Symplectic leaves of Coulomb branches

Ben Webster

University of Waterloo & Perimeter Institute

October 30, 2026

Slides: `https://bwebste.github.io/Sydney-2026.pdf`

Fix a reductive group G and a representation $\mathbf{N}$. The **Coulomb branch** $\mathcal{M}(G, \mathbf{N})$ is a symplectic variety constructed via quantum field theory.

Two main questions

1. How can we describe $\mathcal{M}(G, \mathbf{N})$ in a more geometric way, which mathematicians might have a chance of understanding?
2. One basic question: $\mathcal{M}(G, \mathbf{N})$ has symplectic singularities, so it has finitely many **symplectic leaves**. What is their classification, and what are their transverse slices?

An important 3d theory: given G a connected reductive group and $\mathbf{N}$ a representation of G, we have a theory $\mathcal{T}(G,\mathbf{N})$ given by gauging the 3d σ-model with target $T^*\mathbf{N}$.

We 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 **Coulomb** and **Higgs** branches $\mathfrak{M}_C$ and $\mathfrak{M}_H$.

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 $T^*\mathbf{N}/\!\!/\!\!/G$.

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

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

Taylor series $\mathsf{C} = \mathbb{C}[[t]]$	$\mathsf{G} = G[[t]]$	$\mathsf{N} = \mathbf{N}[[t]]$
Laurent series $\mathcal{C} = \mathbb{C}((t))$	$\mathcal{G} = G((t))$	$\mathcal{N} = \mathbf{N}((t))$

Relevant spaces:

$$\mathsf{Y} = \mathsf{N}/\mathsf{G} \qquad \mathcal{Y} = \mathcal{N}/\mathcal{G}$$

These can be interpreted as spaces of principal G-bundles with a section of the associated $\mathbf{N}$-bundle on D and D^*.

Thus, the fiber product $\mathsf{Y} \times_{\mathcal{Y}} \mathsf{Y}$ is the space of such bundles on the "raviolo" gluing two copies of D along D^*.

Definition (BFN)

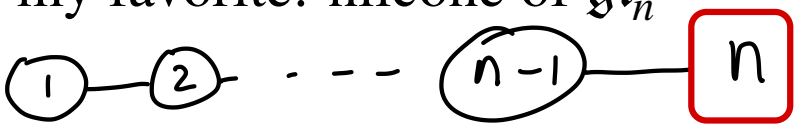
The **Coulomb branch** *is the spectrum* $\mathcal{M} = \operatorname{Spec} H^{BM}_*(\mathsf{Y} \times_{\mathcal{Y}} \mathsf{Y})$.

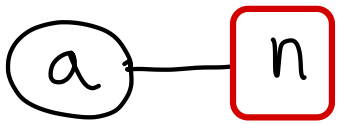
Most important case: let Γ be a quiver, and $\mathbf{v}\colon I = \mathcal{V}(\Gamma) \to \mathbb{Z}_{\geq 0}$ a dimension vector. Let

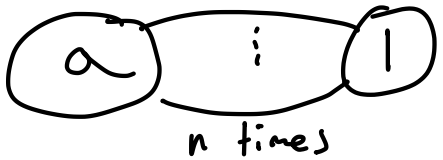
$$G_{\mathbf{v}} = (\prod_i \mathrm{GL}_{v_i})/\mathbb{C}^\times \qquad \mathbf{N}_{\mathbf{v}} = \bigoplus_{i\to j} \mathrm{Hom}(\mathbb{C}^{v_i}, \mathbb{C}^{v_j})$$

These recover interesting varieties, always symplectic:

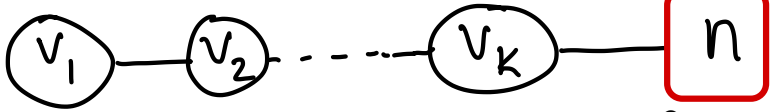
- my favorite: nilcone of $\mathfrak{gl}_n$



