The black hole stability problem – an introduction and results

András Vasy (joint work with Peter Hintz, and in part with Dietrich Häfner)

January 2021, Cambridge by Zoom

These lectures are about the stability of black holes, which are certain Lorentzian manifolds solving Einstein's equation – the first lecture will mostly discuss what these black holes are and what stability means.

We adopt the convention that Lorentzian metrics on an n-dimensional manifold have signature $(1, n-1)$. For instance, the Minkowski metric on $\mathbb{R}^4 = \mathbb{R}^{1+3}$, with coordinates z_0, z_1, z_2, z_3, is

$$g = dz_0^2 - dz_1^2 - dz_2^2 - dz_3^2.$$

Here z_0 is 'time', (z_1, z_2, z_3) 'space', but there are many other timelike and spacelike coordinate functions on it! Here f timelike means $g^{-1}(df, df) > 0$, spacelike means $g^{-1}(df, df) < 0$.

In 4 dimensions Einstein's equation in vacuum is an equation for the metric tensor of the form

$$\mathrm{Ric}(g) + \Lambda g = 0,$$

where Λ is the (given) cosmological constant, and $\mathrm{Ric}(g)$ is the Ricci curvature of the metric. If there were matter present, there would be a non-trivial right hand side of the equation, given by (a modification of) the matter's stress-energy tensor.

In local coordinates, the Ricci curvature is a non-linear expression in up to second derivatives of g; thus, this is a partial differential equation. Only a few properties of Ric matter for our purposes; we point these out later.

The type of PDE that Einstein's equation is most similar to (with issues!) are (tensorial, non-linear) wave equations. The typical formulation of such a wave equation is that one specifies 'initial data' at a spacelike hypersurface, such as $z_0 = C$, C constant, in Minkowski space. For linear wave equations $\Box u = f$ on spaces like $\mathbb{R}^{1+3}$, where $\Box = d^* d = D_{z_0}^2 - D_{z_1}^2 - D_{z_2}^2 - D_{z_3}^2$, the solution u for given data exists globally and is unique.

The analogue of the question how solutions of Einstein's equation behave is: if one has a solution u_0 of $\Box u = 0$, say $u_0 = 0$ with vanishing data, we ask how the solution u changes when we slightly perturb data (to be still close to 0). For instance, does u stay close to u_0 everywhere? Does it perhaps even tend to u_0 as $z_0 \to \infty$? This is the question of stability of solutions.

Since one cannot expect that the universe is given by some explicit solution of Einstein's equation, even if it is close to it, answering this question is of great importance.

Now, Ric is diffeomorphism invariant, so if Ψ is a diffeomorphism, and g solves Einstein's equation, then so does $\Psi^* g$. This means that if there is one solution, there are many (even with same IC). In practice (duality) this means that it may not be easy to solve the equation at all!

Thus, Einstein's equation is not quite a wave equation, but it can be turned into one by imposing extra gauge conditions. Concretely, imposing that the local coordinates solve wave equations enabled Choquet-Bruhat to show local well-posedness in the 1950s. A closely related version, is DeTurck's trick – more on this later.

Turning to global questions: the first stability results were obtained for Minkowski space and de Sitter space, respectively, and are due to Christodoulou and Klainerman (1990s), later simplified by Lindblad and Rodnianski (2000s) (and extended by Bieri and Zipser), resp. Friedrich (1980s). In the late 2010s Hintz-V. gave a different proof that provided a full asymptotic expansion (polyhomogeneous, with logarithmic terms) of the metric.

The first main result, joint with Peter Hintz, in this lecture is the *global non-linear asymptotic stability* of the Kerr-de Sitter family for the initial value problem for small angular momentum a ($\Lambda > 0$). The second main result with Dietrich Häfner and Peter Hintz is the analogus *linearized* stability of the Kerr family for small a ($\Lambda = 0$).

These are a family of metrics depending on two parameters, called mass m and angular momentum a (as well as the cosmological constant Λ), whose geometric features we explore at first.

We will discuss both $\Lambda = 0$ and $\Lambda > 0$, though we focus on the latter as the results are more complete then. We remark that the observed accelerating expansion of the universe is consistent with a positive cosmological constant, which plays the role of a positive vacuum energy density; indeed, in theoretical physics $\Lambda > 0$ *is* what plays a dominant role.

Roughly, $\Lambda > 0$ is the geometer's problem, as it has all the interesting black hole features without serious analytic complications, while $\Lambda = 0$ is the analyst's problem as most of the additional difficulties are ultimately of analytic nature.

While physically $\Lambda > 0$ is small, on the scale of stability, i.e. 'time tends to ∞' behavior, there is no such thing as small Λ: on the relevant time scale the only relevant distinction is whether Λ is zero, or it is positive.

The simplest solution of Einstein's equation with $\Lambda = 0$ is Minkowski space, which is of course flat: it is the Lorentzian version of Euclidean space.

Its counterpart in $\Lambda > 0$ is de Sitter space. This is a symmetric space; it is a Lorentzian version of hyperbolic space.

A simple description is in terms of Minkowski space of one higher dimension: n-dimensional de Sitter space dS^n is the hyperboloid

$$z_0^2 - (z_1^2 + \ldots + z_n^2) = -1$$

in $\mathbb{R}^{n+1}$ with the Minkowski metric $dz_0^2 - (dz_1^2 + \ldots + dz_n^2)$. Pulling back the metric to dS^n one obtains a signature $(1, n-1)$ Lorentzian manifold which solves Einstein's equation with cosmological constant $\frac{(n-1)(n-2)}{2}$. (Cf. hyperbolic space!)

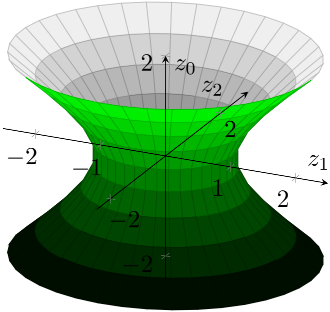

Figure: As a manifold, $\mathrm{dS}^n = \mathbb{R}_{t^*} \times \mathbb{S}^{n-1}$, where t_* is given by an explicit expression in terms of z_0 (roughly $\log z_0$ for $z_0 \gg 2$); the metric is then $dt_*^2 - (\cosh^2 t_*)\, h$, h the round metric. Here $n = 2$.

Another family of explicit solutions to Einstein's equations with $\Lambda \geq 0$ (in $1+3$ dimensions here) is the Schwarzschild, resp. Schwarzschild-de Sitter (SdS) family of metrics depending on a parameter, called mass $m > 0$:

$$g = \mu(r)\,dt^2 - \mu(r)^{-1}\,dr^2 - r^2\,h,\ \ \mu(r) = 1 - \frac{2m}{r} - \frac{\Lambda r^2}{3},$$

h the metric on the 2-sphere, $m > 0$ a parameter.

- $\Lambda = 0$ gives the Schwarzschild metric, discovered about a month after Einstein's 1915 paper; $\Lambda > 0$ is the SdS metric.
- Depending on Λ, $m = 0$ gives the Minkowski/de Sitter metric in different coordinates.
- Thus, this family describes a black hole in Minkowski/de Sitter space in a certain sense.

