Novel Magnetic Confinement Topologies for Fusion Energy

A Research History and Status

A record of what was tried, what failed, and why. The result of this program to date is largely a record of disproof – which is itself the finding. Diagnosis is reported here in full, including the configurations that did not work and the reasons they could not.

Premise

Tokamaks dominate fusion research because they confine a plasma stably. They also inherit four constraints: the large physical scale required for adequate confinement time, complex superconducting magnet systems, pulsed operation requiring periodic current drive, and a geometry that obstructs maintenance access. The question this program asked was never *how to optimize the tokamak* but *which of these constraints are fundamental physics, and which are artifacts of the historical development path* – and whether a simpler, steady-state, more maintainable topology could confine a fusion-relevant plasma at all.

The method has been constraint-first and simulation-validated. Rather than starting from an existing reactor and tuning parameters, each configuration began from the physics constraints – pressure balance, MHD stability, particle and energy confinement – and asked what field geometry satisfies them. A single test-particle and collisional simulation engine (a Boris pusher with binary-collision and Bosch–Hale fusion operators, GPU-accelerated) was carried forward across every configuration, so results are directly comparable between versions.

The honest summary of the work to date: **the open magnetic mirror does not close, in any variant tried, for reasons that are now well understood and largely fundamental.** The materials available cannot build the one mirror geometry that would be stable, and even if they could, an open mirror is capped below the energy gain a reactor requires by the loss cone – a topological property no field strength or conductor removes. The active line abandons the mirror entirely for a closed-field-line configuration, and its viability rests on a single stability question that is not yet answered.

The configuration sequence

The work proceeds by numbered configurations. Each was retained only as long as it survived the physics.

V1–V3 – the dynamic mirror

A three-coil "magnetic bottle" with time-varying currents and a central null ring, on the premise that field modulation would enhance scattering and confinement. It confined roughly 0.9% of a test population. Critically, 78% of the losses were *radial* – driven by grad-B and curvature drift in a weak (~0.1 T) midplane field at a mirror ratio of only ~4.4 – not axial loss through the mirror throats. Every dynamic variant made this worse: rocking the mirror axis (V3) dropped survivors to single digits out of 100,000.

The lesson, recorded early and confirmed repeatedly afterward: a near-marginal mirror has no field margin to modulate. Any time-variation periodically drops the field below the confinement threshold, and any tilt converts the otherwise-azimuthal grad-B drift into radial loss. The "dynamic" premise was the original mistake.

V4–V6 – the mirror taken to its limit

Rotating compression layers, rotating-field mirrors, and Ioffe bars were each tried as fixes; each failed. The Ioffe result is instructive: minimum-B stabilization had no measurable effect on confinement, because a single-particle model contains no collective interchange mode for a minimum-B field to stabilize. The stabilizer had nothing to stabilize and merely perturbed orbits.

The redesign that finally worked geometrically was **V5**, a nine-coil graduated end-plug mirror – a central solenoid at ~5 T with four-coil plugs reaching a ~17 T throat, mirror ratio ~21, loss-cone half-angle ~12.6°. It eliminated radial loss entirely and confined 81% of ions. But the failures hidden inside V5 are the important part:

- The true field minimum is *off-axis* (~0.83 T at the plug shoulder), a weak point and a bad-curvature region.
- The headline confinement time was an extrapolation. Adding a Takizuka–Abe Coulomb collision operator collapsed it: measured energy confinement fell to ~15 ms, and the device did not break even at any density tested.
- The 3.5 MeV alphas required for self-heating are not adiabatic in this field – their gyroradius (~300 mm) exceeds the bore, and only ~12% are confined. Self-heating is structurally unavailable.

V6 committed the coils to superconductor to make the power balance meaningful and confirmed the verdict: self-sustainment would demand reactor-grade density ($\sim 10^{21}$–10^{22} m^{-3}) and fields no available conductor survives under neutron load.

> *Note on numbers: an early V5 collisional analysis reported a more favorable breakeven density and confinement scaling. Those figures were computed before a cross-section units error was found and corrected (the fusion cross-section had been ~1000× too high). They are invalid and are not used here. The honest anchor is the corrected V7 deuterium sweep below.*

V7 – high-field axisymmetric mirror

A thick-REBCO nine-coil mirror – B_min ~1.7 T, B_max ~21 T, mirror ratio ~12. The most thoroughly documented configuration, and the source of the program's central finding.

