Dissertation

submitted to the

Combined Faculty of Mathematics, Engineering and Natural Sciences
of Heidelberg University, Germany

for the degree of

Doctor of Natural Sciences

Put forward by

Michael Bleher

born in: Schwäbisch Hall

Oral examination: 13.11.2023

Haydys-Witten Instantons and Symplectic Khovanov Homology

Referees: Prof. Dr. Johannes Walcher
Priv.-Doz. Dr. Jan Swoboda

‘*The road to wisdom? – Well, it’s plain*
and simple to express:
Err
and err
and err again
but less
and less
and less.’

– Piet Hein

Acknowledgments

I would like to take this opportunity to express my sincere gratitude to all the people that have accompanied me in one way or another throughout the time I was working on this dissertation.

I would like to express my appreciation to my supervisor Johannes Walcher for guidance and advice over the years. Also, many thanks to Jan Swoboda for helpful comments, crucial corrections, and his exceptional effort into a close reading of this thesis. I also appreciate the help I have received in securing financial support through the HGSFP, the STRUCTURES cluster of excellence, and various teaching assignments. Thanks in particular to Daniel Roggenkamp, Anna Wienhard, Hendrik Kasten, Thomas Dimpfl, and Anna Marciniak-Czochra.

I am very grateful for the company and help of my colleagues and friends at the Bel Étage: Ingmar, Felipe, Sebastian, Lukas, Fabio, Fabian, Steffen S., and Raphael. A particular thanks to Andreas Ott and Lukas Hahn for joining a refreshing side project.

Above all, I want to thank my friends and family for their constant support, understanding, and patience during this challenging process. I would like to extend special thanks: to Jascha, Nadia, Steffen A., Fabian, and Thomas J. for shared passions; to Thomas K., for numerous crucial distractions and many formative experiences throughout the years; to my brother and collaborator Johannes, for many enthusiastic discussions and for looking out for me; to my parents, for their constant and loving support; and finally, to Katha, for all of the above and more.

Abstract

An influential conjecture by Witten states that there is a Floer theory based on Haydys-Witten instantons that provides a gauge theoretic approach to Khovanov homology. This thesis explores a novel approach towards a potential proof of this claim. One of the key insights is the existence of a Hermitian Yang-Mills structure for a 'decoupled' version of the Haydys-Witten and Kapustin-Witten equations. It is shown that, in favourable circumstances, any Haydys-Witten solution is already a solution of the decoupled equations. This utilizes a dichotomy that is proved to be satisfied by θ-Kapustin-Witten solutions on any ALE or ALF space, generalizing a corresponding result on $\mathbb{R}^4$. The Hermitian Yang-Mills structure gives rise to a Kobayashi-Hitchin-like correspondence. It is proposed that solutions are classified by intersections of Lagrangian submanifolds in the moduli space of solutions of the extended Bogomolny equations. In that interpretation, Haydys-Witten instantons are in correspondence with pseudo-holomorphic discs, leading to a conjectural equivalence with a Lagrangian intersection Floer homology. A physically motivated argument suggests that the latter is fully determined in a finite-dimensional model space, given by a Grothendieck-Springer resolution of the nilpotent cone inside the underlying Lie algebra. This provides a relation to symplectic Khovanov homology, which is known to be isomorphic to a grading-reduced version of Khovanov homology.

Zusammenfassung

Eine einflussreiche Vermutung von Witten besagt, dass es eine Floer-Theorie auf Grundlage von Haydys-Witten Instantonen gibt, die einen eichtheoretischen Zugang zur Khovanov-Homologie bietet. In dieser Arbeit wird ein neuer Ansatz für einen möglichen Beweis dieser Behauptung untersucht. Eine der wichtigsten Erkenntnisse dieser Arbeit ist die Existenz einer hermitschen Yang-Mills-Struktur für eine „entkoppelte" Version der Haydys-Witten- und Kapustin-Witten-Gleichungen. Es wird gezeigt, dass unter günstigen Umständen jede Haydys-Witten-Lösung bereits eine Lösung der entkoppelten Gleichungen ist. Dies macht sich eine Dichotomie zunutze, die von θ-Kapustin-Witten-Lösungen auf jedem ALE- oder ALF-Raum erfüllt wird. Dieses Ergebnis verallgemeinert ein bereits zuvor bekanntes Resultat im Fall des $\mathbb{R}^4$. Die hermitsche Yang-Mills-Struktur führt zu einer Kobayashi-Hitchin-ähnlichen Korrespondenz. Eine Behauptung wird aufgestellt, derzufolge Lösungen durch Schnittpunkte von Lagrangeschen Untermannigfaltigkeiten im Moduli-Raum von Lösungen der erweiterten Bogomolny Gleichungen klassifiziert werden. In dieser Interpretation werden Haydys-Witten Instantonen durch pseudoholomorphe Kurven beschrieben, was zu einer mutmaßlichen Äquivalenz mit einer Floer-Homologie Lagrangescher Schnitte führt. Ein physikalisch motiviertes Argument legt nahe, dass letztere vollständig durch einen endlichdimensionalen Modellraum bestimmt ist, der aus einer Grothendieck-Springer-Auflösung des nilpotenten Kegels innerhalb der zugrunde liegenden Lie-Algebra hervorgeht. Dies stellt eine Beziehung zur symplektischen Khovanov-Homologie her, die bekanntermaßen isomorph zu einer Version der Khovanov-Homologie mit reduzierter Graduierung ist.

Contents

1 Introduction

Some of the most exciting discoveries in contemporary mathematics have been influenced by developments in physics. One area that has been particularly informed by physical intuition is the intersection of geometric analysis and geometric topology. The underlying phenomenon is an intricate connection between analytic and topological properties of manifolds: It seems to be a general rule that instead of studying properties of a manifold directly, it is more insightful to investigate moduli spaces of solutions of partial differential equations on the manifold. Insights from classical and quantum gauge theories have proven particularly fruitful in that regard.

The mathematical study of classical and quantum gauge theories was initially driven by the necessity to develop a framework that captures observations in particle physics, prompted by the detection of a surprisingly versatile zoo of particles starting in the 1950s. There was a series of notable breakthroughs in the theory of classical gauge theories during the 1980s. Many of these results, perhaps most notably Uhlenbeck's compactness theorem for Yang-Mills connections [Uhl82b; Uhl82a], resulted in profound insights into the properties of low-dimensional smooth manifolds. An influential example is Donaldson-Floer theory, an infinite-dimensional analogue of Morse theory that constructs a homological smooth invariant for a given three-manifold X^3. The underlying chain complex is spanned by flat connections on X^3 and the differential is determined by the number of anti-self dual connections on the cylindrical four-manifold $\mathbb{R}_s \times X^3$ that interpolate between flat connections at the cylindrical ends $s \to \pm\infty$ [Flo88a; Flo89].

At roughly the same time, physicists realized that supersymmetric quantum theories are strongly restricted by their large amount of symmetry; so much so that certain subsectors of the theory may only 'see' a small part of the structure that is present on the underlying manifold. For example, although the laws of physics are ultimately governed by local differential equations, observables in 'topological' quantum field theories[1] only depend on the global topology of the underlying manifold [Wit82; Wit88a]. In this way, the connection between geometric analysis and topological invariants might be viewed as an incarnation of a fundamental property of physics – or vice versa, depending on preference.

Perhaps one of the most striking examples of such a topological theory was discovered by Witten in the late 1980s, who realized that the partition function and observables of three-dimensional Chern-Simons theory are topological invariants [Wit89]. The only gauge-invariant observables in pure Chern-Simons theory are Wilson line operators that are supported along a knot K, the image of an embedding $S^1 \hookrightarrow X^3$. Two knots are considered topologically

[1]The observables of such quantum theories typically depend on the smooth structure of the manifold.

equivalent when there is an ambient isotopy that takes one knot into the other. Witten showed that, if one interprets the partition function in $G = SU(N)$ Chern-Simons theory at level k with Wilson line inclusions as a Laurent polynomial in the parameter $q = \exp\big(2\pi i/(N+k)\big)$, then it satisfies the relation

$$q^{N/2}\, Z_{\mathrm{CS}}\big(\quad\big) - q^{-N/2}\, Z_{\mathrm{CS}}\big(\quad\big) = \big(q^{1/2} - q^{-1/2}\big)\; Z_{\mathrm{CS}}\big(\quad\big)\ .$$

This type of equation is known as skein relation. It provides a connection between the expectation value of Wilson lines that are located on three topologically distinct knots that only differ in the vicinity of a single crossing as depicted. Witten also showed that the normalization of the partition function is uniquely determined by the expectation value of an unknotted Wilson line inside of $X^3 = S^3$ and is given by

$$Z_{\mathrm{CS}}\big(\bigcirc\big) = \frac{q^{N/2} - q^{-N/2}}{q^{1/2} - q^{-1/2}}\ .$$

Polynomial knot invariants are uniquely determined by their skein relation and normalization for the unknot. For $G = SU(2)$, the Chern-Simons partition function turns out to be equivalent to the Jones Polynomial [Jon85]

$$Z_{\mathrm{CS}}(S^3, K; q) = J(K; q)\ .$$

More generally, for $G = SU(N)$, the Chern-Simons partition function reproduces the quantum $\mathfrak{sl}(N)$ polynomial invariant of K, which in turn are one-variable specializations of the three-variable HOMFLY-PT polynomial $\mathcal{P}_K(x = q^{N/2}, y = -q^{N/2}, z = q^{1/2} - q^{-1/2})$ [Fre+85; PT88].

Intriguingly, from the gauge theory point of view the invariance of the partition function $Z_{\mathrm{CS}}(S^3, K; q)$ under ambient isotopies is inherent to the topological nature of Chern-Simons theory. Moreover, since the theory is exactly solvable, moving from S^3 to invariants of knots in a general three-manifold X^3 is relatively easy in gauge theory, but calculating the associated partition functions becomes challenging [LR89; LLR91]. In contrast, the definition of the Jones polynomial relies on a projection of the knot to two-dimensions and invariance under Reidemeister moves must be shown a posteriori.

In the late 1990s Khovanov observed that there is a categorification of the Jones polynomial, now known as Khovanov homology [Kho99]. Khovanov homology associates to a knot K in S^3 a finite-dimensional bigraded vector space $Kh^{\bullet,\bullet}(K)$. The Jones polynomial is the graded Euler characteristic of this invariant:

$$J(K; q) = \sum_{i,j\in\mathbb{Z}} (-1)^i q^j \dim Kh^{i,j}(K).$$

The graded vector space $Kh^{\bullet,\bullet}(K)$ is the cohomology of a bigraded chain complex $C^{\bullet,\bullet}(K)$ with respect to a $(1,0)$-graded differential Q (a linear map that satisfies $Q^2 = 0$). While Khovanov homology is invariant under ambient isotopies and only depends on the knot K, the larger chain complex depends on various choices in its construction.

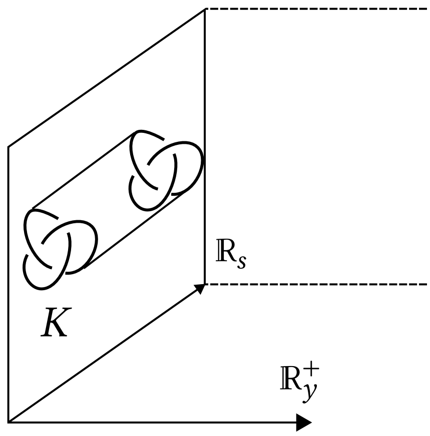

Figure 1.1 The five-manifold $\mathbb{R}_s \times S^3 \times \mathbb{R}_y^+$ together with an embedded surface $\Sigma_K = \mathbb{R}_s \times K$ that plays a central role in Witten's gauge theoretic approach to Khovanov homology.

Since the inception of Khovanov homology, many more knot homologies have been constructed, including but not limited to: knot Floer homology, independently discovered by Ozsváth-Szabó and Rasmussen [OS03; Ras03]; symplectic Khovanov homology, introduced for $\mathfrak{sl}(2)$ by Seidel-Smith and generalized to $\mathfrak{sl}(N)$ by Manolescu [SS04; Man07]; and Khovanov-Rozansky homology, a generalization of Khovanov homology that categorifies quantum $\mathfrak{sl}(N)$ invariants and the HOMFLY-PT polynomial [KR08a; KR08b].

It was soon realized that Khovanov's homological invariant must also have a physical interpretation in the context of Chern-Simons theory. One such interpretation was described by Gukov, Schwarz and Vafa, who utilized an equivalence between Chern-Simons theory and topological string theory to relate Khovanov-Rozansky homology to the Hilbert space of BPS states of open topological strings [GSV05]. This work sparked a significant body of research into the structural relationships between various different homological knot invariants and topological string theory. Amongst others, this includes a categorification of the Kauffman polynomial by generalizing the constructions to Lie algebras $\mathfrak{so}(N)$ and $\mathfrak{sp}(N)$ and the conjectured existence of a triply- or even quadruply-graded homological invariant at large N that categorifies the HOMFLY-PT polynomial [DGR05; GW05; GS11; AS12; GGS13].

Witten introduced an alternative approach that constructs Khovanov homology purely from gauge theory [Wit11a]. This provides a conceptually cleaner explanation of categorification. Starting from a Wilson line along K in Chern-Simons theory and using a series of dualities between equivalent realizations of the same physical system, Witten arrives at the Hilbert space of BPS states $\mathcal{H}_{\mathrm{BPS}}(K)$ in five-dimensional $\mathcal{N} = 2$ supersymmetric Yang-Mills (SYM) theory on the manifold with boundary $\mathbb{R}_s \times S^3 \times \mathbb{R}_y^+$. The space of BPS states $\mathcal{H}_{\mathrm{BPS}}(K)$ is the cohomology of a nilpotent supersymmetry charge Q that acts on the Hilbert space of states $\mathcal{H}(K)$ of the theory. The five-dimensional SYM theory exhibits a $U(1) \times U(1)$ symmetry, which translates into a bigrading of the Hilbert space with respect to the symmetry generators F and P, satisfying $[F, Q] = Q$ and $[P, Q] = 0$. While the larger Hilbert space $\mathcal{H}^{\bullet,\bullet}(K)$ depends on various parameters of the theory, the subsector of BPS states is protected by supersymmetry and depends only on the topology of the knot K. Crucially, Witten concludes that the following equality must hold:

$$Z_{\mathrm{CS}}(S^3, K; q) = \mathrm{Tr}_{\mathcal{H}_{\mathrm{BPS}}}(-1)^F q^P \,.$$

This equation is the physical analogue of the relation between the Jones polynomial and Khovanov homology. Witten proposed that the Hilbert space of BPS states of the five-dimensional theory $\mathcal{H}^{\bullet,\bullet}_{\mathrm{BPS}}(K)$ is a homological knot invariant and that it coincides with Khovanov homology.

The relation between Witten's approach and the earlier explanations in terms of topological strings has been addressed in [Das+16] under the umbrella of M-theory. There is also a recent endeavour by Aganagic to further unite these approaches from a slightly different perspective [Aga20; Aga21].

In Witten's approach, the vector space $\mathcal{H}^{\bullet,\bullet}_{\mathrm{BPS}}(K)$ is calculated as follows. Start with the vector space $CF^{\bullet,\bullet}(S^3, K)$ spanned by all supersymmetric *classical* ground states of the theory. The latter are given by $\mathbb{R}_s$-invariant solutions of the equations of motion, a set of partial differential equations known as Kapustin-Witten equations. BPS states (supersymmetric *quantum* ground states) are obtained from the classical ones by taking into account quantum-mechanical instanton corrections. Instanton corrections in this context arise from solutions of the Haydys-Witten equations on $\mathbb{R}_s \times S^3 \times \mathbb{R}^+_y$ that interpolate between two Kapustin-Witten solutions at $s \to \pm\infty$. Counting Haydys-Witten instantons defines a differential Q on $CF^{\bullet,\bullet}$, such that the space of BPS states is given by an infinite-dimensional version of Morse theory

$$\mathcal{H}^{\bullet,\bullet}_{\mathrm{BPS}}(K) = HF^{\bullet,\bullet}(K) := H^\bullet(CF^{\bullet,\bullet}(K), Q)\,.$$

Since this is an analogue of Donaldson-Floer theory, we refer to this as Haydys-Witten Floer theory.

Expanding on this recipe, Gaiotto and Witten laid out a program to calculate Khovanov homology directly from the partial differential equations in gauge theory, without relying on dualities from string theory and quantum field theory [GW12]. One of the key ideas is to construct solutions of the four-dimensional Kapustin-Witten equations on $\mathbb{R}_t \times \Sigma \times \mathbb{R}^+_y$, where Σ is a Riemann surface, by adiabatically braiding solutions of the three-dimensional specialization known as extended Bogomolny equations (EBE) along a braid representation of K extended in the direction of $\mathbb{R}_t$, see Figure 1.2.

The EBE exhibit a Hermitian Yang-Mills structure that makes part of the equations invariant under *complexified* gauge transformations. This suggests that there is a classification of solutions in terms of holomorphic data along the lines of the Donaldson-Uhlenbeck-Yau theorem [Don85; UY86]. Gaiotto and Witten conjectured that the relevant solutions of the extended Bogomolny equations on $\Sigma \times \mathbb{R}^+_y$ are in one-to-one correspondence with Higgs bundles on Σ together with the additional structure of a distinguished line subbundle that captures the position of knot punctures in Σ. Moreover, they suggest that one can glue EBE-solutions along a

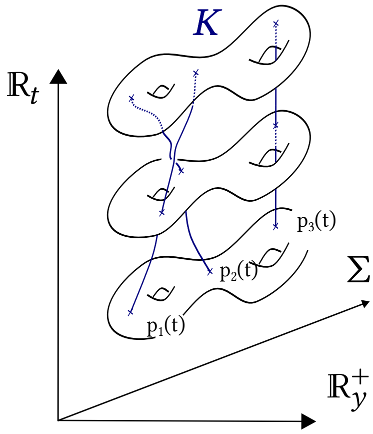

Figure 1.2 In the three-dimensional context of the extended Bogomolny equations (EBE) on $\Sigma \times \mathbb{R}_y^+$, the knot K reduces to a collection of points on Σ. The adiabatic approach treats these points as slowly varying functions and assumes that, at any given time, the fields are in a three-dimensional ground state that is well-approximated by a solution of the EBE.

stretched, slowly varying knot to obtain approximate solutions of the Kapustin-Witten equations, which in turn can be used as proxies for the classical ground states of the gauge theory. Gaiotto and Witten showed that this procedure works in principle, by reconstructing the skein relations of the Jones polynomial.

Over the past few years the Gaiotto-Witten program has been subject to close scrutiny from a mathematical point of view. By now it has been proven that the solutions of the extended Bogomolny equations on $\Sigma \times \mathbb{R}_y^+$ are indeed classified by Higgs bundles with additional structure [HM19c; HM20], or by Opers in case of a twisted version of the EBE [HM19b]. More recently [Dim22b] and [Sun23] have made further progress in proving variations of these claims that play important roles at various points in Gaiotto and Witten's work. However, it remains an open problem to classify Kapustin-Witten and Haydys-Witten solutions along these lines, let alone to prove that Haydys-Witten Floer homology coincides with Khovanov homology.

The present work reports new advances in that direction. Chapter 2 introduces Haydys-Witten Floer theory and, in doing so, also provides some previously unreported results. The main contributions of this work are presented in Chapters 3 - 5. The key insights of Chapter 3 and Chapter 4 take the form of vanishing results for the Kapustin-Witten and Haydys-Witten equations and might be of independent interest. In Chapter 5 these insights are combined with Gaiotto and Witten's approach of adiabatically braiding solutions of the extended Bogomolny equations. This leads to a conjectural relation between Haydys-Witten Floer theory and a version of symplectic Khovanov homology. The arguments of Chapter 5 are ultimately based on physical intuition, and while a proof along these lines currently seems out of reach, the results presented in this work lay out a new strategy to confirm that Witten's gauge theoretic categorification of the Jones polynomial produces Khovanov homology.

Overview

Chapter 2 begins with a brief account of the underlying physics, after which the main focus is put on the ingredients of Haydys-Witten Floer Theory. The Haydys-Witten equations for a pair of (bosonic) fields (A, B) are defined for any Riemannian five-manifold (M^5, g) that is equipped with a preferred non-vanishing vector field v. The interplay between cylindrical ends of M^5 and the vector field v is important for understanding many of the details that arise in the description of Floer theory. In particular, boundary conditions at both compact and non-compact ends of M^5 play an important role in this regard. For this reason, a detailed introduction to various dimensional reductions of the Haydys-Witten equations is provided. This includes, on the one hand, the previously unreported result that dimensional reduction produces the full one-parameter family of θ-Kapustin-Witten equations, where the parameter $\theta \in [0, \pi]$ receives a geometric interpretation as the angle between the vector field v and the direction of dimensional reduction. On the other hand, it is shown that the incidence angle of v at compact ends in general determines the use of tilted Nahm pole boundary conditions, necessitating a discussion of their elliptic properties that, so far, has not been part of the literature. The chapter closes with the definition of Haydys-Witten Floer homology groups and the relation with Witten's approach to Khovanov homology.

The main results of Chapter 3 are given by Theorem A and Theorem B, which establish a dichotomy for θ-Kapustin-Witten solutions (A, ϕ), between the growth of the average of $\|\phi\|^2$ on geodesic spheres and the vanishing of $[\phi \wedge \phi]$ on ALX spaces. This generalizes a similar result by Taubes that holds on Euclidean space $\mathbb{R}^n$. As an immediate consequence, Corollary C confirms a conjecture by Nagy and Oliveira, stating that on asymptotically locally Euclidean (ALE) and asymptotically locally flat (ALF) manifolds, and under the assumption of finite energy, any $G = SU(2)$ Kapustin-Witten solution satisfies $F_A = 0$, $\nabla^A \phi = 0$, and $[\phi \wedge \phi] = 0$. This vanishing theorem has important implications for the properties of Haydys-Witten instantons that approach Kapustin-Witten solutions at non-compact cylindrical ends.

In Chapter 4, the Haydys-Witten equations for the pair (A, B) are investigated on five-manifolds that additionally admit an almost Hermitian structure on the vector bundle $\ker g(v, \cdot)$. In this situation, there exists a simplified version of the equations in which certain components of B decouple from the curvature two-form F_A. These decoupled equations exhibit a Hermitian Yang-Mills structure that is closely related to the analogous structure of the EBE. In particular, the decoupled equations exhibit an enlarged invariance under complex gauge transformations, resulting in a considerable reduction in complexity of Haydys-Witten solutions. It is shown that the relation between the full Haydys-Witten equations and their decoupled version is captured by a Weitzenböck formula that states that the difference is governed by the asymptotic behaviour of solutions near boundaries and non-compact ends. This implies Theorem D, which reports necessary conditions for the asymptotic geometry of the manifold and the boundary conditions for (A, B), under which the Haydys-Witten equations reduce to the decoupled equations. Regarding the analysis near boundaries, a detailed analysis of the polyhomogeneous expansion of Haydys-Witten solutions with twisted Nahm pole boundary conditions is presented, generalizing work of Siqi He in the untwisted case. The corresponding analysis at non-compact ends relies on the results of Chapter 3, specifically the vanishing theorem Corollary C.

In Chapter 5, the decoupled equations of Chapter 4 are investigated in the context of Gaiotto-Witten's adiabatic approach on $M^5 = C \times \Sigma \times \mathbb{R}^+_y$, where C is either $\mathbb{R}_s \times S^1_t$ or $\mathbb{R}_s \times \mathbb{R}_t$. The Hermitian Yang-Mills structure of the decoupled equations provides the existence of homotopies between field configurations that are associated with knot singularities along isotopic knots. Using this insight, and inspired by Gaiotto-Witten's approach of adiabatically braiding EBE-solutions along a knot $K = S^1_t \times \{p_a(t)\}$, it is proposed that there is an equivalence between solutions of the decoupled version of the Kapustin-Witten equations and non-vertical paths in the moduli space of EBE-solutions. Physical intuition suggests that there is a finite dimensional model for the moduli space of solutions, given by a Grothendieck-Springer resolution of the Lie algebra $\mathfrak{sl}(N, \mathbb{C})$. These arguments lead to Conjecture E, stating that the number of solutions to the decoupled Kapustin-Witten equations on $S^1_t \times \mathbb{C} \times \mathbb{R}^+_y$ is bounded from below by the number of intersection points of the Grothendieck-Springer fiber and its parallel transport along K. A stronger claim is made in Conjecture F in the context of compact knots in $\mathbb{R}_s \times \mathbb{R}_t \times \mathbb{C} \times \mathbb{R}^+_y$, namely that there is an isomorphism between Haydys-Witten Floer homology and symplectic Khovanov-Rozansky homology. For $G = SU(2)$, symplectic Khovanov homology is known to be isomorphic to a grading-reduced version of Khovanov homology, such that the contents of this thesis establish a novel proof strategy for Witten's approach to Khovanov homology.

Chapter 6 provides a brief summary of the main results and discusses possible future directions.

2 Haydys-Witten Floer Theory

This chapter provides a review of Witten's gauge theoretic approach to Khovanov homology. We adopt the point of view of an instanton Floer theory associated to the Haydys-Witten equations on Riemannian five-manifolds (M^5, g) equipped with a non-vanishing unit vector field v.

This Haydys-Witten instanton Floer theory is a five-dimensional analogue of the well-known four-dimensional Donaldson-Floer theory. The latter calculates topological invariants of a three-manifold X^3 by counting the number of flat connections on it. The naive count is corrected by identifying flat connections for which there is a suitable flow line that interpolates between them. As it turns out, the relevant notion of flow lines is that of (anti-)self-dual connections on the four-dimensional cylinder $W^4 = \mathbb{R}_s \times X^3$, also known as Yang-Mills instantons.

From the point of view of physics, Yang-Mills instantons have an interpretation as quantum mechanical corrections to the classical ground states of Chern-Simons theory. In that interpretation Donaldson-Floer theory arises from four-dimensional $\mathcal{N} = 2$ super Yang-Mills (SYM) theory on $\mathbb{R}_s \times X^3$ that is coupled to three-dimensional Chern-Simons theory on X^3 at infinity. The coupling of Chern-Simons and SYM theory is by virtue of the fact that the Chern-Simons functional coincides with the boundary term[1] of the SYM action. From that point of view, Floer showed that the space of quantum mechanical ground states provides topological invariants of X^3, now known as Floer groups [Flo88a; Flo89].

Analogously, Haydys-Witten Floer theory can be viewed as the topological twist of a five-dimensional $\mathcal{N} = 2$ SYM theory on $\mathbb{R}_s \times W^4$ that is coupled to a four-dimensional $\mathcal{N} = 4$ SYM theory on W^4. Classical ground states in this theory are given by θ-Kapustin-Witten solutions on W^4, while instanton corrections correspond to Haydys-Witten solutions on $\mathbb{R}_s \times W^4$ that interpolate between two solutions. In that way, for any four-manifold W^4, Haydys-Witten Floer theory constructs from the Morse-Smale-Witten complex $CF_\theta(W^4)$ of θ-Kapustin-Witten solutions the homology groups

$$HF_\theta(W^4) = H(CF_\theta(W^4), d_v) \ .$$

Here, assuming there is a non-vanishing vector field w on W^4, the differential d_v counts solutions of the Haydys-Witten equations on the cylinder $\mathbb{R}_s \times W^4$ with respect to the fixed vector field

[1] In fact, the Chern-Simons functional is generally only defined with respect to a corresponding bounding four-manifold.

$v = \cos\theta\partial_s + \sin\theta w$:

$$d_v[x] = \sum_{\mu(x,y)=1} \#\mathcal{M}(x,y)[y] \ .$$

Since there is a physical incarnation as Hilbert space $\mathcal{H}_{\mathrm{BPS}}(W^4)$ of a topologically twisted theory, it is expected that these homology groups are topological invariants of W^4. In particular, if the four-manifold is the manifold with boundary $W^4 = S^3 \times \mathbb{R}^+$ together with a knot in its boundary $S^1 \hookrightarrow \partial W^4 = S^3 \times \{0\}$, an influential conjecture by Witten states that this topological invariant coincides with Khovanov homology.

While most results presented in this chapter have previously appeared in the literature, the definition of Haydys-Witten homology groups $HF_\theta(W^4)$ has so far not been spelled out explicitly. The definition of a one-parameter family of such homology groups utilizes the observation that dimensional reduction of the Haydys-Witten equations generally results in θ-Kapustin-Witten equations, where the value of θ is determined by the angle between v and the direction of dimensional reduction. In a closely related manner, the incidence angle of v at a boundary imposes the use of *twisted* (or tilted) Nahm pole boundary conditions and we also include a discussion of their analytic properties, with special focus on elliptic regularity of Haydys-Witten and θ-Kapustin-Witten solutions.

This chapter is structured as follows. We first recall some background from physics, starting with a short overview of $4d\ \mathcal{N} = 4$ SYM theory with boundaries and line operators in Section 2.1. Then we review topological twists in Section 2.2 and specify the twists relevant for Haydys-Witten theory. Section 2.3 explains that the partition function of the four-dimensional theory admits a categorification in terms of a Hilbert space of BPS states in a 5d $\mathcal{N} = 2$ SYM theory. These considerations explain the origin of the Haydys-Witten and Kapustin-Witten equations, and why they are expected to give rise to an interesting Floer theory. In Section 2.4 we give short individual introductions to a slightly confusing number of differential equations in various dimensions. These equations are relevant at different points throughout this thesis. It is shown in Section 2.5 that they can all be viewed as dimensional reductions of the Haydys-Witten equations. The Kapustin-Witten equations exhibit surprisingly restrictive vanishing results on closed manifolds and it seems to be an important aspect of the theory to consider field configurations with singular boundary conditions. This motivates the introduction of Nahm pole boundary conditions with knot singularities, which we review in Section 2.6. The chapter concludes with a definition of Haydys-Witten Floer theory and an explanation how this captures the overarching philosophy behind Witten's proposal in Section 2.7.

2.1 Supersymmetric Yang-Mills Theory with Boundaries

To set the stage and in view of later sections, we start by introducing some notation for general manifolds. Let G be a Lie group, and let $E \to W^4$ be a principal G-bundle with gauge connection $A \in \mathcal{A}(E)$ over an oriented Riemannian four-manifold (W^4, g). We denote the Lie algebra of G by $\mathfrak{g}$ and the adjoint bundle $E \times_{\mathrm{Ad}} \mathfrak{g}$ by $\mathrm{ad}\, E$.

There are spinor bundles $S^{\pm}$ associated to the $SO(4)$-frame bundle over W^4, respectively of positive and negative chirality. The underlying Weyl spinor representations are complex conjugates $\overline{S^+} = S^-$. Write V for the complexified vector representation of $SO(4)$. A standard construction[2] identifies $S := S^+ \oplus S^-$ with $\wedge^\bullet L = \wedge^{\text{even}} L \oplus \wedge^{\text{odd}} L$ for some choice of maximal totally isotropic subspace L of (V, g). The Clifford algebra $Cl(V, g)$, viewed as a 'deformation quantization' of the exterior algebra of V, acts naturally on $\wedge^\bullet L$. In particular, since $V \subset Cl(V, g)$, this induces a complex linear map $cl : V \otimes S^{\pm} \to S^{\mp}$ called Clifford multiplication. Moreover, there is a $\mathbb{C}$-valued inner product (s, t) on S, defined by restriction to the top-degree component of the element $s \wedge t \in \wedge^\bullet L$. The inner product pairs each of $S^{\pm}$ with itself. In combination with Clifford multiplication it induces a bilinear map $\Gamma : S^+ \otimes S^- \to V$, defined by duality: $g(\Gamma(s, t), v) := (s, cl(v, t))$.

The gauge connection induces a covariant exterior derivative d_A on $\Omega^\bullet(W^4, \operatorname{ad} E)$ and a Dirac operator for spinors, given by composing the covariant derivative with Clifford multiplication

$$D^A : \Gamma(W^4, S^{\pm} \otimes \operatorname{ad} E) \xrightarrow{\nabla^A} \Gamma(W^4, T^*W^4 \otimes S^{\pm} \otimes \operatorname{ad} E) \xrightarrow{cl \circ \sharp} \Gamma(W^4, S^{\mp} \otimes \operatorname{ad} E) \,.$$

Let $\operatorname{Tr}(\cdot)$ denote the trace on $\mathfrak{g}$. For $\alpha, \beta \in \Omega^k(W^4, \operatorname{ad} E)$ introduce the density-valued inner product $\langle \alpha, \beta \rangle := \operatorname{Tr} \alpha \wedge \star \beta$ with associated norm $\|\alpha\|^2 = \langle \alpha, \alpha \rangle$. For $s, t \in \Gamma(W^4, S^{\pm} \otimes \operatorname{ad} E)$ we similarly write $\langle s, t \rangle := \operatorname{Tr}\,(s, t)\, \mu_{W^4}$ where $\mu_{W^4} = \sqrt{g}\, dx^1 \dots dx^n$ denotes the volume form.

With that notation in place, we now specify the Lagrangian of $d = 4$ $\mathcal{N} = 4$ super Yang-Mills theory. The field content consists of a gauge connection $A \in \mathcal{A}(E)$, six scalar fields $\phi_i \in \Omega^0(W^4, \operatorname{ad}_E)$, $i = 1, \dots, 6$, and in total four Weyl spinors $\psi_a \in \Gamma(W^4, S^+ \otimes \operatorname{ad} E)$, $a = 1, \dots, 4$. The action is the sum of a kinetic and a topological term

$$S = \frac{1}{g_{\text{YM}^2}} \int_{W^4} \mathcal{L}_{\text{kin}} + \frac{\theta_{\text{YM}}}{32\pi^2} \int_{W^4} \operatorname{Tr} F_A \wedge F_A \,.$$

For the rest of this section, we assume that W^4 is a region in Euclidean space $\mathbb{R}^4$. In that case the kinetic Lagrangian is given by the sum of the following two parts, where the first is purely bosonic and the second contains all contributions that involve fermions

$$\begin{aligned}
\mathcal{L}^{A,\phi}_{\text{kin}} &= \tfrac{1}{2} \|F_A\|^2 + \sum_{i=1}^{6} \|d_A \phi_i\|^2 + \tfrac{1}{2} \sum_{i,j=1}^{6} \|[\phi_i, \phi_j]\|^2 \,, \\
\mathcal{L}^{\psi}_{\text{kin}} &= \sum_{a=1}^{4} \langle \bar{\psi}_a, D^A \psi_a \rangle + \sum_{\substack{i=1,\dots,6 \\ a,b=1,\dots,4}} C^{iab} \langle \psi_a, [\phi_i, \psi_b] \rangle \,.
\end{aligned}$$

The coefficients C^{iab} are related to the structure constants of $SU(4)_R$, the R-symmetry of the $\mathcal{N} = 4$ super-Poincaré algebra explained in more detail below.

[2] Since the Clifford algebra acts on $\wedge^\bullet L$, the latter carries a representation of $\mathfrak{so}(4) \simeq \wedge^2 V \subset Cl(V)$. Spinor representations are usually *defined by* this representation, see for example [Del+99].

The definitions of both the classical and quantum theory rely on a well-defined variational principle for the action. For this it is necessary to specify boundary conditions. For $W^4 = \mathbb{R}^4$, one typically assumes that fields and their derivatives fall off sufficiently fast at infinity. In general, however, admissible boundary conditions are determined by the requirement that any boundary terms that arise in a variation of S vanish.

Under a variation $\delta\psi_a$ of one of the spinor fields, the variation of the action contains a boundary term $\int_{\partial W^4} \Gamma^{\perp}(\bar{\psi}_a, \delta\psi_a)$, where $\Gamma^{\perp}$ denotes the post-composition of Γ with projection to the direction perpendicular to the boundary. This boundary term vanishes as long as any non-zero part of ψ_a is orthogonal to its variations $\delta\psi_a = 0$ with respect to $\Gamma^{\perp}(\bar{\cdot}, \cdot)$. Put differently, admissible boundary conditions for ψ_a are determined by a choice of totally isotropic subspace $\mathfrak{S}_a \subset S^+$ and imposing $\psi_a|_{\partial W^4} \in \Gamma(\partial W^4, \mathfrak{S}_a \otimes \operatorname{ad} E)$. A similar analysis for the connection and scalar fields shows that their boundary conditions are in general given by Robin-type conditions, which relate normal derivatives and boundary values. All in all, admissible boundary conditions are given by configurations that satisfy

$$\begin{aligned} (F_A)_{y\mu} + \frac{\theta_{\mathrm{YM}} g_{\mathrm{YM}}^2}{32\pi^2} \epsilon_{y\mu\nu\lambda} (F_A)^{\nu\lambda} &= 0 \ , \\ \nabla_y^A \phi_i - \sum_{j=1}^{6} c_{ij} \phi_j &= 0 \ , \\ \psi_a|_{\partial W^4} &\in \Gamma(\partial W^4, \mathfrak{S}_a) \ . \end{aligned} \tag{2.1}$$

Note that the boundary conditions for the gauge field are completely fixed by a combination of the θ_{YM}-angle and coupling constant g_{YM}. In contrast, the coefficient functions c_{ij} are generally not restricted, though a generic choice will break invariance under the action of $SU(4)_R \simeq SO(6)$ that rotates the scalar fields ϕ_i into linear combinations of each other.

For $W^4 = \mathbb{R}^4$, the theory is invariant under the action of the $4d\ \mathcal{N} = 4$ super-Poincaré algebra. This is the $\mathbb{Z}/2\mathbb{Z}$-graded Lie algebra $A = A^0 \oplus A^1$ with bosonic and fermionic part given by

$$\begin{aligned} A^0 &= (V \rtimes \mathfrak{so}(V)) \times \mathfrak{su}(4)_R \ , \\ A^1 &= (S^+ \otimes \mathbb{C}^4) \oplus (S^- \otimes (\mathbb{C}^4)^*) \ . \end{aligned}$$

The inclusion of $\mathbb{C}^4$ in the fermionic part provides $\mathcal{N} = 4$ copies of the minimal super-Poincaré algebra in four dimensions. The $\mathfrak{su}(4)_R$ part in the bosonic part is the Lie algebra of R-symmetry, which is defined to be any transformation that is represented non-trivially on A^1 and commutes with the action of the Lorentz group $SO(V)$. With respect to $SU(4)_R$, the four spinors ψ_a are in the defining representation $\mathbb{C}^4$ and the six scalars ϕ_i in the six-dimensional vector representation.

The super-Poincaré algebra is equipped with a $\mathbb{Z}/2\mathbb{Z}$ graded Lie bracket

$$[x, y] = (-1)^{|x||y|+1} [y, x] \ ,$$

where $|x|$ denotes the degree of homogeneous elements. As a consequence A^1 carries a representation of A^0, while on A^1 the Lie bracket $[\cdot,\cdot] \ : \ A^1 \times A^1 \to A^0$ yields an intertwiner of

A^0-representations. For the spinorial part of A^1 this intertwiner is given by $\Gamma : S^+ \times S^- \to V$ extended by zero to all of S, together with the natural pairing on the $\mathbb{C}^4 \times (\mathbb{C}^4)^*$ factor. It follows that the anti-commutator of fermionic generators $[Q_1, Q_2]$ is always an element of V, corresponding to the common adage that supersymmetry squares to translations.

Consider now $W^4 = \mathbb{R}^3 \times \mathbb{R}^+$. The factor $\mathbb{R}^+ = [0, \infty)$ introduces a spacetime boundary and explicitly breaks translation invariance in the direction perpendicular to the boundary. Accordingly, only fermionic symmetries in A^1 that don't square to a translation in the normal direction $\mathbb{R}^\perp \subset V$ can be preserved. To that end, observe that $[Q_1, Q_2]|_{\mathbb{R}^\perp}$ is non-isotropic on $\mathbb{C}^4 \times (\mathbb{C}^4)^*$ and reduces to $\Gamma^\perp(\cdot, \cdot)$ on $S^+ \times S^-$. It follows that unbroken fermionic symmetries are elements of a totally isotropic subspace of the form $(\mathfrak{S} \otimes \mathbb{C}^4) \oplus (\overline{\mathfrak{S}} \otimes (\mathbb{C}^4)^*) \subset A^1$. As a consequence, the remaining supersymmetry algebra can contain at most half of the original fermionic generators.

A generic choice of boundary conditions satisfying (2.1) will not be invariant under the remaining super-Poincaré algebra. A complete classification of 'half-BPS' boundary conditions, those that are invariant under the remaining half of the fermionic generators, is described in [GW09b; GW09a]. Here, we only cite the result most relevant to us.

It turns out that for any choice of maximal totally isotropic subspace $\mathfrak{S} \subset S^+$, there exists a unique half-BPS boundary condition for (A, ϕ_i, ψ_a) that preserves the remaining half of the super-Poincaré algebra and full gauge symmetry. There is a basis of S^+ in which the choice of $\mathfrak{S}$ is equivalent to fixing a generator of the form $Q = \begin{pmatrix} 1 \\ t \end{pmatrix}$ with $t \in \mathbb{R}$, possibly infinite (interpreted as vanishing of the top component). Invariance of the boundary values of ψ_a under the action of Q fully determines the choice of what was called $\mathfrak{S}_a \subset S^+$ earlier. Namely, for each fermion we must have $\psi_a = \lambda_a \begin{pmatrix} t \\ 1 \end{pmatrix}$ for some $\lambda_a \in \mathbb{C}$. Invariance under the remaining Lorentz group $SO(3)$, R-symmetry $SO(3) \times SO(3) \subset SU(4)_R$, and supersymmetry then also fixes the boundary conditions of the bosonic fields.

$$
\begin{aligned}
(F_A)_{y\mu} + \frac{t}{1 - t^2} \epsilon_{y\mu\nu\lambda} F^{\nu\lambda} &= 0 \\
\nabla^A_y \phi_i - \frac{t}{1 + t^2} \epsilon_{ijk} [\phi_j, \phi_k] &= 0 \,, \quad i, j, k = 1, 2, 3 \\
\phi_{i+3} &= 0 \,, \quad i = 1, 2, 3 \\
\psi_a &= \lambda_a \otimes \begin{pmatrix} t \\ 1 \end{pmatrix}
\end{aligned}
$$

By comparison with the admissible boundary conditions of (2.1), it is clear that $\frac{t}{1-t^2} = \frac{\theta_{\mathrm{YM}} g^2_{\mathrm{YM}}}{32\pi^2}$. This condition has two roots, so for any value of SYM parameters $(g_{\mathrm{YM}}, \theta_{\mathrm{YM}})$ there are two half-BPS boundary conditions that preserve the full gauge symmetry and which differ only by a reversal of orientation. Conversely, for any choice of preserved supersymmetries $\mathfrak{S} \subset S^+$, there exists a set of SYM parameters $(g_{\mathrm{YM}}, \theta_{\mathrm{YM}})$ for which the equations above determine half-BPS boundary conditions.

Remark. The half-BPS equations for the scalar fields form the basis of a set of conditions known as Nahm pole boundary conditions. These conditions allow the inclusion of 't Hooft operators along a knot $K \subset \partial W^4$, by encoding a monodromy of A and a certain singular behaviour of ϕ_i near the boundary. Since the boundary conditions encode the presence of a knot K, this plays a crucial role in the gauge theoretic approach to Khovanov homology and is discussed in detail in Section 2.6.

2.2 The Kapustin-Witten Twists and Localization

While the topological term $\int_{W^4} \operatorname{Tr} F_A \wedge F_A$ in the action functional is manifestly independent of the metric on spacetime and only depends on the topology and smooth structure of the gauge bundle, this is not true for $\mathcal{L}_{\text{kin}}$. As a consequence, the partition function and observables of $4d$ $\mathcal{N} = 4$ super Yang-Mills theory are generally not topological invariants. However, in supersymmetric theories it is often possible to restrict to a subsector of the theory that depends only topologically on spacetime by a procedure known as twisting.

Below we shortly recall the twisting procedure and subsequently present the relevant twists of $4d$ $\mathcal{N} = 4$ SYM theory. Importantly, half-BPS boundary conditions retain enough supersymmetry that the theory still admits a topological twist. The section concludes with an explanation of how twisting leads to the Kapustin-Witten equations.

Topological Twists of Supersymmetric Theories Topological twists are a standard tool to extract topological field theories from supersymmetric ones [Wit88a; Wit88b; Wit91]. Here we closely follow expositions in [Ell13; ESW18]; for a more thorough introduction see [Von05].

First, recall from Noether's theorem that for every continuous symmetry of the action functional there is an associated conserved current j. Applying this to translation symmetry of a field theory on $\mathbb{R}^n$ results in a conserved current for each element of $\mathbb{R}^n$. Choosing a basis x^μ for $\mathbb{R}^n$ the components of the associated conserved current are $j_\mu = T_{\mu\nu} dx^\nu$. The 2-tensor T is the (canonical) energy-momentum tensor associated to the field theory and is related to variations of the metric by $\delta_g S = \int_{\mathbb{R}^4} T_{\mu\nu} \delta_{g_{\mu\nu}} \mathcal{L}$. The observables of a theory are independent of the metric – hence, topological invariants in the sense of physics – if the energy-momentum tensor vanishes, so translations must act trivially.

Even if a given field theory generally depends on the metric, it is still possible that certain protected subsectors of the theory are purely topological. This happens, for example, if there is a non-anomalous symmetry Q with the following two properties.

- Q is nilpotent (more precisely, square-zero): $Q^2 = 0$.
- T is Q-exact: there exists a functional V such that $T = [Q, V]$.

Now recall the general fact that in any quantum theory the expectation value of a symmetry-exact operator $\langle[Q,\mathcal{O}]\rangle$ vanishes. This can be seen, at least formally, by the following argument. If Q is a non-anomalous symmetry, i.e. the path-integral measure is invariant under transformations induced by Q and $[Q,S]=0$, then the following expression must be independent of the infinitesimal parameter ϵ

$$\int \mathcal{D}\Phi \exp(\epsilon Q)\big(\mathcal{O}\exp(-S)\big) = \int \mathcal{D}\Phi \exp(-S)\big(\mathcal{O}+\epsilon[Q,\mathcal{O}]\big)\ .$$

This can only be the case if the second term vanishes, which is exactly the statement $\langle[Q,\mathcal{O}]\rangle = 0$. As a particular consequence we observe that when $Q^2 = 0$, the subsector given by Q-cohomology, i.e. the set of Q-invariant operators modulo those operators that are invariant for the uninteresting reason of being Q-exact, determines a viable quantum field theory in its own right. Moreover, the fact that the energy-momentum tensor is Q-exact implies that the expectation value of any Q-closed observable is topologically invariant. This follows from studying the effect of a continuous variation of the metric on the expectation value of a Q-closed operator $\mathcal{O}$:

$$\delta_g\langle\mathcal{O}\rangle = \int \mathcal{D}\Phi\ \mathcal{O}\ \delta_g \exp(iS) = \int \mathcal{D}\Phi\ \mathcal{O}[Q,V]\ \exp(iS) = \langle[Q,\mathcal{O}V]\rangle = 0\ .$$

In summary, if Q satisfies the two properties stated above, Q-cohomology provides a topological subsector of the theory.

Since supersymmetric theories are invariant under the super-Poincaré algebra $A = A^0 \oplus A^1$, which contains nilpotent symmetries, one can ask if there are $Q \in A^1$ that satisfy the two properties described above. With regards to the first property, the set of nilpotent fermionic elements of any $\mathbb{Z}/2\mathbb{Z}$-graded algebra forms a projective variety $Y := \{Q^2 = 0\} \subset A^1$, known as the nilpotence variety [ESW18]. The nilpotence variety of the super-Poincaré algebra depends only on the spacetime dimension and quantity of supersymmetry $\mathcal{N}$. In view of the discussion above, every $Q \in Y(d,\mathcal{N})$ gives rise to an associated Q-invariant subsector of the theory.

For this to be a topological sector the energy momentum tensor must be Q-exact, which it cannot be, because Q is in the spinor representation of the Lorentz group. This problem can sometimes be circumvented by 'twisting' the action of the Lorentz algebra $\mathfrak{so}(d)$ in such a way that Q can be reinterpreted as a scalar operator. To that end note that upon restriction to Q-cohomology the symmetry algebra is broken to the stabilizer of the line spanned by Q:

$$I(Q) = \{x \in A \mid [x,Q] \propto Q\}\ .$$

Other elements of A do not preserve the kernel and image of Q, and therefore do not act on Q-cohomology. With respect to $I(Q)$, Q is tautologically a scalar (perhaps up to some currently irrelevant $U(1)$ charges). Since the Lorentz symmetry and the semi-simple part of R-symmetry act non-trivially on Q, $I(Q)$ can't contain generators of the corresponding subalgebras. However, it can contain combinations of the two, where the action of the Lorentz algebra is undone by the action of R-symmetry. Indeed, $I(Q)$ may contain a bosonic subalgebra $\mathfrak{so}'(d)$ that is isomorphic to the Lorentz algebra and is embedded in A^0 as follows

$$\mathfrak{so}(d) \xrightarrow{1\times\alpha} \mathfrak{so}(d) \times R \subset A^0\ .$$

The map α is (the derivative of) a non-trivial homomorphism from the Lorentz group to the R-symmetry group and is commonly known as the twisting map. Note that for any $Q \in Y$, there might exist several twisting maps and conversely the graph of a fixed twisting map might appear in the unbroken symmetries $I(Q)$ of several Q's.

When such a twisted Lorentz symmetry exists one can view the Q-invariant subsector as a field theory with $\mathfrak{so}'(d)$ invariance in its own right. Crucially, in this reinterpretation Q is a Lorentz scalar and the energy momentum tensor can be Q-exact. This in turn is the case if all translations are Q-exact, i.e. if the subalgebra $E(Q) := [Q, A^1]$ is all of V.

In summary, to twist a supersymmetric theory means to restrict to the subsector of a theory given by Q-cohomology of some nilpotent supercharge $Q \in Y$, together with a choice of twisting homomorphism that identifies a Lorentz subgroup $\mathfrak{so}'(d) \subset I(Q)$. If $E(Q) = V$, the result is a topological theory and is referred to as topological twist of the original theory. Topologically twisted theories can manifestly be defined on general Riemannian manifolds M^n, while their supersymmetric 'parents' may only make sense on very special manifolds: e.g. $\mathbb{R}^n$, or manifolds that admit covariantly constant spinors. Although the metric of the Riemannian manifold is needed in the definition of the action, the existence of the nilpotent symmetry Q makes sure that the theory is independent of the metric.

Kapustin-Witten Twists $\mathcal{N} = 4$ super Yang-Mills theory on $\mathbb{R}^4$ admits several interesting twisting maps [Yam88; Mar95]; also see [LL97] for a more complete discussion. The topological twist that is relevant to Haydys-Witten Floer theory and Khovanov homology is commonly dubbed Kapustin-Witten or Geometric Langlands twist. Here we cite the most relevant results; more detailed explanations can be found in [KW07; Wit11a].

The four-dimensional Lorentz algebra is isomorphic to $\mathfrak{so}(4) \simeq \mathfrak{su}(2)_\ell \times \mathfrak{su}(2)_r$, where the subscripts ℓ and r stand for *left* and *right* chiral part. The Kapustin-Witten twisting homomorphism is given by a diagonal embedding

$$\alpha : \mathfrak{so}(4) \simeq \mathfrak{su}(2)_\ell \times \mathfrak{su}(2)_r \hookrightarrow \begin{pmatrix} \mathfrak{su}(2)_\ell & 0 \\ 0 & \mathfrak{su}(2)_r \end{pmatrix} \subset \mathfrak{su}(4)_R \, .$$

The centralizer of the graph of α in $\mathfrak{so}(4) \times \mathfrak{su}(4)_R$ is an additional $U(1)$ that acts on the two blocks by multiplication with $e^{\pm i\omega}$, respectively. This gives rise to a $U(1)$ symmetry in the twisted theory and an associated $\mathbb{Z}$-grading on the twisted fields. It will play the role of a fermion number F and will be denoted $U(1)_F$ and $\mathbb{Z}_F$ from now on.

The full nilpotence variety in $4d$ $\mathcal{N} = 4$ SYM theory is an 11-dimensional stratified variety [ESW18]. Of those there is a $\mathbb{CP}^1$-family of supercharges for which on the one hand $I(Q)$ contains the (fixed) twisted Lorentz algebra $\mathfrak{so}'(4)$ and on the other hand $E(Q) = V$, such that Q-cohomology becomes topological after twisting.

From the point of view of the $SO'(4)$ Lorentz group, the fields transform in the following representations

- The gauge connection A transforms trivially under R-symmetry and thus remains unchanged.
- Four of the scalar fields ϕ_i combine into a vector representation ϕ with $F = 0$, while the remaining two are $SO(4)'$ scalars that are rotated into each other by $U(1)_F \simeq SO(2)$. Combining the latter two into $\sigma, \bar{\sigma} = \phi_5 \pm i\phi_6$ makes these into $\mathfrak{g}_{\mathbb{C}}$-valued Lorentz scalars with fermion number $F = \pm 2$.
- The sixteen real components of the four Weyl fermions ψ_a distribute into two vector representations λ_1, λ_2 of $SO(4)'$ with $F = 1$, an antisymmetric representation χ with $F = -1$, and two Lorentz scalars η_1, η_2 with $F = -1$.

On a general Riemannian manifold the field content of the twisted theory thus arranges into a $\mathbb{Z}$-graded chain complex, where Q acts as differential of degree 1.

$F = -2$	-1	0	1	2
$\bar{\sigma} \in \Omega^1(W^4, \operatorname{ad} E_{\mathbb{C}})$	$\eta_1, \eta_2 \in \Omega^0(W^4, \operatorname{ad} E)$ $\chi \in \Omega^2(W^4, \operatorname{ad} E)$	$A \in \mathcal{A}(E)$ $\phi \in \Omega^1(W^4, \operatorname{ad} E)$	$\lambda_1, \lambda_2 \in \Omega^1(W^4, \operatorname{ad} E)$	$\sigma \in \Omega^1(W^4, \operatorname{ad} E_{\mathbb{C}})$

Recall that in the presence of a boundary, a totally isotropic subspace $(\mathfrak{S} \otimes \mathbb{C}^4) \oplus (\overline{\mathfrak{S}} \otimes (\mathbb{C}^4)^*) \subset A^1$ of the fermionic symmetries can be preserved by choosing half-BPS boundary conditions as described in Section 2.1. The preserved supersymmetry determined by $\mathfrak{S}$ and the $\mathbb{CP}^1$-family of nilpotent charges that admit a compatible Kapustin-Witten twist intersect in a single point Q [Wit11a].

Localization of Topologically Twisted Partition Function Topologically twisted $\mathcal{N} = 4$ super Yang-Mills theory is defined on a general Riemannian four-manifold with boundary W^4. As before, the action is given by the expression

$$S = \frac{1}{g_{\mathrm{YM}^2}} \int_{W^4} \mathcal{L}_{\mathrm{kin}} + \frac{\theta_{\mathrm{YM}}}{32\pi^2} \int_{W^4} \operatorname{Tr} F_A \wedge F_A \ ,$$

where the Lagrangian arises from the one on flat Euclidean space by rewriting it in terms of the $SO'(4)$ fields on W^4, and adding curvature terms as necessary to make the action Q-invariant. For example the part of the Lagrangian that contains the gauge connection A and the one-form $\phi \in \Omega^1(W^4, \operatorname{ad} E)$ is given by

$$\mathcal{L}^{A,\phi}_{\mathrm{kin}} = \tfrac{1}{2} \|F_A\|^2 + \|d_A \phi\|^2 + \tfrac{1}{2} \|[\phi \wedge \phi]\| + \langle \operatorname{Ric} \phi, \phi \rangle \ ,$$

where Ric denotes the Ricci tensor, viewed as linear map on $\Omega^1(W^4)$.

Let us now investigate the partition function of the topologically twisted theory, possibly including the insertion of a 't Hooft operator supported on $K \subset \partial W^4$ as specified by the BPS boundary conditions of Section 2.1. We denote the partition function of this theory by

$$Z^Q_{\mathrm{SYM}}(W^4, K) = \int \mathcal{D}\Phi \exp(-S) \ .$$

It turns out that the path integral localizes on field configurations that obey $[Q, \Psi] = 0$, for all Ψ with odd fermion number F. Among the fields with fermion number $F = -1$ is the two-form χ, with (anti-)self-dual parts $\chi^{\pm}$, and a scalar η_1. The action of Q on these fields produces

$$[Q, \chi^+] = (F_A - [\phi \wedge \phi] + t d_A \phi)^+ \ , \quad [Q, \chi^-] = (F_A - [\phi \wedge \phi] - t^{-1} d_A \phi)^- \ , \quad [Q, \eta_1] = d_A^{\star_4} \phi \ .$$

Since the topologically twisted theory only depends on Q-cohomology, we can modify the action by adding Q-exact terms without changing the partition function.

$$\begin{aligned} S' &= S - \left[Q, \ \frac{1}{\epsilon} \int_{W^4} \left(\langle \chi^+, [Q, \chi^+] \rangle + \langle \chi^-, [Q, \chi^-] \rangle + \langle \eta, [Q, \eta] \rangle \right) \right] \\ &= S - \frac{1}{\epsilon} \int_{W^4} \left(\left\| [Q, \chi^+] \right\|^2 + \| [Q, \chi^-] \|^2 + \| [Q, \eta] \|^2 + \ldots \right) \end{aligned}$$

The omitted terms on the right are fermion bilinears and can be neglected, since in a supersymmetric ground state fermionic fields vanish. Taking $\epsilon \to 0$, the action diverges except when $[Q, \chi^{\pm}] = [Q, \eta] = 0$, such that the path integral localizes on field configurations that satisfy these equations. These are known as Kapustin-Witten equations and will be discussed in much more detail in Section 2.4.2. There is a similar localization argument for the complex scalar σ and its complex conjugate that provides the equations $d_A \sigma = [\phi, \sigma] = [\sigma, \bar{\sigma}] = 0$. These equations imply that σ and $\bar{\sigma}$ vanish, at least as long as (A, ϕ) are irreducible solutions of the Kapustin-Witten equations.

A standard argument relates the expected dimension of the moduli space of solutions of the localization equations to the index of a Dirac-like operator on W^4. In the path-integral formalism this index is equal to the anomaly in the fermion number F, so the partition function of the topologically twisted theory on W^4 vanishes except when the expected dimension is zero. As a result, the partition function reduces to a sum over classical solutions of the localization equations on W^4 that satisfy the supersymmetric boundary conditions described in Section 2.1 if $\partial W^4 \neq \emptyset$.

Consider now the contribution to the partition function from a given solution of the localization equations. In expanding the path integral around the solution we assume that there are no bosonic or fermionic zero modes and the solution fully breaks gauge symmetry, which is the expected case if the index is zero. Taking $g_{\mathrm{YM}} \to 0$, the calculation reduces to a one-loop approximation. On the one hand this results in a factor of $\exp(-\theta_{\mathrm{YM}} P) =: q^P$ from the topological part of the action. Here we abbreviated $P = \frac{1}{32\pi^2} \int_{W^4} F \wedge F$, which is the integral of the second Chern class of the gauge bundle E and is known as instanton number. On the other hand, we pick up the ratio of the fermion and boson determinant, which due to supersymmetry are equal up to a sign $(-1)^F$, where F is interpreted as the fermion number of the solution. It follows that any classical solution of the localization equations contributes to the partition function with a term $(-1)^F q^P$.

Writing $\mathcal{M}^{\mathrm{KW}}$ for the set of classical solutions of the Kapustin-Witten equations, we find that the topological partition function is given by

$$Z^Q_{\mathrm{SYM}}(W^4, K; q) = \sum_{\Phi \in \mathcal{M}^{\mathrm{KW}}} (-1)^{F(\Phi)} q^{P(\Phi)} = \sum_{n \in \mathbb{Z}} a_n q^n \ .$$

In rewriting the sum in the second equation, we assume that there are only finitely many solutions and denote by a_n the number of solutions with $P = n$. Since the presence of a 't Hooft operator along a knot $K \subset \partial W^4$ is encoded in the boundary conditions, it affects the coefficients a_n.

To conclude this section, we cite a further key result of [Wit10; Wit11a], which states that the action of SYM theory on $W^4 = S^3 \times \mathbb{R}^+$ is equivalent to

$$S = [Q, \cdot] + i\psi\, CS(A + i\omega\phi) ,$$

where ψ and ω are determined by θ_{YM} – or, depending on preference, the choice of preserved supersymmetry. As was the case for the instanton number P, the definition of the Chern-Simons functional in that equation needs a bit of care, see [Wit10; Wit11a; Wit11b].

The twisted partition function is independent of Q-exact operators, so topologically twisted SYM theory on $W^4 = S^3 \times \mathbb{R}^+$ together with a 't Hooft operator along $K \subset \partial W^4$ calculates the Jones polynomial of K:

$$Z^Q_{SYM}(S^3 \times \mathbb{R}^+, K) = Z_{CS}(S^3, K) = J(K) .$$

2.3 Hilbert Space of BPS States in Five Dimensions

The Jones polynomial admits a categorification known as Khovanov homology [Kho99]. Categorification involves replacing a classical mathematical object with a richer, more structured object 'one step up' in category theory. The new object usually captures more information and often allows for more powerful tools and insights.

In the case of Khovanov homology and the Jones polynomial, categorification replaces a polynomial invariant of knots and links with a homology theory that assigns to each knot or link K a bigraded vector space $Kh(K) = \bigoplus_{i,j} Kh^{i,j}(K)$. The Jones polynomial is obtained from Khovanov homology as the Euler characteristic of the graded vector space

$$J(K) = \sum_{i,j} (-1)^i q^j \dim Kh^{i,j}(K) .$$

A similar phenomenon exists in supersymmetric quantum field theories when the localization procedure reduces the path integral to a sum over supersymmetric vacua. In such situations the partition function can often be expressed as a trace over the Hilbert space of BPS states in one dimension higher.

Indeed, it is well-known that $4d\,\mathcal{N} = 4$ super Yang-Mills theory can be viewed as the low-energy effective theory of a $5d\,\mathcal{N} = 2$ super Yang-Mills theory[3] compactified on a small circle S^1. The Nahm pole boundary condition lifts to the five-dimensional theory by a translation-invariant continuation in the direction of the circle. In particular, any 't Hooft operator supported on a knot K in ∂W^4 lifts to an S^1-invariant surface operator supported on $\Sigma_K = S^1 \times K$ in the boundary of the five-manifold. The Nahm pole boundary condition and 't Hooft operator preserve the same supersymmetry generator as in the four-dimensional theory. Furthermore, the Kapustin-Witten twisting homomorphism of the four-dimensional theory corresponds to an analogous topological twist in five dimensions. In particular, the topological supercharge Q that defined the topological subsector in four dimensions, remains a nilpotent symmetry when the model is lifted to five dimensions.

In the five-dimensional theory on $S^1 \times W^4$ one constructs the one-particle Hilbert space $\mathcal{H}$ by canonical quantization on a 'temporal' slice, i.e. on a codimension one submanifold $\{s\} \times W^4$, $s \in S^1$, in the background determined by the boundary conditions. The partition function of the topologically twisted four-dimensional theory on W^4 is then equivalent to a trace in $\mathcal{H}$ with certain operator insertions:

$$Z^Q_{\mathrm{SYM}}(q) = \mathrm{Tr}_{\mathcal{H}}(-1)^F q^P \, .$$

As before F and P denote fermion and instanton number, but in the five-dimensional theory are interpreted as operators on Hilbert space. P is given in terms of the classical integral of the four-dimensional fields and promoted to an operator by classical quantization. The operator F in the expansion around a classical solution is given by summing over the F-eigenvalues of the filled Dirac sea, i.e. the total fermion number of the combination of all negative energy states.

In this approach $\mathcal{H}$ plays the role of the chain complex underlying Khovanov cohomology. Indeed, since Q is a nilpotent fermionic symmetry, it satisfies $Q^2 = 0$ and accordingly acts as a differential on the Hilbert space. The cohomology of Q is commonly known as the space of BPS states $\mathcal{H}_{\mathrm{BPS}} := H^\bullet(\mathcal{H}, Q)$ and, in the current situation, is spanned by supersymmetric ground states. While the physical Hilbert space generally depends on various parameters and perhaps on choices during quantization, BPS states are protected by supersymmetry and are independent of continuous deformations of the theory. As a consequence $\mathcal{H}_{\mathrm{BPS}} = \mathcal{H}_{\mathrm{BPS}}(W^4, K)$ is a knot invariant and one expects that it is a gauge theoretic manifestation of Khovanov homology.

In a first approximation, the quantum ground states of the five-dimensional theory are determined by time-independent classical solutions of the equations of motion. The construction outlined above results in the fact that these correspond to solutions of the localization equations of the four-dimensional theory. For simplicity assume that there is a finite, non-degenerate set of solutions $\mathcal{M}^{\mathrm{KW}}$, in particular assume that after gauge fixing the solutions do not have bosonic zero modes. Then an expansion around a given solution $\Phi \in \mathcal{M}^{\mathrm{KW}}$ produces a single perturbative ground state $|\Phi\rangle$ of zero energy.

To move beyond perturbation theory, consider the space $\mathcal{H}_0$ spanned by all perturbative supersymmetric ground states. States in $\mathcal{H}_0$ can fail to be true quantum mechanical ground states if

[3] Five-dimensional super Yang-Mills theory is not UV-complete and from the physics point of view it may be more satisfying to consider compactifications of the $6d\ \mathcal{N} = (2,0)$ superconformal theory or of $10d\ \mathcal{N} = 1$ super Yang-Mills theory, see [Wit11a]. From the Floer theory point of view, as adopted here, the five-dimensional interpretation is more natural.

they are lifted from zero energy by some non-perturbative process. It is a well-known property of supersymmetric quantum theories that eigenstates of the Hamiltonian with non-zero energy must appear in Bose-Fermi pairs that only differ with respect to their spin statistics, which in the current situation is given by $\mathbb{Z}_F/2\mathbb{Z}$. Recall that $\mathcal{H}_0$ carries a $\mathbb{Z}_P\times\mathbb{Z}_F$ grading[4] with respect to the instanton and fermion numbers. Since the fermion operator satisfies $[F,Q]=Q$, the fermion numbers of such a Bose-Fermi pair differ by exactly one, while their instanton numbers coincide.

In the present context, quantum corrections can only arise by tunneling from one classical (perturbative) solution to another. To understand the fundamentally quantum mechanical process of tunneling in the context of field theories, it is helpful to slightly change perspective. The result is the well-known application of Morse theory to Yang-Mills theory pioneered by Floer in the context of flat connections on three-manifolds and anti-self-dual connections (Yang-Mills instantons) on four-manifolds [Flo88a; Flo89]. For this, one reinterprets super Yang-Mills theory on $\mathbb{R}_s\times W^4$ as supersymmetric quantum mechanics on the space of field configurations over W^4, where the real coordinate s takes the place of the time coordinate in quantum mechanics. If one interprets the action functional of super Yang-Mills theory as potential energy, one finds that a tunneling event corresponds to a solution of gradient flow equations [Wit82]. Following this line of argument out in the current context and formulating everything in terms of five-dimensional fields, one arrives at the Haydys-Witten equations on $\mathbb{R}_s\times W^4$ [Wit11a]. We will provide a detailed description of these equations and their specializations in Section 2.4.

The space of exact BPS ground states $\mathcal{H}_{\mathrm{BPS}}$ is now given by Q-cohomology of the approximate space $\mathcal{H}_0$, where Q acts by instanton corrections as explained above. Since Q has F-degree $+1$, or by the Bose-Fermi pair argument from above, only gradient flows that interpolate between solutions whose fermion numbers f_i and f_j differ by 1 can contribute to the correction. The action of Q is then given by

$$Q\,|\Phi_i\rangle = \sum_{\substack{\Phi_j\in\mathcal{M}^{\mathrm{KW}}\\ |f_i-f_j|=1}} n_{ij}\,|\Phi_j\rangle \ ,$$

where n_{ij} is the signed count of instanton solutions that interpolate between the solutions Φ_i and Φ_j.

Conceptually, the approximate Hilbert space $\mathcal{H}_0$ corresponds to the Morse-Smale-Witten complex, quantum corrections are given by the instanton Floer differential associated to the Haydys-Witten equations, and the Hilbert space of BPS states is given by the associated Floer cohomology of the manifold. The mathematical formulation of this instanton Floer theory is standard and will be summarized in Section 2.7. The key insight provided by physics is that this Floer theory gives rise to interesting topological invariants of four-manifolds. In particular, it suggests that there is a Floer theory interpretation of Khovanov homology that arises if one considers a manifold with boundary, where boundary conditions encode a 't Hooft operator insertion.

[4] On general five-manifolds there is no $U(1)_F$ symmetry, but there is still a $\mathbb{Z}/2\mathbb{Z}$ grading by spin-statistics.

2.4 The Haydys-Witten Equations and their Specializations

The instanton equations of the preceding section are conveniently summarized in a formulation due to Haydys [Hay15]. In this section we first introduce Haydys' geometric setup and the Haydys-Witten equations. This provides a covariant lift of anti-self-dual equations in four dimensions to five-manifolds M^5 that are equipped with a non-vanishing vector field v.

Moreover, we review the definitions and some relevant properties of several closely related differential equations on lower dimensional manifolds. Namely, the θ-Kapustin-Witten and Vafa-Witten equations on four-manifolds, the twisted extended Bogomolny equations on three-manifolds, and Nahm's equations on one-dimensional manifolds. All these equations can be viewed as dimensional reductions of the Haydys-Witten equations. We dedicate Section 2.5 to a detailed discussion of this fact.

2.4.1 The Haydys-Witten Equations

The Haydys-Witten Equations are a set of partial differential equations on Riemannian five-manifolds that are equipped with a nowhere-vanishing vector field v. The equations were introduced by Haydys on general five-manifolds [Hay15] and at roughly the same time discovered independently by Witten in the special case $\mathbb{R}_s \times X^3 \times \mathbb{R}^+_y$ with $v = \partial_y$ [Wit11a]. This section closely follows the original exposition of [Hay15].

Let (M^5, g) be a Riemannian five-manifold and v be a nowhere-vanishing vector field of pointwise unit norm. Consider a principal G-bundle $E \to M^5$ for G a compact Lie group, write $\mathcal{A}(E)$ for the space of gauge connections, and denote by $\operatorname{ad} E$ the adjoint bundle associated to the Lie algebra $\mathfrak{g}$ of G. Furthermore, for a gauge connection $A \in \mathcal{A}(E)$ we denote the associated covariant derivative by ∇^A and the exterior covariant derivative by d_A.

Write $\eta = g(v, \cdot) \in \Omega^1(M)$ for the one-form dual to the vector field v and observe that the pointwise linear map

$$T_\eta \, : \, \Omega^2(M) \to \Omega^2(M), \quad \omega \mapsto \star_5(\omega \wedge \eta)$$

has eigenvalues $\{-1, 0, 1\}$, such that $\Omega^2(M)$ decomposes into the corresponding eigenspaces:

$$\Omega^2(M) = \Omega^2_{v,-}(M) \oplus \Omega^2_{v,0}(M) \oplus \Omega^2_{v,+}(M) \, .$$

Below we will use the notation ω^+ to denote the part of ω that lies in $\Omega^2_{v,+}(M)$.

At every point $p \in M^5$, the fiber $\Omega^2_{v,+}(M)|_p$ is a three-dimensional vector space and thus carries a natural Lie algebra structure given by the usual cross product $(\cdot \times \cdot)$ of $\mathbb{R}^3$, unique up to orientation. The map $\sigma(\cdot, \cdot) = \frac{1}{2}(\cdot \times \cdot) \otimes [\cdot, \cdot]_{\mathfrak{g}}$ determines a corresponding cross product on $\Omega^2_{v,+}(M, \operatorname{ad} E) \simeq \mathbb{R}^3 \otimes \mathfrak{g}$.

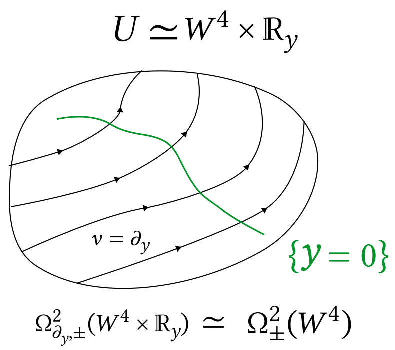

Figure 2.1

Example. To parse these constructions consider their incarnation in a small neighbourhood U of a point $p \in M$, see Figure 2.1. Choose orthonormal coordinates $(x_i, y)_{i=0,1,2,3}$ based at p in such a way that $v = \partial_y$. An explicit basis of $\Omega^2_{v,+}(U)|_p$ is given by

$$e_1 = dx^0 \wedge dx^1 + dx^2 \wedge dx^3 \,, \quad e_2 = dx^0 \wedge dx^2 - dx^1 \wedge dx^3 \,, \quad e_3 = dx^0 \wedge dx^3 + dx^1 \wedge dx^2 \,.$$

Observe that this is also a basis of the self-dual two-forms with respect to $\star_4$ acting on the orthogonal complement of v in $\Omega^2(U)|_p$. We can extend the e_i to a local frame over all of U such that an element of $\Omega^2_{v,+}(U, \operatorname{ad} E)$ is of the form $B = e_1 \otimes \phi_1 + e_2 \otimes \phi_2 + e_3 \otimes \phi_3$ for some $\mathfrak{g}$-valued functions ϕ_a, $a = 1, 2, 3$. Equivalently, the non-vanishing components of B are $B_{0a} = \phi_a$, $B_{ab} = \epsilon_{abc}\phi_c$. Finally, the cross product of B with itself is given by

$$\sigma(B, B) = e_1 \otimes [\phi_2, \phi_3] + e_2 \otimes [\phi_3, \phi_1] + e_3 \otimes [\phi_1, \phi_2] \,,$$

which in 2-form components corresponds to $\sigma(B, B)_{\mu\nu} = \frac{1}{4} g^{\rho\tau}[B_{\mu\rho}, B_{\nu\tau}]$.

Example. Consider $(M^5, v) = (W^4 \times \mathbb{R}_y \,, \partial_y\,)$ with product metric and denote by $i \;:\; W^4 \hookrightarrow W^4 \times \mathbb{R}_y$ inclusion at $y = 0$. Then

$$\Omega^2_\pm(W^4) \simeq i^* \left(\Omega^2_{\partial_y,\pm}(W^4 \times \mathbb{R}_y) \right) \,,$$

where $\Omega^2_\pm(W^4)$ denotes (anti-)self-dual two-forms on W^4 with respect to $\star_4$. Indeed, whenever the metric is of product type, the Hodge star operator factorizes as

$$\star_5(\alpha \wedge \beta) = (-1)^{k\ell} \star_{W^4} \alpha \wedge \star_{\mathbb{R}_y} \beta \,,$$

when $\alpha \in \Omega^k(W^4)$, $\beta \in \Omega^\ell(\mathbb{R}_y)$ and this gives rise to the stated isomorphisms.

The two examples demonstrate that Haydys' setup provides a covariant lift of (anti-)self-dual 2-forms into 5-manifolds. Given the relevance of gauge theory and anti-self-dual connections in the study of 4-manifolds, Haydys then suggests to consider a closely related set of equations in five-dimensional gauge theory. The equations make use of the following differential for Haydys-self-dual two-forms

$$\delta^+_A \;:\; \Omega^2_{v,+}(M, \operatorname{ad} E) \xrightarrow{\nabla^A} T^*M \otimes \Omega^2_{v,+}(M, \operatorname{ad} E) \simeq TM \otimes \Omega^2_{v,+}(M, \operatorname{ad} E) \xrightarrow{-\iota} \Omega^1(M, \operatorname{ad} E) \,.$$

Here, as usual, ∇^A is the covariant derivative with respect to both, the gauge connection and the Levi-Civita connection, while we use ι to denote contraction $\iota(u\otimes\omega) := \iota_u\omega$. In normal coordinates $(x^\mu, y)_{\mu=0,1,2,3}$ with $v = \partial_y$, the action of this differential is given by $\delta_A^+ B = -\sum_{\mu=0}^3 \nabla_\mu^A \, \iota_\mu B$, since by construction $\iota_{\partial_y} B = 0$.

Definition 2.1 (Haydys-Witten Equations)**.** Let (M^5, v) be a Riemannian 5-manifold together with a preferred vector field and $E \to M^5$ a principal G-bundle. With notation as above, consider a pair $(A, B) \in \mathcal{A}(E) \times \Omega^2_{v,+}(\operatorname{ad} E)$. The Haydys-Witten equations for (A, B) are given by:

$$\begin{aligned} F_A^+ - \sigma(B,B) - \nabla_v^A B &= 0 \\ \iota_v F_A - \delta_A^+ B &= 0 \end{aligned} \tag{2.2}$$

We denote the corresponding differential operator by

$$\mathbf{HW}_v \; : \mathcal{A}(E) \times \Omega^2_{v,+}(\operatorname{ad} E) \to \Omega^2_{v,+}(\operatorname{ad} E) \times \Omega^1(\operatorname{ad} E) \, .$$

If $B = 0$ the Haydys-Witten equations provide a five-dimensional analogue of the anti-self-dual equations that underlie Donaldson-Floer theory on 4-manifolds. This special case was studied from a slightly different perspective in [Fan96].

As mentioned already above, the perspective from supersymmetric Yang-Mills theory suggests that the Haydys-Witten equations have an interpretation in the six-dimensional $\mathcal{N} = (2,0)$ SCFT. In fact they can be obtained via dimensional reduction in several, closely related ways: from six-dimensional equations [Wit11a, sec. 5], from octonionic monopole equations on seven-manifolds with G_2 holonomy, or from Spin(7)-instantons on eight-manifolds [Che15]. The octonionic structure is closely related to the 10d $\mathcal{N} = 1$ super-Poincaré algebra.

2.4.2 The Kapustin-Witten Equations

The Kapustin-Witten equations are a one-parameter family of partial differential equations on Riemannian four-manifolds W^4. They were first described by Kapustin and Witten in the context of the geometric Langlands program and its interpretation in quantum field theory, in which case one considers a product of two Riemann surfaces $W^4 = \Sigma \times C$ [KW07]. A few years later the equations resurfaced in the context of the gauge theoretic approach to Khovanov homology on manifolds of the form $W^4 = X^3 \times \mathbb{R}^+$ [Wit11a]. Since then their study has grown into an active field of research, for an incomplete list of recent developments see e.g. [GU12; Tau13; He15; He19; Tau17a; Tan19; NO19b].

Let (W^4, g) be a smooth Riemannian four-manifold and G a compact Lie group. Consider a principal G-bundle $E \to W^4$ together with a principal connection A, and denote by $\operatorname{ad} E$ the adjoint bundle associated to the Lie algebra $\mathfrak{g}$ of G. Furthermore, consider an $\operatorname{ad} E$-valued one-form $\phi \in \Omega^1(W^4, \operatorname{ad} E)$, usually called the Higgs field. The θ-Kapustin-Witten equations for the pair (A, ϕ) and an angle $\theta \in [0, 2\pi]$ are the following family of differential equations.

