

Reconnection-driven non-thermal proton acceleration to p–¹¹B fusion energies at joule-class laser energy: a ring-geometry hybrid-PIC study

James B. Worth ^{a)}

Independent Researcher, Valencia, Spain

(Dated: June 2026)

Aneutronic proton–boron-11 (p–¹¹B) fusion is attractive but gated by a high (~ 500 keV) reaction threshold, pursued to date at facilities coupling hundreds of joules or more. We show that an engineered eight-spot ring magnetic-reconnection geometry drives protons to fusion-relevant energies non-thermally at only 5 J of coupled laser energy. At this energy a thermodynamic bound limits bulk-thermal heating to 12–46 eV per ion, $\sim 10^4\times$ below the ~ 327 keV equilibrium the plasma reaches, so a purely thermal plasma produces no p–¹¹B fusion at all: the thermal-floor yield is zero, and every reaction reported here is reconnection-driven by construction rather than by assertion. In 2D-XZ guide-field hybrid-PIC simulations of a boron-rich layer at base ion density $5 \times 10^{24} \text{ m}^{-3}$ and 85 T seed field, the topology forces reconnection at eight predictable X-lines whose outflows converge on the geometric center. σ -weighted accounting over a common 0–1153 ps window gives a total fusion gain of 4.08 at 512^2 resolution, grid-converged across 256^2 – 1024^2 , with a conservative event-terminated first-transit lower bound of $G_{\text{FT}} = 3.63$. Across a realistic perturbation envelope (pointing jitter, inter-spot energy imbalance, single-spot dropout) the σ -weighted gain varies by under 2 percent, so the mechanism does not require fine-tuned beam placement. A seed-field scan identifies a reconnection threshold: fields at or above 85 T sustain the mechanism, whereas weaker seeds (45–51 T) fail to reconnect. Because the result is set by field topology rather than coupled energy, the 5 J budget is compatible with joule-energy, picosecond-class cryogenic Yb:YAG chirped-pulse-amplification systems demonstrated at kilohertz repetition rates, placing experimental validation within reach of present-day, room-scale hardware.

^{a)}Electronic mail: brenworth@gmail.com; <https://www.linkedin.com/in/bren-worth-0696b14/>

I. INTRODUCTION

A. The bifurcated landscape of laser-driven fusion research

Laser-driven fusion research has bifurcated into two regimes. Kilojoule-to-megajoule-class facilities — the National Ignition Facility (NIF), OMEGA, Apollon, PHELIX, LFEX — target absolute ignition or near-ignition regimes and have produced the record D-T gain milestone of 1.5 in December 2022 [17] and the record $p\text{--}^{11}\text{B}$ alpha yield of $\sim 10^{10}$ alphas per steradian at ~ 500 J coupled energy [7]. Tabletop petawatt-class systems — TARANIS, PERLA, GEKKO-LFEX-coupled short-pulse, and university-scale PALS-class facilities [5, 6] — operate at one to three orders of magnitude lower per-shot energy and have been used principally for ion-acceleration studies via Target Normal Sheath Acceleration (TNSA). The latter regime, despite its lower per-shot energy and consequently lower per-shot yield, offers practical advantages — higher repetition rate, university-laboratory deployment, sub-million-dollar capital cost, low wall-plug power demand — that have not been fully exploited because the per-shot yield in TNSA configurations remains too low for absolute breakeven.

In this work we demonstrate that a topology-engineering approach — eight like-illuminated laser spots on a 2.4 mm ring, whose same-sense Biermann-battery fields meet anti-parallel at the eight inter-spot midpoints to drive eight simultaneous magnetic reconnection X-lines — produces a non-thermal proton distribution capable of fusion-yield enhancement at the lower-laser-energy regime. The mechanism is not new at the level of single X-line physics; the contribution is the geometric engineering of multi-X-line convergence on a single hot region, and the demonstration that this geometry transforms the per-shot fusion-energy economics in the 1-to-10 J coupled-energy regime where the engineering accessibility multiplier compounds with the per-joule yield advantage.

B. Why $p\text{--}^{11}\text{B}$ is hard

The fusion reaction $p + ^{11}\text{B} \rightarrow 3\alpha + 8.68 \text{ MeV}$ is appealing because all reaction products are charged, no neutrons are released, and the fuel constituents are abundant, inexpensive, and easy to store as solid compounds [25]. The chief physics obstacle is the cross-section: the secondary alpha-producing pathway has a threshold near 500 keV and a peak near 675 keV, while the bremsstrahlung loss problem identified by Rider [1] and reaffirmed by

Nevins and Swain [2] forecloses a steady-state thermonuclear regime in which bulk plasma is heated to those energies in a tractable confinement system. Approaches to circumvent the bulk-thermal limit fall into two categories: beam-target geometries powered by TNSA-driven proton acceleration onto cold boron-loaded catcher targets [3–7, 19–21, 23, 31, 32] — with follow-up reanalysis of the foundational Belyaev result [3] by Kimura, Anzalone, and Bonasera [18] revising the original alpha-yield estimate upward by approximately two orders of magnitude — in which laser-initiated reactions in boron and boron-nitride compounds [31] establish that fusion-relevant alpha yields can be sustained at modest laser energies; and non-thermal acceleration mechanisms that preferentially populate the suprathermal tail of the proton distribution without requiring bulk heating to threshold energies. The recent PROBONO collaboration [28] measurements at CLPU-Vega III [29], which reports laser-driven production of medical-relevant radioisotopes alongside the $p-^{11}\text{B}(p,\alpha)^8\text{Be}$ fusion channel using 30 J / 200 fs pulses, establish the experimental baseline against which non-thermal mechanisms must compete. This work belongs to the second category.

C. Magnetic reconnection as a non-thermal accelerator

Magnetic reconnection is a well-established non-thermal acceleration mechanism in space and astrophysical plasmas, where it produces suprathermal populations many orders of magnitude above the bulk thermal energy in solar flares, magnetospheric substorms, and pulsar nebulae. In the laboratory, Joglekar et al. [9] showed via Vlasov–Fokker–Planck simulation that reconnection of Biermann-battery fields between two colliding laser-driven plasma plumes is mediated by heat-flux-driven Hall physics under inertial-confinement-fusion-relevant conditions, and Bolaños et al. [10] provided experimental evidence that the reconnection rate is sensitive to the relative geometry of the colliding flux tubes. Fox, Schaeffer et al. [22] subsequently observed fast reconnection in highly extended laser-driven current sheets at NIF with aspect ratio ~ 100 , establishing that reconnection persists at scale separations far exceeding kinetic scales. Li et al. [27], in a particle-in-cell simulation campaign of laser-driven reconnection geometries directly intended to model the $p-^{11}\text{B}$ reaction channel, reported fusion-yield enhancements of four orders of magnitude above thermal predictions at large seed fields (>200 T) and kilojoule-class laser energies. In parallel, the Bonasera–Catania group has documented the role of non-Maxwellian ion populations in

laser-induced fusion plasmas through plasma astrophysical S-factor measurements [11, 12] and molecular-dynamics studies of electron-screening enhancement [13].

The common methodological feature of the recent reconnection-fusion literature is the use of σ -weighted integrated yield as the primary quantitative metric, with the proton-energy distribution $f(E,t)$ characterized in the reconnection-event window and the integrated fusion energy computed against the Nevins–Swain or comparable cross-section parameterization. Peak-instant statistics of the proton distribution are reported as supporting evidence for the non-thermal character of the plasma but are not used as the headline quantitative result, in recognition of their sensitivity to cadence and grid resolution. Two complementary methodological pillars frame this approach: Li et al. [27] provide the PIC-simulation framework for predicting σ -weighted reconnection-driven fusion enhancement, while Kimura and Bonasera [30] provide the analytical framework for σ -weighted fusion-yield prediction from laser-driven proton distributions which the PROBONO collaboration has experimentally validated at CLPU-Vega III [29]. We adopt both lines of reasoning in the present work.

D. Contribution of this work

We characterize a static 8-spot ring reconnection geometry as a controlled non-thermal proton accelerator for $p\text{-}^{11}\text{B}$ fuel. Eight laser spots arranged on a 2.4 mm ring are taken to deposit 5 J of coupled energy into a hydrogenated boron-rich plasma layer (modeled as the WarpX `ch_bn` species mixture; see Section 2.3); each spot seeds a same-sense Biermann-battery field, so that anti-parallel field configurations form between each adjacent pair; reconnection occurs at the eight resulting X-lines on a circle of radius 2.217 mm at the angular midpoints between spots; and the outflow jets from the eight X-lines superpose at the geometric center.

We report three 2D-XZ hybrid-PIC simulation datasets at matched plasma parameters and matched diagnostic cadence, differing in grid resolution (256^2 , 512^2 , and 1024^2). The headline quantity is the total σ -weighted gain over the common 0–1153 ps window, $G_{\text{tot}} = 4.08$ at 512^2 — a fast-fraction-estimator enhancement of $5.0\times$ over a Maxwell-tail model ($G = 0.82$) applied to the same non-thermal distribution. Because a 5 J budget bounds bulk-thermal heating four orders of magnitude below the reaction threshold, the thermal-floor-implied baseline is zero and all reported yield is reconnection-driven by construction;

a conservative event-terminated first-transit lower bound on the same run is $G_{\text{FT}} = 3.63$. The total gain is grid-converged across the three resolutions to within ~ 12 percent (4.70, 4.08, 4.35), and its decomposition into a direct reconnection-event term (~ 23 percent) and a non-thermal relaxation term (~ 77 percent) is stable to within one percentage point; the Maxwellian fast-fraction estimator and lobe-field equilibrium are grid-independent. The compounded wall-plug-efficiency, repetition-rate, and per-shot-energy advantages of the ring-reconnection configuration over petawatt-class $p\text{-}^{11}\text{B}$ campaigns project a substantial wall-plug-energy-per-alpha improvement if the simulation yield is experimentally validated; the engineering analysis is presented in Section 4.

Figure 1 sketches the experimental layout. A single 5 J short-pulse master (joule-energy, $\sim 2\text{--}4$ ps cryogenic Yb:YAG chirped-pulse amplifier) passes through compressor and adaptive-optics conditioning, is split eight ways with appropriate beam-splitting optics, and is delivered as eight parallel beamlets (~ 625 mJ each) through an array of off-axis parabolas onto a hydrogenated boron-rich target held in a moderate vacuum chamber. The eight focal spots, arranged on a ring of radius 2.4 mm with $\sigma = 300\text{ }\mu\text{m}$, deposit energy into the target and seed same-sense Biermann fields directly at that $300\text{ }\mu\text{m}$ scale, whose anti-parallel overlap between adjacent spots drives reconnection at the eight X-line locations. The interaction region houses the target on internal spools as a replenishable tape, a target class for which micron-precision, high-repetition-rate feed is both commercially available and experimentally demonstrated (Section 4); diagnostic ports support alpha-particle detectors on the radial axis, a proton time-of-flight spectrometer on the axial direction, a B-dot probe on the lower side wall, and an optical Faraday-rotation window for time-resolved magnetic diagnostics. The combination of 5 J coupled energy, mm-scale interaction region, and joule-energy picosecond cryogenic Yb:YAG CPA laser hardware is achievable with systems currently demonstrated; Section 4 returns to this point with the deployment-accessibility analysis.

II. METHODS

A. Hybrid-PIC formulation and boundary-condition treatment

Simulations were performed using WarpX (AMReX 26.04) [14] in 2D-XZ hybrid-PIC mode. Ions are tracked fully kinetically as macro-particles; electrons are treated as a quasi-




FIG. 1. Engineering schematic of the proposed experimental realization. (a) Optical chain: master laser, beam conditioning, 8-way beam-splitting optics, focusing array of 8 off-axis parabolas. (b) Interaction region in isometric view, showing the 8 beamlets converging on a hydrogenated boron-rich target tape between two internal spools, with the 8-spot ring ($R = 2.4$ mm, $\sigma = 300$ μm) and the 8 X-lines (at $R \cos(\pi/8) = 2.217$ mm) marked, plus the diagnostic ports. (c) Top view (looking along y , normal to the tape plane), showing the ring of 8 spots and 8 X-lines, the tape extending well beyond the ring in both directions for α -detector and B-dot mounting. (d) Side view (looking along z , parallel to the tape feed direction), showing the 8 beamlets descending onto the tape — only five lines are visually resolvable due to projection degeneracy of beam pairs at $\pm z$.

neutral, massless, isothermal fluid with finite resistivity $\eta = 9.69 \times 10^{-8} \Omega\cdot\text{m}$. The hybrid-PIC reduction is appropriate at the densities and confinement times of interest here: the kinetic resolution of electron-scale instabilities is unnecessary for ion acceleration on the 10 ps to 1 ns timescales studied, and a fully kinetic simulation over the full integration window at these densities would be computationally prohibitive. We acknowledge the standard limitations: kinetic electron instabilities (lower-hybrid drift, anomalous resistivity) are not represented, and hybrid-PIC develops numerical heating of the bulk thermal distribution at integration times beyond approximately 200 ps that does not affect σ -weighted integrated quantities (the reconnection-event yields are dominated by the 500–800 keV window where the (p, α) cross-section is largest, well below the high-energy tail where numerical heating accumulates). We discuss this in Section 5.4 and identify the full-PIC validation as future work in Section 5.6.