- **In deuterium it never breaks even.** Energy gain plateaus at Q ≈ 0.15, scaling as Q ∝ n^0.63, with extrapolated breakeven at $\sim 1.6\times10^{25}$ m^{-3} – the classic Pastukhov end-loss limit of open mirrors, and physically infeasible.
- **The apparent deuterium–tritium Q ≈ 21 is an artifact.** That operating point sits at β ≈ 7,500 – four orders of magnitude beyond MHD equilibrium. The simulator holds the field fixed and cannot feel plasma pressure, so it reports a gain for a plasma that would expel the field before the run began. No conductor rescues this: pulling β down to a stable ~0.05 by raising field alone would require fields of order thousands of tesla.
- **Stabilizing it destroys it.** The bare axisymmetric mirror is flute/interchange unstable above β ≈ 0.01–0.05 because of bad midplane curvature. Adding Ioffe bars for minimum-B stabilization raised the loss rate roughly 100-fold, because the stabilizing field varies on the gyroradius scale and breaks adiabaticity – particles scatter into the loss cone.

This produced the dilemma that defines the mirror line: **there is no operating point that is simultaneously stable and confining.** Stable implies not confining; confining implies not stable. With realistic refueling and ICRH heating, the honest steady-state gain is Q ≈ 0.001. V7 is closed.

Two bookkeeping errors were found and corrected during V7 – the 1000× cross-section units error noted above, and a ~50× transient-versus-steady-state population error that had inflated an early Q estimate. The corrected, far less favorable numbers are the ones this program stands behind. The self-correction is the most credible quality signal in the record; any uncorrected high-Q figure should be distrusted by default.

V8 — minimum-B (Yin-Yang), and the materials wall

The attempt to escape V7's stability/confinement dilemma with a true minimum-B geometry — stable by construction rather than by added Ioffe fields. Two coil topologies were studied; one reproduced bad curvature (the V7 failure mode again), the other produced near-zero field on axis (infinite gyroradius, adiabaticity lost).

The viable topology was disqualified by the conductor — but not, as first assumed, by insufficient current density. REBCO carries enough current at the required ~20 T and 10–15 K *at the tape level*. It fails **mechanically**. A minimum-B coil is a three-dimensional baseball-seam saddle, and forming brittle coated tape along that twisting path drives the combined bend-and-torsion strain to three-to-five times REBCO's ~0.3% multi-axial limit at compact scale. The strain only relaxes if the coil radius grows past ~1.2 m — and at that scale the amp-turns needed to hold the field well drive the Lorentz forces several times beyond an engineering-tractable regime (of order 9× the MFTF-B class). The geometry cannot be shrunk to raise the field without cracking the conductor, and cannot be enlarged to relieve strain without the forces and the required mega-amp cable format becoming the new wall.

This is the single point in the program where a materials advance would change the verdict — but the advance required is **not a denser superconductor; it is a differently-shaped one.** The bar is a conductor delivering ≳150 A/mm² at the winding-pack level in a 3D saddle geometry with bend-torsion strain below 0.3%, at ~20 T and 10–15 K. REBCO meets the current density but cannot be formed into that saddle without exceeding its strain limit several-fold, and no cable format exists at the 1–7 MA strand count required. The credible candidate is an isotropic round-wire conductor that tolerates the saddle strain a flat tape cannot — BSCCO-2212 round wire, demonstrated in laboratory high-field magnets but not yet produced as the mega-amp-class cable this geometry needs. Higher current density, colder operation, and advanced pinning all improve margin and none reopen V8, because its wall is the *shape* of the conductor, not the current it carries. V8 was closed at the engineering gate before any confinement simulation ran.

V9 — field-reversed configuration (the active line)

With the mirror family exhausted, the program pivoted to a higher-β, inherently *closed-field-line* topology. This is the decisive move: a field-reversed configuration has no loss cone. The topological leak that caps every open mirror does not exist in a closed-field-line plasma. (See *Why the mirror cannot close*, below, for why this — not stability and not materials — is the real reason to leave the mirror.)

A full reactor design exists on paper: rotating-magnetic-field current drive, ICRH, neutral-beam injection held in reserve as stability insurance, and direct energy conversion of the axial exhaust. The design claims an engineering gain of order 1.6 net-positive. It is explicitly **gated, not validated.** At the design's kinetic parameter, the classical n = 1 tilt mode grows roughly four orders of magnitude faster than the target confinement time, and the fast-ion stabilization that might suppress it sits marginally at the known threshold. Whether the configuration is stable is the single question on which the entire program now turns.

V10 — counter-rotating collider (closed)

A side branch: two counter-rotating ion beams set to overlap inside a magnetic mirror, the mirror's role being not to confine a thermal plasma but to recapture whatever left the beams — deflected particles, scattering products, and fusion products — and return them to the circulating tracks. The intent was that ions passing through the overlap region without colliding or fusing would be retained pass over pass.

It failed at the recapture step. The beams deflected just enough that off-track particles struck the collider walls instead of re-entering their tracks. Rather than maintaining the non-fused population for another pass, the colliders emptied — the loss channel was wall impact, not burn-up — and beam density collapsed. The lesson rhymes with the rest of the program: a magnetic mirror asked to perform precise orbit recapture, rather than coarse statistical confinement, runs into the same drift-and-deflection physics that defeated the thermal mirrors. The field bends particles, but not onto the trajectory you need. V10 is closed.

