rightarrow \hat{k} + 1$. We can interpret this action by making use of the quantum parity anomaly. This states that upon integrating out a particle of mass $m > 0$ with charge ± 1 the CS level shifts as $\hat{k} \rightarrow \hat{k} + 1$. In terms of its action of CS levels, the operation of adding a massive particle is therefore identical to the desired T operation. In the above, we have associated the charged particles to the crossings in the braid diagram, that is to the action of the braid group generator $b^{\pm 1}$. Sticking to this principle, means that we simply identify the action of the $SL(2, \mathbb{Z})$ element T with the insertion of b^{-1} at the conclusion of the braid, as in Figure 4.26. In order to only modify the CS level, the particles




Figure 4.26: The definition of the operator T . In (a) we see a generic braid. In (b) the action of T on this braid.

inserted by the T transformation should be interpreted as having parametrically large mass. This is natural if we view T as acting on boundary conditions of the theory. Then, the closure of the braid associated to T^p acting on the basic tetrahedron braid in Figure 4.21a is again an unknot. However, the integer p is physical and keeps track of the background CS level. Thus, although all such examples produce covers \widetilde{M}_c which are topologically S^3 's there is a physical integer ambiguity, namely the CS level, which is resolved by the braid diagram.

Having defined the generator T let us now turn to the generator S . In our braid diagrams time flows from left to right vertical slices define the notion of space. The operator T respects this partition into space and time directions since it preserves the pairs of endpoints that appear as initial and final points of the braid. By contrast, the operator S will not respect this partition into space and time and mixes what were originally the initial and final endpoints of the braid. Specifically, our definition of S is to permute the endpoints of the braid as shown in Figure 4.27.



Figure 4.27: The definition of the operator S . In (a) we see a generic braid. In the gray region the strands braid over each other in an arbitrary fashion. In (b) the action of S on this braid.

Given that S creates no new crossings in the diagram, we will not associate the creation of new chiral particles with its action. However, the operator S does have the desired effect of gauging the background $U(1)$. To illustrate this fact consider the comparison of the closure of the trivial braid with the closure of the braid defined by S as shown in Figure 4.28. In the case of the trivial braid, the closure forms a connected unknot. However, in the case of inserting S , the

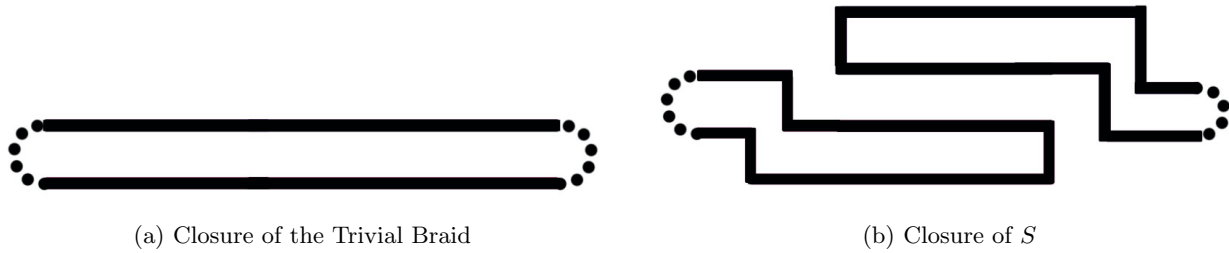


Figure 4.28: The operator S changes the gauging prescription. In (a), the closure of a trivial braid leads to an unknot. In (b) the closure of S leads to two unlinked circles. This changes the topology of \widetilde{M} by increasing b_1 .

closure defines two unlinked circles. In the first case, the cycle γ encircling the two components of the braid, has become contractible at infinity and the associated $U(1)$ is not gauged. Meanwhile in the case of the S braid, γ remains as a homologically non-trivial one-cycle and hence in this theory the $U(1)$ is gauged. Topologically, the compactified double cover geometry has changed to $\widetilde{M}_c \cong S^1 \times S^2$.

From these two definitions of S and T , we may now see that they satisfy the required relations to generate an action of $SL(2, \mathbb{Z})$. This means that S^2 must be a central element whose square is the identity (sometimes written as $S^2 = -1$), and further that $(ST)^3 = 1$. To begin consider the action of S^2 shown in Figure 4.29. As compared to the original braid, the action of S^2




Figure 4.29: The action of S^2 . This acts as time reversal on the braid.

has been to reverse the direction of time flow by changing the initial versus final conditions. Thus, S^2 is simply time reversal along the R-flow and hence acts centrally. Since reversing time twice is the identity operation, we conclude that $S^4 = 1$.

Similarly, we may consider the action of $(ST)^3$ illustrated in Figure 4.30. One can see,

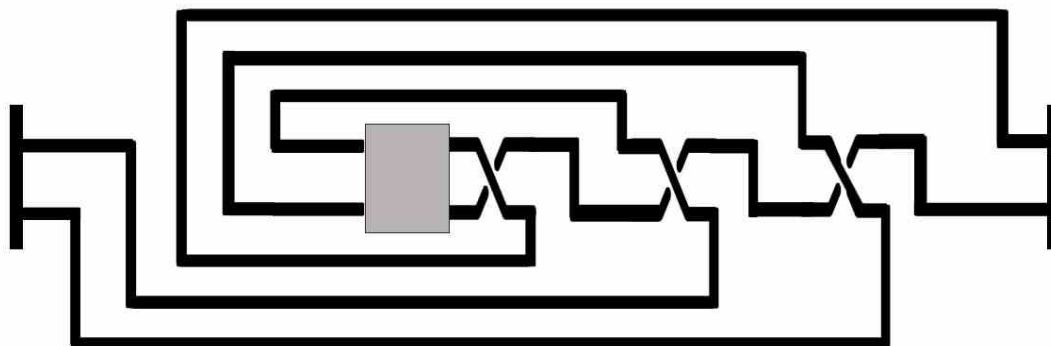


Figure 4.30: The action of $(ST)^3$. On braids this is the identity operator.

manifestly from the above, that the operator $(ST)^3$ acts as the identity on the braid. This completes the verification of the $SL(2, \mathbb{Z})$ group structure.

Given that we have completely specified our choices of boundary data, we may now ask quite generally: what are the possible IR geometries $\widetilde{\mathcal{M}}_c$ which we obtain by these methods? Since the geometry is determined by the resulting closed knot obtained from capping off the braid diagram, we may alternatively ask: what is the set of knots that we can obtain from the trivial braid by repeated action of S and T ? The answer to this question is exactly the set of *rational knots*. They are completely classified by their so-called Conway fraction, z , which is valued in $\mathbb{Q} \cup \{\infty\}$. To define this fraction, we first normalize z by setting its value for the link defined by the closure of S shown in Figure 4.28b to be 0. Then, given any rational knot \mathcal{K}_1 , constructed by action of $\rho \in SL(2, \mathbb{Z})$ from the rational knot \mathcal{K}_2 , we set

$$z(\mathcal{K}_1) = \rho(z(\mathcal{K}_2)). \quad (4.5.5)$$

Where in the above the action of ρ on z is the usual action of $SL(2, \mathbb{Z})$ as fractional linear transformations.

The result of this construction is thus an invariant fraction $z = p/q$ associated to each rational link. We demand that the integers p and q are coprime. Consider two such rational knots with Conway fractions $z_1 = p_1/q_1$ and $z_2 = p_2/q_2$. Then, a theorem due to Schubert asserts that the resulting knots are isotopic (that is equal as knots) if and only if

$$p_1 = p_2, \quad q_1 \equiv q_2^{\pm 1} \pmod{p_i}. \quad (4.5.6)$$

This is exactly the same arithmetic conditions that occur in the classification of lens spaces $L(p, q)$. This is not a coincidence. The double branched cover of the S^3 , branched over the rational knot with Conway fraction p/q , is $L(p, q)$. Thus we recover our original answer. The IR geometries \widetilde{M}_c for the tetrahedron theory are exactly the lens spaces.

Finally, let us note that these methods allow us to fully specify the basic theory associated to closing the tetrahedron braid shown in Figure 4.31a. Indeed, the closed knot shown there is the unknot and thus there is no gauged $U(1)$. However, the chiral particle is still charged under a flavor $U(1)$. To fully specify the resulting theory, it remains to determine the background CS level \hat{k} for this flavor $U(1)$. If \hat{k} is non-vanishing then, upon gauging the background $U(1)$, that is acting with the operator S , we obtain a three-manifold cover \widetilde{M}_c which has $H_1(\widetilde{M}_c, \mathbb{Z}) \cong \mathbb{Z}_{\hat{k}}$. Meanwhile, if $\hat{k} = 0$ then acting with S produces a geometry with non-vanishing first Betti number.

Now, we know that T acts to change the CS level by one unit, and hence the operator $ST^{-\hat{k}}$ must act on Figure 4.31a to produce a cover geometry with $b_1(\widetilde{M}_c) = 1$. However, as is clear from Figure 4.31b, the action of ST on the tetrahedron braid produces topologically two unlinked circles. The double cover of S^3 branched over two unlinked circles is precisely $S^2 \times S^1$ which indeed has $b_1 = 1$. Therefore we conclude that the basic tetrahedron theory defined by the closed knot in Figure 4.31a has CS level $\hat{k} = -1$. This is identical to the definition of the theory given in [16].

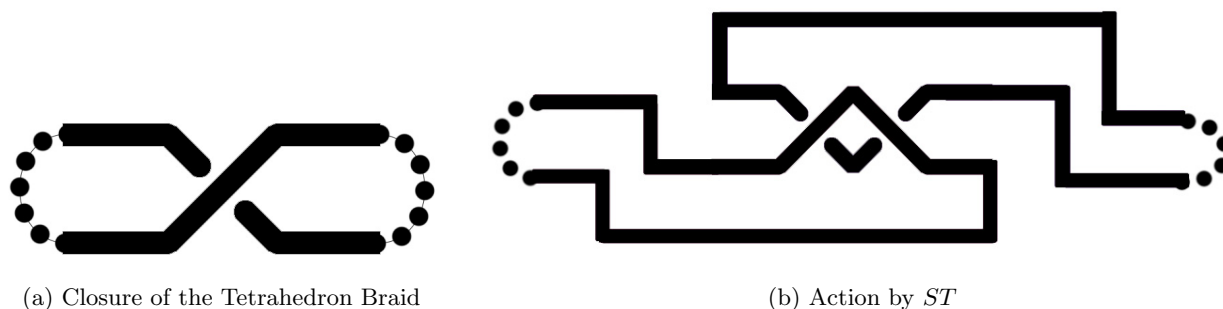


Figure 4.31: Computing the CS level for the background $U(1)$ in the tetrahedron theory. In (a) we see the closure of the basic tetrahedron braid. In (b), the action of ST on this braid changes the topology of the cover.

Doubled Tetrahedron as a Special Lagrangian in \mathbb{C}^3

We have now described a class of IR geometries relevant for the study of the A_1 domain walls. These are special Lagrangians presented as double covers of the tetrahedron, and are given by the Joyce-Harvey flow of the SW geometry $y^2 = x^2 - m$ in \mathbb{C}^3 . It is therefore natural to try and identify these special Lagrangian subspaces of \mathbb{C}^3 more explicitly. As already mentioned, before imposing boundary conditions, the IR special Lagrangian geometry \widetilde{M} is a non-compact solid torus. Furthermore, the geometry supports a unique BPS state described by an M2 brane ending on an S^1 inside the solid torus. In other words, we should be looking for a non-compact special Lagrangian in \mathbb{C}^3 which has the topology of $T^2 \times \mathbb{R}_+$ where at the origin of \mathbb{R}_+ one of the two circles of T^2 shrinks to a point. The M2-brane boundary is then also located at the origin of \mathbb{R}_+ , and is supported on the non-contracted $S^1 \subset T^2$. Precisely such special Lagrangian submanifolds have been constructed by Joyce [74], and figure prominently in the study of open string mirror symmetry [104]. Here we will recall some facts about this class of special Lagrangians.

Let z_i for $i = 1, 2, 3$ denote the three complex coordinates of \mathbb{C}^3 . Then, the special

Lagrangians of interest can be depicted as follows:

$$|z_3|^2 = |z_2|^2 = |z_1|^2 - \frac{m_0}{2\pi}, \quad \theta_1 + \theta_2 + \theta_3 = 0 \quad (4.5.7)$$

where $m_0 > 0$. Another way to characterize this subspace is as the locus where

$$\bar{z}_3 = \frac{z_1 z_2}{|z_1|}, \quad |z_1|^2 = |z_2|^2 + \frac{m_0}{2\pi}. \quad (4.5.8)$$

From the second description, we see that this subspace has the topology of $\mathbb{C} \times S^1$, parameterized by $\{z_2, \theta_1\}$. We can view this as a T^2 fibration over \mathbb{R}_+ , where the torus is made of the angles θ_1, θ_2 , \mathbb{R}_+ is parameterized by $|z_2|$, and at the origin of \mathbb{R}_+ , the θ_2 circle shrinks. The projection of this special Lagrangian on the base of the toric representation of \mathbb{C}^3 , given by $(|z_1|^2, |z_2|^2, |z_3|^2)$ is shown in Figure 4.32. Note that this special Lagrangian supports a unique M2-brane [104], which

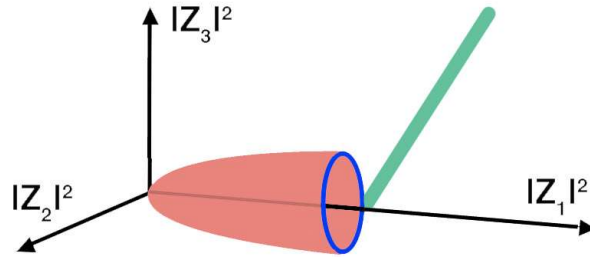


Figure 4.32: Toric special Lagrangian in \mathbb{C}^3 . The green ray denotes the toric projection of the $T^2 \times \mathbb{R}_+$ special Lagrangian. The red disc is the worldvolume of a BPS M2-brane in \mathbb{C}^3 which ends on the special Lagrangian on the blue circle, and gives rise to a BPS particle in $\mathbb{R}^{1,2}$.

ends on the θ_1 circle at the origin of z_1 space. In fact, one can show that if $M \subset \mathbb{C}^3$ is any special Lagrangian submanifold, then the harmonic form λ is

$$\lambda = \sum_i |z_i|^2 d\theta_i|_M, \quad (4.5.9)$$

and from equation (4.5.7) we get

$$\lambda = \frac{m_0}{2\pi} d\theta_1, \quad (4.5.10)$$

and the mass of the corresponding BPS state is m_0 .

This geometry thus has all the characteristics we expect for the special Lagrangian corresponding to the double cover of the tetrahedron, and we conjecture that they are equal. In fact, given our explicit description we can see how the double cover works: it is simply given by the complex conjugation action on \mathbb{C}^3

$$z_i \rightarrow \bar{z}_i. \quad (4.5.11)$$

This is clearly a symmetry of the space defined by (4.5.8). Furthermore, the fixed locus of this geometry are two strands given by

$$(z_2 \in \mathbb{R}, \theta_1 = 0) \quad \text{and} \quad (z_2 \in \mathbb{R}, \theta_1 = \pi). \quad (4.5.12)$$

A further check for the identification of this subspace with the double cover \widetilde{M} , is that if we compactify the theory on S^1 then the moduli space of the special Lagrangian submanifold is given by the mirror geometry defined by a pair of complex variables (u, v) subject to the relation

$$e^u + e^v = 1. \quad (4.5.13)$$

This is exactly the moduli space of $SL(2, \mathbb{C})$ Chern-Simons theory on the tetrahedron. As we will explain in section 7, this is to be expected and demystifies some of the observations in [105], and explains why the partition function of $SL(2, \mathbb{C})$ Chern-Simons on a tetrahedron should be the same as that of the special Lagrangian brane on \mathbb{C}^3 . Furthermore, this shows why the partition function of the $SL(2, \mathbb{C})$ Chern-Simons on a tetrahedron should be that of the open topological string for this special Lagrangian A-brane.