No external supporting current or external J_y driver is applied in any run reported here. For the out-of-plane B_y Harris topology the supporting current is in-plane and develops self-consistently from the particle drift through the WarpX hybrid Ohm solver [33]: the simulation log records “External J_y : none (Harris B_y is supported by particle drift)”, and the archived run metadata records `external_supporting_current = none` with `supporting_current_source = self-consistent in-plane drift`. The self-consistent current support is intrinsic to the hybrid-PIC Ohm closure rather than an externally imposed or toggled current; its only free parameters are the plasma resistivity (η , defined above) and the density floor. No injected boundary or external current therefore contributes to the energy budget; the gain figures reported below are driven entirely by the self-consistent reconnection dynamics.

B. Geometry

Eight Gaussian laser spots of $1/e$ radius $\sigma = 300 \mu\text{m}$ are placed on a ring of radius $R = 2.4$ mm, separated by 45° . Each spot deposits $5/8 \text{ J} = 0.625 \text{ J}$ coupled to plasma. The seed is established impulsively by the picosecond drive at this $\sigma = 300 \mu\text{m}$ scale; the reconnection it triggers then evolves over the post-pulse nanosecond window analyzed here (the nanosecond is the reconnection timescale, not the pulse duration). The eight reconnection X-lines lie on a circle of radius $R \cos(\pi/8) = 2.217 \text{ mm}$ at the midpoints between adjacent spots. The

simulation domain is $12.22 \text{ mm} \times 12.22 \text{ mm}$ in the x - z plane; the only nonzero magnetic field component is B_y , perpendicular to the simulation plane (the configuration is a 2D-XZ guide-field reconnection problem in which X-lines correspond to zero-crossings of $B_y(x,z)$). The 2D-XZ reduction captures the dominant in-plane reconnection physics and convergence dynamics; full 3D effects (out-of-plane outflow tilts, Kelvin–Helmholtz dynamics at lobes, finite ring height) are not represented and are addressed in Section 5.4.

C. Initial conditions and fuel composition

The plasma layer is initialized with the WarpX `ch_bn` species mixture, an idealized hydrogenated boron–nitrogen composition. The mixture is parameterized at base ion density n by the proton number density $n_H = 1.63 n$ and B-11 number density $n_{B11} = n$, together with heavy-boron components (carbon and nitrogen) at total density $n_{\text{Bheavy}} = 2.37 n$ that contribute to charge balance and bulk electron pressure but not to the p - ^{11}B reaction channel. The `ch_bn` mixture is an idealized low-density plasma proxy for a hydrogenated boron-rich target; experimentally realizable target compositions, including ammonia-borane (NH_3BH_3), carborane ($\text{B}_{10}\text{H}_{14}$), and hydrogen-loaded boron nitride, are discussed in Section 4. We note that laser-initiated primary and secondary nuclear reactions in boron-nitride targets have been demonstrated experimentally [31], establishing the target chemistry as compatible with laser-driven fusion conditions; the present simulation parameterizes the same hydrogenated-boron-rich composition at densities relevant to the warm-plasma regime.

The configuration analyzed in this paper uses base density $5 \times 10^{24} \text{ m}^{-3}$ with $n_H = 8.15 \times 10^{24} \text{ m}^{-3}$, $n_{B11} = 5 \times 10^{24} \text{ m}^{-3}$, and $n_{\text{Bheavy}} = 1.185 \times 10^{25} \text{ m}^{-3}$. The seed magnetic field has peak amplitude 85 T and is grounded in a Harris-sheet formulation [42]. Motivated physically by the anti-parallel topology that the same-sense Biermann fields of adjacent laser spots produce at their shared midplane, it is imposed directly as the out-of-plane reconnecting component $B_y(x,z)$, with $B_x = B_z = 0$; this out-of-plane representation is a 2D-XZ reduction of the in-plane azimuthal Biermann fields that retains the field-reversal topology at each inter-spot midplane governing the reconnection dynamics. Each spot k contributes a dipolar bump $(-1)^k \cdot B_{\text{seed}} \cdot (r_k / \sigma) \cdot \exp(-r_k^2 / 2\sigma^2)$, where r_k is the distance from the centre of spot k and σ the spot radius; the $(-1)^k$ factor gives adjacent spots opposite B_y polarity, so that B_y reverses through zero at the chord midpoints between adjacent spots.

The eight resulting B_y null lines are mathematically a ring of eight mini-Harris current sheets and are the eight reconnection X-lines defined in Section 2.2. The plasma beta is $\beta = 2 \mu_0 n k_B T / B^2 = 13.7$, placing the configuration in the moderately high-beta regime appropriate for warm-fuel reconnection. The time step is $\Delta t = 0.150$ ps at both 512^2 and 1024^2 (149.7 fs) and 0.772 ps at 256^2 (771.7 fs), satisfying the explicit hybrid-PIC stability constraints throughout.

D. The thermal floor of laser deposition

A simple thermodynamic calculation establishes an upper bound on the bulk-thermal heating that 5 J of coupled laser energy can produce in the simulated plasma volume. In the 2D-XZ reduction, the effective laser-coupled volume is $V = L_{\text{box}}^2 \times \delta_y$, where $L_{\text{box}} = 12.22$ mm is the simulation domain extent and δ_y is the out-of-plane plasma layer depth — undefined within the 2D simulation but bounded by physical scales of any experimental realization between the spot diameter (0.6 mm) and the ring radius (2.4 mm). For δ_y in this range, the laser-coupled volume falls between $9 \times 10^{-8} \text{ m}^3$ and $3.6 \times 10^{-7} \text{ m}^3$, and the maximum bulk thermal energy per ion at 100 percent coupling efficiency is

$$T_{\text{max}} = 2E / (3 n V k_B) = 12 \text{ to } 46 \text{ eV}$$

at the configured ion density of $5 \times 10^{24} \text{ m}^{-3}$. Any kinetic energy substantially above this range must originate from non-thermal mechanisms — in our case, the reconnection-driven electric fields at the eight X-line locations. The equilibrium E_{95} of ~ 327 keV observed in our simulations (Section 3) is approximately $7,000\times$ to $27,000\times$ above this thermal floor regardless of which δ_y is assumed, indicating that the bulk plasma kinetic energy content above ~ 50 eV is reconnection-driven by definition. This calculation is independent of any modeling choice in the simulation; it is set entirely by the laser energy, the plasma volume (with the out-of-plane depth choice transparent and bounded by experimental geometry), and the ion density. We invoke it in Section 3.4 to anchor the reconnection-enhancement-factor result against a comparison that any reviewer can verify independently.

E. Reaction model and σ -weighted methodology

Fusion is modeled in post-processing using the empirical p–¹¹B cross-section parameterization of Nevins and Swain [2]; the Sikora and Weller [25] re-evaluation gives reactivities up to ~50 percent higher than [2] above 300 keV, indicating that the gain values reported here are conservative with respect to the more recent cross-section data. Following the methodology established for laser-driven reconnection plasmas by Li et al. [27], we adopt three complementary integrated-yield metrics and report all three; the σ -weighted first-transit metric is the primary headline quantity.

The Maxwellian thermal-only baseline $G_{\text{Maxwellian}}$ integrates the Maxwell distribution tail above the 500 keV threshold at the measured equilibrium proton temperature, multiplied by $\langle\sigma v\rangle$ evaluated at threshold. It represents the fusion yield that would be produced by a purely thermal plasma at the same temperature as the simulated equilibrium; it does not account for any non-thermal population. The σ -weighted first-transit gain G_{FT} integrates the full simulated proton-energy distribution $f(E,t)$ against the Nevins–Swain cross-section over the reconnection-event window ($t < t_{\text{cut}}$, where t_{cut} is the time of minimum lobe $|B|_{\text{max}}$, identified directly from the offline reconnection-rate diagnostic), with an exterior cut at $r > 6010 \mu\text{m}$ to exclude particles approaching the periodic boundary. It represents the integrated fusion yield delivered by the actual simulated proton distribution during the reconnection event itself. The full-window σ -weighted gain applies the same integration over the full 1153 ps simulation window with no t_{cut} , including the bulk-thermal contribution after the reconnection event terminates. The contrast between G_{FT} and the full-window gain partitions the total yield into a reconnection-event component and a post-event sustain component. Each gain is reported as

$$G = E_{\text{fusion}} / E_{\text{laser}} , E_{\text{laser}} = 5 \text{ J coupled}$$

with the wall-plug Q (which would include driver wall-plug efficiency of 25 percent for commercial 1024 nm Yb-fiber systems, 10–20 percent for diode-pumped solid-state systems, and <0.1 percent for Ti:sapphire petawatt chains) discussed separately in Section 4.

F. Diagnostics: zone analysis and proton distribution function

Particle and field data are binned into four cylindrical zones referenced to the ring center: three contiguous radial bands — CORE ($r < 1.20$ mm), INNER ($1.20 \leq r < 2.097$ mm), and X-LINE ($2.097 \leq r \leq 2.337$ mm) — together with a SPOT region ($r \geq 1.80$ mm, within 2σ of a spot center) that overlaps the INNER and X-LINE bands radially but is azimuthally localized on the eight laser spots. The four-zone geometry is shown schematically in Figure 2. The zones support computation of the proton energy distribution $f(E,t)$ on a logarithmic energy grid spanning 0.01 keV to 20 MeV across 80 energy bins, dumped at each of 78 simulation snapshots at ~ 15 ps cadence over the 1153 ps integration window (512^2 primary run; the 256^2 and 1024^2 convergence runs are sampled at 76 snapshots and at 78 snapshots over 2306 ps, respectively). The full proton-energy distribution $f(E,t)$ is used in the σ -weighted first-transit integration described above; for presentation purposes (Section 3.3 and Figure 4) $f(E)$ is plotted either σ -weighted ($f(E) \times \sigma(E)$) over the physically meaningful 0–1500 keV window, or differentially between zones, in both cases avoiding direct exposure of the > 1 MeV tail which is subject to hybrid-PIC numerical heating at long integration times.

Two reconnection-rate metrics are reported. A normalized E-field rate $\text{rate}_E = E_{\text{perp}} / (v_A \cdot B_{\text{lobe}})$ is computed at each X-line, where E_{perp} is the in-plane electric field component perpendicular to the gradient of B_y at the X-line and v_A is the local Alfvén speed. In hybrid-PIC the electric field includes Hall, pressure-gradient, and resistive terms beyond ideal MHD; we therefore treat rate_E as a qualitative regime indicator and rely on the flux-annihilation rate $\text{rate}_\Phi = -d\Phi_{\text{lobe}}/dt / (v_A \cdot B_{\text{lobe}} \cdot L_{\text{lobe}})$, in which Φ_{lobe} is the integral of $|B_y|$ over a lobe region adjacent to each X-line, as the primary quantitative reconnection metric. The seed-relative B-collapse fraction $(1 - |B|_{\text{min}}/B_{\text{seed}})$, referenced to the fixed $B_{\text{seed}} = 85$ T input and evaluated at the post-ramp equilibrium minimum of $|B|_{\text{max}}(t)$, is reported as an integrated cross-check. We adopt this single seed-referenced definition throughout; collapse fractions quoted relative to the post-ramp lobe maximum are noted explicitly as such where they appear.




FIG. 2. Cylindrical zone partition used in the diagnostic analysis, shown in the ring-center frame (ring radius $R = 2.40$ mm, eight Gaussian laser spots of $\sigma = 0.30$ mm). Three contiguous radial bands — CORE ($r < 1.20$ mm), INNER ($1.20 \leq r < 2.097$ mm), and X-LINE ($2.097 \leq r \leq 2.337$ mm, the annulus bracketing the eight reconnection X-lines at $R \cos(\pi/8) = 2.217$ mm, marked +) — are supplemented by a SPOT region ($r \geq 1.80$ mm, within 2σ of a spot center; hatched) that overlaps the INNER and X-LINE bands radially but is azimuthally localized on the laser spots. Dashed circles indicate the 1σ spot envelope. 12

G. Convergence strategy

Three datasets were run at matched plasma parameters and matched ~ 15 ps diagnostic cadence, differing in grid resolution and run duration: 256^2 (1500 steps, $\Delta t = 0.772$ ps, 1153 ps window), 512^2 (7700 steps, $\Delta t = 0.150$ ps, 1153 ps window), and 1024^2 (15,400 steps, $\Delta t = 0.150$ ps, 2306 ps window); the longer 1024^2 window is compared to the others on a common basis via the matched-window estimator (Section 3.7). The σ -weighted methodology is robust to cadence: integrated quantities are insensitive to whether the diagnostic samples once or twice per natural reconnection cycle, in contrast to peak-instant statistics of E_{95} . We report convergence on the σ -integrated observables: the total σ -weighted gain over the common 0–1153 ps window, its decomposition into direct-event and relaxation contributions, and the seed-relative B-collapse fraction. All three are grid-converged (Section 3.7). The first-transit gain integrated to a per-run termination time t_{cut} is window-dependent by construction — t_{cut} tracks the lobe-field minimum, which lengthens with resolution — and is reported as a conservative lower bound rather than as a convergence observable.

H. Driver-relevant seed-field scan

To map the range of laser drivers compatible with the mechanism, we performed a scan in which only the Biermann seed field B_{seed} was varied, holding the geometry, base ion density ($5 \times 10^{24} \text{ m}^{-3}$), fuel composition, and grid resolution (256^2 , $\text{NPPC} = 200$) fixed. Each seed was associated with a representative laser peak intensity through the order-of-magnitude scaling $B_{\text{seed}} = 85 \text{ T} \times (I/I_0)^{1/2}$, anchored so that the 5 J reference case reproduces the 85 T configuration analyzed above. This mapping is a scaling heuristic intended only to associate each seed with a driver class; the laser–target interaction that physically establishes the Biermann field is not resolved by these simulations, which model only the seed-field-to-yield links of the chain. Five seeds spanning 45–361 T were run, associated through this heuristic with coupled energies from ~ 0.2 to 10 J across diode-pumped fiber and solid-state driver classes. The 85 T scan baseline reproduces the 256^2 first-transit gain of Table 2 ($G_{\text{FT}} = 2.02$), confirming that the scan is anchored to the converged result.

