Practical Interplanetary Propulsion via One-Plug Collimated Nonthermal p-^{11}B Fusion Exhaust:

A Concept Paper on Proton-Boron (p-^{11}B) Propulsion

for Near-Term Research Prioritization

Tyler Martin

Lobster Labs

Madrid, Spain

April 2026

Abstract

This concept paper identifies a performance regime for fusion propulsion that has not been systematically explored: direct magnetic collimation of non-thermal proton-boron (p-^{11}B) fusion products, exploiting the asymmetric entropy structure of the three-alpha final state. We argue that p-^{11}B products occupy a region of phase space where entropy is concentrated primarily in angular distribution, while momentum magnitude remains narrowly bounded at the MeV scale. Under Liouville's theorem, magnetic nozzle geometries can trade angular dispersion for spatial dispersion without reducing phase space volume. Because propulsion depends only on axial momentum and is indifferent to the spatial spread of the exhaust, this trade is favorable in principle.

We revise the reference architecture to a one-plug open magnetic nozzle: the p-^{11}B alpha source is placed just aft of a high-field magnetic plug and at the inception of a monotonic aft-diverging nozzle. Aft-born products enter the nozzle directly; forward-born products outside the plug loss cone are mirrored aft; the small loss-cone leakage is intercepted by a localized alpha direct-conversion assembly for auxiliary electrical power. This topology avoids the two-mirror trapped-particle population that arises if products are born between a forward mirror and an aft nozzle throat.

Adiabatic analysis indicates that collimation is physically effective at aft mirror ratios up to approximately 10, where particles undergo non-adiabatic detachment with most perpendicular energy already converted to parallel motion. We introduce a total propulsive energy fraction ψ that includes capture, nozzle collimation, leakage, thermalization, neutralization, radiation, and recirculation losses. For a 10-day Earth-Jupiter brachistochrone transfer, ψ values of 0.7, 0.4, and 0.3 correspond to estimated propellant mass fractions of approximately 30%, 37%, and 42%, respectively. These estimates define a target design space, not a demonstrated performance level. The paper does not propose or validate a p-^{11}B ignition scheme; it asks whether, given a source of nonthermal charged fusion products, a magnetic exhaust architecture can convert them into useful directed momentum.

Keywords: *proton-boron fusion, aneutronic propulsion, phase space, Liouville's theorem, magnetic nozzle, adiabatic invariant, directed exhaust, interplanetary propulsion, autonomous research*

1. Introduction

1.1 The Propulsion Gap

The distances between the inner planets are not large by the standards of nuclear energy. Earth to Mars at opposition is roughly 0.5 AU; Earth to Jupiter is roughly 5 AU. Fusion reactions release approximately six orders of magnitude more energy per unit mass than chemical combustion. In principle, fusion-powered spacecraft could reduce interplanetary transit times from months to days, transforming space transport from expedition logistics to routine operations.

No fusion propulsion system has been built or tested at any scale. The reasons are partly engineering (plasma confinement, ignition thresholds, materials) and partly thermodynamic. This paper addresses the thermodynamic obstacle, which has received less systematic attention.

1.2 Why This Matters Now

Several companies are actively pursuing p-^{11}B fusion: TAE Technologies (field-reversed configurations), HB11 Energy (laser-driven approaches), and ENN (spherical torus). No existing system has demonstrated net-energy p-^{11}B operation, and ignition remains the dominant unresolved problem. The question this paper addresses is therefore narrower: assuming a source of nonthermal p-^{11}B alpha products exists, what performance regime does direct charged-product exhaust open for propulsion, and is that regime worth targeting?

The answer matters for infrastructure planning. If p-^{11}B directed exhaust can deliver the performance estimated here, the downstream engineering (magnetic nozzle design, thermal management, fuel handling) should be studied in parallel with ignition research, not deferred until ignition is demonstrated. The exhaust architecture and the fusion source are partially separable: exhaust physics can be studied first, but source-nozzle integration must eventually be co-designed.

1.3 The Strategic Context

Multiple actors are building toward sustained human presence beyond Earth orbit. SpaceX's Starship architecture targets launch costs enabling large-scale orbital and cislunar construction. Proposals for lunar and Martian industrial operations are advancing from concept to early engineering. The missing element is fast, efficient interplanetary transport. Chemical propulsion can reach Mars, but transit times (6–9 months) and mass ratios (>90% propellant) constrain mission architecture to the point where permanent industrial presence requires enormous logistical overhead.

A propulsion system delivering continuous acceleration at ~0.05–0.4g with propellant fractions below ~40% would change this calculus. Transit times to Mars drop to days or weeks. Jupiter becomes reachable in weeks. The asteroid belt, Jovian moons, and outer solar system become accessible for industrial operations.

The propellant fractions estimated in this paper (roughly 30–42% across the main sensitivity band) are compatible with routine, reusable interplanetary transport in a way that conventional thermal fusion concepts (60–97% propellant) are not. At these

mass fractions, permanent bases, mining operations, and industrial facilities across the solar system become logistically supportable rather than requiring expedition-scale commitment for every resupply.

1.4 Scope and Limitations

This is a concept paper, not a design study. It identifies a performance regime, argues that the exhaust physics permits it, estimates its parametric sensitivity, and specifies what further work is needed. It does not present transport simulations, particle-tracing results, engineering designs, or an ignition solution. Quantitative estimates should be understood as defining a target design space, not predicting system performance.

The central claim is narrow: p-^{11}B fusion products have an entropy structure qualitatively more favorable for direct magnetic collimation than thermal plasmas, and this has not been properly exploited or clearly articulated in the existing literature.

2. Theoretical Framework

2.1 Liouville's Theorem and Propulsion

A system of N particles in three dimensions occupies a volume in 6N-dimensional phase space. Liouville's theorem states that the phase space density is constant along Hamiltonian trajectories:

$$d\rho/dt = \partial\rho/\partial t + \{\rho, H\} = 0 \quad (1)$$

The volume occupied by a collection of particles in phase space is invariant under lossless dynamics. The standard conclusion is that fusion exhaust cannot be collimated without dissipating entropy or discarding particles.

This conclusion is correct about full phase space volume. But propulsion does not require compression of all phase space dimensions. It requires only that momentum vectors align along the thrust axis. The spatial distribution of the exhaust is irrelevant to thrust. This distinction is routine in accelerator physics (beam emittance trades) but has not been systematically applied to fusion propulsion.

