



Brane Constructions and BPS Spectra

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

Citation	Rastogi, Ashwin. 2013. Brane Constructions and BPS Spectra. Doctoral dissertation, Harvard University.
Citable link	http://nrs.harvard.edu/urn-3:HUL.InstRepos:11148286
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

Brane Constructions and BPS Spectra

A dissertation presented

by

Ashwin Rastogi

to

The Department of Physics

in partial fulfillment of the requirements

for the degree of

Doctor of Philosophy

in the subject of

Physics

Harvard University

Cambridge, Massachusetts

May 2013

©2013 - Ashwin Rastogi

All rights reserved.

Thesis advisor

Author

Cumrun Vafa

Ashwin Rastogi

Brane Constructions and BPS Spectra

Abstract

The object of this work is to exploit various constructions of string theory and M-theory to yield new insights into supersymmetric theories in both four and three dimensions. In 4d, we extend work on Seiberg-Witten theory to study and compute BPS spectra of the class of complete $\mathcal{N} = 2$ theories. The approach we take is based on the program of geometric engineering, in which 4d theories are constructed from compactifications of type IIB strings on Calabi-Yau manifolds. In this setup, the natural candidates for BPS states are D3 branes wrapped on supersymmetric 3-cycles in the Calabi-Yau. Our study makes use of the mathematical structure of quivers, whose representation theory encodes the notion of stability of BPS particles. Except for 11 exceptional cases, all complete theories can be constructed by wrapping stacks of two M5 branes on Riemann surfaces. By exploring the connection between quivers and M5 brane theories, we develop a powerful algorithm for computing BPS spectra, and give an in-depth study of its applications. In particular, we compute BPS spectra for all asymptotically free complete theories, as well as an infinite set of conformal $SU(2)^k$ theories with certain matter content.

From here, we go on to apply the insight gained from our 4d study to 3d gauge theories. We consider the analog of the M5 brane construction in the case of 3d

Abstract

$\mathcal{N} = 2$ theories: pairs of M5 branes wrapped on a 3-manifold. Using the ansatz of R-flow, we study 3-manifolds consisting of Riemann surfaces fibered over \mathbb{R} . When the construction is non-singular, the resulting IR physics is described by a free abelian Chern-Simons theory. The mathematical data of a tangle captures the data of the gauge theory, and the Reidemeister equivalences on tangles correspond to dualities of physical descriptions. To obtain interacting matter, we allow singularities in the construction. By extending the tangle description to these singular cases, we find a set of generalized Reidemeister moves that capture non-trivial mirror symmetries of 3d gauge theories. These results give a geometric origin to these well-known 3d dualities.

Contents

Title Page	i
Abstract	iii
Table of Contents	v
Citations to Previously Published Work	vii
Acknowledgments	viii
Dedication	x
1 Introduction	1
2 Formalism of BPS Quivers	8
2.1 Introduction	8
2.2 $\mathcal{N} = 2$ theories in 4d	10
2.3 BPS Quivers and Spectra	11
2.4 Quiver Mutation and Duality	32
2.5 The Mutation Method	48
3 Complete Theories	60
3.1 Introduction	60
3.2 Classification of Complete Theories	62
3.3 BPS Quivers of Complete Theories	67
3.4 Theories with Finite Chambers	100
3.5 Exceptional Complete Theories	121
4 3d Superconformal Theories and Mirror Symmetry	126
4.1 Introduction	126
4.2 Abelian Chern-Simons Theory and Tangles	134
4.3 Particles, Singularities, and Superpotentials	168
4.4 Dualities and Generalized Reidemeister Moves	190
4.5 R-flow	205
4.6 Applications	224

Bibliography	236
---------------------	------------

Citations to Previously Published Work

The content of most of this thesis has appeared in preprints on the **arXiv**, and are currently in various stages of publication. Ch. 2 is based on material in the preprint:

“ $\mathcal{N} = 2$ Quantum Field Theories and Their BPS Quivers”, M. Alim, S. Cecotti, C. Cordova, S. Espahbodi, A. Rastogi, C. Vafa. arXiv:1112.3984 [hep-th]. To appear in Communications in Mathematical Physics.

Ch. 3 is based on the preprint:

“BPS Quivers and Spectra of Complete $\mathcal{N} = 2$ Quantum Field Theories”, M. Alim, S. Cecotti, C. Cordova, S. Espahbodi, A. Rastogi, C. Vafa. arXiv:1109.4941 [hep-th]. Submitted to Journal of High Energy Physics.

Ch. 4 is based on the preprint:

“Tangles, Generalized Reidemeister Moves, and Three-Dimensional Mirror Symmetry”, C. Cordova, S. Espahbodi, B. Haghighat, A. Rastogi, C. Vafa. arXiv:1211.3730 [hep-th]. Submitted to Advances in Theoretical and Mathematical Physics.

Acknowledgments

This dissertation represents the culmination of the most challenging endeavor I have undertaken thus far in my life. I owe the success of this work and my years in graduate school to many who have shaped my academic path and supported me throughout:

- My advisor, Cumrun Vafa, who introduced me to compelling research problems and offered his invaluable guidance and powerful intuition throughout my research work.
- My dissertation and orals committee members, Frederik Denef, Arthur Jaffe, and Melissa Franklin, for offering their time and insights to improve my research.
- My collaborators, Murad Alim, Clay Cordova, Sergio Cecotti, and Babak Haghighat, for sharing their enthusiasm about this work and helping to break the amorphous and intimidating challenge of research into manageable steps of progress and understanding.
- Matt Schwartz and David Simmons-Duffin, who guided me through my early study of quantum field theory and provided me with two very distinct perspectives on this crucial subject; and my classmates in that course, Chi-Ming Chang, Daniel Flassig, Dilani Kahawala, and Arnav Tripathi.
- My undergraduate professors, Chris Carone, Kostas Orginos, Vladimir Bolotnikov, Ilya Spitkovsky whose careful pedagogy and passion sparked my curiosity and inspired my own ambitions in academic research.
- My family, who instilled in me the values and character that were necessary to

Acknowledgments

succeed in my graduate studies.

- My comrades in suffering: Tarek Anous, Sam Espahbodi, and Nick Vanmeter, for sharing their moments of confusion and frustration as much as those rarer joys of insight and revelation; those who came before us: Dionysios Anninos, Mboyo Esole, and Tongyan Lin, for helping to light the way; and our erstwhile fearless leader, Frederik Denef, for teaching me that happiness is a choice.
- My wife, Lauren Merrill, for whom any blurb of an acknowledgement would be an understatement of incredible proportions. Her companionship on this journey has been simply indispensable.

Dedicated to my dearest friend and closest ally, Lauren Nicole Merrill.

Chapter 1

Introduction

One of the primary and enduring features of string theory is the appearance of extra dimensions and the need for compactification. Perturbative supersymmetric string theory naturally exists in 10 dimensions; to make contact with familiar 4d physics of any kind, it is necessary to dispose of these extra dimensions in some way. A wide array of such constructions have been explored. This work has led to important progress with applications both to our understanding of abstract field theory, and to models of phenomenology. This course of study has also deepened our knowledge of string theory itself, as it has illuminated a whole web of dualities within string theory.

In particular, tools from string theory have found immense success in the study of $\mathcal{N} = 2$ supersymmetric 4d gauge theories. This class of theories is known to be non-phenomenological; if the universe contains any 4d low-energy supersymmetry at all, it must be a minimal $\mathcal{N} = 1$ symmetry. Nonetheless, these theories are extremely rich, exhibit a wide array of non-perturbative features, and have offered deep insight into

phenomenological physics. The application of string theoretic techniques to $\mathcal{N} = 2$ physics has become collectively known in the literature as *geometric engineering*. Compactification of type IIB superstrings on certain Calabi-Yau 3-manifolds yields low-energy $\mathcal{N} = 2$ 4d gauge theories. From here, additional tools from string theory and geometry allow us to obtain a variety of non-perturbative results in $\mathcal{N} = 2$ field theory. The Seiberg-Witten solution [1, 2] is in many ways the foundational result in this area. Although the Seiberg-Witten work was originally understood without reference to string theory, the geometric construction of that result from string theory gives a more complete perspective on the role of the Seiberg-Witten Riemann surface, and its relation to a string compactification.

The Seiberg-Witten solution gives the extreme low-energy, non-perturbative solution for the dynamics of massless fields in $\mathcal{N} = 2$ 4d gauge theory. There is, of course, much more to study in these theories. In particular, the Seiberg-Witten solution says nothing about the spectrum or dynamics of massive states. Yet, string theory has a great deal of insight to offer into the massive sector of the theory as well. The string theory construction includes the full UV dynamics of the gauge theory, including all massive modes and their interactions. Of course, in general these features are susceptible to perturbative and non-perturbative corrections. As a consequence of the special geometric and string theoretic characteristics of these constructions, there are many attributes of the resulting 4d theories that are protected from corrections. These results in Calabi-Yau compactification of type II strings have been aggressively pursued in the literature.

The BPS spectrum is one such protected feature. BPS states saturate a lower-

bound on mass, given by $M \geq |Z|$, where M is the mass of the state, and Z is the central charge. These states form short supermultiplets, and are protected from decay by the mass bound. From the string perspective, we have a natural set of candidates for BPS states - namely, the extended D-branes that arise in type II superstrings. To produce massive BPS states, the D-branes should wrap non-vanishing supersymmetric cycles in the 6-dimensional Calabi-Yau manifold, and thus we will make use of geometric features of the compactification.

It has long been understood that BPS spectra play a pivotal role in the study of quantum field theories with extended supersymmetry. This was especially underscored in Seiberg-Witten solution to $\mathcal{N} = 2$ 4d supersymmetric Yang-Mills theory. In that work, finding singularities in moduli space where various BPS states became massless was the fundamental insight that led to the full solution of the IR dynamics of the theory. However, in spite of the general consensus about the importance of BPS states and a wealth of recent progress in the subject, there are relatively few explicit examples, and hence little known about the general structure of BPS spectra.

The first two chapters of this thesis develop a framework for studying BPS spectra, and compute spectra for a new infinite class of examples. Our methods begin in chapter 2 from the perspective of quiver representation theory. Quivers are directed, 2-acyclic graphs; the representation theory of these structures has been studied in the mathematics literature, and found powerful applications in the computation of BPS states of physics. In our general discussion of BPS quivers, we make a key observation connecting mutation, a transformation on quivers that preserves the resulting representation theory, and the physical duality between particle and anti-particle.

Essentially, the specification of a BPS quiver associated to a particular theory depends on some choice of which states to denote particles and which states to denote anti-particles. By changing these artificial designations, we implement a mutation on the resulting quiver. While this observation seems straightforward, it has important and unforeseen implications for the resulting spectrum and the phenomenon of BPS wall-crossing. The most powerful result, however, is the algorithm we denote as the *mutation method*, which, in many cases, allows for the computation of a full BPS spectrum by a finite and easily implementable procedure.

In chapter 3, we will narrow our focus to the class of *complete theories*, introduced and classified in [3], where our program and techniques are quite powerful. Aside from 11 exceptional cases, all complete theories can be constructed from a stack of two M5-branes wrapped on a Riemann surface. These include general $SU(2)^k$ asymptotically free or conformal gauge theories, but no higher $SU(n)$ gauge groups. For these theories, our methods correspond elegantly to certain geometric structures on the Riemann surface. As noted before, we expect BPS particles in the 4d theory to correspond to some extended BPS objects in the UV string theory; the geometry becomes relevant because these higher extended BPS objects may be wrapped in non-trivial ways along the compactification manifold. In this case, the extended BPS objects are M2 branes stretching between the two M5 branes on the Riemann surface. Thus we would expect some correspondence between the BPS quivers we study and the Riemann surfaces in the geometry; indeed, there is a beautiful connection between quivers and triangulated Riemann surfaces. We are able to give a completely geometric description for the construction of the quiver as well as the implementation of the mutation method.

Following our techniques, we compute BPS spectra for any asymptotically free theory in this class, as well as an infinite subset of conformal $SU(2)^k$ theories.

This connection to M5 compactifications leads to chapter 4, where we explore features of the $\mathcal{N} = 2$ 3d gauge theories arising from stacks of two M5-branes wrapped on a 3d surface, with the intention of exploring features of both 3d gauge theory and the six-dimensional superconformal theory of M5 branes. The $(2,0)$ superconformal field theories in six dimensions, in particular the theory of N parallel M5-branes, are among the most important quantum systems, and yet they remain poorly understood. Their importance stems not only from the fact that they represent the highest possible dimension in which superconformal field theories can exist, but also from the observation that their compactifications to lower dimensions yield a rich class of quantum field theories whose dynamics are encoded by geometry. Following the perspective of chapter 3, four-dimensional $\mathcal{N} = 2$ theories arise upon compactification on a Riemann surface and provide a geometric explanation for Seiberg-Witten theory.

It is natural to expect that more general compactifications will provide more information about these mysterious six-dimensional theories. One way to do this is to increase the dimension of the compactification geometry. Thus, the next cases of interest would be compactifications with dimensions $d \geq 3$ resulting at low-energies in effective quantum field theories in dimensions $6 - d$. The aim of this chapter is to focus on the situation where $d = 3$ with $\mathcal{N} = 2$ supersymmetry. We will focus on the situation of two M5-branes wrapping some ultraviolet 3d geometry. In such constructions, the infrared dynamics of the system is described by a single recombined brane that can be viewed as a double cover of the original compactification manifold.

This infrared geometry is captured by describing the branching strands for the cover which in general are knotted. When the branching strands collide the cover becomes singular, and on that locus an M2-brane of vanishing size can end on the M5-branes, leading to massless charged matter fields. The goal of the final chapter of this thesis is to explore the dictionary between 3d geometry and 3d physics for these constructions; that is, to find the correspondence between the knotted branch locus encoding the geometry of the double cover and the underlying $\mathcal{N} = 2$ quantum field theory.

In this study, we can make use of our results from chapter 3 by applying the ansatz of *R*-flow in our construction of the 3d compactification manifold [4]. The complete theories of chapter 3 were constructed by wrapping two M5 branes on Riemann surfaces; to consider 3-manifolds, we may take the Riemann surface fibered over an additional real line \mathbb{R} . The Riemann surface includes certain puncture singularities, which are now lifted to the branching strands described above. If we allow the moduli of the Riemann surface vary along the fibration, then we can generate some arbitrary knotting behavior.¹ R-flow is given by enforcing certain technical constraints to preserve supersymmetry in the flow of the moduli, and makes contact with the manipulation of BPS charges induced by the tuning of Riemann surface moduli. In this ansatz, each 4d BPS state appears as a 3d massless particle localized in the 3d geometry.

We will discover that certain dualities, particularly 3d mirror symmetry, are made manifest in this construction. Our dictionary from geometry to a classical Lagrangian involves a certain geometric choice. Different choices will produce distinct Lagrangian

¹Of course, the resulting knot structure is highly dependent on the boundary conditions we take to close off the 3-manifold

descriptions of identical physics, and thus exhibit duality. In the context of R-flow, certain 3d mirror symmetries are related to wall-crossing of 4d BPS states, allowing us to make use of our results from the previous chapters. Continuing further in our study, we disassemble these dualities into underlying geometric rules that can be used to explore interesting, non-trivial dualities. The results offer a novel perspective and geometric origin for 3d mirror symmetries and dualities.

Chapter 2

Formalism of BPS Quivers

2.1 Introduction

This chapter will serve to collect background material on the subject of BPS quivers, which will be employed heavily throughout the rest of the thesis. Quiver techniques for BPS spectra of 4d $\mathcal{N} = 2$ field theories have been studied in a variety of contexts, though traditionally from the framework of geometric engineering [5–11]. Although string theory gives the mathematical inspiration for these techniques, we will take the perspective of 4d gauge theory in studying the applications of BPS quivers, and avoid string theoretic derivations whenever possible. The point is that these techniques are quite powerful for a huge class of 4d gauge theories, and perhaps have some underlying purely 4d justification.

The direction connection between BPS states and quivers comes through quiver representation theory. Representations of the quiver correspond to candidate BPS particles, while a stability condition on representations known as Π -stability is known

to correspond to the preservation of supersymmetry associated with BPS particles. Section 2.2 will be dedicated to discussing the basics of this representation theory, its relevance to the physics being studied, and some simple examples. Unfortunately, the representation theory quickly becomes quite convoluted, even for relatively simple quivers.

The valuable insights come from studying various dualities and transformations on quivers and spectra in Section 2.4. BPS quivers come equipped with a set of transformations known as *mutations*. For each node of a BPS quiver, there is a procedure to construct a new quiver mutated at that node. In the course of our discussion, we will see that mutation is related to the physical duality of particle and anti-particle. To specify a BPS quiver, we must make some designation of which states are particles and which ones are antiparticles. We choose to do this by simply picking a half-plane in the complex-plane of central charges \mathcal{Z} . Of course, there are infinitely many different half-planes one could choose, corresponding to different designations of particle and anti-particle. We find that reversing one of these designations is equivalent to a certain mutation. The mutated quiver produces an equivalent representation theory, and thus an equivalent spectrum of BPS states.

The point is that each quiver makes some BPS states (or equivalently, some stable representations) completely trivial to compute; however, these trivial states change as we apply various mutations. By carefully tabulating trivial states as we mutate, one can resolve the question of BPS stability for many states very efficiently. In fact, for theories with a finite BPS spectrum, we can even simply tune through all possible choices of particle half-plane and compute all the BPS states by this method.

Since the mutations transformation is easily implementable by hand or by computer programming, this method allows for many new spectra to be studied. We explore the details and implications of this algorithm at the end of Section 2.4.

2.2 $\mathcal{N} = 2$ theories in 4d

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.¹ 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 particles is more subtle. According to the $\mathcal{N} = 2$ superalgebra,

¹Here 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

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.

2.3 BPS 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 2.3.1 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 2.3.3 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 2.3.2. It is through these methods that we can hope in turn to discover previously unknown spectra.

2.3.1 Formal Construction of 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

$$\gamma = \sum_{i=1}^{2r+f} n_i \gamma_i. \quad n_i \in \mathbb{Z}^+ \quad (2.3.1)$$

We emphasize that the basis need not span Γ , but only the subset of occupied states

in Γ . We will see in section 2.3.3 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

$$\tilde{\gamma}_i = n_{ij} \gamma_j; \quad \gamma_i = (n^{-1})_{ij} \tilde{\gamma}_j. \quad (2.3.2)$$

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:

$$\begin{aligned} \text{Vector multiplet } W - \text{boson} : (2, 0), \\ \text{Hypermultiplet dyons} : (2n, 1), (2n + 2, -1) \quad n \geq 0. \end{aligned} \tag{2.3.3}$$

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.




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.

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 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 ,

$$\gamma = \sum_i n_i \gamma_i \implies \mathcal{Z}_u(\gamma) = \sum_i n_i \mathcal{Z}_u(\gamma_i). \quad n_i \geq 0 \quad (2.3.4)$$

Since $\mathcal{Z}(\gamma_i)$ all lie in the upper half-plane, (2.3.4) 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, as we will see in chapter 3. 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.

2.3.2 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:²

- Building on the original orbifold construction of quiver gauge theories of [5] refs. [6–8, 13] provided the identification of the quiver nodes with a basis of BPS states obtained from fractional branes, these BPS quivers were further explored in [9, 10].
- The relation of the $4d$ quivers with the soliton spectrum in $2d$ [14] was studied in various places, see for example [15–17], more recently this $2d/4d$ correspondence and the associated construction of BPS quivers was discussed in [18].
- The toric methods of [19, 20] and the relation to dimer models [21] were used in [22] to construct a large class of quivers, their construction using mirror symmetry was studied in [23].
- Based on the geometric study of BPS states in SW theories pioneered in [24] and further studied in [25, 26], the BPS quivers can be obtained from triangulations of Riemann surfaces as described in [3, 27] using the relation of triangulations and quivers of [28]. 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. These techniques have since been generalized to larger stacks of M5 branes [29].

²See also [12] 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.

2.3.3 Quiver Quantum Mechanics

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

$$\gamma = \sum_i n_i \gamma_i \quad (2.3.5)$$

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 $i \rightarrow j$,

$$\text{Gauge Group} = \prod_{\text{nodes } i} U(n_i), \quad \text{Matter} = \bigoplus_{\text{arrows } a} B_{ij}^a. \quad (2.3.6)$$

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 asses 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

$$\theta_i = |\mathcal{Z}_u(\gamma_i)| \left(\arg(\mathcal{Z}_u(\gamma_i)) - \arg(\mathcal{Z}_u(\gamma)) \right). \quad (2.3.7)$$

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

$$\sum_{\substack{\text{arrows} \\ \text{starting at } i}} |B_{ij}^a|^2 - \sum_{\substack{\text{arrows} \\ \text{ending at } i}} |B_{ki}^a|^2 = \theta_i. \quad (2.3.8)$$

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.³ Later in section 2.3.4 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 2.4 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

$$\frac{\partial \mathcal{W}}{\partial B_{ij}^a} = 0. \quad (2.3.9)$$

Having fully fixed the quantum mechanics, we now turn to the moduli space of supersymmetric ground states with charge γ , \mathcal{M}_γ .⁴ This space is simply the solution

³Alternatively 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.

⁴From now on, whenever we refer to supersymmetric ground states of the quiver quantum me-

to the equations of motion described above, quotiented by the action of the unitary gauge groups.

$$\mathcal{M}_\gamma = \left\{ B_{ij}^a \left| \frac{\partial \mathcal{W}}{\partial B_{ij}^a} = 0, \sum_{\substack{\text{arrows} \\ \text{starting at } i}} |B_{ij}^a|^2 - \sum_{\substack{\text{arrows} \\ \text{ending at } i}} |B_{ki}^a|^2 = \theta_i \right. \right\} / \prod_i U(n_i). \quad (2.3.10)$$

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,

$$\text{Spin} = \text{Lefschetz} \otimes \left(\left[\frac{1}{2} \right] + 2[0] \right). \quad (2.3.11)$$

Equation (2.3.11) 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

chanics, we will always mean on the Higgs branch. The Coulomb branch can also be studied and gives rise to equivalent results for BPS spectra. [11]

shift by $1/2$ in (2.3.11) simply takes into account the intrinsic spin contribution.

In practice the most important application of (2.3.11) 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 multiplets of length two, the canonical example of which is $\mathcal{M}_\gamma \cong \mathbb{P}^1$. In complete generality the formula (2.3.11) 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

$$d = \sum_{B_{ij}^a} (n_i n_j) - \sum_{\text{nodes } i} n_i^2 - (\# \text{F-term constraints}) + 1. \quad (2.3.12)$$

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.

2.3.4 Quiver Representation Theory

While this supersymmetric quantum mechanics 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 (2.3.10)).
- As the solution to the F-term equations modulo the action of the complexified gauge group $\prod_i \text{Gl}(n_i, \mathbb{C})$, 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 $B_{ij}^a : \mathbb{C}^{n_i} \longrightarrow \mathbb{C}^{n_j}$ 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 [12] 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:

$$\begin{array}{ccc} \mathbb{C}^{n_i} & \xrightarrow{B_{ij}^a} & \mathbb{C}^{n_j} \\ \uparrow & & \uparrow \\ \mathbb{C}^{m_i} & \xrightarrow{b_{ij}^a} & \mathbb{C}^{m_j} \end{array} \quad (2.3.13)$$

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,⁵

$$\mathcal{Z}_u(R) \equiv \mathcal{Z}_u(\gamma_R) = \sum_i n_i \mathcal{Z}_u(\gamma_i). \quad (2.3.14)$$

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

⁵When 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.

S other than R and zero, one has

$$\arg(\mathcal{Z}_u(S)) < \arg(\mathcal{Z}_u(R)). \quad (2.3.15)$$

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 (2.3.15) reduces to the D-term equations of motion presented earlier [7, 30].

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.

$$\mathcal{M}_\gamma = \left\{ R = \{B_{ij}^a : \mathbb{C}^{n_i} \rightarrow \mathbb{C}^{n_j}\} \left| \frac{\partial \mathcal{W}}{\partial B_{ij}^a} = 0, \text{ } R \text{ is } \Pi - \text{stable} \right. \right\} / \prod_i \text{Gl}(n_i, \mathbb{C}). \quad (2.3.16)$$

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.

2.3.5 Walls of Marginal Stability

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 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 proper-

ties 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.

2.3.6 Examples

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 [18]. We will denote by \mathcal{Z}_i the central charges of the two basis particles,

$$\textcircled{1} \longrightarrow \textcircled{2} \quad (2.3.17)$$

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

$$\mathbb{C}^{n_1} \xrightarrow{B} \mathbb{C}^{n_2} \quad (2.3.18)$$

To determine stability we investigate subrepresentations. Let's start with a subrep-

resentation of the form

$$\begin{array}{ccc}
 \mathbb{C}^{n_1} & \xrightarrow{B} & \mathbb{C}^{n_2} \\
 \uparrow & & \uparrow \\
 0 & \xrightarrow{0} & \mathbb{C}
 \end{array} \tag{2.3.19}$$

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

$$\arg(\mathcal{Z}_2) < \arg(n_1 \mathcal{Z}_1 + n_2 \mathcal{Z}_2) \implies \arg(\mathcal{Z}_2) < \arg(\mathcal{Z}_1). \tag{2.3.20}$$

Next we consider a similar decay involving the first basis particle

$$\begin{array}{ccc}
 \mathbb{C}^{n_1} & \xrightarrow{B} & \mathbb{C}^{n_2} \\
 \uparrow & & \uparrow \\
 \mathbb{C} & \xrightarrow{0} & 0
 \end{array} \tag{2.3.21}$$

If this is a subrepresentation, then stability demands that $\arg(\mathcal{Z}_1) < \arg(\mathcal{Z}_2)$, so (2.3.20) 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 (2.3.21) 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

$$\begin{array}{ccc}
 \mathbb{C}^{n_1} & \xrightarrow{B} & \mathbb{C}^{n_2} \\
 \uparrow & & \uparrow \\
 \mathbb{C} & \xrightarrow{b} & \mathbb{C}
 \end{array} \tag{2.3.22}$$

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

$$\arg(\mathcal{Z}_1 + \mathcal{Z}_2) < \arg(n_1 \mathcal{Z}_1 + n_2 \mathcal{Z}_2). \quad (2.3.23)$$

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 (2.3.22) 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 [25]. This basic 2-3





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.

decay process is known in various contexts as a primitive decay [31]. In formalism of Kontevich and Soibelman this wall-crossing gives rise to the pentagon identity of quantum dilogarithms.

Example: 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



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

$$\mathcal{W} = \alpha_3 \alpha_2 \alpha_1. \quad (2.3.25)$$

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

$$\alpha_2 \circ \alpha_1 = 0, \quad \alpha_3 \circ \alpha_2 = 0, \quad \alpha_1 \circ \alpha_3 = 0. \quad (2.3.26)$$

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

$$\mathbb{C}^{n_1} \xrightarrow{\alpha_1} \mathbb{C}^{n_2} \xrightarrow{\alpha_2} \mathbb{C}^{n_3} \quad (2.3.27)$$

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 (2.3.26). 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 $(1, 1, 0)$.

$$\begin{array}{ccccc}
 & & \alpha_3 & & \\
 & \curvearrowleft & & \curvearrowright & \\
 \mathbb{C}^{n_1} & \xrightarrow{\alpha_1} & \mathbb{C}^{n_2} & \xrightarrow{\alpha_2} & \mathbb{C}^{n_3} \\
 i \uparrow & & j \uparrow & & 0 \uparrow \\
 \mathbb{C} & \xrightarrow{\beta_1} & \mathbb{C} & \xrightarrow{\beta_2} & 0 \\
 & \curvearrowright & & \curvearrowleft & \\
 & 0 & & &
 \end{array} \tag{2.3.28}$$

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

$$\arg(\mathcal{Z}_1 + \mathcal{Z}_2) < \arg(n_1 \mathcal{Z}_1 + n_2 \mathcal{Z}_2 + n_3 \mathcal{Z}_3). \tag{2.3.29}$$

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 superpotnetial 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.

2.4 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 [32] 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 2.5 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

$$\mathcal{Z}_u(\gamma_i) \in \mathbb{R}. \quad (2.4.1)$$

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.




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 \mathcal{Z}_1 changes its identity to an antiparticle. The cone of particles jumps discontinuously and a new quiver description is required.

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 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 \mathcal{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 2.4.1 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 2.5 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.

2.4.1 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 . [15, 16, 33–37]. The new basis is given by

$$\tilde{\gamma}_1 = -\gamma_1 \tag{2.4.2}$$

$$\tilde{\gamma}_j = \begin{cases} \gamma_j + (\gamma_j \circ \gamma_1)\gamma_1 & \text{if } \gamma_j \circ \gamma_1 > 0 \\ \gamma_j & \text{if } \gamma_j \circ \gamma_1 \leq 0. \end{cases} \tag{2.4.3}$$

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

$$B_{i1}^a B_{1j}^b \longrightarrow \tilde{B}_{ij}^c. \quad (2.4.4)$$

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

$$\tilde{B}_{i1}^a \longrightarrow \tilde{B}_{1i}^a; \quad \tilde{B}_{1j}^a \longrightarrow \tilde{B}_{j1}^a. \quad (2.4.5)$$

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)

$$\tilde{B}_{1i}^a \tilde{B}_{ij}^c \tilde{B}_{j1}^b. \quad (2.4.6)$$

Add to the superpotential all such three cycles.

As a simple example of this procedure we consider the A_3 quiver of section 2.3.6 shown on the left and its mutation at node 1 shown on the right.



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

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

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

$$\frac{\partial \tilde{W}}{\partial \tilde{B}_{ij}^a} = 0. \quad (2.4.8)$$

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.

$$(2) \longrightarrow (1) \longrightarrow (3) \quad (2.4.9)$$

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 the vast majority of BPS quivers related to quantum field theories that we discuss in this paper the potential will turn out to be non-degenerate. However exceptions to this general rule do arise. For example the quiver for the \mathcal{T}_2 theory defined by a free trifundamental half-hypermultiplet of a flavor group $SU(2) \times SU(2) \times SU(2)$ involves a quiver with canceling arrows and a potential which is too degenerate to integrate out all the associated by bifundamental fields. In the following unless otherwise stated we avoid this complication and assume that all of our quivers involve non-degenerate superpotentials. However, even when this is not the case one may still apply the mutation rules written above. Mutation at a node supporting a pair of canceling arrows then results in adjoint fields at the mutated node.

