The Geometry of the Space of Celestial CFTs

Dan Kapec

Harvard University, Center of Mathematical Sciences and Applications

AdS vs. AFS

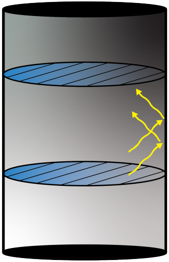

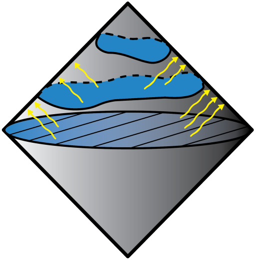

Many important qualitative differences. The boundary is null rather than timelike, we get balance laws rather than conservation laws along $\mathcal{I}$:

$$\partial_u m_B = -T_{uu} - N_{ab}N^{ab} + \cdots$$

Something like AdS/CFT had to hold if we believe black hole thermo.

The idea that led to the AdS/CFT correspondence was to study the response of a near-extremal black hole/brane to low energy probes.

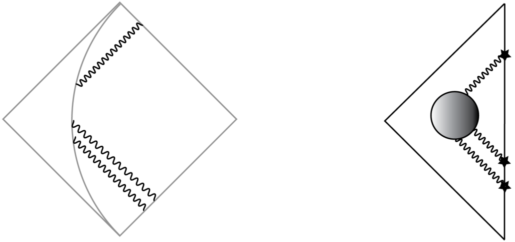

Replace the probes with sources at the boundary of AdS, do the calculation in the throat geometry, and then match back onto the far region.

But if the BH is really a quantum system, and the dynamics is occurring right outside the black hole, then you can do the calculation another way by studying the response of the quantum system to these sources.

In the best understood examples we can do the calculation on both sides and they match. In the other cases we use this idea to learn about the QM.

In global AdS, black holes give us a window into the generic properties of high energy eigenstates in the dual quantum mechanics.

The fact that the thermodynamics of a large AdS black hole resembles that of a non-gravitational CFT is a huge clue to the correspondence.

Black holes behave very differently in flat space.

The Bekenstein-Hawking entropy is super-Hagedorn, the specific heat is negative, the black holes are really long-lived metastable resonances.

So it is possible that a holographic dual to flat space is going to look very different than the quantum mechanics of AdS, if it exists at all.

However, the black hole entropy still scales like the area, and the gravitational Hamiltonian is still a boundary term.

We don't expect any precisely defined local observables, so some form of holography might still apply to QG with flat asymptotics.

But we are far from a precise statement of the correspondence.

Simple kinematic observation: massless particles moving in asymptotically flat spacetime pass through $\mathcal{I}^{\pm}$ at isolated points on the celestial sphere.

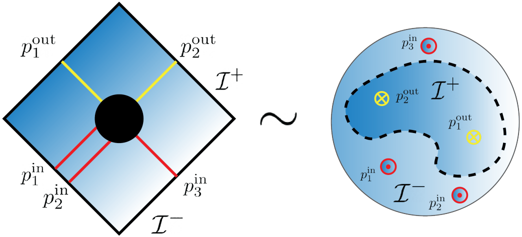

This is reflected in momentum space since we can parameterize

$$p^\mu(\omega, x) = \omega \hat{p}^\mu(x)$$

$$\hat{p}^\mu(x) = \left(\frac{1+x^2}{2}, x^a, \frac{1-x^2}{2}\right) , \quad \omega \geq 0 , \quad x^a \in \mathbb{R}^d .$$

The Lorentzian inner product is then given by

$$-2\hat{p}(x_1) \cdot \hat{p}(x_2) = (x_1 - x_2)^2 \ .$$

So the short distance expansion on $\mathcal{M}_d$ is the collinear expansion:

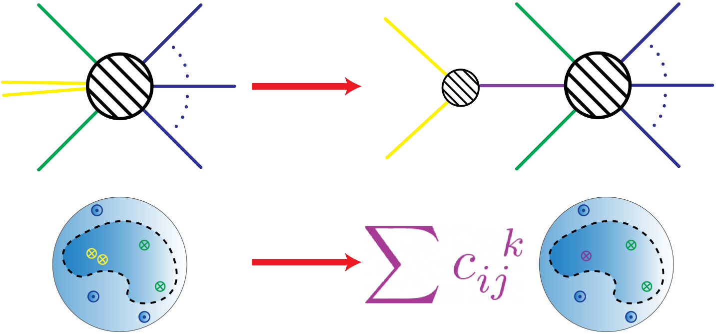

Collinear factorization of amplitudes looks a lot like the OPE on $\mathcal{M}_d$.

