Two forced finite-time singularities, and the mechanism behind each

A friendly reading of Alpöge–Buckmaster (Boussinesq) and OpenAI (Navier–Stokes)

Notes for the `finite-time-blow-up` repository

September 2026

Abstract

These notes explain, through the equations, how two recent papers manufacture finite-time singularities for fluid equations with a *smooth, compactly supported external force*:

- **Boussinesq (2D, inviscid).** L. Alpöge and T. Buckmaster, *Blowup for the Boussinesq equations with smooth forcing* (`papers/boussinesq_tristan.pdf`). A multiscale cascade of oscillatory layers, each amplified by a Rayleigh–Taylor instability of the layers before it. The temperature gradient and the vorticity blow up.
- **Navier–Stokes (3D, any viscosity).** OpenAI, *Finite time blowup for Navier–Stokes* (`papers/navier-stokes_openai.pdf`). A self-similar collapsing vortex column; oscillatory pulses grown by centrifugal instability supply, through their Reynolds stress, the force that the collapse would otherwise need at the singular time. The velocity blows up with bounded kinetic energy.

Both live in the "forced" program of Córdoba and Martínez-Zoroa. Section 1 states the philosophy they share, Sections 2 and 3 walk through each mechanism, and Section 4 puts them side by side.

Contents

1 The shared philosophy: choose the solution, let the force be the residual

Take any evolution equation of the form $\partial_t U + N(U) = f$. For *any* smooth field U on $[0, T_*)$ one can *define* the force $f := \partial_t U + N(U)$, and then U solves the forced equation by construction. So the following statement is trivially true: “there is a force for which the solution blows up”. The content of these theorems lies entirely in the regularity of the force. The game is:

The residual game

Construct $U(t)$ that becomes singular as $t \uparrow T_*$, such that the residual $f = \partial_t U + N(U)$ *and all of its space-time derivatives* stay bounded, with compact support, and extend smoothly through $t = T_*$. Individual terms of the residual are allowed to diverge; their sum must not.

Both papers organize the residual into a *leading part*, which is made to vanish exactly (or to every order) by a clever ansatz, and an *error part*, which is pushed to higher and higher order by an infinite sequence of corrections. In both, the leading part is controlled by the same piece of classical fluid mechanics:

Exact waves on affine backgrounds (Craik–Criminale)

Let the “old” velocity be affine, $u_{\mathrm{old}}(x,t) = D(t)x$ with $\operatorname{tr} D = 0$. Add a plane wave $w = a(t)\,\phi(\xi(t)\cdot x)$ with polarization $a \perp \xi$ (so $\operatorname{div} w = 0$). Then

$$(w\cdot\nabla)w = a\,(a\cdot\xi)\,\phi\phi' = 0,$$

so *the wave does not advect itself*: the quadratic term drops out and the wave obeys a *linear* equation. Transporting the phase by the background,

$$\dot\xi = -D^\top \xi,$$

kills the $Dx\cdot\nabla$ term, and what remains is an ODE for the amplitude $a(t)$ in which the background strain D can *feed energy into the wave.* Explicitly, for Euler, insert $u = Dx + w$ into $\partial_t u + u\cdot\nabla u + \nabla p = 0$ and subtract the background equation. The surviving terms are $\dot a\,\phi$ (time derivative, after the phase transport), $Da\,\phi$ (the wave advecting the background, $(w\cdot\nabla)(Dx) = Dw$), and a pressure $p = q(t)\,\Phi(\xi\cdot x)$ with $\Phi' = \phi$, whose gradient $q\,\xi\,\phi$ is along ξ. The scalar q is fixed by keeping $a\cdot\xi = 0$ in time, which gives $q = -2\,\xi\cdot Da/|\xi|^2$, and hence

$$\boxed{\dot\xi = -D^\top\xi, \qquad \dot a = -Da + 2\,\frac{\xi\,(\xi\cdot Da)}{|\xi|^2} = -\Big(I - \frac{2\,\xi\otimes\xi}{|\xi|^2}\Big)Da.} \tag{1}$$