2.4.2 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 2.3.6. 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 2.3.6. 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 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


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 $Z_i + Z_j$.

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

equivalent with directed arrows. For example we define the expression

$$Q_i \longrightarrow \tilde{Q}_j , \quad (2.4.10)$$

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.



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

2.4.3 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. In this subsection we fill in this gap.⁶ 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.

⁶The 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 .



Let us describe the leftmost bound state supported by these two node quivers. In the case on the right of (2.4.12), 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 (2.4.12), where γ_1 appears as a source, bound states can exist. We consider a general representation of the form

$$\begin{array}{ccc}
 & B_1 & \\
 & \swarrow \downarrow \searrow & \\
 \mathbb{C}^n & & \mathbb{C}^m
 \end{array} \quad (2.4.13)$$

B_1
 B_2
 \vdots
 B_{k_i}

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

$$\begin{array}{ccc}
 \mathbb{C}^n & \xrightarrow{B_j} & \mathbb{C}^m \\
 \uparrow & & \uparrow \\
 \mathbb{C} & \xrightarrow{0} & 0
 \end{array} \quad (2.4.14)$$

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

$$\text{dimension} \left(\bigcap_{j=1}^{k_i} \ker(B_j) \right) \geq n - k_i m. \quad (2.4.15)$$

And so in particular when the right-hand side of the above is positive, the subrepresentation (2.4.14) exists and hence the bound state is unstable. Thus we learn that

stability requires

$$\frac{n}{m} \leq k_i. \quad (2.4.16)$$

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

$$\gamma_i + k_i \gamma_1. \quad (2.4.17)$$

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

$$\gamma_R = n\gamma_1 + \sum_{\gamma_i \circ \gamma_1 > 0} m_i \gamma_i + \sum_{\gamma_j \circ \gamma_1 \leq 0} l_j \gamma_j \quad (2.4.18)$$

Let us focus in on the representation R near the node γ_1 . There are now many nodes connected to the node γ_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 (2.4.14).

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 (2.4.15) 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

$$\text{dimension} \left(\bigcap_{\gamma_i \circ \gamma_1 > 0} \bigcap_{j=1}^{k_i} \ker(B_j) \right) \geq n - \sum_{\gamma_i \circ \gamma_1 > 0} k_i m_i. \quad (2.4.19)$$

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

$$n \leq \sum_{\gamma_i \circ \gamma_1 > 0} k_i m_i. \quad (2.4.20)$$

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

$$\begin{aligned}\arg(\mathcal{Z}(R)) &= \arg \left(nZ_1 + \sum_{\gamma_i \circ \gamma_1 > 0} m_i Z_i + \sum_{\gamma_j \circ \gamma_1 \leq 0} l_j Z_j \right) \quad (2.4.21) \\ &\leq \arg \left(\sum_{\gamma_i \circ \gamma_1 > 0} m_i (k_i Z_1 + Z_i) + \sum_{\gamma_j \circ \gamma_1 \leq 0} l_j Z_j \right).\end{aligned}$$

But the final expression in (2.4.21) 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:

$$\tilde{\gamma}_1 = -\gamma_1 \quad (2.4.22)$$

$$\tilde{\gamma}_j = \begin{cases} \gamma_j + (\gamma_j \circ \gamma_1) \gamma_1 & \text{if } \gamma_j \circ \gamma_1 > 0 \\ \gamma_j & \text{if } \gamma_j \circ \gamma_1 \leq 0 \end{cases} \quad (2.4.23)$$

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.

2.5 The Mutation Method

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 2.3.4. 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 2.3.4, 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} .⁷ 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

⁷From now on we will abuse verbiage slightly and simply say that “ γ_1 is left-most.”

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 $\mathcal{H} \rightarrow e^{2\pi i}\mathcal{H}$.⁸ The key point is every 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.⁹ If we record every state γ we mutate on as \mathcal{H} is rotated, and then add all antiparticles $-\gamma$,

⁸Recall that for a given choice of \mathcal{H} , the quiver description is actually unique - there is a unique positive integral 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.

⁹By a similar argument as above, the final quiver will have nodes $-\gamma_i$.

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 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,

●. 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.¹⁰ 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

¹⁰If you try to label nodes and keep track of them, which the drawings may subliminally suggest you do, in general you will return to $(-1) \times$ permutation.

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 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 $\mathcal{N} = 2^* \text{ } SU(2)$ (the mass deformed $\mathcal{N} = 4$ theory), it appears that the method

isn't sophisticated enough to exhaust the spectrum.¹¹

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 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 concluding this discussion, 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

¹¹Of course we can always produce some arbitrarily large subset of states of the theory by mutating until exhaustion (of the mutator, that is).

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 [3]. In the next chapter we study the application of these techniques to the class of complete theories. In the more general case of non-complete theories, existence of the desired changes in the

physical moduli space must be checked by hand.

2.5.1 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 [3, 32, 38]. 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.

$$Y_i Y_j = q^{-\gamma_i \circ \gamma_j} Y_j Y_i, \quad (2.5.1)$$

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:

$$Y_\gamma \equiv N[Y_1^{n_1} \cdot Y_2^{n_2} \cdots Y_m^{n_m}]. \quad (2.5.2)$$

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

$$Y_\alpha \rightarrow \Psi(Y_\gamma, q) Y_\alpha \Psi(Y_\gamma, q)^{-1}. \quad (2.5.3)$$

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

$$\mu_k(Y_i) = \begin{cases} Y_k^{-1} & \text{if } i = k \\ Y_i & \text{if } \gamma_i \circ \gamma_k > 0 \\ Y_{\gamma_i + (\gamma_k \circ \gamma_k) \gamma_k} & \text{if } \gamma_i \circ \gamma_k < 0 \end{cases} \quad (2.5.4)$$

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

$$\mathcal{Q}_k = \text{Ad}(\Psi(Y_k, q)) \circ \mu_k. \quad (2.5.5)$$

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

$$\mu_{k(s)} \cdots \mu_{k(2)} \mu_{k(1)} Q = Q. \quad (2.5.6)$$

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

$$\mathcal{Q}_{k(s)} \cdots \mathcal{Q}_{k(2)} \mathcal{Q}_{k(1)}. \quad (2.5.7)$$

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 (2.5.6) 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 [18, 39, 40]. In this way we recover the full KS monodromy operator from ordered mutation sequences.

Chapter 3

Complete Theories

3.1 Introduction

The goal of this chapter is to explore applications to complete theories, where the framework described in chapter 2 is most powerful. 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 [41–44]. As determined in [3], such examples have an additional remarkable property: their BPS spectrum can be encoded by a *BPS quiver* [3,5,7–11,18,38]. This dramatically simplifies the problem of finding BPS states. In place of some tedious weak coupling physics or intractable strong coupling dynamics, the BPS spectrum is

governed by a quantum mechanics problem encoded in this quiver.

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 and, when possible, to solve the associated quantum mechanics problem and determine the BPS spectrum. We begin in section 3.2 with a brief review of the classification of complete theories given in [3]. The classification makes use of the notion of finite mutation-type quivers, which we will also explore here.

To accomplish our first goal of determining the BPS quiver, in section 3.3 we reconstruct these complete theories via geometric engineering in type IIB string theory on a local Calabi-Yau threefold. [24, 45, 46]. 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. [5] 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.

Having accomplished our first task, in section 3.4 we focus on determining what

structure the resulting spectra possess and on computing some explicit examples. The problem of determining the BPS spectrum is computationally most tractable in a chamber where there are finitely many BPS states; we restrict our attention to this case. One of our most interesting results is a determination of an infinite class of theories which have such a finite chamber. Indeed, as we prove in section 3.4, theories with finite chambers include all asymptotically free examples, Argyres-Douglas models, and theories defined by punctured spheres and tori. The latter examples are particularly interesting: They are conformal field theories where the only breaking of scale invariance is that introduced by adding bare mass terms. In all such cases the spectrum can be calculated explicitly and algorithmically using the techniques developed herein.

Finally, in section 3.5 we undertake a brief investigation of complete theories with BPS quivers which do not come from Gaiotto type constructions. In [3] 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 Classification of Complete Theories

Perhaps one of the most unexpected applications of quiver mathematics in 4d $\mathcal{N} = 2$ gauge theory was the classification of complete theories given by Cecotti and Vafa [3]. This classification followed through two key steps; first, the observation of a 4d/2d correspondence, which connected BPS quivers of 4d gauge theories to soliton quivers of 2d Landau Ginzburg theories; second, the connection to finite mutation-

type quivers.

3.2.1 4d/2d Correspondence

The easiest approach to this concept is by considering the type IIB construction of 4d $\mathcal{N} = 2$ gauge theories. These are geometrically engineered from the space-time point of view by compactifying 10d string theory on a 6d Calabi-Yau, which may be defined as the vanishing locus of some polynomial $W(z_i)$ in four complex dimensional space. On the other hand, we may consider the worldsheet definition of perturbative strings in the context of Calabi-Yau compactification. In this case, the 2d $\mathcal{N} = 2$ linear sigma model of string theory is replaced with a supersymmetric Landau-Ginzburg theory of four chiral superfields, whose superpotential is simply given by the defining equation of the Calabi-Yau manifold, W .

For our purposes, the substance of the correspondence comes from identifying the data of the 4d BPS quiver with the data of the 2d solitons. For the Landau-Ginzburg, classical vacua occur at points with z' satisfying

$$\left. \frac{\partial W}{\partial z_i} \right|_{z'} = 0. \quad (3.2.1)$$

Around such points we expand

$$W = W_0 + \sum_i z_i^2. \quad (3.2.2)$$

Thus we see that the vanishing locus of W is approximated near this point as

$$\sum_i z_i^2 = -W_0, \quad (3.2.3)$$

a 6-sphere, containing a special Lagrangian 3-sphere given by

$$\sum_i \text{Re}(z_i)^2 = -W_0, \quad (3.2.4)$$

Recall that nodes of the 4d quiver correspond to some basis of non-trivial special Lagrangian 3-cycles in the Calabi-Yau, which we can now identify with vacua of this 2d theory. Meanwhile, through a careful analysis of the 2d BPS equations, one finds that the number of solitons (signed with a supersymmetric index) running between any two vacua correspond precisely to the signed intersection number between the corresponding special Lagrangian 3-cycles in the Calabi-Yau. So we find that the 4d BPS quiver is neatly encoded in the soliton data of the associated 2d theory. It is important to note that the 4d/2d correspondence is a broader statement than this. In particular, it is expected that there should be a matching of all 4d and 2d $\mathcal{N} = 2$ theories, not just those engineered via string theory. Furthermore, the quiver-soliton equivalence we've found here is only a manifestation of the correspondence. It is believed that such a matching persists even, say, for the class of 4d theories that do not admit a BPS quiver.

3.2.2 Classification of 2d $\mathcal{N} = 2$ theories

The soliton structure of 2d theories has been studied in-depth in [14], with many powerful results discussed there. The primary conclusion of that study was a classification of possible 2d soliton structures. It was found that the matrix of vacua-soliton adjacency (that is, the matrix consisting of the number of solitons between any pair of vacua) is constrained to be a solution to a certain Diophantine equation. This classification program was also explored in [14]. Here we only need to reference a single result: for a system with two vacua, there can be either 0, 1, or 2 solitons between the two vacua, but no more.

This conclusion holds very interesting consequences for the structure of BPS quivers. Of course, it trivially rules out almost all two-node quivers as possible BPS quivers. However, by restricting our attention to the subclass of *complete theories*, we can accomplish much more. A theory is denoted complete if it has sufficient parameters (masses, couplings, and moduli) to tune each primitive BPS charge independently. In other words, it must have as many parameters as dimensions of its charge lattice. Consider the example of a gauge theory with matter. The charge lattice has dimension $2r + f$ where r is the rank of the gauge group (each gauge charge has an associated electric and magnetic charge), and f is the rank of the flavor group. The tunable parameters are the f masses of the matter fields, the n coupling constants of the gauge theory (where $n < r$ is the number of simple factors in the total gauge group), and the r Coulomb branch parameters. Therefore, the only gauge theories that are complete are those with $n = r$, that is $SU(2)^r$ gauge theories.

The condition of completeness allows us to manipulate these theories in very

useful ways. We can tune any one BPS charge without affecting the others at all. In fact, we can tune any particular BPS mass to infinity, and thereby decouple it completely. In the BPS quiver, this means we can select any single node and delete it. We can also tune the phases of any particular node, and thus induce any specific mutation we choose. It is important to note that these possibilities are unique to complete theories. In a general, non-complete theory, as we try to tune BPS charges, we may find that the other charges move in some compensating way, preventing some particular mutation or decoupling. In fact, this is known to happen in many cases.

3.2.3 Quivers of Finite Mutation-Type

Armed with these tools, we can identify some features that quivers of complete theories must exhibit. Such a quiver must have no mutation form that contains a triple (or higher) arrow. If such a mutation form existed, we could immediately decouple all nodes except the two adjacent to the triple arrow, and then by the 4d/2d correspondence, we have arrived at a forbidden soliton structure for the 2d theory. Actually this feature implies something much more interesting: all quivers of complete theories must have only finitely many distinct quivers in their associated mutation class. This follows trivially from the previous observation; a quiver has a fixed set of nodes, and if the number of arrows between any two nodes is constrained to be less than k , then there are only finitely many possible combinations. We say that such quivers are finite mutation-type.

Finite mutation-type quivers were fully classified by mathematicians, [47]. The classification consists of the following:

- Quivers with at most two nodes;
- Quivers associated to ideal triangulations of surfaces with interior punctures, boundaries, and marked points on the boundary;
- Quivers mutation equivalent to the nine E-type Dynkin diagrams
 - finite: $E_6, E_7, E_8,$
 - affine: $\hat{E}_6, \hat{E}_7, \hat{E}_8,$
 - elliptic: $\hat{\hat{E}}_6, \hat{\hat{E}}_7, \hat{\hat{E}}_8;$
- Two Derksen-Owen mutation classes X_6, X_7 , [48].

The elliptical E-type quivers and Derksen-Owen quivers are shown in Figure 3.15.

For the study of 4d BPS states, we simply discard the quivers of the first type with more than three arrows. The second set of quivers will be the main focus of the rest of this chapter; we will proceed studying the physics of the relevant theories, and discover the associated quivers along the way. The remaining 11 exceptional cases will be discussed in 3.5, without delving into the details of the related physical theories.

The physics of those theories is explored further in ??.

3.3 BPS Quivers of 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

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.

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.

3.3.1 BPS States from Geometric Engineering

We will construct these models using geometric engineering [24, 45, 46, 49] 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

$$K_{\mathcal{C}} \oplus K_{\mathcal{C}} \oplus K_{\mathcal{C}} \rightarrow \mathcal{C}, \quad (3.3.1)$$

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

$$\phi'(x') = \phi(x) \left(\frac{dx}{dx'} \right)^2. \quad (3.3.2)$$

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

$$\phi(x) \sim \frac{1}{x^{k_i+2}} dx^2 + \text{less singular terms.} \quad (3.3.3)$$

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 3.3.3.²

Given this data our Calabi-Yau threefold is then defined by introducing local coordinates (u, v, y) on the fiber of the vector bundle (3.3.1) and solving the following equation

$$uv = y^2 - \phi(x). \quad (3.3.4)$$

The associated holomorphic three-form Ω is given by

$$\Omega = \frac{du}{u} \wedge dy \wedge dx. \quad (3.3.5)$$

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.

$$\Sigma = \{(x, y) | y^2 = \phi(x)\}; \quad \lambda = \int_{S^2(x)} \Omega = y dx = \sqrt{\phi}. \quad (3.3.6)$$

By varying the quadratic differential we obtain a family of Seiberg-Witten curves,

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.

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 (3.3.3).

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 [41], 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 (3.3.4). 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 [25, 50]. These are compact lagrangian three-manifolds N on which the holomorphic three-form has a constant phase

$$\Omega|_N = e^{i\theta}|\Omega|. \quad (3.3.7)$$

The central charge of such a brane is given by

$$\mathcal{Z}_u(N) = \int_N \Omega, \quad (3.3.8)$$

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

$$\theta = \arg \mathcal{Z}(N). \quad (3.3.9)$$

Now one of the key observations of [24] is that, in the geometries described by (3.3.4), 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 (3.3.1) 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 .

The shape of η in \mathcal{C} is constrained by the special Lagrangian condition (3.3.7) on N . Explicitly if we let $t \in \mathbb{R}$ parametrize η then the condition of constant phase Ω




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.

reduces to

$$\sqrt{\phi}|_\eta = e^{i\theta} dt. \quad (3.3.10)$$

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 (3.3.10).

3.3.2 BPS States from M-theory

The above conclusions regarding BPS states for complete theories can be obtained by an alternative means, beginning instead with the M5 brane world-volume theory. We consider wrapping stacks of two M5 branes on a Riemann surface to produce an $\mathcal{N} = 2$ 4d theory in the IR, following [41, 42]. We briefly review the salient features of this M5-brane theory, and the structure of BPS states from this perspective.

We begin with the same Riemann surface \mathcal{C} along with its canonical line bundle $K_{\mathcal{C}}$. We also take the same quadratic differential ϕ , and consider a stack of two parallel M5 branes on the surface \mathcal{C} , intersecting at zeroes of the quadratic differential ϕ . More precisely, we take the two M5 branes to be embedded in the total space of $K_{\mathcal{C}}$, at the loci given by

$$y^2 = \phi(x). \quad (3.3.11)$$

This equation describes a surface \mathcal{C}' given as a double-cover of \mathcal{C} . Over a generic point $x \in \mathcal{C}$, there are positive and negative roots of this equation, giving the position of the two M5 branes, while at zeros of ϕ , the two roots degenerate and the two branes intersect.

A standard twisting procedure ensures that we preserve 8 supersymmetries from the 6d (2,0) supersymmetric theory of the M5 branes, yielding a 4d $\mathcal{N} = 2$ gauge theory as desired. Further details of the construction, and its correspondence to type II setups are discussed in [42]. For our purposes, we would like to understand the BPS objects in this M5 brane theory.

The only natural candidates in this case are M2 branes ending on the wrapped stack of M5 branes. The M2 branes are discs with boundary on \mathcal{C}' , and volume-

minimizing within homology. Again as in Figure 3.1, there are only two possible forms for the M2 projected to \mathcal{C} : a segment interpolating between two zeros of ϕ , or a loop encircling a singularity in ϕ . The resulting M2 brane geometries are that of a disc and a cylinder. The condition of volume-minimization allows us to restrict our attention to M2 branes that fill the y direction perpendicular to \mathcal{C} ; thus, the brane is fully specified by some trajectory η on \mathcal{C} . Computing the volume, we find

$$V = \int_{\eta} |2ydx| = \int_{\eta} |2\sqrt{\phi(x)}| = \int_{\eta} 2|\lambda|. \quad (3.3.12)$$

However $\int_{\eta} \lambda$ is a homological invariant, with

$$\int_{\eta} \lambda \leq \int_{\eta} |\lambda|. \quad (3.3.13)$$

Thus, the volume minimization condition is precisely equivalent to (3.3.10),

$$\lambda|_{\eta} = \sqrt{\phi}|_{\eta} = e^{i\theta}dt, \quad (3.3.14)$$

which requires that λ have constant phase along η . We have reduced the question of BPS states in the M5 brane construction of our 4d theories to the same mathematical problem found in our type IIB setup.

3.3.3 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 (3.3.10). This is a problem which is well-studied in mathematics [51] and has received much attention in the present physical context [25, 26, 52–54]. 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 (3.3.10) takes the local form

$$\sqrt{w}dw = e^{i\theta}dt \implies w(t) = \left(\frac{3}{2}e^{i\theta}t + w_0^{3/2}\right)^{2/3}. \quad (3.3.15)$$

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 given by

$$\frac{dw}{w^{1+k_i/2}} = e^{i\theta}dt. \quad (3.3.16)$$




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.

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 [41]. 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.

$$m \frac{dw}{w} = e^{i\theta} dt. \quad (3.3.17)$$

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 (3.3.17).

The solution with initial condition w_o takes the form

$$w(t) = w_o \exp(m^{-1}e^{i\theta}t). \quad (3.3.18)$$

Assume that the BPS angle θ has been chosen so that $m^{-1}e^{i\theta}$ is not purely imaginary. Then the solution (3.3.18) 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 (3.3.16):

$$w(t) = \left(\frac{-2e^{i\theta}}{k_i} t + \frac{1}{w_o^{k_i/2}} \right)^{-2/k_i}. \quad (3.3.19)$$

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

$$\sqrt{\phi} = e^{i\theta} dt. \quad (3.3.20)$$

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.

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




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.

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 3.3.6.

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

$$\sqrt{\phi} = e^{i\theta} dt. \quad (3.3.21)$$

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.

3.3.4 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 [28]. 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 [3] 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 (3.3.10). 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 [28], 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 [3]. 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 3.5 [47].

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

$$(\mathcal{C}, \theta) \longrightarrow Q_\theta = \text{BPS Quiver.} \quad (3.3.22)$$

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 [28] 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 [26]. 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 [28], 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 [55] 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. As mentioned in section 2.3.4, 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

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.

potential for these theories does indeed generate all possible gauge invariant mass terms and therefore simplifies the resulting quivers considerably.

3.3.5 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 [56–58]. 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 2.4, 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 [56]. 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.




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} .

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 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 the it automatically ensures that all of our superpotentials will be compatible with mutation. That is, just as in equation (3.3.22), 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 [56] 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 2.4 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

$$|\mathcal{Z}(\delta)| \longrightarrow \infty. \quad (3.3.23)$$

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 triangu-




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 .

lations 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 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 3.3.4. 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, 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.

3.3.6 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 [28] 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 3.3.4 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 $B_{\alpha_1\alpha_2} \dots B_{\alpha_{n-1}\alpha_n}B_{\alpha_n\alpha_1}$.

3.3.7 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 [55], 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.



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 2.4. 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.



It is amusing to note that 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.

3.4 Theories with Finite Chambers

In this section we will identify a subset of complete $\mathcal{N} = 2$ theories for which there exists some chamber containing only finitely many BPS states. In particular, we will show that all asymptotically free $SU(2)^n$ gauge theories, Argyres-Douglas models, and conformal theories with genus zero and genus one surfaces and sufficiently many punctures, meet this criterion. Our main motivation for studying theories with finite chambers is that they are especially well-adapted to the mutation method. As described in [55], the mutation method is most straightforward for computing BPS spectra which consist of only finitely many states. Additionally, as was mentioned in

subsection ??, finite chambers have BPS spectra which consist exclusively of multiplicity one hypermultiplets.

Complete theories also have especially well-behaved wall-crossing phenomena. It is a fact that the quiver of any complete theory has at most two arrows between any two nodes.⁴ Consider some wall crossing of two adjacent hypermultiplet states p, q , and choose the half-plane for the quiver such that p is just outside of the half-plane on the left and q is just inside the half-plane. This situation is illustrated in Figure 3.10. The quiver must contain both q and $-p$ as nodes since they form the boundary of the cone of positive states. Since we are studying a complete theory, we must have $|p \circ q| \leq 2$. The hypermultiplet wall-crossing is completely straightforward and explicit for any of the three possibilities.

- $|p \circ q| = 0$: there is no change in the spectrum across the wall,
- $|p \circ q| = 1$: pentagon identity, which gives two states p, q on one side of the wall and three states $p, p + q, q$ on the other side of the wall,
- $|p \circ q| = 2$: $SU(2)$ identity, which gives two states p, q on one side of the wall, and the vector $p + q$ with infinite tower of dyons $(n + 1)p + nq, np + (n + 1)q$ for $n \geq 0$ on the other.

While the hypermultiplet wall-crossings are highly simplified, we should point out

⁴This can be understood via the triangulation construction. Two diagonals can share at most two triangles between them, and therefore the resulting quiver can have at most two arrows between any two nodes.




Figure 3.10: Here we illustrate a choice of half-plane that forces $q, -p$ to be nodes of the quiver, for any arbitrary adjacent hypermultiplet BPS states p, q . The grey region indicates the choice of particle half-plane, \mathcal{H} , while red vectors are BPS charges of particles, and blue vectors are BPS charges of anti-particles

that it is still possible to have wall crossing of vector multiplets in a complete theory.

This may produce some wild behavior involving infinitely many vectors, which is not so explicitly understood.

Of course, for complete theories the central charges for a basis of states can all be varied independently by tuning parameters; thus in principle, all chambers found via the wall-crossings described above should be physically realized in parameter space. Combining the mutation method and the wall crossing formulae above, explicit computation of BPS spectra for any complete theory with a finite chamber is now reduced to a completely algorithmic procedure for a large region of parameter space.

We devote the rest of this section to finding finite chambers of complete theories. The result of this study will produce finite chambers for the following theories:

- Conformal Argyres-Douglas type theories,
- Asymptotically free $SU(2)^n$ gauge theories,

- Conformal $SU(2)^n$ gauge theories with bifundamentals charged under the i th and $i+1$ th $SU(2)$ s for $i = 1, \dots, k$, and 2 additional fundamentals each for the first and last $SU(2)$,
- Conformal $SU(2)^n$ gauge theories with bifundamentals charged under the i th and $i+1$ th $SU(2)$ s for $i = 1, \dots, k$, and a bifundamental charged under the first and last $SU(2)$.

The first two classes of theories arise from surfaces with boundary, which will be the main focus of the abstract arguments to follow. For the third and fourth class, which correspond to boundaryless spheres and tori with arbitrary punctures, some ad hoc techniques are applied to find finite chambers. Of the complete theories associated to Riemann surfaces, we have failed to find finite chambers for boundaryless $g \geq 2$ surfaces.⁵ We note that there is another distinguishing feature of these boundaryless higher genus theories, namely, that they contain some matter fields in half-hypermultiplets, which cannot be given masses. As a result, it is impossible to take various decoupling limits with large masses. It would be interesting to understand if this fact somehow precludes the existence of finite chambers for such theories.

3.4.1 Examples


Before we study the abstract arguments to prove existence of various finite chambers, we will present some explicit examples in this subsection to illustrate the objective of this program. The examples will also illustrate the three classes of theories

⁵Among the exceptional theories, discussed in 3.15, we will find finite chambers for all except one, X_7

which we will explore in this section. Our main tool here is the mutation method. We recall that, when applying the mutation method to complete theories, we are free to simply choose any ordering of central charges we wish. In the examples below, we demonstrate the existence of the finite chamber by providing an ordering of central charges that yields finitely many mutations in the mutation method; completeness guarantees that a corresponding region of parameter space exists.

- *Argyres-Douglas D_4 theory.*

The BPS structure of Argyres-Douglas A_n theories was studied systematically in [25]. There exist analogues of Argyres-Douglas theory associated to ADE Dynkin diagrams, which were studied in [18, 44]. The quivers of these theories are precisely their associated Dynkin diagrams.⁶ Here we study the Argyres-Douglas theory associated to D_4 . The Gaiotto curve of this theory is given by a sphere with one regular puncture, and one puncture with $k = 4$. The resulting surface with boundary and quiver are given below.




It is quite easy to identify a finite chamber for this theory via the mutation

⁶The underlying graph of the quiver, where we ignore orientation of arrows, exactly agrees with the associated Dynkin diagram. It can be checked that all orientations of arrows for such quivers are mutation equivalent.

method. For example, take $\arg \mathcal{Z}(\gamma_1) > \arg \mathcal{Z}(\gamma_3) > \arg \mathcal{Z}(\gamma_2) > \arg \mathcal{Z}(\gamma_4)$; then we mutate on $1, 3, 2, 4$ in that order. This gives a chamber whose BPS stable states are precisely those associated to nodes of this quiver, without any additional bound states. In fact, a chamber with just the nodes themselves always exists for any acyclic quiver: choose an ordering on the nodes so that $\arg \mathcal{Z}(\gamma_i) > \arg \mathcal{Z}(\gamma_j)$ if and only if $\gamma_j \circ \gamma_i \geq 0$. That such a choice is possible is due to the fact that the quiver has no oriented cycles. Then we can see that the resulting chamber will have only its nodes as Π -stable representations, via either the mutation method or directly from quiver representation theory.

- $SU(2)^2$, one bifundamental hypermultiplet.

This theory corresponds to the surface and quiver shown below.



The gauge groups and matter content can be read off directly from the quiver. Each $SU(2)$ corresponds to a two-node $SU(2)$ subquiver, and the bifundamental field corresponds to the node which is attached to each $SU(2)$ in the same way as the third node of the $SU(2), N_f = 1$ quiver. The theory is asymptotically free. A finite chamber can be found via the mutation method; for example, we find the chamber $\{\gamma_3, \gamma_1 + \gamma_3, \gamma_1 + \gamma_2 + \gamma_3, \gamma_1 + \gamma_2 + 2\gamma_3 + \gamma_4, \gamma_5, \gamma_3 + \gamma_4, \gamma_2, \gamma_1 + \gamma_3 + \gamma_4, \gamma_1, \gamma_4\}$,

in decreasing phase order. This follows from the following mutation sequence: $3, 1, 2, 4, 5, 2, 1, 3, 2, 4$. This chamber includes the nodes themselves along with several bound states; because this quiver contains cycles, there is no chamber without bound states, as there was for Argyres-Douglas. Nonetheless, we have exhibited a finite chamber for this theory.

