BPS geometry

Andrew Neitzke

Preliminary notes

These are preliminary notes for a Les Houches lecture series. Please send corrections or improvements to andrew.neitzke@yale.edu.

1. Lecture 1

These talks are about geometry associated to $\mathcal{N} = 2$ supersymmetric field theory, mostly related to the notion of *BPS state*.[1]

A rough plan of the talks:

(1) Coulomb branches of $\mathcal{N} = 2$ theories; BPS states; wall-crossing formula; the examples of class $\mathcal{S}$
(2) Interpretation of the wall-crossing formula in terms of algebras of line defects
(3) Line defect vevs as distinguished cluster-type coordinates; the TBA for line defect vevs; consequences for hyperkahler metrics
(4) q-deformation (briefly)
(5) Conformal blocks

This first lecture concerns some old facts and almost-facts about $\mathcal{N} = 2$ supersymmetric quantum field theories (QFT) in four dimensions. I'll give my understanding of the "standard" picture of the physics on their Coulomb branch, as developed by very many authors beginning from [1, 2]. A useful general pedagogical reference is [3].

1.1. Data that define $\mathcal{N} = 2$ field theories.

A remark: Mostly we will be talking about some concrete geometric structures, which are downstream from the QFT. So to get the main content of the lectures it's not strictly necessary to know what an $\mathcal{N} = 2$ supersymmetric QFT in four dimensions is. But at least you should know what kind of data determines one. The idea is that, starting from any of these data, we will get all the geometric structures we are going to discuss.

Example 1.1 (Pure nonabelian gauge theory). Fix a compact simple group G, say $G = \mathrm{SU}(N)$, and a parameter $\Lambda \in \mathbb{C}^\times$ with dimensions of mass. These data determine an $\mathcal{N} = 2$ supersymmetric theory, in a conventional way (write down a space of fields and a Lagrangian, then "quantize.") The space of fields includes a G-connection on the 4-dimensional spacetime as in Nekrasov's lectures, plus new ingredients found only in the supersymmetric theory: a $\mathfrak{g}_\mathbb{C}$-valued scalar field Φ, and fermion fields which I won't write.

The example of pure nonabelian gauge theory with $G = \mathrm{SU}(2)$ will be a running example throughout these notes.

Example 1.2 (Pure abelian gauge theory). Fix a compact abelian group $T = U(1)^r$, and a coupling matrix $(\tau_{IJ})^r_{I,J=1}$, symmetric, with $\mathrm{Im}\,\tau$ positive definite. These data also determine an

[1] "BPS" stands for Bogomolny-Prasad-Sommerfield.

$\mathcal{N} = 2$ supersymmetric theory, similar to the above. (The coupling matrix enters the action as

$$S = \frac{\tau_{IJ}}{4\pi} \int F^{I+} \wedge F^{J+} + \cdots$$

where F^I means the curvature of the I-th $U(1)$ gauge field, and $+$ means the self-dual part.)

You should have the idea that the abelian theory is much simpler than the nonabelian one.

Not every $\mathcal{N} = 2$ theory is determined by a Lagrangian. Here are a few other kinds of data which are supposed to determine $\mathcal{N} = 2$ theories.

Example 1.3 (Class $\mathcal{S}$ theory). Fix a nonsingular Riemann surface C, and a Lie algebra $\mathfrak{g}$ of ADE type. (Also fix a Lagrangian subgroup of $H_1(C, Z(G))$ where G is the simply connected form of $\mathfrak{g}$, but this is a detail better ignored most of the time.) These data determine an $\mathcal{N} = 2$ supersymmetric theory, called a "theory of class $\mathcal{S}$."

The class $\mathcal{S}$ theories with $\mathfrak{g} = A_1$ will supply another of our main running examples.

Example 1.4 (Calabi-Yau compactification). Fix a complete non-compact Calabi-Yau 3-fold X. (And perhaps some additional discrete data.) This also determines an $\mathcal{N} = 2$ supersymmetric theory, in two different ways, by "geometric engineering."

The various classes overlap quite a bit; when they do, sometimes the different constructions lead to different insights into the theory.

1.2. **Outputs: executive summary.** In the rest of this lecture + beginning of the following one, I am going to describe the following *outputs* of an $\mathcal{N} = 2$ supersymmetric QFT:

- Coulomb branch: a complex manifold $\mathcal{B}$, with an open dense subset $\mathcal{B}_{\text{reg}}$.
- Charge lattice: a local system $\Gamma \to \mathcal{B}_{\text{reg}}$ of lattices, with a skew pairing $\langle \cdot, \cdot \rangle$.
- Central charge functions: homomorphisms $Z : \Gamma_u \to \mathbb{C}$, varying holomorphically with $u \in \mathcal{B}_{\text{reg}}$.
- BPS indices: functions $\Omega : \Gamma \to \mathbb{Z}$, obeying the *wall-crossing formula*.

1.3. **Coulomb branch.** The physics of a QFT in $\mathbb{R}^4$ (e.g. correlation functions, Hilbert space) depends on an additional choice: a point u of the *moduli space of vacua*. In the case of $\mathcal{N} = 2$ theories, this moduli space can be described rather concretely. We will focus on a subspace called the *Coulomb branch* $\mathcal{B}$,[2] and on an open dense subset $\mathcal{B}_{\text{reg}} \subset \mathcal{B}$ where the physics is simplest. $\mathcal{B}_{\text{reg}}$ is a complex manifold. Let $r = \dim \mathcal{B}_{\text{reg}}$.

Example 1.5. In the pure $\mathrm{SU}(2)$ theory, the Coulomb branch is $\mathcal{B} \simeq \mathbb{C}$ (so $r = 1$). (Explicitly it is parameterized by the expectation value $u = \langle \mathrm{Tr}\, \Phi^2 \rangle$, where Φ is the $\mathfrak{su}(2)$-valued complex scalar field in the theory.) The regular locus is $\mathcal{B}_{\text{reg}} = \mathcal{B} \setminus \{\pm 2\Lambda^2\}$.

[2] We could define it as the subspace of the full moduli space which is fixed by the R-symmetry $\mathrm{SU}(2)_R$.

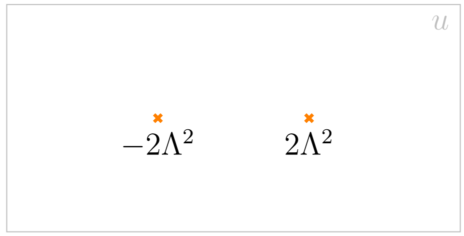

For a generic $u \in \mathcal{B}$, the long-distance (sometimes also called "IR" or "low-energy") physics in vacuum u is given by pure abelian gauge theory, where $T = U(1)^r$, with a symmetric coupling matrix $\tau_{IJ}(u)$ depending holomorphically on u, with $\operatorname{Im}\tau$ positive definite.[3] Practically speaking, this is a big win: the original theory may be very complicated, but its long-distance physics is simple.[4]

There is one more important point. $\tau_{IJ}(u)$ is allowed to be *multivalued* as a function of u, with monodromies

$$\tau \to \tau' = (A\tau + B)(C\tau + D)^{-1}$$

where $\begin{pmatrix} A & B \\ C & D \end{pmatrix} \in \mathrm{Sp}(2r, \mathbb{Z})$. The reason this is consistent is that the pure abelian gauge theories with coupling τ' and τ are actually equivalent (electric-magnetic duality). Said otherwise τ is not canonically defined, only defined after we make an extra choice.

Example 1.6. A local model, in $r = 1$ case: $\tau(u) = \tau_0 + \frac{1}{2\pi \mathrm{i}} \log u$, multivalued with $\tau' = \tau + 1$, corresponding to the matrix $\begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}$. If $\operatorname{Im}\tau_0 > 0$, then this model has $\operatorname{Im}\tau > 0$ for $|u|$ small enough. But it cannot extend to the whole complex plane.