2.2 Phase Space Structure of p-^{11}B Products

The proton-boron fusion reaction:

$$p + {}^{11}B \rightarrow 3\alpha + 8.7\ MeV \quad (2)$$

produces three alpha particles (rest mass 3727.4 MeV/c² each) sharing 8.7 MeV of kinetic energy. The reaction proceeds predominantly through an intermediate ^{8}Be* state: one alpha is emitted promptly with high energy (~4–5 MeV), and the remaining ^{8}Be* decays into two alphas at lower energy (~1.5 MeV each). The resulting energy distribution is not Gaussian but bimodal, with structure visible in the Dalitz plot [16, 17]. The mean kinetic energy per alpha is approximately 2.9 MeV, but individual alphas span a range from roughly 1.5 to 5 MeV.

This bimodal structure has a specific consequence for collimation: the high-energy prompt alpha (~5 MeV) has a larger Larmor radius than the ^{8}Be* decay alphas (~1.5 MeV) by a factor of roughly $\sqrt{(5/1.5)} \approx 1.8$. It will therefore reach the adiabatic breakdown threshold ($\varepsilon \approx 1$) earlier in the nozzle expansion, detaching at a lower

effective mirror ratio with a wider pitch angle. The exhaust is not a single population with ±25% spread; it is two populations with different collimation characteristics. The directional efficiency estimates in Section 4 should be understood as averages over this bimodal distribution, and the actual performance will depend on the relative weighting of the prompt and sequential channels. Proper quantification requires particle tracing using the measured differential cross-sections [16, 17] as initial conditions.

The mean velocity for a 2.9 MeV alpha is:

$$v_a = c\sqrt{(1 - 1/(1 + T/m_a c^2)^2)} \approx 0.039c \approx 11{,}700 \text{ km/s} \quad (3)$$

Despite the bimodal energy structure, the velocity spread remains qualitatively different from a Maxwellian plasma: the alphas occupy a bounded energy range (1.5–5 MeV) rather than the factor-of-several spread typical of thermal distributions. The angular distribution is approximately isotropic (4π steradians) in the lab frame.

The occupied phase space decomposes schematically as:

$$\Gamma \sim \Delta V \times \Delta p_r \times \Delta\Omega \quad (4)$$

where Δp_r (momentum magnitude spread) is narrow relative to the mean, while ΔΩ (solid angle) approaches its maximum. The entropy is dominated by the angular term. This is the key observation: the phase space is concentrated in the angular subspace, which is exactly the subspace magnetic fields act on.

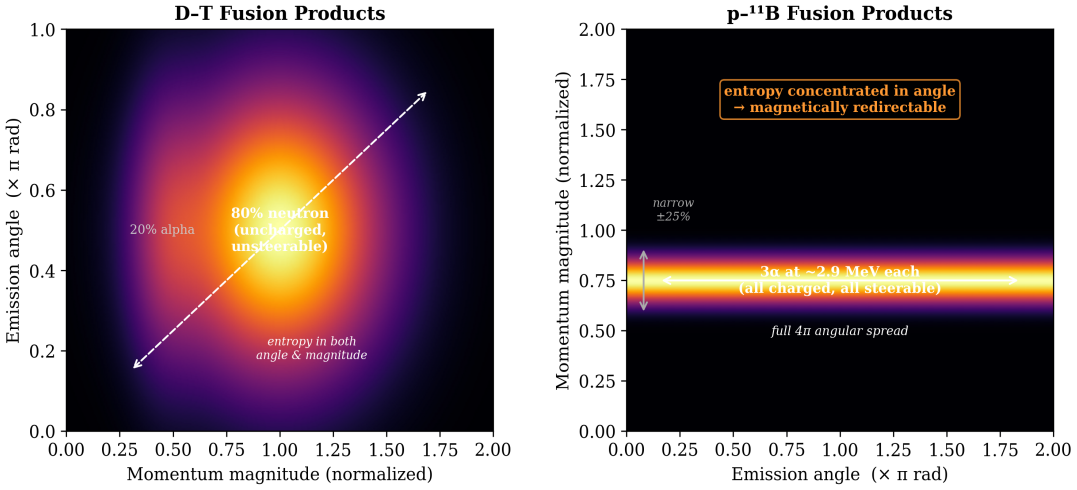

> **Figure 2.** *Phase space distributions of D–T and p-¹¹B fusion products. Left: D–T products occupy a broad region in both momentum magnitude and emission angle, with 80% of the energy in uncharged neutrons. Right: p-¹¹B products form a narrow band in momentum magnitude (±25% spread from three-body kinematics) but span the full 4π solid angle. The entropy is concentrated entirely in the angular dimension, which is the dimension magnetic fields can act on.*

2.3 The Collimation-Position Trade

A diverging magnetic field redirects charged particle trajectories through conservation of the adiabatic invariant (magnetic moment $\mu = mv_\perp^2/2B$). As particles enter weaker field regions, transverse velocity converts to parallel velocity while the spatial distribution broadens:

$$\Delta p_\perp \times \Delta x_\perp \geq constant \quad (5)$$

For propulsion this trade is favorable: thrust depends on total axial momentum flux integrated over the exhaust cross-section, and a spatially broad but well-collimated exhaust delivers the same thrust as a narrow one.

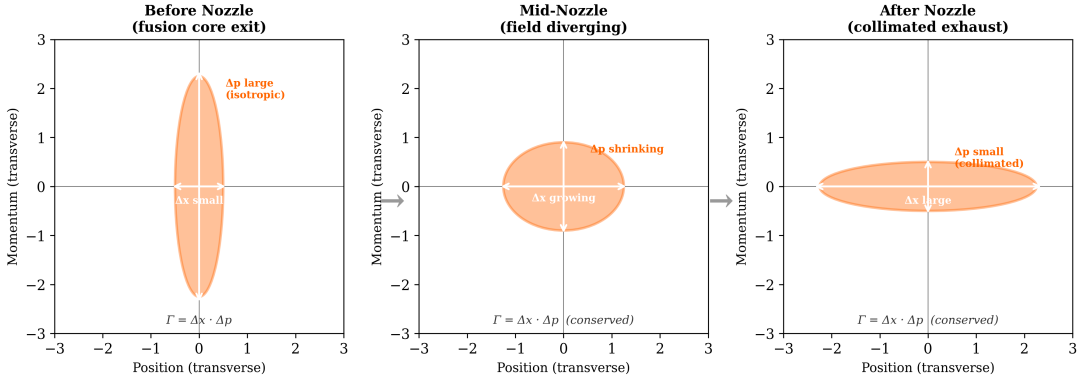

Liouville's theorem: phase space volume Γ is conserved. The nozzle reshapes it, trading angular spread for spatial spread.