The matrix in front of Da is minus the reflection across the plane $\xi^\perp$: the pressure removes the component of $-Da$ along ξ and, because ξ itself rotates, adds it back with the opposite sign. A wave whose wavevector is being compressed by the strain ($\xi\cdot D\xi < 0$) sees its amplitude stretched. Whether a grows is therefore a question about the stability of the background, read off from the pair (1). This pair is the whole content of the "exact wave" idea; the two papers add one ingredient each. In Boussinesq the wave also carries a temperature amplitude Θ, fed by $-a \cdot G$ where $G = \nabla\theta_{\text{old}}$, and Θ feeds back on a through buoyancy; in two dimensions with $a \parallel J\xi$ this becomes the system (6) below, written there for the vorticity amplitude Ω instead of a. In Navier–Stokes the same equation acquires the damping $-\nu|\xi|^2 a$, which is what eventually kills each pulse, equation (15) below.

The two papers use this in opposite roles:

- In **Boussinesq**, the amplified waves *are* the singularity. Each layer's exponentially grown gradient becomes the "affine background" for the next, finer, layer. Infinitely many layers are squeezed into finite time.
- In **Navier–Stokes**, the singularity is a smooth *self-similar* collapse that one writes down first. The amplified waves are a tool: their averaged momentum flux $\langle w \otimes w\rangle$ (a Reynolds stress) cancels the singular part of the collapse's residual.

2 Boussinesq: a cascade of Rayleigh–Taylor layers

2.1 The equations and the theorem

The forced inviscid Boussinesq system on $\mathbb{R}^2$ is

$$\partial_t\theta + u \cdot \nabla\theta = f_\theta, \qquad \partial_t u + u \cdot \nabla u + \nabla p = \theta e_2 + f_u, \qquad \operatorname{div} u = 0, \tag{2}$$

where θ is a temperature anomaly (positive θ = lighter fluid = upward buoyancy θe_2). With $J(x_1, x_2) = (-x_2, x_1)$, $\nabla^\perp = J\nabla$, $\omega = \partial_1 u_2 - \partial_2 u_1$, and the Biot–Savart law $u = \nabla^\perp\Delta^{-1}\omega$, the vorticity form is

$$\partial_t\omega + u \cdot \nabla\omega = \partial_1\theta + \operatorname{curl} f_u. \tag{3}$$

The only source of vorticity is the *horizontal* temperature gradient $\partial_1\theta$: a horizontal density contrast is a torque.

Theorem (Alpöge–Buckmaster, Thm. 1.1, paraphrased)**.** *Take $\theta_{\text{in}}(x) = -\frac{A_0}{\lambda_0}\sin(\lambda_0 x_2)\chi_0(x)$, $u_{\text{in}} = 0$. There are odd forces $f_\theta \in C_c^\infty(\mathbb{R}^2 \times \mathbb{R})$, $f_u \in C_c^\infty(\mathbb{R}^2 \times \mathbb{R}; \mathbb{R}^2)$, supported in one fixed ball for all time, and a $T_* < \infty$ such that the unique smooth solution on $[0, T)$ for every $T < T_*$ satisfies*

$$\sup_{t<T_*}\|\theta(t)\|_\infty < \infty, \qquad \lim_{t\uparrow T_*}\|\nabla\theta(t)\|_\infty = \infty, \qquad \limsup_{t\uparrow T_*}\|\omega(t)\|_\infty = \infty.$$

The datum is a smooth, Rayleigh–Taylor *unstable* stratification: near the origin $\partial_2\theta_{\text{in}}(0) = -A_0$, so cold heavy fluid sits above warm light fluid. Push a parcel up and buoyancy pushes it further. That is the engine.

2.2 One layer: the exact wave calculation

Ignore cutoffs for a moment and suppose that near the origin the fields built so far are exactly affine:

$$u_{\text{old}}(x, t) = D(t)x \ (\operatorname{tr} D = 0), \qquad \theta_{\text{old}}(x, t) = G(t) \cdot x,$$

so $\nabla\theta_{\rm old} = G$ is the current temperature gradient and $\omega_{\rm old}$ is spatially constant. Add a temperature wave and a vorticity wave with a common phase

$$s = \lambda\,\zeta(t)\cdot x, \qquad \vartheta = \Theta(t)\sin s, \qquad \varpi = \Omega(t)\cos s, \tag{4}$$

where $\lambda \gg 1$ is a fixed frequency and $\zeta(t)$ is a time-dependent direction, so the physical wavevector is $\lambda\zeta$. The streamfunction and velocity of the vorticity wave are

$$\psi = -\frac{\Omega}{\lambda^2|\zeta|^2}\cos s, \qquad v = \nabla^\perp\psi = \frac{\Omega}{\lambda|\zeta|^2}\,J\zeta\,\sin s.$$