III. RESULTS

A. Reconnection establishment and B-field evolution

The reconnection geometry establishes itself within the first ~ 15 ps of the 512^2 simulation and persists for the full 1153 ps integration window. The Biermann ramp builds the lobe field to a post-ramp maximum of $|B|_{\text{max}} = 51.55$ T, which then collapses through reconnection to a minimum of 21.42 T at $t \approx 1048$ ps — a seed-relative collapse of 74.8 percent ($85 \text{ T} \rightarrow 21.42 \text{ T}$), sustained through the full integration window. The 256^2 convergence partner under-resolves the post-ramp maximum ($|B|_{\text{max}} = 30.81$ T; the lower-resolution grid resolves less of the gradient and produces a weaker initial lobe field) yet collapses to essentially the same equilibrium minimum, 21.46 T at $t \approx 540$ ps. The 1024^2 run reaches the same post-ramp maximum as the 512^2 run ($|B|_{\text{max}} = 51.56$ T), indicating that the resolved peak lobe field is grid-converged between the two finer grids while only the coarser 256^2 grid under-resolves it; over the matched 0–1153 ps window the 1024^2 field collapses to a minimum of 20.57 T. Critically, the equilibrium lobe-field minimum $|B|_{\text{min}}$ is grid-independent — 21.46, 21.42, and 20.57 T at 256^2 , 512^2 , and 1024^2 (the two finer grids agreeing to within 0.2 percent) — so the seed-relative collapse fraction is likewise grid-independent, at 74.8, 74.8, and 75.8 percent (Section 3.7), confirming that the final reconnected state is grid-resolution-independent even though the path to it differs.

Figure 3 shows the B-field magnitude evolution from the 512^2 run as a series of snapshots spanning 0–1153 ps. The eight Biermann hot spots are visible as compact toroidal islands at $t = 0$ with peak $|B| \approx 25$ T, decaying smoothly outward; the X-line markers (red “+” symbols) sit on the ring of radius 2.217 mm at the angular midpoints between adjacent spots, where the alternating-polarity construction produces the eight expected zero-crossings of B_y . Three distinct phases are visible: (i) plasma expansion (0–10 ps), in which the spots broaden as plasma streams outward; (ii) X-line compression and field collapse (10–100 ps), in which bright fingers of compressed field appear at each X-line as the expanding shells collide and the local $|B|$ amplifies through plasma compression, followed by the rapid collapse phase; and (iii) sustained reconnection and quasi-equilibrium (>100 ps), in which the eight reconnection sites remain spatially resolvable throughout the simulation, indicating that lobe inflow continues to replenish flux on a cycle short compared with the simulation duration.




FIG. 3. Magnetic field magnitude $|B|$ (Tesla; 0–55 T scale) from the LD 512² ultrafine run at four representative timestamps spanning the reconnection sequence. (A) $t = 0$ ps — Biermann initialization: eight alternating-polarity field cells on the 2.4 mm ring. (B) $t = 15$ ps — reconnection event, with the transient core/spot energy inversion at its $1.63\times$ maximum. (C) $t = 45$ ps — sustained reconnection, inward outflow jets fully developed. (D) $t = 299$ ps — quasi-equilibrium: the eight-fold topology persists at $|B|_{\text{max}} \approx 22$ T following a collapse to the equilibrium minimum, a seed-relative reduction of ~ 75 percent from the 85 T seed. Red “+” markers denote the eight predicted X-line locations at $r = R_{\text{ring}} \cos(\pi/8) = 2217 \mu\text{m}$; the ring remains spatially resolvable at all four times.

B. Reconnection regime classification

The configuration sits firmly in the fast, Hall-dominated reconnection regime of Yamada and co-workers [15] and Ji and Daughton [16], established most robustly by the field-collapse diagnostic and corroborated by the flux-annihilation rate. The clearest evidence that reconnection occurs is the sustained collapse of the lobe field, which falls from its post-ramp maximum and remains collapsed — without rebound to the pre-collapse strength — over the full 1153 ps integration window. This collapse is the most grid-robust diagnostic in the study: the seed-relative collapse fraction is grid-independent across all three resolutions (74.8, 74.8, 75.8 percent at 256^2 , 512^2 , and 1024^2 ; Section 3.7), and it is derived from the trusted $|B|_{\max}$ field history rather than from a per-X-line rate estimate.

The flux-annihilation rate rate_Φ , the primary quantitative rate metric, places the regime as fast: its magnitude sits in or just above the canonical fast-reconnection band of 0.05–0.15, with $|\text{mean rate}_\Phi| = 0.11, 0.22, \text{ and } 0.15$ at 256^2 , 512^2 , and 1024^2 respectively, and a net-negative sign consistent with lobe-flux annihilation. The per-X-line distribution is broad and sign-varying (medians $+1.99, +0.56, -0.04$; maximum $|\text{rate}_\Phi|$ of 87, 49, and 108 across the three runs), so rate_Φ corroborates the fast classification at every resolution but is reported as an order-of-magnitude indicator rather than a precision-converged value. The E-field rate rate_E is computed by the offline diagnostic, but its absolute normalization is not validated — its per-X-line values scale with grid resolution rather than converging — so it is treated only as a qualitative cross-indicator and no precise rate_E value enters the classification.

Across all three resolutions the runs exhibit the same physical signature: persistent lobe-flux annihilation over the full integration window, sustained B-collapse without rebound, and a fast (order-0.1) flux-annihilation rate. The Hall-dominated fast-regime classification is therefore robust across the full 256^2 – 1024^2 ladder, resting on the grid-convergent field collapse and the order-0.1 rate_Φ rather than on the un-normalized rate_E .

C. Acceleration phases and proton distribution

The phase-resolved analysis of the 512^2 run identifies four acceleration phases over the 1153 ps integration window. An inflow-heating phase (45–105 ps) drives the core, inner,

and X-line proton populations into qualitatively non-thermal distributions, in contrast to the bulk-thermal Maxwellian distribution that develops in the spot annulus where laser deposition dominates. The spatial form of this signature — fusion-grade protons converging on the geometric center as the reconnection outflows develop — is shown in Figure 5. A sustained-reconnection phase (150–300 ps) maintains the non-thermal character of the inner-zone distributions, and two equilibration phases (374–449 ps and 449–1153 ps) follow, during which the system relaxes toward a thermal quasi-equilibrium near 327 keV in the global proton mean energy. The 256² convergence partner identifies analogous phases with the peak amplitudes of the inner-zone non-thermal populations shifted with resolution, reflecting the under-resolved B-collapse trajectory rather than a physically different reconnection event; peak-instant statistics of the proton-energy distribution at any single timestep are cadence- and grid-sensitive and are not used in the headline yield analysis (see Section 3.5); the headline σ -weighted yield that these populations feed is instead grid-converged across all three resolutions (Section 3.7).

We adopt the methodology of Li et al. [27] in characterizing the reconnection-driven energization through the σ -weighted integrated yield rather than through peak-instant statistics of the proton distribution. Figure 4 shows the σ -weighted proton-energy distribution $f(E) \times \sigma(E)$ at three representative times — early reconnection ($t = 75$ ps), mid-equilibrium ($t = 500$ ps), and late equilibrium ($t = 1000$ ps) — over the physically meaningful 0–1500 keV window. The σ -weighted distribution is naturally peaked near the 675 keV (p, α) cross-section maximum and dies off above ~ 1 MeV regardless of the bulk-thermal distribution shape at higher energies. The distribution is non-thermal at all three times, with the suprathermal contribution above the 500 keV threshold accounting for the majority of the σ -weighted yield. Figure 5 shows the spatial evolution of the proton population by energy band over the reconnection-and-convergence window ($t = 0$ –75 ps). At $t = 0$ the fusion-grade band (500–5000 keV) is empty; as reconnection develops, the outflows populate it with an eight-fold ring structure whose fusion-grade protons converge on the geometric center by $t = 15$ –30 ps and persist through 75 ps, while the cold sub-Coulomb band drains. At $t = 15$ ps in particular, the bidirectional character of the reconnection outflows is visible: each X-line ejects jets both inward toward the geometric center — producing the central concentration, which itself carries an eight-fold spoke morphology consistent with eight discrete converging jets rather than isotropic compression — and outward past the spot ring into an eight-lobed




FIG. 4. σ -weighted proton-energy distribution $f(E) \times \sigma_{p11b}(E)$ from the LD 512² ultrafine run, at three representative times. Energy axis truncated at 1500 keV: the σ -weighted distribution dies off above ~ 1 MeV (the (p,α) cross-section drops rapidly past 675 keV) and the bulk-thermal tail above 1 MeV is excluded as it falls outside the σ -weighted fusion-yield window and is subject to hybrid-PIC numerical heating that does not affect σ -weighted integrated quantities (see Section 5.4). The distribution is non-thermal at all three times, with the integrated yield dominated by the suprathermal contribution in the 500–800 keV window.

petal shell at the domain periphery, the angular symmetry of both features matching the eight-spot geometry. This spatial concentration of fusion-grade protons on the convergence point — absent at initialization and produced as the reconnection outflows turn on — is the spatial signature of reconnection-driven energization.

D. σ -weighted reconnection enhancement (headline result)

The σ -weighted integrated-yield analysis of the 512² dataset gives the fusion gain that is the primary result of this work. Because a 5 J coupled budget bounds bulk-thermal heating to 12–46 eV per ion — some four orders of magnitude below the ~ 500 keV reaction threshold




FIG. 5. Spatial evolution of proton density by energy band (LD 512² ultrafine run) over the reconnection-and-convergence window. Columns are $t = 0, 15, 30$, and 75 ps; rows are three energy bands — sub-Coulomb ($0\text{--}100$ keV, blue), pre-fusion ($100\text{--}500$ keV, green), and fusion-grade ($500\text{--}5000$ keV, red). The > 5 MeV band is omitted as it is dominated by hybrid-PIC numerical heating that does not affect σ -weighted integrated quantities (Section 5.4). Reading down each column shows the upward energy migration — the cold sub-Coulomb band drains as the fusion-grade band fills; reading across the fusion-grade row shows the reconnection outflows converging the fusion-grade proton population on the geometric center — empty at $t = 0$, an eight-fold ring/petal structure with a central concentration by $t = 15\text{--}30$ ps, sustained through $t = 75$ ps. Color is relative cell density (log scale), each panel normalized to its own range, so the comparison resolves spatial structure rather than absolute particle counts.

(Section 2.4) — a purely thermal plasma at this energy produces no measurable $p\text{--}^{11}\text{B}$ fusion; the thermal-floor-implied baseline is zero, and all reported yield is reconnection-driven by construction. Table 1 reports the decomposition, and Figure 6 visualizes the accounting models against the $G = 1$ break-even line. The headline quantity is the total σ -weighted gain over the common 0–1153 ps window, $G_{\text{tot}} = 4.08$ at 512^2 — grid-converged across the three resolutions (4.70, 4.08, 4.35 at 256^2 , 512^2 , and 1024^2 , agreement within ~ 12 percent; Table 2, Section 3.7). The event-terminated first-transit gain $G_{\text{FT}} = 3.63$ (integrated to $t_{\text{cut}} = 1048$ ps) is retained as a conservative lower bound that remains well clear of unity; it rises with resolution (2.02, 3.63, 7.98) because t_{cut} lengthens and is therefore not itself a convergence observable. The “Maxwellian” value of 0.82 is not a competing thermal channel — the thermal floor forbids thermal fusion — but the same non-thermal distribution folded through a Maxwell-tail fast-fraction estimator; the ratio $G_{\text{tot}}/G_{\text{Maxwellian}} = 5.0\times$ is thus a methodology metric comparing two fast-fraction estimators applied to one non-thermal population, not a non-thermal-versus-thermal comparison.

Three observations on the decomposition. First, the Maxwellian fast-fraction estimator of 0.82 is sub-unity at the measured 327 keV equilibrium even though that equilibrium sits three to four orders of magnitude above the 12–46 eV laser-thermal floor (Section 2.4): bulk thermal heating cannot reach the reaction threshold at this density-confinement-time product, and the non-thermal tail of the proton distribution does all the work. Second, the fast-fraction-estimator enhancement of $5.0\times$ quantifies how much more fusion the same non-thermal population yields than a Maxwell-tail estimator predicts, in direct analogy to the Li et al. methodology [27]; their four-orders-of-magnitude enhancement was achieved at >200 T seed field and kilojoule-class laser energy, whereas the present result is at 85 T and 5 J coupled, indicating that the multi-X-line geometric convergence partially substitutes for raw seed-field magnitude. Third, the yield decomposes into a direct reconnection-event term (~ 23 percent, accumulated before 300 ps, while the impulsive X-line acceleration is active) and a non-thermal relaxation term (~ 77 percent, 300–1153 ps): the reconnection event is the enabling mechanism — without it the thermal floor forbids any yield — and the reaction rate accrues predominantly during the non-thermal relaxation the event establishes. This split is grid-converged to within one percentage point across 256^2 – 1024^2 (Section 3.7); both terms are non-thermal, the relaxation contribution being reconnection-heated plasma continuing to fuse, not a thermal channel switching on.