> **Figure 3.** *The collimation-position trade under Liouville's theorem. The phase space volume $\Gamma = \Delta x \cdot \Delta p$ is conserved through the nozzle. At the fusion core exit (left), the distribution is narrow in position but broad in transverse momentum (isotropic emission). The diverging magnetic field reshapes the distribution (center) until, at the nozzle exit (right), the transverse momentum spread is small (collimated) and the position spread is large. The total phase space area is unchanged. Since propulsion depends only on momentum alignment, not spatial distribution, the constraint falls on an irrelevant degree of freedom.*

2.4 Adiabatic Limits and Detachment

The adiabatic invariant is conserved only when the field varies slowly over a single Larmor orbit. The adiabaticity parameter is:

$$\varepsilon = r_L / L_B \quad (6)$$

where r_L is the Larmor radius and $L_B = |B/\nabla B|$ is the field gradient scale length. Adiabaticity requires $\varepsilon \ll 1$. For a 2.9 MeV alpha (He^{2+}), the magnetic rigidity gives:

$$r_L \approx 0.24/B \text{ meters} \quad (7)$$

where B is in Tesla. At a nozzle throat field of 10 T, $r_L \approx 2.4$ cm and the particle is well-confined. As the field drops through the expansion, the Larmor radius grows. For a nozzle with gradient scale length $L_B \approx 2$ m:

At 10 T: $r_L \approx 0.024$ m, $\varepsilon \approx 0.012$. Fully adiabatic.

At 1 T: $r_L \approx 0.24$ m, $\varepsilon \approx 0.12$. Adiabaticity weakening but collimation still effective.

At 0.1 T: r_L ≈ 2.4 m, ε ≈ 1.2. Adiabatic invariant broken. Particle detaches.

This analysis yields an effective mirror ratio of approximately R ≈ 10 (from 10 T throat to ~1 T detachment region) before non-adiabatic effects dominate. Beyond this point, particles go ballistic with whatever pitch angle they have at detachment.

2.5 Detachment Physics

The detachment process has two important characteristics that refine the idealized picture:

Gyrophase dependence. At ε ≈ 1, the particle's final pitch angle depends on its gyrophase at detachment. The result is a statistical spread around the mean detachment angle, not a clean cutoff. However, the particle retains the bulk of the parallel momentum it accumulated in the strong-field region. Detachment does not reverse collimation; it freezes it at a non-ideal but useful point.

Velocity-dependent smearing. Particles born with different ratios of perpendicular to parallel velocity detach at different field strengths. Low-v⊥ particles (born nearly parallel to the axis) have small Larmor radii and remain adiabatic deeper into the expansion, achieving excellent collimation. High-v⊥ particles detach earlier at lower effective mirror ratios. The exhaust angular distribution is therefore smeared: a superposition of well-collimated and less-well-collimated populations. For an initially isotropic distribution, this smearing degrades the directional efficiency below the idealized single-particle estimate.

Both effects reduce the directional efficiency from the idealized adiabatic prediction but do not eliminate the collimation advantage. The net effect is to bound the realistic directional efficiency to approximately 0.70–0.90, compared to the idealized 0.90–0.95 from the simple mirror formula.

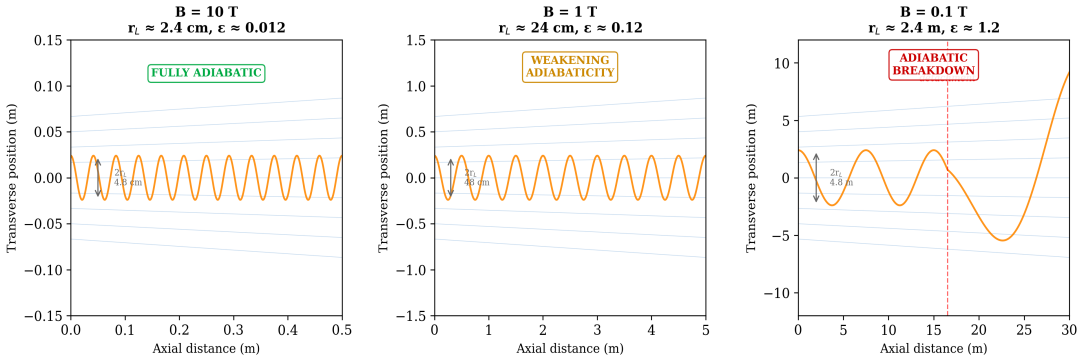

As B decreases through the nozzle expansion, the Larmor radius grows by 100×. At ε ≈ 1, the adiabatic invariant breaks and the particle detaches.

> **Figure 4.** *Adiabatic breakdown of a 2.9 MeV alpha particle trajectory at three magnetic field strengths. Left: at 10 T (nozzle throat), the Larmor radius is 2.4 cm and the particle spirals tightly along field lines (ε ≈ 0.012, fully adiabatic). Center: at 1 T, the Larmor radius grows to 24 cm and adiabaticity weakens (ε ≈ 0.12). Right: at 0.1 T, the Larmor radius reaches 2.4 m and the adiabatic invariant breaks (ε ≈ 1.2), causing the particle to detach from the field lines and go ballistic. The effective mirror ratio is bounded at R ≈ 10 by this breakdown.*

3. System Architecture (Conceptual)

3.1 Fusion Source Assumption

The fusion source must produce nonthermal p-^{11}B alpha products at propulsion-relevant power density. No existing confinement or beam-target concept has demonstrated net energy gain from p-^{11}B. Several approaches are under active development, including field-reversed configurations, laser-driven approaches, and spherical torus experiments, but this paper does not depend on any one of them.

The central assumption is narrower: a compact source produces p-^{11}B alpha products fast enough that their initial MeV-scale, nonthermal phase-space structure is available to the exhaust system before equilibration with a bulk plasma. If the products thermalize before entering the nozzle, the phase-space advantage is lost.

For an order-of-magnitude comparison, the Spitzer slowing-down time for a fast ion in a background plasma is approximately:

$$\tau_s \approx (3\sqrt{\pi}/4)(m_\alpha/m_e)^{1/2}(T_e/E_\alpha)^{3/2}\,\tau_e \quad (A)$$

For 2.9 MeV alphas in a p-^{11}B plasma at electron temperature $T_e \approx 300$ keV and density $\sim 10^{20}$ m^{-3}, the thermalization timescale is of order 10–100 μs. The transit time for a 0.039c alpha through a ~1 m source region is approximately 0.1 μs. This comparison suggests a possible two-to-three-order-of-magnitude extraction margin, but the margin is not a proof: it depends on density, path length, field topology, wave activity, and the details of the ignition scheme.

This paper therefore treats ignition and product generation as an upstream source problem. The exhaust architecture below is evaluated as a conditional system: given nonthermal alpha products at the nozzle inception region, can magnetic geometry convert their angular entropy into useful aft-directed momentum?

3.2 Magnetic Nozzle

The magnetic nozzle converts angular spread to axial alignment. The mirror force:

$$F_{\parallel} = -\mu(\partial B/\partial s) \quad (8)$$

acts on particles with transverse velocity. Under the adiabatic approximation, the limiting half-angle is:

$$\sin(\theta) \approx 1/\sqrt{R} \quad (9)$$

As established in Section 2.4, the effective mirror ratio for 2.9 MeV alphas is bounded at approximately $R \approx 10$ for nozzle gradient scale lengths of ~2 m, giving a mean half-angle of ~18° and directional efficiency of ~0.90 before accounting for smearing effects.

