The geometry of Coulomb branches

Ben Webster

University of Waterloo
Perimeter Institute

June 25, 2025

Usual (2d) mirror symmetry comes from considering 2d σ-models.

- If the target space M is Kähler, we obtain $\mathcal{N} = (2,2)$ supersymmetry.
- This gives us two topological twists Q_A, Q_B, and two theories are mirror dual if they are equivalent in a way that swaps $Q_A \leftrightarrow Q_B$.

This gives us 2 extended TQFTs Z_A and Z_B, such that $Z_A \cong Z_B^\vee$. In particular:

X	$Z_A(X)$	$Z_B(X)$	interpretation
S^1	$QH^*(M)$	$HH^*(M)$	local operators
pt	$\mathrm{Fuk}(M)$	$\mathrm{QCoh}(M)$	boundary conditions

For 3d mirror symmetry, we just consider 3d σ-models instead. In order to have enough supersymmetry to get a topological twist, we need $\mathcal{N} = 4$ (i.e. 8 supercharges, which form 4 copies of the spinor representation of $SO(3, \mathbb{C})$) so the target must be hyperkähler.

We again have two natural topological twists Q_A, Q_B. We call theories **3-d mirror dual** if they are equivalent in a way that swaps $Q_A \leftrightarrow Q_B$.

Since theories are 3d, their local operators are commutative rings equipped with Poisson brackets. We think of these as functions on algebraic varieties called the **Higgs** and **Coulomb** branches $\mathfrak{M}_H$ and $\mathfrak{M}_C$.

X	$Z_A(X)$	$Z_B(X)$	interpretation
S^2	$\mathbb{C}[\mathfrak{M}_C]$	$\mathbb{C}[\mathfrak{M}_H]$	local operators
S^1			line operators
pt			boundary conditions

I don't know about you. I don't have a bunch of hyperkähler manifolds behind my couch.

There is one that everybody knows: $\mathbb{H}^n$. We can spice up this boring example by gauging by the action of a compact group G which preserves the hyperkähler structure.

Equivalently, think of $\mathbb{H}^n \cong T^*\mathbb{C}^n$ as a complex symplectic vector space, and look for G preserving the symplectic form. "G-gauge theory with matter content $\mathbb{H}^n$."

We'll focus on the case where $\mathbf{N} = \mathbb{C}^n$ has a G-representation, and we consider the induced action on $T^*\mathbf{N} = \mathbb{H} \otimes_{\mathbb{C}} \mathbf{N}$.

Call the resulting theory $\mathcal{T}(G, \mathbf{N})$.

Important examples:

- $\mathbf{N} = 0$: pure gauge theory
- $\mathbf{N} = \mathfrak{g}$
- abelian G (so we can assume $G \subset U(1)^n \circlearrowright \mathbb{C}^n$).
- quiver gauge theory:

$$G_c = \prod_i U(v_i) \qquad \mathbf{N} = \bigoplus_{i \to j} M_{v_j \times v_i}(\mathbb{C})$$

What do points in $\mathfrak{M}_H$ and $\mathfrak{M}_C$ mean?

The moduli of vacua $\mathfrak{M}_H(G, \mathbf{N}) = \operatorname{Spec} Z_B(S^2)$ for the B-twist (the **Higgs branch**) was studied by Hitchin, Karlhede, Lindström and Roček and interpreted as a hyperkähler quotient $\mathbb{H}^n ////G$.

We interpret $\|(\mu_I(x), \mu_J(x), \mu_K)\|^2$ as an energy functional (as in AO's talk), so points of $\mathfrak{M}_H$ are G-orbits of points in $T^*\mathbf{N}$ where μ_I, μ_J, μ_K vanish.

Of course, not all orbits are the same. Some are free (complete symmetry breaking), some are not. The stabilizer G_x of a point x gives an invariant of the orbit (up to conjugacy).

"Higgs mechanism"

Physically speaking, we see this stabilizer from the fact that its Lie algebra gives the massless gauge fields.