Why the mirror cannot close

Across V1–V8 and V10, every available lever was pulled — geometry, field strength, minimum-B stabilization, and exhaust-energy recovery — and the open mirror still does not reach the energy gain a reactor requires. The reasons separate cleanly, and only one of them is fixable:

1. **The loss cone is topological.** An open mirror leaks ions through a fixed solid angle in velocity space regardless of how stable, how strong, or how well-built it is. The loss-cone fraction is set by the mirror ratio and is refilled by collisions. This is the binding constraint, and it is independent of the other three.

2. **A simple mirror is capped near Q ≈ 1.** This is a textbook result (Post; Pastukhov): in a single-cell mirror the ion energy lost out the loss cone per confinement time is comparable to the fusion energy produced, so the gain cannot rise much above unity even ignoring all engineering. An independent estimate for a *stable, high-β minimum-B* mirror — the best case this program could

build if the materials existed — lands at a deuterium-tritium physics gain of Q ≈ 0.2–0.5, reaching ~1 only at fields of 15–20 T and β ≈ 1, which is precisely the configuration that died at the conductor wall. In deuterium the figure is ~0.02–0.05. Raising β (which minimum-B permits) raises the density ceiling and helps the gain, but β and the loss cone are orthogonal axes; β cannot defeat the loss cone.

3. **Minimum-B stability is not buildable with present conductor.** The geometry that removes the flute instability (V8) cannot be wound from coated tape without exceeding its mechanical strain limit several-fold. This is the one fixable constraint — it waits on a round-wire superconductor that does not yet exist at the required scale.
4. **Exhaust-energy recovery does not change the physics.** Direct energy conversion captures the kinetic energy of escaping ions at 70–85% efficiency. It improves the *engineering* energy balance — it can swing a device with a physics gain of Q ≈ 1–5 from net-negative to net-positive — but it does nothing to the *confinement*. The particles still leave; the energy confinement time is unchanged. Because a stable minimum-B mirror tops out at a physics gain of ~0.2–0.5, below the ~1 floor recovery needs to pay off, direct conversion never had a viable mirror to attach to. Recovering exhaust energy cannot manufacture the fusion power the loss cone prevented.

The historical escape from the loss cone was the *tandem mirror* — adding electrostatic end-plugs with ambipolar potentials to confine the loss cone (MFTF-B, TMX-U). That is a different architecture, not a conductor upgrade and not energy recovery. This program took neither the tandem route nor a further mirror iteration. It changed topology.

Current status

The active task is not reactor design. It is answering V9's stability question before any further design proceeds. A dedicated three-dimensional **hybrid simulator** was built for exactly this — fluid MHD for the bulk plasma coupled to a particle-in-cell kinetic treatment of the fast ions — to test whether the field-reversed configuration survives the tilt mode at the design point, or whether it requires neutral-beam stabilization.

The simulator is in good engineering shape: GPU-accelerated, broadly unit-tested across its subsystems, validated against standard reference problems (a Brio–Wu shock tube and others), and numerically stable for long runs. Recent work has been performance and robustness — a roughly 2.6× GPU speedup and a hardware migration — rather than new physics.

The physics question the simulator exists to answer is still open. The benchmark coupling validation against the standard reference case has not yet produced a measurable tilt-mode growth rate; the runs complete stably but return no extractable result, so they have never been compared to the benchmark curve. An earlier validation step was reduced in scope when it emerged that the implemented equilibrium did not yet have true field reversal. A numerical drift in the MHD solver was suppressed with added dissipation rather than a corrected scheme, and the energy budget is acknowledged as not fully closed. The downstream steps — the actual V9 stability prediction and the fast-ion stabilization threshold scan — cannot begin until the benchmark gate passes.

Where it stands, plainly

- The open magnetic mirror is disproven as a reactor path, across every variant tried, for understood and largely fundamental reasons: open-mirror end loss, the stability/confinement dilemma, alpha non-confinement, a materials ceiling on minimum-B fields, and the simple-mirror gain ceiling near $Q \approx 1$.
- The collider variant, which used a mirror for beam recapture rather than confinement, is also closed: the beams deflected into the collider walls and the circulating population could not be sustained.
- The field-reversed configuration is the sole surviving candidate. It is the one topology among those studied with no loss cone — which is the actual reason to leave the mirror behind. It is a paper design awaiting a single verdict.
- That verdict requires the hybrid-MHD simulator to first reproduce a known benchmark, which it currently does not. This is the one blocking item.

Every approach that relied on a magnetic mirror — for thermal confinement or for beam recapture — has failed, for reasons that are now understood rather than mysterious. The program's one open path abandons the mirror entirely. The next defensible claim it can make, in either direction, is gated on getting the benchmark validation to return a real tilt-mode growth rate.

Cultivated Progression · cultivatedprogression.com · Las Vegas, NV