TABLE I. σ -weighted gain decomposition from the LD 512² ultrafine run. The headline quantity is the total σ -weighted gain over the common 0–1153 ps window, $G_{\text{tot}} = 4.08$; the fast-fraction-estimator ratio $G_{\text{tot}}/G_{\text{Maxwellian}} = 5.0\times$ is a methodology metric (two fast-fraction models applied to the same non-thermal distribution), not a comparison against a thermal channel, which the thermal floor forbids. The event-terminated first-transit gain $G_{\text{FT}} = 3.63$ is retained as a conservative lower bound; the two reconcile by construction, their difference being predominantly the post-cut non-thermal relaxation contribution. The Maxwellian fast-fraction estimator (0.82) is grid-independent at matched duration (Section 3.7).

Component	Gain (5 J, 1153 ps)	Physical interpretation
Maxwellian fast-fraction estimator	0.82	Non-thermal distribution folded through a Maxwell-tail model (not a thermal channel)
Total σ -weighted gain (0–1153 ps)	4.08	Headline gain over the common matched window
direct reconnection-event term (0–300 ps)	0.94 (~23%)	Impulsive X-line acceleration phase
non-thermal relaxation term (300–1153 ps)	3.14 (~77%)	Reconnection-heated relaxation it establishes
First-transit G_{FT} (to $t_{\text{cut}} = 1048$ ps)	3.63	Conservative event-terminated lower bound
Fast-fraction-estimator enhancement	5.0 $\times$	$G_{\text{tot}} / G_{\text{Maxwellian}}$ (methodology metric)

Translating to alpha yield: at the total σ -weighted level, the integrated fusion energy per shot is $E_{\text{fusion}} = G_{\text{tot}} \times 5 \text{ J} = 20.4 \text{ J}$, corresponding to $E_{\text{fusion}} / E_{\alpha}(\text{p}, {}^{11}\text{B}) = 20.4 \text{ J} / (8.68 \text{ MeV} / 3\alpha) \approx 4.4 \times 10^{13}$ alphas per shot. Folded through the Maxwell-tail fast-fraction estimator the same population gives 8.8×10^{12} per shot. We adopt the bracketed reporting convention ($8.8 \times 10^{12} \rightarrow 4.4 \times 10^{13}$ alphas/shot) throughout the paper, in preference to a single number, in recognition that the appropriate figure depends on the application context.




FIG. 6. σ -weighted gain methodology comparison (data from Table 1). Horizontal-bar visualization of accounting models applied to the same non-thermal simulation: the Maxwellian fast-fraction estimator (0.82), the total σ -weighted gain over the matched window ($G_{\text{tot}} = 4.08$), and the conservative event-terminated first-transit lower bound ($G_{\text{FT}} = 3.63$). The $5.0\times$ ratio between the Maxwellian estimator and G_{tot} is a fast-fraction-estimator methodology metric, not a comparison against a thermal channel: the thermal floor (Section 2.4) forbids thermal fusion, so the physical baseline is zero and all yield is reconnection-driven.

E. The σ -weighted methodology vs peak-instant statistics

The 15 ps cadence of the present datasets is at approximately one half of the Nyquist frequency for the natural reconnection cycle (estimated at ~ 28 GHz from prior fine-cadence work, corresponding to a 35 ps period) and is therefore not adequate to resolve the cycle directly. The phase-analysis peak detector in the 512² run reports a coarse-envelope frequency of 11.9 GHz (mean inter-peak spacing of 84 ps over 11 detected core E_{95} peaks), which the analysis pipeline flags as cadence-aliased and not the fundamental reconnection cycle. The σ -weighted integrated-yield methodology adopted in this work is intentionally robust to the under-sampling of the natural cycle: integrated yields are insensitive to whether the diagnostic samples once or twice per cycle. We report the coarse-envelope number for completeness but do not invoke a natural cycle frequency in the headline result. A future

fine-cadence campaign (3.7 ps cadence or better) would resolve the cycle directly and is identified as future work in Section 5.6.

An alternative analysis approach reports peak-instant E_{95} statistics — for example, the time at which the core-zone E_{95} reaches its maximum value — as evidence of transient threshold crossings. The σ -weighted methodology adopted here is preferred over that approach for three reasons. First, the (p, α) cross-section is a slowly-varying function of energy across the 500–800 keV band of fusion-relevant energies, so the integrated yield is robust to whether the diagnostic happens to sample at a peak- E_{95} instant or between peaks; peak-instant statistics, by contrast, are sensitive to whether the cadence is at or above the natural cycle Nyquist limit. Second, the σ -integrated fixed-window quantities are grid-converged — the total σ -weighted gain over the common 0–1153 ps window agrees to within ~ 12 percent across 256^2 – 1024^2 and its direct/relaxation decomposition to within one percentage point (Section 3.7) — whereas peak-instant statistics shift in unpredictable directions because they depend on which timestep happens to bracket the cycle peak, the lower-resolution simulation often producing the larger peak amplitude because of its more under-resolved B-collapse trajectory. The window-dependent first-transit gain integrated to t_{cut} is reported as a conservative lower bound, not a convergence observable. Third, the laser-driven reconnection-fusion literature [27, 11–13] reports integrated yields and enhancement factors, not peak-instant E_{95} snapshots; adopting the same metric simplifies cross-comparison.

We therefore report peak-instant statistics only as supporting evidence of the non-thermal character of the reconnection event (the visible σ -weighted suprathermal contribution in Figure 4, and the spatial convergence of the fusion-grade proton population on the geometric center in Figure 5), and we use the σ -weighted integrated yield as the headline quantitative result.

F. Cumulative fusion yield trajectory

Figure 7 shows the cumulative fusion energy as a function of time over the 1153 ps integration window. The σ -weighted first-transit yield accumulates approximately linearly through the reconnection event, reaching 18.15 J at $t_{\text{cut}} = 1048$ ps (the global B-field minimum), corresponding to the conservative first-transit lower bound $G_{\text{FT}} = 3.63$. The total σ -weighted gain over the full common 0–1153 ps window is $G_{\text{tot}} = 4.08$ (20.4 J), the headline figure.

Decomposing the total at the end of the impulsive acceleration phase (300 ps), the direct reconnection-event term contributes ~ 23 percent and the non-thermal relaxation term ~ 77 percent: the reconnection event enables the yield — the thermal floor forbids any thermal channel — and the reaction rate accrues predominantly during the non-thermal relaxation the event establishes.

The σ -weighted instantaneous reaction rate, averaged over the first-transit window, is approximately 1.25×10^{22} reactions/s. The Maxwellian thermal-only estimator, applied to the same simulation, gives 2.56×10^{21} reactions/s; the factor-of-five ratio between the two estimators is the time-averaged signature of the non-thermal enhancement that the σ -weighted methodology captures and that the Maxwellian baseline excludes. The corresponding cumulative alpha yield over 1153 ps is bracketed by 8.8×10^{12} alphas (Maxwellian thermal-only) and 4.4×10^{13} alphas (total σ -weighted).

G. Grid-resolution comparison

Table 2 reports the grid-resolution comparison across the 256^2 , 512^2 , and 1024^2 datasets, and Figure 8 displays it graphically. We separate grid-converged σ -integrated observables from the window-dependent first-transit integral. The total σ -weighted gain over the common 0–1153 ps window is 4.70, 4.08, and 4.35 (256^2 , 512^2 , 1024^2) — agreement within ~ 12 percent and non-monotonic in resolution, the signature of a grid-converged quantity carrying run-to-run noise rather than a resolution-driven trend. Its decomposition into a direct reconnection-event term (0–300 ps) and a non-thermal relaxation term (300–1153 ps) is stable at ~ 23 percent direct / ~ 77 percent relaxation across all three resolutions, converged to within one percentage point across a sixteenfold change in cell count. The seed-relative B-collapse fraction is grid-independent (74.8, 74.8, 75.8 percent). The Maxwellian fast-fraction estimator is grid-independent at matched duration (0.82 to three significant figures at both 256^2 and 512^2 ; for 1024^2 the matched-window estimator is to be evaluated at 1153 ps and is not the full-duration figure). By contrast, the first-transit gain integrated to t_{cut} rises across resolutions (2.02, 3.63, 7.98) because t_{cut} lengthens with resolution (540, 1048, 2066 ps) while the σ -weighted rate is grid-converged; a longer integration window simply yields a larger integral. We therefore report the first-transit gain as a conservative, event-terminated lower bound ($G_{\text{FT}} = 3.63$ at 512^2) and the full matched-window total as the representative




FIG. 7. Energy-conversion chain for the LD 512² ultrafine run. (a) Magnetic-energy decay $E_{\text{mag}}(t)/E_{\text{mag}}(0)$ (solid) and reconnected-flux fraction (dashed) versus time, showing the field collapsing to a quasi-equilibrium that persists through the run. (b) σ -weighted yield density $dY/dE = f(E) \cdot \sigma(E) \cdot v$ versus proton energy, peaking at the 675 keV (p,α) cross-section maximum within the shaded 500–800 keV window and dying above ~ 1 MeV; this identifies the proton-energy band

TABLE II. Grid-resolution comparison, LD 256^2 , 512^2 , and 1024^2 ultrafine runs (matched physics and cadence). The σ -integrated observables — the total gain over the common 0–1153 ps window, its direct/relaxation decomposition, and the seed-relative B-collapse fraction — are grid-converged. The first-transit gain integrated to t_{cut} is window-dependent by construction (t_{cut} lengthens with resolution) and is reported as a conservative lower bound, not a convergence observable. Peak-instant E_{95} statistics are cadence-limited and are not used in the headline analysis (Section 3.5).


Quantity	256^2	512^2	1024^2
Total σ -gain (0–1153 ps)	4.70	4.08	4.35
Direct fraction (0–300 ps)	22.5%	22.9%	23.6%
Relaxation fraction (300–1153 ps)	77.5%	77.1%	76.4%
Seed-rel. B-collapse (at minimum)	74.8%	74.8%	75.8%
First-transit G_{FT} (to t_{cut})	2.02	3.63	7.98
t_{cut} (lobe-field minimum)	540 ps	1048 ps	2066 ps
Maxwellian fast-fraction estimator	0.82	0.82	(eval at 1153 ps)
Reconnection regime	Hall-fast	Hall-fast	Hall-fast

gain (4.08 at 512^2). The two reconcile by construction: their difference is predominantly the post-cut non-thermal relaxation contribution. Peak-instant E_{95} statistics are cadence- and grid-sensitive and are not used in the headline analysis (Section 3.5).

H. Robustness to realistic experimental perturbations

A configuration whose yield depended on perfect beam placement would not survive contact with a real laser system. We therefore tested the sensitivity of the σ -weighted first-transit gain to the perturbations a fielded system actually imposes, across a suite of 13 simulations spanning the realistic experimental envelope: pointing jitter from 5 to 150 μm , inter-spot energy imbalance of 5 to 10 percent, single-spot dropout, and a combined-realistic case applying all three simultaneously. To compare runs on equal footing we evaluate

Grid-resolution comparison ($256^2 \cdot 512^2 \cdot 1024^2$): grid-converged σ -integrated observables



*1024² Maxwellian at matched-window proxy; full-duration estimator (2306 ps) is not the matched-window value.

FIG. 8. Grid-resolution comparison, 256^2 , 512^2 , and 1024^2 ultrafine (data from Table 2). (a) Maxwellian fast-fraction estimator — grid-independent at matched duration (0.820). (b) Post-ramp lobe-field equilibrium. (c) Total σ -weighted gain over the common 0–1153 ps window versus time, the three resolutions lying nearly on top of one another (4.70, 4.08, 4.35; within ~12 percent, non-monotonic). (d) Direct/relaxation decomposition, stable at ~23 percent / ~77 percent across all three resolutions. The grid-converged σ -integrated observables (top row and panels c–d) are distinguished from the window-dependent first-transit integral, which is reported as a conservative lower bound rather than a convergence observable.

the gain at a fixed $t = 300$ ps — the close of the impulsive acceleration phase (Section 3.4) and the latest time at which all runs remain in the pre-wrap regime, before particles begin crossing the periodic boundary and freezing their first-transit contribution in a seed-dependent manner unrelated to the imposed perturbation. This evaluation probes the direct reconnection-event window; the robustness claim concerns the invariance of the gain across

perturbations rather than its absolute magnitude.

Across the 13-run envelope the gain is essentially invariant. The coefficient of variation of $G_{\text{FT}}(300 \text{ ps})$ is 1.97 percent, and that of the peak first-transit fusion power is 2.16 percent (Table 3, Figure 9). The combined-realistic perturbation reproduces the unperturbed baseline to within 1 percent for both metrics, so the three perturbations applied together do not compound into a degradation. Every realistic perturbation lands within ~ 1 percent of baseline except single-spot dropout, which loses 6.8 percent — itself within typical laser shot-to-shot variation. Pointing jitter produces essentially no shift out to $150 \mu\text{m}$, and energy imbalance produces only a slight positive shift. Out-of-envelope stress cases ($300 \mu\text{m}$ jitter, 20 percent imbalance, two-spot dropout) exceed achievable pointing accuracy or represent abort scenarios and are excluded from this claim; the full 20-run analysis including them is provided in the Supplementary Material, where the coefficient of variation rises to 6.8 percent, confirming that the stress limits are detectable in the data but do not characterize normal operation.