In asymptotically flat $(d+2)$-dimensional spacetimes, the Lorentz group $SO(d+1,1)$ is realized at $\mathcal{I}^{\pm}$ as conformal transformations of S^d.

So the S-matrix in $(d+2)$-dimensions can be re-expressed as a correlation function of (local?) operators on S^d, and **Lorentz invariance guarantees that these operators transform like those in a Euclidean CFT$_d$.**

$$\langle p_1(\omega_1, x_1), \cdots | \cdots p_n(\omega_n, x_n)\rangle \to \langle \mathcal{O}_1(\Delta_1, x_1) \cdots \mathcal{O}_n(\Delta_n, x_n)\rangle \ .$$

So we expect the massless S-matrix to display some (but not all) of the features of a d-dimensional Euclidean conformal field theory.

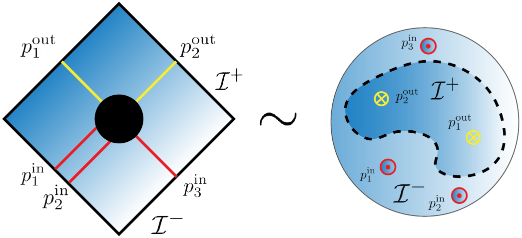

Universal Properties of CCFT: Global Symmetries

Conformal field theories have a small set of universal local operators.

For instance, every local CFT comes with a stress tensor $T_{ab}(x)$.

The existence of this operator in CCFT is guaranteed by the subleading soft graviton theorem. We can explicitly define an operator satisfying all the appropriate Ward identities in any dimension [DK, Mitra '18]. Spin-2 ✓

When the CFT has a global symmetry, the local operator spectrum includes a conserved current $J_a(x)$. These operators come from the leading soft photon, soft gluon, and soft graviton theorems. Spin-1 ✓

The soft expansion encodes universal information, and the Mellin transform turns Laurent expansions into poles with residues given by soft operators

$$\langle O_{ab}(\omega, x) \cdots \rangle \sim \frac{S_{ab}^{(n)}(x)}{\omega^n} \langle \cdots \rangle \iff \langle O_{ab}^{\Delta}(x) \cdots \rangle \sim \frac{S_{ab}^{(n)}(x)}{\Delta - n} \langle \cdots \rangle$$

The "leading soft-photon operator"

$$S_a(x) = \oint \frac{d\omega}{2\pi i}\mathcal{O}_a(\omega, x) \sim \lim_{\omega\to 0} \omega\mathcal{O}_a(\omega, x) \sim \lim_{\Delta\to 1}(\Delta-1)\int \frac{d\omega}{\omega}\omega^{\Delta}\mathcal{O}_a(\omega, x)$$

is a conformal primary operator with $(\Delta, s) = (1, 1)$. Its matrix elements are completely controlled by the soft-photon theorem:

$$\begin{aligned}\langle S_a(x)\mathcal{O}_1(\omega_1, x_1)\ldots\mathcal{O}_n(\omega_n, x_n)\rangle &= 2\sum_{k=1}^{n} Q_k \frac{(x-x_k)_a}{(x-x_k)^2}\langle\mathcal{O}_1(\omega_1, x_1)\ldots\mathcal{O}_n(\omega_n, x_n)\rangle \\ &= \partial_a \sum_{k=1}^{n} Q_k \log\left[(x-x_k)^2\right]\langle\mathcal{O}_1(\omega_1, x_1)\ldots\mathcal{O}_n(\omega_n, x_n)\rangle.\end{aligned}$$

A conserved current $J_a(x)$ is a primary operator with $(\Delta, s) = (d-1, 1)$ satisfying the Ward identity

$$\langle\partial^b J_b(y)\mathcal{O}_1(\omega_1, x_1)...\mathcal{O}_n(\omega_n, x_n)\rangle = \sum_{k=1}^{n} Q_k \delta^{(d)}(y-x_k)\langle\mathcal{O}_1(\omega_1, x_1)...\mathcal{O}_n(\omega_n, x_n)\rangle.$$

This is not quite the soft photon theorem, but it is closely related.