Recall:

$$g = \mu(r)\,dt^2 - \mu(r)^{-1}\,dr^2 - r^2\,h,\ \ \mu(r) = 1 - \frac{2m}{r} - \frac{\Lambda r^2}{3}.$$

- $\mu(r) = 0$ has two positive solutions r_+, r_- if $m, \Lambda > 0$ and $9\Lambda m^2 < 1$ (SdS); if $\Lambda = 0$, $m > 0$, the only root is $r_- = 2m$ (Schw); if $m = 0$, $\Lambda > 0$, the only root is $r_- = \sqrt{3/\Lambda}$ (dS).
- In this form the metric makes sense where $\mu > 0$: $\mathbb{R}_t \times (r_-, \infty)_r \times \mathbb{S}^2$ ($\Lambda = 0$), resp. $\mathbb{R}_t \times (r_-, r_+)_r \times \mathbb{S}^2$ ($\Lambda > 0$).
- It is spherically symmetric,
- ∂_t is a Killing vector field, i.e. translation in t preserves the metric.

It is not hard to see that $r = r_\pm$ are coordinate singularities.

A better coordinate than t is, with $c_\pm$ smooth,

$$t_* = t - F(r),\ \ F'(r) = \pm(\mu(r)^{-1} + c_\pm(r)) \text{ near } r = r_\pm.$$

In the coordinates (t_*, r, ω), the metric makes sense (as a Lorentzian metric) on

$$\mathbb{R}_{t_*} \times (0,\infty)_r \times \mathbb{S}^2_\omega,$$

thus for $r \le r_-$ and $r \ge r_+$ as well.

$r = r_-$ is called the *event horizon*, $r = r_+$ the *cosmological horizon* (if $\Lambda > 0$); the geometry of the spacetime is very similar at these.

The Schwarzschild/SdS metric fits into an even bigger family discovered by Kerr and Carter in the 1960s: the *Kerr/Kerr-de Sitter* family of metrics depending on 2 parameters, called mass m and angular momentum a; $a = 0$ gives the Schwarzschild/Schwarzschild-de Sitter metric.

Without specifying the general Kerr(-de Sitter) metric, we just mention that the underlying manifold is *still* $\mathbb{R}_{t_*} \times (0, \infty)_r \times \mathbb{S}^2$, and ∂_{t_*} is a Killing vector field, i.e. translation in t_* preserves the metric. These metrics are axisymmetric around the axis of rotation; in the case $a = 0$ they are spherically symmetric (like the de Sitter metric). There are restrictions on a to preserve the geometric features; if $\Lambda = 0$, this is $|a| < m$; if $\Lambda > 0$ they are more complicated.

To better understand the relationship between the spaces for $\Lambda > 0$, it is useful to conformally compactify $\mathrm{dS}^4 = \mathbb{R} \times \mathbb{S}^3$ by compactifying $\mathbb{R}$ to an interval. Here we concentrate on $z_0 \geq 1$; then $\tau = z_0^{-1}$, adding $\tau = 0$ as infinity, the metric is roughly like

$$\tau^{-2}(d\tau^2 - h),$$

h a metric on the sphere, and τ being essentially e^{-t_*}. ('Conformally compact'; cf. the Riemannian analogue, the Poincaré model of hyperbolic space.)

A nice feature is that null-geodesics (geodesics with null, i.e. $g(V, V) = 0$, tangent vectors V, geodesics are similar to the Riemannian setting) are simply reparameterized by such a conformal factor, i.e. they are essentially the same as those of $d\tau^2 - h$.

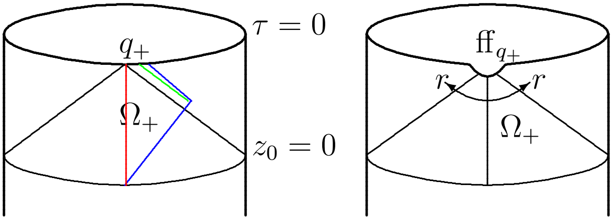

Figure: Left: the conformal compactification of de Sitter space dS^n, $n = 2$, with the backward light cone (null-geodesics) Ω_+ from q_+. The red line is the path of an observer (or particle) who tends to q_+. The blue line is that of another who leaves Ω_+... then no matter how desperately she/he/it tries, cannot get back to it. Even the green flashlight signal cannot make it back!!!
Right: the blow up of de Sitter space at q_+. This desingularizes the tip of the light cone, and the interior of the light cone inside the front face ff_{q_+} can be identified with a ball, which itself is a conformal compactification of hyperbolic space $\mathbb{H}^{n-1}$. The radial variable r for the SdS presentation of dS is that of the ball; $r = 1$ is the light cone.

The interior of the backward light cone from a point at $\tau = 0$ (future infinity) can be identified with $\mathbb{R}_{t_*} \times \mathbb{B}^3$; in the coordinates $(0, \infty)_r \times \mathbb{S}^2$ above, *singular* at $r = 0$, this is $r < 1$, often called the static (region of) de Sitter space.

Notice that dS has the feature that if a forward timelike or lightlike curve leaves the backward light cone, it can never return. Thus, the lightcone, $r = 1$, acts as a *horizon*; it is called the cosmological horizon.

Notice that nothing drastic happens at the horizons though; the manifold and the metric continue smoothly across it!

For KdS then we consider an analogue of this region, or rather that of its slight enlargement $r < 1 + \epsilon$.

Kerr-de Sitter space has two such horizons, at $r = r_\pm$, with r_+ called the cosmological horizon, r_- the black hole event horizon. They are extremely similar: once one leaves, one cannot return along timelike or lightlike curves.

There is one more relevant null-feature of KdS: there are some trapped null-geodesics in the exterior region $r \in (r_-, r_+)$, i.e. null-geodesics that do not cross either horizon. (This does not happen in dS.) This is the photonsphere in SdS, deformed in KdS.

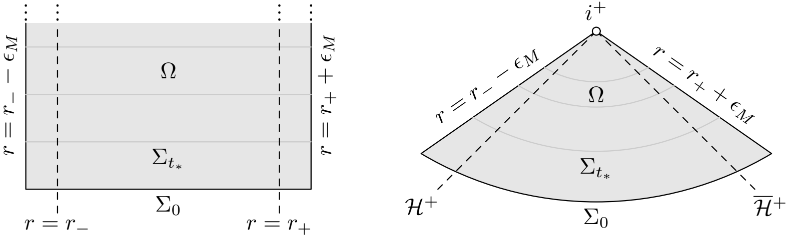

Figure: Setup for the initial value problem for perturbations of a Schwarzschild–de Sitter spacetime (M, g_{b_0}), showing the Cauchy surface Σ_0 of Ω and a few translates (level sets of the nonsingular time t_*) Σ_{t_*}; here $\epsilon_M > 0$ is small. *Left:* Product-type picture, illustrating the stationary nature of g_{b_0}. *Right:* Penrose diagram of the same setup. The event horizon is $\mathcal{H}^+ = \{r = r_-\}$, the cosmological horizon is $\overline{\mathcal{H}}^+ = \{r = r_+\}$, and the (idealized) future timelike infinity is i^+.

Kerr behaves completely analogously to KdS near the event horizon. The key difference is the presence of the Minkowski infinity. For this purpose it is useful to have a time function $\mathfrak{t}$ that is equal to t_* near the event horizon (i.e. r close to r_-), and is equal to the standard t for r large. Then the underlying manifold is still $\mathbb{R}_{\mathfrak{t}} \times (0,\infty)_r \times \mathbb{S}^2$.

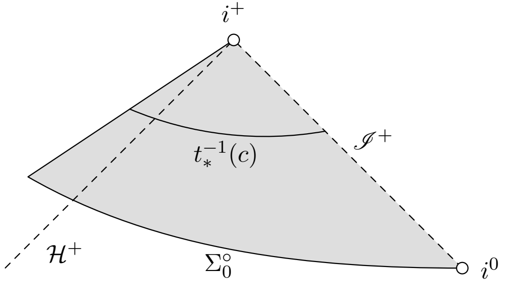

Figure: Part of the Penrose diagram of a Kerr spacetime: the event horizon $\mathcal{H}^+$, null infinity $\mathscr{I}^+$, timelike infinity i^+ and spacelike infinity i^0. We show the domain $\{\mathfrak{t} \geq 0\}$ inside of M° in gray, the Cauchy surface $\Sigma_0^\circ = \mathfrak{t}^{-1}(0)$, and a level set of t_*; $t_* = t - (r + 2m\log(r - 2m))$, r large.