1. For $G \subset U(1)^n \circlearrowright \mathbb{C}^n$, just depends on which coordinates in x are non-zero.
2. For quiver gauge theory, think of matrices as representation of doubled quiver; stabilizer captures dimension vectors of simple subrepresentations (in particular, orbit is free iff representation is simple).

The Higgs branch as a whole is singular, but the set of points with fixed stabilizer are smooth (and in fact, hyperhähler). They form a "symplectic leaf."

These leaves aren't a special feature of Higgs branches; any well-behaved singular hyperkähler (symplectic singularity) will have finitely many leaves.

In physics, we can think of these leaves as phases: each point is a choice of vacuum, and moving around the leaf, the theory changes in a continuous way, whereas if we move from one leaf to another, the changes will be discontinuous.

The moduli of vacua for the A-twist (the **Coulomb branch**) proved much more mysterious. Denote this $\mathfrak{M}_C(G, \mathbf{N}) = \operatorname{Spec} Z_A(S^2)$.

The space $\mathfrak{M}_C(G, \mathbf{N})$ is birational to $T^* T^\vee / W$, but with a modification of the multiplication of functions on this space.

Work of Cremonesi–Hanany–Zaffaroni, Bullimore–Dimofte–Gaiotto, and Braverman–Finkelberg–Nakajima gradually brought this picture into clearer view:

- $\mathbb{C}[\mathfrak{M}_C(G, \mathbf{N})]$ is a subalgebra of the fraction field $\mathbb{C}(T^* T^\vee)^W$,
- obtained by replacing m_λ, a character on $T^\vee$, by multiplying it by certain rational functions (roughly with denominator depending on G and numerator depending on $\mathbf{N}$).

Note, this allows us to prove that at least the $\mathfrak{M}_H(\mathcal{T}) \cong \mathfrak{M}_C(\mathcal{T}^!)$ part of duality rigorously.

Theorem

- *If G is a torus acting diagonally and faithfully on $\mathbb{C}^n$, then $\mathfrak{M}_C(G, \mathbb{C}^n) \cong \mathfrak{M}_H(G^!, \mathbb{C}^n)$ where $G^! = \ker({}^L D \to {}^L G)$. This was first proposed by Kapustin and Strassler.*
- *For an affine type A quiver gauge theory (satisfying a dominance condition), the dual is another theory of the same type. This result was derived from brane diagrams by Hanany and Witten.*

For example, in the case of $G = U(1)$, we obtain well-known varieties:

- $\mathfrak{M}_C(U(1), 0) \cong T^*\mathbb{C}^\times = \mathrm{Spec}(\mathbb{C}[t, m, m^{-1}])$
- If $\mathbf{N}$ is a representation of $U(1)$ and $q_\pm$ is the sum of the positive and negative weights, $\mathbb{C}[\mathfrak{M}_C(U(1), \mathbf{N})]$ is the subalgebra generated by $t, m_+ = t^{q_+}m, m_- = t^{-q_-}m^{-1}$ with the relations

$$m_+ m_- = t^{q_+ - q_-}$$

so $\mathfrak{M}_C(U(1), \mathbf{N}) \cong \mathbb{C}^2/\mathbb{Z}_Q$ for $Q = q_+ - q_-$.

If $G = PSU(2)$ or $SU(2)$, then $\mathfrak{M}_C(G, \mathbf{N}) \approx \mathfrak{M}_C(U(1), \mathbf{N})/W$ for $U(1) \subset G$. Note that in this case we have $q = q_+ = -q_-$. The group $W = \{1, s\}$ acts by $s \cdot m_\pm = (-1)^q m_\mp$ and $s \cdot t = -t$.

- If $G = PSU(2)$, the ring $\mathbb{C}[\mathfrak{M}_C(PSU(2), \mathbf{N})]$ is generated by

$$t^2, m_\bullet = \frac{m_+ - s \cdot m_+}{t}, m_\circ = m_+ + s \cdot m_+$$

  modulo the relation $m_\circ^2 = t^2 m_\bullet^2 + (-1)^q \cdot 4t^{2q}$.

Theorem (Bielawski-Foscolo)

$\mathfrak{M}_C(PSU(2), \mathbf{N})$ is the space of W-invariant ideals I in $A = \mathbb{C}[\mathfrak{M}_C(G, \mathbf{N})]$ such that A/I is 2-dimensional and generated by t.