Multiply both sides of the Ward identity by $\int d^d y \partial_a \log[(x-y)^2]$

$$\int d^d y \partial_a \log[(x-y)^2] \langle \partial^b J_b(y) \mathcal{O}_1(\omega_1, x_1) \dots \mathcal{O}_n(\omega_n, x_n) \rangle$$
$$= \partial_a \sum_{k=1}^{n} Q_k \log[(x-x_k)^2] \langle \mathcal{O}_1(\omega_1, x_1) \dots \mathcal{O}_n(\omega_n, x_n) \rangle$$
$$= \langle S_a(x) \mathcal{O}_1(\omega_1, x_1) \dots \mathcal{O}_n(\omega_n, x_n) \rangle \ .$$

Integrating by parts expresses $S_a(x)$ as an integral transform of $J_a(x)$:

$$S_a(x) = \int d^d y \partial_a \log[(x-y)^2] \partial^b J_b(y) = 2 \int d^d y \frac{\mathcal{I}_{ab}(x-y)}{(x-y)^2} J^b(y) \ ,$$

where $\mathcal{I}_{ab}(x-y)$ is the conformally covariant tensor

$$\mathcal{I}_{ab}(x-y) = \delta_{ab} - 2 \frac{(x-y)_a (x-y)_b}{(x-y)^2} \ .$$

Shadow Transform

This nonlocal relationship between the $\Delta = 1$ primary S_a and the $\Delta = d-1$ primary J_a is known as a shadow transform.

For a spin-s operator of dimension Δ, the shadow operator is given by

$$\widetilde{\mathcal{O}}_{\mathcal{R}}(x) = \int d^d y \frac{\mathcal{R}(\mathcal{I}_{ab})}{[(x-y)^2]^{d-\Delta}} \mathcal{O}_{\mathcal{R}}(y) \ .$$

The shadow transform maps conformal primary operators with (Δ, s) onto conformal primary operators with $(d-\Delta, s)$.

The shadow transform is, up to normalization, its own inverse

$$\widetilde{\widetilde{O}}_{a_1 \dots a_s}(x) = c(\Delta, s) \mathcal{O}_{a_1 \dots a_s}(x) \ ,$$

Using this, we can immediately write

$$S_a(x) = 2 \widetilde{J}_a(x) \ , \qquad J_a(x) = \frac{1}{2c(1,1)} \widetilde{S}_a(x) \ .$$

Similarly the bulk subleading soft-graviton operator

$$B_{ab}(x) = \oint \frac{d\omega}{2\pi i} \frac{\mathcal{O}_{ab}(\omega, x)}{\omega} \sim \lim_{\Delta \to 0} \Delta \int \frac{d\omega}{\omega} \omega^{\Delta} \mathcal{O}_{ab}(\omega, x)$$

is related to the boundary stress tensor:

$$B_{ab}(x) = -\widetilde{T}_{ab}(x) \ .$$

This could have been guessed based on the dimensions of B_{ab} and T_{ab}.

We can then invert the shadow transform to find an operator

$$T_{ab}(x) = -\frac{1}{c(0,2)} \widetilde{B}_{ab}(x)$$

that satisfies

$$\langle \partial^d T_{dc}(y) \mathcal{O}_1(\omega_1, x_1) \dots \mathcal{O}_n(\omega_n, x_n) \rangle = -\sum_{k=1}^{n} \delta^{(d)}(y - x_k) \partial_{x_k^c} \langle \mathcal{O}_1(\omega_1, x_1) \dots \mathcal{O}_n(\omega_n, x_n) \rangle$$

along with the all the other Ward identities for a stress tensor.