We now return to the stability questions for Einstein's equation. Recall that one subtlety is the diffeomorphism invariance of the equation, causing non-uniqueness; this invariance is the only cause of non-uniqueness locally.

On the flipside, one cannot specify the initial data completely arbitrarily: they need to satisfy certain equations, called the constraint equations, implied by Einstein's equation.

In general, for a manifold M with Σ_0 a codimension 1 hypersurface, the initial data are a Riemannian metric h and a symmetric 2-cotensor k which satisfy the constraint equations, and one calls a Lorentzian metric g on M a solution of Einstein's equation with initial data (Σ_0, h, k) if the pull-back of g to Σ_0 is $-h$, and k is the second fundamental form of Σ_0 in M.

For instance, a very roughly (and weakly!) stated version of stability of Minkowski space $\mathbb{R}^4_z$, $\Sigma_0 = \{0\} \times \mathbb{R}^3$, due to Christodoulou and Klainerman, is that given initial data (h, k) close to $(g_{\mathrm{Eucl}}, 0)$ in an appropriate sense (in particular decaying), there is a global solution of Einstein's equation on $\mathbb{R}^4$, and it tends to g_{Mink} as $|z| \to \infty$.

The KdS stability is simplest phrased by considering a fixed background Schwarzschild-de Sitter metric, g_{b_0}, $b_0 = (m, \mathbf{0})$, where we use $\mathbf{a} \in \mathbb{R}^3$ as the angular momentum parameter instead of the scalar a. Let Σ_{t_*} denote the translate of Σ_0 by the ∂_{t_*} flow. Let

$$\Omega = \cup_{t_* \geq 0} \Sigma_{t_*}.$$

Theorem (Stability of the Kerr–de Sitter family for small a; informal version, Hintz-V., arXiv 2016, published 2018)

Suppose (h, k) are smooth initial data on Σ_0, satisfying the constraint equations, which are close to the data (h_{b_0}, k_{b_0}) of a Schwarzschild–de Sitter spacetime in a high regularity norm. Then there exist a solution g of Einstein's equation in Ω attaining these initial data at Σ_0, and black hole parameters b which are close to b_0, so that

$$g - g_b = \mathcal{O}(e^{-\alpha t_*})$$

for a constant $\alpha > 0$ independent of the initial data; that is, g decays exponentially fast to the Kerr–de Sitter metric g_b. Moreover, g and b are quantitatively controlled by (h, k).

The strongest $\Lambda = 0$ nonlinear black hole result is the very recent work of Klainerman and Szeftel under polarized axial symmetry assumptions.

What the theorem states is that the metric ‘settles down to’ a Kerr-de Sitter metric at an exponential rate. Note that even if we perturb a Schwarzschild-dS metric, we get a KdS limit!

This ‘settling down’ means that gravitational waves are being emitted; far away observers (hopefully us!) can see this ‘tail’. LIGO exactly aimed (successfully!) at capturing these waves, using numerical computations as a template to see what one would expect from the merger of binary black holes.

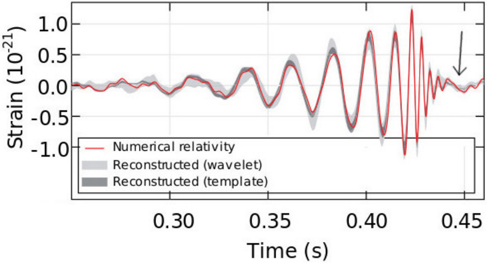

Figure: LIGO/Virgo collaboration 2016

In $\Lambda = 0$ in full generality (no additional symmetry assumptions) only linearized results are available. For the following statement recall that at the linearized level pullbacks by diffeomorphisms correspond to Lie derivatives along vector fields.

Theorem (Linearized stability of the Kerr family for small a; informal version, Häfner-Hintz-V., arXiv 2019, to appear)

Let $b = (m, a)$ be close to $b_0 = (m_0, 0)$; let $\alpha \in (0, 1)$. Suppose $\dot{h}, \dot{k} \in C^\infty(\Sigma_0^\circ; S^2 T^*\Sigma_0^\circ)$ satisfy the linearized constraint equations, and decay according to $|\dot{h}(r, \omega)| \le C r^{-1-\alpha}$, $|\dot{k}(r, \omega)| \le C r^{-2-\alpha}$, together with their derivatives along $r\partial_r$ and ∂_ω (spherical derivatives) up to order 8. Let $\dot{g}$ denote a solution of the linearized Einstein vacuum equations on Ω which attains the initial data $\dot{h}, \dot{k}$ at Σ_0°. Then there exist linearized black hole parameters $\dot{b} = (\dot{m}, \dot{a}) \in \mathbb{R} \times \mathbb{R}^3$ and a vector field V on Ω, lying in a 6-dimensional space, consisting of generators of spatial translations and Lorentz boosts, such that

$$\dot{g} = \dot{g}_b(\dot{b}) + \mathcal{L}_V g_b + \dot{g}',$$

where for bounded r the tail $\dot{g}'$ satisfies the bound $|\dot{g}'| \le C_\eta \mathfrak{t}^{-1-\alpha+\eta}$ for all $\eta > 0$.

There are a number of closely related linearized $\Lambda = 0$ black hole results: linearized Schwarzschild, plus Teukolsky in the slowly rotating case: Dafermos, Holzegel and Rodnianski (2016, 2017), as well as the linearized stability result of Andersson, Bäckdahl, Blue and Ma (2019) also in the slowly rotating case, with also a more restricted general result, under a strong asymptotic assumption.

There has been extensive research in the area, including works by (in addition to the authors already mentioned) Wald, Kay, Finster, Kamran, Smoller, Yau, Tataru, Tohaneanu, Marzuola, Metcalfe, Sterbenz, Donninger, Schlag, Soffer, Sá Barreto, Wunsch, Zworski, Wang, Bony, Dyatlov, Luk, Ionescu, Shlapentokh-Rothman...

Back to the nonlinear setting: in *DeTurck's gauge*, one fixes a background metric $\mathfrak{g}_0$, and requires that the identity map $(M, g) \to (M, \mathfrak{g}_0)$ be a wave map (solve a wave equation). This is *implemented* by working with the equation (called a gauge fixed, or reduced, Einstein's equation)

$$\mathrm{Ric}(g) + \Lambda g - \Phi(g, \mathfrak{g}_0) = 0,$$

where

$$\Phi(g, \mathfrak{g}_0) = \delta_g^* \Upsilon(g), \ \Upsilon(g) = g \mathfrak{g}_0^{-1} \delta_g G_g \mathfrak{g}_0.$$

Here δ_g^* is the symmetric gradient mapping one-forms to symmetric 2-cotensors, δ_g its adjoint (negative divergence), G_g is the trace-reversal operator $G_g r = r - \frac{1}{2}(\mathrm{tr}_g\, r)g$, and $\Upsilon(g)$ is the gauge one-form, whose vanishing is equivalent to the wave map condition.

The point is that this *is* a (quasilinear) wave-type equation, so the problems with diffeomorphism invariance have been eliminated, thus at least one has local existence and uniqueness near the initial surface Σ_0!

To see that for given initial data solving the gauged Einstein's equation actually gives a solution of the original, ungauged, problem, one constructs Cauchy data for the gauged problem for g which give rise to the required initial data and moreover solve $\Upsilon(g) = 0$ at Σ_0 (Υ is a first order differential operator, so this is determined by Cauchy data).