Because $v \parallel J\zeta \perp \zeta$ while $\nabla\vartheta, \nabla\varpi \parallel \zeta$,

$$v\cdot\nabla\vartheta = v\cdot\nabla\varpi = 0, \qquad v\cdot\nabla\omega_{\rm old} = 0.$$

Subtracting the equations of the old fields, the *unforced* wave equations are exactly linear:

$$(\partial_t + Dx\cdot\nabla)\vartheta + v\cdot G = 0, \qquad (\partial_t + Dx\cdot\nabla)\varpi = \partial_1\vartheta. \tag{5}$$

The first says: the wave's velocity v carries the old temperature gradient G, creating new temperature perturbation. The second says: the horizontal derivative of the new temperature creates new vorticity. Now require the phase to be transported, $(\partial_t + Dx\cdot\nabla)s = \lambda(\dot\zeta + D^\top\zeta)\cdot x = 0$, and match coefficients of $\sin s$ and $\cos s$:

$$\boxed{\dot\zeta = -D^\top\zeta, \qquad \dot\Theta = -\frac{J\zeta\cdot G}{\lambda|\zeta|^2}\,\Omega, \qquad \dot\Omega = \lambda\,\zeta_1\,\Theta.} \tag{6}$$

This is Lemma 3.1 of the paper (proved for a general odd periodic profile $F(s)$ in place of $\sin s$, with $F(s) = s$ near 0 so that the wave is *exactly affine* near the origin).

What each equation does

- $\dot\zeta = -D^\top\zeta$: the old velocity turns and stretches the wavevector.
- $\dot\Theta \propto \Omega$: the new velocity (from the vorticity wave) moves the old temperature.
- $\dot\Omega = \lambda\zeta_1\Theta$: the *horizontal* component of the new temperature gradient creates vorticity. It uses the laboratory component ζ_1 because gravity has a fixed direction.

2.3 Growth: the Rayleigh–Taylor rate is independent of the wavelength

Freeze $D = 0$, $G = -Ae_2$ ($A > 0$: unstable stratification), and write $\zeta = r\,e(\varphi)$ with $e(\varphi) = (\sin\varphi, \cos\varphi)$, so φ is the tilt of the wavevector away from vertical. Then $J\zeta\cdot G = -Ar\sin\varphi$ and $\zeta_1 = r\sin\varphi$, and (6) becomes

$$\frac{d}{dt}\begin{pmatrix}\Theta\\ \Omega\end{pmatrix} = \begin{pmatrix} 0 & \dfrac{A\sin\varphi}{\lambda r} \\ \lambda r\sin\varphi & 0\end{pmatrix}\begin{pmatrix}\Theta\\ \Omega\end{pmatrix}, \qquad \text{eigenvalues } \pm\sqrt{A}\,\sin\varphi. \tag{7}$$

On the growing eigenline $\Omega = (\lambda r/\sqrt{A})\Theta$ both amplitudes are multiplied by $e^{\sqrt{A}\sin\varphi\, t}$.

Why short waves are the right tool

The growth rate $\sqrt{A}\sin\varphi$ does *not* depend on λ. But the gradient does: at the origin

$$\nabla\vartheta(0,t) = \lambda\,\Theta\,\zeta, \qquad Dv(0,t) = \Omega\,Je(\varphi)\otimes e(\varphi).$$

So a wave with tiny amplitude Θ and huge frequency λ has a gradient $\lambda\Theta|\zeta|$ that can be *large* while $\|\theta\|_\infty$ stays small. That is exactly the shape of the theorem: bounded temperature, unbounded gradient. The price is the shear Ω that comes along for the ride.