Summary

The shadow transforms of the leading soft photon operator and the subleading soft graviton operator

$$J_a(x) = \frac{1}{2c_{1,1}}\widetilde{S}_a(x) \, , \qquad T_{ab}(x) = -\frac{1}{c_{0,2}}\widetilde{S}_{ab}(x) \, ,$$

define operators in celestial CFT that obey all the Ward identities of a conserved current and stress tensor:

$$\langle \partial^a J_a(x) \mathcal{O}_1 \cdots \mathcal{O}_n \rangle = \sum_k Q_k \delta^{(d)}(x - x_k) \langle \mathcal{O}_1 \cdots \mathcal{O}_n \rangle \, ,$$
$$\langle \partial^a T_{ab}(x) \mathcal{O}_1 \cdots \mathcal{O}_n \rangle = \sum_k \delta^{(d)}(x - x_k) \partial_{x_k^b} \langle \mathcal{O}_1 \cdots \mathcal{O}_n \rangle \, .$$

Just as in AdS, bulk gauge fields $\implies$ boundary currents and the graviton yields a stress tensor.

Universal Properties of CCFT: Conformal Manifolds

What about spin-0 ? What would a "universal" spin-0 operator be?

Spin zero operators are operators that you can use to deform the model.

In the context of CCFT, global conformal symmetry corresponds to Lorentz invariance and cannot be violated: **the most interesting deformations are those that preserve conformal invariance**.

So we would like to understand the space of exactly marginal deformations, or equivalently the **conformal manifold**, of CCFT [DK, Law, Narayanan '22].

This is the class of universal spin-0 operators we get from the soft theorem

In AdS/CFT, marginal deformations in CFT map to continuous spaces of vacua in the bulk gravitational theory. This is true for AFS/CCFT as well.

The dual description is simple because the vacuum manifold is defined at spatial infinity and is explicitly a boundary quantity.

Sigma Models and Moduli Spaces of Vacua

In AFS, the vacuum is determined by boundary conditions (vevs) at i^0, and long-wavelength fluctuations about these vevs are described by a sigma model with target space given by the vacuum manifold $\mathcal{M}$.

$$S = \frac{1}{2} \int d^{d+2}x \, G_{IJ}(\Phi) \partial_\mu \Phi^I \partial^\mu \Phi^J \ , \qquad \Phi : \mathbb{R}^{d+1,1} \to \mathcal{M} \ .$$

$\mathcal{M}$ comes with an intrinsic geometry: the curvature of the metric on the moduli space of vacua. This will be the geometry of the space of CCFTs.

In order to perform perturbative calculations in the sigma model, we expand the fields about their vevs v^I at spatial infinity, $\Phi^I = v^I + \phi^I$, and path integrate over the normalizable fluctuations. The action becomes

$$S = \frac{1}{2} \int \partial \phi^I \partial \phi_I + \frac{1}{3} R_{IKLJ} \partial \phi^I \partial \phi^J \phi^K \phi^L + \frac{1}{6} \nabla_K R_{ILMJ} \partial \phi^I \partial \phi^J \phi^K \phi^L \phi^M + \dots$$

Infinitesimal variations along the bulk moduli space are captured by long wavelength (soft) scalars, whose S-matrix elements are universal and controlled by the moduli space geometry.

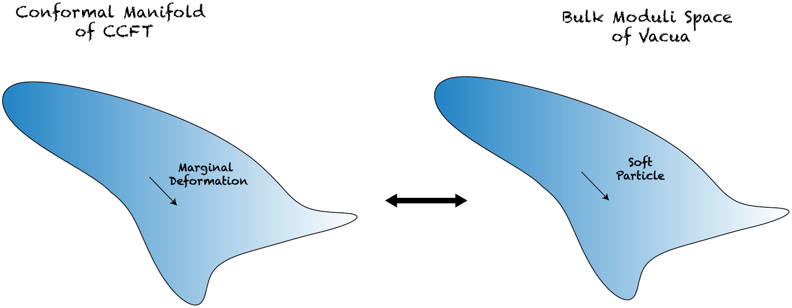

These soft scattering states define distinguished operators in CCFT, whose role is to generate marginal deformations along the conformal manifold.