Example 1.7. How could we get a $\tau(u)$ which does the job in the pure $\mathrm{SU}(2)$ theory? Seiberg-Witten proposed (by educated guesswork) the picture we drew above, and figured out what the monodromies would have to be: around each of $u = \pm 2\Lambda^2$ it has to look like the local model above. Then they pointed out that if we have a family of elliptic curves Σ_u over $\mathcal{B}$, which are *smooth* for $u \in \mathcal{B}_{\mathrm{reg}}$ and *nodal* for $u = \pm 2\Lambda^2$, the modulus $\tau(u)$ will have all the requisite properties.

So: let

$$\Sigma_u = \{(y, z) \in \mathbb{C}^2 \mid y^2 = \Lambda^2 z^{-1} + u + \Lambda^2 z\}\,.$$

Then for $u \in \mathcal{B}_{\mathrm{reg}}$, Σ_u is a twice-punctured smooth torus. Its closure $\overline{\Sigma}_u = \mathbb{C}/[\mathbb{Z} \oplus \tau(u)\mathbb{Z}]$; that determines $\tau(u)$. For $u = \pm 2\Lambda^2$, Σ_u is a twice-punctured nodal torus. (Exercise: check this!)

So the IR physics here is determined by some algebraic geometry, which at first looks completely alien to the QFT we started with. We'll discuss interpretations of it later.

---

[3] The long-distance theory may also have flavor symmetries.

[4] Loosely we could say the exact physics in vacuum u is given by pure abelian gauge theory with a coupling matrix $\tau_{IJ}(u)$, deformed by irrelevant operators. To all orders in an expansion in distances, these irrelevant operators are determined by the Taylor expansion of τ_{IJ} around u.

1.4. **Charge lattices.** For $u \in \mathcal{B}_{\text{reg}}$, there is a lattice Γ_u^{EM} of electromagnetic charges. It has a skew pairing ("DSZ")

$$\langle \cdot, \cdot \rangle : \Gamma_u^{\text{EM}} \times \Gamma_u^{\text{EM}} \to \mathbb{Z}$$

determined by the Poynting vector, or more concretely $\langle \gamma_e, \gamma_m \rangle = 1$. The lattices Γ_u^{EM} form a local system over $\mathcal{B}_{\text{reg}}$, with monodromies valued in $\mathrm{Sp}(2r, \mathbb{Z})$ — the same ones we discussed above.

In general there may also be flavor symmetries which give rise to a finer notion of charge. We need to pay particular attention to the abelian ones; these give an extension of the charge lattice,

$$0 \to \Gamma^{\text{flavor}} \to \Gamma_u \to \Gamma_u^{\text{EM}} \to 0\,.$$

The full monodromy of Γ is valued in $\mathrm{Sp}(2r, \mathbb{Z}) \ltimes (\Gamma^{\text{flavor}})^{2r}$. (The extra monodromies reflect the possibility of shifting the flavor charge by a multiple of the electromagnetic charge.)

Example 1.8. In the pure $\mathrm{SU}(2)$ theory, $\Gamma^{\text{flavor}} = 0$, and $\Gamma_u = \Gamma_u^{\text{EM}} = H_1(\overline{\Sigma}_u, \mathbb{Z})$. The pairing $\langle \cdot, \cdot \rangle$ is the intersection pairing.

The Hilbert space of the theory on $\mathbb{R}^4$ is decomposed as

$$\mathcal{H}_u = \bigoplus_{\gamma \in \Gamma_u} \mathcal{H}_{\gamma, u}\,.$$

1.5. **Central charges.** Inside the $\mathcal{N} = 2$ supersymmetry algebra $\mathcal{S}$ there is a central generator Z, which acts as a constant in each charge sector, and is additive in the charges: thus there is a homomorphism

$$Z : \Gamma_u \to \mathbb{C}$$

the *central charge*. Thus we have a global function $Z : \Gamma \to \mathbb{C}$. Z is holomorphic. It is critically important for the rest of the story.

Example 1.9. In the pure $\mathrm{SU}(2)$ theory, we have a holomorphic 1-form $\lambda = y/z\,\mathrm{d}z$ on Σ_u, and $Z : \Gamma_u \to \mathbb{C}$ is

$$Z(\gamma) = \oint_\gamma \lambda = \oint_\gamma \sqrt{\Lambda^2 z^{-1} + u + \Lambda^2 z}\,\frac{\mathrm{d}z}{z}\,.$$

(It's well defined even though λ has poles at the punctures of Σ_u; that's because those poles have zero residue.) Concretely, the $Z(\gamma)$ are given by hypergeometric functions of u, with branching around the points $u = \pm 2\Lambda^2$.

1.6. **The class $\mathcal{S}$ examples.**

Example 1.10. In a class $\mathcal{S}$ theory with $\mathfrak{g} = \mathfrak{sl}_2$, taking C to be a surface of genus g_C, the Coulomb branch is the space of holomorphic quadratic differentials, $\mathcal{B} = H^0(C, K_C^2)$. (Recall that a holomorphic quadratic differential is an object written locally as $\phi_2(z) = f(z)\,\mathrm{d}z^2$, with $f(z)$ holomorphic.) For each $u = \phi_2 \in \mathcal{B}$ there is a corresponding *spectral curve*

$$\Sigma_u = \{(z \in C, y \in T_z^* C) \mid y^2 + \phi_2(z) = 0\} \subset T^*C\,.$$

The projection map $\pi : \Sigma_u \to C$ is a branched double cover, branched at the zeroes of ϕ_2. The curve Σ_u is smooth and reduced just if all zeroes of ϕ_2 are simple.

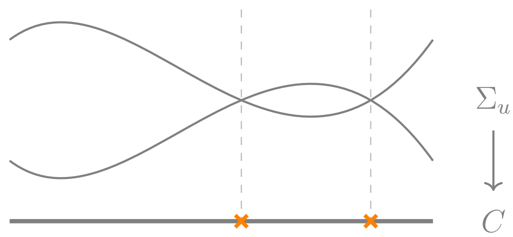

The regular locus is

$$\mathcal{B}_{\text{reg}} = \{\phi_2 \in H^0(C, K_C^2) \mid \Sigma_u \text{ is smooth}\}\,.$$

It has dimension $r = 3g_C - 3$.

Σ_u has an action of $\mathbb{Z}_2$ given by $\sigma(z, y) = (z, -y)$. Then σ also acts on $H_1(\Sigma_u, \mathbb{Z})$. Let $H_1(\Sigma_u, \mathbb{Z})^-$ be the σ-odd part. The charge lattice Γ_u is an extension of $H_1(\Sigma_u, \mathbb{Z})^-$ by $(\mathbb{Z}/2\mathbb{Z})^{g_C}$. (We won't have to worry much about this finite extension; almost everything we do will just use classes which are in $H_1(\Sigma_u, \mathbb{Z})^-$.)

2. Lecture 2

[first, quick review of the previous lecture]

2.1. **BPS states in a baby example.** I'll start with a baby example, to get oriented.