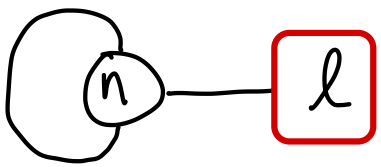




- more generally, Slodowy slices in type A



- symmetric power $\mathrm{Sym}^n(\mathbb{C}^2/\mathbb{Z}_\ell)$



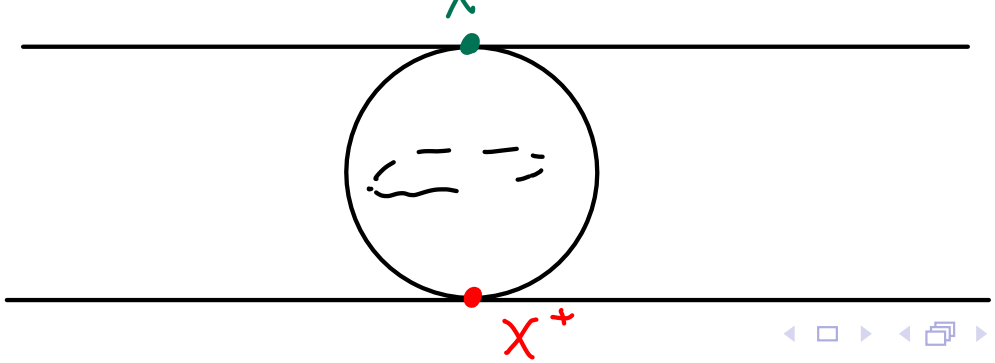
The important idea is to think about the interplay between characters e^λ on $T^\vee$ and the weights of $\mathbf{N}$ as functions on $\mathfrak{t}$.

Most important case: $G = \mathbb{G}_m$. Weights of $\mathbf{N}$ are $n_1, \dots, n_k \in \mathbb{Z}$. Let

$$E_\pm = \sum_{\pm n_i > 0} |n_i|.$$

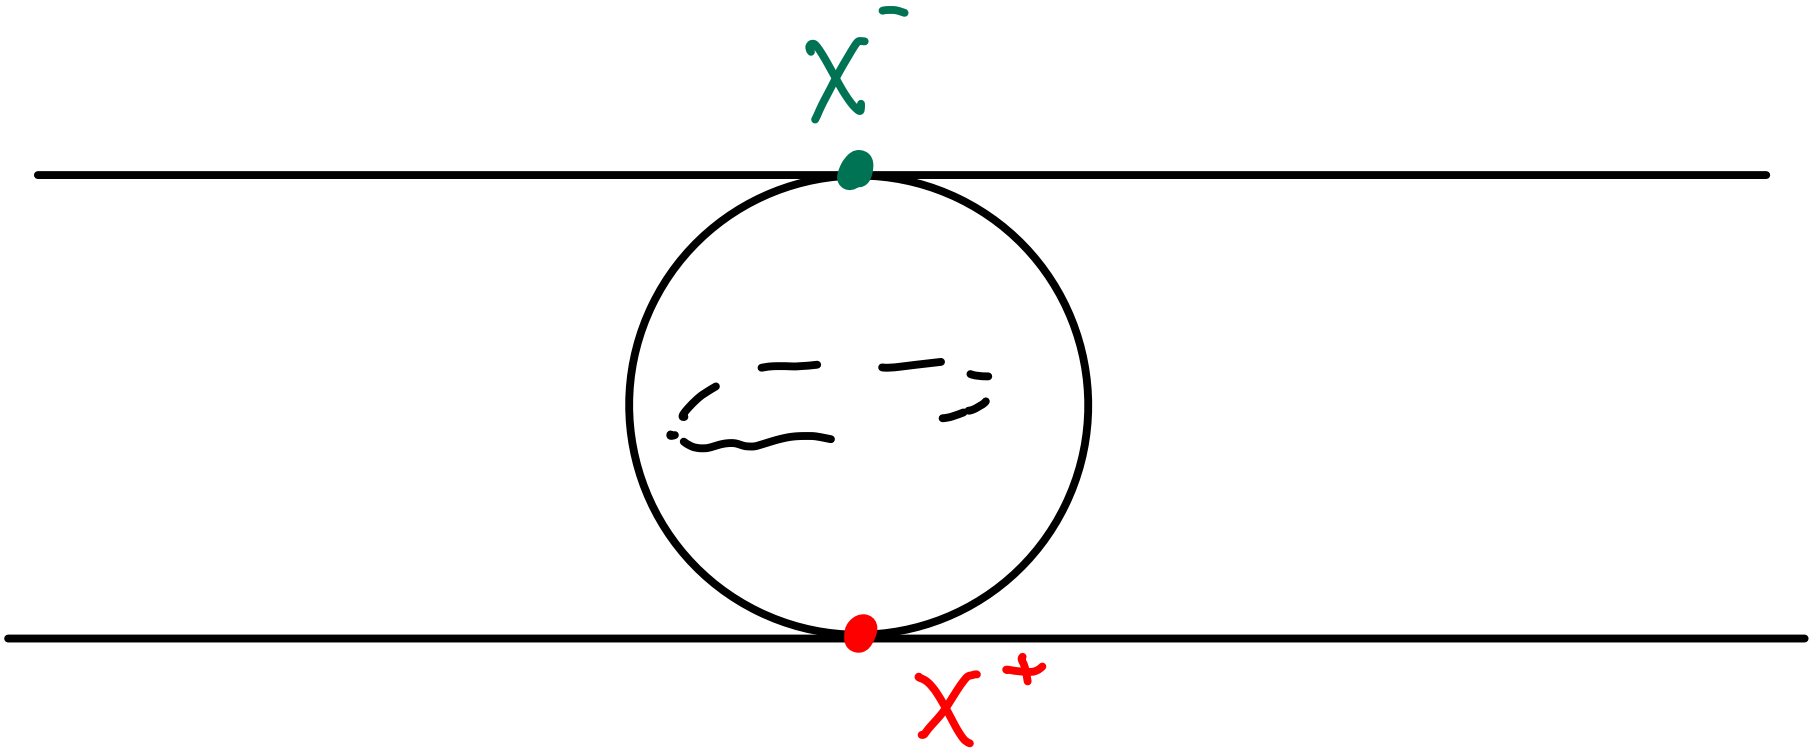
$\mathcal{M}$ is birational to $\mathbb{A}^1 \times \mathbb{G}_m \subset \mathbb{A}^1_\tau \times \mathbb{P}^1_z$. Consider the subschemes in $\mathbb{A}^1 \times \mathbb{P}^1_z$ defined by