Solving the gauged Einstein equation with these data (which can be done locally since this is a wave equation), the constraint equations show that the normal derivative of $\Upsilon(g)$ at Σ_0 also vanishes...

...then applying $\delta_g G_g$ to the gauged Einstein's equation, in view of the second Bianchi identity,

$$\delta_g G_g \mathrm{Ric}(g) = 0,$$

true for any metric g, gives

$$\Box_g^{\mathrm{CP}} \Upsilon(g) = 0,\ \Box_g^{\mathrm{CP}} = 2\delta_g G_g \delta_g^*,$$

so by the vanishing of the Cauchy data for $\Upsilon(g)$ we see that $\Upsilon(g)$ vanishes identically.

While any choice of $\mathfrak{g}_0$ works for this local theory, for the global solvability $\mathfrak{g}_0$ makes a difference; it is natural to choose $\mathfrak{g}_0 = g_{b_0}$.

The analytic framework we use:

- non-elliptic linear global analysis with coefficients of finite Sobolev regularity,
- with a simple Nash-Moser iteration to deal with the loss of derivatives corresponding to both non-ellipticity and trapping,

gives global solvability for quasilinear wave equations like the gauged Einstein's equation provided

- certain dynamical assumptions are satisfied (only trapping is normally hyperbolic trapping, with an appropriate subprincipal symbol condition) and
- there are no exponentially growing modes (with a precise condition on non-decaying ones), i.e. non-trivial solutions of the linearized equation at g_{b_0} of the form $e^{-i\sigma t_*}$ times a function of the spatial variables r, ω only, with $\operatorname{Im} \sigma > 0$.

Unfortunately, in the DeTurck gauge, while the dynamical assumptions are satisfied, there *are* growing modes, although only a finite dimensional space of these. The key to proving the theorem (given the analytic background, discussed in the next lecture) is to overcome this issue.

Typically when solving a non-linear equation any growing modes of the linearization destroy stability; even non-decaying ones typically do.

For instance, for the ODE $u' = u^2$ with initial condition at 0, $u \equiv 0$ is a solution, which is stable on $[0, T]$ for any T, but for any positive initial condition ϕ the solution $u = \phi/(1 - t\phi)$ blows up in finite time, so there cannot be any stability on $[0, \infty)$.

Here the linearized operator is $v \mapsto v'$, which has the non-decaying mode $v \equiv 1$ (i.e. $\sigma = 0$).

The Kerr-de Sitter family automatically gives rise to non-decaying modes with $\sigma = 0$, but as these correspond to non-linear solutions, one may expect these not to be a problem with some work.

However, in the DeTurck gauge there are even growing modes, which are definitely problematic!

The reason this problem can be overcome is that the PDE is not fixed: one can modify $\Phi(g, \mathfrak{g}_0)$ as long as it gives a wave-type equation which asymptotically behaves like a Kerr-de Sitter wave equation.

In spite of this gauge freedom, we actually cannot arrange a gauge in which there are no non-decaying modes, even beyond the trivial Kerr-de Sitter family induced ones.

However, we can arrange for a partial success: we can modify Φ by changing δ_g^* by a 0th order term:

$$\widetilde{\delta}^*\omega = \delta_{g_0}^*\omega + \gamma_1\, dt_* \otimes_s \omega - \gamma_2 g_0 \operatorname{tr}_{g_0}(dt_* \otimes_s \omega),$$

$$\Phi(g, \mathfrak{g}_0) = \widetilde{\delta}^*\Upsilon(g).$$

For suitable choices of $\gamma_1, \gamma_2 \gg 0$, this preserves the dynamical requirements, and while the gauged Einstein's equation does still have growing modes, it has a new feature:

$$\widetilde{\Box}_g^{\mathrm{CP}} = 2\delta_g\, G_g \widetilde{\delta}^*,\ g = g_{b_0}$$

has no non-decaying modes! It should not be a surprise that such a change is useful: there is no reason to expect that the DeTurck gauge is well-behaved in any way except in a high differential order sense, relevant for the local theory!

We call this property SCP, or stable constraint propagation; by a general feature of our analysis, this property is preserved under perturbations of the metric around which we linearize.

Such a change to the gauge term, called 'constraint damping', has been successfully used in the numerical relativity literature by Pretorius and others, following the work of Gundlach et al, to damp numerical errors in $\Upsilon(g) = 0$; here we show rigorously why such choices work well.

SCP is useful because it means that, for $g = g_{b_0}$, any non-decaying mode h of the linearized gauge fixed Einstein equation is a solution of $D_g(\mathrm{Ric}(g) + \Lambda g)h = 0$.

Indeed this follows by applying $\delta_g G_g$ to the gauge fixed Einstein's equation, using Bianchi's second identity, giving that $\widetilde{\Box}_g^{\mathrm{CP}}(D_g \Upsilon)h$ and thus $(D_g \Upsilon)h$ vanish. Thus, properties of a gauge dependent equation are reduced to those of one independent of the gauge!

Growing modes are disastrous for non-linear equations, such as Einstein's, so we also need a statement that the above modes are actually pure gauge modes, i.e. given by linearized diffeomorphisms, so of the form $\delta_g^*\omega$ for a one-form ω, corresponding to infinitesimal diffeomorphisms. We call this, together with a precise treatment of the zero modes, UEMS, ungauged Einstein mode stability.

UEMS is actually well-established in the physics literature in a form that is close to what one needs for a precise theorem; this goes back to Regge-Wheeler, Zerilli and others; the invariant form we use is due to Ishibashi, Kodama and Seto.

Now, without the KdS-family zero modes (we call such a setting UEMS*, which holds for dS), we could easily have a framework to show global stability: namely consider

$$\Phi(g,\mathfrak{g}_0;\theta)=\widetilde{\delta}^*(\Upsilon(g)-\theta),$$

where θ is an unknown, lying in a finite dimensional space Θ of gauge terms of the form $D_{g_{b_0}}\Upsilon(\delta^*_{g_{b_0}}(\chi\omega))$, where $\chi\equiv 1$ for $t_*\gg 1$, $\chi\equiv 0$ near $t_*=0$, and such that $\delta^*_{g_{b_0}}\omega$ is a non-decaying resonance of the gauged Einstein operator.

As $D_{g_{b_0}}\Upsilon(\delta^*_{g_{b_0}}(\omega))=0$ by SCP, $D_{g_{b_0}}\Upsilon(\delta^*_{g_{b_0}}(\chi\omega))$ is compactly supported, away from Σ_0, i.e. elements of Θ are also such.

Then we could solve

$$\mathrm{Ric}(g) + \Lambda g - \Phi(g, \mathfrak{g}_0; \theta) = 0$$

for g and θ, with $g - g_{b_0}$ in a decaying function space. So crucially θ is also treated as an unknown.

This can be seen by solving the linearized equation without θ in a space which is decaying apart from finitely many non-decaying resonant modes, and then subtracting away cut off versions of these resonant terms and checking the equation they satisfy.

The full KdS version is not much harder... more details coming after the break.

Some interesting questions:

- Is the Kerr-de Sitter family stable even if a is not small? The biggest issue here is checking UEMS, which is harder due to the lack of symmetry.
- Cosmic censorship: what's going on farther in the black hole ($r < r_-$)? Recent work of Dafermos-Luk in the $\Lambda = 0$ setting, giving a conditional result, should have an unconditional analogue, and the modification of the Dafermos-Luk argument should not be too hard.
- Can we see show an expansion of the solution in terms of decaying modes? This would mathematically justify the ringdown.
- Last, but certainly not least: can we extend the non-linear stability results to the case $\Lambda = 0$?