Finally, let us note that the identification of the IR geometry \widetilde{M} as an explicit special Lagrangian in \mathbb{C}^3 yields yet one more way to see the $SL(2, \mathbb{Z})$ action on boundary conditions, and to recover the fact that the compactified geometries \widetilde{M}_c are lens spaces. Specifically, we can consider toric compactifications of the subspace (4.5.8). For example, we can complete the $\mathbb{C} \times S^1$ geometry

to $S^2 \times S^1$, which corresponds to having a locus where θ_2 shrinks, depicted torically in Figure 4.33a. Note that here the special Lagrangian has a modulus corresponding to ‘sliding’ it along the $|z_1|$ axis. Thus in this phase the $U(1)$ is gauged, and coupled to a charged chiral field described by BPS M2-brane. Suppose instead we want to have the geometry of S^3 . This corresponds to shrinking the θ_1 circle, which is depicted in Figure 4.33b. In this case we have no gauged $U(1)$ but we still have a chiral field living on the M5-brane, again described by the M2-brane ending on the special Lagrangian. Similarly we can obtain lens space geometries. For example, $L(p, 1)$ is obtained by having the circle $p[S^1]_1 + [S^1]_2$ shrink at infinity.




Figure 4.33: Toric compactifications of the special Lagrangian. In (a) the compactified geometry is $S^2 \times S^1$ and has a modulus which is described by sliding it along the $|z_1|$ axis. In (b) the compactified geometry is S^3 and the special Lagrangian is rigid.

Black-White Duality, Mirror Symmetry and Geometric Transitions

To summarize the results of the previous sections, we have obtained a class of 3d theories which are described in the IR by a single M5-brane on a lens space $L(p, q)$ together with a single BPS M2-brane charged under a gauged or global $U(1)$ symmetry of the theory. To conclude our discussion of these theories, in this section we discuss simple examples of mirror symmetries.

Let us revisit the basic tetrahedron theory. We equip the resulting knot with a checkerboard coloring as shown in Figure 4.34a. As explained in section 2, the checkerboard provides an

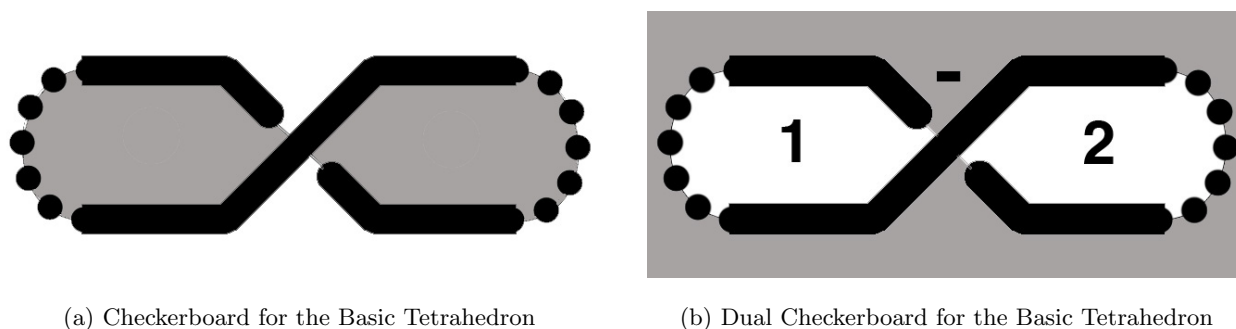


Figure 4.34: Black-White duality for the tetrahedron theory. In (a) we see a description of the theory with no gauge group. In (b) there is a $U(1)$ with $\hat{k} = 1$. The two theories are dual.

algorithmic way to read off the data of the gauge multiplets on $\mathbb{R}^{1,2}$. In Figure 4.34a, we see one white region, and hence no gauge field. However, we may alternatively consider the dual checkerboard for the same knot shown in Figure 4.34b. Now, there are two white regions and hence the gauge group is $U(1)$. Further there is now a crossing connecting the white regions labeled 1 and 2 and correspondingly, the CS level for the $U(1)$ is $\hat{k} = 1$.

Thus, without changing any data about the knot, and hence without changing the field theory, we have found two distinct descriptions of the basic tetrahedron theory:

- A free chiral multiplet coupled to a background flavor $U(1)$ with level $\hat{k} = -1$.
- A chiral multiplet coupled to a gauged $U(1)$ with level $\hat{k} = 1$.

Consistency of our formalism demands that these two descriptions are equivalent, and this is indeed a known mirror symmetry [103].

We can further investigate this basic duality by noting that the second description of the theory involving a gauged $U(1)$ is in fact identical to the action of ST^2 on the first description of the theory. Thus, we can alternatively study this mirror symmetry by acting with the operator ST^2 on the knot in Figure 4.34a. This produces the checkerboard shown in Figure 4.35. Of course, geometrically one can clearly see that the knot defined by Figure 4.35 is equivalent to that of Figure




Figure 4.35: The action of ST^2 on the basic tetrahedron theory. This recovers the black-white duality of the theory as invariance under the operator ST^2 .

4.34b as two of the crossings in the diagram are redundant and can be eliminated. Nevertheless, it is still instructive to see that the algorithmic procedure of extracting the IR field content from the checkerboard produces the correct duality. This is easily verified. The two white regions yield one gauge field. Summing over the crossings connecting the regions with the indicated sign then gives $\hat{k} = 1$ and hence reproduces the black-white mirror symmetry above.

Finally, we can also describe this duality in terms of a geometric transition. We consider the basic tetrahedron theory encoded in Figure 4.34a and ask what happens as the mass m_0 of the chiral particle is smoothly taken through 0 to $-m_0$. As studied in section 3, under this process the strands of the braid reconnect as illustrated in Figure 4.36.

In terms of the braid diagrams used throughout this section, we can describe this reconnection as follows. First, as $m_0 \rightarrow 0$ the braid develops a self-crossing. Then, as m_0 becomes negative the original overcross is exchanged with an undercross. This means that the theory has been acted on by the operator T^2 . Second, the strands reconnect. This changes the identification of endpoints which occurs at infinity. To see this, we compare the topology of the knot obtained by identifying the upper endpoints and the lower endpoints of Figure 4.36a, with the same identification performed in Figure 4.36e. This changes the topology of the knot which is the signature of the operator S . We conclude that the entire reconnection process is described by acting on the

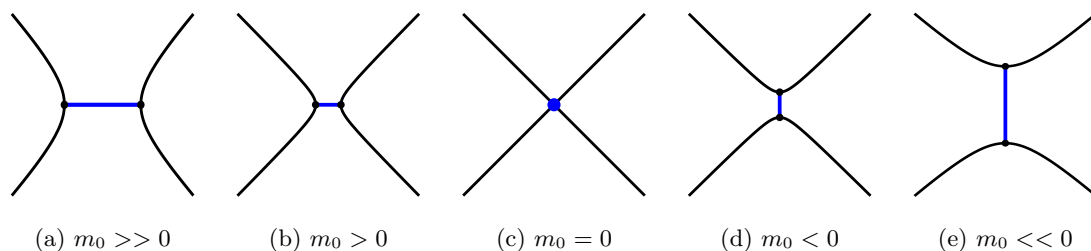


Figure 4.36: The reconnection process. The strands are illustrated in black and the blue line indicates the projection of the boundary of the BPS M2 brane. In (c), when the mass m_0 of the particle vanishes, the two strands touch and their individual identity is ambiguous. As m_0 becomes negative the strands reconnect.

theory with ST^2 , and thus reproduces the black-white duality.

In fact, the above line of reasoning, that is the study of BPS particles with vanishing masses, in some sense explains why it is possible to encode particles in a braid diagram to begin with. The basic point is simply that when the mass is zero the strands must cross, and a braid diagram is simply a resolution of this situation to account for non-zero masses.

4.5.2 A_2 Domain Walls: The Bipiramide

Having investigated the simplest possible example of domain walls in the A_1 model, we now turn to the A_2 theory. This is the simplest 4d theory that exhibits the wall-crossing phenomenon. In one chamber, there are two particles, while in the second chamber there are three particles. This fact has dramatic implications for domain walls. The spectrum in 3d is given by trapped particles from the ambient 4d theory. Thus the different chambers in 4d yield 3d theories with distinct spectra. For each such domain wall we must again specify boundary conditions. A key observation is that the boundary conditions for the two domain walls are related in a non-trivial manner. Thus, fixing a choice of boundary data in one wall, uniquely fixes the boundary data on the other wall,

and hence fully specifies two 3d field theories. As we will illustrate by example in this section, such pairs of 3d field theories, which are connected by 4d wall-crossing, are mirror pairs. Thus, equivalence of the parent 4d theory under wall-crossing explains 3d mirror symmetry.

The Two Chambers and the Pachner Move

We begin our analysis with a discussion of the geometry of the manifold M . The A_2 theory is described by triangulations of a pentagon, and hence this is the front and back face of M as shown in Figure 4.18. In terms of the triangulation on its boundary, the manifold M is therefore a *bipiramid*, that is topologically a solid ball whose boundary is triangulated into six triangles.

As in the discussion of the tetrahedron, as we flow through time, the triangulation on the front face evolves by a sequence of flips to the triangulation of the back face. However because the A_2 theory exhibits wall-crossing, there are now two distinct ways in which the time evolution can occur. One possibility is that in the course of time evolution, the triangulation will undergo two flips, and hence the 4d theory will support two BPS particles. The other possibility is that the flow through time produces three flips, and hence three particles. In each of these cases, a flip encodes a solid tetrahedron and a trapped BPS particle on the wall. The two possible sequences of flips thus describe two distinct ways of decomposing the bipiramid into tetrahedra as illustrated in Figure 4.37.

Above and beyond simply indicating the number of tetrahedra, the sequence of flips on the triangulation completely specifies how the tetrahedra are to be glued together to form the manifold M . Let us illustrate this feature for the case of the bipiramid. We label the triangles in the front and back faces by F_i and B_i , and let I_l, J_l denote the possible triangles appearing in the interior of M for the two chambers respectively. Then, each flip of an edge $E \rightarrow E'$ is associated to four triangles: the two triangles adjacent to E which appear before the flip and the two triangles adjacent to E' which appear after the flip. These give the four sides of each tetrahedron. The




Figure 4.37: The two chambers of the A_2 theory give rise to two ways of decomposing the bipiramid. In the center we see the bipiramid. The pentagon of the A_2 theory is given by the black edges. The red diagonals flip as one flows from the front to the back face. On the left, it is decomposed into two tetrahedra glued along a face. On the right, it is decomposed into three tetrahedra glued along the green edge.

complete sequences of flips then describes all the faces of all of the tetrahedra as illustrated in Figure 4.38.

From this sequence of flips we can then easily extract the tetrahedra and their labelled faces. Thus in the two particle chamber, the tetrahedra are

$$F_2F_3B_3I_1 \quad \text{and} \quad F_1I_1B_1B_2. \quad (4.5.14)$$

While in the three particle chamber they are

$$F_1F_2J_1J_2 \quad \text{and} \quad J_1F_3J_3B_1 \quad \text{and} \quad J_2J_3B_2B_3. \quad (4.5.15)$$

The gluing is then specified by simply identifying the shared faces. As one can easily check, this reproduces the decompositions of the bipiramid shown in Figure 4.37. One can view the entire sequence of flips as giving rise to a 'holographic' view of the 3d geometry by drawing all edges that are flipped as in Figure 4.39.

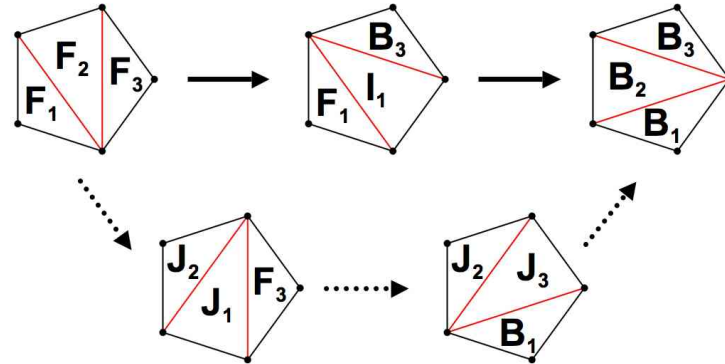


Figure 4.38: The flips of the A_2 theory reinterpreted as gluing data for a bipyramid. In the above, the upper left pentagon with triangles F_i is the front face. The upper-right pentagon with triangles B_i is the back face. As flips happen in time, tetrahedra are created. Following the solid arrows we see two tetrahedra, while following the dashed arrows we see three tetrahedra.

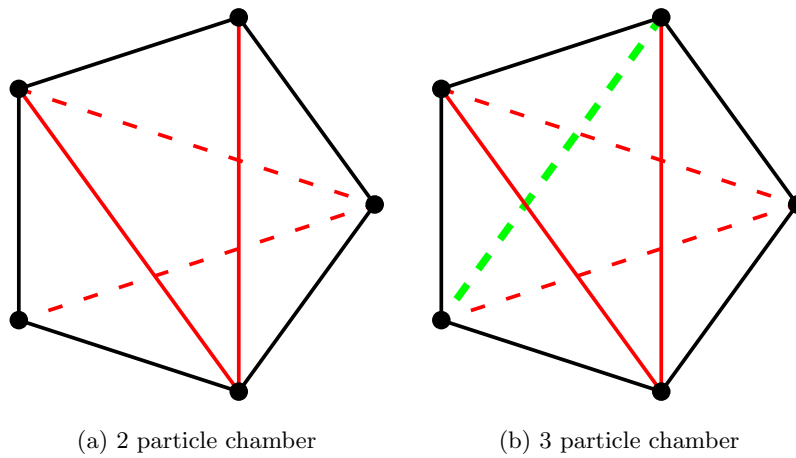


Figure 4.39: The two chambers of the A_2 pentagon model viewed holographically from the front face. The solid red lines are diagonals on the front face, while the dashed red lines are diagonals on the back face. The dashed green line in (b) is the internal edge where the three tetrahedra are glued together.

Quite generally in the study of triangulated three-manifolds, the operation that we have just described where two tetrahedra glued along a face are replaced with three tetrahedra glued along an edge is known as a 2-3 *Pachner mover*. In our physical context, the 2-3 Pachner move is a geometric manifestation of the basic 2-3 wall-crossing of the A_2 theory. As we will discuss later in this section, in generalizing to the A_n model all wall crossing that we encounter is exactly of this 2-3 sort and is thus completely captured in 3d by the Pachner move.

Boundary Conditions and Braids

Now that we have addressed the UV geometry, we turn to the solution of the model as encoded by the IR geometry or equivalently its associated braid. As usual \widetilde{M} is a double cover of the solid ball. The fastest way to understand its topology is to focus on the double cover of its boundary S^2 . This cover is branched over exactly the six points which are the zeros of the SW differential on the front and back face of M . Since a double branched cover of the sphere branched at six points is a Riemann surface of genus two, we conclude that \widetilde{M} must be a smooth three-manifold whose boundary is a surface of genus two.