$$X^+ = Z(\tau^{E_+}) \times \{0\}, \quad X^- = Z(\tau^{E_-}) \times \{\infty\}.$$



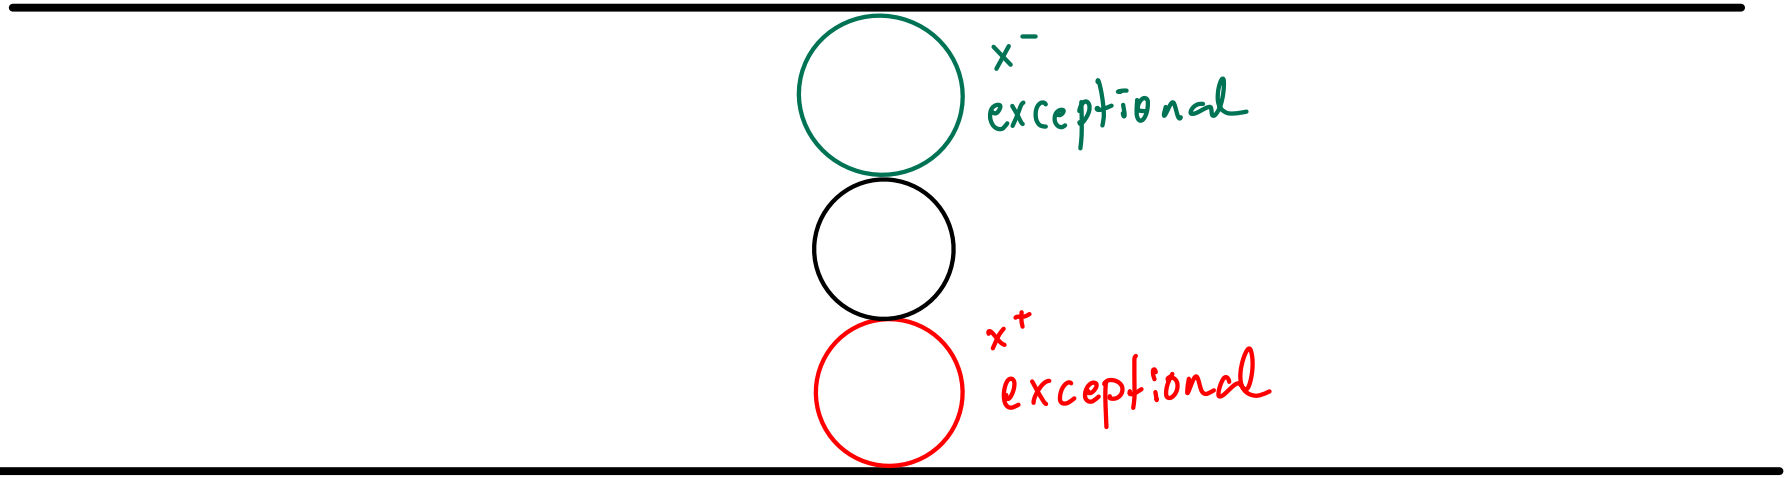
Theorem

$\mathcal{M}(\mathbb{G}_m, \mathbf{N})$ is the affinization of the blowup of $\mathbb{A}^1_\tau \times \mathbb{P}^1_z$ at $X^+ \cup X^-$ with strict transforms of $\mathbb{A}^1_\tau \times \{0\}$ and $\mathbb{A}^1_\tau \times \{\infty\}$ removed.



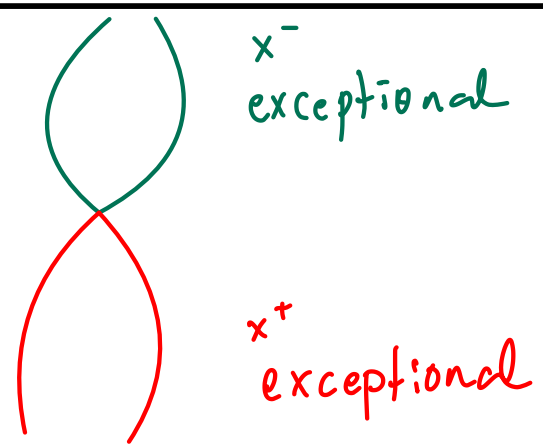
Theorem

$\mathcal{M}(\mathbb{G}_m, \mathbf{N})$ is the affinization of the blowup of $\mathbb{A}^1_\tau \times \mathbb{P}^1_z$ at $X^+ \cup X^-$ with strict transforms of $\mathbb{A}^1_\tau \times \{0\}$ and $\mathbb{A}^1_\tau \times \{\infty\}$ removed.



Theorem

$\mathcal{M}(\mathbb{G}_m, \mathbf{N})$ is the affinization of the blowup of $\mathbb{A}^1_\tau \times \mathbb{P}^1_z$ at $X^+ \cup X^-$ with strict transforms of $\mathbb{A}^1_\tau \times \{0\}$ and $\mathbb{A}^1_\tau \times \{\infty\}$ removed.



This is $\mathbb{A}^2/\mathbb{Z}/\ell\mathbb{Z}$ $= \operatorname{Spec} \mathbb{K}[T^{\ell-z}, T^{\ell+z^{-1}}, t]$

$E = E^+ + E^-$

But birational map important!

General case: need to do a similar blowup for each of the lines $\ell_1, \ldots, \ell_k$ spanned by weights of $\mathbf{N}$.

Teleman gluing

If the weights of $\mathbf{N}$ are generic enough, then all the new divisors can be written as the images of $\ell_i^\perp \times T^\vee$ under a birational map, so $\mathcal{M}(T, \mathbf{N})$ is the affinization of $\mathfrak{t} \times T^\vee$ glued to itself along this birational map.