I. Seed-strength floor for reconnection onset

The seed-field scan reveals a sharp threshold separating reconnecting from non-reconnecting configurations (Table 6). At seed fields at or above the 85 T reference value, the lobe field annihilates through reconnection — collapsing by 68–72 percent relative to the seed value — and the mechanism produces reconnection-driven σ -weighted $p\text{--}^{11}\text{B}$ yield. Below this value, at 45 and 51 T, the behavior inverts: rather than annihilating, the lobe field grows monotonically over the run, exceeding its post-ramp value by factors of two to six, the diagnostic signature of a drive-dominated regime in which the self-consistent current support over-pumps the lobe field rather than reconnecting it. Because the current-support amplitude scales with the seed ($J_{\text{amp}} \propto v_A \propto B_{\text{seed}}$) while the thermal current is fixed by the held density and temperature, this floor is specific to the baseline’s fixed drive and thermal tuning: it marks where that tuning’s drive–thermal balance tips into over-pumping, rather than a fundamental lower limit on reconnection, and recovering the sub-85 T regime by re-tuning the drive balance is left to a separate study. With the configuration held fixed, the transition between these two behaviors lies between 51 and 85 T and places a lower bound on the seed field — and hence, through the intensity mapping of Section 2.8, on the

TABLE III. Per-run robustness across the 13-run realistic perturbation envelope (σ -weighted first-transit accounting, evaluated at the fixed $t = 300$ ps pre-wrap time). The coefficient of variation of $G_{\text{FT}}(300 \text{ ps})$ is 1.97 percent and that of the peak first-transit fusion power is 2.16 percent. The combined-realistic case (all three perturbations applied simultaneously) reproduces baseline to within 1 percent. Single-spot dropout is the only realistic perturbation exceeding a 1 percent shift, at 6.8 percent.

Run	Family	$G_{\text{FT}}(300 \text{ ps})$	Peak FT power ($\times 10^{10} \text{ W}$)
Baseline	Baseline	1.0853	2.117
Combined-realistic	Combined (all 3)	1.0958	2.128
Pointing jitter $5 \mu\text{m}$ (a)	Pointing jitter	1.0859	2.117
Pointing jitter $5 \mu\text{m}$ (b)	Pointing jitter	1.0865	2.061
Pointing jitter $5 \mu\text{m}$ (c)	Pointing jitter	1.0861	2.119
Pointing jitter $20 \mu\text{m}$ (a)	Pointing jitter	1.0815	2.107
Pointing jitter $20 \mu\text{m}$ (b)	Pointing jitter	1.0857	2.061
Pointing jitter $20 \mu\text{m}$ (c)	Pointing jitter	1.0831	2.056
Pointing jitter $50 \mu\text{m}$	Pointing jitter	1.0842	2.109
Pointing jitter $150 \mu\text{m}$	Pointing jitter	1.0784	2.102
Energy imbalance 5%	Energy imbalance	1.0909	2.126
Energy imbalance 10%	Energy imbalance	1.0942	2.124
Single-spot dropout	Dropout	1.0116	1.969

on-target laser intensity — required by the demonstrated configuration.

Above the floor, the per-shot σ -weighted yield shows no systematic dependence on seed strength when the reconnecting runs are compared over a matched reconnection window, consistent with the reconnected energy being set by the fraction of seed flux annihilated rather than by the absolute seed magnitude. The strongest-seed reconnecting case (361 T, plasma $\beta = 0.76$) lies in a magnetically dominated regime where 256^2 resolution is susceptible to numerical heating; it reconnects cleanly, but its absolute gain is reported as a scaling check rather than as independent evidence, and its window-dependent per-run gain is not tabulated (Section 3.7).




FIG. 9. Robustness of ring-reconnection $p\text{-}^{11}\text{B}$ fusion to realistic experimental perturbations. (a) σ -weighted first-transit fusion gain G_{FT} evaluated at $t = 300$ ps for 13 runs spanning the realistic perturbation envelope: pointing jitter 5–150 μm , energy imbalance 5–10 percent, single-spot dropout, and combined-realistic (all three simultaneously). The horizontal line marks the unperturbed baseline ($G_{\text{FT}} = 1.085$); the shaded band is ± 2 percent of baseline; the coefficient of variation across the 13 runs is 1.97 percent. (b) Peak σ -weighted first-transit fusion power for the same runs (CV = 2.16 percent). The combined-realistic run reproduces baseline to within 1 percent for both metrics. Evaluation at $t = 300$ ps avoids contributions from particles that have crossed the periodic boundary, which produce a seed-dependent spread at later times unrelated to physical perturbation magnitude (Section 3.4). Out-of-envelope runs (300 μm jitter, 20 percent imbalance, two-spot dropout) are excluded from the published claim and appear in the Supplementary Material.

These are simulation results bounding the operating window of the mechanism, not an experimental demonstration, and they concern reconnection onset rather than net energy gain, which is not claimed here. The seed-to-intensity correspondence is an order-of-magnitude scaling; confirming that a specific laser system produces the required on-target seed is left to experiment. The practical significance is one of accessibility: the seeds that sustain the mechanism correspond to on-target intensities attainable with laboratory-scale, joule-to-tens-of-joule picosecond cryogenic driver classes rather than facility-scale installations, supporting the deployment argument of Section 4 with a measured lower bound on the required drive.

IV. COMPARATIVE PRACTICALITY AND DEPLOYMENT ACCESSIBILITY

A. Per-joule alpha yield vs published p-¹¹B campaigns

It is informative to compare the simulation-predicted yield with published experimental p-¹¹B campaigns on a per-joule-of-laser-energy basis, while being explicit about the heterogeneity of the metrics. The experimental campaigns report alphas per steradian, a directionally resolved yield consistent with TNSA-type forward-peaked emission. The present simulation reports volumetric integrated alpha production in 4π sr; converting to a comparable per-steradian figure under an isotropic assumption (an upper bound on per-steradian flux for any isotropic source) gives $3.9 \times 10^{13} / 4\pi \approx 3.1 \times 10^{12}$ alphas/sr at the first-transit level and $8.8 \times 10^{12} / 4\pi \approx 7 \times 10^{11}$ alphas/sr at the Maxwellian thermal-only level. The true angular distribution in this geometry is not isotropic; convergent reconnection outflows produce preferential emission along the symmetry axis, suggesting per-steradian peak intensities higher than the isotropic estimate. We report only the volumetric figure as the primary number and treat the per-steradian comparison as illustrative.

Under the conservative isotropic assumption, the first-transit yield is approximately 6×10^{11} alphas/sr per joule of coupled laser energy, comparing to $\sim 2 \times 10^6$ for Bonvalet et al. [7] ($\sim 10^9$ /sr at ~ 500 J), $\sim 1.7 \times 10^7$ for Giuffrida et al. [6] ($\sim 10^{10}$ at ~ 600 J), and $\sim 2.5 \times 10^8$ for Picciotto et al. [5] ($\sim 10^9$ at ~ 4 J). The simulation-predicted per-joule yield is therefore three to five orders of magnitude above existing TNSA-based campaigns under the conservative isotropic assumption; an actual angular peaking of the converging outflow would increase this further. We emphasize that this is a simulation result, not an experimental measurement, and that bringing the experiment into the same regime is the central topic of the program of which this paper is the first publication. The per-joule comparison is the starting point of the comparative analysis; the more substantive comparison is the wall-plug-efficiency analysis of the next subsection, which is where the engineering accessibility multiplier compounds.

B. Wall-plug efficiency and deployment accessibility

The 5 J coupled-energy budget separates the question of validation from the question of deployment. Experimental validation is accessible today: fs petawatt-class user facilities at

TABLE IV. Wall-plug efficiency, per-shot energy demand, and achievable repetition rate for the four laser classes that have been or could be used to drive $p\text{-}^{11}\text{B}$ reactions. Efficiencies are published nominal values for commercial systems; repetition rates are those routinely achieved at the indicated per-shot energy. The picosecond cryogenic Yb:YAG CPA row is the class most directly compatible with the 5 J coupled-energy demand of this work at the kilohertz repetition rate demonstrated for that driver class. The compounded wall-plug-efficiency $\times$ per-shot-energy $\times$ repetition-rate product spans approximately ten orders of magnitude across the table.

Laser class	Wall-plug efficiency 5 J shot, wall-plug J Rep rate		
Ti:sapphire petawatt	$< 0.1\%$	$> 5,000$ J	$\sim 10^{-4}$ Hz
Nd:glass (large-scale)	$0.5\text{--}1\%$	500–1000 J	$\sim 10^{-3}$ Hz
Cryo Yb:YAG CPA (ps)	$\sim 5\text{--}10\%$	50–100 J	up to ~ 1 kHz (joule-class)
Yb-fiber (1024 nm)	$25\text{--}30\%$	n/a at 5 J	kHz (mJ-class only)

comparable per-shot energy already exist — the CLPU Vega II facility (6 J, 30 fs, rated to 10 Hz) is one example — and can deliver the per-shot energy required to test the configuration — though, as noted in Section 5.5, its 30 fs pulse overshoots the seed required at the $300\text{ }\mu\text{m}$ spot by about an order of magnitude and would need the spot defocused accordingly — independent of their wall-plug efficiency, which is irrelevant to a single validation campaign. Deployment economics, by contrast, are set by the driver class chosen for a fielded system, where the joule-energy picosecond cryogenic Yb:YAG CPA regime at $\sim 5\text{--}10\%$ wall-plug efficiency is the relevant class. Three engineering quantities determine the deployment accessibility of any laser-driven fusion configuration: the driver wall-plug efficiency (electrical-power-to-coupled-laser-energy conversion ratio), the per-shot coupled energy, and the achievable repetition rate. The product of the three sets the time-averaged wall-plug power demand per unit of laser-coupled energy, and ultimately the per-alpha wall-plug-energy cost of any application built on the system. Published wall-plug efficiencies for laser classes used in the $p\text{-}^{11}\text{B}$ literature span three orders of magnitude, as summarized in Table 4:

The compounded engineering advantage of operating in the joule-energy picosecond cryogenic CPA regime, compared to the petawatt-class $p\text{-}^{11}\text{B}$ record of Bonvalet et al. [7], decomposes into two independent components. The hardware-and-scale component is robust

to spec-sheet verification and is independent of any simulation result. The per-shot coupled energy factor is $500 \text{ J} / 5 \text{ J} = 100$, the reconnection-driven configuration being driven at one one-hundredth of the published-record laser energy. The wall-plug efficiency factor is ~ 7 percent / 0.1 percent ≈ 70 (picosecond cryogenic CPA relative to Ti:sapphire PW), giving a further ~ 70 -fold reduction in the electrical-power input required to deliver each shot. The combined hardware-and-scale wall-plug-energy-per-shot is therefore approximately $50\text{--}100 \text{ J}$ vs 500 kJ — a factor of $\sim 7,000\times$ (order 10^4) advantage per shot before any consideration of per-shot alpha yield.

Between this hardware-and-scale lower bound and the yield-dependent upper bound lies a third figure that is robust to the simulation’s numerical limitations: the absolute wall-plug energy gain. The σ -weighted first-transit fusion energy is $G_{\text{FT}} \times 5 \text{ J} = 18.15 \text{ J}$, while the wall-plug electrical demand of a 5 J shot on a $\sim 5\text{--}10\%$ efficient picosecond cryogenic CPA driver is $50\text{--}100 \text{ J}$ (wall-plug table above). The wall-plug energy gain is therefore $Q_{\text{wp}} = 18.15 / (50\text{--}100) \approx 0.2\text{--}0.4$. Unlike the per-alpha advantage, this quantity depends only on G_{FT} — the σ -weighted gain, which de-weights the numerically heated high-energy tail (Section 5.4) and is independent of the absolute particle normalization — and on published driver efficiencies; it carries no dependence on the absolute alpha count. The deployment case therefore rests on three figures of increasing simulation dependence: a hardware-and-scale per-shot advantage of $\sim 7,000\times$ that uses no simulation output at all; a wall-plug energy gain $Q_{\text{wp}} \approx 0.2\text{--}0.4$ that depends only on the robust σ -weighted gain; and a projected per-alpha advantage approaching $10^7\times$ (below) that is conditional on experimental validation of the simulated absolute yield.

The per-shot alpha-yield component is the simulation-predicted factor of $\sim 4000\times$ (3.9×10^{13} vs $\sim 10^{10}$), which couples to the hardware-and-scale component multiplicatively to give a projected wall-plug-energy-per-alpha advantage approaching $10^7\times$ at the headline simulation yield. We emphasize that this figure compounds a simulation-predicted per-shot alpha yield with an experimentally measured yield from the Bonvalet campaign; the actual advantage will be established by experimental validation of the simulated yield at any of the PROBONO partner facilities (Section 5.5), and will sit between the $\sim 7,000\times$ hardware-and-scale lower bound (which is robust to spec-sheet verification independent of the simulation) and the $\sim 10^7\times$ upper bound (which compounds the simulation-predicted yield at face value): hardware-and-scale wall-plug advantage $\sim 7,000\times$ per shot (robust);

projected wall-plug-energy-per-alpha advantage $\sim 10^5 \times$ to $\sim 10^7 \times$ (simulation-conditional).

