Why Galaxy Mergers Often Kill Star Formation

Published August 08, 2026 · 18 min read · https://thesleepyphysicist.com

The assumption feels almost unavoidable. Two galaxies, each containing hundreds of billions of stars, drawn together across billions of years by nothing but gravity, finally collide in a violent union of gas and dust and dark matter. Surely new stars must pour from that collision in enormous numbers. Surely the compression of all that cold molecular gas must produce the greatest furnace of stellar birth the universe can manage. The assumption is wrong. Galaxy mergers, the largest gravitational events in the observable universe, more often suppress star formation than ignite it. Understanding why requires following the gas through every stage of the encounter, from the first gravitational approach to the sterile, aging remnant left behind when everything has run its course.

The Collision That Does Not Happen

Start with the word most people use. Collision implies contact, impact, two things meeting with force and consequence that arrives in an instant. Galaxy mergers have none of that character. They unfold across timescales of 0.5 to 2 billion years,[1] and the stars inside each galaxy almost never make physical contact. The distances between individual stars are so enormous that two entire galaxy disks can pass through each other without a single stellar impact.

If the sun were the size of a grapefruit, the nearest neighboring star would be roughly 25,000 kilometers away. Two of those grapefruit-sized objects, separated by most of the Earth's circumference, constitute neighboring stars. Galaxy disks containing hundreds of billions of such objects can interpenetrate without any two of them meeting. What interacts is gravity, gas, and dark matter. The drama is written entirely in those three things, and it unfolds at a pace that has no human analog.

During the first close approach, tidal forces stretch both galaxies across their extent.

The differential gravitational pull across each system generates long arms of stars that arc outward into the surrounding dark, and a bridge of material forms between the two systems. Gas in both disks experiences shocks where material from each system meets. Some star formation is triggered: a modest pulse, a flicker above the galactic baseline. Studies of galaxy pairs caught at close separations consistently show star formation rate enhancements of roughly a factor of two to three above normal rates. That is not a starburst. It is an uptick.

After pericentre, the galaxies separate again, slowing under their shared gravitational potential before reversing and falling back inward for the second pass. The entire sequence from first approach to final coalescence of the two nuclei spans roughly one billion years, with most of that time spent in the quiet interval between encounters. Nothing about the process is sudden, and nothing about it guarantees what happens to the gas.

Why the Starburst Expectation Made Sense

The assumption that mergers reliably trigger starbursts was not naive. It came from a genuine observational discovery made when infrared astronomy became technically possible in the 1980s. Certain galaxies were radiating extraordinary amounts of thermal energy, far beyond what normal stellar populations could produce. These luminous and ultraluminous infrared galaxies, abbreviated as LIRGs and ULIRGs, were outputting energies that rivaled quasars at some wavelengths.

When astronomers examined them at optical wavelengths, the pattern was unmistakable. A large fraction of the most infrared-luminous galaxies were morphologically disturbed. They had irregular shapes, tidal tails, double nuclei, and all the structural signatures of ongoing or recent mergers. The correlation was strong enough to support a causal inference: galaxy mergers were somehow driving extreme star formation.

The physical mechanism seemed clear. Tidal torques during a merger strip angular momentum from the cold molecular gas in both galactic disks, allowing gas that was previously in stable circular orbits to fall inward on radial trajectories. Nuclear gas density rises. Dense gas collapses into stars. The logic was internally consistent, the early simulations reproduced it, and the observational sample appeared to confirm it.

What the early picture missed was selection bias. The luminous infrared galaxies that anchored the assumption were the brightest, most extreme mergers in the observable sky. They were found because they were exceptional. The quieter mergers, where the gas fell in and nothing much ignited, were harder to detect and easier to overlook in flux-limited surveys. The sample was not representative of galaxy mergers in general. It was a sample of galaxy mergers where star formation was already exceptional, which is a very different thing.

The Specific Conditions Required to Ignite a Starburst

A genuine merger-driven starburst requires a specific assembly of conditions that must arrive together. Examining what those conditions are is the clearest way to understand why they so rarely coincide.

The first requirement is gas. Both merging galaxies must carry substantial reservoirs of cold molecular gas, the raw material from which stars form. Gas fractions of 30 to 50 percent of total baryonic mass or higher are needed for the most intense events. At stellar masses above roughly $10^{10.7}$ solar masses, galaxies in the modern universe are predominantly gas-poor. Their mergers cannot produce starbursts because there is nothing available to burn.