Two additional effects modify the single-particle picture. First, electrons remain magnetized much longer than alphas (their Larmor radius is smaller by a factor of ~85), creating an ambipolar electric field as the charge populations separate. This is the well-studied electron detachment problem in magnetic nozzle physics [23, 24]. The ambipolar field accelerates ions and decelerates electrons until they detach together, introducing corrections of order a few percent to the single-particle thrust estimates. Second, the plasma beta ($\beta = P_{kinetic}/P_{magnetic}$) varies along the nozzle. At the 10 T throat, $\beta \approx 0.06$ and the magnetic field dominates, validating the single-particle adiabatic analysis. As the field drops through the expansion, β rises

above unity in the detachment zone, consistent with the transition from magnetically guided to ballistic exhaust. The high-β regime at detachment does not invalidate the concept; it partially defines the detachment surface.

3.3 One-Plug Open Magnetic Nozzle

The reference architecture is a one-plug open magnetic nozzle. The alpha source is placed just aft of a high-field magnetic plug, at the inception of a monotonic aft-diverging field. The spacecraft and payload sit forward of the plug; the nozzle opens aft.

This field topology changes the capture problem. Aft-born alpha products immediately enter decreasing magnetic field and are collimated by the nozzle. Forward-born products move toward increasing magnetic field; all particles outside the plug loss cone are reflected aft by conservation of magnetic moment. Perpendicular-born products are no longer trapped between two high-field ends. In a monotonic aft-decreasing field, the mirror force accelerates high-$v_\perp$ particles toward weaker field, so the perpendicular population is pushed into the nozzle rather than bounced indefinitely between a forward mirror and an aft throat.

For a plug ratio $R_p = B_{plug}/B_{source}$, the forward loss-cone half-angle satisfies:

$$\sin^2\theta_c = 1/R_p \quad (9a)$$

For an isotropic alpha distribution, the fraction of all products that leak forward through the plug is:

$$f_{leak} = (1 - \cos\theta_c)/2 \quad (9b)$$

This gives approximately 2.6% leakage at $R_p = 10$, 1.3% at $R_p = 20$, and 0.5% at $R_p = 50$. The loss cone is not zero, but it is a bounded parasitic term rather than a 50% loss or a large trapped population. The leakage should be intercepted deliberately by a localized alpha interception and direct-conversion assembly located behind the plug, shielded from the main exhaust path. At propulsion power levels this subsystem is extreme: 1% of a 25 TW alpha source is 250 GW before conversion losses. It is therefore not a minor component, but it is a bounded megavolt, high-current power-conversion problem rather than an existential thrust-cancellation problem.

The one-plug architecture also avoids the double-pass thermalization problem of a two-ended mirror. Forward-born alphas that mirror from the plug do not traverse a long confinement core twice; they reverse near the high-field plug and enter the aft nozzle. The relevant dwell time is therefore set by local plug reflection and aft transit, not repeated bounce confinement.

The hard requirement shifts from recapture to source integration: the burn or alpha-production region must coexist with an open aft extraction channel and a high-field plug without flushing the reacting fuel or quenching the plasma. That problem belongs to the ignition/source architecture, not to the exhaust phase-space argument. It is nevertheless a first-order integration constraint and must be part of any follow-on design study.

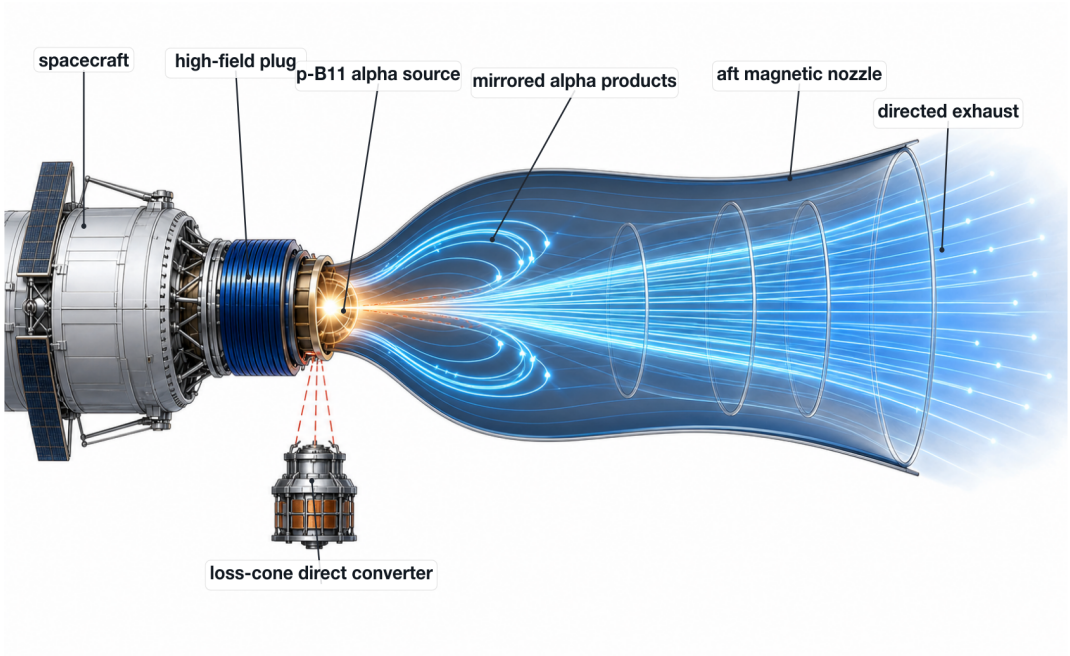

> **Figure 1.** *Conceptual one-plug p-^{11}B directed exhaust topology. The alpha source sits immediately aft of a high-field magnetic plug. Forward-born products outside the plug loss cone mirror aft into a monotonic diverging magnetic nozzle; aft-born products enter the nozzle directly. The narrow forward loss cone is intercepted by a localized alpha direct-conversion assembly.*

3.4 Contrast with Prior Work

The closest prior work in the academic literature is Tarditi's NASA NIAC Phase I study [12]. Tarditi assumed an aneutronic source and investigated beam conditioning for propulsion. His architecture deliberately thermalizes the products: alphas are injected non-adiabatically into a magnetic duct, transfer energy into gyro-motion, and heat a denser propellant expanded through a magnetic nozzle. This is a fusion-heated thermal rocket.