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

The quiver of this theory is associated to a sphere with four regular punctures, and was given along with the appropriate superpotential in subsection 3.3.7. It is well known that this theory is conformal. Again the mutation method yields a finite chamber: in decreasing phase order, $\{\gamma_3, \gamma_4, \gamma_5, \gamma_6, \gamma_1 + \gamma_4 + \gamma_6, \gamma_2 + \gamma_3 + \gamma_5, \gamma_2 + \gamma_3, \gamma_1 + \gamma_4, \gamma_2 + \gamma_5, \gamma_1 + \gamma_6, \gamma_1, \gamma_2\}$. The mutation sequence for this chamber is $3, 4, 5, 6, 1, 2, 3, 4, 5, 6, 1, 2$. This finite chamber is particularly interesting because it occurs in the moduli space of a conformal theory. If we tune to the conformal point, by turning off all the masses of the flavor fields, it is expected that the BPS structure becomes highly intricate, respecting some large conformal duality group. In spite of this, we have exhibited a region of moduli space where the BPS spectrum is very simple, and consists of 12 hypermultiplet states.

These three cases are neatly representative of the types of theories for which we will find finite chambers. As described above, the existence of finite chambers for Argyres-Douglas theories is already clear, since they all correspond to acyclic Dynkin diagrams. The discussion below will extend this to all complete theories associated to surfaces with boundary; this class includes, in particular, Argyres-Douglas theories,

as well as all complete asymptotically free $SU(2)^k$ gauge theories. We will also find finite chambers for the conformal $SU(2)^k$ theories associated to spheres and tori.

3.4.2 Quiver Glueing Rule

Consider two quivers, A, B which separately have finite chambers $\mathcal{G}_A, \mathcal{G}_B$; in each quiver, choose a distinguished node, a, b respectively. We will consider the composite quiver $A \oplus_a^b B$ which is given by drawing one arrow from $a \rightarrow b$. More generally, we might choose several nodes from each quiver, $\{a_i\}, \{b_i\}$ (where we allow repeats in the chosen nodes), and consider the composite quiver $A \oplus_{\{a_i\}}^{\{b_i\}} B$ formed by drawing arrows between pairs of nodes, $a_i \rightarrow b_i$. Note that all arrows must point from A to B . The resulting quiver will contain a finite chamber whose BPS states are precisely the union of the BPS states $\mathcal{G}_A \cup \mathcal{G}_B$. To specify such a chamber, we simply consider the ordering within each quiver A, B to be given by the known finite chambers $\mathcal{G}_A, \mathcal{G}_B$, and in addition we require for any nodes $\alpha \in A, \beta \in B$ we have $\arg \mathcal{Z}(\alpha) < \arg \mathcal{Z}(\beta)$.

The representation theory makes this fact completely transparent. Consider any representation of the composite quiver. It is given by some representations \mathcal{A}, \mathcal{B} respectively of quivers A, B along with a set of maps $\phi_i : V_{a_i} \rightarrow V_{b_i}$ corresponding to the arrows $a_i \rightarrow b_i$. We will denote this rep as $R = (\mathcal{A}, \mathcal{B}, \{\phi_i\})$. Let \mathcal{A}, \mathcal{B} be nonzero. Now we may consider the subrep $S = (0, \mathcal{B}, \{0_i\})$. This is always a valid subrep, as can be seen by the following commutative diagram:

$$\begin{array}{ccc} \mathcal{A} & \xrightarrow{\phi_i} & \mathcal{B} \\ \uparrow 0 & & \uparrow \text{id} \\ 0 & \xrightarrow{0} & \mathcal{B} \end{array}$$

Note that by our choice of chamber, $\arg\mathcal{Z}(S) > \arg\mathcal{Z}(R)$, so that this is automatically a destabilizing subrep. Consequently, any representation that has support on both subquivers A, B will be unstable, leaving only the stable reps of the subquiver A, B separately. This rule can be checked as a simple exercise using the mutation method. Note that we have made no reference to A, B being quivers of complete theories. The glueing rule is completely general and can be applied to any pair of quivers that are known to have finite chambers.

As a first application of the glueing rule, we study acyclic quivers. Any acyclic quiver can be built up by glueing in one-node quivers, one at a time. Simply pick an ordering of the nodes consistent with the arrows - this is possible because the quiver is acyclic. Then we may glue the nodes to each other one-by-one in the given ordering. Since each one node quiver has only the node itself as a BPS state, we can build up a finite chamber which consists only of the nodes of the quiver. This immediately confirms the claim in subsection 3.4.1, and allows us to conclude that all Argyres-Douglas theories have such chambers. In fact, acyclic finite mutation type quivers were classified by [59], and consist precisely of usual ADE and affine $\widehat{A}\widehat{D}\widehat{E}$ Dynkin diagrams. These are the only complete theories containing a minimal chamber in which only the nodes of the quiver are stable BPS states.

3.4.3 Triangulation Glueing Rule

The relation between complete theories and triangulated surfaces allows us to translate the above quiver glueing rule to a glueing rule at the level of the triangulation. First, we define an *augmented quiver* associated to the triangulation of a

surface, in which we include nodes corresponding to the boundary edges in the triangulation, and draw arrows as given by the rules of section 3.3.4, treating boundary edges and interior diagonals on equal footing. The nodes corresponding to boundary edges will be referred to as *augmented nodes*. Then when we glue together two triangulations along their boundaries, the new augmented quiver of the full surface is given by identifying some pair of augmented nodes in the augmented quivers of the two surfaces.

Notice that if the augmented quiver has a finite chamber, then so does the usual, unaugmented quiver: the usual quiver is a subquiver of the augmented one, and the finiteness of a chamber is preserved by taking subquivers. This can be seen via representation theory. Stability for a representation of a subquiver is equivalent to stability for the same rep considered in the full quiver, since in either case we need to study the same set of destabilizing subreps. So the BPS spectrum of a subquiver is just the restriction of the BPS spectrum of the full quiver to states that have support only on the subquiver of interest.

Consider two triangulated surfaces A, B , each with at least one boundary component. We will use the same symbols A, B to denote the associated *augmented* quivers. To achieve the glueing of quivers described above, we consider glueing the two triangulated surfaces along one component of their respective boundary components to two sides of a triangle, as in Figure 3.11. Let us denote by a, b the augmented nodes corresponding to the glued boundary edges of A, B respectively, and let c be the augmented node corresponding to the unglued edge of the triangle. The augmented quiver of the full surface is given in Figure 3.11 as well.




Figure 3.11: General glueing rule for triangulations. A and B indicate surfaces with boundary, glued along one component of their respective boundaries to a triangle. Red lines indicate interior diagonals, which give nodes of the adjacency quiver. Black lines indicate boundary edges which give augmented nodes in the augmented quiver.




Figure 3.12: Mutated form of quiver shown in Figure 3.11, obtained by mutating at node c

Note that c is an augmented node, so that the unaugmented quiver is already a subquiver of $A \oplus_a^b B$. Hence if A and B have finite chambers, then so does the resulting unaugmented composite quiver corresponding to the glueing described. However, in order to induct and continue glueing more pieces to this composite quiver, we would like to check that the *augmented* quiver also has a finite chamber. This will be the case if both A, B have finite chambers, and there is a finite chamber of A (or B) such that no bound state has coefficient of a (resp. b) greater than 1.

To see this, begin by mutating on node c . We find a quiver $A \oplus_a^c \{c\} \oplus_c^b B$ (Fig-

ure 3.12), which has a finite chamber consisting of $\mathcal{G}_A \cup \{c\} \cup \mathcal{G}_B$, with $\arg \mathcal{Z}(b_i) > \arg \mathcal{Z}(c) > \arg \mathcal{Z}(a_j)$ for all $a_i \in A, b_j \in B$, as described above. Now if we do a sequence of wall crossings to let $\arg \mathcal{Z}(c) > \arg \mathcal{Z}(b_j)$, then we will be in a region covered by the quiver form of Figure 3.11. This can be seen by the mutation algorithm: c is now the left-most node, so we mutate at c first, away from the direct sum form in Figure 3.12, resulting Figure 3.11. We then see that we are in a region of moduli space covered by the quiver Figure 3.11. As long as this wall-crossing procedure only goes through pentagon-type crossings, we will only generate finitely many new bound states. Since c only has inner product with b in B , the condition is just that there are no bound states in \mathcal{G}_B with more than one b . A similar argument with inverse mutation yields an analogous conclusion for A .

To reiterate, we have developed a glueing rule for triangulations, depicted in Figure 3.11. The glueing rule provides a finite chamber for the composite triangulated surface, given finite chambers for the two separate triangulated surfaces, subject to an additional mild conditions that there be no bound states of multiple a 's or b 's.

3.4.4 Surfaces with Boundary


In this section we will explore the quiver glueing rule and its implication for triangulations, to attempt to build up a large class of Riemann surfaces whose quivers contain a finite chamber. In fact, we will find that any surface with boundary has a quiver with finite chamber. Recall from [3] that surfaces with boundary correspond to asymptotically free theories along with the conformal Argyres-Douglas theories. Aside from the Argyres-Douglas cases, these theories have negative beta function

because they are constructed by taking certain decoupling limits of the conformal theories that correspond to boundaryless Riemann surfaces.

A surface in this context is characterized completely by its genus g , number of punctures n , and number of boundary components b , along with some number of marked points $k_i \geq 1$ for every boundary component, $i = 1 \dots b$. The k_i are identified with the orders of poles as given in section 3.3.3. In order to build up new surfaces, we will glue triangulated pieces B to some existing surface A with finite chamber, as in Figure 3.11, all while making sure to preserve the finite chamber. Suppose we have some surface $(g, n, b, \{k_i\}_{i=1}^b)$ whose quiver, A , has a finite chamber. There are four types of operations we will need to consider:

- *Add a marked point on the boundary*


An unpunctured triangle glued to boundary component i of the surface A will increase the number of marked points on i by one ($k_i \rightarrow k_i + 1$) and leave the other parameters of the surface unchanged.



On the augmented quivers, this adds an oriented three-node cycle with one node identified with an existing node on the quiver A . This is just the general triangulation glueing described above, in which the surface B is empty and the quiver B is only the node b itself. So this glueing preserves the finite chamber.

- *Add a puncture*


To add a puncture, we take B to be a once-punctured monogon. This takes $n \rightarrow n + 1$, leaving everything else unchanged.



The quiver B is just two copies of the node b . Here we have encountered a self-folded triangle in the triangulation, so we must refer to the extended rules given in the appendix 3.3.6. The quiver has a finite chamber by the general glueing rule.

- *Add a boundary component*

For this we let B be the annulus with one marked point on each boundary component. This glueing adds one boundary with one marked point, and leaves everything else fixed. That is, $b \rightarrow b + 1$ and $k_{b+1} = 1$.




B as a quiver is the $SU(2)$ $N_f = 2$ quiver. By hand, we can check that B has

a finite chamber in which there is no bound state with multiple b 's, using the mutation method. For example, we find a chamber with states in decreasing phase order $\{b, \gamma_3, \gamma_1 + b + \gamma_3, \gamma_2, \gamma_1 + b, \gamma_1 + \gamma_3, \gamma_2\}$. Thus, the full augmented quiver also has a finite chamber.

- *Increase genus*

We may increase the genus of the surface by taking B to be a torus with boundary with one marked point. This gives $g \rightarrow g+1$ with all other parameters fixed.



Note that we have only drawn B , the torus with boundary, which must be glued into the surface A as in Figure 3.11. Again we can check by hand that B contains a finite chamber with no bound states of multiple b 's. For example, the mutation method gives a finite chamber with states in decreasing phase order $\{b, \gamma_3, \gamma_1, \gamma_2 + \gamma_3, \gamma_1 + \gamma_4 + b, \gamma_1 + \gamma_4, \gamma_2, \gamma_4 + b, \gamma_4\}$. So the resulting augmented quiver has a finite chamber.

Finally, we need to check that we have sufficient base cases in order to build up all possible surfaces with boundary. Again we will be parameterizing surfaces as $(g, n, b, \{k_i\})$. The following are the base cases we need: once-punctured monogon $(0, 1, 1, \{1\})$, unpunctured triangle $(0, 0, 1, \{3\})$, annulus with one marked point on

each boundary $(0, 0, 2, \{1, 1\})$, torus with one boundary component and one marked point $(1, 0, 1, \{1\})$. It is straightforward to see any surface not generated by increasing the four parameters $(g, n, b, \{k_i\})$ starting from one of these base cases is either a surface without boundary or a surface that cannot be triangulated. For example, if we try to reduce n in punctured monogon $(0, 1, 1, \{1\})$, we see that the unpunctured monogon, $(0, 0, 1, \{1\})$ cannot be triangulated. Notice that the base cases are precisely the pieces that we used in the glueings above, so we have already checked that the corresponding augmented quivers all contain the desired finite chambers. So we conclude that all surfaces with boundary (and thus all asymptotically free complete theories) have at least one chamber in their parameter space with finitely many states.

Using the glueing rule and the wall-crossing formulae given at the beginning of this section, computing explicit spectra for these theories is now a completely algorithmic process. For any surface with boundary, we take a decomposition into the pieces used above: punctured monogon, unpunctured annulus, and torus with boundary. The pieces should all be glued together using unpunctured triangles as in Figure 3.11. The choice of decomposition will specify the mutation form of the quiver we must study, along with a point of parameter space, fixed by the ordering of central charges compatible with the glueing rule. Now we simply take the union of the finite spectra associated to each of these pieces; this gives the resulting spectrum of the total surface, according to the glueing rule. Finally, we can use wall-crossing formulae to move to other points in parameter space.

3.4.5 Conformal Theories

For surfaces without boundary (that is, conformal theories), there seems to be an essential complication in trying to decompose these quivers using the techniques above. Very generally, quivers for boundaryless surfaces have large cyclic structures that prevent such a decomposition. In particular, no node for a boundaryless surface can be a sink or source; consequently, the quiver glueing rule is of little use.

Nonetheless, some progress has been made in searching for finite chambers using the mutation method. We have extracted a finite chamber for genus $g = 0, 1$ with arbitrary punctures, which we give below. First, we recall some reasoning introduced in [41], which allows us to deduce a Lagrangian description for these theories. For any of these rank 2 Gaiotto-type theories, we can understand the gauge groups and matter contained in the theory as follows. Take a pair-of-pants decomposition of the boundaryless Riemann surface \mathcal{C} . Each pair of pants corresponds to a half-hypermultiplet charged under 3 $SU(2)$'s, where each $SU(2)$ is represented by one of the boundary components of the pair of pants. Each glueing of a pair of pants identifies the corresponding $SU(2)$'s and gauges that $SU(2)$ symmetry. Given a boundaryless surface, one can use this recipe to deduce the gauge group and matter content of the corresponding theory.

Sphere with $n \geq 4$ Punctures

The sphere with $n < 3$ punctures cannot be triangulated; for $n = 3$ punctures, it corresponds to three nodes with no arrows, which yields no interesting structure. The sphere with $n \geq 4$ punctures has a Lagrangian description as an $SU(2)^{n-3}$ theory




Figure 3.13: Triangulation and quiver for the sphere with $n \geq 4$ punctures. The triangulation is drawn on a plane with the point at infinity omitted. Note there are self-folded triangles formed by the interior of a and the exterior of f (see appendix 3.3.6). In both the triangulation and the quiver, the dots indicate repetition of the 3-node structure, $c_i d_i e_i$. The sphere with n punctures has $n - 4$ such pieces, and $3n - 6$ nodes.

with bifundamentals charged under the i th and $i + 1$ th $SU(2)$ s for $i = 1, \dots, n - 4$, and 2 additional fundamentals each for the first and last $SU(2)$.

A triangulation and quiver of a sphere with $n \geq 4$ punctures is given in Figure 3.13. In fact, the Lagrangian description can be read off directly from this quiver, forgetting the surface and triangulation. Each two-node structure $c_i d_i$ is precisely a pure $SU(2)$ subquiver, and thus indicates an independent $SU(2)$. The nodes a, b, f, g appear just as the flavor nodes in subsection 3.3.7, and correspond to flavors charged under the

first and last $SU(2)$. Finally the nodes e_i appear as flavor nodes for two adjacent $SU(2)$ s, and thus correspond to bifundamental flavors. So, we have reconstructed the description of the gauge group and matter given above. This type of reasoning was discussed further in [3].

A finite chamber for $n \geq 4$ is given by the following sequence of states, in decreasing phase order:

$$\begin{aligned}
 & a, b, a + b + c_1, d_1, a + c_1, b + c_1, \\
 & d_1 + e_1, c_1, d_2, e_1 + d_2, c_1 + d_1 + e_1 + c_2, d_1 + e_1 + c_2, e_1, e_1 + c_2, \\
 & \quad \vdots \\
 & d_k + e_k, c_k, d_{k+1}, e_k + d_{k+1}, c_k + d_k + e_k + c_{k+1}, d_k + e_k + c_{k+1}, e_k, e_k + c_{k+1}, \\
 & \quad \vdots \\
 & d_{n-4} + e_{n-4}, c_{n-4}, d_{n-3}, e_{n-4} + d_{n-3}, c_{n-4} + d_{n-4} + e_{n-4} + c_{n-3}, d_{n-4} + e_{n-4} + c_{n-3}, \\
 & e_{n-4}, e_{n-4} + c_{n-3}, \\
 & f + d_{n-3}, g + d_{n-3}, f + g + d_{n-3}, c_{n-3}, f, g
 \end{aligned}$$

which is a chamber with $8n - 20$ states. This can be verified by applying the mutation

method with the following mutations, in order:

$$\begin{aligned}
 & a, b, c_1, d_1, a, b, \\
 & e_1, c_1, d_2, d_1, c_2, c_1, d_2, e_1, \\
 & e_2, d_2, d_3, d_1, c_3, d_2, d_3, e_2, \\
 & \quad \vdots \\
 & e_k, d_k, d_{k+1}, d_1, c_{k+1}, d_k, d_{k+1}, e_k, \\
 & \quad \vdots \\
 & e_{n-4}, d_{n-4}, d_{n-3}, d_1, c_{n-3}, d_{n-4}, d_{n-3}, e_{n-4}, \\
 & f, g, d_1, d_{n-3}, f, g
 \end{aligned}$$

Torus with $n \geq 2$ Punctures

The torus with one puncture is the $\mathcal{N} = 2^*$ theory, which has no finite chamber; this theory is explored further in [55]. The torus with $n \geq 2$ punctures has a Lagrangian description as an $SU(2)^n$ gauge theory with a bifundamental between the i th and $i + 1$ th $SU(2)$ for $i = 1, \dots, n - 1$ and a bifundamental between the first and last $SU(2)$.

A triangulation and quiver for the torus with n punctures is given in Figure 3.14. Again from the triangulation the gauge group and matter content can be directly read off. We have n $SU(2)$ subquivers, giving gauge group $SU(2)^n$, with bifundamental matter arranged cyclically between every adjacent pair of $SU(2)$ s.

A finite chamber for this theory is given by the following sequence of states, in




Figure 3.14: Triangulation and quiver for the torus with $n \geq 2$ punctures. The triangulation is drawn on a rectangle with opposite sides identified. In both the triangulation and the quiver, the dots indicate repetition of the 3-node structure $a_i b_i c_i$. The torus with n punctures has n sets of double arrows, and $3n$ nodes. Note that the two nodes labelled a_n should be identified, producing a quiver with cyclic symmetry.

decreasing phase order:

$$\begin{aligned}
 & a_1, a_1 + b_1, a_1 + b_1 + c_1, 2a_1 + b_1 + c_1 + b_2, c_2, a_1 + b_2, c_1, a_1 + b_1 + b_2, \\
 & a_2 + c_2, b_2, a_2, 2a_2 + b_2 + c_2 + b_3, c_3, a_2 + b_2 + c_2 + b_3, a_2 + c_2 + b_3, a_2 + b_3, \\
 & \quad \vdots \\
 & a_k + c_k, b_k, a_k, 2a_k + b_k + c_k + b_{k+1}, c_{k+1}, a_k + b_k + c_k + b_{k+1}, a_k + c_k + b_{k+1}, a_k + b_{k+1}, \\
 & \quad \vdots \\
 & a_{n-1} + c_{n-1}, b_{n-1}, a_{n-1}, 2a_{n-1} + b_{n-1} + c_{n-1} + b_n, c_n, a_{n-1} + b_{n-1} + c_{n-1} + b_n, \\
 & a_{n-1} + c_{n-1} + b_n, a_{n-1} + b_n, \\
 & a_n + c_n + c_1, b_n, a_n + c_1, 2a_n + b_n + c_n + c_1, b_1, a_n + b_n + c_n, a_n + c_n, a_n
 \end{aligned}$$

which is a chamber with $8n$ states.

This can be verified by applying the mutation method with the following mutations, in order:

$$\begin{aligned}
 & a_1, b_1, c_1, b_2, c_2, c_1, b_1, a_1, a_2, b_2, c_2, b_3, c_3, c_2, b_2, a_2, \dots, a_k, b_k, c_k, b_{k+1}, c_{k+1}, c_k, b_k, a_k, \dots, \\
 & a_n, b_n, c_n, b_1, c_1, c_n, b_n, a_n.
 \end{aligned}$$

3.5 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. As described above in section 3.2, there are 11 additional exception quivers associated to complete theories. These quivers are not the adjacency quivers of any triangulated

surface.

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.⁷

3.5.1 $E_n, \widehat{E}_n, \widehat{\widehat{E}}_n$

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 [18]. These quivers are acyclic, and thus have no superpotential. As described in section 3.4.2, 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 [3]. They are given by glueing linear acyclic quivers to the quiver of $SU(2)$, $N_f = 3$, (see Figure 3.16). 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 3.3.5. 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

⁷After completing the manuscript, we were informed that these potentials (excluding X_7) were independently obtained in [60] from slightly different considerations.




Figure 3.15: The three elliptic E -type Dynkin diagrams oriented as to give finite mutation quivers, and the two Derksen–Owen quivers.




Figure 3.16: 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 $\widehat{\widehat{E}}_n$ quivers, as shown in Fig. 3.15. A decoupling argument indicates that this gives the correct superpotential for studying the $\widehat{\widehat{E}}_n$ quivers.

in Figure 3.16. Since the quivers involved in the glueing (i.e. A_n linear quivers and $SU(2)$, $N_f = 3$) have finite chambers⁸ we conclude that the $\widehat{\widehat{E}}_n$ quivers also have finite chambers.

3.5.2 X_6, X_7

The corresponding theories to the Derksen-Owen quivers were also studied in [3]. 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 3.17. 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.

⁸We 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 3.17: 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 .

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 [48]; 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

3d Superconformal Theories and Mirror Symmetry

4.1 Introduction

In the two preceding chapters, we have gained a great deal of insight along with computational technology by combining different perspectives on the celebrated Seiberg-Witten 4d $\mathcal{N} = 2$ gauge theories. We first followed a careful study of quiver techniques, which were originally motivated from the point of view of D-brane stability in type IIB strings. We then moved on to study the class of complete theories, and noted a powerful connection between our quiver techniques and the M-theory constructions of these same 4d theories [41, 42]. By studying the correspondence between these dual constructions, we were led to the novel results of chapter 3.

One natural direction is to pursue an analogous M5 brane construction in different dimensions, hoping that again the lift to M-theory will yield insight into the

lower-dimensional physics. A conceptual slogan for the program of this chapter is that we will investigate a three-dimensional analog of Seiberg-Witten theory. In the ultraviolet, one may envision an unknown non-Abelian three-dimensional field theory arising from the interacting theory of two M5-branes on \mathbb{R}^3 with suitable boundary conditions at infinity. Moving onto the moduli space of this theory is accomplished geometrically by allowing the pair of M5-branes to fuse together into a single three-manifold M . The long-distance Abelian physics can then be directly extracted from the geometry of M . The situation we have described should be compared with the case of four-dimensional $\mathcal{N} = 2$ theories whose infrared moduli space physics can be extracted from a Seiberg-Witten curve. In that case, charged matter fields are described by BPS states and can be constructed in M-theory from M2-branes. The case of an interacting conformal field theory can arise when the M2-brane particles become massless and the Seiberg-Witten curve develops a singularity.

One would hope to use the results of the previous chapters in a concrete way, rather than merely as a source of conceptual inspiration. To do so, we consider stacks of two M5 branes on 3-manifolds of the form $\Sigma_t \times \mathbb{R}_t$, where the Riemann surface Σ varies in complex structure along the line parameterized by t . Then, after brane recombination, we have a single M5 brane sitting on a double cover of this flow manifold. These examples are closely connected to the four-dimensional quantum field theories of chapter 3. At a fixed value of t , the situation is that of an M5-brane on Σ whose resulting IR physics and BPS spectra were studied in detail in the previous chapter. As t varies, this field theory moves in its parameter space and hence describes a kind of domain wall in four dimensions. When equipped with suitable

boundary conditions, this geometry can engineer a three-dimensional $\mathcal{N} = 2$ theory.

In the context of such examples, one may utilize the machinery of Seiberg-Witten theory and BPS state counting to determine the resulting three-dimensional physics. When the variation of Σ takes a particularly natural form, known as *R-flow*, the spectrum of three-dimensional chiral multiplets is in one-to-one correspondence with the BPS states of the underlying four-dimensional model in a particular chamber. The mass of the chiral multiplet is set by the minimal mass attained by the associated BPS state over the flow in moduli space. Additionally, as the moduli of the four-dimensional theory are varied, one may cross walls of marginal stability and hence find distinct spectra of chiral multiplets in three-dimensions. Remarkably, the resulting three-dimensional theories are mirror symmetric. In this way, the geometry provides a striking confluence between two fundamental quantum phenomena: wall crossing of BPS states, and mirror symmetry.

Having made these observations, we can generalize our exploration to a broader class of 3-manifolds, and draw connections with geometric structures. The primary relevant feature of the flow construction above is that branch points of the original Riemann surface Σ are stretched out into branching lines. When we introduce some fixed boundary conditions on the flow manifold, these branch lines will connect into some arbitrary tangle of branch lines sitting in the flow manifold. A tangle is a mathematical generalization of a knot that allows for fixed open ends. After brane recombination, we are left with a double cover of this flow manifold, branched over this tangle. Thus, for the manifold where the infrared M5-brane resides, we take an arbitrary double cover of \mathbb{R}^3 branched along a tangle. The reduction of the theory

of a single M5-brane along M will result in the three-dimensional quantum field theories under investigation. The simplest class of examples are associated to non-singular tangles. In this situation M is a smooth manifold and a single M5-brane on M constructs a free Abelian $\mathcal{N} = 2$ Chern-Simons theory in the macroscopic dimensions. Light matter, appearing in chiral multiplets in three dimensions, arises in the theory from M2-brane discs which end along M .

To generate non-trivial interacting conformal theories, we must relax some constraints of this construction. If we allow the branching lines to actually intersect, then the light matter arising from the associated M2 disc will become massless. We describe these intersecting branching lines geometrically as a *singular tangle*. In this situation, we generate a non-trivial interacting conformal field theory with massless matter. We would like to uncover the relationship between three-dimensional $\mathcal{N} = 2$ supersymmetric conformal field theories and singular tangles. Examples of singular tangles are illustrated in Figure 4.1.




Figure 4.1: Examples of tangles and their singularities. In (a) a tangle in the \mathbb{R}^3 . In (b) a singular tangle where the strands have merged at various points. The singularities are modeled in quantum field theory by charged massless matter.

An important feature of the constructions carried out in this chapter, familiar from many constructions of field theories by branes, is that non-trivial quantum properties of field theories are mapped to simpler geometric properties of the compactification

manifold. In the case of $\mathcal{N} = 2$ Abelian Chern-Simons matter theories the quantum features which are apparent in geometry are the following.

- $Sp(2F, \mathbb{Z})$ Theory Multiplets:

The set of three dimensional theories with $\mathcal{N} = 2$ supersymmetry and $U(1)^F$ flavor symmetry is naturally acted on by the group $Sp(2F, \mathbb{Z})$ [61, 62]. This group does not act by dualities. It provides us with a simple procedure for building complicated theories out of simpler ones by a sequence of shifts in Chern-Simons levels and gauging operations.

- Anomalies:

In three dimensions, charged chiral multiplets have non-trivial parity anomalies. This means that upon integrating out a massive chiral field the effective Chern-Simons levels are shifted by half-integral amounts [63].

- Dualities:

Three dimensional $\mathcal{N} = 2$ conformal field theories enjoy mirror symmetry dualities. Thus, distinct $\mathcal{N} = 2$ Abelian Chern-Simons matter theories may flow in the infrared to the same conformal field theory. In the case of three-dimensional Abelian Chern-Simons matter theories, there are essentially three building block mirror symmetries that we may compose to engineer more complicated dualities:

- Equivalences amongst pure CS theories. These theories are free and characterized by a matrix of integral levels K . It may happen that two distinct classical theories given by matrices K_1 and K_2 nevertheless give rise

to equivalent correlation functions and hence are quantum mechanically equivalent.

- Gauged $U(1)$ at level 1/2 with a charge one chiral multiplet is mirror to the theory of a free chiral multiplet [62].
- Super-QED with one flavor of electron is mirror to a theory of three chiral multiplets, no gauge symmetry, and a cubic superpotential [64, 65].

One way non-trivial dualities appear stems from the fact that the M5-brane theory reduced on M does not have a preferred classical Lagrangian. To obtain a Lagrangian description of the dynamics requires additional choices. In our context such a choice is a Seifert surface, which is a Riemann surface with boundary the given tangle. For any given tangle there exist infinitely many distinct choices of Seifert surfaces each of which corresponds to a distinct equivalent Lagrangian description of the physics. This fact is closely analogous to the choice of triangulation appearing in the approach of [66] for studying the same theories, as well as the choice of pants decomposition required to provide a Lagrangian description of M5-branes on Riemann surfaces [41].

Throughout this chapter, our discussion of duality will be guided by a particular invariant of the infrared conformal field theory, the squashed three-sphere partition function

$$\mathcal{Z}_b(x_1, \dots, x_F). \quad (4.1.1)$$

This is a complex-valued function of a squashing parameter b (which we frequently suppress in notation) as well as F chemical potentials x_i . It is an invariant of a field theory with prescribed couplings to $U(1)^F$ background flavor fields. This partition

function gives us a strong test for two theories to be mirror and as such it is useful to build into the formalism techniques for computing \mathcal{Z} .