We can be more specific about \widetilde{M} by making use of the SW curve of the A_2 theory. For the A_2 model, the SW geometry is a double branched cover of the complex x plane described by the equation

$$y^2 = x^3 + ax + b. \quad (4.5.16)$$

Where in the above $a, b \in \mathbb{C}$ are parameters of the theory. This SW curve is a punctured torus, i.e. topologically a torus minus a disc. Then, the IR geometry \widetilde{M} is a thickening of this Riemann surface and is therefore a “torus bottle,” with boundary a genus two Riemann surface as shown in Figure 4.40. This picture allows us to determine the geometry more precisely. Given the boundary $\partial\widetilde{M}$, we choose generating homology classes A_1, A_2, B_1, B_2 with canonical symplectic relations

$$A_i \cdot A_j = B_i \cdot B_j = 0, \quad A_i \cdot B_j = \delta_{ij}. \quad (4.5.17)$$




Figure 4.40: The IR geometry \widetilde{M} for the bipiramide theory. The SW curve is a punctured torus. \widetilde{M} is a thickening of this to a torus bottle. The boundary of \widetilde{M} is composed of the exterior torus less a disc which appears on the outside of the bottle, along with the interior of the bottle which is again a torus less a disc. These are glued together to form a Riemann surface of genus two.

Then, the filling \widetilde{M} is specified at the level of homology by choosing a pair of cycles in $\partial\widetilde{M}$ and declaring them to be contractible in the interior. In our case, as is manifest from Figure 4.40, the cycles which become contractible in the interior of \widetilde{M} are $A_1 - A_2$ and $B_1 - B_2$.

The fact that \widetilde{M} has boundary given by a surface of genus two makes clear the fact that the 3d field theories that we obtain from such domain walls will naturally be acted on by $Sp(4, \mathbb{Z})$. Indeed, to completely specify our theory we must now impose boundary conditions on the manifold \widetilde{M} . This means that we must complete this IR geometry to a manifold without boundary. As \widetilde{M} has boundary a surface of genus two, to remove the boundary we must glue \widetilde{M} to another manifold with boundary a genus two surface. Now to specify the gluing, we choose an element of the mapping class group of $\partial\widetilde{M}$ and glue the boundaries together. On considering the action of this mapping

class group element in homology of $\partial\widetilde{M}$, we obtain the desired action of $Sp(4, \mathbb{Z})$.¹¹

From the point of view of the domain wall construction, the action of $Sp(4, \mathbb{Z})$ that we are describing is physically natural. After the decoupling limit, the 3d wall theory comes equipped with a coupling to two $U(1)$'s, which are the electric and magnetic gauge fields that propagate in the bulk. As in the case of the tetrahedron model, our choice of boundary conditions involves a specification of whether or not these $U(1)$ are gauged and what their background CS level is. Then, $Sp(4, \mathbb{Z})$ acts naturally on this data with the various S transformations inducing gaugings and the T transformations changing CS levels.

To really pin down the IR physics, we now turn to a more detailed description of the geometry as defined by its associated braid. As with our tetrahedral example, the geometric significance of this braid is that it is the branching locus for the double cover. Since each face of the three-manifold M has three triangles, the braid will be composed of three strands. However, since the 4d A_2 theory exhibits wall-crossing there are two distinct braids that we can associate with these domain walls corresponding to the two chambers of the 4d theory. The first has two particles and hence two braid moves, while the second has three particles and hence three braid moves. These are shown in Figure 4.41.

The braids shown in Figure 4.41 have a number of significant properties that demand explanation. To do that we recall from section 4 that we may describe the BPS spectrum by a sequence of mutations on the A_2 quiver. The nodes of this quiver are cycles in the SW curve and hence are associated to a pair of branch points, or equivalently a pair of strands in the braid. Thus for example in the quiver

$$\begin{array}{ccc}
 \bigcirc & \longrightarrow & \bigcirc \\
 u_i & & v_i
 \end{array} \tag{4.5.18}$$

¹¹ Note that the choice of boundary data intrinsically involves the mapping class group as opposed to merely $Sp(4, \mathbb{Z})$. It would be interesting to discover a physical phenomenon which is sensitive to the more refined data of the mapping class group element

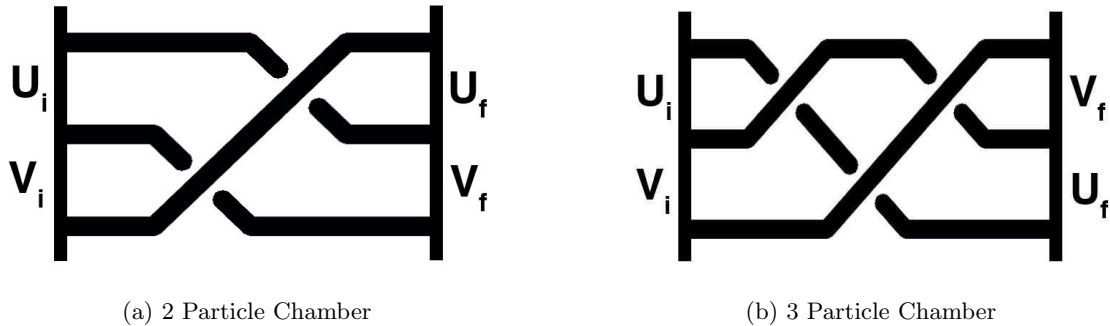


Figure 4.41: The two braids associated to the bipiramid theory. In (a) we see the two particle braid with two moves. In (b) the three particle chamber with three braid moves. The labels u_i v_i denote cycles in the initial and final Riemann surface. Notice that in the case of the three particle chamber the final cycles differ from the initial ones by a permutation.

the node labeled u_i is associated to the cycle defined by the first and second strand as $t \rightarrow -\infty$. Similarly, the node labeled v_i is associated to the cycle defined by the second and third strand as $t \rightarrow -\infty$.

Now, for each mutation in the sequence defining the BPS spectrum in 4d, we obtain a chiral particle in 3d and hence a braid move. In the case of a three strand braid relevant for our current example, we first label the strands as 1,2 ,3 going down the page. Then, the braid group is generated by two elements b_{12} and b_{23} where b_{ij} moves strand i under strand j in time order. To determine which braid move we do, we look at which node of the quiver is being mutated. Specifically:

- If node 1 is mutated do the braid move b_{12} .
- If node 2 is mutated do the braid move b_{23} .

This determines completely how the sequence of mutations describing the BPS spectrum is mapped to the sequence of braid moves encoding the 3d geometry. As described in section 4, in the two

particle chamber the we recall that the mutation sequence is given by node 2 followed by node 1, and yields the braid shown in Figure 4.41a, while in the case of the three particle chamber the mutation sequence is 1, 2, 1 and determines the braid shown in Figure 4.41b.

The fact that the braids are determined by mutation sequences with nodes corresponding to cycles also explains another crucial feature of Figure 4.41. We recall from section 4.4.2, that at the conclusion of a sequence of mutations describing a 4d BPS spectrum the cycles have in general undergone a non-trivial permutation χ . As we saw there, this permutation element depends on the BPS chamber. In the two particle chamber of A_2 , we found that χ is the identity, while in the three particle chamber, χ was the non-trivial element of S_2 . This explains the labeling of cycles that we have made in Figure 4.41. In the two particle chamber the initial basis of cycles denoted u_i, v_i , agrees with the final basis of cycles denoted u_f, v_f . Meanwhile in the three particle chamber the initial and final basis of cycles disagree, having been acted on by the permutation χ .

Duality

Now we are equipped to study how boundary conditions are imposed on the braids, and thus how we can use the result to extract an explicit Lagrangian description of the resulting field theories. As in our analysis of the theory of a single tetrahedron, boundary conditions at the level of the braid are a specification of choices for how the braid is closed into a knot. Then, this knot is the branching locus for the compactified IR geometry \widetilde{M}_c presented as a double cover of S^3 . The most general set of boundary conditions thus involves choosing three pairs of the six endpoints of the braid to glue together. Then, given any fixed gluing prescription, $Sp(4, \mathbb{Z})$ acts to produce another one by performing various S transformations which change the gauging prescription, and T transformations which act as additional braid moves creating ultra-massive BPS particles and changing the CS levels. This generates an interesting family of knots in S^3 which classify in full generality the IR field theories that we obtain from domain walls in the A_2 theory.

Rather than investigate the general case of such knots, we instead note that our construction of these field theories as domain walls, naturally singles out a simple subclass of boundary conditions which respect the order of time flow. Indeed, at each boundary $t = \pm\infty$, there exists a pair of cycles. At $t = -\infty$ these are u_i, v_i , while at $t = +\infty$ these are u_f and v_f . To impose boundary conditions in general, means to choose cycles to be contractible. Doing this in a way which preserves the time ordering implies that we choose one cycle from the initial set and one cycle from the final set and declare them to be contractible. Thus, for example, we may choose u_i and u_f . Then, given such a choice there is an action on such boundary conditions not by the full $Sp(4, \mathbb{Z})$ group, but rather by the subgroup $Sp(2, \mathbb{Z}) \times Sp(2, \mathbb{Z})$ acting on the initial and final trivialized cycles. Explicitly, given $(g_i, g_f) \in Sp(2, \mathbb{Z}) \times Sp(2, \mathbb{Z})$ the action is

$$u_i \rightarrow g_i u_i, \quad u_f \rightarrow g_f u_f. \quad (4.5.19)$$

As with the general action, this induces changes in the gauging prescription and adds various CS levels, but it does so in a way which respects the time order defined by the flow on the geometry.

Let us now see some examples of such boundary conditions applied to the braids of Figure 4.41. We declare that u_i and v_f are contractible and connect the corresponding strands of the braid without introducing additional twists and CS levels. Further, we give a checkerboard coloring to the resulting knots shown in Figure 4.42.

Since we have fully specified boundary conditions, we have now fully specified the compactified IR geometry \widetilde{M}_c and thus we may now read off the field content and determine the resulting 3d field theories obtained for each chamber. Let us consider first the theory determined by Figure 4.42a. We read off gauge structure by making use of the Seifert surface defined by the checkerboard coloring. There are two white regions, labeled 1 and 2 in the Figure. Correspondingly, there is 1 $U(1)$ gauge field in 3d theory on $\mathbb{R}^{1,2}$. The resulting level \hat{k} of this $U(1)$ is vanishing, since the net number of crossings between regions 1 and 2 vanishes. Meanwhile, as we will derive later in this

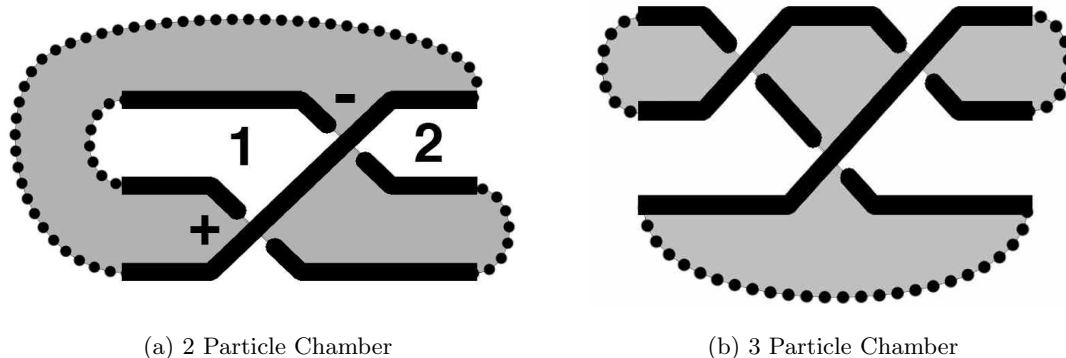


Figure 4.42: Closed links and checkerboards for the gluing determined by trivializing u_i and v_f . In (a), we see the result from the two particle chamber. In (b), the result from the three particle chamber. The theory on the left is $N_f = 1$ SQED. The theory on the right is the XYZ model.

section, the two particles in the theory carry opposite $U(1)$ charges. Thus, the theory encoded by the diagram of Figure 4.42a is exactly $N_f = 1$ SQED.

Similarly, we can read off the IR geometry and field content for the theory encoded by Figure 4.42b. There is now only one white region in the checkerboard and hence there are no gauge fields. The theory supports three BPS chiral multiplets X , Y , and Z encoded by the crossings in the braid diagram. However, this theory has one additional crucial feature. The triangular region of the knot diagram, bounded by the three crossings, is exactly the kind of geometry described in section 2 in which BPS M2-branes yield contributions to the superpotential. This triangular region should be contrasted with other discs with boundary along the knot that are apparent in the diagram. As we have previously explained, dashed regions of the knot encode boundary conditions *at infinity*. Thus, every disc which has its boundary along a dashed component of the knot has infinite volume and hence supports no M2-brane contributions to the superpotential. However, the triangular region in question has all solid boundaries and hence supports a finite disc. Thus, this

theory also has a superpotential coupling its chiral fields as

$$W = XYZ. \quad (4.5.20)$$

So defined, the theory described by Figure 4.42b exactly the so-called XYZ model.

Now we observe a striking fact. The two theories that we have produced via this construction are a mirror pair! Both $N_f = 1$ SQED and the XYZ model have the same IR dynamics near their conformal fixed points [15]. This example illustrates a general phenomenon. The two open braid diagrams of Figure 4.41 admit many different choices of boundary conditions. However, if we fix boundary conditions in the two particle chamber of Figure 4.41a, then those of Figure 4.41b are also fixed automatically by simply demanding that the same cycles are contractible. Thus, fixing *one* choice of boundary conditions determines *two* theories, and the resulting models are always mirror pairs. For another familiar example, we may consider trivializing u_i and u_f which has as a result Figure 4.43.

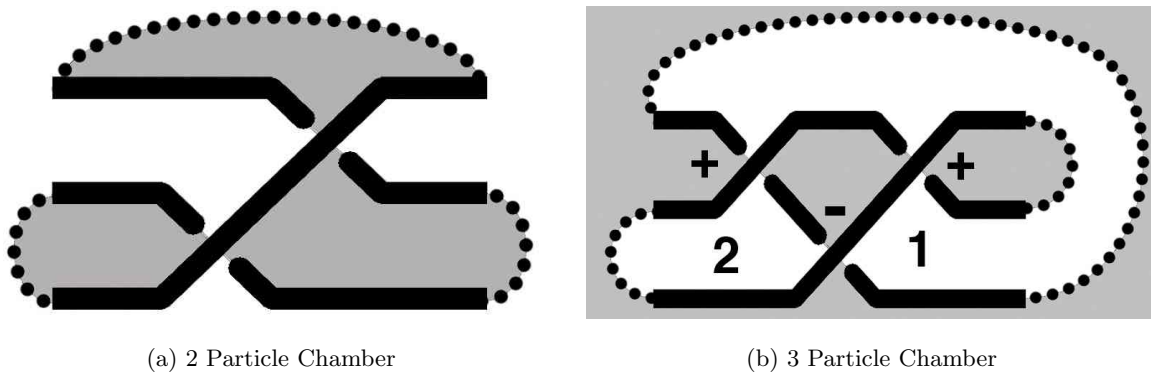


Figure 4.43: Closed links for the gluing determined by trivializing u_i and u_f . In (a), we see the result from the two particle chamber. In (b), the result from the three particle chamber. The theory on the left is that of two neutral chirals. The theory on the right is $\mathcal{N} = 4 U(1)$ Yang-Mills.

What theories are these? The answer is again obtained by a trivial application of the now familiar rules. In Figure 4.43a, we see a theory with no gauge group and two uncharged particles

say X and Y . Meanwhile in Figure 4.43b, we see a theory with a $U(1)$ gauge group, vanishing level, two particles of opposite charge Q and \tilde{Q} and a neutral particle Φ with superpotential coupling

$$W = \Phi Q \tilde{Q}. \quad (4.5.21)$$

The latter theory is thus exactly the $\mathcal{N} = 4$ $U(1)$ gauge theory coupled to a fundamental hypermultiplet. That this theory is mirror to a theory with just two neutral scalars is in fact the paradigmatic example of three-dimensional mirror symmetry [14].