Fix a Riemannian manifold M. Then consider

$$\mathcal{H} = \Omega^*_{L^2}(M)\,.$$

It is acted on by the Hermitian operator

$$H = \frac{1}{2}\Delta\,.$$

But there are more operators around than just H: $\mathcal{H}$ is a unitary $\mathbb{Z}/2\mathbb{Z}$-graded representation of a Lie superalgebra $\mathcal{S}$, defined as follows.

$$\mathcal{S} = \mathcal{S}^0 \oplus \mathcal{S}^1, \quad \mathcal{S}^0 = \mathbb{C}\cdot H, \quad \mathcal{S}^1 = \mathbb{C}\cdot Q \oplus \mathbb{C}\cdot \overline{Q},$$

i.e. $\mathcal{S}$ has 2 odd generators $Q, \overline{Q}$ and one even generator H, with the brackets[5]

$$[Q, \overline{Q}] = 2H, \quad [Q, Q] = 0, \quad [\overline{Q}, \overline{Q}] = 0, \quad [Q, H] = 0, \quad [\overline{Q}, H] = 0.$$

A $\mathbb{Z}/2\mathbb{Z}$-*graded* representation of $\mathcal{S}$ is a representation of $\mathcal{S}$ on a $\mathbb{Z}/2\mathbb{Z}$-graded vector space $\mathcal{H} = \mathcal{H}^0 \oplus \mathcal{H}^1$, where $\mathcal{S}^i$ maps $\mathcal{H}^j \to \mathcal{H}^{j+i}$. A *unitary* representation of $\mathcal{S}$ is a representation in which $\mathcal{H}$ is a Hilbert space, H acts by a formally self-adjoint operator, and $Q, \overline{Q}$ act by operators which are formally adjoint to one another.

[5]Our convention is that $[,]$ means the *graded* bracket, i.e. for objects x, y which are in grade n_x, n_y respectively, $[x, y] = xy - (-1)^{n_x n_y} yx$. So this bracket is a commutator unless both x and y are odd, in which case it is an anticommutator.

To realize $\mathcal{H}$ as a representation of $\mathcal{S}$ we just take

$$Q = \mathrm{d}, \qquad \overline{Q} = \mathrm{d}^*, \qquad H = \frac{1}{2}\Delta.$$

Then $\mathcal{H}$ is a unitary representation of $\mathcal{S}$. It is also $\mathbb{Z}/2\mathbb{Z}$-graded:

$$\mathcal{H} = \mathcal{H}^0 \oplus \mathcal{H}^1, \quad \mathcal{H}^0 = \bigoplus_k \Omega^{2k}_{L^2}(M), \quad \mathcal{H}^1 = \bigoplus_k \Omega^{2k+1}_{L^2}(M).$$

Now let us explore a bit of the unitary representation theory of $\mathcal{S}$, following [4]. The unitarity implies that all eigenvalues of H are nonnegative, because

$$2\langle\psi, H\psi\rangle = \langle\psi, Q\overline{Q}\psi\rangle + \langle\psi, \overline{Q}Q\psi\rangle = \|Q\psi\|^2 + \|\overline{Q}\psi\|^2 \geq 0.$$

Moreover the norm $\|\cdot\|$ is nondegenerate, so we conclude that

$$H\psi = 0 \quad \Leftrightarrow \quad Q\psi = 0, \overline{Q}\psi = 0.$$

In particular, each state with $H\psi = 0$ generates a 1-dimensional (trivial) representation of $\mathcal{S}$. We call this representation V_0^0 or V_0^1 depending whether the single state is in $\mathcal{H}^0$ or $\mathcal{H}^1$. These representations are called "short." The only other possibility for a unitary irreducible $\mathbb{Z}/2\mathbb{Z}$-graded representation is a 2-dimensional representation, with one state in $\mathcal{H}^0$ and one in $\mathcal{H}^1$, both with $H\psi = E\psi$ for some $E > 0$; these representations are called "long."

If M is compact, then $\mathcal{H}$ is a countable orthogonal direct sum of unitary irreducible representations of $\mathcal{A}$. (We say $\mathcal{H}$ "contains only discrete spectrum.") See the figure below, where each dot represents one state; note that the states with $E > 0$ come paired up into long representations, while those with $E = 0$ are in short representations by themselves.

[fig]

The short and long representations have very different character, which we see clearly if we consider deformations of the representation $\mathcal{H}$, e.g. by varying the Riemannian metric on M. As we deform $\mathcal{H}$, the nonzero eigenvalues $E > 0$ of H can change continuously: the long representations are not rigid. The eigenvalues $E = 0$ have a harder time changing, because the representations V_0^i are rigid. This fact helps to "protect" the ground states. However, it doesn't protect them absolutely: the reducible representation $V_0^0 \oplus V_0^1$ is not rigid, since it can deform to a long representation with $E = \epsilon > 0$. See the figure below.

With this deformation process in mind we consider the following quantity.

$$\chi(\mathcal{H}) = \mathrm{Tr}_{\mathcal{H}}(-1)^F = (\text{\# copies of } V_0^0 \text{ in } \mathcal{H}) - (\text{\# copies of } V_0^1 \text{ in } \mathcal{H}).$$

The key property of the index is that it is invariant under deformations of $\mathcal{H}$, as long as $\mathcal{H}$ contains only discrete spectrum: if inside $\mathcal{H}$ a copy of $V_0^0 \oplus V_0^1$ deforms into a long representation, then χ changes by $1 - 1 = 0$.

The main lesson here is: while the full Hilbert space $\mathcal{H}$ depends strongly on every little detail of the system, by using a little bit of the representation theory of the supersymmetry algebra $\mathcal{A}$ — looking at representations which are particularly rigid — we are able to extract a more robust and invariant quantity. A second lesson is that the rigid representations are the ones which are smaller than usual, by virtue of being annihilated by part of $\mathcal{A}$ (in this case actually all of $\mathcal{A}$).

2.2. **BPS states in $\mathcal{N}=2$ theories.** Now let's return to our context of $\mathcal{N}=2$ theories.

Inside the Hilbert space $\mathcal{H}_{u,\gamma}$ there is the subspace $\mathcal{H}^1_{u,\gamma}$ of *1-particle states*.[6]

Like the full Hilbert space, it is a unitary representation of the $\mathcal{N}=2$ supersymmetry algebra $\mathcal{S}$. So we want to discuss the representations of this algebra.

$\mathcal{S}$ is a super Lie algebra extending $\mathfrak{iso}(3,1)$. It has 8 odd generators, described as follows. $\mathrm{Spin}(3,1)$ has 2 inequivalent spin representations $S^{\pm}$, both complex and 2-dimensional. Each of $S^{\pm}$ is equipped with an invariant antisymmetric pairing $\cdot$, and there is an intertwiner $\Gamma : S^+ \otimes S^- \to V$, with $V = \mathbb{R}^{3,1}$ the vector representation. The odd generators of $\mathcal{A}$ are Q^1, Q^2 valued in S^+ and $\overline{Q}^1$, $\overline{Q}^2$ valued in S^-.[7] The odd bracket relations are

$$[Q^I(s), Q^J(s')] = (s\cdot s')\epsilon^{IJ} Z\,, \quad [\overline{Q}^I(s), \overline{Q}^J(s')] = (s\cdot s')\epsilon^{IJ}\overline{Z}\,,$$
$$[Q^I(s), \overline{Q}^J(s')] = \delta^{IJ}\Gamma(s,s')P\,.$$

Suppressing the spinor and vector indices, we write these more schematically as

$$[Q^I, Q^J] = \epsilon^{IJ} Z\,, \qquad [\overline{Q}^I, \overline{Q}^J] = \epsilon^{IJ}\overline{Z}\,, \qquad [Q^I, \overline{Q}^J] = \delta^{IJ} P\,.$$

Now we can describe the 1-particle representations. In fact, it's sufficient to consider the subspace $\mathcal{H}^{1,\mathrm{rest}}_M$ where the translation generator P acts by the character $(M,0,0,0)$ for some $M>0$ (we only want massive particles, and no tachyons, so don't allow $M \le 0$). This subspace is a unitary representation of the "little algebra" $\mathcal{S}^{\mathrm{rest}}$ generated by Q^I, $\overline{Q}^I$, Z, P and generators of $\mathrm{Spin}(3) \subset \mathrm{Spin}(3,1)$. Now we want to classify these.