I is the ideal that vanishes on a free W-orbit, or on a single W-fixed point with a vector. Generation by t means that t has non-zero value or non-zero directional derivative.

If $G = PSU(2)$ or $SU(2)$, then $\mathfrak{M}_C(G, \mathbf{N}) \approx \mathfrak{M}_C(U(1), \mathbf{N})/W$ for $U(1) \subset G$. Note that in this case we have $q = q_+ = -q_-$. The group $W = \{1, s\}$ acts by $s \cdot m_\pm = (-1)^q m_\mp$ and $s \cdot t = -t$.

- If $G = SU(2)$, then the ring $\mathbb{C}[\mathfrak{M}_C(SU(2), \mathbf{N})]$ is generated by

$$t^2, m_\bullet = \frac{m_+ + s \cdot m_+}{t^2}, m_\circ = \frac{m_+ - s \cdot m_+}{t}$$

  modulo the relation $m_\circ^2 = t^2 m_\bullet^2 - (-1)^{q-1} \cdot 4t^{2q-2}$.

Theorem

$\mathfrak{M}_C(SU(2), \mathbf{N}) \cong \mathfrak{M}_C(PSU(2), \mathbf{N})/C_2$ *where* $C_2 = \{1, c\}$ *acts by* $c \cdot m = -m, c \cdot t = t$.

What is the "Coulomb mechanism"?

There are natural scalars from the vector multiplet that define a map $\mathfrak{M}_C(G, \mathbf{N}) \to \mathfrak{t} \otimes \mathrm{im}(\mathbb{H})/W$. In the mathematics literature, it's more common to look at the map $\mathfrak{M}_C(G, \mathbf{N}) \to \mathfrak{t}_{\mathbb{C}}/W$ induced by two components.

Basic idea:

Leaves in the Coulomb branch are mostly, but not perfectly, determined by which of these scalars are massless.

We've already seen the reason I have to say "mostly": $\mathfrak{M}_C(PSU(2), \mathbb{C}^3) \cong T^*\mathbb{C}^\times/W$ has 2 zero-dimensional leaves, at

$$(m,t) = (1,0) \qquad (m,t) = (-1,0).$$

The scalar t can't distinguish these.

In general, any leaf which is a single point must have all scalars vanish, but there might be more than one.

Theorem (Bellamy, Weekes)

$\mathfrak{M}_C(G, \mathbf{N})$ *has symplectic singularities and thus finitely many symplectic leaves.*

We call $\mathfrak{M}_C(G, \mathbf{N})$ **good** if its natural $\mathbb{R}_{\geq 0}$-action contracts to a single point, which is a leaf, **ugly** if contracting, but fixed point is not a leaf, **bad** otherwise

- If $\mathfrak{M}_C(G^1_\lambda, \mathbf{N}_0)$ is good, then life is simple: you get exactly one zero-dimensional leaf, defined by the vanishing of all monopole operators and scalar fields.
- If $\mathfrak{M}_C(G^1_\lambda, \mathbf{N}_0)$ is ugly, then there are no zero-dimensional leaves.
- If $\mathfrak{M}_C(G^1_\lambda, \mathbf{N}_0)$ is bad, it could go either way: sometimes no zero-dimensional leaves, but sometimes more than one. $\mathfrak{M}_C(G, \mathfrak{g}) \cong T^*\check{T}/W$ with G simple gives lots of examples.

I have some hope that the adjoint representation will be the only interesting bad case, but it remains unclear.

Consider any $\lambda \in \mathfrak{t}_{\mathbb{C}}$. Let

- $\mathbf{N}_\lambda$ be the fixed locus of λ ("the massless hypermultiplets")
- G_λ be the centralizer of λ in G
- G_λ^0, G_λ^1 be the kernel and image of G_λ acting on $\mathbf{N}_\lambda$ (G_λ^0 is our "unbroken symmetry.")