We can also investigate the role of black-white duality in these theories. For example, we study first the case of SQED shown in Figure 4.44. In terms of the physical content of the theories

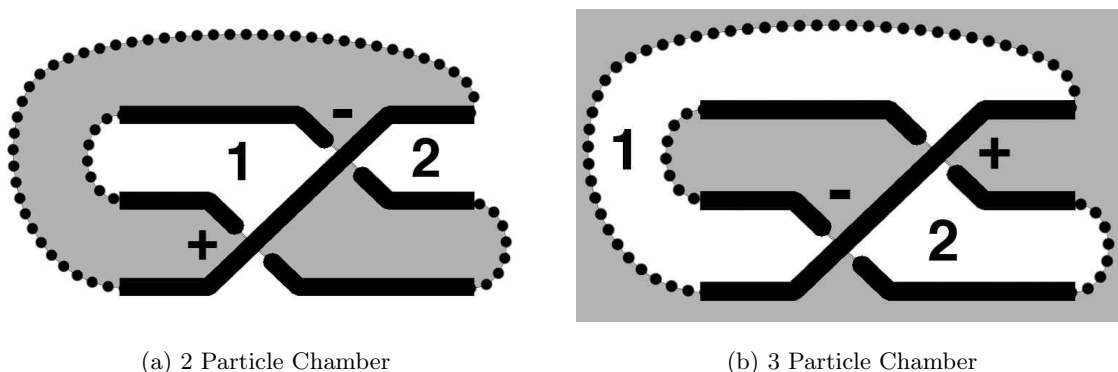


Figure 4.44: Black-White duality for $U(1)$ QED. The theory is self-dual.

defined by the knot diagram, the two theories are identical. Thus under black-white duality, $U(1)$ QED is self-dual

A more interesting case is given by the XYZ model illustrated in Figure 4.45. In Figure 4.45b, there are now four white regions and hence three $U(1)$ gauge fields in the theory on $\mathbb{R}^{1,2}$. If we take the generators to correspond to the regions labeled 1, 2, 3, then the matrix of level is given

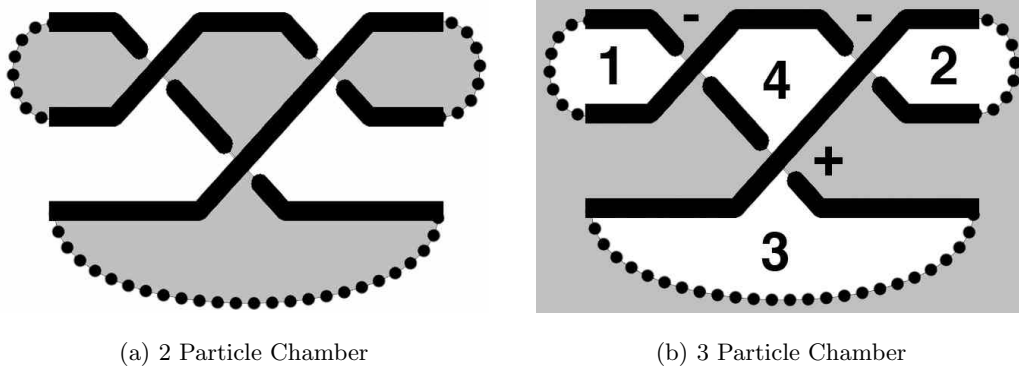


Figure 4.45: Black-White duality for the XYZ model. The theory on the right involves a non-local superpotential which contains the monopole operators.

by

$$\hat{k}_{ij} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & -1 \end{pmatrix}. \quad (4.5.22)$$

There are three charged particles, X_i charged only under $U(1)_i$, with charges $(-1, -1, +1)$. Again the region 4 bounded by undashed components of the knot encodes a superpotential term. However, now we find a novelty. Since region 4 is white, the superpotential also couples to the monopole operators $\mathcal{M}_i = \exp(\sigma_i + i\gamma_i)$ of the corresponding $U(1)$. Thus, in this case

$$W = (\mathcal{M}_1 X_1)(\mathcal{M}_2 X_2)(\mathcal{M}_3 X_3). \quad (4.5.23)$$

Note that this is gauge invariant because \mathcal{M}_i carries a $U(1)_i$ charge \hat{k}_i , due to the Chern-Simons terms. We can see that this is a valid duality by simply invoking the black-white duality for the A_1 theory three times, once for each of the fields X_i , replacing $\mathcal{M}_i X_i$ with the dual field \tilde{X}_i .

4.5.3 General A_n Walls

Having investigated the two most basic examples, we now state our proposal for the general structure of domain walls in the A_n theories.

Five-Brane Geometry

First, there is the UV five-brane geometry M . This is described by a solid ball whose front and back face are triangulations of the $(n + 3)$ -sided polygon. As time evolves, the triangulation of the front face will evolve into the triangulation of the back face by a sequence of flips. Each flip describes a BPS state of the 4d theory and thus each state gives rise to a trapped 3d particle on the wall. On the other hand, each flip naturally describes a solid tetrahedron in M . As in previous sections, the full sequence of flips then encodes a complete tetrahedral decomposition of M . Thus, we have the natural identifications

$$\text{Flip} \leftrightarrow \text{4d BPS particle} \leftrightarrow \text{Trapped 3d BPS particle} \leftrightarrow \text{Tetrahedron.}$$

In particular, the induced tetrahedral decomposition of M completely captures the 3d BPS spectrum of chiral multiplets. In the IR, these are all of the matter particles that are physically relevant.

Just as in the A_2 example, the general A_n 4d field theory can exhibit wall-crossing in its BPS spectrum. This means that there are different chambers of 4d BPS states which in turn describe different possible spectra of trapped BPS states living on the domain wall. According to our discussion above, this implies that the tetrahedral decomposition of the manifold M is not fixed. Rather, distinct chambers are related by the primitive 2-3 wall-crossing where in crossing the wall, a single hypermultiplet disappears from the spectrum. The geometric manifestation of this in M is precisely the 2-3 Pachner move. In one BPS chamber there are three particles encoded by three tetrahedra glued along an edge. In the second BPS chamber one of the particles disappear and the three tetrahedra glued along an edge are replaced by two tetrahedra glued along a face.

Next, we may describe the IR geometry \widetilde{M} and the way in which it encodes the solution of the model. The manifold \widetilde{M} is a branched double cover of the (infinite) solid ball, with branching locus given by the braid determined by the zeros of λ or equivalently the evolving SW differential. At the asymptotic boundary of the ball, there is thus a sphere with $2n + 2$ zeros of λ describing the initial and final terminal points of the $n + 1$ strands in the braid. Thus, the boundary of \widetilde{M} can be described as a double cover of the sphere branched over $2n + 2$ points and is therefore a hyperelliptic Riemann surface of genus n . It follows that \widetilde{M} is a filling in of this Riemann surface to a three-manifold.

In fact, our knowledge of the SW curve of the ambient A_n theory allows us to be more precise and to specify exactly which filling in is prescribed by the time flow. Indeed, the SW curve is given by a certain polynomial of degree $n + 1$ in x as

$$y^2 = P_{n+1}(x). \quad (4.5.24)$$

There are then two cases determined by the parity of n .

- n even:

$\widetilde{\Sigma}$ is a surface of genus $n/2$ which has been made non-compact by removing a single disc. Then, \widetilde{M} is a thickening of this to a genus $n/2$ bottle. If we label the cycles on the initial surface $\widetilde{\Sigma}_i$ on the outside of the bottle as (A_i, B_i) , and those on the final surface $\widetilde{\Sigma}_f$ on the inside of the bottle as $(\tilde{A}_i, \tilde{B}_i)$, for $i = 1, \dots, n/2$, the relations defining the filling in of the boundary surface of genus n to make the manifold \widetilde{M} are

$$A_i = \tilde{A}_i, \quad B_i = \tilde{B}_i. \quad (4.5.25)$$

- n odd:

$\widetilde{\Sigma}$ is a surface of genus $(n - 1)/2$ which has been made non-compact by removing two discs. Then, \widetilde{M} is a thickening of this to a genus $(n - 1)/2$ pipe. Label the symplectically paired

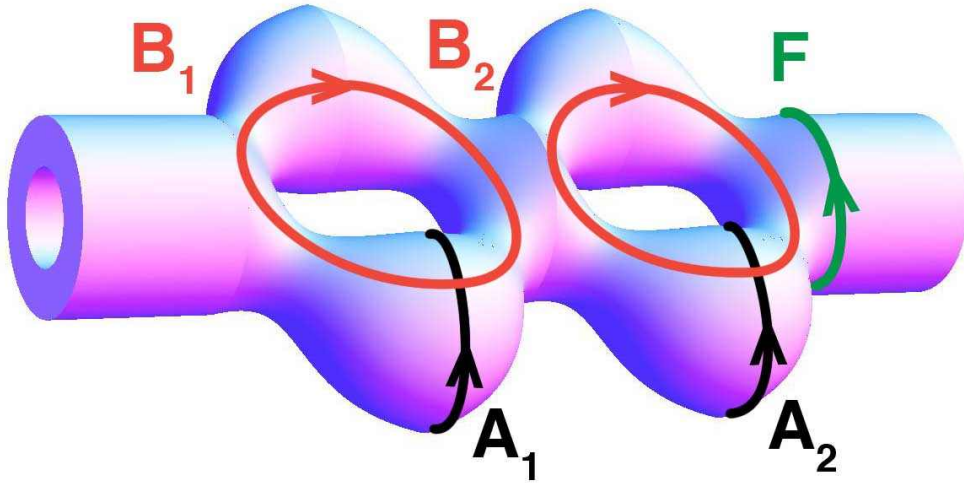


Figure 4.46: A thickened genus two pipe describing the IR geometry \widetilde{M} in the case of A_5 . The SW geometry is a genus two curve with two discs removed. The cycles A_i, B_i, F , span the homology of the SW curve $\widetilde{\Sigma}$ at $t \rightarrow -\infty$. The boundary of \widetilde{M} is a surface of genus five.

cycles on the initial $\widetilde{\Sigma}_i$ on the outside of the pipe as (A_i, B_i) , and those on the final $\widetilde{\Sigma}_f$ on the inside of the pipe as $(\tilde{A}_i, \tilde{B}_i)$, as for $i = 1, \dots, (n-1)/2$. The remaining two cycles on the boundary are unpaired cycles F, \tilde{F} on the initial and final surfaces. The relations defining how \widetilde{M} is filled in are then

$$A_i = \tilde{A}_i, \quad B_i = \tilde{B}_i, \quad F = \tilde{F}. \quad (4.5.26)$$

This geometry is illustrated in Figure 4.46.

Just as in the explicit examples we have studied thus far, to fully specify a 3d theory we must impose boundary conditions. These are defined by taking the manifold \widetilde{M} and gluing it to another copy of itself to determine a compact manifold \widetilde{M}_c with no boundary. As the boundary of \widetilde{M} is a Riemann surface of genus n , there is a natural action of $Sp(2n, \mathbb{Z})$ acting on our choices of boundary conditions and hence the class of theories constructed in this manner. Such actions can again be interpreted as changing the gauging prescription and the CS levels of the model.

A_n Braids and Lagrangians

To determine the detailed structure of the theory, including charges of fields, superpotentials, and CS levels, we proceed as in our examples to construct a braid canonically associated to each chamber. The structure of this braid is completely fixed by R-flow. Indeed, R-flow is specified by the evolution of the 4d central charges which are the periods of the Seiberg-Witten differential, ϕ . Then, given the evolution of these central charges, we can in principle invert the period map to determine the evolution of the loci where $\phi = 0$. The strands in M swept out by these zeros during the flow, are then exactly the strands of the braid. However, even for the simple case of the A_n model, inverting the period map explicitly is a non-trivial task. Nevertheless, for these A_n R-flows, we will see that the structure of the braid, and its detailed 3d physical interpretation, can essentially be determined by simple consistency conditions. Of course, it would still be desirable to invert the period map and verify our results directly.

First, we address how particles in the theory are visible from the braid diagram. As we have seen in our analysis of the A_1 and A_2 examples, before closing the braid (which may involve T transformations), there is a one-to-one correspondence between braid moves and 3d particles. In fact, this correspondence holds generally for those 4d chambers, where all the mutating nodes of the 4d BPS quiver have either, all incoming arrows, i.e. *sinks*, or all outgoing arrows, i.e. *sources*. Note that the A_1 and A_2 examples are both of this type. For more general sequences of mutations which involve nodes which are neither sources nor sinks, what we find is a kind of ‘non-planar’ structure, where each particle corresponds to a specific crossing, but not all crossings correspond to particles.

To begin the investigation, note that BPS states can be viewed as segments connecting a pair of strands in the braid. This observation provides the basic link between particles and braid moves: when the particle becomes *massless*, the associated pair of strands must meet. Thus, if we give the particle a small finite mass, we simply resolve the intersection of the strands into a

braid move. It follows that, up to an overcross/undercross prescription to be determined, each 3d particle will be associated with a braid move.

From this basic fact, we can already deduce why source/sink mutation sequences result in a one-to-one correspondence between 3d particles and braid moves. Indeed, each node of the quiver labels a cycle encircling a pair of strands, and if the corresponding nodes have an arrow between them, then the corresponding pairs share a strand. Suppose we focus on three adjacent nodes of our quiver, which we label α, β, γ . Let us consider the mutation of the node β . If the node β is a sink, i.e., the arrow structure of the quiver is $\alpha \rightarrow \beta \leftarrow \gamma$, then after mutation the quiver has changed to $\alpha \leftarrow -\beta \rightarrow \gamma$. Therefore, up to the ambiguous overcross/undercross, the braid and the associated cycles would appear as in Figure 4.47. Thus, in this case there is one particle and one

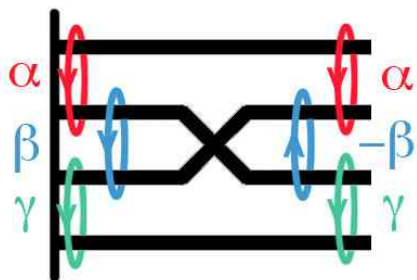


Figure 4.47: The effect of mutation at a sink node β on the strands. The undercross/overcross is ambiguous.

braid move.

Meanwhile, if the node β is a source, i.e. the quiver $\alpha \leftarrow \beta \rightarrow \gamma$, then after mutation the classes of the nodes have changed to $(\alpha + \beta) \rightarrow -\beta \leftarrow (\beta + \gamma)$, and thus the corresponding braid and cycles would appear, up to overcross/undercross, as in Figure 4.48a. This may look like a complicated structure. Indeed, if say the next node to mutate is $(\beta + \gamma)$, this means that the second and fourth strand should cross, which cannot be done without an additional crossing involving the third strand, which does not correspond to a physical particle. To avoid this, we

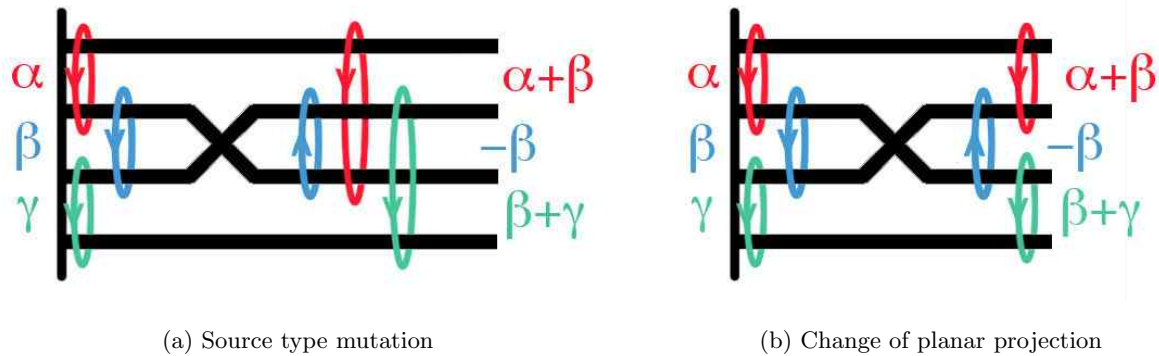


Figure 4.48: A source type mutation on the node β . In (a) we see the resulting change of basis on the cycles. In (b), the planar projection is changed by rotating strand two and three.

rotate the second and third strand after the mutation, resulting in Figure 4.48b. Then, after this change in our planar projection of the braid, the source type mutation looks the same as the sink type. The only difference is that the classes we associate to cycles between nearby strands have changed in correspondence with the labeling of charges on the nodes of the quiver. So again, in this case we see that there is one particle and one braid move.