$\mathrm{Spin}(3) \simeq \mathrm{SU}(2)$ has only one spin representation S, which is complex and 2-dimensional, with an invariant antisymmetric pairing. So $S^+ \simeq S^- \simeq S$ when considered as representations of $\mathrm{Spin}(3)$. After fixing an isomorphism, we can write the odd brackets in $\mathcal{S}^{\mathrm{rest}}$ acting on $\mathcal{H}^{1,\mathrm{rest}}_M$ as

$$[Q^I, Q^J] = \epsilon^{IJ} Z, \qquad [\overline{Q}^I, \overline{Q}^J] = \epsilon^{IJ}\overline{Z}, \qquad [Q^I, \overline{Q}^J] = \delta^{IJ} M.$$

Now, we consider the generators

$$Q_\vartheta = \frac{1}{\sqrt{2}}(\mathrm{e}^{\mathrm{i}\vartheta/2}Q^1 + \mathrm{e}^{-\mathrm{i}\vartheta/2}\overline{Q}^2).$$

Now for any $\Psi \in \mathcal{H}^{1,\mathrm{rest}}_M$ we must have

$$\langle \Psi, [Q_\vartheta, \overline{Q_\vartheta}]\Psi\rangle = \|Q_\vartheta\Psi\|^2 + \|\overline{Q_\vartheta}\Psi\|^2 \ge 0$$

[6] Warning: this definition is not as transparent as it might sound. To define $\mathcal{H}^1_{u,\gamma}$ we need to be able to separate the 1-particle states from the continuum of multiparticle states. Even for generic u, there are at least two sources of such continua: massless photons, which are always present, and also possible decays if γ is not primitive (if $\gamma = k\mu$ and $k_1 + k_2 = k$ then 2-particle states of charges $k_1\mu$ and $k_2\mu$ could mix with 1-particle states of charge γ.) Nevertheless it seems that this does not cause a problem in practice. From now on I'll assume that $\mathcal{H}^1_{u,\gamma}$ is indeed a well defined representation of the supersymmetry algebra.

[7] What we mean by "Q is valued in S^+" is that for any $s \in S^+$, there is a corresponding operator $Q(s)$, depending linearly on s. $S^+ \oplus S^-$ admits a conjugate-linear involution exchanging the factors, so given $s \in S^{\pm}$ we can write $\overline{s} \in S^{\mp}$; then the adjointness condition means that in unitary representations $Q^I(s)$ is adjoint to $\overline{Q}^I(\overline{s})$, i.e. $\overline{Q(s)} = \overline{Q}(\overline{s})$.

and on the other hand

$$[Q_\vartheta, \overline{Q_\vartheta}] = \mathrm{e}^{\mathrm{i}\vartheta} Z + \mathrm{e}^{-\mathrm{i}\vartheta}\overline{Z} + 2M$$

so we conclude that

$$M \geq \mathrm{Re}(\mathrm{e}^{\mathrm{i}\vartheta})Z\,.$$

Since this holds for *every* ϑ it follows that

$$M \geq |Z|\,.$$

Moreover, if we have equality $M = |Z|$, then Q_ϑ and $\overline{Q_\vartheta}$ both annihilate the whole representation $\mathcal{H}^{1,\mathrm{rest}}_M$. The representations with $M = |Z|$ are called short or BPS while those with $M > |Z|$ are called long or non-BPS.

Classification of short representations of $\mathcal{S}^{\mathrm{rest}}$:

$$V_n = (\mathbf{1} \oplus \mathbf{1} \oplus \mathbf{2}) \otimes \mathbf{n}$$

for each $n = 1, 2, 3, \cdots$ There are also long representations:

$$W_n = (5 \cdot \mathbf{1} \oplus 4 \cdot \mathbf{2} \oplus \mathbf{3}) \otimes \mathbf{n}$$

This also gives a classification of the representations of $\mathcal{S}$ containing the massive 1-particle states.

2.3. **The BPS index.** We count how many short representations occur, by a *BPS index* (second helicity supertrace)

$$\Omega(\gamma; u) = -\frac{1}{2}\,\mathrm{Tr}_{\mathcal{H}^{1,\mathrm{BPS,rest}}_{u,\gamma}} (-1)^{2J_3}(2J_3)^2$$

This index has the virtue that it vanishes in all long representations (exercise!) Thus it is invariant under processes where the representation $\mathcal{H}^{1,\mathrm{BPS,rest}}_{u,\gamma}$ varies continuously (including processes where short representations pair up into long representations.) Moreover we compute directly that each copy of V_n in $\mathcal{H}^{1,\mathrm{BPS,rest}}$ contributes $(-1)^{n+1}n$ to $\Omega(\gamma; u)$ (exercise!)

What are these indices in practice?

Example 2.1. In the pure $\mathrm{SU}(2)$ theory, for sufficiently large $|u|$, one finds by studying the classical theory

$$\begin{aligned} \Omega(\pm\gamma_m + 2n\gamma_e; u) &= +1 \quad \text{for all } n \in \mathbb{Z} \text{ (massive hypermultiplets, } V_1)\,, \\ \Omega(\pm 2\gamma_e; u) &= -2 \quad \text{(massive vectormultiplets, } V_2)\,. \end{aligned}$$

Example 2.2. In a theory of class $\mathcal{S}$ of type A_1, for any $u = \phi_2 \in \mathcal{B} = H^0(C, K_C^2)$, and any phase ϑ, we consider ϑ-trajectories: this means paths on C where the 1-form $\mathrm{e}^{-\mathrm{i}\vartheta}\sqrt{-\phi_2}$ is real. The ϑ-trajectories form a singular foliation of C. We are interested in finite ϑ-trajectories, of two types: *saddle connections*, which are isolated, and *closed loops*, which occur in ring domains.

[insert a figure here]

Each finite ϑ-trajectory has an associated charge γ, as indicated in the figure, and one has $\vartheta = \arg Z_\gamma$ (exercise: prove it!) Then, according to [5, 6], the BPS index $\Omega(\gamma; u)$ is a weighted count of finite-length geodesics of phase $\vartheta = \arg Z_\gamma$:

$$\Omega(\gamma; u) = (\#\text{ saddle connections of charge }\gamma) - 2(\#\text{ ring domains of charge }\gamma).$$

One way to explore the BPS indices in a class $\mathcal{S}$ theory: fix ϑ and draw the ϑ-trajectories emanating from the branch points. As we vary ϑ, the picture exhibits discontinuous jumps when we cross the phase of a BPS particle.

[show animations at this point]

3. Lecture 3

[short review of previous lectures]

3.1. **Remarks and connections.** To connect with Nekrasov's lectures: sometimes it is convenient to choose a basis $(\gamma_e^I, \gamma_{m,I}, \gamma_f^A)$ of Γ_u, where $\gamma_f^A \in \Gamma^{\text{flavor}}$ and $\langle \gamma_e^J, \gamma_{m,I} \rangle = \delta_I^J$. Then the functions $a^I = Z(\gamma_e^I)$ give local coordinates on a neighborhood of u in $\mathcal{B}_{\text{reg}}$. Also define $a_{D,I} = Z(\gamma_{m,I})$. These aren't independent: they can be expressed as functions of the a^I. Then the couplings $\tau_{IJ}(u)$ are determined by

$$\tau_{IJ} = \frac{\partial a_{D,J}}{\partial a^I}\,,$$

and locally we can write $\tau_{IJ} = \partial_{a^I} \partial_{a^J} \mathcal{F}(a)$.

To connect with Gukov's lectures: Gukov discussed a BPS index which is (*roughly*) counting BPS (ground) states of a 2-d system in finite volume (space = S^1). In our story we are discussing a system in infinite volume (space = $\mathbb{R}^3$), and studying 1-particle states rather than ground states.