Moduli Scalar Soft Theorems

The soft limit of a moduli scalar $\mathcal{O}_I(\omega, x)$ takes the form [Cheung, ... '21]

$$\lim_{\omega\to 0} \langle \mathcal{O}_I(\omega, x)\mathcal{O}_1(\omega_1, x_1)...\mathcal{O}_n(\omega_n, x_n)\rangle_v = \nabla_I \langle \mathcal{O}_1(\omega_1, x_1)...\mathcal{O}_n(\omega_n, x_n)\rangle_v$$

The subscript $\langle\cdot\rangle_v$ means that the S-matrix is computed with the boundary conditions $\langle \Phi^I\rangle_{i^0} = v^I$, and ∇ acts on the S-matrix as a function of $\mathcal{M}$.

This formula says that the zero-mode of $\phi(x)$ is the vev, so exciting the zero mode shifts the vev infinitesimally. The "leading soft moduli" operator

$$S_I(x) \equiv \oint \frac{d\omega}{2\pi i}\omega^{-1}\mathcal{O}_I(\omega, x)$$

has $\Delta = 0$. We interpret it as the shadow transform of a marginal operator $M_I(x)$ with $\Delta = d$

$$S_I = \int d^dx\, M_I(x) \ , \qquad \langle S_I(x)\mathcal{O}_1\cdots\mathcal{O}_n\rangle_v = \nabla_I \langle\mathcal{O}_1\cdots\mathcal{O}_n\rangle_v \ .$$

Turning on a coherent state of soft radiation amounts to parallel transporting the S-matrix around the space of vacua.

For an infinitesimal deformation this is

$$\langle \mathcal{O}_1 \cdots \mathcal{O}_n \rangle_{v-\lambda} = \langle \mathcal{O}_1 \cdots \mathcal{O}_n \exp\left[-\lambda^I S_I\right] \rangle_v$$
$$\equiv \langle \mathcal{O}_1 \cdots \mathcal{O}_n \exp\left[-\lambda^I \int d^d x M_I(x)\right] \rangle_v .$$

If the vacuum manifold is curved, then there is path dependence in the transport. The antisymmetric double soft limit corresponds to parallel transporting around an infinitesimal loop and computes the curvature

$$\left[\lim_{q_I \to 0}, \lim_{q_J \to 0}\right] A_{n+2}^{K_1 \cdots K_n I J} = \left[\nabla^I, \nabla^J\right] A_n^{K_1 \cdots K_n}$$

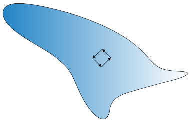

Correspondence with Conformal Perturbation Theory

$$\langle \mathcal{O}_1(x_1)\ldots\mathcal{O}_n(x_n)\rangle_{v-\lambda} = \sum_{N=0}^{\infty}\frac{(-1)^N}{N!}\langle \mathcal{O}_1(x_1)\ldots\mathcal{O}_n(x_n)(\lambda^I S_I)^N\rangle_v$$

looks a lot like conformal perturbation theory.

In fact, independently of CCFT, you can think of the shadow transform of a marginal operator as the integrated deformation that you add to the action

$$S(\lambda) = S_0 + \lambda^I S_I\ , \qquad\qquad S_I = \int d^d x\, M_I(x)\ .$$

Correlators in the deformed model are defined by conformal pert. theory

$$\langle \mathcal{O}_1(x_1)\ldots\mathcal{O}_n(x_n)\rangle_{v-\lambda} = \langle \mathcal{O}_1(x_1)\ldots\mathcal{O}_n(x_n)e^{-\lambda^I S_I}\rangle_v$$

Do the expansion and the first term tells you about the derivative.

$$\langle \mathcal{O}_1\ldots\mathcal{O}_n S_I\rangle_v = \nabla_I\,\langle \mathcal{O}_1\ldots\mathcal{O}_n\rangle_v\ .$$

The full exponential is like a coherent state of soft radiation.

If we do finite deformations in multiple directions and the conformal manifold is curved, then there is path dependence at higher orders.

Parallel transport around an infinitesimal closed loop computes the leading nontrivial holonomy in terms of the curvature of the conformal manifold $[\nabla, \nabla]$, which matches the antisymmetric double soft theorem.

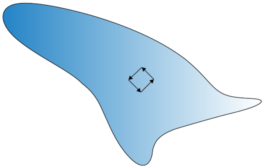

Conclusion: Bulk vacuum moduli space = boundary conformal manifold
Lesson for CCFT: Non-commuting soft limits are possible because soft operators are really integrated operators. Important for Yang-Mills CCFT.