In the informal literature, MatterBeam's ToughSF analysis [22] is the most detailed public attempt to explain the Epstein Drive from The Expanse using real physics. That analysis uses D-He3 fuel with inertial confinement (laser-ignited pellets) and a modified VISTA architecture, achieving approximately 75% thrust efficiency through shaped fusion charges and magnetic redirection. The ToughSF treatment is valuable as an engineering feasibility sketch but differs from the present work in three respects: it uses D-He3 rather than p-^{11}B (retaining some neutron production from D-D side reactions), it does not address the phase space structure of the products or the Liouville constraint, and it does not quantify the adiabatic limits on magnetic collimation.

The approach proposed here is distinct from both: preserve the non-thermal character of p-^{11}B products and redirect them directly, avoiding both thermalization (Tarditi) and shaped-charge inertial schemes (ToughSF). The phase space argument in Section 2 provides the theoretical basis for why direct redirection is favorable for p-^{11}B specifically. To our knowledge, this argument has not been made in the prior literature.

4. Parametric Estimates

The following estimates define the target design space. They are not predictions. Each depends on assumptions that require validation through simulation.

4.1 Effective Exhaust Velocity

The useful performance variable is the total propulsive energy fraction ψ: the fraction of fusion-product kinetic energy that becomes aft-directed axial kinetic energy of the exhaust. It includes all major losses:

$$\psi = \eta_{nozzle} \times (1 - f_{leak} - f_{trap} - f_{therm} - f_{rad} - f_{recirc} - \dots) \quad (10)$$

where η_nozzle is the magnetic collimation efficiency, f_leak is forward plug loss-cone leakage, f_trap is any trapped or non-extracted population, f_therm is product energy lost to thermalization before extraction, f_rad is radiation and side-reaction loss, and f_recirc is recirculated power required to maintain the source or distribution. This definition deliberately folds uncertain penalties into one quantity so that the mission sensitivity can be shown without hiding the losses.

The equivalent exhaust velocity per unit fuel consumed is:

$$v_e = v_a \times \sqrt{\psi} \quad (11)$$

For $v_a = 0.039c$:

Table 1: Effective Exhaust Velocity vs. Total Propulsive Fraction

Case	ψ	v_e
Idealized one-plug	0.90	0.037c
High-performance design	0.70	0.033c
Conservative design	0.50	0.028c
Aft-hemisphere fallback	0.40	0.025c
Severe degradation	0.30	0.021c
Marginal case	0.20	0.017c

4.2 Comparison with D-T

For propulsion, the relevant metric is directed energy: how much reaction energy converts to steerable axial momentum.

D-T releases 17.6 MeV, but 80% (14.1 MeV) goes to an uncharged neutron. The steerable alpha carries 3.5 MeV. In terms of rest mass, the directed fraction is approximately 0.20 × 0.378% ≈ 0.076%.

p-^{11}B releases 8.7 MeV, 100% in charged products. The directed fraction is 100% of 0.078% ≈ 0.078%.

The two fuels produce comparable directed energy per unit fuel mass. But the system-level implications diverge dramatically. A D-T propulsion system inherits the full burden of neutron management:

Neutron activation and structural life. 14.1 MeV neutrons bombard every structural component in line of sight of the core. The atoms in the ship's walls, magnets, and support structure absorb neutrons and transmute into radioactive isotopes. After sustained operation, the engine structure itself becomes radioactive waste. Maintenance requires remote handling. Component lifetime is determined by neutron damage (displacement per atom), not mechanical wear. A D-T propulsion system has a finite structural life that cannot be extended without replacing irradiated components.

Tritium breeding. Tritium does not exist in nature in useful quantities (half-life 12.3 years). A D-T system must breed its own fuel by surrounding the core with lithium blankets that capture neutrons. This is hundreds of tonnes of lithium structure, plus extraction plumbing, plus containment for a radioactive gas. On a spacecraft, this mass competes directly with payload and propellant.

Biological shielding. 14.1 MeV neutrons are lethal to crew and penetrate most materials. Meters of shielding (typically hydrogenous materials or borated polyethylene) are required between the engine and any crewed volume. For a 1000-tonne ship, shielding mass alone could be 50–200 tonnes depending on geometry and standoff distance.

Thermal conversion chain. 80% of D-T energy is in the neutron. To use it for thrust, the neutron energy must be captured thermally, converted through a heat engine (Carnot-limited at ~30–40% efficiency), converted to electricity, and then used to power a secondary thruster (ion drive, MPD, or similar). This is a nuclear power plant bolted to a rocket: heat exchangers, turbines, generators, power conditioning, secondary propulsion. Each component adds mass, complexity, and failure modes.

Waste heat rejection. At 30–40% thermal conversion efficiency, 60–70% of the captured neutron energy becomes waste heat requiring radiators. For a high-power D-T system, radiator areas of 10^4–10^5 m² may be required, with corresponding mass in the tens to hundreds of tonnes.

A p-^{11}B system greatly reduces these requirements. It does not require tritium breeding, and the primary reaction energy is carried by charged alphas rather than 14.1 MeV neutrons. Secondary neutrons from side reactions, impurities, and beam/plasma interactions are not zero and must be quantified, but the neutron burden is orders of magnitude smaller than D-T for the same useful directed energy. Direct charged-particle exhaust also avoids making the thermal conversion chain the primary thrust path.

This is not merely a difference in raw reaction yield. It is the difference between a neutron-dominated power plant that must survive its own radiation environment and a charged-product exhaust system whose primary energy carrier can, in principle, be magnetically directed. The interplanetary logistics argument in Section 1.3 depends on routine, maintainable, reusable operations. A fair system mass comparison (identified as future work in Section 5) must include shielding, breeding blankets, thermal conversion hardware, radiators, direct converters, coils, structure, and replacement intervals.

Table 2: Energy Partition Comparison

Parameter	D-T	p-^{11}B

Total energy	17.6 MeV	8.7 MeV
Mass-energy conversion	0.378%	0.078%
Charged fraction	20%	100%
Directed energy (% rest mass)	~0.076%*	~0.078%†
Neutron shielding	Significant	Much smaller
Thermal conversion	Required for neutron energy	Not primary thrust path

> ** Before thermal conversion losses. Actual directed fraction after Carnot cycle is lower.*
>
> *† Before nozzle collimation and system losses. Actual directed fraction depends on ψ.*

4.3 Mass Ratio Estimates

$m_0/m_f = \exp(\Delta v/v_e)$ (12)

For a 10-day Earth-Jupiter brachistochrone (5 days acceleration, 5 days deceleration) across approximately 5 AU:

$\Delta v \approx 2 \times d/t_{accel} = 2 \times (7.5 \times 10^{11}\ m)/(4.3 \times 10^5\ s) \approx 3.5 \times 10^6\ m/s \approx 0.012c$ (12a)

where d ≈ 5 AU ≈ 7.5×10^{11} m is the Earth-Jupiter distance and t_{accel} = 5 days ≈ 4.3×10^5 s is the acceleration phase duration. The factor of 2 accounts for both acceleration and deceleration phases. Using this Δv:

Table 3: Mass Ratio Sensitivity

Case	ψ	v_e	Propellant %
Idealized one-plug	0.90	0.037c	27%
High-performance design	0.70	0.033c	30%
Conservative design	0.50	0.028c	34%
Aft-hemisphere fallback	0.40	0.025c	37%
Severe degradation	0.30	0.021c	42%
Marginal case	0.20	0.017c	49%

The concept is robust to moderate degradation because ψ enters through a square root and the rocket equation is logarithmic in v_e. This is stronger than a nozzle-only robustness claim: the same table absorbs forward leakage, direct-conversion penalties, thermalization, imperfect neutralization, and recirculation power. The architecture becomes operationally unattractive as ψ approaches ~0.2, but it does not fail catastrophically at the first departure from ideal collimation.