The growth stage of layer q runs until $|\Theta_q| = \lambda_q^{-7/8}$, starting from a seed $\Theta_q^{\text{seed}} = -\lambda_q^{-k_q-6}$ (activated by an exponentially small force). The gradient amplitude at that point is

$$A_q := \lambda_q|\zeta_q||\Theta_q| \;\approx\; \lambda_q^{1/8} \to \infty.$$

2.4 Steering: get rid of the vorticity, keep the gradient

The wave leaves behind $\Omega \neq 0$, i.e. a shear Dv that would interfere with the next layer. The paper's trick is a *common rotation*: rotate the background gradient G and the wavevector ζ together by an angle $\alpha(t)$ (implemented as a rigid rotation $\dot{\alpha}J$ added to D). A common rotation preserves $J\zeta\cdot G$ and $|\zeta|$, hence preserves the coupling $\dot{\Theta} \propto \Omega$, but it *changes the sign of the laboratory component* ζ_1, and therefore of $\dot{\Omega} = \lambda\zeta_1\Theta$. With $\Theta < 0$, briefly making $\zeta_1 < 0$ drives Ω back to 0. The rotation is chosen (a one-parameter shooting problem, Theorem 3.2(i)) to end with

$$\Omega_q = 0 \quad \text{and} \quad \zeta_{q,1} = 0$$

while the temperature amplitude gain is retained. Then *both* derivatives in (6) vanish for as long as $\zeta_{q,1}$ is held at 0: the layer is frozen ("holding") and the next layer can grow on top of it. Each stage is therefore: **grow** $\to$ **transition** $\to$ **steer** $\to$ **hold**.

2.5 The cascade and why the time is finite

After stage q the gradient at the origin is

$$\nabla\theta(0,t) = -e_0(t) - \sum_j w_j A_j\, e_j(t), \qquad \omega(0,t) = 2\dot{\alpha}(t) + \sum_j w_j \Omega_j(t), \tag{8}$$

where $e_j = \zeta_j/|\zeta_j|$ and $w_j \in [0,1]$ are activation weights (eq. (10.1) of the paper). Define the *growth scale* $\sigma_q^2 := |\nabla\theta(0,t_q^b)| \approx \lambda_q^{1/8}$. By (7) the next layer grows at rate $\approx \sigma_q$, so the length of stage $q+1$ is $\approx$ (required log-gain)$/\sigma_q$. With the super-exponential frequency schedule

$$\lambda_q = \lambda_{q-1}^{Q_*+q}, \qquad \ell_q = \lambda_{q-1}^{-3} \text{ (envelope radius)}, \qquad |[t_q^{\text{in}}, t_q^b]| \lesssim \sigma_{q-2}^{-1}, \qquad \sigma_q \geq \sigma_{q-1}^3,$$

the stage lengths are summable, so $T_* = \lim_q t_q^b < \infty$ and infinitely many layers fit before T_*.

Why the gains don't cancel

All gradient directions $e_j(t)$ are kept inside one acute cone: $e_j \cdot e_1 \geq \cos s_* > 0$. Since every $w_j A_j \geq 0$, projecting (8) onto $-e_1$ gives

$$\|\nabla\theta(t)\|_\infty \geq c\, A_n(t) \geq c\, \lambda_n^{1/8} \quad (t \geq t_n^a),$$

which is the *full limit* $\|\nabla\theta\|_\infty \to \infty$. For vorticity, evaluate at the end of steering t_n^b: the newest $\Omega_n = 0$, and the older ones have a common sign and size $\gtrsim \sigma_{n-1}$, so $\omega(0, t_n^b) \geq c_W \sigma_{n-1} \to \infty$. Between these times a growing layer can carry vorticity of the opposite sign, which is why that conclusion is only a $\limsup$.

2.6 Why the force is smooth

Three ingredients.

1. **Nested affine regions.** The profile $F(s) = s$ near 0 and the transported envelope $g \equiv 1$ near the origin make every layer *exactly* affine there, and each finer layer is supported inside the affine region of all coarser ones ($\ell_q \ll \ell_{q-1}$). So the exact ODE (6) is not an approximation near the origin: the residual there is identically zero.
2. **High-order corrections outside.** Where the envelope varies, the wave ansatz leaves oscillatory errors. These are cancelled by explicit corrections up to order $J_q = 2k_q + 8$, with $k_q \to \infty$ as $q \to \infty$ (Sections 6–7). Increasing the correction order with the layer index is the Córdoba–Martínez-Zoroa principle that makes the force series converge with every derivative.
3. **Summability.** Each layer's contribution to the force is bounded by a fixed power of λ_q^{-1} (after all derivatives of order $\leq k_q$), and $\sum_q \lambda_q^{-c} < \infty$ trivially for the schedule above. The force is defined on the closed slab $[0, T_*]$ and extended smoothly beyond.