Definition 2.2 (Kapustin-Witten Equations).

$$
\begin{aligned}
\left(\cos\tfrac{\theta}{2}\, (F_A - \tfrac{1}{2}[\phi\wedge\phi]) - \sin\tfrac{\theta}{2}\, d_A\phi \right)^+ &= 0 \\
\left(\sin\tfrac{\theta}{2}\, (F_A - \tfrac{1}{2}[\phi\wedge\phi]) + \cos\tfrac{\theta}{2}\, d_A\phi \right)^- &= 0 \\
d_A^{\star_4}\phi &= 0
\end{aligned}
\tag{2.3}
$$

The corresponding differential operator is denoted

$$
\mathbf{KW}_\theta \;:\; \mathcal{A}(E)\times\Omega^1(X^4,\operatorname{ad}E) \to \Omega^2_+(X^4,\operatorname{ad}E)\times\Omega^2_-(X^4,\operatorname{ad}E)\times\Omega^0(X^4,\operatorname{ad}E)\,.
$$

For $\theta \neq 0 \pmod{\pi}$ the self-dual and anti-self-dual parts in (2.3) can be combined into the following single expression

$$
F_A - \tfrac{1}{2}[\phi\wedge\phi] + \cot\theta d_A\phi - \csc\theta \star_4 d_A\phi = 0\,. \tag{2.4}
$$

Furthermore, as observed already by Kapustin and Witten [KW07] and discussed in more detail by Gagliardo and Uhlenbeck [GU12], this is equivalent to phase-shifted conjugate anti-self-dual equations for the $G_\mathbb{C}$-gauge connection $A + i\phi$:

$$
F_{A+i\phi} = e^{i(\pi-\theta)} \star_4 \overline{F_{A+i\phi}}\;\;. \tag{2.5}
$$

This point of view suggests the applicability of several powerful results from the theory of self-dual Yang-Mills connections and geometric analysis in general. For example it is apparent from (2.5) that in the case $\phi = 0$ the one-parameter family of Kapustin-Witten equations interpolates between Donaldson's anti-self-dual equations ($\theta = 0$) and self-dual equations ($\theta = \pi$).

Remark. In the literature it is more common to parametrize the family of Kapustin-Witten equations by $\theta_{GU} = \theta/2 \in [0, \pi/2]$. However, when viewed as dimensional reduction of the Haydys-Witten equations as explained in Section 2.5, the parameter θ obtains a geometric interpretation as angle between the non-vanishing vector field v and the direction of invariance. This motivates the slightly non-standard choice of normalization used throughout this thesis. The naturality of θ, as opposed to θ_{GU}, can also be seen from equations (2.4) and (2.5), where our normalization avoids an additional factor of two.

At the midpoint $\theta = \pi/2$, the equations are usually referred to as 'the' Kapustin-Witten equations and written more succinctly as

$$
\begin{aligned}
F_A - \tfrac{1}{2}[\phi\wedge\phi] - \star_4 d_A\phi &= 0 \\
d_A^{\star_4}\phi &= 0
\end{aligned}
$$

In Witten's approach to Khovanov homology these are considered on $X^3\times\mathbb{R}^+_y$ and their solutions are generators of the Morse-Smale-Witten complex of BPS states.

Moduli spaces of Kapustin-Witten solutions play a fundamental role in Haydys-Witten theory. In their original article, Kapustin and Witten realized that on closed manifolds the study of solutions is relatively simple.

Theorem 2.3 ([KW07; GU12]). *Let $E \to W^4$ be an $SU(2)$ principal bundle over a compact manifold without boundary. Assume (A, ϕ) satisfies the θ-Kapustin-Witten equations with $\theta \in (0, \pi)$. If $E \to W^4$ has non-zero Pontryagin number then A and ϕ are identically zero. Otherwise $A + i\phi$ is a flat $PSL(2, \mathbb{C})$ connection; equivalently $F_A = [\phi \wedge \phi]$ and $\nabla^A \phi = 0$.*

As a consequence, one mostly concentrates on open manifolds or manifolds with boundary. Recently, solutions of Kapustin-Witten equations have been studied on $\mathbb{R}^4$ and it turns out that also in this case solutions are also remarkably constrained.

Theorem 2.4 (Taubes' dichotomy [Tau17a]). *Let $W^4 = \mathbb{R}^4$, $G = SU(2)$, and define $\kappa^2(r) := r^{-3} \int_{\partial B_r} \|\phi\|^2$. Assume that (A, ϕ) is a solution of the Kapustin-Witten equations, then either there is an $a > 0$ such that $\liminf_{r \to \infty} \kappa / r^a > 0$ or $[\phi \wedge \phi] = 0$.*

In a similar way, the Kapustin-Witten energy

$$E_{\mathrm{KW}} = \int_{W^4} \left(\|F_A\|^2 + \left\|\nabla^A \phi\right\|^2 + \|[\phi \wedge \phi]\|^2 \right)$$

offers some degree of control over the asymptotic behaviour of the Higgs field ϕ on manifolds with controlled asymptotic volume growth (ALX spaces).

Theorem 2.5 (Kapustin-Witten solutions on ALE and ALF spaces [NO21, Main Theorem 1]). *Assume W^4 is an ALE or ALF gravitational instanton and (A, ϕ) is a finite energy solution of the θ-Kapustin-Witten equations, $\theta \neq 0 \pmod{\pi}$. Then ϕ has bounded L^2-norm.*

Since bounded L^2-norm in particular implies bounded L^2-average on spheres, it follows that for finite energy solutions on $\mathbb{R}^4$ and $S^1 \times \mathbb{R}^3$ the function $\kappa \to 0$. In combination with Taubes' dichotomy this yields the following corollary.

Corollary 2.6. *On $\mathbb{R}^4$, $\mathbb{R}^3 \times S^1$ and compact manifolds, solutions of the θ-Kapustin-Witten equations with finite positive energy are such that A is flat, $\nabla^A \phi = 0$, and $[\phi \wedge \phi] = 0$.*

In Chapter 3, we prove a generalization of Taubes' dichotomy to any ALE or ALF space. It then follows that Corollary 2.6 actually applies to ALE, ALF, and compact manifolds.

2.4.3 The Vafa-Witten Equations

The Vafa-Witten equations are partial differential equations on a four-manifold W^4. They were first discovered by Vafa and Witten in the context of 4d $\mathcal{N} = 2$ super Yang-Mills theory [VW94]. Their solutions give rise to topological invariants of four-manifolds. The equations have since been subject to close scrutiny both in physics and mathematics, see [Mar10; Tan17; Tau17b; Tan19; GSY22; OT22] and references therein.

Let (W^4, g) be a smooth Riemannian four-manifold and G a compact Lie group. Consider a principal G-bundle $E \to W^4$ together with a principal connection A, and denote by $\operatorname{ad} E$ the adjoint bundle associated to the Lie algebra $\mathfrak{g}$ of G. Let $B \in \Omega^2_+(W^4)$ be a self-dual two-form with respect to the Hodge star operator and $C \in \Omega^0(W^4, \operatorname{ad} E)$ a section of $\operatorname{ad} E$. In complete analogy to the situation in Haydys' setup, there is a three-dimensional cross product $\sigma(\cdot,\cdot) = \frac{1}{2}(\cdot \times \cdot) \otimes [\cdot,\cdot]_{\mathfrak{g}}$ on $\Omega^2_+(W^4, \operatorname{ad} E)$.

The Vafa-Witten equations for the triple (A, B, C) are the following partial differential equations.

Definition 2.7 (Vafa-Witten Equations)**.**

$$\begin{aligned} F_A^+ - \sigma(B,B) - \tfrac{1}{2}[C,B] &= 0 \\ d_A^{\star_4} B + d_A C &= 0 \end{aligned} \tag{2.6}$$

We denote the associated differential operator by $\mathbf{VW}(A, B, C)$.

The Vafa-Witten equations are closely related to the Kapustin-Witten equations equations. For one, we will later describe in detail that both arise as dimensional reduction of the Haydys-Witten equations. Furthermore, it is well known that on Euclidean space the $\theta = 0$ version of the Kapustin-Witten equations and the Vafa-Witten equations are equivalent. The correspondence arises by combining the components of $B = \sum_{i=1}^3 \phi_i(dx^0 \wedge dx^i + \frac{1}{2}\epsilon_{ijk} dx^j \wedge dx^k)$ and C into a one-form $\phi = C dx^0 + \phi_i dx^i$.

2.4.4 The Extended Bogomolny Equations

The extended Bogomolny equations (EBE) are a set of partial differential equations on a three-manifold X^3. As the name suggests, they are an extension of Bogomolny's magnetic monopole equations on three-manifolds [Bog76; JT80; AH88]. The latter are a dimensional reduction of the self-dual Yang–Mills equations from four dimensions to three dimensions. From a purely three-dimensional point of view, the EBE can be viewed as complexification of the Bogomolny equations [NO19a].

In the context of Haydys-Witten theory, the EBE – in fact a one-parameter family called twisted extended Bogomolny equations (TEBE) – appear as dimensional reduction of the Haydys-Witten (or Kapustin-Witten) equations to three dimensions. They play an important role in the definition of the Nahm pole boundary conditions in the presence of knots, because a 't Hooft operator, upon reduction to three dimensions, reduces to the insertion of a monopole-like solution of the EBE.

Let (X^3, g) be a Riemannian three-manifold and $E \to X^3$ a G-principal bundle. Consider a gauge connection A, a one-form $\phi \in \Omega^1(X^3, \operatorname{ad} E)$, and two functions $c_1, c_2 \in \Omega^0(X^3, \operatorname{ad} E)$. The extended Bogomolny equations are the following set of differential equations for the tuple (A, ϕ, c_1, c_2):

Definition 2.8 (Extended Bogomolny Equations).

$$\begin{aligned} F_A - \tfrac{1}{2}[\phi \wedge \phi] + \star_3 (d_A c_1 - [c_2, \phi]) &= 0 \\ d_A c_2 + [c_1, \phi] - \star_3 d_A \phi &= 0 \\ d_A^{\star_3} \phi - [c_1, c_2] &= 0 \end{aligned} \tag{2.7}$$

We write **EBE** (A, ϕ, c_1, c_2) for the associated differential operator.

As a side remark, note that the equations admit several specializations, each of which had important reverberations in contemporary mathematics.

- If $\phi = c_2 = 0$ these are the Bogomolny equations for a magnetic monopole (A, c_1) on a three-manifold X^3.
- If $X^3 = \Sigma \times \mathbb{R}_y$, $c_1 = c_2 = 0$, and (A, ϕ) are independent of y, the equations reduce to Hitchin's equations for a Higgs bundle over Σ [Hit87a]. Note that solutions to the y-invariant equations provide natural stationary boundary conditions at non-compact ends.
- If $X^3 = \Sigma \times \mathbb{R}_y$, $A = c_2 = 0$, and ϕ is independent of Σ, they reduce to Nahm's equations [Nah83]. These will be described in more detail in the upcoming Section 2.4.5.

Since the extended Bogomolny equations can be obtained by a dimensional reduction of the $\theta = \pi/2$ version of the Kapustin-Witten equations (cf. Section 2.5), it's clear that there should also be a one-parameter family of EBEs, defined by dimensional reduction for any value of $\theta \in [0, \pi]$. The result is known as θ-twisted extended Bogomolny equations (TEBE) [Dim22a].

Definition 2.9 (Twisted Extended Bogomolny Equations).

$$\begin{aligned} F_A - \tfrac{1}{2}[\phi \wedge \phi] + \cot\theta \, d_A \phi + \csc\theta \star_3 (d_A c_1 - [c_2, \phi]) &= 0 \\ d_A c_2 + [c_1, \phi] + \cot\theta \, (d_A c_1 - [c_2, \phi]) - \csc\theta \star_3 d_A \phi &= 0 \\ d_A^{\star_3} \phi - [c_1, c_2] &= 0 \end{aligned} \tag{2.8}$$

We write $\textbf{TEBE}_\theta (A, \phi, c_1, c_2)$ for the associated differential operator.

Observe that for $\theta = \pi/2$, one obtains the untwisted EBE of equation (2.7).

Let us also mention the work of Nagy and Oliveira in [NO19a; NO19b], where the TEBE are investigated at $\theta = 0$ and $\pi/2$, respectively. From their point of view, the two sets of equations arise from two ways to extend the Hodge star operator to the complexification of the principal bundle. From our point of view, these are special points of a one-parameter family of extended Bogomolny equations, and their findings are analogues in three dimensions of the interpretation of the Kapustin-Witten equations as shifted anti-self dual equations (2.5) of a complex connection in four dimensions, as explained by Gagliardo and Uhlenbeck [GU12] and summarized in Section 2.4.2.

Sometimes the EBE are defined only for a triple (A, ϕ, c_1). The reason for this is the following result.

Proposition 2.10. *Assume (A, ϕ, c_1, c_2) is an irreducible solution of the extended Bogomolny equations on a Riemannian three-manifold X^3. If X^3 is closed, or if it has ends at which the fields satisfy boundary conditions such that $\int_{X^3} d\,\mathrm{Tr}\,(c_1 \wedge d_A c_2 - [c_1, c_2] \wedge \star_3 \phi) = 0$, then $c_2 = 0$.*

Proof. The proof was originally outlined in [Wit11a, p. 58]. The following, more explicit presentation of the argument is taken from [HM19c].

If we set

$$\begin{aligned} I_1 &= \int_{X^3} \left\| F_A - \tfrac{1}{2}[\phi \wedge \phi] + \star_3 d_A c_1 \right\|^2 + \left\| \star_3 d_A \phi - [c_1, \phi] \right\|^2 + \left\| d_A^{\star_3} \phi \right\|^2 , \\ I_2 &= \int_{X^3} \left\| [c_2, \phi] \right\|^2 + \left\| d_A c_2 \right\|^2 + \left\| [c_1, c_2] \right\|^2 , \\ I_3 &= \int_{X^3} d\,\mathrm{Tr}\,(c_1 \wedge d_A c_2 - [c_1, c_2] \wedge \star_3 \phi) , \end{aligned}$$

then there is a Weitzenböck formula

$$\int_{X^3} \left\| \mathbf{EBE}\ (A, \phi, c_1, c_2) \right\|^2 = I_1 + I_2 + I_3 .$$

By assumption $I_3 = 0$, either because $\partial M = \emptyset$ or because the boundary conditions on (c_1, c_2) are exactly such that I_3 vanishes. The remaining terms in the Weitzenböck formula are non-negative, so any solution of the extended Bogomolny equations also satisfies $I_1 = I_2 = 0$. If $c_2 \neq 0$, vanishing of the terms in I_2 implies that the pair (A, ϕ) is reducible, which is in contradiction to our assumption and concludes the proof. □

An astoundingly important fact about the EBE is that over three-manifolds of the form $X^3 = \Sigma \times \mathbb{R}$, for some Riemann surface Σ, they admit a Hermitian Yang-Mills structure [Wit11a]. To see this one introduces three differential operators

$$\mathcal{D}_1 = \nabla_1^A + i\nabla_2^A , \qquad \mathcal{D}_2 = [\phi_1, \cdot] + i[\phi_2, \cdot] , \qquad \mathcal{D}_3 = \nabla_3^A + i[\phi_3, \cdot] .$$

The extended Bogomolny equations are then equivalent to

$$[\mathcal{D}_i, \mathcal{D}_j] = 0 , \qquad \sum_{i=1}^{3} [\ \overline{\mathcal{D}_i}\ , \mathcal{D}_i] = 0 .$$

These equations are equivalent to the Hermitian Yang-Mills equation for a Hermitian connection on a holomorphic vector bundle over $\Sigma \times \mathbb{R}$.

More generally, the θ-TEBE exhibit a Hermitian Yang-Mills structure if $c_2 - \tan\beta/3\ c_1 = 0$, where β denotes the complementary angle $\beta = \pi/2 - \theta$ [GW12]. Note that for $\theta = \pi/2$ this condition is automatically satisfied for all irreducible solutions due to Proposition 2.10.

The Hermitian Yang-Mills structure is an important tool in the classificatio of solutions of the EBE and TEBE. It has been used to prove a Kobayashi-Hitchin correspondence between solutions of the EBE and Higgs bundles with certain extra structure [GW12; HM19c; HM20]. First analogous results regarding the classification of TEBE-solutions have been achieved in [HM19b; Dim22a]. We will discuss this in much more detail in Chapter 5, where we utilize a similar Hermitian Yang-Mills structure to determine certain solutions of the Haydys-Witten and Kapustin-Witten equations, respectively.

2.4.5 Nahm's equations

Nahm's equations are ordinary, non-linear differential equations for a vector bundle of $\mathfrak{g}$-valued functions over a one-dimensional interval I. The equations rely on the existence of a cross product on the vector bundle. As is well known, real vector spaces with a (bilinear) cross product are in correspondence with the imaginary part of normed division algebras. Correspondingly, the associated equations are referred to as complex, quaternionic, and octonionic Nahm equations, depending on the rank of the vector bundle.

The equations were originally introduced in the quaternionic case by Nahm and play an important role in the classification of monopoles [Nah83; Hit83; Don84]. They can be viewed as dimensional reductions either from anti-self dual Yang-Mills equations on four-manifolds or equivalently from Bogomolny's monopole equations on three-manifolds.

The octonionic Nahm equations first appear in [GT93]. They have recently attracted renewed attention in the context of Spin(7)-instantons on eight-dimensional manifolds and monopoles on seven-dimensional manifolds with G_2 holonomy [Che15; CN22]. Also see the recent article by He [He20] for a discussion of the moduli space of solutions for the octonionic Nahm equations.

In the context of Haydys-Witten theory, a dimensional reduction of the Haydys-Witten equations from five to one dimension in general gives rise to a twisted version of the octonionic Nahm equations. Since the Haydys-Witten equations can be viewed as a lift of self-dual Yang-Mills equations to five dimensions, it's not too surprising that a variant of Nahm's equations appears. The twisted octonionic Nahm equations play an intermittent but important role in the definition of the Nahm pole boundary conditions.

Below, we first provide the general definition of Nahm's equations and state explicit formulae for each case, mostly following a similar exposition in [He20]. Subsequently, the twisted octonionic equations are introduced and we explain the underlying structure from the point of view of octonionic multiplication.

Let G be a compact Lie group with Lie algebra $\mathfrak{g}$ and I a real interval with coordinate y. Consider a trivial G-principal bundle E over I with connection $A = A_y dy$ and denote the associated covariant derivative by $\nabla_y^A = \frac{d}{dy} + [A_y, \cdot]$. Let $\mathbb{V}$ be one of the normed division algebras $\mathbb{C}, \mathbb{H}, \mathbb{O}$.

Write $(\operatorname{Im}\mathbb{V}, \times)$ for its imaginary part together with the cross product induced by multiplication.

We consider the trivial bundle $\operatorname{Im}\mathbb{V} \otimes \operatorname{ad} E \to I$. Let the tuple $\mathbf{X} = (X_1, \dots, X_k)$ denote a section of this bundle, where $k = \dim_{\mathbb{R}} \mathbb{V} - 1$. In particular, the components X_i are $\mathfrak{g}$-valued functions over I. This bundle admits a cross product, induced by the cross product on $\operatorname{Im}\mathbb{V}$ and the Lie bracket on $\mathfrak{g}$ and given by

$$\mathbf{X} \times \mathbf{X} = \sum_{i,j,k} f_{ijk}\, e_i \otimes [X_j, X_k] \ ,$$

where f_{ijk} are the structure constants of $\mathbb{V}$.

Let E be a G-principal bundle over an interval I with gauge connection A and denote a section $\mathbf{X} \in \Gamma(I, \operatorname{Im}\mathbb{V} \otimes \mathfrak{g})$. The Nahm equations associated to $\mathbb{V}$ are the following system of non-linear, ordinary differential equations for the pair $(A, \mathbf{X})$

Definition 2.11 (Nahm Equations)**.**

$$\nabla_y^A \mathbf{X} + \mathbf{X} \times \mathbf{X} = 0$$

We will occasionally write $\mathbf{Nahm}_{\mathbb{V}}(A, \mathbf{X})$ for the associated differential operator.

Remark. The division algebras only appear fiberwise, in the multiplication of sections. In particular, the underlying differential equations are based in real analysis, as opposed to complex, quaternion, or octonion analysis. Furthermore, Nahm's equations only make use of pairwise products, such that the non-associativity of octonions does not play a role.

For $\mathbb{V} = \mathbb{C}$ the section $\mathbf{X}$ corresponds to a single $\operatorname{ad} E$-valued function X_1, while the cross product on $\operatorname{Im}\mathbb{C}$ is the zero map. Hence, the complex Nahm equation is the single equation

$$\nabla_y^A X_1 = 0 \ ,$$

which is just the statement that X_1 is covariantly constant along I.

For $\mathbb{V} = \mathbb{H}$ the section consists of three components $\mathbf{X} = (X_1, X_2, X_3)$ and the structure constants are the completely anti-symmetric tensor ϵ_{ijk}, for $i, j, k \in \{1, 2, 3\}$, which is 1 when $ijk = 123$.

$$\nabla_y^A X_i + \frac{1}{2}\epsilon_{ijk}[X_j, X_k] = 0$$

The quaternionic Nahm equations are typically simply known as 'the' Nahm equations. An important set of solutions are Nahm pole solutions $(A, \mathbf{X}) = (0, X_1, X_2, X_3)$, with $X_i = \frac{\mathfrak{t}_i}{y}$, where $\mathfrak{t}_i \in \mathfrak{g}$ satisfy $\mathfrak{su}(2)$ commutation relations.

For $\mathbb{V} = \mathbb{O}$ the section has seven components $\mathbf{X} = (X_1, \dots, X_7)$. A possible choice of structure constants[5] for the octonions is given by the completely antisymmetric tensor f_{ijk}, $i, j, k \in$

$\{1,\dots 7\}$, that is $+1$ when ijk is any of 123, 145, 176, 246, 257, 347, or 365. Explicitly, the octonionic Nahm equations are given by

$$
\begin{aligned}
\nabla_y^A X_1 + [X_2,X_3] + [X_4,X_5] - [X_6,X_7] &= 0\\
\nabla_y^A X_2 - [X_1,X_3] + [X_4,X_6] + [X_5,X_7] &= 0\\
\nabla_y^A X_3 + [X_1,X_2] + [X_4,X_7] - [X_5,X_6] &= 0\\
\nabla_y^A X_4 - [X_1,X_5] - [X_2,X_6] - [X_3,X_7] &= 0\\
\nabla_y^A X_5 + [X_1,X_4] - [X_2,X_7] + [X_3,X_6] &= 0\\
\nabla_y^A X_6 + [X_1,X_7] + [X_2,X_4] - [X_3,X_5] &= 0\\
\nabla_y^A X_7 - [X_1,X_6] + [X_2,X_5] + [X_3,X_4] &= 0
\end{aligned}
$$

Note that solutions of the quaternionic Nahm equations give rise to solutions of the octonionic Nahm equations: If (A,X_1,X_2,X_3) is a solution of the quaternionic equations, then $(A,X_1,X_2,X_3,0,0,0,0)$ is a solution of the octonionic equations. For more aboute the moduli space of the octonionic Nahm equations see [He20].

We will also come across a twisted version of the octonionic equations. To explain this, it is convenient to first rename components as $\mathbf{X} = (X_1,X_2,X_3,Y,Z_1,Z_2,Z_3)$. We can then express the octonionic Nahm equations in terms of ϵ_{ijk} with $i,j,k \in \{1,2,3\}$ and $\epsilon_{123} = 1$ as

$$
\begin{array}{lclcl}
\nabla_y^A X_i & + & [Y,Z_i] & + & \epsilon_{ijk}\left(\frac{1}{2}[X_j,X_k] - \frac{1}{2}[Z_j,Z_k]\right) & = 0\\
\nabla_y^A Y & - & [X_i,Z_i] & & & = 0\\
\nabla_y^A Z_i & - & [Y,X_i] & - & \epsilon_{ijk}[X_j,Z_k] & = 0
\end{array}
$$

In this notation the twisted equations arise from mixing the terms that appear in the last column by a rotation of angle $\beta \in [0,\pi]$, as follows:

$$
\begin{array}{lclcl}
\nabla_y^A X_i & + & [Y,Z_i] & + & \epsilon_{ijk}\Big(\ \ \cos\beta\left(\frac{1}{2}[X_j,X_k] - \frac{1}{2}[Z_j,Z_k]\right) + \sin\beta\,[X_j,Z_k]\Big) & = 0\\
\nabla_y^A Y & - & [X_i,Z_i] & & & = 0\\
\nabla_y^A Z_i & - & [Y,X_i] & - & \epsilon_{ijk}\Big(-\sin\beta\left(\frac{1}{2}[X_j,X_k] - \frac{1}{2}[Z_j,Z_k]\right) + \cos\beta\,[X_j,Z_k]\Big) & = 0
\end{array}
$$

The twisted equations may be viewed as the result of deforming the cross product on $\mathrm{Im}\,\mathbb{O}$. To explain this, first recall that by the Cayley-Dickson construction octonions can be viewed as a product of the quaternions equipped with a particular multiplication. Explicitly, if we denote the basis of the quaternions by $(1,i,j,k)$, we can identify the basis elements of $\mathbb{O}$ with $1=(1,0)$, $e_1=(i,0)$, $e_2=(j,0)$, $e_3=(k,0)$, $h=(0,1)$, $f_1=(0,i)$, $f_2=(0,j)$, $f_3=(0,k)$. As real vector spaces, the imaginary octonions can then be identified with the direct sum $\mathrm{Im}\,\mathbb{O} = \mathrm{Im}\,\mathbb{H} \oplus \mathbb{R} \oplus \mathrm{Im}\,\mathbb{H}$. Observe that this corresponds precisely to the previous renaming of components $\mathbf{X} = (\vec{X},Y,\vec{Z})$. Octonionic multiplication can be summarized as follows: 1 is the unit element,

[5] We use a slightly different convention than [He20]; the two choices differ by $e_7 \mapsto -e_7$.

for any other basis element set $(x_i)^2 = -1$, and for the remaining imaginary products specify $x \times y := \frac{1}{2}(xy - yx) = xy$ to be given by $e_i \times f_i = -h$ and when $i \neq j$ by

$$e_i \times e_j = \epsilon_{ijk}\, e_k \qquad f_i \times f_j = -\epsilon_{ijk}\, e_k \qquad e_i \times f_j = -\epsilon_{ijk}\, f_k \qquad h \times e_i = -f_i \qquad h \times f_i = e_i$$

Clearly, octonionic multiplication does not preserve the decomposition as real vector spaces. In particular, h maps one of the $\operatorname{Im}\mathbb{H}$ factors into the other. This is used to deform the cross product by a rotation between the two factors of $\operatorname{Im}\mathbb{H}$. More precisely, whenever the product has values in one of the $\operatorname{Im}\mathbb{H}$'s, we post-compose it with the left-action of $\cos\beta\, 1 + \sin\beta\, h$, which corresponds to adjusting the following multiplications.

$$\begin{aligned}
e_i \times_\beta e_j &= (\cos\beta\, 1 + \sin\beta\, h)\,(e_i \times e_j) = \epsilon_{ijk}\,(\cos\beta\, e_k - \sin\beta\, f_k) \\
f_i \times_\beta f_j &= (\cos\beta\, 1 + \sin\beta\, h)\,(f_i \times f_j) = -\epsilon_{ijk}\,(\cos\beta\, e_k - \sin\beta\, f_k) \\
e_i \times_\beta f_j &= (\cos\beta\ 1 + \sin\beta\ h)\,(e_i \times f_j) = -\epsilon_{ijk}\,(\sin\beta\, e_k + \cos\beta\ f_k)
\end{aligned}$$

Here $i, j, k \in \{1, 2, 3\}$ and we sum over repeated indices. All other products remain unchanged.

With this deformation the twisted version of the octonionic Nahm equations can be succinctly defined by the following equations.

Definition 2.12 (Twisted Octonionic Nahm Equations)**.**

$$\nabla^A_y \mathbf{X} + \mathbf{X} \times_\beta \mathbf{X} = 0$$

More explicitly, writing $\mathbf{X} = (X_1, X_2, X_3, Y, Z_1, Z_2, Z_3)$ as above, the β-twisted octonionic equations are

$$\begin{aligned}
\nabla^A_y X_i + [Y, Z_i] + \epsilon_{ijk}\Big(\ \ \cos\beta\big(\tfrac{1}{2}[X_j, X_k] - \tfrac{1}{2}[Z_j, Z_k]\big) + \sin\beta[X_j, Z_k]\Big) &= 0 \\
\nabla^A_y Z_i - [Y, X_i] - \epsilon_{ijk}\Big(-\sin\beta\big(\tfrac{1}{2}[X_j, X_k] - \tfrac{1}{2}[Z_j, Z_k]\big) + \cos\beta[X_j, Z_k]\Big) &= 0 \\
\nabla^A_y Y - [X_i, Z_i] \quad &= 0
\end{aligned} \tag{2.9}$$

We will occasionally denote the associated differential operator by $\mathbf{Nahm}_{\mathbb{O},\beta}(A, \mathbf{X})$.

The embedding of quaternionic solutions into the moduli space of octonionic solutions carries over to the twisted case by rotating $\vec{X}$ into $\vec{Z}$: If $(A, \vec{X}) = (A, X_1, X_2, X_3)$ is a solution of the quaternionic Nahm equations, then

$$(A,\ \cos\beta\,\vec{X},\ 0,\ \sin\beta\,\vec{X}^\tau) := (A,\ \cos\beta\, X_1,\ \cos\beta\, X_2,\ \cos\beta\, X_3,\ 0,\ \sin\beta\, X_1,\ \sin\beta\, X_3,\ \sin\beta\, X_2)$$

is a solution of the β-twisted octonionic Nahm equations. Note that the identification of $\vec{X}$ with $\vec{Z}$ involves an anti-cyclic permutation of components, denoted by $\tau = (132)$.

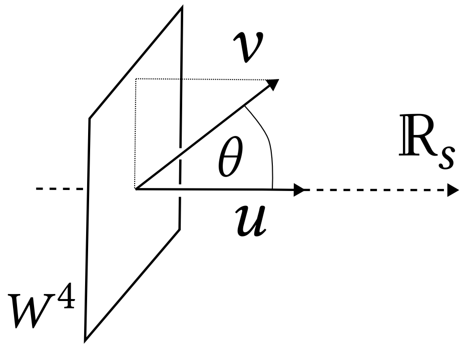

Figure 2.2

2.5 Dimensional Reductions of the Haydys-Witten Equations

As mentioned before, each of the equations presented in Section 2.4 can be viewed as a dimensional reduction of the Haydys-Witten equations. Here we explicitly perform the reduction step and explain how the various one-parameter families arise naturally from Haydys' five-dimensional geometry.

Throughout this section, we denote the Haydys-Witten fields by $(\hat{A}, \hat{B})$ to explicitly distinguish them from four-dimensional fields A, B and ϕ and three-dimensional fields $\tilde{A}$ and $\tilde{\phi}$. For convenience let us repeat the Haydys-Witten equations, as defined in Equation 2.2.

$$\begin{aligned} F_{\hat{A}}^{+} - \sigma(\hat{B}, \hat{B}) - \nabla_v^{\hat{A}} \hat{B} &= 0 \\ \iota_v F_{\hat{A}} - \delta_{\hat{A}}^{+} \hat{B} &= 0 \end{aligned} \tag{2.10}$$

Dimensional reduction on $M^5 = \mathbb{R}^k \times Y^{5-k}$ assumes that the fields $(\hat{A}, \hat{B})$ and all gauge transformations are independent of the position in $\mathbb{R}^k$. Equivalently, if we write u_i, $i = 1, \dots, k$, for a set of orthonormal coordinate vector fields on $\mathbb{R}^k$, then $\hat{A}$ and $\hat{B}$ are invariant under the action of all u_i's. This only makes sense if the angles $g(u_i, v) = \cos\theta_i$ are constant, since otherwise already the equations depend inherently on the position in $\mathbb{R}^k$. Below, we discuss dimensional reduction for $k = 1, 2$ and 4, which leads to the θ-Kapustin-Witten equations, θ-twisted extended Bogomolny equations, and β-twisted octonionic Nahm equations, respectively.

2.5.1 $\mathbb{R}$-invariant Solutions

Consider a product space $M^5 = \mathbb{R}_s \times W^4$ equipped with a product metric and denote the inclusion of W^4 at $s = 0$ by $i : W^4 \hookrightarrow \mathbb{R}_s \times W^4$. Let $u := \partial_s$ be the unit vector field tangent to $\mathbb{R}_s$ and assume that $g(u, v) = \cos\theta$ is constant. The angle θ can either be thought of as the glancing angle between v and the direction of invariance $\mathbb{R}_s$, or equivalently as incidence angle of v on the hypersurface W^4, see Figure 2.2.

As explained in Section 2.4.1, Haydys' geometric setup provides a lift of four-dimensional (anti-)self-dual two-forms into five dimensions. Loosely speaking, this was achieved by identifying the orthogonal complement of v with the tangent space of four-manifolds W^4 that foliate M^5. When u and v are aligned, dimensional reduction simply recovers the self-dual two forms. However, in general u and v are not necessarily aligned and the consequences of imposing u-invariance depend on the interplay between the orthogonal complements of u and v. To that end observe that u and v define a distribution $\Delta_{(u,v)} \subset TM$. Since u and v are non-vanishing and θ is constant, $\Delta_{(u,v)}$ is regular, i.e. a vector bundle. We now have to distinguish two cases: $\theta \equiv 0 \pmod{\pi}$ and $\theta \neq 0$.

If $\theta \equiv 0 \pmod{\pi}$ or equivalently if $u = \pm v$, the vector bundle $\Delta_{(u,v)}$ has rank one. In this case the orthogonal complements of u and v are identical, perhaps up to a reversal of orientation, and dimensional reduction is fairly straightforward. The gauge connection splits into $\hat{A} = \hat{A}_s ds + A$, where $A = i^*\hat{A}$ is a connection over W^4. The component $\hat{A}_s$ can be reinterpreted as an $\operatorname{ad} E$-valued function $C \in \Omega^0(\operatorname{ad} E)$, since gauge transformations are assumed to be u-invariant[6].

Regarding $\hat{B}$, recall from Section 2.4.1 that there is an isomorphism $i^*\big(\Omega^2_{\partial_s,+}(\mathbb{R}_s \times W^4)\big) \simeq \Omega^2_+(W^4)$. As a result $(\hat{A},\hat{B})$ pull back to a triple of fields (A,B,C) on the four-manifold W^4, where A is a gauge connection, B is a self-dual two-form, and C is an $\operatorname{ad} E$-valued function. With these identifications in place, note that $\nabla^{\hat{A}}_v \hat{B} = [C \wedge B]$, $\iota_v F_{\hat{A}} = d_A C$, and $\delta^+_{\hat{A}} \hat{B} = d_A^{\star_4} B$. Plugging this into the Haydys-Witten equations for $(\hat{A},\hat{B})$ immediately yields the Vafa-Witten equations (2.6) for the triple (A,B,C).

If $\theta \not\equiv 0 \pmod{\pi}$, the vector bundle $\Delta_{(u,v)}$ has rank two. The tangent bundle splits into orthogonal complements $TM = \Delta_{(u,v)} \oplus \Delta^\perp_{(u,v)}$ and this induces a similar decomposition for one-forms as sections of $(\Delta_{(u,v)})^* \oplus (\Delta^\perp_{(u,v)})^*$. Moreover, since $\Delta_{(u,v)}$ admits the two linearly independent global sections u and v, it is trivial. It follows that there is a non-vanishing vector field w that together with u provides an orthonormal basis of $\Delta_{(u,v)}$. In this basis v is given as $v = \cos\theta u + \sin\theta w$ and we can define $v^\perp := -\sin\theta u + \cos\theta w$, which is the unique (up to a sign) unit vector field in $\Delta_{(u,v)}$ that is orthogonal to v.

Crucially, contraction with $v^\perp$ provides an isomorphism between u-invariant self-dual two-forms and sections of $(\Delta^\perp_{(u,v)})^*$. One way to see this is to observe that locally the following two-forms provide a basis of $\Omega^2_{v,+}$ (cf. Section 2.4.1):

$$e_i = \eta^\perp \wedge dx^i + \tfrac{1}{2}\epsilon_{ijk} dx^j \wedge dx^k \quad , \; i = 1,2,3 \; ,$$

Here $\eta^\perp$ is the (global) one-form dual to $v^\perp$ and dx^i are (local) sections of $(\Delta^\perp_{(u,v)})^*$. If we write $\hat{B} = \sum_i \phi_i e_i$, contraction with $v^\perp$ yields a one-form $\iota_{v^\perp}\hat{B} = \sum_i \phi_i dx^i$. Using this isomorphism, the u-invariant fields $(\hat{A},\hat{B})$ on $\mathbb{R}_s \times W^4$ can be reinterpreted as a gauge connection A and an $\operatorname{ad} E$-valued one-form ϕ on W^4.

[6] In general $\iota_u\hat{A}$ transforms as $\iota_u\hat{A} \mapsto g^{-1}\iota_u\hat{A}g + g^{-1}\nabla^{\hat{A}}_u g$ under gauge transformations. If the gauge transformation g is u-invariant the second term vanishes and $\iota_u\hat{A}$ is equivalent to an $\operatorname{ad} E$-valued function.

To make this explicit, consider for the moment the example of Euclidean space $\mathbb{R}_s \times \mathbb{R}^4$ with Cartesian coordinates $(s, t, x^a)_{a=1,2,3}$, chosen in such a way that

$$u = \partial_s \ , \ w = \partial_t \ , \ \text{and} \ v = \cos\theta\, \partial_s + \sin\theta\, \partial_t \ .$$

The gauge connection can be split into $\hat{A} = \hat{A}_s ds + A$, where A is the part of the connection on $\mathbb{R}^4$. Furthermore, we can combine the remaining component $\hat{A}_s$ with the 3 components ϕ_i of B into a one-form $\phi = \hat{A}_s dt + \sum_i \phi_i dx^i$ on $\mathbb{R}^4$.

The definitions in the Euclidean case are the local model underlying the following identifications for general manifolds $\mathbb{R}_s \times W^4$:

$$A := i^* \hat{A} \ , \qquad \phi := \hat{A}_s \wedge w^\flat + \iota_{v^\perp} \hat{B} \ .$$

Remark. To shed some more light on the definition of ϕ, it might be helpful to directly compare the situations for $\theta = 0$ and $\theta \neq 0$. In both cases, the pullback of $\hat{A}$ provides a gauge connection on the pullback bundle $i^*E \to W^4$ and projects out the component $\hat{A}_s$, which on its own can be viewed as an $\operatorname{ad} E$-valued function. The difference arises in the reinterpretation of $\hat{B}$. If $\theta \neq 0$ and we pull back $\hat{B}$ to a two-form on W^4, as we do when $\theta = 0$, we would on the one hand project out any components that annihilate u and on the other hand wouldn't obtain a generic (self-dual) two-form on W^4. Instead, we consider the contraction $\iota_{v^\perp}\hat{B}$ as a section of $(\Delta^\perp_{(u,v)})^*$, which contains neither $u^\flat$- nor $w^\flat$-components. The absence of terms proportional to $u^\flat$ $(= ds)$ ensures that the pullback is injective, i.e. does not project out any components of $\iota_{v^\perp}\hat{B}$. Using $\hat{A}_s$ as the missing $w^\flat$-component we then obtain a generic one-form ϕ on W^4.

To determine the reduction of the Haydys-Witten equations in terms of (A, ϕ) on W^4, it is sufficient to investigate the differential equations (2.10) in arbitrarily small neighbourhoods of a point x. Hence, choose normal coordinates $(s, t, x^i)_{i=1,2,3}$ at x such that $u = \partial_s$, $w = \partial_t$, and $v = \cos\theta\partial_s + \sin\theta\partial_t$. Due to u-invariance and after setting $\hat{A}_s = \phi_t$, we find

$$\nabla^{\hat{A}}_{v^\perp} = -\sin\theta[\phi_t, \cdot] + \cos\theta\nabla^A_t \ , \ (F_{\hat{A}})_{s\mu} = -\nabla^A_\mu \phi_t \ .$$

The second of the Haydys-Witten equations (2.10) thus becomes

$$\begin{aligned}
0 &= \iota_v F_{\hat{A}} - \delta^+_{\hat{A}} \hat{B} \\
&= (\cos\theta\, \iota_{\partial_s} + \sin\theta\, \iota_{\partial_t}) F_{\hat{A}} + (\nabla^{\hat{A}}_{v^\perp} \iota_{v^\perp} + \sum \nabla^{\hat{A}}_i \iota_{\partial_i}) \hat{B} \\
&= \left(-\nabla^A_t \phi_t - \sum_{i=1}^3 \nabla^A_i \phi_i \right) \eta^\perp \\
&\quad + \sum_{(ijk)} \Big(\sin\theta\, (F_{ti} - [\phi_t, \phi_i]) + \cos\theta\, (\nabla^A_t \phi_i - \nabla^A_i \phi_t) - (\nabla^A_j \phi_k + \nabla^A_k \phi_j) \Big)\, dx^i \ ,
\end{aligned}$$

where the sum in the last line is over cyclic permutations of (123). The $\eta^\perp$-component of this equation states $d^{\star_4}_A \phi = 0$, which is the θ-independent part of the Kapustin-Witten equations (2.3). The dx^i-components imply vanishing of the ti-components of the Kapustin-Witten equations as given in equation (2.4).

For the evaluation of the first of the Haydys-Witten equations (2.10) we expand

$$F^+_{\hat{A}} = \sum_{(ijk)} (-\sin\theta F_{si} + \cos\theta F_{ti} + F_{jk})(\eta^\perp \wedge dx^i + dx^j \wedge dx^k)$$

and similarly for $\sigma(\hat{B},\hat{B}) = \sum_{(ijk)}[\phi_j,\phi_k](\eta^\perp \wedge dx^i + dx^j \wedge dx^k)$. The equations then become

$$\begin{aligned} 0 &= F^+_{\hat{A}} - \sigma(\hat{B},\hat{B}) - \nabla^{\hat{A}}_v \hat{B} \\ &= \sum_{(ijk)} (\sin\theta\, \nabla^A_i \phi_t + \cos\theta\, F_{ti} + F_{jk} - [\phi_j,\phi_k] - \cos\theta\, [\phi_t,\phi_i] - \sin\theta \nabla^A_t \phi_i)(\eta^\perp \wedge dx^i + dx^j \wedge dx^k) \end{aligned}$$

This implies that also the ij-components of the Kapustin-Witten equations (2.4) vanish. To see this multiply by $\sin\theta$ and use the dx^i-component of the earlier equation to replace $\sin\theta(F_{ti} - [\phi_t,\phi_i])$.

In summary, the key result of this section is the following statement.

Proposition 2.13. *Let $M^5 = \mathbb{R} \times W^4$ equipped with a product metric and a non-vanishing unit vector field v. Write u for the unit vector field along $\mathbb{R}$ and assume $g(u,v) = \cos\theta$ is constant. Let $(\hat{A},\hat{B})$ be Haydys-Witten fields on M^5, write $A = i^*\hat{A}$ for the pullback connection on W^4, and depending on the value of θ define fields on W^4 as follows*

$$\begin{aligned} \theta = 0 :&\quad B = i^*\hat{B}\,, \quad C = \hat{A}_s \\ \theta \neq 0 :&\quad \phi = \hat{A}_s w^\flat + \iota_{v^\perp}\hat{B} \end{aligned}$$

Then u-invariant Haydys-Witten equations for $(\hat{A},\hat{B})$ are either equivalent to the Vafa-Witten equations for (A,B,C) if $\theta \equiv 0 \pmod{\pi}$, or to the θ-Kapustin-Witten equations for (A,ϕ) otherwise.

$$\mathbf{HW}_v\left(\hat{A},\hat{B}\right) \overset{\mathbb{R}\text{-}inv.}{\rightsquigarrow} \begin{cases} \mathbf{VW}(A,B,C) & \theta \equiv 0 \pmod{\pi}, \\ \mathbf{KW}_\theta(A,\phi) & \textit{else} \end{cases}$$

Let us stress that dimensional reduction is not continuous at $\theta = 0$. For general four-manifolds, the Vafa-Witten equations and $\theta = 0$ version of the Kapustin-Witten equations are not equivalent. From the perspective of the Haydys-Witten equations, it should be expected that there is a possibly non-trivial relation between solutions of these equations whenever the four-manifold admits a non-vanishing vector field. This is well-known for $W^4 = \mathbb{R}^4$, where the Vafa-Witten and $\theta = 0$ Kapustin-Witten equations are equivalent by identifying $\phi = C dx^0 + B_{0i} dx^i$.

The dimensional reductions of the Haydys-Witten equations have previously been carried out for the cases $\theta = 0$ and $\pi/2$ by Witten [Wit11a], and independently for $\theta = 0$ and in slightly more generality by Haydys [Hay15]. Concretely, Witten considered the case where the five-manifold is $M^5 = \mathbb{R}_s \times X^3 \times \mathbb{R}^+_y$ and $v = \partial_y$. On the one hand he investigates $\mathbb{R}_s$-invariant solutions, i.e. dimensional reduction with respect to $u = \partial_s$. In this situation the glancing angle is $\theta = \pi/2$ (or $3\pi/2$; the two cases differ only by a reversal of orientation). Witten explains that setting $\phi = \hat{A}_s dy + B_{si} dx^i$, the u-invariant Haydys-Witten equations are equivalent to the $\theta = \frac{\pi}{2}$ version of the Kapustin-Witten equations. This is in accordance with the general definition $\phi = \iota_u \hat{A} \wedge w^\flat + \iota_{v^\perp} \hat{B}$, because in the current situation $u = \partial_s$, $w^\flat = dy$, and $v^\perp = -\partial_s$. On the other hand Witten also shortly inspected $\mathbb{R}^+_y$-invariant solutions, which provide non-trivial boundary conditions at $y \to \infty$. Since in that case $u = \partial_y$ coincides with v, the glancing angle is $\theta = 0$ and the equations reduce to the Vafa-Witten equations. This was observed more generally by Haydys for $\mathbb{R}_y$-invariant solutions on $M^5 = \mathbb{R}_y \times W^4$ with $v = \partial_y$ [Hay15, sec. 4.1].

Let us remark that Witten explains in great detail that the full Haydys-Witten equations on $\mathbb{R}_s \times X^3 \times \mathbb{R}^+_y$ represent antigradient flow equations (with respect to a conveniently chosen functional) that interpolate between $\theta = \pi/2$ Kapustin-Witten solutions at $s \to \pm\infty$. Haydys similarly explains that the full Haydys-Witten equations on $\mathbb{R}_s \times W^4$ represent antigradient flow equations that interpolate between Vafa-Witten solutions at $s \to \pm\infty$.

Both of these statements can be viewed as special case of the more general fact that on $\mathbb{R}_s \times W^4$ the equations $\mathbf{HW}_v(\hat{A}, \hat{B}) = 0$ take the form of flow equations that interpolate between $\mathbb{R}_s$-invariant solutions at $s \to \pm\infty$. In particular, for $\theta \neq 0, \pi/2$, the Haydys-Witten equations on $\mathbb{R}_s \times W^4$ are equivalent to the following equations.

$$
\begin{aligned}
\nabla^A_s A &= -\iota_{\partial_s}\left(d_A\phi - \star_4\left(\cos\theta\,(F_A - [\phi\wedge\phi]) + \sin\theta\, d_A\phi\right)\right)\\
\nabla^A_s \phi &= -\iota_{\partial_s}\left(F_A - \star_4\left(-\sin\theta\,(F_A - [\phi\wedge\phi]) + \cos\theta\, d_A\phi\right)\right)\\
d^{\star_4}_A \phi &= 0
\end{aligned}
$$

2.5.2 $\mathbb{R}^2$-invariant Solutions

Consider a product space $M^5 = \mathbb{R}^2 \times X^3$ equipped with a product metric and denote by $i : X^3 \hookrightarrow \mathbb{R}^2 \times X^3$ inclusion at the origin of $\mathbb{R}^2$. Let s_1 and s_2 be Cartesian coordinates on $\mathbb{R}^2$ and u_1, u_2 the associated coordinate vector fields. Assume that $g(u_1, v) = \cos\theta_1$ and $g(u_2, v) = \cos\theta_2$ are constant. In fact we are free to choose coordinates in such a way that u_2 and v are orthogonal, fixing $\theta_2 = \pi/2$, see Figure 2.3. The dimensional reduction only depends on the remaining parameter $\theta := \theta_1$, which is the glancing angle between v and $\mathbb{R}^2$.

Instead of imposing $\mathbb{R}^2$-invariance from scratch, we may proceed by iteration: First utilize the results of Section 2.5.1 to determine the dimensional reduction along one of the directions and afterwards additionally demand invariance in the second direction. Due to the results of the previous section it's clear that we need to distinguish between the cases $\theta \equiv 0$ and $\theta \not\equiv 0$ (mod π).

In the case where $\theta = 0$ (the case $\theta = \pi$ follows from this by a reversal of orientation), we start by imposing u_2-invariance. By Proposition 2.13 this leads to a pair (A, ϕ) on $\mathbb{R}_{s_1} \times X^3$ that satisfies the $\theta_2 = \pi/2$ version of the Kapustin-Witten equations. The Higgs field is given by $\phi = \hat{A}_{s_2} ds_1 - \iota_{u_2}\hat{B}$, while A is the pullback connection.

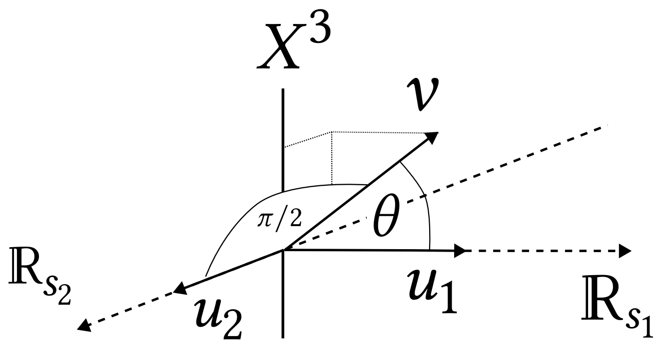

Figure 2.3

In the case where $\theta \neq 0$, it is more convenient to first consider u_1-invariance, which leads to the θ-Kapustin-Witten equations for a (different) pair (A, ϕ). The Higgs field is given by $\phi = \hat{A}_{s_1} w^\flat + \iota_{v^\perp} \hat{B}$, where w and $v^\perp$ are the sections of $\Delta_{(u_1,v)}$ that were introduced in Section 2.5.1.

Still for $\theta \neq 0$, it is helpful to have a closer look at the structure of $\Omega^2_{v,+}(\mathbb{R}^2 \times X^3)$, to simplify the subsequent dimensional reduction along u_2. Observe that w is a non-vanishing vector field that is orthogonal to both u_1 and u_2, i.e. it is the pushforward of a non-vanishing vector field on X^3. Let us generalize the notation from the previous section and denote by $\Delta_{(u_1,u_2,v)}$ the regular distribution spanned by the three vector fields u_1, u_2, and v. Since $\theta \neq 0$, this distribution is a trivial subbundle of TM^5 of rank 3 that admits an orthonormal basis of sections (u_1, u_2, w). The orthogonal complement $\Delta^\perp_{(u_1,u_2,v)}$ is a rank 2 subbundle of TX^3 orthogonal to w. It follows that the tangent space TX^3 splits into a trivial one-dimensional part, spanned by w, and a two-dimensional part that we denote $\Delta^\perp_{(u_1,u_2,v)}$. This provides a splitting of $\Omega^1(\mathbb{R}_{s_2} \times X^3)$ into sections of $C^\infty(M^5) ds_2 \oplus C^\infty(M^5) w^\flat \oplus (\Delta^\perp_{(u_1,u_2,v)})^*$, which in turn induces the existence of a global section $e_1 := \frac{1}{2}(1 + T_\eta)\,(\eta^\perp \wedge ds_2)$ of $\Omega^2_{v,+}$. As a consequence $\hat{B}$ splits globally into

$$\hat{B} = \phi_1 e_1 + \varphi \,.$$

In the usual local basis, the two-form φ is given by $\varphi = \phi_2 e_2 + \phi_3 e_3$. With respect to this expression of $\hat{B}$ and the splitting of $\Omega^1(\mathbb{R}_{s_2} \times X^3)$, the Higgs field is defined as $\phi = \phi_1 ds_2 + \hat{A}_{s_1} w^\flat + \iota_{v^\perp} \varphi$. In local coordinates $(s_2, x^i)_{i=1,2,3}$ of $\mathbb{R}_{s_2} \times X^3$ this reduces to $\phi = \phi_1 ds_2 + \hat{A}_{s_1} dx^1 + \phi_2 dx^2 + \phi_3 dx^3$.

We are now ready to perform dimensional reduction along the second direction. In either case, denote the remaining direction of invariance by s, i.e. $s = s_1$ if $\theta = 0$ and $s = s_2$ if $\theta \neq 0$. We are in the situation of either the $\pi/2$- or θ-version of the Kapustin-Witten equations for (A, ϕ) on $\mathbb{R}_s \times X^3$. Imposing invariance in the second direction thus corresponds to a fairly straightforward dimensional reduction of the Kapustin-Witten equations: Regardless of the five-dimensional origin of ϕ and A they split on $\mathbb{R}_s \times X^3$ as $\phi = c_1 ds + \tilde{\phi}$ and $A = c_2 ds + \tilde{A}$. As usual, we view the ds components c_1 and c_2 as $\operatorname{ad} E$-valued functions over X^3. Choose an orientation, say $ds \wedge \mu_{X^3}$, and determine the individual terms in the Kapustin-Witten equations (2.4) in terms of $(\tilde{A}, \tilde{\phi}, c_1, c_2)$, whilst dropping any derivatives in the direction of s:

$$\begin{aligned}
F_A &= & d_{\tilde{A}} c_2 \wedge ds &\quad + & F_{\tilde{A}} \\
\tfrac{1}{2}[\phi \wedge \phi] &= & -[c_1, \tilde{\phi}] \wedge ds &\quad + & \tfrac{1}{2}[\tilde{\phi} \wedge \tilde{\phi}] \\
d_A \phi &= & -([c_2, \tilde{\phi}] - d_{\tilde{A}} c_1) \wedge ds &\quad + & d_{\tilde{A}} \tilde{\phi} \\
\star_4 d_A \phi &= & \star_3 d_{\tilde{A}} \tilde{\phi} \wedge ds &\quad + & \star_3([c_2, \tilde{\phi}] - d_{\tilde{A}} c_1)
\end{aligned}$$

Plugging these expressions into (2.4) yields the first two lines of the TEBE (2.8). The third equation of the TEBE follows from the remaining constraint $0 = d_A^{\star_4} \phi = -[c_1, c_2] + d_{\tilde{A}}^{\star_3} \tilde{\phi}$.

Proposition 2.14. *Let $M^5 = \mathbb{R}^2 \times X^3$ equipped with a product metric and a non-vanishing unit vector field v. Write u_i, $i = 1, 2$, for Cartesian vector fields on $\mathbb{R}^2$, chosen such that $g(u_1, v) = \cos\theta$ and $g(u_2, v) = 0$, and assume both angles are constant. Let $(\hat{A}, \hat{B})$ be Haydys-Witten fields on M^5, write $\tilde{A} = i^* \hat{A}$ for the pullback connection on X^3, and depending on the value of θ define fields on X^3 as follows.*

$$\begin{aligned}
\theta = 0 : &\quad \tilde{\phi} = -\iota_{u_2} \hat{B}\,, &\quad c_1 = \hat{A}_{s_2}\,, &\quad c_2 = \hat{A}_{s_1} \\
\theta \neq 0 : &\quad \tilde{\phi} = \hat{A}_{s_1} w^\flat + \iota_{v^\perp} \varphi\,, &\quad c_1 = \phi_1, &\quad c_2 = \hat{A}_{s_2}
\end{aligned}$$

Then (u_1, u_2)-invariant Haydys-Witten equations for $(\hat{A}, \hat{B})$ are equivalent to the EBE if $\theta = 0$ or to the θ-twisted extended Bogomolny equations for $(\tilde{A}, \tilde{\phi}, c_1, c_2)$.

$$\mathbf{HW}_v(\hat{A}, \hat{B}) \overset{\mathbb{R}^2\text{-}inv.}{\rightsquigarrow} \begin{cases} \mathbf{EBE}\ (\tilde{A}, \tilde{\phi}, c_1, c_2) & \theta \equiv 0 \pmod{\pi}\,, \\ \mathbf{TEBE}_\theta(\tilde{A}, \tilde{\phi}, c_1, c_2) & else \end{cases}$$

The dimensional reduction of the Haydys-Witten equations to three-manifolds inherits the discontinuity at $\theta = 0$ that is already present in the reduction to four-manifolds. In particular, in the limit $\theta \to 0$ dimensional reduction does not lead to the $\theta = 0$ version of the TEBE, but instead to the 'untwisted' $\pi/2$ version. As before, this behaviour is encoded in the rank of the regular distribution $\Delta_{(u_1, u_2, v)}$ spanned by u_1, u_2 and v. If $\theta \equiv 0 \pmod{\pi}$, i.e. if v is orthogonal to X^3, $\Delta_{(u_1, u_2, v)}$ is of rank two and dimensional reduction leads to the (untwisted) EBE. If $\theta \not\equiv 0 \pmod{\pi}$, the distribution $\Delta_{(u_1, u_2, v)}$ has rank three, there exists a non-vanishing vector field w on X^3, and dimensional reduction produces θ-TEBE.

Another special situation arises when v is parallel to X^3, since then dimensional reduction leads to the $\pi/2$-TEBE, which we recall are just the (untwisted) EBE. However, the existence of the vector field w provides additional structure that allows us to continuously deform the $\pi/2$-TEBE to generic θ-TEBE by rotating $v \mapsto v = \cos\theta u_1 + \sin\theta w$. Such a continuous deformation does not exist if the EBE arise from a dimensional reduction for which $\Delta_{(u_1, u_2, v)}$ is of rank 2. In that case any small deformation of $v \mapsto v + \epsilon w$, lifting v off the plane spanned by u_1 and u_2, leads to a discontinuous jump from $\pi/2$-EBE to θ_ϵ-TEBE, where θ_ϵ is the corresponding (small) glancing angle.

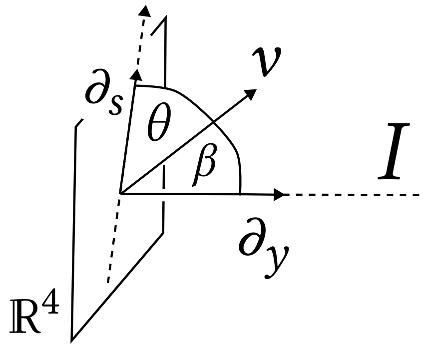

Figure 2.4

The idea of deforming (or twisting) the EBE away from $\theta = \pi/2$ is used to great effect by Gaiotto and Witten in [GW12]. In their setup, they consider the dimensional reduction of the Haydys-Witten equations from

$$M^5 = \mathbb{R}_s \times \mathbb{R}_t \times \Sigma \times \mathbb{R}^+_y \to \Sigma \times \mathbb{R}^+_y = X^3 .$$

The three vector fields of interest are $u_1 = \partial_s$, $u_2 = \partial_t$ and $w = \partial_y$. When $v = \partial_y$, dimensional reduction results in the EBE. However, we are in the situation where $\Delta_{(u_1,u_2,v)}$ is of rank three and we can deform the EBE equations away from $\theta = \pi/2$, e.g. by considering $v = \cos\theta u_1 + \sin\theta w$.

2.5.3 $\mathbb{R}^4$-invariant Solutions

Consider now the case $M^5 = \mathbb{R}^4 \times I$, where I is a real interval, and let ∂_y be the coordinate vector field along I. Assume Haydys' preferred vector field v is such that the incidence angle between v and $\mathbb{R}^4$ – determined by $g(v, \partial_y) = \cos\beta$ – is constant. On $\mathbb{R}^4$ fix Cartesian coordinates (s, x^i), $i = 1, 2, 3$, where ∂_s is the vector field that satisfies $g(v, \partial_s) = \sin\beta$, while $g(v, \partial_i) = 0$. In these coordinates $v = \sin\beta\partial_s + \cos\beta\partial_y$, see Figure 2.4. Notice that in the current situation β is the incidence angle between v and the hyperplane of invariant directions. This is in contrast to the preceding discussions, where it was more convenient to use the glancing angle $\theta = \pi/2 - \beta$.

Write $\hat{A} = A_s ds + \sum_i A_i dx^i + A_y dy$ and $\hat{B} = \sum_i \phi_i e_i$ with $e_i = \frac{1}{2}(1 + T_\eta)(\eta^\perp \wedge dx^i)$. Under the assumption that $\hat{A}$ and $\hat{B}$ are independent of $\mathbb{R}^4$, the components A_s, A_i and ϕ_i, $i = 1, 2, 3$, become a collection of seven $\mathfrak{g}$-valued functions on I, while $A = A_y dy$ provides a gauge connection over I. As a result dimensional reduction of the Haydys-Witten equations is relatively straightforward. The i-th component (with respect to the basis $\{e_1, e_2, e_3\}$ of $\Omega^2_{v,+}$) of the Haydys-Witten equations (2.2) is given by

$$\begin{aligned} 0 &= \left(F^+_{\hat{A}} - \sigma(\hat{B}, \hat{B}) - \nabla^{\hat{A}}_v \phi\right)_i \\ &= \left(\cos\beta F_{si} - \sin\beta F_{yi} + \tfrac{1}{2}\epsilon_{ijk}F_{jk}\right) - \tfrac{1}{2}\epsilon_{ijk}[\phi_j, \phi_k] - \cos\beta \nabla^A_y \phi_i - \sin\beta \nabla^A_s \phi_i \\ &= -\sin\beta(\nabla^A_y A_i + [A_s, \phi_i]) - \cos\beta(\nabla^A_y \phi_i - [A_s, A_i]) - \tfrac{1}{2}\epsilon_{ijk}([\phi_j, \phi_k] - [A_j, A_k]) \end{aligned}$$

Meanwhile, the second of the Haydys-Witten equations becomes

$$
\begin{aligned}
0 &= \iota_v F_{\hat{A}} - \delta^+_{\hat{A}} \hat{B} \\
&= (\cos\beta \iota_{\partial_y} + \sin\beta \iota_{\partial_s}) F_{\hat{A}} + (\nabla^{\hat{A}}_s \iota_s + \nabla^{\hat{A}}_y \iota_y + \sum \nabla^{\hat{A}}_i \iota_i) B \\
&= \left(\nabla^A_y A_s + [A_i, \phi_i]\right) \eta^\perp + \left(\cos\beta \left(\nabla^A_y A_i + [A_s, \phi_i]\right) - \sin\beta \left(\nabla^A_y \phi_i - [A_s, A_i]\right) - \epsilon_{ijk}[A_j, \phi_k]\right) dx^i
\end{aligned}
$$

The component proportional to $\eta^\perp$ is exactly the third of the twisted octonionic Nahm equations (2.9), while the remaining equations are just a linear combination of the first two lines of these equations. To see this, multiply the i-th components of the two equations with $\sin\beta$ and $\cos\beta$, respectively, and add them up (and vice versa with subsequent subtraction). More explicitly, the reduced Haydys-Witten equations are thus rearranged to read

$$
\begin{aligned}
\nabla^A_y \phi_i - [A_s, A_i] + \epsilon_{ijk} \Big(\quad \cos\beta \, \tfrac{1}{2}([\phi_j, \phi_k] - [A_j, A_k]) + \sin\beta [\phi_j, A_k] \Big) &= 0 \\
\nabla^A_y A_i + [A_s, \phi_i] - \epsilon_{ijk} \Big(-\sin\beta \, \tfrac{1}{2}([\phi_j, \phi_k] - [A_j, A_k]) + \cos\beta [\phi_j, A_k] \Big) &= 0 \\
\nabla^A_y A_s + [A_i, \phi_i] &= 0
\end{aligned}
$$

These are exactly the β-twisted octonionic Nahm equations for $\mathbf{X} = (\phi_1, \phi_2, \phi_3, -A_s, A_1, A_2, A_3)$.

Proposition 2.15. *Let $M^5 = \mathbb{R}^4 \times I$, where I is a connected one-dimensional manifold, equipped with a product metric and a preferred non-vanishing unit vector field v. Write u for the unit vector field on I. Assume $g(u, v) = \cos\beta$ is constant and $(\hat{A}, \hat{B})$ are invariant under translations in $\mathbb{R}^4$ and arranged into the pair $(A, \mathbf{X})$ as specified above. Then the Haydys-Witten equations reduce to the β-twisted octonionic Nahm equations.*

$$
\mathbf{HW}_v(\hat{A}, \hat{B}) \overset{\mathbb{R}^4-inv.}{\rightsquigarrow} \mathbf{Nahm}_{\mathbb{O},\beta}(A, \mathbf{X})
$$

2.6 The Nahm Pole Boundary Condition

The Nahm pole boundary conditions with knot singularities play a fundamental role in the relation between Haydys-Witten theory and Khovanov homology. They prescribe an asymptotic equivalence of the fields (A, B) with a certain set of singular model solutions near the boundary. Which model solutions to use depends on whether one is in the vicinity of a knot or not.

Recall from Section 2.3 that in the five-dimensional setting a knot is supported on a two-dimensional surface Σ_K inside the four-dimensional boundary of M^5; one direction is parallel to the longitude of the original, one-dimensional knot K, while the other direction arises from extending K along the additional flow direction of Floer theory. Typically Σ_K arises in this way from either a compact knot or a collection of infinitely extended strands embedded in X^3, so for all intents and purposes Σ_K is either an embedding of $\mathbb{R} \times S^1$ or a disjoint union of $\mathbb{R}^2$'s. Nevertheless, Nahm pole boundary conditions are defined for general embedded surfaces, including link cobordisms and knotted surfaces.

The singular structure of the model solutions comes in two flavours. First, denoting by y a boundary-defining function, the fields diverge as y^{-1}. Second, near the position of a knot they display a monopole-like behaviour associated to a singularity inside of ad E.

A careful analysis of the Nahm pole boundary conditions with knot singularities is described in great detail in [MW14; MW17]. These articles make use of Melrose's b-calculus [Mel90]; or rather a variant of it that was introduced by Mazzeo [Maz91] and further developed by Mazzeo-Vertman [MV13]. In this context it is natural to consider the geometric blowup of M^5 along Σ_K. This is the manifold with corners, denoted by $[M^5; \Sigma_K]$, whose underlying set of points is the disjoint union of M^5 and the inward-pointing unit normal bundle of Σ_K. Many analytic properties of differential operators and their solutions become more apparent when viewed as conormal distributions on the blowup. As long as an explicit distinction between the original manifold and its blowup is irrelevant, we'll simply denote the blowup by M^5, the original boundary component after removing Σ_K by $\partial_0 M^5$, and the newly introduced boundary at Σ_K by $\partial_K M^5$.

The definition of Nahm pole boundary conditions features an aspect that arises naturally in the context of the blowup $[M^5; \Sigma_K]$. Namely, the boundary conditions for the two types of boundary $\partial_0 M$ and $\partial_K M$ have individual descriptions. Near $\partial_0 M$ the Haydys-Witten pair (A, B) is locally modeled on maximally symmetric, $\mathbb{R}^4$-invariant Nahm-pole solutions of the Haydys-Witten equations on $\mathbb{R}^4 \times \mathbb{R}^+_y$, while near $\partial_K M$ the model solution is that of an EBE-monopole on $\mathbb{R}^2 \times \mathbb{R}^+_y$.

Below we provide descriptions of these two distinct model solutions, followed by a definition of the Nahm pole boundary conditions with knot singularities on general manifolds. We conclude with the investigation of the analytic properties of the Haydys-Witten equations with β-twisted Nahm pole boundary conditions, which has not yet appeared in the literature.

2.6.1 Model Solutions without Knot Singularity

Consider Euclidean half-space $\mathbb{R}^4 \times \mathbb{R}^+_y$ and denote Cartesian coordinates by $(s, x^i, y)_{i=1,2,3}$. Assume that $v = \sin\beta \partial_s + \cos\beta \partial_y$, where the incidence angle β between v and the boundary is constant. This geometry is invariant under translations parallel to the boundary, and accordingly, we demand that the model solutions are independent of the position in the boundary. Proposition 2.15 states that they must then be solutions of the β-deformed octonionic Nahm equations on $\mathbb{R}^+$.

Let $\rho : \mathfrak{su}(2) \to \mathfrak{g}$ be a Lie algebra homomorphism and denote by $(\mathfrak{t}_i)_{i=1,2,3}$ the image of the standard basis of $\mathfrak{su}(2)$ under ρ. Furthermore, let us fix the anti-cyclic permutation $\tau = (132)$. As mentioned in Section 2.4.5, one easily checks that the following is a solution of the β-twisted octonionic Nahm equations (2.9)

$$A_i = \frac{\sin\beta \ \mathfrak{t}_{\tau(i)}}{y} \ , \qquad \phi_i = \frac{\cos\beta \ \mathfrak{t}_i}{y} \ , \qquad A_s = A_y = 0 \ . \tag{2.11}$$

Since the fields exhibit a pole at $y = 0$ these are called (β-twisted or tilted) Nahm pole solutions. We say the Nahm pole is *regular*, if ρ is a principal embedding in the sense of Kostant, i.e. if the commutant of $\mathfrak{su}(2)$ in $\mathfrak{g}$ is a Cartan subalgebra. These solutions are the local model for the Nahm pole boundary condition near $\partial_0 M$.

If $v = \partial_y$ is orthogonal to the boundary, i.e. $\beta = 0$, the gauge field A vanishes and the model solution coincides with the standard, untwisted Nahm pole solution described in [Wit11a; MW14]. The twisted Nahm pole model solutions have previously appeared in the context of supersymmetric boundary conditions in [GW09b, sec. 4] and their role in calculating the Jones polynomial via gauge theory was described in [GW12].

2.6.2 Model Solutions with Knot Singularity

Consider, again, Euclidean half-space $\mathbb{R}^4 \times \mathbb{R}^+_y$, but now assume that we additionally include a 't Hooft operator supported on a single, infinitely extended 'strand'. More precisely, in five-dimensions this corresponds to the inclusion of a distinguished two-dimensional plane $\Sigma_K = \mathbb{R}^2$ in the boundary of $\mathbb{R}^4 \times \mathbb{R}^+$. Denote Cartesian coordinates (s, t, x^2, x^3, y), where Σ_K extends along the (s, t)-plane. For simplicity assume that $v = \cos\theta\partial_s + \sin\theta\partial_y$. To make contact with notation in Section 2.5.2: the orthonormal coordinate vector fields parallel to Σ_K coincide with $u_1 = \partial_s$ and $u_2 = \partial_t$, while the unit normal vector $w = \partial_y$ plays the role of a distinguished global vector field on the remaining three-manifold $X^3 = \mathbb{R}^2_{x^2,x^3} \times \mathbb{R}^+_y$. For reasons that will become clear momentarily, we only consider $\theta \not\equiv 0 \pmod{\pi}$.

As a first step we demand that the model solutions are invariant with respect to translations along Σ_K. Due to Proposition 2.14, this means that the relevant model is a solution of the θ-TEBE (2.8) for three-dimensional fields $(\tilde{A}, \tilde{\phi}, c_1, c_2)$ on $\mathbb{R}^2_{x^2,x^3} \times \mathbb{R}^+_y$. Since $\theta \neq 0$ we are in the situation where the Haydys-Witten fields are expressed as $A = A_s ds + A_t dt + \tilde{A}$ and $B = \phi_1 e_1 + \varphi$, where the two-form $\varphi = \phi_2 e_2 + \phi_3 e_3$ is such that $\iota_{v^\perp}\varphi = \phi_2 dx^2 + \phi_3 dx^3$. Disentangling the general definitions of Proposition 2.14 in this way specifies the three-dimensional fields in terms of the components of (A, B) as follows:

$$\tilde{A} = A_2 dx^2 + A_3 dx^3 + A_y dy \ , \qquad \tilde{\phi} = \phi_2 dx^2 + \phi_3 dx^3 + A_s dy \ , \qquad c_1 = \phi_1, \qquad c_2 = A_t \ .$$

For the case $\theta = \pi/2$ and $G = SU(2)$ the relevant model solutions of the $\pi/2$-TEBE (which are simply the untwisted EBE) were described by Witten [Wit11a]. Introduce (hemi-)spherical coordinates (R, ψ, ϑ) on $\mathbb{R}^2_{x^2,x^3} \times \mathbb{R}^+_y \simeq [0, \infty)_R \times H^2_{\vartheta,\psi}$, where $R \in [0, \infty)$, $\psi \in [0, \pi/2]$ and $\vartheta \in [0, 2\pi]$ are given by

$$R^2 = x_2^2 + x_3^2 + y^2 \ , \qquad \cos\psi = \frac{y}{R} \ , \qquad \cos\vartheta = \frac{x_2}{x_2^2 + x_3^2} \ .$$

Let $(\mathfrak{t}_i)_{i=1,2,3}$ denote a standard basis of $\mathfrak{su}(2)$ and view $\mathfrak{t}_1$ as the generator of a fixed Cartan subalgebra. Introduce, by abuse of notation, the $\mathfrak{sl}(2, \mathbb{C})$-valued function $\varphi = \phi_2 - i\phi_3$ that conveniently combines the components ϕ_2 and ϕ_3 of the two form φ of the same name. Similarly, denote by $E = \mathfrak{t}_2 - i\mathfrak{t}_3$, $H = \mathfrak{t}_1$, and $F = \mathfrak{t}_2 + i\mathfrak{t}_3$ the elements of an $\mathfrak{sl}(2,\mathbb{C})$-triple (E,H,F). Finally, express the components of the three-dimensional gauge connection in terms of spherical coordinates $\tilde{A} = A_R\, dR + A_\psi\, d\psi + A_\vartheta\, d\vartheta$. The knot singularity solutions of the EBE with charge $\lambda \in \mathbb{Z}$ in terms of the components of (A,B) are given by the following expressions.

$$
\begin{aligned}
A_\vartheta &= -(\lambda+1)\cos^2\psi \frac{(1+\cos\psi)^\lambda - (1-\cos\psi)^\lambda}{(1+\cos\psi)^{\lambda+1} - (1-\cos\psi)^{\lambda+1}}\, H \\
\phi_1 &= -\frac{\lambda+1}{R}\frac{(1+\cos\psi)^{\lambda+1} + (1-\cos\psi)^{\lambda+1}}{(1+\cos\psi)^{\lambda+1} - (1-\cos\psi)^{\lambda+1}}\, H \\
\varphi &= \frac{(\lambda+1)}{R}\frac{\sin^\lambda\psi\ \exp(i\lambda\vartheta)}{(1+\cos\psi)^{\lambda+1} - (1-\cos\psi)^{\lambda+1}}\, E \\
A_s &= A_t = A_R = A_\psi = 0
\end{aligned}
$$

These solutions exhibit a singular behaviour in several distinct ways. First, the components of B diverge with rate R^{-1} as $R \to 0$. Second, whenever $R \neq 0$ the solution is asymptotically equivalent to the (untwisted) Nahm pole solution as we approach the original boundary component $\psi \to \pi/2$. This compatibility will be relevant in the definition of the boundary conditions on general manifolds. Third, and most importantly, the solutions exhibit a monopole-like singularity. This is characterized by a nontrivial monodromy of the connection as one moves around the origin in the (x_2,x_3)-plane. The monodromy is supported by a non-trivial behaviour of φ, which picks up an extra factor[7] of $e^{2\pi i\lambda}$ when ϑ increases by 2π, and vanishes on the half-line $\psi = 0$ that sits over the origin in the (x^2,x^3)-plane. An insightful way to view the behaviour of φ is to observe that it takes values in the nilpotent cone $\mathcal{N} \subset \mathfrak{g}_\mathbb{C} = \mathfrak{sl}(2,\mathbb{C})$ and that it approaches the cone singularity of $\mathcal{N}$ for $\psi \to 0$.

Analogous solutions for the more general case $\theta = \pi/2$ and $G = SU(N)$ have been constructed by Mikhaylov in [Mik12]. In this case the solution is labeled by an element of the co-character lattice $\lambda \in \Gamma^\vee_{\mathrm{ch}} = \mathrm{Hom}(\mathbb{C}^\times, G_\mathbb{C})$, or equivalently, by a representation of the Langlands dual group $G^\vee_\mathbb{C}$. From the physics perspective λ corresponds to a choice of magnetic charge. In this generalization the (untwisted) Nahm pole divergence of order R^{-1} remains unchanged. However, the singular behaviour within the nilpotent cone $\mathcal{N} \subset \mathfrak{sl}(N,\mathbb{C})$ has a richer structure, since the nilpotent cone has various singularities, see for example [CM93]. The co-character λ determines which of these singularities φ approaches as $\psi \to 0$.

With regard to a generalization by twisting, Gaiotto and Witten describe in [GW12] that it is sometimes beneficial to consider the θ-TEBE for $\theta \neq \pi/2$. They predicted that there are analogous knot singularity solutions also in these cases. This has recently been confirmed for $G = SU(2)$ by Dimakis [Dim22a], who utilized a continuity argument to prove the existence of knot singularity solutions for the θ-TEBE for any $\theta \in (0,\pi)$. The deformation of θ away from $\pi/2$ to $\pi/2 - \beta$ has the effect that the Nahm pole divergence of order R^{-1} that appears in B is rotated into A, in very much the same way as is the case for the twisted Nahm pole solution

[7] φ remains single-valued, since for $G = SU(N)$ only integer values appear, while for $G = SO(N)$ half-integer values may appear, but then the Lie algebra is really $\mathfrak{psl}(2,\mathbb{C})$, where multiplication by -1 is modded out.

in (2.11). As we have seen in Proposition 2.14, the dimensional reduction of the Haydys-Witten equations for $\theta = 0$ is generally not continuously connected to the $\theta \neq 0$ reductions, such that continuity methods break down at $\theta = 0$. It is not currently known if there are knot singularity models for $\theta \equiv 0 \pmod{\pi}$.

Both of these generalizations are given by less explicit descriptions than the model solutions above and the exact formulas, where available, do not provide additional insights. For our purposes it will suffice to assume that model solutions exist for any $\theta \in (0, \pi)$ and $G = SU(N)$, and are labeled by a magnetic charge $\lambda \in \Gamma^{\vee}_{\text{char}}$.