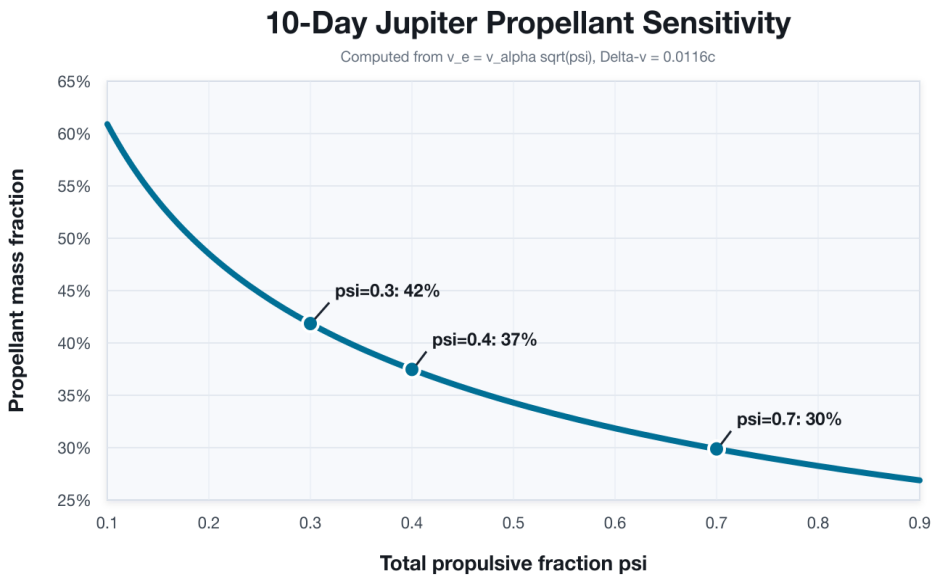

> **Figure 5.** *Sensitivity of propellant mass fraction to total propulsive fraction ψ for the reference mission (Δv = 0.012c, v0 = 0.039c). The curve is nearly flat across ψ = 0.7–0.9 and remains below ~40% propellant at ψ ≈ 0.4. This insensitivity arises from two layers of mathematical compression: ψ enters the exhaust velocity through a square root, and exhaust velocity enters the mass ratio through a logarithm.*

4.4 Power Requirements

The power requirement scales directly with the mission acceleration profile. For a 1000-tonne spacecraft:

30-day Earth-Jupiter transfer (~0.046g continuous): This near-term mission profile requires thrust of ~4.5×10^5 N. For ψ = 0.7, the directed jet power is ~2.2 TW and total fusion power is ~3.1 TW. For ψ = 0.4, total fusion power rises to ~4.1 TW. This is an enormous power level by current standards, but it is an order of magnitude less demanding than the 10-day reference case.

10-day Earth-Jupiter transfer (~0.41g continuous): The reference mission used throughout this paper requires thrust of ~4.0×10^6 N. For ψ = 0.7, directed jet power is ~20 TW and total fusion power is ~28 TW. For ψ = 0.4, total fusion power is ~37 TW. This represents the asymptotic performance limit of the concept, not a near-term design point.

F = ma; P_directed = Fv_e/2; P_fusion = P_directed/ψ (15)

These power levels are inherent to the mission profiles, not specific to any propulsion technology. Any system delivering continuous high-g thrust at relativistic exhaust velocities must sustain comparable output. The propulsion architecture described in this paper does not change the power requirement; it changes the fraction of that power that becomes useful thrust.

Table 4: Propellant Fractions by Propulsion System (10-day Jupiter)

System	Propellant Fraction
Chemical (LOX/LH2)	>99% (infeasible)
Nuclear thermal	>95% (infeasible)
D-T fusion (thermal)	60–97%
p-^{11}B directed exhaust (this paper)	30–42% (ψ = 0.7–0.3)
Antimatter (theoretical limit)	~2%

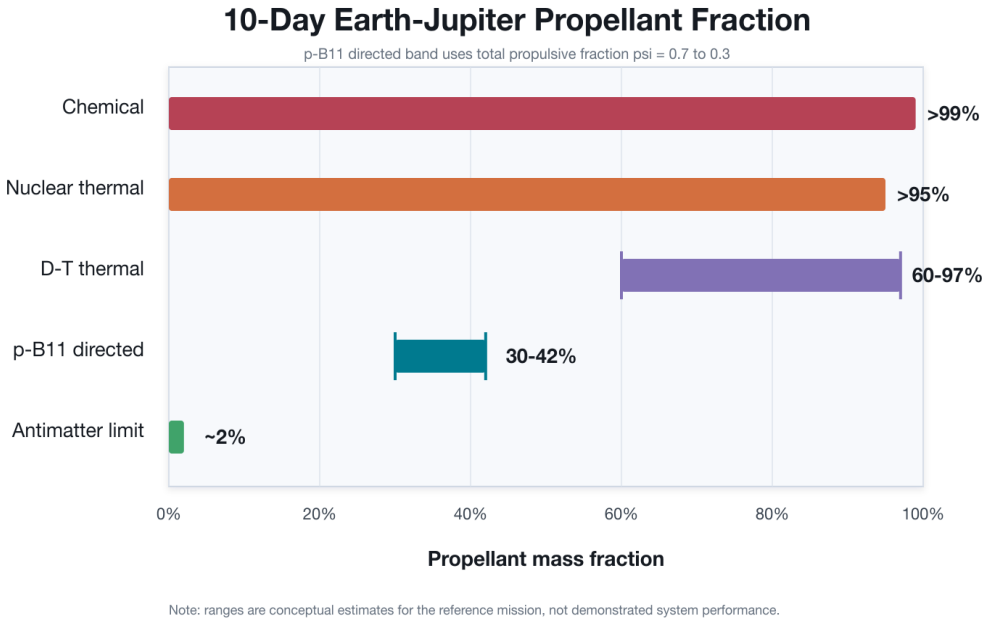

> **Figure 6.** *Propellant mass fractions for a 10-day Earth–Jupiter brachistochrone transfer across propulsion technologies. Chemical and nuclear thermal systems require propellant fractions exceeding 95%, rendering the mission infeasible. D–T fusion with thermal exhaust spans a wide range (60–97%) depending on assumptions about thermal conversion efficiency. The p-^{11}B directed exhaust architecture proposed in this paper spans roughly 30–42% across ψ = 0.7–0.3. Only antimatter propulsion (~2%) offers superior performance, but requires fuel that does not exist in macroscopic quantities.*