One method of explicit computation is provided by supersymmetric localization formulas. At the classical level, an Abelian Chern-Simons matter theory coupled to background flavor fields is determined by the following data:

- Integers G and F specifying that the theory in question has a $U(1)^G$ gauge group and a $U(1)^F$ flavor group,
- An $(G + F) \times (G + F)$ matrix of Chern-Simons levels, K
- A set of chiral multiplets Φ_a , with $(G + F)$ dimensional charge vectors q_a ,
- A superpotential $W(\Phi_a)$; a holomorphic function of chiral fields.

Given such data, the three-sphere partition function for the infrared conformal field theory can be presented as a finite dimensional integral¹ [67, 68]

$$\mathcal{Z}(x_i) = \int d^G y \exp \left(-\pi i (y \cdot x) K \begin{pmatrix} y \\ x \end{pmatrix} \right) \prod_a E(q_a \cdot (y \cdot x)). \quad (4.1.2)$$

In the above, $E(x)$ denotes a certain transcendental function, the so-called non-compact quantum dilogarithm, which will be discussed in detail in section 4.3. The superpotential W enters the discussion only in so far as it restricts the flavor symmetries of the theory. The real integration variables y appearing in the formula can

¹In the following formula, certain details about R-charge assignments are suppressed. These will be dealt with more fully in section 4.3.

be interpreted as parameterizing fluctuations of the real scalars in the $\mathcal{N} = 2$ vector multiplets.

We will be interested in computation of \mathcal{Z} up to multiplication by an overall phase independent of all flavor variables. This means in particular that throughout this work we will ignore all framing anomalies of Chern-Simons terms. We will see that the partition function in (4.1.2) can be usefully viewed as a wavefunction in a certain finite dimensional quantum mechanics and develop this interpretation throughout.

Although our motivation, as outlined in this introduction, follows from the R-flow constructions, we will proceed from a more general standpoint, and later explore the details of the specific R-flow setup that connects to the material in previous chapters. Indeed, R-flow is rather one implementation of the more general 3d constructions that we shall study. In section 4.2 we explain how free Abelian Chern-Simons theories arise from tangles, and how their partition functions are encoded in a simple quantum mechanical setup. In section 4.3 we show how the data of massless chiral fields is encoded in terms of singular tangles where branch loci collide. Each such singularity can be geometrically resolved in one of three ways, matching the expected deformations of the field theory. Upon fixing a Seifert surface, a surface with boundary on the tangle, we are able to extract a Lagrangian description of the theory associated to the singular tangle including superpotential couplings. In section 4.4 we generalize to arbitrary singular tangles, and explore physical redundancy in the geometry. As a consequence of mirror symmetries, distinct singular tangles can give rise to the same superconformal theory. These equivalences on field theories can be described geometrically by introducing a set of generalized Reidemeister moves acting

on singular tangles. On deforming away from the critical point by activating relevant deformations of the field theory, we find that the generalized Reidemeister moves resolve to the ordinary Reidemeister moves familiar from elementary knot theory. The appearance of Reidemeister moves clarifies the relationship between quantum dilogarithm functions and braids first observed by [69]. In section 4.5 we describe how three-dimensional mirror symmetries can be understood from the perspective of four-dimensional $\mathcal{N} = 2$ parent theories via R-flow. Finally, in section 4.6 we describe three-dimensional $U(1)$ SQED with arbitrary N_f .

4.2 Abelian Chern-Simons Theory and Tangles

In this section we explore the simplest class of examples: Abelian $\mathcal{N} = 2$ Chern-Simons theories without matter fields. Such theories are free and hence of course conformal. We find that such models are usefully constructed via reduction of the M5-brane on a non-singular manifold which is conveniently viewed as a double cover of \mathbb{R}^3 branched over a tangle, and describe the necessary geometric technology for elucidating their structure. In addition we describe a finite dimensional quantum mechanical framework for evaluating their partition functions. Throughout we will study the theories with $U(1)^F$ flavor symmetries and couple them to F non-dynamical vector multiplets. The set of such theories is acted upon by $Sp(2F, \mathbb{Z})$ and we describe this action from various points of view.

4.2.1 Chern-Simons Actions, $Sp(2F, \mathbb{Z})$, and Quantum Mechanics

Consider a classical $\mathcal{N} = 2$ Abelian Chern-Simons theory. Let G denote the number of $U(1)$ gauge groups, and F the number of $U(1)$ flavor groups.² The Lagrangian of the theory coupled to F background vector multiplets is specified by a $(G + F) \times (G + F)$, symmetric matrix of levels

$$K = \begin{pmatrix} k_G & k_M \\ k_M^T & k_F \end{pmatrix}, \quad k_G = k_G^T, k_F = k_F^T. \quad (4.2.1)$$

Here, k_G denotes the ordinary Chern-Simons levels of the $U(1)^G$ gauge group, k_M indicates the $G \times F$ matrix of mixed gauge-flavor levels, and k_F the $F \times F$ matrix of flavor levels. The action for the theory is

$$\sum_{\alpha\beta} \frac{K_{\alpha\beta}}{4\pi} \int d^3x \ A_\alpha \wedge dA_\beta + \dots, \quad \alpha, \beta = 1, 2, \dots, G + F. \quad (4.2.2)$$

Where in the above the terms “ \dots ” indicate the supersymmetrization of the Chern-Simons Lagrangian. $K_{\alpha\beta}$ is integrally quantized with minimal unit one. The first G vector multiplets are dynamical variables in the path integral while the last F are non-dynamical background fields.³

²Here and in the following *flavor* symmetries refer to non- R symmetries except when explicitly indicated otherwise.

³The normalization of the Chern-Simons levels appearing in (4.2.2) indicates that these are *spin* Chern-Simons theories [70] whose definition depends on a choice of spin structure on spacetime. Since all the models we consider are supersymmetric and hence contain dynamical fermions, this is no restriction.

It is worthwhile to note that one might naively think that the matrix K does not completely specify an $\mathcal{N} = 2$ Chern-Simons theory. Indeed, since such theories are conformal they contain a distinguished flavor symmetry, $U(1)_R$, whose associated conserved current appears in the same supersymmetry multiplet as the energy-momentum tensor. One might therefore contemplate Chern-Simons couplings involving background $U(1)_R$ gauge fields. However, such terms while supersymmetric violate conformal invariance. Thus, as our interest here is superconformal field theories, we are justified in ignoring these couplings.

Already in this simple context of Abelian Chern-Simons theory, we can see the action of $Sp(2F, \mathbb{Z})$ specified as operations on the level matrix K defined in equation (4.2.1). For later convenience, it is useful to use a slightly unconventional form of the symplectic matrix J

$$J = \begin{pmatrix} 0 & 1 & 0 & 0 & \cdots & 0 & 0 \\ -1 & 0 & 1 & 0 & \cdots & 0 & 0 \\ 0 & -1 & 0 & 1 & \cdots & 0 & 0 \\ \vdots & \vdots & \ddots & \ddots & \ddots & \vdots & \vdots \\ 0 & 0 & 0 & 0 & \cdots & 0 & 1 \\ 0 & 0 & 0 & 0 & \cdots & -1 & 0 \end{pmatrix}. \quad (4.2.3)$$

In this basis, the integral symplectic group is conveniently generated by $2F$ generators σ_n with $n = 1, 2, \dots, 2F$ whose matrix elements are given as

$$(\sigma_n)_{i,j} = \delta_{i,j} + \delta_{i,n}\delta_{n+1,j} - \delta_{i,n}\delta_{n-1,j}, \quad i, j = 1, 2, \dots, 2F \quad (4.2.4)$$

To define an action of the symplectic group $Sp(2F, \mathbb{Z})$ on this class of theories, it therefore suffices to specify the action of the generators σ_n .

The action of the generators with odd labels σ_{2n-1} preserves the number of gauge groups and shifts the levels of the n -th background field

$$\sigma_{2n-1} : \quad (k_G)_{i,j} \rightarrow (k_G)_{i,j}, \quad (k_M)_{i,j} \rightarrow (k_M)_{i,j}, \quad (k_F)_{i,j} \rightarrow (k_F)_{i,j} + \delta_{i,n} \delta_{n,j}. \quad (4.2.5)$$

The action of the even generators, σ_{2n} , is more complicated and performs a change of basis in the flavor symmetries while at the same time increasing the number of gauge groups by one. Explicitly, σ_{2n} can be factored as $\sigma_{2n} = g_n \circ c_U$ where c_U is a change of basis operation

$$c_U : \quad k_G \rightarrow k_G, \quad k_M \rightarrow k_M U, \quad k_F \rightarrow U^T k_F U \quad (4.2.6)$$

where in the above, the $F \times F$ matrix U is given by

$$(U)_{i,j} = \delta_{i,j} - \delta_{i-1,j}. \quad (4.2.7)$$

And the gauging operation g_n is given by

$$\begin{aligned}
 k_G &\rightarrow \begin{pmatrix} k_G & (k_M)_{i,n} \\ \hline (k_M^T)_{n,i} & (k_F)_{n,n} - 1 \end{pmatrix}, \\
 k_M &\rightarrow \begin{pmatrix} (k_M)_{i,1} & (k_M)_{i,2} & \cdots & (k_M)_{i,n-1} & 0 & (k_M)_{i,n+1} & \cdots & (k_M)_{i,F} \\ \hline (k_F)_{n,1} & (k_F)_{n,2} & \cdots & (k_F)_{n,n-1} & 1 & (k_F)_{n,n+1} & \cdots & (k_F)_{n,F} \end{pmatrix}, \\
 k_F &\rightarrow \begin{pmatrix} (k_F)_{1,1} & (k_F)_{1,2} & \cdots & (k_F)_{1,n-1} & 0 & (k_F)_{1,n+1} & \cdots & (k_F)_{1,F} \\ (k_F)_{2,1} & (k_F)_{2,2} & \cdots & (k_F)_{2,n-1} & 0 & (k_F)_{2,n+1} & \cdots & (k_F)_{2,F} \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \cdots & \vdots \\ (k_F)_{n-1,1} & (k_F)_{n-1,2} & \cdots & (k_F)_{n-1,n-1} & 0 & (k_F)_{n-1,n+1} & \cdots & (k_F)_{n-1,F} \\ 0 & 0 & \cdots & 0 & -1 & 0 & \cdots & 0 \\ (k_F)_{n+1,1} & (k_F)_{n+1,2} & \cdots & (k_F)_{n+1,n-1} & 0 & (k_F)_{n+1,n+1} & \cdots & (k_F)_{n+1,F} \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \cdots & \vdots \\ (k_F)_{F,1} & (k_F)_{F,2} & \cdots & (k_F)_{F,n-1} & 0 & (k_F)_{F,n+1} & \cdots & (k_F)_{F,F} \end{pmatrix}.
 \end{aligned} \tag{4.2.8}$$

Straightforward calculation using Gaussian path integrals may be used to verify that these operations satisfy the defining relations of $Sp(2F, \mathbb{Z})$. Notice that, while these relations are simple to prove, they nevertheless involve *quantum* field theory in an essential way. If w is any word in the generators σ_i which is equal to the identity element by a relation in the symplectic group, then the action of w on a given matrix of levels K produces a new matrix $w(K)$ which in general is not equal, as a matrix, to K . Nevertheless, the path integral performed with the matrices K and $w(K)$ produce identical correlation functions. Thus, the relations in $Sp(2F, \mathbb{Z})$ provide us with elementary, provable examples of duality in three-dimensional conformal field

theory.

Let us now turn our attention to the partition function \mathcal{Z} for this class of models. Since Abelian Chern-Simons theory is free, an application of the localization formula (4.1.2) reduces the computation to a simple Gaussian integral which is a function of an F -dimensional vector x of chemical potentials for the $U(1)^F$ flavor symmetry

$$\mathcal{Z}(x) = \int d^G y \exp \left[-\pi i (y \cdot x) K \begin{pmatrix} y \\ x \end{pmatrix} \right]. \quad (4.2.9)$$

The integral is trivially done to obtain⁴

$$\mathcal{Z}(x) = \frac{1}{\sqrt{|\det(k_G)|}} \exp \left[-\pi i x^T \tau x \right], \quad \tau \equiv k_F - k_M^T k_G^{-1} k_M. \quad (4.2.10)$$

From the resulting formula we see that the partition function is labeled by two invariants

$$|\det(k_G)| \in \mathbb{N}, \quad \tau \in \mathfrak{gl}(F, \mathbb{Q} \cup \{\infty\}), \quad \tau^T = \tau. \quad (4.2.11)$$

The possibility that the matrix τ may have infinite entries is included to allow for non-invertible k_G . In that case, the associated vector in the kernel of k_G describes a massless $U(1)$ vector multiplet and the flavor variable coupling to this multiplet is interpreted as a Fayet-Illiopoulos parameter. At the origin of this flavor variable the vector multiplet in question has a non-compact cylindrical Coulomb branch. This flat direction is not lifted when computing the path integral on S^3 because the R -charge

⁴As remarked in the introduction, we are only interested in \mathcal{Z} up to overall phases independent of all flavor variables. Thus in the following formulas we neglect such phases.

assignments do not induce conformal mass terms. This implies that the partition function \mathcal{Z} has a divergence. Meanwhile, away from the origin the non-zero FI parameter breaks supersymmetry and \mathcal{Z} vanishes. In total then, the partition function is proportional to a delta function in the flavor variable, and the narrow width limit of the Gaussian, when entries of τ are infinite, with infinite coefficient, $\det(k_G) \rightarrow 0$, should be interpreted as such a delta function.

The partition function formula (4.2.10) provides another context to illustrate the symplectic group $Sp(2F, \mathbb{Z})$ on conformal field theories, in this case, via its action on the invariants (4.2.11). A general symplectic matrix can be usefully written in terms of $F \times F$ blocks as

$$R \begin{pmatrix} A & B \\ C & D \end{pmatrix} R^T, \quad (4.2.12)$$

where R is certain invertible matrix which transforms the standard symplectic form to our choice (4.2.3) whose precise form is not important. Then, the action of symplectic transformations on τ is simply the standard action of the symplectic group on the Siegel half-space

$$\tau \rightarrow (A\tau + B)(C\tau + D)^{-1}. \quad (4.2.13)$$

Meanwhile, $\det(k_G)$ transforms as a modular form

$$\det(k_G) \rightarrow \det(C\tau + D) \det(k_G). \quad (4.2.14)$$

Thus the symplectic action on field theories reduces, at the level of partition functions, to the more familiar symplectic action on Gaussian integrals.

Before moving on to additional methods for studying these theories, let us revisit the issue of Chern-Simons couplings involving a background $U(1)_R$ gauge field. As remarked above such couplings are forbidden by superconformal invariance. Nevertheless, to elucidate the physical content of $\mathcal{Z}(x)$ as well as the partition functions on interacting field theories appearing later in this paper it is useful to examine exactly how such spurious terms would enter the result.

The squashed three-sphere partition functions under examination are Euclidean path integrals on the manifold

$$b|z_1|^2 + \frac{1}{b}|z_2|^2 = 0, \quad (z_1, z_2) \in \mathbb{C}^2, \quad b \in \mathbb{R}_+. \quad (4.2.15)$$

This geometry is labelled by a parameter b , a positive real number; however the symmetry under $b \rightarrow 1/b$ allows us to restrict our attention to the parameter

$$c_b \equiv \frac{i}{2} \left(b + \frac{1}{b} \right). \quad (4.2.16)$$

In this geometry preservation of supersymmetry requires one to turn on background values for scalars in the supergravity multiplet. While these fields are normally real, like the real mass variables x_i coupling to the ordinary flavors, in this background they are imaginary and proportional to c_b . As a result $R-R$ Chern-Simons levels, and R -flavor Chern-Simons levels appear as Gaussian prefactors in the partition function of the form

$$\exp(i\pi k_{RR}(c_b)^2 + 2\pi i k_{RF}(c_b)x). \quad (4.2.17)$$

From the above, we note that the $R-R$ Chern-Simons levels appear as multiplicative

constants independent of the flavor variables x . Since we are interested in computation of partition functions up to overall multiplication by phases such terms are not relevant for this work. On the other hand, the $R - F$ Chern-Simons terms appear as linear terms in x in the exponent. One can easily see why such terms violate superconformal invariance. The round three-sphere partition function for the conformal field theory in the absence of background fields is given by evaluating $\mathcal{Z}(x)$ at vanishing x and $c_b = i$. The first derivative with respect to x evaluated at the round three-sphere and vanishing x therefore computes the one-point function of the associated current

$$\partial_x \mathcal{Z}(x)|_{x=0, c_b=i} \sim k_{RF} \sim \langle j_F \rangle. \quad (4.2.18)$$

As the three-sphere is conformal to flat space, conformal invariance means that this one point function vanishes implying that k_{RF} must also vanish.

Quite generally throughout this paper we encounter examples of partition functions of interacting CFTs where the naive value of k_{RF} , as extracted from the first derivative of $\mathcal{Z}(x)$ evaluated at the conformal point, does not vanish. Superconformal invariance can always be restored in such examples by explicitly including ultraviolet counterterm values for k_{RF} to cancel the spurious contributions [71]. Thus, from now on we write expressions for partition functions with non-vanishing first derivatives, always keeping in mind that the true physical partition function of the conformal theory is only obtained by including suitable counterterms.

Quantum Mechanics and Partition Functions

The partition function calculations and $Sp(2F, \mathbb{Z})$ action described in the previous section can be phrased in a useful structure analogous to elementary quantum mechanics. We consider the Hilbert space of complex valued functions of F real variables and aim to interpret $\mathcal{Z}(x)$ as a wavefunction.⁵

First, introduce position and momentum operators acting on wavefunctions and consistent with the symplectic matrix J introduced in (4.2.3)

$$\hat{x}_i \rightarrow x_i, \quad \hat{p}_j \rightarrow -\frac{i}{2\pi} \frac{\partial}{\partial x_j} + \frac{i}{2\pi} \frac{\partial}{\partial x_{j-1}}, \quad [\hat{x}_i, \hat{p}_j] = \frac{i}{2\pi} (\delta_{i,j} - \delta_{i,j+1}). \quad (4.2.19)$$

We use Dirac bra-ket notation for states, and let $|y\rangle$ denote a normalized simultaneous eigenstate of the position operators

$$\hat{x}_i |y\rangle = y_i |y\rangle, \quad \langle x | y \rangle = \delta(x - y), \quad \mathbf{1} = \int dy |y\rangle \langle y|. \quad (4.2.20)$$

For convenience we also note that the wavefunction of a momentum eigenstate takes the form

$$\langle y | p \rangle = \exp [2\pi i (y_1 p_1 + y_2 (p_1 + p_2) + \cdots + y_F (p_1 + p_2 + \cdots + p_F))]. \quad (4.2.21)$$

On this Hilbert space there is a natural unitary representation of $Sp(2F, \mathbb{Z})$. This

⁵As is typical in quantum mechanical settings, we have need of wavefunctions which are not square integrable. Indeed all the partition functions associated to pure Chern-Simons theories are non-normalizable.

representation is defined using the generators (4.2.4) as follows⁶

$$\sigma_{2j-1} \mapsto \exp(-i\pi\hat{x}_j^2), \quad \sigma_{2j} \mapsto \exp(-i\pi\hat{p}_j^2). \quad (4.2.22)$$

One important feature of this representation is that its action by conjugation on position and momentum operators produces quantized canonical transformations. Explicitly, if M is any symplectic transformation we have

$$M \left(\sum_{j=1}^F a_{2j-1} \hat{x}_j + a_{2j} \hat{p}_j \right) M^{-1} = \sum_{j=1}^F \sum_{k=1}^{2F} (M_{2j-1,k} a_k) \hat{x}_j + (M_{2j,k} a_k) \hat{p}_j. \quad (4.2.23)$$

This fact underlies the significance of this representation in all that follows.

We now wish to show that we may interpret the partition function of a theory Ψ as a wavefunction of an associated state $|\Psi\rangle$

$$\mathcal{Z}_\Psi(x) = \langle x | \Psi \rangle. \quad (4.2.24)$$

Of course both wavefunctions and partition functions are complex-valued functions of a F real variables x_i so we are free to make the identification appearing in (4.2.24). The non-trivial aspect of this identification is that the $Sp(2F, \mathbb{Z})$ action on quantum field theories, defined by the operations appearing in (4.2.5)-(4.2.6) can be achieved at the level of the partition function by the action of the operators of the same name defined by the representation given in (4.2.22). To see that these quantum mechanics

⁶Technically speaking, the operators above must be multiplied by a certain overall (operator independent) phase. However, since we are ignoring phases in our partition functions, we will also ignore overall phases in quantum mechanics matrix elements.

operators behave correctly, note that given any arbitrary state $|\Psi\rangle$ we have

$$\langle x|\sigma_{2j-1}|\Psi\rangle = \exp(-i\pi x_j^2)\langle x|\Psi\rangle. \quad (4.2.25)$$

Thus, if the state $|\Psi\rangle$ corresponds to a quantum field theory with partition function $\langle x|\Psi\rangle$, then the integral definition of the partition function given in equation (4.2.9) implies that σ_{2j-1} shifts the background Chern-Simons level for the j -th flavor by one unit as expected. We can similarly see that the quantum mechanical σ_{2j} operator acts as required. We have

$$\langle x|\sigma_{2j}|\Psi\rangle = \int dy \prod_{k \neq j} \delta(x_k - y_k) e^{i\pi(y_j - x_j)^2} \Psi(Uy). \quad (4.2.26)$$

This is exactly the action expected for the S operation at the level of partition functions. It performs a change of basis on the flavors, given by the U matrix, and introduces a single new gauge group with specified Chern-Simons levels.

$SL(2, \mathbb{Z})$ Examples As a sample application of the above ideas, we present here a simple set of calculations based on $SL(2, \mathbb{Z})$, relevant for the case of a single flavor symmetry. Our symplectic transformations acting on quantum field theories are generated by the familiar operators S and T subject to the relations⁷

$$S^2 = (ST)^3 = 1. \quad (4.2.27)$$

⁷As usual we ignore phases in \mathcal{Z} and hence the central element S^2 can be set to the identity.

T acts on theories by increasing the Chern-Simons level of the flavor

$$T : k_G \rightarrow k_G, \quad k_M \rightarrow k_M, \quad k_F \rightarrow k_F + 1, \quad (4.2.28)$$

while the S generator acts to gauge the flavor symmetry and introduces a new flavor which is dual to the original symmetry

$$S : k_G \rightarrow \begin{pmatrix} k_G & k_M \\ k_M^T & k_F \end{pmatrix}, \quad k_M^T \rightarrow \begin{pmatrix} 0 & 0 & \dots & 0 & 1 \end{pmatrix}, \quad k_F \rightarrow 0. \quad (4.2.29)$$

The relevant quantum mechanics is now single variable for the single $U(1)$ flavor symmetry with standard commutation relations

$$[\hat{x}, \hat{p}] = \frac{i}{2\pi}. \quad (4.2.30)$$

And the representation of symplectic transformations is given by

$$T \rightarrow \exp(-i\pi\hat{x}^2), \quad STS^{-1} \rightarrow \exp(-i\pi\hat{p}^2). \quad (4.2.31)$$

A simple class of theories is defined starting from the trivial theory Ω . This theory has no gauge groups and vanishing flavor Chern-Simons levels. Its partition function is unity

$$\mathcal{Z}_\Omega(x) = \langle x | \Omega \rangle = 1. \quad (4.2.32)$$

More interesting theories can be generated by starting with the trivial theory Ω and acting with S and T . For a general $SL(2, \mathbb{Z})$ element \mathcal{O} we have the following result

for the partition function⁸

$$\mathcal{O} = \begin{pmatrix} r & t \\ s & u \end{pmatrix} \in SL(2, \mathbb{Z}) \implies \mathcal{Z}_{\mathcal{O}}(x) = \langle x | \mathcal{O} | \Omega \rangle = \frac{1}{\sqrt{|u|}} \exp(-i\pi x^2 t/u). \quad (4.2.33)$$

The answer thus takes the general form (4.2.10) with associated invariants

$$|\det(k_G)| = u, \quad \tau = t/u. \quad (4.2.34)$$

Notice that, consistent with our general discussion, a particular element \mathcal{O} defines a particular *quantum* Abelian Chern-Simons theory, not a classical Lagrangian presentation of such a theory. To obtain such a Lagrangian presentation, one must pick a word in the generators S and T which is equal to the given element M . Different words in the generators which are equal to the same fixed \mathcal{O} provide examples of dual theories.

Doubled Flavor Variables, Operator Multiplication, and Gauging The $SL(2, \mathbb{Z})$ examples described above can be readily extended to the case of more flavor symmetry. For any F we consider the F -variable quantum mechanics described in section 4.2.1 and introduce a trivial theory Ω with unit partition function. Then, if \mathcal{O} is any element of $Sp(2F, \mathbb{Z})$ we can consider a quantum theory generated by acting with \mathcal{O} on the trivial theory. The resulting partition function can be expressed as the wavefunction obtained by acting on the vacuum state $|\Omega\rangle$, a normalized momentum

⁸Formula (4.2.33) is correct when u is non-vanishing. In the special case where $u = 0$ the result is just $\mathcal{Z}_{\mathcal{O}}(x) = \delta(x)$.

eigenstate with eigenvalue zero

$$\mathcal{Z}_{\mathcal{O}}(x_1, \dots, x_F) = \langle x_1, \dots, x_F | \mathcal{O} | \Omega \rangle. \quad (4.2.35)$$

As in the case of a single flavor symmetry discussed above, the resulting quantum field theory and partition function depends only on the element \mathcal{O} in $Sp(2F, \mathbb{Z})$, while a particular Lagrangian realization of the theory requires a choice of word in the generators σ_n which represents \mathcal{O} .

This quantum mechanical setup naturally suggests additional quantities to compute. Rather than considering the wavefunction of \mathcal{O} acting on the trivial state $|\Omega\rangle$, we may instead double the flavor variables and compute the complete matrix element of \mathcal{O}

$$\mathcal{Z}_{\mathcal{O}}^{\text{Op}}(x_1, \dots, x_F, y_1, \dots, y_F) \equiv \langle x_1, \dots, x_F | \mathcal{O} | y_1, \dots, y_F \rangle. \quad (4.2.36)$$

Where in the above the superscript ‘Op’ for operator, is used to distinguish from the partition functions introduced in (4.2.35). For $\mathcal{O} \in Sp(2F, \mathbb{Z})$ a symplectic operator, the matrix element $\mathcal{Z}_{\mathcal{O}}^{\text{Op}}$, is the partition function of an Abelian Chern-Simons theory now coupled to $2F$ background flavor fields.

The construction of (4.2.36) is not limited to the case of symplectic operators. Indeed, in section 4.5 we will see that an interesting class of non-symplectic operators \mathcal{O} have matrix elements which are identified with partition functions of interacting three-dimensional conformal field theories coupled to $2F$ flavor fields. In general, such matrix element partition functions have the following features.

- If $\mathcal{Z}_{\mathcal{O}}^{Op}(x, y)$ is known then the partition function $\mathcal{Z}_{\mathcal{O}}(x)$ is determined,

$$\mathcal{Z}_{\mathcal{O}}(x) = \int dy \mathcal{Z}_{\mathcal{O}}^{Op}(x, y). \quad (4.2.37)$$

In the physical interpretation we have developed, the integration over the y variables is the gauging of the associated flavor variables at vanishing values of the associated FI parameters.

- More generally, the quantum-mechanical operation of operator multiplication can be interpreted in field theory. A product of operators can always be decomposed into a convolution by an insertion of a complete set of states

$$\mathcal{Z}_{\mathcal{O}_1 \mathcal{O}_2}^{Op}(x, y) = \int dz \mathcal{Z}_{\mathcal{O}_1}^{Op}(x, z) \mathcal{Z}_{\mathcal{O}_2}^{Op}(z, y). \quad (4.2.38)$$

Again, the integration is physically interpreted as gauging. We consider the two theories, whose partition functions are given by the matrix elements of \mathcal{O}_i , we identify flavors as indicated in (4.2.38) and gauge with no FI-term.

- $\mathcal{Z}_{\mathcal{O}}^{Op}(x, y)$ is a partition function of a theory coupled to $2F$ background flavor fields. A general theory of this type is acted on by the symplectic group $Sp(4F, \mathbb{Z})$, however a matrix element is acted on only by the subgroup

$$Sp(2F, \mathbb{Z}) \times Sp(2F, \mathbb{Z}) \subset Sp(4F, \mathbb{Z}) \quad (4.2.39)$$

which does not mix the x and y variables. The geometrical and physical interpretation of this splitting will be explained in section 4.5.

4.2.2 Tangles

Our goal in this section is to give a geometric counterpart to the field theory and partition function formalism developed in the previous analysis. A natural way to develop such an interpretation is to engineer the Abelian Chern-Simons theory by compactification of the M5-brane on a three-manifold M . In six dimensions, the worldvolume of the M5-brane supports a two-form field B with self-dual three-form field strength [72]. When reduced on a three-manifold, the modes of B may engineer an Abelian Chern-Simons theory. We review aspects of this reduction and explain the three-dimensional geometry required to understand the $Sp(2F, \mathbb{Z})$ action.

Reduction of the Chiral Two-Form

Consider the free Abelian M5-brane theory reduced on a three-manifold M . To formulate the theory of a chiral two-form, M must be endowed with an orientation which we freely use throughout our analysis. The effective theory in the three macroscopic dimensions is controlled by the integral homology group $H_1(M, \mathbb{Z})$. The simplest way to understand this fact is to note that a massive probe particle in the theory arises from an M2-brane which ends on a one-cycle γ in M . In particular the homology class of $\gamma \in H_1(M, \mathbb{Z})$ labels the charge of the particle.