3 Navier–Stokes: a collapsing vortex fed by Reynolds stress

3.1 The equations and the theorem

$$\partial_t u + (u \cdot \nabla)u - \nu\Delta u + \nabla p = f, \qquad \operatorname{div} u = 0, \qquad u(\cdot, 0) = 0. \tag{9}$$

Theorem (OpenAI, Thm. 1.1, paraphrased)**.** *For every $\nu > 0$ there exist $f \in C_c^\infty(\mathbb{R}^3 \times (0,\infty); \mathbb{R}^3)$, a compact $K \subset \mathbb{R}^3$, and smooth (u, p) on $\mathbb{R}^3 \times [0,1)$ solving (9), supported in K, with*

$$\sup_{0 \leq t < 1} \|u(t)\|_{L^2} < \infty, \qquad \limsup_{t \uparrow 1} \|u(t)\|_{L^\infty} = \infty.$$

Consequently no smooth solution with the same force and datum and uniformly bounded kinetic energy exists on $\mathbb{R}^3 \times [0,\infty)$ (alternative (C) of Fefferman's Millennium statement; (D) on $\mathbb{T}^3$ by compact support).

Scaling in ν is exact, $u_\nu(x,t) = \sqrt{\nu}\, u(x/\sqrt{\nu}, t)$, so one works at $\nu = 1$. The residual is

$$\mathrm{R}(u,p) := \partial_t u + (u \cdot \nabla)u - \Delta u + \nabla p, \qquad f := \mathrm{R}(u,p). \tag{10}$$

Caveat

The Boussinesq paper is by two mathematicians who describe (Section 2 of the paper) iterating the proof with Claude and Codex and Lean-verifying an earlier blowup solution. The Navier–Stokes paper is authored as "OpenAI". These notes summarize the mechanism as written in each paper; they are not an independent verification of either proof.

3.2 The leading flow: a self-similar column that spins up

Use cylindrical coordinates (r, θ, z) and remaining time $\tau = 1-t$. The leading field is axisymmetric with all three components, and self-similar with *anisotropic* scales:

$$\ell_r \asymp \tau^{1/2}, \qquad \ell_z \asymp \tau^{1/2-h}, \qquad |u_\theta|, |u_z| \asymp \tau^{-1/2-h}, \qquad |u_r| = O(\tau^{-1/2}), \qquad 0 < h < \tfrac{1}{100}. \tag{11}$$

Precisely, with $A = \frac{1}{2} + h$, $D = \frac{1}{2} - h$, similarity variables $X = r^2/(2q)$, $\eta = z/q^D$, and a concentration scale $q \asymp \tau$,

$$\begin{aligned} u_\theta^{(0)} &= q^{-A} E(X, \eta), \qquad u_z^{(0)} = q^{-A} U(X, \eta), \qquad r u_r^{(0)} = V_0(X, \eta), \\ p^{(0)} &= q^{-2A} \Pi(X, \eta), \qquad \partial_r p^{(0)} = \frac{(u_\theta^{(0)})^2}{r} \quad \text{(centripetal balance).} \end{aligned} \tag{12}$$

The picture

Fluid spirals *inward* toward the axis and is expelled *axially* on either side of the plane $z = 0$ (incompressibility: what comes in must go out). Inward motion conserves angular momentum $r u_\theta$ per parcel, so the swirl speeds up as the core contracts: $u_\theta \sim \tau^{-1/2-h}$. Radial viscous diffusion acts on exactly the core scale, $\nu/\ell_r^2 \asymp \tau^{-1}$, which is why $\ell_r = \tau^{1/2}$ is forced: the radial Reynolds number $\mathrm{Re}_r = |u_r| \ell_r/\nu = O(1)$ stays bounded while the angular one $\mathrm{Re}_\theta = |u_\theta| \ell_r/\nu \asymp \tau^{-h} \to \infty$. The column is increasingly slender, $\ell_r/\ell_z \asymp \tau^h \to 0$, so axial diffusion is weaker by τ^{2h}: this small parameter drives the background expansion in powers q^{2nh}. The kinetic energy of the core is $\ell_r^2 \ell_z \, |u|^2 \asymp \tau^{1/2-3h} \to 0$: unbounded speed, vanishing energy.