2.6.3 From Model Solutions to Boundary Conditions

Let M^5 be a Riemannian manifold with a single boundary component, together with a preferred non-vanishing unit vector field v that approaches the boundary at a constant angle. We take this to mean that there is a tubular neighbourhood of the boundary $U = \partial M^5 \times [0, \epsilon)_y$, on which the incidence angle $\cos\beta = g(v, \partial_y)$ is constant. Furthermore, let $\Sigma_K \subseteq \partial M^5$ be an embedded surface and assume that also the glancing angle $\cos\theta = \min_{u \in T\Sigma_K} g(u, v)/\,\|u\|$ is constant. To simplify the discussion we assume that the incidence and glancing angle are related by $\theta = \pi/2 - \beta$. This is equivalent to the assumption that there is a neighbourhood of $\partial_K M$ where $v = \cos\theta u + \sin\theta \partial_y$ for some non-vanishing unit vector field $u \in T\Sigma_K$.

As explained in the introductory paragraphs of this section, we promote M^5 to the geometric blowup along Σ_K, such that there are two boundary components, $\partial_0 M$ and $\partial_K M$. In the preceding sections we have described the model solutions that shall describe the local behaviour of the fields at each of the two boundaries. A complete boundary condition requires a global specification of the fields on $\partial M = \partial_0 M \sqcup \partial_K M$. Since the model solutions diverge at the boundary, this involves additional data regarding the leading order on tubular neighbourhoods of $\partial_0 M$ and $\partial_K M$, respectively. The descriptions on these neighbourhoods must of course be compatible on intersections.

We start by fixing the global boundary data on $\partial_0 M$. Consider a tubular neighbourhood $U = \partial_0 M \times [0, \epsilon)_y$ on which $g(\partial_y, v) = \cos\beta$ is constant. The components of A and B in the Nahm pole model (2.11) are given by the same expression, up to a relative rotation with respect to the incidence angle β. For this reason it is helpful to recall that there is a relation between one-forms and self-dual two forms on U. Hence, note that whenever the incidence angle $\beta \neq 0$, v induces a non-vanishing vector field u parallel to $\partial_0 M$. This leads to a splitting of the tangent space $TU \simeq \Delta_{(u,v)} \oplus \Delta^{\perp}_{(u,v)}$, where the orthogonal complement $\Delta^{\perp}_{(u,v)}$ is a vector bundle of rank three. As in Section 2.5, contraction with $v^{\perp}$ then provides an isomorphism between $\Omega^2_{v,+}(U)$ and sections of $(\Delta^{\perp}_{(u,v)})^*$. In coordinates $(s, x^i, y)_{i=1,2,3}$ where $v = \sin\beta\partial_s + \cos\beta\partial_y$, the isomorphism identifies $\sum \phi_i e_i \mapsto \sum \phi_i dx^i$, where e_i denotes the usual basis of $\Omega^2_{v,+}(U)$. It follows that any $\operatorname{ad} E$-valued self-dual two-form over U is equivalent to an $\operatorname{ad} E$-valued one-form on a subbundle of TU:

$$\Omega^2_{v,+}(U, \operatorname{ad} E) \simeq \operatorname{Hom}(\Delta^{\perp}_{(u,v)}\ , \operatorname{ad} E)\ .$$

To keep notation at a minimum, this identification will be used implicitly in the formulas below.

Regardless of the value of β, the three-dimensional cross product on $\Omega^2_{v,+}(U)$ (cf. Section 2.4.1) provides each fiber with the Lie algebra structure of $\mathfrak{su}(2)$. Thus, at every point in the tubular neighbourhood U, any $\operatorname{ad} E$-valued self-dual two-form ϕ_ρ that satisfies $\phi_\rho - \sigma(\phi_\rho, \phi_\rho) = 0$ gives rise to a Lie algebra homomorphism $\rho : \mathfrak{su}(2) \to \mathfrak{g}$. If we denote the image of the standard basis of $\mathfrak{su}(2)$ under ρ by $(\mathfrak{t}_i)_{i=1,2,3}$, then $\phi_\rho = \sum_{i=1}^3 \mathfrak{t}_i e_i$ in the usual local basis.

Conversely, a smooth family of homomorphisms $\{\rho_p : \mathfrak{su}(2) \to \mathfrak{g}\}_{p\in U}$ determines a unique two-form $\phi_\rho \in \Omega^2_{v,+}(U, \operatorname{ad} E)$ that satisfies $\phi_\rho - \sigma(\phi_\rho, \phi_\rho) = 0$. Moreover, $\{\rho_p\}$ also induces another two form ϕ^τ_ρ, related to ϕ_ρ by a change of orientation of $\Omega^2_{v,+}(U)$ from (e_1, e_2, e_3) to (e_1, e_3, e_2). In a local basis $\phi^\tau_\rho = \sum_i \mathfrak{t}_{\tau(i)} e_i$ where τ is the anti-cyclic permutation (132) from earlier. Since $\sigma(\cdot, \cdot)$ is defined with respect to the original orientation on $\Omega^2_{v,+}(U, \operatorname{ad} E)$, ϕ^τ_ρ satisfies $\phi^\tau_\rho + \sigma(\phi^\tau_\rho, \phi^\tau_\rho) = 0$.

Hence, let $\rho : \mathfrak{su}(2) \to \operatorname{ad} E$ be a Lie algebra homomorphism and consider the Haydys-Witten fields on U that are given by

$$A^{\rho,\beta} = \frac{\sin\beta\ \phi^\tau_\rho}{y}\ , \qquad B^{\rho,\beta} = \frac{\cos\beta\ \phi_\rho}{y}\ .$$

Locally $(A^{\rho,\beta}, B^{\rho,\beta})$ coincide with the Nahm pole solutions, perhaps up to conjugation in $\mathfrak{g}$. The main take-away is that the boundary data at $\partial_0 M$ is fully determined by a choice of y-independent two-form $\phi_\rho \in \Omega^2_{v,+}(U, \operatorname{ad} E)$ in a tubular neighbourhood U of $\partial_0 M$.

Moving on to the boundary component $\partial_K M$, denote by $V = \partial_K M \times [0, \epsilon]_R$ a tubular neighbourhood. V is the product of Σ_K and the filled hemisphere $H^2_{\psi,\vartheta} \times [0, \epsilon]_R$, where the latter admits global coordinates (ψ, ϑ, R). We wish to impose that at leading order R^{-1} the behaviour of (A, B) is described completely by the knot singularity model solution. The only degree of freedom is the choice of $\mathfrak{su}(2)$ generators $\mathfrak{t}_i \in \mathfrak{g}$ at each point in V. However, except for the points at $\psi = 0$ or $R = 0$, every $p \in V$ is also contained in the tubular neighbourhood U of $\partial_0 M$, where ϕ_ρ already determines a triple $(\mathfrak{t}_i)_{i=1,2,3}$. Since the hemisphere at $R = 0$ corresponds to a single point on the original manifold and the line $\psi = 0$ is of codimension two, the two-form $\phi_\rho \in \Omega^2_{v,+}(U, \operatorname{ad} E)$ extends uniquely to all of V.

It remains to note that since $\theta \neq 0$, the two-form decomposes on V into $\phi_\rho = (\phi_\rho)_1 e_1 + \varphi_\rho$. This splitting provides a natural distinction between H and E in the model solutions, by identifying the $(\phi_\rho)_1$ component with the Cartan element H. Using this to replace the generators $\mathfrak{t}_i$ in the knot singularity solutions by $(\phi_\rho)_i$, i.e. pointwise by the image of the induced map $\rho : \mathfrak{su}(2) \to \mathfrak{g}$, determines a unique field configuration $(A^{\lambda,\theta}, B^{\lambda,\theta})$ of order $\mathcal{O}(R^{-1})$ on all of V.

In the more general situation with $\theta \neq \pi/2 - \beta$, the same discussion goes through with minor modifications when identifying coordinates and field components over U and V, respectively. We can now state the definition of the regular Nahm pole boundary conditions.

Definition 2.16 (Regular Nahm Pole Boundary Conditions with Knot Singularities)**.** Assume v approaches $\partial_0 M$ at a constant incidence angle $\beta\,(\neq \pi/2)$ and has glancing angle $\theta\,(\neq 0)$ with Σ_K. Let $\{\rho_p : \mathfrak{su}(2) \to \mathfrak{g}\}$ be a smooth family of principal embeddings on a tubular neighbourhood of the boundary and ϕ_ρ, ϕ_ρ^τ the associated self-dual two-forms. The Haydys-Witten pair (A, B) satisfies the regular Nahm pole boundary conditions at ∂M, with knot singularity of weight $\lambda \in \Gamma^\vee_{\text{char}}$ along Σ_K, if for some $\epsilon > 0$

(i) near $\partial_0 M$: $(A, B) = \left(A^{\rho,\beta},\, B^{\rho,\beta}\right) + \mathcal{O}\left(y^{-1+\epsilon}\right)$

(ii) near $\partial_K M$: $(A, B) = \left(A^{\lambda,\theta},\, B^{\lambda,\theta}\right) + \mathcal{O}(R^{-1+\epsilon})$

and such that the leading orders are compatible at the corner $R = y = 0$. This means that in spherical coordinates, where $y = R\cos\psi$, the expansion is of product type $(A, B) = (A^{\lambda,\theta}, B^{\lambda,\theta}) + \mathcal{O}(R^{-1+\epsilon}\cos\psi^{-1+\epsilon})$.

Remark. There is an analogous definition associated to arbitrary embeddings $\rho : \mathfrak{su}(2) \to \mathfrak{g}$. However, throughout this thesis we only consider the regular Nahm pole boundary conditions and omit a discussion of the more general case.

2.6.4 Elliptic Theory of Nahm Pole Boundary Conditions

In this section, we summarzie some relevant analytic properties of the Haydys-Witten equations. The fundamental questions include under which conditions $\mathbf{HW}_v$, acting on appropriate function spaces, is Fredholm and to analyze the regularity of solutions of $\mathbf{HW}_v(A, B) = f$. These properties are controlled by the fact that the Haydys-Witten and Kapustin-Witten operators are elliptic.

As is common for gauge theoretic equations, the Haydys-Witten and Kapustin-Witten equations on their own are not elliptic 'on the nose'. But they become elliptic after choosing a reference connection A^0 and imposing additionally that the linearization of the gauge action vanishes. Since ellipticity depends only on the principal symbol, we are free to add terms in subleading orders of derivatives. In the context of Nahm pole boundary conditions it is convenient to include, in this way, the leading order term B^{NP} that captures the Nahm pole behaviour of B. The gauge fixing equation we use is

$$d^\star_{A^0}(A - A^0) + \sigma(B^{\text{NP}}, B - B^{\text{NP}}) = 0\ . \tag{2.12}$$

From now on we always assume that the Haydys-Witten (and Kapustin-Witten) equations include this equation.

On closed manifolds, standard elliptic theory provides answers to many of the relevant analytic questions. On manifolds with boundary, however, these considerations are complicated by the choice of boundary conditions. In particular, under the assumption of Nahm pole boundary conditions with knot singularities, the associated differential operators are known to be 'depth-two incomplete iterated edge (iie) operators'. The study of such operators is part of the larger framework of geometric microlocal analysis and may be viewed as a variant of Melrose's b-calculus. We refer to [MW17, Sec. 8 & 9] for a very clear, if concise, account of the relevant ideas. Also see [Maz91; MV13] for a more detailed discussion of much of the relevant background.

The results we discuss here, as well as most of the necessary background, have previously been described in great detail in the context of the $\theta = \pi/2$ version of the Kapustin-Witten equations [MW14; MW17]. In these articles Mazzeo and Witten proved that the Nahm pole boundary conditions with knot singularities amend the Kapustin-Witten operator to an elliptic system. At the heart of the analysis lies the determination of formal rates of growth for homogeneous solutions of the linearization of the Kapustin-Witten equations. As is usually the case, ellipticity is accompanied by a regularity theorem, showing that these formal growth rates provide the building blocks for an asymptotic expansion of solutions of $\mathbf{HW}_v(A, B) = f$ near the boundaries of $[M; \Sigma_K]$.

In this context, regularity is described in terms of polyhomogeneous functions, which are defined by the existence of asymptotic expansions with respect to scale functions of the form $y^\alpha(\log y)^k$. To make this more precise, let us call $\Delta \subset \mathbb{C} \times \mathbb{N}$ an *indicial set* if it is a countable subset that is 'bounded from the left' in $\mathbb{C}$ and contains only 'finite towers' in $\mathbb{N}$. In other words, for any $\alpha_0 \in \mathbb{R}$, there are only finitely many elements $(\alpha, k) \in \Delta$ with $\operatorname{Re} \alpha \leq \alpha_0$ and if $(\alpha, k) \in \Delta$ then so is $(\alpha, k-1)$. A function f is polyhomogeneous at a submanifold $\{y = 0\}$ if there is an indicial set Δ such that $f \sim \sum_{(\alpha,k)\in\Delta} f_{\alpha,k}\, y^\alpha(\log y)^k$ as $y \to 0$, where the functions $f_{\alpha,k}$ are independent of y.

Theorem 2.17 (Elliptic Regularity, cf. [MW14, Prop. 5.9], [MW17, Thm. 9.6])**.** *The Haydys-Witten and Kapustin-Witten operators, together with β-twisted Nahm pole boundary conditions with knot singularities, are elliptic iie operators. Assume that (A, B) is a solution of the Haydys-Witten equations that satisfies the β-twisted Nahm pole boundary conditions with θ-twisted knot singularities. Denote by A^{NP} and B^{NP} the leading terms of (A, B) at the boundary, choose a reference connection $A^0 = A^{\mathrm{NP}} + \omega$ where ω is a connection on the restriction of E to the boundary, and write $A = A^{\mathrm{NP}} + \omega + a$ and $B = B^{\mathrm{NP}} + b$. Then a and b are polyhomogeneous on $[M; K]$, i.e. there are asymptotic expansions*

$$a \sim \textstyle\sum_{(\alpha,k)\in\Delta_0} a_{\alpha,k}\, y^\alpha(\log y)^k\,, \qquad b \sim \textstyle\sum_{(\alpha,k)\in\Delta_0} b_{\alpha,k}\, y^\alpha(\log y)^k \qquad (y \to 0)$$

$$a \sim \textstyle\sum_{(\beta,m)\in\Delta_K} a_{\beta,m}\, R^\beta(\log R)^m\,, \quad b \sim \textstyle\sum_{(\beta,m)\in\Delta_K} b_{\beta,m}\, R^\beta(\log R)^m \quad (R \to 0)$$

and corresponding product-type expansions near the corner $y = R = 0$, that are compatible with the fact that $y = R\cos\psi$. Moreover, the indicial sets Δ_0 and Δ_K are bounded from the left by $\alpha \geq 1$ and $\beta \geq 0$, respectively.

A complete proof for the case with $\beta = 0$ (equivalently $\theta = \pi/2$) was provided by Mazzeo and Witten in [MW14; MW17]. The proof for other values of β is completely analogous, so we refrain from reproducing the full details here. Instead, we only quote the main line of arguments and concentrate on the calculations of formal growth rates to determine the indicial sets Δ_0 and Δ_K.

Consider the Haydys-Witten operator on the stratified space $(M^5 \setminus \partial M^5) \sqcup (\partial M^5 \setminus \Sigma_K) \sqcup \Sigma_K$ as a depth 2 iie operator of order 1. This means that in a neighbourhood of each stratum its linearization takes a certain iterative form, that combines a differential operator on the stratum with an iie operator (of smaller depth) on the link. For example, in coordinates $(s, t, R, \psi, \vartheta)$ in a neighbourhood of a point on the lowest (depth 2) stratum Σ_K, the linearization must look like

$$L = (L_s\partial_s + L_t\partial_t) + L_R\partial_R + \frac{1}{R}\left(L_\psi\partial_\psi + L_\vartheta\partial_\vartheta + \frac{1}{\cos\psi}L_0\right) ,$$

where each L_i is a smooth (or polyhomogeneous) endomorphism of $\mathfrak{g}$. Note that for $\psi \to \pi/2$ (i.e. $\cos\psi \to 0$) one approaches points in the middle (depth 1) stratum $\partial M^5 \setminus \Sigma_K$. The operator that is multiplied with R^{-1} is itself an iie operator of depth 1, exemplifying the iterative nature of the definition. Linearizing the Haydys-Witten operator around $(A^{\mathrm{NP}}, B^{\mathrm{NP}})$ indeed yields an operator of that form.

The definition of full ellipticity as iie operator then involves three properties. The first property is invertibility of the 'iie symbol', which is a suitable analogue of the principal symbol of $\mathbf{HW}_v$, while the second and third property are the iterative invertibility of certain model operators, called 'normal operators', at points on strata of increasing depth.

To determine the iie symbol, one considers the non-singular operator $R\cos\psi L$ and replaces derivatives by covectors according to a rule that takes into account the depth of the associated stratum. In the specific case above, one replaces $R\cos\psi\partial_s$ and $R\cos\psi\partial_t$ by $-i\xi_1$ and $-i\xi_2$, respectively, while $\cos\psi\partial_\psi$ and $\cos\psi\partial_\vartheta$ are similarly replaced by $-i\xi_3$ and $-i\xi_4$, and subleading orders of differentiation (here L_0 of order 0) are discarded. Invertibility of the iie symbol of the Haydys-Witten equations immediately carries over from the same result for the Kapustin-Witten equations.

The two normal operators are determined as follows. A tubular neighborhood in M^5 of either of the two strata $S = \partial M^5 \setminus \Sigma_K$ or $S = \Sigma_K$ is diffeomorphic to a bundle of cones $C(Z)$, where Z is either a point or the hemisphere H^2, respectively. For any $p \in S$ the normal operator $N_p(L)$ is defined as the scale- and translation-invariant operator on $T_pS \times C(Z)$, that is induced by freezing the coefficient functions L_i to their values at the point p. In general, the properties of $N_p(L)$ may depend on $p \in S$ as a parameter, but this is fortunately not the case in our situation.

Invertibility of $N_p(L)$, acting on appropriately defined 'iterated edge Sobolev spaces' on the blowup $[M; S]$, depends on the formal rates of growth for solutions of $Lu = 0$. These rates are called the indicial roots of $N_p(L)$ and are determined by solving the condition $N_p(L)(\rho^\lambda u) = \mathcal{O}(\rho^\lambda)$ for λ, where ρ denotes a boundary defining function of the (blown-up) stratum under consideration. The key property we need to show is that, under the assumption of regular Nahm pole boundary conditions, there are no indicial roots in an interval $(\underline{\lambda}, \overline{\lambda})$ that contains -1. Once this is established, $N_p(L)$ is invertible on function spaces associated to the scale function ρ^μ for any weight $\mu \in (-1, \overline{\lambda})$. As a consequence, the Haydys-Witten operator is an elliptic iie operator and the indicial sets Δ_0 and Δ_K in the polyhomogeneous expansion of a and b are bounded from the left by $\overline{\lambda}$.

As it turns out, the relevant indicial roots $\underline{\lambda}$ and $\overline{\lambda}$ are independent of the twisting angle β. Abstractly, the reason for this is that the Haydys-Witten equations with β-twisted Nahm pole boundary conditions are locally equivalent to Haydys-Witten equations with untwisted Nahm pole boundary conditions, but where instead the corrections a and b are rotated into each other by an angle $-\beta$. In what follows, this statement is explained in full detail for the indicial equations at the depth 1 stratum $\partial M^5 \setminus \Sigma_K$. Subsequently, we also comment on the indicial roots of the depth 2 stratum Σ_K, where indicial roots are known only implicitly.

Let $p \in \partial M^5 \setminus \Sigma_K$, in which case the normal operator $N_p(L)$ acts on pairs of functions (a,b) over $T_p(\partial M^5 \setminus \Sigma_K) \times C(\{\text{pt.}\}) \simeq \mathbb{R}^4 \times \mathbb{R}^+_y$. Assume that (a,b) are $\mathbb{R}^4$-invariant and, moreover, that they are of the form $a = y^\lambda a_0$ and $b = y^\lambda b_0$ for some y-independent, $\mathfrak{g}$-valued differential forms a_0, b_0. Since the Nahm pole terms are proportional to y^{-1} (and their contributions at order y^{-2} vanish by construction), the leading order of $N_p(L)(a,b)$ is $y^{\lambda-1}$. Terms at this order arise from the action of ∂_y, as well as from commutators with A^{NP} or B^{NP}. The condition $N_p(L)(y^\lambda a_0, y^\lambda b_0) = \mathcal{O}(y^\lambda)$ then corresponds to equations that arise from setting to zero the terms proportional to $y^{\lambda-1}$. These are to be interpreted as equations for λ and determine the indicial roots.

More concretely, since the indicial equations invoke $\mathbb{R}^4$-invariance, it is clear from Proposition 2.15 that the indicial equations are equivalent to the linearization of the β-twisted octonionic Nahm equations around $(A^{\mathrm{NP}}, B^{\mathrm{NP}})$. In the current situation the leading order terms are given by the model solutions $A_i^{\mathrm{NP}} = y^{-1} \sin\beta\, \mathfrak{t}_{\tau(i)}$ and $B_i^{\mathrm{NP}} = y^{-1}\sin\beta\, \mathfrak{t}_i$. Plugging $A^{\mathrm{NP}} + y^\lambda a$ and $B^{\mathrm{NP}} + y^\lambda b$ into these equations and extracting the terms at order $y^{\lambda-1}$ leads to the following set of indicial equations

$$\begin{aligned} &\lambda b_i + \sin\beta\,[\mathfrak{t}_{\tau(i)}, a_s] - \cos\beta\,[\mathfrak{t}_i, a_y] \\ &\quad + \epsilon_{ijk}\Big(\cos^2\beta\,[\mathfrak{t}_j, b_k] + \sin\cos\beta\,[\mathfrak{t}_j, a_k] - \sin\cos\beta\,[\mathfrak{t}_{\tau(j)}, a_k] + \sin^2\beta\,[\mathfrak{t}_{\tau(j)}, b_k]\Big) = 0 \end{aligned} \tag{2.13}$$

$$\begin{aligned} &\lambda a_i - \cos\beta\,[\mathfrak{t}_i, a_s] - \sin\beta\,[\mathfrak{t}_{\tau(i)}, a_y] \\ &\quad - \epsilon_{ijk}\Big(\cos^2\beta\,[\mathfrak{t}_j, a_k] - \sin\cos\beta\,[\mathfrak{t}_j, b_k] + \sin\cos\beta\,[\mathfrak{t}_{\tau(j)}, b_k] + \sin^2\beta\,[\mathfrak{t}_{\tau(j)}, a_k]\Big) = 0 \end{aligned} \tag{2.14}$$

$$\lambda a_s - \cos\beta\,[\mathfrak{t}_i, a_i] + \sin\beta\,[\mathfrak{t}_{\tau(i)}, b_i] = 0 \tag{2.15}$$

$$\lambda a_y + \sin\beta\,[\mathfrak{t}_{\tau(i)}, a_i] + \cos\beta\,[\mathfrak{t}_i, b_i] = 0 \tag{2.16}$$

The first three equations arise from the β-twisted octonionic Nahm equations (2.9), while the last is the gauge fixing condition (2.12).

When $\beta = 0$, and thus $\theta = \pi/2$, the equations decouple into two quaternionic Nahm-like equations for $(a_y, \vec{b} = (b_1, b_2, b_3))$ and $(a_s, \vec{a} = (a_1, a_2, a_3))$, respectively. Up to a reinterpretation of a_s, these are the indicial equations for the $\theta = \pi/2$ version of the Kapustin-Witten equations that were analyzed by Mazzeo and Witten. As will be explained momentarily, an analogous decoupling also exists for $\beta \neq 0$, and this will ultimately lead to the conclusion that the indicial roots at $\partial M^5 \setminus \Sigma_K$ do not depend on β at all. In the upcoming discussion we closely follow the exposition for the case $\beta = 0$ in [MW14, Sec. 2.3].

The Lie algebra $\mathfrak{su}(2)$ acts on the fields in (2.13)-(2.16) in various ways and this can be exploited to simplify the equations. Below, we denote by V_j the $2j+1$ dimensional representations of $\mathfrak{su}(2)$, where j is a non-negative half-integer, commonly called spin.

First, consider the subalgebra $\mathfrak{su}(2)_{\mathfrak{t}} \subseteq \mathfrak{g}$ that is spanned by $(\mathfrak{t}_i)_{i=1,2,3}$. It is the image of $\rho : \mathfrak{su}(2) \to \mathfrak{g}$ and depends on the choice of ϕ_ρ in the Nahm pole boundary condition. Since we are only concerned with regular Nahm pole boundary conditions, ρ is a principal embedding and $\mathfrak{su}(2)_{\mathfrak{t}}$ is a regular subalgebra. Under the adjoint action of $\mathfrak{su}(2)_{\mathfrak{t}}$, the Lie algebra $\mathfrak{g}$ then decomposes[8] into a direct sum of non-zero integer spin representations V_j. Observe that none of the terms in the indicial equations mixes components with values in distinct V_j's (the action of $\mathfrak{t}_i$ preserves V_j by definition). This means that in solving the equations we can from now on assume that a_s, a_y, $\vec{a}$ and $\vec{b}$ all take values in the same representation V_j.

Second, there is an $\mathfrak{su}(2)$ action on the vector degrees of freedom of $\vec{a}$ and $\vec{b}$, which we denote $\mathfrak{su}(2)_{\mathfrak{s}}$ with generators $(\mathfrak{s}_i)_{i=1,2,3}$. For this, recall from Section 2.4.5 that in the quaternionic Nahm equations we think of the vectors $\vec{a}$ and $\vec{b}$ as elements of $\mathfrak{g} \otimes \operatorname{Im}\mathbb{H}$. The imaginary quaternions $\operatorname{Im}\mathbb{H}$ naturally form an $\mathfrak{su}(2)$ representation, induced by acting with commutators (= cross product) on themselves. Slightly abusing notation, this action is represented on $\mathbb{R}^3 \simeq \operatorname{Im}\mathbb{H}$ by multiplication with the 3×3-matrices $(\mathfrak{s}_i)_{jk} = -\epsilon_{ijk}$.

If $\beta \neq 0$, we need to take into account that the two quaternionic parts combine into $(a_y, \vec{b}, a_s, \vec{a}) \in \mathfrak{g} \otimes \mathbb{O}$. There is an analogous $\mathfrak{su}(2)_{\mathfrak{s}}$ action on each of the vector space summands of $\mathbb{O} \simeq \mathbb{R} \oplus \operatorname{Im}\mathbb{H} \oplus \mathbb{R} \oplus \operatorname{Im}\mathbb{H}$, which is again induced by taking commutators with elements of the first $\operatorname{Im}\mathbb{H}$-factor (and subsequently discarding any terms that land outside the summand in question). Specifically, we define $\mathfrak{su}(2)_{\mathfrak{s}}$ by the action of its generators $(\mathfrak{s}_i)_{i=1,2,3}$ as follows. The action on a_y and a_s is trivial, i.e. generators are represented by multiplication with $\mathfrak{s}_i = 0$. Meanwhile, the action on $\vec{b}$ is represented by the 3×3 matrices $(\mathfrak{s}_i)_{jk} = -\epsilon_{ijk}$ as in the quaternionic case. The action on $\vec{a}$ is similar but comes with an additional subtlety, since octonionic multiplication introduces an additional sign in the action of $\operatorname{Im}\mathbb{H}$. In this case the (octonionic) action of the imaginary quaternions is represented by the 3×3-matrices $(A_i)_{jk} = +\epsilon_{ijk}$, which satisfy the commutation relations $[A_i, A_j] = -\epsilon_{ijk} A_k$. This only provides an $\mathfrak{su}(2)_{\mathfrak{s}}$ representation if we set $\mathfrak{s}_i = A_{\tau(i)}$ with an additional anti-cyclic permutation $\tau = (132)$. In conclusion, a_s and a_y take values in the trivial representation V_0, while $\vec{a}$ and $\vec{b}$ are elements of three-dimensional representations V_1, where the generators $\mathfrak{s}_i$ act as described above.

Third, the indicial equations are invariant under $\mathfrak{su}(2)_{\mathfrak{f}}$, generated by the action of $\mathfrak{f}_i := \mathfrak{t}_i \otimes 1 + 1 \otimes \mathfrak{s}_i$ on $\mathfrak{g} \otimes \mathbb{O}$. If the fields take values in $V_j \subset \mathfrak{g}$, they decompose under the action of $\mathfrak{su}(2)_{\mathfrak{f}}$ as follows.

$$
\begin{aligned}
a_s, a_y &\in V_j \otimes V_0 = V_j^0 \\
\vec{a}, \vec{b} &\in V_j \otimes V_1 = \bigoplus_{\eta \in \{-,0,\}} V_j^\eta
\end{aligned}
$$

[8] For a generic $\mathfrak{su}(2)$-subalgebra this decomposition exists only for the complexification $\mathfrak{g}_{\mathbb{C}}$, in which case the decomposition may also involve half-integer spins.

Here we have introduced $V_j^\eta \simeq V_{j+\eta}$ to denote representations with total spin $j+\eta$. It's worth pointing out that for fixed j, according to the first line, the components a_s and a_y can only be non-zero when $\eta = 0$.

Now, to better understand the indicial equations, consider the 'spin-spin' operator

$$\mathfrak{J} := \mathfrak{t} \cdot \mathfrak{s} = \sum_{i=1}^{3} \mathfrak{t}_i \otimes \mathfrak{s}_i \,.$$

This operator yields the expressions in (2.13) and (2.14) that contain ϵ_{ijk}. Indeed, denoting $\vec{a}^\tau = (a_1, a_3, a_2)$ and $\vec{b}^\tau = (b_1, b_3, b_2)$, the action of $\mathfrak{J}$ is given by

$$\left(\mathfrak{J}\,\vec{a}\,\right)_i = -\epsilon_{ijk}[t_{\tau(j)}, a_k] \qquad \left(\mathfrak{J}\,\vec{b}\,\right)_i = \epsilon_{ijk}[t_j, b_k]$$
$$\left(\mathfrak{J}\,\vec{a}^{\,\tau}\right)_{\tau(i)} = \epsilon_{ijk}[t_j, a_k] \qquad \left(\mathfrak{J}\,\vec{b}^{\,\tau}\right)_{\tau(i)} = -\epsilon_{ijk}[t_{\tau(j)}, b_k]$$

The action of $\mathfrak{J}$ on elements of V_j^η is determined by the quadratic Casimir operators of the three $\mathfrak{su}(2)$ actions. In general, the quadratic Casimir operator of $\mathfrak{su}(2)$ with basis $\mathfrak{c}_i$ is defined by $C^2 = -\sum_{i=1}^{3} \mathfrak{c}_i^2$. On a spin J representation it takes the constant value $C^2 = J(J+1)$. In our case there are three such operators $C_\mathfrak{t}^2$, $C_\mathfrak{s}^2$ and $C_\mathfrak{f}^2$, associated to the three $\mathfrak{su}(2)$ actions on $\mathfrak{g} \otimes \mathbb{O}$. The values $C_\mathfrak{t}^2 = j(j+1)$ and $C_\mathfrak{s}^2 = 2$ are fixed, while $C_\mathfrak{f}^2$ depends on V_j^η and takes values $(j+\eta)(j+\eta+1)$. The relevance of this is that the spin-spin operator is equivalent to

$$\mathfrak{J} = -\frac{1}{2}\left(C_\mathfrak{f}^2 - C_\mathfrak{t}^2 - C_\mathfrak{s}^2\right) ,$$

such that V_j^η, $\eta = -1, 0, 1$, are eigenspaces of $\mathfrak{J}$ with eigenvalues $j+1$, 1, and $-j$, respectively.

Note that orientation reversal via τ does not preserve the total spin: if $\vec{a} \in V_j^\eta$ then $\vec{a}^\tau$ does not have definite spin with respect to $\mathfrak{su}(2)_\mathfrak{f}$, but is instead given by some non-trivial linear combination in $\oplus_\eta V_j^\eta$. Since the indicial equations (2.13)-(2.16) contain contributions from $\vec{a}$, $\vec{a}^\tau$, $\vec{b}$, and $\vec{b}^\tau$, its not possible to restrict the equations to V_j^η. However, by taking suitable linear combinations of (2.13) and (2.14), one can rewrite these as a set of decoupled equations in $\sin\beta\,\vec{a} + \cos\beta\,\vec{b}$ and $\cos\beta\,\vec{a} - \sin\beta\,\vec{b}$.

On the one hand, if we restrict to $\eta \neq 0$, the terms containing a_s and a_y vanish. In this case the indicial equations are equivalent to

$$\lambda\left(\sin\beta\,\vec{a} + \cos\beta\,\vec{b}\right) + \mathfrak{J}\left(\sin\beta\,\vec{a} + \cos\beta\,\vec{b}\right) = 0$$
$$\lambda(\cos\beta\,\vec{a}^{\,\tau} - \sin\beta\,\vec{b}^{\,\tau}) - \mathfrak{J}\left(\cos\beta\,\vec{a}^{\,\tau} - \sin\beta\,\vec{b}^{\,\tau}\right) = 0$$

On the other hand, if $\eta = 0$, we can replace the terms containing a_s and a_y by utilizing that equations (2.15) and (2.16) are solved by

$$a_i = \frac{\lambda}{j(j+1)}\left(\cos\beta\,[\mathfrak{t}_i, a_s] + \sin\beta\,[\mathfrak{t}_{\tau(i)}, a_y]\right)$$
$$b_{\tau(i)} = \frac{\lambda}{j(j+1)}\left(-\sin\beta\,[\mathfrak{t}_i, a_s] + \cos\beta[\mathfrak{t}_{\tau(i)}, a_y]\right)$$

Solving for $[\mathfrak{t}, a_y]$ and $[\mathfrak{t}, a_s]$ and plugging it into the linear combinations of (2.13) and (2.14) yields

$$\left(\lambda - \frac{j(j+1)}{\lambda}\right)\left(\sin\beta\,\vec{a} + \cos\beta\,\vec{b}\right) + \mathfrak{J}\left(\sin\beta\,\vec{a} + \cos\beta\,\vec{b}\right) = 0$$
$$\left(\lambda - \frac{j(j+1)}{\lambda}\right)\left(\cos\beta\,\vec{a}^{\tau} - \sin\beta\,\vec{b}^{\tau}\right) - \mathfrak{J}\left(\cos\beta\,\vec{a}^{\tau} - \sin\beta\,\vec{b}^{\tau}\right) = 0$$

In any case, the indicial equations reduce to an eigenvalue problem for the spin-spin operator $\mathfrak{J}$. Using the fact that V_j^η are eigenspaces of $\mathfrak{J}$, where the eigenvalue is determined by Casimir operators, leads to the following table of indicial roots:

$$(\sin\beta\,\vec{a} + \cos\beta\,\vec{b}) \in V_j^\eta \;:\; \lambda = \begin{cases} j & \eta = 1 \\ -(j+1), j & \eta = 0 \\ -(j+1) & \eta = -1 \end{cases}$$

$$(\cos\beta\,\vec{a} - \sin\beta\,\vec{b})^\tau \in V_j^\eta \;:\; \lambda = \begin{cases} -j & \eta = 1 \\ j+1, -j & \eta = 0 \\ j+1 & \eta = -1 \end{cases}$$

This concludes the evaluation of indicial roots at points in the depth 1 stratum $\partial M \setminus \Sigma_K$. Crucially, the list of indicial roots coincides with the one for the $\pi/2$-version of the Kapustin-Witten equations determined in [MW14]. Note, in particular, that there are no indicial roots in $[\underline{\lambda}, \overline{\lambda}] = [-2, 1]$, such that the indicial set Δ_0 is bounded from the left by 1.

Moving on to a short descrpition of the depth 2 stratum, let $p \in \Sigma_K$. In this case the normal operator $N_p(L)$ acts on functions over $T_p\Sigma_K \times C(H^2) \simeq \mathbb{R}^2_{s,t} \times [0,\infty)_R \times H^2_{\psi,\vartheta}$. The indicial equations now arise from considering $\mathbb{R}^2_{s,t}$-invariant functions of the form $(R^\lambda a, R^\lambda b)$, where a, b are independent of R. Equivalently, according to Proposition 2.14, these are determined by plugging in $A = A^{\theta,\lambda} + R^\lambda a$ and $B^{\theta,\lambda} + R^\lambda b$ into the θ-TEBE and extracting the terms at leading order $R^{\lambda-1}$.

The evaluation of these equations is somewhat more involved than before. The indicial roots of θ-twisted knot singularities near $\partial_K M$ were determined for the θ-Kapustin-Witten equations by Dimakis.

Lemma 2.18 ([Dim22a, Lemma 3.5]). *If $G = SU(2)$ and $(A, \phi) \sim (A^{\lambda,\theta}, \phi^{\lambda,\theta})$ as $R \to 0$, then the set of indicial roots at $\psi = \pi/2$ is $\{-1, 2\}$ in accordance with the Nahm pole boundary condition, at $\psi = 0$ is $\{-\lambda - 1, 0, 0, \lambda + 1\}$, and there are no indicial roots in the interval $(-2, 0)$ at $R = 0$.*

Importantly, the indicial roots at $\psi = \pi/2$ are compatible with the indicial roots at the depth 1 stratum, such that the depth 2 normal operator is iteratively invertible (roughly: it is invertible on certain rescaled versions of the function spaces on which the depth 1 normal operator is invertible). We conclude that the Haydys-Witten operator is an elliptic iie operator, since up to a reinterpretation of field components its normal operator at glancing angle θ coincides with the normal operator of the θ-Kapustin-Witten operator. In particular, the indicial set Δ_K is bounded from the left by $\overline{\lambda} = 0$.

2.7 Haydys-Witten Homology

We now have all ingredients at hand to qualitatively define Haydys-Witten Homology, which assigns a Floer-type instanton homology $HF(W^4)$ to any Riemannian four-manifold. The construction is a standard application of the ideas of Floer theory and is summarized in Section 2.7.1. In fact, if W^4 admits a non-vanishing unit vector field w, there is a one-parameter family of such homology groups $HF_\theta(W^4)$, $\theta \in [0, \pi]$. Moreover, the construction is functorial: Any cobordism (M^5, v), where v is a non-vanishing vector field on M^5, provides a linear map between the homology theories associated to its boundaries. In particular, there are natural linear maps between homology groups for different values of θ.

It's to a large extend unclear under what conditions Haydys-Witten homology has a fully rigorous meaning. Currently the most important missing parts are compactness and gluing results for the moduli space of Kapustin-Witten and Haydys-Witten solutions. There have been some important advances in this direction, mostly due to Taubes [Tau13; Tau17b; Tau18; Tau19; Tau21], but also see [Tan19; He19].

We conclude this section with a short explanation of Witten's proposal regarding Khovanov homology from the perspective of Haydys-Witten Floer theory in Section 2.7.2. This has attracted a lot of attention and is also the topic that lies at the heart of this thesis. The fact that there is a relation to an already existing homology theory provides an important testing ground for the general ideas of Haydys-Witten Floer theory and hints at the information that is measured by the homology. Conversely, one may hope to 'read off' properties of Khovanov homologies on general three-manifolds from general properties of Floer-like theories.

2.7.1 An Instanton Floer Homology for the Haydys-Witten Equations

The way things have been set up, it is convenient to explain the construction of Haydys-Witten homology from the perspective of the five-dimensional Haydys-Witten geometry. Let M^5 be a non-compact Riemannian manifold with corners, G a simply connected compact Lie group, and $E \to M^5$ a principal G-bundle. Assume M^5 is equipped with a non-vanishing unit vector field v that approaches ends at constant angles. The standard example to keep in mind is $M^5 = \mathbb{R}_s \times X^3 \times \mathbb{R}^+_y$ with $v = \sin\theta\partial_s + \cos\theta\partial_y$. Note that M^5 may have 'corners at infinity', commonly called poly-cylindrical ends, that separate non-compact ends at which v has different incidence angles. If we wish to include a 't Hooft operator supported on some embedded surface $\Sigma_K \subset \partial M$ in one of the boundary components, then we implicitly take M^5 to be the blowup along Σ_K and label the newly introduced boundary component $\partial_K M$ with a magnetic charge $\lambda \in \Gamma^\vee_{\mathrm{char}}$.

Denote by $\mathfrak{B}$ a complete collection of boundary conditions for Haydys-Witten fields (A, B) on M^5. We think of this as a set that contains for each end of M^5 the information of the type of boundary condition and any choices associated to it. For example, this might be the choice of β-twisted Nahm pole boundary conditions, which involves boundary data $\phi_\rho \in \Omega^2_{v,+}([0,\epsilon)_y \times W^4)$. Similarly, if the boundary arises from the blowup of a surface Σ_K and is labeled with a non-zero charge $\lambda \in \Gamma^\vee_{\mathrm{char}}$, then $\mathfrak{B}$ associates the description of a knot singularity within a Nahm pole boundary conditions. At non-compact ends we generally demand that the fields approach $\mathbb{R}$-invariant solutions, so $\mathfrak{B}$ specifies a choice of θ-Kapustin-Witten solution (or Vafa-Witten solution if $\theta = 0$). We write $\mathcal{M}^{\mathrm{HW}}(M^5, v\,;\,\mathfrak{B})$ for the space of Haydys-Witten solutions that satisfy the boundary conditions determined by $\mathfrak{B}$, modulo gauge transformations that act trivially at boundaries and non-compact ends.

Let us now consider five-manifolds of the form $M^5 = \mathbb{R}_s \times W^4$, where W^4 is a smooth Riemannian manifold with corners, not necessarily compact. M^5 always admits the non-vanishing vector field $v = \partial_s$, which approaches the ends at $s = \pm\infty$ with incidence angle $\theta = 0$. Whenever W^4 admits a non-vanishing unit vector field w, there is a natural one-parameter family of non-vanishing vector fields $v = \cos\theta\partial_s + \sin\theta w$, with $\theta \in [0, \pi]$, that interpolates between ∂_s and w. There are, of course, many other possible choices of v on $\mathbb{R}_s \times W^4$; in particular, θ could vary along $\mathbb{R}_s$, such that the angles at $s = \pm\infty$ need not coincide. We will come back to this later, in the more general context of cobordisms (M^5, v) between four-manifolds with associated incidence angles (W^4, θ) and $(\tilde{W}^4, \tilde{\theta})$. For now consider the cylinder $M^5 = \mathbb{R}_s \times W^4$ and fix $v = \cos\theta\partial_s + \sin\theta w$ for some constant θ.

The boundary conditions at $s \to \pm\infty$ are classified by solutions of the θ-Kapustin-Witten equations on W^4. Denote the moduli space of Kapustin-Witten solutions modulo gauge transformations by $\mathcal{M}^{\mathrm{KW}}(W^4, \theta)$. Given $x, x' \in \mathcal{M}^{\mathrm{KW}}(W^4, \theta)$, a complete set of boundary conditions on $M^5 = \mathbb{R}_s \times W^4$ is given by additionally specifying boundary conditions for each of the remaining ends. Let us denote this by

$$\mathfrak{B} = \mathfrak{b} \sqcup \Big\{ \lim_{s\to-\infty} (A, B) = x \Big\} \sqcup \Big\{ \lim_{s\to+\infty} (A, B) = x' \Big\} .$$

The choices collected in $\mathfrak{b}$ have to be be compatible with x and x' at corners of M^5. This means that $\mathfrak{b}$ has to be chosen in such a way that it interpolates between the boundary conditions that x and x' satisfy. One can think of this as a collection of four-dimensional instantons, one for each end of M^5.

A simple, yet non-trivial example of such a boundary instanton arises for $M^5 = \mathbb{R}_s \times X^3 \times \mathbb{R}^+_y$ at $y \to \infty$. First note that a natural boundary condition for a Kapustin-Witten solution $x = (A, \phi)$ on $X^3 \times \mathbb{R}^+_y$ is that it approaches a flat connection $(A^\sigma, 0)$ as $y \to \infty$, specified by a choice of a group homomorphism $\sigma : \pi_1(X^3) \to G$. Hence, assume that $x, x' \in \mathcal{M}^{\mathrm{KW}}(X^3 \times \mathbb{R}^+_y, \theta)$ approach flat connections associated to σ and σ', respectively. A consistent boundary condition at the non-compact end $y \to \infty$ of $M^5 = \mathbb{R}_s \times X^3 \times \mathbb{R}^+_y$ must then be a solution of β-Kapustin-Witten equations on $\mathbb{R}_s \times X^3$ that interpolates between the two flat connections A^σ and $A^{\sigma'}$. In the special case where $\beta = 0$ and $\phi = 0$, this is equivalent to a choice of self-dual connection, i.e. a Donaldson-Floer instanton, on the four-manifold $\mathbb{R}_s \times X^3$ that sits at $y = \infty$.

Let us now define the Morse-Smale-Witten complex that underlies Haydys-Witten homology. For simplicity assume that there is only a finite set[9] of Kapustin-Witten solutions on W^4 and consider the free abelian group generated by these solutions:

$$CF_\theta = \bigoplus_{x \in \mathcal{M}^{\mathrm{KW}}(W^4,\theta)} \mathbb{Z}\,[x] \ .$$

Note that, by definition, elements of CF_θ are stationary solutions of the Haydys-Witten equations – or equivalently, critical points of an appropriate Kapustin-Witten functional. This coincides with the usual construction of the Morse-Smale-Witten complex in Morse theory.

The Morse-Smale-Witten complex is equipped with a differential d_v that counts Haydys-Witten instantons that interpolate between x and y. To make this precise, consider the moduli space of Haydys-Witten solutions where the fields (A,B) approach x and y as $s \to \pm\infty$, respectively, and moreover satisfy some fixed boundary conditions $\mathfrak{b}$ at the remaining boundaries and non-compact ends. This moduli space always admits an $\mathbb{R}$-action by translation $s \mapsto s + c$ along the flow direction $\mathbb{R}_s$, which maps one solution to an equivalent one that differs only by the parametrization of $\mathbb{R}_s$. To count instantons we thus consider the quotient of the moduli space by this action. Also, we need to take into account that several instantons at the boundary might provide a consistent choice of boundary conditions $\mathfrak{b}$. As a result Haydys-Witten instantons that interpolate from x to y are classified by

$$\mathcal{M}(x,y) := \bigcup_{\mathfrak{b}} \mathcal{M}^{\mathrm{HW}}\Big(\mathbb{R}_s \times W^4,\ v\,;\ \mathfrak{B} = \mathfrak{b} \sqcup \big\{ \lim_{s\to-\infty} (A,B) = x \big\} \sqcup \big\{ \lim_{s\to+\infty} (A,B) = y \big\}\Big)\Big/ \mathbb{R} \ .$$

On grounds of general properties of elliptic differential operators, this is expected to be a smooth oriented manifold.

Note that the boundary conditions $\mathfrak{b}$ are sometimes classified by an analogous moduli space of instanton solutions in one dimension less. The disjoint union of possible boundary conditions then is equivalent to a product of smooth manifolds. For example, in the context of boundary instantons at $y = \infty$ on the five-manifold $\mathbb{R}_s \times X^3 \times \mathbb{R}^+_y$, the moduli space is of the form

$$\mathcal{M}(x,x') = \mathcal{M}^{\mathrm{HW}}\left(\lim_{s\to\pm\infty} (A,B) = \begin{cases} x, & \lim_{y\to\infty} x = \sigma \\ x', & \lim_{y\to\infty} x' = \sigma' \end{cases} \right) \times \mathcal{M}^{\mathrm{asd}}(\sigma,\sigma') \ ,$$

where $\mathcal{M}^{\mathrm{asd}}(\sigma,\sigma')$ is the moduli space of anti-self-dual connections on $\mathbb{R}_s \times X^3$ that interpolate between A^σ and $A^{\sigma'}$.

In the definition of d_v we rely on the dimension of $\mathcal{M}(x,y)$. In Morse theory, i.e. on finite dimensional manifolds, the Morse-Smale-Witten complex carries a natural $\mathbb{Z}$-grading by the Morse index, defined by the number of negative eigenvalues of the Hessian at a given critical point. Since this determines the number of unstable flow directions in the vicinity of a critical point, the difference between the index of distinct critical points determines the dimension of the moduli space of flows $\mathcal{M}(x,y)$. In the infinite dimensional setting the Morse index does

[9] This is equivalent to the statement that the moduli space is a compact manifold of dimension zero, which is something one would ultimately like to prove.

not make sense; the linearization of the Kapustin-Witten operator typically has infinitely many negative eigenvalues[10]. Observe, however, that the difference of Morse indices only depends on the relative change of negative eigenvalues along a flow line. This is known as the spectral flow of an operator and has an analogue in Floer theory. Thus, as is common in Floer theory, we define a relative index $\mu(x, y)$ for any pair of generators $x, y \in \mathcal{M}^{\mathrm{KW}}(W^4, \theta)$, by the spectral flow of the Kapustin-Witten differential operator along a Haydys-Witten instanton. This, in turn, coincides with the index of the Haydys-Witten differential operator when it acts on fields that are subject to the complete set of boundary conditions $\mathfrak{B}$.

$$\mu(x, y) := \operatorname{ind} \mathbf{HW}_v|_{\mathfrak{B}}$$

The moduli space of Haydys-Witten instantons $\mathcal{M}(x, y)$ is expected to have dimension $\mu(x, y)-1$. Notably, it is zero-dimensional whenever $\mu(x, y) = 1$, in which case we denote by $\#\mathcal{M}(x, y)$ the signed count of its (oriented) elements.

The Floer differential is the linear map $CF_\theta \to CF_\theta$ defined by

$$d_v[x] = \sum_{\mu(x,y)=1} \#\mathcal{M}(x, y) \cdot [y] \,.$$

One expects that $d_v^2 = 0$, such that (CF_θ, d_v) is indeed a cochain complex. The standard proof in Floer theory relies on compactness and gluing theorems for the flow equations. More precisely, consider the compactification of the moduli space of gradient flows with relative index 2, which is an oriented manifold of dimension 1. If the Haydys-Witten equations with boundary conditions $\mathfrak{B}$ are well-behaved, the compactification is fully determined by adding broken flow lines. The latter are exactly what we need to count when calculating d_v^2. Oriented manifolds of dimension one are either circles, which don't have boundary components and can't contribute to d_v^2, or intervals with boundary components of opposite orientation. It follows that contributions to d_v^2 always arise in pairs of opposite orientation and consequently add up to 0.

Definition 2.19 (Haydys-Witten Homology)**.** The Haydys-Witten homology associated to a four-manifold W^4 is the homology of the chain complex (CF_θ, d_v):

$$HF_\theta(W^4) := H(CF_\theta, d_v) \,.$$

It might be helpful to emphasize that only the $\theta = 0$ version of this homology theory exists for arbitrary Riemannian four-manifolds and Proposition 2.13 states that in that case the Morse-Smale-Witten complex is generated by Vafa-Witten solutions. The one-parameter family of homology groups only exists if W^4 admits a non-vanishing vector field w. Since any non-compact manifold automatically admits a non-vanishing vector field, this is only an obstruction for compact manifolds. Recall, that according to paragraph 4.4.3, finite energy solutions of

[10] We have encountered a similar situation in the path-integral description, where we mentioned a grading by the 'fermion number of the filled Dirac sea'. This makes sense only after choosing a reference vacuum for which the fermion number is *defined* to be zero.

the θ-Kapustin-Witten equations on compact manifolds are trivial whenever $\theta \neq 0$, so in this case one could arguably define $HF_\theta(W^4)$ to be the trivial group. Therefore, the only situation where $HF_\theta(W^4)$ with $\theta \neq 0$ remains ambiguous is in the context of infinite energy solutions on compact manifolds, which for example appear in connection with Nahm pole boundary conditions.

Instanton Grading The Morse-Smale-Witten complex CF_θ is naturally $\mathbb{Z}$-graded by the instanton number of the principal bundle $E \to W^4$, which up to a constant is the integral of the first Pontryagin class $p_1(\operatorname{ad} E) = \frac{1}{8\pi^2} \operatorname{Tr} F_A \wedge F_A$. The complex decomposes into submodules CF_θ^k, spanned by Kapustin-Witten solutions with instanton number $k \in \mathbb{Z}$:

$$CF_\theta^\bullet = \bigoplus_{k \in \mathbb{Z}} CF_\theta^k \ .$$

To understand the interaction between the instanton grading and the differential d_v, observe that $p_1(\operatorname{ad} E)$ is a conserved four-form current. We can consider the associated charge at any time $s \in \mathbb{R}_s$:

$$P(s) = \frac{1}{32\pi^2} \int_{\{s\} \times W^4} \operatorname{Tr} F_A \wedge F_A \ .$$

Although the integrand is conserved, current density may disappear at boundaries and non-compact ends of W^4 as we follow the flow along $\mathbb{R}_s$. The difference in instanton number between the start and end point $s \to \pm\infty$ of a Haydys-Witten instanton is given by Stokes' theorem.

$$\Delta P := \lim_{s \to \infty} \big(P(s) - P(-s) \big) = \sum \frac{1}{32\pi^2} \int_{\partial_i M} \operatorname{Tr} F_A \wedge F_A$$

The right hand side is a sum over all boundaries and non-compact ends of M^5, except the ones at $s = \pm\infty$ (which appear on the left hand side). Each end contributes with its own instanton number, or more precisely the instanton number associated to the pullback of E to $\partial_i M$.

In conclusion, the differential d_v generally doesn't preserve the grading on $CF_\theta^\bullet$ and consequently there is no $\mathbb{Z}$-grading on $HF_\theta(W^4)$. However, the topology of the pullback bundles at ends of M^5 may arguably be viewed as part of the boundary data $\mathfrak{b}$, so the change in P-grading is ultimately controlled by the interplay of all the boundary conditions that are imposed on the Haydys-Witten instantons. For example, one could choose to only take into account Haydys-Witten instantons for which ΔP is fixed, such that d_v has a fixed degree ΔP.

Cobordisms Let us shortly comment on the functorial properties of Haydys-Witten Floer theory. Assume (M^5, v) is a cobordism that interpolates between four-manifolds (W^4, θ) and $(\tilde{W}^4, \tilde{\theta})$, where $\theta, \tilde{\theta}$ denote the incidence angles between v and the boundaries. We promote the boundaries to non-compact ends by gluing in cylindrical ends $(-\infty, 0]_s \times W^4$ and $[0, \infty)_s \times \tilde{W}^4$, respectively. The vector field v extends to a unique constant vector field on the cylinders, since we assume that v is already constant in some tubular neighbourhood of the boundaries.

To each end associate the corresponding Morse-Smale-Witten complex of boundary conditions $CF_\theta(W^4)$ and $CF_{\tilde{\theta}}(\tilde{W}^4)$. We can proceed exactly as before to define a linear map

$$\Phi_{(M^5,v)} \,:\, CF_\theta(W^4) \to CF_{\tilde{\theta}}(\tilde{W}^4)\,,\ [x] \mapsto \sum_{\mu(x,y)=1} \#\mathcal{M}(x,y)[y]\,.$$

The only difference is that we now count Haydys-Witten solutions on (M^5, v) instead of the cylinder $(\mathbb{R}_s \times W^4, v = \partial_y)$.

Under appropriate compactness and gluing assumptions for the Haydys-Witten equations, the induced map $\Phi_{(M^5,v)}$ is a chain map:

$$\Phi_{(M^5,v)} \circ d_v = d_v \circ \Phi_{(M^5,v)}\,.$$

One way to see this is to realize that the concatenation of $\Phi_{(M^5,v)}$ and d_v is determined by the number of broken flow lines of index 2 on M^5, since we glue the instantons described by d_v to either the initial or final cylindrical end of M^5. As before, these broken flow lines are in correspondence with the boundary components of the moduli space of Haydys-Witten instantons of index 2. Since the relevant moduli space is the same, regardless of the order of $\Phi_{(M^5,v)}$ and d_v, these counts coincide.

It follows that $\Phi_{(M^5,v)}$ induces a linear map on homology:

$$\left(\Phi_{(M^5,v)}\right)_* \,:\, HF_\theta(W^4) \to HF_{\tilde{\theta}}(\tilde{W}^4)\,.$$

Hence, Haydys-Witten homology is a functor from the category of five-dimensional cobordisms, equipped with a non-vanishing vector field, to the category of groups. It is, therefore, a topological quantum field theory (TQFT) in the sense of the Atiyah-Segal axioms. From the point of view of physics it is the TQFT that arises by a topological twist of $5d\,\mathcal{N} = 2$ super Yang-Mills theory on $\mathbb{R}_s \times W^4$ coupled to $4d\,\mathcal{N} = 4$ super Yang-Mills theory at $s = \pm\infty$.

2.7.2 Relation to Khovanov Homology

Haydys-Witten Floer theory was introduced by Witten to describe Khovanov homology in terms of quantum field theory. Witten showed that there is a relation between Haydys-Witten Floer homology and Chern-Simons theory on X^3 – and thus knot invariants – if one considers four-manifolds of the form $W^4 = X^3 \times \mathbb{R}^+_y$ with Nahm pole boundary conditions at $y = 0$ [Wit10; Wit11a]. Under this correspondence, a knot carries over to a magnetically charged 't Hooft operator embedded in the boundary $\partial W^4 = X^3$. As explained in Section 2.6, this setup is geometrized by considering the blowup $[W^4; K]$ and imposing a certain singular behaviour at the blown up boundary $\partial_K W$.

Hence, let us associate to a pair (X^3, K) the Haydys-Witten homology of $[X^3 \times \mathbb{R}^+_y; K]$. The vector field $w = \partial_y$ provides a non-vanishing vector field on W^4, so there is a one-parameter family of Haydys-Witten homologies with respect to $v = \cos\theta\partial_s + \sin\theta\partial_y$, $\theta \in [0,\pi]$. The associated Morse-Smale-Witten complex CF_θ is spanned by solutions of the θ-Kapustin-Witten equations that satisfy suitable Nahm pole boundary conditions with knot singularities.

The differential d_v counts Haydys-Witten instantons on the cylinder $M^5 = \mathbb{R}_s \times [W^4; K]$. This manifold is equivalent to the one obtained by first lifting the knot to the $\mathbb{R}_s$-invariant surface $\Sigma_K = \mathbb{R}_s \times K \times \{0\}$ inside the boundary of $\mathbb{R}_s \times W^4$ and blowing up afterwards, i.e.

$$M^5 = \mathbb{R}_s \times [W^4; K] = [\mathbb{R}_s \times W^4; \Sigma_K] .$$

As always, we leave the blowup mostly implicit and simply write $M^5 = \mathbb{R}_s \times X^3 \times \mathbb{R}^+_y$ with original boundary $\partial_0 M$ at $y = 0$ (with Σ_K removed) and blown up boundary $\partial_K M$.

To fully determine d_v it remains to specify which kind of boundary conditions $\mathfrak{b}$ the Haydys-Witten instantons on (M^5, v) shall satisfy:

- At $\partial_0 M^5$ the fields satisfy the β-twisted regular Nahm pole boundary condition, where the incidence angle is given by $\beta = \pi/2 - \theta$. The boundary data ϕ_ρ of the five-dimensional Nahm pole boundary condition is the unique $\mathbb{R}_s$-invariant continuation of some fixed four-dimensional Nahm pole boundary condition at $\partial_0 W^4$.
- At $\partial_K M^5$, the fields exhibit a knot singularity and are otherwise consistent with the surrounding Nahm pole boundary conditions. Since the glancing angle between v and Σ_K is θ, the knot singularity is modeled on solutions of the θ-TEBE.
- At $y \to \infty$ the fields approach an $\mathbb{R}_y$-invariant finite energy solution of the Haydys-Witten equations. This corresponds to a solution of the β-Kapustin-Witten equations, where $\beta = \pi/2 - \theta$.
- At any non-compact end or boundary of X^3, the fields approach maximally symmetric, stationary solutions of the Haydys-Witten equations that are compatible with the boundary conditions at adjacent boundaries and independent of the flow direction $\mathbb{R}_s$. What exactly this means is best described on a case-by-case basis.

The first two items just spell out the Nahm pole boundary conditions with knot singularity, as described more thoroughly in Section 2.6. For $y \to \infty$ there might be non-trivial boundary instantons, classified by solutions of β-Kapustin-Witten solutions on $\mathbb{R}_s \times X^3$. Note that for $X^3 = S^3$ or $\mathbb{R}^3$ there are no non-trivial Kapustin-Witten solutions with finite energy, because of the vanishing result of Corollary 2.6. Since the rest of the boundary conditions are explicitly $\mathbb{R}_s$-invariant, the differential d_v preserves the instanton grading and Haydys-Witten homology is $\mathbb{Z}$-graded.

By definition, Haydys-Witten homology $HF_\theta([X^3 \times \mathbb{R}^+_y; K])$ is given by solutions of the θ-Kapustin-Witten equations, subject to Nahm pole boundary conditions with knot singularities at $y = 0$, modulo Haydys-Witten instantons. Witten originally described the case where $v = \partial_y$, in which case $CF_{\pi/2}$ is spanned by solutions of the $\theta = \pi/2$ version of the Kapustin-Witten equations, in which case boundary instantons at $y \to \infty$ are given by Vafa-Witten solutions on $\mathbb{R}_s \times X^3$. As mentioned earlier, the deformation to $\theta \neq \pi/2$ was considered by Gaiotto and Witten soon afterwards.

Witten's gauge theoretic approach to Khovanov homology is now summarized by the following statement.

Conjecture ([Wit11a]). *Haydys-Witten homology $HF_\theta([X^3 \times \mathbb{R}^+_y; K])$ is a topological invariant of the pair (X^3, K). In particular, if $X^3 = S^3$ (or $\mathbb{R}^3$) and $\theta = \pi/2$, this invariant is $\mathbb{Z}$-graded and coincides with Khovanov homology:*

$$HF^\bullet_{\pi/2}([S^3 \times \mathbb{R}^+_y; K]) = Kh^\bullet(K) \, .$$

Moreover, any knot cobordism Σ induces a map on Haydys-Witten homology via the five-dimensional cobordism $M^5 = [\mathbb{R}_s \times X^3 \times \mathbb{R}^+_y; \Sigma]$, and this coincides with the corresponding map on Khovanov homology.

3 Growth of the Higgs Field for Kapustin-Witten Solutions on ALE and ALF Gravitational Instantons

Let $G = SU(2)$ and consider a principal G-bundle E over a complete Riemannian manifold (W^n, g) of dimenson n. Throughout, we assume that W^n is an ALX manifold. Suffice it to say for now that we take this to mean W^n is a non-compact manifold with fibered ends such that the k-dimensional fibers have bounded volume. Consequently, the volume of geodesic balls asymptotically grows like r^{n-k}.

Denote by $(A, \phi) \in \mathcal{A}(E) \times \Omega^1(W^n, \operatorname{ad} E)$ a pair consisting of a connection on E and an $\operatorname{ad} E$-valued one-form. We write $\star$ for the Hodge star operator and equip $\Omega^k(W^n, \operatorname{ad} E)$ with the density-valued inner product $\langle a, b\rangle = \operatorname{Tr} a \wedge \star b$. Upon integration this provides the usual L^2-product $\langle a, b\rangle_{L^2(W)} = \int_{W^n} \langle a, b\rangle$ on $\Omega^k(W^n, \operatorname{ad} E)$. Throughout, we assume that A and ϕ have enough derivatives and are locally square-integrable.

In this chapter we report on a property of the pair (A, ϕ) whenever it satisfies the second order differential equation

$$\nabla^{A\dagger}\nabla^{A}\phi + \frac{1}{2} \star [\star[\phi \wedge \phi] \wedge \phi] + \operatorname{Ric} \phi = 0 \, . \tag{3.1}$$

Here $\nabla^{A\dagger}$ is the formal adjoint of ∇^A with respect to the L^2-product and the Ricci curvature is viewed as an endomorphism of $\Omega^1(W, \operatorname{ad} E)$.

The differential equation (3.1) is of particular relevance in the context of the Kapustin-Witten equations. To see this, consider for the moment the case of a four-manifold W^4 and define the Laplace-type differential operator on $\Omega^1(W, \operatorname{ad} E)$:

$$\tilde{\Delta}_A(\phi) = -d_A d_A^\star \phi + \star 2 d_A (d_A \phi)^- \, ,$$

where $d_A^\star = \star d_A \star$ is the usual codifferential and $(\cdot)^\pm$ denotes the (anti-)self-dual part of a given two-form. Compare this operator with the θ-Kapustin-Witten equations for (A, ϕ) (cf. Section 2.4.2), which are given by

$$\begin{aligned}
\left(\cos\tfrac{\theta}{2}\, (F_A - \tfrac{1}{2}[\phi \wedge \phi]) - \sin\tfrac{\theta}{2}\, d_A\phi \right)^+ &= 0 \, , \\
\left(\sin\tfrac{\theta}{2}\, (F_A - \tfrac{1}{2}[\phi \wedge \phi]) + \cos\tfrac{\theta}{2}\, d_A\phi \right)^- &= 0 \, , \\
d_A^\star \phi &= 0 \, .
\end{aligned}$$

Clearly, if (A, ϕ) is a solution of the $\theta = 0$ version of the Kapustin-Witten equations, then $\tilde{\Delta}_A \phi = 0$. Moreover, using a Bochner-Weitzenböck identity that relates $\tilde{\Delta}_A$ and the Bochner Laplacian $\nabla^{A^\dagger} \nabla^A$, as well as the remaining part of the 0-Kapustin-Witten equations $F_A^+ = [\phi \wedge \phi]^+$, one finds that harmonicity of ϕ with respect to $\tilde{\Delta}_A$ is equivalent to equation (3.1). In fact a very similar argument shows that the same is true if (A, ϕ) is a solution of the θ-Kapustin-Witten equations [Tau13; Tau17a; NO21].

Let us now return to general n-manifolds. In what follows, we are guided by the intuition that if ϕ satisfies (3.1), then it is harmonic with respect to some well-behaved Laplace-type operator. In particular, one should expect that it satisfies an appropriate analogue of the mean-value principle. Hence, fix some point $p \in W^4$, denote by B_r the closed geodesic ball of radius r centered at p, and consider the non-negative function κ on $[0, \infty)$ that is defined by

$$\kappa^2(r) = \frac{1}{r^{n-k-1}} \int_{\partial B_r} \|\phi\|^2 \ . \tag{3.2}$$

As a consequence of the asymptotic volume growth on W^n, $\kappa(r)$ is related to the average value of $\|\phi\|$ on geodesic spheres ∂B_r with large radius r. The mean-value principle for Laplace-type differential operators then suggests that ϕ should satisfy an inequality of the form $\|\phi(p)\| \leq \kappa(r)$ for $r > 0$. Although contributions from non-trivial curvature in the interior of B_r in general preclude this naive mean-value inequality, the controlled asymptotics of ALX spaces retains enough control to deduce analogous bounds for points that are far away from p.

A classical consequence of the mean-value principle is a relation between the asymptotic behaviour of κ at large radius and the values of ϕ in the interior of the ball. For example, if the naive mean-value inequality was satisfied at every point $p \in W^n$ and $\kappa(r) \to 0$ as $r \to \infty$, then ϕ would be identically zero everywhere.

For $W^n = \mathbb{R}^n$, Theorem 2.4 by Taubes generalizes this kind of statement to a dichotomy between the growth of $\kappa(r)$ at infinity and the vanishing of $[\phi \wedge \phi]$ on all of W^n [Tau17a]. Here we prove that this dichotomy holds more generally if W^n is an ALX gravitational instanton (this is Theorem 3.19 below):

Theorem A. *Let W^n be a complete, Ricci flat ALX manifold of dimension $n \geq 2$ with asymptotic fibers of dimension $k \leq n-1$ and sectional curvature bounded from below. Consider (A, ϕ) as above and assume the pair satisfies the second-order differential equation (3.1). Then either*

(i) there is an $a > 0$ such that $\liminf_{r \to \infty} \frac{\kappa(r)}{r^a} > 0$, or

(ii) $[\phi \wedge \phi] = 0$.

If the fields (A, ϕ) are solutions of the θ-Kapustin-Witten equations and have square-integrable field strength we can say slightly more (cf. Theorem 3.20).

Theorem B. *Let W^4 be a complete, Ricci flat ALX manifold of dimension 4 with asymptotic fibers of dimension $k \leq 3$ and sectional curvature bounded from below. Assume (A, ϕ) are solutions of the θ-Kapustin-Witten equations and satisfy $\int_{W^4} \|F_A\|^2 < \infty$, then either*

(i) *there is an* $a > 0$ *such that* $\liminf_{r\to\infty} \frac{\kappa(r)}{r^a} > 0$, *or*

(ii) $[\phi\wedge\phi] = 0, \nabla^A\phi = 0$, *and* A *is self-dual if* $\theta = 0$, *flat if* $\theta \in (0, \pi)$, *and anti-self-dual if* $\theta = \pi$.

As an immediate consequence of Theorem B and Theorem 2.5 we are able to generalize Corollary 2.6, confirming a conjecture of Nagy and Oliveira. For this recall the definition of Kapustin-Witten energy

$$E_{\mathrm{KW}} = \int_{W^4} \left(\|F_A\|^2 + \|\nabla^A\phi\|^2 + \|[\phi\wedge\phi]\|^2\right) .$$

Corollary C (Nagy-Oliveira Conjecture [NO21])**.** *Let* (A, ϕ) *be a finite energy solution of the* θ*-Kapustin-Witten equations with* $\theta \neq 0 \pmod{\pi}$ *on an ALE or ALF gravitational instanton and let* $G = SU(2)$. *Then* A *is flat,* ϕ *is* ∇^A*-parallel, and* $[\phi\wedge\phi] = 0$.

Proof. Under the given assumptions, the main result of Nagy and Oliveira [NO21, Main Theorem 1] states that ϕ has bounded norm and thus, in particular, bounded average over spheres. It follows that $\liminf_{r\to\infty} \frac{\kappa(r)}{r^a} \to 0$ for any $a > 0$, while the finite energy condition subsumes square-integrability of F_A. Therefore, Theorem B implies that $[\phi\wedge\phi] = 0$, $\nabla^A\phi = 0$, and that A is flat. □

Remark. The preceding argument was used by Nagy and Oliveira to establish Corollary 2.6, which holds for $W^4 = \mathbb{R}^4$ and $S^1 \times \mathbb{R}^3$ [NO21, Corollary 1.3]. Nagy and Oliveira relied on a version of Theorem B that applies to $W^4 = \mathbb{R}^4$ and was provided by Taubes alongside the original dichotomy [Tau17a]. Their conjecture stemmed from the expectation that Taubes' results can be extended to ALX spaces in general.

The main insight reported in this chapter is that Taubes' proof strategy for the special case $W^n = \mathbb{R}^n$ carries over to general ALX spaces. This is a consequence of the well-behaved asymptotic volume growth, where problems that arise from non-zero curvature in the interior can be excised. Hence, the proof of Theorem A closely follows the one provided by Taubes in [Tau17a].

We proceed as follows: In Section 3.1 we collect the relevant definitions and recall several classical results that will be used throughout. Then, in Section 3.2, we investigate the derivative of κ and introduce the relevant analogue of Almgren's frequency function, as well as a function that captures contributions from the mean curvature of the geodesic sphere. The key finding of that section is that κ is asymptotically almost non-decreasing, which is a prerequisite for most of the heavy lifting in subsequent sections. In Section 3.3, we present a somewhat unusual version of unique continuation that is satisfied by κ. The main insight is the content of Section 3.4, where we explain that slow asymptotic growth of κ results in bounds for the frequency function. All these results are refined with respect to the components of the one-form $\phi = \sum_i \phi_i dx^i$ by introducing in Section 3.5 what we call the correlation tensor. Using a second line of arguments, we also determine a priori bounds of the type $\|\phi(x)\| \leq \kappa(r)$ in Section 3.6, which are the anticipated analogues of the mean value inequalities mentioned already above. Finally, all these ingredients are combined into a proof of Theorem A in Section 3.7, while the proof of Theorem B occupies Section 3.8.

3.1 ALX Manifolds and Classical Results in Riemannian Geometry

For the purposes of this thesis, an ALX space is a non-compact complete Riemannian manifold that asymptotically looks like a fibration of closed manifolds, where the fibers have bounded volume. This is made precise in the following definition.

Definition 3.1 (ALX$_k$ Manifold)**.** Let (W^n, g) be a complete Riemannian manifold of dimension n and fix $p \in W^n$. Let $\pi_Y : Y^{n-1} \to B^{n-k-1}$ be a fibration over an $n-k-1$-dimensional closed Riemannian base (B, g_B) with k-dimensional closed Riemannian fibers (X, g_X). Equip $(0,\infty)_r \times Y$ with the model metric $g_\infty = dr^2 + g_X + r^2 g_B$. We say W^n is an ALX$_k$ manifold if its end is modeled on $(0,\infty) \times Y$, that is, if there exists $R > 0$ such that there is a diffeomorphism $\varphi : W^n \setminus B_R(p) \to (R,\infty) \times Y^{n-1}$ that satisfies for $j = 0, 1$, and 2

$$\lim_{r\to\infty} r^j \left\|(\nabla^{\mathrm{LC}})^j (g - \varphi^* g_\infty)\right\|_{L^\infty(\partial B_r)} = 0\,. \tag{3.3}$$

Proposition 3.2. *If (W^n, g) is an ALX$_k$ space then*

$$\operatorname{vol} B_r(p) \sim r^{n-k} \operatorname{vol} X \quad (r \to \infty)\,.$$

Definition 3.3. We call an ALX space W^n a *gravitational instanton* if it is Ricci flat, sectional curvature is bounded from below, and the Riemann curvature tensor satisfies the decay condition $|\mathrm{Rm}|(x) \le d(p,x)^{-2-\epsilon}$ for some $\epsilon > 0$.

Remarks.

- ALX spaces are usually considered in the context of four-manifolds and the 'X' is a placeholder for the following cases: ALE or Asymptotically Locally Euclidean ($k = 0$), ALF or Asymptotically Locally Flat ($k = 1$), ALG ($k = 2$), ALH ($k = 3$), where the last two are named by induction.
- We do not demand that ALX gravitational instantons are hyperkähler, while we add the possibly non-standard condition of bounded sectional curvature.

Examples.

- The prototypical example of an ALE manifold is Euclidean space $\mathbb{R}^n$. In this case there is the obvious diffeomorphism $\mathbb{R}^n \setminus \{0\} \simeq (0,\infty) \times S^{n-1}$ via spherical coordinates, the metric is $\varphi_* g = dr^2 + r^2 g_{S^{n-1}} = g_\infty$, and the volume of balls grows with the radius as r^n.

- A prototypical example of an ALF space is $S^1 \times \mathbb{R}^{n-1}$ with the product metric $g = dt^2 + g_{\mathbb{R}^{n-1}}$. Again, spherical coordinates on the $\mathbb{R}^{n-1}$ factor provide a diffeomorphism to $(0,\infty)\times S^1 \times S^{n-2}$ with metric $\varphi_* g = dr^2 + dt^2 + r^2 g_{S_r^{n-2}}$, such that S^1 has constant size, while S_r^{n-2} is the sphere of radius r centered at $0 \in \mathbb{R}^{n-1}$. Once the volume of S^1 is 'filled', the volume of a geodesic ball approaches a growth of order r^{n-1}.

- Famous examples of less trivial four-dimensional ALF gravitational instantons are (multi-centered) Taub-NUT spaces. These are S^1-fibrations over $\mathbb{R}^3$, where the S^1-fiber has asymptotically finite volume.

- Motivated by an influential conjecture by Gaiotto, Moore, and Neitzke about the asymptotic geometry of the moduli space of $SU(N)$ Higgs bundles [GMN11], it has recently been proved that the moduli space of $SL(2,\mathbb{C})$-Higgs bundles on the 4-punctured sphere is an example of a four-dimensional hyperkähler ALG gravitational instanton [Fre+20].

Theorem 3.4 (Global Laplacian Comparison Theorem [Cal58])**.** *If* $\mathrm{Ric} \geq (n-1)K$ *and* $r(x) = d(p,x)$ *denotes the geodesic distance function based at a point p, then*

$$\Delta r \leq (n-1)\Delta_K r \,,$$

where on the right hand side Δ_K *is the Laplacian on the unique complete, n-dimensional, simply connected space of constant sectional curvature K.*

Remark. Since r is not necessarily differentiable the global Laplacian comparison must be understood in a weak sense, e.g. in the weak sense of barriers as in the work of Calabi [Cal58]. However, for our purposes it is sufficient to consider the smooth locus of r, where the inequality holds as stated.

Proposition 3.5 (Mean Curvature Comparison on ALX spaces)**.** *Let* W^n *be an* ALX_k *space and fix a point* $p \in W^n$*. The Laplacian of the distance function* $r(x) = d(p,x)$*, or equivalently the mean curvature of the geodesic sphere of radius r based at x, has the following asymptotic behaviour.*

$$\Delta r \sim \begin{cases} \frac{n-1}{r} & (r \to 0) \\ \frac{n-k-1}{r} & (r \to \infty) \end{cases}$$

Furthermore, if $\mathrm{Ric} \geq 0$*, then it is bounded from above by*

$$\Delta r \leq \frac{n-1}{r} \,.$$

Proof. For a start, note that $r(x)$ is smooth on $M \setminus \{p, \mathrm{Cut}(p)\}$, where $\mathrm{Cut}(p)$ is the cut locus of p. It is a standard result that the cut locus on a complete Riemannian manifold has measure zero, so r is differentiable almost everywhere. The Gauss lemma tells us that $\nabla^{\mathrm{LC}} r = \partial_r$ is the radial vector field of unit norm and is normal to geodesic spheres. As an aside, note that

$\Delta r = \mathrm{tr}(\nabla^{\mathrm{LC}})^2 r$ is the trace of the second fundamental form of the geodesic sphere and as such is identical to its mean curvature.

The asymptotic behaviour for $r \to 0$ follows e.g. by a direct calculation in Riemann normal coordinates. In particular, use $g_{ij} = \delta_{ij} + \mathcal{O}(r^2)$ and $\Gamma^i_{jk} = \mathcal{O}(r)$ and then observe that at leading order the result is identical to the Euclidean case, while higher order corrections are $\mathcal{O}(r)$:

$$\Delta r = \frac{n-1}{r} + \mathcal{O}(r) .$$

When $r \to \infty$ the ALX_k condition (3.3) implies that $\Delta r \sim \Delta_\infty r$, where Δ_∞ denotes the Laplacian associated to g_∞ on $(0,\infty) \times Y^{n-1}$. Under the diffeomorphism to $(0,\infty) \times Y^{n-1}$ the distance function is identified with the coordinate on the first factor. Since the model metric is block diagonal and only depends on r via the r^2 factor in front of g_B, we can calculate $\Delta_{g_\infty} r$ explicitly. Let $(e_i)_{i=1,\dots,n}$ be an orthonormal frame of $(0,\infty) \times Y$ such that $e_1 = \partial_r$, $e_2, \dots, e_{k+1}$ are tangent to the fibers, and $e_{k+2}, \dots, e_n$ are tangent to the base. Write ∇ for the Levi-Civita connection associated to g_∞. By a direct calculation $\nabla_{e_i} \partial_r = \frac{1}{r} e_i$ for $i = k+2, \dots, n-1$ and zero otherwise. Hence,

$$\Delta_{g_\infty} r = \mathrm{tr}\, \nabla^2 r = g^{ij} g(\nabla_{e_i} \partial_r , e_j) = \frac{\mathrm{tr}\, g_B}{r} = \frac{n-k-1}{r} .$$

Finally, the upper bound in the case that Ric is non-negative follows directly from the Laplacian Comparison Theorem (Theorem 3.4). □

Theorem 3.6 (Bishop-Gromov's Volume Comparison). *Let (M, g) be a complete Riemannian manifold and assume $\mathrm{Ric} \geq (n-1)K$. Denote by $\mathrm{vol}\, B_r(p)$ the volume of the geodesic ball of radius r based at $p \in M$. Similarly write $\mathrm{vol}_K\, B_r(p_K)$ for the volume of a geodesic ball with the same radius inside the unique complete, n-dimensional, simply connected space of constant sectional curvature K at an arbitrary point p_K. Then the function defined by*

$$r \mapsto \frac{\mathrm{vol}\, B_r(p)}{\mathrm{vol}_K\, B_r(p_K)}$$

is non-decreasing and approaches 1 as $r \to 0$. This implies, in particular, $\mathrm{vol}\, B_r(p) \leq \mathrm{vol}_K\, B_r(p_K)$ and $\mathrm{vol}\, \partial B_r(p) \leq \mathrm{vol}_K\, \partial B_r(p_K)$.