Thank you!

Analysis and geometry in the black hole stability problem

András Vasy

January 2021, Cambridge by Zoom

The claim is that under UEMS* (i.e. ignoring the KdS-family zero modes, which e.g. would be the case for dS) we can solve

$$\mathrm{Ric}(g) + \Lambda g - \Phi(g, \mathfrak{g}_0; \theta) = 0,$$

for g and θ, with $g - g_{b_0}$ in a decaying function space. Here $\Phi(g, \mathfrak{g}_0; \theta) = \widetilde{\delta}^*(\Upsilon(g) - \theta)$.

This can be seen by solving the linearized equation without θ in a space which is decaying apart from finitely many non-decaying resonant modes, and then subtracting away cut off versions of these resonant terms and checking the equation they satisfy.

Concretely:

The linearization of the gauged Einstein equation at $(g_{b_0}, 0)$, in (g, θ) (with the linearized change of g denoted by r, that in θ is still denoted by θ since the equation is linear in θ) is

$$(D_{g_{b_0}}\mathrm{Ric} + \Lambda)r - \widetilde{\delta}^*(D_{g_{b_0}}\Upsilon)r + \widetilde{\delta}^*\theta = 0.$$

This *can* be solved in a decaying function space.

Indeed, with slight imprecision, dropping the θ term, the equation can be solved with solution $\tilde{r}$ with

$$\tilde{r} = \sum_j r_j + r'$$

r' in a decaying function space, r_j finitely many non-decaying terms, given by the resonances, which satisfy the linearized gauged Einstein equation (but of course not the initial conditions).

We have $r_j = \delta^*_{g_{b_0}} \omega_j$ (UEMS*!), so as

$$(D_{g_{b_0}} \mathrm{Ric} + \Lambda)\delta^*_{g_{b_0}} \omega = 0$$

for any one-form ω, due to g_{b_0} solving Einstein's equation and the diffeomorphism invariance of Ric, the tensor

$$r = \tilde{r} - \sum_j \delta^*_{g_{b_0}}(\chi\omega_j)$$

satisfies

$$(D_{g_{b_0}} \mathrm{Ric} + \Lambda)r - \widetilde{\delta}^*(D_{g_{b_0}} \Upsilon)r = \sum_j \widetilde{\delta}^*(D_{g_{b_0}} \Upsilon)\delta^*_{g_{b_0}}(\chi\omega_j),$$

which is exactly of the form given above!

Analytically, the point is that the operator

$$L_{b_0} = (D_{g_{b_0}} \mathrm{Ric} + \Lambda) - \widetilde{\delta}^*(D_{g_{b_0}} \Upsilon)$$

is not surjective between appropriate decaying function spaces, though the range is closed with a finite dimensional complement.

So we need to add a finite dimensional complementary subspace W so that

$$L_{b_0} r = f$$

for given f is replaced by...

$$L_{b_0} r = f + h,$$

$h \in W$ undetermined, for this equation to become solvable in these function spaces.

For us, $W = \widetilde{\delta}^* \Theta$, and it is important that this lies in the range of $\widetilde{\delta}^*$ because this assures (much like without the θ term) that the solution of the gauged Einstein equation actually gives a solution of the ungauged one!

An important point is that the analytic framework is stable under perturbations, so if one has a metric g which is close to g_{b_0} in the appropriate sense then for the gauged Einstein's equation, linearized at g,

$$L_g = (D_g\mathrm{Ric} + \Lambda) - \widetilde{\delta}^*(D_g\Upsilon),$$

$L_g r = f + h$ is also solvable with h in the same space W. In particular, this holds for Kerr-de Sitter metrics with small a (and their perturbations!).

This assures that the non-linear equation is also solvable for perturbations of the initial data, near g_{b_0}, in the same decaying function spaces, which then gives (under UEMS*) the non-linear stability result.

We interpret as saying that solving the equation finds the gauge, $\Upsilon(g) = \theta$, in which the solution is stable as well as the actual solution of Einstein's equation.

Now, UEMS* does not hold for the KdS family (exactly because it is a family) but it does hold for de Sitter space, giving a new proof of its stability.

However, it is not hard to actually deal with the full KdS family by modifying our equation by adding another term to it which corresponds to the family and somewhat enlarging the space Θ.

The result is that for an appropriate finite dimensional space $\overline{\Theta}$ the nonlinear equation

$$(\mathrm{Ric}(g)+\Lambda g)-\widetilde{\delta}^*(\Upsilon(g)-\Upsilon(g_{b_0,b})-\theta)=0$$

with prescribed initial condition is solvable for g, θ, b with $\theta \in \overline{\Theta}$, b near b_0, and $g-g_b$ exponentially decaying; here $g_{b_0,b}=(1-\chi)g_{b_0}+\chi g_b$ is the asymptotic Kerr-de Sitter metric with parameter b. Thus, *both b and θ are found along with g in the nonlinear iteration!* This proves the nonlinear stability of the KdS family with small a.

Recall that the analytic framework we use:

- non-elliptic linear global analysis with coefficients of finite Sobolev regularity,
- with a simple Nash-Moser iteration to deal with the loss of derivatives corresponding to both non-ellipticity and trapping

gives global solvability for quasilinear wave equations like the gauged Einstein's equation provided

- certain dynamical assumptions are satisfied (only trapping is normally hyperbolic trapping, with an appropriate subprincipal symbol condition) and
- there are no exponentially growing modes (with a precise condition on non-decaying ones), i.e. non-trivial solutions of the linearized equation at g_{b_0} of the form $e^{-i\sigma t_*}$ times a function of the spatial variables r, ω only, with $\operatorname{Im} \sigma > 0$.

This fits into a much broader class of linear and non-linear problems from general relativity, QFT, dynamical systems and inverse problems.

As mentioned, the non-linear aspects can be reduced to a precise understanding of underlying linear problems, via linearization and an iteration such as Picard, Newton or Nash-Moser, or 'pseudolinearization'.

In all of these problems one solves the linear (and non-linear) problems *globally* on a certain underlying 'physical space'. Here 'globally' still leaves us freedom in deciding what region of perhaps an even bigger physical space we care about, but *once we decide this*, we need to work globally in the region.

The non-linear aspects usually simply mean that the linear analysis needs to be 'done right', so I ignore these for the general discussion. To reiterate, the linear problems are solved globally and this is used to solve the non-linear problems so as well, rather than using the finite time non-linear solvability and attempting to control it uniformly as time goes to infinity.

Thus, one decides on an underlying 'physical space' (often a complete manifold for QFT, possibly a region bounded by horizons in space-time in GR, a domain for the inverse problems) M, and one considers an operator P on M.

More precisely, one needs to specify some function spaces (usually with considerable freedom) X, Y, and consider the continuous map

$$P : X \to Y.$$

In spite of the considerable freedom, it is *crucial* to be able to fix these spaces. Note also that while many choices may be equivalent, other choices may result in very different operators (cf. boundary conditions)!

- Solving equations amounts to a surjectivity statement for P
- Inverse problems/rigidity amount to an injectivity statement for P.

Since function spaces are infinite dimensional, we also need estimates: these are (semi-)Fredholm estimates.

The almost-injectivity estimate is

$$\|u\|_X \leq C(\|Pu\|_Y + \|u\|_{Z_1})$$

and the almost surjectity estimate is

$$\|v\|_{Y^*} \leq C(\|P^* v\|_{X^*} + \|u\|_{Z_2}),$$

where the inclusion maps $X \to Z_1$ and $Y^* \to Z_2$ are compact.