Outside a fixed annulus in X, the flow is a pure swirl $K(r, \tau) e_\theta$ that solves the radial heat equation $-\partial_\tau K = \partial_r^2 K + r^{-1} \partial_r K - r^{-2} K$ *exactly* (the "heat exterior"); its residual is zero and it has smooth limits at $\tau = 0$ away from the axis.

3.3 The problem: the join costs a singular force

Inside the core the profiles are chosen to satisfy the leading tangential momentum balance; the exterior is exact. But joining them leaves a nonzero residual in an intermediate annulus $X_a < X < X_b$. Written as a stress divergence,

$$\mathrm{R}_\theta^{(0)} = -\Big(\partial_r + \frac{2}{r}\Big) T_{r\theta}, \qquad \mathrm{R}_z^{(0)} = -\Big(\partial_r + \frac{1}{r}\Big) T_{rz}, \qquad T \neq 0 \text{ only in the annulus,} \tag{13}$$

and by the scaling (11) this residual is of size $\tau^{-3/2-h}$: it diverges. If we simply called it the force, the theorem would be empty.

Let the fluid supply its own force

Add a divergence-free perturbation w (with pressure π). The residual changes by the *exact* identity

$$\mathrm{R}(u_B + w,\, p_B + \pi) = \mathrm{R}(u_B, p_B) + \underbrace{\partial_t w + (u_B \cdot \nabla) w + (w \cdot \nabla) u_B - \Delta w + \nabla\pi}_{L_{u_B}(w,\pi)} + \nabla\cdot (w \otimes w). \tag{14}$$

Choose w oscillatory so that (i) $L_{u_B}(w, \pi) \approx 0$: the pulses evolve on their own, seeded by an exponentially small force and grown by the background shear; and (ii) the *averaged* momentum flux cancels the singular stress:

$$\begin{pmatrix} \langle w_r w_\theta \rangle \\ \langle w_r w_z \rangle \end{pmatrix} = T + \text{higher order.}$$

Then $\nabla \cdot \langle w \otimes w \rangle$ cancels the leading $-\nabla \cdot T$ in $\mathrm{R}(u_B, p_B)$. This is a Reynolds stress: the fluid redistributes momentum by itself, and no external force is needed at leading order.

A single plane wave has $(w \cdot \nabla)w = 0$, so the stress comes from the slowly varying envelope and the cylindrical geometry, not from the carrier. And the sign works out because products are even: "outward motion carrying an azimuthal surplus and inward motion carrying a deficit both increase outward angular-momentum flux". Since T has two components, one needs two pulse families with linearly independent flux directions v_1, v_2, and the profiles are built so that $T = c_1 v_1 + c_2 v_2$ with $c_1, c_2 > 0$ (the *admissible stress cone* condition, Fig. 4 of the paper); c_i become squared amplitudes, so they must be positive.

3.4 Pulse dynamics: centrifugal growth, then viscous decay

Locally a pulse is $w \approx a \cos(\xi \cdot x + \phi)$ with $a \cdot \xi = 0$. In a normalized chart with base swirl $F = V/R$ (angular velocity) and radial shear $g = (RF_R,\, G_R)$ of the tangential velocity, the amplitude t_m of harmonic m obeys (eq. (7.5))

$$t_m' + \mathcal{K}\, t_m + m^2 d\, t_m + ikm\, n_\Phi\, \pi_m = -f_m, \qquad n_\Phi \cdot t_m = 0, \qquad d = \varepsilon k^2 |n_\Phi|^2, \tag{15}$$

where $'$ is the derivative along the pulse, $\mathcal{K}$ is the background shear/rotation matrix (with the cylindrical connection terms $2F$), d is viscous damping, and $n_\Phi = \nabla\Phi$ is the phase gradient, transported by the background exactly as $\dot{\zeta} = -D^\top \zeta$ in Boussinesq. Projected onto the plane $n_\Phi^\perp$ and written in a moving frame (Lemma 7.1), the undamped part is

$$\mathrm{diag}(\lambda, -\lambda) + E, \qquad \lambda_0^2 = -2F_0 N_\theta \big(2F_0 N_\theta + |g_0|\big) > 0, \qquad |E| \lesssim S_*^{-1}.$$