Lemma 3.7. *For any point x in the interior of $B_r(p)$ there is a smooth, positive Green's function G_x for the Dirichlet-Laplace problem on $B_r(p)$ with singularity at x, i.e. $\Delta G_x(y) = \delta_x(y)$ and $G_x(\partial B_r(p)) = 0$. If W^n is a Ricci non-negative ALX_k space with effective dimension $n-k > 2$, then for any $\epsilon > 0$ there is a distance D such that whenever $d(x,y) > D$ the Green's function is bounded by*

$$G_x(y) \leq \frac{(1+\epsilon)\, c}{\mathrm{vol}\, X\; d(x,y)^{n-k-2}} ,$$

where the constant c depends only on n.

Proof. The existence of a positive Green's function on compact, connected manifolds with boundary is standard. The bound follows immediately from Theorem 5.2 in Li-Yau's seminal work [LY86]. Their theorem states

$$G_x(y) \leq c \int_{r^2}^{\infty} \frac{1}{\operatorname{vol} B_{\sqrt{t}}(x)} dt ,$$

where $r = d(x, y)$ denotes the Riemannian distance between x and y and the constant c depends only on n.

Let $\epsilon > 0$. By Proposition 3.2 there is a distance $R \geq 0$, such that whenever $r \geq R$ we find

$$G_x(y) \leq c \int_{r^2}^{\infty} \frac{1+\epsilon}{\operatorname{vol} X \, t^{(n-k)/2}} dt = \frac{(1+\epsilon)\, c}{\operatorname{vol} X \, r^{n-k-2}} .$$

□

3.2 The Frequency Function

As explained in the introduction, we are interested in the growth of the function κ defined in (3.2) and which captures the L^2-mean of the Higgs field ϕ on large geodesic spheres. On our way to show that κ must have some minimal asymptotic growth, the first step is to realize that for large enough radii it can only decrease at an arbitrarily small rate. To make this more precise, let us introduce the following notation.

Definition 3.8. For $\epsilon > 0$ we say κ is *ϵ-almost non-decreasing* whenever its derivative is bounded from below according to $\frac{d\kappa}{dr} \geq -\frac{\epsilon\kappa}{r}$. Furthermore, in this context κ is called *asymptotically almost non-decreasing* if for any $\epsilon > 0$ there is some (large) radius R, such that κ is ϵ-almost non-decreasing on $[R, \infty)$, i.e. its derivative satisfies $\frac{d\kappa}{dr} \geq -\frac{\epsilon\kappa}{r}$ for all $r \geq R$.

To learn more about the potential growth and decay rates of κ, the upcoming proposition establishes its derivative and provides first estimates in the limit of small and large radius, respectively. The function $N(r)$ that arises in that context is an analogue of the frequency function as introduced by Almgren [Alm79] and we will refer to it by that name. The function $D(r)$ will be called mean curvature deviation, as it is related to the average deviation of the mean curvature of the geodesic sphere from its limit at infinity.

Proposition 3.9. *Assume the pair (A, ϕ) satisfies (3.1). Whenever κ is non-zero its derivative is*

$$\frac{d\kappa}{dr} = \frac{(N+D)\,\kappa}{r} ,$$

where N and D are given by

$$N(r) = \frac{1}{r^{n-k-2}\kappa^2}\int_{B_r}\left(\left\|\nabla^A\phi\right\|^2 + \|[\phi\wedge\phi]\|^2 + \langle \mathrm{Ric}\,\phi,\phi\rangle\right) \tag{3.4}$$

$$D(r) = \frac{1}{2r^{n-k-2}\kappa^2}\int_{\partial B_r}\left(\Delta r - \frac{n-k-1}{r}\right)\|\phi\|^2\ . \tag{3.5}$$

Moreover, if $\mathrm{Ric}\geq 0$*, then N is non-negative, D is bounded from above by* $k/2$*,* $\lim_{r\to 0} D = k/2$*, and* $\lim_{r\to\infty} D = 0$*. As a consequence, if* κ *is not identically zero near* $r=0$*, then it is increasing on small enough neighbourhoods of* 0*. Similarly, if* κ *is not asymptotically zero as* $r\to\infty$*, then it is asymptotically almost non-decreasing.*

Proof. The derivative of κ can be calculated from $\frac{d\kappa}{dr} = \frac{1}{2\kappa}\frac{d}{dr}\kappa^2$. Thus, denote by X the radial unit vector field on B_r and observe that

$$\kappa^2(r) = \frac{1}{r^{n-k-1}}\int_{\partial B_r}\|\phi\|^2 = \frac{1}{r^{n-k-1}}\int_{B_r}\mathcal{L}_X\ \|\phi\|^2\ .$$

By the product and Leibniz' integral rule the derivative is then given by

$$\frac{d}{dr}\kappa^2(r) = -\frac{n-k-1}{r}\kappa^2 + \frac{1}{r^{n-k-1}}\int_{B_r}\mathcal{L}_X\circ\mathcal{L}_X\ \|\phi\|^2\ .$$

We can write the integral on the right hand side equivalently as an integral over the trace of the (asymmetric) second Lie derivative $\mathcal{L}^2_{Y,Z} := \mathcal{L}_Y\circ\mathcal{L}_Z$. To see this denote by (r,θ_i) polar normal coordinates on B_r and note that in these coordinates the metric is block-diagonal, i.e. $g = dr^2 + g_{S^{n-1}}$. Since for any top-form ω the pullback of $\imath_{\partial_{\theta_i}}\omega$ to the boundary of the geodesic ball is zero, one finds

$$\int_{B_r}\mathcal{L}^2_{X,X}\|\phi\|^2 = \int_{B_r}(\mathcal{L}^2_{X,X} + g^{ij}_{S^{n-1}}\mathcal{L}^2_{\partial_{\theta_i},\partial_{\theta_j}})\|\phi\|^2 = \int_{B_r}\mathrm{tr}_{TM}\mathcal{L}^2\ \|\phi\|^2\ .$$

Next, for any vector field Y and top-form ω we may express the action of the Lie derivative in terms of the Levi-Civita connection as $\mathcal{L}_Y\omega = \nabla_Y\omega + \mathrm{div}\,Y\ \omega$. Using this we can write the second Lie derivative as

$$\mathcal{L}^2_{Y,Z}\ \|\phi\|^2 = (\nabla_Y + \mathrm{div}\,Y)\,\nabla_Z\|\phi\|^2 + \mathcal{L}_Y(\mathrm{div}\,Z\ \ \|\phi\|^2)\,.$$

Furthermore, we use ad-invariance and metric compatibility to write $\nabla_Y\langle\phi,\phi\rangle = 2\langle\phi,\nabla^A_Y\phi\rangle$, and use that the formal adjoint is given by $\nabla^A_Y + \mathrm{div}\,Y = -(\nabla^A_Y)^\dagger$. This leads to

$$\begin{aligned}\int_{B_r}\mathrm{tr}_{TM}\mathcal{L}^2\ \|\phi\|^2 &= 2\int_{B_r}\left(\left\|\nabla^A\phi\right\|^2 - \langle\phi,\nabla^{A^\dagger}\nabla^A\phi\rangle\right) + \int_{B_r}\mathrm{tr}_{TM}\left(\mathcal{L}_\cdot\,\mathrm{div}(\cdot)\ \|\phi\|^2\right)\\ &= 2\int_{B_r}\left(\left\|\nabla^A\phi\right\|^2 + \|[\phi\wedge\phi]\|^2 + \langle\phi,\mathrm{Ric}\,\phi\rangle\right) + \int_{\partial B_r}\Delta r\ \ \|\phi\|^2\ ,\end{aligned}$$

where we used the second order differential equation (3.1) in the first term and that the only non-zero contribution in the second term contains the mean curvature of the geodesic sphere since $\mathrm{div}\,X = \Delta r$.

All in all, as long as $\kappa \neq 0$, the derivative is given by

$$\frac{d\kappa}{dr} = \frac{1}{\kappa r^{n-k-1}} \int_{B_r} \left(\left\| \nabla^A \phi \right\|^2 + \| [\phi \wedge \phi] \|^2 + \langle \phi, \operatorname{Ric} \phi \rangle \right) + \frac{1}{2\kappa r^{n-k-1}} \int_{\partial B_r} \left(\Delta r - \tfrac{n-k-1}{r} \right) \| \phi \|^2 \ ,$$

which yields the desired result upon identification of the terms on the right hand side with the definitions of N and D in (3.4) and (3.5), respectively.

Now assume $\operatorname{Ric} \geq 0$. On the one hand, N is then clearly non-negative. On the other hand, the results for Δr from Proposition 3.5 immediately provide both the global upper bound and the limits of D.

Combining these facts with the formula for $\frac{d\kappa}{dr}$ leads to the conclusion that κ is (almost) non-decreasing at both ends: Since D is continuous and $\lim_{r \to 0} D = k/2$, D must be positive on some small interval $[0, s)$. Thus, if κ is non-zero somewhere in that interval then it is increasing. The asymptotic bound for $r \to \infty$ works out similarly, where according to Proposition 3.5 for any $\delta > 0$ there is an interval $[R, \infty)$ on which

$$\Delta r \geq \frac{1}{1+\delta} \frac{n-k-1}{r} \ .$$

After a suitable choice of δ this yields the desired bound $\frac{d\kappa}{dr} \geq -\frac{\epsilon\kappa}{r}$ for any $\epsilon > 0$, which concludes the proof. □

In the preceding proposition we already encountered lower bounds for $\frac{d\kappa}{dr}$ near $r = 0$ and $r \to \infty$. But when we keep track of N it becomes clear that $\frac{d\kappa}{dr}$ satisfies stronger bounds than recorded so far. This is the content of the following two corollaries. The first records a global growth limitation, while the second determines asymptotic lower and upper bounds, both in dependence of the frequency function N.

Corollary 3.10. *Assume $\kappa \neq 0$ on $[r_0, r_1]$, then*

$$\kappa(r_1) \ \leq \ \kappa(r_0) \ \exp \int_{r_0}^{r_1} \frac{N(t) + k/2}{t} dt \ .$$

Proof. Since $\kappa \neq 0$, we can rely on Proposition 3.9. In particular, using the global bound $D \leq k/2$, the derivative of κ is bounded by $\frac{d\kappa}{dr} \leq \frac{(N+k/2)}{r}\kappa$. By Grönwall's inequality, κ then can't become larger than a solution of the underlying differential equation, which is the stated bound. □

Corollary 3.11. *Let $\epsilon > 0$ and R be such that $|D| \leq \epsilon$ on $[R, \infty)$. If κ has no zeroes in $[r_0, r_1] \subset [R, \infty)$ then it is bounded at r_1 from both sides as follows*

$$\kappa(r_0) \ \exp \int_{r_0}^{r_1} \frac{N(t) - \epsilon}{t} dt \ \leq \ \kappa(r_1) \ \leq \ \kappa(r_0) \ \exp \int_{r_0}^{r_1} \frac{N(t) + \epsilon}{t} dt \ . \tag{3.6}$$

Consequently, bounds for the frequency function on $[r_0, r_1]$ have the following effect:

- *if* $a \le N$ *then* $\kappa(r_1) \ge \left(\frac{r_1}{r_0}\right)^{a-\epsilon} \kappa(r_0)$
- *if* $N \le b$ *then* $\kappa(r_1) \le \left(\frac{r_1}{r_0}\right)^{b+\epsilon} \kappa(r_0)$

Proof. Since $|D| \le \epsilon$ and κ is non-zero on $[r_0, r_1]$, its derivative is bounded in both directions as follows

$$\frac{(N-\epsilon)\,\kappa}{r} \le \frac{d\kappa}{dr} \le \frac{(N+\epsilon)\,\kappa}{r} \,.$$

Grönwall's inequality states that κ is then bounded in either direction by the corresponding solutions of the underlying differential equations, which is exactly (3.6).

Clearly, if the frequency function is bounded below by some $a \ge 0$ the first inequality in (3.6) reduces to

$$\kappa(r_1) \ge \kappa(r_0)\ \exp \int_{r_0}^{r_1} \frac{a-\epsilon}{s}\, ds = \left(\frac{r_1}{r_0}\right)^{a-\epsilon} \kappa(r_0)\,.$$

The same argument, but based on the second inequality, yields the corresponding upper bound if N is bounded from above. □

We will later also need the derivative of N, which is given directly by the product and Leibniz' integral rule.

$$\frac{d}{dr}N = \frac{1}{r^{n-k-2}\kappa^2} \int_{\partial B_r} \left(\|\nabla^A \phi\|^2 + \|[\phi \wedge \phi]\|^2 + \langle \mathrm{Ric}\, \phi, \phi \rangle \right) - \big(n-k-2+2(N+D)\big)\frac{N}{r} \quad (3.7)$$

As an immediate consequence, we see that if N is ever small, then it can't have been very much larger at nearby smaller radii $s < r$. This observation is recorded more precisely in the following proposition.

Proposition 3.12. *Assume* $\mathrm{Ric} \ge 0$. *If* $N \le 1$ *on some interval* $[r_0, r_1]$, *then* $N(r_0) \le \left(\frac{r_1}{r_0}\right)^n N(r_1)$. *Moreover, whenever* $N(r) < 1$ *at some* $r \in (0, \infty)$ *then* $N \le 1$ *on the interval* $[N(r)^{1/n}\, r, r]$.

Proof. Since Ric is non-negative the same is true for the first term in (3.7). Moreover, in that case $D \le \frac{k}{2}$ by Proposition 3.9. Assume now that $N \le 1$ on all of $[r_0, r_1]$. Then N satisfies the following differential inequality for any $r \in [r_0, r_1]$

$$\frac{dN}{dr} \ge -\frac{n\,N}{r}\,.$$

Grönwall's inequality states that then for any pair $s \le r$ in $[r_0, r_1]$ the following inequality holds

$$N(r) \ge \left(\frac{s}{r}\right)^n N(s)\,,$$

which proves the first part of the statement.

Now assume $N(r) < 1$ for some $r \in (0, \infty)$. By continuity, r must be contained in some interval $[r_0, r_1]$ on which $N \le 1$, so we can use the preceding inequality in the form $N(s) \le (r/s)^n N(r)$. The right hand side is less than or equal to 1 as long as $s \ge N(r)^{1/n}\, r$. □

3.3 Asymptotically Unique Continuation

In this section, we observe that unless κ is the zero function, it cannot have zeroes in (R, ∞) for some large enough radius R. This rests on Aronszajn's unique continuation theorem, which states that a non-trivial solution of an elliptic second order differential equation cannot exhibit zeroes of infinite order [Aro57]. As it turns out, Aronszajn's theorem applies[1] to the elliptic differential equation (3.1). This was already observed by Taubes in several closely related situations, see e.g. [Tau13, Sec. 3a], [Tau17a, Sec. 2c], and [Tau18, Sec. 2e]. It follows that whenever the Higgs field ϕ vanishes on an open subset of M, then it must be identically zero. This, in turn, implies that whenever κ is non-trivial it can't vanish on an open subset of $[0, \infty)$.

Now note that any real function on $[0, \infty)$ that is both non-negative and non-decreasing has the following property: if it is non-zero at a particular point r_0, it will remain non-zero for any subsequent point $r > r_0$. In combination with a unique continuation property, such a function is then clearly either identically zero or strictly positive on all of $(0, \infty)$. Although κ is not non-decreasing, there is some large radius R beyond which it is 'non-decreasing enough' to draw the same conclusion on $[R, \infty)$.

To make this precise, recall that κ is continuous, non-negative, and ϵ-almost non-decreasing at large radius. With respect to the last property, fix some $\epsilon > 0$ with associated radius $R \geq 0$. Assume $\kappa(r_0) \neq 0$ for some $r_0 \in [R, \infty)$. Since κ is ϵ-almost non-decreasing, the first bullet of Corollary 3.11 with respect to the bound $0 \leq N$ provides a strictly positive lower bound for any larger radius $r \geq r_0$

$$\kappa(r) \geq \left(\frac{r_0}{r}\right)^{\epsilon} \kappa(r_0) \,,$$

which prevents κ from vanishing at any larger radius, at least as along as $r_0 \neq 0$. Note that, if the mean curvature deviation D as defined in equation (3.5) is bounded from below, then there is a possibly large choice of ϵ for which $R = 0$, such that the conclusion holds for any $r_0 \in (0, \infty)$.

Now assume there is a $t > R$ at which $\kappa(t) = 0$. According to the result above, this can only be the case if κ vanishes on the interval $[R, t]$. It follows that ϕ vanishes on the open set $B_t(p) \setminus B_R(p)$ and consequently is identically zero due to the unique continuation property of (3.1). In conclusion, we arrive at the following result.

Lemma 3.13. *Assume* $\mathrm{Ric} \geq 0$ *and that sectional curvature is bounded from below. There is a radius* $R \geq 0$*, such that* κ *is either strictly positive on* (R, ∞) *or vanishes on all of* $[0, \infty)$*. Moreover, one can choose* $R = 0$ *if the mean curvature deviation* D *is bounded from below.*

[1] Equation (3.1) is non-linear, while Aronszajn's theorem pertains to linear differential equations. Assuming, as we do, that ϕ is regular enough, say C^3, we may pass to the linearization of (3.1) and deduce the desired unique continuation property from there.

Remark. It is also instructive to investigate the behaviour of κ near one of its zeroes $\kappa(t) = 0$, where the growth rate of κ is controlled by the frequency function N and the mean curvature deviation D. Observe that if $t \neq 0$ and N is bounded from above on $(t, r_0]$, then the assumptions $\kappa(t) = 0$ and $\kappa(r_0) \neq 0$ lead to a contradiction. To see this use Corollary 3.10 with the upper bound $N \leq b$, which yields

$$\kappa(r) \geq \left(\frac{r}{r_0}\right)^{b+k/2} \kappa(r_0) \quad \text{for all } r \in (t, r_0]\,.$$

The right hand side is strictly positive and prevents κ from vanishing at r_0. Thus, a version of Lemma 3.13 also follows from a proof that N is a priori bounded on any interval $(t, r_0]$ in which κ does not have zeroes.

3.4 Slow Growth and Bounded Frequency

We have previously seen that an upper bound for N leads to bounded growth of κ. The goal of this section is to show that the converse is true when $r \to \infty$. More precisely, we show that whenever κ grows slower than $\mathcal{O}(r^a)$ between two large radii, then N must have been bounded from above on an interval leading up to the violation. Note that such violations must occur for arbitrarily large radii when κ is not asymptotically bounded below by r^a, which is the situation of the second alternative in Theorem A. Accordingly, the upcoming lemma and its refinement in Section 3.5 play a crucial role in the proof of the main theorem.

Lemma 3.14. *Assume κ is not asymptotically zero. Fix an $\epsilon > 0$ and denote by R the radius beyond which $|D| \leq \epsilon$. If there is a pair of radii $r_0 \leq r_1$ in $[R, \infty)$ and an $a \geq 0$ such that $\kappa(r_1) \leq \left(\frac{r_1}{r_0}\right)^{a-\epsilon} \kappa(r_0)$, then there exists a radius $t \in [r_0, r_1]$ such that*

(i) $N(t) \leq a$.

Moreover, if $a < 1$ the following holds on the interval $[\tilde{R}, t]$, where $\tilde{R} = \max\big(a^{\frac{1}{2n}}\, t,\ R\big)$.

(ii) $N < \sqrt{a}$,

(iii) $\kappa \geq a^{\frac{\sqrt{a}+\epsilon}{2n}} \kappa(t)$.

Proof. To see *(i)* assume to the contrary that $N > a$ on all of $[r_0, r_1]$. Then the first bullet in Corollary 3.11 states that

$$\kappa(r_1) > \left(\frac{r_1}{r_0}\right)^{a-\epsilon} \kappa(r_0)\,,$$

which violates the assumption that $\kappa(r_1)$ satisfies exactly the opposite inequality. Hence, there is a $t \in [r_0, r_1]$ at which $N(t) \leq a$.

For (*ii*), assume that $a < 1$ and note that then the same is true for $N(t)$. In that case we conclude via the second part of Proposition 3.12 that $N \leq 1$ on $[N(r)^{1/n}\, t, t]$. Since $N(t) \leq a < \sqrt{a}$, this interval contains as a subinterval $[a^{1/2n}\,t, t]$. Then the first part of Proposition 3.12 for $s \in [a^{1/2n}\,t, t]$ yields

$$N(s) \le \left(\frac{r}{s}\right)^{n-k} N(t) \le \frac{1}{\sqrt{a}} N(t) \le \sqrt{a}\,,$$

which proves that the same is true on the possibly smaller interval $[\tilde{R}, t]$ where the lower bound is determined by $\tilde{R} := \max\big(a^{1/2n}\,t, R\big)$.

Finally, (*iii*) follows from $N \le \sqrt{a}$ on $[\tilde{R}, t]$ and the second bullet of Corollary 3.11, which provides the bound

$$\kappa(t) \le \left(\frac{r}{t}\right)^{\sqrt{a}+\epsilon} \kappa(r) \le a^{-\frac{\sqrt{a}+\epsilon}{2n}} \kappa(r)\,,$$

where in the last step we used $a^{1/2n}\,t \le \tilde{R} \le r$ for any $r \in [\tilde{R}, t]$. □

3.5 The Correlation Tensor

There is an $\Omega^1_p \otimes \Omega^1_p$-valued function T that refines κ^2 to the effect that it resolves the behaviour of the components of ϕ. Note that, being a one-form, ϕ can be evaluated in particular on the unit vector field on $B_r(p)$ that is defined by parallel transport of a unit vector $v \in T_pW$ along radial geodesics. The output is a smooth function ϕ_v on $B_r(p) \setminus \mathrm{Cut}(p)$ that captures the evolution of the v-component of ϕ along the geodesics emanating from p. This allows the definition of what we will call the correlation tensor $T : W^4 \times (0,\infty) \to \Omega^1_p \otimes \Omega^1_p$. Its value on $v, w \in T_pW$ is defined by

$$T(p,r)(v,w) = \frac{1}{r^{n-k-1}} \int_{\partial B_r(p)} \langle \phi_v, \phi_w \rangle\,. \tag{3.8}$$

Note, in particular, that $\mathrm{tr}_{T_pW}\, T(p,r) = \kappa^2(p,r)$, while the induced quadratic form $\kappa_v^2 := T(v,v)$ returns a version of the function κ^2 that is based on the component ϕ_v. Just like $\kappa_v(p,r)$ has an interpretation as the mean value of ϕ_v on the geodesic sphere, the value of $T(v,w)$ measures the correlation between the components ϕ_v and ϕ_w on the geodesic sphere; hence its name.

An important observation is that κ_v satisfies essentially the same properties as κ. In fact, if ϕ satisfies the main assumption (3.1), then ϕ_v satisfies the following analogous second order equation

$$\nabla^{A\dagger}\nabla^{A}\phi_v + \frac{1}{2} \star [\star[\phi \wedge \phi_v] \wedge \phi] + \mathrm{Ric}(\phi)(v) + \mathrm{tr}_{TM} \nabla^{A\dagger}\left(\phi(\nabla v)\right) = 0\,. \tag{3.9}$$

As a consequence, the derivative of κ_v comes with its own version of the frequency function and mean curvature deviation, denoted N_v and D_v.

$$\frac{d\kappa_v}{dr} = \frac{N_v + D_v}{r} \kappa_v$$

The functions N_v and D_V are given by essentially the same expressions as before, but with ϕ replaced by ϕ_v as follows.

$$N_v = \frac{1}{r^{n-k-2}\kappa_v^2} \int_{B_r} \left(\|\nabla^A \phi_v\|^2 + \|[\phi \wedge \phi_v]\|^2 + \langle \mathrm{Ric}(\phi)(v), \phi_v \rangle \right) \tag{3.10}$$

$$D_v = \frac{1}{r^{n-k-1}\kappa_v^2} \int_{\partial B_r} \left(\Delta r - \frac{n-k-1}{r} \right) \|\phi_v\|^2 + \mathrm{tr}_{T\partial B_r(p)} \left((\nabla^A \phi)(\nabla v) \right) + \phi(\nabla^\dagger \nabla v)$$

Proposition 3.15. *Let $v \in T_pW$. All previous results hold verbatim when we replace κ and N by κ_v and N_v, respectively.*

There are in fact analogous results for the correlation tensor T. To see this define the (Frobenius) norm of T with respect to the inner product on $\Omega^1_p \otimes \Omega^1_p$ induced by the metric, i.e.

$$\|T\|^2 = g^{\mu\rho} g^{\nu\sigma} T_{\mu\nu} T_{\rho\sigma} \, .$$

In this expression the metric is evaluated at the point p. The norm of T satisfies $\frac{1}{c}\kappa^2 \leq \|T\| \leq \kappa^2$ for some constant c. The tensor T is differentiable with respect to r and there is then a (possibly larger) c such that the following inequality holds.

$$\left\| \frac{dT}{dr} \right\| \leq c \, \frac{N + D}{r} \|T\|$$

Below we will also use the notation $T' := \frac{dT}{dr}$.

From now on view T as linear map from $T_pW \to T_pW$. If T has zero as an eigenvalue at some radius r and v denotes the associated eigenvector, then $\kappa_v^2(r) = 0$. As a consequence of the unique continuation property in Lemma 3.13, there is an $R > 0$ such that κ_v must be identically zero if r is larger than R. This in turn implies that the component ϕ_v vanishes on all of W^n. Therefore, in the definition of T we will restrict ourselves to the subspace of T_pW that is orthogonal to the zero eigenspace of T at infinity. The correlation tensor then has strictly positive eigenvalues on an interval (R, ∞) for some large enough R.

Denote by $\lambda : (0, \infty) \to \mathbb{R}$ the function that assigns to a radius r the smallest eigenvalue of $T(r)$. By the preceding paragraph we may assume that this function is non-zero on all of (R, ∞). Clearly $\lambda \leq \kappa^2$, as the latter is the trace of T, so it should be noted that under this assumption κ is necessarily non-zero on (R, ∞).

Since the eigenvalues of a differentiable one-parameter family of matrices may coincide and cross each other, the smallest eigenvalue function is generally not differentiable, see e.g. [Kat95, §6.4]. In absence of exact differentiability, the following proposition provides estimates for the finite differences of λ and in effect states that λ is nearly differentiable.

Proposition 3.16. *The finite differences of the smallest-eigenvalue function λ at any $r \in (0, \infty)$ satisfy the following bounds: Denote by v a unit length eigenvector of $T(r)$ with eigenvalue $\lambda(r)$ and let $\Delta > 0$.*

$$\lambda(r + \Delta) - \lambda(r) \leq \langle v, T'(r)\, v\rangle\, \Delta + \mathcal{O}(\Delta^2)$$

$$\lambda(r) - \lambda(r - \Delta) \geq \langle v, T'(r)\, v\rangle\, \Delta + \mathcal{O}(\Delta^2)$$

Moreover, λ is locally Lipschitz continuous on $(0, \infty)$.

Proof. For any $r \in (0, \infty)$ denote by v_r an eigenvector of $T(r)$ with eigenvalue $\lambda(r)$. For $\Delta \geq 0$ we then have

$$\lambda(r + \Delta) - \lambda(r) = \langle v_{r+\Delta}, T(r + \Delta)v_{r+\Delta}\rangle - \langle v_r, T(r)v_r\rangle \leq \langle v_r, \big(T(r + \Delta) - T(r)\big)v_r\rangle\ .$$

To get the upper bound on the right hand side note that the smallest eigenvalue at $r + \Delta$ is $\lambda(r + \Delta)$, so the evaluation of $T(r + \Delta)$ on v_r can only ever result in the same or larger values. The upper bound for the forwards finite difference is then a consequence of Taylor's theorem. Up to a minus sign the lower bound for the backwards finite difference in the second inequality follows in exactly the same way.

The Lipschitz property of λ follows by paying closer attention to the remainder in Taylor's theorem. Consider a pair $s < r \in (0, \infty)$ and Taylor's theorem at zeroth order

$$T(r) = T(s) + R_0(r)\ .$$

A standard estimate for the remainder states $\|R_0(r)\| \leq \sup_{(r,s)} \|T'\|\ (r - s)$, as long as the derivative is bounded on the given interval. As mentioned before, the derivative of T at any radius $\tilde{r}$ is bounded by a multiple of $\frac{N+D}{\tilde{r}}\ \|T(\tilde{r})\|$. Recall from Proposition 3.9 that $D \leq \frac{k}{2}$. Also, as was discussed in the context of the unique continuation property, N is bounded on any compact interval on which κ is non-zero. It follows that λ is Lipschitz on any compact $[s, t] \subset (0, \infty)$, since

$$|\lambda(r) - \lambda(s)| \leq \left|\langle v_s, \big(T(r) - T(s)\big)v_s\rangle\right| \leq c\,|r - s|\ .$$

Equivalently (since $(0, \infty)$ is locally compact): T is locally Lipschitz. □

We can use the preceding proposition to show that, similar to κ, the function λ is asymptotically almost non-decreasing. The idea here is that λ consists of piecewise smooth segments, each coinciding with a function κ_v^2 for some v. Any κ_v is eventually ϵ-almost non-decreasing, so an analogues statement must be true for λ. Observe from the proof of Proposition 3.9 that asymptotic bounds for the mean curvature deviation D_v only depend on the asymptotics of the mean curvature Δr, which is independent of v. Thus, if D is bounded from below by $-\epsilon$, then the same is true for D_v.

Hence, let R be such that $|D| \leq \epsilon$ for any larger radius and consider points $r > s \geq R$. Define an equidistant partition of the interval $[s, r] = \bigcup_{k=0}^{M-1} [r_k, r_{k+1}]$ where $r_k = s + \frac{k}{M}(r - s)$. Then, denoting by v_k an eigenvector with eigenvalue $\lambda(r_k)$, the two estimates in Proposition 3.16 imply

$$\lambda(r) = \lambda(s) + \sum_{k=1}^{M} \left(\langle v_k, T'(r_k)\, v_k \rangle \frac{r-s}{M} + \mathcal{O}\left(\frac{1}{M^2}\right) \right).$$

This can be rewritten by using $\langle v, T'v \rangle = \frac{d}{dr}\kappa_v^2$ and the fact that each v_k is an eigenvector with the smallest eigenvalue at r_k, i.e. $\kappa_{v_k}^2(r_k) = \lambda(r_k)$.

$$\lambda(r) = \lambda(s) + 2\sum_{k=1}^{M} \lambda(r_k) \frac{\left(N_{v_k}(r_k) + D_{v_k}(r_k)\right)}{r_k} \frac{r-s}{M} + \mathcal{O}\left(\frac{1}{M}\right) \tag{3.11}$$

Since $N_{v_k} \geq 0$ and $\left|D_{v_k}\right| \leq \epsilon$ we find that λ satisfies the following inquality on $[R, \infty)$

$$\frac{\lambda(r) - \lambda(s)}{r-s} \geq -\frac{2\epsilon}{M} \sum_{k=1}^{M} \frac{\lambda(r_k)}{r_k} + \mathcal{O}\left(\frac{1}{M}\right).$$

This is the finite difference analogue of the differential inequality satisfied by κ when it is ϵ-almost non-decreasing. We will correspondingly say that λ is ϵ-almost non-decreasing when it satisfies this inequality.

Remark. To make the relation to an ϵ-almost non-decreasing function as given in Definition 3.8 more apparent, observe that the right hand side of the last inequality contains the arithmetic mean of the function λ/r for the given partition of $[s, r]$. If one takes $M \to \infty$ this expression approaches the mean value of λ/r on the given interval. Seeing that λ is continuous, it is then a consequence of the (integral) mean value theorem that there is a point $t \in [s, r]$ at which the right hand side is given by $\lambda(t)/t$, such that

$$\frac{\lambda(r) - \lambda(s)}{r-s} \geq -\frac{2\epsilon\lambda(t)}{t}.$$

If λ is differentiable at r, the limit $s \to r$ exactly recovers the differential inequality of a 2ϵ-almost non-decreasing function.

The fact that λ is asymptotically ϵ-almost non-decreasing allows us to extend Lemma 3.14 in such a way that it also provides bounds for N_v and κ_v, where v is a unit eigenvector of $T(t)$ associated to the smallest eigenvalue $\lambda(t)$ at some distinguished $t \in [r_0, r_1]$. This is the content of the next lemma.

Lemma 3.17. *Fix $\epsilon > 0$ and denote by R the radius beyond which $|D| \leq \epsilon$. Let $r_0, r_1 \in [R, \infty)$ be a pair of radii such that r_1 is larger than any of $4r_0$, $\left(\kappa^2(r_0)/\lambda(r_0)\right)^{\frac{1}{2a}} r_0$, and $r_0^{1/(1-100\sqrt{a})}$. If $\kappa(r_1) \leq \left(\frac{r_1}{r_0}\right)^{a-\epsilon} \kappa(r_0)$ and a is sufficiently small (e.g. $a^{1/4n} < 0.1$), then there exists a radius $t \in [r_1^{1-100\sqrt{a}}, r_1]$ of the following significance:*

Let v be an eigenvector of $T(t)$ associated to the smallest eigenvalue $\lambda(t)$ and set $\tilde{R} := \max\left(a^{\frac{1}{8n}} t, R\right)$. On all of $[\tilde{R}, t]$

(i) $N \leq a^{1/4}$ *and* $\kappa \geq a^{\frac{a^{1/4+\epsilon}}{4n}} \kappa(t)$

(ii) $N_v \leq a^{1/4}$ *and* $\kappa_v \geq a^{\frac{a^{1/4+\epsilon}}{4n}} \kappa_v(t)$

Proof. Let $\epsilon > 0$, denote by R the radius beyond which $|D| \leq \epsilon$, and assume $\sqrt{a} < 1$. The proof proceeds by establishing the existence of regions in $[r_0, r_1]$ where both N and N_v are less or equal to $\sqrt{a}$ at the same time. The stated bounds then follow from Lemma 3.14 (with a replaced by $\sqrt{a}$ everywhere).

First, regarding the condition $N \leq \sqrt{a}$, observe that the set

$$I = \left\{ r \in [\log r_0, \log r_1] \mid N(\exp r) > \sqrt{a} \right\}$$

makes up at most $\sqrt{a}$ of the length of the surrounding interval. To see this, write $I = \coprod(a_k, b_k)$ and then go from $\kappa(r_0)$ to $\kappa(r_1)$ by iteratively using Corollary 3.11, with bounds $N > \sqrt{a}$ on each (a_k, b_k) and $N \geq 0$ on the intervals $[b_k, a_{k+1}]$ in between. This leads to

$$\kappa(r_1) \geq \left(\frac{r_1}{r_0}\right)^{-\epsilon} \prod_k \left(\frac{b_k}{a_k}\right)^{\sqrt{a}} \kappa(r_0)\,.$$

This inequality is only compatible with the assumption that $\kappa(r_1) \leq (r_1/r_0)^{a-\epsilon} \kappa(r_0)$ if

$$\sum_k (\log b_k - \log a_k) \leq \sqrt{a}\,(\log r_1 - \log r_0)\,.$$

Equivalently, if $|I| \leq \sqrt{a}\,|[\log r_0, \log r_1]|$.

An analogous statement for points that satisfy the condition $N_v \leq \sqrt{a}$ makes use of a slightly longer argument. As before, we set out to investigate the measure of the set on which $N_v > \sqrt{a}$. Denote by L the largest integer such that $2^L r_0 < r_1$ and consider the sequence $\{2^\ell r_0\}_{\ell=0,1,\ldots,L-1}$. For each pair of neighbours in this sequence we can apply equation (3.11) in the form

$$\lambda(2s) \geq \lambda(s) + \frac{1}{M} \sum_{k=1}^{M} \lambda(s_k)\left(N_{v_k}(r_k) - \epsilon\right) + \mathcal{O}\left(\frac{1}{M}\right)\,, \tag{3.12}$$

where $s_k = (1 + k/M)\,s$ and v_k denotes an eigenvector of $T(s_k)$ with eigenvalue $\lambda(s_k)$. Next, note that Corollary 3.11 with the lower bound $N \geq 0$ provides

$$\lambda(s_k) = \kappa_{v_k}^2(s_k) \geq \left(\frac{s_k}{s}\right)^{-2\epsilon} \kappa_{v_k}^2(s) \geq \frac{1}{4}\lambda(s)\,.$$

Here the last inequality holds by virtue of $s_k/s < 2$ and because $\lambda(s)$ is the smallest eigenvalue of T at s, such that $\kappa_{v_k}^2(s) = T(s)(v_k, v_k)$ can't be smaller. Also introduce for the moment the notation $\langle N_v \rangle_\ell$ for the average of $\{N_{v_k}(s_k)\}_{k=1,\ldots,M}$ in the interval $[2^\ell r_0, 2^{\ell+1} r_0]$. The inequality in (3.12) becomes

$$\frac{\lambda(2^{\ell+1} r_0)}{\lambda(2^\ell r_0)} \geq 1 + \frac{1}{4}\Big(\langle N_v \rangle_\ell - \epsilon\Big) + \mathcal{O}\left(\frac{1}{M}\right)\,.$$

In the last of these pairings we can use r_1 as endpoint of the interval instead of $2^L r_0$ without changing the inequality, since $r_1 - 2^{L-1}r_0 > 2^{L-1}r_0$ and $r_1 \leq 2^{L+1}r_0$. Hence, the product of all these ratios yields

$$\frac{\lambda(r_1)}{\lambda(r_0)} \geq \prod_{\ell=0}^{L-1}\left(1 + \frac{\langle N_v\rangle_\ell - \epsilon}{4} + \mathcal{O}\left(\frac{1}{M}\right)\right) .$$

Let f denote the fraction of $\{0, 1, \ldots, L-1\}$ for which $\langle N_v\rangle_\ell + \mathcal{O}(1/M) > \sqrt{a}$. Then the product simplifies to

$$\frac{\lambda(r_1)}{\lambda(r_0)} \geq \left(1 + \frac{\sqrt{a} - \epsilon}{4}\right)^{fL}\left(1 - \frac{\epsilon}{4} + \mathcal{O}\left(\frac{1}{M}\right)\right)^{(1-f)L} .$$

On the left use that $\lambda(r_1) \leq \kappa^2(r_1) \leq \left(\frac{r_1}{r_0}\right)^{2(a-\epsilon)}\kappa^2(r_0)$ to replace $\lambda(r_1)$. Upon taking logarithms on both sides the resulting inequality reads

$$2(a-\epsilon)\log\left(\frac{r_1}{r_0}\right) + \log\left(\frac{\kappa^2(r_0)}{\lambda(r_0)}\right) \geq fL\log\left(1 + \frac{\sqrt{a}-\epsilon}{4}\right) + (1-f)L\log\left(1 - \frac{\epsilon}{4} + \mathcal{O}\left(\frac{1}{M}\right)\right) .$$

This can be simplified in several ways: (1) Because $2^{L+1}r_0 > r_1$ and as long as we make sure that $r_1 \geq 4r_0$ we can assume $L > \frac{1}{2}\log r_1/r_0$. (2) Since $\sqrt{a} - \epsilon < 1$ the term $\log\left(1 + (\sqrt{a}-\epsilon)/4\right)$ is no smaller than $(\sqrt{a}-\epsilon)/8$. (3) By taking M large, we can make sure that $\log(1 - \epsilon/4 + \mathcal{O}(1/M))$ is larger than $-\epsilon$. Plugging everything in and solving for f leads to the following upper bound

$$f < 32\sqrt{a}\left(1 + \frac{\log\left(\kappa^2(r_0)/\lambda(r_0)\right)}{2a\log(r_1/r_0)}\right) .$$

In particular, if $r_1 > \left(\kappa^2(r_0)/\lambda(r_0)\right)^{\frac{1}{2a}} r_0$ we find that $f < 64\sqrt{a}$.

Consider the subset $J \subset [\log r_0, \log r_1]$ that consists of those intervals $[\log 2^\ell r_0, \log 2^{\ell+1} r_0]$ on which $\langle N_v\rangle_\ell$ is larger than $\sqrt{a}$. The inequality for f now tells us that the length of J can be at most $64\sqrt{a}\ |[\log r_1, \log r_0]|$. This means that the union of J with the interval I from the first part has bounded measure

$$|I \cup J| < 65\sqrt{a}\ |[\log r_0, \log r_1]| .$$

We conclude that whenever $\sqrt{a}$ is sufficiently small (for example $\sqrt{a} < 0.01$), then there exists a point $\log(t_1)$ that is neither in I nor in J. In fact, if $r_0 < r_1^{1-100\sqrt{a}}$, we can moreover find a t_1 that is larger than $r_1^{1-100\sqrt{a}}$, since the measure of $[r_1^{1-100\sqrt{a}}, r_1]$ is still larger than I and J combined.

Choose and fix a t_1 with these properties. On the one hand this means that $N(t_1) \leq \sqrt{a}$ and Lemma 3.14 implies that on the interval $[\tilde{R}_1, t_1]$ we have $N < a^{1/4}$ and $\kappa \geq a^{\frac{a^{1/4}+\epsilon}{4n}}\kappa(r)$. Here, $\tilde{R}_1$ is the larger of $a^{1/4n}t_1$ and R. On the other hand it means that t_1 is from an interval $[2^\ell r_0, 2^{\ell+1}r_0]$ on which the average value $\langle N\rangle_\ell$ is less than $\sqrt{a}$. Consequently, the partition of that interval contains a radius t_2 at which $N_v(t_2) < \sqrt{a}$. We again deduce that $N_v < a^{1/4}$ and $\kappa_v \geq a^{\frac{a^{1/4}+\epsilon}{4n}}\kappa_v(r)$ on $[\tilde{R}_2, t_2]$, where $\tilde{R}_2$ is the larger of $a^{1/4n}t_2$ and R.

Finally, $[\tilde{R}_1, t_1]$ and $[\tilde{R}_2, t_2]$ intersect whenever a is small enough. To see this explicitly, assume w.l.o.g that $t_1 > t_2$ and note that $t_1/t_2 \leq 2$ since both are contained in an interval of the form $[2^\ell r_0, 2^{\ell+1} r_0]$. It follows that $\tilde{R}_1/t_2 < 2a^{1/4n}$, which is less than one if for example $a^{1/4n} < 0.1$. The upper bound of 0.1 is convenient, because we then also have $\tilde{R}_1 < a^{1/8n} t_2$, such that choosing $t = t_2$ and $\tilde{R} = \max(a^{1/8n} t, R)$ concludes the proof. □

3.6 A Priori Bounds

The final ingredient for the proof of Theorem A are pointwise a priori bounds for $\|\phi\|$ in terms of κ, as well as estimates for the volume of the subsets of B_r on which $\|\phi\|$ is small compared to that upper bound. These are the promised analogues of the mean value inequality.

The proof relies mainly on the standard approach to the mean-value inequality and the fact that according to Proposition 3.9, κ can't decrease too rapidly at infinity. However, due to the asymptotic nature of the latter statement and the related lack of control in the interior, these bounds only hold for points that lie in some large spherical geodesic shell surrounding p.

Below we use the notation $S(r_1, r_2)$ for the closed spherical geodesic shell based at p with inner radius r_1 and outer radius r_2, i.e. all points that satisfy $r_1 \leq d(p, x) \leq r_2$. Moreover, we write $f_{\Omega_i} := \frac{\operatorname{vol}\Omega_i}{\operatorname{vol}\Omega}$ for the fraction that a subset Ω_i occupies within its corresponding surrounding set Ω.

The following result holds verbatim if we replace ϕ, κ and N by the versions ϕ_v, κ_v and N_v associated to a unit vector $v \in T_p W$.

Lemma 3.18. *Let W^n be an ALX$_k$ space of dimension n, $p \in W^n$, $\operatorname{Ric} \geq 0$, and assume (A, ϕ) is a solution of (3.1). For any $\epsilon \in (0, 1/2)$ there exists a radius $R \geq 0$ and constants $c_i > 1$ that only depend on n and k, such that for sufficiently large $r \geq R$ the following holds.*

(i) $\|\phi(x)\| \leq c_0 \frac{\kappa(r)}{\sqrt{\operatorname{vol} X}}$ *for any* $x \in S\left(R + \frac{r-R}{8}, r - \frac{r-R}{8}\right)$,

(ii) *Let* $t \in [R + \frac{6}{8}(r-R), R + \frac{7}{8}(r-R)]$ *and assume that* $N \leq 1$ *on* $[t, r]$. *Then the subset* $\Omega_1 \subseteq \partial B_t$ *on which* $\|\phi(x)\| \leq \frac{1}{2} \frac{\kappa(r)}{\sqrt{\operatorname{vol} X}}$ *has relative volume* $f_{\Omega_1} \leq c_1 \epsilon + c_2 \sqrt{N}$.

Proof. Throughout the proof we fix some $\epsilon \in (0, 1/2)$ and choose R to be some large enough radius, such that on the one hand κ is ϵ-almost non-decreasing as provided by Proposition 3.9 and on the other hand $\operatorname{vol} B_r \leq (1+\epsilon) \operatorname{vol} X \, r^{n-k}$ for any $r \geq R$ as in Proposition 3.2. Assume for now simply that $r > R$ is some fixed outer radius. We will collect conditions on r as we go and will find that they can always be met by choosing some larger r to start with.

ad (i)

In this part of the proof we use a specific bump function β on W^4 with compact support inside the geodesic shell $S(R, r)$, constructed as follows. Denote by $\beta_{\mathbb{R}}$ a non-increasing function on

$\mathbb{R}$ that is equal to 1 on $\left(-\infty, \frac{14}{16}\right)$ and equal to 0 on $\left(\frac{15}{16}, \infty\right)$. Use this to define a corresponding function on W^4 by the rule

$$\beta(x) = \beta_{\mathbb{R}}\left(\left|\frac{d(p,x) - \frac{r+R}{2}}{\frac{r-R}{2}}\right|\right) .$$

The function β vanishes on the inner ball $B_{R+\frac{r-R}{32}}(p)$, is equal to 1 on the geodesic shell with inner radius $R + \frac{r-R}{16}$ and outer radius $r - \frac{r-R}{16}$, and is zero again outside of $B_{r-\frac{r-R}{32}}(p)$. Note in particular that β and its derivatives have compact support in the interior of $S(R, r)$. Furthermore, the gradient and Laplacian of β are bounded as follows

$$|d\beta| \leq \frac{32}{(r-R)} \qquad\qquad |\Delta\beta| \leq \frac{32}{(r-R)^2} .$$

Denote by G_x the positive Dirichlet Green's function of the Laplacian on $B_r(p)$ based at $x \in B_r(p)$. Recall from Lemma 3.7 that, due to the volume growth of W^n, for large enough distances the Green's function and its derivative are bounded from above as follows

$$G_x(y) \leq \frac{(1+\epsilon)c}{\operatorname{vol} X\, d(x,y)^{n-k-2}} \qquad\qquad |dG_x(y)| \leq \frac{(1+\epsilon)c}{\operatorname{vol} X\, d(x,y)^{n-k-1}} ,$$

where c depends only on the dimension n.

With that in mind we now note that contracting equation (3.1) with $\langle \cdot, \phi\rangle$ leads to

$$\frac{1}{2}\Delta_B \|\phi\|^2 + \|\nabla^A\phi\|^2 + \|[\phi\wedge\phi]\|^2 + \langle \operatorname{Ric}\phi, \phi\rangle = 0 , \tag{3.13}$$

where $\Delta_B = \nabla^\dagger\nabla$ is the Bochner Laplacian associated to the Levi-Civita connection. This equation implies $\Delta_B \|\phi\|^2 \leq 0$, so the function $\|\phi\|^2$ is subharmonic[2] and accordingly satisfies a version of the mean-value inequality. To see this directly, multiply (3.13) by βG_x and integrate over $B_r(p)$

$$\int_{B_r(p)} \beta G_x \left(\Delta_B \|\phi\|^2 + \|\nabla^A\phi\|^2 + \|[\phi\wedge\phi]\|^2 + \langle \operatorname{Ric}\phi, \phi\rangle\right) = 0 .$$

Upon integration by parts in the first term, using that $\beta G_x = 0$ on $\partial B_r(p)$, and assuming that x is an element of the geodesic shell on which $\beta = 1$, we find

$$\begin{aligned} \|\phi(x)\|^2 &+ \int_{B_r(p)} \left(g^{-1}(d\beta, dG_x) + (\Delta_B\beta)G_x\right)\|\phi\|^2 \\ &+ \int_{B_r(p)} \beta G_x \left(\|\nabla^A\phi\|^2 + \|[\phi\wedge\phi]\|^2 + \langle \operatorname{Ric}\phi, \phi\rangle\right) = 0 . \end{aligned}$$

Since $\operatorname{Ric} \geq 0$, the last term on the left hand side is non-negative, so the equation provides the following estimate.

$$\|\phi(x)\|^2 \leq \left|\int_{B_r(p)} \left(g^{-1}(dG_x, d\beta) + G_x\Delta_B\beta\right)\|\phi\|^2\right|$$

[2]Note that the Bochner and connection Laplacian differ by a sign: $\Delta_B = -\operatorname{tr}\nabla^2$.

Assume $x \in S(R + \frac{r-R}{8}, r - \frac{r-R}{8})$, such that the distance from x to the support of $d\beta$ and $\Delta\beta$ is greater or equal to $\frac{r-R}{16}$. Furthermore, make r large enough that the previously mentioned bounds on $G_x(y)$ hold for all points with distance $d(x,y) \geq \frac{r-R}{16}$. Then, using the Cauchy-Schwarz inequality on the first term and the estimates for G, $|dG|$, $|d\beta|$ and $|\Delta_B \beta|$, we arrive at

$$\|\phi(x)\|^2 \leq \frac{32(1+\epsilon)c}{\operatorname{vol} X\,(r-R)^{n-k}} \int_{S(R,r)} \|\phi\|^2 \ .$$

As final step use that in the given domain of integration κ is ϵ-almost non-decreasing. Hence, Corollary 3.11 with $N \geq 0$ provides the estimate $\kappa(t) \leq \left(\frac{r}{t}\right)^{\epsilon} \kappa(r)$ for all $t \leq r$. The integral in the last equation is thus bounded from above by

$$\int_{S(R,r)} \|\phi\|^2 = \int_R^r dt\, t^{n-k-1} \kappa^2(t) \leq \frac{r^{n-k}}{n-k-2\epsilon}\, \kappa^2(r) \ .$$

Plugging this in, assuming $r > 2R$, and using $\epsilon < 1/2$ leads to

$$\|\phi(x)\|^2 \leq \frac{2^{n-k+6}c}{(n-k-1)\operatorname{vol} X} \kappa^2(r) \leq c_0^2 \frac{\kappa^2(r)}{\operatorname{vol} X} \ .$$

We note that c_0 only depends on the dimension of W^n and the dimension of the fibers at infinity.

ad (ii)

Fix some radius $t \in [R + \frac{6}{8}(r-R), R + \frac{7}{8}(r-R)]$ and consider the associated geodesic sphere ∂B_t. We are interested in the volume of the subset of ∂B_t where $\|\phi\|$ is small compared to the bound in *(i)*.

Hence, write $\Omega_1 \subseteq \partial B_t$ for the points where $\|\phi(x)\| \leq \frac{\kappa(r)}{2\sqrt{\operatorname{vol} X}}$. Also introduce $\Omega_2 \subseteq \partial B_t$ to denote the set of points at which $\|\phi(x)\| \leq (1 + \sqrt{N(r)})\frac{\kappa(r)}{\sqrt{\operatorname{vol} X}}$. Since $N \geq 0$, Ω_2 contains Ω_1.

Split up the contributions to $\kappa^2(t)$ that arise from integration over Ω_1, $\Omega_2 \setminus \Omega_1$, and their complement $\Omega_3 = \partial B_t \setminus \Omega_2$.

$$\kappa^2(t) = \frac{1}{t^{n-k-1}} \left(\int_{\Omega_1} \|\phi\|^2 + \int_{\Omega_2 \setminus \Omega_1} \|\phi\|^2 + \int_{\Omega_3} \|\phi\|^2 \right)$$

On Ω_1 the integrand is bounded by $\frac{\kappa^2(t)}{4 \operatorname{vol} X}$, on $\Omega_2 \setminus \Omega_1$ we use $(1+\sqrt{N})^2 \frac{\kappa^2(r)}{\operatorname{vol} X}$, and the integral over Ω_3 can't be larger than $\kappa^2(t)$ in any case.

With regard to the last of these bounds we now make use of the fact that κ is almost non-decreasing, which allows us to compare $\kappa^2(t)$ to $\kappa^2(r)$. To that end consider the derivative of κ^2.

$$\left.\frac{d\kappa^2}{dr}\right|_{\tilde{r}} = 2\frac{N+D}{\tilde{r}} \kappa^2(\tilde{r}) \geq -2\epsilon \frac{\tilde{r}^{1+2\epsilon}}{r^{2+2\epsilon}} \kappa^2(r)$$

For the estimate on the right hand side we have used on the one hand that $N \geq 0$ and $D \geq -\epsilon$, and on the other hand that $N \leq 1$ such that Corollary 3.11 implies that $\kappa^2(\tilde{r}) \geq (\tilde{r}/r)^{2+2\epsilon}\kappa^2(r)$. Integration from t to r leads to

$$\kappa^2(t) \leq \frac{\epsilon}{1+\epsilon}\left(2 - \left(\frac{t}{r}\right)^{2+2\epsilon}\right)\kappa^2(r)\,.$$

Since $t \leq r$, we may as well use the somewhat simpler statement $\kappa^2(t) \leq 2\epsilon\kappa^2(r)$.

Plugging in the corresponding bounds on each of the Ω_i, writing $\operatorname{vol}\Omega_i = f_{\Omega_i}\operatorname{vol}\partial B_t$, $f_{\Omega_2} \leq 1$, and using $\operatorname{vol}\partial B_t \leq (1+\epsilon)\operatorname{vol} X t^{n-k-1}$ thus leads to

$$\begin{aligned}\kappa^2(t) &\leq \left(\frac{1}{4}f_{\Omega_1} + (1-f_{\Omega_1})\,(1+\sqrt{N(r)}\,)^2 + 2\epsilon\right)(1+\epsilon)\kappa^2(r)\\ &\leq \left(1+2\epsilon - \frac{3}{4}f_{\Omega_1}\right)(1+\sqrt{N(r)}\,)^2(1+\epsilon)\kappa^2(r)\,,\end{aligned}$$

which can be rearranged to

$$f_{\Omega_1} \leq \frac{4}{3}\left(1+2\epsilon - \frac{\kappa^2(t)}{\left(1+\sqrt{N(r)}\right)^2(1+\epsilon)\kappa^2(r)}\right). \tag{3.14}$$

To make this expression more useful we'll now also need a lower bound for $\kappa^2(t)$ in terms of $\kappa^2(r)$. This can again be achieved by considering the derivative of κ^2. In this case we observe that the following function is non-decreasing in $\tilde{r}$.

$$\tilde{r} \mapsto \tilde{r}^{n-k-2}\kappa^2(\tilde{r})N(\tilde{r}) = \int_{B_{\tilde{r}}(p)} \|\nabla^A\phi\|^2 + \|[\phi\wedge\phi]\|^2 + \langle \operatorname{Ric}(\phi),\phi\rangle$$

Moreover, since κ is ϵ-almost non-decreasing for $\tilde{r} > R$, we have $|D| \leq N + \epsilon$. This yields the following upper bound for the derivative of κ^2:

$$\left.\frac{d\kappa^2}{dr}\right|_{\tilde{r}} = 2\frac{N+D}{\tilde{r}}\kappa^2(\tilde{r}) \leq 4\,\frac{r^{n-k-2}}{\tilde{r}^{n-k-1}}N(r)\kappa^2(r) + 2\epsilon\,\frac{r^{2\epsilon}}{\tilde{r}^{1+2\epsilon}}\kappa^2(r)\,.$$

Integration[3] from t to r now leads to

$$\kappa^2(r) - \kappa^2(t) \leq \frac{4}{n-k-2}\left(\frac{r}{t}\right)^{n-k-2}N(r)\kappa^2(r) + \left(\left(\frac{r}{t}\right)^{2\epsilon} - 1\right)\kappa^2(r)\,.$$

Recall that $r/t < 4/3$ and observe that thus $(r/t)^{2\epsilon} - 1 < \epsilon$ for any choice of $\epsilon \in (0, 1/2)$. It follows that

$$\kappa^2(t) \geq \left(1 - \epsilon - cN(r)\right)\kappa^2(r)\,,$$

where the constant c only depends on n and k. Plugging this lower bound for $\kappa^2(t)$ into (3.14) and using that $\epsilon, N < 1$, we conclude that

$$f_{\Omega_1} \leq c_2\epsilon + c_3\sqrt{N(r)}\,,$$

where c_2 and c_3 only depend on n and k. □

[3] For notational simplicity we assume here that $k \neq n-2$. However, the result holds similarly for the case $k = n-2$, where the only difference is that upon integration the formulas contain logarithms.

3.7 Proof of Taubes' Dichotomy on ALX spaces

We are now in a position to prove Theorem A. As advertised before, the arguments are in complete analogy to Taubes' original proof [Tau17a].

Theorem 3.19. *Let W^n be an ALX_k gravitational instanton of dimension $n \geq 2$ with asymptotic fibers of dimension $k \leq n-1$ and fix a point $p \in W^n$. Assume (A, ϕ) satisfies the second-order differential equation (3.1). Then*

(i) either there is an $a > 0$ such that $\liminf_{r\to\infty} \frac{\kappa(r)}{r^a} > 0$,

(ii) or $[\phi \wedge \phi] = 0$.

Proof. It is sufficient to consider the case where κ is not asymptotically zero, since otherwise κ – and thus ϕ – vanish identically due to the unique continuation property of Lemma 3.13. By the same reasoning one also finds that $\phi_v = 0$ whenever κ_v has compact support. As a consequence it is sufficient to consider the components of ϕ that are in the complement of the zero eigenspace of T at infinity, such that the smallest eigenvalue function satisfies $\lambda > 0$ at large enough radii (see the related discussion in Section 3.5). Also, note that when only a single component ϕ_v is non-zero at infinity, then $[\phi \wedge \phi] = 0$ everywhere. Hence, we assume from now on that T acts on a vector space of dimension at least 2.

To prove the dichotomy stated in the theorem, assume that κ is not asymptotically bounded below by any positive power of r. This means that for any small $a > 0$ (say, for example, small enough that $a^{1/8n} < 0.1$) we can do the following: Set $\epsilon = a/2$ and let $R > 1$ be a radius beyond which $|D| \leq \epsilon$, $\operatorname{vol} B_r(p) \leq (1+\epsilon) \operatorname{vol} X r^{n-k}$, and such that the smallest eigenvalue function $\lambda > 0$ on all of (R, ∞). Then we can find an arbitrarily large $r_1 \in [R, \infty)$ such that $\kappa(r_1) \leq \left(\frac{r_1}{R}\right)^{a-\epsilon} \kappa(R)$. In particular we may choose some r_1 that is larger than each of the four numbers $4R$, $\left(\kappa^2(R)/\lambda(R)\right)^{1/2a} R$, $\kappa(R)^{-1/a}$, and $R^{1/(1-100\sqrt{a})}$. We are then in the situation in which we can rely on Lemma 3.17. The arguments in the upcoming six parts show that this leads to a contradiction if $N \neq 0$ and a is too small.

Part 1 Recall that Lemma 3.17 provides the existence of a distinguished radius

$$t \in [r_1^{1-100\sqrt{a}}, r_1] \subseteq [R, r_1] \, .$$

In this part we collect our previous results and slightly expand on our knowledge about the eigenvalues of T at and below t.

Write u and v for unit eigenvectors associated to the *largest* and *smallest* eigenvalue of $T(t)$, respectively. Recall that these eigenvalues coincide at t with the values of $\kappa_u^2(t)$ and $\kappa_v^2(t)$, so they satisfy $\kappa_v^2(t) \leq \kappa_u^2(t)$. If $\kappa_v^2(t) \neq \kappa_u^2(t)$ the eigenvectors v and u are guaranteed to be orthogonal. Otherwise $T(t)$ is a multiple of the identity matrix and we choose an arbitrary pair of orthonormal vectors.

Denote by $\tilde{R}$ the larger of $a^{\frac{1}{8n}}t$ and R. Lemma 3.17 establishes that on the interval $I := [\tilde{R}, t]$ the frequency functions N and N_v are bounded from above by $a^{1/4}$ and provides associated lower bounds for κ and κ_v. As we show now, analogous estimates hold for N_u and κ_u.

First observe that the largest eigenvalue satisfies $\kappa_u^2(t) \geq \frac{1}{n}\kappa^2(t)$ since κ^2 is the trace of T. From this and the definition of N and its N_u version (cf. Proposition 3.9 and (3.10), respectively) it follows that $N_u(t) \leq nN(t)$. As a consequence $N_u \leq na^{1/4}$ and, as long as we make sure that $a^{1/4} < 1/n$, a small variation of the second part of Proposition 3.12 finds that $N_u \leq a^{1/8}$ on all of I. Since N_u is bounded from above, we can now deduce bounds on κ_u as usual via Corollary 3.11.

$$\kappa_u \geq a^{\frac{a^{1/8}+\epsilon}{8n}} \kappa_u(t) \, .$$

In the current situation with $\epsilon = a/2$ this bound and its analogue for κ_v can be simplified considerably. For that observe that $a^\epsilon = a^{a/2}$ is certainly larger than $1/2$, and similarly $a^{\frac{a^{1/4}}{4n}} > (1/2)^{\frac{4}{4n}} > 1/2$ and $a^{\frac{a^{1/8}}{8n}} > 1/2$. Applying these observations to the bounds for κ_u and κ_v results in the main conclusions of this part.

In conclusion, the following estimates hold on all of $I = [\tilde{R}, t]$:

- $N_v \leq a^{1/4}$ and $\kappa_v \geq \frac{1}{4}\kappa_v(t)$
- $N_u \leq a^{1/8}$ and $\kappa_u \geq \frac{1}{4}\kappa_u(t)$

Part 2 We now focus on the upper half of $I = [\tilde{R}, t]$ and investigate the values of the correlation tensor T evaluated on the two unit eigenvectors u and v from Part 1. Since u and v are orthogonal, we know that $T(t)(u, v) = 0$. As a result, for any $s \in \left[\frac{\tilde{R}+t}{2}, t\right]$ the norm of $T(u, v)$ is bounded by

$$|T(s)(u,v)| \leq \int_s^t \left| \frac{dT(u,v)}{dr} \right| dr \, .$$

The derivative of T follows from differentiation of (3.8) and works out in complete analogy to the derivative of κ in Proposition 3.9. In particular, upon use of equation (3.9), we find

$$\begin{aligned} \frac{dT(u,v)}{dr} = \frac{1}{r^{n-k-1}} \int_{B_r} \Big(\langle \nabla^A \phi_u, \nabla^A \phi_v \rangle + \langle [\phi \wedge \phi_u], [\phi \wedge \phi_v] \rangle + \langle \mathrm{Ric}(\phi)(u), \phi_v \rangle + \langle \phi_u, \mathrm{Ric}(\phi)(v) \rangle \Big) \\ + \frac{1}{r^{n-k-1}} \int_{\partial B_r} \left(\Delta r - \frac{n-k-1}{r} \right) \langle \phi_u, \phi_v \rangle \, . \end{aligned}$$

At this point we assume that W^n is Ricci flat since then the terms that involve $\mathrm{Ric}\,\phi$ vanish and we can use the Cauchy-Schwarz inequality in the first integral. Also, on (R, ∞) the absolute value of the bracket in the second integral is less than ϵ and upon use of this bound we can also use the Cauchy-Schwarz inequality there. This provides the upper bound

$$\left| \frac{dT(u,v)}{dr} \right| \leq \frac{1}{r} \kappa_u(r) \kappa_v(r) \sqrt{N_u(r) N_v(r)} + \epsilon \kappa_u(r) \kappa_v(r) \, .$$

From Part 1 we know that both N_u and N_v are less or equal to $a^{1/8}$, while we have set things up such that $\epsilon < a < a^{1/8}$. Since on the present interval both κ_u and κ_v are ϵ-almost non-decreasing, we can use the estimate $\kappa_u(r) \leq \left(\frac{t}{r}\right)^{\epsilon} \kappa_u(t)$ and the analogous expression for κ_v. Integration from s to t leads to

$$|T(s)(u,v)| \leq a^{1/8}\kappa_u(t)\kappa_v(t)\left(\frac{t}{s}\right)^{1+2\epsilon} \leq 4\kappa_u(t)\kappa_v(t)a^{1/8} \,, \tag{3.15}$$

where in the last inequality we have used that $t/s \leq 2$ and $\epsilon < 1/2$. The main result of this part then is that $|T(u,v)|$ is smaller by a factor of $\mathcal{O}(a^{1/8})$ than $\kappa_u(t)\kappa_v(t)$ on all of $[\frac{\tilde{R}+t}{2}, t]$.

Part 3 The goal of this part is to show that there exist points in $J = \left[\tilde{R} + \frac{6}{8}(t-\tilde{R}), \tilde{R} + \frac{7}{8}(t-\tilde{R})\right]$ for which the integrals that appear in N_u and N_v are both small. This interval is of significance, because it corresponds to radii that are contained in the geodesic shell that appears in item *(iii)* of Lemma 3.18.

Denote by $v_u := \left(\|\nabla^A\phi_u\|^2 + \|[\phi\wedge\phi_u]\|^2\right)$ the integrand in N_u and analogously for v_v. Consider the following sets of radii in J.

$$J_u := \left\{ s\in J \,\middle|\, \int_{\partial B_s} v_u \leq \frac{1}{2\,|J|}\; t^{n-k-2}\; \kappa_u^2(t)\; a^{1/16} \right\} \tag{3.16}$$

$$J_v := \left\{ s\in J \,\middle|\, \int_{\partial B_s} v_v \leq \frac{1}{2\,|J|}\; t^{n-k-2}\; \kappa_v^2(t)\; a^{1/16} \right\} \tag{3.17}$$

Note for later that $|J| = |I|/8$, where $I = [\tilde{R}, t]$, so the fact that $a^{1/8n}t \leq \tilde{R}$ and using $a^{1/8n} < 0.5$ yields $|J| \geq \frac{t}{16}$.

The measure of J_u is bounded from below, since

$$t^{n-k-2}\kappa_u^2(t)N_u(t) \geq \int_J ds \int_{\partial B_s} v_u \geq (|J| - |J_u|)\;\frac{1}{2|J|}\; t^{n-k-2}\;\kappa_u^2(t)\; a^{1/16} \,,$$

and similarly for J_v. Since both $N_u(t), N_v(t) \leq a^{1/8}$, we find

$$|J_u|\,, |J_v| > \left(1 - 2a^{1/16}\right)|J| \;.$$

Since $2a^{1/16} < 0.5$ (recall that $a^{1/8n} < 0.1$ and $n \geq 2$ in any case), it follows that J_u and J_v must intersect. Hence, choose and fix from now on a point $r \in J$ at which both (3.16) and (3.17) are satisfied.

Part 4 As a next step we establish an L^2-bound for the function $\operatorname{Tr}\phi_u\phi_v$ on ∂B_r, where r is the fixed radius we found in Part 3. For this we view ∂B_r with the induced metric as a compact Riemannian manifold and rely on a Poincaré inequality. Note that $\operatorname{Tr}\phi_u\phi_v$ is the integrand that defines the correlation tensor $T(r)(u,v)$, so up to normalization the latter yields the average value of $\operatorname{Tr}\phi_u\phi_v$ on ∂B_r.

First note that the derivative of $\operatorname{Tr}\phi_u\phi_v$ is bounded by

$$\|d\operatorname{Tr}\phi_u\phi_v\| \leq \|\nabla^A\phi_u\|\,\|\phi_v\| + \|\phi_u\|\,\|\nabla^A\phi_v\|\,.$$

Because r lies inside the geodesic shell of Lemma 3.18 we know that $\|\phi_u\| \leq c_0\kappa_u(t)$ and $\|\phi_v\| \leq c_0\kappa_v(t)$. In light of the fact that both (3.16) and (3.17) are satisfied at r and using the Cauchy-Schwarz inequality, we find that

$$\int_{\partial B_r}\|d\operatorname{Tr}\phi_u\phi_v\|^2 \leq 16c_0^2 t^{n-k-3}\kappa_u^2(t)\kappa_v^2(t)a^{1/16}\,.$$

Consider now the function $G = \operatorname{Tr}\phi_u\phi_v - \frac{T(r)(u,v)}{\alpha\operatorname{vol}X}$ on the geodesic sphere ∂B_r, where the normalization factor α is chosen such that G captures the deviations of $\operatorname{Tr}\phi_u\phi_v$ from its average value on ∂B_r. It follows that G satisfies $\int_{\partial B_r} G = 0$ and $dG = d\operatorname{Tr}\phi_u\phi_v$. Since ∂B_r is a compact Riemannian manifold without boundary and $\operatorname{Ric}\geq 0$, the following Poincaré inequality holds (see e.g. [Li82, Thm. 3])

$$\int_{\partial B_r}|G|^2 \leq \frac{2d^2}{\pi^2}\int_{\partial B_r}\|dG\|^2\,.$$

Here d denotes the diameter of the geodesic sphere ∂B_r, which is bounded by some constant multiple of its radius $d \leq \tilde{c}r$. The constant $\tilde{c}$ can be chosen uniformly for all geodesic spheres based at $p \in W^n$ and only depends on n and k, because by Definition 3.3 curvature approaches zero at infinity such that the geometry of geodesic spheres effectively becomes that of the standard sphere S_r^{n-k-1} embedded in $\mathbb{R}^{n-k}$. Together with (3.15) and the usual volume comparison for ∂B_r, this leads to the desired L^2-bound for $\operatorname{Tr}\phi_u\phi_v$. Specifically, there is a constant $C > 1$ such that

$$\int_{\partial B_r}|\operatorname{Tr}\phi_u\phi_v|^2 \leq \tilde{c}^2r^2\int_{\partial B_r}\|dG\|^2 + \int_{\partial B_r}\left(\frac{T(r)(u,v)}{\alpha\operatorname{vol}X}\right)^2 \leq Ct^{n-k-1}\kappa_u^2(t)\kappa_v^2(t)a^{1/16}\,. \tag{3.18}$$

From now on we allow the value of C to increase from one equation to the next.

Part 5 Our next goal is to derive a closely related L^2-bound for the function $|\phi_u|\cdot|\phi_v|$ on ∂B_r, where $|\phi_u|$ denotes pointwise application of the norm induced by the Killing form. It is a property of $\mathfrak{su}(2)$ that

$$|[\phi_u,\phi_v]|^2 = 4\,|\phi_u|^2\,|\phi_v|^2 - 4\operatorname{Tr}(\phi_u\phi_v)^2\,.$$

Moreover, since u and v are orthonormal $\|[\phi\wedge\phi_v]\|^2 \geq |[\phi_u,\phi_v]|^2\,\mu_{W^n}$, so we can bound the following integral with the help of (3.17).

$$\int_{\partial B_r}|\phi_u|^2\,|\phi_v|^2 - \int_{\partial B_r}|\operatorname{Tr}\phi_u\phi_v|^2 = \frac{1}{4}\int_{\partial B_r}|[\phi_u,\phi_v]|^2 \leq \frac{1}{4}\int_{\partial B_r}\nu_v \leq 2\,t^{n-k-3}\,\kappa_v^2(t)\,a^{1/16}$$

Together with (3.18), the main result of Part 4, this leads to

$$\int_{\partial B_r}|\phi_u|^2\,|\phi_v|^2 \leq \left(C + \frac{2}{t^2\kappa_u^2(t)}\right)t^{n-k-1}\,\kappa_u^2(t)\,\kappa_v^2(t)\,a^{1/16}\,.$$

Observe that things have been set up in such a way that $t^2\kappa_u^2(t) > 1$: First, κ_u is almost non-decreasing, so $\kappa_u^2(t) \geq (t/R)^{2\epsilon}\kappa_u^2(R)$. Second, we know that $\kappa_u^2(R) \geq \lambda(R)$ since the latter is the smallest eigenvalue of $T(R)$. Third, we have previously chosen r_1 large enough that $\lambda(R) > (R/r_1)^{2a}\kappa^2(R)$. Fourth, t is larger than $r_1^{1-\sqrt{100a}}$ while $R > 1$. Plugging everything together yields

$$t^2\kappa_u^2(t) \geq r_1^{2(1-100\sqrt{a}-4a)}\kappa^2(r_1) > r_1^{2a}\kappa^2(r_1)\,,$$

where the last step follows via $a^{1/4n} < 0.1$ and $n \geq 2$. Finally, recall that we have also explicitly chosen r_1 to be large enough that the rightmost expression is larger than 1.

The upshot of this part is that there exists a (larger) constant $C > 1$ that only depends on n and k, such that

$$\int_{\partial B_r} |\phi_u|^2\,|\phi_v|^2 \leq Ct^{n-k-1}\,\kappa_u^2(t)\,\kappa_v^2(t)\,a^{1/16}\,. \tag{3.19}$$

Part 6 In this final part, we combine the results of the previous five parts with item *(ii)* of Lemma 3.18. Thus, write Ω_1 for the subset of ∂B_r where $\|\phi_u\| \leq \frac{\kappa_u(t)}{2\operatorname{vol}X}$.

The inequality in (3.19) remains true if we restrict the domain of integration to $\partial B_r \setminus \Omega_1$, such that

$$\frac{\kappa_u^2(t)}{4\operatorname{vol}X}\int_{\partial B_r\setminus\Omega_1}|\phi_v|^2 \leq \int_{\partial B_r\setminus\Omega_1}|\phi_u|^2\,|\phi_v|^2 \leq Ct^{n-k-1}\,\kappa_u^2(t)\,\kappa_v^2(t)\,a^{1/16}\,.$$

More concisely, this leads to the inequality

$$\int_{\partial B_r\setminus\Omega_1}|\phi_v|^2 \leq C\operatorname{vol}Xt^{n-k-1}\kappa_v^2(t)a^{1/16}\,.$$

Meanwhile, item (i) of Lemma 3.18 provides the upper bound $|\phi_v| \leq c_0^2\kappa_v^2(t)$. Writing $\operatorname{vol}\Omega_1 = f_{\Omega_1}\operatorname{vol}B_r$ this leads to

$$\int_{\Omega_1}|\phi_v|^2 \leq f_{\Omega_1}(1+\epsilon)\operatorname{vol}X\,r^{n-k-1}c_0^2\kappa_v^2(t)\,.$$

Item (ii) of Lemma 3.18 states that $f_{\Omega_1} \leq c_1\epsilon + c_2\sqrt{N_u(t)}$. Recall from Part 1 that $N_u(t) \leq a^{1/8}$ while $\epsilon = a/2$, such that f_{Ω_1} is bounded by some multiple of $a^{1/16}$.

It follows that the sum of the two integrals satisfies

$$\int_{\partial B_r}|\phi_v|^2 \leq Ct^{n-k-1}\kappa_v^2(t)a^{1/16}\,.$$

This is equivalent to the statement that $\kappa_v^2(r) \leq Ca^{1/16}\kappa_v^2(t)$. Finally, combining this with the bound $\kappa_v^2(r) \geq \frac{1}{4}\kappa_v^2(t)$ from Part 1 culminates in the inequality

$$\kappa_v^2(t) \leq 4Ca^{1/16}\kappa_v^2(t)\,,$$

which is absurd, as we are free to choose a arbitrarily small and in particular such that $a^{1/16} < \frac{1}{4C}$.

□

3.8 Proof of Taubes' Dichotomy for Kapustin-Witten Solutions

In this section we prove Theorem B, which enhances Theorem 3.19 for solutions of the Kapustin-Witten equations on four-manifolds. We again closely follow Taubes' arguments, who proved an analogous statement in case the four-manifold is Euclidean space $\mathbb{R}^4$.

Theorem 3.20. *Let W^4 be an ALX gravitational instanton of dimension 4, with asymptotic fibers of dimension $k \leq 3$, and such that sectional curvature is bounded from below. Assume (A, ϕ) are solutions of the θ-Kapustin-Witten equations and if $\theta \not\equiv 0, \pi$ also assume that $\int_{W^4} \|F_A\|^2 < \infty$, then*

(i) either there is an $a > 0$ such that $\liminf_{r\to\infty} \frac{\kappa(r)}{r^a} > 0$,

(ii) or $[\phi \wedge \phi] = 0$, $\nabla^A \phi = 0$, and A is self-dual if $\theta = 0$, flat if $\theta \in (0, \pi)$, and anti-self-dual if $\theta = \pi$.

Proof. Since solutions of the Kapustin-Witten equations also satisfy equation (3.1), the dichotomy of Theorem 3.19 holds. It remains to show that in the case where $[\phi \wedge \phi]$ is identically zero, also $\nabla^A \phi$ vanishes and A is either (anti-)self-dual or flat as stated. Hence, assume from now on that $[\phi \wedge \phi] = 0$. At points where ϕ is non-zero the Higgs field can then be written as $\phi = \omega \otimes \sigma$, where $\omega \in \Omega^1(M)$ and $\sigma \in \Gamma(M, \operatorname{ad} E)$, normalized such that $\|\sigma\| = 1$.

Consider first the case where $\theta = 0$ (this also covers the case $\theta = \pi$ by a reversal of orientation). The Kapustin-Witten equations then state on the one hand $F_A^+ = 0$, so A is anti-self-dual, and on the other hand $(d_A\phi)^- = 0$ and $d_A \star \phi = 0$. The latter two equations can only be satisfied if the constituents of ϕ satisfy $\nabla^A \sigma = 0$ and $(d\omega)^- = 0 = d \star \omega$. In particular σ is guaranteed to be covariantly constant at points where $\phi \neq 0$.