Finally, consider a mutation on a node which is neither a source nor sink, say mutation of the node β for the quiver $\alpha \rightarrow \beta \rightarrow \gamma$. Then, the mutated quiver would become $\alpha \rightarrow \overrightarrow{-\beta} \rightarrow \beta + \gamma$. The corresponding braid looks, up to overcross/undercross as in Figure 4.49. Suppose next we

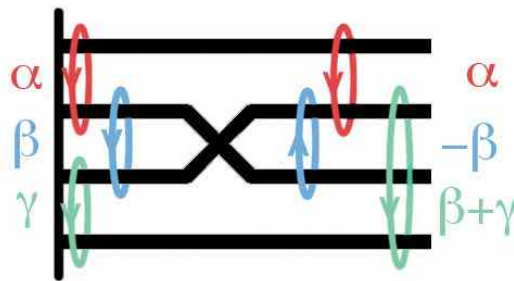


Figure 4.49: The effect of mutation at a mixed node β on the strands.

wish to mutate on $\beta + \gamma$. This cannot be done without an extra crossing. However, unlike the source/sink case where we could change our planar projection to avoid the unnecessary crossing, this is not possible to do by any rotation of the strands after mutation. Continuing with further mutations we will get a ‘non-planar’ braid, for which some crossings will be unphysical in the sense that they do not correspond to chiral particles.

Thus, for precisely those mutation sequences which involve only sources and sinks, we can achieve a planar projection of the braid where each crossing corresponds to a 3d chiral particle. For this reason, we restrict our analysis in the remainder of this section to these sink/source chambers.¹² For these braids, the structure is completely determined by the sequence of mutations of the A_n quiver describing the parent 4d BPS spectrum. We label the strands of the braid as $1, 2, \dots, n + 1$ going down the page. Then, the i -th node of the quiver labels a pair of adjacent strands $(i, i + 1)$, and mutation on the node i corresponds to a crossing involving the pair of strands $(i, i + 1)$. To fully specify the braid we must now fix an overcross/undercross rule. As we will argue later in this section, this rule is determined by consistency to be that the strand $i + 1$ always overcrosses i . We take this as our definition of the braid group generator $b_{i,i+1}$, and thus for source/sink sequences, the braid is completely fixed by:

$$\text{mutation at node } i \leftrightarrow b_{i,i+1}. \tag{4.5.27}$$

Next, we specify boundary conditions for the theory by closing the braid with dashed regions encoding the fact that the closure occurs at infinity. The simplest example of such boundary conditions are those which do not introduce any additional crossings in the diagram. This preserves the feature that all crossings can be associated to finite mass dynamical 3d particles, and we confine ourselves to such simple examples. Then, to extract the IR Lagrangian we draw a checkerboard coloring of the resulting knot. This fixes the gauge multiplet sector of the theory as:

¹²A more general example of a non-planar braid is considered in section 8.6.

- The number of $U(1)$'s is one less than the number of white regions in the coloring.
- The matrix of CS terms is given by computing the Goeritz form for these white regions.

Finally, we must fix the superpotentials and charges of particles. As a consequence of our source/sink assumption, and the simple choice of boundary conditions, each crossing in the diagram corresponds to a particle. Let i index the white regions in the checkerboard. Between regions i and j there are some number α of crossings c_{ij}^α and associated chiral fields Φ_{ij}^α . The field Φ_{ij}^α carries charge ± 1 under the $U(1)_i$ and the $U(1)_j$ and vanishing charges for the remaining gauge groups. We will now determine the sign of these charges by demanding that all apparent superpotential terms are gauge invariant.

To study the superpotential, note that for a given checkerboard coloring of the projected link we have both black and white regions. For each finite region, black or white, whose boundary does not include any dashed portions of the knot that arise from boundary conditions, we expect a superpotential contribution to our theory. This superpotential derives from M2-brane instantons ending on the M5-brane. Thus, each one of the regions corresponds to a superpotential term. However, depending on the color of the region, white or black, the interpretation is different for the gauge theory on $\mathbb{R}^{1,2}$. We discuss each of these in turn.

- *White Regions*

Since white regions R_i are associated to gauge groups of the theory, the i -th finite white region in the checkerboard describes a superpotential which is proportional to the monopole operator for $U(1)_i$

$$\mathcal{M}_i^{\pm 1} = \exp[\pm(\sigma_i + i\gamma_i)], \quad (4.5.28)$$

where γ_i denotes the corresponding scalar dual to photon, and the sign $\epsilon_i = \pm 1$ in the exponent of \mathcal{M}_i depends on the sign conventions for the gauge field yet to be determined. In addition, the M2-brane instanton for white region i will contribute a monomial given by the

chiral fields associated to each of the crossings of that region with other white regions. Thus, each finite white region contributes a term

$$W_i = \mathcal{M}_i^{\epsilon_i} \prod_{j,\alpha} \Phi_{ij}^\alpha \in \mathcal{W}. \quad (4.5.29)$$

- *Black Regions*

For each finite black region B , we also get a superpotential term. But this time there is no associated gauge cycle. Indeed, the one-cycles surrounding the white regions have trivial intersection with any of the black regions (including the neighboring ones), and this implies that the two-cycles defined by the black regions carry no monopole charge. Thus, for each finite black region we simply get the contribution of the fields at the crossings on the boundary of the region

$$W_B = \prod_{(ij,\alpha) \in \partial B} \Phi_{ij}^\alpha \in \mathcal{W}. \quad (4.5.30)$$

Now, we fix the charges of the fields Φ_{ij}^α by demanding that the superpotential terms (4.5.29) and (4.5.30) are gauge invariant. Let $q_{k,ij}^\alpha$ denote the $U(1)_k$ charge of the field Φ_{ij}^α corresponding to the crossing c_{ij}^α . Each of these fields corresponds to a basic crossing, and hence is $q_{k,ij}^\alpha = \pm 1$ for $k = i, j$ and $q_{k,ij}^\alpha = 0$ for $k \neq i, j$. Note also that the monopole field $\mathcal{M}_i^{\epsilon_i}$ carries $U(1)_j$ charge given by $\epsilon_i \hat{k}_{ij}$ induced from the CS term.

Then, the $U(1)_j$ gauge invariance of the white region contribution W_i implies that

$$\epsilon_i \hat{k}_{ij} + \sum_{\alpha} q_{j,ij}^\alpha = 0, \quad (4.5.31)$$

and similarly, the $U(1)_i$ invariance of W_i implies that

$$\epsilon_i \hat{k}_{ii} + \sum_{\alpha,j} q_{i,ij}^\alpha = 0. \quad (4.5.32)$$

Using the $U(1)_i$ invariance of W_j the latter equation we learn

$$\epsilon_i \hat{k}_{ii} - \sum_j \epsilon_j \hat{k}_{ji} = 0. \quad (4.5.33)$$

And finally, multiplying this by ϵ_i we see that

$$\hat{k}_{ii} - \sum_j \epsilon_i \epsilon_j \hat{k}_{ij} = 0 \quad (4.5.34)$$

This is compatible with the definition of the CS matrix (4.2.39) only if, for each pair of i, j with non-vanishing \hat{k}_{ij} , we have

$$\epsilon_i \epsilon_j = -1. \quad (4.5.35)$$

This implies that it must be possible for this class of gauge theories to assign a parity to each $U(1)$ factor of the gauge group, defined by the sign ϵ_i , such that gauge fields which have non-vanishing \hat{k}_{ij} have opposite parities. As a result, we learn that, after deleting a single ungauged node, the Tait graph, defined by the checkerboard coloring as in section 2, is a bipartite graph for which we can assign opposite ± 1 to vertices which are connected. This turns out to be true for all the graphs which arise in our constructions for the sink-source sequence of mutations.

Furthermore, note that equations (4.5.31) and (4.2.39) can be combined to express the charges of the fields in terms of the sign $\zeta(c_{ij}^\alpha)$ associated to the crossing

$$\sum_\alpha (\epsilon_i \zeta(c_{ij}^\alpha) + q_{i,ij}^\alpha) = 0. \quad (4.5.36)$$

This suggests the canonical solution

$$q_{i,ij}^\alpha = -\epsilon_i \zeta(c_{ij}^\alpha) \quad (4.5.37)$$

Equation (4.5.37) is the key final result which specifies the charges of the theory and completes our description of these models. Together with the fact that $\epsilon_i \epsilon_j = -1$ for connected vertices, it implies that the each of fields Φ_{ij}^α , charged under white regions i and j , are bifundamentals which carry opposite charges under $U(1)_i$ and $U(1)_j$

$$q_{i,ij}^\alpha = -q_{j,ij}^\alpha. \quad (4.5.38)$$

This means that each of the links in the Tait graph, which corresponds to a crossing c_{ij}^α and hence a field Φ_{ij}^α , can be oriented by making use of the bipartite structure. If we make the convention

that the link c_{ij}^α points out of the node associated to $U(1)_i$ if the field Φ_{ij}^α carries charge +1 under $U(1)_i$, then this makes the Tait graph into the quiver for the resulting gauge theory.

Finally, we observe that these equations uniquely fix the charges in terms of the Chern-Simons levels, up to the choice of ϵ_i . However, there are only two global choices of ϵ_i depending on which nodes we assign as even and which one as odd. A change of an overall sign of ϵ_i simply flips the overall sign of the charges, which gives an equivalent theory, by replacing all gauge fields by their opposites, $A_i \rightarrow -A_i$ (which does not affect the Chern-Simons level matrix). The reverse is also true: if we assign arrows to the links of the Tait diagram, thus fixing $q_{k,ij}^\alpha$, we can read off the $\zeta(c_{ij}^\alpha)$ from equation (4.5.37) and hence determine the associated overcross/undercross. This provides a strong consistency check on our proposal for the charges of the fields, and our identification of mutations with the basic braid move $b_{i,i+1}$.

Cookbook

Let us summarize the rules derived in the preceding section into a recipe for extracting the 3d theory. We confine our description to the simplest examples where the boundary conditions respect the order of time flow, and no additional T transformations are performed.

- Given an A_n BPS quiver, identify a source-sink sequence of mutations describing a chamber of the 4d theory.
- Construct a braid on $n + 1$ strands by reinterpreting the mutation sequence as a sequence of braid moves. When the node i is mutated do the braid move $b_{i,i+1}$.
- Impose boundary conditions by choosing cycles to be contractible. If n is odd, this means contracting $\frac{n+1}{2}$ initial and final cycles. If n is even, this means contracting $\frac{n}{2}$ initial and final cycles.
- Equip the resulting knot with a checkerboard coloring and draw the associated Tait graph.

This graph is bipartite except for the presence of one auxiliary framing node. By framing node, we mean the node that is not gauged. For each other node i in the graph, assign a parity $\epsilon_i = \pm 1$ in such a way that nodes connected by a link have opposite parity.

- Orient the links in the Tait graph by making use of the parity of the nodes and the parity of the links. Specifically:
 - If both nodes are not the framing node and the link has parity $+1$ orient the link by pointing it from the node with parity -1 to the node with parity $+1$.
 - If both nodes are not the framing node and the link has parity -1 orient the link by pointing it from the node with parity $+1$ to the node with parity -1 .
 - If one node is the framing node, orient the link by having it point out of the framing node if the product of ϵ_i and the link orientation is $+1$ and point into the framing node if the product of ϵ_i and the link orientation is -1 .
- The oriented Tait graph can now be interpreted as the quiver describing the field content and gauge group of the resulting theory on $\mathbb{R}^{1,2}$. Thus, each node other than the framing node yields a $U(1)$. Each oriented link defines a bifundamental field. And the matrix of CS levels \hat{k}_{ij} is determined by computing the Goeritz form the sign \pm assigned to each of the links in the Tait graph.
- Finally, the superpotential of the theory is given by summing over contributions from finite white and black regions in the checkerboard

$$\mathcal{W} = \mathcal{W}_{\text{Black}} + \mathcal{W}_{\text{White}}. \quad (4.5.39)$$

At the level of the Tait graph this means the following:

- For each finite black region B , we obtain a contribution to \mathcal{W} in the form of a monomial in elementary fields. Specifically, each such region B defines a cell in the Tait graph,

and we add the cycle in the Tait graph defined by ∂B to the superpotential

$$\mathcal{W}_{\text{Black}} = \sum_B \prod_{\partial B} \Phi_{ij}^\alpha. \quad (4.5.40)$$

- For each finite white region W , we add an associated term in the superpotential involving the monopole operator \mathcal{M}_i associated to the i -th gauged node which corresponds to that white region. Specifically, we take the product over fields charged under the node in question

$$\mathcal{W}_{\text{White}} = \sum_W \mathcal{M}_i^{\epsilon_i} \prod_i \Phi_{ij}^\alpha. \quad (4.5.41)$$

The algorithm defined above can be applied to any source-sink mutated chamber of the 4d A_n theories. To compare the theories defined by two distinct chambers, we keep track permutation χ which acts on the nodes of the quiver. Let χ_i for $i = 1, 2$ denote the two permutations. We impose boundary conditions at $t = +\infty$ on a given braid by contacting some set of cycles γ_i . Then, to compare to the second braid we contract the cycles $\chi_2 \circ \chi_1^{-1}(\gamma_i)$. Thus, one choice of boundary conditions, fixed for the braid defined by one chamber, determines boundary conditions for the braids defined by all other chambers. Extracting the physics from the resulting knot as above we obtain a class of mirror 3d theories. In the next section, we will use the procedure to give new examples of dual pairs.

4.5.4 A Final Example: Alternating A_{2n}

As an example application of these rules we will consider domain wall theories in the general A_{2n} model. We consider an alternating orientation of the quiver.

$$\begin{array}{ccccccc}
 \bigcirc & \longrightarrow & \bigcirc & \longleftarrow & \bigcirc & \longrightarrow & \bigcirc & \longleftarrow & \cdots & \longleftarrow & \bigcirc & \longrightarrow & \bigcirc \\
 1 & & \dot{1} & & 2 & & \dot{2} & & \cdots & & n & & \dot{n}
 \end{array} \quad (4.5.42)$$

This quiver corresponds to the zig-zag triangulation of an $(2n + 3)$ -gon shown in Figure 4.50a. This is the triangulation present on the front and back face of the three-manifold, and as usual,

its evolution determines a decomposition of the solid ball into tetrahedra viewed holographically in Figure 4.50b.