Compactness of these maps typically comes from the Z_j being weaker in the sense of *derivatives*, and if M is non-compact, or has a degenerate structure, then from in addition the Z_j being weaker in the sense of *decay*. (Invertibility amounts to being able to drop the relatively compact terms.)

Global analysis is standard for elliptic PDE, like Laplace's equation: one cannot solve an elliptic PDE by solving it locally!

The simplest example for Fredholm theory is elliptic operators P on compact manifolds without boundary M. (There is a similar theory for M with smooth boundary with boundary conditions.)

Recall that for $P \in \mathrm{Diff}^m(M)$, $m \in \mathbb{N}$, the *principal symbol* $\sigma_m(P)$ captures the leading terms. In local coordinates, if $P = \sum_{|\alpha| \leq m} a_\alpha(x) D_x^\alpha$, $\sigma_m(P)(x, \xi) = \sum_{|\alpha|=m} a_\alpha(x)\xi^\alpha$. Ellipticity is the statement that this does not vanish (is invertible) if $\xi \neq 0$. Then

- $X = H^s = H^s(M)$, $Y = H^{s-m}(M)$, $s \in \mathbb{R}$,
- so $X^* = H^{-s}(M)$, $Y^* = H^{-s+m}(M)$,
- $Z_1 = H^{-N}(M)$, $Z_2 = H^{-N}(M)$, N large.

The Fredholm property follows from the elliptic estimate

$$\|\phi\|_{H^r} \leq C(\|L\phi\|_{H^{r-m}} + \|\phi\|_{H^{-N}}),$$

with $L = P$, $r = s$, resp. $L = P^*$, $r = -s + m$. Note that the choice of s is irrelevant here (elliptic regularity).

The non-elliptic problems we consider are problems in which the elliptic estimate is replaced by estimates of the form

$$\|u\|_{H^s} \leq C(\|Pu\|_{H^{s-m+1}} + \|u\|_{H^{-N}}),$$

i.e. with a loss of one derivative relative to the elliptic setting, and

$$\|v\|_{H^{s'}} \leq C(\|P^* v\|_{H^{s'-m+1}} + \|v\|_{H^{-N'}}),$$

with $s' = -s + m - 1$ being the case of interest. These are often proved by *propagation* estimates using microlocal analysis.

Such estimates imply that $P : X \to Y$ is Fredholm if

$$X = \{u \in H^s : \ Pu \in H^{s-m+1}\}, \ Y = H^{s-m+1}.$$

It is easy to see that C^∞ is still dense in X.

Our non-elliptic problems are usually more complicated (except: dymanical systems!) as there *is* an infinity, which means that the Sobolev spaces will only have a compact inclusion if the error term is in a weaker *weighted* space; this is how resonances enter.

In these cases with infinity there is typically a 2-step process of obtaining estimates on weighted Sobolev spaces. Thus, one works with spaces $H^{s,\ell} = e^{-\ell t_*} H^s$, and the estimates to prove are

$$\|u\|_{H^{s,\ell}} \leq C(\|Pu\|_{H^{s-m+1,\ell}} + \|u\|_{H^{-N,-N}}). \tag{1}$$

In Step 1 one proves an estimate

$$\|u\|_{H^{s,\ell}} \leq C(\|Pu\|_{H^{s-m+1,\ell}} + \|u\|_{H^{-N,\ell}}). \tag{2}$$

Thus, the error term is lower order in the differential sense but not in the decay sense, and hence the inclusion from $H^{s,\ell}$ into $H^{-N,\ell}$ is *not* compact. Again, this step is often proved using microlocal analysis. Then, *in the simplest settings* in Step 2 one proves an estimate for a model operator at infinity

$$\|u\|_{H^{s',\ell}} \leq C\|P_0 u\|_{H^{s'-m+1,\ell}}. \tag{3}$$

Applying this to u (with $s' \geq -N$) or its localized to large t_* version, and using that $P - P_0$ has decaying coefficients, thus maps into a more decaying space, one gets (1).

For nonlinear stability problems, P is the linearization of the nonlinear problem at a current step in the iteration, and P_0 is the linearization at a metric *to which* the current iterate is asymptotic. Thus, for KdS stability, this would be a KdS space.

In order to have the P_0 estimate (3), one conjugates it by the Fourier transform to obtain a family $\hat{P}_0(\sigma)$ where σ is the (complex!) Fourier dual of $-t_*$; this is where the stationarity is used. One then automatically has a(n analytic) Fredholm theory for $\hat{P}_0(\sigma)$, corresponding to the Step 1 estimate. Thus, the question is invertibility, i.e. whether $\hat{P}_0(\sigma)$ has a non-trivial nullspace (index 0 follows from large σ considerations); this is *how* the resonances play a role. The net result is that as long as $-\ell$ is not the imaginary part of a resonance σ, one has the desired estimate (3), and so the Fredholm theory for P for all but a discrete set of weights (at least as long as the Step 1 theory allows this: trapping!).

For the stability problems, we want *forward* solutions, which is to say we want to work on distributions supported in the future of our Cauchy hypersurface.

On these spaces we get invertibility for sufficiently negative ℓ (growing space!) due to the absence of resonances in an upper half space, and indeed in more decaying spaces up to having to allow finitely many terms corresponding to the more growing resonances in the solution.

For Kerr, the fixed σ theory already has a similar aspect due to the Minkowski end. It turns out that an analogous, but more subtle, theory also works in this case.

We still need to prove the estimates with a gain in differential order (2). The tool we use is *microlocal analysis*. Microlocal analysis is local in phase space, T^*M, which is locally $\mathbb{R}^n_z \times \mathbb{R}^n_\zeta$, with ζ the momentum variables. A key point is that this is both *perturbation stable*, and (for GR) *works in a limited regularity setting*, with work on this going back to Beals and Reed in the 1980s.

More precisely, corresponding to the differential order gain we are after, we are interested in what happens as $|\zeta| \to \infty$, referred to as 'fiber infinity'. This is encoded by using dilations in the fibers (i.e. in ζ) (or a compactification), so the phase space can be considered as $T^*M \setminus o$ modulo dilations in the fibers, i.e. $S^*M = (T^*M \setminus o)/\mathbb{R}^+$.

For instance, one can say where, i.e. at which point z and which codirection ζ is a distribution in a Sobolev space or is C^∞ (wave front set) or in H^s. E.g. $\mathrm{WF}((z_1 + i0)^{-1}) = \{(0, z', \zeta_1, 0) : \zeta_1 > 0\}$.

The idea is that one proves estimates *locally in* S^*M, i.e. *microlocally*. This microlocalization is carried out by *pseudodifferential operators*, but for now all that matters is that they can in particular be associated to functions b on S^*M so that the associated $B = \mathrm{Op}(b) \in \Psi^0(M)$ is localizing to $\operatorname{supp} b$, and is non-degenerate, namely *elliptic* on $\{b > 0\}$; b is the *principal symbol* of B, extending the notion for differential operators.

For a wave operator $P = \Box_g$, the principal symbol is given by the dual metric, i.e. the inverse $G = (g^{ij})$ of g, which we think of as an 'energy function' on phase space: $p(z, \zeta) = \sum g^{ij}(z)\zeta_i\zeta_j$.

Now P is elliptic at $\alpha \in S^*M$ if the homogeneous function p is non-zero at α, i.e. at non-null covectors. Near such points we have elliptic estimates:

$$\|B_1 u\|_{H^s} \leq C(\|B_3 P u\|_{H^{s-m}} + \|u\|_{H^{-N}}),$$

$B_j \in \Psi^0$, provided $b_3 \neq 0$ on $\operatorname{supp} b_1$ and $p \neq 0$ on $\operatorname{supp} b_1$.