5. What This Paper Does Not Establish

The following are not established by the arguments in this paper:

One-plug nozzle performance from simulation. No particle-tracing or PIC simulation has been performed. The adiabatic analysis in Section 2.4 provides order-of-magnitude bounds but does not capture non-adiabatic transition dynamics, plug loss-cone leakage, finite-source-size effects, gyrophase-dependent detachment, or the full velocity-dependent smearing of the detachment surface. The magnitude of performance degradation from these effects is unknown.

Ignition and source physics. This paper does not propose or validate a p-^{11}B ignition scheme. It assumes a source of nonthermal alpha products and analyzes the exhaust architecture downstream of that source. Crystalline-boron beam-target ignition, pulsed laser block ignition, field-reversed configurations, spherical torus

approaches, and other source concepts are separate research programs. Their repetition rate, target survival, driver wall-plug efficiency, debris handling, and recirculating power requirements are not solved here.

Bremsstrahlung budget and Rider's objection. Proton-boron plasmas suffer significant bremsstrahlung losses due to boron's atomic number ($Z=5$) and the high temperatures required for ignition. Rider's 1995 thesis and 1997 analysis argued that near-Maxwellian p-^{11}B systems lose more power to bremsstrahlung than they generate in fusion, and that maintaining non-equilibrium distributions can itself cost more power than the fusion system produces. This paper does not refute Rider. Any viable source architecture must specify the non-equilibrium mechanism, its maintenance power, and the experimental or kinetic-simulation evidence supporting it. Candidate mitigations include low-electron-temperature operation, alpha channeling or alpha power capture, nonthermal ion distributions, and direct conversion of escaping alpha power, but these are not interchangeable labels; each requires its own power balance.

Three-body energy distribution. The alpha energy spread is characterized here as ~±25%. The actual distribution is governed by the Dalitz plot including ^{8}Be resonance structure. Proper treatment requires the measured differential cross-section as initial conditions for particle tracing.

Thermalization avoidance. The concept depends on extracting products before thermalization. Section 3.1 estimates a timescale margin of 100–1000× under a representative density and length scale, but this is not a substitute for source-specific kinetic modeling. The one-plug topology is intended to avoid repeated trapped bounces, but product dwell time near the plug and source must still be modeled.

System mass comparison. The directed-energy comparison (Table 2) is suggestive but incomplete. A fair comparison with D-T requires a full system mass model including shielding, conversion hardware, radiators, and structural mass.

Beam neutralization (critical). This may be the most consequential exhaust-transport open question. At a 10-day reference mass flow of order 0.4 kg/s of He^{2+}, the ion flux is ~6×10^{25} ions/s, corresponding to a beam current of order 20 MA. The Alfvén-Lawson current limit for an unneutralized relativistic beam is approximately $I_A = 17\ \text{kA} \times (\beta\gamma)$, which for 0.039c alphas ($\beta\gamma \approx 0.039$) gives ~660 A. The proposed beam current exceeds the single-particle limit by roughly four orders of magnitude. A naked alpha beam is physically impossible. The exhaust must remain quasi-neutral, with co-moving electrons suppressing self-field defocusing. Ambipolar nozzle physics makes this plausible in a plasma exhaust, but its sufficiency at propulsion-relevant current density requires self-consistent PIC simulation and is a prerequisite for viability, not a secondary correction.

Plug leakage and alpha direct conversion. The one-plug architecture replaces forward recapture with finite loss-cone leakage. For $R_p = 10–20$, the leakage is only ~1–3% of the alpha population, but at multi-terawatt propulsion power this is still tens to hundreds of gigawatts. The proposed alpha interception and direct-conversion assembly is therefore a major subsystem. Its voltage, current, segmentation, shielding, heat rejection, and failure modes are not designed here.

Magnet and structural feasibility. The one-plug concept requires a large-bore high-field plug and a shaped aft nozzle. Fields of 10–20 T over meter-class apertures are beyond present flight hardware and impose large magnetic stresses, cryogenic

loads, radiation shielding requirements, vibration constraints, and assembly challenges. This is a specific magnet and structure R&D gap, not a solved engineering detail.

Fuel and operations. Boron-11 is abundant as an isotope fraction of natural boron, but propulsion-scale operations require purified feedstock, high-throughput injection, ash handling, target or plasma replenishment, and orbital assembly of large magnetic structures. These are industrialization problems, not fundamental objections, but they determine whether the concept is a vehicle or only a physics demonstration.

Each item above is a well-defined computational, analytical, or engineering problem: particle tracing, radiation transport, kinetic simulation, system mass modeling, magnet design, and high-throughput source integration. The exhaust-side problems can be attacked before ignition is solved. The ignition and bremsstrahlung problems cannot be made to disappear by invoking the exhaust architecture; they define the boundary between this concept paper and the follow-on reactor/source papers.

6. Proposed Research Program

If the concept is to be validated or falsified, the following work is required:

6.1 Particle-Tracing Simulation (Highest Priority)

A Boris-pusher particle-tracing simulation using the measured three-body energy and angular distribution through the one-plug nozzle geometry. The minimum study should launch correlated three-alpha events from a finite source volume just aft of the plug, propagate them through a shaped magnetic field, and report ψ, aft axial momentum distribution, plug loss-cone leakage, dwell-time distribution, detachment radius, plume divergence, wall interception, and sensitivity to B_{plug}/B_{source} and nozzle gradient length. This single study would convert the conceptual argument into a quantitative result or reveal that the one-plug estimates are unachievable.

6.2 Bremsstrahlung and Radiation Budget

A self-consistent power balance model including bremsstrahlung, synchrotron radiation, secondary neutron production from side reactions, alpha direct conversion, driver recirculation, and heat rejection. This work must engage Rider's objection explicitly: the source model must identify the non-equilibrium mechanism, estimate the power required to maintain it, and show whether the residual fusion power available to the exhaust remains large enough for ψ in the 0.3–0.7 band.

6.3 Space Charge Analysis

Self-consistent PIC modeling of quasi-neutral alpha-electron exhaust transport at propulsion-relevant flux densities. The key falsification criterion is whether co-moving electrons remain sufficiently coupled through detachment to suppress self-field defocusing. A naked alpha beam is impossible; a quasi-neutral plasma exhaust may be viable.