The zero locus of ϕ coincides with the zero locus of ω. Since the one-form ω satisfies the first order differential equations from above, it is an example of what Taubes calls a $\mathbb{Z}/2$ harmonic spinor in [Tau14]. In that article he investigates the zero locus of such $\mathbb{Z}/2$ harmonic spinors in general and Theorem 1.3 of [Tau14] states that the zero locus has Hausdorff dimension 2. The relevance for us is that the complement of the zeroes of ω in any given ball in W^4 is path connected. This means that σ can be defined at points where $\omega = 0$ by parallel transport along paths where ω is non-zero. Since A is smooth and σ is ∇^A-parallel, parallel transport along two different paths results in the same value.

Since σ is defined everywhere, the same is true for $\omega = \frac{1}{2}\operatorname{Tr}(\phi\sigma)$. The elliptic differential equations $(d\omega)^- = 0 = d \star \omega$ imply that $\|\omega\|^2$ is subharmonic. Thus, by a classical result of Yau [Yau76, Theorem 3 & Appendix (ii)], either $\|\omega\|^2$ is constant or $\lim_{r\to\infty} r^{-1} \int_{B_r} \|\omega\|^2 > 0$. The latter is precluded by our assumptions and we conclude that $\nabla^A \phi = 0$.

In the case that $\theta \not\equiv 0 \pmod{\pi}$, neither of the terms $(d_A\phi)^\pm$ vanishes automatically. However, if we additionally assume that F_A is square-integrable, we can employ Uhlenbeck's compactness theorem for (anti-)self-dual connections to deduce that A must be flat. To see this, first note that we can find a coefficient $c(\theta)$ such that the connection $A + c(\theta)\phi$ satisfies the $\pi/2$ version of the Kapustin-Witten equations, so it is sufficient to consider the case $\theta = \pi/2$.

With that in mind and since $[\phi \wedge \phi] = 0$, the Kapustin-Witten equations state that $F_A = \star d_A \phi$. As a consequence the connection $\hat{A} := A + \phi$ is self-dual. Since $\int_{W^4} \|F_{\hat{A}}\|^2 < \infty$, this connection $\hat{A}$ is the pullback of a smooth, regular connection on the one-point compactification of W^4. In fact, by Uhlenbeck's removable singularities theorem [Uhl82b, Thm. 4.1 & Cor. 4.2], the field strength at large radius falls off as $\|F_{\hat{A}}\| \leq \frac{c}{r^4}$. Keeping this in mind, note that $\nabla^{\hat{A}}\phi = \nabla^A \phi$, such that $F_{\hat{A}} = 2(d_{\hat{A}}\phi)^+$ and

$$\int_{B_r} \|F_{\hat{A}}\|^2 = 2 \int_{B_r} \operatorname{Tr} F_{\hat{A}} \wedge d_{\hat{A}} \phi = 2 \int_{\partial B_r} \operatorname{Tr} F_{\hat{A}} \wedge \phi \,,$$

where we have used Stokes' theorem and the Bianchi identity in the last equality. With the bound on $\|F_{\hat{A}}\|$ the integral on the right is bounded by a multiple of $\frac{\kappa}{r^{k+1}}$, which approaches 0 for $r \to \infty$, so $\hat{A}$ is flat. From here we are back in the situation where $(d_A \phi)^+ = 0 = d_A \star \phi$ and the same argument as before leads to $\nabla^A \phi = 0$.

$\square$

4 The Decoupled Haydys-Witten Equations and a Weitzenböck Formula

Throughout this chapter, let (M^5, g) be an oriented five-manifold with poly-cylindrical ends, where ends may be located at either finite or infinite geodesic distance. Assume M^5 admits a non-vanishing unit vector field v and that the subbundle $\ker g(v, \cdot) \subset TM^5$ admits an almost Hermitian structure. This means that there is an almost complex structure $J : \ker g(v, \cdot) \to \ker g(v, \cdot)$ that is compatible with the metric, i.e. $g(J\cdot, J\cdot) = g(\cdot, \cdot)$.

In this chapter we investigate the Haydys-Witten equations on manifolds (M^5, g, v, J). The existence of J provides a specialization of the equations that we call *decoupled* Haydys-Witten equations. Crucially, the decoupled version of the equations exhibits a Hermitian Yang-Mills structure, which provides additional tools in solving the equations. This structure becomes most apparent in the $4\mathcal{D}$-formulation of the Haydys-Witten equations, an extension of Witten's $3\mathcal{D}$-formulation of the extended Bogomolny equations (EBE) [Wit11a], which is introduced below and used throughout the introduction. The main contribution consists in working out conditions under which the Haydys-Witten equations reduce to the decoupled version.

Curiously, in the context of Witten's gauge theoretic approach to homological knot invariants [Wit11a], manifolds are generally equipped with the additional structure (g, v, J). In that situation one considers the Haydys-Witten equations on five-manifolds of the form $M^5 = \mathbb{R}_s \times X^3 \times \mathbb{R}^+_y$, equipped with a product metric g, and sets $v = \partial_y$. The subbundle $\ker g(v, \cdot)$ is then simply the tangent space of $\mathbb{R}_s \times X^3$, which always admits an almost complex structure, as it is an open and orientable four-manifold. Conjecturally, when $X^3 = \mathbb{R}^3$ or S^3 and in the presence of a magnetically charged knot $K \subset X^3$, the homology groups obtained from θ-Kapustin-Witten solutions (stationary Haydys-Witten solutions) at $s \to \pm\infty$ modulo Haydys-Witten instantons reproduces a knot invariant known as Khovanov homology. Hence, the results presented here may offer a fresh perspective on the gauge theoretic approach to Khovanov homology.

In the following we briefly introduce the $4\mathcal{D}$-formulation of the Haydys-Witten equations and the decoupled version of the equations, see Section 4.1 for a more detailed, global discussion. For this, let G denote a compact Lie group, $G_{\mathbb{C}}$ its complexification, $E \to M^5$ a G-principal bundle, and $E_{\mathbb{C}}$ the associated $G_{\mathbb{C}}$-principal bundle. Furthermore, let $\mathcal{A}(E)$ denote gauge connections and write $\Omega^2_{v,+}(M^5)$ for Haydys' self-dual two-forms with respect to v [Hay15].

Given a pair $(A, B) \in \mathcal{A}(E) \times \Omega^2_{v,+}(M^5, \operatorname{ad} E)$ and an almost complex structure J, one can locally define four differential operators $\mathcal{D}_\mu$ that act on sections of $\operatorname{ad} E_{\mathbb{C}}$. To that end, consider normal coordinates $(x^i, y)_{i=0,1,2,3}$ near a point p, chosen in such a way that $v = \partial_y$ and that J takes the canonical form with respect to the coordinate vector fields ∂_i at p. In these coordinates, $B = \sum_{a,b,c=1}^{3} \phi_a(dx^0 \wedge dx^a + \frac{1}{2}\epsilon_{abc}dx^b \wedge dx^c)$. The four differential operators are defined by

$$\mathcal{D}_0 = \nabla_0^A + i\nabla_1^A \qquad \mathcal{D}_1 = \nabla_2^A + i\nabla_3^A$$
$$\mathcal{D}_2 = \nabla_y^A - i[\phi_1, \cdot] \qquad \mathcal{D}_3 = [\phi_2, \cdot] - i[\phi_3, \cdot]$$

There is a complex conjugation, induced from ad $E_{\mathbb{C}}$, that we denote by $\overline{\mathcal{D}_\mu}$. Furthermore, $G_{\mathbb{C}}$-valued gauge transformations act on the operators by conjugation, i.e. $\mathcal{D}_\mu \mapsto g\mathcal{D}_\mu g^{-1}$.

In this formulation, the Haydys-Witten equations $\mathbf{HW}_v(A, B) = 0$ are given by

$$[\overline{\mathcal{D}_0}, \overline{\mathcal{D}_i}] - \tfrac{1}{2}\varepsilon_{ijk}[\mathcal{D}_j, \mathcal{D}_k] = 0 \, , \qquad i = 1, \ldots, 3 \, , \tag{4.1}$$

$$\sum_{\mu=0}^{3}[\mathcal{D}_\mu, \overline{\mathcal{D}_\mu}] = 0 \, . \tag{4.2}$$

A typical approach in solving equations of this type is to utilize their symmetry properties. Since there is an action by complex gauge transformation, one natural idea is to use a Donaldson-Uhlenbeck-Yau type approach, where one first extracts some underlying holomorphic data from $G_{\mathbb{C}}$-invariant parts of the equations, and subsequently hopes to find a complex gauge transformation that ensures also the remaining equations are satisfied.

Unfortunately, while the Haydys-Witten equations are invariant under G-valued gauge transformations and the action lifts naturally to $G_{\mathbb{C}}$, neither (4.1) nor (4.2) are invariant under $G_{\mathbb{C}}$-valued gauge transformations. There is, however, a subset of solutions for which the three equations in (4.1) decompose into their $G_{\mathbb{C}}$-invariant parts. This is given by solutions that satisfy the following equations:

$$\begin{aligned} [\mathcal{D}_\mu, \mathcal{D}_\nu] &= 0 \, , \qquad \mu, \nu = 0, \ldots, 3 \, , \\ \sum_{\mu=0}^{3}[\mathcal{D}_\mu, \overline{\mathcal{D}_\mu}] &= 0 \, . \end{aligned} \tag{4.3}$$

We refer to (4.3) as *decoupled Haydys-Witten equations* and denote them by $\mathbf{dHW}_{v,J}(A, B) = 0$. A global version of these equations is provided in Section 4.1. The commutativity equations are $G_{\mathbb{C}}$-invariant and can be interpreted as a complex moment map condition in a hyper-Kähler reduction, while the remaining equation provides the real moment map condition. Put differently, the decoupled equations exhibit a Hermitian Yang-Mills structure, such that a Donaldson-Uhlenbeck-Yau type approach and other powerful tools become available.

Clearly, whenever $\mathbf{dHW}_{v,J}(A, B) = 0$, then $\mathbf{HW}_v(A, B) = 0$. Here, we prove that in certain situations the reverse is true, such that the Haydys-Witten equations reduce to the decoupled Haydys-Witten equations. This is controlled, on the one hand by the (asymptotic) geometry, and on the other hand by the asymptotic behaviour of the fields (A, B), at boundaries and cylindrical ends of M^5. Crucially, we need to assume that (a) boundaries of M^5 are flat and (A, B) satisfies Nahm pole boundary conditions with knot singularities, and (b) non-compact ends of M^5 are asymptotically locally Euclidean (ALE) or flat (ALF) gravitational instantons

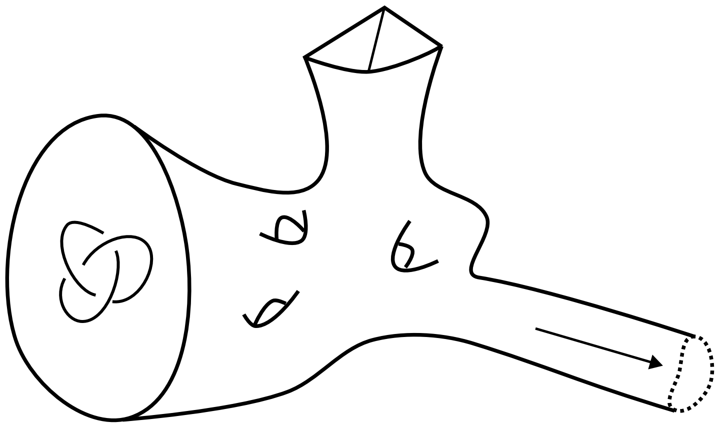

Figure 4.1 A polycylindrical five-manifold with boundaries and corners. In Haydys-Witten Floer theory one includes embeddings of knotted surfaces Σ_K as additional boundary faces by a geometric blowup. The blown-up boundaries intersect the original ambient boundary in corners of depth two.

and (A, B) approaches a finite energy solution of the θ-Kapustin-Witten equations. For the sake of brevity, we omit further technical conditions for now and instead refer to assumptions (A1) - (A4) in Section 4.5 for details. With that understood, the main result is as follows (this is Theorem 4.11 below).

Theorem D. *Let $G = SU(2)$, M^5 a manifold with poly-cylindrical ends, v a non-vanishing vector field that approaches ends at a constant angle, and J an almost Hermitian structure on* $\ker g(v, \cdot)$. *Assume $\mathbf{HW}_v(A, B) = 0$ and that assumptions* (A1) - (A4) *are satisfied, then $\mathbf{dHW}_{v,J}(A, B) = 0$.*

The proof of Theorem D is based on a Weitzenböck formula of the form

$$\int_{M^5} \|\mathbf{HW}_v(A, B)\|^2 = \int_{M^5} \|\mathbf{dHW}_{v,J}(A, B)\|^2 + \int_{M^5} d\chi \ . \tag{4.4}$$

From this it's clear that whenever $\mathbf{HW}_v(A, B) = 0$ and any boundary contributions (including contributions from non-compact ends) in $\int_{M^5} d\chi$ vanish, then also $\mathbf{dHW}_{v,J}(A, B) = 0$. The key insights of this chapter lie in determining conditions under which all boundary contributions vanish, if one imposes Nahm pole boundary conditions at finite distances and asymptotically stationary solutions at infinity. Due to the two different flavours of boundary conditions, we rely on two distinct facts: Elliptic regularity of the Nahm pole boundary conditions, on the one hand, and a vanishing theorem for solutions of θ-Kapustin-Witten solutions on ALE and ALF spaces, on the other. Let us shortly explain how these two facts appear in the proof.

First, the Nahm pole boundary conditions for (A, B) state, in particular, that at order y^{-1} the fields satisfy the extended Bogomolny equations (EBE), which in the $4\mathcal{D}$ formalism correspond to

$$\mathcal{D}_0 = 0 \ , \qquad [\mathcal{D}_i, \mathcal{D}_j] = 0 \quad i, j = 1, \dots, 3 \ , \qquad \sum_{i=1}^{3} [\mathcal{D}_i, \overline{\mathcal{D}_i}] = 0 \ .$$

This means that the leading order terms are already solutions of the decoupled Haydys-Witten equations. As will be discussed in detail below, elliptic regularity of the Haydys-Witten equations states that deviations from the EBE-solutions can only appear at order $y^{1+\delta}$, for some $\delta > 0$ [MW14; MW17; He18]. This, in turn, implies that χ only involves terms of order y^{δ}. As a consequence, any contributions to $\int d\chi$ from boundaries with Nahm pole boundary conditions vanish.

We expect that Theorem D is also generally true in the presence of knot singularities. Unfortunately, the twisted knot singularity solutions are only known implicitly (by a continuation argument) [Dim22a], such that extracting information from elliptic regularity is difficult. Although we currently have no proof for this extension of the result, we include a discussion of the relevant boundary conditions and state a necessary condition that is known to be satisfied in the untwisted case.

Second, and perhaps more surprisingly, the boundary terms also vanish at asymptotic ends when the fields approach stationary solutions of the Haydys-Witten equations, or equivalently solutions of the θ-Kapustin-Witten equations, with finite energy. This relies on the vanishing theorem given in Corollary C, originally conjectured by Nagy and Oliveira [NO21] and proven in Chapter 3. The vanishing theorem states that for a finite energy solution of the θ-Kapustin-Witten equations on an ALE or ALF gravitational instanton A is flat, ϕ is ∇^A-parallel, and $[\phi \wedge \phi] = 0$.

This chapter is structured as follows: In Section 4.1 we summarize Haydys' geometry and the Haydys-Witten equations, define the decoupled Haydys-Witten equations, and establish the promised Weitzenböck formula. In Section 4.2 we further specify the basic geometric setting that is necessary to specify boundary conditions and which is used in evaluating the integral of the exact term in the Weitzenböck formula. A key step in this discussion is an investigation of the polyhomogeneous expansion of a twisted Nahm pole solution of the Haydys-Witten equations, which is presented in Section 4.3. Subsequently, in Section 4.4, we determine the asymptotic behaviour of χ at the various boundaries and ends, filling in the details of the remaining boundary conditions as we go. Finally, in Section 4.5, we bring everything together and show that in certain situations the boundary term in (4.4) vanishes, which immediately implies Theorem D.

4.1 The Decoupled Haydys-Witten Equations and a Weitzenböck Formula

In this section, we introduce the global version of the Haydys-Witten equations and their decoupled version, and establish a Weitzenböck formula that relates it to the full Haydys-Witten equations.

Let (M^5, g) be an oriented complete Riemannian manifold. Consider a principal G-bundle $E \to M^5$ for some compact Lie group G, write $\mathcal{A}(E)$ for the space of gauge connections, and denote by $\operatorname{ad} E$ the adjoint bundle associated to the Lie algebra $\mathfrak{g}$ of G.

Assume M^5 admits a non-vanishing unit vector field v and write $\eta = g(v, \cdot) \in \Omega^1(M)$ for its dual one-form. Since v is non-vanishing, the distribution $\ker \eta \subset TM^5$ is regular, i.e. a vector subbundle, of rank 4. Observe that the pointwise linear map

$$T_\eta \, : \, \Omega^2(M) \to \Omega^2(M), \quad \omega \mapsto \star_5(\omega \wedge \eta)$$

has eigenvalues $\{-1, 0, 1\}$, such that $\Omega^2(M)$ decomposes into the corresponding eigenspaces:

$$\Omega^2(M) = \Omega^2_{v,-}(M) \oplus \Omega^2_{v,0}(M) \oplus \Omega^2_{v,+}(M) \, .$$

Here $\Omega^2_{v,0}(M)$ has rank 4, while the (anti-)self-dual parts have rank 3. We use the notation ω^+ to denote the part of ω that lies in $\Omega^2_{v,+}(M)$.

Since $\Omega^2_{v,+}(M)$ has rank 3, the usual cross product $(\cdot \times \cdot)$ of $\mathbb{R}^3$ provides it with a Lie algebra structure. The map $\sigma(\cdot, \cdot) = \frac{1}{2}(\cdot \times \cdot) \otimes [\cdot, \cdot]_{\mathfrak{g}}$ then determines a corresponding cross product on $\Omega^2_{v,+}(M, \operatorname{ad} E)$. This map is unique up to a choice of orientation and is locally given by the product on $\mathbb{R}^3 \otimes \mathfrak{g}$.

Consider a pair $(A, B) \in \mathcal{A}(E) \times \Omega^2_{v,+}(\operatorname{ad} E)$, consisting of a gauge connection A and a self-dual two form B. Let ∇^A denote the covariant derivative on $\Omega^2_{v,+}(\operatorname{ad} E)$ induced by A and the Levi-Civita connection, and define $\delta^+_A \, : \, \Omega^2_{v,+}(M, \operatorname{ad} E) \to \Omega^1(M, \operatorname{ad} E)$ by the composition $\delta^+_A B = -g^{\mu\rho} \iota_\mu \nabla^A_\rho B$. The Haydys-Witten equations are defined by:

$$\begin{aligned} F_A^+ - \sigma(B, B) - \nabla^A_v B &= 0 \, , \\ \iota_v F_A - \delta^+_A B &= 0 \, . \end{aligned} \tag{4.5}$$

Assume now that there is an almost complex structure J on the vector bundle $\ker \eta \to M^5$ that is compatible with the metric, i.e. $g(J\cdot, J\cdot) = g(\cdot, \cdot)$. Note that J lifts, via the tensor product $J \otimes J$, to a map on $\Omega^2_{v,+}(M^5)$. At each point this map has eigenvalues $\{+1, -1, -1\}$. To see this, let $p \in M^5$ and $(x^\mu, y)_{\mu=0,\ldots,3}$ be coordinates in a neighbourhood of p, such that $v = \partial_y$ and chosen in such a way that dx^μ is a canonical basis of J, i.e. $Jdx^0 = dx^1$ and $Jdx^2 = dx^3$. The two-forms $e_i = dx^0 \wedge dx^i + \frac{1}{2}\epsilon_{ijk} dx^j \wedge dx^k$, $i = 1, 2, 3$, are a local basis of $\Omega^2_{v,+}$. One easily sees that $Je_1 = +e_1$ and $Je_{2/3} = -e_{2/3}$.

The decoupled Haydys-Witten equations arise if one sets to zero the -1-eigenparts of the expressions in the Haydys-Witten equations (4.5) that contain B, independently from the remaining terms:

Definition 4.1 (Decoupled Haydys-Witten Equations)**.**

$$\begin{aligned} \tfrac{1+J}{2}\left(\sigma(B, B) + \nabla^A_v B\right) &= F_A^+ & \qquad \tfrac{1-J}{2}\left(\sigma(B, B) + \nabla^A_v B\right) &= 0 \\ \delta^+_A \tfrac{1+J}{2} B &= \iota_v F_A & \qquad \delta^+_A \tfrac{1-J}{2} B &= 0 \end{aligned}$$

We denote the associated differential operator by $\mathbf{dHW}_{v,J}$.

In these equations the part of the B-terms that is located in the negative eigenspaces is decoupled from the gauge curvature, hence the name. A direct calculation in coordinates shows that the decoupled Haydys-Witten equations are locally equivalent to (4.3) in the $4\mathcal{D}$-formalism.

To find a relation between $\mathbf{HW}_v$ and $\mathbf{dHW}_{v,J}$, split the terms in the Haydys-Witten equations that involve B into their positive and negative eigenparts with respect to J. To simplify notation, write $J^\pm := \frac{1\pm J}{2}$ as shorthand for the projection to the ± 1-eigenspaces of J in $\Omega^2_{v,+}$.

Denote by $\star$ the Hodge star operator and equip $\Omega^k(M^5, \operatorname{ad} E)$ with the density-valued inner product $\langle a, b\rangle = \operatorname{Tr} a \wedge \star b$. Upon integration this provides the usual L^2-product $\langle a, b\rangle_{L^2(W)} = \int_{W^n} \langle a, b\rangle$ on $\Omega^k(W^n, \operatorname{ad} E)$. The (density-valued) L^2-norm of the Haydys-Witten operator can be rewritten as follows

$$
\begin{aligned}
\|\mathbf{HW}_v(A,B)\|^2 &= \left\|F_A^+ - \left(J^+ + J^-\right)\left(\sigma(B,B) + \nabla_v^A B\right)\right\|^2 + \left\|\iota_v F_A - \delta_A^+ \left(J^+ + J^-\right) B\right\|^2 \\
&= \left\|\mathbf{dHW}_{v,J}(A,B)\right\|^2 - 2\left\langle F_A^+, J^-\left(\sigma(B,B) + \nabla_v^A B\right)\right\rangle - 2\left\langle \iota_v F_A - \delta_A^+ J^+ B, \delta_A^+ J^- B\right\rangle
\end{aligned}
$$

Here we used that the ± 1 eigenspaces of J in $\Omega^2_{v,+}$ are orthogonal with respect to $\langle\cdot,\cdot\rangle$ to remove one of the mixed terms on the right hand side. In the next step we will observe that the extra terms on the right hand side are in fact total derivatives.

First, notice that

$$
\begin{aligned}
d\operatorname{Tr}\left(F_A \wedge J^- B\right) &= \operatorname{Tr}\left(F_A \wedge d_A J^- B\right) = \operatorname{Tr}\left(F_A^+ \wedge d_A J^- B + F_A^0 \wedge d_A J^- B\right) \\
&= \operatorname{Tr}\left(F_A^+ \wedge \nabla_v^A J^- B \wedge \eta + \iota_v F_A \wedge d_A J^- B \wedge \eta\right) \\
&= \left\langle F_A^+, \nabla_v^A J^- B\right\rangle + \left\langle \iota_v F_A, \delta_A^+ J^- B\right\rangle \ .
\end{aligned}
$$

Here we used the following steps: On the first line we use the Bianchi identity $d_A F_A = 0$ and then decompose $F_A = F_A^+ + F_A^0 + F_A^-$, utilizing that $F_A^- \wedge B = 0$; for the second line note that only $\nabla_v^A \eta \wedge$ contributes to the wedge product; and for the last line we use that $B \wedge \eta = \star B$ and that the action of $d_A^\star$ on $J^- B$ coincides with δ_A^+.

Second, there is a similar equality for the remaining terms, as can be seen in local coordinates $(x^\mu, y)_{\mu=0,\dots,3}$ by expanding both sides in the local basis e_i, $i = 1,2,3$, of $\Omega^2_{v,+}(M)$ that was introduced above.

$$
d\operatorname{Tr}\left(\delta_A^+ J^+ B \wedge J^- B \wedge \eta\right) = \left\langle F_A^+, J^- \sigma(B,B)\right\rangle - \left\langle \delta_A^+ J^+ B, \delta_A^+ J^- B\right\rangle \ .
$$

With these identifications we arrive at the following lemma.

Lemma 4.2. *There is a Weitzenböck formula adapted to the decoupled Haydys-Witten equations*

$$
\int_{M^5} \|\mathbf{HW}_v(A,B)\|^2 = \int_{M^5} \left(\left\|\mathbf{dHW}_{v,J}(A,B)\right\|^2 + d\chi\right) \ ,
$$

where the exact term $\chi = \chi_1 + \chi_2$ is given by

$$
\begin{aligned}
\chi_1 &= -2\operatorname{Tr}\left(F_A \wedge J^- B\right) \\
\chi_2 &= -2\operatorname{Tr}\left(\delta_A^+ J^+ B \wedge J^- B \wedge \eta\right)
\end{aligned}
$$

4.2 Poly-Cylindrical Ends and Boundary Conditions

In Witten's approach to Khovanov homology, one considers manifolds of the form $\mathbb{R}_s \times X^3 \times \mathbb{R}^+_y$, which possess both a boundary and non-compact ends. Moreover, the non-vanishing vector field $v = \partial_y$ differentiates between the non-compact directions $\mathbb{R}_s$ and $\mathbb{R}^+_y$, by way of the glancing angles $\theta = \pi/2$ and $\beta = 0$, respectively. This structure is well-described by manifolds with poly-cylindrical ends and boundaries.

In the literature, poly-cylindrical ends are typically assumed to be located at geodesic infinity. In this case poly-cylindrical manifolds are identified with the interior of a manifold with corners that is equipped with a metric, such that any boundary point lies at infinity and the metric is of product type within some small tubular neighborhood. Sometimes the definition is further relaxed by assuming that the metric on the manifold with corners is an *exact b-metric* in the sense of Melrose [Mel95; Mel96]. In this latter case the metric on M approaches a product metric exponentially fast, or put differently, its asymptotic expansion in $\{x^k\}_{k\in\mathbb{Z}}$ vanishes to all orders. We expect that our results can be generalized to this situation, though we don't further pursue this here.

In our situation, we also want to include boundaries and corners at finite distance. Therefore, when we talk about a (poly-)cylindrical end of a manifold we always take this to encompass boundaries (and corners) at both, finite and infinite geodesic distance. We will now set up notation that will be used in the rest of this chapter and provide definitions for the underlying geometry. After that, we also provide a first description of the relevant boundary conditions.

Let M be a manifold with corners. For any $p \in M$ choose a chart (U, ϕ) with $\phi(x) = p$ and define the *depth* of p by the number of components of x that are zero. This coincides with the number of boundary faces of M that contain p and is independent of the choice of chart. Define the depth-k stratum of M to be the points

$$S^k(M) := \{\, p \in M \mid \; \mathrm{depth}(p) = k \,\} \; .$$

The interior of M corresponds to $S^0(M)$, while all boundary faces of codimension one are contained in the boundary stratum $\partial M := S^1(M)$, and more generally points that lie on corners of codimension k are collected in $S^k(M)$. Clearly $M = \sqcup_{k=0}^n S^k(M)$, providing a stratification of M. Each stratum $S^k(M)$ is a manifold of dimension $n - k$, since corners of higher codimension are explicitly excluded.

The boundary stratum is a disjoint union of connected components $\partial M = \sqcup_{i\in I}\partial_i M$, where for simplicity we assume that I is finite. Let us emphasize that each boundary face $\partial_i M$ is an open manifold without boundary. We typically denote a boundary defining function of the boundary face $\partial_i M$ by s_i.

Definition 4.3. A manifold with poly-cylindrical ends is a complete Riemannian manifold (M, g) that is diffeomorphic to a submanifold (with corners) of a compact manifold with corners M_0, equipped with a metric g_0 that pulls back to g, and such that the following conditions are satisfied: For each boundary component $\partial_i M_0$ there exists a tubular neighbourhood $[0,\epsilon)_{s_i} \times \partial_i M_0$, on which g_0 is either

$$g_0 = ds_i^2 + h_{\partial_i M} \qquad \text{or} \qquad g_0 = \frac{ds_i^2}{s_i^2} + h_{\partial_i M}\,,$$

where $h_{\partial_i M}$ is a complete metric on $\partial_i M$. Moreover, this is compatible at corners, i.e. there exists a neighbourhood $[0,\epsilon)^m \times X^{n-m}$ of each connected component of $S^k(M)$, where

$$g_0 = \sum_{k\in K} \frac{ds_k^2}{s_k^2} + \sum_{i\in I\setminus K} ds_i^2 + h_{X^{n-m}}$$

Example. Let $M^5 = \mathbb{R}_s \times X^3 \times \mathbb{R}_y^+$, together with the product metric $g = ds^2 + g_{X^3} + dy^2$. Consider the compact manifold with corners $M_0 = [-1,1]_{s_0} \times X^3 \times [0,1]_{y_0}$, equipped with $g_0 = \big(\frac{ds_0}{1-|s_0|}\big)^2 + h_{X^3} + \big(\frac{dy_0}{1-y_0}\big)^2$. Then the map that takes $s_0 \mapsto s = -\operatorname{sgn}(s_0)\log(1-|s_0|)$ and $y_0 \mapsto y = -\log(1-y_0)$ defines an isometry between the submanifold $U = (-1,1)_{s_0} \times X^3 \times [0,1)_{y_0}$ and (M^5, g).

A poly-cylindrical manifold admits a convenient compact exhaustion by a family of manifolds with corner M_ϵ that mirrors the poly-cylindrical structure of M. To make this precise, note that $h := \big(\prod_{k\in K} s_k^2\big) \cdot g_0 : TM_0 \times TM_0 \to \mathbb{R}$ defines a Riemannian metric on M_0. Let $d_h(x,y)$ denote the induced distance function and define $M_\epsilon = \{\, x \in M_0 \mid d_h(x, \partial_i M_0) \geq \epsilon \,\}$. We equip M_ϵ with the restriction of g, making it into a compact Riemannian manifold with poly-cylindrical ends that approaches M for $\epsilon \to 0$. We will use the manifolds M_ϵ to regularize the integrals in the Weitzenböck formula Lemma 4.2. In particular, since M_ϵ is compact, the following version of Stokes' theorem holds.

Theorem 4.4 (Stokes' Theorem on Manifolds with Corners [Whi57])**.** *Let M_ϵ be a compact manifold with corners, then*

$$\int_{M_\epsilon} d\chi = \sum_{i\in I} \int_{\partial_i M_\epsilon} \chi\,,$$

where the sum on the right hand side is over all boundary faces of codimension one.

Each of the boundaries of M_ϵ can be thought of as an 'ϵ-displacement' of the corresponding end of M into the interior. By this we mean that for each end $[0,1)_s \times W$ of M, the corresponding end of M_ϵ can be identified with an embedding of $[\epsilon,1]_s \times W \hookrightarrow [0,1)_s \times W$. At non-compact ends, the distance between points with $s = \epsilon$ and $s = 1$ is finite and proportional to ϵ^{-1}, while points at $s = 0$ reside at infinity.

From now on let (M^5, g) denote a poly-cylindrical five-manifold and write (M_0, g_0) for a suitable ambient manifold. Throughout, we take ∂M^5 to include both, its boundary components and non-compact cylindrical ends, i.e. we identify the boundary faces $\partial_i M^5$ with the corresponding boundary components of M_0.

Let $[0,\epsilon)_s \times W^4$ be a tubular neighbourhood of a cylindrical end and denote the inward-pointing unit-normal vector field by u. In case of a boundary this means that $u = \partial_s$, while for a non-compact end we have $u = s\partial_s$. We say v approaches a cylindrical end at constant angle if there is a tubular neighbourhood[1] on which the incidence angle $g(u,v) = \cos\theta$ is constant. This angle determines the natural boundary conditions at a given cylindrical end.

We will now give a first overview of the relevant boundary conditions, first for boundaries and then for non-compact ends. The classes of boundary conditions serve as organizational structure for the rest of the chapter. A detailed discussion, in particular for the Nahm pole and knot singularity model solutions, will be given later, see Section 4.3 and Section 4.4.2, respectively.

At a boundary with incidence angle $g(u,v) = \cos\beta$, $\beta \neq \pi/2$, we impose β-twisted Nahm pole boundary conditions, which fix the maximal rate of growth of (A,B) as one approaches the boundary. The condition specifies that (A,B) are asymptotic to a model configuration $(A^{\rho,\beta}, B^{\rho,\beta})$ of order $\mathcal{O}(y^{-1})$. The model configuration is associated to a choice of Lie algebra homomorphism $\rho : \mathfrak{su}(2) \to \mathfrak{g}$.

It is also possible to include knot singularities of weight $\lambda \in \Gamma^{\vee}_{\text{char}}$ in the co-character lattice of $\mathfrak{g}$. In that case one considers the geometric blowup of a surface $\Sigma_K \subset \partial M$, which introduces an additional boundary $\partial_K M \simeq \Sigma_K \times H^2$, where H^2 denotes the two-dimensional hemisphere. While the incidence angle of v at this new boundary can no longer be constant, we assume that v maintains a constant glancing angle θ with Σ_K. At the knot boundary $\partial_K M$, the fields must then be asymptotic to another model configuration $(A^{\lambda,\theta}, B^{\lambda,\theta})$ of order $\mathcal{O}(R^{-1})$, where R denotes the radial distance to Σ_K.

The Haydys-Witten pair (A,B) satisfies β-twisted Nahm pole boundary conditions at ∂M, with knot singularity of weight $\lambda \in \Gamma^{\vee}_{\text{char}}$ along Σ_K, if for some $\epsilon > 0$

(i) near $\partial_0 M$: $(A,B) = \left(A^{\rho,\beta},\ A^{\rho,\beta}\right) + \mathcal{O}\left(y^{-1+\epsilon}\right)$

(ii) near $\partial_K M$: $(A,B) = (A^{\lambda,\theta}, B^{\lambda,\theta}) + \mathcal{O}(R^{-1+\epsilon})$

Consider now the case of a non-compact end $[0,\epsilon)_s \times W^4$ with incidence angle $g(u,v) = \cos\theta$. Natural boundary conditions at non-compact ends are given by stationary (s-invariant) solutions of the underlying differential equations. It turns out that the specialization to s-invariant Haydys-Witten equations is given either by the Vafa-Witten equations or the θ-Kapustin-Witten, depending on the value of θ:

$$\mathbf{HW}_v(A,B) \overset{s\text{-inv.}}{\rightsquigarrow} \begin{cases} \mathbf{VW}(\tilde{A},\tilde{B},C) & \theta \equiv 0 \pmod{\pi}, \\ \mathbf{KW}_\theta(\tilde{A},\phi) & \text{else} \end{cases}$$

If $\theta \neq 0 \pmod{\pi}$, there must be a non-vanishing vector field w on W^4, such that $v = \cos\theta u + \sin\theta w$. Moreover, we can define an orthogonal vector field $v^\perp = -\sin\theta u + \cos\theta w$. With this

[1] This is equivalent to an asymptotic equivalence of $g(u,v) - \cos\theta \sim 0$ as $s_i \to 0$. It seems reasonable that this condition can be slightly weakened to $g(u,v) \sim \cos\theta$.

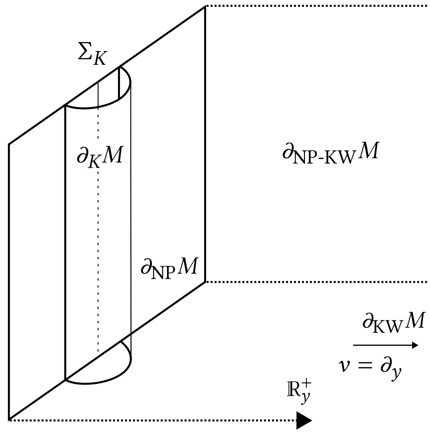

Figure 4.2 Witten's setup is captured by a polycylindrical manifold $M^5 = \mathbb{R}_s \times X^3 \times \mathbb{R}_y^+$, blown-up along a surface embedding $\Sigma_K = \mathbb{R}_s \times K$ in the boundary at $y = 0$. For each of the four boundary classes, the incidence angles of the vector field $v = \partial_y$ at a given boundary component determines the natural boundary condition.

notation, the four-dimensional fields above are given in terms of (A, B) as follows

$$\tilde{A} := i^* A \,, \qquad \begin{aligned} &\tilde{B} = i^* B \,, \; C = A_s \\ &\phi := i^* \left(\iota_{v^\perp} B + A_s \wedge w^\flat \right) \end{aligned}$$

We will consider stationary solutions that either satisfy a finite energy condition, or that exhibit a $\beta = (\pi/2 - \theta)$-twisted Nahm pole at one of the adjacent corners of M^5.

We arrange the ends of M^5, in accordance with these boundary conditions, into the following four classes:

(1) Nahm pole boundaries $\partial_{\text{NP}} M$,
(2) Knot boundaries $\partial_K M$,
(3) Kapustin-Witten ends with finite energy $\partial_{\text{KW}} M$, and
(4) Kapustin-Witten ends with Nahm poles, denoted by $\partial_{\text{NP-KW}} M$.

Let us shortly exemplify this in connection with the gauge theoretic approach to Khovanov homology. In this situation one considers the five-manifold $M^5 = \mathbb{R}_s \times X^3 \times \mathbb{R}_y^+$ together with the vector field $v = \partial_y$. In presence of a knot $K \subset X^3$, one then moves to the geometric blowup $[M^5; \Sigma_K]$, where $\Sigma_K = \mathbb{R}_s \times K$. The blown-up manifold has two types of boundary components. First, a Nahm pole boundary $\partial_{\text{NP}} M$, corresponding to those original boundary points at $y = 0$ that are not part of the Σ_K. Second, a knot boundary $\partial_K M^5$, given by the points of the unit normal bundle over the surface Σ_K. M^5 also has several cylindrical ends. A Kapustin-Witten end appears at $y \to \infty$, where the angle between v and the inward pointing unit normal vector $-\partial_y$ is $\theta = 0$, such that the fields approach a solution of the Vafa-Witten equations (hence, this should really be called a Vafa-Witten end). Kapustin-Witten ends with additional Nahm poles at a corner of M^5 appear at $s \to \pm\infty$, where fields approach a Kapustin-Witten solutions with a Nahm pole located at the corner $\{y = 0, s = \pm\infty\}$.

Remark. In case (3) the fields (A, B) approach a stationary, finite energy solution of the Kapustin-Witten equations, such that one can determine the behaviour of χ with help of the vanishing theorem for θ-Kapustin-Witten solutions. In the situation of (4) we similarly ask for a stationary solution, but the fields exhibit a Nahm pole at an adjacent corner. In this situation the fields diverge and, in particular, do not have finite energy.

4.3 Polyhomogeneous Expansion of Twisted Nahm Pole Solutions

In this section we investigate the asymptotic behaviour of a Haydys-Witten solution (A, B) in the vicinity of a Nahm pole boundary $\partial_{\mathrm{NP}} M$, i.e. under the assumption that (A, B) satisfy twisted Nahm pole boundary conditions without knot singularities. The goal of this analysis is to extract the leading order behaviour of F_A and $\delta_A^+ J^+ B$ near a Nahm pole boundary, which is later used to determine the asymptotics of χ.

The analysis goes back to an essentially identical discussion for the $\theta = \pi/2$ version of the Kapustin-Witten equations with untwisted Nahm pole boundary conditions by Mazzeo and Witten in [MW14]. There, a Weitzenböck formula on $\mathbb{R}^3 \times \mathbb{R}^+$ very similar to Lemma 4.2 – in fact a more restrictive one – was established and then utilized to determine the analytic properties of the Nahm pole boundary conditions in the first place. While Mazzeo and Witten already touched on the leading and constant orders for general four-manifolds, Siqi He has later expanded on the constraints that arise for subleading orders on general four-manifolds and unveiled a deep relation to the geometry of the boundary [He18].

To set the stage, let $U = W^4 \times [0, 1)_y$ be a tubular neighbourhood of a Nahm pole boundary $\partial_{\mathrm{NP}} M$. Write $u = \partial_y$ for the inward-pointing, unit normal vector field. As always, assume that the incidence angle determined by $g(u, v) = \cos\beta$ is constant on all of U. Whenever $\beta \neq 0$ there is a non-vanishing unit vector field w on W^4, such that $v = \sin\beta w + \cos\beta u$.

Let us shortly recall some additional details about the Nahm pole boundary conditions on $\partial_{\mathrm{NP}} M$, see [MW14; MW17]. Let $\phi_\rho \in \Omega^2_{v,+}(U, \operatorname{ad} E)$ be some fixed boundary configuration associated to the Nahm pole boundary conditions on U. This means that it satisfies $\phi_\rho - \sigma(\phi_\rho, \phi_\rho) = 0$, such that its components $(\phi_\rho)_i = \mathfrak{t}_i$, $i = 1, 2, 3$, span an $\mathfrak{su}(2)$ subalgebra inside of $\mathfrak{g}$. There is a unique corresponding two-form that satisfies $\phi_\rho^\tau + \sigma(\phi_\rho^\tau, \phi_\rho^\tau) = 0$. Here $\tau = (132)$ denotes an anticyclic permutation that acts on components by $(\phi_\rho^\tau)_i = (\phi_\rho)_{\tau(i)}$ and is related to a reversal of orientation on $\Omega^2_{v,+}(U)$. If $\beta \neq 0$, there is an isomorphism by which we can view elements of $\Omega^2_{v,+}(U)$ as the one-forms dual to a rank 3 subbundle $\Delta^\perp_{(u,v)} \subset TU$. Here, $\Delta^\perp_{(u,v)}$ is the orthogonal complement of the subbundle spanned by u and v. With that in mind, the Nahm pole boundary condition states that as $y \to 0$, A and B are asymptotic to

$$A^{\rho,\beta} = \sin\beta \, \frac{\phi_\rho^\tau}{y} \,, \qquad B^{\rho,\beta} = \cos\beta \, \frac{\phi_\rho}{y} \,.$$

Near a Nahm pole, the Haydys-Witten equations are amended by the following gauge fixing condition

$$d^{\star_5}_{A^0}(A - A^0) + \sigma(B^{\mathrm{NP}}, B - B^{\mathrm{NP}}) = 0 ,$$

making the equations into an elliptic system of differential equations.

Ultimately, we want to investigate the leading order of the functions $\chi_1 = -2\,\mathrm{Tr}(F_A \wedge J^- B)$ and $\chi_2 = -2\,\mathrm{Tr}(\delta_A^+ J^+ B \wedge J^- B \wedge \eta)$ near $\partial_{\mathrm{NP}} M$. Since this involves derivatives and products of A and B, we need to study the subleading orders in a polyhomogeneous expansion of A and B as $y \to 0$.

Write $A = y^{-1} \sin\beta\, \phi_\rho^\tau + \omega + a$ and $B = y^{-1} \cos\beta\, \phi_\rho + b$, where by assumption of the Nahm pole boundary condition $a, b \in \mathcal{O}(y^{-1+\epsilon})$, while ω is a gauge connection on the restriction of E to W^4, smoothly extended into the bulk. If (A, B) is a solution of the Haydys-Witten equations that satisfies Nahm pole boundary conditions, then an elliptic regularity theorem Mazzeo and Witten [MW14] states that there exist polyhomogeneous expansions

$$a \sim \sum_{(\alpha,k)\in\Delta_0} a_{\alpha,k}(x^\mu)\, y^\alpha (\log y)^k , \qquad b \sim \sum_{(\alpha,k)\in\Delta_0} b_{\alpha,k}(x^\mu)\, y^\alpha (\log y)^k \qquad (y \to 0) ,$$

where the indicial set $\Delta_0 \subset \mathbb{C} \times \mathbb{N}$ is bounded from the left by 1 (in particular $a_{0,k} = b_{0,k} = 0$).

Plugging $A = y^{-1} \sin\beta\, \phi_\rho^\tau + \omega + a$ and $B = y^{-1} \cos\beta\, \phi_\rho + b$ into the Haydys-Witten equations then determines a sequence of constraints on the functions $a_{\alpha,k}$ and $b_{\alpha,k}$. For our purposes it will suffice to determine χ up to $\mathcal{O}(y^\delta)$, $\delta > 0$, so we only need to consider an expansion up to order $y^{2+\delta}$. We will see below that for generic values of β only the following terms in the expansion are non-zero.

$$\begin{aligned} a &= a_{1,1}\, y \log y + a_{1,0}\, y + a_{2,1}\, y^2 \log y + a_{2,0}\, y^2 + \mathcal{O}(y^{5/2}) \\ b &= b_{1,1}\, y \log y + b_{1,0}\, y + b_{2,1}\, y^2 \log y + b_{2,0}\, y^2 + \mathcal{O}(y^{5/2}) \end{aligned}$$

Moreover, the coefficient functions satisfy additional constraints that link them to the Riemannian curvature of the manifold. In particular, the $\log y$-terms vanish if sectional curvature is flat in any plane that contains w.

To make this more precise, we need to establish analogues of the tools that were used in the discussion of indicial roots in [MW14] and [He18]. First, the fibers of $\mathrm{ad}\, E$ decompose under the action of $(\phi_\rho)_i = \mathfrak{t}_i \in \mathfrak{su}(2)_\mathfrak{t}$ as a direct sum of spin j representations. We denote the associated vector bundles by V_j, $j \in J$. Since we consider regular Nahm pole conditions, only positive integers j appear; in fact for $G = SU(N)$ one finds $J = \{1, \dots, N-1\}$ [MW14]. Second, there is an action by rotations $\mathfrak{so}(3) \simeq \mathfrak{su}(2)_\mathfrak{s}$ on $\Omega^2_{v,+}(W^4)$ and $(\Delta^\perp_{(u,v)})^*$. For fixed j, the bundles $\Omega^2_{v,+}(W^4, \mathrm{ad}\, E)$ and $\mathrm{Hom}(\Delta^\perp_{(u,v)}, V_j) \subset \Omega^1(W^4, \mathrm{ad}\, E)$ thus further decomposes under a combined $\mathfrak{su}(2)_\mathfrak{f}$ action, generated by $\mathfrak{f}_i = \mathfrak{t}_i + \mathfrak{s}_i$, into a direct sum of bundles V_j^-, V_j^0, and V_j^+ of total spin $j-1$, j and $j+1$, respectively. To simplify notation, we will furthermore write $V_j^w :=$

$C^\infty(W^4, V_j)\, w^\flat$, where we let $\mathfrak{su}(2)_\mathfrak{s}$ acts trivially. All in all, the relevant $\operatorname{ad} E$-valued differential forms decompose as

$$\Omega^1(W^4, \operatorname{ad} E) = \bigoplus_{j\in J} \left(V_j^w \oplus V_j^- \oplus V_j^0 \oplus V_j^+ \right)$$
$$\Omega^2_{v,+}(W^4, \operatorname{ad} E) = \bigoplus_{j\in J} \left(V_j^- \oplus V_j^0 \oplus V_j^+ \right)$$

This decomposition is helpful, because the leading order terms $A^{\mathrm{NP}} = y^{-1}\sin\beta\phi_\rho^\tau$ and $B^{\mathrm{NP}} = y^{-1}\cos\beta\phi_\rho$ are invariant under the action of $\mathfrak{su}(2)_\mathfrak{f}$ and, as a consequence, the constraints for $a_{\alpha,k}$ and $b_{\alpha,k}$ decompose into their V_j^η-valued components. However, there is an additional subtlety regarding invariance of the Nahm pole terms under the action of $\mathfrak{su}(2)_\mathfrak{f}$. To see this, let rotations $\mathfrak{su}(2)_\mathfrak{s}$ act on $\Omega^2_{v,+}$ with the 'standard' orientation on (e_1, e_2, e_3) and assume this basis is identified with $(dx^1, dx^2, dx^3) \in (\Delta^\perp_{(u,v)})^*$. Then $\phi_\rho = \sum_i \mathfrak{t}_i e_i$ or equivalently $\sum_i \mathfrak{t}_i dx^i$ is an element of the trivial representation $V_1^- \subset \Omega^2_{v,+}(W^4, \operatorname{ad} E) \simeq \operatorname{Hom}(\Delta^\perp_{(u,v)}, \operatorname{ad} E)$, such that B^{NP} is invariant. Unfortunately, under this identification $\phi_\rho^\tau = \sum_i \mathfrak{t}_{\tau(i)} dx^i = \mathfrak{t}_1 dx^1 + \mathfrak{t}_2 dx^3 + \mathfrak{t}_3 dx^2$ is not invariant, since orientation reversal $\phi_\rho \mapsto \phi_\rho^\tau$ does not preserve V_1^-. To remedy this, we need rotations to act on the $\Delta^\perp_{(u,v)}$-part of A with the opposite orientation (dx^1, dx^3, dx^2). With that choice ϕ_ρ^τ is an element of the trivial representation $\tilde{V}_1^-$, with respect to the analogous (but different) decomposition

$$\Omega^1(W^4, \operatorname{ad} E) = \bigoplus_{j\in J} \left(\tilde{V}_j^w \oplus \tilde{V}_j^- \oplus \tilde{V}_j^0 \oplus \tilde{V}_j^+ \right) .$$

A^{NP} becomes invariant under $\mathfrak{su}(2)_\mathfrak{f}$ only with that choice. We denote the restriction of a given differential form α to one of the subspaces V_j^η or $\tilde{V}_j^\eta$ in this decomposition by $(\alpha)^{j,\eta}$ or $(\alpha)^{\widetilde{j,\eta}}$, respectively.

Example. In a local basis $(e_i)_{i=1,2,3}$ of $\Omega^2_{v,+}(U, \operatorname{ad} E)$ the subbundles $V_1^\eta \subset \Omega^2_{v,+}(W^4, \operatorname{ad} E)$ are given by:

$$\begin{aligned} V_1^- &= \operatorname{span}\{\mathfrak{t}_1 e_1 + \mathfrak{t}_2 e_2 + \mathfrak{t}_3 e_3\} , \\ V_1^0 &= \operatorname{span}\{\mathfrak{t}_i e_j - \mathfrak{t}_j e_i\}_{i\neq j} , \\ V_1^+ &= \operatorname{span}\{\mathfrak{t}_1 e_1 - \mathfrak{t}_2 e_2,\ \mathfrak{t}_1 e_1 - \mathfrak{t}_3 e_3,\ \mathfrak{t}_i e_j + \mathfrak{t}_j e_i\}_{i\neq j} . \end{aligned}$$

Substituting e_i by dx^i provides the corresponding basis under the isomorphism of $\Omega^2_{v,+}(W^4, \operatorname{ad} E)$ and $\operatorname{Hom}(\Delta^\perp_{(u,v)}, \operatorname{ad} E) \subset \Omega^1(W^4, \operatorname{ad} E)$. This is the correct decomposition for the action of $\mathfrak{su}(2)_\mathfrak{f}$ on B and its expansions ϕ_ρ, $b_{\alpha,k}$. In contrast, for the decomposition of $\Omega^1(W^4, \operatorname{ad} E)$ with respect to the reverse orientation, a corresponding basis of $\tilde{V}_1^\eta$, $\eta \in \{-, 0, +\}$, is obtained from the one above by replacing e_i with $dx^{\tau(i)}$. For example $\tilde{V}_1^- = \operatorname{span}\{\mathfrak{t}_1 dx^1 + \mathfrak{t}_3 dx^2 + \mathfrak{t}_2 dx^3\}$. This is the correct basis for the action of $\mathfrak{su}(2)_\mathfrak{f}$ on A and its expansions ϕ_ρ^τ, $a_{\alpha,k}$.

The upcoming lemma and its proof specify the algebraic constraints the Haydys-Witten equations put on the coefficient functions in the polyhomogeneous expansion up to order y^2. See

[He18] for a similar and more extensive investigation of these constraints in the case of $\pi/2$-Kapustin-Witten solutions with Nahm pole boundary conditions.

As becomes clear during the proof, we have to exclude a finite set of angles at which certain spin-j modes of $\iota_w a$, i.e. the components $(a_{\alpha,k})^{j,w}$, become free parameters. Explicitly, these values are given by:

$$\cos 2\beta \in \left\{-\tfrac{2j+3}{(j+1)^2},\ \tfrac{2j-1}{j^2},\ -\tfrac{3j^2-4j+3}{j(j+1)^2},\ -\tfrac{3j^2+2j-4}{j^2(j+1)}\right\}_{j\in J}$$

The statement of the following lemma holds whenever β is not of that form.

Lemma 4.5. *Assume $A = y^{-1}\sin\beta\,\phi_\rho^\tau + \omega + a$ and $B = y^{-1}\cos\beta\,\phi_\rho + b$ are a solution of the Haydys-Witten equations on $W^4 \times \mathbb{R}^+_y$ with respect to the non-vanishing vector field $v = \sin\beta w + \cos\beta\partial_y$. Then ω pulls back under ϕ_ρ to the Levi-Civita connection on TW^4 and $\nabla^\omega_w \phi_\rho = 0$. Furthermore, if β is generic, then $a_{1,k} = b_{1,k} = a_{2,k} = b_{2,k} = 0$ for all $k \geq 2$. The solution is smooth up to the boundary (i.e. $a_{1,1} = b_{1,1} = 0$), if and only if the Riemannian curvature satisfies $(\cos\beta\iota_w F_\omega + \iota_w \star_4 F_\omega)|_{V_1^+/\tilde{V}_1^0} = 0$. Finally, if $F_\omega = 0$, the polyhomogeneous expansion up to order $y^{2+\delta}$, $\delta > 0$, reduces to*

$$\begin{aligned} a &= y\sin\beta\,(C^{1,+})^\tau + y^2\left(\sin\beta\,C^{2,+} + \cos\beta\,D^{1,-}\right)^\tau + \mathcal{O}(y^{2+\delta}) \\ b &= y\cos\beta\,C^{1,+} + y^2\left(\cos\beta\,C^{2,+} - \sin\beta\,D^{1,-}\right) + \mathcal{O}(y^{2+\delta}) \end{aligned}$$

where $C^{1,+} \in V_1^+/\tilde{V}_1^0$, $D^{1,-} \in V_1^-$, and $C^{2,+} \in V_2^+/\tilde{V}_2^0$ remain unconstrained.

Corollary 4.6. *Denote by $i_y : W^4 \hookrightarrow W^4 \times \mathbb{R}^+_y$ inclusion at y. For generic β and if $F_\omega = 0$, the total field strength pulls back to*

$$\begin{aligned} i_y^*(F_A) = \;& y^{-2}\left(\sin^2\beta\,[\phi_\rho^\tau \wedge \phi_\rho^\tau]\right) + y^0\left(\sin^2\beta\,[\phi_\rho^\tau \wedge (C^{1,+})^\tau]\right) \\ &+ y^1\left(\sin\beta\,d_\omega(C^{1,+})^\tau + \sin^2\beta[\phi_\rho^\tau \wedge (C^{2,+})^\tau] + \sin\beta\cos\beta\,[\phi_\rho^\tau \wedge (D^{1,-})^\tau]\right) + \mathcal{O}(y^{1+\delta}) \end{aligned}$$

Meanwhile, the dx^2- and dx^3- components of $\delta_A^+ J^+ B$ are specified by

$$\begin{aligned} i_y^*(\delta_A^+ J^+ B) \propto\;& y^{-2}\sin\beta\cos\beta\left([(\phi_\rho^\tau)_3, (\phi_\rho)_1]dx^2 - [(\phi_\rho^\tau)_2, (\phi_\rho)_1]dx^3\right) \\ &+ y^0\sin\beta\cos\beta\,\Big(\left([(\phi_\rho^\tau)_3, (C^{1,+})_1] + [(C^{1,+})^\tau_3, (\phi_\rho)_1]\right)dx^2 \\ &\qquad - \left([(\phi_\rho^\tau)_2, (C^{1,+})_1] + [(C^{1,+})^\tau_2, (\phi_\rho)_1]\right)dx^3\Big) \\ &+ y^1\Big(\left(\cos\beta\,\nabla_3^\omega(C^{1,+})_1 + \sin\beta\cos\beta\,[(\phi_\rho^\tau)_3, (C^{2,+})_1] - \sin^2\beta\,[(\phi_\rho^\tau)_3, (D^{1,-})_1]\right)dx^2 \\ &\qquad - \left(\cos\beta\,\nabla_2^\omega(C^{1,+})_1 + \sin\beta\cos\beta[(\phi_\rho^\tau)_2, (C^{2,+})_1] - \sin^2\beta\,[(\phi_\rho^\tau)_2, (D^{1,-})_1]\right)dx^3\Big) \\ &+ \mathcal{O}(y^{1+\delta}) \end{aligned}$$

The rest of this section is occupied with the proof of Lemma 4.5.

Proof. The proof is analogous to the analysis in [He18], which covers the case $\beta = 0$. Hence, assume that $\beta \neq 0$, in which case there is an isomorphism $\Omega^2_{v,+}(U, \mathrm{ad}\, E) \simeq \mathrm{Hom}(\Delta^\perp_{(u,v)}, \mathrm{ad}\, E)$. We use this equivalence to first re-express the Haydys-Witten equations (4.5) purely in terms of one-forms. This simplifies the evaluation of the formulas later on.

Fix a gauge in which $A_y = 0$ and choose as reference connection $A^0 = y^{-1} \sin\beta\, \phi^\tau_\rho + \omega$, where ω is the pullback of a connection on $E \to W^4$. Note that when we view B as an element of $\mathrm{Hom}(\Delta^\perp_{(u,v)}, \mathrm{ad}\, E) \subset \Omega^1(W^4 \times \mathbb{R}^+_y, \mathrm{ad}\, E)$, it does not have a dy-component: in coordinates $(s, x^i, y)_{i=1,2,3}$ with $v = \sin\beta \partial_s + \cos\beta \partial_y$, the two-form $B = \sum_{i=1}^3 \phi_i e_i$ is identified with the one-form $\sum_{i=1}^3 \phi_i dx^i$. Since the coefficients ϕ^τ_ρ, ϕ_ρ, $a_{\alpha,k}$ and $b_{\alpha,k}$ are independent of y and as they don't have dy-components, we can view them as elements of $\Omega^1(W^4, \mathrm{ad}\, E)$. Motivated by this, we slightly abuse notation in the upcoming version of the Haydys-Witten equations (4.6) - (4.9) and treat A and B as one-forms on $W^4 \times \{y\}$. (Specifically, we implicitly take the pullback and drop any dy-components in F_A and $d_A B$.) With that understood, the Haydys-Witten equations can be brought into the form

$$\partial_y A = -\iota_w (d_A B - \sin\beta\ \star_4 (F_A - [B \wedge B]) - \cos\beta\ \star_4 d_A B) \tag{4.6}$$

$$\partial_y B = \iota_w (F_A + \cos\beta\ \star_4 (F_A - [B \wedge B]) - \sin\beta\ \star_4 d_A B) \tag{4.7}$$

$$\partial_y(\iota_w A) = d_A^{\star_4} B \tag{4.8}$$

$$0 = d_\omega^{\star_4} a + \sin\beta \star_4 [\phi^\tau_\rho \wedge \star_4 y^{-1} a] + \cos\beta \star_4 [\phi_\rho \wedge \star_4 y^{-1} b] \tag{4.9}$$

Plugging the polyhomogeneous expansions of A and B into the terms that appear in these equations yields:

$$\begin{aligned} F_A &= y^{-2}\big(\sin^2\beta\, [\phi^\tau_\rho \wedge \phi^\tau_\rho] + \sin\beta\, \phi^\tau_\rho \wedge dy\big) \ +\ y^{-1}\big(\sin\beta\, d_\omega \phi^\tau_\rho\big) \ +\ y^0\big(F_\omega\big) \\ &\quad + \sum_{(\alpha,k)\geq 0} y^\alpha (\log y)^k \Big(d_\omega a_{\alpha,k} + \sin\beta[\phi^\tau_\rho \wedge a_{\alpha+1,k}] - (\alpha+1) a_{\alpha+1,k} \wedge dy - (k+1) a_{\alpha+1,k+1} \wedge dy\Big) \end{aligned}$$

$$[B \wedge B] = y^{-2}\Big(\cos^2\beta\, [\phi_\rho \wedge \phi_\rho]\Big) + \sum_{(\alpha,k)\geq 0} y^\alpha (\log y)^k \Big(\cos\beta [\phi_\rho \wedge b_{\alpha+1,k}]\Big)$$

$$\begin{aligned} d_A B &= y^{-2}\Big(\sin\beta\cos\beta\, [\phi^\tau_\rho \wedge \phi_\rho] + \cos\beta\, \phi_\rho \wedge dy\Big) + y^{-1}\Big(\cos\beta\, d_\omega \phi_\rho\Big) \\ &\quad + \sum_{(\alpha,k)\geq 0} y^\alpha (\log y)^k \Big(d_\omega b_{\alpha,k} + \sin\beta\, [\phi^\tau_\rho \wedge b_{\alpha+1,k}] + \cos\beta\, [\phi_\rho \wedge a_{\alpha+1,k}] \\ &\qquad\qquad - (\alpha+1) b_{\alpha+1,k} \wedge dy - (k+1) b_{\alpha+1,k+1} \wedge dy\Big) \end{aligned}$$

$$\begin{aligned} d_A^{\star_4} B &= y^{-1}\Big(\cos\beta\, d_\omega^{\star_4} \phi_\rho\Big) + \sum_{(\alpha,k)\geq 0} y^\alpha (\log y)^k \Big(d_\omega^{\star_4} b_{\alpha,k} + \sin\beta\ \star_4 [\phi^\tau_\rho \wedge \star_4 b_{\alpha+1,k}] \\ &\qquad\qquad - \cos\beta\ \star_4 [\phi_\rho \wedge \star_4 a_{\alpha+1,k}]\Big) \end{aligned}$$

Here we write $(\alpha, k) \geq 0$ for pairs of exponents with $\mathrm{Re}\, \alpha \geq 0$. Using $(\partial_y A)_i = -(F_A)_{iy}$ and

$(\partial_y B)_i = -(d_A B)_{iy}$, we can now read off the contributions to the Haydys-Witten equations from each of these expressions, order by order in $y^\alpha(\log y)^k$.

$\mathcal{O}(y^{-2})$

At order $\mathcal{O}(y^{-2})$ only equations (4.6) and (4.7) are non-trivial. Using that $\iota_w(\star_4[\phi_\rho^\tau \wedge \phi_\rho]) = \frac{1}{2}\iota_w(\star_4[\phi_\rho^\tau \wedge \phi_\rho^\tau] + \star_4[\phi_\rho \wedge \phi_\rho])$, the Haydys-Witten equations then state:

$$\begin{aligned} \phi_\rho^\tau &= -\iota_w\big(\star_4[\phi_\rho^\tau \wedge \phi_\rho^\tau]\big) \\ \phi_\rho &= \iota_w\big(\star_4[\phi_\rho \wedge \phi_\rho]\big) \end{aligned}$$

This is satisfied, since under the isomorphism $\Omega^2_{v,+}(U, \operatorname{ad} E) = \operatorname{Hom}(\Delta^\perp_{(u,v)}, \operatorname{ad} E)$ the expression $\iota_w(\star_4[\phi_\rho \wedge \phi_\rho])$ is identified with $\sigma(\phi_\rho, \phi_\rho)$. By construction of the Nahm pole boundary condition $\phi_\rho - \sigma(\phi_\rho, \phi_\rho) = 0$, while the same is true for ϕ_ρ^τ with the opposite sign.

$\mathcal{O}(y^{-1})$

The first additional constraint arises at $\mathcal{O}(y^{-1})$, where we find

$$\begin{aligned} 0 &= \iota_w\big(\cos\beta\, d_\omega\phi_\rho - \sin^2\beta\ \star_4 d_\omega\phi_\rho^\tau - \cos^2\beta\ \star_4 d_\omega\phi_\rho\big) \\ 0 &= -\iota_w\big(\sin\beta\, d_\omega\phi_\rho^\tau + \sin\beta\cos\beta\ \star_4 d_\omega\phi_\rho^\tau - \sin\beta\cos\beta\ \star_4 d_\omega\phi_\rho\big) \\ 0 &= \cos\beta\, d_\omega^{\star_4}\phi_\rho \end{aligned}$$

These equations are satisfied if the tensor[2] $\nabla^\omega\phi_\rho$ vanishes. This means that ϕ_ρ intertwines the Levi-Civita and gauge connection on TW^4: On the one hand, ω must pull back under $\phi_\rho : \Delta^\perp_{(u,v)} \to \operatorname{ad} E$ to the Levi-Civita connection on $\Delta^\perp_{(u,v)}$. This also implies that the restriction of ω to $\Delta^\perp_{(u,v)} \subset TW^4$ is valued in $\rho(\mathfrak{su}(2)) = V_1 \subset \operatorname{ad} E$. On the other hand, the $w^\flat$-component of ω must be such that $\nabla^\omega_w\phi_\rho = 0$, i.e. ϕ_ρ extends along w by parallel transport. Since the Levi-Civita connection acts trivially on $\operatorname{ad} E$, it preserves $V_1 \subset \operatorname{ad} E$ and we find that also $\iota_w\omega \in V_1$. We note that $\iota_w\star_4 F_\omega$ and $\iota_w F_\omega$ are then elements of $V_1^- \oplus V_1^0 \oplus V_1^+$ and are determined by the Riemannian curvature of W^4. Furthermore, the exterior covariant derivative satisfies the following mapping properties (cf. [Hen12, p.5]):

$$\begin{aligned} \iota_w\star_4 d_\omega &: V_j^w \to 0\,, \quad V_j^- \to V_j^0\,, \quad V_j^0 \to V_j^- \oplus V_j^0 \oplus V_j^+\,, \quad V_j^+ \to V_j^0 \oplus V_j^+ \\ \iota_w d_\omega &: V_j^w \to V_j^- \oplus V_j^0 \oplus V_j^+\,, \quad V_j^\eta \to V_j^\eta,\ \eta \in \{-,0,+\} \end{aligned} \tag{4.10}$$

With a view at the expansions of F_A and $d_A B$, it becomes clear that these identities will be helpful in the upcoming analysis, starting at order $\mathcal{O}(y^0(\log y)^k)$.

$\mathcal{O}(y^0(\log y)^k)$

Examining the higher order expansions of the right hand side of (4.6) - (4.9), one finds that the equations will from now on involve terms of the form $\iota_w\star_4[\phi_\rho\wedge\cdot]$ and its analogue with respect

[2] Here, as always, ∇^ω is the product connection on $\Omega^1(W^4 \times \mathbb{R}^+_y, \operatorname{ad} E)$ that is induced by the Levi-Civita connection and the gauge connection ω.

to ϕ_ρ^τ. These are global versions of the spin-spin operator $\mathfrak{J} := \mathfrak{s} \cdot \mathfrak{t}$, as they act pointwise by $\sum_{a=1}^3 \mathfrak{s}_a \otimes \mathfrak{t}_a$. Indeed, taking into account that rotations act on $a_{\alpha,k}$ and $b_{\alpha,k}$ with opposite orientation one finds

$$\mathfrak{J} a_{\alpha,k} = -\iota_w \star_4 [\phi_\rho^\tau \wedge a_{\alpha,k}] \qquad\qquad \mathfrak{J} b_{\alpha,k} = \iota_w \star_4 [\phi_\rho \wedge b_{\alpha,k}]$$

The vector bundles V_j^η, $\eta \in \{-,0,+\}$, are eigenspaces of the spin-spin operator with eigenvalues $j+1$, 1, and $-j$. Meanwhile, $\mathfrak{J}$ sends V_j^w to zero. We will also use that $\iota_w[\phi_\rho \wedge \cdot]$ maps $V_j^w \to V_j^0$ and $V_j^\eta \to 0$, $\eta \in \{-,0,+\}$.

Equipped with this, we can now proceed to the constraints that arise for the tower of $\log y$ terms at order y^0. Hence, consider equations (4.6) - (4.9) at $\mathcal{O}\big((\log y)^k\big)$, where k is the largest integer such that $(1,k) \in \Delta_0$. The existence of such a maximal value is part of the defining properties of a polyhomogeneous expansion. For the moment assume that $k \geq 2$, we will treat the cases $k=1$ and $k=0$ in the end.

We now introduce the following linear combinations of $V_j^- \oplus V_j^0 \oplus V_j^+$-parts of $a_{\alpha,k}$ and $b_{\alpha,k}$:

$$\begin{aligned} c_{\alpha,k} &:= \sin\beta\, a_{\alpha,k}^\tau + \cos\beta\, b_{\alpha,k}\,, \\ d_{\alpha,k} &:= \cos\beta\, a_{\alpha,k}^\tau - \sin\beta\, b_{\alpha,k}\,. \end{aligned}$$

Note, in particular, that $c_{\alpha,k}$ and $d_{\alpha,k}$ do not include the component $(a_{1,k})^{j,w}$. Appropriate linear combinations of (4.8) and (4.9) then state

$$\begin{aligned} \sin\beta\, \iota_w a_{1,k} &= \cos\beta\ \star_4 [\phi_\rho \wedge \star_4 c_{1,k}] + \sin\beta\ \star_4 [\phi_\rho \wedge \star_4 d_{1,k}] \\ \cos\beta\, \iota_w a_{1,k} &= \sin\beta\ \star_4 [\phi_\rho^\tau \wedge \star_4 c_{1,k}] - \cos\beta\ \star_4 [\phi_\rho^\tau \wedge \star_4 d_{1,k}] \end{aligned}$$

These equations can be brought in a slightly more helpful form. Note that the operator $[\phi_\rho \wedge \star_4[\phi_\rho \wedge \star_4 \cdot\,]]$ annihilates $V_j^\pm$ and acts on V_j^0 as $\sum_i [\mathfrak{t}_i, [\mathfrak{t}_i, \cdot]]$, which is just (the negative of) the quadratic Casimir on V_j. Applying $[\phi_\rho \wedge \cdot]$ to the equations thus simplifies the operation on the right hand side to a simple multiplication by $-j(j+1)$. In complete analogy $[\phi_\rho^\tau \wedge \star_4[\phi_\rho^\tau \wedge \star_4 \cdot\,]]$ annihilates $\tilde{V}_j^\pm$ and acts on $\tilde{V}_j^0$ by multiplication with $-j(j+1)$. This results in (4.13) and (4.14) below. Including appropriate linear combinations of (4.6) and (4.7), the coefficient functions in the polyhomogeneous expansion have to be compatible with the following four conditions.

$$c_{1,k} + \mathfrak{J} c_{1,k} = \sin\beta\cos\beta\, \Big(\iota_w[\phi_\rho^\tau \wedge a_{1,k}] - \iota_w[\phi_\rho \wedge a_{1,k}]\Big) \tag{4.11}$$

$$d_{1,k} - \mathfrak{J} d_{1,k} = -\cos^2\beta\, \iota_w[\phi_\rho^\tau \wedge a_{1,k}] - \sin^2\beta\, \iota_w[\phi_\rho \wedge a_{1,k}] \tag{4.12}$$

$$j(j+1)\big(\cos\beta\, c_{1,k} + \sin\beta\, d_{1,k}\big)^{j,0} = \sin\beta\, \iota_w[\phi_\rho \wedge a_{1,k}] \tag{4.13}$$

$$j(j+1)\big(\sin\beta\, c_{1,k} - \cos\beta\, d_{1,k}\big)^{\widetilde{j,0}} = \cos\beta\, \iota_w[\phi_\rho^\tau \wedge a_{1,k}] \tag{4.14}$$

Observe that the terms on the left only depend on $c_{1,k}$ and $d_{1,k}$ and thus only on the $V_j^- \oplus V_j^0 \oplus V_j^+$ components of $a_{1,k}$ and $b_{1,k}$. Meanwhile, the right hand side only depends on the remaining V_j^w component $(a_{1,k})^{j,w}$. Moreover, all terms on the right are necessarily located in $V_j^0 + \tilde{V}_j^0$ (fiberwise sum of subspaces). To solve the equations we distinguish between the following six cases:

(I) $V_j^-/\tilde{V}_j^0$ (III) $V_j^- \cap \tilde{V}_j^0$ (V) $V_j^0/\tilde{V}_j^0$

(II) $V_j^+/\tilde{V}_j^0$ (IV) $V_j^+ \cap \tilde{V}_j^0$ (VI) $V_j^0 \cap \tilde{V}_j^0$

In fact, a combination of the metric and Killing form provides an inner product on each V_j^η, such that we can (and do) identify the quotient spaces with the orthogonal complement of $\tilde{V}_j^0$ inside of V_j^η (but reflecting this in the notation only adds unnecessary complexity). In the upcoming paragraphs we will write $\alpha|_{(\mathrm{I})}$ to denote restriction to the subspace specified in case (I), et cetera. Let us note at this point that $j = 1$ is slightly special, since $V_1^- \cap \tilde{V}_1^0 = \{0\}$.

When we restrict (4.11) - (4.14) to either of the subspaces in (I) or (II), all terms on the right hand side vanish. In particular, (4.11) and (4.12) reduce to simple eigenvalue equations. Since $\mathfrak{J}$ acts on V_j^- with eigenvalue $j+1$ (which is strictly larger than 1), we immediately find $c_{1,k}|_{(\mathrm{I})} = 0 = d_{1,k}|_{(\mathrm{I})}$ for any j. Similarly, since $\mathfrak{J}$ acts on V_j^+ with eigenvalue $-j$, we get $c_{1,k}|_{(\mathrm{II})} = 0$, except perhaps when $j = 1$, while $d_{1,k}|_{(\mathrm{II})} = 0$ in any case. For the $j = 1$ component of $c_{1,k}|_{(\mathrm{II})}$ we can rely on the analogue of (4.11) at the subleading order $\mathcal{O}\big((\log y)^{k-1}\big)$, which states

$$kc_{1,k} + c_{1,k-1} + \mathfrak{J}c_{1,k-1} = \sin\beta\cos\beta\ \big(\iota_w[\phi_\rho^\tau \wedge a_{1,k-1}] - \iota_w[\phi_\rho \wedge a_{1,k-1}]\big)$$

Upon restriction to $V_1^+/\tilde{V}_1^0$ this reduces to $kc_{1,k}|_{(\mathrm{II})} = 0$.

Restriction to the subspaces in (III) and (IV) works out slightly differently, because there the restriction of $\iota_w[\phi_\rho^\tau \wedge a_{1,k}]$ doesn't vanish (though we still have vanishing of $\iota_w[\phi_\rho \wedge a_{1,k}]$). Concentrate on case (III) first, where J has eigenvalue $j+1$. We find that (4.11) and (4.12) determine $c_{1,k}|_{(\mathrm{III})} = \frac{\sin\beta\cos\beta}{j+2}(\iota_w[\phi_\rho^\tau \wedge a])^{j,-}$ and $d_{1,k}|_{(\mathrm{III})} = \frac{\cos^2\beta}{j}(\iota_w[\phi_\rho^\tau \wedge a])^{j,-}$. However, when plugging this into (4.14), we get the following consistency condition:

$$j(j+1)\left(\frac{\sin^2\beta\cos\beta}{j+2} - \frac{\cos^3\beta}{j}\right)(\iota_w[\phi_\rho^\tau \wedge a_{1,k}])^{j,-} = -\cos\beta(\iota_w[\phi_\rho^\tau \wedge a_{1,k}])^{j,-}$$

This can only be satisfied if $\cos 2\beta = -\frac{2j+3}{(j+1)^2}$ or if $(\iota_w[\phi_\rho^\tau \wedge a_{1,k}])^{j,-} = 0$. Interestingly, for any $j \geq 2$ there exists a single value of $\beta \in (0, \pi/2)$ for which the first equation is satisfied. Therefore, whenever β is related to one of the spins j in this way, both $(c_{1,k})^{j,-}$ and $(d_{1,k})^{j,-}$ are given as stated above by a multiple of $(\iota_w[\phi_\rho^\tau \wedge a_{1,k}])^{j,-}$. For generic values of β, however, we find that $(\iota_w[\phi_\rho^\tau \wedge a_{1,k}])^{j,-}$ vanishes, and thus also $c_{1,k}|_{(\mathrm{III})} = d_{1,k}|_{(\mathrm{III})} = 0$.

For case (IV), recall that J acts with eigenvalue $-j$ on V_j^+. If $j = 1$, equation (4.11) directly yields $\iota_w[\phi_\rho^\tau \wedge a_{1,k}]|_{(\mathrm{IV})} = 0$, such that $d_{1,k}|_{(\mathrm{IV})} = 0$ by (4.12) and $c_{1,k}|_{(\mathrm{IV})} = 0$ by (4.14). Otherwise, for $j \geq 2$, equations (4.11) and (4.12) specify $c_{1,k}|_{(\mathrm{IV})} = \frac{\sin\beta\cos\beta}{1-j}\iota_w[\phi_\rho^\tau \wedge a_{1,k}]|_{(\mathrm{IV})}$ and $d_{1,k}|_{(\mathrm{IV})} = -\frac{\cos^2\beta}{1+j}\iota_w[\phi_\rho^\tau \wedge a_{1,k}]|_{(\mathrm{IV})}$. Plugging these expressions into (4.14) now leads to the condition $\cos 2\beta = \frac{2j-1}{j^2}$ if $j \geq 2$. In this condition, we can also allow $j = 1$ as input, since the corresponding solution is $\beta = 0$, which we exclude anyway. If β coincides with one of these special values, $(c_{1,k})^{j,+}$ and $(d_{1,k})^{j,+}$ are determined as above in terms of $(\iota_w[\phi_\rho^\tau \wedge a_{1,k}])^{j,+}$. For generic values of β, we have $(\iota_w[\phi_\rho^\tau \wedge a_{1,k}])^{j,+} = c_{1,k}|_{(\mathrm{IV})} = d_{1,k}|_{(\mathrm{IV})} = 0$

Case (V) is comparatively simple. Observe that $\iota_w[\phi_\rho^\tau \wedge a_{1,k}]|_{(\mathrm{V})} = 0$ and that J acts with eigenvalue 1. This means that (4.12) immediately provides $\iota_w[\phi_\rho \wedge a_{1,k}]|_{(\mathrm{V})} = 0$, such that (4.11) states $c_{1,k}|_{(\mathrm{V})} = 0$ and as a consequence (4.13) yields $d_{1,k}|_{(\mathrm{V})} = 0$.