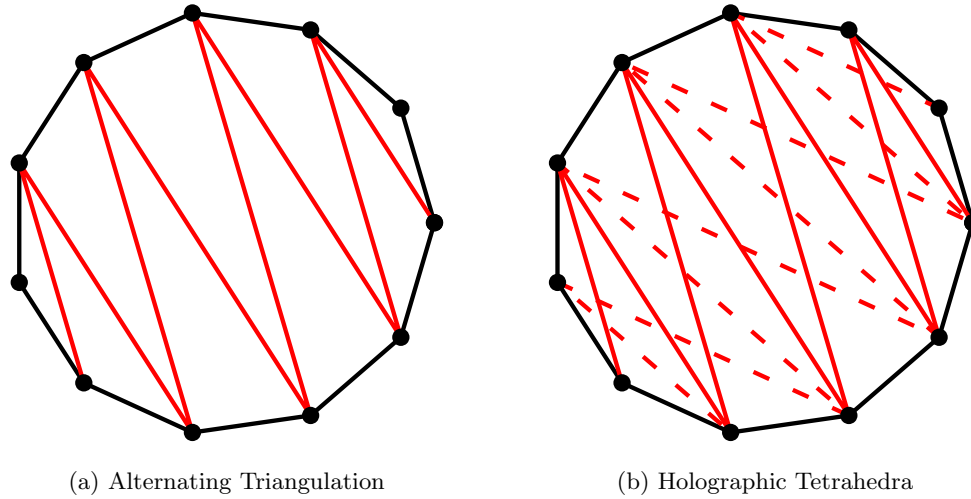


Figure 4.50: Triangulations for the alternating A_n quiver in the case of A_8 . In (a) we see the triangulation of the front face. In (b) the holographic view of the tetrahedra in the case of the minimal chamber.

There are two simple chambers of these theories described by their mutation sequences as

- *Minimal Chamber*

There are $2n$ states. The mutation sequence proceeds by first mutating on all dotted nodes, and then mutating on all undotted nodes. The associated permutation element χ is the identity.

- *Maximal Chamber*

The theory has $n(2n + 1)$ states. The mutation sequence proceeds by mutating on all all undotted nodes, then all dotted nodes, then all undotted nodes, etc. for a total of $2n^2 + n$ mutations. The associated permutation element is

$$\chi = (1, \dot{n}) (\dot{1}, n) (2, n - \dot{1}) (\dot{2}, n - 1) \cdots . \quad (4.5.43)$$

We construct the braid by identifying mutations with braid moves as described in the previous section: when an undotted node m is mutated we do the braid move $b_{2m-1,2m}$, when a dotted node \dot{m} is mutated we do the braid move $b_{2\dot{m},2\dot{m}+1}$. This leads to braids of the form shown in Figure 4.51.

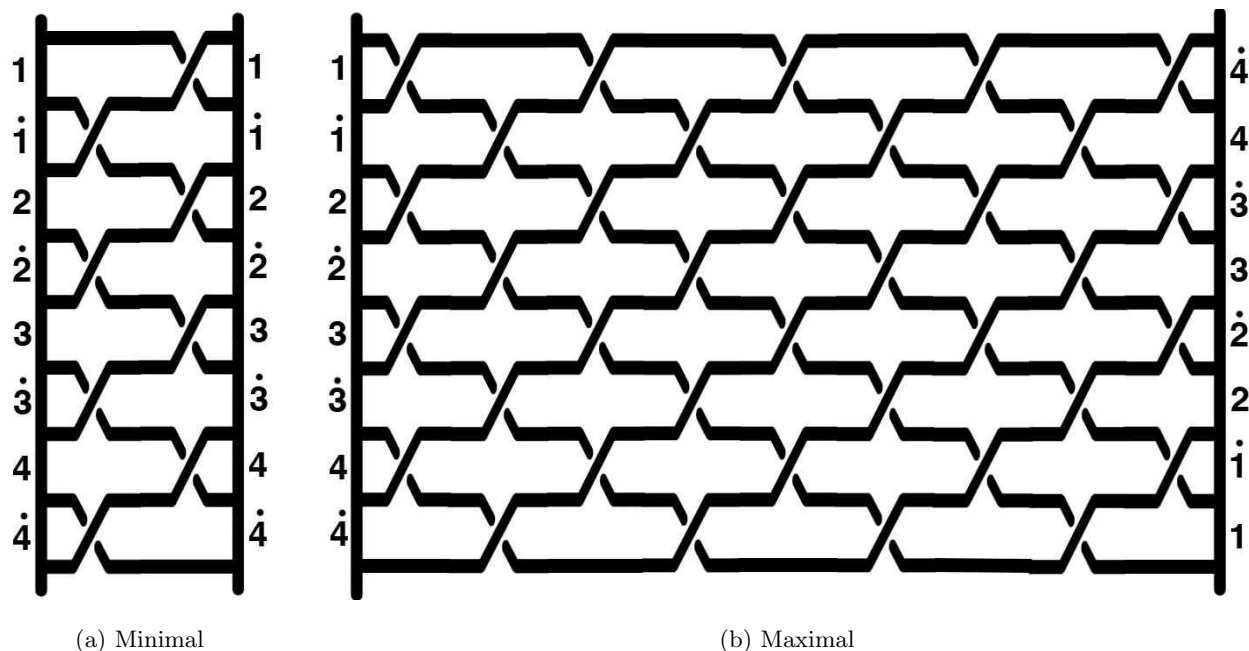


Figure 4.51: Minimal and maximal braids for the A_8 alternating quiver.

Next, to determine a 3d field theory we impose boundary conditions. As an illustrative example, we choose to trivialize the undotted cycles at the initial time, and the dotted cycles at the final time. Of course in doing so, we must also take into account the non-trivial permutation χ in the maximal chamber.

Minimal Chamber

First, we investigate the physics of the minimal chamber. We follow the general instructions of the previous section. We draw a checkerboard coloring of the resulting knot, and its associated Tait graph. Then, we identify a framing node in the graph which will be ungauged. All

Next, we consider the maximal chamber. The checkerboard coloring of the knot, and its associated Tait graph are shown in Figure 4.53. From the graph, we read off that there are $n(n-1)$ gauge groups, and $2n^2 + n$ particle, $2n$ of which are gauge neutral and encoded in the links connecting the framed node to itself.¹³ Moreover, there is now a superpotential consisting of:

- Black Terms.

There are n^2 finite black regions. Of these, $n^2 - n$ yield quartic monomials in \mathcal{W} , and n yield cubic contributions. These can be off from from cells in the Tait graph.

- White Terms.

All white regions are finite and hence yield monopole contributions to the superpotential.

There are $n(n-1)$ such contributions.

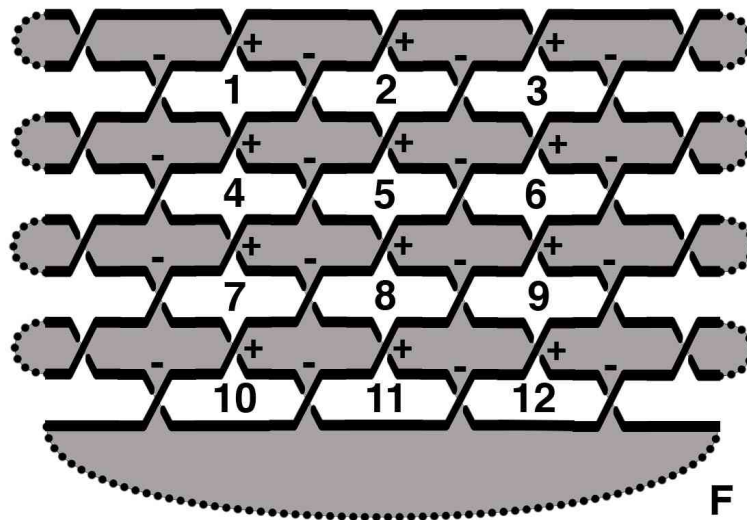
The remaining data in the model, such as the charges of the fields and the CS levels, are all encoded by the Tait graph.

In section 8 we will check this proposed duality by comparing partition functions of these two theories.

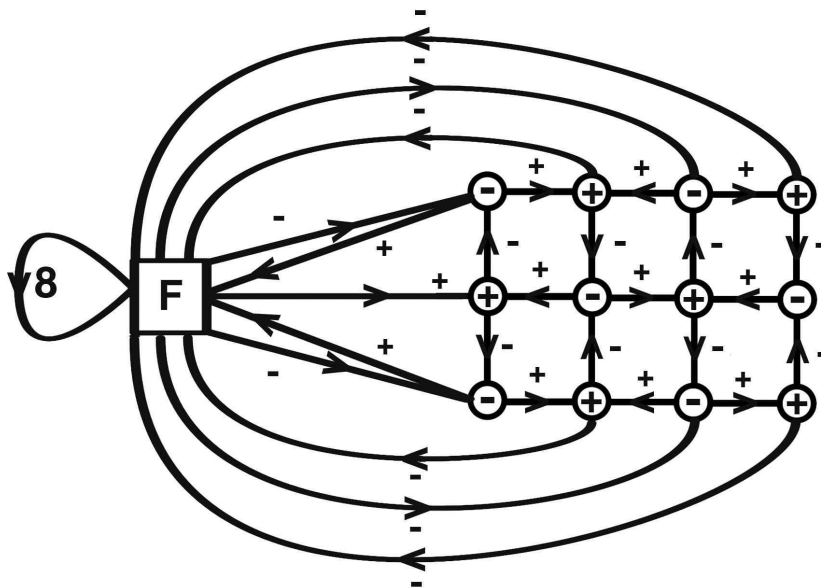
4.6 Flows of General 4d $\mathcal{N} = 2$ Theories

The examples described in the previous section, illustrate domain walls in the simplest possible context of the A_n Argyres-Douglas models. However, the general procedure of extracting a 3d theory from an R-flow of a parent 4d theory can be carried out for an arbitrary $\mathcal{N} = 2$ model. For example, the E_n case which does not correspond to 3d geometry will be discussed in section 8. One could perhaps also consider the R-flow of other $\mathcal{N} = 2$ theories which are not complete, by relaxing the constraint of UV finiteness, though we will not provide examples of that in this

¹³ The neutral links correspond to crossing connecting the framed region with itself. These do not contribute to the CS levels and hence the links do not have an associated sign.



(a) Checkerboard for the maximal chamber of alternating A_8



(b) Tait graph for the maximal chamber of alternating A_8

Figure 4.53: Checkerboard coloring and associated Tait graph for the maximal chamber of the alternating A_8 theory. The white region denoted by F corresponds to the square framed node in the graph. The generalization to A_{2n} is an $n(n - 1)$ grid Tait graph.

paper. Here, we will focus on the case where 4d gauge theory is defined by wrapping a pair of M5 branes on a punctured Riemann surface C of arbitrary genus $g(C)$. Such 4d models have a number of interesting geometric features which translate into properties of the resulting three-manifold M which is given as a thickening of C . For example, if we consider the punctures of the surface C there are two basic types [13, 36]:

- *Irregular Punctures*

These are equivalent to boundary components of the Riemann surface. For example, the A_n Argyres-Douglas model is equipped with such a puncture. In three dimensions, the boundary data for these punctures is fixed for all time along the R-flow and hence these boundary components are identified on the front and back face of M . Thus, in three dimensions, irregular punctures do not give rise to boundary components, but instead map to pure topology of M .

- *Regular Punctures*

These encode mass parameters of the 4d theory and hence describe first order poles in the SW differential. If we consider a one-parameter family of such punctures then we obtain a line of cusp singularities in \widetilde{M}

Topologically, the manifold M is given by a thickening of C modulo the relation that the boundary components of C , defined by the irregular punctures, are identified for all time. It has annular cusp singularities for each regular puncture. Further, if C has at least one boundary component, then ∂M is a connected Riemann surface obtained by gluing two copies of C along their common boundary. Specifically, if C has $b \geq 1$ boundary components, then the genus of ∂M is determined by a simple computation to be

$$g(\partial M) = 2g(C) + b - 1. \quad (4.6.1)$$

The manifold M is then a certain filling in of the boundary ∂M .

Given any such surface C , its BPS data may be encoded in an ideal triangulation as described in section 4. As we flow through time, the triangulation of the front face will evolve to the triangulation of the back face and this determines a decomposition of M into tetrahedra. Each tetrahedron encodes a 3d chiral particle in the theory, and finite 2-3 wall-crossings describe 2-3 Pachner moves on the 3d triangulation.

Next, we state some general facts about the resulting IR geometries \widetilde{M} . These are branched double covers of M and their structure is encoded by the evolving zeros of the Seiberg-Witten differential ϕ . For each triangle in the front face of the triangulation we obtain a zero of ϕ . As the zeros evolve, they determine an open knot composed of strands, whose endpoints are fixed at the front and back face of M . In principle, we can find the geometry of these strands using the R-flow. Indeed, given the evolution of the 4d central charges, which are the periods of ϕ , we can invert the period map and find the geometry of the branch point flow. This is quite similar to the case of the A_n models studied in detail in section 5. However, unlike the the examples there, where these strands moved and were tangled in a space with trivial topology of a ball B^3 , now the strands evolve in a space M with non-trivial topology which has as boundary the surface of higher genus (4.6.1). Further, the strands may also become braided around the annular cusp singularities descending from the regular punctures in the surface C .

Thus, the result of a general R-flow on a punctured Riemann surface is a potentially complicated topological configuration, and some of the technology that we developed for the A_n case will be need to be enhanced to study this situation. Nevertheless, we can still see that some of our general observations hold. For example, to impose boundary conditions on the resulting theory, we close the open knot in M into an honest knot and this fixes the compactified geometry \widetilde{M}_c . From this description, it is also clear that the resulting 3d theories will be acted on by $Sp(2g(\partial M), \mathbb{Z})$ and, as in our general discussion, this action is physically realized by changing CS levels, and gauging or ungauging some $U(1)$'s.

4.6.1 Effective 3d Gauge Theories with Infinite Dimensional Representations

A major novel feature of the general flows outlined above is the presence of BPS chambers of the 4d theory with infinitely many BPS states. Indeed, the main examples we have considered involve 3d theories whose 4d parents have finite number of BPS states. However this is not the typical situation. For example, the pure $SU(2)$ gauge theory in 4d has infinitely many BPS states in the weak coupling chamber. For $SU(n)$ theories, not only can we have infinitely many BPS states, but in addition, we may have chambers which support BPS states with arbitrarily high spin. It is thus natural to ask: what would the interpretation of the reductions of such chambers to 3d, and their equivalence to chambers with finitely many states imply?

To gain some insight to what implications these chambers and dualities may have in 3d, let us consider the example of pure $SU(2)$ gauge theory in 4d. This theory has two chambers. In the strong coupling regime, we have two states given by (electric, magnetic) charges $(2, -1)$ and $(0, 1)$, and in three dimensions this R-flows to a 3d theory with two chiral multiplets much as in our analysis of the A_n models. Meanwhile, in the weak coupling chamber of 4d, there are infinitely many BPS particles: the monopole and its dyonic descendants, with charges $(2n, 1)$, and the vector W-boson which carries charge $(2, 0)$. Consider the R-flow of the weak coupling chamber, where we take the projection to be along the electric charge direction. In this way, all the 4d dyons will have the same real projection defining the 3d supersymmetry, and hence all the trapped 3d dyons will have equal finite masses. In addition, the 4d BPS W-boson will result in a *massless* trapped 3d particle.

From the fact that these 3d particles arise from trapped 4d bulk fields, we can make a number of observations. First of all, the W-boson must carry vector quantum numbers. Therefore, we in fact have a massless 3d vector particle. There are only two possibilities for a non-abelian 3-vector theories: either we have an $SU(2)$ gauge theory in 3d, or an $SL(2, \mathbb{R})$ gauge theory in 3d. The first option may look more natural from the 4d perspective, where in the infinitely weak

coupling regime we recover an $SU(2)$ gauge theory. However, we believe the $SL(2, \mathbb{R})$ is what is realized in 3d, for the following reason: the infinite number of trapped 3d dyons have the same mass, and this strongly suggests that they form one irreducible object. Furthermore, note that the W-boson from the bulk can bind to any of these trapped dyons, transforming one to the other. Since the vector particles should form either $SU(2)$ or $SL(2, \mathbb{R})$, and since $SU(2)$ has no infinite dimensional unitary representations with finite Casimir, we conclude that we must have an $SL(2, \mathbb{R})$ theory and that the 3d dyons form a single irreducible representation of $SL(2, \mathbb{R})$ as illustrated in Figure 4.54. Note also that if we tilt the angle of projection to 3d, so that the W-boson has a tiny mass ϵ , the infinite tower of 3d dyonic states will have BPS masses given by $|m + n\epsilon|$ where m is the real mass and $n \in \mathbb{Z}$. This can be interpreted as a deformation to the Coulomb branch of the $SL(2, \mathbb{R})$ theory by giving an expectation value to the adjoint scalar field in the $SL(2, \mathbb{R})$ gauge multiplet.