Key motivation: You can use other cohomology theories to get K-theoretic and elliptic Coulomb branches. Construction is the same with:

- $\mathfrak{t}$ replaced by T or $E \otimes_{\mathbb{Z}} X_*(T)$.
- the weights and roots replaced by Euler classes in the appropriate cohomology theory.

Passing from T to G requires a further birational modification: if no root lies in ℓ_i, then we blow up the subscheme in $\mathfrak{t} \times T^\vee$ given by $Z(\alpha, e^{\alpha^\vee} - 1)$ for each root α, and remove the strict transform of $Z(\alpha)$.

Theorem

The Coulomb branch $\mathcal{M}(G, \mathbf{N})$ is the affinization of this birational modification of $\mathfrak{t} \times T^\vee$ modulo W.

If roots do lie in ℓ_i, then fixable, but need to be more careful.

Put differently: $\mathcal{M}(G, \mathbf{N})$ is a normal variety, so we can describe it by giving Weil divisors. These are only different from $(\mathfrak{t} \times T^\vee)/W$ over the divisors where roots or weights vanish, and the modification is obtained from those divisors in a uniform geometric way.

- Over a weight divisor where no root vanishes, we have the two divisors coming from the blowups as in the abelian case.
- Over a root divisor, we usually only have one divisor, since the reflection combines them (though there are a couple of exceptions where there are two).