The basic non-elliptic estimates relate to the Hamiltonian formulation of classical mechanics, which encodes geodesics as follows.

- The lifted geodesics are integral curves of the *Hamilton vector field* given by the symplectic structure on T^*M; locally

$$H_p = \sum_{j=1}^{n} \frac{\partial p}{\partial \zeta_j}\frac{\partial}{\partial z_j} - \frac{\partial p}{\partial z_j}\frac{\partial}{\partial \zeta_j},$$

  i.e. $H_p = (\partial_\zeta p, -\partial_z p)$; these are also called *bicharacteristics*.
- Each point in T^*M gives rise to a unique lifted geodesic (the integral curve through the point).

The 'quantum' or analytical, version, then propagates H^s-estimates in the characteristic set $\mathrm{Char}(P) = \{p = 0\}$ along bicharacteristics (Hörmander).

This is an estimate of the form

$$\|B_1 u\|_{H^s} \leq C(\|B_2 u\|_{H^s} + \|B_3 P u\|_{H^{s-m+1}} + \|u\|_{H^{-N}}),$$

$B_j \in \Psi^0$, provided $\operatorname{supp} b_1 \subset \{b_3 \neq 0\}$, and all bicharacteristics from points in $\operatorname{supp} b_1 \cap \mathrm{Char}(P)$ reach $\{b_2 \neq 0\}$ of while remaining in $\{b_3 \neq 0\}$. This is usually proved via a positive commutator estimate, which is a microlocal version of an energy estimate.

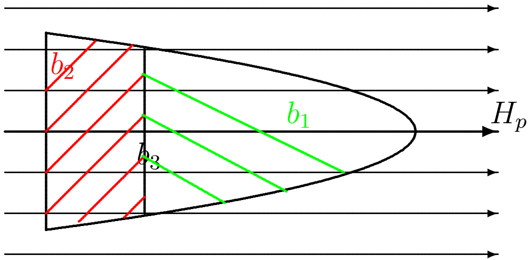

A key question is how one *starts* the propagation estimate, i.e. how one controls the $B_2 u$ term.

For wave equations, one option is Cauchy hypersurfaces (in the base manifold M); this gives the usual finite time formulation of wave propagation: a support condition makes this term trivial.

Another possibility is to have a structured bicharacteristic flow: we need that there are submanifolds L of S^*M which act as sources/sinks in the normal direction: it turns out that on high regularity spaces, one can get an estimate in which the B_2 term can be dropped. This plays a key role in scattering theory, where it was introduced by Melrose in the 1990s, though has a long history in a non-microlocal way, and around 2009 Faure and Sjöstrand also introduced this to dynamical systems.

Then one can propagate them along the flow, from say the source, and then eventually into the sink, on low regularity spaces. Often thus one needs variable order (s), or anisotropic, spaces. Crucially, these also give estimates for the adjoint on dual spaces.

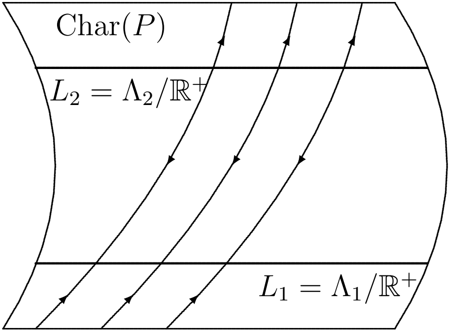

In this case there is a threshold, s_Λ, which depends on m and the imaginary part of the subprincipal symbol at Λ.

- If $s \geq s_0 > s_\Lambda$, then

$$\|B_1 u\|_{H^s} \leq C(\|B_3 P u\|_{H^{s-m+1}} + \|u\|_{H^{s_0}}),$$

$B_j \in \Psi^0$ elliptic on L, provided $\operatorname{supp} b_1 \subset \{b_3 \neq 0\}$, and all bicharacteristics from points in $\operatorname{supp} b_1 \cap \mathrm{Char}(P)$ tend to L while remaining in $\{b_3 \neq 0\}$.

- If $s < s_\Lambda$ then

$$\|B_1 u\|_{H^s} \leq C(\|B_2 u\|_{H^s} + \|B_3 P u\|_{H^{s-m+1}} + \|u\|_{H^{-N}}),$$

$B_j \in \Psi^0$ elliptic on L, provided $\operatorname{supp} b_1 \subset \{b_3 \neq 0\}$, and all bicharacteristics from points in $(\operatorname{supp} b_1 \cap \mathrm{Char}(P)) \setminus L$ reach the elliptic set $\{b_2 \neq 0\}$ of B_2 while remaining in $\{b_3 \neq 0\}$.

Replacing P by P^* changes s_Λ, and it naturally leads to estimates on the required dual spaces.

As a consequence, if there are radial sets L_1, L_2 such that all bicharacteristics in $\mathrm{Char}(P) \setminus (L_1 \cup L_2)$ escape to L_1 in one of the directions along the bicharacteristics and to L_2 in the other, one has the required Fredholm estimate provided one can arrange the Sobolev spaces so that

- at L_1 the Sobolev order is above the threshold for P,
- at L_2 the Sobolev order is above the threshold for P^*.

Typically this requires *variable order* Sobolev spaces, i.e. the order s is a function on S^*M, in which case we also need

- the Sobolev order is monotone decreasing from L_1 to L_2,

for the real principal type estimates are valid in that case.

Namely,

$$\|u\|_{H^s} \leq C(\|Pu\|_{H^{s-m+1}} + \|u\|_{H^{-N}}),$$

$$\|v\|_{H^{s'}} \leq C(\|P^*v\|_{H^{s'-m+1}} + \|v\|_{H^{-N'}}),$$

with $s' = -s + m - 1$.

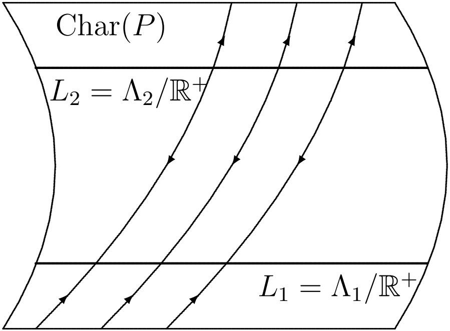
Char(P)
$L_2 = \Lambda_2/\mathbb{R}^+$
$L_1 = \Lambda_1/\mathbb{R}^+$

- A frequent place these arise is *radial sets*, i.e. points in T^*M where H_p is tangent to the fiber dilation orbits; propagation provides no information here as in S^*M the induced vector field vanishes.
- In non-degenerate settings, i.e. when H_p is non-zero, the biggest possible dimension of a radial set is that of M, in which case it is a conic Lagrangian submanifold of T^*M.
- In this case, they act as source or sink within $\mathrm{Char}(P)$; in the source case H_p flows to the zero section within Λ, in the sink case from the zero section: red shift/blue shift.
- More generally there may be a non-trivial flow within the sources/sinks; this is the case for dynamical systems as well as rotating black holes, where these are at the conormal bundles of the horizons.

The just described theory works for instance for the Fourier transformed family $\hat{P}_0(\sigma)$ of the model. For the operator P itself it is best to think of M as a compact manifold with boundary; the new boundary is an ideal boundary added to infinity with defining function $e^{-t_*} = \tau$. The microlocal analysis is then on the phase space of this compactified space, which is ${}^bS^*M$; here 'b' encodes the compactification used.

The normal sinks/sources are then replaced by *saddle points* of the Hamilton flow, at which a similar theory to the normal sources/sinks holds, with the point being the ability to propagate estimates through radial points. Again this corresponds to red/blue shifts.