In the effective theory in three dimensions, massive charged probes are described by Wilson lines. Let C denote a one-cycle in the non-compact Minkowski space. A general Wilson line can be written as

$$\exp \left(iq_\alpha \oint_C A_\alpha \right). \quad (4.2.40)$$

If the theory in question has G gauge fields and F flavor fields, then the charge vector q_α has $G + F$ components and integral entries. However, in the presence of non-vanishing Chern-Simons levels, the charge vector q is in general torsion valued. Thus, distinct values of the integral charge vector q may be physically equivalent. The allowed distinct values of the charge vector are readily determined by examining the two-point function of Wilson loops in Abelian Chern-Simons theory coupled to background vectors. The results are summarized as follows. Let $\mathbb{Z}^G \subset \mathbb{Z}^{G+F}$ be the subset of charges uncharged under the flavor group $U(1)^F$. We view the level matrix as specifying a map

$$K : \mathbb{Z}^G \longrightarrow \mathbb{Z}^{G+F}, \quad (4.2.41)$$

and those charge vectors in the image of this map are physically equivalent to no charge at all.

Since we have determined that possible Wilson lines encode the homology of M it follows that

$$H_1(M, \mathbb{Z}) \cong \mathbb{Z}^{G+F}/\mathfrak{S}(K). \quad (4.2.42)$$

Equation (4.2.42) encodes the appropriate generalization of Kaluza-Klein reduction to the case of torsion valued charges. The fact that we study Chern-Simons theories up to possible framing anomalies (equivalently, overall phases in the partition function) means that the entire theory is characterized by the group (4.2.42). However, the homology of M , and hence the underlying physics has no preferred description via a classical Lagrangian. Indeed as we will illustrate in the remainder of this section, distinct classical theories, with the same group of Wilson line charges, can in fact

arise from compactification on the same underlying manifold M . Thus, already in this elementary discussion of reduction of the two-form we see the important fact that compactification of the M5-brane theory produces a specific *quantum field theory* not, as one might naively expect, a specific Lagrangian presentation of a classical theory which we subsequently quantize. It is for this reason that our geometric constructions of field theories are powerful, for in these constructions, quantum dualities are manifest.

Finally, before moving on to discuss explicit examples we remark on the geometry associated to flavor symmetries. These arise when the manifold M is allowed to become non-compact. Suppose that M develops cylindrical regions near infinity which take the form of $\mathbb{R} \times \mathbb{R}_+ \times S^1$. Then on the asymptotic S^1 cycle we may reduce the two-form field to obtain another gauge field

$$A = \int_{S^1} B. \quad (4.2.43)$$

However, unlike the compact cycles in the interior of M , the cycle S^1 has no compact Poincaré dual and hence A is a non-dynamical background field; it provides the effective theory in three dimensions with a $U(1)$ flavor symmetry. Moreover, since the boundary behavior of A must be specified to obtain a well-defined theory in three dimensions, the resulting theory is of the type we have considered in the introduction: a theory with flavor symmetries and a specified coupling to background gauge fields. As a result the partition function $\mathcal{Z}(x)$ is a well-defined observable of the theory. The number of flavor variables on which the result depends is the number of homologically independent cylindrical ends of M . For F flavors we require $F + 1$ cylindrical ends.

Double Covers From Tangles

The specific class of geometries that we will study are conveniently presented as double covers over the non-compact space \mathbb{R}^3 , branched over a one-dimensional locus L

$$\begin{array}{ccc} \mathbb{Z}_2 & \longrightarrow & M \\ & & \downarrow \\ & & L \subset \mathbb{R}^3 \end{array} \quad (4.2.44)$$

Topologically L is simply the union of $F + 1$ lines, however its embedding in \mathbb{R}^3 is constrained. On the asymptotic two-sphere at the boundary of three-space, we mark $2F + 2$ distinct points p_1, \dots, p_{2F+2} . The $2F + 2$ ends of L at infinity are the points p_i . Meanwhile, in the interior of \mathbb{R}^3 the components of L may be knotted. Such an object is known as an $(F + 1)$ -tangle. An example in the case of $F = 1$ is illustrated in Figure 4.2. Given two distinct tangles L_1 and L_2 , they are considered to be equal




Figure 4.2: A tangle. The four endpoints of L extend forever towards the points at infinity.

topologically when one can be deformed to the other by isotopy in the interior of \mathbb{R}^3 which keeps the ends at infinity fixed.

In the following we will also need to be more precise about the behavior near the asymptotes p_i . Let $\overline{B}_r \subset \mathbb{R}^3$ denote the exterior of a closed ball of radius r centered at the origin. We view \overline{B}_r topologically as $S^2 \times I$ where I is an open interval. For large r the portion of the tangle $L \cap \overline{B}_r$ contained in \overline{B}_r consists of $2F + 2$ arcs. We constrain the behavior of these arcs by requiring that the pair $(\overline{B}_r, L \cap \overline{B}_r)$ is homeomorphic to the trivial pair $(S^2 \times I, \{p_1, p_2, \dots, p_{2F+2}\} \times I)$ where the p_i are points in S^2 . This constraint implies that the knotting behavior of the tangle eventually stops as we approach infinity. In practice it means that any planar projection of the tangle L appears at sufficiently large distances as $2F + 2$ disjoint semi-infinite line segments which undergo no crossings.

For most of the remainder of this section, we will argue that the class of three-manifolds obtained as double covers branched over tangles have exactly the correct properties to engineer the Abelian Chern-Simons theories coupled to a background flavor gauge field which we have discussed in the previous section. As a first step, observe that such geometries do indeed support F flavor symmetries. Group the asymptote points into $F + 1$ pairs $\{p_{2i-1}, p_{2i}\}$. The double cover of \mathbb{R}^3 branched over the two straight arcs emanating from $\{p_{2i-1}, p_{2i}\}$ yields the anticipated cylindrical ends of M required to support flavor symmetry. The fact that there are $2F + 2$ asymptotes ensures that there are exactly F independent flavor symmetries; note that one linear combination of the asymptotic cycles can be contracted in the interior of the manifold.

In section 4.2.3 we explain how to extract a Lagrangian for an Abelian Chern-Simons theory from the geometric data of a tangle. As we have previously described,

the M5-brane on M does not provide a preferred Lagrangian. Consistent with this fact, we find that a Lagrangian description of the field theory associated to a particular tangle requires additional geometric choices. In this case the choice is a *Seifert surface*, a surface whose boundary is the given tangle. For any fixed L there are infinitely many such surfaces each giving rise to a distinct Lagrangian presentation of the same underlying physics.

Finally we argue that tangles, and hence the class of three-manifolds described as double covers branched over tangles, enjoy a natural action by $Sp(2F, \mathbb{Z})$. To illustrate this action, we draw a generic tangle with $F + 1$ strands as in Figure 4.3. Then, the




Figure 4.3: A generic tangle L in \mathbb{R}^3 . The ellipsis indicate that the strands continue to infinity with no additional crossings. In the interior of the box, the strands are in general knotted in an arbitrary way.

action of the symplectic group is defined by the generators σ_j where $j = 1, \dots, 2F$, which act on the tangles by braid moves in a neighborhood of the asymptotes p_i . Several examples are illustrated in Figure 4.4.

One way to understand this three-dimensional geometry is to note that the boundary at infinity of M is a double cover of S^2 branched over $2F + 2$ points, and hence is a Riemann surface of genus F . The action defined in Figure 4.4 is a surgery on M which in general changes its topology. This surgery is induced by mapping class group transformations in a neighborhood of the boundary of M . In particular, as is clear




Figure 4.4: The symplectic action on the tangles. In (a) a typical odd generator σ_{2n+1} . In (b) an even generator σ_{2n} .

from the illustrations, what we have defined is not, a priori, an action of the symplectic group, but rather an action of the braid group, B_{2F+1} on $2F+1$ strands [73].⁹ The braid group and the symplectic group are related by a well-known exact sequence

$$1 \rightarrow \mathcal{T}_{2F+1} \rightarrow B_{2F+1} \rightarrow Sp(2F, \mathbb{Z}) \rightarrow 1, \quad (4.2.45)$$

where \mathcal{T}_{2F+1} is the Torelli group. To make contact with our discussion of field theories, we wish to illustrate that the action of the braid group defined by Figure 4.4 reduces to an action of the symplectic group on the associated field theories. This implies that any two elements of B_{2F+1} that differ by multiplication by a Torelli element must give rise equivalent actions on the field theories extracted from an arbitrary tangle. More bluntly, the Torelli group generates dualities. One of the outcomes of this section is a proof of this fact.

⁹The ‘last strand’ appearing at the bottom of the diagram in Figure 4.3 is stationary under all braid moves. Alternatively one may work with the spherical braid group and impose additional relations. For simplicity we stick with the more familiar planar braids.

4.2.3 Seifert Surfaces

To understand the physics encoded by a tangle we need control over the homology of the cover manifold M . The appropriate tool for this task is a Seifert surface. In general given any knot¹⁰, a Seifert surface Σ for the knot is a connected Riemann surface with boundary the given knot. An example is illustrated in Figure 4.5. In the mathematics literature it is common to impose the additional requirement that Σ be oriented. In our context there is no natural orientation for Σ and hence we proceed generally allowing possibly non-orientable Seifert surfaces.




Figure 4.5: A sample Seifert surface. In (a) a pretzel knot in the three-sphere. In (b) a Seifert surface for the knot.

For any knot, there exist infinitely many distinct Seifert surfaces; alternatively, given a knot diagram there are a number of simple algorithms to construct a Σ [74]. We describe one useful algorithm in section 4.2.3. Seifert surfaces are relevant for our discussion because, if one wishes to construct a double cover branched over a knot, then a choice of Σ is equivalent to a choice of branch sheet. As such, features of the homology of the branched cover M can be extracted from a knowledge of a

¹⁰In this paper the term *knot* will be used broadly to include both knots and multicomponent links.

Seifert surface. However, the resulting three-manifold M depends only on the branch locus L and hence the homology and ultimately the associated physical theory are independent of the choice of Σ . In the following we explain how any fixed choice of Seifert surface allows us to extract a set of gauge and flavor groups and a matrix of Chern-Simons levels from the geometry.

To begin, we assume for simplicity that we are dealing with a knot in S^3 , as opposed the non-compact tangles in \mathbb{R}^3 needed to support flavor symmetry. The generalizations to the present non-compact situation will then be straightforward. The detailed statements that we require are as follows. Any cycle in $H_1(M, \mathbb{Z})$ can be thought of as a cycle on the base S^3 which encircles Σ . This can be viewed as a direct parallel with the theory of branched covers of the two-sphere.¹¹ Thus, we deduce that there is a surjective map

$$H_1(S^3 - \Sigma, \mathbb{Z}) \rightarrow H_1(M, \mathbb{Z}) \rightarrow 0. \quad (4.2.46)$$

Meanwhile, there is a linking number pairing between cycles in $H_1(S^3 - \Sigma, \mathbb{Z})$ and cycles in $H_1(\Sigma, \mathbb{Z})$. This linking number pairing is perfect and hence we may extend (4.2.46) to

$$H_1(\Sigma, \mathbb{Z}) \cong H_1(S^3 - \Sigma, \mathbb{Z}) \rightarrow H_1(M, \mathbb{Z}) \rightarrow 0. \quad (4.2.47)$$

Our task is thus reduced to determining which cycles on the Seifert surface correspond to trivial cycles in the homology of M .

¹¹In making this comparison it is crucial that the branch locus is connected.

To this end, we define a symmetric bilinear form, the so-called *Trotter* form

$$K : H_1(\Sigma, \mathbb{Z}) \times H_1(\Sigma, \mathbb{Z}) \rightarrow \mathbb{Z}. \quad (4.2.48)$$

Our choice of notation is intentional: we will see that the Trotter form defines the Chern-Simons levels. To extract K we let $\alpha \in H_1(\Sigma, \mathbb{Z})$, and set $\tilde{\alpha}$ to be the cycle in S^3 obtained from locally pushing α off of Σ in both directions. The cycle $\tilde{\alpha}$ is a two-to-one cover of α . If Σ is orientable then $\tilde{\alpha}$ consists of two disconnected cycles each on a given side of Σ (as determined by the orientation), however in general $\tilde{\alpha}$ may be connected. The definition of the Trotter form is

$$K(\alpha, \beta) = lk_{\#}(\tilde{\alpha}, \beta), \quad (4.2.49)$$

where $lk_{\#}$ denotes the linking number pairing of cycles in S^3 . A simple calculation illustrates that K is symmetric. A slightly less trivial argument shows that the image of K (as a map from cycles to cocycles) is exactly the set of cycles on Σ which are trivial in M . Thus, the completion of the sequence (4.2.47) is

$$0 \rightarrow \text{Im}(K) \rightarrow H_1(\Sigma, \mathbb{Z}) \cong H_1(S^3 - \Sigma, \mathbb{Z}) \rightarrow H_1(M, \mathbb{Z}) \rightarrow 0. \quad (4.2.50)$$

In particular we conclude that $H_1(M, \mathbb{Z}) \cong H_1(\Sigma, \mathbb{Z})/\text{Im}(K)$.

Double covers of S^3 branched over knots are exactly the geometries we expect to engineer Abelian Chern-Simons theories without flavor symmetries and we may relate theorem (4.2.50) to physics as follows:

- A choice of Seifert surface Σ and a set of generators of homology $\alpha_1, \dots, \alpha_G$ determines a set of G Abelian gauge fields.
- The Trotter form pairing on cycles in $H_1(\Sigma, \mathbb{Z})$ is equal to the Chern-Simons levels matrix on the associated gauge fields.

Distinct choices of Seifert surfaces are physically related by duality transformations. This fact is easy to verify directly. For example, distinct choices of Σ which differ by gluing in handles or Möbius bands add new gauge cycles and compensating levels to keep the underlying physics unmodified.

Finally, we generalize our discussion of Seifert surfaces and homology to the case of non-compact geometries required to discuss flavor symmetries. Let L denote a tangle in \mathbb{R}^3 . We extend to the case of non-compact Seifert surfaces Σ , again defined by the condition that they are connected surfaces with boundary L . However, now to compute flavor data we must fix a compactification of both L and Σ . We achieve this by identifying the points p_i in pairs and glueing in arcs near infinity as illustrated in Figure 4.6.




Figure 4.6: The asymptotic geometry of a Seifert surface for a generic tangle. The shaded blue region indicates the interior of Σ . The arcs at infinity indicate the compactification of L and Σ . The non-compact cycles on Σ give rise to flavor symmetries.

Let δ indicate the union of the arcs at infinity, and Σ_c the compactified Seifert surface including δ . The surface Σ_c should be viewed as embedded inside S^3 , the one-

point compactification of \mathbb{R}^3 , and all calculations of linking numbers take place inside S^3 . For simplicity in future diagrams we often leave the compactification data of the Seifert surface implicit by setting the convention that whenever a non-compact Seifert surface consists of strips extending to infinity in \mathbb{R}^3 the intended compactification is the one where the strips are capped off with arcs as in Figure 4.6.

With these preliminaries about compactifications fixed, we may now state the required generalization of the sequence (4.2.50)

$$0 \rightarrow \text{Im}(K) \rightarrow H_1(\Sigma_c, \delta, \mathbb{Z}) \cong H_1(\mathbb{R}^3 - \Sigma, \mathbb{Z}) \rightarrow H_1(M, \mathbb{Z}) \rightarrow 0. \quad (4.2.51)$$

Note that in addition to the boundaryless cycles in Σ_c which give rise to gauge groups, $H_1(\Sigma_c, \delta, \mathbb{Z})$ also contains F cycles with boundary in δ . In the uncompactified Seifert surface these cycles are non-compact and illustrated in Figure 4.6. They correspond physically to the $U(1)^F$ flavor symmetry. To complete the construction, it thus remains to extend the definition of the Trotter form. For boundaryless cycles in Σ_c , the definition is as before. Meanwhile to evaluate the Trotter form on cycles with boundary, we again push them out locally in both directions from Σ_c and compute the local linking number from the interior of Σ . Alternatively, one may simply think of the pair of points in the boundary of a flavor cycle in Σ_c as formally identified. In this way we obtain a closed cycle in S^3 and we compute its Trotter pairings as before. In this way we obtain a bilinear form K defined on $H_1(\Sigma_c, \delta, \mathbb{Z})$, and the image of this form restricted to the boundaryless cycles in $H_1(\Sigma_c, \delta, \mathbb{Z})$ defines the term $\text{Im}(K)$ appearing in (4.2.51).

To summarize, given any tangle L in \mathbb{R}^3 , we extract a Lagrangian description of

the effective Abelian Chern-Simons theory as follows:

- A choice of Seifert surface Σ and a set of generators of the relative homology $H_1(\Sigma_c, \delta, \mathbb{Z})$, $\alpha_1, \dots, \alpha_{G+F}$, determines a set of Abelian vector fields. Generators corresponding to boundaryless one-cycles correspond to gauged $U(1)$'s while those corresponding to one-cycles with boundary in δ are background flavor fields.
- The Trotter form pairing on cycles in $H_1(\Sigma_c, \delta, \mathbb{Z})$ is equal to the Chern-Simons levels pairing on the associated vector fields. We denote by $\text{Im}(K)$ the image of this pairing restricted to the subset of boundaryless cycles in Σ_c , and we have

$$\{\text{Wilson Line Charges}\} = H_1(\Sigma_c, \delta, \mathbb{Z})/\text{Im}(K) = H_1(M, \mathbb{Z}). \quad (4.2.52)$$

Checkerboards

The previous discussion of Seifert surfaces is complete but abstract. For computations with explicit examples, it is useful to have a fast algorithm for computing the relevant linking numbers and hence extracting a set of Chern-Simons levels from geometry. One such method, described in this section, is provided by so-called checkerboard Seifert surfaces.

To begin, fix a planar projection of the tangle $L \subset \mathbb{R}^3$. In such a planar diagram, information about the knotting behavior of L is contained in the crossings in the diagram. Each crossing locally divides the plane into four quadrants. We construct a Seifert surface for L by coloring the two of the four quadrants at each crossing in checkerboard fashion and extending consistently to all crossings. The colored region

then defines Σ . Note that each crossing c in the diagram is endowed with a sign $\zeta(c) = \pm 1$. The sign is given by the orientation of the crossing c (over- or under-cross) relative to the Seifert surface itself. Imagine rotating the under-strand through the surface to align with the over-strand. If this rotation is counter-clockwise, the sign is $+1$, if clockwise, the sign is -1 ; see Figure 4.7.



Figure 4.7: Checkerboard colorings and their associated signs. In (a), a positive crossing. In (b), a negative crossing.

To compute the Trotter form, we first assume that Σ_c appears compactly in the plane.¹² Then, there is a natural basis of boundaryless cycles in Σ_c associated to the compact uncolored regions of the plane. We orient these cycles counterclockwise. Similarly, in the diagram of Σ , non-compact white regions may be associated to flavor cycles. These cycles are again canonically oriented “counterclockwise,” i.e. the cross-product of the tangent vector to the cycle with the outward normal pointing into the associated non-compact uncolored region must be out of the plane.¹³ The Trotter pairing on these cycles is determined by the summing over crossings involving a given pair of cycles weighted by the sign of the crossing. Explicitly, for α and β a pair of

¹²This assumption cannot in general be relaxed. Indeed when Σ_c is non-compact in the plane one must take into account the fact that in the compactification procedure, the plane becomes an embedded S^2 inside S^3 and hence may endow Σ_c with additional topology.

¹³There is one linear relation among the flavor cycles obtained in this way. So a given Σ will have $F + 1$ non-compact uncolored regions and F independent flavor cycles.

generators as defined above we have

$$K(\alpha, \beta) = \begin{cases} + \sum_{\alpha, \beta \in c} \zeta(c) & \text{if } \alpha \neq \beta, \\ - \sum_{\alpha \in c} \zeta(c) & \text{if } \alpha = \beta. \end{cases} \quad (4.2.53)$$

Equation (4.2.53) provides a convenient way to read off Chern-Simons levels for a given tangle and will be utilized heavily (although often implicitly) throughout the remainder of this work.

4.2.4 The Torelli Group of Dualities

We are now equipped to investigate the symplectic action on tangles. In particular, we wish to prove that the action of the braid group B_{2F+1} on tangles, reduces to an action of the symplectic group $Sp(2F, \mathbb{Z})$ when considered as an action on the corresponding physical theories.

To prove this statement, we proceed in the most direct way possible. We compute the action of the braid group generators σ_n , illustrated in Figure 4.4, on the Chern-Simons levels extracted from any Seifert surface associated to the tangle. We show that this action matches exactly the previously defined action (4.2.5)-(4.2.6). Since the later action is symplectic this implies that the former is as well. In particular, this suffices to prove that the Torelli group acts trivially on the underlying quantum field theory.

To begin, we fix a Seifert surface with definite compactification data δ . As we have previously described, δ is a union of $F + 1$ arcs δ_i with $i = 1, \dots, F + 1$. We

draw diagrams such that the arcs are ordered down the page, with δ_1 appearing at the top, δ_2 next and so on. A basis of flavor cycles in $H_1(\Sigma_c, \delta, \mathbb{Z})$ is given by F cycles α_i each of which begins at δ_{F+1} and terminates at δ_i . This geometry is shown in Figure 4.6. With these conventions, the braid moves act as in Figure 4.8.




Figure 4.8: The action of braid moves on linking numbers. In (a), all linking numbers are unmodified except for those of the flavor cycle α_1 which runs from δ_{F+1} to δ_1 (illustrated in red), whose self-linking number is increased by one. In (b), we first change basis of flavor cycles from α_j to β_j . Then we gauge β_1 , shown in green, and introduce a new flavor cycle, shown in red, linked with the gauged cycle.

Consider first the odd braid moves σ_{2j-1} illustrated in Figure 4.8a. According to formula 4.2.53, the effect of such a move is to modify the Trotter form by increasing $K(\alpha_j, \alpha_j)$ by one while leaving all other entries invariant. This is exactly the expected action given by (4.2.5) on Chern-Simons levels for this transformation.

Similarly we may consider the braid moves with even index σ_{2j} illustrated in Figure 4.8b. To understand this transformation we first change basis on flavor cycles to β_i which run from δ_i to δ_{i+1} . The transformation from the basis α_i to the basis β_i is given by the matrix U in (4.2.6)-(4.2.7). Then, the braid move σ_{2j} gauges β_j and introduces a new flavor cycle $\tilde{\beta}_j$. Finally, we update the Trotter form to account for

the new linking numbers apparent in Figure 4.8b

$$\delta K(\beta_j, \beta_j) = -1, \quad \delta K(\tilde{\beta}_j, \tilde{\beta}_j) = -1, \quad \delta K(\beta_j, \tilde{\beta}_j) = 1. \quad (4.2.54)$$

This is exactly the gauging operation of equation (4.2.8). Thus we have completed the verification of the symplectic action.

As a result of this analysis we conclude that the Torreli group \mathcal{T}_{2F+1} acts via dualities on Abelian Chern-Simons theories. Given any tangle one may act on it with a Torreli element to obtain a new geometry. Fixing Seifert surfaces, the two geometries in general will have distinct classical Lagrangian descriptions yet their underlying quantum physics is identical.

Moreover, as we see in section 4.3 and beyond, the technology of this section generalizes immediately to the more complicated geometries required for constructing interacting field theories. In particular, the symplectic action we have described arises from braid moves near infinity and hence is enjoyed by any geometry with the same asymptotics.

4.2.5 Geometric Origin of Quantum Mechanics

To conclude our discussion of Abelian Chern-Simons theories we briefly comment on the origin of the quantum mechanical framework for partition function calculations discussed in section 4.2.1. We fix an Abelian Chern-Simons theory $\mathcal{T}(M)$ engineered by reduction of the M5-brane on a three-manifold M . The three-sphere partition function of this theory then has an underlying six-dimensional origin as the M5-brane

partition function on the product manifold $M \times S^3$,

$$\mathcal{Z}_{S^3}^{\mathcal{T}(M)} = \mathcal{Z}_{M \times S^3}^{\text{M5}}. \quad (4.2.55)$$

Thus far, we have viewed M as small and interpreted the long-distance physics as an Abelian Chern-Simons theory coupled to flavors which we subsequently compactify on S^3 . However, an alternative point of view is to consider S^3 to be small, and obtain another effective three-dimensional description which is subsequently compactified on M . As S^3 has vanishing first homology, the resulting three-dimensional description is one with no Wilson line observables. From the point of view of this paper, which studies partition functions on compact manifolds up to multiplication by overall factors, we cannot distinguish the result from the trivial theory.

However, a standing conjecture is that in fact the reduction on S^3 gives rise to a $U(1)$ Chern-Simons theory at level one. Assuming the veracity of this statement, we then arrive at a beautiful physical interpretation of the quantum mechanical calculations in section 4.2.1.

Recall that M is not a compact manifold, but rather has non-compact cylindrical ends required to support flavor symmetry. One may equivalently view M as a manifold with boundary at infinity and with specified boundary conditions supplied by the background flavor gauge fields. On general grounds, the path-integral of $U(1)$ level one Chern-Simons theory on M produces a state in the boundary Hilbert space determined by the quantization of Chern-Simons theory on ∂M . In this case, as a consequence of the conjecture, one is quantizing a space of $U(1)$ flat connections on a Riemann surface with $2F$ independent cycles. The Hilbert space thus consists of

wavefunctions of F real variables x_1, \dots, x_F , which are interpreted as the holonomies of a flat connection around a maximal collection of F non-intersecting homology classes in ∂M . The symplectic action is then the standard action in this Hilbert space induced by the action on the homology of the genus F Riemann surface ∂M .

Thus, the quantum mechanical framework which emerged abstractly from supersymmetric localization formulas in section 4.2.1, takes on a natural physical interpretation when the associated field theories are geometrically engineered. In particular, the viewpoint of the partition function $\mathcal{Z}_{S^3}^{\mathcal{T}(M)}(x)$ as a wavefunction in a Hilbert space is a simple consequence of the six-dimensional origin of the computation and leads to a correspondence of partition functions

$$\mathcal{Z}_{S^3}^{\mathcal{T}(M)}(x) = \mathcal{Z}_M^{U(1)_1}(x). \quad (4.2.56)$$

This identification is reminiscent to the one studied in [75] and was obtained in the case of three-manifolds from different perspectives by [76, 77].

4.3 Particles, Singularities, and Superpotentials

In this section we exit the realm of free Abelian Chern-Simons theories and enter the world of interacting quantum systems. We study conformal field theories described as the terminal point of renormalization group flows from Abelian Chern-Simons matter theories. Thus, in addition to the vector multiplets describing gauge fields, our field theories will now have charged chiral multiplets. We will find that, in close analogy with the study of $\mathcal{N} = 2$ theories in four-dimensions, such theories can

be geometrically encoded by studying the M5-brane on a *singular* manifold. In the context of three-manifolds branched over tangles the natural class of singularities are those where strands of the tangle collide and lose their individual identity. We refer to such objects as *singular* tangles. Our main aim in this section is to give a precise description of these objects and explain how they encode non-trivial conformal field theories. In the process we will also describe how the geometry encodes superpotentials. A summary of results in the form of a concise set of rules for converting singular tangles to physics appears in section 4.3.4.

4.3.1 Singularities and Special Lagrangians

We begin with a discussion of the geometric meaning of chiral multiplets and their associated wavefunctions in the three-sphere partition function. In our M-theory setting, the three-manifold M is embedded in an ambient Calabi-Yau Q , and massive particles arise from M2-branes which end along M on a one-cycle. In the simplest case of a spinless BPS chiral multiplet, supersymmetry implies that M is a special-Lagrangian and the M2-brane is a holomorphic disc as illustrated in Figure 4.9 [78,79]. The mass of the BPS particle is proportional to the area of the disc, and hence in the massless limit the cycle on which the M2-brane ends collapses.

Thus, when a particle becomes massless the three-manifold M develops a singularity. A local model for this geometry is a *special Lagrangian cone* on T^2 in \mathbb{C}^3 . Such




Figure 4.9: A particle represented by an M2-brane disc ending M . The M2-brane is the red disc located in the ambient space Q , while the dark blue circle represents the cycle of the three-manifold on which it ends.

a cone is defined to be the subset L_0 in \mathbb{C}^3 obeying [50]

$$L_0 = \{(z_1, z_2, z_3) \in \mathbb{C}^3 : |z_1|^2 = |z_2|^2 = |z_3|^2, \quad \text{Im}(z_1 z_2 z_2) = 0, \quad \text{Re}(z_1 z_2 z_3) \geq 0\}. \quad (4.3.1)$$

When the mass of the M2-brane is restored, the singularity is resolved. This can be done in three distinct ways [50]. Let $m > 0$, then the resolutions are

$$\begin{aligned} L_m^1 &= \{(z_1, z_2, z_3) \in \mathbb{C}^3 : |z_1|^2 - m = |z_2|^2 = |z_3|^2, \quad \text{Im}(z_1 z_2 z_2) = 0, \quad \text{Re}(z_1 z_2 z_3) \geq 0\}, \\ L_m^2 &= \{(z_1, z_2, z_3) \in \mathbb{C}^3 : |z_1|^2 = |z_2|^2 - m = |z_3|^2, \quad \text{Im}(z_1 z_2 z_2) = 0, \quad \text{Re}(z_1 z_2 z_3) \geq 0\}, \\ L_m^3 &= \{(z_1, z_2, z_3) \in \mathbb{C}^3 : |z_1|^2 = |z_2|^2 = |z_3|^2 - m, \quad \text{Im}(z_1 z_2 z_2) = 0, \quad \text{Re}(z_1 z_2 z_3) \geq 0\}. \end{aligned} \quad (4.3.2)$$

The resulting spaces are special Lagrangian three-manifolds in \mathbb{C}^3 [79] diffeomorphic to $S^1 \times \mathbb{R}^2$. They differ by the orientation of a closed holomorphic disc in \mathbb{C}^3 with

area πm which represents the M2-brane. In the case of L_m^1 this disc is given by

$$D_m^1 = \{(z_1, 0, 0) : z_1 \in \mathbb{C}, |z_1|^2 \leq m\}. \quad (4.3.3)$$

The other cases, D_m^2 and D_m^3 , are analogous. We see that the boundary of the disc is an oriented S^1 in L_m^1 whose homology class generates $H_1(L_m^1, \mathbb{Z}) \cong \mathbb{Z}$. In the other cases the boundary is given by an oriented circle around the origin of z_2 and z_3 respectively. One can thus see that the difference between the three ways the disc appears is determined by the orientation of its central axis in \mathbb{C}^3 .