3.2. **Wall crossing.** The arguments I gave so far suggest that $\Omega(\gamma; u)$ would be locally independent of u — we have a family of Hilbert spaces $\mathcal{H}^{1,\text{BPS}}_{u,\gamma}$, all representations of $\mathcal{S}$, and the index is supposed to be invariant under deformations of representations.

But this would lead to a contradiction: we saw in the examples I showed that the $\Omega(\gamma; u)$ in fact *do* depend on u. Moreover, the $\Omega(\gamma; u)$ we described in the pure $\mathrm{SU}(2)$ theory are not invariant under monodromy in the full u-plane.

The correct statement is that the $\Omega(\gamma; u)$ depend on u in a *piecewise* constant way. The technical problem is that the continuum of multiparticle states interacts with the 1-particle states, so that at some values of u the space $\mathcal{H}^{1,\text{BPS}}_{u,\gamma}$ is actually ill defined: this violates our argument for deformation invariance. (The analogue in our baby example of the Witten index would be to study a "family" of compact Riemannian manifolds which at some moment loses compactness: then the index can jump.)

Imagine watching a BPS particle of charge γ as we vary u. We are worried that maybe this particle can decay into particles of charges μ_1, μ_2 with $\mu_1 + \mu_2 = \gamma$, but μ_1, μ_2 not multiples of γ. Can it happen? Conservation of energy would require that the rest masses obey

$$M_1 + M_2 \leq M\,.$$

But $M = |Z(\gamma)|$ since the particle is BPS, while the decay constituents must at least have $M_i \geq |Z(\mu_i)|$, and thus we can extend the above to

$$|Z(\mu_1)| + |Z(\mu_2)| \leq M_1 + M_2 \leq M = |Z(\gamma)|$$

On the other hand $\mu_1 + \mu_2 = \gamma$, which means that

$$Z(\mu_1) + Z(\mu_2) = Z(\gamma)\,,$$

so the triangle inequality gives

$$|Z(\mu_1)| + |Z(\mu_2)| \geq |Z(\gamma)|\,.$$

The only way to make these relations consistent is to have

$$Z(\mu_1) \parallel Z(\mu_2)$$

i.e. these two complex numbers have the same phase.

The upshot is that this decay can occur *only when* there exist BPS particles with charges μ_1, μ_2 and $Z(\mu_1) \parallel Z(\mu_2)$. For a generic u, the condition $Z(\mu_1) \parallel Z(\mu_2)$ will not be satisfied; rather, it will be satisfied for u on some codimension-1 wall in $\mathcal{B}_{\text{reg}}$. The loci where this happens are called *potential walls of marginal stability.*

So, we expect that $\Omega(\gamma; u) \in \mathbb{Z}$ is piecewise constant, with jumps only at potential walls of marginal stability. The places where a jump actually occurs are called *walls of marginal stability.*

Example 3.1. In the pure $\mathrm{SU}(2)$ theory, there is a single potential wall of marginal stability.

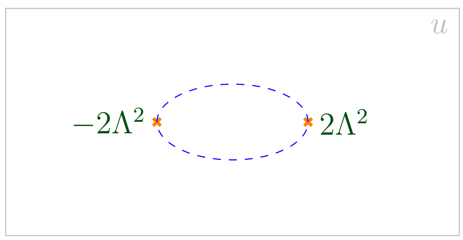

(Exercise: using our description of Z, explain why this is the right picture.)

It turns out this is an actual wall of marginal stability, not only a potential one. We have already discussed the spectrum outside the wall: there are infinitely many BPS particles. Inside the wall, there are BPS particles with charges $\{\pm\gamma_m, \pm(2\gamma_e - \gamma_m)\}$, each with $\Omega = 1$ (hypermultiplet). So when we cross the wall from outside to inside, almost all of these infinitely many particles decay, leaving just four.

3.3. **Wall-crossing formula.** If we know the spectrum of BPS particles on one side of a wall of marginal stability, how do we determine what it will be on the other side? It turns out that this doesn't depend on any UV details of the theory: it involves the long-range forces holding together very weakly bound states, i.e. it is a completely IR phenomenon.

We draw a diagram in $\mathcal{B} \times S^1$: a point (u, ϑ) is marked on the diagram if the theory at $u \in \mathcal{B}$ has a BPS particle with $\arg Z = \vartheta$. This gives a collection of codimension-1 walls, one for each BPS particle. Call it the *BPS scattering diagram* for this theory. I emphasize that the walls in this diagram are not the walls of marginal stability for bulk BPS particles which we discussed before. Rather each wall represents a bulk BPS particle.[8]

[8]I use the word "bulk" to emphasize that these are particles which propagate freely in the bulk spacetime $\mathbb{R}^4$, and to distinguish from some other kind of BPS states which will show up momentarily.

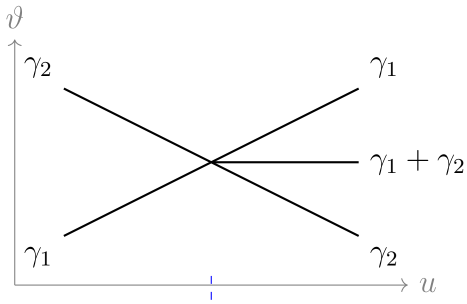

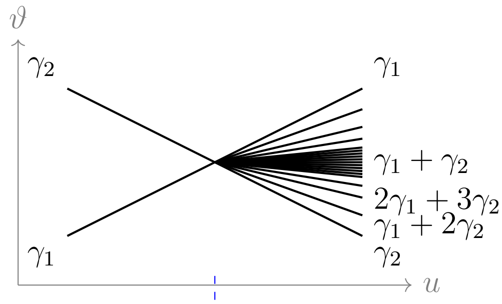

Now I can explain the *Kontsevich-Soibelman wall-crossing formula* of [7]. It was originally written down in the context of algebraic geometry (generalized Donaldson-Thomas invariants) but Denef-Moore soon realized it should apply also to BPS particles in $\mathcal{N} = 2$ theories. It was the endpoint of a long sequence of works which addressed many special cases.

To write down the wall-crossing formula, we first consider a field $\mathcal{A}_u$ generated by formal variables X_γ, $\gamma \in \Gamma_u$, with the product law

$$X_\gamma X_\mu = (-1)^{\langle \gamma, \mu \rangle} X_{\gamma+\mu} \,.$$

Define an automorphism of $\mathcal{A}_u$ by

$$\mathcal{K}_\gamma(X_\mu) = (1 - X_\gamma)^{\langle \gamma, \mu \rangle} X_\mu \,. \tag{3.1}$$

Now, consider a contractible oriented loop ℓ in $\mathcal{B} \times S^1$. Every time we cross a wall in the BPS scattering diagram, we include a factor $\mathcal{K}_\gamma^{\pm\Omega(\gamma)}$; the $\pm$ keeps track of whether we cross in the direction of increasing or decreasing ϑ. The wall-crossing formula says that

$$\prod \mathcal{K}_\gamma^{\pm\Omega(\gamma)} = \mathbf{1} \,. \tag{3.2}$$

Example 3.2. Suppose we have two BPS particles with charges γ_1 and γ_2 with inner product $\langle \gamma_1, \gamma_2 \rangle = 1$, and $\arg Z(\gamma_2; u) > \arg Z(\gamma_1; u)$.