A suggestive formula in pure Yang-Mills theory

In pure Yang-Mills, the soft gluon theorem takes the form

$$\langle S_a^I(x)\mathcal{O}_1\cdots\mathcal{O}_n\rangle = i\sum_{k=1}^{n}\partial_a\log(x-x_k)^2\ \mathsf{ad}_{T^I}\ \langle\mathcal{O}_1\cdots\mathcal{O}_n\rangle$$

and the antisymmetric double soft theorem is

$$\left[\lim_{q_I(x)\to 0}\omega,\ \lim_{q_J(y)\to 0}\omega'\right]A_{n+2}^{IJ,ab} = -2g^2\sum_{k=1}^{n}\frac{\mathcal{I}_{ac}(x-x_k)\mathcal{I}^c{}_b(y-x_k)}{(x-y)^2}\mathsf{ad}_{[T^I,T^J]}A_n$$

Meanwhile, on a compact Lie group with a bi-invariant metric, the covariant derivative of left-invariant vector fields is $\nabla_X Y = \frac{1}{2}[X,Y]$ and the curvature is $R(X,Y)Z = \frac{1}{4}\mathsf{ad}_{[X,Y]}Z$

$$\nabla_{T^I} = \frac{1}{2}\mathsf{ad}_{T^I} \qquad\qquad R(T^I,T^J) = \frac{1}{4}\mathsf{ad}_{[T^I,T^J]}$$

In other words, we seem to have formulas of the form [DK '22]

$$\langle S_I^a(x)\mathcal{O}_1\cdots\mathcal{O}_n\rangle = 2i\sum_{k=1}^{n}\partial^a\log(x-x_k)^2\ \nabla_I\ \langle\mathcal{O}_1\cdots\mathcal{O}_n\rangle$$

$$\left[\lim_{q_I\to 0}\omega,\ \lim_{q'_J\to 0}\omega'\right]A_{n+2}^{IJ,ab} = -8g^2\sum_{k=1}^{n}\frac{\mathcal{I}_{ac}(x-x_k)\mathcal{I}^c{}_b(y-x_k)}{(x-y)^2}R(T^I,T^J)A_n$$

These formulas seem very similar to the formulas from the sigma model

$$\langle S_I(x)\mathcal{O}_1\cdots\mathcal{O}_n\rangle_v = \nabla_I\ \langle\mathcal{O}_1\cdots\mathcal{O}_n\rangle_v$$

$$\left[\lim_{q_I\to 0},\ \lim_{q_J\to 0}\right]A_{n+2}^{K_1\cdots K_n IJ} = [\nabla^I,\nabla^J]A_n^{K_1\cdots K_n}$$

Why? Related to an older question about the Yang-Mills S-matrix.

Question: is there a "right definition" of multi-soft limits in Yang-Mills?

The answer has consequences for CCFT, since the soft gluon limits are related to conserved non-Abelian currents.

Clue: There is a claim that gauge theory and gravity in asymptotically flat space have an infinite-dimensional space of vacua arising from supertranslations or gauge transformations with non-compact support.

Idea: the amplitude $A_{n+2}^{K_1 \cdots K_n; IJ, ab}$ is actually a tensor on a bundle over an infinite-dimensional space of vacua $\mathcal{M}$ labeled by flat (trivial) $G-$connections $\mathcal{C}_a(x)$ on the celestial sphere.

Soft photon/gluon/graviton limits can then be interpreted as functional derivatives of the S-matrix on this infinite-dimensional space of vacua.

Answer: there is no "right definition" for multi-soft limits in non-Abelian gauge theory, just as there is no unique parallel transport between two points in a curved space.

Insertions of soft gluons in the S-matrix enact parallel transport about this space, and the antisymmetric double soft limit computes the holonomy around an infinitesimal closed curve

$$\left[\lim_{q_I(x)\to 0}\omega, \lim_{q_J(y)\to 0}\omega'\right] A_{n+2}^{K_1\cdots K_n;IJ,ab} = R\left(\frac{\delta}{\delta\widetilde{\mathcal{C}}_a^I(x)}, \frac{\delta}{\delta\widetilde{\mathcal{C}}_b^J(y)}\right) A_n^{K_1\cdots K_n} \ .$$

So the fact that the antisymmetric double-soft limit does not vanish simply means that **the vacuum manifold is curved**.