The second requirement is orbital geometry. A prograde encounter, where each galaxy's disk rotation is aligned with the orbital angular momentum of the pair, builds gravitational resonances that efficiently strip angular momentum from the gas and drive strong inflows to the nuclear region. A retrograde encounter produces far weaker torques and correspondingly weaker inflows. The mass ratio matters too: mergers between galaxies of comparable mass, roughly 3:1 or tighter, allow both galaxies to perturb each other strongly. In highly unequal mergers, the more massive galaxy barely responds to the gravitational influence of the smaller companion.

The third requirement is timing. The most intense starburst typically occurs at or near final coalescence, when gas that has been accumulating in both nuclear regions across the full duration of the encounter is finally compressed together. Gas consumed in earlier passes is unavailable for that final event. The Antennae galaxies, the nearest major wet merger in the sky at approximately 45 million light-years, illustrate this timing problem directly. At current star formation rates, the system has a gas consumption timescale of roughly 700 million years, which is comparable to the remaining merger timescale.[2] The starburst happening now may exhaust the fuel that would otherwise power the final coalescence event.

When all conditions are met, the results are real. In the most intensely starbursting merger remnants, star formation rates reach hundreds of solar masses per year. The Milky Way currently forms roughly one to two solar masses per year. But only approximately 12 percent of major merging galaxies are genuinely starbursting at any observed moment. The majority show modest enhancement, flat rates, or in some cases active suppression.

The merger does not choose between a starburst and a black hole feeding event. It triggers both. And then the black hole begins to undo the starburst.

The Turbulence That Suppresses Its Own Trigger

One of the more counterintuitive mechanisms in this story involves the same physical process that drives what star formation does occur. Merger-induced turbulence in the interstellar medium operates simultaneously as a trigger and a suppressor, and which role dominates depends on scale and timing.

In the pre-merger galaxy, the interstellar medium has a velocity dispersion of roughly 10 km/s. Gas clouds sitting in this relatively quiescent medium can collapse under their own gravity where the local density is sufficient. During a merger, velocity dispersions climb to 30 to 40 km/s as tidal forces, shocks from disk interpenetration, and stellar feedback all inject disordered kinetic energy into the gas. This elevated turbulent pressure contributes to the effective pressure resisting gravitational collapse.

The Toomre stability criterion captures this balance. For a gas disk to fragment into star-forming clouds, the dimensionless parameter Q must fall below unity:

$$Q = \frac{c_s \kappa}{\pi G \Sigma}$$

where c_s is the effective sound speed (including turbulent pressure), κ is the epicyclic frequency related to disk rotation, G is the gravitational constant, and Σ is the gas surface density. When turbulence elevates c_s, Q rises. Regions that would have been gravitationally unstable in the pre-merger galaxy are now stable. The clouds that would have formed do not form. Gas that was capable of collapsing into stars is held in a diffuse, disordered state by the kinetic energy the merger itself injected.

At small scales, local shocks can compress individual clouds past the collapse threshold, producing the off-nuclear star formation observed in systems like the Antennae. At large scales, the elevated dispersion prevents the assembly of the giant molecular cloud complexes that the most productive star-forming regions require. The merger energizes the interstellar medium enough to trigger isolated events while simultaneously preventing the organized, widespread star formation that would constitute a true starburst across most of the system.

When the Black Hole Wakes Up

Every massive galaxy hosts a supermassive black hole at its center, typically quiescent between significant accretion events. The same nuclear inflows that accumulate gas for the starburst also deliver low-angular-momentum material to the vicinity of the black hole.

The accretion rate onto the black hole spikes during and after coalescence. Studies using the EAGLE, Illustris, and IllustrisTNG cosmological simulations consistently find that recently coalesced post-merger galaxies host elevated black hole accretion rates by factors of 2 to 5 relative to non-merging control galaxies.[3] The AGN luminosity climbs with the accretion rate, and in gas-rich major mergers it can approach quasar-level output.

The AGN accretion peak and the star formation peak occur nearly simultaneously at coalescence. Both are driven by the same underlying cause: gas falling into the nuclear region as the two systems merge. The merger does not choose between producing a starburst and activating the AGN. It produces both. And the activated AGN then begins depositing energy back into the surrounding gas.

The AGN continues accreting for approximately 100 million years after star formation has effectively ended, sustained by the small quantities of low-angular-momentum gas that continue trickling inward. During this interval, the energy output from the AGN is being deposited into the gas that remains. How efficiently that energy couples to the gas determines whether the galaxy recovers some star formation capacity after the merger or remains permanently quenched.

Thermal Feedback Versus Kinetic Feedback

AGN feedback arrives in two distinct modes, and the difference between them is the difference between temporary regulation and permanent quenching.

Thermal feedback occurs when a rapidly accreting black hole operating near its Eddington limit irradiates the surrounding gas. The gas heats, expands, and temporarily resists collapse. But hot gas cools. In the dense nuclear regions where thermal feedback is strongest, the cooling timescale can be short enough that the gas re-condenses and star formation resumes after the AGN episode dims. Thermal feedback can hold star formation down during the brightest AGN phase. It cannot remove the cold gas reservoir.