We can further argue why we may have obtained an $SL(2, \mathbb{R})$ instead of $SU(2)$ by observing that the main difference between these two cases is the sign of the kinetic term for the W-bosons. How could the sign of the kinetic term for the W have flipped? This actually has a simple explanation: the W-boson can never become massless in 4d. No matter how weak we make the 4d coupling, as we come close to making the W massless (by taking the 4d scalar vev to zero) we cross the curve of marginal stability, rendering the W unstable. However, if we did analytically continue to the region where the W-boson is unstable, it is known that the kinetic term for the W will flip sign [106, 107], which is the signature of an $SL(2, \mathbb{R})$ gauge theory. Of course, the CS level must be non-zero, otherwise we would end up with a non-unitary theory, and the existence of CS terms would render the gauge particles massive in the IR and make their wrong sign kinetic term irrelevant.

As in our general discussion at the beginning of this section, the precise theory we get from these R-flows of the weak and strong coupling chambers of $SU(2)$, depends on the boundary

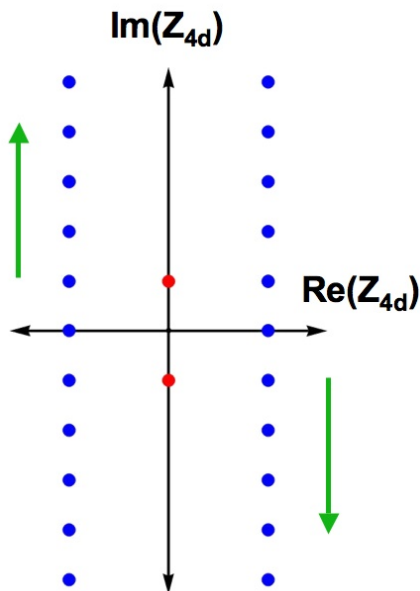


Figure 4.54: BPS spectrum and R-flow for weak coupling $SU(2)$. The blue dots denote the 4d dyons, the red dots the 4d W-bosons. The green arrows specify the direction of R-flow. The 3d trapped W-bosons are massless and give rise to an $Sl(2, \mathbb{R})$ gauge symmetry. The 3d trapped dyons all have equal mass and form an infinite dimensional representation of $Sl(2, \mathbb{R})$.

conditions. However, what is clear is that this construction produces a 3d theory with infinitely many chiral fields, corresponding to the weak coupling chamber, which is mirror to a theory with only two chiral fields, corresponding to the strong coupling chamber. As for the higher $SU(N)$ gauge theories in 4d, the chambers which support higher spin BPS states, can lead under R-flow to 3d theories with massless (or nearly massless) particles of higher spin. One might speculate that this suggests a 3d structure of a higher spin gauge theory. These are clearly exciting possibilities, and are the subject of active investigation.

4.6.2 Accumulating Tetrahedra

One feature of the resulting 3d geometry which we can see directly from the correspondence between tetrahedra and 3d BPS particles, is that in infinite chambers of the 4d theory, the UV three-manifold M will be partitioned into *infinitely* many tetrahedra. Further, at the accumulation rays in the 4d BPS spectrum which describe vector multiplets, the resulting tetrahedra will also accumulate.

As an explicit example, we can consider the case of pure $SU(2)$ described above. Then, the 4d $\mathcal{N} = 2$ Riemann surface is given by an annulus with one marked point on each boundary. The three-manifold M obtained from a flow of this data is therefore a filling of two annuli glued along their common boundary. A trivial application of the genus formula (4.6.1) shows that ∂M is a torus, and hence M is a solid torus. The front and back face of this torus are equipped with triangulations, each with two triangles as shown in Figure 4.55.

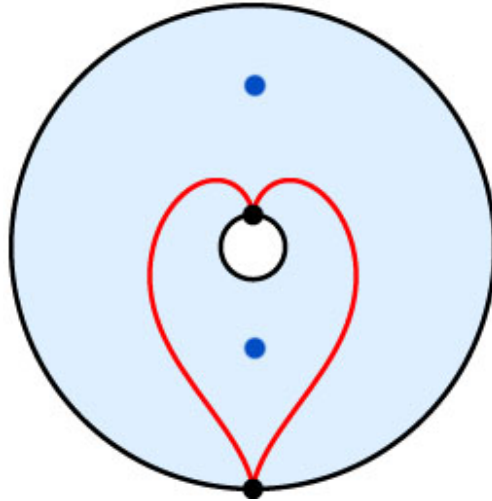


Figure 4.55: The $SU(2)$ triangulation of an annulus. The red lines denote the two internal edges whose flips describe BPS states. The dark blue dots are the zeros of the SW differential.

As we flow through time, the triangulation of the front face will evolve to that of the back face. This results in a tetrahedral decomposition of M . In the strong coupling chamber

there are two states, and this determines a decomposition of the solid torus into two tetrahedra, much as in the previous section's description of the A_n theories. However, in the weak coupling chamber, the geometry is much more novel. Due to the presence of infinitely many BPS particles in 4d, the triangulation undergoes an infinite sequence of flips. Along these flips the internal edges begin to accumulate as shown in Figure 4.56, and in the limit of infinitely many flips the W-boson appears. Now we can reinterpret this sequence of flips as describing a decomposition of the solid torus into infinitely many tetrahedra which degenerate. Similar structures have been studied by mathematicians [85].




Figure 4.56: The sequence of flips for the weak coupling spectrum of $SU(2)$. There are infinitely many BPS states which give rise to flips that accumulate. In the corresponding 3d theory these are reinterpreted as gluing data for accumulating tetrahedra.

Just as the 2-3 Pachner move is the geometric manifestation of 2-3 wall crossing in 4d, the relationship described above, between two tetrahedra, and infinitely many, can be seen as the geometric version of the wall-crossing governing the decay of a BPS vector multiplet. The physical meaning, and mathematical consequences of this phenomenon demand further illumination.

Appendix A

Self-Folded Triangles

In our discussion above we have left out a minor technicality involving *self-folded* triangles. A self-folded triangle is one in which two sides become identified, resulting in the degenerate structure seen below.



We will call the edge labeled *ext* exterior, and the edge labeled *int* interior. The framework of triangulations above requires allowance of self-folded triangles. In particular, some triangulations obtained from special lagrangian flows will require self-folded triangles, and similarly, some flips will force self-folded triangles to occur.

To properly include these structures, we must slightly augment the rules for obtaining a quiver Q and superpotential \mathcal{W} from a triangulation \mathcal{T} . First, it is useful to note that self-folded triangles, while necessary for the formalism, are a bit of an extraneous complication. It is a theorem from [54] that every surface admits a triangulation without self-folded triangles. Thus, having

carefully understood the map from triangulations and quivers, which maps flips to mutations, the rules for self-folded triangles can be derived from the rules given in the body of the paper. We would simply apply flips of the triangulation to remove all self-folded triangles, use the given rules to obtain Q and \mathcal{W} , and then invert the flips with the appropriate inverse mutations on the quiver. For completeness, we give the relevant rules here.

To obtain the quiver Q , we apply the usual rules as given in section 3.2.2 to all diagonals, except for interior edges of self-folded triangles. For the interior edge of each self-folded triangle, we draw a node corresponding to it, and draw arrows that duplicate the arrows of the node corresponding to the exterior edge of the same self-folded triangle. For clarity, let us define a function e on diagonals δ : if δ is an interior edge, $e(\delta)$ is the exterior edge of the self-folded triangle whose interior edge is δ ; otherwise, $e(\delta)$ is simply δ . Similarly, we define $i(\delta)$ to give the associated interior edge if δ is an exterior one. Thus the full rules are:

- For each diagonal δ in the triangulation, draw exactly one node of the quiver.
- For each pair of diagonals δ_1, δ_2 find all triangles for which $e(\delta_1), e(\delta_2)$ are both edges. Then for each such triangle draw one arrow from δ_1 to δ_2 if $e(\delta_1)$ immediately precedes $e(\delta_2)$ going counter-clockwise around the triangle.

Similarly, we should also extend the superpotential to include self-folded triangles. We use $\alpha, \beta, \gamma \dots$ to denote both the diagonals and their respective nodes in the quiver, and $B_{\alpha\beta}$ to denote both an arrow from α to β and the associated bifundamental matter field. The full rules are as follows:

- For each internal, non-self-folded triangle $\alpha\beta\gamma$, we add the associated three cycle $B_{\alpha\beta}B_{\beta\gamma}B_{\gamma\alpha}$.
- For each internal, non-self-folded triangle $\alpha\beta\gamma$ adjacent to exactly two self-folded triangles enclosed by α, β respectively, we add an additional three cycle $B_{i(\alpha)i(\beta)}B_{i(\beta)\gamma}B_{\gamma i(\alpha)}$.

- For each internal, non-self-folded triangle $\alpha\beta\gamma$ adjacent to exactly three self-folded triangles, we add three additional terms $B_{i(\alpha)i(\beta)}B_{i(\beta)\gamma}B_{\gamma i(\alpha)} + B_{i(\alpha)\beta}B_{\beta i(\gamma)}B_{i(\gamma)i(\alpha)} + B_{\alpha i(\beta)}B_{i(\beta)i(\gamma)}B_{i(\gamma)\alpha}$.
- For each internal, regular puncture adjacent to exactly one internal diagonal α , we must have a self-folded triangle. The diagonal $e(\alpha)$ occurs in at most one non-self-folded triangle. If that triangle is internal, $e(\alpha)\beta\gamma$, we add the three cycle $B_{\alpha\beta}B_{\beta\gamma}B_{\gamma\alpha}$.
- For each internal, regular puncture adjacent to more than one internal diagonal, we remove all the exterior edges of self-folded triangles incident on the puncture. Now let n be the number of remaining diagonals incident on the puncture. The quiver must have an n cycle $\alpha_1 \dots \alpha_n$; we add the term $B_{\alpha_1\alpha_2} \dots B_{\alpha_{n-1}\alpha_n}B_{\alpha_n\alpha_1}$.

Bibliography

- [1] M. Kontsevich and Y. Soibelman, “Stability structures, motivic Donaldson-Thomas invariants and cluster transformations,” *ArXiv e-prints* (Nov., 2008) , [arXiv:0811.2435](#) [[math.AG](#)].
- [2] D. Gaiotto, G. W. Moore, and A. Neitzke, “Four-dimensional wall-crossing via three-dimensional field theory,” *Commun.Math.Phys.* **299** (2010) 163–224, [arXiv:0807.4723](#) [[hep-th](#)].
- [3] S. Cecotti and C. Vafa, “BPS Wall Crossing and Topological Strings,” [arXiv:0910.2615](#) [[hep-th](#)].
- [4] S. Cecotti, A. Neitzke, and C. Vafa, “R-Twisting and 4d/2d Correspondences,” [arXiv:1006.3435](#) [[hep-th](#)].
- [5] A. Sen, “Equivalence of Three Wall Crossing Formulae,” [arXiv:1112.2515](#) [[hep-th](#)].
- [6] M. R. Douglas and G. W. Moore, “D-branes, Quivers, and ALE Instantons,” [arXiv:hep-th/9603167](#).
- [7] M. R. Douglas, B. Fiol, and C. Römelsberger, “Stability and BPS branes,” *JHEP* **0509** (2005) 006, [arXiv:hep-th/0002037](#) [[hep-th](#)].

-
- [8] M. R. Douglas, B. Fiol, and C. Römelsberger, “The Spectrum of BPS branes on a noncompact Calabi-Yau,” *JHEP* **0509** (2005) 057, [arXiv:hep-th/0003263](#) [hep-th].
- [9] N. Seiberg and E. Witten, “Electric - magnetic duality, monopole condensation, and confinement in N=2 supersymmetric Yang-Mills theory,” *Nucl.Phys.* **B426** (1994) 19–52, [arXiv:hep-th/9407087](#) [hep-th].
- [10] N. Seiberg and E. Witten, “Monopoles, duality and chiral symmetry breaking in N=2 supersymmetric QCD,” *Nucl.Phys.* **B431** (1994) 484–550, [arXiv:hep-th/9408099](#) [hep-th].
- [11] A. Klemm, W. Lerche, P. Mayr, C. Vafa, and N. P. Warner, “Selfdual strings and N=2 supersymmetric field theory,” *Nucl.Phys.* **B477** (1996) 746–766, [arXiv:hep-th/9604034](#) [hep-th].
- [12] S. Cecotti and C. Vafa, “Classification of complete N=2 supersymmetric theories in 4 dimensions,” [arXiv:1103.5832](#) [hep-th].
- [13] D. Gaiotto, G. W. Moore, and A. Neitzke, “Wall-crossing, Hitchin Systems, and the WKB Approximation,” [arXiv:0907.3987](#) [hep-th].
- [14] K. Intriligator and N. Seiberg, “Mirror Symmetry in Three-Dimensional Gauge Theories,” [arXiv:9607.207](#) [hep-th].
- [15] O. Aharony, A. Hanany, K. Intriligator, N. Seiberg, and M. Strassler, “Aspects of N=2 Supersymmetric Gauge Theories in Three Dimensions,” [arXiv:9703.110](#) [hep-th].
- [16] T. Dimofte, D. Gaiotto, and S. Gukov, “Gauge Theories Labelled by Three-Manifolds,” [arXiv:1108.4389](#) [hep-th].

- [17] F. Denef, “Supergravity flows and D-brane stability,” *JHEP* **0008** (2000) 050, [arXiv:hep-th/0005049](#) [hep-th].
- [18] F. Denef and G. W. Moore, “Split states, entropy enigmas, holes and halos,” [arXiv:hep-th/0702146](#) [hep-th].
- [19] D. Joyce and Y. Song, “A Theory of generalized Donaldson-Thomas invariants,” [arXiv:0810.5645](#) [math.AG].
- [20] T. Dimofte and S. Gukov, “Refined, Motivic, and Quantum,” *Lett.Math.Phys.* **91** (2010) 1, [arXiv:0904.1420](#) [hep-th].
- [21] T. Dimofte, S. Gukov, and Y. Soibelman, “Quantum Wall Crossing in N=2 Gauge Theories,” *Lett.Math.Phys.* **95** (2011) 1–25, [arXiv:0912.1346](#) [hep-th].
- [22] D. Gaiotto, G. W. Moore, and A. Neitzke, “Framed BPS States,” [arXiv:1006.0146](#) [hep-th].
- [23] E. Andriyash, F. Denef, D. L. Jafferis, and G. W. Moore, “Wall-crossing from supersymmetric galaxies,” [arXiv:1008.0030](#) [hep-th].
- [24] E. Andriyash, F. Denef, D. L. Jafferis, and G. W. Moore, “Bound state transformation walls,” [arXiv:1008.3555](#) [hep-th].
- [25] J. Manschot, B. Pioline, and A. Sen, “Wall Crossing from Boltzmann Black Hole Halos,” *JHEP* **1107** (2011) 059, [arXiv:1011.1258](#) [hep-th].
- [26] D. Gaiotto, G. W. Moore, and A. Neitzke, “Wall-Crossing in Coupled 2d-4d Systems,” [arXiv:1103.2598](#) [hep-th].
- [27] S. Cecotti and M. Del Zotto, “On Arnold’s 14 ‘exceptional’ N=2 superconformal gauge theories,” *JHEP* **1110** (2011) 099, [arXiv:1107.5747](#) [hep-th].