This is Rayleigh's centrifugal instability

For a pure swirl ($N = e_\theta$, no axial shear) one has $\lambda_0^2 = -2F(2F + RF_R)$, and $\frac{d}{dr}(r^2 F)^2 = 2r^3 F(2F + rF_R)$. So $\lambda_0^2 > 0$ exactly when the angular momentum $r^2 F = r u_\theta$ *decreases outward*: a parcel displaced outward keeps its angular momentum, has a swirl surplus over its new neighbours, feels excess centrifugal force, and moves further out. The profiles are built with this steep decrease in the annulus; the axial shear G_R helps near $z = 0$ where rotation alone is too weak.

Meanwhile the shear tilts the wavevector: its radial component grows linearly along the pulse, which (a) weakens the growth rate, $\lambda(v) = \lambda_0/\sqrt{1+s(v)^2}$, and (b) increases $|n_\Phi|^2$ and hence the damping d. Choosing the carrier frequency $k = \lceil \varepsilon^{-1/2} \rceil$ keeps $\varepsilon k^2 = O(1)$, so viscosity is in the leading balance. The net envelope is

$$P(v) = \exp\Big(\int_{L_s/2}^{v} \big(\lambda(w) - d_{\mathrm{ref}}(w)\big)\,\mathrm{d}w \Big), \qquad e^{-C(v-L_s/2)^2/L_s} \le P(v) \le e^{-c(v-L_s/2)^2/L_s}, \quad (16)$$

a Gaussian bump: amplification wins first, damping wins later. The Gaussian tails are what allow the pulses to be cut off in time with errors that vanish to all orders at the singularity.

3.5 Why the scales match

The pulses have amplitude and wavelength

$$A_{\mathrm{wave}} \asymp q^{-1/2-h/2}, \qquad \ell_{\mathrm{wave}} \asymp q^{1/2+h/2},$$

both smaller than the background by the factor $q^{h/2}$, so the pulses are a genuine perturbation. But the stress they produce,

$$A_{\mathrm{wave}}^2 \asymp q^{-1-h} \asymp \frac{|u_\theta^{(0)}|}{q^{1/2}}, \qquad \frac{A_{\mathrm{wave}}^2}{\ell_r} \asymp q^{-3/2-h},$$

is exactly the size of the singular residual $\mathrm{R}^{(0)} \asymp \tau^{-3/2-h}$. That is the whole point: a small oscillation has an $O(1)$-relative Reynolds stress at the singular scale because the stress is quadratic and divided by the shrinking length.

3.6 Corrections, summation, and the smooth force

After adding the leading pulses, the residual still contains: higher angular harmonics, angular means that depend on the auxiliary phase, curl and cutoff errors, and defects in five radial integrals (pressure, angular momentum, axial flux, and two conservation constraints). The paper runs a fixed four-operation cycle (Prop. 9.6): solve the inhomogeneous pulse equation for the harmonics, add signed amplitude increments whose cross-covariance with the leading pulses fixes the averaged stress, invert the fast auxiliary-time derivative for the angular means, and solve the five moment equations. Each cycle improves the residual exponent by a fixed amount,

$$\sigma_{j+1} = \sigma_j + \tfrac{1}{10}, \qquad |\partial_x^\alpha \partial_t^b \,\mathrm{R}\,| \lesssim q^{h\sigma_j - K_m} \text{ (log factors)},$$

so after j cycles the residual vanishes to order $j/10$ at the singular point. Summing all cycles with cutoffs that shrink toward $(0,1)$ gives a locally finite sum with a residual that is *flat*: every derivative is $O(q^N)$ for every N. Multiplying the vector potential and pressure by cutoffs and taking curls afterwards preserves $\operatorname{div} u = 0$ and compact support; the residual then extends by hand from $t < 1$ to a force in $C_c^\infty(\mathbb{R}^3 \times (0,\infty))$. Finally, uniqueness in the bounded-energy smooth class (Lemma 10.5) turns the constructed blowup into the non-existence statement of the theorem.