The Eddington luminosity that governs the maximum thermal feedback output scales linearly with black hole mass:

$$L_\text{Edd} = \frac{4 \pi G M_\text{BH} m_p c}{\sigma_T}$$

where M_BH is the black hole mass, m_p is the proton mass, c is the speed of light, and σ_T is the Thomson scattering cross-section. A more massive black hole can sustain higher thermal feedback luminosities, but the fundamental limitation remains: the gas cools, and what was heated can become available for star formation again.

Kinetic feedback, which operates when the black hole is accreting at a lower fraction of its Eddington limit, injects momentum rather than heat. Jets and wide-angle winds driven from the nucleus propagate outward through the disk and into the circumgalactic medium, accelerating gas in their path. Some of this gas achieves escape velocity and leaves the galaxy. What enters the halo is deposited at higher temperatures, extending the cooling time of the circumgalactic gas.

This is where permanent quenching becomes physically possible. The circumgalactic medium is the galaxy's long-term fuel reserve, the envelope of diffuse multiphase gas extending out to the virial radius from which the disk draws cold gas through cooling inflows. During coalescence, between 27 and 51 percent of the star-forming gas in a merger originates from the CGM rather than the disk.[4] When kinetic feedback heats the CGM above the cooling threshold, that replenishment stops. The disk consumes what remains without replacement arriving.

IllustrisTNG, which currently reproduces observed quiescent galaxy populations better than previous simulation generations, implements kinetic feedback as the dominant quenching mechanism. Quenching in TNG requires the central black hole to exceed approximately $10^{8.2}$ solar masses, above which the kinetic feedback is sufficient to both eject gas from the inner galaxy and prevent CGM re-cooling.

The Geometry of Quenching

A third quenching mechanism operates independently of gas removal and independently of AGN activity. It was demonstrated through cosmological zoom-in simulations and remains one of the more unexpected results in this area of research.

When a major merger transforms two disk galaxies into a single spheroid-dominated system, the change in stellar morphology itself suppresses star formation. The Toomre Q parameter for the gas disk rises when the galaxy's stellar mass distribution transitions from disk-like to spheroidal. A spheroidal stellar potential provides a smooth, extended gravitational background that suppresses the growth of local overdensities in the gas, preventing the fragmentation into giant molecular clouds that star formation requires.

The result is morphological quenching: gas continues to accrete from the halo and settle onto the disk in a cold, molecular state, yet does not form stars efficiently because the surrounding stellar structure prevents the instabilities that fragmentation needs. Simulations have shown galaxies transitioning to the red sequence while gas continues to arrive and the halo remains well-supplied.[5] The fuel is present. The geometry of the remnant has rendered it inaccessible.

The Halo Mass Threshold That Was Always There

For the most massive galaxies, a separate and independent mechanism was operating long before any merger completed. When a dark matter halo exceeds a critical mass of approximately 10^{12} solar masses, infalling intergalactic gas is shock-heated at the virial radius to the halo virial temperature:

$$T_\text{vir} \approx \frac{\mu m_p G M_\text{halo}}{2 k_B r_\text{vir}}$$

where μ is the mean molecular weight, m_p the proton mass, G the gravitational constant, k_B the Boltzmann constant, and r_vir the virial radius. For halos above the critical mass, the virial temperature drives the shock-heated gas to temperatures where the cooling time exceeds the halo dynamical time. The gas cannot efficiently cool and fall. Cold mode accretion, the mechanism by which lower-mass galaxies receive a continuous supply of fresh cold gas from the intergalactic medium, is suppressed.

This critical halo mass corresponds closely to the stellar mass above which galaxies are predominantly quiescent in the local universe. The correspondence is not accidental. The most massive galaxies were structurally predisposed to lose cold mode accretion when their halos crossed this threshold, independently of whether a merger occurred. The merger, when it happens among galaxies in halos of this mass, arrives into a system where the external fuel supply was already being cut off by a feature of the gravitational potential that predates the merger by billions of years in many cases.

The Simulation Gap No One Has Closed

Observational data from approximately 500 post-merger galaxies identified in the UNIONS survey shows that rapid quenching is 30 to 60 times more common in post-merger galaxies than in matched non-merging control populations.[6] This is the observational benchmark against which all theoretical models are measured.

IllustrisTNG finds a factor of approximately two. EAGLE finds eleven. Illustris, the predecessor to TNG, finds no statistically significant excess at all. Three state-of-the-art simulations produce three wildly different answers, none of which matches the observational value.