To make contact with our discussion of tangles we view this local model for the singularity as a double cover over \mathbb{R}^3 . The special Lagrangians L_m^a are acted on by the involution

$$z_i \mapsto \bar{z}_i. \quad (4.3.4)$$

The quotient space is parametrized by the triple $(x_1, x_2, x_3) \in \mathbb{R}^3$ where $x_i = \text{Re}(z_i)$. Locally the x_i provide coordinates on L_m^a , but the global structure of the special Lagrangian is a double cover. The branch locus is the fixed points of (4.3.4) and is composed of two strands explicitly given by

$$\begin{aligned} L_m^1 : & x_1 = \sqrt{t^2 + m}, \quad x_2 = t, \quad x_3 = t, \quad \text{and} \quad x_1 = -\sqrt{t^2 + m}, \quad x_2 = t, \quad x_3 = -t, \\ L_m^2 : & x_1 = t, \quad x_2 = \sqrt{t^2 + m}, \quad x_3 = t, \quad \text{and} \quad x_1 = t, \quad x_2 = -\sqrt{t^2 + m}, \quad x_3 = -t, \\ L_m^3 : & x_1 = t, \quad x_2 = t, \quad x_3 = \sqrt{t^2 + m}, \quad \text{and} \quad x_1 = t, \quad x_2 = -t, \quad x_3 = -\sqrt{t^2 + m}, \end{aligned} \quad (4.3.5)$$

where $t \in \mathbb{R}$ provides a coordinate along the strands.

One way to see that the branched cover is an equivalent description of the original topology is to slice \mathbb{R}^3 into planes labelled by a time direction. The coordinate t on the branch lines in (4.3.5) provides such a foliation and increasing time defines a notion of *flow*. Each slice is a Riemann surface which is a double-cover of the plane branched over two points and is thus a cylinder. Therefore, including time, we see that topologically the cover is $\mathbb{R}^2 \times S^1$. We pursue this perspective on local flows in M and connect them to four-dimensional physics in section 4.5.

Returning to our analysis of the special Lagrangian cone, we note that when viewed as a double cover it is easy to see how the three different resolutions L_m^a are realized in terms of the configurations of the branch lines (4.3.5). We fix a planar projection of the geometry by declaring \hat{x}_3 to be the oriented perpendicular direction. Then, we can depict the geometry as in Figure 4.10. Note that Figure 4.10c only




Figure 4.10: The three different resolutions together with the M2-brane disc represented by a dashed line. The pictures are drawn in the projection to the (x_1, x_2) -plane. In (a) we see the branch locus giving rise to the special Lagrangian L_m^1 . In (b), the branch locus underlying L_m^2 . In (c) the branch locus underlying L_m^3 .

shows the *overcross*. The other choice, where the strand from upper left to lower right goes *under* the second strand, called the *undercross*, does not occur. This

is an artifact of the planar projection which we use to visualize the configuration. Indeed, exchanging the oriented normal \hat{x}_3 to $-\hat{x}_3$ exchanges the overcross for the undercross. By contrast, changing the normal direction from \hat{x}_3 to \hat{x}_1 or \hat{x}_2 permutes the resolutions appearing in 4.10 but leaves the triple, as a set, invariant.

In the limit $m \rightarrow 0$ the branch lines collide and we recover the singularity (4.3.1). In \mathbb{R}^3 , this appears as four branch half-lines all emanating from the origin. These half-lines approach infinity in four distinct octants and hence specify the vertices of a tetrahedron. In this way, we see the tetrahedral geometry of [66] emerge from the structure of special Lagrangian singularities.

Having thoroughly analyzed the local model, we may now introduce a precise definition of the concept of a *singular* tangle. It is simply a tangle where we permit pairs of strands to touch at a finite number of points. The local structure of the cover manifold M at each such point is that of the singular special Lagrangian cone discussed above, and the global identification of strands in the tangle indicates how these local models are glued together. In specifying the gluing we must keep track of additional pieces of discrete data.

- We draw singular tangles in planar projections of \mathbb{R}^3 . Hence each singularity is equipped with an oriented normal vector $\pm\hat{x}_3$. Varying the sign of the normal vector changes whether the overcross or undercross appears upon resolution.
- Fix a sheet labeling 1, and 2, at each singularity. Then in the gluing we must specify whether the identified sheets are the same or distinct. Varying between these two choices alters the relative signs of the charges of the particles as determined by the orientation of the M2-branes.

Both of the data described above have only a relative meaning: for a single singularity they are convention dependent while for multiple singularities they may be compared. All told then, if we draw singular tangles in a plane, each singularity is one of four possible types. We encode the four possibilities graphically with a thickened arrow on one of the strands passing through the singularity as in Figure 4.11. The thickened




Figure 4.11: Two different singularities. In (a) we see how an overcross singularity resolves after applying Figure 4.10c. In (b) the corresponding resolution is shown for the undercross singularity. In both cases the two other resolutions of Figure 4.10 are also present but not depicted.

strand always resolves out of the page while the direction of the arrow encodes the charge of the massless M2-brane residing at the origin of the singularity.

4.3.2 Wavefunctions and Lagrangians

Our next task is to explain in general how to extract a Lagrangian description of the physics defined by a singular tangle. As in the case of the free Abelian Chern-Simons theories studied in section 4.2, there is no unique Lagrangian but rather for each choice of Seifert surface we obtain a distinct dual presentation. In the case of singular tangles, we will see that these changes in Seifert surfaces are related by non-trivial mirror symmetries.

To begin, let us recall the data associated to a chiral multiplet in an Abelian Chern-Simons matter theory.

- A charge vector $q_\alpha \in \mathbb{Z}^{G+F}$ indicating its transformation properties under $U(1)^G \times U(1)^F$ gauge and flavor rotations. In all of our examples the vector q_α will be primitive meaning that the greatest common divisor of the integers q_α is one.
- A parity anomaly contribution. If a chiral multiplet is given a mass m , it may be integrated out leaving a residual contribution to the Chern-Simons levels of fields. The shift in the levels is given by

$$\delta k_{\alpha\beta} = \frac{1}{2} \text{sign}(m) q_\alpha q_\beta. \quad (4.3.6)$$

For primitive charge vectors the above shift has at least one non-integral entry. This implies that the ultraviolet levels are subject to a shifted half-integral quantization law. We take the associated shift to be part of the definition of the chiral multiplet.

- An R charge indicating the scaling dimension of the associated chiral operator in the conformal field theory. This data is fixed by a maximization principle once a superpotential is specified, and hence is not an additional data in the geometry [67]. This will be addressed in section 4.3.3.

To encode the partition function of such chiral multiplets we must introduce a new class of wavefunctions depending on these data. Each is given by a non-compact

quantum dilogarithm of the form

$$\begin{aligned} E_+(z - c_b(1 - R)) &\equiv e^{-i\frac{\pi}{2}z^2} s_b(-z + c_b(1 - R)), \\ E_-(z - c_b(1 - R)) &\equiv e^{i\frac{\pi}{2}z^2} s_b(z + c_b(1 - R)), \end{aligned} \quad (4.3.7)$$

where c_b is the imaginary constant given in (4.2.16), and the function $s_b(x)$, defined as

$$s_b(x) = e^{-i\frac{\pi}{2}x^2} \frac{\prod_{n=0}^{\infty} (1 + e^{(2n+1)\pi i b^2 + 2\pi b x})}{\prod_{n=0}^{\infty} (1 + e^{-(2n+1)\pi i b^{-2} + 2\pi b^{-1} x})}, \quad (4.3.8)$$

was obtained through a localization computation on the squashed three-sphere in [80] where the numerator and denominator come from vortex partition functions on the two half-spheres [81]. The physical interpretation of this function is read from the variables as follows.

- The subscript of E_{\pm} encodes the fractional ultraviolet Chern-Simons level $\pm\frac{1}{2}$ assigned to the particle.
- The variable z indicates the linear combination of gauge and flavor fields under which the chiral multiplet is charged. For E_{\pm} the charge is $z = \pm q \cdot (y \ x)$.
- The variable R denotes the R -charge.

Thus, we see that the physical data of a chiral multiplet is completely encoded by the wavefunctions (4.3.7). It follows that to assign a definite matter content to a singular tangle, as well as extract the associated contributions to the partition function \mathcal{Z} , it suffices to assign a quantum dilogarithm to each singularity. To proceed,

we introduce a singular Seifert surface Σ for a singular tangle L . As explained in section 4.2.3, from the homology of Σ we extract a basis of gauge and flavor cycles under which particles may be charged. Let α be such a cycle. Utilizing the sequence (4.2.51), we may view α equivalently as a cycle in the cover M . An M2-brane disc D ending on M has a charge determined by its linking numbers

$$q_\alpha = lk_\#(\alpha, \partial D). \quad (4.3.9)$$

The extension of this formula to the case of singular M is then depicted in our graphical notation in Figure 4.12.




Figure 4.12: The dilogarithm assignments for singularities. The particles are charged under the $U(1)$ super-field associated to the cycle α indicated in red, and x_α is the associated scalar. The overcross vs. undercross resolution encodes the distinction between E_\pm and specifies the fractional part of Chern-Simons levels. The orientation of the arrow on the thickened strand relative to α determines the sign of the particle's charge.

These dilogarithm assignments completely determine the matter content of a sin-

gular tangle. However, the assignments require a choice of Seifert surface. This surface is a choice of branch sheet for the double cover and varying it does not alter the underlying geometry. As a consequence, our rules are subject to the crucial test: *the underlying quantum physics must be independent of the choice of Seifert surface.*

Given the dualities between free Abelian Chern-Simons theories already described in section 4.2, independence of the choice of Seifert surface is ensured provided we have the equality shown in Figure 4.13. There, we see that one and the same singularity



Figure 4.13: A duality results from changing the Seifert surface. In (a) a singularity contributing $E_+(x_\alpha)$ to the partition function. In (b) the Seifert surface is changed and the same singularity contributes $E_-(x_\beta)$.

may make different contributions to an ultraviolet Lagrangian depending on the choice of Seifert surface. At the level of partition functions, this means that a singularity which contributes as $E_+(x_\alpha)$ with one choice of branch sheet can contribute with $E_-(x_\beta)$ with a different choice. Thus, we see that consistency of our analysis requires a mirror symmetry under which the same underlying conformal field theory may arise from ultraviolet theories with distinct matter content.

To understand the nature of the duality implied by Figure 4.13 we analyze its impact on the local model of the singular tangle involving a single singularity. Equality in more complicated examples follows from the locality of our constructions. The singular tangle together with its dual choices of Seifert surface and fixed compactifi-

cation data δ_i are shown in Figure 4.14. The ultraviolet field content in each case is




Figure 4.14: The duality between a free chiral multiplet and a $U(1)$ gauge field with a charged chiral field. In (a) we see the free chiral field coupled to the flavor cycle α . In (b) we see the gauge cycle β and flavor cycle α of the dual theory.

given by the following.

- Figure 4.14a: There is a background $U(1)$ flavor symmetry associated to the cycle α and no propagating gauge fields. Associated to the singularity there is a chiral multiplet with charge 1 under the flavor symmetry. This particle contributes $+\frac{1}{2}$ to the Chern-Simons level. The scalar x_α in the background $U(1)$ multiplet is the real mass of the chiral field.
- Figure 4.14b: There is a $U(1)$ flavor symmetry associated to the cycle α and a $U(1)$ gauge symmetry associated to the cycle β . Associated to the singularity is a chiral multiplet uncharged under the flavor symmetry and with charge -1 under the gauge symmetry. The level matrix, including classical contributions from the Trotter pairing as well as the fractional contributions of the particles is given by

$$K(\beta, \beta) = -\frac{1}{2}, \quad K(\alpha, \alpha) = 0, \quad K(\alpha, \beta) = 1. \quad (4.3.10)$$

The off-diagonal portion of the level implies that the scalar x_α is the FI-parameter of the gauged $U(1)$.

These two field theories are indeed known to form a mirror pair [62]. At the level of partition functions this equivalence is represented by a quantum dilogarithm identity, known as the *Fourier transform identity* [69]

$$E_+(x_\alpha - c_b) = \int dx_\beta e^{-2\pi i x_\alpha x_\beta} E_-(x_\beta). \quad (4.3.11)$$

The fact that our geometric description of conformal field theories provides a framework where this duality is manifest is a satisfying outcome of our analysis.

To gain further insight into this duality we now study resolutions of the singularity in both theories and interpret these from the viewpoint of three-dimensional physics. These resolutions correspond to motion onto the moduli space of the conformal field theory. From the perspective of the ultraviolet Lagrangians, the various branches of the moduli space can be described as Coulomb or Higgs branches, and the effect of the mirror symmetry is to exchange the two descriptions.¹⁴

The three different resolutions (4.3.5) have the following effect on the geometry of branch lines, see Figure 4.15. Let us start with the case (c). One can clearly see that the self-Chern-Simons level of the field α , as determined by the Trotter pairing, is one. This has a simple explanation from the point of view of field theory. Resolving the singularity means making the M2-brane massive with a mass $m \gg 0$. Thus the IR physics is obtained by integrating out this massive field which according to (4.3.6)

¹⁴Here and in the following the term *Coulomb branch* will be used generally to include the expectation values of scalars in both dynamical and background vector multiplets.




Figure 4.15: The three resolutions of the free Chiral field singularity. Part (a) corresponds to the Higgs branch. Part (b) and (c) represent motion onto the two sides of the Coulomb branch.

gives rise to a shift

$$\delta k_{\alpha\alpha} = \frac{1}{2} \text{sign}(m) q_\alpha q_\alpha = \frac{1}{2}. \quad (4.3.12)$$

Thus, as the ultraviolet Chern-Simons level was already one-half, the effective level is one exactly as the geometry of resolution (c) predicts. There is yet another way to see this. The limiting behavior of the quantum dilogarithm is as follows

$$E_+(m) \xrightarrow{m \rightarrow \infty} e^{-i\pi m^2}, \quad (4.3.13)$$

which again gives CS-level one in the effective theory as in our case $m = x_\alpha$. Resolution (b) corresponds to the other extreme where we take $m \ll 0$. This gives rise to

$$\delta k_{\alpha\alpha} = -\frac{1}{2}, \quad (4.3.14)$$

which results in an effective Chern-Simons level $k_{\alpha\alpha} = 0$. This is in complete accord with the geometry as cycle α has no self-linking after push-off in Figure 4.15b. Equivalently, this can be again seen in the limiting behaviour of the quantum dilogarithm

$$E_+(m) \xrightarrow{m \rightarrow -\infty} 1. \quad (4.3.15)$$

The two resolutions we have studied thus correspond to motion onto the Coulomb branch of the theory parameterized by the real mass m .

Now let us come to resolution (a) which is of a different nature. In order to understand what is happening we follow a path in the moduli space of the Joyce special Lagrangian starting from a point which corresponds to a resolution (b) or (c) to a point of resolution type (a). Along such a path the absolute value of the mass of the particle shrinks, as the volume of the M2-brane disc shrinks, until the field becomes massless at the singularity. As long as the field is massive it is not possible to turn on a vacuum expectation value for the scalar ϕ of the chiral multiplet as this would lead to an infinite energy potential. However, when we sit at the CFT point and the field is massless we can deform the theory onto the Higgs branch by activating an expectation value for ϕ . We draw the three branches of the theory schematically in Figure 4.16. We claim that motion onto the Higgs branch corresponds to resolution




Figure 4.16: Moduli space of a free chiral field.

(a) on the geometry. In order to see how this comes about we flip the Seifert surface to obtain the resolutions of the dual description of the theory as shown in Figure

4.17. In this dual theory resolution (a) arises from choosing $x_\beta \ll 0$ as can be seen




Figure 4.17: Resolutions of the theory dual to a free Chiral field.

from the limiting behavior of the negative parity quantum dilogarithm

$$E_-(x_\beta) \xrightarrow{x_\beta \rightarrow -\infty} 1. \quad (4.3.16)$$

Thus in the dual channel this resolution is obtained by giving a vev to the scalar part of a vector multiplet and therefore corresponds to a point on the Coulomb branch of the dual theory. But then the D-term equation of the dual theory requires that x_α be set to zero due to the Chern-Simons coupling of the two fields. Translating back to the original theory we indeed see that $m = x_\alpha = 0$ and that we have a propagating massless field. We therefore conclude that we are capturing the correct effective description of the physics on the Higgs branch. For completeness we note that the dual theory is on the Higgs branch for resolution (b) and on the Coulomb branch for resolution (c). This can be easily seen by noting the limiting behavior of the negative parity quantum dilogarithm for $x_\beta \gg 0$

$$E_-(x_\beta) \xrightarrow{x_\beta \rightarrow \infty} e^{i\pi x_\beta^2}. \quad (4.3.17)$$

The fact that resolutions of singular tangles capture motion onto the moduli space

of the corresponding conformal field theories is a general feature of our constructions which will be pursued in more detail in section 4.4.2.

4.3.3 Superpotentials From Geometry

There is one more ingredient in defining a three-dimensional theory with $\mathcal{N} = 2$ supersymmetry that we have yet to address: the superpotential. In this section we fill this gap. As with previous constructions, we find that the precise form of the superpotential as an explicit expression involving fields depends on a choice of Seifert surface used to construct a Lagrangian description.

The superpotential itself has a straightforward geometric interpretation in terms of M2-brane instantons, as described in [4]. Here we will briefly review that discussion. Consider some collection of massless chiral fields, X_i . Our M5-brane resides on a three-manifold M , which is a double cover of \mathbb{R}^3 branched over a singular tangle L . Meanwhile, the entire construction is embedded in an ambient Calabi-Yau Q . As studied above, each of the particles X_i corresponds to a singularity of the tangle L .

Given this setup, a superpotential interaction for the chiral fields X_i may arise from an instanton configuration of an M2-brane. This is a three-manifold C in Q , whose boundary ∂C is a two-cycle in M that intersects the particle singularities X_i . Consider the projection of the instanton $M2$ to one sheet of the double cover, ∂C_{\pm} . This must be a polygon bounded by the tangle L with vertices given by the singularities of X_i . A volume-minimizing configuration of this three-cycle will correspond to an interaction generated by a supersymmetric M2 instanton. This object is precisely of the correct geometric form to generate a superpotential term of the schematic form $\mathcal{W} = \prod_i X_i$.




Figure 4.18: Projections of BPS M2-brane instanton to the base. A portion of the branching tangle L is shown in black. The tangle has singular self-intersections supporting massless particles shown in blue. In (a), the interior of the polygon, shown in green, is the projection to M , of the boundary of an M2 instanton. In (b), we see the lift of the M2 instanton to the ambient manifold Q . Its boundary is doubled to an S^2 presented as two hemispheres glued along L . In the interior, this S^2 is filled in to make a three-ball.

To sharpen this discussion, there are several further considerations.

- The coefficient of the interaction is controlled by the instanton action, which is proportional to e^{-V} , where V is the volume of the supersymmetric three-manifold C . To generate a non-zero interaction, we need the three-manifold to have finite volume. Since our framework allows a non-compact manifold M with L going off to infinity, we must restrict our superpotential polygons on ∂C_{\pm} to be compact.
- The instanton action gets a contribution of $\exp(i \int_{\partial C} B)$, from the boundary of the M2 ending on the M5-brane. If $\partial C = 0$, that is, the boundary of the M2 is a trivial two-cycle, then this term is irrelevant. However in general, ∂C is a

non-trivial homology class and we find

$$\exp\left(i \int_{\partial C} B\right) = \exp(i\gamma), \quad (4.3.18)$$

where γ is a scalar field dual to a photon. This indicates the presence of a monopole operator $\mathcal{M}_j = \exp(\sigma + i\gamma)$ in the superpotential. So in this situation, we find a superpotential $\mathcal{W} = \mathcal{M} \prod_i X_i$. Of course, more generally ∂C is some integer linear combination of homology basis elements and so we might find that the superpotential contains a product of several monopole operators.

- The invariance of \mathcal{W} under all gauge symmetries apparent in the homology of the Seifert surface implies a compatibility condition on the discrete data living at the singularities bounding the associated polygonal region. To analyze the charge, we make use of the fact that the exact quantum corrected charge of the monopole operator is

$$q_\beta(\mathcal{M}_\alpha) = k_{\alpha\beta} - \frac{1}{2} \sum_{\text{Chirals } X_i} |q_\alpha^i| q_\beta^i, \quad (4.3.19)$$

where $k_{\alpha\beta}$ is the Chern-Simons level including both the integral part from the Trotter form, and the fractional contribution from particle singularities.

Given the above discussion, the next step is to analyze the explicit geometry of supersymmetric M2-brane instantons and determine which possible contributions in fact occur. This problem is important, but beyond the scope of this work. For our purposes we simply take as an ansatz that every possible gauge invariant contribution to the superpotential arising in the geometry as a polygon bounded by singularities

is in fact present.

With this hypothesis, to extract the superpotential in complete generality, we analyze a candidate contribution by expressing the boundary two-cycle ∂C in a basis of two-cycles $\{\beta_a\}$ dictated by the Seifert surface

$$\partial C = \sum_a c_a \beta_a, \quad c_a \in \mathbb{Z}. \quad (4.3.20)$$

For example, when utilizing the planar checkerboard Seifert surfaces discussed in section 4.2.3, the sum a ranges over compact un-colored regions, associated to gauge cycles, as non-compact un-colored regions associated to flavor cycles. Then, the term in question is

$$\prod_a \mathcal{M}_a^{c_a} \prod_{i \in \partial C} X_i. \quad (4.3.21)$$

We include such a term in the superpotential provided it is gauge invariant as dictated by the charge formula (4.3.19). The full superpotential is then a sum over all gauge invariant terms associated to all polygonal regions present in the tangle diagram of L .

Although it may seem cumbersome to explicitly calculate which polygons yield gauge invariant contributions to \mathcal{W} , in practice there is a simple sufficient, but not necessary, graphical rule which ensures gauge invariance that applies to the simplest class of contributions to the superpotential, namely polygons which lie entirely in the plane of a given projection of the Seifert surface. This rule is simply that the arrows on the singularities must circulate all in one direction around the gauge cycle in question. It may be easily derived from formula (4.3.19) as well as the charge

assignments of particles dictated by Figure 4.12. Examples of this type are shown in Figure 4.19. We encounter more general ‘non-planar’ superpotential terms in our




Figure 4.19: Projections of BPS M2-brane instanton, with the singular tangle in black. The particles X_i are indicated by the location of the black arrows, the Seifert surface is shaded in blue, and the projection of the instanton is shown in green. In (a), the M2 instanton projects to a trivial 2-cycle in M , and therefore has no monopole contribution. We find $\mathcal{W} = X_1 X_2 X_3$. In (b), the M2 projects to the non-trivial 2-cycle dual to the 1-cycle y shown on the Seifert surface. This contributes a monopole operator, yielding $\mathcal{W} = \mathcal{M}_y X_1 X_2 X_3$. The different shades of green are used here to indicate whether or not the polygon projection coincides with the Seifert surface.

analysis of examples in section 4.6.1.

4.3.4 Physics From Singular Tangles: A Dictionary

To conclude our discussion of singularities, we briefly summarize the algorithm for extracting an ultraviolet Lagrangian description of the physics associated to a singular tangle L .

- Pick a Seifert surface Σ . The homology $H_1(\Sigma_c, \delta, \mathbb{Z})$ specifies a basis of gauge and flavor cycles. Boundaryless cycles are dynamical gauge variables, while cycles with boundary are background flavor fields.

- Compute the Chern-Simons levels by computing the Trotter form on the homology $H_1(\Sigma_c, \delta, \mathbb{Z})$. In this procedure the singularities make fractional contributions to linking numbers. The singularities of plus type, illustrated in Figures 4.12a and 4.12b, contribute $1/2$. The singularities of minus type, illustrated in Figures 4.12c and 4.12d, contribute $-1/2$.
- Assign to each singularity a chiral field X_i . The field is charged under cycles on Σ passing through the singularity. The charge is $+1$ (-1) if the singularity is of plus type and the cycle is oriented with (against) the arrow at the singularity. The charge is -1 ($+1$) if the singularity is of minus type and the cycle is oriented with (against) the arrow at the singularity.
- Compute the superpotential by summing over gauge invariant contributions from closed polygonal regions in L . Each monomial entering in \mathcal{W} contains a product of chiral fields dictated by the vertices of the polygon, and possibly various monopole operators determined by expressing the polygon in a basis of two-cycles dual to $H_1(\Sigma_c, \delta, \mathbb{Z})$. Gauge invariance of the contribution of a given polygon is determined by application of the quantum corrected charge formula for monopole operators (4.3.19).

The physical theory associated to L is the infrared fixed point determined by this ultraviolet Lagrangian data. Varying the choice of Seifert surface, provides mirror ultraviolet Lagrangians, but does not alter the underling infrared dynamics.

In general the resulting theory is a strongly interacting system which enjoys a $U(1)^F$ flavor symmetry. The action of $Sp(2F, \mathbb{Z})$ on this conformal field theory is determined geometrically by the braid group action studied in section 4.2.4. The


three-sphere partition function \mathcal{Z} is an invariant of the theory which is extracted from this ultraviolet Lagrangian by generalizing the quantum-mechanical framework of section 4.2.1 and assigning to each singularity the quantum dilogarithm wavefunctions dictated by Figure 4.12.


In the remainder of this paper we apply these rules to further analyze the geometric description of mirror symmetries, and explore applications of the framework.

4.4 Dualities and Generalized Reidemeister Moves

In the previous sections we have developed a technique for extracting conformal field theories from singular tangles. However, there is still non-trivial redundancy in our description: as a consequence of mirror symmetry, two distinct singular tangles may give rise to equivalent quantum field theories. In this section, we determine the equivalence relation implied on singular tangles by mirror symmetries, and explore their geometric content.

In searching for such relationships, one may take inspiration from the case of non-singular tangles. In that case, the basic relations are the Reidemeister moves shown below.





These moves are local and may be applied piecewise in any larger tangle diagram. Further, these moves are a generating set for equivalences: any two tangles which are isotopic may be related to one another by a sequence of Reidemeister moves.

In the case of singular tangles, we find similar structure. Basic mirror symmetries determine relations on singular tangles which take the form of *generalized* Reidemeister moves. They are related to the moves presented above by replacing some crossings by singularities. Further, each of these equivalences is local, and hence they may be applied piecewise in a larger singular tangle to engineer more complicated relations. It is natural to conjecture that these generalized Reidemeister moves, together with the Torelli dualities of section 4.2.4 provide a complete set of quantum equivalence relations on singular tangles.

In section 4.4.1 we present a detailed description of the generalized Reidemeister moves as well as the associated quantum dilogarithm identities that result from application of these moves to partition functions. In section 4.4.2 we show how deformations away from the conformal fixed point result resolve generalized Reidemeister moves into the ordinary Reidemeister moves.

4.4.1 Generalized Reidemeister Moves

In this section we present the list of generalized Reidemeister moves. Each takes the form of a graphical identity involving two singular tangles. The precise form of these equalities depends on the discrete data living at the singularities. There are two

things to note about this dependence which follow immediately from our analysis of the local model in section 4.3.1.

- If we flip arrows at all singularities by 180 degrees on both sides of an identity, it still holds. Indeed, such a flip is equivalent to reflecting the sign of all $U(1)$ gauge and flavor groups. Geometrically, this is equivalent to globally changing the labeling of sheets from 1 to 2 in the double cover.
- If we exchange all overcross and undercross of all singular and non-singular crossings on both sides of an identity, the identity still holds. This is true because each of our diagrams is drawn in a fixed projection with oriented normal vector \hat{x}_3 . Globally reflecting $\hat{x}_3 \rightarrow -\hat{x}_3$ generates the indicated transformation on diagrams, as shown for example in Figure 4.20.




Figure 4.20: Reflection in the projection plane transforms an overcross singularity to an undercross singularity.

In the following, we take these two principles into account and thereby present a reduced set of generalized Reidemeister moves. Additional dualities may be generated by changing the discrete data at the singularities as above.

Rules Descending from Move 1

Here, we consider a singular version of the first Reidemeister move. Populating the singular tangles with a Seifert surface generates partition function identities. We

will look at two such choices of Seifert surface differing by black-white duality.



With a choice of planar Seifert surface, we have the following two interpretations.


Singularity Transform



$$E_+(x - c_b(1 - R))e^{i\pi x^2} = E_-(-x + c_b(1 - R))$$

In quantum mechanics language this is equivalent to starting with a quantum dilogarithm and applying a T -transformation. This does not involve any integrals, as the quantum dilogarithm is an eigenstate of the T -operator. Hence there is also no gauge group in the 3d gauge theory interpretation. The only effect on the theory is a decrease in the background Chern-Simons level by one unit.

Fourier Transform



$$\int dz E_-(z - x + c_b(1 - R))e^{-i\pi(z-y)^2} = E_+(y - x - c_b(1 + R)/2)$$

This represents a duality containing a $U(1)$ gauge field on the one side but no gauge field on the other. This rule is equivalent to the Fourier transformation identity

discussed in section 4.3.1, and is another singular-tangle representation of that duality. Here, the theory of one $U(1)$ gauge field at level one-half together with a charged chiral particle is mirror to a free chiral field.


Rules Descending from Move 2

The second Reidemeister move can be generalized to give rise to an identity between singular tangles where neighbouring singularities cancel pairwise such that on the other side of the identity there is no singularity at all. Therefore, we denote these identities with the term *pairwise cancellation of singularities*. We will also examine a partition function identity inherited from the tangle identity for one choice of Seifert surface. The relevant singular tangle identities are the following.


The top diagram shows a tangle with two strands that cross and then cancel each other out, resulting in a simple loop. The bottom diagram shows a tangle with two strands that cross and then cancel, resulting in two parallel strands.

From the perspective of the 3d gauge theory these can be understood as follows. We have a closed polygonal region bounded by two singularities. As discussed in section 4.3.3 this gives rise to a superpotential with the two chiral fields. Thus the particles are given mass and make no contribution to the infrared physics. The dual theory then contains no particles, but depending on the UV Chern-Simons levels it can contain background Chern-Simons levels.