Important consequence: removing a divisor has the same effect as removing the corresponding weight from $\mathbf{N}$ or root from G.

Given $\lambda \in \mathfrak{t}$, we have the centralizer $G_\lambda = C_G(\lambda)$ with Weyl group W_λ and fixed representation $\mathbf{N}^\lambda \subset \mathbf{N}$.

Let $V \subset \mathfrak{t}$ be the complement of the divisors corresponding to roots or weights which don't vanish at λ.

Theorem

The varieties $\mathcal{M}(G, \mathbf{N})$ and $\mathcal{M}(G_\lambda, \mathbf{N}^\lambda)$ over V/W_λ are isomorphic.

$$\begin{array}{ccccc} \mathcal{M}(G,\mathbf{N}) & \longleftarrow & \mathcal{M}^\circ(G_\lambda,\mathbf{N}^\lambda) & \hookrightarrow & \mathcal{M}(G_\lambda,\mathbf{N}^\lambda) \\ \downarrow & & \downarrow & & \downarrow \\ \mathfrak{t}/W & \longleftarrow & V/W_\lambda & \hookrightarrow & \mathfrak{t}/W_\lambda \end{array}$$

A **symplectic singularity** has a canonical stratification into finitely many **symplectic leaves**: locally closed smooth symplectic subvarieties. These are the orbits of the Hamiltonian vector fields.

Fundamental questions for $\mathcal{M}(G, \mathbf{N})$:

1. Index the leaves (how many are there? how to describe them?).
2. Identify the **transverse slice** to each leaf (another symplectic singularity).

In the abelian case $G = T$, the answer is well-known from the Higgs branch description:

- Leaves are in bijection with **co-loop-free flats** $F \subset \mathfrak{t}$.
  - A **flat** F is an intersection of hyperplanes H_φ for weights φ of $\mathbf{N}$; this is precisely the stabilizer in $\mathfrak{t}$ of a point in $\mathbf{N}$.
  - A **co-loop** of F is a weight φ vanishing on F such that there exists a cocharacter ξ of T with $\xi(\varphi) = 1$ and $\xi(\psi) = 0$ for all other weights ψ vanishing on F.
- If $\lambda \in F$ is generic and $T_F \subset T$ is the subtorus with Lie algebra F, then the slice to this leaf is $\mathcal{M}(T/T_F, \mathbf{N}^\lambda)$.

If a flat has a co-loop, it still gives a subspace, but not a symplectic leaf: consider $\mathcal{M}(\mathbb{G}_m, \mathbb{C}) \cong \mathbb{C}^2$.

Quite suggestive, but much simpler than the general case.

The distinction between good, bad and ugly theories will be important for us:

- **good**: the scaling action contracts to a unique fixed point, which is a 0-dimensional symplectic leaf. Example: $\mathcal{M}(\mathbb{G}_m, \mathbb{C}^2) \cong \mathbb{C}^2/\mathbb{Z}_2$, abelian theories with finite kernel, no co-loops.
- **ugly**: the scaling action contracts to a fixed point, but this fixed point lies in a positive-dimensional leaf (in fact a symplectic vector space). Example: $\mathcal{M}(\mathbb{G}_m, \mathbb{C}) \cong \mathbb{C}^2$, abelian theories with finite kernel and co-loops.
- **bad**: the scaling action doesn't contract to a single fixed point. Example: $\mathcal{M}(\mathbb{G}_m, 0) \cong \mathbb{C} \times \mathbb{C}^\times$, abelian theories with infinite kernel.

Abelian case: a leaf appears exactly when $(T/T_F, \mathbf{N}^\lambda)$ is good.

There is a natural **integrable system**:

$$\varrho\colon \mathcal{M}(G,\mathbf{N}) \longrightarrow \mathfrak{t}/W,$$

induced by the map from the pure Coulomb branch $\mathcal{M}(G)$, which is a group scheme over $\mathfrak{t}/W$.

Over the open dense set of $\mathfrak{t}/W$ where no weights or roots vanish, the fiber of ϱ is a torsor for the dual torus $T^\vee$.

Can break the classification of leaves into two steps:

1. Classify the possible images of leaves in $\mathfrak{t}/W$.
2. For each image, classify the leaves lying over it.

Definition

A **flat** *of the generalized weight hyperplane arrangement in* $\mathfrak{t}$ *is an intersection of a subset of the hyperplanes* H_φ *defined by weights of* $\mathbf{N}$*; these are exactly the stabilizers in* $\mathfrak{t}$ *of points in* $\mathbf{N}$.