Finally, we arrive at case (VI), where neither $\iota_w[\phi_\rho^\tau \wedge a_{1,k}]|_{(\mathrm{VI})}$ nor $\iota_w[\phi_\rho \wedge a_{1,k}]|_{(\mathrm{VI})}$ vanish. Let us first consider generic values of β, where we already know that $(\iota_w[\phi_\rho^\tau \wedge a_{1,k}])^{j,\pm} = 0$. Using $(\iota_w[\phi_\rho \wedge a_{1,k}])^\tau = \iota_w[\phi_\rho^\tau \wedge a_{1,k}]$ and that acting with τ twice is the identity map, we can then deduce that $\iota_w[\phi_\rho \wedge a_{1,k}]$ decomposes into a sum of τ-eigenvectors inside of $V_j^0 \cap \tilde{V}_j^0$ with eigenvalues ± 1. Plugging this decomposition into (4.11) - (4.14) then shows that there can only be non-trivial solutions if $j = 1$ and $\beta = \pi/2$. Since we excluded $\beta = \pi/2$ already in the definition of the Nahm pole boundary condition, we find that $(a_{1,k})^{j,w} = (c_{1,k})^{j,0} = (d_{1,k})^{j,0} = 0$ for all j.

Still in case (VI), but when β is one of the special values $\cos 2\beta = -\frac{2j+3}{(j+1)^2}$ or $\frac{2j-1}{j^2}$, it is no longer assured that $(\iota_w[\phi_\rho^\tau \wedge a_{1,k}])^{j,\pm}$ vanishes. As a result we can't decompose it into $\pm$-eigenparts of τ within $V_j^0 \cap \tilde{V}_j^0$. In this situation equations (4.11) - (4.14) instead provide $(c_{1,k})^{j,0} = -\tan\beta\, \iota_w[\phi_\rho \wedge a_{1,k}]$ and $(d_{1,k})^{j,0} = \iota_w[\phi_\rho \wedge a_{1,k}]$.

By way of an intermediate conclusion, it can thus be stated that for generic β and $k \geq 2$, the functions $a_{1,k}$ and $b_{1,k}$ vanish. Meanwhile, if β is one of the (finitely many) special angles determined by $\cos 2\beta = -\frac{2j+3}{(j+1)^2}$ or $\frac{2j-1}{j^2}$, their spin j components are not necessarily zero and are determined – via $a^\tau_{1,k} = \sin\beta\, c_{1,k} + \cos\beta\, d_{1,k}$ and $b_{1,k} = \cos\beta\, c_{1,k} - \sin\beta d_{1,k}$ – by the following expressions

$$
\begin{aligned}
c_{1,k} &= \frac{\sin\beta\cos\beta}{j+2}\left(\iota_w[\phi_\rho^\tau \wedge a_{1,k}]\right)^{j,-} - \tan\beta\, \left(\iota_w[\phi_\rho \wedge a_{1,k}]\right)^{j,0} - \frac{\sin\beta\cos\beta}{j-1}\left(\iota_w[\phi_\rho^\tau \wedge a_{1,k}]\right)^{j,+} \\
d_{1,k} &= \frac{\cos^2\beta}{j}\left(\iota_w[\phi_\rho^\tau \wedge a_{1,k}]\right)^{j,-} + \left(\iota_w[\phi_\rho \wedge a_{1,k}]\right)^{j,0} - \frac{\cos^2\beta}{j+1}\left(\iota_w[\phi_\rho^\tau \wedge a_{1,k}]\right)^{j,+}
\end{aligned}
\tag{4.15}
$$

The lower order terms $a_{1,k'}$ and $b_{1,k'}$, $0 \leq k' < k$ are then determined by induction. Let us stress that the possible values of $j \geq 2$ are determined by the (finitely many) spins $j \in J$ that appear in the decomposition of $\mathrm{ad}\, E$ under the action of $\mathfrak{su}(2)_\mathfrak{s}$.

As an aside, note that all components are determined by the free parameters $(a_{1,k})^{j,w}$, $(a_{1,k-1})^{j,w}$, and so on, which describe the y-dependence of $\iota_w A$, i.e. the $w^\flat$-component of the gauge connection. While our current analysis cannot determine the maximal power of $(\log y)^k$, this is in general controlled by the geometry and topology of $E \to M^5$. For example, if the boundary is of the form $W^4 = \mathbb{R} \times X^3$ and w is the vector field parallel to the real line, then upon dimensional reduction $\iota_w A$ would be identified with the dy-component of a one-form ϕ that satisfies the Kapustin-Witten equations. In that situation a well-known vanishing theorem states that $\iota_w A$ vanishes if $A + i\phi$ approaches an irreducible flat connection at $y \to \infty$. In that way, the regularity of twisted Nahm pole solutions is further controlled by global data, away from the boundary.

The arguments above break down at $k = 1$, where in case (II) it is no longer possible to deduce that the $V_1^+/\tilde{V}_1^0$ part of $c_{1,1}$ vanishes. This is because for this component we needed to rely on the subleading order of (4.6) and (4.7). For $c_{1,1}|_{(\mathrm{II})}$ this means that we have to look at the equations at order $\mathcal{O}(1)$, which receive an additional contribution from F_ω:

$$\begin{aligned} c_{1,1} + c_{1,0} + \mathfrak{J}c_{1,0} &= \cos\beta\, \imath_w F_\omega + \imath_w \star_4 F_\omega + \sin\beta\cos\beta\big(\imath_w[\phi_\rho^\tau \wedge a_{1,0}] - \imath_w[\phi_\rho \wedge a_{1,0}]\big) \\ d_{1,0} - \mathfrak{J}d_{1,0} &= -\sin\beta\, \imath_w F_\omega^\tau - \cos^2\beta\, \imath_w[\phi_\rho^\tau \wedge a_{1,0}] - \sin^2\beta\, \imath_w[\phi_\rho \wedge a_{1,0}] \end{aligned}$$

Since the terms that contain F_ω are located in $V_1^- \oplus V_1^0 \oplus V_1^+$, only these parts of $c_{1,1}$, $c_{1,0}$ and $d_{1,0}$ are affected by curvature contributions. Note that the special values of β are related to spins $j \geq 2$, so there are no additional contributions to the $V_1^- \oplus V_1^0 \oplus V_1^+$ part of the equations. By restriction of the $\mathcal{O}(1)$ equations to V_1^-, V_1^0, and V_1^+ (remembering that $V_1^- \cap \tilde{V}_1^0 = \{0\}$) we thus obtain

$$c_{1,1} = (\cos\beta\, \imath_w F_\omega + \imath_w \star_4 F_\omega)|_{V_1^+/\tilde{V}_1^0}$$

$$\begin{aligned} c_{1,0} = &\ \tfrac{1}{3}\left(\cos\beta\, \imath_w F_\omega + \imath_w \star_4 F_\omega\right)^{1,-} \\ &+ \tfrac{1}{2}\left(\cos\beta\, \imath_w F_\omega + \imath_w \star_4 F_\omega + \sin\beta\cos\beta\big(\imath_w[\phi_\rho^\tau \wedge a_{1,0}] - \imath_w[\phi_\rho \wedge a_{1,0}]\big)\right)^{1,0} \\ &- \tfrac{1}{2}\left(\cos\beta\, \imath_w F_\omega^\tau + \sin\beta\cos\beta\, \imath_w[\phi_\rho^\tau \wedge a_{1,0}]\right)|_{V_1^+ \cap \tilde{V}_1^0} + C^{1,+} \end{aligned}$$

$$d_{1,0} = (\sin\beta\, \imath_w F_\omega^\tau)^{1,-} + \tfrac{1}{2}\left(\imath_w[\phi_\rho \wedge a_{1,0}] - \cot\beta\, c_{1,0}\right)^{1,0} - \tfrac{1}{2}\left(\sin\beta\, \imath_w F_\omega^\tau + \cos^2\beta\, \imath_w[\phi_\rho^\tau \wedge a_{1,0}]\right)^{1,+}$$

Here, $C^{1,+} \in V_1^+/\tilde{V}_1^0$ is some undetermined 'integration constant' that cannot be fixed by the current analysis. In contrast, $\imath_w[\phi_\rho \wedge a_{1,0}]$ and $\imath_w[\phi_\rho^\tau \wedge a_{1,0}]$ are further restricted by the $\mathcal{O}(1)$ versions of (4.13) and (4.14) (which remain unchanged) and consequently depend explicitly on F_ω. Indeed, if the F_ω contributions to the $\mathcal{O}(1)$ equations vanish, our earlier arguments show that, apart from $c_{1,0} = C^{1,+}$, all coefficient functions vanish.

Finally, if β is one of the special angles where $\cos 2\beta = -\frac{2j+3}{(j+1)^2}$ or $\frac{2j-1}{j^2}$, then the results for $c_{1,1}$, $c_{1,0}$ and $d_{1,0}$ differ only by additionally including the terms induced by (4.15), which only appear for spins $j \neq 1$.

The main take away of this discussion is that, for generic β and in absence of curvature contributions, the functions $a_{1,k}$ and $b_{1,k}$ vanish for all $k \geq 1$, while they are proportional to $C^{1,+} \in V_1^+/\tilde{V}_1^0$ for $k = 0$.

$\underline{\mathcal{O}(y^1(\log y)^k)}$

For the constraints at order $\mathcal{O}(y(\log y)^k)$ the analysis is very similar. The main difference is that the equations now incorporate terms of the form $d_\omega c_{1,k}$. As before, let k be the largest integer for which $(2,k) \in \Delta_0$. The Haydys-Witten equations (4.6) - (4.9) now become

$$\begin{aligned} \left(2c_{2,k} + \mathfrak{J}c_{2,k}\right) - \imath_w d_\omega d_{1,k}^\tau - \imath_w \star_4 d_\omega a_{1,k} &= \sin\beta\cos\beta\, \left(\imath_w[\phi_\rho^\tau \wedge a_{2,k}] - \imath_w[\phi_\rho \wedge a_{2,k}]\right) \\ \left(2d_{2,k} - \mathfrak{J}d_{2,k}\right) + \imath_w d_\omega c_{1,k}^\tau - \left(\imath_w \star_4 d_\omega b_{1,k}\right)^\tau &= -\cos^2\beta\, \imath_w[\phi_\rho^\tau \wedge a_{2,k}] - \sin^2\beta\, \imath_w[\phi_\rho \wedge a_{2,k}] \\ j(j+1)\left(\cos\beta\, c_{2,k} + \sin\beta\, d_{2,k}\right)^{j,0} &= \sin\beta\, \imath_w[\phi_\rho \wedge a_{2,k}] \\ j(j+1)\left(\sin\beta\, c_{2,k} - \cos\beta\, d_{2,k}\right)^{\widetilde{j,0}} &= \cos\beta\, \imath_w[\phi_\rho^\tau \wedge a_{2,k}] \end{aligned}$$

We have again arranged this such that the right hand side is contained in $V_j^0 + \tilde{V}_j^0$. Moreover, we know that for generic values of β and any $k \geq 2$, the functions $c_{1,k}$ and $d_{1,k}$ vanish completely, such that the derivatives on the left hand side disappear. In that situation essentially the same arguments as before show that $(a_{2,k})^{j,w} = c_{2,k} = d_{2,k} = 0$, as long as we additionally exclude the special angles determined by $\cos 2\beta = -\frac{3j^2-4j+3}{j(j+1)^2}$ or $-\frac{3j^2+2j-4}{j^2(j+1)}$.

In carrying out these arguments, we need to rely on the subleading equations at $\mathcal{O}(y(\log y)^{k-1})$ again. Here, the subleading equations provide conditions for the $V_2^+/\tilde{V}_2^0$ part of $c_{2,k}$ and the $V_1^-/\tilde{V}_1^0$ $(= V_1^-)$ part of $d_{2,k}$, which are projected out from the equations at order $\mathcal{O}(y(\log y)^k)$. Specifically, the equation for the c-terms at order $\mathcal{O}(y(\log y)^{k-1})$ is

$$\begin{aligned}\left(kc_{2,k} + 2c_{2,k-1} + \mathfrak{J}c_{2,k-1}\right) - \iota_w d_\omega d^\tau_{1,k-1} - \iota_w \star_4 d_\omega a_{1,k-1}& \\ = \sin\beta\cos\beta\left(\iota_w[\phi_\rho^\tau \wedge a_{2,k-1}] - \iota_w[\phi_\rho \wedge a_{2,k-1}]\right)&\end{aligned}$$

The $V_2^+/\tilde{V}_2^0$ part of this equation simply states $kc_{2,k}|_{(\mathrm{II})} = 0$ for all $k \geq 1$ (for $k = 1, 2$ notice that $c_{1,k}$ and $a_{1,k}$ are elements of $V_1^- \oplus V_1^0 \oplus V_1^+$). Similarly, the equation for the d terms at order $\mathcal{O}(y(\log y)^{k-1})$ states

$$\begin{aligned}\left(kd_{2,k} + 2d_{2,k-1} - \mathfrak{J}d_{2,k-1}\right) + \iota_w d_\omega c^\tau_{1,k-1} - (\iota_w \star_4 d_\omega b_{1,k-1})^\tau& \\ = -\cos^2\beta\, \iota_w[\phi_\rho \wedge a_{2,k-1}] - \sin^2\beta\, \iota_w[\phi_\rho^\tau \wedge a_{2,k-1}]&\end{aligned}$$

For $k > 2$ the V_1^- part of this equation simply states $(kd_{2,k})^{1,-} = 0$. For $k = 2$ we instead have $(2d^\tau_{2,2} + \iota_w d_\omega c_{1,1} - \iota_w \star_4 d_\omega b_{1,1})^{1,-} = 0$. Note that $b_{1,1} = \cos\beta\, c_{1,1}$ is an element of $V_1^+/\tilde{V}_1^0$. By comparison with the mapping properties (4.10), we conclude that also $(d_{2,2})^{1,-} = 0$.

The equations at order $\mathcal{O}(y(\log y))$ and $\mathcal{O}(y)$ involve derivatives of $c_{1,1}$, $c_{1,0}$ and $d_{1,0}$, which are all contained in $V_1^- \oplus V_1^0 \oplus V_1^+$, so only these parts of $c_{2,1}$, $d_{2,1}$, $c_{2,0}$ and $d_{2,0}$ may be non-zero. By restricting the equations to the various possible subspaces (and since $V_1^- \cap \tilde{V}_1^0 = \{0\}$) one obtains

$$\begin{aligned}c_{2,1} =\ & \tfrac{1}{3}\left(\sin\beta\, \iota_w \star_4 d_\omega c_{1,1} + \sin\beta\cos\beta\left(\iota_w[\phi_\rho^\tau \wedge a_{2,1}] - \iota_w[\phi_\rho \wedge a_{2,1}]\right)\right)^{1,0} \\ +\ & \left(\sin\beta\, \iota_w \star_4 d_\omega c_{1,1} + \sin\beta\cos\beta\, \iota_w[\phi_\rho^\tau \wedge a_{2,1}]\right)^{1,+} \\ d_{2,1} =\ & \left(-\iota_w d_\omega c^\tau_{1,0} + (\iota_w \star_4 d_\omega b_{1,0})^\tau\right)^{1,-} \\ & + \left(-\iota_w d_\omega c^\tau_{1,1} + (\iota_w \star_4 d_\omega b_{1,1})^\tau - \cos^2\beta\, \iota_w[\phi_\rho^\tau \wedge a_{2,1}] - \sin^2\beta\, \iota_w[\phi_\rho \wedge a_{2,1}]\right)^{1,0} \\ & + \tfrac{1}{3}\left(-\iota_w d_\omega c^\tau_{1,1} + (\iota_w \star_4 d_\omega b_{1,1})^\tau - \sin^2\beta\, \iota_w[\phi_\rho^\tau \wedge a_{2,1}]\right)^{1,+}\end{aligned}$$

Since the $\mathcal{O}(y\log y)$ terms are non-zero, they appear in the $\mathcal{O}(y)$ equations and, accordingly, in the following result for $c_{2,0}$ an $d_{2,0}$:

$$\begin{aligned}
c_{2,0} = \ & \tfrac{1}{4}\left(\iota_w d_\omega d^\tau_{1,0} + \iota_w \star_4 d_\omega a_{1,0}\right)^{1,-} \\
& + \tfrac{1}{3}\left(\iota_w d_\omega d^\tau_{1,0} + \iota_w \star_4 d_\omega a_{1,0} + \sin\beta\cos\beta\ \left(\iota_w[\phi^\tau_\rho \wedge a_{1,0}] - \iota_w[\phi_\rho \wedge a_{1,0}]\right) - c_{2,1}\right)^{1,0} \\
& + \ \left(\iota_w d_\omega d^\tau_{1,0} + \iota_w \star_4 d_\omega a_{1,0} + \sin\beta\cos\beta\, \iota_w[\phi^\tau_\rho \wedge a_{1,0}] - c_{2,1}\right)^{1,+} + C^{2,+} \\
d_{2,0} = \ & D^{1,-} + \left(-\iota_w d_\omega c^\tau_{1,0} + \left(\iota_w \star_4 d_\omega b_{1,0}\right)^\tau - \cos^2\beta\, \iota_w[\phi^\tau_\rho \wedge a_{2,0}] - \sin^2\beta\, \iota_w[\phi_\rho \wedge a_{2,0}] - c_{2,1}\right)^{1,0} \\
& + \tfrac{1}{3}\left(-\iota_w d_\omega c^\tau_{1,0} + \left(\iota_w \star_4 d_\omega b_{1,0}\right)^\tau - \cos^2\beta\, \iota_w[\phi^\tau_\rho \wedge a_{2,0}] - c_{2,1}\right)^{1,+}
\end{aligned}$$

Here, $C^{2,+} \in V_2^+/\tilde{V}_2^0$ and $D^{1,-} \in V_1^-$ remain undetermined. While we do not reproduce explicit formulas, it is easy to see that $\iota_w[\phi_\rho \wedge a_{2,k}]$ and $\iota_w[\phi^\tau_\rho \wedge a_{2,k}]$ are fully determined by the bottom two Haydys-Witten equations.

The main conclusion of this part is that for generic β and in absence of curvature contributions to the $\mathcal{O}(1)$ equations (such that $a_{1,k} = b_{1,k} = 0$ for $k \geq 1$) all components of $c_{2,k}$ and $d_{2,k}$ vanish, except for the functions $C^{2,+}$ and $D^{1,-}$ that appear at $k = 0$. □

4.4 Asymptotics of the Boundary Term

We now determine the asymptotic behaviour of $\chi = \chi_1 + \chi_2$ at cylindrical ends of M^5, where we recall from Lemma 4.2 that

$$\begin{aligned}
\chi_1 &= -2\,\mathrm{Tr}\left(F_A \wedge J^- B\right)\ , \\
\chi_2 &= -2\,\mathrm{Tr}\left(\delta^+_A J^+ B \wedge J^- B \wedge \eta\right)\ .
\end{aligned}$$

In doing so, we assume that we evaluate χ for a Haydys-Witten solution, which is the situation we are interested in for the Weitzenböck formula.

Since we eventually need to pullback χ to the boundary, it is important to understand the interplay between the vector subbundles $\ker\eta$ and $T(\partial_i M)$, with a particular focus on the properties of J. Thus, let $U = [0,1)_{s_i} \times \partial_i M^5$ be a cylindrical end and let u denote the inward-pointing unit normal vector. This means that near a compact end $u = \partial_{s_i}$, while at a non-compact end $u = s_i \partial_{s_i}$. Assume that $g(u,v) = \cos\theta$ is constant on all of U.

If $\theta = 0$, u and v are parallel and $\ker\eta = T\partial_i M$. In this situation J simply corresponds to an almost Hermitian structure on $\partial_i M$. If $\theta \neq 0$, u and v are linearly-independent, non-vanishing vector fields on the tubular neighbourhood and span a regular distribution $\Delta_{(u,v)} \subset TU$ of rank two. Since $\ker\eta$ has rank four, the two distributions intersect in a line bundle $L \subset TU$. At each point, L specifies a direction in the (u,v)-plane that is perpendicular to v. We denote the generating unit vector field by $v^\perp$ and fix its orientation such that $g(u, v^\perp) = -\sin\theta$, such that it points into the interior. The vector fields are related by $v = \sin\theta w + \cos\theta u$ and $v^\perp = \cos\theta w - \sin\theta u$.

Starting from $v^{\perp}$, we can always find a local basis of $\ker \eta$ that interacts nicely with the almost complex structure J. It is given by $\{v^{\perp}, w_1 = Jv^{\perp}, w_2, w_3 = Jw_2\}$, where w_2 is some local section of $T\partial_i M$ that is orthogonal to v, $v^{\perp}$, and $Jv^{\perp}$. This induces an associated basis of $\Omega^2_{v,+}(U)$, given by

$$e_i = \eta^{\perp} \wedge w_i^{\flat} + \frac{1}{2}\epsilon_{ijk} w_j^{\flat} \wedge w_k^{\flat}$$

and that satisfies $Je_1 = +e_1$ and $Je_{2/3} = -e_{2/3}$.

We start our investigation of the asymptotics of χ with Nahm pole boundaries $\partial_{\mathrm{NP}}M$ in Section 4.4.1. Though we only establish Theorem D in the context of pure Nahm pole boundary conditions, we also include a short discussion of the expected behaviour of χ near knot boundaries $\partial_K M$ in Section 4.4.2. Subsequently, we discuss Kapustin-Witten ends with either finite energy or Nahm poles in Section 4.4.3.

4.4.1 Nahm Pole Boundaries

Consider a Nahm pole boundary component $\partial_{\mathrm{NP}}M^5$. Let $W^4 \times [0,1)_y$ denote a tubular neighbourhood, where W^4 is a complete Riemannian manifold without boundary, and write μ_{W^4} for the induced volume form. Write $u = \partial_y$ for the inward-pointing unit normal vector field and assume $g(u, v) = \cos\beta$ is constant, with $\beta \in [0, \pi/2)$. Furthermore, denote by $i_y : W^4 \hookrightarrow W^4 \times [0,1)_y$ inclusion of W^4 at y.

Recall from Section 4.1 that there is a local basis $\{v^{\perp}, w_i\}_{i=1,2,3}$ of $\ker\eta$, where each w_i is parallel to TW^4, and associated to it a basis $\{e_i\}_{i=1,2,3}$ of $\Omega^2_{v,+}(W^4 \times [0,1)_y)$ for which $J^- B = B_2 e_2 + B_3 e_3$. Let $(s, x^i, y)_{i=1,2,3}$ be coordinates that make these into coordinate vector fields. With that choice we have $v = \sin\beta\partial_s + \cos\beta\partial_y$, $v^{\perp} = \cos\beta\partial_s - \sin\beta\partial_y$, $w = \partial_s$, and $w_i = dx^i$.

Proposition 4.7. *Assume (A, B) is a β-twisted regular Nahm pole solution of the Haydys-Witten equations. If β is generic, $F_\omega = 0$ and $\nabla^{\omega}_s (C^{1,+})_2 + \nabla^{\omega}_1 (C^{1,+})_3 = \nabla^{\omega}_s (C^{1,+})_3 - \nabla^{\omega}_1 (C^{1,+})_2 = 0$, then there is a $\delta > 0$ and a constant C such that*

$$i_y^* \chi \sim C y^{\delta} \mu_{W^4} \quad (y \to 0)$$

Proof. Consider first the term $\chi_1 = -2\,\mathrm{Tr}(F_A \wedge J^- B)$. Postponing the discussion of the ds-components for the moment, let's concentrate on the contributions from $(F_A)_{ij} dx^i \wedge dx^j$. Observe that e_i pulls back to $\cos\beta ds \wedge dx^i + \frac{1}{2}\epsilon_{ijk} dx^j \wedge dx^k$. Fixing an orientation in which $\mu_{W^4} = \sqrt{g} ds \wedge dx^1 \wedge dx^2 \wedge dx^3$, we find that these contributions are given by $\big((F_A)_{12} B_3 - (F_A)_{13} B_2\big)\cos\beta\mu_{W^4}$.

Plugging the expansions of Lemma 4.5 and Corollary 4.6 into this formula, we encounter the following terms at orders y^{-3} and y^{-1}

$$y^{-3} i_y^*\ \mathrm{Tr}([\phi_\rho^\tau \wedge \phi_\rho^\tau] \wedge J^- \phi_\rho) \propto \mathrm{Tr}([\mathfrak{t}_1^\tau, \mathfrak{t}_2^\tau]\mathfrak{t}_3) - \mathrm{Tr}([\mathfrak{t}_1^\tau, \mathfrak{t}_3^\tau]\mathfrak{t}_2) = -2\,\mathrm{Tr}(\mathfrak{t}_2\mathfrak{t}_3) = 0$$

$$\begin{aligned} y^{-1} i_y^*\ &\mathrm{Tr}([\phi_\rho^\tau \wedge \phi_\rho^\tau] \wedge J^- C^{1,+}) + y^{-1} i_y^*\ \mathrm{Tr}([\phi_\rho^\tau \wedge (C^{1,+})^\tau] \wedge J^- \phi_\rho) \\ &\propto \mathrm{Tr}\left(-\mathfrak{t}_2 (C^{1,+})_3 - \mathfrak{t}_3 (C^{1,+})_2\right) + \mathrm{Tr}\left(\mathfrak{t}_2 (C^{1,+})_3 + \mathfrak{t}_3 (C^{1,+})_2]\right) = 0 \end{aligned}$$

Meanwhile, at order $\mathcal{O}(1)$, we have the following contributions

$$i_y^* \operatorname{Tr}([\phi_\rho^\tau \wedge \phi_\rho^\tau] \wedge J^- C^{2,+}) = 0 \qquad i_y^* \operatorname{Tr}([\phi_\rho^\tau \wedge \phi_\rho^\tau] \wedge J^- D^{1,-}) = 0$$
$$i_y^* \operatorname{Tr}([\phi_\rho^\tau \wedge (C^{2,+})^\tau] \wedge J^- \phi_\rho) = 0 \qquad i_y^* \operatorname{Tr}([\phi_\rho^\tau \wedge (D^{1,-})^\tau] \wedge J^- \phi_\rho) = 0$$

On the left we used that the subspaces $V_1, V_2 \subset \mathfrak{g}$ are orthogonal with respect to the trace (keeping in mind that the action of $\phi_\rho \in V_1$ preserves V_2), while the equations on the right similarly use $D^{1,-} \in V_1^-$ insofar that $D^{1,-} \propto \phi_\rho$, such that the trace vanishes by the same calculation as at order $\mathcal{O}(y^{-3})$.

According to Corollary 4.6, the only ds-component of F_A can arise from $d_\omega C^{1,+}$, which appears at $\mathcal{O}(1)$. As we have seen, all other contributions vanish, so we are left with

$$\begin{aligned}
i_y^*\chi_1 &= -2 i_y^* \operatorname{Tr}(\sin\beta\cos\beta\, d_\omega (C^{1,+})^\tau \wedge J^-\phi_\rho) \\
&= -2\sin\beta\cos^2\beta\,\mu_{W^4}\ \operatorname{Tr}\left((\nabla_s^\omega (C^{1,+})_2 + \nabla_1^\omega (C^{1,+})_3)\mathfrak{t}_3 + (\nabla_s^\omega (C^{1,+})_3 - \nabla_1^\omega (C^{1,+})_2)\mathfrak{t}_2\right) \\
&\quad - 2\sin\beta\cos^2\beta\,\mu_{W^4}\ \operatorname{Tr}\left(\nabla_3^\omega (C^{1,+})_1\mathfrak{t}_2 - \nabla_2^\omega (C^{1,+})_1\mathfrak{t}_3\right) + \mathcal{O}(y^\delta)
\end{aligned}$$

The first line of the result vanishes by assumption, while the remaining terms cancel in combination with χ_2.

Hence, moving on to $\chi_2 = -2\operatorname{Tr}(\delta_A^+ J^+ B \wedge J^- B \wedge \eta)$. Note that $i_y^*\eta = \sin\beta ds$ and as a result only the dx^2- and dx^3-components of $\delta_A^+ J^+ B$ can contribute to this expression. Specifically, we get $i_y^*\chi_2 = 2\operatorname{Tr}\left((\delta_A^+ J^+ B)_2 B_2 + (\delta_A^+ J^+ B)_3 B_3\right)\sin\theta\mu_{W^4}$. Upon plugging in Lemma 4.5 and Corollary 4.6 one encounters essentially the same expressions as for χ_1. The only non-zero terms are again those that contain derivatives of $C^{1,+}$, which are now given by

$$i_y^*\chi_2 = 2\sin\beta\cos^2\beta\mu_{W^4}\ \operatorname{Tr}\left(\nabla_3^\omega (C^{1,+})_1\mathfrak{t}_2 - \nabla_2^\omega (C^{1,+})_1\mathfrak{t}_3\right) + \mathcal{O}(y^\delta)$$

These cancel the remnants of $i_y^*\chi_1$, which concludes the proof. □

4.4.2 Knot boundaries

Consider a boundary face of type $\partial_K M$ and remember that it arises by blowup of a knot surface Σ_K. Let $\Sigma_K \times H^2 \times [0,1)_R$ denote a tubular neighbourhood, where H^2 is the two-dimensional hemisphere. Moreover, assume that the glancing angle θ between v and Σ_K is constant. If $\theta \neq \pi/2$, there is a non-vanishing vector field u parallel to Σ_K such that $g(u,v) = \sin\theta$.

In analyzing the behavior of the function χ near $\partial_K M$, we encounter some difficulties. The polyhomogeneous expansion of (a,b) in this case starts at y^0. While the leading order of χ remains essentially unchanged, given by straightforward analogues of the y^{-3}-terms from earlier, the additional y^0-modes of a and b lead to new contributions, beginning at $\mathcal{O}(y^{-2})$.

To determine these contributions, we need to repeat the analysis of the polyhomogeneous expansion near $\partial_K M$, but now based on the knot singularity model solutions. In this case, one considers model solutions $(A^{\lambda,\theta}, B^{\lambda,\theta})$ that are solutions of a θ-twisted version of the extended

Bogomolny equations (TEBE). Unfortunately, unlike the Nahm pole solutions, these knot singularity solutions are to the most part known only implicitly. For $G = SU(N)$ and $\theta = \pi/2$ the solutions have been described in [Mik12], while the existence of solutions for $G = SU(2)$ and general $\theta \in (0, \pi/2]$ is described by a continuation argument due to [Dim22a].

For the specific case with $G = SU(2)$ and $\theta = \pi/2$, the solutions have been described in [Wit11a] and can be given explicitly. To do so, introduce coordinates $(s, t, \psi, \vartheta, R)$ on $\Sigma_K \times H^2 \times [0,1)_R$, (s,t) are local coordinates on the surface Σ_K that we will also collectively refer to as z, while $R \in [0,1)$, $\psi \in [0, \pi/2]$ and $\vartheta \in [0, 2\pi]$ are global coordinates on the filled hemisphere. Let $(\mathfrak{t}_i)_{i=1,2,3}$ denote a standard basis of $\mathfrak{su}(2)$ and view $\mathfrak{t}_1$ as the generator of a fixed Cartan subalgebra. Introduce the $\mathfrak{sl}(2,\mathbb{C})$-valued function $\varphi = \phi_2 - i\phi_3$ as a convenient combination of the components ϕ_2 and ϕ_3 of $B = \sum_1^3 \phi_i e_i$. Similarly, denote by $E = \mathfrak{t}_2 - i\mathfrak{t}_3$, $H = \mathfrak{t}_1$, and $F = \mathfrak{t}_2 + i\mathfrak{t}_3$ the elements of an $\mathfrak{sl}(2,\mathbb{C})$-triple (E, H, F). The knot singularity solutions with charge $\lambda \in \mathbb{Z}$ are given by

$$\begin{aligned}
A_\vartheta &= -(\lambda+1)\cos^2\psi \frac{(1+\cos\psi)^\lambda - (1-\cos\psi)^\lambda}{(1+\cos\psi)^{\lambda+1} - (1-\cos\psi)^{\lambda+1}}\, H \\
\phi_1 &= -\frac{\lambda+1}{R} \frac{(1+\cos\psi)^{\lambda+1} + (1-\cos\psi)^{\lambda+1}}{(1+\cos\psi)^{\lambda+1} - (1-\cos\psi)^{\lambda+1}}\, H \\
\varphi &= \frac{(\lambda+1)}{R} \frac{\sin^\lambda\psi\ \exp(i\lambda\vartheta)}{(1+\cos\psi)^{\lambda+1} - (1-\cos\psi)^{\lambda+1}}\, E \\
A_s &= A_t = A_R = A_\psi = 0
\end{aligned}$$

We refer to these as the untwisted knot singularity model solutions.

The higher orders in a polyhomogeneous expansion of (A, B) around these model solutions was previously investigated by He [He18; HM19a]. Originally this was for Nahm pole solutions of the $\theta = \pi/2$ Kapustin-Witten equations, i.e. in four dimensions. However, the $\theta = \pi/2$ Kapustin-Witten equations are equivalent to a dimensional reduction of the Haydys-Witten equations along a direction perpendicular to v, and along which K extends to Σ_K. After an inconsequential reinterpretation of field components, the polyhomogeneous expansion carries over to the Haydys-Witten equations with knot singularity. Indeed, in that context the glancing angle between v and Σ_K is $\theta = \pi/2$, such that the untwisted knot singularity solutions are natural boundary conditions.

Theorem 4.8 ([He18; HM19a]). *Let (A, B) be a solution of the Haydys-Witten equations that satisfies Nahm pole boundary conditions with a knot singularity of weight λ and with glancing angle $\theta = \pi/2$. Correspondingly, write $A = A^{\lambda,\pi/2} + a$ and $B = B^{\lambda,\pi/2} + b$ with $a, b \in \mathcal{O}(R^{-1+\epsilon}s^{-1+\epsilon})$. Then a and b are polyhomogeneous in R and s, with non-negative exponents. In particular they satisfy the estimates*

$$\left|\nabla_z^\ell \nabla_R^m \nabla_{\cos\psi}^n\, a\right|_{C^0} \leq C_{\ell,m,n} R^{-\epsilon-m} (\cos\psi)^{2-\epsilon-n}\ , \quad \left|\nabla_z^\ell \nabla_R^m \nabla_{\cos\psi}^n\, b\right|_{C^0} \leq C_{\ell,m,n} R^{-\epsilon-m} (\cos\psi)^{1-\epsilon-n}$$

for any $\epsilon > 0$ and $\ell, m, n \in \mathbb{N}$.

We can use this to determine the leading order of χ in the case of $\theta = \pi/2$ knot singularities along $\partial_K M$. For this, let $i_R : \Sigma_K \times H^2 \hookrightarrow \Sigma_K \times H^2 \times [0,1)_R$ denote inclusion at radius R and write $\mu_{\Sigma_K \times H^2_R}$ for the pullback of the volume form under i_R, i.e. to the hemisphere (or cylinder) of radius R. First, since the dR-components of F_A drop out under pullback and the remaining components contain at most one derivative $\nabla_{\cos\psi}$, only components of F_A of order $\mathcal{O}(R^{-\epsilon}(\cos\psi)^{1-\epsilon})$ contribute to χ_1. Furthermore, $J^- B = \mathcal{O}(R^{-1}(\cos\psi)^{-1})$, such that $\chi_1 = -2\operatorname{Tr}(F_A \wedge J^- B) \sim CR^{-1-\epsilon}(\cos\psi)^{-\epsilon})\mu_{\Sigma_K \times H^2_R}$ Second, $\delta_A^+ J^+ B$ does not contain ∇_R and the leading order contributions of the product $J^+ B \wedge J^- B$ at $R^{-2}(\cos\psi)^{-2}$ vanish by construction of the underlying Nahm pole ($\operatorname{Tr}(HE) = 0$). We conclude that also $\chi_2 = -2\operatorname{Tr}(\delta_A^+ J^+ B \wedge J^- B \wedge \eta) \sim CR^{-1-\epsilon}(\cos\psi)^{-\epsilon})\mu_{\Sigma_K \times H^2_R}$. In fact, when the assumptions of Section 4.4.1 are satisfied, the exponent of $\cos\psi$ can be improved to $(\cos\psi)^\delta$, for some $\delta > 0$. This follows from the fact that $R\cos\psi = y$ and for fixed $R \neq 0$, the expansion in $\cos\psi$ must be consistent with the expansion in y^δ given in Proposition 4.7.

All in all, if (A, B) are Haydys-Witten solutions that satisfy (untwisted) Nahm pole boundary conditions near $\partial_K M$, then χ is of order $R^{-1-\epsilon}(\cos\psi)^\delta$, for some $\delta > 0$. It is natural to expect that for general $\theta \in (0, \pi/2]$ the behavior of the polyhomogeneous expansions of A and B near $\partial_K M$ follows a similar pattern to the one we observed for the pure β-twisted Nahm pole solutions at $\partial_{\mathrm{NP}} M$. Specifically, if $A = A^{\lambda,\theta} + a$ and $B = B^{\lambda,\theta} + b$ with respect to the twisted model solutions, it is still true that a and b admit polyhomogeneous expansions of order $\mathcal{O}(R^{-\epsilon}(\cos\psi)^{-\epsilon})$. In analogy to the cancellations in the case of Nahm pole boundary conditions, we then expect that a and b are restricted in a way that makes any contributions to χ below $R^{-1-\epsilon}(\cos\psi)^\delta$ vanish.

In any case, as we will see below, the following asymptotic behaviour of χ near a knot boundary $\partial_K M$ is sufficient to extend Theorem D to situations with knot singularities:

$$\chi \sim CR^{-2+\delta}(\cos\psi)^\delta \mu_{\Sigma_K \times H^2_R} \quad (R \to 0)$$

It's likely that such a result only holds under certain conditions on the topology and geometry of $E \to M^5$ and $\Sigma_K \subset \partial_{\mathrm{NP}} M$.

4.4.3 Kapustin-Witten ends

We now consider Kapustin-Witten ends $\partial_{\mathrm{KW}} M^5$ and $\partial_{\mathrm{NP-KW}} M^5$. As explained in Section 4.2, we take this to mean non-compact ends at which fields converge to a Kapustin-Witten solution with either finite energy or with Nahm pole at a corner of M^5, respectively. We let $[0,1)_s \times W^4$ be a tubular neighbourhood and denote by $i_s : W^4 \hookrightarrow [0,1)_s \times W^4$ inclusion of W^4 at s. Write $u = s\partial_s$ for the inward-pointing unit normal vector field and assume $g(u, v) = \cos\theta$ is constant. As usual, if $\theta \neq 0$, let w be the non-vanishing unit vector field on W^4 with respect to which $v = \cos\theta u + \sin\theta w$.

We begin with a general result about the pullback of χ in dependence of the limiting field configuration. Hence, assume (A, B) approaches a solution of the θ-Kapustin-Witten equations

$(\tilde{A}, \phi)$ on W^4. In a local basis $(s, t, x^i)_{i=1,2,3}$ with $w = \partial_t$, the pullback of $\chi = 2\,\mathrm{Tr}(F_A \wedge J^- B + \delta_A^+ J^+ B \wedge J^- B \wedge \eta)$ is then given by

$$\frac{1}{2} \lim_{s\to 0} i_s^* \chi = \cos\theta \,\mathrm{Tr}\left((F_{\tilde{A}}^+)_{t2}\phi_2 + (F_{\tilde{A}}^+)_{t3}\phi_3\right)\mu_{W^4} + \sin\theta\,\mathrm{Tr}\left(\nabla_2^{\tilde{A}}\phi_1\,\phi_3 - \nabla_3^{\tilde{A}}\phi_1\,\phi_2\right)\mu_{W^4} \quad (4.16)$$

Here, $F_{\tilde{A}}^+$ denotes the self-dual part with respect to the four-dimensional Hodge operator $\star_{W^4}$.

Finite Energy Solutions. Kapustin and Witten observed that Kapustin-Witten solutions on closed manifolds are highly restricted [KW07]. A similar result was established by Nagy and Oliveira for finite energy solutions on ALE and ALF gravitational instantons [NO21]. The energy in question, usually called Kapustin-Witten energy, is defined by the functional

$$E_{\mathrm{KW}} = \int_{W^4} \|F_{\tilde{A}}\|^2 + \left\|\nabla^{\tilde{A}}\phi\right\|^2 + \|[\phi\wedge\phi]\|^2$$

Nagy and Oliveira observed that for finite energy solutions the norm of the Higgs field is bounded, which yields surprisingly strong restrictions for solutions on $\mathbb{R}^4$ and $S^1 \times \mathbb{R}^3$, when combined with a result by Taubes [Tau17a]. As proposed by Nagy and Oliveira, we were able to generalize this to arbitrary ALE and ALF gravitational instantons in Chapter 3. The situation is summarized by the following two theorems.

Theorem ([KW07; GU12]). *Let $E \to W^4$ be an $SU(2)$ principal bundle over a compact manifold without boundary. Assume (A, ϕ) satisfies the θ-Kapustin-Witten equations with $\theta \in (0, \pi)$. If $E \to W^4$ has non-zero Pontryagin number, then A and ϕ are identically zero. Otherwise $A + i\phi$ is a flat $PSL(2, \mathbb{C})$ connection; equivalently $F_A = [\phi \wedge \phi]$ and $\nabla^A \phi = 0$.*

Theorem (Corollary C). *Let (A, ϕ) be a finite energy solution of the θ-Kapustin-Witten equations with $\theta \neq 0 \pmod{\pi}$ on an ALE or ALF gravitational instanton and let $G = SU(2)$. Then A is flat, ϕ is ∇^A-parallel, and $[\phi\wedge\phi] = 0$.*

In combination with equation (4.16), we immediately arrive at the following result.

Proposition 4.9. *Let $G = SU(2)$ and $\theta \not\equiv 0 \pmod{\pi}$. Assume (A, B) approaches a finite energy solution of the θ-Kapustin-Witten equations on W^4 as $s \to 0$. If W^4 is an ALE or ALF gravitational instanton, or if W^4 is compact and either (i) $E \to W^4$ has non-zero Pontryagin number or (ii) $\theta = \pi/2$, then $\lim_{s\to 0} i_s^*\chi = 0$.*

Proof. If W^4 is an ALE or ALF gravitational instanton, both terms in (4.16) vanish directly. The first since $\tilde{A}$ is flat and the second since ϕ is $\nabla^{\tilde{A}}$ parallel.

If W^4 is compact, A and ϕ can only be non-zero if the Pontryagin number of $E \to W^4$ is zero and, furthermore, in that case $F_{\tilde{A}} = [\phi\wedge\phi]$ and $\nabla^{\tilde{A}}\phi = 0$. It follows that $\lim_{s\to 0} i_s^*\chi = 2\cos\theta\,\mathrm{Tr}([\phi_1, \phi_2]\phi_3)\mu_{W^4}$, which vanishes if $\theta = \pi/2$. □

We make two observations. First, if $\theta \equiv 0 \pmod{\pi}$, the natural boundary condition for (A, B) are finite energy solutions $(\tilde{A}, B, C)$ of the Vafa-Witten equations on W^4. On $\mathbb{R}^4$ the Vafa-Witten and $\theta = 0$ Kapustin-Witten equations are equivalent by an identification of $\phi = Cdx^0 + B_{0i}dx^i$. According to the result by Taubes mentioned earlier, any solution with bounded $\|\phi\|$ satisfies $\sigma(B, B) = [B, C] = 0$ and $\nabla^A B = \nabla^A C = 0$ [Tau17a].

It is not currently known if Vafa-Witten solutions on ALX spaces have a similar property. However, the Vafa-Witten equations are still closely related to the $\theta = 0$ version of the Kapustin-Witten equations, and Taubes' result has been established for the latter in Chapter 3. One might thus expect that, at least in certain situations, the vanishing results carry over to Vafa-Witten solutions. Whenever this is the case, the Vafa-Witten equations $F^+_{\tilde{A}} = \sigma(B, B) + [B, C]$ reduce to anti-self-dual equations for $\tilde{A}$, in which case $\lim_{s\to 0} i_s^* \chi = F^+_{\tilde{A}} \wedge i_s^* J^- B = 0$, as we found above.

Second, the fact that χ converges to zero holds at two ends of a spectrum of asymptotic volume growth of W^4. On the one hand, in the case of compact manifolds with bounded volume or equivalently asymptotic volume growth of order r^0, and on the other hand, on ALF and ALE manifolds with asymptotic volume growth of order r^3 and r^4, respectively. For ALG and ALH gravitational instantons, the proof strategy of Nagy and Oliveira result doesn't work, because they rely on the existence of a positive Green's function for the Laplacian. The approach to Taubes' dichotomy fails for ALH manifolds for much the same reason.

To the best of the authors knowledge it is currently not known if an analogue of the results of Nagy and Oliveira should be expected to be true or false on ALG and ALH gravitational instantons. However, from the physics perspective there is no obvious reason to single out intermediate volume growth like that. On the contrary, in Witten's approach to Khovanov homology, where one considers $M^5 = \mathbb{R}_s \times X^3 \times \mathbb{R}^+_y$, it is natural to expect simplifications whenever X^3 has additional substructure.

For example, one expects that Khovanov homology arises for $X^3 = S^3$, in which case M^5 has an ALH end $W^4 = \mathbb{R}_s \times S^3$ at $y \to \infty$. Also, Gaiotto and Witten recovered the Jones polynomial purely by adiabatically braiding solutions of the EBE under the assumption that $X^3 = \mathbb{R}_t \times \Sigma^2$, in which case one encounters an ALG manifold at $y \to \infty$. It thus seems plausible to postulate that the results of this section hold, more generally, whenever W^4 is a complete Ricci-flat Riemannian manifold with sectional curvature bounded from below and (A, B) approaches a finite energy solution of the θ-Kapustin-Witten equations ($\theta \neq 0, \pi$) or the Vafa-Witten equations ($\theta = 0$).

Nahm Poles at Corners. We now discuss the behaviour of χ at non-compact ends of the class $\partial_{\mathrm{NP-KW}} M^5$. As before, we denote by $[0, 1)_s \times W^4$ a tubular neighbourhood. In this situation we still demand that (A, B) converges to a θ-Kapustin-Witten solution $(\tilde{A}, \phi)$ on W^4 as $s \to 0$. But in contrast to earlier, we now assume that $(\tilde{A}, \phi)$ satisfies β-twisted Nahm pole boundary conditions at an adjacent corner of M^5 – or equivalently at a boundary of W^4. Observe that $(\tilde{A}, \phi)$ cannot have finite energy if it exhibits a Nahm pole, such that Proposition 4.9 doesn't apply.

In the absence of a finite energy condition, we now also have to specify boundary conditions at non-compact ends $[0,1)_{s'} \times X^3$ of W^4. We demand that $\tilde{A} + i\phi$ converges to a flat $G_{\mathbb{C}}$ connection on X^3 as $s' \to 0$. From now on we refer to these configurations simply as Nahm pole solutions of the θ-Kapustin-Witten equations.

Remark. Note that a non-compact end $[0,1)_{s'} \times X^3$ of W^4 corresponds to a 'corner at infinity' $[0,1)_s \times [0,1)_{s'} \times X^3$ of M^5. This corner is adjacent to two non-compact ends of M^5, at which (A, B) converges to a corresponding solution of the θ- or θ'-Kapustin-Witten equation, respectively. The two associated asymptotic boundary conditions, which demand that both $\tilde{A} + i\phi$ and $\tilde{A}' + i\phi'$ are flat connections on X^3, have to be consistent with the fact that both Kapustin-Witten solutions arise from the common five-dimensional fields (A, B). Put differently: If we view B as a one-form by the usual isomorphism, the pullback of $A + iB$ converges to a flat connection on X^3 as $s, s' \to 0$.

We similarly need to ensure that the boundary conditions are compatible at corners $[0,1)_s \times X^3 \times [0,1)_y$ that are adjacent to a θ-Kapustin-Witten end as $s \to 0$ and a β-Nahm pole as $y \to 0$. This is only the case if $\beta = \pi/2 - \theta$, since otherwise the β-twisted Nahm pole model solutions are not solutions of the θ-Kapustin-Witten equations.

Proposition 4.10. *Assume (A, B) approaches a solution of the θ-Kapustin-Witten equations on W^4. Then*

$$\lim_{s \to 0} i_s^* \chi = \frac{2}{3} i_0^* \operatorname{Tr}(\sigma(B,B) \wedge B) + 2 \frac{\sin\theta}{\cos^2\theta} \, i_0^* d \operatorname{Tr}\left(\iota_w(J^+ B) \wedge J^- B\right) .$$

and the expression is exact if $\theta = \pi/2$.

Proof. This is a rewriting of equation (4.16). Since $(\tilde{A}, \phi)$ satisfy the θ-Kapustin-Witten equations, we can replace the field strength by $F_{\tilde{A}} = \frac{1}{2}[\phi \wedge \phi] - \cot\theta \, d_{\tilde{A}}\phi + \csc\theta \star_4 d_{\tilde{A}}\phi$. After a short calculation (and slightly miraculous cancellations), we find

$$\lim_{s \to 0} i_s^* \chi = 2\cos\theta \operatorname{Tr}([\phi_1, \phi_2]\phi_3)\mu_{W^4} + \sin\theta \;\; \operatorname{Tr}(\nabla_2^A(\phi_1\phi_3) - \nabla_3^A(\phi_1\phi_2))\mu_{W^4} \, ,$$

which is a local representation of the expression above and shows that the right hand side is exact if $\cos\theta = 0$ □

4.5 Vanishing of the Boundary Term

We can now show that the contributions from the exact term in the Weitzenböck formula of Lemma 4.2 vanishes when the various conditions we have encountered in Section 4.3 and Section 4.4 are satisfied. In summary we make the following assumptions:

(A1) At $\partial_{\mathrm{NP}} M^5$ the fields satisfy regular β-twisted Nahm pole boundary conditions for some generic β. Writing $A = y^{-1} \sin\beta \phi_\rho^\tau + \omega + a$ and $B = y^{-1} \cos\beta \phi_\rho + b$, assume that $F_\omega = 0$ and that $\nabla_s^\omega b_2 + \nabla_1^\omega b_3 = \mathcal{O}(y^2)$ and $\nabla_s^\omega b_3 - \nabla_1^\omega b_2 = \mathcal{O}(y^2)$ (cf. Section 4.4.1).

(A2) At $\partial_K M^5$ the fields are asymptotic to knot singularity models and there is some $\delta > 0$ such that $i_R^* \chi \sim CR^{-2+\delta}(\cos\psi)^\delta \mu_{\Sigma_K \times H_R^2}$ as $R \to 0$ (cf. Section 4.4.2).

(A3) At $\partial_{\mathrm{KW}} M^5$ the fields approach a finite energy solution of the θ-Kapustin-Witten equations. The boundary face $\partial_{\mathrm{KW}} M^5$ is either a) an ALE or ALF gravitational instanton, b) a compact manifold on which the bundle $E \to \partial_{\mathrm{KW}} M^5$ has non-zero Pontryagin number, or c) a compact manifold with incidence angle $\theta = \pi/2$ (cf. Section 4.4.3).

(A4) At $\partial_{\mathrm{NP-KW}} M^5$ the incidence angle is $\theta = \pi/2$ and the fields approach Kapustin-Witten solutions with Nahm poles at boundaries. Moreover, at non-compact cylindrical ends of $\partial_{\mathrm{NP-KW}} M^5$, the combination $A + iB$ converges to a flat $G_{\mathbb{C}}$ connection that satisfies $J^- \sigma(B, B) = 0$. (cf. Section 4.4.3).

Theorem 4.11. *Let $G = SU(2)$, M^5 a manifold with poly-cylindrical ends, v a non-vanishing vector field that approaches ends at a constant angle, and J an almost Hermitian structure on* $\ker\eta$. *Assume* $\mathbf{HW}_v(A, B) = 0$ *and that* (A1) - (A4) *are satisfied, then* $\mathbf{dHW}_{v,J}(A, B) = 0$.

Proof. Our starting point is a regularized version of the Weitzenböck formula of Lemma 4.2, obtained by restricting the domain of integration to the compact submanifold with corner M^5_ϵ introduced in Section 4.2 and taking $\epsilon \to 0$.

$$\int_{M^5} \|\mathbf{HW}_v(A, B)\|^2 = \int_{M^5} \left\|\mathbf{dHW}_{v,J}(A, B)\right\|^2 + \lim_{\epsilon \to 0} \int_{M^5_\epsilon} d\chi$$

According to Stokes' theorem, the contributions of the exact term are now determined by

$$\lim_{\epsilon \to 0} \int_{M_\epsilon} d\chi = \sum_{i \in I} \lim_{\epsilon \to 0} \int_{\partial_i M_\epsilon} \chi \ .$$

We address the integrals for each of the four boundary classes independently.

Nahm Pole Boundaries: Let $W^4 \times [0, 1)_y$ be a tubular neighbourhood of a Nahm pole boundary $\partial_{\mathrm{NP}} M^5$. The boundary face $\partial_{NP} M_\epsilon$ is a subset of the ϵ-displacement $\{\epsilon\} \times W^4 \hookrightarrow [0, 1)_s \times W^4$. We know from Proposition 4.7 that there is some $\delta > 0$ such that $i_y^* \chi \sim Cy^\delta \mu_{W^4}$ as $y \to 0$. This provides the estimates

$$\lim_{\epsilon \to 0} \left| \int_{\partial_{\mathrm{NP}} M_\epsilon} i_\epsilon^* \chi \right| \le \lim_{\epsilon \to 0} \int_{\partial_{\mathrm{NP}} M_\epsilon} |i_\epsilon^* \chi| \le \lim_{\epsilon \to 0} \int_{W^4 \times \{\epsilon\}} |C| \, \epsilon^\delta \mu_{W^4} = 0 \ .$$

We have used that the leading order of the pullback of χ extends unchanged to all of W^4, due to compatibility of boundary conditions at corners.

Knot Boundaries: Let $\Sigma_K \times H^2_{\psi,\vartheta} \times [0,1)_R$ be a tubular neighbourhood of a knot boundary $\partial_K M^5$, where $H^2_{\psi,\vartheta}$ denotes the two-dimensional hemisphere, parametrized by $\psi \in [0,\pi/2]$ and $\vartheta \in [0,2\pi]$. The associated boundary of the regularized manifold $\partial_K M^5_\epsilon$ is contained in the ϵ-displacement $\Sigma_K \times H^2 \times \{\epsilon\} = \Sigma_K \times H^2_\epsilon$. The pullback of the volume form to the hemisphere of radius $R = \epsilon$ is given by $\mu_{\Sigma_K \times H^2_\epsilon} = \epsilon^2 d\psi d\vartheta \mu_{\Sigma_K}$. Assuming explicitly that $i_R^*\chi \sim cR^{-2+\delta}(\cos\psi)^\delta \mu_{\Sigma_K \times H^2_R}$ as $R \to 0$, we find

$$\lim_{\epsilon\to 0}\left|\int_{\partial_K M^5_\epsilon} i_\epsilon^*\chi\right| \le \lim_{\epsilon\to 0}\int_{\Sigma_K\times H^2\times\{\epsilon\}} |i_\epsilon^*\chi| \le \lim_{\epsilon\to 0}\int_{\Sigma_K}\int_{H^2} |C|\,\epsilon^{-2+\delta}\epsilon^2 d\psi d\vartheta \mu_{\Sigma_K} = 0\ ,$$

where we have used the Fubini-Tonelli theorem to split off integration along Σ_K. In extending the integral from $\partial_K M_\epsilon$ to all of $\Sigma_K \times H \times \{\epsilon\}$, we have used the compatibility of knot singularities with the pure Nahm pole boundary conditions (which are the only boundary conditions that are adjacent to knot singularities) at the corner $\cos\psi \to 0$.

Kapustin-Witten Ends: Let $[0,1)_s \times W^4$ be a tubular neighbourhood of a Kapustin-Witten end $\partial_{\mathrm{KW}} M^5$. Proposition 4.9 states that $\lim_{s\to 0} i_s^*\chi = 0$. Since $\partial_{\mathrm{KW}} M_\epsilon \subset W^4$ and W^4 is ALE or ALF, its volume grows asymptotically at most with ϵ^{-4}. Looking back at (4.16), the rate of decay of χ is determined by how fast F_A and $\nabla^A\phi$ approach zero as $s \to 0$. Since the Haydys-Witten equations represent flow equations of the Kapustin-Witten equations, a typical solution is expected to decay exponentially fast towards the stationary solution. We conclude that in the limit $\epsilon \to 0$:

$$\lim_{\epsilon\to 0}\left|\int_{\partial_{\mathrm{KW}} M^5_\epsilon} \chi\right| \le \lim_{\epsilon\to 0}\int_{\partial_{\mathrm{KW}} M^5_\epsilon} |i_\epsilon^*\chi| = 0\ .$$

Kapustin-Witten Ends with Nahm Poles: Let $[0,1)_s \times W^4$ be a non-compact end of M^5 and assume that $g(\partial_s, v) = 0$, i.e. $\theta = \pi/2$. As before, denote by $i_s : W^4 \hookrightarrow [0,1)_s \times W^4$ inclusion of W^4 at s.

Assume (A,B) converges to a Kapustin-Witten solution $(\hat{A},\phi)$ that exhibits a Nahm pole at a boundary of W^4. Since $\theta = \pi/2$, Proposition 4.10 states that $\lim_{s\to 0} i_s^*\chi = d\omega$, where $\omega = \mathrm{Tr}\left(\iota_w(J^+B) \wedge J^-B\right)$. It follows that

$$\lim_{s\to 0}\int_{\partial_{\text{NP-KW}} M_\epsilon} i_s^*\chi = \int_{W^4} d\omega$$

and it remains to determine the integral of $d\omega$ over W^4.

For this, let $B_r(p)$ denote the (four-dimensional) ball of radius r centered at some point $p \in W^4$ and $\mu_{B_r(p)}$ its volume form. A classic result by Yau then states the following.

Theorem ([Yau76, Theorem 3 & Appendix (ii)]). *If* $\liminf_{r\to\infty} r^{-1}\int_{B_r(p)} |\omega|\, \mu_{B_r(p)} = 0$, *then* $\int_{W^4} d\omega = 0$.

Near the boundary we can rely on Lemma 4.5. Applied to the Kapustin-Witten solution $(\tilde{A}, \phi)$ with $\beta = \pi/2 - \theta = 0$, this yields in the limit $y \to 0$

$$\omega = y^{-2}\,\mathrm{Tr}(\mathfrak{t}_1\mathfrak{t}_2 + \mathfrak{t}_1\mathfrak{t}_3) + y^0\,\mathrm{Tr}(\mathfrak{t}_1(C^{1,+})_2 + \mathfrak{t}_1(C^{1,+})_3 + (C^{1,+})_1\mathfrak{t}_2 + (C^{1,+})_1\mathfrak{t}_3) + \mathcal{O}(y^{\delta})\,.$$

The term proportional to y^{-2} vanishes and the term at constant order is assumed to be integrable on X^3, such that contributions from the Nahm pole to the integral $\int_{B_r(p)} |\omega|$ are harmless.

At a non-compact end $[0,1)_s \times X^3$ of W^4, the asymptotic boundary condition states that $A + iB$ approaches a flat $SL(2,\mathbb{C})$-connection on X^3 that satisfies $J^-\sigma(B,B) = 0$. Equivalently, the component ϕ_1 of J^+B commutes with the components $\phi_{2/3}$ of J^-B. This is for example the case if X^3 is a product $S^1 \times \Sigma$, where Σ is a Riemann surface and the almost complex structure J is the direct sum of complex structures on the cylinder $[0,1)_s \times S^1$ and Σ. In any case, since $[\phi_1, \phi_2] = [\phi_1, \phi_3] = 0$ we find that $\omega = \mathrm{Tr}(\phi_1\phi_2) + \mathrm{Tr}(\phi_1\phi_3) \to 0$. Assuming $A + iB$ converges to the flat connection faster than the volume of the geodesic balls $B_r(p)$ grows, we conclude that $\int_{W^4} d\omega = 0$.

Conclusion: Since all boundary contributions to the exact term vanish in the limit $\epsilon \to 0$, we arrive at

$$\int_{M^5} \|\mathbf{HW}_v(A,B)\|^2 = \int_{M^5} \left\|\mathbf{dHW}_{v,J}(A,B)\right\|^2 + \lim_{\epsilon\to 0}\int_{M^5_\epsilon} d\chi = \int_{M^5} \left\|\mathbf{dHW}_{v,J}(A,B)\right\|^2\,.$$

Seeing that the integrands on both sides are non-negative, we find that whenever $\mathbf{HW}(A,B) = 0$ also $\mathbf{dHW}(A,B) = 0$, which concludes the proof.

□

5 Comoving Higgs Bundles and Symplectic Khovanov Homology

Let $M^5 = C \times \Sigma \times \mathbb{R}^+_y$, where C and Σ are Riemann surfaces, and assume M^5 is equipped with a product metric g and the non-vanishing unit vector field $v = \partial_y$. Write $\eta = g(v, \cdot)$ and observe that $\ker \eta$ coincides with the tangent space of $C \times \Sigma$. Throughout, we let J be the almost complex structure on $\ker \eta$ that is induced by the complex structures on C and Σ, respectively. We consider a principal G-bundle $E \to M^5$ for $G = SU(N)$ and denote by $\mathrm{ad}\, E$ its adjoint bundle.

In this chapter we investigate the decoupled Haydys-Witten equations on M^5 with respect to v and J, as introduced in Chapter 4 (also see Section 2.4.1 for more detailed definitions). These are equations for a pair of connection $A \in \mathcal{A}(E)$ and Haydys' self-dual two-form $B \in \Omega^2_{v,+}(M^5, \mathrm{ad}\, E)$. J lifts to a map on $\Omega^2_{v,+}(M^5)$ with eigenvalues ± 1. The decoupled Haydys-Witten equations (dHW) are defined by

$$\begin{aligned} \tfrac{1+J}{2}\left(\sigma(B,B) + \nabla^A_v B\right) &= F^+_A & \qquad \tfrac{1-J}{2}\left(\sigma(B,B) + \nabla^A_v B\right) &= 0 \\ \delta^+_A \tfrac{1+J}{2} B &= \iota_v F_A & \qquad \delta^+_A \tfrac{1-J}{2} B &= 0 \end{aligned} \tag{5.1}$$

We are interested in solutions of these equations and focus on their role in the context of Witten's gauge theoretic approach to Khovanov homology [Wit11a]. As described in much more detail in Chapter 2, one constructs a Floer cochain complex out of solutions of the Kapustin-Witten equations on the four-manifold $W^4 = X^3 \times \mathbb{R}^+_y$, subject to certain singular boundary conditions with monopole-like behaviour along a knot $K \subset \partial W^4 = X^3$. Its cohomology with respect to the Floer differential, which counts the number of solutions of the (full) Haydys-Witten equations on $M^5 = \mathbb{R}_s \times W^4$ that interpolate between the Kapustin-Witten solutions at $s \to \pm\infty$, is expected to be a topological invariant. Witten conjectures that for $X^3 = S^3$ or $\mathbb{R}^3$, this topological invariant coincides with Khovanov homology.

The results obtained in Chapter 4 raise hope that, on nice enough manifolds, every solution of the Haydys-Witten (and Kapustin-Witten) equations is already a solution of the decoupled version of the equations. This is advantageous, because the decoupled equations exhibit a Hermitian Yang-Mills structure that simplifies their analysis considerably. In particular, they contain as a subset the extended Bogomolny equations (EBE), and for the latter it is possible to exploit the corresponding Hermitian Yang-Mills structure to establish a classification in terms of Higgs bundles over Σ with certain extra structure [HM19c; HM20; HM19b; Dim22b; Sun23].

This classification of EBE-solutions follows an earlier conjecture of Gaiotto and Witten from their highly influential article [GW12]. In that work, they further propose to determine solutions of the Kapustin-Witten equations by way of an ‘adiabatic braiding’ of EBE-solutions. While they were able to show that these ideas lead to an action of the braid group on the Poincaré polynomial of the Floer complex in terms of the Jones representation of Virasoro conformal blocks, a direct calculation on the level of Floer homology remains an open problem.

It should be noted that Gaiotto and Witten also laid out an Atiyah-Floer type program to calculate the invariants associated to Haydys-Witten Floer theory, where instanton Floer theory is replaced by a Lagrangian intersection Floer theory. For this, fix a Heegard splitting $W^4 = H_1 \cup_\Sigma H_2$ and suppose that we stretch the metric transversely to $H_1 \cap H_2$ such that the two handlebodies are joined by a long neck of the form $[-L, L]_t \times \Sigma$, $L >> 1$. Position the knot K such that the portion of the knot in the long neck consists of a set of parallel straight lines $[-L, L]_t \times \{p_j\}$, intersecting $\Sigma \times \{0\}$ in a finite collection of points. If $\mathcal{M}_\Sigma$ denotes the $\mathcal{G}_\mathbb{C}$ character variety of Σ, then the character varieties of the H_i are Lagrangians $L_1, L_2 \subset \mathcal{M}_\Sigma$. The moduli space of Kapustin-Witten solutions over $[-L, L]_t \times \Sigma \times \mathbb{R}^+_y$ that are invariant in the direction of t provides a third Lagrangian L_3. In the absence of knots, the Atiyah-Floer conjecture states that Lagrangian intersection Floer homology of L_1 and L_2 is an invariant of W^4 and that this invariant coincides with the original instanton Floer cohomology, see [AM20; DF17] for recent progress. Including the effect of knots and counting instead holomorphic triangles that span between L_1, L_2, L_3 in $\mathcal{M}_\Sigma$ conjecturally yields the coefficients of the Jones polynomial. We refer to [Guk+17] for more details and advances in this approach.

In the present work, however, we remain on the side of instanton Floer theory and investigate the implications of the adiabatic approach on solutions of the decoupled Haydys-Witten and Kapustin-Witten equations directly. The adiabatic condition can be viewed as the assumption that on a given slice $[-L, L]_t \times \Sigma \times \mathbb{R}^+_y$ the knot position varies only slightly in t, such that one obtains a Kapustin-Witten solution from a smooth family of EBE-solutions that ‘remain in the ground state’ when one moves from $t = -L$ to L. This means that Kapustin-Witten solutions should be related to certain well-behaved paths in the moduli space of EBE-solutions.

Let us explain this in more detail. Consider the decoupled Kapustin-Witten equations on $S^1_t \times \mathbb{C}$ and a knot of the form $K = \bigsqcup_{a=1}^k \{(t, p_a(t), 0)\}$. The collection of trajectories $\{p_a(t)\}_{a=1,\dots,k}$ can be viewed as a loop $\beta : S^1_t \to \mathrm{Conf}_k\,\mathbb{C}$ in the configuration space of k distinct, ordered points in $\mathbb{C}$. Write $\mathcal{M}^{\mathrm{dKW}}_K$ for the moduli space of dKW-solutions subject to Nahm pole boundary conditions with knot singularities along K. We can view $\mathcal{M}^{\mathrm{dKW}}_K$ as fiber of a bundle $\mathcal{M}^{\mathrm{dKW}} \to \Omega\,\mathrm{Conf}_k\,\mathbb{C}$, where the fiber map sends each solution to the knot at which the solution exhibits a knot singularity. There is an analogous fiber bundle for solutions of the extended Bogomolny equations $\mathcal{M}^{\mathrm{EBE}} \to \mathrm{Conf}_k\,\mathbb{C}$ that sends a solution to the position of knot singularities $D = \{p_a\}_{a=1,\dots,k} \subset \Sigma$.

The adiabatic approach can now be interpreted as the statement that one expects that there is a bundle map of the form

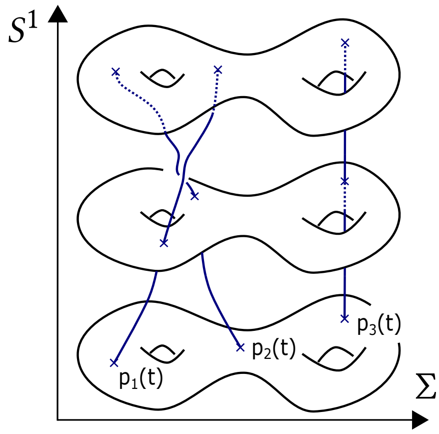

Figure 5.1 A general knot K in the boundary of $W^4 = S^1_t \times \Sigma \times \mathbb{R}^+_y$ varies with time. The adiabatic approach can be viewed as stretching the size of S^1_t, such that at any given time t, the fields (A, ϕ) are well-approximated by a solution of the extended Bogomolny equations.

$$\begin{tikzcd} \mathcal{M}^{\mathrm{dKW}} \arrow[rr, dashed] \arrow[dr] & & \Omega\mathcal{M}^{\mathrm{EBE}} \arrow[dl] \\ & \Omega\,\mathrm{Conf}_k\,\Sigma & \end{tikzcd}$$

The following result by He and Mazzeo establishes the existence of such a map for the case of S^1-invariant Kapustin-Witten solutions.

Theorem ([HM19a]). *Every solution of the EBE on $\Sigma \times \mathbb{R}_+$ with Nahm pole boundary condition and knot singularities at a divisor $D = \{z_a\}_{a=1,\dots,k} \subset \Sigma$ lifts to an S^1-invariant solution of the KW-equations with knot singularities along the S^1-invariant knot $K = S^1 \times D$. Moreover, every S^1-invariant solution of the KW-equations is given by such a lift and there is a bijection between moduli spaces*

$$\mathcal{M}^{\mathrm{KW}}_{S^1\times D} \to \mathcal{M}^{\mathrm{EBE}}_{D}\,.$$

Note that, as dimensional reduction of the Kapustin-Witten equations, any EBE-solution provides an S^1-invariant solution of the Kapustin-Witten equations. So the content of the theory is that also the reverse is true and any S^1-invariant solution is necessarily a lift of an EBE-solution. This also means that the target space of the bundle map must be the space of non-vertical loops, which we will denote by $\Omega_h\mathcal{M}^{\mathrm{EBE}}$.

The proof of He and Mazzeo's theorem is based on a Weitzenböck formula that equates the full Kapustin-Witten equations with the EBE up to certain boundary terms. A prerequisite for this Weitzenböck formula is the assumption that the position of knot singularities is S^1_t-invariant. This assumption will be dropped in the analysis presented here, such that solutions can have a richer structure.

Instead of the Weitzenböck formula of He and Mazzeo, we rely on the Hermitian Yang-Mills structure of the decoupled equations[1]. This structure is most easily described in the $4\mathcal{D}$-formalism, where one uses the Haydys-Witten fields (A, B) to define four $\operatorname{ad} E_{\mathbb{C}}$-valued differential operators $\mathcal{D}_\mu$, $\mu = 0, 1, 2, 3$ (see Section 5.1). In terms of these operators, the decoupled Haydys-Witten equations are equivalent to

$$[\mathcal{D}_\mu, \mathcal{D}_\nu] = 0,\ \ \mu, \nu = 0, 1, 2, 3\ , \qquad \sum_{\mu=0}^{3} [\mathcal{D}_\mu, \overline{\mathcal{D}_\mu}]\ .$$

Crucially, the set of equations on the left is invariant under complex gauge transformations $\mathcal{G}_{\mathbb{C}}$, while the remaining equation on the right can be viewed as a real moment map condition. As a consequence, it's possible to first solve the easier $\mathcal{G}_{\mathbb{C}}$-invariant part of the equations and subsequently solve for a complex gauge transformation that solves the moment map condition.

In a first step, following the adiabatic approach of Gaiotto and Witten, we consider a situation where the knot K is a small deformation of an S^1-invariant one K'. More precisely, we assume that the two knots are connected by a well-behaved isotopy $\beta_\bullet$, where $\beta_0 = K'$ and $\beta_1 = K$. Since the underlying theory is topological, such an isotopy can not have a large effect on the solution. We then introduce an Ansatz that solves the $\mathcal{G}_{\mathbb{C}}$-invariant part of the decoupled Haydys-Witten equations and exhibits a knot singularity along the knot trajectories determined by $\beta_q(t)$ at each point $q \in [0, 1]$ of the isotopy. Put differently, the isotopy provides a deformation from an initially S^1-invariant Ansatz with knot singularities along the S^1-invariant knot K', to a t-dependent Ansatz with knot singularities along the t-dependent knot K. We propose that this isotopy Ansatz provides a way to transport a solution of the decoupled Haydys-Witten equations with S^1-invariant knot to a solution with t-dependent knot.

Another main tool that we make extensive use of is a physically motivated reduction from the infinite-dimensional moduli space of Kapustin-Witten solutions (before gauge fixing) to a finite-dimensional model space. The space in question is related to a partial Grothendieck-Springer resolution of $\mathfrak{sl}(N)$ that is naturally fibered over the configuration space of points. Motivated by these considerations, we formulate Conjecture E stating that on $S^1_t \times \Sigma \times \mathbb{R}^+_y$, the number of intersection points of the Grothendieck-Springer fiber and its parallel transport along S^1_t determines a lower bound for the number of solutions to the decoupled Kapustin-Witten. Moreover, taking this line of argument to its logical conclusion, we state a stronger version Conjecture F, claiming that Haydys-Witten Floer theory is isomorphic to symplectic Khovanov-Rozansky homology defined by Seidel, Smith and Manolescu [SS04; Man07]. Since symplectic Khovanov homology (associated to $G_{\mathbb{C}} = SL(2, \mathbb{C})$) is isomorphic to a grading reduced version of Khovanov homology [AS19], the arguments developed in this chapter provide a novel approach to prove Witten's conjecture, complementary to the Atiyah-Floer approach pursued in the Gaiotto-Witten program.

This chapter is arranged into two parts.

[1] In a sense, the Weitzenböck formula of He and Mazzeo is replaced by the one of Chapter 4 that establishes the decoupling of the equations on $M^5 = C \times \Sigma \times \mathbb{R}^+_y$.

Sections 5.1 - 5.4 provide a review of the relevant background and we use the opportunity to fix some of the notation used in the remainder of this chapter. Specifically, Section 5.1 spells out the Hermitian Yang-Mills structure of the decoupled Haydys-Witten equations; Section 5.2 provides some notation regarding Lie algebras and adjoint orbits; Section 5.3 introduces the Nahm pole boundary conditions with knot singularities in a form that is adjusted to the Hermitian Yang-Mills structure; and Section 5.4 summarizes the results of He and Mazzeo for solutions of EBE-solutions [HM19c; HM20], highlighting certain aspects that will have close analogues in subsequent discussions for solutions of the decoupled Haydys-Witten and Kapustin-Witten equations.

Sections 5.5 - 5.10 formalize the adiabatic approach. We introduce the isotopy Ansatz and an associated expansion of the decoupled Haydys-Witten equations in Section 5.5 and subsequently describe in Section 5.6 a strategy to obtain solutions of the decoupled Kapustin-Witten equations for any null-isotopic single-stranded knot by use of a continuity argument. In Section 5.7 we explain the relation between solutions of the decoupled Haydys-Witten equations and paths in the moduli space of Higgs bundles $(\mathcal{E}, \varphi)$ that are equipped with the additional structure of a distinguished line subbundle L. Based on physical intuition, we propose that for our purposes the Grothendieck-Springer fibration can be used to model the moduli space of triples $(\mathcal{E}, \varphi, L)$ in Section 5.8 and using these insights, we formulate Conjecture E. In Section 5.9 we explain that one naturally obtains Lagrangian submanifolds of the fibers when the S^1_t-factor is decompactified to $\mathbb{R}_t$ and studies compact knots as braid closures. Finally, Section 5.10 incorporates Haydys-Witten instantons into the setting, which leads to Conjecture F that Haydys-Witten instanton Floer homology coincides with symplectic Khovanov homology.

5.1 Hermitian Yang-Mills Structure of the Decoupled Haydys-Witten Equations

Let $M^5 = C \times \Sigma \times \mathbb{R}^+_y$, where C and Σ are Riemann surfaces, equipped with a product metric g and fix the non-vanishing unit vector field $v = \partial_y$. In this situation $\ker \eta$ coincides with the tangent space of $C \times \Sigma$. Let J be the almost complex structure on $\ker \eta$ that is induced by the complex structures on C and Σ.

In the context of Haydys-Witten Floer theory, we always assume that C contains the non-compact flow direction and correspondingly is either $C \simeq \mathbb{R}_s \times \mathbb{R}_t$ or $\mathbb{R}_s \times S^1_t$ with the corresponding standard complex structure. In contrast, Σ might generally be an arbitrary Riemann surface. In the end, we will be most interested in the special case $\Sigma = \mathbb{C}$, because in that case Haydys-Witten Floer theory can be related to Khovanov homology. We let (w, z) denote holomorphic coordinates on $C \times \Sigma$ and will also write $w = s + it$ and $z = x^2 + ix^3$ in terms of real coordinates.

Consider a principal bundle E over M^5 for $G = SU(N)$ and let $G_{\mathbb{C}}$ and $E_{\mathbb{C}}$ denote their complexifications. Let $A \in \mathcal{A}(E)$ be a gauge connection and $B \in \Omega^2_{v,+}(M^5, \operatorname{ad} E)$ an element of Haydys' self-dual two-forms with respect to v (cf. Section 2.4.1). In holomorphic coordinates B is locally determined by

$$B = i\phi_1(dw \wedge d\bar{w} + dz \wedge d\bar{z}) + \varphi\, dw \wedge dz + \overline{\varphi}\, d\bar{w} \wedge d\bar{z}.$$

In real coordinates, this corresponds to $B = \sum_{a=1}^{3} \phi_a(dx^0 \wedge dx^a + \frac{1}{2}\epsilon_{abc}dx^b \wedge dx^c)$, where components are related by $\varphi = \phi_2 - i\phi_3$, $\overline{\varphi} = \phi_2 + i\phi_3$.

Introduce the following differential operators $\mathcal{D}_\mu$ that act on sections of $\operatorname{ad} E_{\mathbb{C}}$.

$$\begin{aligned} \mathcal{D}_0 &= 2\nabla^A_{\bar{w}} = \nabla^A_0 + i\nabla^A_1 & \mathcal{D}_1 &= 2\nabla^A_{\bar{z}} = \nabla^A_2 + i\nabla^A_3 \\ \mathcal{D}_2 &= \nabla^A_y - i[\phi_1, \cdot] & \mathcal{D}_3 &= [\varphi, \cdot] = [\phi_2, \cdot] - i[\phi_3, \cdot] \end{aligned} \tag{5.2}$$

The complex structure of $\operatorname{ad} E_{\mathbb{C}}$ induces a complex conjugation that we will denote by $\overline{\mathcal{D}_\mu}$. Furthermore, there is an action of $G_{\mathbb{C}}$-valued gauge transformations $g(x) \in \mathcal{G}_{\mathbb{C}}(M^5)$ by conjugation $\mathcal{D}_\mu \mapsto g(x)^{-1}\mathcal{D}_\mu g(x)$.

The Haydys-Witten equations are equivalent to

$$[\overline{\mathcal{D}_0}, \overline{\mathcal{D}_i}] - \tfrac{1}{2}\epsilon_{ijk}[\mathcal{D}_j, \mathcal{D}_k] = 0\,,\ i,j,k = 1,2,3\,, \qquad \sum_{\mu=0}^{3}[\overline{\mathcal{D}_\mu}, \mathcal{D}_\mu] = 0\,,$$

while the decoupled Haydys-Witten equations correspond to the specialization

$$[\mathcal{D}_\mu, \mathcal{D}_\nu] = 0\,,\ \mu,\nu = 0,1,2,3\,, \qquad \sum_{\mu=0}^{3}[\overline{\mathcal{D}_\mu}, \mathcal{D}_\mu] = 0\,.$$

The decoupled Haydys-Witten equations on $C \times \Sigma \times \mathbb{R}^+_y$ are an extension of the EBE on $\Sigma \times \mathbb{R}^+_y$, which are given by exactly the same equations and additionally satisfy $\mathcal{D}_0 = 0$.