Picking a Seifert surface these rules translate to the following quantum dilogarithm identities.



$$E_+(-x + c_b(1 - R))E_+(x - c_b(1 - R)) = e^{-i\pi x^2}$$




$$E_+(x - c_b(1 - R))E_-(x - c_b(1 - R)) = 1$$

From this perspective, the underlying identity of pairwise cancellation of singularities is equation (16) in the appendix of reference [69].


Rules Descending from Move 3

The most important rule arises from singularization of the third Reidemeister move. This rule is called the 3-2 move and encodes a non-trivial three-dimensional mirror symmetry. In this section we will clarify its relation to the third Reidemeister move by singularizing all crossings on one side of the identity and only two on the other side. Apart from the 3-2 move, the third Reidemeister move can be singularized by adding only one singularity on both sides. This application follows from the previously identified Fourier transform identity and hence does not represent an independent mirror symmetry. Nevertheless, the simple application is useful when moving between Seifert surfaces in the examples of section 4.5 and 4.6. We will turn to this simple application first and then discuss the 3-2 move.

Change of Branch sheet Applying the Fourier transform identity of Figure 4.13 locally, we obtain a generalization of the third Reidemeister move. On one side of the duality we have a theory with a chiral particle charged under a $U(1)$ gauge field which in turn couples to two background gauge fields. The duality relates this theory to one with no gauge group, a chiral multiplet and two flavor fields. The partition function equality is again an application of Figure 4.13.




$$\int dw E_-(w-x+c_b(1-R)) e^{-i\pi(w-z)^2+i\pi(w-y)^2} = E_+(y-z-c_b(1+R)/2) e^{-i\pi(z-x)^2+i\pi(x-y)^2}$$




$$\int dw E_-(w-x+c_b(1-R)) e^{-i\pi(w-y)^2+i\pi(w-z)^2} = E_+(z-y-c_b(1+R)/2) e^{-i\pi(y-x)^2+i\pi(x-z)^2}$$

The 3-2 move The relevant singular tangle identity is depicted below.



We clearly see that this identity relates a theory with three chiral fields to the one with just two chiral fields. Such theories are known to come in mirror pairs in three dimensions [62, 64, 65, 82]. Examining the left-hand-side we notice the presence of a closed polygonal region bounded by three singularities and hence the existence of a superpotential. To extract the physical content we choose Seifert surfaces as shown below.



$$\begin{aligned}
 & E_+(z - y - c_b(1 - r)) E_-(z - x + c_b(1 - s)) E_-(x - y - c_b(1 - r - s)) \\
 & = \int dw E_+(y - w - c_b) E_+(z + w - c_b(1 - r)) e^{i\pi(x-w)^2}
 \end{aligned}$$

The physical theories are then read off:

- LHS: A theory with three chiral fields X, Y, Z no gauge symmetry and a cubic superpotential $\mathcal{W} = XYZ$, known as the XYZ -model.
- RHS: A theory with a gauged $U(1)$ with vanishing self Chern-Simons level and two oppositely charged chiral fields Q and \tilde{Q} , known as $U(1)$ super-QED with $N_f = 1$.

These theories are known to form a mirror pair [65]. At the level of partition functions this duality is the *pentagon identity* for quantum dilogarithms [69].

4.4.2 Resolutions of Singularities

In this section we make the connection between generalized Reidemeister moves and ordinary Reidemeister moves precise. We show that by moving onto the moduli space of the conformal field theories on both sides of a generalized Reidemeister move, we find the identity resolved into an ordinary Reidemeister move. To achieve this we will choose a particular Seifert surface such that all the resolutions in question are obtained as a motion onto the Coulomb branch. In general such a deformation gives masses to all chiral fields and in the infrared they can be integrated out. Generically, this leads to a fractional shift in the Chern-Simons levels of the form [65]

$$(K_{IJ})_{\text{eff}} = K_{IJ} + \frac{1}{2} \sum_{a=1}^{N_f} (q_a)_I (q_a)_J \text{sign}(m_a) \in \mathbb{Z}, \quad I, J = 1, \dots, G + F, \quad (4.4.1)$$

where we have noted that the effective levels are integral in order to ensure gauge invariance. These effective levels are depicted in Figure 4.21 as applied to a single singularity as studied in section 4.3.2.




Figure 4.21: Resolution of singularities by turning on Fayet-Iliopoulos or Coulomb branch parameters. In both parts, (a) and (b), m is the argument of the relevant quantum dilogarithm.

In applying this logic to study resolutions of singular tangles, one must take care to remain in a supersymmetric vacuum. In other words the F - and D -term equations have to be satisfied. This will be dealt with next.

F- and D-term Equations

Let us elaborate the Coulomb branch resolutions from the viewpoint of the 3d gauge theory. The singular tangle describes the CFT at the origin of the Coulomb and Higgs branches. If we discuss only resolutions which remain at the origin of the Higgs branch then the resulting resolutions correspond to different leaves of the Coulomb branch parameterized by Fayet-Iliopoulos parameters and scalar fields in vector multiplets.

In order to determine which resolutions are possible in a complicated singular tangle we need to solve the D- and F-term equations of the relevant 3d gauge theory. The potential \mathcal{V} for the theory is a sum of a D-term and an F-term contributions of the form

$$\mathcal{V} = \mathcal{V}_D + \mathcal{V}_F. \quad (4.4.2)$$

In a supersymmetric vacuum this potential must vanish. As both \mathcal{V}_D and \mathcal{V}_F are non-negative, both must vanish separately.

Let us first consider the F-term potential which reads

$$\mathcal{V}_F = \sum_{a=1}^{N_f} \left| \frac{\partial \mathcal{W}}{\partial \phi_a} \right|^2, \quad (4.4.3)$$

where \mathcal{W} is the superpotential of the theory and ϕ_a is the scalar component of the

chiral field X_a . In our geometric examples, \mathcal{W} arises from a sum over polygons and hence each monomial in \mathcal{W} has degree larger than one in the ϕ_a fields. It follows that if we remain at the origin of the Higgs branch $\phi_a = 0$ the F-term potential is trivially minimized.

Let us next turn to the D-term potential. In the following we will drop the subscript eff from all Chern-Simons levels and assume that the IR limit has been taken. The D-term potential is then given by

$$\begin{aligned} \mathcal{V}_D = & \sum_{i,j} \frac{k_{ij}}{2\pi} D_i y_j + \sum_{i,\lambda} \frac{k_{\lambda i}}{2\pi} x_\lambda D_i \\ & + \sum_{a,i} q_{a,i} D_i |\phi_a|^2 + \sum_{a,i} |q_{a,i} y_i|^2 |\phi_a|^2, \end{aligned} \quad (4.4.4)$$

where the summation is over $i, j = 1, \dots, G$ for the gauge indices, and $\lambda = 1, \dots, F$ for the Fayet-Illiopoulos parameters x_λ . The associated D-term equation then reads

$$\frac{\partial \mathcal{V}_D}{\partial D_i} = \sum_j \frac{k_{ij}}{2\pi} y_j + \sum_\lambda \frac{k_{i\lambda}}{2\pi} x_\lambda + \sum_a q_{a,i} |\phi_a|^2 = 0. \quad (4.4.5)$$

On the Coulomb branch we have that $\phi_a = 0$ which simplifies the above equation considerably. Defining

$$K_{IJ} := \begin{pmatrix} k_{ij} & k_{i\lambda} \\ k_{\lambda i} & k_{\lambda\mu} \end{pmatrix}, \quad \Sigma_J := \begin{pmatrix} y_i \\ x_\lambda \end{pmatrix}, \quad (4.4.6)$$

it is possible to write equation (4.4.5) in the compact form

$$K_{iJ} \Sigma_J = 0, \quad (4.4.7)$$

for $i = 1, \dots, G$.

Equation (4.4.7) is our desired result. It implies that, restricting attention to Coulomb branch deformations, we can determine which deformations are allowed by searching for null-vectors of the effective level matrix K .

Resolution of Move Descending from Rule 1

Here, we examine how a particular resolution on the two sides of our first generalized Reidemeister move yields the ordinary Reidemeister move of first kind. In order to proceed, we need to pick a particular Seifert surface which allows us to obtain the relevant resolution as motion onto the Coulomb branch. We will pick the second Seifert surface corresponding to the dilogarithm identity

$$\int dz E_-(z - x + c_b(1 - R)) e^{-i\pi(z-y)^2} = E_+(y - x - c_b(1 + R)/2). \quad (4.4.8)$$

The limit we take is the following

$$y \rightarrow -\infty, \quad x \rightarrow \infty, \quad z \rightarrow -\infty, \quad (4.4.9)$$

resulting in

$$E_-(y - x - c_b(1 + R)/2) \rightarrow 1, \quad E_+(z - x + c_b(1 - R)) \rightarrow 1. \quad (4.4.10)$$

The effective Chern-Simons levels of the left-hand-side become

$$k_{zz} = 1, \quad k_{zy} = -1, \quad (4.4.11)$$

which in turn lead to the D-term equation (4.4.7)

$$z - y = 0. \quad (4.4.12)$$

As this is consistent with the limit prescribed, we are indeed looking at a valid resolution satisfying the equations of motion of the gauge theory. The pictorial representation is shown in Figure 4.22. We clearly see that the resolution reproduces the ordinary first Reidemeister move as claimed.




Figure 4.22: Resolution of the first generalized Reidemeister move.

Resolution of Moves Descending From Rule 2

Next, we look at resolutions of the second generalized Reidemeister move. This rule consists of two parts and we shall examine both of them. Again we have to pick a Seifert surface which we choose to be the same as in section 4.4.1. The relevant quantum dilogarithm identity for the first subrule is

$$E_+(-x + c_b(1 - R))E_+(x - c_b(1 - R)) = e^{-i\pi x^2}. \quad (4.4.13)$$

Here we can consider the following limit

$$x \rightarrow \infty : \quad E_+(-x + c_b(1 - R))E_+(x - c_b(1 - R)) \rightarrow 1 \cdot e^{-i\pi x^2} = e^{-i\pi x^2}. \quad (4.4.14)$$

As the limit gives the right hand side of the identity trivially there is nothing to be checked. Therefore, this resolution does not involve any Reidemeister moves.

Let us now move to the second subrule. The relevant quantum dilog identity is

$$E_+(x - c_b(1 - R))E_-(x - c_b(1 - R)) = 1. \quad (4.4.15)$$

Taking the limit $x \rightarrow \infty$ the left-hand-side becomes

$$E_+(x - c_b(1 - R))E_-(x - c_b(1 - R)) \rightarrow e^{-i\pi x^2}e^{i\pi x^2}. \quad (4.4.16)$$

The pictorial representation of this resolution is the second Reidemeister rule, as shown in figure 4.23.




Figure 4.23: Resolution of the second generalized Reidemeister move.

Resolution of Move Descending From Rule 3

Let us now come to our last and most involved case, namely the 3-2-move. The relevant identity here is

$$\begin{aligned} E_+(z - y - c_b(1 - r))E_-(z - x + c_b(1 - s))E_-(x - y - c_b(1 - r - s)) \\ = \int dw E_+(z - w - c_b)E_+(w - y - c_b(1 - r))e^{i\pi(x-w)^2}. \end{aligned} \quad (4.4.17)$$

Defining

$$x - y \equiv c_1, \quad z - x \equiv c_2, \quad z - y \equiv c_3, \quad (4.4.18)$$

we will consider the limit

$$c_i \gg 0 \quad \text{for } i = 1, 2, 3. \quad (4.4.19)$$

As the above equation set implies the relation $c_1 = -c_2 + c_3$ we find that

$$c_3 - c_2 \gg 0. \quad (4.4.20)$$

Setting $w \equiv c_3$ ensures that we have the effective Chern-Simons-levels

$$k_{wx} = 1, \quad k_{wy} = -1, \quad k_{zw} = -1, \quad k_{ww} = 1. \quad (4.4.21)$$

The D-term equation (4.4.7) thus gives

$$x - y + w - z = c_3 - c_2 + c_3 + c_2 - 2c_3 = 0, \quad (4.4.22)$$

and hence confirms that we are on the Coulomb-branch. The pictorial representation




Figure 4.24: Resolution of the 3-2-move.

of the limit discussed is the third Reidemeister move as shown in Figure 4.24.

4.5 R-flow

We have seen how singular tangles capture the content of a 3d conformal field theory with four supercharges, and that resolutions of such objects describe dynamics on the moduli space of the same theory. This is very similar to the Seiberg-Witten description of the Coulomb branch of 4d gauge theories with eight supercharges. In fact, the similarity goes even further. In the Seiberg-Witten case, the multicover of a complex curve with punctures captures all the information about the BPS states of the four-dimensional gauge theory [24, 41, 42, 83]. In our case a multicover (more specifically, a double cover) of \mathbb{R}^3 with specified boundary conditions captures the content of a three-dimensional theory. The connection of these two descriptions can be made precise by looking at specific class of examples where the three-manifolds

in question arise from *flows* of a Seiberg-Witten curve of a 4d theory. By this we mean that there exists a *slicing* of the three-manifold along a ‘*time*’ direction such that each slice represents a SW-curve. This construction is known as *R-flow* and has been studied in [4, 18]. This section is devoted to the definition and properties of R-flow. R-flow is defined on the space of central charges of certain 4d $\mathcal{N} = 2$ theories and describes a domain wall solution which has the interpretation of a 3d $\mathcal{N} = 2$ theory [84–86].

4.5.1 Definition of R-Flow

R-flow is a motion in the space central charges of four-dimensional theories with eight supercharges. We focus on complete theories [3], in which deformations in the space of central charges are locally equivalent to deformations of branch points of the Seiberg-Witten curve. Then R-flow takes the following form

$$\frac{d}{dt} Z_i = i \text{Re} Z_i, \quad (4.5.1)$$

where Z_i is the central charge of the i -th charge in the $\mathcal{N} = 2$ 4d theory. Central charges flow in time along straight lines preserving their real parts while their imaginary parts move at a rate which is proportional to their real parts. As a consequence of this flow equation, the phase ordering of central charges is preserved and hence the entire evolution takes place in a fixed BPS chamber. In summary, we can say that *phase ordering* is *time ordering* and depict this in a graph shown in Figure 4.25. This describes a three-dimensional theory as a domain-wall solution of the four-dimensional parent theory where each 4d BPS state gives rise to a 3d BPS state whose mass is

given by the real part of Z_i .




Figure 4.25: R-flow for an example with three central charges.

4.5.2 A_n flow and the KS-operator

In this paper, we are particularly interested in flows of 4d gauge theories which arise from wrapping a M5-brane on a Riemann surface of the type A_n describing Argyres-Douglas CFTs [41, 87]. These surfaces are double covers of the \mathbb{C} -plane, in the form

$$y^2 = (x - a_1)(x - a_2) \cdots (x - a_n)(x + a_{n+1}), \quad (4.5.2)$$

where $a_{n+1} = \sum_{i=1}^n a_i$. The Seiberg-Witten differential is given by the square root of the quadratic differential

$$\phi = (x - a_1)(x - a_2) \cdots (x - a_n)(x + a_{n+1})dx^2, \quad (4.5.3)$$

i.e. $\lambda_{\text{SW}} = \sqrt{\phi}$. Having established the above definitions, it is straightforward to write down the central charges of the theory:

$$\begin{aligned} Z_1 &= \int_{a_1}^{a_2} \sqrt{\phi}, \\ Z_2 &= \int_{a_2}^{a_3} \sqrt{\phi}, \\ &\vdots \\ Z_n &= \int_{a_n}^{a_{n+1}} \sqrt{\phi}. \end{aligned} \quad (4.5.4)$$

Now, choosing an ordering of the phases of the central charges, one arrives in a specific chamber of the moduli space, with some particular set of BPS particles. For the choice

$$\arg Z_1 < \arg Z_2 < \cdots < \arg Z_n, \quad (4.5.5)$$

we obtain the so called *minimal chamber* with exactly n stable particles. On the other hand, the *maximal chamber* is defined for the configuration

$$\arg Z_n < \arg Z_{n-1} < \cdots < \arg Z_1. \quad (4.5.6)$$

Here the number of stable BPS particles is $\frac{1}{2}n(n + 1)$ [25]. There are also various

intermediate chambers, and we shall denote the number of states in a given chamber by N . Note that each of these states is a linear combination of the primitive basis given in (4.5.4), and its central charge is correspondingly a linear combination of the central charges computed therein. We next assign to each central charge ordering a Kontsevich-Soibelman operator of the following form [32, 52, 53]:

$$\mathbb{K}(q) = \prod_i^N E_+(\hat{\gamma}_i), \quad (4.5.7)$$

where E_+ is the non-compact quantum dilogarithm, and the $\hat{\gamma}_i$ label the stable BPS states. The $\hat{\gamma}$'s can be interpreted as phase space variables of the quantum Hilbert space which differ by actions of $Sp(n, \mathbb{Z})$ if n is even and $Sp(n-1, \mathbb{Z})$ if n is odd. From the point of view of the A_n curve the $\hat{\gamma}_i$ represent cycles determined by two branch points a_k and a_l . In particular, from the point of view of the quantum mechanics description of section 4.2.1, they are linear combinations of \hat{x}_i and \hat{p}_i and are mapped to each other by actions of the generators

$$\sigma_{2j-1} = \exp(-i\pi\hat{x}_j^2), \quad \sigma_{2j} = \exp(-i\pi\hat{p}_j^2). \quad (4.5.8)$$

We can assign to each KS-operator a quantum mechanical matrix element of the form

$$\mathcal{Z}_{\mathbb{K}} = \langle x | \mathbb{K} | y \rangle, \quad (4.5.9)$$

which have an interpretation as partition functions of 3d theories as discussed in section 4.2.1. These partition functions enjoy a $Sp(n, \mathbb{Z}) \times Sp(n, \mathbb{Z})$ action which

has the interpretation of the braid group action on the two ends of a braid with $n + 1$ strands. Recalling the constructions of sections 4.2, 4.3, we can thus assign a singular braid $B_{\mathbb{K}}$ to the matrix element $\mathcal{Z}_{\mathbb{K}}$. This is depicted schematically in Figure 4.26. As also indicated there, the braid naturally defines a time direction that we can understand as follows. Each line of the braid describes the flow of a branch point of the A_n -curve along the time direction. At the singularities, these branch points collide, thereby loosing their individual identities.




Figure 4.26: For each KS-operator there is an associated singular braid $B_{\mathbb{K}}$.

Let us zoom into the braid $B_{\mathbb{K}}$ to see how the strands approach each other for an isolated singularity. To this end, we rewrite the partition function as a gluing of three braids according to the formalism developed in section 4.2.1

$$\mathcal{Z}_{\mathbb{K}} = \int dx' dy' \langle x | \cdots | y' \rangle \langle y' | E_+(\hat{\gamma}_{kl}) | x' \rangle \langle x' | \cdots | y \rangle, \quad (4.5.10)$$

where $\hat{\gamma}_{kl}$ represents the contribution of the 4d BPS state whose central charge is




Figure 4.27: Braid realization of a local singularity. The relevant branch points come close to each other until they collide in the singularity and loose their individual identities. After that they depart again until they reach their original positions in the braid.

given by

$$Z_{kl} = \int_{a_l}^{a_k} \sqrt{\phi}. \quad (4.5.11)$$

Zooming into the braid we then have the local representation for an isolated singularity shown in Figure 4.27. Resolving the singularity means turning the points at which the branch points touch to near misses. As we have seen, for each singularity there are exactly three ways to do this. R-flow, as a flow of branch points of the Seiberg-Witten curve, is equivalent to choosing the resolution of Figure 4.10 (b) for all singularities. Said differently, the singular braid $B_{\mathbb{K}}$ is obtained from the flow defined by equation (4.5.1) in the limit in which all near misses are replaced by singularities.

Let us now come to the justification of this picture. The initial condition of R-flow is determined by the chamber in which the flow starts. Furthermore, as the flow continues one stays in the initial chamber due to the phase-order preservation of the

flow. As central charges cross the real axis something special happens. Recall that a 4d BPS hypermultiplet has an interpretation as a geodesic on the complex plane between branch points of the Riemann surface [26, 27]. These geodesics obey the equation

$$\sqrt{\phi} = e^{i\theta_m} dt, \quad (4.5.12)$$

where θ_m , $m = 1, \dots, N$ is the phase of the m th BPS state, i.e.

$$\theta_m = \arg Z_m. \quad (4.5.13)$$

There are two remarks in order here. First, R-flow describes a motion on the Coulomb branch (including mass parameters) of the four-dimensional gauge theory. On the other hand, the flow equation (4.5.12) is a flow on the \mathbb{C} -plane at a *fixed* point in the moduli space. The Seiberg-Witten curve, being a double-branched cover of the \mathbb{C} -plane, is not subject to change under the flow (4.5.12). Therefore, in order to relate the two motions, we have to choose a fixed angle θ_m corresponding to a line in the complex plane of central charges. Secondly, the geometry of R-flow predestines exactly such a line, namely the real axis which defines a mirror axis for the flow. Each time a central charge crosses the real axis there is a geodesic solution with minimal length. Thus at such points the pair of branch points corresponding to the BPS bound state whose central charge crosses the real axis are closest. Away from the real axis, the pair of branch points repel in a symmetric way, thus producing the near-miss behavior in Figure 4.10 (b).

4.5.3 Examples

In this section present some examples of R-flow. We start with the simplest case and proceed to increasing complexity. Already in the very first example, the A_1 flow, we will find that R-flow gives insight into the behavior of branch lines near local singularities.

A_1 flow

As a first example we will consider the most simple case of R-flow. This is the theory corresponding to the curve

$$y^2 = x^2 + \epsilon, \quad (4.5.14)$$

with a single central charge, denoted by Z_1 , given by

$$Z_1 = \int_{-\sqrt{\epsilon}}^{\sqrt{\epsilon}} \sqrt{x^2 + \epsilon} dx = -\frac{\pi i}{2} \epsilon. \quad (4.5.15)$$

We will find that this theory has significant importance for the resolution of arbitrary singular tangles as it produces the various possible local resolutions of an isolated singularity via Fayet-Iliopoulos parameters. Let us describe how this comes by. First of all, note that we can parametrize ϵ as $\epsilon = -\frac{2}{\pi}(-im + t)$ with m real and positive so that

$$Z_1 = m + it, \quad m > 0, \quad (4.5.16)$$

obeys the flow equation (4.5.1)¹⁵. The motion of the branch points of the curve are then given by the law

$$a_1 = \alpha\sqrt{m+it}, \quad a_2 = -\alpha\sqrt{m+it}, \quad (4.5.17)$$

where α is a proportionality constant. We can now view this motion from two perspectives. The first is as a motion on the \mathbb{C} -plane which forms the base of the double cover. The second perspective is obtained by looking at the motion of the two branch points as giving rise to branch lines in $\mathbb{C} \times \mathbb{R}$ where \mathbb{R} is the *time*-direction parametrized by t . As the square root behaviour of (4.5.17) is fairly simple we can depict the two perspectives easily as shown below in Figure 4.28. A very interesting phenomenon




Figure 4.28: $m > 0$. Part a) depicts the motion of branch points of the A_1 -curve on the \mathbb{C} -plane, where the arrows indicate direction of flow with increasing time t . Part b) depicts the motion as branch lines in $\mathbb{C} \times \mathbb{R}$.

happens when we flip the sign of the real part of the central charge, i.e. if we choose $m < 0$ instead. Fixing the projection plane, we now obtain the following picture for the branch-point flow (Figure 4.29).

¹⁵Note that the derivative with respect to t is not equal to m . This is no problem however, as this condition was imposed initially to maintain the order of the central charges along the flow. But as the A_1 curve has just one central charge, we just demand that the real part stays constant.




Figure 4.29: $m < 0$. Part a) depicts the motion of branch points of the A_1 -curve on the \mathbb{C} -plane while part b) describes the motion as branch lines in $\mathbb{C} \times \mathbb{R}$.

We see that this analysis exactly mirrors two of the three possible resolutions described in section 4.3.1, namely resolutions 4.10 (b) and (c). Note that resolution (a) cannot be obtained in this formalism as it breaks time-flow or equivalently keeps the mass parameter m at zero but deforms the theory onto the Higgs branch.

A_2 flow

We now turn to our next example, the A_2 curve. It is, apart from the A_1 case, the most important flow example as it provides insight into three-dimensional mirror symmetry in terms of flows of four-dimensional theories. In order to illustrate this we consider the two central charge orderings of this theory which provide two BPS chambers with distinct particle content. We have a 2-particle chamber,

$$\arg Z_1 < \arg Z_2 < 0, \quad (4.5.18)$$

and a three-particle chamber,

$$\arg Z_2 < \arg Z_1 < 0, \quad (4.5.19)$$

where the third state is the one with charge $Z_1 + Z_2$. Looking at the Kontsevich-Soibelman operator we see that in the first case it is given by

$$E_+(\hat{p})E_+(\hat{x}), \quad (4.5.20)$$

while in the second case one has

$$E_+(\hat{x})E_+(\hat{x} + \hat{p})E_+(\hat{p}). \quad (4.5.21)$$

The crucial point here is that these two operators are actually equal if we impose the commutator

$$[\hat{x}, \hat{p}] = \frac{i}{2\pi}, \quad (4.5.22)$$

as was first proven in [69]. This is the underlying equality leading to the 3-2-move discussed in section 4.4.1. Therefore, the 3-2-move can actually be thought of as arising from R-flow of the A_2 curve. However, note that the 3-2 move is obtained by looking at matrix elements $\langle x|\mathbb{K}|p\rangle$, that is position/momentum matrix elements, whereas R-flow is equivalent to matrix elements of the form $\langle x|\mathbb{K}|y\rangle$, namely position/position matrix elements. Furthermore, there are many braid realizations of these matrix elements differing by the other various dualities discussed in section 4.4.1. In this section we will look at representations which are obtained from the prescription described in Figure 4.27. That is, we will now look at the above KS-operators and their braid realizations from the perspective of branch-point flow.

Let us start with the minimal particle chamber. Using the identity

$$E_+(\hat{p}) = e^{i\pi\hat{x}^2} e^{i\pi\hat{p}^2} e^{i\pi\hat{x}^2} E_+(\hat{x}) e^{-i\pi\hat{x}^2} e^{-i\pi\hat{p}^2} e^{-i\pi\hat{x}^2}, \quad (4.5.23)$$

we obtain

$$\begin{aligned} \langle x | E_+(\hat{p}) E_+(\hat{x}) | y \rangle &= \langle x | e^{i\pi\hat{x}^2} e^{i\pi\hat{p}^2} e^{i\pi\hat{x}^2} \hat{x} e^{-i\pi\hat{x}^2} e^{-i\pi\hat{p}^2} e^{-i\pi\hat{x}^2} E_+(y) | y \rangle \\ &= \int dx' \langle x | \sigma_1^{-1} \sigma_2^{-1} E_+(x') | x' \rangle \langle x' | \sigma_2 \sigma_1 E_+(y) | y \rangle. \end{aligned} \quad (4.5.24)$$

This way we have rewritten the partition function in terms the σ_i which describe actions of the braid group. The braid representation of the right-hand side of the above identity is shown in Figure 4.30¹⁶. The single integration variable in (4.5.24)




Figure 4.30: The singularized braid of the A2 flow in the minimal chamber.

corresponds to a $U(1)$ gauge group manifest as a compact white region in Figure 4.30. Furthermore, we have used that σ_1 and σ_1^{-1} commute with $E_+(\hat{x})$ and therefore cancel each other. Note that the theory described by the braid in Figure 4.30 is related to $U(1)$ SQED by changing the branch sheet as discussed in section 4.4.1¹⁷.

¹⁶We have suppressed the R-charges of the singularities as these are not relevant for the present discussion.

¹⁷We also need to apply an S -transformation to the boundary condition in order to switch from position boundary to momentum boundary.

We will not discuss this here and rather turn our attention to a particular resolution of the singular braid. Applying resolution rule (b) of Figure 4.10 to both singularities we obtain the Figure 4.32. It is also possible to explicitly solve equation (4.5.1) and compute the flow of branch points in the minimal chamber. The result is shown in the second part of Figure 4.32. We see that the resolved braid and the flow of branch points are topologically equivalent and just differ by change of projection plane. That is, the location of particles is represented in both pictures by cusps at which the same strands come approach.

Next, we turn to the maximal chamber. Here, we need further the following identity

$$E_+(\hat{x} + \hat{p}) = e^{i\pi\hat{p}^2} E_+(\hat{x})e^{-i\pi\hat{p}^2}, \quad (4.5.25)$$

which allows us to rewrite the partition function as

$$\begin{aligned} & \langle x | E_+(\hat{x})E_+(\hat{x} + \hat{p})E_+(\hat{p}) | y \rangle \\ &= \langle x | E_+(\hat{x})e^{i\pi\hat{p}^2} E_+(\hat{x})e^{-i\pi\hat{p}^2} e^{i\pi\hat{x}^2} e^{i\pi\hat{p}^2} e^{i\pi\hat{x}^2} E_+(\hat{x})e^{-i\pi\hat{x}^2} e^{-i\pi\hat{p}^2} e^{-i\pi\hat{x}^2} | y \rangle \end{aligned} \quad (4.5.26)$$

We depict the corresponding braid representation in Figure 4.31. One can immedi-




Figure 4.31: The singular braid of the A_2 flow in the maximal chamber.

ately extract from this singular braid the presence of three $U(1)$ gauge fields corresponding to the three white regions in the diagram. Moreover, we see that two of the

chiral fields are charged under two $U(1)$'s. Again, by a change of branch sheet (see section 4.4.1) and S-duality at the boundary, we can transform this picture to the one corresponding to the XYZ model discussed in section 4.4.1. We will not discuss this here but will rather analyse the connection to R-flow as branch point flow. This connection is established by looking at the particular resolution of the singular braid which corresponds to R-flow of branch points. This resolution is depicted below in Figure 4.33. The second part of Figure 4.33 shows the flow of branch points obtained by explicitly solving equation (4.5.1) in the maximal chamber. Again we see that the two figures are topologically identical.