The leading candidate explanation involves resolution. Current cosmological simulations cannot resolve the sub-parsec scales at which AGN outflows are physically launched. Instead, they use subgrid prescriptions: parameterized models of the net effect of unresolved physics, calibrated to reproduce global galaxy properties. If the real coupling between AGN outflows and surrounding gas is more efficient than these prescriptions assume, the simulations will consistently underestimate the quenching rate. Another candidate involves the clumpy structure of the interstellar medium. Dense, irregular gas interacts with outflowing material very differently than the smooth medium most simulations represent, potentially entraining and ejecting cold gas far more efficiently in reality.

The gap between a factor of two and a factor of 30 to 60 is not an uncertainty to be dismissed. It is a structural disagreement telling us that something physically real is missing from the best models currently available.

The Dragon Effect: External Quenching

An additional mechanism, identified observationally in 2025, suggests that the accounting of quenching processes in mergers has been systematically incomplete in a specific way.

In a system observed at redshift 1.61, a host galaxy containing an X-ray AGN showed signatures of an ongoing starburst while a companion galaxy at a projected separation of approximately 8 kiloparsecs showed spectral signatures of recent rapid quenching with no active internal AGN to explain the shutdown. The proposed mechanism, named the Dragon Effect by the authors, is that the AGN-driven outflow from the host galaxy propagated outward across the 8-kiloparsec gap and deposited sufficient energy into the companion galaxy's cold gas reservoir to suppress or terminate star formation before the two systems had coalesced.

If this mechanism operates in a meaningful fraction of close galaxy pairs, particularly at cosmic noon around redshift 2 when both merger rates and AGN activity were near their peak, then the standard framework of treating each galaxy's quenching as driven by its own internal processes is incomplete. One galaxy's black hole can quench its neighbor before the merger is finished.

What We Actually Know

Galaxy mergers are not starburst engines. They are occasions: gravitational events that concentrate multiple quenching mechanisms into a brief interval and make the outcome of a long process abrupt enough to leave a clear observational signature.

The cold gas that once filled both progenitor disks is gone from the remnant. Not destroyed. Locked. In aging stars that formed during the starburst. In a hot circumgalactic halo that kinetic AGN feedback has raised above the cooling threshold. In a disk geometry too spheroidally stabilized to fragment. In a halo virial temperature that was always going to prevent cold mode accretion from continuing indefinitely.

The galaxies that quench rapidly after mergers were, in most cases, already predisposed to quench before the encounter began. They arrived with low gas fractions, massive central black holes, and halos approaching or past the virial shock threshold. The merger met them partway through their transition and compressed what might have been a slow decline into a transition fast enough to be clearly visible in the galaxy population statistics.

What the observations show is a factor of 30 to 60 enhancement in quenching rates among post-merger galaxies relative to controls. What the best simulations produce is a factor of 2 to 11. The distance between those numbers is where the current frontier of galaxy formation physics actually lives. Closing it will require either higher-resolution simulations that can resolve the real coupling between AGN feedback and multiphase gas, or new physics prescriptions that better capture what the parameterized models are missing.

The universe's most massive galaxies are also among its quietest. They grew large through mergers and accretion during an earlier cosmic epoch when cold gas was abundant and star formation was at its peak. That epoch ended for them, accelerated by the very processes their growth made possible. They now drift through a universe that is itself slowly and structurally running low on the cold gas that makes new stars possible, each merger among them adding stellar mass in silence, without a single new star to show for the collision.

Notes and Sources

1. Total merger duration estimates from Lotz et al. (2008), MNRAS 391, and synthesis across multiple simulation suites. The range 0.5 to 2 billion years reflects variation in mass ratio and orbital parameters.
2. Gas consumption timescale in the Antennae estimated in Zhang et al. (2001), arXiv:astro-ph/0105174. Current star formation rate approximately 20 solar masses per year from Spitzer IRAC observations.
3. SMBH accretion rate enhancement in post-mergers: Quai et al. (2023), MNRAS 519, 2119. Comparison across EAGLE, Illustris-1, and TNG100-1.
4. CGM contribution to merger-era star formation: Sparre et al. (2022), MNRAS 509, 2720. Result derived from cosmological zoom-in simulations tracking gas origin tags through the merger sequence.
5. Morphological quenching mechanism: Martig et al. (2009), ApJ 707, 250. Demonstrated using cosmological zoom-in simulations with an efficient technique for following individual galaxy evolution in full cosmological context.
6. Observational post-merger quenching excess: Ellison et al. (2022), arXiv:2209.07613. Sample of approximately 500 post-mergers from the UNIONS survey, compared against a matched control sample of non-interacting galaxies using multiple post-starburst classification methods.