The order-of-magnitude character of the advantage is preserved under all reasonable interpretations of these caveats; the precise numerical figure depends on the experimental validation of the simulation yield. The repetition-rate factor compounds further: at ~ 1 kHz — the repetition rate demonstrated for joule-energy cryogenic Yb:YAG CPA drivers — the time-averaged alpha output would be approximately 3.9×10^{16} alphas/s at the first-transit simulation yield (8.8×10^{15} alphas/s at the Maxwellian thermal-only level), compared to the Bonvalet-class time-averaged output of approximately 3×10^6 alphas/s at one shot per hour. Wall-plug efficiency for the picosecond-CPA class is lower than the diode-pumped solid-state figure first considered ($\sim 5\text{--}10\%$ rather than $10\text{--}20\%$), but the kilohertz repetition rate raises the time-averaged alpha output by $\sim 100 \times$ over the 10 Hz case while the per-alpha wall-plug cost rises only $\sim 2 \times$, remaining orders of magnitude below petawatt-class campaigns. The kilohertz figure is anchored to a demonstrated driver capability rather than to rep-rated operation in the ring geometry, which remains an engineering goal. We note the further caveats: target replenishment at this repetition rate is not the limiting factor: spooled-tape targetry has been demonstrated in published laser-plasma experiments at multi-joule pulse energy and multi-hertz operation, including 70,000-shot continuous runs at the Rutherford Appleton Laboratory [40] and kilo-shot bursts at the CLPU PROBONO node [41]. These demonstrations are single-spot rather than eight-spot-ring; the genuine open engineering item is therefore not target supply but the extension of rep-rated tape targetry to the eight-spot ring geometry at kilohertz repetition, together with management of the ~ 5 kW time-averaged drive power that 5 J at ~ 1 kHz implies. The latter is eased by the present configuration operating well below the petawatt peak-intensity regime of the demonstrated tape systems, which reduces per-shot ablation, debris, and electromagnetic-pulse loading. The order-of-magnitude character of the advantage is preserved under all reasonable interpretations of these caveats.

C. Footprint, cost, and per-channel power projection

A facility built on the ring-reconnection architecture would require no petawatt or megajoule infrastructure: the interaction-region volume is on the order of a few cubic millimeters, the vacuum chamber requirement is modest, and the diagnostic suite is comparable to that

of a university-scale plasma laboratory. We project a hardware capital cost in the low to mid seven figures (USD), dominated by the multi-channel picosecond cryogenic Yb:YAG CPA laser system (a room-scale eight-channel laser hall with its cryogenic plant), vacuum chamber, target replenishment subsystem, and diagnostics. This figure is a projection rather than a built-system cost and excludes site, personnel, and facility-engineering costs that may dominate at scale. At ~ 1 kHz repetition rate, the average drive power is ~ 5 kW and, if the simulation G_{FT} is borne out experimentally, the fusion-power equivalent would be on the order of ~ 18 kW per single ring (Maxwellian conservative) or ~ 80 kW per ring (first-transit) — with the caveat that this extrapolation assumes shot-to-shot repeatability together with the management of multi-kilowatt average drive power and kilohertz target replenishment, none of which has yet been demonstrated. The relevant integrated reactor architecture (multi-column scaling) is the subject of a separate paper in preparation.

D. Comparison across the p- ^{11}B literature

Table 5 summarizes the comparison; we caution against direct numerical comparison across rows that mix per-steradian and volumetric yields.

E. Caveats and where this comparison oversimplifies

Several limitations are intrinsic to the comparisons above. First, the simulation G_{FT} includes only the σ -weighted yield of the simulated proton distribution; real plasmas exhibit additional loss channels (radiation transport, target hydrodynamic disassembly, edge cooling) not represented in the hybrid-PIC model. Second, the volumetric vs per-steradian yield comparison is not formally apples-to-apples and is presented as illustrative; a forthcoming experimental campaign would provide the direct comparison. Third, the projected facility cost excludes target injection rep-rate engineering and beam delivery from the master oscillator that, at scale, may dominate. Fourth, the rep-rated operation claim assumes target replenishment between shots; this is a known engineering challenge for solid-target laser fusion and is the subject of separate work. Fifth, the wall-plug-energy-per-alpha advantage figure compounds a simulation-predicted alpha yield with experimentally measured yields; sustained operation requires that all components (laser, target, vacuum, diagnostics) scale

TABLE V. Comparison across representative $p\text{--}^{11}\text{B}$ campaigns and fusion contexts. The yield row mixes per-steradian (Bonvalet, Picciotto) and volumetric (this work) values and should not be compared numerically across rows; see Section 4.1 for the per-joule normalization. The wall-plug-J-per- α row is the single most directly comparable engineering metric across the table; the this-work value is simulation-projected and pending experimental validation, while the experimental and NIF values are measured.

Property	NIF (D–T)	Bonvalet PW	Picciotto	This work
Fuel	D–T cryogenic	$p\text{--}^{11}\text{B}$	$p\text{--}^{11}\text{B}$ (B:Si)	$p\text{--}^{11}\text{B}$ (sim.)
Laser energy (coupled)	2.05 MJ	~ 500 J	~ 4 J	5 J
Driver class	Nd:glass	Ti:Sa PW	Iodine	ps cryo Yb:YAG CPA
Wall-plug efficiency	$\sim 0.5\%$	$< 0.1\%$	$\sim 1\%$	$\sim 5\text{--}10\%$
Wall-plug J per shot	~ 2 GJ	~ 500 kJ	~ 400 J	~ 70 J
Yield per shot	n/a (D–T)	$\sim 10^{10}$ α/sr	$\sim 10^9$ α/sr	3.9×10^{13} α (4π)
Wall-plug J / α (projected)	—	$\sim 5 \times 10^{-5}$	$\sim 4 \times 10^{-7}$	$\sim 2 \times 10^{-12}$ (sim.)
Repetition rate	months	$\sim$ hour	$\sim$ day	~ 1 kHz
Aneutronic?	No	Yes	Yes	Yes
Mechanism	ICF	TNSA+B	TNSA	Magnetic reconnection

together and is bounded by whichever subsystem becomes rate-limiting first. The order-of-magnitude character of the advantage is preserved under all reasonable interpretations of these caveats; the precise numerical figure is not.

V. DISCUSSION

A. Why the ring geometry helps

The ring topology converts a collection of independent two-spot reconnection events into a single system with predictable X-line locations and a geometrically guaranteed convergence point. Eight X-lines arranged on a circle of radius 2.217 mm produce eight outflow jets directed radially inward; these jets superpose at the center, where the local proton distribution receives contributions from all eight reconnection sites. The geometric convergence is what produces the spatial signature of preferential energization — the fusion-grade proton population concentrating on the geometric center as the reconnection outflows develop (Figure 5) — that distinguishes reconnection-driven acceleration from bulk thermal heating. The mechanism is not new at the level of single X-line physics; the contribution is the engineering of the convergence and its observed manifestation as a fusion-yield-enhanced non-thermal proton distribution.

Experiments in multi-site laser-driven reconnection have reported that adding interaction sites can degrade rather than enhance reconnection: in three-plume configurations, Bolaños et al. [43] observed that flux annihilation slowed relative to the canonical two-plume case, attributing the mitigation to disruption of the quadrupolar Hall-field structure where opposite-polarity Hall fields meet on the inner side of the configuration, and to accumulation of trapped outflow plasma between the sites; the companion analysis of obliquely interacting flux tubes develops the Hall-field mechanism in detail [10]. Those configurations are open arrangements of a few strongly interacting plumes, and the authors note that the trapped-plasma (“cork”) contribution is amplified by the two-dimensional geometry of their simulations and is expected to be weaker in three dimensions. The geometry studied here differs structurally: a closed, even-N ring places X-lines symmetrically at chord midpoints, and the σ -weighted gain is robust (under 2 percent variation) across pointing jitter, energy imbalance, and single-spot dropout, with reconnection completeness (seed-relative B-collapse) preserved even when an entire hemisphere of spots is removed (Section 3.8). This suggests the configuration sensitivity reported for open few-site arrangements does not transfer directly to the closed ring. We emphasize, however, that the present perturbation suite varies an eight-spot ring and does not scan down to the two- and three-site counts at

which the mitigation was reported; a direct spot-count comparison at matched parameters is the appropriate test of whether the ring escapes the mitigation regime or merely operates above it, and is identified as a target for follow-up.

B. Mechanism interpretation in the Li et al. framework

The σ -weighted enhancement factor of approximately $5\times$ reported here is consistent with the framework established by Li et al. [27] for laser-driven reconnection-fusion plasmas. In their PIC simulations, the (p, ^{11}B) reaction rate was enhanced by approximately four orders of magnitude above thermal predictions at substantially larger seed fields (>200 T) and kilojoule-class laser energies. The reduced enhancement factor in the present work is consistent with the smaller seed field (85 T vs >200 T) and the lower coupled laser energy (5 J vs kJ-class), at the level at which a comparable PIC framework would predict; the qualitative signature — a non-thermal proton distribution producing fusion yield several times larger than a Maxwellian plasma at the same equilibrium temperature — is preserved. The enhancement of approximately $5\times$ (with a conservative event-terminated lower bound of $4.4\times$ on the first-transit gain) is expected for the present configuration parameters, and represents the partial substitution of multi-X-line geometric convergence for raw seed-field magnitude as the driver of non-thermal energization.

The acceleration mechanism is consistent with direct E-field acceleration in the X-line outflow combined with geometric focusing at the convergence point. The X-line zone reaches its highest σ -weighted contribution during the inflow-heating phase (45–105 ps), and the core zone receives the focused superposition of outflows from all eight X-lines through this and the subsequent sustained-reconnection phase. The persistent oscillating quasi-equilibrium of the B-field at ~ 22 T, rather than monotonic decay, indicates that the system is processing flux in a quasi-steady fashion: lobe inflow continues to replenish the field on a timescale comparable to the natural cycle period. This is the regime that supports multi-pass Fermi-type acceleration in principle [24]; it is not directly demonstrated in the present simulation but is consistent with the cycle-averaged proton energy floor staying well above the laser-thermal floor of 12–46 eV.

C. The thermal-floor argument

The 12–46 eV laser-thermal floor calculated in Section 2.4 forecloses one class of reviewer objection: the possibility that the energization is bulk-thermal in origin rather than reconnection-driven. At 5 J coupled to a plasma volume bounded between $9 \times 10^{-8} \text{ m}^3$ and $3.6 \times 10^{-7} \text{ m}^3$ at $5 \times 10^{24} \text{ m}^{-3}$ ion density, the maximum possible bulk-thermal energy per ion is bounded at 46 eV regardless of the out-of-plane depth assumption; the observed equilibrium E_{95} of 327 keV is approximately $7,000\times$ to $27,000\times$ above this floor. The discrepancy is a direct measurement of the energy that has been transferred from laser-deposited bulk-thermal heating into the magnetic-field free energy via the Biermann battery, and subsequently into non-thermal proton kinetic energy via reconnection. The arithmetic is independent of any modeling choice in the simulation and can be performed by any reviewer who accepts the published values of the laser energy, the plasma volume, and the ion density. This is the strongest argument against a thermal interpretation of the result, and we recommend it as the entry point for reviewer engagement with the central claim.

D. Limitations

The principal limitations are: (i) 2D-XZ geometry, which captures the dominant in-plane reconnection dynamics but omits 3D effects including out-of-plane outflow tilts, kink-like instabilities of the X-lines, and finite ring height; (ii) hybrid-PIC electron treatment, which represents electrons as a quasi-neutral isothermal fluid with finite resistivity but cannot resolve electron-scale instabilities such as the lower-hybrid drift instability or anomalous resistivity; (iii) hybrid-PIC numerical heating of the bulk distribution at long integration times, which generates a high-energy proton tail above $\sim 1 \text{ MeV}$ that is not physical; the σ -weighted integrated-yield methodology adopted in this work is intentionally robust to this effect because the (p, α) cross-section dies off rapidly above 1 MeV, but quantities sensitive to the bulk distribution (the proton spectrum above 10 MeV, the alpha kinematic spectrum, and angular distributions) require full-PIC validation and are deferred as future work; (iv) the 1153 ps simulation duration, which is short compared with the hydrodynamic disassembly timescale of a typical solid target and assumes that the plasma layer remains in a regime where the hybrid model applies; (v) the operating-point specificity of the result — plasma beta of 13.7 and

85 T seed field define one point in a multi-dimensional parameter space, and related work in preparation addresses the density and field-strength dependence; (vi) the hybrid Ohm closure that supplies the self-consistent in-plane supporting current (no external current is imposed; see Section 2.1 and the archived run metadata, `external_supporting_current = none`), a standard but non-trivial modeling choice for thermal-dominated regimes whose sensitivity is governed by the plasma-resistivity and density-floor parameters and could be quantified by varying them; (vii) the event-terminated first-transit G_{FT} is window-dependent by construction and rises with resolution rather than asymptoting ($G_{\text{FT}} = 2.02, 3.63, 7.98$ at $256^2, 512^2$, and 1024^2 , as t_{cut} lengthens with resolution), so the convergence claim rests not on G_{FT} but on the fixed-window total σ -weighted gain, which is grid-converged to within ~ 12 percent across the three resolutions (Section 3.7); $G_{\text{FT}} = 3.63$ is reported as a conservative lower bound rather than a converged value; and (viii) the hybrid-PIC model carries no Coulomb collision operator (electrons are a resistive fluid; ions are collisionless kinetic macro-particles), appropriate to the simulated regime in which the ion mean-free-path exceeds the $\sigma = 300 \mu\text{m}$ spot scale by about an order of magnitude (collisionless transit); this underwrites the first-transit accounting and the negligible-secondary-reaction assumption and correspondingly restricts the model’s validity to this low-density, collisionless operating point. Concrete extensions that would relax these limitations are enumerated in Section 5.6.