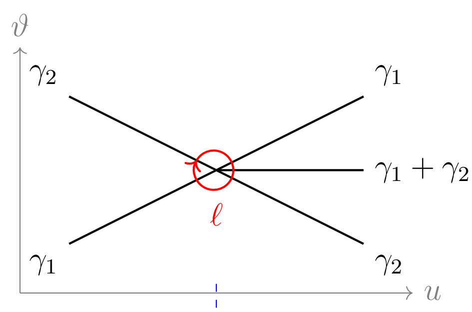

Now what happens if we vary u to a region where $\arg Z(\gamma_2; u) < \arg Z(\gamma_1; u)$? Then we are crossing a potential wall of marginal stability. To see what happens on the other side we use the algebraic identity

$$\mathcal{K}_{\gamma_2} \mathcal{K}_{\gamma_1} = \mathcal{K}_{\gamma_1} \mathcal{K}_{\gamma_1+\gamma_2} \mathcal{K}_{\gamma_2} \,.$$

(Exercise: prove this identity!) The right side is the unique product of the form $\mathcal{K}_{\gamma_1} \cdots \mathcal{K}_{\gamma_2}$ which equals the left side, where the $\cdots$ consists of charges of the form $m\gamma_1 + n\gamma_2$, $m, n > 0$.

Thus, applying (3.2) to the loop shown in the BPS scattering diagram above, we conclude that (as shown in the diagram) on the other side of the wall there is an additional particle with charge $\gamma_1 + \gamma_2$, and no other particles with charges $m\gamma_1 + n\gamma_2$, $m, n > 0$.

Example 3.3. Now suppose everything is as before except that $\langle \gamma_1, \gamma_2 \rangle = 2$ instead of 1.

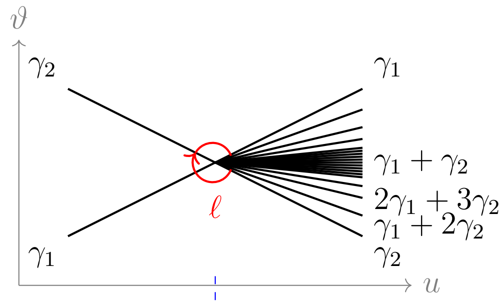

In this case we have a different identity:

$$\mathcal{K}_{\gamma_2}\mathcal{K}_{\gamma_1} = \left(\mathcal{K}_{\gamma_1}\mathcal{K}_{2\gamma_1+\gamma_2}\mathcal{K}_{3\gamma_1+2\gamma_2}\cdots\right)\mathcal{K}_{\gamma_1+\gamma_2}^{-2}\left(\cdots\mathcal{K}_{2\gamma_1+3\gamma_2}\mathcal{K}_{\gamma_1+2\gamma_2}\mathcal{K}_{\gamma_2}\right) .$$

This identity is harder to formulate than the previous one, since it involves an infinite product. Still, one can introduce a suitable topology in which the product converges, and then the identity is indeed true. One useful way to think of it is as follows. We study $\mathcal{A}_u$ through its *spectrum*: a point of $\operatorname{Spec}\mathcal{A}_u$ is given by actual numbers $x_\gamma \in \mathbb{C}^\times$ obeying (3.1). The operators $\mathcal{K}_\gamma$ have an induced action on $\operatorname{Spec}\mathcal{A}_u$. Then acting on a point with $|x_{\gamma_1+\gamma_2}| < 1$, both sides make sense and are equal. (Warning: the order of composition is reversed when we consider the action on $\operatorname{Spec}\mathcal{A}_u$, so to check the identity this way, you act first with the leftmost operator rather than the rightmost.) A proof is in [8].

The conclusion is that on the other side of the wall we have infinitely many charges with $\Omega = 1$, and a single charge with $\Omega = -2$.

Note that this second example "explains" the wall-crossing behavior we described in the pure $\mathrm{SU}(2)$ theory. At least, it would explain it, if we knew why the wall-crossing formula is true!

So: why is the wall-crossing formula true? By now there are various arguments. I'll describe a physical argument which appeared in [9], with the advantage that it makes clear where the algebra $\mathcal{A}_u$ comes from. (See [10] for another physical argument, in some ways more direct, just using the form of the long-distance interaction between BPS particles.)

4. Lecture 4

[review]

4.1. **Interpretations.** BPS particles in 4d $\mathcal{N} = 2$ theories have many different interpretations. To be concrete, for the BPS particles in the pure $SU(2)$ theory which we discussed, we have at least six interpretations. I mention them here, just in case one of them reminds you of something!

- 't Hooft-Polyakov solutions of Bogomolny equations in $\mathbb{R}^3$. ($SU(2)$ gauge theory).
- Saddle connections and closed loops. (strings of 5d Yang-Mills) [fig]

- Special Lagrangian surfaces in $T^*\mathbb{C}$ ending on $\Sigma_u = \{y^2 + \phi_2(z) = 0\}$. (M2-branes ending on M5-brane in M theory) [fig]
- Special Lagrangian 3-spheres in a Calabi-Yau manifold $X = \{x^2 + y^2 + w^2 + \phi_2(z) = 0\}$ (terms and conditions apply: more precise statement would be "stable objects in Fukaya category of a symplectic manifold with appropriate Bridgeland stability condition") (D3-branes of Type IIB string theory on X) [fig]
- Cohomology of moduli spaces of stable representations of the 2-Kronecker quiver. (quiver quantum mechanics) [fig]
- Split attractor flows: certain networks drawn on $\mathcal{B}$. (effective dynamics on Coulomb branch) [fig]

4.2. **The algebra.** We consider *supersymmetric line defects* in our $\mathcal{N} = 2$ field theory. Fix coordinates in $\mathbb{R}^4$. We'll consider line defects which are extended in the x^3 direction. For each $\zeta \in \mathbb{C}^\times$ with $|\zeta| = 1$ there is a $\frac{1}{2}$-BPS subalgebra[9]

$$\mathcal{S}_\zeta \subset \mathcal{S}$$

which contains the translations in the x^3 direction, thus is a candidate to be preserved by a line defect. $\mathcal{S}_\zeta$ is generated by the operators Q_ϑ, $\overline{Q}_\vartheta$ from the previous lecture, where $\zeta = \mathrm{e}^{\mathrm{i}\vartheta}$.

After continuation to Minkowski signature, it is sometimes useful to think of an $\mathcal{S}_\zeta$-invariant line defect as representing the world-line of a very heavy particle with $\arg Z = \vartheta$.

Example 4.1. In the pure $\mathrm{U}(1)$ theory, let Φ denote the complex scalar field and A the connection 1-form. Then there is a supersymmetric Wilson line operator defined by

$$L = \exp \int \zeta^{-1}\Phi + \mathrm{i}A + \zeta\overline{\Phi}\,.$$

Similarly one can define supersymmetric Wilson lines in pure nonabelian gauge theory.

For each fixed ζ, the supersymmetric line defects form some kind of tensor category, not yet fully understood. We will work with a certain reduction of this category. First, we pass to some kind of K-theory of the category — considering line defects up to deformation equivalence. This gives an algebra. Then we further simplify this algebra by taking coinvariants for the $\mathrm{SO}(3)$ action by rotations.

Applying this procedure to the original (UV) $\mathcal{N} = 2$ theory, we get a single algebra $\mathcal{A}_\zeta$ of line defects. We can also apply it to low energy (IR) theories on the Coulomb branch, and get a family of algebras $\mathcal{A}^{\mathrm{IR}}_{u,\zeta}$. The UV algebra can be complicated, but the IR algebras are relatively simple. We have one defect L^{IR}_γ for each $\gamma \in \Gamma_u$. They have an OPE given by

$$L^{\mathrm{IR}}_\gamma L^{\mathrm{IR}}_\mu = (-1)^{\langle\gamma,\mu\rangle} L^{\mathrm{IR}}_{\gamma+\mu}$$

(Up to the tricky sign, this expresses the linearity of the theory: two charged particles behave like one particle with the combined charge. The sign has to do with the shift in fermion number

[9] By $\frac{1}{2}$-BPS we mean it contains *half* of the odd generators in $\mathcal{S}$.

associated to the angular momentum.) So this is the physical realization of the $\mathcal{A}_u$ that appeared in the wall-crossing formula.