The additional issue is the trapping, corresponding to bicharacteristics that even in a generalized sense (going through the saddle points) do not propagate to the initial/final 'Cauchy hypersurfaces'.

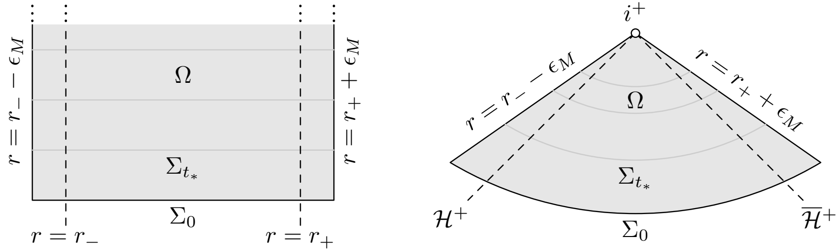

Thus, $M = [T_0, \infty)_{t_*} \times [r_- - \epsilon_M, r_+ + \epsilon_M]_r \times \mathbb{S}^2$, $\epsilon_M > 0$, where $r = r_-, r_+$ are the event and cosmological horizons.

Here $r = r_+ + \epsilon_M$ and $r = r_- - \epsilon_M$ are final Cauchy hypersurfaces, $t_* = T_0$ is the initial Cauchy hypersurface (all space-like). (These are *artificial boundaries*: we choose them. While they are important for the complete framework, they do behave as for finite time problems. But once chosen: work globally!)

Compactifying, by making $\tau = e^{-t_*}$ the boundary defining function and adding in $\tau = 0$ as ideal boundary, the flow structure is:

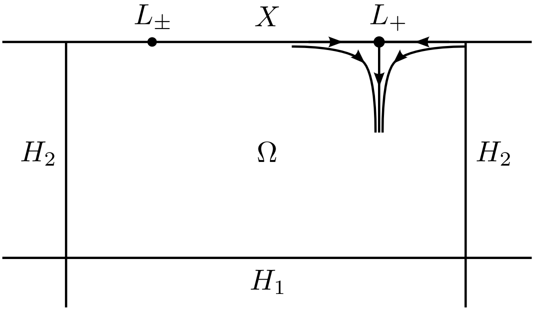

So one propagates estimates from H_1 through the radial saddle points $L_\pm$ at the horizons to H_2: in this context, the radial estimates are called red-shift (forward)/blue-shift (adjoint, backward) estimates. (There's also normally hyperbolic trapping, not shown.)

Thank you!

For the sake of perspective, let us consider an even simpler problem: the linear wave equation on a Lorentzian spacetime (M, g): $\Box_g u = f$ (f given). Typically the Cauchy problem is considered (data at an (embedded) spacelike hypersurface S). Then

- there is a unique local solution (near S), and
- if (M, g) is globally hyperbolic, i.e. each maximally extended time-like curve intersects S exactly once, or equivalently there is a global time function t, there is a unique global solution.

The Cauchy problem is equivalent to a forcing (or inhomogeneous) problem: solve $\Box_g u = f$ where f is supported in $t \geq t_0$, by finding u which is supported in $t \geq t_0$, together with its analogue where $\geq$ is replaced by $\leq$.

The solution operator $\Box_{g,R}^{-1} : f \mapsto u$ is the forward, or retarded solution operator. If one replaces $\geq$ by $\leq$, one obtains the backward, or advanced, solution operator, $\Box_{g,A}^{-1}$.

Question

What are the natural inverses of $\Box_g$? Are the inverses beyond the advanced/retarded ones?

With Gell-Redman, Haber and Wrochna we show that in reasonable (but quite general geometric) settings, there are two more natural inverses, the Feynman and anti-Feynman propagators (introduced by Feynman in the Minkowski setting!).

Idea: encode propagators via the choice of function spaces (the inverse depends on the choice!) on which $\Box_g$ is Fredholm: in terms of the source/sink picture where the regularity is high vs. low.

In a 'parametrix' sense (modulo smoothing errors) this was analyzed by Duistermaat and Hörmander: 'distinguished parametrix' for each choice of a direction in each connected component of the characteristic set. But: 'smoothing errors' are weak in non-compact settings; we set up Fredholm problems.

In fact, our discussion is not really specific for the wave equation, rather it is a general non-elliptic phenomenon.

But back in the setting of second order PDE, another place where Feynman and anti-Feynman propagators arise is ultrahyperbolic PDE such as $\sum_{j=1}^{k} D_{x_j}^2 - \sum_{j=k+1}^{n} D_{x_j}^2$, $k, n-k \geq 2$. These are in fact *very much* like the wave equation *except* for the Cauchy problem — but our approach of constructing inverses works just as well!

There has been much work in mathematical quantum field theory on Feynman propagators. The closest works in terms of general (non-algebraic) outlook have been due to Dereziński, Gérard, Häfner, Siemssen and Wrochna. Some others in the field are Bär, Brunetti, Dappiaggi, Fredenhagen, Köhler, Moretti, Pinamonti, Strohmaier...

- Let $\tilde{M} = \mathbb{R}^{n+1}$ with the Minkowski metric and $\Box$ be the wave operator; there are natural generalizations.
- Let ρ be a homogeneous degree 1 positive function, e.g. a Euclidean distance from the origin. (Analogue of τ^{-1} above.)
- The conjugate of $\rho^2\Box$ by the Mellin transform along the dilation orbits gives a family of operators P_σ, σ the Mellin dual parameter, on $\mathbb{S}^n$ (smooth transversal to the dilation orbits).
- P_σ is elliptic inside the light cone, but Lorentzian outside the light cone.
- The conormal bundle of the light cone consists of radial points.
- The characteristic set has two components, and there are four components of the radial set: a future and a past component within each component of the characteristic set.

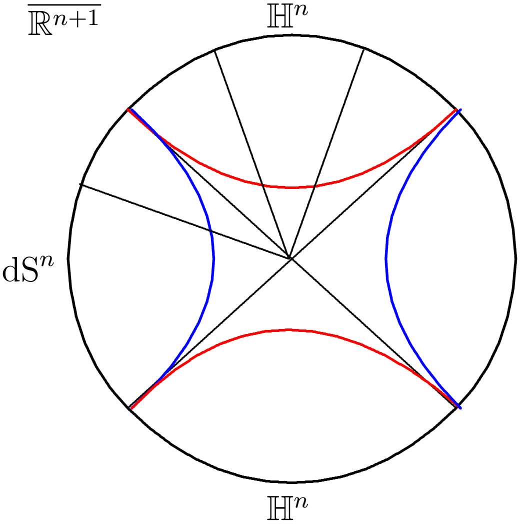
$\overline{\mathbb{R}^{n+1}}$
$\mathbb{H}^n$
dS^n
$\mathbb{H}^n$

In one component Σ_+ of the characteristic set, the bicharacteristics go from the past component of the radial set L_{+-} to the future one L_{++}; in the other component Σ_- they go from the future component of the radial set L_{-+} to the past one L_{--}.

In this case the interior of the light cone is naturally identified with hyperbolic space, while the exterior with de Sitter space.

Reasonable choices of Fredholm problems:

- Make the Sobolev spaces high regularity at the past radial sets and low at the future radial sets: this is the *forward propagator*.
- Make the Sobolev spaces low regularity at the past radial sets and high at the future radial sets: this is the *backward propagator*.
- Make the Sobolev spaces high regularity at the sources L_{+-} and L_{-+} and low regularity at the sinks, or vice versa. These are the Feynman propagators, and they propagate estimates for P_σ in the direction of the Hamilton flow in the first case, and against the Hamilton flow in the second.
- Note that the adjoint of these inverses always propagates estimates in the *opposite* direction!

Thank you again!