In this language, the vanishing of the antisymmetric double-soft limit in gravity is a statement about flatness of the space of supertranslation vacua, and a similar statement holds for Abelian gauge theory.

The geometric interpretation makes it clear that the space of vacua is detectable using standard perturbative calculations.

Feynman diagram calculations not only know about the Yang-Mills vacuum manifold, they can also be used to compute its curvature.

Simplest Example: Abelian Gauge Theory

Finite energy boundary conditions for abelian gauge theory in AFS allow for a leading trivial flat connection at infinity [Strominger '13; DK, Lysov, Strominger '14]

$$\langle A_a \rangle_{\mathcal{I}} = \mathcal{C}_a(x) \ , \qquad \partial_{[a}\mathcal{C}_{b]}(x) = 0 \ .$$

This means that near null infinity

$$A_a(u, r, x^a) \sim \mathcal{C}_a(x) + \frac{N_a(u,x)}{r^{\#}} + \dots$$

The subleading term is the radiation term. Finite energy requires $\mathcal{C}_a(x)$ has no time dependence, so it is usually set to zero. But just setting it to zero is itself a boundary condition.

This boundary condition is invariant under global gauge transformations, but spontaneously breaks the others with non-compact support.

The space of boundary conditions is the orbit of $\mathcal{C}_a = 0$ under the action of $\mathcal{G} =$Map$(S^d, U(1))$. The moduli space is $\mathcal{M} = \mathcal{G}/U(1)$.

S-matrix elements $\langle \cdot \rangle_{\mathcal{C}}$ carry an extra label corresponding to the boundary condition but calculations are usually performed with $\mathcal{C} = 0$.

The infrared sector of abelian gauge theory is exactly solvable. The dependence of $\langle \mathcal{O}_1 \cdots \mathcal{O}_n \rangle_{\mathcal{C}}$ on $\mathcal{C}$ is known explicitly and is very simple [Campiglia, Laddha '15; He, Mitra '20; DK, Mitra '21]. Using known formulas one can check

$$\langle S_a(x) \mathcal{O}_1 \cdots \mathcal{O}_n \rangle_{\mathcal{C}=0} = -2ic_{1,1} \frac{\delta}{\delta \widetilde{\mathcal{C}}^a(x)} \langle \mathcal{O}_1 \cdots \mathcal{O}_n \rangle_{\mathcal{C}} \Big|_{\mathcal{C}=0} \, .$$

In the sigma model, changing the boundary condition for the bulk scalar turned on a source for the marginal operator dual to the soft scalar $(\lambda^I S_I)$

The same is true in gauge theory

$$\langle \mathcal{O}_1 \cdots \mathcal{O}_n \rangle_{\mathcal{C}_a - \delta \mathcal{C}_a} = \langle \mathcal{O}_1 \cdots \mathcal{O}_n e^{\frac{-i}{2c_{1,1}} \int d^d x \, \delta \widetilde{\mathcal{C}}_a(x) S^a(x)} \rangle_{\mathcal{C}_a} \, .$$

This resembles the sigma model deformation if we think of the index I as running over polarization **and position** in the infinite-dimensional case.

To determine if this parallel transport is path dependent, we could calculate the curvature of $\mathcal{G}/G$ using the antisymmetric double-soft theorem. It vanishes, and so does the curvature:

$$\left[\lim_{q(x)\to 0}\omega, \lim_{q'(y)\to 0}\omega'\right] A_{n+2}^{K_1\cdots K_n;ab} = R\left(\frac{\delta}{\delta\widetilde{\mathcal{C}}_a(x)}, \frac{\delta}{\delta\widetilde{\mathcal{C}}_b(y)}\right) A_n^{K_1\cdots K_n} = 0\ .$$