Robustness of the two-dimensional reduction. The 2D-XZ reduction (limitation i) omits the out-of-plane dynamics — flux-rope formation, oblique tearing, and drift-kink of the X-lines — that distinguish three-dimensional reconnection from its two-dimensional idealization. The existing three-dimensional kinetic-reconnection literature nonetheless gives specific reason to expect the integrated, σ -weighted gain reported here to survive the extension to three dimensions, while identifying which features require direct verification. First, the fast-reconnection rate on which the regime classification of Section 3.2 rests is not a two-dimensional artifact: the normalized rate of order 0.1 is set by a force balance at the magnetohydrodynamic scale and is largely insensitive to the dissipation mechanism and to the upstream opening angle [34]. Second, and most directly, the time-integrated energy-conversion quantities are robust across the two-to-three-dimensional transition. In matched two- and three-dimensional kinetic layers, the three-dimensional layer develops turbulence and coherent flux ropes absent in two dimensions, yet the energy-conversion timescale remains very close to the two-dimensional value [35]; the σ -weighted first-transit gain G_{FT} is

precisely such a time-integrated energy quantity. The comparison is sharper for a geometry close to ours: in fully kinetic three-dimensional simulations of parallel flux ropes merging via the coalescence instability, the three-dimensional reconnection timescale agrees with the matched two-dimensional simulation to within a factor of about two, with the departure governed by the growth of the drift-kink mode [36]. The eight merging field islands of the present ring are the multi-site analog of this coalescence geometry.

Because the configuration reported here was constructed for numerical stability within the available compute rather than for optimal suppression of the drift-kink mode, we adopt the conservative factor-of-two bound for this work and do not claim the tighter agreement (3–8 percent) that the same study finds when the kink is suppressed by a guide field [36]. Establishing where an optimized, guide-field configuration of the ring falls within this range is identified as part of the three-dimensional follow-up (Section 5.6, item ii).

Third, three-dimensional reconnection organizes multiple, initially independent reconnection sites into a single coherent diffusion region rather than fragmenting them [37, 38] — the behavior the ring assumes when its eight X-lines feed a common convergence point, at the cost of additional secondary-flux-rope and turbulent structure not represented here. Finally, particle-in-cell simulations of laser-driven reconnection in the non-relativistic high-energy-density regime show that direct acceleration by the reconnection electric field at the X-points produces a non-thermal ion population once the ions are sufficiently magnetized [39], establishing the laser-coupled, X-line ion-acceleration mechanism invoked here as one already realized in reconnection simulations, while leaving its efficiency in the present small-system regime to the experimental and three-dimensional follow-ups of Section 5.6.

These results bound rather than eliminate the three-dimensional risk. The principal genuinely three-dimensional effect — drift-kink of the X-lines — would lengthen the reconnection timescale by at most a factor of about two in the unfavorable limit [36], and is the specific quantity the planned three-dimensional run (Section 5.6, item ii) is designed to measure. The integrated σ -weighted yield is expected to be the more robust of the reported quantities.

E. Path to experimental validation

The combination of 5 J coupled energy, mm-scale interaction region, and joule-energy picosecond cryogenic Yb:YAG CPA laser hardware places experimental validation within reach of facilities and driver systems currently operating in the EU and globally. CLPU in Salamanca, Spain is the natural primary candidate: the Vega III petawatt facility has an established record in laser-driven nuclear-reaction measurements [29], and the lower-energy Vega II system delivers 6 J / 30 fs at up to 10 Hz [29] — matching the present work’s 5 J coupled-energy budget, although its 30 fs pulse overshoots the 85 T seed at the 300 μm spot by about an order of magnitude, so it serves as a validation analog only with the focal spot defocused to recover the required intensity (a picosecond-class driver matches the seed directly). Further European laser-fusion facilities include ELI Beamlines (Margarone, Giuffrida and co-workers [19, 20, 29]) which hosts complementary petawatt and high-repetition-rate beamlines, PALS [5, 6], LNS-INFN Catania (Bonasera–Catania group [11–13, 29]), CELIA Bordeaux, and Queen’s University Belfast, all of which are co-authors on recent PROBONO experimental campaigns and share the methodological infrastructure required to replicate or extend the present work. HB11 Energy [29] is a further industrial partner within the PROBONO collaboration with direct interest in laser-driven $p\text{--}^{11}\text{B}$ fusion; their non-thermal block-ignition roadmap is reviewed in Hora et al. [8]. The minimum experimental signature would be (a) the eight X-line locations themselves, observable via proton or alpha emission patterns; (b) a time-of-flight spectrum showing a non-thermal high-energy tail above 500 keV in the converging-flow direction; (c) a measured alpha yield in the range projected by Table 5 normalized for facility-specific coupling efficiency.

F. Compute-gated extensions and future work

Several extensions of this work are tractable within the present code base (WarpX/AMReX 26.04) and have been gated by available compute access rather than by methodology. We list them here as concrete invitations for collaboration.

- (i) Full-PIC validation of the bulk proton distribution. The σ -weighted G_{FT} and the integrated alpha yield reported in this work are robust to hybrid-PIC numerical heating because the (p,α) cross-section dies off above 1 MeV. The bulk proton distribution

above 1 MeV — together with the alpha kinematic spectrum and angular distributions — is sensitive to this effect and requires full-PIC cross-validation. A reduced-scope full-PIC pilot on a subset of the geometry over the reconnection-event window (0–200 ps) is the natural test.

- (ii) Three-dimensional hybrid-PIC. A fully 3D simulation of the LD ring at the same plasma parameters would resolve out-of-plane outflow tilts, X-line kink instability, and finite ring-height effects that the 2D-XZ reduction omits. A reduced-scope 3D pilot at $256 \times 32 \times 256$ over ~ 200 ps is single-GPU tractable and would directly bound these effects.
- (iii) Fine-cadence $f(E,t)$ campaign. The present 15 ps cadence is at approximately one half of the Nyquist frequency for the natural reconnection cycle. A 3.7 ps cadence over the reconnection-event window (0–200 ps) would resolve the cycle directly and enable cycle-resolved $f(E,t)$ histograms. The σ -weighted G_{FT} is not expected to change materially (it is integrated and cadence-robust), but the cycle-resolved diagnostics would support the mechanism-interpretation discussion of Section 5.2.
- (iv) Density and B-seed parameter sweeps. The present LD configuration represents one point on a $(n, B\text{-seed})$ plane. A controlled sweep mapping the dependence of G_{FT} and the enhancement factor across the $(n, B\text{-seed})$ plane, including extensions to higher seed fields where the Li et al. limit becomes relevant, is the natural successor study.
- (v) Cross-code validation. Cross-validation of the σ -weighted G_{FT} methodology against an independent PIC code (OSIRIS, EPOCH, LSP, or SMILEI [26]) would close the simulation-methodology question and is identified as a community deliverable; we invite collaboration with groups operating these codes.

Each item above is a tractable problem given the current code base and existing algorithms; compute access has been the binding constraint. We invite collaboration on any of these extensions. The simulation input files, analysis pipeline, and derived datasets are openly archived (see Data Availability) to lower the barrier to independent replication and extension.

TABLE VI. Driver-relevant seed-field scan (256^2 , $p^{-11}\text{B}$, base density $5 \times 10^{24} \text{ m}^{-3}$). Only the seed field is varied; each seed is associated with a representative driver class through the order-of-magnitude intensity mapping of Section 2.8. The coupled-energy column is deliberately non-monotonic in seed: the seed tracks peak intensity, not pulse energy. “Final B-collapse” is the change in lobe $|B|_{\text{max}}$ relative to each run’s own seed value, evaluated at end-of-window (a scan-specific convention; the convergence results of Section 3.7 instead evaluate the seed-relative collapse against the fixed 85 T reference at the equilibrium minimum, so the 85 T entry here need not equal the 256^2 value in Table 2); positive denotes flux annihilation (reconnection), negative denotes field growth (drive-dominated amplification, no reconnection). Configurations at or above 85 T reconnect; the two below do not, placing the reconnection threshold between 51 and 85 T. *The 361 T case reconnects but lies at $\beta = 0.76$, where 256^2 numerical heating makes its absolute gain a scaling check rather than evidence. Per-laser gains are not tabulated: they are window-dependent (Section 3.7) and not comparable as raw values.

Driver class	Coupled E	Seed (T)	β	Final B-collapse	Reconnects?
Q-switched Nd:YAG	$\sim 1 \text{ J}$	45	48.9	−301% (grew)	No
Yb-fiber MOPA	$\sim 0.2 \text{ J}$	51	38.1	−93% (grew)	No
Hybrid	5 J	85	13.7	+72%	Yes
Yb:fiber/Nd:YAG (ref.)					
Hybrid	10 J	120	6.9	+68%	Yes
Yb:fiber/Nd:YAG					
ps Nd:YAG	$\sim 1 \text{ J}$	361	0.76	+68%	Yes*

VI. CONCLUSIONS

A static 8-spot ring magnetic reconnection geometry, seeded with 5 J of laser energy coupled to an idealized hydrogenated boron–nitrogen composition at base ion density $5 \times 10^{24} \text{ m}^{-3}$, produces a non-thermal proton distribution whose σ -weighted total fusion gain over the common 0–1153 ps window of $G_{\text{tot}} = 4.08$ at 512^2 — grid-converged across 256^2 – 1024^2 (4.70, 4.08, 4.35; within ~ 12 percent) — a factor of $5.0\times$ above a Maxwell-tail fast-fraction estimator (0.82) applied to the same non-thermal distribution. A conservative event-terminated first-transit lower bound on the same run is $G_{\text{FT}} = 3.63$. The yield decomposes into a

direct reconnection-event term (~ 23 percent) and a non-thermal relaxation term (~ 77 percent), grid-converged to within one percentage point. The total alpha yield is bracketed by 8.8×10^{12} alphas per shot (Maxwellian estimator) and 4.4×10^{13} alphas per shot (total σ -weighted). A simple thermodynamic calculation bounds the laser-thermal heating contribution at 12–46 eV per ion, three to four orders of magnitude below the observed equilibrium proton temperature of 327 keV, establishing that the bulk plasma kinetic energy content above ~ 50 eV is reconnection-driven by definition; the thermal-floor-implied baseline is zero.

Grid-resolution comparison across matched 256^2 , 512^2 , and 1024^2 ultrafine runs shows the total σ -weighted gain grid-converged to within ~ 12 percent (4.70, 4.08, 4.35), the direct/relaxation decomposition stable to within one percentage point, the Maxwellian fast-fraction estimator grid-independent at matched duration, and the seed-relative B-collapse fraction grid-independent (~ 75 percent). The window-dependent first-transit gain is reported as a conservative lower bound rather than a convergence observable. The reconnection regime is Hall-dominated and fast across all resolutions, with sustained flux annihilation and B-collapse persisting over the full 1153 ps integration window.

A scan over driver-relevant seed fields locates a reconnection threshold between 51 and 85 T: seeds at or above the 85 T reference sustain the mechanism while weaker seeds fail to reconnect, placing a lower bound on the required on-target intensity that falls within the reach of laboratory-scale picosecond cryogenic driver classes.

The compounded engineering accessibility of the configuration — 5 J coupled, mm-scale interaction region, ~ 5 –10 percent wall-plug efficiency for joule-energy picosecond cryogenic Yb:YAG CPA laser systems, and a ~ 1 kHz repetition rate demonstrated for that driver class — gives a hardware-and-scale wall-plug-energy-per-shot advantage of approximately $7,000\times$ over petawatt-class $p\text{--}^{11}\text{B}$ campaigns, robust to spec-sheet verification independent of the simulation. Compounding this with the simulation-predicted per-shot alpha yield gives a projected wall-plug-energy-per-alpha advantage approaching $10^7\times$ at the headline yield, with the actual figure to be established by experimental validation of the simulation prediction at any of the PROBONO partner facilities (Section 5.5). The σ -weighted integrated quantities reported here (G_{FT} , total alpha yield, reconnection-event timing) are robust to known hybrid-PIC numerical limitations; the kinematic alpha spectrum and the proton spectrum above 10 MeV are deferred quantities awaiting full-PIC validation. The compute-gated extensions of Section 5.6 are tractable within the current code base and

offered as concrete invitations for collaboration.

ACKNOWLEDGMENTS

The author thanks the WarpX and AMReX development teams for the open-source code base on which this work is built. This work was funded personally by the author. No external funding was received and no conflicts of interest are reported.

DATA AVAILABILITY

The derived datasets supporting this work are openly archived on Zenodo at DOI 10.5281/zenodo.20528212⁴⁴. The deposit spans the full grid-resolution ladder reported in Section 3.7 and Table 2 — the 256², 512², and 1024² runs, together with the reduced-macroparticle 1024² convergence variant — and provides, for each run, the first-transit and zone-resolved fusion-rate time series, per-X-line reconnection-rate diagnostics, run metadata, and the associated analysis reports, with per-file checksums. The analysis pipeline and simulation input scripts are archived at DOI 10.5281/zenodo.20540413⁴⁵, and the containerized software environment used to produce the results is archived at DOI 10.5281/zenodo.20541777⁴⁶; the three records are cross-referenced. Per-zone proton energy distribution histograms are included for the 256² and 512² runs, and are deposited with a documented caveat: the absolute particle normalization is unreliable owing to hybrid-PIC weighting, whereas the σ -weighted and percentile-based quantities (G_{FT} , E_{95}) on which all reported results depend are robust to this effect (Sections 2.5, 3.4, and 5.4).