-
- [28] M. Alim, S. Cecotti, C. Cordova, S. Espahbodi, A. Rastogi, and C. Vafa, “BPS Quivers and Spectra of Complete $N=2$ Quantum Field Theories,” [arXiv:1109.4941](#) [[hep-th](#)].
- [29] M. Del Zotto, “More Arnold’s $N = 2$ superconformal gauge theories,” *JHEP* **1111** (2011) 115, [arXiv:1110.3826](#) [[hep-th](#)].
- [30] D.-E. Diaconescu and J. Gomis, “Fractional branes and boundary states in orbifold theories,” *JHEP* **0010** (2000) 001, [arXiv:hep-th/9906242](#) [[hep-th](#)].
- [31] B. Fiol and M. Marino, “BPS states and algebras from quivers,” *JHEP* **0007** (2000) 031, [arXiv:hep-th/0006189](#) [[hep-th](#)].
- [32] B. Fiol, “The BPS spectrum of $N = 2$ $SU(N)$ SYM and parton branes,” [arXiv:hep-th/0012079](#).
- [33] F. Denef, “Quantum quivers and Hall / hole halos,” *JHEP* **0210** (2002) 023, [arXiv:hep-th/0206072](#) [[hep-th](#)].
- [34] P. C. Argyres and M. R. Douglas, “New phenomena in $SU(3)$ supersymmetric gauge theory,” *Nucl.Phys.* **B448** (1995) 93–126, [arXiv:hep-th/9505062](#) [[hep-th](#)].
- [35] E. Witten, “Solutions of four-dimensional field theories via M theory,” *Nucl.Phys.* **B500** (1997) 3–42, [arXiv:hep-th/9703166](#) [[hep-th](#)].
- [36] D. Gaiotto, “ $N=2$ dualities,” [arXiv:0904.2715](#) [[hep-th](#)].
- [37] D. Gaiotto and J. Maldacena, “The Gravity duals of $N=2$ superconformal field theories,” [arXiv:0904.4466](#) [[hep-th](#)].
- [38] F. Benini, S. Benvenuti, and Y. Tachikawa, “Webs of five-branes and $N=2$ superconformal field theories,” *JHEP* **0909** (2009) 052, [arXiv:0906.0359](#) [[hep-th](#)].

- [39] O. Chacaltana and J. Distler, “Tinkertoys for Gaiotto Duality,” *JHEP* **1011** (2010) 099, [arXiv:1008.5203 \[hep-th\]](#).
- [40] D. Xie, “Network, Cluster coordinates and N=2 theory I,” [arXiv:1203.4513 \[hep-th\]](#).
- [41] N. Seiberg, “Electric - magnetic duality in supersymmetric nonAbelian gauge theories,” *Nucl.Phys.* **B435** (1995) 129–146, [arXiv:hep-th/9411149 \[hep-th\]](#).
- [42] P. S. Aspinwall, T. Bridgeland, A. Craw, M. Douglas, M. Gross, A. Kapustin, G. W. Moore, G. Segal, B. Szendrői, and P. Wilson, *Dirichlet Branes and Mirror Symmetry, (Clay mathematics monographs. Volume 4)*. American Mathematical Society, Clay Mathematics Institute, 2009.
- [43] D.-E. Diaconescu, M. R. Douglas, and J. Gomis, “Fractional branes and wrapped branes,” *JHEP* **9802** (1998) 013, [arXiv:hep-th/9712230 \[hep-th\]](#).
- [44] S. Cecotti and C. Vafa, “On classification of N=2 supersymmetric theories,” *Commun. Math. Phys.* **158** (1993) 569–644, [arXiv:hep-th/9211097](#).
- [45] K. Hori, A. Iqbal, and C. Vafa, “D-branes and mirror symmetry,” [arXiv:hep-th/0005247 \[hep-th\]](#).
- [46] F. Cachazo, B. Fiol, K. A. Intriligator, S. Katz, and C. Vafa, “A Geometric unification of dualities,” *Nucl.Phys.* **B628** (2002) 3–78, [arXiv:hep-th/0110028 \[hep-th\]](#).
- [47] B. Feng, A. Hanany, Y. H. He, and A. Iqbal, “Quiver theories, soliton spectra and Picard-Lefschetz transformations,” *JHEP* **0302** (2003) 056, [arXiv:hep-th/0206152 \[hep-th\]](#).
- [48] B. Feng, A. Hanany, and Y.-H. He, “D-brane gauge theories from toric singularities and toric duality,” *Nucl.Phys.* **B595** (2001) 165–200, [arXiv:hep-th/0003085 \[hep-th\]](#).

- [49] B. Feng, A. Hanany, and Y.-H. He, “Phase structure of D-brane gauge theories and toric duality,” *JHEP* **0108** (2001) 040, [arXiv:hep-th/0104259](#) [hep-th].
- [50] A. Hanany and K. D. Kennaway, “Dimer models and toric diagrams,” [arXiv:hep-th/0503149](#) [hep-th].
- [51] S. Franco, A. Hanany, K. D. Kennaway, D. Vegh, and B. Wecht, “Brane dimers and quiver gauge theories,” *JHEP* **0601** (2006) 096, [arXiv:hep-th/0504110](#) [hep-th].
- [52] B. Feng, Y.-H. He, K. D. Kennaway, and C. Vafa, “Dimer models from mirror symmetry and quivering amoebae,” *Adv.Theor.Math.Phys.* **12** (2008) 3, [arXiv:hep-th/0511287](#) [hep-th].
- [53] A. D. Shapere and C. Vafa, “BPS structure of Argyres-Douglas superconformal theories,” [arXiv:hep-th/9910182](#) [hep-th].
- [54] S. Fomin, M. Shapiro, and D. Thurston, “Cluster algebras and triangulated surfaces. Part I: Cluster complexes,” *Acta Mathematica* **201** (Aug., 2006) 83–146, [arXiv:math/0608367](#).
- [55] A. King, “Moduli of representations of finitedimensional algebras,” *Quart. J. Mat. Oxford* **45** (1994) 515–530.
- [56] P. Gabriel, “Unzerlegbare Darstellungen,” *Manuscripta Mathematica* **6** (1972) 71–103.
- [57] I. N. Bernstein, I. M. Gel’fand, and V. A. Ponomarev, “Coxeter Functors and Gabriel’s Theorem,” *Russian Mathematical Surveys* **28** (Apr., 1973) 17–32.
- [58] C. E. Beasley and M. Plesser, “Toric duality is Seiberg duality,” *JHEP* **0112** (2001) 001, [arXiv:hep-th/0109053](#) [hep-th].
- [59] B. Feng, A. Hanany, Y.-H. He, and A. M. Uranga, “Toric duality as Seiberg duality and brane diamonds,” *JHEP* **0112** (2001) 035, [arXiv:hep-th/0109063](#) [hep-th].

- [60] D. Berenstein and M. R. Douglas, “Seiberg duality for quiver gauge theories,”
[arXiv:hep-th/0207027](#).
- [61] B. Keller, “Cluster algebras, quiver representations and triangulated categories,”
[arXiv:0807.1960 \[math.RT\]](#).
- [62] B. Keller, “On cluster theory and quantum dilogarithm identities,” *ArXiv e-prints* (Feb., 2011), [arXiv:1102.4148 \[math.RT\]](#).
- [63] F. Ferrari and A. Bilal, “The Strong coupling spectrum of the Seiberg-Witten theory,”
Nucl.Phys. **B469** (1996) 387–402, [arXiv:hep-th/9602082 \[hep-th\]](#).
- [64] A. Bilal and F. Ferrari, “Curves of marginal stability, and weak and strong coupling BPS spectra in N=2 supersymmetric QCD,” *Nucl.Phys.* **B480** (1996) 589–622,
[arXiv:hep-th/9605101 \[hep-th\]](#).
- [65] A. Bilal and F. Ferrari, “The BPS spectra and superconformal points in massive N=2 supersymmetric QCD,” *Nucl.Phys.* **B516** (1998) 175–228, [arXiv:hep-th/9706145 \[hep-th\]](#).
- [66] S. H. Katz, A. Klemm, and C. Vafa, “Geometric engineering of quantum field theories,”
Nucl.Phys. **B497** (1997) 173–195, [arXiv:hep-th/9609239 \[hep-th\]](#).
- [67] S. Katz, P. Mayr, and C. Vafa, “Mirror symmetry and exact solution of 4-D N=2 gauge theories: 1.,” *Adv.Theor.Math.Phys.* **1** (1998) 53–114, [arXiv:hep-th/9706110 \[hep-th\]](#).
- [68] W. Lerche, “On a boundary CFT description of nonperturbative N=2 Yang-Mills theory,”
[arXiv:hep-th/0006100 \[hep-th\]](#).
- [69] A. Klemm, W. Lerche, S. Yankielowicz, and S. Theisen, “Simple singularities and N=2

- supersymmetric Yang-Mills theory,” *Phys.Lett.* **B344** (1995) 169–175,
[arXiv:hep-th/9411048](#) [hep-th].
- [70] P. C. Argyres and A. E. Faraggi, “The vacuum structure and spectrum of N=2 supersymmetric SU(n) gauge theory,” *Phys.Rev.Lett.* **74** (1995) 3931–3934,
[arXiv:hep-th/9411057](#) [hep-th].
- [71] A. Klemm, W. Lerche, and S. Theisen, “Nonperturbative effective actions of N=2 supersymmetric gauge theories,” *Int.J.Mod.Phys.* **A11** (1996) 1929–1974,
[arXiv:hep-th/9505150](#) [hep-th].
- [72] W. Lerche, “Introduction to Seiberg-Witten theory and its stringy origin,”
Nucl.Phys.Proc.Suppl. **55B** (1997) 83–117, [arXiv:hep-th/9611190](#) [hep-th].
- [73] O. Chacaltana and J. Distler, “Tinkertoys for the D_N series,” [arXiv:1106.5410](#) [hep-th].
- [74] D. Joyce, “On counting special Lagrangian homology 3-spheres,” *Contemp. Math.* **314** (2002) 125–151, [arXiv:hep-th/9907013](#).
- [75] K. Strebel, *Quadratic Differentials*. Springer Verlag, 1984.
- [76] A. Felikson, M. Shapiro, and P. Tumarkin, “Skew-symmetric cluster algebras of finite mutation type,” [arXiv:0811.1703](#) [math.CO].
- [77] M. Alim, S. Cecotti, C. Cordova, S. Espahbodi, A. Rastogi, and C. Vafa, “N=2 Quantum Field Theories and Their BPS Quivers,” [arXiv:1112.3984](#) [hep-th].
- [78] D. Labardini-Fragoso, “Quivers with potentials associated to triangulated surfaces, Part II: Arc representations,” [arXiv:0909.4100](#) [math.RT].
- [79] D. Labardini-Fragoso, “Quivers with potentials associated to triangulated surfaces, Part II: Arc representations,” [arXiv:0909.4100](#) [math.RT].

- [80] G. Cerulli Irelli and D. Labardini-Fragoso, “Quivers with potentials associated to triangulated surfaces, Part III: tagged triangulations and cluster monomials,” [arXiv:1108.1774](#) [math.RT].
- [81] H. Derksen and T. Owen, “New Graphs of Finite Mutation Type,” [arXiv:0804.0787](#) [math.CO].
- [82] S. Ladkani, “Mutation classes of certain quivers with potentials as derived equivalence classes,” *ArXiv e-prints* (Feb., 2011) , [arXiv:1102.4108](#) [math.RT].
- [83] N. Drukker, D. Gaiotto, and J. Gomis, “The Virtue of Defects in 4D Gauge Theories and 2D CFTs,” [arXiv:1003.1112](#) [hep-th].
- [84] Y. terashima and M. Yamazaki, “SL(2,R) Chern-Simons, Liouville, and Gauge Theory on Duality Walls,” [arXiv:1103.5748](#) [hep-th].
- [85] T. Jorgensen, “On pairs of once-punctured tori. Klienian groups and hyperbolic 3-manifolds,” *London Math. Soc. Lecture Notes* **299** (2003) 183–207.
- [86] F. Gueritaud, “Triangulated cores of punctured-torus groups,” *J. Differential Geom.* **81** (2009) 91–142.
- [87] R. Dijkgraaf and E. Witten, “Topological Gauge Theories and Group Cohomology,” *Commun. Math. Phys.* **129** (1990) 393.
- [88] R. Mclean, “Deformations of Calibrated Submanifolds,” *Commun. Analy. Geom.* **6** (1996) 705.
- [89] N. Hitchin, “The moduli space of special Lagrangian sub manifolds,” [arXiv:9711.002](#) [math.DG].

-
- [90] N. Seiberg, “The Power of Holomorphy-Exact Results in 4D SUSY Field Theories,”
[arXiv:9408.013 \[hep.th\]](#).
- [91] A. Polyakov, “Fermi-Bose transmutations induced by gauge fields,” *Mod. Phys. Lett.* **A3**
(1988) 325.
- [92] T. Banks and N. Seiberg, “Symmetries and Strings in Field Theory and Gravity,”
[arXiv:1011.5120 \[hep-th\]](#).
- [93] D. Belov and G. Moore, “Classification of Abelian Spin Chern-Simons Theories,”
[arXiv:0505.235 \[hep-th\]](#).
- [94] C. Cecotti, C. Cordova, J. Heckman, and C. Vafa, “T-Branes and Mondromy,”
[arXiv:1010.5780 \[hep-th\]](#).
- [95] R. Lickorish, *An Introduction to Knot Theory*. Springer-Verlag, 1997.
- [96] P. Griffiths and J. Harris, *Principles of Algebraic Geometry*. Wiley, 1978.
- [97] E. Witten, “ $SL(2, \mathbb{Z})$ Action on Three-Dimensional Conformal Field Theories With Abelian
Symmetry,” [arXiv:0307.041.002 \[hep.th\]](#).
- [98] D. Joyce, “Constructing special Lagrangian m -folds in C^m by evolving quadrics,”
[arXiv:0008.155 \[math.DG\]](#).
- [99] D. Joyce, “Evolution equations for special Lagrangian 3-folds in C^3 ,” [arXiv:0010.036](#)
[\[math.DG\]](#).
- [100] R. Harvey, *Spinors and Calibration*. Academic Press, 1990.
- [101] D. Gaiotto and E. Witten, “Janus Configurations, Chern-Simons Couplings, And the
Theta-Angle in $N=4$ Super Yang-Mills Theory,” [arXiv:0804.2907 \[hep.th\]](#).

-
- [102] D. Gaiotto and E. Witten, “S-Duality of Boundary Conditions In N=4 Super Yang-Mills Theory,” [arXiv:0807.3720](#) [[hep.th](#)].
- [103] A. Kapustin and M. Strassler, “On Mirror Symmetry in Three Dimensional Abelian Gauge Theories,” [arXiv:9902.033](#) [[hep.th](#)].
- [104] M. Aganagic and C. Vafa, “Mirror Symmetry, D-Branes, and Counting Holomorphic Discs,” [arXiv:0012.041](#) [[hep.th](#)].
- [105] T. Dimofte, “Quantum Riemann Surfaces in Chern-Simons Theory,” [arXiv:1102.4847](#) [[hep-th](#)].
- [106] U. Lindstrom and M. Rocek, “A note on the Seiberg-Witten solution of N=2 Super Yang-Mills Theory,” [arXiv:9503.012](#) [[hep-th](#)].
- [107] A. Giveon and M. Rocek, “Effective Actions and Gauge Field Stability,” [arXiv:9508.043](#) [[hep-th](#)].