Theorem

Around a point in the fiber over λ, the variety $\mathfrak{M}_C(G, \mathbf{N})$ looks like $\mathfrak{M}_C(G_\lambda^{(1)}, \mathbf{N}_\lambda) \times \mathbb{C}^{2\,\mathrm{rk}(G_\lambda^0)}$.

Theorem (W.)

The symplectic leaves of $\mathfrak{M}_C(G, \mathbf{N})$ are in bijection with pairs $(\mathbf{N}_0, L)$ modulo the action of $\mathcal{Z}^\vee$ where

- $\mathbf{N}_0 \subset \mathbf{N}$ *is the fixed locus of a cocharacter λ (generic with this property)*
- *L is a zero-dimensional leaf of $\mathfrak{M}_C(G^1_\lambda, \mathbf{N}_0)$.*
- *$\mathcal{Z}^\vee$ is a finite group acting on $\mathfrak{M}_C(G^1_\lambda, \mathbf{N}_0)$ arising from failure of $G_\lambda \to G^1_\lambda$ to split.*

A neighborhood of the whole fiber over λ is a fiber bundle over $T^*\check{T}^0_\lambda$ with fiber $\mathfrak{M}_C(G^1_\lambda, \mathbf{N}_0)$, and $\mathcal{Z}^\vee$ captures the monodromy of this bundle.

Examples:

- If G is abelian and acts faithfully of $\mathbf{N}$ then $\mathfrak{M}_C(G, \mathbf{N}_0)$ is good or ugly, good if "coloop-free" (i.e. dual action is faithful). Recovers classification of leaves in terms of coloop-free flats.
- In the quiver case ($W_i = 0$), λ gives a decomposition $\mathbf{v} = \mathbf{v}^{(1)} + \cdots + \mathbf{v}^{(r)}$ in terms of the eigenspaces of λ. Good if and only if $\exists$ a simple module over preprojective algebra with dimension vector $\mathbf{v}^{(i)}$ for each i.

General expectation

Since a smaller leaf in the Coulomb branch has massless hypermultiplets and only partly breaks the symmetry, it lies in a mixed branch that gives a leaf in the Higgs branch.

If you believe I used the words “unbroken symmetry” and “massless hypermultiplets” correctly before, I shouldn’t have to explain what leaf it is.

One hope

There is a point in the Higgs branch where G_λ^0 is exactly the stabilizer G_x (or maybe just connected component G_x°).

Theorem

If $G_\lambda^0 = G_x^\circ$, the normalization of $\mathfrak{M}_H(G_\lambda^1, \mathbf{N}_0)$ is the affinization of a cover of a stratum in $\mathfrak{M}_H(G, \mathbf{N})$.

Importance of matching slices on one side to covers of strata on the other has shown up in work of Losev–Mason Brown–Matvieievskyi on the case of nilpotent cones in types B/C/D.

The picture works beautifully for a finite type quiver (with square vertices):

- The stabilizers for closed orbits are $GL_{v_1'} \times \cdots \times GL_{v_r'}$ with $\Lambda' = \Lambda - \sum(r_i - r_i')\alpha_i$ dominant.
- The groups G^0_λ for good slices are exactly the same.

Theorem

For a finite type quiver (with square vertices), there is an order reversing bijection between strata in $\mathfrak{M}_H(G,N)$ and $\mathfrak{M}_C(G,N)$ such that strata are mirror to slices in both directions.

Consider $(G = GL_n, N = \mathfrak{gl}_n \oplus \mathbb{C}^n)$.

In this case, $\mathfrak{M}_C(G,N) \cong \mathfrak{M}_H(G,N) \cong \mathrm{Sym}^n(\mathbb{C}^2)$. Leaves are given by partitions $\lambda = (\lambda_1 \geq \lambda_2 \geq \cdots)$

- The stabilizers for closed orbits are $GL_{\lambda_1} \times \cdots \times GL_{\lambda_r}$.
- The groups G^0_λ for good slices are $(GL_1)^r$ for $r \leq n$.

Only overlap is $\lambda_1 = \cdots = \lambda_r = 1$, so matching all of $\mathfrak{M}_C(G,N)$ with all of $\mathfrak{M}_H(G,N)$.

Life is complicated!