Consider the submanifold $C \times \Sigma \times \{y\}$ for some fixed y. The structure we have described above becomes that of a Kähler manifold $(C \times \Sigma, \omega)$, where the Kähler form is defined by $\omega = g(J\cdot, \cdot)$, together with a complex vector bundle $\operatorname{ad} E_{\mathbb{C}}$. This vector bundle is equipped with a Hermitian metric h, with respect to which $\mathcal{A}$ is Hermitian.

We can view the four operators $\mathcal{D}_\mu$ and their complex conjugates as holomorphic and anti-holomorphic components of a *complexified* covariant derivative associated to the Hermitian connection $\mathcal{A}$ on $\operatorname{ad} E_{\mathbb{C}}$. By this we mean the $\mathbb{C}$-linear map $\nabla^{\mathcal{A}} : \Gamma(\operatorname{ad} E_{\mathbb{C}}) \to \Gamma(T^*_{\mathbb{C}}M \otimes \operatorname{ad} E_{\mathbb{C}})$ that is locally given by $\nabla^{\mathcal{A}}_{\partial_\mu} = \mathcal{D}_\mu$ and $\nabla^{\mathcal{A}}_{\overline{\partial_\mu}} = \overline{\mathcal{D}_\mu}$. Here we denote by $\overline{\partial_\mu} = \hat{J}\partial_\mu$ a local orthonormal frame of $T_{\mathbb{C}}M$ with respect to the standard complex structure $\hat{J}$. The covariant derivative has the property that it vanishes automatically in the direction of the vector field v, i.e. $\mathcal{D}_v = 0$. Let $F_{\mathcal{A}} \in \Omega^2_{\mathbb{C}}(M^5, \operatorname{ad} E_{\mathbb{C}})$ be its curvature two-form, defined by

$$(F_{\mathcal{A}})_{\mu\nu}\, s = [\mathcal{D}_\mu, \mathcal{D}_\nu]$$

Denote by $F^{p,q}_{\mathcal{A}}$ the (p,q)-part of the field strength. Since $\mathcal{D}_v = 0$, the field strength $F_{\mathcal{A}}$ is an element of the subbundle $\Omega^2_{v,-}(M^5, \operatorname{ad} E_{\mathbb{C}}) \oplus \Omega^2_{v,+}(M^5, \operatorname{ad} E_{\mathbb{C}})$, where the map $T_\eta := \star_5(\eta \wedge \cdot)$ acts on the summands with eigenvalues ± 1, respectively. The anti-self-dual anti-holomorphic part of $F_{\mathcal{A}}$ is then given by

$$\frac{1}{2}\left(F_{\mathcal{A}}^{2,0} - T_\eta \hat{J} F_{\mathcal{A}}^{0,2}\right) = \sum \left((F_{\mathcal{A}})_{0i} - \tfrac{1}{2}\epsilon_{ijk}\,(\overline{F_{\mathcal{A}}})_{jk}\right)\left(dx^0 \wedge dx^i + \tfrac{1}{2}\epsilon_{ijk}\, dx^j \wedge dx^k\right)$$

Furthermore, there is an inner product $\Lambda_\omega : \Omega^{1,1}_{\mathbb{C}}(C \times \Sigma) \to \Omega^0_{\mathbb{C}}(C \times \Sigma)$, induced by the Kähler form $\omega = g(J\cdot,\cdot)$ and normalized such that in coordinates and if the metric on $C \times \Sigma$ is flat $\omega = i/2(dx^0 \wedge d\bar{x}^0 + \ldots + dx^3 \wedge d\bar{x}^3)$. Application of Λ_ω to $F_{\mathcal{A}}$ corresponds to the trace of its mixed part $[\overline{\mathcal{D}_\mu}, \mathcal{D}_\nu]$. With this the Haydys-Witten equations become anti-holomorphic anti-self-duality equations

$$F_{\mathcal{A}}^{2,0} - T_\eta \hat{J} F_{\mathcal{A}}^{0,2} = 0\ , \quad \Lambda_\omega F_{\mathcal{A}} = 0\ ,$$

while the decoupled Haydys-Witten equations are then equivalent to the Hermitian Yang-Mills equations

$$F_{\mathcal{A}}^{2,0} = 0\ , \quad \Lambda_\omega F_{\mathcal{A}} = 0\ .$$

The first equation, $F_{\mathcal{A}}^{2,0} = 0$, is invariant under complex gauge transformations $\mathcal{G}_{\mathbb{C}}(M^5)$, while the equation $\Lambda_\omega F_{\mathcal{A}} = 0$ is only invariant under the subgroup of unitary gauge transformations with respect to the metric h. More explicitly: those gauge transformation that satisfy $\overline{g}hg = h$. Accordingly, the second equation describes a real moment map condition. This extends the Hermitian-Yang-Mills structure of the EBE.

The work of Donaldson [Don85; Don87a] and Uhlenbeck-Yau [UY86] shows that the geometric data of solutions to the $G_{\mathbb{C}}$-invariant equations play an important role in understanding the solutions of the full equations. The main underlying idea is that one can first solve the easier $\mathcal{G}_{\mathbb{C}}$-invariant equations and subsequently try to find a complex gauge connection such that the solution satisfies the remaining real moment map condition. Indeed, the model solutions for the Nahm pole boundary conditions are constructed in this way, and this is what will be discussed in the next section.

5.2 Adjoint Orbits and Slodowy Slices

This section introduces the relevant notation and certain standard constructions for the Lie algebra $\mathfrak{sl}(N,\mathbb{C})$ that will be used in the rest of this chapter. Throughout, we choose and fix a Cartan subalgebra $\mathfrak{h}$ and a Chevalley basis $\{H_i, E_i^\pm\}_{i\in\{1,\ldots,N-1\}}$ of $\mathfrak{sl}(N,\mathbb{C})$ with Cartan matrix A_{ij}. The Lie bracket satisfies

$$[H_i, H_j] = 0, \qquad [H_i, E_j^\pm] = \pm A_{ji} E_j^\pm, \qquad [E_i^+, E_j^-] = \delta_{ij} H_i\ .$$

An element $X \in \mathfrak{sl}(N,\mathbb{C})$ is called regular if the dimension of its centralizer $Z_{\mathfrak{sl}(N,\mathbb{C})}(X) = \{Y \in \mathfrak{sl}(N,\mathbb{C}) | [Y,X] = 0\}$ is minimal, i.e. it is equal to the dimension of the Cartan subalgebra $\mathfrak{h}$. An element $E \in \mathfrak{sl}(N,\mathbb{C})$ is nilpotent if there is a positive integer such that $(\mathrm{ad}_E)^n = 0$. Let $\pi_1, \ldots, \pi_s$

be a collection of positive integers that satisfy $\pi_1 + \ldots + \pi_s = N$. Denoting by $J_{\pi_i}(\lambda)$ a Jordan block of size π_i with eigenvalue λ, the Jordan normal form of a nilpotent element is given by

$$E_\pi = \begin{pmatrix} J_{\pi_1}(0) & & \\ & \ddots & \\ & & J_{\pi_s}(0) \end{pmatrix}$$

More generally, we use the notation $\pi = [{\pi_1}^{\nu_1} \ldots {\pi_s}^{\nu_s}]$ for partitions of N, where ν_i denote multiplicities such that $N = \nu_1\pi_1 + \ldots + \nu_s\pi_s$, and entries are ordered according to $\pi_1 \geq \pi_2 \geq \ldots \pi_s$. The Jordan normal form of a regular nilpotent element is associated to the partition $\pi = [N]$, while the Jordan normal form of 0 corresponds to $\pi = [1^N]$.

The orbit of an element $X \in \mathfrak{sl}(N, \mathbb{C})$ under the adjoint action of $SL(N, \mathbb{C})$ is denoted by

$$\mathcal{O}(X) = \left\{ gXg^{-1} \;\middle|\; g \in SL(N, \mathbb{C}) \right\} .$$

The adjoint orbits of nilpotent elements are classified by partitions of N. For any nilpotent element E, the associated partition π is determined by its Jordan normal form. For a given partition π, we write $\mathcal{O}_\pi = \mathcal{O}(E_\pi)$ and any orbit of nilpotent elements is of that form.

The orbit $\mathcal{O}_{[N]}$ corresponds to the adjoint orbit of regular nilpotent elements and will also be denoted by $\mathcal{O}_{\mathrm{reg}}$. The closure of the regular nilpotent orbit $\mathcal{N} := \overline{\mathcal{O}_{\mathrm{reg}}}$ is called the nilpotent cone of $\mathfrak{sl}(N, \mathbb{C})$. The nilpotent cone is an algebraic variety that contains all nilpotent orbits of $\mathfrak{sl}(N, \mathbb{C})$ as lower dimensional strata as its singular loci.

Nilpotent orbits come with a partial order, defined by setting $\mathcal{O}_\pi \leq \mathcal{O}_\rho$ if $\mathcal{O}_\pi \subseteq \overline{\mathcal{O}_\rho}$. This partial order is equivalent to the dominance order on the set of partitions of N, where $\pi \leq \rho$ if and only if $\pi_1 + \ldots + \pi_k \leq \rho_1 + \ldots + \rho_k$ for all k.

$$\pi \leq \rho \iff \mathcal{O}_\pi \leq \mathcal{O}_\rho .$$

The Jacobson-Morozov theorem states that for any non-zero nilpotent element E there exists an $\mathfrak{sl}_2$-triple (E, H, F), i.e. elements that satisfy the $\mathfrak{sl}(2, \mathbb{C})$ commutation relations $[H, E] = 2E$, $[E, F] = H$, $[H, F] = -2F$. Moreover, H and F are unique up to conjugation by elements of the centralizer $Z_{SL(N,\mathbb{C})}(E) = \{g \in SL(N, \mathbb{C}) | gEg^{-1} = E\}$.

Let π be a partition of N and fix a nilpotent element $E \in \mathcal{O}_\pi$. Choose a completion to an $\mathfrak{sl}_2$-triple (E, H, F). The affine subspace of $\mathfrak{sl}(N, \mathbb{C})$ defined by

$$\mathcal{S}_E := E + \ker \mathrm{ad}_F$$

is called the Slodowy slice to $\mathcal{O}_\pi$ at E. Slodowy slices are transversal slices in $\mathfrak{sl}(N, \mathbb{C})$. This means that that they have transverse intersections with all adjoint orbits in $\mathfrak{sl}(N, \mathbb{C})$. We will also write $\mathcal{S}_\pi$ for the Slodowy slice to $\mathcal{O}_\pi$ at our favourite element of the orbit, the Jordan normal form E_π.

5.3 The Hermitian Version of the Nahm Pole Boundary Conditions

The Nahm pole boundary conditions are modeled on certain singular solutions of Nahm's equations over $\mathbb{R}^+_y$ and, in the presence of knots, on monopole solutions of the extended Bogomolny equations (EBE) over $\mathbb{C} \times \mathbb{R}^+_y$ [Wit11a; MW14; MW17]. Both of these equations are dimensional reductions of the (decoupled) Haydys-Witten equations (5.1) (cf. Section 2.5). The EBE arise by setting $\mathcal{D}_0 = 0$, while Nahm's equations correspond to the case $\mathcal{D}_0 = \mathcal{D}_1 = 0$. Note that both equations retain the Hermitian Yang-Mills structure of the decoupled Haydys-Witten equations.

In this section we first provide a short review of the derivation of the model solutions. This was described for $SU(2)$ by Witten and later for $SU(N)$ by Mikhaylov [Wit11a; Mik12]. In Section 5.5 we extend the original Ansatz of Witten and Mikhaylov to the situation of the decoupled Haydys-Witten equations with $\mathcal{D}_0 \neq 0$. The section concludes with a definition of the Nahm pole boundary conditions of Section 2.6, viewed as a condition on a complex gauge transformation. This was described in a very similar way by He and Mazzeo in [HM19c; HM20].

For the rest of this section assume that $\Sigma = \mathbb{C}$ with holomorphic coordinate $z = x^2 + ix^3$. We will also use polar coordinates $z = re^{i\vartheta}$ on $\mathbb{C}$ and (hemi-)spherical coordinates (R, ψ, ϑ) on $\mathbb{C} \times \mathbb{R}^+_y$, where $R = r^2 + y^2$, $\cos\psi = \frac{y}{R}$, and ϑ is the azimuthal angle in the complex plane. Note that in spherical coordinates the boundary $y = 0$ corresponds to points with $\psi = \pi/2$.

The Hermitian Yang-Mills structure of the decoupled Haydys-Witten equations suggests a way to solve the equations: Following the ideas of Donaldson-Uhlenbeck-Yau [Don85; UY86; Don87a], one starts from holomorphic data that satisfies the $\mathcal{G}_\mathbb{C}$-invariant equations $[\mathcal{D}_i, \mathcal{D}_j] = 0$ and then determines a gauge transformation $g \in \mathcal{G}_\mathbb{C}(\mathbb{C} \times \mathbb{R}^+_y)$ that solves the moment map condition $\sum_{i=1}^3 [\overline{\mathcal{D}_i}, \mathcal{D}_i] = 0$ [Wit11a; Mik12; GW12]. We start with a description of Nahm pole solutions and afterwards discuss monopole solutions that additionally incorporate knot singularities.

Nahm Pole Solutions Let $E \in \mathcal{N}$ be a nilpotent element and consider as an initial Ansatz

$$A^0 = 0\,, \qquad \varphi^0 = E\,, \qquad \phi_1^0 = 0\,. \tag{5.3}$$

Using the definitions in (5.2), this field configuration clearly satisfies Nahm's equations, which are given by $[\mathcal{D}_2, \mathcal{D}_3] = 0$. Since $\varphi = \phi_2 + i\phi_3$, the complex conjugate $\overline{\varphi}^0 = \phi_2 - i\phi_3 =: F$ determines a unique $\mathfrak{sl}_2$-triple (E, H, F).

We now ask for a complex gauge transformation $g_0 \in \mathcal{G}_\mathbb{C}(\mathbb{R}^+_y)$ that maps the fields of this initial Ansatz to a solution of the real moment map $\sum_{i=2}^3 [\overline{\mathcal{D}}_i, \mathcal{D}_i] = 0$. Since the unitary part of the gauge transformation drops out, we can assume that g_0 takes values in $\exp i\mathfrak{g}$ with respect to the Cartan decomposition $G_\mathbb{C} = G \cdot \exp i\mathfrak{g}$. Writing $g_0 = \exp\psi$, assuming that $\psi \in i\mathfrak{h}$ and only depends on y, and plugging the transformed operator $g_0\mathcal{D}_i g_0^{-1}$ into the moment map equation leads to

$$\partial_y^2\psi + \frac{1}{2}[\overline{\varphi}^0, e^{2\psi}\varphi^0 e^{-2\psi}] = 0 \,. \tag{5.4}$$

This equation has a simple solution[2], given by

$$g_0 = \exp(-\log y\, H) \,. \tag{5.5}$$

The action of g_0 on $\mathcal{D}_\mu$ transforms the initial Ansatz (5.3) into the Nahm pole solution

$$A = 0 \,, \qquad \varphi = \frac{E}{y} \,, \qquad \phi_1 = \frac{H}{y} \,.$$

Observe that the choice of nilpotent element $E \in \mathcal{N}$ in the initial Ansatz (5.3) uniquely determines the Nahm pole solution. We call the solution a regular Nahm pole if $E \in \mathcal{O}_{\text{reg}}$ and in that case there always exists a constant $g \in G_{\mathbb{C}}$ such that $gEg^{-1} = E_{[N]} = \sum_{i=1}^{N-1} E_i^+$.

Monopole Solutions To include the presence of knots, we additionally want to add a monopole-like behaviour near the points p_a at which K intersects Σ. Monopoles are characterized by the fact that they exhibit a monodromy of magnetic charge $\lambda \in \Gamma^\vee_{\text{char}}$ around the origin in $\mathbb{C}$. The monodromy is carried by the behaviour of φ when moving in a circle around the origin $z = re^{i\vartheta} \mapsto re^{i(\vartheta+2\pi)}$. This is encoded in the following knot singularity Ansatz

$$A^\lambda = 0 \,, \qquad \varphi^\lambda = \sum_{i=1}^{N-1} z^{\lambda_i} E_i^+ \,, \qquad \phi_1^\lambda = 0 \,. \tag{5.6}$$

This provides an initial solution of the $\mathcal{G}_{\mathbb{C}}$-invariant part of the EBE, given by $[\mathcal{D}_i, \mathcal{D}_j] = 0$, $i = 1, 2, 3$.

In this Ansatz φ^λ is an element of $\mathcal{O}_{\text{reg}}$ for all $z \neq 0$, but exactly at $z = 0$ it is an element of some subordinate orbit $\mathcal{O}_\pi$. The partition π is given by the Jordan blocks of $\varphi^\lambda|_{z=0}$. It can be determined from the weight λ by moving through the entries of $\lambda = (\lambda_1, \ldots, \lambda_{N-1})$, counting the number of consecutive λ_i with value 0, and shifting the resulting counts by +1. For example, if the knot is labeled by a co-character that corresponds to either the fundamental or anti-fundamental representation, one finds:

$$\lambda = (\underbrace{1}_{0}, \underbrace{0, \ldots, 0}_{N-2}) \quad \text{or} \quad (\underbrace{0, \ldots, 0}_{N-2}, \underbrace{1}_{0}) \quad \rightsquigarrow \quad \pi = [(N-1)\; 1] \,.$$

In that case $\varphi^\lambda|_{z=0}$ is an element of the (unique) subregular nilpotent orbit $\mathcal{O}_{\text{subreg}} = \mathcal{O}_{[N-1\;1]}$. Note that the same is true when the knot is labeled by any symmetric or anti-symmetric representation of $\mathfrak{sl}(N, \mathbb{C})$, since these correspond to the co-characters $\lambda = (n, 0, \ldots, 0)$ and $(0, 0, \ldots, n)$. Higher (anti-)symmetric representations can be distinguished from the fundamental ones because the associated monodromies have higher winding number.

[2] Use Lie's expansion formula (also attributed to Campbell and Hadamard): $e^X Y e^{-X} = \sum [X, Y]_k / k!$, where $[X, Y]_k = [X, [X, Y]_{k-1}]$ and $[X, Y]_0 = Y$.

Let us note for later, that the Ansatz in (5.6) is of the form $\varphi^\lambda = E + K(z,\lambda)$ for some basepoint $E \in \mathcal{O}_\pi$ together with a choice of $\mathfrak{sl}_2$-completion (E,H,F), and where $K(z,\lambda) \in \ker F \cap \mathcal{N}$ is a holomorphic function that vanishes at $z = 0$. Put differently, φ^λ is a map from $\mathbb{C}$ to the Slodowy slice $S_E \cap \mathcal{N}$ that sends $z = 0$ to the basepoint E and monodromy prescribed by λ.

Consider now a complex gauge transformation $g_\lambda \in \mathcal{G}_{\mathbb{C}}(\mathbb{C} \times \mathbb{R}^+_y)$ and assume it is of the form $g_\lambda = \exp\psi$ for some $\psi \in i\mathfrak{h}$. The moment map condition becomes

$$(\Delta_{z,\bar z} + \partial_y^2)\,\psi + \frac{1}{2}[\overline{\varphi}^\lambda, e^{2\psi}\varphi^\lambda e^{-2\psi}] = 0 \tag{5.7}$$

Here $\Delta_{z,\bar z} = 4\partial_z\partial_{\bar z}$ denotes the Laplacian on $\mathbb{C}$. Note that this is simply the three-dimensional version of (5.4). Since any solution of (5.7) gives rise to a solution of the extended Bogomolny equations, we will abbreviate this equation by **EBE** $(\psi) = 0$.

Mikhaylov proved that, for $G = SU(N)$ and any weight $\lambda \in \Gamma^\vee_{\mathrm{char}}$, there exists a unique solution g_λ that is compatible with the Nahm pole solution at boundary points away from $z = 0$. The explicit formulae are unfortunately somewhat unwieldy; we refer to [Mik12] for a detailed description of the general case.

For $G = SU(2)$, the Cartan subalgebra is spanned by a single element H and λ is a single non-negative half-integer. In that case g_λ is comparatively simple. In spherical coordinates (R,ψ,ϑ) on $\mathbb{C}\times\mathbb{R}^+_y$ and using $s = \sin(\pi/2 - \psi)$ as boundary defining function on $\mathbb{C}\setminus\{0\}\times\mathbb{R}^+_y$, g_λ is given by

$$g_\lambda = \exp\left(\log\frac{\lambda+1}{R^{\lambda+1}s s_\lambda}\,H\right). \tag{5.8}$$

Here we used the abbreviation $s_\lambda = \sum_{k=0}^{\lambda}(1+s)^{n-k}(1-s)^k = \frac{(1+s)^{\lambda+1}-(1-s)^{\lambda+1}}{2s}$. The action of g_λ on $\mathcal{D}_\mu$ yields the field configuration:

$$\begin{aligned}
A_\vartheta &= -(\lambda+1)\cos^2\psi\frac{(1+\cos\psi)^\lambda-(1-\cos\psi)^\lambda}{(1+\cos\psi)^{\lambda+1}-(1-\cos\psi)^{\lambda+1}}\,H\,,\\
\phi_1 &= -\frac{\lambda+1}{R}\frac{(1+\cos\psi)^{\lambda+1}+(1-\cos\psi)^{\lambda+1}}{(1+\cos\psi)^{\lambda+1}-(1-\cos\psi)^{\lambda+1}}\,H\,,\\
\varphi &= \frac{\lambda+1}{R}\frac{\sin^\lambda\psi\ \exp(i\lambda\vartheta)}{(1+\cos\psi)^{\lambda+1}-(1-\cos\psi)^{\lambda+1}}\,E\,,\\
A_s &= A_t = A_R = A_\psi = 0\,.
\end{aligned}$$

These are exactly the knot singularity model solutions of the EBE that were already described in Section 2.6.

The Nahm Pole Boundary Conditions We are now ready to provide a definition of the Nahm pole boundary conditions that is convenient for the discussions in this chapter.

Definition 5.1 (Nahm Pole Boundary Condition)**.** Consider three-manifolds of the form $\Sigma \times \mathbb{R}^+_y$. Let g_0 and g_λ be the singular gauge transformations given in equations (5.5) and (5.8), respectively. The fields (A, φ, ϕ_1) satisfy regular Nahm pole boundary conditions with knot singularities at $D = \bigsqcup \{p_a\} \subset \Sigma$ if

- in a neighbourhood of a boundary point $(p, 0) \in (\Sigma \setminus D) \times \mathbb{R}^+_y$ away from knot insertions, there exists a $G_{\mathbb{C}}$-valued gauge transformation of the form $g = g_0 e^u$ with $|u| + |ydu| < Cy^\epsilon$, such that

$$(A, \varphi, \phi_1) = g \cdot (0, \sum E_i^+, 0)$$

- in a neighbourhood of a knot insertion $(p_a, 0) \in \Sigma \times \mathbb{R}^+_y$ of weight λ, there exists a gauge transformation of the form $g = g_\lambda e^u$ with $|u| + |Rs\, du| \leq CR^\epsilon s^\epsilon$, such that

$$(A, \varphi, \phi_1) = g \cdot (0, \sum z^{\lambda_i} E_i^+, 0)$$

Up to an inconsequential reinterpretation of field components, this definition lifts to Nahm pole boundary conditions for Kapustin-Witten fields (A, ϕ) on four-manifolds $X^3 \times \mathbb{R}^+_y$ and Haydys-Witten fields (A, B) on five-manifolds $W^4 \times \mathbb{R}^+_y$. In higher dimensions, knot singularities are supported along a knot $K \subset X^3 \times \{0\}$ or a surface $\Sigma_K \subset W^4 \times \{0\}$, respectively.

5.4 EBE-Solutions and Higgs bundles

The Hermitian Yang-Mills structure of the extended Bogomolny equations on $\Sigma \times \mathbb{R}^+_y$ suggests that there is a deep relation between full solutions and the simpler holomorphic data that underlies the initial knot singularity Ansatz (5.6). Indeed, there is a Kobayashi-Hitchin correspondence between solutions of the extended Bogmolny equations on $\Sigma \times \mathbb{R}^+_y$ and Higgs bundle data on Σ. This correspondence was originally proposed by Gaiotto and Witten in [GW12] and has since been proven by He and Mazzeo [HM19c; HM20] (also see [HM19b; Dim22b; Sun23] for variations of correspondence). Here we repeat the relevant definitions, review parts of the proof that will be of particular relevance to us, and use the opportunity to fix some notation.

We are interested in solutions of the EBE on $\Sigma \times \mathbb{R}^+_y$ that, on the one hand, satisfy Nahm pole boundary conditions with knot singularities at the boundary, and for which, on the other hand, $A + i\varphi$ approaches a flat $G_{\mathbb{C}}$ connection as $y \to \infty$. Let $D = \{(p_a, \lambda_a)\}_{a=1,\ldots,k}$ denote a collection of points $p_a \in \Sigma$ that are decorated with weights $\lambda_a \in \Gamma^\vee_{\mathrm{char}}$ in the co-character lattice of $\mathfrak{g}$. The moduli space of EBE-solutions with knot singularity data D will be denoted by:

Definition 5.2 (Moduli Space of EBE Solutions)**.**

$$\widehat{\mathcal{M}}_D^{\mathrm{EBE}} := \Big\{ \mathbf{EBE}\ (A, \varphi, \phi_1) = 0 \;\Big|\; (A, \varphi, \phi_1) \text{ satisfies Nahm-pole boundary conditions as } y \to 0, \\ \text{has knot singularities at } D, \text{ and approaches an irreducible flat } SL(N, \mathbb{C}) \text{ connection as } y \to \infty \Big\}$$

The absence of a knot will be denoted by $D = \emptyset$, in which case $\mathcal{M}_{\emptyset}^{\mathrm{EBE}}$ is the moduli space of pure Nahm pole solutions.

Let $\mathcal{G}_0(\Sigma \times \mathbb{R}_y^+)$ be the subset of (real) gauge transformations that vanish at the boundary. We also define the moduli space given by the quotient

$$\mathcal{M}_D^{\mathrm{EBE}} = \widehat{\mathcal{M}}_D^{\mathrm{EBE}} \big/ \mathcal{G}_0(\Sigma \times \mathbb{R}_y^+)\,.$$

Remark. We only mod out $\mathcal{G}_0(\Sigma \times \mathbb{R}_y^+)$ because the configuration of the fields at the boundary is in principle a physical observable; relatedly, from the perspective of mathematics, the boundary configuration is part of the input of the variational principle.

We have seen in the previous section, for $\Sigma = \mathbb{C}$, that once the data of an initial solution of $[\mathcal{D}_i, \mathcal{D}_j] = 0$ is fixed, the remaining real moment map equation determines a unique complex gauge transformation, g_0 or g_λ, that transforms the Ansatz into a proper solution of the extended Bogomolny equations. The Kobayashi-Hitchin correspondence states that this generalizes to arbitrary Riemann surfaces Σ, where the initial solution is replaced by certain holomorphic data over Σ. In the following, we introduce the relevant geometric objects over Riemann surfaces Σ.

Let $(\mathcal{E}, \bar\partial_{\mathcal{E}})$ be a holomorphic vector bundle over Σ. Denote the canonical bundle[3] over Σ by K and write $\Omega^{1,0}(\mathrm{End}\,\mathcal{E}) = H^0(\mathrm{End}\,\mathcal{E} \otimes K)$ for the space of holomorphic one-forms with values in the endomorphism bundle of $\mathcal{E}$.

Definition 5.3 (Higgs Bundle)**.** A Higgs bundle $(\mathcal{E}, \varphi)$ is a holomorphic vector bundle $(\mathcal{E}, \bar\partial_{\mathcal{E}})$ over Σ, together with a holomorphic one-form $\varphi \in \Omega^{1,0}(\mathrm{End}\,\mathcal{E})$ called the Higgs field.

We will always assume that $\det \mathcal{E} = \mathcal{O}_\Sigma$, the sheaf of holomorphic functions on Σ, and that $\deg \mathcal{E} = 0$. In that situation, a Higgs bundle is called stable if $\deg V < 0$ for any holomorphic subbundle V that satisfies $\varphi(V) \subset V \otimes K$ and polystable if it is a direct sum of stable Higgs bundles.

We denote the moduli space of Higgs bundles by $\widehat{\mathcal{M}}_{\mathrm{Higgs}}$ and as above write $\mathcal{M}_{\mathrm{Higgs}}$ for its quotient by $SL(N, \mathbb{C})$-valued gauge transformations $\mathcal{G}_{\mathbb{C}}(\Sigma)$. By a famous result of Hitchin, for any Higgs bundle there exists an irreducible solution of the Hitchin equations if and only if it is stable (and a reducible solution if and only if polystable) [Hit87b]. Instead of introducing the Hitchin equations, we use that any solution of the Hitchin equations is associated to a flat $SL(N, \mathbb{C})$ connection, which are classified by $\rho : \pi_1(\Sigma) \to SL(N, \mathbb{C})$. A series of articles by Hitchin, Donaldson, Simpson, and Corlette famously culminated in a proof that there is a diffeomorphic equivalence between the moduli spaces of stable Higgs bundles, irreducible solutions of the Hitchin equations, and irreducible flat connections [Don87b; Cor88; Hit87a; Sim90; Sim92].

[3] This is the dual of the complex vector bundle $T\Sigma$, where fibers are given by complex linear maps on the tangent space of Σ.

In the study of Higgs bundles an important role is played by the Hitchin fibration. It is locally built from the adjoint quotient map, which sends a matrix in $\mathfrak{sl}(N,\mathbb{C})$ to its generalized eigenvalues:

$$\chi : \mathfrak{sl}(N,\mathbb{C}) \to \mathfrak{h}/\mathcal{W} , \ A \mapsto (c_2(A), \dots , c_N(A)) ,$$

where $\mathfrak{h}$ is the Cartan subalgebra, $\mathcal{W}$ the Weyl group, and $c_j(A)$ denote the invariant polynomials of $\mathfrak{sl}(N,\mathbb{C})$ as determined by $\det(\lambda 1 - A) = \sum \lambda^{N-j}(-1)^j c_j(A)$. The Hitchin fibration lifts this to the level of Higgs bundles [Hit87b].

$$\begin{aligned} \mathcal{M}_{\text{Higgs}} &\to \bigoplus_{i=1}^{n} H^0(\Sigma, K^{i+1}) \\ (\mathcal{E},\varphi) &\mapsto (c_2(\varphi), \dots , c_N(\varphi)) \end{aligned}$$

The Hitchin component (or Hitchin section) $\mathcal{M}_{\text{Hit}}$ is the section of this fibration that associates to the invariants $(q_2, \dots q_N)$ the gauge orbit of the following Higgs bundle

$$\mathcal{E} = K^{-\frac{N-1}{2}} \oplus K^{-\frac{N-1}{2}+1} \oplus \dots \oplus K^{\frac{N-1}{2}} , \qquad \varphi = \begin{pmatrix} 0 & * & 0 & \\ & 0 & \ddots & 0 \\ & & \ddots & * \\ q_N & \dots & q_2 & 0 \end{pmatrix}$$

Given a Higgs bundle $(\mathcal{E},\varphi)$ and a holomorphic line subbundle $L \subset \mathcal{E}$, there is an associated divisor $\mathfrak{d} = \mathfrak{d}(L,\varphi)$ that encodes the linear dependencies between the 'subbundles' L, $\varphi(L)$, ..., $\varphi^n(L)$ of $\mathcal{E}$. To make this precise, introduce the following maps, which capture these linear dependencies by virtue of being antisymmetric

$$f_i := 1 \wedge \varphi \wedge \dots \wedge \varphi^i : L^i \to K^{\frac{i(i+1)}{2}} \otimes L^{-(i+1)} \otimes \wedge^{i+1}\mathcal{E} .$$

The divisor associated to (L,φ) is constructed from the zeroes of f_i and their vanishing orders.

$$\mathfrak{d}(L,\varphi) := \sum_{\substack{i=1,\dots,n \\ p\in\Sigma}} \operatorname{ord}_p f_i \cdot p$$

Following [HM19c; HM20], we call $\mathfrak{d}(L,\varphi)$ effective if at each point p the tuple

$$\lambda = \big(\operatorname{ord}_p f_1 , \operatorname{ord}_p f_1 - \operatorname{ord}_p f_2 , \dots , \operatorname{ord}_p f_{n-1} - \operatorname{ord}_p f_n \big)$$

has only non-negative entries. In that case we also write $\mathfrak{d}(L,\varphi) = \{(p_a, \lambda_a)\}_{a=1,\dots,k}$.

Definition 5.4 (Effective Triples). An effective triple $(\mathcal{E},\varphi,L)$ consists of a stable Higgs bundle $(\mathcal{E},\varphi)$ together with a holomorphic line bundle L such that the divisor $\mathfrak{d}(L,\varphi)$ is effective.

The moduli space of effective triples will be denoted $\widehat{\mathcal{M}}_{(\mathcal{E},\varphi,L)}$ and we also write $\mathcal{M}_{(\mathcal{E},\varphi,L)}$ for its quotient by $\mathcal{G}_{\mathbb{C}}(\Sigma)$.

We have now collected all ingredients to state the Kobayashi-Hitchin correspondence for solutions of the extended Bogomolny equations on $\Sigma \times \mathbb{R}^+_y$.

Theorem 5.5 ([HM19c; HM20]). *There are bijections*

$$\mathcal{M}^{\mathrm{EBE}}_{\emptyset} \xrightarrow{I_{NP}} \mathcal{M}_{\mathrm{Hit}}\,, \qquad \mathcal{M}^{\mathrm{EBE}}_{D} \xrightarrow{I_{NPK}} \mathcal{M}_{(\mathcal{E},\varphi,L)}$$

For recent variants of this result on $\Sigma = \mathbb{C}$ we also refer to [Dim22b], who elaborated on the situation for nilpotent Higgs fields, as well as [Sun23], for the case that ϕ_1 approaches a non-zero value at $y \to \infty$, a situation that is also known as 'real symmetry breaking' [GW12]. In the remainder of this section we provide a brief review of the construction of I_{NP} and I_{NPK}.

As a start, one observes that any Nahm pole solution of the EBE equations determines an effective triple. This was originally explained by Gaiotto and Witten and can be seen as follows. For the moment, write V for the associated vector bundle of the N-dimensional fundamental representation of $SL(N,\mathbb{C})$ (the gauge group of the complexified principal bundle $E_{\mathbb{C}}$). Denote by V_y the restriction of V to $\Sigma \times \{y\}$. Observe that $\mathcal{D}_1$ provides a $\bar\partial$ operator on V_y that satisfies $\bar\partial^2 = 0$. By the Newlander-Nirenberg theorem, this makes V_y into a holomorphic vector bundle that we will denote by $\mathcal{E}_y$. Next, $\mathcal{D}_3 = \mathrm{ad}_\varphi$ can be interpreted as a K_Σ-valued endomorphism of $\mathcal{E}_y$ and, moreover, $[\mathcal{D}_1, \mathcal{D}_3] = 0$ implies that $\mathcal{D}_1\varphi = 0$, so this endomorphism is holomorphic. Put differently, if we let φ_y be the restriction of φ to $(\mathrm{ad}\, E_{\mathbb{C}})_y$, we obtain a family of Higgs bundles $(\mathcal{E}_y, \varphi_y)$ over Σ. Finally, $\mathcal{D}_2 = \nabla^A_y - i[\phi_1, \cdot]$ provides a notion of parallel transport in the y-direction of $\Sigma \times \mathbb{R}^+_y$. The equations $[\mathcal{D}_1, \mathcal{D}_2] = [\mathcal{D}_2, \mathcal{D}_3] = 0$ then imply that the family of Higgs bundles is parallel with respect to $\mathcal{D}_2$. The boundary conditions at $y \to 0$ and $y \to \infty$ provide two additional points of data.

On the one hand, the asymptotic boundary condition at $y \to \infty$, namely that (A, φ, ϕ_1) converges to an irreducible flat $SL(N,\mathbb{C})$ connection, is equivalent to the statement that the one-parameter family $(\mathcal{E}_y, \varphi_y)$ consists of stable Higgs bundles. Since the one-parameter family is parallel with respect to $\mathcal{D}_2$, it is then fully determined by specifying the limiting stable Higgs bundle $(\mathcal{E}_\infty, \varphi_\infty)$.

On the other hand, the Nahm pole boundary condition at $y = 0$ with knot singularity data $D = \{(p_a, \lambda_a)\}_{a=1,\dots,k}$ determines a distinguished line bundle $L \to \Sigma\times\mathbb{R}^+_y$ as follows. Let $\{U_a\}_{a=0,\dots k}$ be a collection of open disks, with U_0 an open set that does not contain any p_a and the remaining U_a centered at p_a, respectively, and such that $\bigcup_{a=0}^k U_a = \Sigma$. Use spherical coordinates $(R_a, \vartheta_a, \psi_a)$ with boundary defining function $s = \sin(\pi/2 - \psi)$ on $U_a \times \mathbb{R}^+$.

$$\begin{aligned} L := \Big\{ u \in \Gamma(\Sigma \times \mathbb{R}^+_y, \mathrm{ad}\, E_{\mathbb{C}}) \;\Big|\; \mathcal{D}_2 u = 0\,,\ \lim_{y\to 0} |y^{-(N-1)/2+\epsilon} u| = 0 \text{ on } U_0, \\ \text{and } \lim_{s\to 0} |\psi_a^{-(N-1)/2+\epsilon} u| = 0 \text{ on } U_a,\ a = 1, \dots, k,\ \epsilon > 0 \Big\}\,. \end{aligned}$$

This is a line bundle by definition of the Nahm pole boundary condition. To see this note that $\mathcal{D}_2 u = \partial_y u - i\phi_1 u = 0$. Away from knot insertions $\phi_1 = H/y$ has eigenvalues $(N-1)/2y$, $(N-2)/2y$, ..., $-(N-1)/2y$, such that there is only a single component of u whose parallel transport vanishes at the maximal possible rate $y^{(N-1)/2}$. At a knot insertion p_a, the same argument with the version of ϕ_1 for a monopole solution shows that the maximal vanishing rate is $\psi_a^{-(N-1)/2}$. L is commonly called the vanishing line bundle. For each $y \in \mathbb{R}^+_y$ it is a holomorphic subbundle of $\mathcal{E}_y$.

In summary, the vanishing line bundle over $\Sigma \times \mathbb{R}^+$ determines a unique holomorphic line subbundle of $\mathcal{E}$. Moreover, the divisor $\mathfrak{d}(L, \varphi)$ is effective and at each point p the zeroes of f_i are exactly such that their vanishing orders are related to the co-character by $\lambda = (\mathrm{ord}_p(f_1), \mathrm{ord}_p(f_2) - \mathrm{ord}_p(f_1), \dots)$. In absence of knot singularities the divisor is empty and one can show that the Higgs bundle is an element of the Hitchin section.

Proving that each effective triple $(\mathcal{E}, \varphi, L)$ gives rise to a solution of the extended Bogomolny equations, involves more work. This part of the theorem is due to He and Mazzeo and we refer to [HM20, Sec. 7] for more details. Since we later find that similar arguments should apply to solutions of the decoupled Haydys-Witten equations, we summarize the basic approach while omitting the necessary analytic prerequisites.

Assume we are given an effective triple $(\mathcal{E}, \varphi, L)$. From the discussion above it's clear that the effective divisor $\mathfrak{d}(L, \varphi)$ determines the position and charges $\{(p_a, \lambda_a)\}_{a=1,\dots,k}$ of knot singularities. The main insight is that the effective triple provides enough information to construct a field configuration (A, φ, ϕ_1) on $\Sigma \times \mathbb{R}^+_y$ that satisfies the Nahm pole boundary conditions with knot singularities and is a solution of the EBE at leading order in y^{-1}. This approximate solution can subsequently be improved, order by order in y, to a unique, proper solution of the EBE. For the construction we choose an open cover of Σ, consisting of an open set U_0 that does not contain any of the points p_a and non-intersecting open disks U_a centered at p_a. This is used to construct field configurations on each of the U_j independently, which are then glued over U_0 to an approximate solution on Σ. The construction proceeds in four steps.

First, dropping the index a for the moment, we restrict to a small disk U centered at p and extract from $(\mathcal{E}, \varphi, L)$ a field configuration on $U \times \mathbb{R}^+_y$ that looks like the knot singularity Ansatz $(0, \varphi = \sum_i z^{\lambda_a} E_i^+, 0)$. Let z be a holomorphic coordinate on U and use spherical coordinates (R, ψ, ϑ) on $U \times \mathbb{R}^+_y$. Write $\varphi = \varphi_z dz$ and set $L_1 := L$ and $L_{i+1} := \varphi_z(L_i)$. Choose a non-vanishing section e_1 of L and extend it to a local holomorphic frame $\{e_1, \hat{e}_2, \dots \hat{e}_{n+1}\}$. Let λ_1 be the order of vanishing of $1 \wedge \varphi_z : L^2 \to \wedge^2 \mathcal{E}$. We can now write $\varphi_z(e_1) = f e_1 + z^{\lambda_1} \sum_{i=2}^{n+1} c_i \hat{e}_i$, where at least one of the c_i is non-vanishing at $z = 0$, because the map $1 \wedge \varphi_z \wedge \dots \wedge \varphi_z^n : L_1 \wedge L_2 \wedge \dots L_{n+1} \to \det \mathcal{E}$ fails to be an isomorphism exactly at $z = 0$. Setting $e_2 := \sum_{i=2}^{n+1} c_i \hat{e}_i$, we have arranged that $\varphi_z(e_1) = f e_1 + z^{\lambda_1} e_2$. Next, let λ_2 be the order of vanishing of $1 \wedge \varphi_z \wedge \varphi_z^2$, such that $\varphi_z(e_2) = f_1 e_1 + f_2 e_2 + z^{\lambda_2} \sum_{i=3}^{n+1} c_i \hat{e}_i$, and proceed by induction. We obtain a frame $\{e_1, \dots, e_{n+1}\}$ for which $\varphi_z(e_i) = z^{\lambda_a} e_{i+1} + \mathrm{span}\{e_1, \dots e_i\}$ Equivalently, when viewed as an $\mathrm{End}\,\mathcal{E}$-valued function, $\varphi_z = \sum z^{\lambda_a} E_i^+$, as required.

In the second step, we act on $(0, \sum z^{\lambda_i} E_i^+, 0)$ with the singular gauge transformation g_λ of Equation 5.8. The result is of order $\mathcal{O}(R^{-1} s^{-1})$, as is required for a field configuration that satisfies Nahm pole boundary conditions with knot singularity at p of weight λ. Note, in particular, that when $R \neq 0$ and $s \to 0$, the gauge transformation is of order $(Rs)^{-1} = y^{-1}$ and is asymptotically equivalent to g_0.

The third step constructs a global gauge transformation g over $\Sigma\times\mathbb{R}^+_y$ from the collection of local frames and gauge transformations $\{g_{\lambda_a}\}$ for each $U_a\times\mathbb{R}^+_y$. For this, choose a holomorphic frame on U_0 and denote by g_0 the gauge transformation defined in equation (5.5). On the overlaps $U_0\cap U_a$, the holomorphic frames are related by an explicit transition function, given by

$$g_{0a} = \exp\big(-\log r\sum_{i=1}^{n}\lambda_{a,j}A^{-1}_{ij}H_i\big)$$

We find that in the limit $s_a\to 0$ each g_{λ_a} is equivalent to g_0e^u with $|u|+|ydu| < Cy^\epsilon$. Gluing the various gauge transformations with help of a partition of unity produces a global gauge transformation g over Σ. By construction, g determines a field configuration that satisfies Nahm pole boundary conditions with knot singularities along K.

In the final step, the approximate solution is improved to a proper solution of the equations. For this, one first improves the approximation near knot singularities for each $U_a\times\mathbb{R}^+_y$ by solving the moment map equation to desired precision, order by order in R. Determining in that way the higher order terms of g_{λ_a}, and if necessary using Borel resummation to find a function that has the given expansion, one obtains a gauge transformation g on all of $\Sigma\times\mathbb{R}^+_y$ for which the solution near p_a vanishes to all orders as $R_a\to 0$. Carrying out an analogous procedure for the higher orders of g in an expansion in y yields a unique solution of the extended Bogomolny equations.

5.5 The Isotopy Ansatz

In this section we return to investigate the decoupled Haydys-Witten equations over $M^5 = C\times\Sigma\times\mathbb{R}^+_y$. In contrast to the analysis of the extended Bogomolny equations reviewed in the preceding sections, we now investigate situations in which $\mathcal{D}_0$ is non-zero. For the time being, we assume that $C=\mathbb{R}_s\times S^1_t$ and restrict ourselves to the investigation of $\mathbb{R}_s$-invariant solutions of the decoupled Haydys-Witten equations in temporal gauge $A_s = 0$. This corresponds to a dimensional reduction of the decoupled Haydys-Witten equations for which $\mathcal{D}_0 = iD_t$. The result are differential equations over the four-manifold $S^1_t\times\Sigma\times\mathbb{R}^+_y$ that correspond to a decoupled version of the Kapustin-Witten equations.

We start here with an analogue of the considerations in Section 5.3, where we described the Nahm pole and monopole model solutions. Let $\Sigma=\mathbb{C}$ with complex coordinate $z = x^2+ix^3$. Consider a single-stranded knot K that extends along S^1_t, and view it as the image of

$$\begin{aligned}\beta\,:\,S^1_t &\to S^1_t\times\mathbb{C}\times\mathbb{R}^+\\ t&\mapsto (t,z_0(t),0)\end{aligned}$$

There always exists an isotopy $\beta_\bullet$ that connects $\beta_0(t)=(t,0,0)$, the S^1-invariant knot centered at the origin in $\mathbb{C}$, to $\beta_1(t)=\beta(t)$. To be specific, let us choose the isotopy

$$\begin{aligned}\beta_\bullet\,:[0,1]_q\times S^1_t &\to S^1_t\times\mathbb{C}\times\mathbb{R}^+_y\\ (q,t)&\mapsto (t,qz_0(t),0)\end{aligned}$$

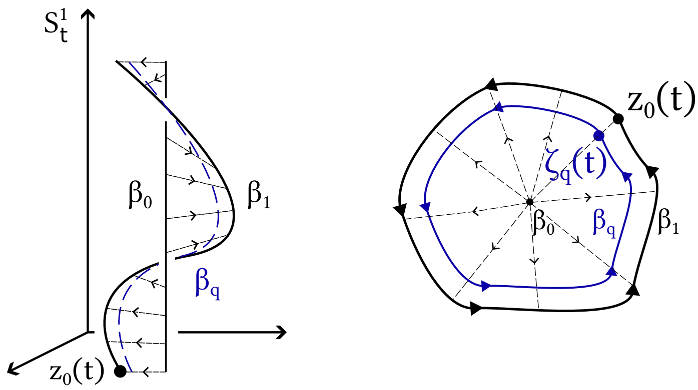

Figure 5.2 Isotopy $\beta_\bullet$ interpolating between the S^1_t-invariant strand β_0 centered at the origin of $\mathbb{C}$ and a non-trivial single stranded knot β_1. The isotopy describes a homotopy of trajectories $\zeta_q(t)$ that interpolate from the constant to the original trajectory $z_0(t)$.

For fixed $q \in [0,1]$, introduce the comoving holomorphic coordinate $\zeta_q(t) = z - qz_0(t) \in \mathbb{C}$. The knot defined by β_q lies at the origin $\{\zeta_q = 0\}$ of the complex plane parametrized by ζ_q. Let us also introduce polar coordinates $\zeta_q = r_q e^{i\vartheta_q}$. We often drop the subscript q from ζ_q, r_q and ϑ_q and only highlight the dependence on $q \in [0,1]$ where necessary.

Our starting point is the following generalization of the Ansatz $(A^\lambda, \varphi^\lambda, \phi_1^\lambda)$ in equation (5.6) that was the starting point in Section 5.3 to find the monopole model solutions:

$$A_t^\beta = \frac{q\dot{z}_0}{\zeta} \sum_{i,j}^{N-1} \lambda_i A_{ij}^{-1} H_j \,, \qquad A_z^\beta = A_y^\beta = 0 \,, \qquad \varphi^\beta = \sum_{i=1}^{N-1} \zeta^{\lambda_i} E_i^+ \,, \qquad \phi_1^\beta = 0 \,. \tag{5.9}$$

For each $q \in [0,1]$, this provides an initial solution of the equations $[\mathcal{D}_\mu, \mathcal{D}_\nu] = 0$. We call $(A^\beta, \varphi^\beta, \phi_1^\beta)$ the isotopy Ansatz for the decoupled Kapustin-Witten equations.

The singular behaviour of φ^β at $\zeta = 0$ exactly encodes the presence of a single stranded 't Hooft operator of charge $\lambda \in \Gamma^\vee_{\text{char}}$ along $\beta_q = (t, qz_0(t))$. In particular, for any fixed $t \in S^1_t$, the terms at order q^0 are equivalent to the initial Ansatz $(A^\lambda, \varphi^\lambda, \phi^\lambda)$ of (5.6) with knot singularity at $z = qz_0(t)$. Corrections due to the t-dependence of the knot singularity arise only at order $\mathcal{O}(q^1)$. Note, in particular, that for $q = 0$ the Ansatz genuinely coincides with $(0, \sum_i z^{\lambda_i} E_i^+, 0) = (A^\lambda, \varphi^\lambda, \phi^\lambda)$, where the knot is S^1-invariant and located at $z = 0$.

We propose that, by a continuity argument along the isotopy parameter $q \in [0,1]$, one can find a complex gauge transformation $g_\beta \in \mathcal{G}_\mathbb{C}(S^1_t \times \mathbb{C} \times \mathbb{R}^+_y)$ that, on the one hand, solves the moment map condition $\sum[\overline{\mathcal{D}}_\mu, \mathcal{D}_\mu] = 0$ and, on the other hand, encodes Nahm pole boundary conditions. Hence, given $(A^\beta, \varphi^\beta, \phi_1^\beta)$, apply a complex gauge transformation $g_\beta = \exp\psi$ with $\psi \in i\mathfrak{h}$. The moment map condition becomes

$$\mathbf{dKW}(\psi) = (\partial_t^2 + \Delta_{z,\bar{z}} + \partial_y^2)\,\psi + \tfrac{1}{2}[\overline{\varphi}^\beta, e^{2\psi}\varphi^\beta e^{-2\psi}] = 0 \,, \tag{5.10}$$

where we have used that $A_z^\beta = A_y^\beta = \phi_1^\beta = 0$ and also that A_t^β and $\psi \in \mathfrak{h}$, such that

$$[\overline{A_t^\beta}, e^{2\psi} A_t^\beta e^{-2\psi}] = 0 \ .$$

Note that (5.10) is just the four-dimensional version of the one-dimensional equation (5.4) and the three-dimensional equation **EBE** $(\psi) = 0$ in (5.7).

To keep notation at a minimum we restrict the following discussion to $G = SU(2)$, though the general case with $G = SU(N)$ is not much different. Assume $\psi = \psi(t, z, y)H$, $\varphi^\beta = \zeta^\lambda E$, and $\overline{\varphi} = \bar{\zeta}^\lambda F$, where (E, H, F) is the standard basis of $\mathfrak{sl}(2, \mathbb{C})$. We can then bring (5.10) into the slightly more explicit form

$$\mathbf{dKW}_q(\psi) = (\partial_t^2 + \frac{1}{2}\Delta_{z,\bar{z}} + \partial_y^2)\psi - r_q^{2\lambda} \exp(2\psi) = 0 \ . \tag{5.11}$$

Drawing inspiration from Gaiotto-Witten's adiabatic approach, we now restrict to functions $\psi(\zeta, y; z_0(t))$ that depend on t only through $z_0(t)$ and its appearance in the comoving coordinate $\zeta = z - qz_0(t)$. The operator $\mathbf{dKW}_\bullet$ describes a homotopy of differential operators and associated boundary conditions. It interpolates between the operator $\mathbf{dKW}_0 = \mathbf{EBE}$ together with an S^1-invariant knot singularity along $K = (t, 0, 0)$ on the one hand, and the decoupled Kapustin-Witten equations $\mathbf{dKW}_{q=1}$ with a knot singularity along $K = (t, z_0(t), 0)$ on the other.

Given the assumption that ψ depends adiabatically on t, we can replace $\Delta_{z,\bar{z}} = \Delta_{\zeta,\bar{\zeta}}$ and split the differentiation with respect to t into its contributions from the comoving coordinate ζ and explicit appearances of $z_0(t)$:

$$\partial_t^2 = \tilde{\partial}_t^2 - q\,(\ddot{z}_0\partial_\zeta + \ddot{\bar{z}}_0\partial_{\bar{\zeta}}) - 2q\,(\dot{z}_0\partial_\zeta + \dot{\bar{z}}_0\partial_{\bar{\zeta}})\,\tilde{\partial}_t + q^2\,(\dot{z}_0\partial_\zeta + \dot{\bar{z}}_0\partial_{\bar{\zeta}})^2 \ .$$

The notation $\tilde{\partial}_t$ on the right hand side shall reflect that this derivative only acts on $z_0(t)$ (and its derivatives).

Observe that equation (5.10) naturally organizes into powers of q. Accordingly, we make the formal Ansatz $\psi = \sum_{n\geq 0} q^n \psi^{(n)}$. Plugging this into (5.10) and expanding the exponential function in powers of q leads to:

$$\mathbf{dKW}_q(\psi) = \mathbf{EBE}_q(\psi^{(0)}) + \sum_{n\geq 1} q^n\, \mathbf{dKW}_q^{(n)}\big(\psi^{(n)};\ \psi^{(0)}, \dots, \psi^{(n-1)}\big) \ . \tag{5.12}$$

We call this the homotopy expansion of $\mathbf{dKW}_\bullet$ induced by the knot isotopy β The operator at zeroth order of the expansion is a comoving, or 'adiabatic', version of the extended Bogomolny equations in equation (5.7):

$$\mathbf{EBE}_q(\psi^{(0)}) = (\Delta_{\zeta,\bar{\zeta}} + \partial_y^2)\,\psi^{(0)} - r_q^{2\lambda} \exp(2\psi^{(0)}) \ .$$

Corrections to the adiabatic operator appear at order q^n with $n \geq 1$ and are given by the following linear second-order operators

$$\mathbf{dKW}_q^{(n)}(\psi^{(n)};\ \psi^{(0)}, \dots, \psi^{(n-1)}) = L_q\psi^{(n)} + K_q^{(n)}(\psi^{(0)}, \dots, \psi^{(n-1)}) \ .$$

The linear differential operator L_q is a comoving version of (5.11) and independent of n

$$L_q = \tilde{\partial}_t^2 + \Delta_{\zeta,\bar{\zeta}} + \partial_y^2 - r^{2\lambda} \exp(2\psi^{(0)}) \, .$$

In contrast, the inhomogeneous terms in $K_q^{(n)}$ depend explicitly on the solutions of the lower order equations. Using the notation $\pi \vdash n$ for a partition of n into $|\pi| = s$ positive integers π_i of multiplicity ν_i, the inhomogeneous terms are given by

$$\begin{aligned} K_q^{(n)} = & - \left(\ddot{z}_0 \partial_\zeta + \ddot{\bar{z}}_0 \partial_{\bar{\zeta}}\right) \psi^{(n-1)} - \left(\dot{z}_0 \partial_\zeta + \dot{\bar{z}}_0 \partial_{\bar{\zeta}}\right) \tilde{\partial}_t \psi^{(n-1)} + \left(\dot{z}_0 \partial_\zeta + \dot{\bar{z}}_0 \partial_{\bar{\zeta}}\right)^2 \psi^{(n-2)} \\ & - r_q^{2\lambda} \exp(2\psi^{(0)}) \sum_{\substack{\pi \vdash n \\ \pi \neq [n]}} \prod_{i=1}^{|\pi|} \frac{1}{\nu_i!} \left(\psi^{(\pi_i)}\right)^{\nu_i} \, . \end{aligned}$$

Since L_q is independent of n, the operators $\mathbf{dKW}_q^{(n)}$ only differ in the inhomogeneous terms $K_q^{(n)}$ determined by the lower order solutions $\psi^{(k)}$, $0 \leq k \leq n-1$. In particular, when $K_q^{(n)}$ is bounded each $\mathbf{dKW}_q^{(n)}$ is a Laplace-type operator and one can rely on the theory of elliptic operators. More generally, since $\psi^{(0)}$ encodes Nahm pole boundary conditions and knot singularities, $\mathbf{dKW}_q^{(n)}$ is expected to be an iterated edge operator.

5.6 The Method of Continuity

We propose that a continuity argument along the homotopy parameter $q \in [0,1]$ guarantees the existence of a solution of the decoupled Kapustin-Witten equations. While a proof is currently out of reach, we sketch a strategy that offers potential avenues for further exploration.

Recall that our current goal is to determine a gauge transformation g_β such that $g_\beta \cdot (A^\beta, \varphi^\beta, \phi_1^\beta)$ satisfies the decoupled Kapustin-Witten equations and exhibits Nahm pole boundary conditions as $y \to 0$ with knot singularities at $\beta_{q=1} = K$. Using the homotopy expansion (5.12) of the decoupled Kapustin-Witten equations induced by the knot isotopy $\beta_\bullet$, it suffices to show that the set

$$\mathcal{I} = \left\{ q \in [0,1] \;\middle|\; \exists\, \psi = \sum_{n=0}^{\infty} q^n \psi^{(n)} \;\text{ s.t. }\; \mathbf{dKW}_q^{(n)}(\psi^{(n)}) = 0,\; n \geq 0 \right\} \subseteq [0,1]$$

is non-empty, open and closed.

In the following we lay out some initial considerations for each of these assertion. Unfortunately, the proof strategy relies on analytic properties of the operators $\mathbf{dKW}_q^{(n)}$ that are currently not available and need a detailed investigation of their iterated edge structure.

$\mathcal{I}$ is non-empty. At $q = 0$ the knot β_0 is the S^1-invariant single-stranded knot located at the origin of the complex plane. In this case the equations reduce to $0 = \mathbf{dKW}_{q=0}(\psi) = \mathbf{EBE}\,(\psi^{(0)})$, a solution of which is provided by the results of [Mik12], so $0 \in \mathcal{I}$.

$\mathcal{I}$ **is open.** Let $q_0 \in \mathcal{I}$. We would like to show that there is an open neighbourhood U_{q_0} of q_0 in $\mathcal{I}$. For this, we first explain that in favourable circumstances there is such an open neighbourhood $U_{q_0}^{(n)}$ for each $\mathbf{dKW}_{q_0}^{(n)}$, individually.

Starting at $\mathcal{O}(q^0)$, let $\psi^{(0)}$ be a solution of $\mathbf{EBE}_{q_0}(\psi^{(0)}) = 0$. This means that at each $t \in S^1_t$ it is given by (5.8) with z replaced by $\zeta_{q_0}(t)$. Explicitly, $\psi^{(0)} = \log \frac{\lambda+1}{R_{q_0}^{\lambda+1} s_{q_0}(s_{q_0})_\lambda}$, where R_{q_0} and s_{q_0} are spherical coordinates based at $(q_0 z_0(t), 0) \in \mathbb{C} \times \mathbb{R}^+_y$. If we replace q_0 in this expression by any other $q \in [0,1]$, the resulting function is again a solution of the extended Bogomolny equations, namely of $\mathbf{EBE}_q(\psi^{(0)}(\zeta_q(t), y)) = 0$. It follows that $U^{(0)} = [0,1]$ is an open neighbourhood of q_0 for which there are solutions of $\mathbf{EBE}_q(\psi^{(0)}) = 0$, as requested.

Moving on to higher orders, let $n \geq 1$ and $(q_0, \psi^{(n)})$ a solution, i.e. $\mathbf{dKW}_{q_0}^{(n)}(\psi^{(n)}) = 0$. Assume that the map $\mathbf{dKW}_{\bullet}^{(n)} : [0,1] \times \mathcal{X} \to \mathcal{Y}$ is Fréchet differentiable and that its linearization $\mathcal{L}^{(n)}_{(q_0, \psi^{(n)})}(0, -)$ at the point $(q_0, \psi_0^{(n)})$ is an isomorphism of Banach spaces (where $\mathcal{X}, \mathcal{Y}$ are some appropriate function spaces that are attuned to Nahm pole boundary conditions with knot singularity). The implicit function theorem then guarantees the existence of an open neighbourhood $U_{q_0}^{(n)} \subset [0,1]$ of q_0 and a function $G^{(n)} : U_{q_0}^{(n)} \to \mathcal{X}$, such that $\mathbf{dKW}_q^{(n)}(G^{(n)}(q)) = 0$ for all $q \in U_{q_0}^{(n)}$.

It remains to show that the size of $U_{q_0}^{(n)}$ is bounded from below by some non-zero radius. In that case the intersection $U_{q_0} := \bigcap_{n \in \mathbb{N}} U_{q_0}^{(n)}$ is open. The size of the open neighbourhoods $U_{q_0}^{(n)}$ can be estimated from Lipschitz properties of the full Fréchet differential at $(q_0, \psi^{(0)})$. Since the physical theory is topological, we can stretch the radius of S^1_t and make $|(d/dt)^n z_0|$ arbitrarily small. We expect that this can be leveraged to gain control over the Fréchet differential. Once one has access to such bounds for each $\mathbf{dKW}_{q_0}^{(n)}$, one can conclude that for any point $q_0 \in \mathcal{I}$ the set U_{q_0} is open.

By construction, for all $q \in U_{q_0}$ there is a sequence $\psi^{(n)}$ that satisfies $\mathbf{dKW}_q^{(n)}(\psi^{(n)}) = 0$. In a last step, one needs to determine a function ψ that has the formal power series $\psi = \sum q^n \psi^{(n)}$ near $q \in U$. Recall that ψ is assumed to be an adiabatic solution, in the sense that for small changes in time $t \to t + \epsilon$ the higher orders $\psi^{(n)}$ provide only miniscule adjustments that keep the configuration 'in equilibrium', i.e. near a solution of the EBE. This suggests that the corrections $\psi^{(n)}$ are small, and we expect they are small enough that the formal power series either converges or might be dealt with by Borel resummation for any q that is close enough to q_0.

$\mathcal{I}$ **is closed.** Consider a sequence $\{q_k\} \subset \mathcal{I}$, together with a corresponding sequence $\{\psi_k\}$, such that $\mathbf{dKW}_{q_k}(\psi_k) = 0$. $\{q_k\}$ converges in $[0,1]$ to some $q := \lim_{k\to\infty} q_k$. We need to show that a subsequence of ψ_k converges to a corresponding ψ, such that $\mathbf{dKW}_q(\psi)$ vanishes to arbitrary order in q.

The desired statement would follow if we knew that the moduli space of solutions is compact. Unfortunately, there is not yet a complete description of the moduli spaces of Kapustin-Witten solutions. Although there has recently been progress for solutions on $X^3 \times \mathbb{R}^+$ when X^3 is compact [Tau18], the theory appears to involve a few subtleties that preclude a naive compactness result (also see related advances in [Tau19; Tau21]). Less is known about solutions of the decoupled Kapustin-Witten equations, but we expect that their Hermitian Yang-Mills structure provides some additional control.

In the situation relevant to us, it might be possible to directly construct a limiting solution, despite the lack of a general compactness theorem. As before, this works almost trivially at order q^0. To see this, consider for each ψ_k the associated formal expansion $\psi_k = \sum q^n \psi_k^{(n)}$. The terms at order q^0 define a sequence of EBE-solutions $\psi_k^{(0)} = \log \frac{\lambda+1}{R_{q_k}^{\lambda+1} s_{q_k}(s_{q_k})_\lambda}$. This is continuous in the isotopy parameter, so in the limit $q_k \to q$ it approaches a limit $\psi^{(0)}$ where q_k is replaced with q. The limit then satisfies $\mathbf{EBE}_q(\psi^{(0)}) = 0$.

At order q^1 we are given the sequence of functions $\psi_k^{(1)}$ that each satisfy $\mathbf{dKW}_{q_k}^{(1)} \psi_k^{(1)} = 0$ with respect to their associated q_k. Evaluating the action of $\mathbf{dKW}_q^{(1)}$ on each of the $\psi_k^{(1)}$ produces an error term

$$\mathbf{dKW}_q^{(1)}(\psi_k^{(1)}; \psi^{(0)}) = L_q \psi^{(1)} + K_q^{(1)}(\psi^{(0)}) \,.$$

The analytic properties of the Laplace-type iterated edge operator L_q are comparatively well-understood and it should be possible to utilize these to determine the size of the associated error term in relation to the distance $q - q_k$. The formula for the inhomogeneities $K_q^{(1)}$ states that it is proportional to $\ddot{z}_0$. It follows that $K_q^{(1)}$ is controlled by the magnitude of $\ddot{z}_0$, which can be made small by further stretching the knot in the direction of t. Although we do not currently have a proof, we expect that one can construct an approximate solution $\psi^{(1)}$ from $\psi_k^{(1)}$ and improve it to arbitrary precision by taking $k \to \infty$.

Moving on from there, one can then proceed order by order in q and in the end use a resummation argument to produce a solution $\psi = \sum q^n \psi^{(n)}$.

5.7 Comoving Higgs Bundles

In this section we move from a single-stranded knot in $S^1_t \times \mathbb{C}$ to the more general case of braids on k strands in $S^1_t \times \Sigma$. We find that the initial holomorphic data underlying the isotopy Ansatz (5.9) is captured by a one-parameter family of effective triples. This is in analogy to the relation between the initial knot singularity Ansatz (5.6) and effective triples reviewed in Section 5.4.

Let $K = \bigsqcup_{a=1}^{k} \{(t, p_a(t), 0)\}$ be a braid in the boundary of $S^1_t \times \Sigma \times \mathbb{R}^+_y$. The collection of trajectories $\{p_a(t)\}_{a=1,\dots,k}$ can be viewed as the image of a map $\beta : S^1_t \to \mathrm{Conf}_k \Sigma$ in the configuration space of k distinct, ordered points in Σ. Equivalently, β is an element of the loop space $\Omega \,\mathrm{Conf}_k \Sigma$. The homotopy classes of loops form what is known as the pure braid group on Σ.

We are interested in the moduli space of solutions of the decoupled Kapustin-Witten equations.

Definition 5.6 (Moduli Space of decoupled Kapustin-Witten solutions)**.**

$$\widehat{\mathcal{M}}_K^{\mathrm{dKW}} := \Big\{ \mathbf{dKW}(A,\phi) = 0 \,\Big|\, (A,\phi) \text{ satisfies Nahm-pole boundary conditions as } y \to 0$$
$$\text{with knot singularities at } K \text{ and converges to a flat } SL(n+1,\mathbb{C}) \text{ connection as } y \to \infty \Big\}$$

As before, we denote by $\mathcal{M}_K^{\mathrm{dKW}}$ the quotient by real gauge transformations $\mathcal{G}_0(S^1_t \times \Sigma \times \mathbb{R}^+_y)$ that vanish at the boundary. However, most of the upcoming discussions will focus on the infinite dimensional moduli spaces before modding out gauge transformations.

We now slightly change perspective and view the knot singularity data as part of the moduli of the problem. Assume, for the sake of simplicity, that all strands of K are labeled by the same weight $\lambda \in \Gamma^\vee_{\mathrm{char}}$. Specifically, we will from now on view the moduli spaces $\widehat{\mathcal{M}}_K^{\mathrm{dKW}}$ as fibers of a bundle

$$\begin{array}{c} \widehat{\mathcal{M}}^{\mathrm{dKW}} \\ \big\downarrow \\ \Omega\,\mathrm{Conf}_k\,\Sigma \end{array}$$

The fiber map sends each solution to the loop β along which the solution exhibits a knot singularity.

There is an analogous fiber bundle $\widehat{\mathcal{M}}^{\mathrm{EBE}} \to \mathrm{Conf}_k\,\Sigma$ for solutions of the extended Bogomolny equations. In this case the fibers are given by $\widehat{\mathcal{M}}_D^{\mathrm{EBE}}$ with fixed knot singularity data $D = \{(p_a,\lambda)\}_{a=1,\dots,k}$ and the fiber map sends a given solution to $\{p_a\}_{a=1,\dots,k} \in \mathrm{Conf}_k\,\Sigma$.

Recall from Section 5.4 that the moduli spaces $\mathcal{M}_D^{\mathrm{EBE}}$ are in bijection with the moduli space of effective triples $\mathcal{M}_{(\mathcal{E},\varphi,L)}$, where the associated divisor $\mathfrak{d}(\varphi,L)$ was required to coincide with the knot singularity data D. Reinterpreting the data of the divisor as part of the moduli of an effective triple, we obtain a corresponding fiber bundle

$$\begin{array}{c} \widehat{\mathcal{M}}_{(\mathcal{E},\varphi,L)} \\ \big\downarrow \\ \mathrm{Conf}_k\,\Sigma \end{array}$$

The fiber map sends each effective triple to the points of its divisor $\mathfrak{d}(L,\varphi) = \{(p_a,\lambda_a)\}_{a=1,\dots k}$.

Remark. More generally, if one includes the information of weights λ_a at each point, the base of the preceding fibrations is the configuration space of labeled points $\mathrm{Conf}_k(\Sigma;\lambda_1,\dots,\lambda_k)$ and its corresponding loop space, respectively.

The same arguments as in Section 5.4 can be used to show that any solution (A, ϕ) of the decoupled Kapustin-Witten equations gives rise to a one-parameter family of 'comoving' effective triples $\{(\mathcal{E}_t, \varphi_t, L_t)\}_{t \in S^1}$. To see this, assume (A, ϕ) is a solution of the decoupled Kapustin-Witten equations with knot singularity along a braid $\beta = \bigsqcup \{p_a(t)\}$. Let $\mathcal{D}_\mu$ denote the operators associated to (A, B) and work in temporal gauge $A_s = 0$, such that $\mathcal{D}_0 = iD_t$.

The decoupled Kapustin-Witten equations contain the equations $[\mathcal{D}_i, \mathcal{D}_j] = 0$, $i, j = 1, 2, 3$. Denote by V the vector bundle associated to the fundamental representation of $SL(N, \mathbb{C})$ and let $V_{(t,y)}$ be its restriction to $\{t\} \times \Sigma \times \{y\}$. As before, $\mathcal{D}_1$ provides a $\bar\partial$ operator, making $V_{(t,y)}$ into a holomorphic vector bundle $\mathcal{E}_{(t,y)}$, while $\mathcal{D}_3 = \mathrm{ad}_\varphi$ is a holomorphic K_Σ-valued endomorphism of $\mathcal{E}_{(t,y)}$. The associated family of Higgs bundles $(\mathcal{E}_{(t,y)}, \varphi_{(t,y)})$ over Σ is parallel with respect to $\mathcal{D}_2$, such that for each $t \in S^1$ the Higgs bundle is determined via parallel transport of a stable reference Higgs bundle $(\mathcal{E}_t, \varphi_t)$ that sits, for example, at $y = \infty$. Finally, for each $t \in S^1$, the boundary condition at $y \to 0$ determines a vanishing line bundle $L_t \subset \mathcal{E}_t$ whose divisor $\mathfrak{d}(L, \varphi) = \{p_a(t), \lambda\}$.

The decoupled Kapustin-Witten equations additionally include the operator $\mathcal{D}_0$, which yields a notion of parallel transport along S^1_t. The equations $[\mathcal{D}_0, \mathcal{D}_i] = 0$ state that $(\mathcal{E}_t, \varphi_t, L_t)$ is parallel. Equivalently, parallel transport via $\mathcal{D}_0$ defines an Ehresmann connection on $\widehat{\mathcal{M}}_{(\mathcal{E},\varphi,L)}$ with respect to which $(\mathcal{E}_t, \varphi_t, L_t)$ is a horizontal lift of the braid β.

Denote by $\Omega_h \widehat{\mathcal{M}}_{(\mathcal{E},\varphi,L)}$ the space of non-vertical loops, i.e. those loops that are horizontal with respect to *some* Ehresmann connection. The discussion above implies that there is the following Kobayashi-Hitchin-like bundle map:

$$
\begin{tikzcd}
\widehat{\mathcal{M}}^{\mathrm{dKW}} \arrow[rr, "\hat{I}_{\mathrm{KH}}"] \arrow[dr] & & \Omega_h \widehat{\mathcal{M}}_{(\mathcal{E},\varphi,L)} \arrow[dl] \\
& \Omega \operatorname{Conf}_k \Sigma &
\end{tikzcd}
$$

We expect that this descends to a corresponding map on quotient spaces

$$
I_{\mathrm{KH}} : \mathcal{M}^{\mathrm{dKW}} \to \Omega_h \widehat{\mathcal{M}}_{(\mathcal{E},\varphi,L)} / \mathcal{G}_{\mathbb{C}}(S^1_t \times \Sigma) \, .
$$

After modding out gauge transformations, He and Mazzeo's classification of S^1-invariant Kapustin-Witten solutions in terms of EBE-solutions corresponds to a bijection of fibers over constant loops in $\Omega \operatorname{Conf}_k \Sigma$. Keeping in mind that according to Theorem 5.5 there is a bijection $I_{NPK} : \mathcal{M}^{\mathrm{EBE}} \to \mathcal{M}_{(\mathcal{E},\varphi,L)}$, this fits into the bundle map picture as follows

$$
\begin{tikzcd}
\mathcal{M}^{\mathrm{dKW}} \arrow[r, "I_{\mathrm{KH}}"] \arrow[dr] \arrow[rr, bend left, "I_{NPK}"] & \Omega_h \mathcal{M}_{(\mathcal{E},\varphi,L)} \arrow[d] & \mathcal{M}_{(\mathcal{E},\varphi,L)} \arrow[l, hook', "\text{const.}"'] \arrow[d] \\
& \Omega \operatorname{Conf}_k \Sigma & \operatorname{Conf}_k \Sigma \arrow[l, hook', "\text{const.}"']
\end{tikzcd}
$$

As a next step, we describe how to obtain a solution of the decoupled Kapustin-Witten equations from a family of effective triples. Hence, suppose we are given a non-vertical loop of effective triples $\gamma(t) = (\mathcal{E}(t), \varphi(t), L(t))$. The knot singularity data associated to $\gamma(t)$ under the fiber map is a braid $\beta(t) = \{p_a(t)\}_{a=1,\dots,k}$. We wish to construct from γ a field configuration (A, B) that exhibits a Nahm pole with knot singularities along β and satisfies the decoupled Kapustin-Witten equations.

Under the assumption that the continuity method of Section 5.6 provides a solution for any single-stranded knot in $\mathbb{C}$, this can be achieved by essentially the same arguments as in the context of the Kobayashi-Hitchin correspondence for EBE-solutions described in Section 5.4. To that end, cover the manifold by open slices $V_i = (t_i - \epsilon, t_i + \epsilon) \times \Sigma \times \mathbb{R}_+$ and write $\beta_i = \{p_{i,a}(t)\}_{a=1,\dots k}$ for the part of the braid that lies in V_i. Choose the slices V_i thin enough that there exists an isotopy that interpolates between β_i and some S^1-invariant braid $\{P_{i,a} = \text{const.}\}_{a=1,\dots,k} \in \text{Conf}_k \Sigma$. Moreover, we demand that this isotopy is of a particularly mild form, namely such that there exist non-intersecting discs $U_{i,a}$, centered at $P_{i,a}$, which each contain the corresponding trajectory $p_{i,a}(t)$ for all $t \in (t_i - \epsilon, t_i + \epsilon)$ and the isotopy that connects $p_{i,a}(t)$ to $P_{i,a}$.

We construct an approximate solution of the decoupled Kapustin-Witten equations according to the following outline. Starting with a single slice V_i, choose an open cover of Σ consisting of an open set $U_{i,0}$ that does not contain any of the trajectories $p_{i,a}(t)$, together with the open discs $U_{i,a}$ described above. Using the method of continuity of Section 5.6, we now that there is a field configuration that satisfies Nahm pole boundary conditions with knot singularity on each $(t_i - \epsilon, t_i + \epsilon) \times U_{i,a} \times \mathbb{R}^+_y$ independently. These configurations are then glued 'horizontally' over $U_{i,0}$ to an approximate solution on V_i, which subsequently can be glued 'vertically' over the intersections of $V_i \cap V_{i+1}$ to produce an approximate solution on $S^1 \times \Sigma \times \mathbb{R}^+_y$. The construction proceeds in five steps.

First, dropping the indices i and a for the moment, consider a small disk U centered at P and containing the trajectory $p(t)$ of a single strand. Let z be a holomorphic coordinate on U, assume P corresponds to $z = 0$, and denote the coordinates of $p(t)$ by $z_0(t)$. We also use coordinates (t, R, ψ, ϑ), with spherical coordinates on the second and third factor of $S^1_t \times U \times \mathbb{R}^+_y$. We may then extract from $\gamma(t) = (\mathcal{E}(t), \varphi(t), L(t))$ a field configuration on $[t_i - \epsilon, t_i + \epsilon] \times U \times \mathbb{R}^+_y$ that looks like the isotopy Ansatz $(A^\beta, \varphi^\beta, \phi^\beta_1)$ as given in (5.9): Recall that for each $t \in (t_i - \epsilon, t_i + \epsilon)$, the effective triple defines a frame $\{e_1, \dots e_{n+1}\}$ of $\mathcal{E}_t$ with respect to which the Higgs field is of the form $\varphi(t) = \sum (z - z_a(t))^{\lambda_i} E^+_i \, dz$. The isotopy that connects P to $p(t)$, and which is contained in U by assumption, is homotopy equivalent to the isotopy given by $\zeta_q(t) = z - qz_a(t)$. Replacing $z - z_a(t)$ by $\zeta_q(t)$ in the expression for $\varphi(t)$ then provides the q^0 part of the isotopy Ansatz. To get the part proportional to q^1, we identify the gauge field component A^β_t with the connection form of some Ehresmann connection with respect to which γ is horizontal. Since $D_t \varphi(t) = 0$, the connection form satisfies $[A_t, \varphi(t)] = -i\partial_t \varphi(t)$. This is solved by setting $A^\beta_t = \frac{q\dot{z}_a}{\zeta} \sum \lambda_i A^{-1}_{ij} H_j$.