For a flat $U \subset \mathfrak{t}$ with generic point λ:

- $G_\lambda = C_G(\lambda)$ is a **Levi subgroup** of G
- $\mathbf{N}^\lambda = \mathbf{N}^\lambda$ is the fixed subspace of λ (a G_λ-representation)
- $G^0_\lambda = \ker^0(G_\lambda \to GL(\mathbf{N}^\lambda))$; $\quad G^1_\lambda = G_\lambda / G^0_\lambda$ almost faithful on $\mathbf{N}^\lambda$
- $\mathfrak{t}^0_\lambda =$ Lie algebra of maximal torus of G^0_λ (of dimension k)

Slice theorem (local structure)

Around the fiber of ϱ over λ, $\mathcal{M}(G, \mathbf{N})$ is étale locally symplectomorphic to $\mathcal{M}(G^1_\lambda, \mathbf{N}^\lambda) \times \mathbb{A}^{2k}$.

We have to be a little more careful about global structure, though.

- There might be a Galois group $\bar{\Gamma}_\lambda$ that appears in the comparison of $\mathcal{M}(G, \mathbf{N})$ and $\mathcal{M}(G_\lambda, \mathbf{N}^\lambda)$ over U.
- There is an isogeny $G_\lambda \to G_\lambda^1 \times \bar{G}_\lambda^0$ for $\bar{G}_\lambda^0$ a finite quotient; let Z_λ be the kernel of this map.

Lemma

The product $\mathcal{M}(G_\lambda^1 \times \bar{G}_\lambda^0, \mathbf{N}^\lambda) \cong \mathcal{M}(G_\lambda^1, \mathbf{N}^\lambda) \times \mathcal{M}(\bar{G}_\lambda^0, 0)$ admits a free action of $Z_\lambda^\vee$, and $\mathcal{M}(G_\lambda, \mathbf{N}^\lambda)$ is the quotient by this action.

The symplectic leaves of $\mathcal{M}(G_\lambda, \mathbf{N}^\lambda)$ are in bijection with $Z_\lambda^\vee$-orbits on the symplectic leaves of $\mathcal{M}(G_\lambda^1, \mathbf{N}^\lambda)$, and the transverse slice to a leaf in $\mathcal{M}(G_\lambda, \mathbf{N}^\lambda)$ is the same as the transverse slice to any leaf in the corresponding orbit in $\mathcal{M}(G_\lambda^1, \mathbf{N}^\lambda)$.

Theorem (W.)

1. *The image of any symplectic leaf of* $\mathcal{M}(G, \mathbf{N})$ *under* ϱ *is the image in* $\mathfrak{t}/W$ *of a flat* U.
2. *The leaves with image* U *are in bijection with* $Z_\lambda^\vee \rtimes \bar{\Gamma}_\lambda$*-orbits on* 0*-dimensional leaves of* $\mathcal{M}(G_\lambda^1, \mathbf{N}^\lambda)$.
3. *If* U *is* ***good*** *(i.e.,* $\mathcal{M}(G_\lambda^1, \mathbf{N}^\lambda)$ *is good), there is a* ***unique*** *leaf with image* U*, and its* ***transverse slice*** *is isomorphic to* $\mathcal{M}(G_\lambda^1, \mathbf{N}^\lambda)$.
4. *If* $\mathcal{M}(G_\lambda^1, \mathbf{N}^\lambda)$ *has no* 0*-dimensional leaves, there is* ***no*** *leaf with image* U*. Important case:* $\mathcal{M}(G_\lambda^1, \mathbf{N}^\lambda)$ *is ugly.*

Lemma

A leaf L with image U lies in the closure of L' with image U' if and only if:

1. *$U \subset U'$, and*
2. *the corresponding closure relation holds inside $\mathcal{M}(G^1_\lambda, \mathbf{N}^\lambda)$.*

In particular, if both U and U' are good, then $L \subset \overline{L'}$ if and only if $U \subset U'$.

This gives a complete description of the **closure poset** on good leaves: inclusion on good flats.

Consider the case $G = \mathrm{GL}_n$, $\mathbf{N} = \mathfrak{gl}_n$.