4 Side by side

	Boussinesq (Alpöge–Buckmaster)	**Navier–Stokes (OpenAI)**
Equation	2D inviscid, θ transported, $\partial_t\omega + u \cdot \nabla\omega = \partial_1\theta$	3D viscous, any $\nu > 0$
What blows up	$\|\nabla\theta\|_\infty \to \infty$ (full limit); $\limsup \|\omega\|_\infty = \infty$; $\|\theta\|_\infty$ bounded	$\limsup \|u\|_\infty = \infty$; kinetic energy bounded (core energy $\to 0$)
Singular object	Infinitely many oscillatory layers at frequencies $\lambda_q = \lambda_{q-1}^{Q_*+q}$, nested near the origin	One self-similar vortex column, $\ell_r \asymp \tau^{1/2}$, $\ell_z \asymp \tau^{1/2-h}$
Instability that drives it	Rayleigh–Taylor: rate $\sqrt{A}\sin\varphi$, independent of λ	Centrifugal (Rayleigh) + axial shear: rate λ_0 with $\lambda_0^2 = -2FN_\theta(2FN_\theta + \|g\|)$
Role of the waves	The waves *are* the singularity; each becomes the affine background of the next	The waves are a tool: their Reynolds stress $\langle w \otimes w\rangle$ cancels the singular residual of the collapse
Exact wave identity	$v \perp \zeta \Rightarrow v \cdot \nabla\vartheta = v \cdot \nabla\varpi = 0$; amplitude ODE (6)	$a \perp \xi \Rightarrow (w \cdot \nabla)w = 0$; amplitude ODE (15)
Phase transport	$\dot{\zeta} = -D^\top\zeta$	n_Φ transported by background; radial tilt grows along pulse
What stops growth	Steering: common rotation flips ζ_1, resets $\Omega = 0$; then hold	Viscous damping $d = \varepsilon k^2\|n_\Phi\|^2$ overtakes as the wavevector tilts (Gaussian envelope)
Why time is finite	Growth rate $\sigma_q \approx \lambda_q^{1/16}$ super-exponential, stage lengths $\lesssim 1/\sigma_{q-2}$ summable	Built in: singular time $t = 1$, self-similar in $\tau = 1 - t$
Small parameter	λ_q^{-1} (frequency)	$\varepsilon = q^h$ (anisotropy ℓ_r/ℓ_z), plus S_*^{-1} (slow scale)
Making the force smooth	Correction order $J_q = 2k_q + 8 \to \infty$ with the layer; force series converges in C^∞	Correction cycle with $\sigma_j = \frac{1}{5} + \frac{j}{10} \to \infty$; residual flat at $(0, 1)$; extend by hand
Force support	One fixed ball, all of $\mathbb{R}$ in time, odd parity	Compact $K \times [0, 2]$, zero near $t = 0$
Initial data	Smooth RT-unstable stratification, $u = 0$	Rest, $u = 0$; everything is switched on by the force
Uniqueness class	Locally Lipschitz, $L_t^\infty L_x^2$, bounded gradients	Smooth with uniformly bounded kinetic energy

The one-sentence version of each

Boussinesq: a short wave on an unstable stratification grows at a rate independent of its wavelength, so it can have a small amplitude and a huge gradient; rotate the background to cancel the vorticity it made, freeze it, and repeat at a much shorter wavelength, infinitely often before a finite time. **Navier–Stokes:** write down a self-similar spinning-up column that conserves angular momentum and balances radial viscosity; it needs a singular force in an annulus, so grow centrifugally unstable pulses there whose averaged momentum flux supplies that force from inside the fluid, and correct the rest to every order.

Where to look in the papers

- Boussinesq: §1.2 (the wave calculation and growth), Lemma 3.1 and eq. (3.2) (exact amplitude ODE), §3.3 (schedule: growth/transition/steering/holding), Theorem 3.2 (controlled stages), §8 (parameter choice), Prop. 10.1 (the two blowup conclusions via eq. (10.1)).
- Navier–Stokes: §2 (physical description), §3.1 (scales and similarity coordinates), §3.2 (stress and cone condition), §3.3 (the increment identity and the pulse scales), Lemma 7.1 and eq. (7.5)–(7.12) (pulse ODE, frame, envelope), §3.4 and Prop. 9.6 (correction cycle), §3.5 and §10 (localization, force extension, energy, uniqueness).