In the second step, we rely on the method of continuity to invoke the existence of a complex gauge transformation g_U that maps the field configuration $(A^\beta, \varphi^\beta, \phi^\beta_1)$ to an approximate solu-

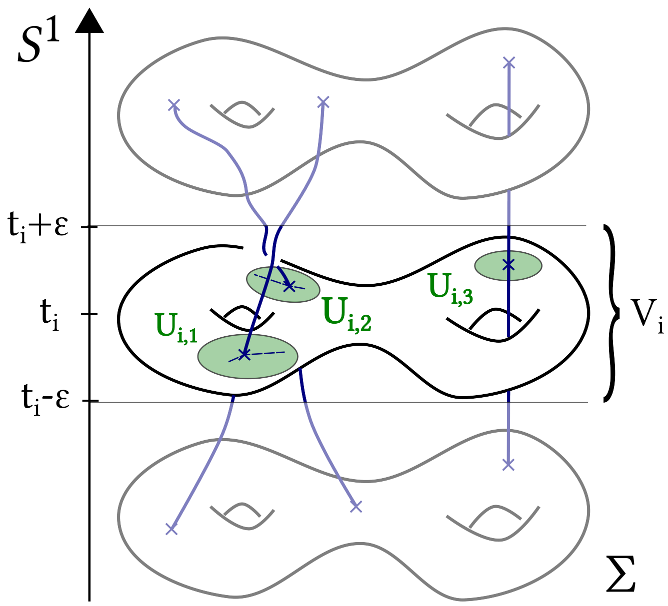

Figure 5.3 Illustration of the covering of $S^1_t \times \Sigma \times \mathbb{R}^+_y$ that is used in the iterative construction of approximate solutions of the decoupled Kapustin-Witten equations for a multi-stranded and time-dependent knot.

tion of the decoupled Kapustin-Witten equations on $(t_i - \epsilon, t_i + \epsilon) \times U \times \mathbb{R}^+_y$. By construction, this satisfies the Nahm pole boundary conditions with knot singularity at $p(t)$.

The third step constructs a gauge transformation g_{V_i} on $V_i = (t_i - \epsilon, t_i + \epsilon) \times \Sigma \times \mathbb{R}^+_y$ from the collection of local frames and gauge transformations $\{g_{\lambda_a}\}$ associated to U_a. To do so, we additionally choose a holomorphic frame on $(t_i - \epsilon, t_i + \epsilon) \times U_0 \times \mathbb{R}^+_y$ and let g_0 be the gauge transformation defined in (5.5). Just as before, on the overlaps $U_0 \cap U_a$, the holomorphic frames are related by the transition functions

$$g_{0a} = \exp\big(- \log r \sum_{i=1}^{n} \lambda_{a,j} A^{-1}_{ij} H_i \big)$$

Gluing the gauge transformations g_0 and g_{U_a} via a partition of unity produces the desired g_{V_i}.

Fourth, the holomorphic frames on $U_{i,0}$ and $U_{i+1,0}$ are equivalent up to gauge transformations, so they can in turn can be glued with $g_{V_{i+1}}$ on the intersection $V_i \cap V_{i+1}$.

The resulting field configuration is a first approximation to a solution of the decoupled Kapustin-Witten equations. It is of order $\mathcal{O}(y^{-1})$ away from the braid β and of order $\mathcal{O}(R^{-1}s^{-1})$ near any of its strands, which means that it satisfies Nahm pole boundary conditions with knot singularities. The approximation can be improved order by order, first to desired precision in R near any point $p \in K$ and afterwards in y.

5.8 Effective Triples, Monopoles, and the Grothendieck-Springer Fibration

We propose that when $\Sigma = \mathbb{C}$, there is a finite-dimensional fiber bundle that encodes enough of the infinite-dimensional moduli space of EBE solutions to provide existence results for solutions of the decoupled Kapustin-Witten equations. The fiber bundle in question is a variant of the Grothendieck-Springer resolution of $\mathfrak{sl}(kN, \mathbb{C})$, which also appears in the definition of symplectic Khovanov homology [SS04; Man07]. The motivation to consider this space as a finite-dimensional model of the problem is ultimately based in physical properties of the system and motivated by observations of Gaiotto and Witten [GW12]. In some sense this is comparable to an ADHM-like approach, where the construction of Yang-Mills instantons is reduced to a finite dimensional problem in linear algebra [Ati+78].

Consider the image of a braid β on k strands in $S^1_t \times \mathbb{C}$ and denote the trajectories of the strands in $\mathbb{C}$ by $z_a(t)$. Our interest is in the time evolution of k-monopole solutions of the extended Bogomolny equations along β. For this, we view the monopole insertions at $(t, z_a(t), 0)$ in $S^1_t \times \mathbb{C} \times \mathbb{R}^+_y$ as 'heavy particles' with magnetic charge λ_a, that trace out fixed, non-colliding trajectories in $\mathbb{C}$. Since the underlying physical theory is topological, these particles don't interact and the multi-particle system is equivalent to the union of k single particles. Correspondingly, the configuration of each individual monopole is fully determined in a small neighbourhood. The assumption that the particles are 'heavy' means that the evolution of each monopole along S^1_t is viewed as an externally determined background in which the quantum theory lives. More specifically, the quantum system starts at an initial time in some ground state and then evolves adiabatically along S^1_t, meaning that at any given time t the system remains in a ground state of the corresponding, fixed background configuration at that time. This picture suggests that, as a first step, we need to determine the moduli of a collection of k individual monopoles.

According to the Kobayashi-Hitchin correspondence discussed in Section 5.4, a k-monopole solution of the extended Bogomolny equations is determined by an effective triple. Hence, assume we are given $(\mathcal{E}, \varphi, L)$ with associated divisor $\mathfrak{d}(L, \varphi) = \{(z_a, \lambda_a)\}_{a=1,\dots,k}$. Due to the physical argument above, we now restrict our attention to the moduli of the Higgs field φ near the points of $\mathfrak{d}(L, \varphi)$. We have seen in Section 5.4 and Section 5.7 that on a small enough disc $U_a \subset \mathbb{C}$ centered at z_a the pair (L, φ) is equivalent to the data contained in the Ansatz $\varphi^\lambda|_{U_a} = \sum_i z^{\lambda_{a,i}} E_i^+$. The latter, in turn, is determined by the choice of

(i) a nilpotent element $E_a \in \mathcal{O}_{\pi_a}$, such that $\varphi^\lambda|_{z=0} = E_a$, and

(ii) a nilpotent element $K_a \in \ker F_a$, such that $\varphi^\lambda|_{z=1} = E_a + K_a \in \mathcal{O}_{\mathrm{reg}}$.

Recall from Section 5.3 that the partition $\pi_a = [\pi_{a,1} \dots \pi_{a,s}]$ is determined by the monopole charge λ_a and given by counting the numbers of consecutive Dynkin labels in $\lambda_a = (\lambda_{a,1}, \dots, \lambda_{a,n})$ that vanish. We conclude that individual monopoles are determined by a choice of element in the intersection of the Slodowy slice $\mathcal{S}_{E_a}$ and the regular nilpotent orbit $\mathcal{O}_{\mathrm{reg}}$ in $\mathfrak{sl}(N, \mathbb{C})$.

A naive approach to keep book of all k monopoles simultaneously is to combine their moduli into a coproduct. In finite dimensions this amounts to putting everything together into a single matrix Y of size kN.

$$Y := \left(\begin{array}{c|c|c} E_1 + K_1 & 0 & 0 \\ \hline 0 & \ddots & 0 \\ \hline 0 & 0 & E_k + K_k \end{array}\right)$$

Put differently, Y is an element of

$$(\mathcal{S}_{E_1} \cap \mathcal{O}_{\text{reg}}) \oplus \ldots \oplus (\mathcal{S}_{E_k} \cap \mathcal{O}_{\text{reg}}) \subset \mathfrak{sl}(kN, \mathbb{C}) .$$

Remember that $\mathcal{O}_{\text{reg}} = \mathcal{O}_{[N]}$ and let $\rho = [N^k]$ be the partition of kN that consists of k copies of N. Similarly, denote the partition of kN that is obtained from concatenating all π_a by $\pi = [\pi_1 \ldots \pi_k]$. There exists a $g \in SL(kN, \mathbb{C})$ that maps $E_1 \oplus \ldots \oplus E_k$ to its Jordan normal form E_π. It follows that the k-monopole configuration Y is always conjugate to an element of $\mathcal{S}_\pi \cap \overline{\mathcal{O}_\rho}$ (recall $\mathcal{S}_\pi := \mathcal{S}_{E_\pi}$). The upcoming constructions only depend on the topology and geometry of Slodowy slices, which are independent of the base point. Accordingly, we from now on consider a k-monopole configuration to be determined by a choice of

$$Y \in \mathcal{S}_\pi \cap \overline{\mathcal{O}_\rho} \subset \mathfrak{sl}(kN, \mathbb{C}) . \tag{5.13}$$

Unfortunately, there are fundamental problems in determining Kapustin-Witten solutions from nilpotent Higgs bundles directly [GW12; Dim22b; Sun23]. Inspired by the 'complex symmetry breaking' suggested by Gaiotto and Witten to circumvent these problems, and also because we wish to encode the positions z_a of the monopoles, we modify the naive approach. Namely, we additionally encode the position z_a of each strand by a deformation of the nilpotent orbit $\overline{\mathcal{O}_\rho}$ that appears in (5.13).

Define the traceless, diagonal $N \times N$-matrix

$$D(z_a) = \operatorname{diag}\left(z_a,\ \ldots,\ z_a,\ -(N-1)z_a\right) \in \mathfrak{sl}(N, \mathbb{C}) .$$

We call z_a 'thick' eigenvalue of $D(z_a)$ as it appears with multiplicity $N-1$, while we refer to the remaining eigenvalue $-(N-1)z_a$ as 'thin'.

For fixed $t \in S^1_t$, denote the position of the strands by $D = \{z_a(t)\}_{a=1,\ldots,k} \in \operatorname{Conf}_k \mathbb{C}$. Write $\mathcal{O}_{\rho,D}$ for the adjoint orbit of $D(z_1) \oplus \ldots \oplus D(z_k)$ in $\mathfrak{sl}(kN, \mathbb{C})$ and consider the map

$$Y \mapsto Y + (D(z_1) \oplus \ldots \oplus D(z_k)) \ .$$

The eigenvalues of the new matrix are given by the positions z_a, each with (generalized) eigenspace of dimension $(N-1)$, and the values $-(N-1)z_a$ with eigenspace of dimension 1. This defines a smooth deformation $\overline{\mathcal{O}_\rho} \rightsquigarrow \mathcal{O}_{\rho,D}$.

We conclude that the moduli space of k monopoles of charge λ and located at $D = \{z_a\}$ is given by

$$\mathcal{Y}_{\pi,\rho,D} := \mathcal{S}_\pi \cap \mathcal{O}_{\rho,D} \subset \mathfrak{sl}(kN, \mathbb{C}) .$$

As an intersection of a Slodowy slice with an adjoint orbit, $\mathcal{Y}_{\pi,\rho,D}$ is a finite dimensional Kähler manifold.

Remark. Going from $\widehat{\mathcal{M}}_{(\mathcal{E},\varphi,L)} \to \mathcal{Y}_{\pi,\rho,D}$ can be interpreted as modding out a class of 'large' gauge transformations that are non-zero at the boundary. A priori, we are not allowed to mod out such gauge transformations, because the associated boundary conditions are physically distinguishable. However, there is some evidence that for the class of large gauge transformations in question the is always a Haydys-Witten instantons that interpolates between the associated two (a priori inequivalent) solutions of the decoupled Haydys-Witten equations. Such a result would provide a good a posteriori justification for replacing the moduli space of effective triples by $\mathcal{Y}_{\pi,\rho,D}$. We further comment on this in Section 5.10.

A variant of the space $\mathcal{Y}_{\pi,\rho,D}$ plays a major role in the definition of symplectic Khovanov homology and symplectic $\mathfrak{sl}(N,\mathbb{C})$-Khovanov-Rozansky homology [SS04; Man07]. Importantly, $\mathcal{Y}_{\pi,\rho,D}$ can be viewed as a fiber of what Manolescu calls a *restricted partial simultaneous Grothendieck resolution* of the adjoint quotient map. For our purposes it will suffice to define the relevant version of the adjoint quotient map 'by hand' and we refer to [Man07] for a more satisfactory exposition of the general construction.

Remember that the adjoint quotient map $\chi : \mathfrak{sl}(kN,\mathbb{C}) \to \mathfrak{h}/\mathcal{W}$ sends a matrix to its generalized eigenvalues. Denote by $\mathfrak{sl}(kN,\mathbb{C})^{[(N-1)^k 1^k]}$ the space of traceless $kN \times kN$ matrices that have k pairs of eigenvalues that cancel each other in the trace, where each pair has one 'thick' eigenvalue of algebraic multiplicity $N-1$ and one 'thin' eigenvalue of multiplicity 1, respectively. A prototypical element of this space is $D(z_1) \oplus \ldots \oplus D(z_k)$, but we explicitly allow non-diagonalizable elements with Jordan blocks of size greater than 1. Define $\tilde{\chi}$ to be the map that sends an element of $\mathfrak{sl}(kN,\mathbb{C})^{[(N-1)^k 1^k]}$ to its thick eigenvalues $(z_1,\ldots,z_k) \in \mathrm{Conf}_k\,\mathbb{C}$:

$$\tilde{\chi} : \mathfrak{sl}(kN,\mathbb{C})^{[(N-1)^k\,1^k]} \to \mathrm{Conf}_k\,\mathbb{C}$$

We will simply refer to this as Grothendieck-Springer fibration. With that understood $\mathcal{Y}_{\pi,\rho,D}$ is equivalent to the following fiber of $\tilde{\chi}$:

$$\mathcal{Y}_{\pi,\rho,D} = \tilde{\chi}|_{S_\pi}^{-1}\big(D(z_1) \oplus \ldots \oplus D(z_k)\big)\ .$$

Our discussion so far shows that every effective triple $(\mathcal{E},\varphi,L)$ with divisor $\mathfrak{d}(L,\varphi) = D$ determines an element $Y \in \mathcal{Y}_{\pi,\rho,D}$. This induces the following bundle map:

$$\begin{tikzcd}
\widehat{\mathcal{M}}_{(\mathcal{E},\varphi,L)} \arrow[rr, "\Upsilon"] \arrow[dr] & & S_\pi \cap \mathfrak{sl}(kN,\mathbb{C})^{[(N-1)^k\,1^k]} \arrow[dl, "\tilde{\chi}|_{S_\pi}"] \\
& \mathrm{Conf}_k\,\mathbb{C} &
\end{tikzcd} \tag{5.14}$$

Υ is clearly not injective, since it maps effective triples that differ in regions of Σ sufficiently far from any monopole insertions to the same element in $\mathcal{Y}_{\pi,\rho,D}$. However, we expect that Υ is surjective. If this holds, we are guaranteed that for every non-vertical path in $\mathfrak{sl}(kN,\mathbb{C})^{[(N-1)^k\,1^k]}$, there exists at least one non-vertical path in $\widehat{\mathcal{M}}_{(\mathcal{E},\varphi,L)}$, which in turn determines a solution of the decoupled Kapustin-Witten equations by the arguments presented in Section 5.7.

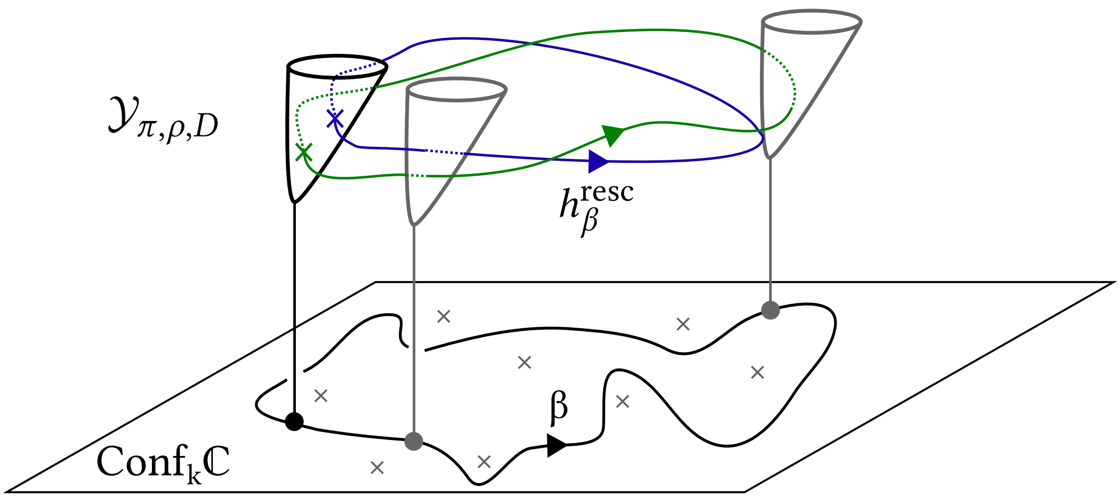

Figure 5.4 The fixed points of rescaled parallel transport $h_\beta^{\mathrm{resc}} : \mathcal{Y}_{\pi,\rho,D} \to \mathcal{Y}_{\pi,\rho,D}$ along a pure braid β determine inequivalent horizontal lifts of β.

As discussed in [SS04, Sec. 4A] and [Man07, Sec. 4.1], since $\mathfrak{sl}(kN,\mathbb{C})^{[(N-1)^k\,1^k]}$ carries a Kähler metric, there exists a suitable notion of (rescaled) parallel transport along paths $\beta : [0,1] \to \mathrm{Conf}_k\,\mathbb{C}$ from $D = \beta(0)$ to $D' = \beta(1)$.

$$h_\beta^{\mathrm{resc}} : \mathcal{Y}_{\pi,\rho,D} \to \mathcal{Y}_{\pi,\rho,D'} \,.$$

Some care is needed in defining h_β^{resc}, since $\mathfrak{sl}(kN,\mathbb{C})^{[(N-1)^k\,1^k]}$ is not compact and $\tilde{\chi}$ may fail to be a proper map. In that situation the integral lines of the vector field $H_\beta = \dot{\beta}\,\nabla\tilde{\chi}/\,\|\nabla\tilde{\chi}\|^2$, which is parallel to β and orthogonal to the fibers, may not exist at all times. To circumvent this problem, one rescales H_β in such a way that its integral lines stay in some controlled compact subset of $\mathfrak{sl}(kN,\mathbb{C})^{[(N-1)^k\,1^k]}$. Strictly speaking, h_β^{resc} is only well-defined on (arbitrarily large) compact subsets.

We can apply parallel transport along a pure braid $\beta : S^1 \to \mathrm{Conf}_k\,\mathbb{C}$, in which case $\mathcal{Y}_{\pi,\rho,D} = \mathcal{Y}_{\pi,\rho,D'}$. Horizontal lifts of β that start and end at the same point in $\mathcal{Y}_{\pi,\rho,D}$ are then in one-to-one correspondence with fixed points of h_β^{resc}. With this we arrive at one of the main claims of this chapter:

Conjecture E. *The number of solutions to the decoupled Kapustin-Witten equations on $S^1_t \times \mathbb{C} \times \mathbb{R}^+_y$ with knot singularities along β of weight λ is bounded from below by the number of fixed points of h_β^{resc}.*

In the upcoming sections the ideas of this section are applied to a slightly different setting. This leads to a considerably stronger version of Conjecture E, providing an interpretation of Witten's conjecture that Haydys-Witten instanton homology is related to Khovanov homology.

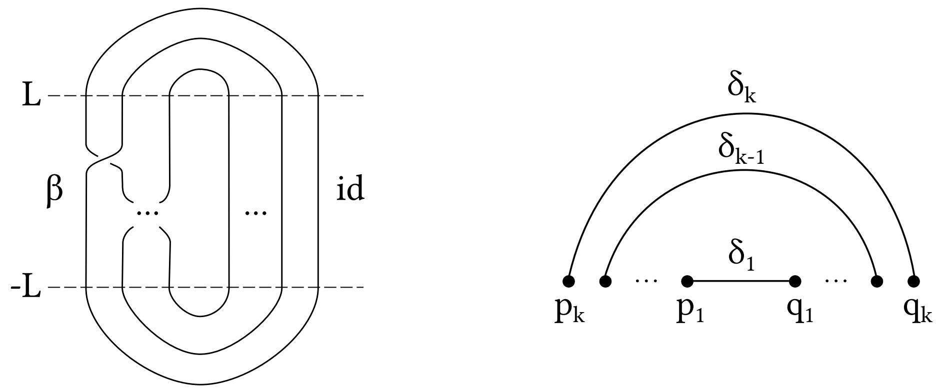

Figure 5.5 Left: Every knot K can be viewed as closure of a bipartite braid $\beta \times \mathrm{id}$ embedded in $[-L, L]_t \times \mathbb{C}$ by gluing in cups and caps. Right: Gluing instructions for cups and caps are captured by a crossingless matching $\mathfrak{m}$ consisting of k disjoint arcs $\delta_i \subset \mathbb{C}$.

5.9 From Braids to Knots

We now apply the ideas of the previous section to the decoupled Kapustin-Witten equations with knot singularities along a compact knot K in the boundary of $W^4 = \mathbb{R}_t \times \mathbb{C} \times \mathbb{R}^+_y$. Since the factor of S^1_t is decompactified, we can no longer identify K with an element of $\Omega \operatorname{Conf}_k \mathbb{C}$. However, by Alexander's theorem any knot K can be represented as the closure of a braid β on k strands.

Let us denote the configuration space of $2k$ points that are partitioned into two individual sets of k points by $\operatorname{Conf}_{k,k} \mathbb{C}$. The closure of a braid β is determined by the following data. First, an embedding of a bipartite braid of the form $\beta \times \mathrm{id} : [-L, L]_t \to \operatorname{Conf}_{k,k} \mathbb{C}$ into $[-L, L] \times \mathbb{C}$, where the identity braid on the second set of points is constant. And second, by a choice of crossingless matching between the two sets of points determined by $\beta \times \mathrm{id}$ at $t = -L$, and analogously for $t = L$. To make this precise, assume the bipartite braid $\beta \times \mathrm{id}$ determines a collection of points $D = \{(p_1, \dots, p_k), (q_1, \dots q_k)\}$ at $t = L$. A crossingless matching of D is a collection of k disjoint embedded arcs $\mathfrak{m} = (\delta_1, \dots, \delta_k)$ in $\mathbb{C}$, with starting points $\delta_i(0) = p_i$ and endpoints $\delta_i(T) = q_i$. The arcs δ_i determine which strands are glued together by cups below $t = -L$ (respectively, by caps above $t = L$) to make the open braid $\beta \times \mathrm{id} : [-L, L]_t \to \operatorname{Conf}_{k,k} \mathbb{C}$ into a closed knot in $\mathbb{R}_t \times \mathbb{C}$.

More abstractly, a crossingless matching determines an extension of $\beta \times \mathrm{id}$ into the singular locus of $\overline{\operatorname{Conf}_{k,k} \mathbb{C}} \simeq \overline{\operatorname{Conf}_{2k} \mathbb{C}}$. Here we view the closure of the configuration space as a stratified space by identifying configurations of k points with two identical points as a configuration of $k - 1$ points, and so on. The associated stratification is given by the inclusions $\operatorname{Conf}_{k-1} \mathbb{C} \subset \overline{\operatorname{Conf}_k \mathbb{C}}$, with lowest stratum $\operatorname{Conf}_0 \mathbb{C} = \emptyset$. An entrance path in a stratified space is a path that starts in a higher-dimensional stratum and can only transition into lower-dimensional strata or remain in the same stratum: it only ever enters lower-dimensional strata.

A single arc δ_i can be viewed as an entrance path $[0, 1] \to \overline{\operatorname{Conf}_{k,k} \mathbb{C}}$ that starts in the top stratum $\mathrm{Conf}_{k,k} \simeq \mathrm{Conf}_{2k}\,\mathbb{C}$ and ends in the lower stratum $\mathrm{Conf}_{2k-1}\,\mathbb{C}$. For this one identifies the arc with the path determined by $p_i(t) = \delta(t/2 \cdot T)$ and $q_i(t) = \delta_i((1-t)/2 \cdot T))$. Moreover, after two points have been matched, the two strands of the braid are closed off and for the purposes of describing the knot closure we can simply forget about the endpoint, further mapping $\mathrm{Conf}_{2k-1}\,\mathbb{C} \to \mathrm{Conf}_{2k-2}\,\mathbb{C}$ which is then identified with $\mathrm{Conf}_{k-1,k-1}\,\mathbb{C}$ in the obvious way. By successively following the arcs in $\mathfrak{m} = (\delta_1, \dots, \delta_k)$, a crossingless matching thus defines an entrance path that starts in $\mathrm{Conf}_{k,k}$ and ends in $\emptyset$.

Example. Consider the case of four points $\{(p_1, p_2), (q_1, q_2)\}$ and an arc δ_1 that connects p_1 and q_1 at $\delta_1(1) = r \in \mathbb{C}$. Then this determines the following entrance path connecting $\mathrm{Conf}_{2,2}\,\mathbb{C}$ to $\mathrm{Conf}_{1,1}\,\mathbb{C}$

$$\begin{array}{ccccc} \overline{\mathrm{Conf}_{2,2}\,\mathbb{C}} \simeq \overline{\mathrm{Conf}_4\,\mathbb{C}} & \supset & \overline{\mathrm{Conf}_3\,\mathbb{C}} & \supset & \overline{\mathrm{Conf}_2\,\mathbb{C}} \simeq \mathrm{Conf}_{1,1}\,\mathbb{C} \\ \{(p_1(t), p_2), (q_1(t), q_2)\} & \to & \{(r, p_2), (r, q_2)\} \simeq \{r, p_2, q_2\} & \to & \{p_2, q_2\} \simeq \{(p_2), (q_2)\} \end{array}$$

Closing off an open braid by cups and caps imposes constraints on k-monopole configurations. From the point of view of two individual monopoles, located at p_i and q_i, the matching arc δ_i specifies that their two configurations must be identical at $\delta_i(1)$. But this also means that they can't be too different for $1 - \epsilon$. In that way a crossingless matching $\mathfrak{m}$ forces monopole configurations near $t = \pm L$ to be elements of a compact Lagrangian subspace $L_{\mathfrak{m}} \subset \mathcal{Y}_{\pi,\rho,D}$. While we approach the problem from a slightly different perspective, the construction of $L_{\mathfrak{m}}$ is equivalent to the one described by Seidel-Smith and Manolescu [SS04; Man07]. We proceed by induction on the number of arcs k in a given matching.

Start with $k = 1$. The matching $\mathfrak{m}$ contains a single arc δ, matching 2 points $\{(p), (q)\} \in \mathrm{Conf}_{1,1}\,\mathbb{C}$ that are labeled by the same partition $\pi_1 = \pi_2 = \pi$. Interpret the arc as an entrance path $\delta : [0, 1] \to \overline{\mathrm{Conf}_{1,1}\,\mathbb{C}}$, starting at $\delta(0) = (p, q) \in \mathrm{Conf}_{1,1}\,\mathbb{C}$ and ending at $\delta(1) = r \in \mathrm{Conf}_1\,\mathbb{C}$. There is nothing special about the choice of r, so we pick $r = 0$ for simplicity. The Grothendieck-Springer fibration $\tilde{\chi} : \mathfrak{sl}(2N, \mathbb{C})^{[(N-1)^2 1^2]} \to \overline{\mathrm{Conf}_{1,1}\,\mathbb{C}}$ has singular fibers over configurations with $p = q$. The singularity in the fibers corresponds to the loci of matrices with Jordan blocks larger than 1.

Use naive (as opposed to rescaled) parallel transport h_δ along δ to define, for sufficiently small $t \in [0, 1]$, the following subsets of the fibers near the singular locus $p = q = 0$:

$$L_{1-t} := \left\{ Y \in \mathcal{Y}_{[\pi^2],[N^2],\delta(1-t)} \;\middle|\; \begin{array}{r} h_{\delta|_{[0,s]}} \text{ is defined in a neighbourhood of } Y \text{ for all } s < 1, \\ h_{\delta(s)}(Y) \xrightarrow{s \to 1} D(r) \oplus D(r) \;= 0 \end{array} \right\}$$

It follows from [Man07, Sec. 4.3] that L_t is diffeomorphic to the direct sum of two copies of $\mathbb{C}P^{N-1}$ when $\pi = [(N-1)1]$ (which corresponds to a magnetic charge $\lambda = (1, 0 \dots, 0)$ or equivalently a strand labeled by the fundamental representation). Each of the two projective spaces arises as a quotient of an ordinary vanishing cycle S^{2n-1} by an S^1 action. For more general magnetic charges λ it is expected that one obtains vanishing Grassmannians instead of vanishing projective spaces.

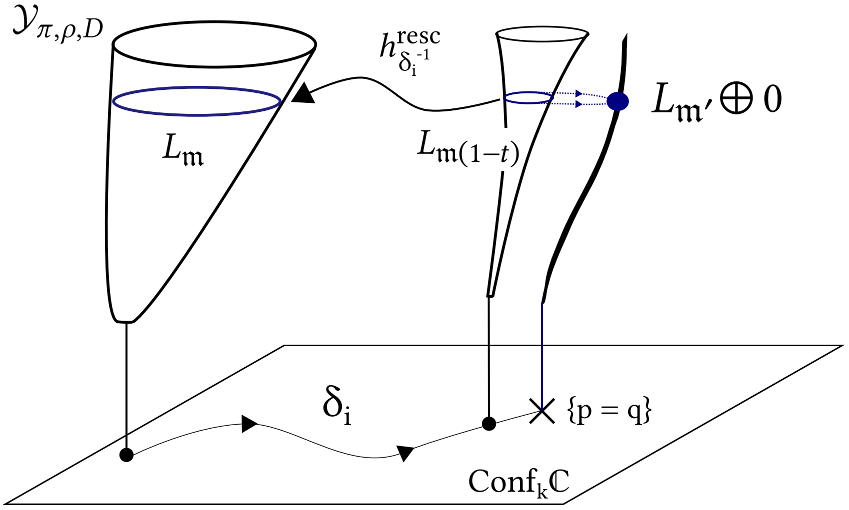

Figure 5.6 Inductive construction of the vanishing cycle $L_{\mathfrak{m}}$. For a given arc δ_i of a crossingless matching $\mathfrak{m} = (\delta_1, \dots \delta_k)$,

In any case, L_{1-t} is a Lagrangian submanifold of $\mathcal{Y}_{\pi,[N^2],\delta(1-t)}$. We then use rescaled parallel transport 'backwards' along δ to move L_{1-t} all the way to a subspace L_0 in the initial fiber $\mathcal{Y}_{[\pi^2],[N^2],\delta(0)}$.

For the induction step, assume we are given the positions of $2k$ strands $D = \{(p_a),(q_a)\}_{a=1,\dots,k} \in \mathrm{Conf}_{k,k}\,\mathbb{C}$, labeled by an admissible partition $\pi = [\pi_{p_1} \dots \pi_{q_1} \dots]$ of $2kN$, and a crossingless matching $\mathfrak{m} = (\delta_1, \dots, \delta_k)$. We are interested in the effect of matching points along $\delta_1 : [0,1] \to \overline{\mathrm{Conf}_{k,k}\,\mathbb{C}}$, so we consider $\mathfrak{m}(t) = (\delta_1(t), \delta_2, \dots \delta_k)$ with $m(0) = D$. Denote by D', π', and $\mathfrak{m}'$ the objects that are obtained from D, π, and $\mathfrak{m}$ after removing (p_1, q_1) along the arc δ_1.

Observe that we can split

$$\mathfrak{sl}(2kN,\mathbb{C})^{[N^{2k}1^{2k}]} = \mathfrak{sl}((2k-2)N,\mathbb{C})^{[N^{2k-2}1^{2k-2}]} \oplus \mathfrak{sl}(N,\mathbb{C})^{[(N-1)1]}\,.$$

The space $\mathcal{Y}_{\pi',[N^{2k-2}1^{2k-2}],D'}$ can be identified with a fiber of

$$\mathfrak{sl}((2k-2)N,\mathbb{C})^{[N^{2k-2}1^{2k-2}]} \oplus \{0\} \xrightarrow{\tilde{\chi}} \mathrm{Conf}_{k-1,k-1}\ .$$

By assumption, we have a Lagrangian $L_{\mathfrak{m}'} \subset \mathcal{Y}_{\pi',[N^{2k-2}1^{2k-2}],D'}$. Using parallel transport, we can move this Lagrangian into the singular locus of $\mathcal{Y}_{\pi,[N^{2k}1^{2k}],\mathfrak{m}(1)}$. By a relative version of the previous construction, one obtains for small t a vanishing Lagrangian $L_{\mathfrak{m}(1-t)} \subset \mathcal{Y}_{\pi,[N^{2k}],\mathfrak{m}(1-t)}$ that is diffeomorphic to $L_{\mathfrak{m}'(1-t)} \oplus L_{1-t}$ Using rescaled parallel transport along $\mathfrak{m}(t)$, we move this to the fiber over $\mathfrak{m}(0)$, which yields the desired Lagrangian $L_{\mathfrak{m}} \subset \mathcal{Y}_{\pi,\rho,D}$.

Using this construction, we attach two Lagrangian vanishing spaces $L_-, L_+ \subset \mathcal{Y}_{\pi,\rho,D}$ to a choice of crossingless matchings at $t = -L$ and $t = L$, respectively. Parallel transport of L_- along the braid provides a horizontal path in the Grothendieck-Springer fibration and any point in the intersection $h^{\mathrm{resc}}_{\beta\times\mathrm{id}} L_- \cap L_+$ corresponds to a path that connects the bottom part of the path to

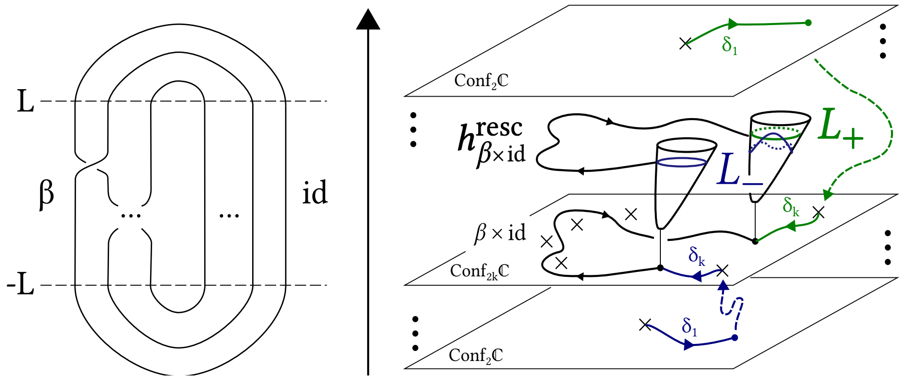

Figure 5.7 Points in the intersection $h^{\text{resc}}_{\beta\times\text{id}} L_- \cap L_+$ determine horizontal lifts of the braid closure $\overline{\beta}$. Each horizontal lift determines a way to glue solutions of the EBE to a global monopole configuration along $K \subset \partial\mathbb{R}_t \times \mathbb{C}$.

the top one. The pre-image of such a path under the bundle map Υ in (5.14) then contains at least one non-vertical family of effective triples, which in turn determines a solution of the decoupled Haydys-Witten equations.

5.10 The Floer Differential and Symplectic Khovanov Homology

In this section we describe implications of our discussion for Witten's gauge theoretic approach to Khovanov homology and formulate a second main conjecture. For this we return to the full decoupled Haydys-Witten equations on $M^5 = C \times \Sigma \times \mathbb{R}^+_y$, where $C = \mathbb{R}_s \times \mathbb{R}_t$ with holomorphic coordinate $w = s + it$, and $\Sigma = \mathbb{C}$ with holomorphic coordinate z.

Let us start with a brief reminder of Haydys-Witten Floer homology $HF_{\pi/2}(W^4)$ in the context of Khovanov homology (see Section 2.7). Assume $W^4 = \mathbb{R}_t \times \mathbb{C} \times \mathbb{R}^+_y$ and that we are given a knot $K \subset \partial W^4 = \mathbb{R}_t \times \mathbb{C}$. The Floer cochain complex is defined to be the abelian group generated by solutions of the $\theta = \pi/2$-version of the Kapustin-Witten equations that satisfy Nahm pole boundary conditions with knot singularities along K:

$$CF_{\pi/2}([W^4; K]) := \bigoplus_{x \in \mathcal{M}^{\text{KW}}(W^4)} \mathbb{Z}[x] \,.$$

Equip $M^5 = \mathbb{R}_s \times W^4$ with the non-vanishing vector field $v = \partial_y$ and extend K to the translation-invariant surface $\Sigma_K = \mathbb{R}_s \times K$ in the boundary of M^5. Then there is an associated Floer cohomology group $HF_{\pi/2}([W^4; K]) := H(CF_{\pi/2}, d_v)$ with respect to the Haydys-Witten Floer differential defined by

$$d_v x = \sum_{\mu(x,y)=1} m_{xy} \cdot y \,.$$

Here m_{xy} is the signed count of Haydys-Witten instantons on M^5, subject to the following conditions

$$\begin{cases} \mathbf{HW}_v(A,B) = 0 \\ \lim_{s\to-\infty}(A,B) = x,\ \lim_{s\to+\infty}(A,B) = y \\ (A,B) \text{ satisfy regular Nahm pole boundary conditions} \\ \qquad \text{with knot singularities along } \Sigma_K \end{cases}$$

Witten's approach to Khovanov homology is the proposal that $HF_{\pi/2}([W^4, K], v = \partial_y)$ coincides with Khovanov homology.

The manifold $M^5 = C \times \Sigma \times \mathbb{R}^+_y$ with $C = \mathbb{C}_w$ and $\Sigma = \mathbb{C}_z$ is special in two important ways. First, there are no non-trivial Vafa-Witten solutions at $y \to \infty$, such that there is an instanton grading $HF^\bullet_{\pi/2}([\mathbb{R}_t \times \mathbb{C} \times \mathbb{R}^+_y, K], v = \partial_y)$. In fact, physical considerations suggest that there actually is a bigrading, where the second grading is related to an absolute version of the Maslov index (cf. Section 2.3). Second, looking at conditions (A1) - (A4) of Theorem D and writing $B = B^{\mathrm{NP}} + b$, with component $b_z = b_2 + ib_3$, we find that in the current situation any Haydys-Witten solution that satisfies $D_{\bar{w}} b_z = \mathcal{O}(y^2)$ is already a solution of the decoupled Haydys-Witten equations. While it is unclear if there are solutions for which the order y^1 term of b_z is not holomorphic, we take this observation as strong motivation to assume that all Haydys-Witten instantons are solutions of the decoupled equations.

In Section 5.9 we have already described that solutions of the decoupled Kapustin-Witten equation on $W^4 = \mathbb{R}_t \times \mathbb{C} \times \mathbb{R}^+_y$ with knot singularities are related to intersections of certain Lagrangian subspaces of $\mathcal{Y}_{\pi,\rho,D}$. To determine the Floer differential it remains to describe Haydys-Witten instantons on M^5 that interpolate between Kapustin-Witten solutions at $s \to \pm\infty$. Observe, again, that the decoupled Haydys-Witten equations contain the extended Bogomolny equations $[\mathcal{D}_i, \mathcal{D}_j] = 0$. As before, the latter define a family of effective triples, now parametrized by C and given by a map

$$u : C \to \mathcal{M}_{(\mathcal{E},\varphi,L)} ,\ w \mapsto (\mathcal{E}(w), \varphi(w), L(w)) .$$

The remaining $\mathcal{G}_{\mathbb{C}}$-invariant equations $[\mathcal{D}_0, \mathcal{D}_i] = 0$ become equivalent to Cauchy-Riemann equations $D_{\bar{w}} u = 0$, so u is a pseudo-holomorphic disc.

We now embark on a short detour and briefly review the definition of Lagrangian Intersection Floer theory [Flo88b] (omitting virtually all technical details). Let (M, ω) be a Kähler manifold and consider closed connected Lagrangian submanifolds $L, L' \subset M$. Let O_x denote the orientation group of the point x, defined as the group that is generated by the two possible orientations of x together with the relation that their sum is zero (thus, O_x is non-canonically isomorphic to $\mathbb{Z}$). The Lagraingian Intersection cochain complex is defined to be the abelian group

$$CF(L, L') := \bigoplus_{x \in L \cap L'} O_x$$

From this one defines Lagragian Intersection Floer cohomology $HF(L, L') := H(CF(L, L'), d_J)$ as cohomology with respect to the differential

$$d_J x = \sum_{y \in L \cap L'} n_{xy} y \,.$$

Here n_{xy} is a signed count of pseudoholomorphic discs, with respect to a generic smooth family of ω-compatible almost complex structures $\{J_t\}$, that are subject to the following conditions.

$$\begin{cases} u : \mathbb{R}_s \times [0,1]_t \to M \\ \partial_s u + J_t(u)\partial_t u = 0 \\ u(s,0) \in L,\ u(s,1) \in L' \\ \lim_{s\to-\infty} u(s,\cdot) = x,\ \lim_{s\to\infty} u(s,\cdot) = y \end{cases}$$

An example of a Lagrangian Intersection Floer theory that is of particular relevance to us is symplectic Khovanov-Rozansky homology of a knot K, developed by Seidel-Smith and Manolescu [SS04; Man07]. The starting point are Lagrangian submanifolds $\tilde{L}_\pm$ of $\mathfrak{sl}(N, \mathbb{C})$ that are constructed in exactly the same way as the vanishing spaces $L_\pm$ above. However, there is a crucial difference between the fibration that is used here and the one defined more carefully by Seidel-Smith and Manolescu. From the perspective of physics it was natural to encode each of the $2k$ monopoles of the braid closure $\bar{\beta} = K$ in an individual $\mathfrak{sl}(N, \mathbb{C})$-block, leading to a construction of $L_\pm$ in $\mathfrak{sl}(2kN, \mathbb{C})$. In contrast, symplectic Khovanov-Rozansky homology is more efficient and utilizes the thin eigenvalues to encode the position of the S^1-invariant part of the braid closure $\bar{\beta}$, such that the associated Lagrangian submanifolds $\tilde{L}_\pm$ live in $\mathfrak{sl}(kN, \mathbb{C})$.

Given $\tilde{L}_\pm$, symplectic Khovanov-Rozansky for a braid β on k strands is defined as the Lagrangian Intersection Floer theory of $h^{\text{resc}}_{\beta\times\text{id}}\tilde{L}_-$ and $\tilde{L}_+ \subset \mathfrak{sl}(kN)$.

Definition 5.7 (Symplectic Khovanov Homology [SS04; Man07])**.**

$$\mathcal{H}^\bullet_{\text{symp. Kh}}(K) := HF^\bullet(h^{\text{resc}}_{\beta\times\text{id}}\tilde{L}_-, \tilde{L}_+)$$

Seidel-Smith proved that symplectic Khovanov homology is a knot invariant for $\mathfrak{sl}(2, \mathbb{C})$ and Manolescu generalized the construction and proof to $\mathfrak{sl}(N, \mathbb{C})$. It is expected that symplectic Khovanov homology coincides with a grading-reduced version of Khovanov-Rozansky homology. In fact, this has been proven for $\mathfrak{sl}(2, \mathbb{C})$ by Abouzaid and Smith:

Theorem 5.8 ([AS19])**.** *Let $\mathfrak{g} = \mathfrak{sl}(2, \mathbb{C})$. Then for any oriented link K one has an absolute grading and an isomorphism*

$$\mathcal{H}^k_{symp.\ Kh}(K) \simeq \bigoplus_{i-j=k} \mathcal{H}^{i,j}_{\text{Kh}}(K)$$

Let us now return to Haydys-Witten instanton Floer theory. Since solutions of the decoupled Kapustin-Witten equations are related to Lagrangian intersections, we conjecture that the Haydys-Witten Floer complex can be replaced by

$$CF_{\pi/2}([\mathbb{C}\times\mathbb{R}^+_y;K]) = CF(L_-,L_+)$$

Moreover, by our explanations above, the Haydys-Witten Floer differential is determined by pseudo-holomorphic discs u that satisfy the following conditions.

$$\begin{cases} u\,:\,\mathbb{R}_s\times[-L,L]_t\to\mathfrak{sl}^{[(N-1)^{2k}\;1^{2k}]} \\ \partial_s u + i\partial_t u = 0 \\ \lim_{s\to-\infty} u(s,\cdot) = x,\;\lim_{s\to\infty} u(s,\cdot) = y \\ u(s,-L)\in L_-,\;u(s,+L)\in L_+ \end{cases}$$

The last condition is satisfied by virtue of $\mathbb{R}_s$-invariance of Σ_K. In conclusion, we claim that Haydys-Witten Floer homology coincides with Lagrangian Intersection homology

$$HF_{\pi/2}\big([\mathbb{R}_t\times\mathbb{C}\times\mathbb{R}^+_y;K]\big) = HF\big(h^{\mathrm{resc}}_{\beta\times\mathrm{id}}L_-,L_+\big)\,.$$

While the right hand side is not symplectic Khovanov-Rozansky homology, the difference is only in handling the constant parts of the knot closure $\bar{\beta}=\beta\times\mathrm{id}$. Motivated by He and Mazzeo's classification of S^1-invariant Kapustin-Witten solutions, we propose that the effect of the k S^1-invariant monopoles supported on id can be neglected. Using this, we arrive at the second main conjecture of this thesis.

Conjecture F. *Haydys-Witten Floer homology of* $[\mathbb{R}_t\times\mathbb{C}\times\mathbb{R}^+_y;K]$ *coincides with symplectic Khovanov-Rozansky homology*

$$HF^\bullet_{\pi/2}\big([\mathbb{R}_t\times\mathbb{C}\times\mathbb{R}^+_y;K]\big) = \mathcal{H}^\bullet_{symp.\;Kh}(K)\,.$$

Theorem 5.9. *If Conjecture F is true, then Haydys-Witten Floer theory with* $G=SU(2)$ *coincides with a grading reduced version of Khovanov homology.*

$$HF^k_{\pi/2}\big([\mathbb{R}_t\times\mathbb{C}\times\mathbb{R}^+_y;K]\big)\simeq\bigoplus_{i-j=k}\mathrm{Kh}^{i,j}(K)\,.$$

6 Conclusion and Outlook

Witten's gauge theoretic approach to Khovanov homology has come across large interest in the mathematical community, as it provides an intriguing example of the connection between geometric analysis and topological invariants of manifolds and knots. So far, a lot of focus was put on Gaiotto and Witten's program to calculate Khovanov homology by adiabatically braiding solutions of the extended Bogomolny equations (EBE). Underlying this is the Hermitian Yang-Mills structure of the EBE, which makes the problem tractable. By now, several conjectures that arose from that work have been rigorously proven, most prominently a Kobayashi-Hitchin correspondence between solutions of the EBE and Higgs bundles with certain additional data.

The main contribution of the present work is a way to extend these ideas to the Kapustin-Witten and Haydys-Witten equations. This builds on insights into the behaviour of Kapustin-Witten and Haydys-Witten solutions in certain geometric settings. It was shown in Theorem A and B of Chapter 3, that solutions of the θ-Kapustin-Witten equations on ALE and ALF gravitational instantons satisfy a dichotomy between either some minimal asymptotic growth of the Higgs field ϕ or vanishing of $[\phi \wedge \phi]$ everywhere. By prior work of Nagy and Oliveira, these results immediately implied the vanishing theorem Corollary C for finite energy solutions. The latter was used in Chapter 4 to obtain control over the boundary behaviour of Haydys-Witten solutions at cylindrical ends at infinity. Similar control at boundaries was achieved by a detailed analysis of polyhomogeneous expansions of θ-Kapustin-Witten solutions with twisted Nahm pole boundary conditions. Taken together, this resulted in a number of sufficient conditions for Theorem D, establishing that, in certain geometric situations, the Haydys-Witten equations simplify to a decoupled version of the equations that exhibit a Hermitian Yang-Mills structure.

While the vanishing results of Chapter 3 and Chapter 4 might be of independent interest, their appearance in this thesis arose due to the associated reduction in complexity of Haydys-Witten instantons in the context of Khovanov homology. This is the content of Chapter 5, which laid out how the Hermitian Yang-Mills structure of the decoupled equations may be exploited to generalize the known Kobayashi-Hitchin equivalence for the EBE to a correspondence between Haydys-Witten instanton homology and symplectic Khovanov homology. Inspired by Gaiotto and Witten's adiabatic braiding of EBE-solutions, a relation between decoupled Kapustin-Witten solutions and non-vertical paths in the moduli space of EBE-solutions was proposed. By a physically motivated argument, it was suggested that one can replace the moduli space of EBE-solutions by a finite-dimensional model fibration known as Grothendieck-Springer resolution of $\mathfrak{sl}(N, \mathbb{C})$. These ideas culminated, on the one hand, in Conjecture E, which claims that the number of decoupled Kapustin-Witten solutions on $S^1_t \times \mathbb{C} \times \mathbb{R}^+_y$ is bounded from below by the number of intersections of the fiber of this space with its parallel transport along the knot. On the other hand, a stronger variant of this conjecture on $\mathbb{R}_t \times \mathbb{C} \times \mathbb{R}^+_y$ was formulated, stating that Haydys-Witten Floer theory is isomorphic to symplectic Khovanov-Rozansky homology. Since for $G = SU(2)$ the latter is known to coincide with a grading-reduced version of Khovanov homology, this provides a novel approach to Witten's conjecture.

Although the results of this thesis add to a rigorous understanding of Haydys-Witten and Kapustin-Witten solutions and provide several novel ideas regarding Gaiotto-Witten's adiabatic approach, many key properties of Haydys-Witten Floer theory and its connection to Khovanov homology remain elusive. Most notably, any problems regarding compactness or gluing theorems, which lie at the very heart of a well-defined instanton Floer theory, were thoroughly ignored. Also, much of the discussion in Chapter 5 remains speculative and passes only at the physics-level of rigor. Consequently, it remains an open problem to classify Kapustin-Witten and Haydys-Witten solutions, let alone calculate Haydys-Witten instanton homology, using the adiabatic approach.

One interesting question arose in Chapter 5 for Haydys-Witten solutions (A, B) on $C \times \Sigma \times \mathbb{R}^+_y$ with $C = \mathbb{C}_w$ and $\Sigma = \mathbb{C}_z$. In describing the relation of Haydys-Witten instanton homology with symplectic Khovanov homology, it was assumed that if one writes $B = B^{\mathrm{NP}} + b$, the function $b_z = b_2 + ib_3 \in \mathcal{O}(y^1)$ is holomorphic in w at leading order, i.e. $D_{\bar{w}} b_z = \mathcal{O}(y^2)$. This condition is a part of assumption (A1) in Theorem D. If this is false, then there might exist Haydys-Witten and Kapustin-Witten solutions that are not captured by symplectic Khovanov homology.

Another interesting avenue in that context would be a detailed investigation of assumption (A2), which is concerned with the asymptotics of β-twisted knot singularity conditions. Although the corresponding model solutions are only known implicitly, it might be possible to use Dimakis' continuity argument to also determine the asymptotic behaviour of χ near $\partial_K M$.

Results like this, which clarify the properties of twisted Nahm pole boundary conditions with knot singularities, would open up an avenue to directly investigate Haydys-Witten homology groups $HF_\theta(W^4)$ for angles $\theta \neq \pi/2$ in the context of knot homologies. More generally, it would be interesting to further investigate the maps $HF_\theta(W^4) \to HF_{\theta'}(W^4)$, since the physical realization of these groups seem to suggest that this is a so-far unknown example of wall-crossing of BPS states.

Clearly, much remains to be done to prove Witten's conjecture and fully understand the connections between Haydys-Witten Floer homology and knot homologies. Nonetheless, the present work also offers avenues for future research in directions beyond the analytical foundations. If one assumes that there is an isomorphism between Haydys-Witten instanton homology and symplectic Khovanov homology, as proposed in Conjecture F, then there is something to be said about aspects of knot homologies that have not yet been addressed in the gauge theoretic approach.

As was mentioned in the introduction, a conjectural realization of Khovanov homology in a physical system was achieved in [GSV05] by using a duality between Chern-Simons theory and topological string theory. The conjecture states that Khovanov homology coincides with the Hilbert spaces of BPS states in topological string theory. In fact, a few years later it was shown by [LM05] that the HOMFLY-PT polynomial arises from this in a large N limit.

Although Witten's approach and the earlier result of [GSV05] are now known to be closely related, as described in [Das+16], there remain several open questions. The most obvious missing link is that it is not at all obvious how to describe HOMFLY-PT homology in Witten's approach. From the gauge theory perspective of Haydsy-Witten Floer theory, the large N limit must be related via the AdS/CFT correspondence to a supergravity theory. Since a categorification on the gauge side must correspond to a categorification on the gravity side, this line of argument offers a way to investigate a categorification of topologically twisted AdS/CFT correspondence (cf. [CL16]). This would lift AdS/CFT correspondence to an isomorphism of Hilbert spaces rather than an equality of partition functions.

Another subject that is raised by the large N behaviour of the topological string BPS states is the emergence of 'stability' at large N. Over the years this has lead to a number of conjectures relating HOMFLY-PT homology to $\mathfrak{sl}_N$-homologies.

Conjecture ([DGR05]). *There is a triply graded homology theory $\mathcal{H}^\lambda(K)$ categorifying the HOMFLY-PT polynomial, coming with differentials $\{\, d_N \,\}_{N\in\mathbb{Z}}$ satisfying the axioms*

(i) Grading: d_N is of degree $(-2, 2N, -1)$ for $N > 0$; d_0 is of degree $(-2, 0, -3)$; d_N is of degree $(-2, 2N, -1 + 2N)$ for $N < 0$.

(ii) Anticommutativity: $\{d_N, d_M\} = 0$ for all $N, M \in \mathbb{Z}$.

(iii) Symmetry: There is an involution $f : \mathcal{H}^\lambda_{i,j,*}(K) \to \mathcal{H}^\lambda_{i,-j,*}(K)$ such that $f d_N = d_{-N} f$.

For all $N > 0$ the homology with respect to d_N is isomorphic to coloured $\mathfrak{sl}(N, \mathbb{C})$ Khovanov-Rozansky homology:

$$H_*\Bigg(\bigoplus_{iN+j=p} \mathcal{H}^\lambda_{i,j,k}(K)\,, d_N\Bigg) = KhR^{\lambda,\mathfrak{sl}_N}(K)\,.$$

Similarly for $N = 0$ the homology of $(\bigoplus_i \mathcal{H}^\lambda_{i,j,k}, d_0)$ is isomorphic to knot Floer homology.

This conjecture was weakened by Rasmussen [Ras16], who argued that a priori one should expect to find spectral sequences instead of differentials. In fact he was able to proof the following theorem.

Theorem ([Ras16]). *For each $N > 0$, there is a spectral sequence $E_k(N)$ that starts at $\mathcal{H}^\lambda(K)$ and converges to $KhR^{\lambda,\mathfrak{sl}_N}(K)$. The higher terms in this sequence are invariants of the knot K.*

Note however, that the spectral sequences of Rasmussen do not explain the negative differentials d_{-N} that are expected to be present as well. Furthermore, in all known examples the spectral sequence abuts at the second page, which is consistent with the existence of differentials d_N.

Assuming Conjecture F is true, there should be a way to recover Rasmussen's spectral sequences in gauge theory. From the physics perspective, going from large N to some finite N is related to 'Higgsing' a large number of the components of the gauge connection, until only $SU(N)$ gauge symmetry is left over. It might be possible to construct in that way differentials d_N, providing a deeper reason for the observation above. However, to make this work one would first need to find a description of the triply-graded homology $\mathcal{H}^{\lambda}(K)$ in gauge theory.

In summary, although the Kapustin-Witten and Haydys-Witten equations have been studied extensively over the last decade, Haydys-Witten Floer theory is still at an early stage and promises many new and deep insights into the relation between geometric analysis and topology.

List of Publications

Versions of Chapter 3 and 4 have been made publicly available as the following preprints.

[Ble23a] Bleher, M. (2023a). *Growth of the Higgs Field for Kapustin-Witten Solutions on ALE and ALF Gravitational Instantons.* arXiv: 2306.17017 [math-DG].

[Ble23b] — (2023b). *The Decoupled Haydys-Witten Equations and a Weitzenböck Formula.* arXiv: 2307.15056 [math-ph].

Bibliography

[Aro57] Aronszajn, N. (1957). 'A Unique Continuation Theorem for Solutions of Elliptic Partial Differential Equations or Inequalities of Second Order'. *J. Math. Pures Appl.* 9.36, pp. 235–249.

[Whi57] Whitney, H. (1957). *Geometric Integration Theory.* Princeton University Press. ISBN: 978-1-4008-7757-7.

[Cal58] Calabi, E. (1958). 'An Extension of E. Hopf's Maximum Principle with an Application to Riemannian Geometry'. *Duke Math. J.* 25.1.

[Bog76] Bogomolny, E. B. (1976). 'Stability of Classical Solutions'. *Sov. J. Nucl. Phys.* 24, p. 449.

[Yau76] Yau, S.-T. (1976). 'Some Function-Theoretic Properties of Complete Riemannian Manifold and Their Applications to Geometry'. *Indiana University Mathematics Journal* 25.7, pp. 659–670.

[Ati+78] Atiyah, M. F., N. J. Hitchin, V. G. Drinfeld and Y. I. Manin (1978). 'Construction of Instantons'. *Phys. Lett. A* 65.3, pp. 185–187.

[Alm79] Almgren Jr, F. J. (1979). 'Dirichlet's Problem for Multiple Valued Functions and the Regularity of Mass Minimizing Integral Currents'. In: *Minimal Submanifolds and Geodesics (Proc. Japan-united States Sem., Tokyo, 1977).* Vol. 1. 6.

[JT80] Jaffe, A. and C. Taubes (1980). *Vortices and Monopoles: Structure of Static Gauge Theories.* Boston: Birkhäuser. ISBN: 978-3-7643-3025-5.

[Li82] Li, P. (1982). 'Poincaré Inequalities on Riemannian Manifolds'. In: *Seminar on Differential Geometry.* Annals of Mathematics Studies no. 102. Princeton University Press, pp. 73–84.

[Uhl82a] Uhlenbeck, K. K. (1982a). 'Connections with L^p Bounds on Curvature'. *Communications in Mathematical Physics* 83.1, pp. 31–42.

[Uhl82b] — (1982b). 'Removable Singularities in Yang-Mills Fields'. *Communications in Mathematical Physics* 83.1, pp. 11–29.

[Wit82] Witten, E. (1982). 'Supersymmetry and Morse Theory'. *J. Differ. Geom.* 17.4, pp. 661–692.

[Hit83] Hitchin, N. J. (1983). 'On the Construction of Monopoles'. *Commun. Math. Phys.* 89, pp. 145–190.

[Nah83] Nahm, W. (1983). *All Self-Dual Multimonopoles for Arbitrary Gauge Groups.* 10 pp.

[Don84] Donaldson, S. K. (1984). 'Nahm's Equations and the Classification of Monopoles'. *Commun.Math. Phys.* 96.3, pp. 387–407.

[Don85] — (1985). 'Anti Self-Dual Yang-Mills Connections Over Complex Algebraic Surfaces and Stable Vector Bundles'. *Proceedings of the London Mathematical Society* s3-50.1, pp. 1–26.

[Fre+85] Freyd, P., D. Yetter, J. Hoste, W. B. R. Lickorish, K. Millett and A. Ocneanu (1985). 'A New Polynomial Invariant of Knots and Links'. *Bulletin (New Series) of the American Mathematical Society* 12.2, pp. 239–246.

[Jon85] Jones, V. F. R. (1985). 'A Polynomial Invariant for Knots via von Neumann Algebras'. *Bull. Amer. Math. Soc.* 12.1, pp. 103–111.

[LY86] Li, P. and S. T. Yau (1986). 'On the Parabolic Kernel of the Schrödinger Operator'. *Acta Math.* 156.0, pp. 153–201.

[UY86] Uhlenbeck, K. and S. T. Yau (1986). 'On the Existence of Hermitian-Yang-Mills Connections in Stable Vector Bundles'. *Communications on Pure and Applied Mathematics* 39.S1, S257–S293.

[Don87a] Donaldson, S. K. (1987a). 'Infinite Determinants, Stable Bundles and Curvature'. *Duke Math. J.* 54.1.

[Don87b] — (1987b). 'Twisted Harmonic Maps and the Self-Duality Equations'. *Proceedings of the London Mathematical Society* s3-55.1, pp. 127–131.

[Hit87a] Hitchin, N. J. (1987a). 'The Self-Duality Equations on a Riemann Surface'. *Proceedings of the London Mathematical Society* s3-55.1, pp. 59–126.

[Hit87b] Hitchin, N. (1987b). 'Stable Bundles and Integrable Systems'. *Duke Math. J.* 54.1.

[AH88] Atiyah, M. F. and N. Hitchin (1988). *The Geometry and Dynamics of Magnetic Monopoles.* Princeton University Press. ISBN: 978-1-4008-5930-6.

[Cor88] Corlette, K. (1988). 'Flat *G*-Bundles with Canonical Metrics'. *Journal of Differential Geometry* 28.3, pp. 361–382.

[Flo88a] Floer, A. (1988a). 'An Instanton-Invariant for 3-Manifolds'. *Communications in Mathematical Physics* 118.2, pp. 215–240.

[Flo88b] — (1988b). 'Morse Theory for Lagrangian Intersections'. *Journal of Differential Geometry* 28.3, pp. 513–547.

[PT88] Przytycki, J. and P. Traczyk (1988). 'Invariants of Links of Conway Type'. *Kobe J. Math.* 4.2, pp. 115–139.

[Wit88a] Witten, E. (1988a). 'Topological Quantum Field Theory'. *Commun. Math. Phys.* 117.3, pp. 353–386.

[Wit88b] — (1988b). 'Topological Sigma Models'. *Commun. Math. Phys.* 118, pp. 411–449.

[Yam88] Yamron, J. P. (1988). 'Topological Actions from Twisted Supersymmetric Theories'. *Phys. Lett. B* 213.3, pp. 325–330.

[Flo89] Floer, A. (1989). 'Witten's Complex and Infinite Dimensional Morse Theory'. *J. Differ. Geom.* 30, pp. 207–221.

[LR89] Labastida, J. M. F. and A. V. Ramallo (1989). 'Operator Formalism for Chern-Simons Theories'. *Physics Letters B* 227.1, pp. 92–102.

[Wit89] Witten, E. (1989). 'Quantum Field Theory and the Jones Polynomial'. *Commun. Math. Phys.* 121.3, pp. 351–399.

[Mel90] Melrose, R. B. (1990). 'Pseudodifferential Operators, Corners and Singular Limits'. *Proc. Int. Congr. Math.*, pp. 217–234.

[Sim90] Simpson, C. T. (1990). 'Nonabelian Hodge Theory'. *Proceedings of the International Congress of Mathematicians*, pp. 747–756.

[LLR91] Labastida, J. M. F., P. M. Llatas and A. V. Ramallo (1991). 'Knot Operators in Chern-Simons Gauge Theory'. *Nuclear Physics B* 348.3, pp. 651–692.

[Maz91] Mazzeo, R. (1991). 'Elliptic Theory of Differential Edge Operators I'. *Commun. Partial Differ. Equations* 16.10, pp. 1615–1664.

[Wit91] Witten, E. (1991). 'Introduction To Cohomological Field Theories'. *Int. J. Mod. Phys. A* 06.16, pp. 2775–2792.

[Sim92] Simpson, C. T. (1992). 'Higgs Bundles and Local Systems'. *Publications Mathématiques de l'Institut des Hautes Scientifiques* 75.1, pp. 5–95.

[CM93] Collingwood, D. H. and W. M. McGovern (1993). *Nilpotent Orbits in Semisimple Lie Algebras*. New York: Van Nostrand Reinhold. 186 pp. ISBN: 978-0-534-18834-4.

[GT93] Grabowski, M. P. and C.-H. Tze (1993). 'On the Octonionic Nahm Equations and Self-Dual Membranes in 9 Dimensions'. In: *Symmetries in Science VI: From the Rotation Group to Quantum Algebras*. Ed. by B. Gruber. Boston, MA: Springer US, pp. 287–297.

[VW94] Vafa, C. and E. Witten (1994). 'A Strong Coupling Test of S-Duality'. *Nuclear Physics B* 431.1-2, pp. 3–77. arXiv: hep-th/9408074.

[Kat95] Kato, T. (1995). *Perturbation Theory for Linear Operators*. Repr. of the 1980 ed. Classics in Mathematics. Berlin Heidelberg: Springer. ISBN: 978-3-540-58661-6.

[Mar95] Marcus, N. (1995). 'The Other Topological Twisting of $\mathcal{N} = 4$ Yang-Mills'. *Nuclear Physics B* 452.1-2, pp. 331–345. arXiv: hep-th/9506002.

[Mel95] Melrose, R. B. (1995). *Geometric Scattering Theory*. Stanford Lectures. Cambridge, Eng. ; New York: Cambridge University Press.

[Fan96] Fan, H. (1996). 'Half De Rham Complexes and Line Fields on Odd-Dimensional Manifolds'. *Trans. Amer. Math. Soc.* 348.8, pp. 2947–2982.

[Mel96] Melrose, R. B. (1996). *Differential Analysis on Manifolds with Corners*. URL: http://math.mit.edu/~rbm/book.html (visited on 30/08/2023).

[LL97] Labastida, J. M. F. and C. Lozano (1997). 'Mathai-Quillen Formulation of Twisted $\mathcal{N} = 4$ Supersymmetric Gauge Theories in Four Dimensions'. *Nuclear Physics B* 502.3, pp. 741–790. arXiv: hep-th/9702106.

[Del+99] Deligne, P., D. Kazhdan, P. Etingof, J. W. Morgan, D. S. Freed, D. R. Morrison, L. C. Jeffrey and E. Witten (1999). *Quantum Fields And Strings: A Course for Mathematicians 1*. Vol. 1. 2 vols. Providence, RI: Oxford University Press. 1501 pp. ISBN: 978-0-8218-2014-8.

[Kho99] Khovanov, M. (1999). 'A Categorification of the Jones Polynomial'. *Duke Math. J.* 101, p. 51.

[OS03] Ozsvath, P. and Z. Szabo (2003). 'Holomorphic Disks and Knot Invariants'. arXiv: math/0209056.

[Ras03] Rasmussen, J. (2003). 'Floer Homology and Knot Complements'. arXiv: math/0306378.

[SS04] Seidel, P. and I. Smith (2004). 'A Link Invariant from the Symplectic Geometry of Nilpotent Slices'. arXiv: math/0405089.

[DGR05] Dunfield, N. M., S. Gukov and J. Rasmussen (2005). 'The Superpolynomial for Knot Homologies'. *Exp. Math.* 15.2, pp. 129–159.

[GSV05] Gukov, S., A. Schwarz and C. Vafa (2005). 'Khovanov-Rozansky Homology and Topological Strings'. *Lett. Math. Phys.* 74.1, pp. 53–74.

[GW05] Gukov, S. and J. Walcher (2005). 'Matrix Factorizations and Kauffman Homology'. arXiv: hep-th/0512298.

[LM05] Labastida, J. M. F. and M. Marino (2005). *Topological Quantum Field Theory and Four Manifolds.* Mathematical Physics Studies 25. Dordrecht : Norwell, MA: Springer. ISBN: 978-1-4020-3058-1.

[Von05] Vonk, M. (2005). 'A Mini-Course on Topological Strings'. *Physics (College. Park. Md).*, p. 120.

[KW07] Kapustin, A. and E. Witten (2007). 'Electric-Magnetic Duality and the Geometric Langlands Program'. *Commun. Number Theory Phys.* 1.1, pp. 1–236.

[Man07] Manolescu, C. (2007). 'Link Homology Theories from Symplectic Geometry'. *Advances in Mathematics* 211.1, pp. 363–416.

[KR08a] Khovanov, M. and L. Rozansky (2008a). 'Matrix Factorizations and Link Homology'. *Fundamenta Mathematicae* 199.1, pp. 1–91.

[KR08b] — (2008b). 'Matrix Factorizations and Link Homology II'. *Geom. Topol.* 12.3, pp. 1387–1425.

[GW09a] Gaiotto, D. and E. Witten (2009a). 'S-Duality of Boundary Conditions in $\mathcal{N} = 4$ Super Yang-Mills Theory'. *Adv. Theor. Math. Phys.* 13.3, pp. 721–896.

[GW09b] Gaiotto, D. and E. Witten (2009b). 'Supersymmetric Boundary Conditions in $\mathcal{N} = 4$ Super Yang-Mills Theory'. *J. Stat. Phys.* 135.5-6, pp. 789–855.

[Mar10] Mares, B. A. (2010). 'Some Analytic Aspects of Vafa-Witten Twisted $\mathcal{N} = 4$ Supersymmetric Yang-Mills Theory'. Thesis. Massachusetts Institute of Technology. URL: https://dspace.mit.edu/handle/1721.1/64488 (visited on 09/02/2022).

[Wit10] Witten, E. (2010). 'Analytic Continuation Of Chern-Simons Theory'. *AMS/IP Stud. Adv. Math.* 50, pp. 347–446.

[GMN11] Gaiotto, D., G. W. Moore and A. Neitzke (2011). 'Wall-Crossing, Hitchin Systems, and the WKB Approximation'. arXiv: 0907.3987 [hep-th].

[GS11] Gukov, S. and M. Stosic (2011). 'Homological Algebra of Knots and BPS States'. arXiv: 1112.0030.

[Wit11a] Witten, E. (2011a). 'Fivebranes and Knots'. *Quantum Topol.* 3.1, pp. 1–137.

[Wit11b] — (2011b). 'Khovanov Homology And Gauge Theory Gradient Flow Of Complex Chern-Simons'. arXiv: 1108.3103.

[AS12] Aganagic, M. and Shakirov (Feb. 2012). 'Refined Chern-Simons Theory and Knot Homology'. URL: http://arxiv.org/abs/1202.2489.

[GU12] Gagliardo, M. and K. Uhlenbeck (2012). 'Geometric Aspects of the Kapustin–Witten Equations'. *J. Fixed Point Theory Appl.* 11.2, pp. 185–198.

[GW12] Gaiotto, D. and E. Witten (2012). 'Knot Invariants from Four-Dimensional Gauge Theory'. *Adv. Theor. Math. Phys.* 16.3, pp. 935–1086.

[Hen12] Henningson, M. (2012). 'Boundary Conditions for Geometric-Langlands Twisted $\mathcal{N} = 4$ Supersymmetric Yang-Mills Theory'. *Phys. Rev. D* 86.8, p. 085003.

[Mik12] Mikhaylov, V. (2012). 'On the Solutions of Generalized Bogomolny Equations'. *J. High Energy Phys.* 2012.5, p. 112.

[Ell13] Elliott, C. (2013). *Twisting Supersymmetric Field Theories.* URL: http://www.ihes.fr/~celliott/Topological_twisting.pdf (visited on 30/08/2023).

[GGS13] Gorsky, E., S. Gukov and M. Stosic (2013). 'Quadruply-Graded Colored Homology of Knots'. arXiv: 1304.3481.

[MV13] Mazzeo, R. and B. Vertman (2013). 'Elliptic Theory of Differential Edge Operators, II: Boundary Value Problems'. arXiv: 1307.2266.

[Tau13] Taubes, C. H. (2013). 'Compactness Theorems for $SL(2;\mathbb{C})$ Generalizations of the 4-Dimensional Anti-Self Dual Equations'. arXiv: 1307.6447 [math].

[MW14] Mazzeo, R. and E. Witten (2014). 'The Nahm Pole Boundary Condition'. *Contemporary Mathematics* The influence of Solomon Lefschetz in geometry and topology.621, p. 171. arXiv: 1311.3167 [hep-th].

[Tau14] Taubes, C. H. (2014). 'The Zero Loci of $\mathbb{Z}/2$ Harmonic Spinors in Dimension 2, 3 and 4'. arXiv: 1407.6206 [math].

[Che15] Cherkis, S. A. (2015). 'Octonions, Monopoles, and Knots'. *Lett. Math. Phys.* 105.5, pp. 641–659.

[Hay15] Haydys, A. (2015). 'Fukaya-Seidel Category and Gauge Theory'. *Journal of Symplectic Geometry* 13.1, pp. 151–207.

[He15] He, S. (2015). 'Rotationally Invariant Singular Solutions to the Kapustin-Witten Equations'. arXiv: 1510.07706 [math].

[CL16] Costello, K. and S. Li (2016). 'Twisted Supergravity and Its Quantization'. arXiv: 1606.00365 [hep-th].

[Das+16] Dasgupta, K., V. Errasti Díez, P. Ramadevi and R. Tatar (2016). 'Knot Invariants and M-Theory I: Hitchin Equations, Chern-Simons Actions, and the Surface Operators'. arXiv: 1608.05128.

[Ras16] Rasmussen, J. (2016). 'Some Differentials on Khovanov–Rozansky Homology'. *Geom. Topol.* 19.6, pp. 3031–3104.

[DF17] Daemi, A. and K. Fukaya (2017). 'Atiyah-Floer Conjecture: A Formulation, a Strategy to Prove and Generalizations'. arXiv: 1707.03924 [math].

[Guk+17] Gukov, S., D. Pei, P. Putrov and C. Vafa (2017). 'BPS Spectra and 3-Manifold Invariants'. arXiv: 1701.06567.

[MW17] Mazzeo, R. and E. Witten (2017). 'The KW Equations and the Nahm Pole Boundary Condition with Knots'. arXiv: 1712.00835 [hep-th].

[Tan17] Tanaka, Y. (2017). 'Some Boundedness Properties of Solutions to the Vafa-Witten Equations on Closed Four-Manifolds'. arXiv: 1308.0862 [math].

[Tau17a] Taubes, C. H. (2017a). 'Growth of the Higgs Field for Solutions to the Kapustin-Witten Equations on $\mathbb{R}^4$'. arXiv: 1701.03072 [math-ph].

[Tau17b] — (2017b). 'The Behavior of Sequences of Solutions to the Vafa-Witten Equations'. arXiv: 1702.04610 [math].

[ESW18] Eager, R., I. Saberi and J. Walcher (2018). 'Nilpotence Varieties'. arXiv: 1807.03766 [hep-th].

[He18] He, S. (2018). 'The Expansions of the Nahm Pole Solutions to the Kapustin-Witten Equations'. arXiv: 1808.03886 [math].

[Tau18] Taubes, C. H. (2018). *Sequences of Nahm Pole Solutions to the* $SU(2)$ *Kapustin-Witten Equations*. arXiv: 1805.02773.

[AS19] Abouzaid, M. and I. Smith (2019). 'Khovanov Homology from Floer Cohomology'. *J. Amer. Math. Soc.* 32.1, pp. 1–79.

[He19] He, S. (2019). 'A Gluing Theorem for the Kapustin-Witten Equations with a Nahm Pole'. *Journal of Topology* 12.3, pp. 855–915. arXiv: 1707.06182 [math].

[HM19a] He, S. and R. Mazzeo (2019a). 'Classification of Nahm Pole Solutions of the Kapustin-Witten Equations on $S^1 \times \Sigma \times \mathbb{R}^+$'. arXiv: 1901.00274.

[HM19b] — (2019b). 'Opers and the Twisted Bogomolny Equations'. arXiv: 1902.11183.

[HM19c] — (2019c). 'The Extended Bogomolny Equations and Generalized Nahm Pole Boundary Condition'. *Geom. Topol.* 23.5, pp. 2475–2517.

[NO19a] Nagy, Á. and G. Oliveira (2019a). 'Complex Monopoles I: The Haydys Monopole Equation'. arXiv: 1906.05432 [math-ph].

[NO19b] — (2019b). 'Complex Monopoles II: The Kapustin–Witten Monopole Equation'. arXiv: 1906.05435 [math].

[Tan19] Tanaka, Y. (2019). 'On the Singular Sets of Solutions to the Kapustin–Witten Equations and the Vafa–Witten Ones on Compact Kähler Surfaces'. *Geom Dedicata* 199.1, pp. 177–187.

[Tau19] Taubes, C. H. (2019). *The* $\mathbb{R}$*-Invariant Solutions to the Kapustin-Witten Equations on* $(0,\infty) \times \mathbb{R}^2 \times \mathbb{R}$ *with Generalized Nahm Pole Asymptotics*. arXiv: 1903.03539.

[AM20] Abouzaid, M. and C. Manolescu (2020). 'A Sheaf-Theoretic Model for $SL(2,\mathbb{C})$ Floer Homology'. *J. Eur. Math. Soc.* 22.11, pp. 3641–3695.

[Aga20] Aganagic, M. (2020). 'Knot Categorification from Mirror Symmetry, Part I: Coherent Sheaves'. arXiv: 2004.14518 [hep-th].

[Fre+20] Fredrickson, L., R. Mazzeo, J. Swoboda and H. Weiss (2020). 'Asymptotic Geometry of the Moduli Space of Parabolic $SL(2,\mathbb{C})$-Higgs Bundles'. arXiv: 2001.03682 [hep-th].

[He20] He, S. (2020). 'On the Moduli Space of the Octonionic Nahm's Equations'. arXiv: 2006.11625 [math].

[HM20] He, S. and R. Mazzeo (2020). 'The Extended Bogomolny Equations with Generalized Nahm Pole Boundary Conditions, II'. *Duke Mathematical Journal* 169.12, pp. 2281–2335. arXiv: 1806.06314.

[Aga21] Aganagic, M. (2021). 'Knot Categorification from Mirror Symmetry, Part II: Lagrangians'. arXiv: 2105.06039 [hep-th].

[NO21] Nagy, Á. and G. Oliveira (2021). 'The Kapustin–Witten Equations on ALE and ALF Gravitational Instantons'. *Lett Math Phys* 111.4, p. 87. arXiv: 1906.05435 [math].

[Tau21] Taubes, C. H. (2021). *The Existence of Instanton Solutions to the* $\mathbb{R}$*-Invariant Kapustin-Witten Equations on* $(0,\infty)\times\mathbb{R}^2\times\mathbb{R}$. arXiv: 2102.04290.

[CN22] Charbonneau, B. and Á. Nagy (2022). 'On the Construction of Monopoles with Arbitrary Symmetry Breaking'. arXiv: 2205.15246 [hep-th, physics:math-ph].

[Dim22a] Dimakis, P. (2022a). 'Model Knot Solutions for the Twisted Bogomolny Equations'. arXiv: 2204.06617 [math-ph].

[Dim22b] — (2022b). 'The Extended Bogomolny Equations on $\mathbb{R}^2\times\mathbb{R}^+$ I: Nilpotent Higgs Field'. arXiv: 2209.11101 [math-ph].

[GSY22] Gukov, S., A. Sheshmani and S.-T. Yau (2022). '3-Manifolds and Vafa-Witten Theory'. arXiv: 2207.05775 [hep-th].

[OT22] Ong, Z.-C. and M.-C. Tan (2022). 'Vafa-Witten Theory: Invariants, Floer Homologies, Higgs Bundles, a Geometric Langlands Correspondence, and Categorification'. arXiv: 2203.17115 [hep-th].

[Sun23] Sun, W. (2023). *The Extended Bogomolny Equations on* $\mathbb{R}^2\times\mathbb{R}^+$ *with Real Symmetry Breaking*. arXiv: 2306.09577 [math-ph].