1. $\mathcal{M}(G, \mathbf{N}) \cong \mathrm{Sym}^n(\mathbb{C} \times \mathbb{C}^\times)$.
2. Flats: indexed by partitions $\lambda \vdash n$ (groups of equal coordinates).
3. For $\lambda = (n_1, \ldots, n_k)$:

$$G_\lambda^1 = \prod PGL_{n_i}, \quad \mathbf{N}^\lambda = \bigoplus \mathfrak{gl}_{n_i}.$$

4. Residual theory $\mathcal{M}(G_\lambda^1, \mathbf{N}^\lambda)$ has $\prod n_i$ zero-dimensional leaves, but $Z^\vee$ acts simply transitively $\Rightarrow$ **one leaf per partition**.

Now modify to $G = \mathrm{GL}_n$, $\mathbf{N} = \mathbb{C}^n \oplus \mathfrak{gl}_n$.

1. $\mathcal{M}(G, \mathbf{N}) \cong \mathrm{Sym}^n(\mathbb{C}^2)$.
2. Flats: indexed by an integer $n_0 \leq n$ (number of coordinates equal to 0) and by partitions $\lambda \vdash (n - n_0)$.
3. Slice described by

$$G^1_\lambda = GL_{n_0} \times \prod PGL_{\lambda_i}, \quad \mathbf{N}^\lambda = \mathbb{C}^{n_0} \oplus \mathfrak{gl}_{n_0} \oplus \bigoplus \mathfrak{gl}_{\lambda_i}.$$

   $(\mathrm{GL}_{n_0}, \mathbb{C}^{n_0} \oplus \mathfrak{gl}_{n_0})$ is **ugly** for all $n_0 > 0$, so $\mathcal{M}(G^1_\lambda, \mathbf{N}^\lambda)$ has no 0-dimensional leaves in this case.
4. For $n_0 = 0$, same flats as above give the same Hasse diagram.

One more: $G = \mathrm{GL}_n$, $\mathbf{N} = (\mathbb{C}^n)^{\oplus \ell} \oplus \mathfrak{gl}_n$ ($\ell > 1$).

1. $\mathcal{M}(G, \mathbf{N}) \cong \mathrm{Sym}^n(\mathbb{C}^2/\mathbb{Z}_\ell)$.
2. Flats indexed by pairs (n_0, λ) with $n_0 \leq n$, $\lambda \vdash (n - n_0)$.
3. Slice described by

$$G^1_\lambda = GL_{n_0} \times \prod PGL_{\lambda_i}, \quad \mathbf{N}^\lambda = (\mathbb{C}^{n_0})^{\oplus \ell} \oplus \mathfrak{gl}_{n_0} \oplus \bigoplus \mathfrak{gl}_{\lambda_i}.$$

4. For $n_0 > 0$, $(\mathrm{GL}_{n_0}, (\mathbb{C}^{n_0})^{\oplus \ell} \oplus \mathfrak{gl}_{n_0})$ is **good**. So there is a leaf for each pair (n_0, λ), given by n_0 points at the origin and $n - n_0$ points in $\mathbb{C}^2/\mathbb{Z}_\ell$ with multiplicities given by λ.

These are special cases of quiver gauge theories.

$$G_{\mathbf{v}} = (\prod_i \mathrm{GL}_{v_i})/\mathbb{C}^\times \qquad \mathbf{N}_{\mathbf{v}} = \bigoplus_{i\to j} \mathrm{Hom}(\mathbb{C}^{v_i}, \mathbb{C}^{v_j})$$

A cocharacter for a quiver gauge theory is a grading of the corresponding quiver representation, so the flat is described by the dimension vectors of the graded pieces.

"Fission and decay"

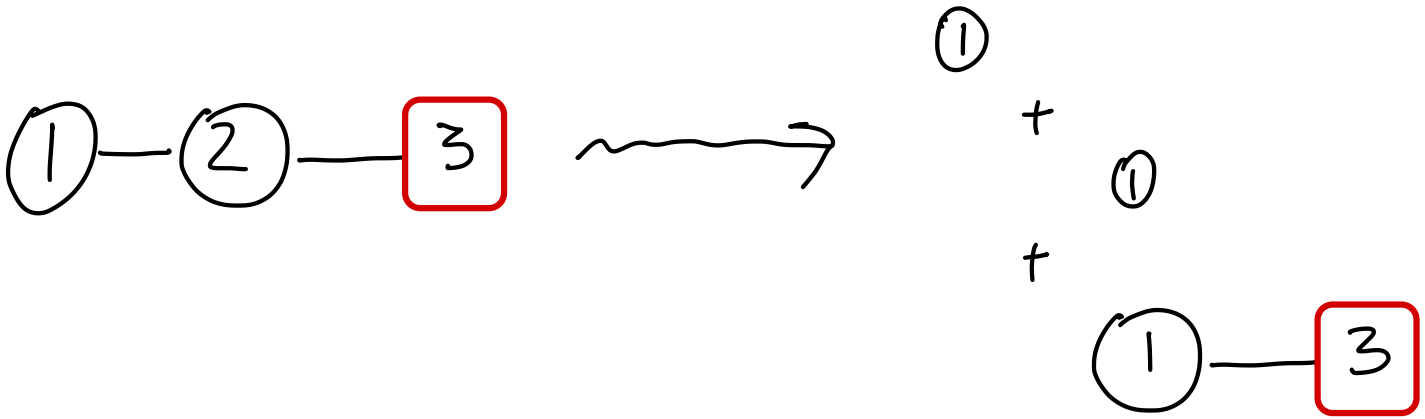
Writing $\mathbf{v} = \sum \mathbf{v}^{(i)}$ as a sum of dimension vectors:

1. "Decay": If $\mathbf{v}^{(i)}$ is a unit vector, trivial contribution to the slice.
2. "Fission": If more than one of the summands $\mathbf{v}^{(i)}$ is not a unit vector, then the slice is the product of contributions from the nontrivial $\mathbf{v}^{(i)}$.

Note that the framing is expressed by adding a new vertex, so it all goes with one summand.

For a framed Dynkin quiver, the only contributions come from dimension vectors of dominant weights, with unit vectors "decayed off."

This matches the results of Muthiah and Weekes.



Let $\mathrm{Leaf}(\mathcal{M}_{\mathsf{Coulomb}})$ and $\mathrm{Leaf}(\mathcal{M}_{\mathsf{Higgs}})$ denote the leaf sets of the Coulomb and Higgs branches.

Definition

Define $\mathbf{CH} \subset \mathrm{Leaf}(\mathcal{M}_{\mathsf{Coulomb}}) \times \mathrm{Leaf}(\mathcal{M}_{\mathsf{Higgs}})$ *by* $(L, L') \in \mathbf{CH}$ *if, for a generic point* $[x] \in L'$*, the image of* L *under* ϱ *matches the flat* $\mathfrak{t}_x$ *(Lie algebra of the torus of the stabilizer* G_x*).*

Physics expects an **order-reversing bijection** between leaves, but several obstructions arise:

1. Multiple Coulomb leaves over the same flat.
2. A flat may not arise as $\mathfrak{t}_x$ for any Higgs leaf.
3. Multiple Higgs leaves with the same stabilizer torus.

Example: $G = \mathrm{GL}_n$, $\mathbf{N} = \mathbb{C}^n \oplus \mathfrak{gl}_n$. Every Higgs leaf has $\mathfrak{t}_x = \mathfrak{t}$, so only the open Coulomb leaf appears.

Conjecture

*There is a natural notion of **special** leaf on each side such that* **CH** *restricts to a bijection between the special leaves of* $\mathcal{M}_{\mathsf{Coulomb}}$ *and* $\mathcal{M}_{\mathsf{Higgs}}$.

The special Coulomb leaves are exactly those whose image under ϱ *is a **good** flat.*

Evidence:

- **Toric case** ($G = T$): all leaves are special; bijection is obvious.
- **Dynkin quivers**: all leaves are special; bijection is obvious.
- $G = \mathrm{GL}_n$, $\mathbf{N} = (\mathbb{C}^n)^{\oplus \ell} \oplus \mathfrak{gl}_n$: special leaves are those with $\lambda = (1, \ldots, 1)$. Order-reversing bijection with special Higgs leaves, but not all leaves.

1. **Zero-dimensional leaves in bad/ugly theories.** When does a bad flat admit a 0-d leaf? Often none exist, but sometimes they do. The example $G = PGL_n$, $\mathbf{N} = \mathfrak{pgl}_n$ shows that any number of zero-dimensional leaves is possible. Bourget *et al.* conjecture that known affine Dynkin quiver examples are the only quiver ones.
2. **Elliptic Coulomb branches.** No consensus definition exists yet. Our construction gives a candidate; comparison with BFN is pending more development in the elliptic literature.
3. **Noncotangent cases.** If $\mathbf{E}$ is a symplectic representation of G that is not of cotangent type, then $\mathcal{M}(G, \mathbf{E})$ can still be defined if certain obstructions vanish. Leaf structure is unknown, but should be similar.