A_3 flow

Next in complexity is the A_3 -flow. For clarity of presentation, we will solely concentrate on the flow in the minimal BPS chamber here. There are three cycles corresponding to the operators $\hat{\gamma}_1 = \hat{x}$, $\hat{\gamma}_2 = \hat{p}$, and $\hat{\gamma}_3 = \hat{x} + c$, which form a central extension of the $SL(2, \mathbb{Z})$ -algebra generated by \hat{x} and \hat{p} with commutators¹⁸

$$[\hat{x}, \hat{p}] = -\frac{i}{2\pi}, \quad [c, \hat{x}] = [c, \hat{p}] = 0. \quad (4.5.27)$$

The KS-operator corresponding to the minimal particle chamber is given by

$$\mathbb{K} = E_+(\hat{x} + c)E_+(\hat{p})E_+(\hat{x}). \quad (4.5.28)$$

¹⁸We have chosen here a different commutator between \hat{x} and \hat{p} compared to the A_2 case. This is merely a convention. We could also have worked with the former commutator.




Figure 4.32: R-flow of A_2 in the minimal chamber. In (a) the resolved braid braid is depicted. The previous singularities appear now as cusps in the braid diagram. These are marked with red arrows. Part (b) shows the flow of branch points obtained by explicitly solving equation (4.5.1).

A partition function can be formed from this operator by considering the wavefunction

$$\mathcal{Z}_{\mathbb{K}} = \langle x | E_+(\hat{x} + c) E_+(\hat{p}) E_+(\hat{x}) | y \rangle. \quad (4.5.29)$$

This partition function now represents a singular braid. In order to extract the braid, we have to rewrite it as a gluing of simple partition functions containing no gauge groups. This is done by using the identity

$$E_+(\hat{p}) = e^{-i\pi\hat{x}^2} e^{-i\pi\hat{p}^2} e^{-i\pi\hat{x}^2} E_+(\hat{x}) e^{i\pi\hat{x}^2} e^{i\pi\hat{p}^2} e^{i\pi\hat{x}^2}, \quad (4.5.30)$$




Figure 4.33: R-flow of A_2 in the maximal chamber. Part (a) shows the resolved braid. The locations of previous singularities are marked with red arrows. Part (b) depicts the flow of branch points as arising from a flow of central charges along straight vertical lines.

which allows us to rewrite $\mathcal{Z}_{\mathbb{K}}$ in the form

$$\mathcal{Z}_{\mathbb{K}} = \int dx' \langle x | E_+(x + c) e^{-i\pi\hat{x}^2} e^{-i\pi\hat{p}^2} e^{-i\pi\hat{x}^2} | x' \rangle \langle x' | E_+(x') e^{i\pi\hat{x}^2} e^{i\pi\hat{p}^2} e^{i\pi\hat{x}^2} E_+(y) | y \rangle. \quad (4.5.31)$$

This partition function can be represented by the singularized braid shown in Figure 4.34. We see again a $U(1)$ gauge group corresponding to the one compact white region. Furthermore, a chiral field is charged under this gauge group while the two other chiral fields are gauge neutral. Applying duality rules we can transform this




Figure 4.34: The singularized braid of the A3 flow in the minimal chamber.

picture to different ones with more or less gauge groups. Applying resolution rule (b) of Figure 4.10 to all singularities we obtain Figure 4.35. This resolved braid can




Figure 4.35: The desingularized braid of the A3 flow in the minimal chamber. The locations where the particles used to be are indicated by red arrows.

again be reproduced by letting the central charges of the A_3 curve R-flow as depicted in Figure 4.25. One can carry out the flow procedure by inverting the central charges as functions of the branch points locally along the flow. The resulting flow of branch points for the minimal chamber is depicted in Figure 4.36. Comparing Figure 4.35 with Figure 4.36 we find that the two are topologically identical in that the strands which come closest at the location of particles are the same in both pictures, i.e. first γ_3 contracts, then γ_2 and at last γ_1 . They merely differ by a change of the projection plane.

We find that this behavior generalizes. That is, associated to the KS-operator corresponding to the A_n theory in a particular chamber, there exists a resolution




Figure 4.36: Flow of branch points of the minimal chamber A_3 by imposing a flow of central charges along vertical straight lines.

which arises as R-flow of the branch points. The prescription for finding the resolution corresponding to R-flow is as follows. Start with the partition function

$$\mathcal{Z}_{A_n} = \langle x | \mathbb{K}(q) | x' \rangle. \quad (4.5.32)$$

Associate to this matrix element the particular braid-representation which contains all particles as singularities *within* the Seifert-surface, where by *within* we mean that the Seifert-surface goes *horizontally* through the singularity as depicted in Figure 4.27. Apply resolution rule of Figure 4.10 (b). Note that it is not possible to obtain other resolutions for the singular braids such as the one of figure 4.34 from R-flow. The reason is that a local flip of the corresponding central charge, as described in the case of A_1 , changes the KS-operator and will thus lead to a completely different picture.

4.6 Applications

In this section we study some further applications of the developed rules. As a first example we examine a more complicated geometry arising from the R-flow prescription. The particular geometry contains a closed non-planar polygon, i.e. a superpotential, which is only partly shaded and thus gives rise to a monopole operator. We will establish that this monopole operator appears in the superpotential. As a second example for the application of the methods developed in this paper we will look at $U(1)$ SQED with $N_f > 1$. This example does not arise from R-flow. However, we will find that the rules presented in section 4.4.1 are powerful enough to establish

mirror symmetry even for these more complicated models geometrically.

4.6.1 Superpotentials from Geometry

In this section we look at an example of a 3d gauge theory which arises from R-flow of an intermediate chamber of the A_4 theory. This example was already analyzed to some extent in [4]. The relevant KS-operator is given by

$$\mathbb{K} = E_+(\hat{x}_1)E_+(\hat{x}_2)E_+(\hat{p}_1 + \hat{x}_2)E_+(\hat{x}_2)E_+(\hat{p}_2), \quad (4.6.1)$$

where the phase space parameters satisfy the following commutation relations

$$[\hat{x}_1, \hat{p}_1] = \frac{i}{2\pi}, \quad [\hat{p}_1, \hat{x}_2] = \frac{i}{2\pi}, \quad [\hat{x}_2, \hat{p}_2] = \frac{i}{2\pi}. \quad (4.6.2)$$

The 3d partition function associated to the KS-operator is now

$$\mathcal{Z}_{\mathbb{K}} = \langle x | E_+(\hat{x}_1)E_+(\hat{x}_2)E_+(\hat{p}_1 + \hat{x}_2)E_+(\hat{x}_2)E_+(\hat{p}_2) | x' \rangle. \quad (4.6.3)$$

Its representation in terms of a singular braid is depicted in figure 4.37. We can




Figure 4.37: R-flow of A_4 in intermediate chamber.

clearly see 4 $U(1)$ gauge groups represented by the four white regions in the braid.

Applying the Fourier transform identity twice and the T -transform rule of section 4.4.1 we obtain the simpler braid depicted in Figure 4.38. This braid represents a




Figure 4.38: R-flow of A_4 in intermediate chamber, second representation. In (a) we see a dual representation of the A_4 braid after application of various dualities to the original R-flow braid. In (b) we see the same dual braid, now with the closed region representing the superpotential highlighted in red. Chiral Fields are indicated by X_i .

dual description of the same quantum field theory. In this description, there is a $U(1)$ gauge group under which two chiral multiplets, denoted by X_3 and X_2 , are charged oppositely. Furthermore, one can clearly see a compact polygonal region bounded by three chiral singularities. This corresponds to a superpotential in the effective 3d gauge theory to which all three chiral multiplets contribute. This theory contains a monopole operator which also participates in the superpotential term. One way to see this, is through the white region contained within the bounded polygonal region. One can check, using the formula (4.3.19) for the charge of the monopole operator discussed in section 4.3.3, that the monopole operator \mathcal{M} is invariant under the $U(1)$ gauge group. This immediately tells us that we can write down a superpotential of the form

$$\mathcal{W} = \mathcal{M} X_2 X_3 X_4, \quad (4.6.4)$$

which is gauge invariant. Furthermore, this superpotential breaks exactly one $U(1)$ flavor symmetry which is consistent as there are five chiral fields but only four non-compact white regions in the geometry.

4.6.2 $U(1)$ SQED with $N_f > 1$

Here, we will demonstrate that our rules for the singular tangles provide a convenient geometric way of encoding general mirror symmetries of 3d $\mathcal{N} = 2$ gauge theories. The example we will use to demonstrate this is the generalization of $U(1)$ SQED/ XYZ mirror symmetry. Start with a 3d $\mathcal{N} = 2$ gauge theory with $U(1)$ gauge group and $N_f > 1$ charged hypermultiplets. This theory has a RG fixed point with a mirror dual description as a $(U(1)^{N_f})/U(1)$ gauge theory with N_f charged hypermultiplets (consisting of chiral multiplets q_i and \tilde{q}_i) and N_f neutral chiral multiplets S_i together with a superpotential [65]

$$W = \sum_{i=1}^{N_f} S_i q_i \tilde{q}_i. \quad (4.6.5)$$

The charge assignments are as follows

	$U(1)_1$	$U(1)_2$	$U(1)_3$	\cdots	$U(1)_{N_f}$	
q_1	1	-1	0	\cdots	0	
\tilde{q}_1	-1	1	0	\cdots	0	
q_2	0	1	-1	\cdots	0	
\tilde{q}_2	0	-1	1	\cdots	0	(4.6.6)
\vdots		\ddots			0	
q_{N_f}	-1	0	0	\cdots	1	
\tilde{q}_{N_f}	1	0	0	\cdots	-1	
S_i	0	0	0	\cdots	0	

The aim will now be to translate both theories into geometric tangles and transform them into each other by using ordinary as well as singularized Reidemeister moves, thereby proofing they are mirror pairs.

$U(1)$ SQED with $N_f = 2$

We will start with the geometry corresponding to $U(1)$ SQED and specialize to the case $N_f = 2$. The relevant diagram describing this gauge theory is depicted in Figure 4.39.

The interior white region represents the $U(1)$ gauge group and each pair of singularities corresponds to a hypermultiplet whose constituents have opposite charges under the $U(1)$. Let us next apply the second Reidemeister move to this diagram. The result is depicted in Figure 4.40. Here we see that there are two extra $U(1)$'s and that




Figure 4.39: Diagram describing $U(1)$ SQED with $N_f = 2$.




Figure 4.40: Application of second Reidemeister move.

two singularities are charged under the first one whereas the second pair is charged under the second. We are now in a position to apply the generalized Reidemeister move known as the 3-2 move. This move can be applied twice, once to the upper

white triangle and once to the lower white triangle, resulting in Figure 4.41. This




Figure 4.41: After applying the 3-2-move twice.

diagram simply shows a $U(1)$ gauge theory with two chiral fields charges positively under it and two fields charges negatively. Moreover, we observe two superpotential terms each combining a neutral field with two oppositely charged fields. These data exactly match those of the mirror dual which confirms the duality.

$U(1)$ SQED with $N_f = 3$

As a second and last example we will consider the more complicated case of $U(1)$ SQED with $N_f = 3$. The relevant diagram is




Figure 4.42: $U(1)$ SQED with $N_f = 3$.

We can see 6 chiral multiplets charged under a $U(1)$ gauge group with the charges of the particles adding up to zero pairwise. The overcross and undercross singularities are arranged such that the net self-Chern-Simons level of the $U(1)$ is zero. We can add a T-transform to turn one type of singularity to another, as shown in Figure 4.43. Next, we do a second Reidemeister move to create a white region.

Performing the 3-2 move we end up with a superpotential and an extra $U(1)$, shown in Figure 4.45.

We now perform the Reidemeister move a second time to create a third white region with two charged fields.

Application of the 3-2 move for a second time leads to the second superpotential term.




Figure 4.43: Adding a T-transform.




Figure 4.44: Applying a Reidemeister move.

As should by now be obvious, we again perform the Reidermeister move with the result shown in Figure 4.48.

The last step is again a 3-2 move leading to the final result depicted in Figure 4.49.




Figure 4.45: After first 3-2 move.




Figure 4.46: After second Reidemeister move.

As one can clearly see the above picture is the diagram describing the mirror dual of our original theory. We have three superpotentials each containing one neutral field and we have three $U(1)$'s under each of which 2 chiral fields are charged. Note




Figure 4.47: After second 3-2 move.




Figure 4.48: After third Reidemeister move.

that the white region in the interior, under which no particle is charged, ensures that the sum of all $U(1)$'s adds up to zero as required by the charge assignments shown in (4.6.6). Thus we see that the diagram captures the theory in all details. The




Figure 4.49: The mirror dual.

constructions we have presented easily generalize to the case of arbitrary N_f .

Bibliography

- [1] 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](https://arxiv.org/abs/hep-th/9407087) [hep-th].
- [2] 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](https://arxiv.org/abs/hep-th/9408099) [hep-th].
- [3] S. Cecotti and C. Vafa, “Classification of complete $N=2$ supersymmetric theories in 4 dimensions,” [arXiv:1103.5832](https://arxiv.org/abs/1103.5832) [hep-th].
- [4] S. Cecotti, C. Cordova, and C. Vafa, “Braids, Walls, and Mirrors,” *ArXiv e-prints* (Oct., 2011) , [arXiv:1110.2115](https://arxiv.org/abs/1110.2115) [hep-th].
- [5] M. R. Douglas and G. W. Moore, “D-branes, Quivers, and ALE Instantons,” [arXiv:hep-th/9603167](https://arxiv.org/abs/hep-th/9603167).
- [6] D.-E. Diaconescu and J. Gomis, “Fractional branes and boundary states in orbifold theories,” *JHEP* **0010** (2000) 001, [arXiv:hep-th/9906242](https://arxiv.org/abs/hep-th/9906242) [hep-th].
- [7] M. R. Douglas, B. Fiol, and C. Römelsberger, “Stability and BPS branes,” *JHEP* **0509** (2005) 006, [arXiv:hep-th/0002037](https://arxiv.org/abs/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](https://arxiv.org/abs/hep-th/0003263) [hep-th].
- [9] B. Fiol and M. Marino, “BPS states and algebras from quivers,” *JHEP* **0007** (2000) 031, [arXiv:hep-th/0006189](https://arxiv.org/abs/hep-th/0006189) [hep-th].
- [10] B. Fiol, “The BPS spectrum of $N = 2$ $SU(N)$ SYM and parton branes,” [arXiv:hep-th/0012079](https://arxiv.org/abs/hep-th/0012079).
- [11] F. Denef, “Quantum quivers and Hall / hole halos,” *JHEP* **0210** (2002) 023, [arXiv:hep-th/0206072](https://arxiv.org/abs/hep-th/0206072) [hep-th].

- [12] 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.
- [13] D.-E. Diaconescu, M. R. Douglas, and J. Gomis, “Fractional branes and wrapped branes,” *JHEP* **9802** (1998) 013, [arXiv:hep-th/9712230](https://arxiv.org/abs/hep-th/9712230) [hep-th].
- [14] S. Cecotti and C. Vafa, “On classification of $N=2$ supersymmetric theories,” *Commun. Math. Phys.* **158** (1993) 569–644, [arXiv:hep-th/9211097](https://arxiv.org/abs/hep-th/9211097).
- [15] K. Hori, A. Iqbal, and C. Vafa, “D-branes and mirror symmetry,” [arXiv:hep-th/0005247](https://arxiv.org/abs/hep-th/0005247) [hep-th].
- [16] 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](https://arxiv.org/abs/hep-th/0110028) [hep-th].
- [17] 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](https://arxiv.org/abs/hep-th/0206152) [hep-th].
- [18] S. Cecotti, A. Neitzke, and C. Vafa, “R-Twisting and 4d/2d Correspondences,” [arXiv:1006.3435](https://arxiv.org/abs/1006.3435) [hep-th].
- [19] 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](https://arxiv.org/abs/hep-th/0003085) [hep-th].
- [20] 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](https://arxiv.org/abs/hep-th/0104259) [hep-th].
- [21] A. Hanany and K. D. Kennaway, “Dimer models and toric diagrams,” [arXiv:hep-th/0503149](https://arxiv.org/abs/hep-th/0503149) [hep-th].
- [22] 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](https://arxiv.org/abs/hep-th/0504110) [hep-th].
- [23] 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](https://arxiv.org/abs/hep-th/0511287) [hep-th].
- [24] 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](https://arxiv.org/abs/hep-th/9604034) [hep-th].

- [25] A. D. Shapere and C. Vafa, “BPS structure of Argyres-Douglas superconformal theories,” [arXiv:hep-th/9910182](https://arxiv.org/abs/hep-th/9910182) [hep-th].
- [26] D. Gaiotto, G. W. Moore, and A. Neitzke, “Wall-crossing, Hitchin Systems, and the WKB Approximation,” [arXiv:0907.3987](https://arxiv.org/abs/0907.3987) [hep-th].
- [27] 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](https://arxiv.org/abs/1109.4941) [hep-th].
- [28] 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](https://arxiv.org/abs/math/0608367).
- [29] D. Xie, “Network, Cluster coordinates and $N=2$ theory I,” *ArXiv e-prints* (Mar., 2012) , [arXiv:1203.4573](https://arxiv.org/abs/1203.4573) [hep-th].
- [30] A. King, “Moduli of representations of finitedimensional algebras,” *Quart. J. Mat. Oxford* **45** (1994) 515–530.
- [31] F. Denef and G. W. Moore, “Split states, entropy enigmas, holes and halos,” [arXiv:hep-th/0702146](https://arxiv.org/abs/hep-th/0702146) [hep-th].
- [32] M. Kontsevich and Y. Soibelman, “Stability structures, motivic Donaldson-Thomas invariants and cluster transformations,” *ArXiv e-prints* (Nov., 2008) , [arXiv:0811.2435](https://arxiv.org/abs/0811.2435) [math.AG].
- [33] P. Gabriel, “Unzerlegbare Darstellungen,” *Manuscripta Mathematica* **6** (1972) 71–103.
- [34] 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.
- [35] C. E. Beasley and M. Plesser, “Toric duality is Seiberg duality,” *JHEP* **0112** (2001) 001, [arXiv:hep-th/0109053](https://arxiv.org/abs/hep-th/0109053) [hep-th].
- [36] 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](https://arxiv.org/abs/hep-th/0109063) [hep-th].
- [37] D. Berenstein and M. R. Douglas, “Seiberg duality for quiver gauge theories,” [arXiv:hep-th/0207027](https://arxiv.org/abs/hep-th/0207027).
- [38] S. Cecotti and M. Del Zotto, “On Arnold’s 14 ‘exceptional’ $N=2$ superconformal gauge theories,” *JHEP* **1110** (2011) 099, [arXiv:1107.5747](https://arxiv.org/abs/1107.5747) [hep-th].

- [39] B. Keller, “Cluster algebras, quiver representations and triangulated categories,” [arXiv:0807.1960](https://arxiv.org/abs/0807.1960) [math.RT].
- [40] B. Keller, “On cluster theory and quantum dilogarithm identities,” *ArXiv e-prints* (Feb., 2011) , [arXiv:1102.4148](https://arxiv.org/abs/1102.4148) [math.RT].
- [41] D. Gaiotto, “N=2 dualities,” [arXiv:0904.2715](https://arxiv.org/abs/0904.2715) [hep-th].
- [42] E. Witten, “Solutions of four-dimensional field theories via M-theory,” *Nuclear Physics B* **500** (Feb., 1997) 3–42, [arXiv:hep-th/9703166](https://arxiv.org/abs/hep-th/9703166).
- [43] O. Chacaltana and J. Distler, “Tinkertoys for Gaiotto Duality,” *JHEP* **1011** (2010) 099, [arXiv:1008.5203](https://arxiv.org/abs/1008.5203) [hep-th].
- [44] O. Chacaltana and J. Distler, “Tinkertoys for the D_N series,” [arXiv:1106.5410](https://arxiv.org/abs/1106.5410) [hep-th].
- [45] S. Katz, P. Mayr, and C. Vafa, “Mirror symmetry and Exact Solution of 4D N=2 Gauge Theories I,” *ArXiv High Energy Physics - Theory e-prints* (June, 1997) , [arXiv:hep-th/9706110](https://arxiv.org/abs/hep-th/9706110).
- [46] S. H. Katz, A. Klemm, and C. Vafa, “Geometric engineering of quantum field theories,” *Nucl.Phys.* **B497** (1997) 173–195, [arXiv:hep-th/9609239](https://arxiv.org/abs/hep-th/9609239) [hep-th].
- [47] A. Felikson, M. Shapiro, and P. Tumarkin, “Skew-symmetric cluster algebras of finite mutation type,” [arXiv:0811.1703](https://arxiv.org/abs/0811.1703) [math.CO].
- [48] H. Derksen and T. Owen, “New Graphs of Finite Mutation Type,” [arXiv:0804.0787](https://arxiv.org/abs/0804.0787) [math.CO].
- [49] 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](https://arxiv.org/abs/hep-th/9706110) [hep-th].
- [50] D. Joyce, “On counting special Lagrangian homology 3-spheres,” *Contemp. Math.* **314** (2002) 125–151, [arXiv:hep-th/9907013](https://arxiv.org/abs/hep-th/9907013).
- [51] K. Strebel, *Quadratic Differentials*. Springer Verlag, 1984.
- [52] 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](https://arxiv.org/abs/0807.4723) [hep-th].
- [53] D. Gaiotto, G. W. Moore, and A. Neitzke, “Framed BPS States,” [arXiv:1006.0146](https://arxiv.org/abs/1006.0146) [hep-th].

- [54] D. Gaiotto, G. W. Moore, and A. Neitzke, “Wall-Crossing in Coupled 2d-4d Systems,” [arXiv:1103.2598](https://arxiv.org/abs/1103.2598) [hep-th].
- [55] M. Alim, S. Cecotti, C. Cordova, S. Espahbodi, A. Rastogi, and C. Vafa, “N=2 Quantum Field Theories and Their BPS Quivers,” *ArXiv e-prints* (Dec., 2011) , [arXiv:1112.3984](https://arxiv.org/abs/1112.3984) [hep-th].
- [56] D. Labardini-Fragoso, “Quivers with potentials associated to triangulated surfaces,” *ArXiv e-prints* (Mar., 2008) , [arXiv:0803.1328](https://arxiv.org/abs/0803.1328) [math.RT].
- [57] D. Labardini-Fragoso, “Quivers with potentials associated to triangulated surfaces, Part II: Arc representations,” [arXiv:0909.4100](https://arxiv.org/abs/0909.4100) [math.RT].
- [58] G. Cerulli Irelli and D. Labardini-Fragoso, “Quivers with potentials associated to triangulated surfaces, Part III: tagged triangulations and cluster monomials,” [arXiv:1108.1774](https://arxiv.org/abs/1108.1774) [math.RT].
- [59] A. Bakke Buan and I. Reiten, “Cluster algebras associated with extended Dynkin quivers,” *ArXiv Mathematics e-prints* (July, 2005) , [arXiv:math/0507113](https://arxiv.org/abs/math/0507113).
- [60] S. Ladkani, “Mutation classes of certain quivers with potentials as derived equivalence classes,” *ArXiv e-prints* (Feb., 2011) , [arXiv:1102.4108](https://arxiv.org/abs/1102.4108) [math.RT].
- [61] E. Witten, “SL(2,Z) Action On Three-Dimensional Conformal Field Theories With Abelian Symmetry,” *ArXiv High Energy Physics - Theory e-prints* (July, 2003) , [arXiv:hep-th/0307041](https://arxiv.org/abs/hep-th/0307041).
- [62] A. Kapustin and M. J. Strassler, “On mirror symmetry in three dimensional Abelian gauge theories,” *Journal of High Energy Physics* **4** (Apr., 1999) 21, [arXiv:hep-th/9902033](https://arxiv.org/abs/hep-th/9902033).
- [63] A. Redlich, “Gauge Noninvariance and Parity Nonconservation of Three-Dimensional Fermions,” *Phys. Rev. Lett.* **52** (1984) 18–21.
- [64] K. Intriligator and N. Seiberg, “Mirror symmetry in three dimensional gauge theories,” *Physics Letters B* **387** (Feb., 1996) 513–519, [arXiv:hep-th/9607207](https://arxiv.org/abs/hep-th/9607207).
- [65] O. Aharony, A. Hanany, K. Intriligator, N. Seiberg, and M. J. Strassler, “Aspects of $N = 2$ supersymmetric gauge theories in three dimensions,” *Nuclear Physics B* **499** (Feb., 1997) 67–99, [arXiv:hep-th/9703110](https://arxiv.org/abs/hep-th/9703110).
- [66] T. Dimofte, D. Gaiotto, and S. Gukov, “Gauge Theories Labelled by Three-Manifolds,” *ArXiv e-prints* (Aug., 2011) , [arXiv:1108.4389](https://arxiv.org/abs/1108.4389) [hep-th].

- [67] D. L. Jafferis, “The exact superconformal R-symmetry extremizes Z,” *Journal of High Energy Physics* **5** (May, 2012) 159, [arXiv:1012.3210](https://arxiv.org/abs/1012.3210) [hep-th].
- [68] A. Kapustin, B. Willett, and I. Yaakov, “Exact results for Wilson loops in superconformal Chern-Simons theories with matter,” *Journal of High Energy Physics* **3** (Mar., 2010) 89, [arXiv:0909.4559](https://arxiv.org/abs/0909.4559) [hep-th].
- [69] L. D. Faddeev, R. M. Kashaev, and A. Y. Volkov, “Strongly Coupled Quantum Discrete Liouville Theory.I: Algebraic Approach and Duality,” *Communications in Mathematical Physics* **219** (2001) 199–219, [arXiv:hep-th/0006156](https://arxiv.org/abs/hep-th/0006156).
- [70] D. Belov and G. W. Moore, “Classification of abelian spin Chern-Simons theories,” *ArXiv High Energy Physics - Theory e-prints* (May, 2005) , [arXiv:hep-th/0505235](https://arxiv.org/abs/hep-th/0505235).
- [71] C. Closset, T. T. Dumitrescu, G. Festuccia, Z. Komargodski, and N. Seiberg, “Contact terms, unitarity, and F -maximization in three-dimensional superconformal theories,” *Journal of High Energy Physics* **10** (Oct., 2012) 53, [arXiv:1205.4142](https://arxiv.org/abs/1205.4142) [hep-th].
- [72] E. Witten, “Five-brane effective action in M-theory,” *Journal of Geometry and Physics* **22** (May, 1997) 103–133, [arXiv:hep-th/9610234](https://arxiv.org/abs/hep-th/9610234).
- [73] J. Birman, *Braids, Links, and Mapping Class Groups*. Princeton University Press, 1974.
- [74] J. Lickorish, *An Introduction to Knot Theory*. Springer-Verlag, 1997.
- [75] L. F. Alday, D. Gaiotto, and Y. Tachikawa, “Liouville Correlation Functions from Four-Dimensional Gauge Theories,” *Letters in Mathematical Physics* **91** (Feb., 2010) 167–197, [arXiv:0906.3219](https://arxiv.org/abs/0906.3219) [hep-th].
- [76] S. Cecotti and C. Vafa, “BPS Wall Crossing and Topological Strings,” [arXiv:0910.2615](https://arxiv.org/abs/0910.2615) [hep-th].
- [77] C. Beem, T. Dimofte, and S. Pasquetti, “Holomorphic Blocks in Three Dimensions,” *ArXiv e-prints* (Nov., 2012) , [arXiv:1211.1986](https://arxiv.org/abs/1211.1986) [hep-th].
- [78] E. Witten, “Phase transitions in M-theory and F-theory,” *Nuclear Physics B* **471** (Feb., 1996) 195–216, [arXiv:hep-th/9603150](https://arxiv.org/abs/hep-th/9603150).
- [79] M. Aganagic and C. Vafa, “Mirror Symmetry, D-Branes and Counting Holomorphic Discs,” *ArXiv High Energy Physics - Theory e-prints* (Dec., 2000) , [arXiv:hep-th/0012041](https://arxiv.org/abs/hep-th/0012041).

Bibliography

- [80] N. Hama, K. Hosomichi, and S. Lee, “SUSY gauge theories on squashed three-spheres,” *Journal of High Energy Physics* **5** (May, 2011) 14, [arXiv:1102.4716 \[hep-th\]](https://arxiv.org/abs/1102.4716).
- [81] S. Pasquetti, “Factorisation of $\{N\} = 2$ theories on the squashed 3-sphere,” *Journal of High Energy Physics* **4** (Apr., 2012) 120, [arXiv:1111.6905 \[hep-th\]](https://arxiv.org/abs/1111.6905).
- [82] A. Hanany and E. Witten, “Type IIB superstrings, BPS monopoles, and three-dimensional gauge dynamics,” *Nuclear Physics B* **492** (May, 1997) 152–190, [arXiv:hep-th/9611230](https://arxiv.org/abs/hep-th/9611230).
- [83] R. Donagi and E. Witten, “Supersymmetric Yang-Mills theory and integrable systems,” *Nuclear Physics B* **460** (Feb., 1996) 299–334, [arXiv:hep-th/9510101](https://arxiv.org/abs/hep-th/9510101).
- [84] N. Drukker, D. Gaiotto, and J. Gomis, “The virtue of defects in 4D gauge theories and 2D CFTs,” *Journal of High Energy Physics* **6** (June, 2011) 25, [arXiv:1003.1112 \[hep-th\]](https://arxiv.org/abs/1003.1112).
- [85] D. Gaiotto and E. Witten, “Janus configurations, Chern-Simons couplings, and the θ -Angle in $\mathcal{N} = 4$ super Yang-Mills theory,” *Journal of High Energy Physics* **6** (June, 2010) 97, [arXiv:0804.2907 \[hep-th\]](https://arxiv.org/abs/0804.2907).
- [86] D. Gaiotto and E. Witten, “S-Duality of Boundary Conditions In $N=4$ Super Yang-Mills Theory,” *ArXiv e-prints* (July, 2008) , [arXiv:0807.3720 \[hep-th\]](https://arxiv.org/abs/0807.3720).
- [87] 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 \[hep-th\]](https://arxiv.org/abs/hep-th/9505062).