Example 4.2. In a class $\mathcal{S}$ theory of type A_1, $\mathcal{A}_\zeta$ is an algebra generated by simple closed curves on C (commutative skein algebra, slightly interesting product law), $\mathcal{A}_{u,\zeta}$ is an algebra generated by simple closed curves on Σ_u (torus algebra, dead simple product law).

4.3. **UV-IR map for line defects.** For a generic (u, ζ), there is a "UV-IR map"

$$\mathbf{RG}_{u,\zeta} : \mathcal{A}_\zeta \to \mathcal{A}^{\mathrm{IR}}_{u,\zeta} \,.$$

It maps a UV line defect L to a linear combination of IR line defects:

$$\mathbf{RG}_{u,\zeta}(L) = \sum_{\gamma \in \Gamma_u} \overline{\underline{\Omega}}(L, \gamma; u, \zeta) L^{\mathrm{IR}}_\gamma$$

where the coefficients $\overline{\underline{\Omega}}(L, \gamma; u, \zeta) \in \mathbb{Z}$.

The coefficients in the UV-IR map also have a "BPS" interpretation, as *framed BPS state counts*. Indeed, the system including the line defect has a Hilbert space $\mathcal{H}_L$, still infinite-dimensional, but now not a representation of $\mathcal{S}$ anymore: rather it is only a representation of the subalgebra $\mathcal{S}_\zeta$ preserved by the line defect. By a *framed BPS state* of charge γ we mean a state of the system including the line defect, which saturates the inequality $E \geq -\operatorname{Re}(Z_\gamma/\zeta)$; such states are also annihilated by all of the odd generators of $\mathcal{S}_\zeta$. We emphasize that unlike the BPS states we originally studied, which describe particles that are free to move in the bulk, a framed BPS state of a line defect L is really to be thought of as a ground state of this fixed object.

Now let $\mathcal{H}^{\mathrm{BPS}}_{L,\gamma,u}$ denote the space of framed BPS states with the line defect L inserted. Then the framed BPS index is

$$\overline{\underline{\Omega}}(L, \gamma; u) = \operatorname{Tr}_{\mathcal{H}^{\mathrm{BPS}}_{L,\gamma,u}} (-1)^{2I_3}$$

where I_3 denotes a generator of the $\mathrm{SU}(2)_R$ symmetry of the $\mathcal{N} = 2$ theory (this symmetry is always preserved by the line defects we consider.)

When $(u, \vartheta = \arg \zeta)$ crosses a wall of the BPS scattering diagram, the UV-IR map jumps. Namely, suppose the wall corresponds to charge γ and BPS index $\Omega(\gamma)$. Then the jump of $\mathbf{RG}_{u,\zeta}$ is by postcomposition with the automorphism $\mathcal{K}_\gamma^{\Omega(\gamma)}$, i.e.

$$\mathbf{RG}_{u_+,\zeta_+} = \mathcal{K}_\gamma^{\Omega(\gamma)} \circ \mathbf{RG}_{u_-,\zeta_-} \,.$$

Example 4.3. In the pure $SU(2)$ theory, take L to be the $\frac{1}{2}$-BPS Wilson line operator in the fundamental representation. Then for any generic (u, ζ), $\mathbf{RG}_{u,\zeta}(L)$ turns out to be a sum of three IR line defects. At large $|u|$ we can describe them: two Wilson lines (electric) and one Wilson-'t Hooft line (dyon). (The dyon is the surprising part: classically we would just expect the two electric states.)

Example 4.4. In class $\mathcal{S}$ theories, $\mathbf{RG}_{u,\zeta}$ is path lifting...

Here is the physical picture, described in [9]. When (u, ζ) is near the wall, there can exist framed BPS states which are very weakly bound: they are described by a "core" which is a framed BPS state of charge μ, and a "halo" made of ordinary BPS particles, of charge γ. Each such particle turns out to have $\Omega(\gamma)\langle\gamma, \mu\rangle$ possible quantum states, each of which can be occupied or not. These framed BPS states exist on one side of the wall and not on the other.

Thus a UV line defect L which has $\mathbf{RG}(L) = L_\mu^{\mathrm{IR}}$ on one side of the wall will be "dressed" by bulk BPS particles and map to $\mathbf{RG}(L) = L_\mu^{\mathrm{IR}}(1 - L_\gamma^{\mathrm{IR}})^{\langle\gamma,\mu\rangle\Omega(\gamma)}$ on the other side. This accounts for the transformation $\mathcal{K}_\gamma^{\Omega(\gamma)}$ of the map $\mathbf{RG}$ at the wall.

4.4. **Return to the wall-crossing formula.** Now how does this relate to the Kontsevich-Soibelman wall-crossing formula? The point is that the $\mathbf{RG}$ map is canonically defined, not subject to any monodromy; it has jumps at the BPS scattering diagram, but it is a perfectly single-valued object. If we travel around a loop ℓ in $\mathcal{B} \times S^1$, beginning and ending at (u, ζ), then we find

$$\mathbf{RG}_{u,\zeta} = \left(\prod_\gamma \mathcal{K}_\gamma^{\Omega(\gamma)}\right) \circ \mathbf{RG}_{u,\zeta}\ .$$

The image of $\mathbf{RG}_{u,\zeta}$ in $\mathcal{A}_{u,\zeta}$ is big enough (at least in many examples, and hopefully always) that this implies

$$\prod_\gamma \mathcal{K}_\gamma^{\Omega(\gamma)} = \mathbf{1}$$

as desired.

4.5. **Reduction to three dimensions.** These algebras might sound a little abstract, so let's try to make them more concrete. We consider the compactification of our 4d $\mathcal{N} = 2$ theory to three dimensions on a circle, of length R. What do we get?

To describe the IR physics, one idea would be to start with the IR description of the 4d theory over $\mathcal{B}_{\mathrm{reg}}$ and compactify *that*. Doing this in the most naive possible way, we obtain a supersymmetric sigma-model: that's a theory whose bosonic part is a theory of maps from 3d spacetime to a Riemannian manifold $\mathcal{M}_{\mathrm{reg}}$. $\mathcal{M}_{\mathrm{reg}}$ is a bundle over the base $\mathcal{B}_{\mathrm{reg}}$, with fiber a compact torus $T_u \simeq (S^1)^{2r}$, a torsor over $\mathrm{Hom}(\Gamma_u^{\mathrm{EM}}, \mathrm{U}(1))$. The fiber coordinates of $\mathcal{M}_{\mathrm{reg}}$ are holonomies of the $U(1)^r$ gauge fields θ_e^I and their magnetic duals $\theta_{m,I}$. The metric g^{sf} is given by an explicit formula:

$$g^{\mathrm{sf}} = R(\mathrm{Im}\,\tau)|\mathrm{d}a|^2 + R^{-1}(\mathrm{Im}\,\tau)^{-1}|\mathrm{d}z|^2\,,$$

where we introduced

$$\mathrm{d}z_I = \mathrm{d}\theta_{m,I} - \tau_{IJ}\mathrm{d}\theta_e^J\,.$$

In fact, g^{sf} is not only Riemannian, its metric is hyperkähler. That is required by the amount of supersymmetry we have here ($\mathcal{N} = 4$ in three dimensions).

Does the naive dimensional reduction tell the full story? Optimistic view in highly supersymmetric theories: things will be exact unless there is a reason not to be. There is one natural candidate source of quantum correction: 3d "instanton" effects associated to BPS particles going around the circle. We can consider a baby example: $\mathrm{U}(1)$ theory coupled to 1 charged hypermultiplet. In this theory we have a single BPS particle. Then we can compute directly: the exact IR theory is still a supersymmetric sigma-model [11, 12], but with a corrected metric g, obeying

$$g - g^{\mathrm{sf}} = O(\mathrm{e}^{-RM})$$

where $M = |Z|$ is the mass of the BPS particle.