REFERENCES

- ¹T. H. Rider, “Fundamental limitations on plasma fusion systems not in thermodynamic equilibrium,” PhD thesis, MIT (1995).
- ²W. M. Nevins and R. Swain, “The thermonuclear fusion rate coefficient for p–¹¹B reactions,” Nucl. Fusion 40, 865 (2000).
- ³V. S. Belyaev, A. P. Matafonov, V. I. Vinogradov, V. P. Krainov, V. S. Lisitsa, A. S. Roussetski, G. N. Ignatyev, and V. P. Andrianov, “Observation of neutronless fusion reactions in picosecond laser plasmas,” Phys. Rev. E 72, 026406 (2005).

- ⁴C. Labaune, C. Baccou, S. Depierreux, C. Goyon, G. Loisel, V. Yahia, and J. Rafelski, “Fusion reactions initiated by laser-accelerated particle beams in a laser-produced plasma,” *Nat. Commun.* 4, 2506 (2013).
- ⁵A. Picciotto, D. Margarone, A. Velyhan, P. Bellutti, J. Krasa, A. Szydlowsky, G. Bertuccio, Y. Shi, A. Mangione, J. Prokupek, et al., “Boron-proton nuclear-fusion enhancement induced in boron-doped silicon targets by low-contrast pulsed laser,” *Phys. Rev. X* 4, 031030 (2014).
- ⁶L. Giuffrida, F. Belloni, D. Margarone, G. Petringa, G. Milluzzo, V. Scuderi, A. Velyhan, M. Rosinski, A. Picciotto, M. Kucharik, et al., “High-current stream of energetic alpha particles from laser-driven proton-boron fusion,” *Phys. Rev. E* 101, 013204 (2020).
- ⁷J. Bonvalet, Ph. Nicolaï, D. Raffestin, E. d’Humières, D. Batani, V. Tikhonchuk, V. Kantarelou, L. Giuffrida, M. Tosca, G. Korn, et al., “Energetic alpha-particle sources produced through proton-boron reactions by high-energy high-intensity laser beams,” *Phys. Rev. E* 103, 053202 (2021).
- ⁸H. Hora, S. Eliezer, G. J. Kirchhoff, N. Nissim, J. X. Wang, P. Lalouis, Y. X. Xu, G. H. Miley, J. M. Martinez-Val, W. McKenzie, et al., “Road map to clean energy using laser beam ignition of boron-hydrogen fusion,” *Laser Part. Beams* 35, 730 (2017).
- ⁹A. S. Joglekar, A. G. R. Thomas, W. Fox, and A. Bhattacharjee, “Magnetic reconnection in plasma under inertial confinement fusion conditions driven by heat flux effects in Ohm’s law,” *Phys. Rev. Lett.* 112, 105004 (2014).
- ¹⁰S. Bolaños, A. Sladkov, R. Smets, S. N. Chen, A. Grisollet, E. Filippov, J.-L. Henares, V. Nastasa, S. Pikuz, R. Riquier, et al., “Laboratory evidence of magnetic reconnection hampered in obliquely interacting flux tubes,” *Nat. Commun.* 13, 6426 (2022).
- ¹¹M. Barbui, W. Bang, A. Bonasera, K. Hagel, K. Schmidt, J. B. Natowitz, R. Burch, G. Giuliani, M. Barbarino, H. Zheng, et al., “Measurement of the plasma astrophysical S factor for the $^3\text{He}(d,p)^4\text{He}$ reaction in exploding molecular clusters,” *Phys. Rev. Lett.* 111, 082502 (2013).
- ¹²D. Lattuada, M. Barbarino, A. Bonasera, W. Bang, H. J. Quevedo, M. Warren, F. Consoli, R. De Angelis, P. Andreoli, S. Kimura, et al., “Model-independent determination of the astrophysical S-factor in laser-induced fusion plasmas,” *Phys. Rev. C* 93, 045808 (2016).
- ¹³S. Kimura and A. Bonasera, “Chaos Driven Fusion Enhancement Factor at Astrophysical Energies,” *Phys. Rev. Lett.* 93, 262502 (2004).

- ¹⁴J.-L. Vay, A. Huebl, A. Almgren, L. D. Amorim, J. Bell, L. Fedeli, L. Ge, K. Gott, D. P. Grote, M. Hogan, et al., “Modeling of a chain of three plasma accelerator stages with the WarpX electromagnetic PIC code on GPUs,” *Phys. Plasmas* 28, 023105 (2021).
- ¹⁵M. Yamada, R. Kulsrud, and H. Ji, “Magnetic reconnection,” *Rev. Mod. Phys.* 82, 603 (2010).
- ¹⁶H. Ji and W. Daughton, “Phase diagram for magnetic reconnection in heliophysical, astrophysical, and laboratory plasmas,” *Phys. Plasmas* 18, 111207 (2011).
- ¹⁷H. Abu-Shawareb, R. Acree, P. Adams, J. Adams, B. Addis, R. Aden, P. Adrian, B. B. Afeyan, M. Aggleton, L. Aghaian, et al. (Indirect Drive ICF Collaboration), “Achievement of target gain larger than unity in an inertial fusion experiment,” *Phys. Rev. Lett.* 132, 065102 (2024).
- ¹⁸S. Kimura, A. Anzalone, and A. Bonasera, “Comment on ‘Observation of neutronless fusion reactions in picosecond laser plasmas’,” *Phys. Rev. E* 79, 038401 (2009).
- ¹⁹D. Margarone, A. Morace, J. Bonvalet, Y. Abe, V. Kantarelou, D. Raffestin, L. Giuffrida, P. Nicolai, M. Tosca, A. Picciotto, et al., “Generation of α -particle beams with a multi-kJ, peta-watt class laser system,” *Front. Phys.* 8, 343 (2020).
- ²⁰D. Margarone, J. Bonvalet, L. Giuffrida, A. Morace, V. Kantarelou, M. Tosca, D. Raffestin, P. Nicolai, A. Picciotto, Y. Abe, et al., “In-target proton–boron nuclear fusion using a PW-class laser,” *Appl. Sci.* 12, 1444 (2022).
- ²¹M. S. Schollmeier, V. Shirvanyan, C. Capper, S. Steinke, A. Higginson, R. Hollinger, J. T. Morrison, R. Nedbailo, H. Song, S. Wang, et al., “Investigation of proton beam-driven fusion reactions generated by an ultra-short petawatt-scale laser pulse,” *Laser Part. Beams* 2022, 2404263 (2022).
- ²²W. Fox, D. B. Schaeffer, M. J. Rosenberg, G. Fiksel, J. Matteucci, H.-S. Park, A. F. A. Bott, K. Lezhnin, A. Bhattacharjee, D. Kalantar, et al., “Fast magnetic reconnection in highly-extended current sheets at the National Ignition Facility,” *arXiv:2003.06351* (2020).
- ²³V. Istokskaia, M. Tosca, L. Giuffrida, J. Psikal, F. Grepl, V. Kantarelou, S. Stancek, S. Di Siena, A. Hadjikyriacou, A. McIlvenny, et al., “A multi-MeV alpha particle source via proton-boron fusion driven by a 10-GW tabletop laser,” *Commun. Phys.* 6, 27 (2023).
- ²⁴F. Belloni, “On a fusion chain reaction via suprathermal ions in high-density H-¹¹B plasma,” *Plasma Phys. Control. Fusion* 63, 055020 (2021).

- ²⁵M. H. Sikora and H. R. Weller, “A new evaluation of the $^{11}\text{B}(\text{p},\alpha)\alpha$ reaction rates,” *J. Fusion Energy* 35, 538–543 (2016).
- ²⁶J. Derouillat, A. Beck, F. Pérez, T. Vinci, M. Chiamello, A. Grassi, M. Flé, G. Bouchard, I. Plotnikov, N. Aunai, et al., “SMILEI: A collaborative, open-source, multi-purpose particle-in-cell code,” *Comput. Phys. Commun.* 222, 351 (2018).
- ²⁷K. Li, Z. Y. Liu, A. D. Liu, S. P. Zhu, X. T. He, and B. Qiao, “A nuclear-reaction-based method for probing the nonthermal ion energy spectrum in high energy density laboratory plasmas,” *Nucl. Fusion* 63, 126021 (2023).
- ²⁸PROBONO COST Action CA21128 — “Proton Boron Nuclear Fusion: From Energy Production to Medical Applications,” www.cost.eu/actions/CA21128 (2021–2025).
- ²⁹M. R. D. Rodrigues, A. Bonasera, M. Scisciò, J. A. Pérez-Hernández, M. Ehret, F. Filippi, P. L. Andreoli, M. Huault, H. Larreur, D. Singappuli, et al., “Radioisotope production using lasers: From basic science to applications,” *Matter Radiat. Extremes* 9, 037203 (2024).
- ³⁰S. Kimura and A. Bonasera, “Deuteron-induced reactions generated by intense lasers for PET isotope production,” *Nucl. Instrum. Methods Phys. Res., Sect. A* 637, 164–170 (2011).
- ³¹C. Labaune, C. Baccou, V. Yahia, C. Neuville, and J. Rafelski, “Laser-initiated primary and secondary nuclear reactions in boron-nitride,” *Sci. Rep.* 6, 21202 (2016).
- ³²D. S. Hey, M. E. Foord, M. H. Key, S. L. Le Pape, A. J. Mackinnon, P. K. Patel, Y. Ping, K. U. Akli, R. B. Stephens, T. Bartal, et al., “Laser-accelerated proton conversion efficiency thickness scaling,” *Phys. Plasmas* 16, 123108 (2009).
- ³³R. E. Groenewald, A. Veksler, F. Ceccherini, A. Necas, B. S. Nicks, D. C. Barnes, T. Tajima, and S. A. Dettrick, “Accelerated kinetic model for global macro stability studies of high-beta fusion reactors,” *Phys. Plasmas* 30, 122508 (2023).
- ³⁴Y.-H. Liu, M. Hesse, F. Guo, W. Daughton, H. Li, P. A. Cassak, and M. A. Shay, “Why does steady-state magnetic reconnection have a maximum local rate of order 0.1?,” *Phys. Rev. Lett.* 118, 085101 (2017).
- ³⁵W. Daughton, T. K. M. Nakamura, H. Karimabadi, V. Roytershteyn, and B. Loring, “Computing the reconnection rate in turbulent kinetic layers by using electron mixing to identify topology,” *Phys. Plasmas* 21, 052307 (2014).
- ³⁶J. P. Sauppe and W. Daughton, “Kinetic-scale flux rope reconnection in periodic and line-tied geometries,” *Phys. Plasmas* 25, 012901 (2018).

- ³⁷L. Yin, W. Daughton, H. Karimabadi, B. J. Albright, K. J. Bowers, and J. Margulies, “Three-dimensional dynamics of collisionless magnetic reconnection in large-scale pair plasmas,” *Phys. Rev. Lett.* 101, 125001 (2008).
- ³⁸W. Daughton, V. Roytershteyn, H. Karimabadi, L. Yin, B. J. Albright, B. Bergen, and K. J. Bowers, “Role of electron physics in the development of turbulent magnetic reconnection in collisionless plasmas,” *Nat. Phys.* 7, 539 (2011).
- ³⁹S. R. Titorica, M. Hoshino, T. Abel, and F. Fiuza, “Nonthermal electron and ion acceleration by magnetic reconnection in large laser-driven plasmas,” *Phys. Plasmas* 27, 112111 (2020).
- ⁴⁰N. Xu, M. J. V. Streeter, O. C. Ettlinger, H. Ahmed, S. Astbury, M. Borghesi, N. Bourgeois, C. B. Curry, S. J. D. Dann, N. P. Dover, et al., “Versatile tape-drive target for high-repetition-rate laser-driven proton acceleration,” *High Power Laser Sci. Eng.* 11, e23 (2023).
- ⁴¹M. Ehret, D. de Luis, J. I. Apiñaniz, J. L. Henares, R. Lera, J. A. Pérez-Hernández, P. Puyuelo-Valdes, L. Volpe, and G. Gatti, “Stability and debris-mitigation of a solid tape target delivery system for intense laser-matter interactions towards high-repetition-rate,” *Plasma Phys. Control. Fusion* 66, 045003 (2024).
- ⁴²E. G. Harris, “On a plasma sheath separating regions of oppositely directed magnetic field,” *Nuovo Cimento* 23, 115 (1962).
- ⁴³S. Bolaños, R. Smets, C. Courtois, N. Blanchot, G. Boutoux, W. Cayzac, S. N. Chen, V. Denis, A. Grisollet, I. Lantuejoul, et al., “Experimental mitigation of fast magnetic reconnection in multiple interacting laser-produced plasmas,” *arXiv:2305.05701* (2023).
- ⁴⁴J. B. Worth (2026). “Simulation results for reconnection-driven non-thermal proton–boron-11 fusion at joule-class laser energy (Paper 1): first-transit and zone-resolved fusion-rate time series, reconnection-rate diagnostics, and analysis reports,” Zenodo. <https://doi.org/10.5281/zenodo.20528212>
- ⁴⁵J. B. Worth (2026). “Ring-geometry magnetic reconnection laser-fusion suite: hybrid-PIC simulation engine, σ -weighted first-transit fusion accounting, and sensitivity reproduction set,” Zenodo. <https://doi.org/10.5281/zenodo.20540413>
- ⁴⁶J. B. Worth (2026). “Reproducibility container for ring-geometry magnetic-reconnection laser-fusion simulations (CUDA-enabled WarpX environment),” Zenodo. <https://doi.org/10.5281/zenodo.20541777>