Similar formulas hold for the space of supertranslation vacua (also flat)

$$\langle \mathcal{O}_1\cdots\mathcal{O}_n\rangle_{\mathcal{C}_{ab}-\delta\mathcal{C}_{ab}} = \langle \mathcal{O}_1\cdots\mathcal{O}_n e^{\frac{i}{16c_{1,2}}\int d^dx\,\delta\widetilde{\mathcal{C}}_{ab}(x)S^{ab}(x)}\rangle_{\mathcal{C}_{ab}}$$

where $\mathcal{C}_{ab}$ is a flat supertranslation connection

$$g_{ab}(u,r,x^a)\sim r^2\delta_{ab}+r\mathcal{C}_{ab}(x)+O(1) \qquad \partial_{[a}\mathcal{C}_{b]c}-\frac{1}{d-1}\delta_{c[a}\partial^d\mathcal{C}_{b]d}=0$$

The position-dependent boundary conditions in gauge theory and gravity $\implies$ the background sources obtained by varying the boundary conditions are position dependent.

This is a peculiar feature of celestial CFT which does not arise in garden-variety Euclidean CFT.

It is a reflection of the subtle interplay between Lorentz invariance and large gauge symmetry in asymptotically flat space.

For gravity in four dimensions, the relevant statement is familiar: "there are an infinite number of copies of $SO(3,1)$ inside the BMS group."

These different copies of $SO(3,1)$, which amount to different definitions of angular momentum, are related by the action of the supertranslations.

For a long time this was known as the "problem of angular momentum in general relativity." Gravity in asymptotically flat space simply has an infinite set of vacua corresponding to different combinations of soft gravitons.

The $SO(3,1)$ subgroup that annihilates the $\mathcal{C}_{ab} = 2(\partial_a \partial_b - \frac{1}{2}\delta_{ab}\partial^c\partial_c)f(x)$ vacuum is related to the $SO(3,1)$ subgroup that annihilates the standard vacuum $\mathcal{C}_{ab} = 0$ through conjugation by the supertranslation $f(x)$.

Classical finite energy boundary conditions in Yang-Mills theory allow the gauge field to approach a flat connection at infinity

$$\langle A_a \rangle_{\mathcal{I}} = \mathcal{C}_a(x) \ , \qquad \mathcal{C}_a(x) = U \partial_a U^{-1} \ .$$

The allowed boundary conditions are the orbit of $\mathcal{C}_a = 0$ under the action of $\mathcal{G} = \mathrm{Map}(S^d, G)$. The moduli space is $\mathcal{M} = \mathcal{G}/G$.

The leading soft gluon operator has matrix elements

$$\langle S_a^I(x) \mathcal{O}_1 \cdots \mathcal{O}_n \rangle_{\mathcal{C}=0} = i \sum_{k=1}^{n} \frac{\hat{p}_k \cdot \varepsilon_a(x)}{\hat{p}_k \cdot \hat{q}(x)} T_k^I \ \langle \mathcal{O}_1 \cdots \mathcal{O}_n \rangle_{\mathcal{C}=0} \ .$$

In contrast to abelian gauge theory and gravity, soft limits in Yang-Mills theory do not commute [Klose '15], so the parallel transport is path dependent

$$\begin{aligned} &\left[\lim_{q_I(x) \to 0} \omega, \lim_{q'_J(y) \to 0} \omega' \right] A_{n+2;ab}^{K_1 \cdots K_n; IJ} \\ &\qquad = -2g^2 \sum_{k=1}^{n} \frac{\mathcal{I}_{ac}(x - x_k) \mathcal{I}^c{}_b(y - x_k)}{(x-y)^2} f^{IJK} T_k^K A_n^{K_1 \cdots K_n} \ . \end{aligned}$$

Summary

The Celestial CFT formulation of flat space quantum gravity shares many structural similarities with AdS/CFT.

In both cases, bulk gauge fields yield boundary currents, and the bulk graviton yields a boundary stress tensor.

In both cases, the bulk moduli space of vacua maps onto the conformal manifold of the CFT dual.

The symmetry structure in flat space is more complicated, and soft photon/gluon/graviton limits can be thought of as functional derivatives in an infinite-dimensional vacuum geometry.

The dual CCFT *is more than just a d-dimensional system with G symmetry.* It also comes with a space of deformations $\mathcal{G}/G$, and that space is curved.

This is the property that is needed to reproduce non-commuting soft limits.