In all but the most trivial examples, though, we have both electrically and magnetically charged particles. Then we can't calculate the corrected metric directly: we need a more powerful method.

4.6. **Twistor space.** Our problem is to describe the low-energy physics of our reduced theory. We expect the answer will be governed by a hyperkähler metric g on the target space $\mathcal{M}$. To explain how we compute g, first we should recall what "hyperkähler" actually means. The usual definition is:

- $\mathcal{M}$ is *hypercomplex*, i.e. it admits three complex structures I_1, I_2, I_3 obeying $I_1 I_2 = I_3$, $I_2 I_1 = -I_3$ and cyclic permutations.
- g is Kähler with respect to each of these complex structures. (Thus it admits three symplectic forms ω_1, ω_2, ω_3.)

Once we have this structure, we get various things for free:

- Even though the definition only involves three Kähler structures (I_i, ω_i), these generate a whole S^2 worth of Kähler structures: any $\vec{s} \in S^2 \subset \mathbb{R}^3$ gives $I_{\vec{s}} = \sum_{i=1}^3 s_i I_i$ and $\omega_{\vec{s}} = \sum_{i=1}^3 s_i \omega_i$. In fact, it's a good idea to view this S^2 as complex itself, i.e. identify it as $\mathbb{CP}^1$, with the usual inhomogeneous coordinate ζ. Then one can assemble all of the complex structures into a single complex manifold $\mathcal{Z}(\mathcal{M})$, the *twistor space* of $\mathcal{M}$. $\mathcal{Z}(\mathcal{M})$ is fibered over $\mathbb{CP}^1$, with the fiber over ζ isomorphic to $(\mathcal{M}, I_\zeta)$.
- For each complex structure I_ζ we have a holomorphic symplectic form Ω_ζ, given by the formula
$$\Omega_\zeta = \zeta^{-1}\frac{\omega_2 + \mathrm{i}\omega_3}{2} - \mathrm{i}\omega_1 + \zeta\frac{\omega_2 - \mathrm{i}\omega_3}{2}$$
for $\zeta \in \mathbb{C}^\times$, and $\Omega_{\zeta=0} = \omega_2 + \mathrm{i}\omega_3$, $\Omega_{\zeta=\infty} = \omega_2 - \mathrm{i}\omega_3$.

To know the hyperkähler metric, it's enough to know the holomorphic symplectic forms Ω_ζ. Strategy:

- Wrapping a line defect $L \in \mathcal{A}_\zeta$ around the circle, we get a local operator, whose vacuum expectation value is a holomorphic function F_L on $(\mathcal{M}, I_\zeta)$.
- These functions decompose into $F_L = \sum_\gamma \overline{\underline{\Omega}}(L, \gamma; u, \zeta)\mathcal{X}_\gamma$ where $\mathcal{X}_\gamma$ are holomorphic Darboux coordinates on $\mathcal{M}$. Thus $(\mathcal{M}, I_\zeta, \Omega_\zeta)$ admits distinguished local Darboux coordinate systems.

[explain why they are Darboux coordinates, by "secondary operations"?] [explain what $\mathcal{M}$ is and what these coordinates are, in the case of class $\mathcal{S}$ theories; cluster-type structure] The functions $\mathcal{X}_\gamma$:

- are piecewise-holomorphic, with jumps determined by the $\mathcal{K}_\gamma^{\Omega(\gamma)}$ at the walls of the BPS scattering diagram

- have asymptotics as $\zeta \to 0$ given by $\mathcal{X}_\gamma \sim \exp(\zeta^{-1} Z_\gamma)$ (when $\theta_\gamma = 0$)
- have a reality condition as $\zeta \to -1/\overline{\zeta}$.

We can write down an integral equation whose solutions would have this property: (again when $\theta_\gamma = 0$)

$$\mathcal{X}_\gamma(u,\zeta) = \mathcal{X}_\gamma^{\mathrm{sf}}(u,\zeta) \exp\left[-\frac{1}{4\pi \mathrm{i}} \sum_\mu \Omega(\mu;u) \langle \gamma, \mu \rangle \int_{\ell_\mu} \frac{\mathrm{d}\zeta'}{\zeta'} \frac{\zeta' + \zeta}{\zeta' - \zeta} \log(1 - \mathcal{X}_\mu(\zeta')) \right],$$

where the contours of integration are the lines of the BPS scattering diagram,

$$\ell_\mu = Z_\mu \mathbb{R}_-.$$

Solving this integral equation in practice we can construct a metric (in very simple examples!) Moreover, experimentally, we can compute Hitchin's metric on the $\theta = 0$ locus inside $\mathcal{M}$ in some class $\mathcal{S}$ examples. The two agree!

[also theorems about asymptotics...]

References

1. Nathan Seiberg and Edward Witten. Electric-magnetic duality, monopole condensation, and confinement in $\mathcal{N} = 2$ supersymmetric Yang-Mills theory. *Nucl. Phys.* **B426** (1994), 19–52. eprint: hep-th/9407087. (Cit. on p. 1).
2. Nathan Seiberg and Edward Witten. Monopoles, duality and chiral symmetry breaking in $\mathcal{N} = 2$ supersymmetric QCD. *Nucl. Phys.* **B431** (1994), 484–550. eprint: hep-th/9408099. (Cit. on p. 1).
3. Yuji Tachikawa. *N=2 supersymmetric dynamics for pedestrians*. Dec. 2013. doi: 10.1007/978-3-319-08822-8. arXiv: 1312.2684 [hep-th] (cit. on p. 1).
4. Edward Witten. Constraints on supersymmetry breaking. *Nucl. Phys.* **B202** (1982), 253. (Cit. on p. 6).
5. Albrecht Klemm, Wolfgang Lerche, Peter Mayr, Cumrun Vafa, and Nicholas P. Warner. Self-Dual Strings and N=2 Supersymmetric Field Theory. *Nucl. Phys.* **B477** (1996), 746–766. eprint: hep-th/9604034. (Cit. on p. 8).
6. Davide Gaiotto, Gregory W. Moore, and Andrew Neitzke. Wall-crossing, Hitchin systems, and the WKB approximation. *Adv. Math.* **234** (2013), 239–403. arXiv: 0907.3987 [hep-th]. (Cit. on p. 8).
7. Maxim Kontsevich and Yan Soibelman. Stability structures, motivic Donaldson-Thomas invariants and cluster transformations (2008). eprint: 0811.2435. (Cit. on p. 11).
8. Davide Gaiotto, Gregory W. Moore, and Andrew Neitzke. Four-dimensional wall-crossing via three-dimensional field theory. *Commun. Math. Phys.* **299** (2010), 163–224. arXiv: 0807.4723 [hep-th]. (Cit. on p. 12).
9. Davide Gaiotto, Gregory W. Moore, and Andrew Neitzke. Framed BPS States. *Adv. Theor. Math. Phys.* **17**(2) (2013), 241–397. arXiv: 1006.0146 [hep-th]. (Cit. on pp. 12, 15).
10. Boris Pioline. Four ways across the wall. Ed. by Alexander Stolin, Viktor Abramov, Jurgen Fuchs, Eugen Paal, Yury Shestopalov, and Sergei Silvestrov. *J. Phys. Conf. Ser.* **346** (2012), 012017. arXiv: 1103.0261 [hep-th]. (Cit. on p. 12).
11. Hirosi Ooguri and Cumrun Vafa. Summing up D-Instantons. *Phys.Rev.Lett.* **77:3296-3298,1996** (Phys.Rev.Lett.77:3296-3298,1996). eprint: hep-th/9608079. (Cit. on p. 16).

12. Nathan Seiberg and Stephen H. Shenker. Hypermultiplet moduli space and string compactification to three-dimensions. *Phys. Lett. B* **388** (1996), 521–523. arXiv: `hep-th/9608086`. (Cit. on p. 16).