6.4 System Mass Model

A comparative model for D-T thermal, D-T with direct conversion, and p-^{11}B directed exhaust, including shielding, tritium breeding, thermal conversion, radiators, direct converters, coils, cryogenics, structure, replacement intervals, and orbital assembly mass. This is required before any system-level claim about D-T versus p-^{11}B can be considered mature.

6.5 Confinement-Extraction Integration

Co-design of the alpha source, high-field plug, aft magnetic nozzle, and forward loss-cone direct converter. The source must produce alpha products immediately aft of the plug, while the plug must reflect forward-born products without trapping ash or preventing fuel sustainment. This interface is the main architecture risk introduced by the one-plug solution.

6.6 Alpha Interception and Direct Conversion

Design of the forward loss-cone interception assembly. Even a 1% leakage fraction at a 25 TW source is 250 GW of alpha power. A passive absorber is unlikely to survive except as a sacrificial or flowing target; the credible design space is segmented electrostatic direct conversion, grazing-incidence interception, liquid-metal or rotating surfaces, and localized heat rejection. The output power could support driver recirculation, RF control, magnets, and auxiliary loads, but only if conversion efficiency and current handling are quantitatively demonstrated.

7. Engineering Research Directions

The physics question (does the phase-space argument hold quantitatively?) is separable from the engineering question (can a system be built that exploits it?). This section catalogs the principal engineering challenges without treating any of them as solved.

7.1 Magnetic Nozzle and Plug Optimization

The adiabatic limit identified in Section 2.4 is not fixed. It depends on the field gradient profile. A simple diverging solenoid produces steep gradients ($\nabla B \sim 1/z^3$ in the far field), forcing early detachment. Shaped field profiles using secondary trim coils downstream of the plug can stretch the gradient scale length L_B, keeping ε small deeper into the expansion and recovering a higher effective mirror ratio.

The trade-off is physical scale. Stretching L_B to ~10 m at 0.05 T requires coil structures and truss extending tens of meters, with bore radii exceeding the local Larmor radius (~5 m at those field strengths). The high-field plug adds a second trade: increasing B_{plug} reduces forward leakage but increases coil stress, mass, stored magnetic energy, shielding, and cryogenic burden. The optimal design likely sits at moderate plug ratios and aft nozzle ratios, not at an extreme 50 T large-bore idealization.

This is a high-dimensional, non-convex optimization over field geometry, coil placement, current profiles, structural mass, direct-converter placement, and source coupling. The output of this optimization is not just η_{nozzle} but ψ: the total propulsive fraction after all loss terms.

7.2 Source and Bremsstrahlung Closure

Bremsstrahlung losses in p-^{11}B plasmas are a fundamental challenge to ignition and sustained operation. A credible source paper must choose a specific mechanism rather than invoking generic non-equilibrium. Examples include maintaining T_e well below ion energy, alpha channeling or alpha power capture, beam-target operation in structured media, or pulsed block ignition. For each mechanism the burden is the same: calculate fusion power, bremsstrahlung power, distribution-maintenance power, driver recirculation power, and usable alpha power.

High-peak-power laser or beam concepts do not automatically solve the average-power problem. A 10 PW, 1 ps pulse carries only 10 kJ, but at 100,000 pulses/s it becomes 1 GW of average driver power and requires target replacement, alignment, debris clearing, and heat removal every 10 μs. For propulsion, repetition rate and target/channel survival are as important as single-shot intensity.

7.3 Alpha Interception and Direct Conversion

The forward plug leakage is small in fraction but large in absolute power. A loss cone of 1–3% corresponds to tens or hundreds of gigawatts for the mission profiles considered here. A direct converter for MeV alpha particles operates at megavolt scale; for example, 250 GW at ~1.5 MV corresponds to order 10^5 A. The practical architecture is therefore likely segmented, distributed, and actively cooled, with strict control of secondary electrons, arcing, sputtering, and radiation damage.

This subsystem is not only a loss handler. If it can recover a useful fraction of plug leakage, it can supply recirculating power for drivers, RF control, superconducting systems, and onboard electrical loads. The condition is quantitative: η_convert f_leak P_fusion must exceed the relevant recirculating loads.

7.4 Thermal Management, Structure, and Assembly

Even with high ψ, residual losses from imperfect collimation, plug leakage, bremsstrahlung, secondary radiation, direct-converter inefficiency, and driver recirculation become heat. At 25–40 TW source power, a few percent loss is a terawatt-scale thermal problem. Radiator area, radiator temperature, shadow shielding, direct-converter heat rejection, and coil cryogenic isolation may become binding constraints.

The plug and nozzle also impose large structural loads. Continuous 0.05–0.4g acceleration on a 1000-tonne vehicle, pulsed thrust smoothing, magnetic stress in high-field coils, thermal cycling, and vibration all require explicit structural dynamics. Zero-gravity assembly and deployment of tens-of-meters magnetic structures is an operational challenge comparable in complexity to major orbital construction projects.

7.5 Fuel Handling and Supply Chain

Boron-11 constitutes roughly 80% of natural boron, and hydrogen is abundant. Neither fuel component is exotic, but propulsion-scale operation is not just fuel availability. A 10-day reference burn consumes fuel at order 0.4–0.5 kg/s depending on ψ, corresponding to hundreds of tonnes across acceleration and deceleration. Sustained operations require boron purification or isotope accounting, high-throughput injection, ash handling, target or plasma replenishment, and orbital logistics.

7.6 Autonomous Research

Autonomous research systems are relevant because several exhaust-side problems are computationally bottlenecked: particle tracing, nozzle/plug optimization, PIC neutralization, direct-converter geometry, shielding layout, and coupled thermal-structural design. These are plausible candidates for large acceleration through simulation, optimization, and automated design iteration.

The acceleration is not limited to computational tasks. The same compression applies to the full research loop – theory, simulation, fabrication, diagnostics, and integration – whenever the binding constraint is human-bandwidth coordination rather than physical law. The development of this paper itself, across multiple adversarial AI reviewers in compressed timeframes, is a demonstration at small scale.

Ignition is different. It is partly a theory and optimization problem, but also a hardware-cycle problem involving targets, magnets, lasers or beams, diagnostics, materials, and repeated experiments. Experimental cycle times are not fixed by physics alone; they reflect current methods, supply chains, fabrication throughput, diagnostic automation, and human-bandwidth coordination. Each is a target for compression, but each must be evaluated against the actual physical bottleneck. This paper therefore makes the modest claim that exhaust architecture should be developed in parallel with ignition research, not that p-^{11}B ignition is imminent.

8. Broader Context

This paper is motivated by a prioritization problem. If p-^{11}B alpha production becomes available, the value of the source depends heavily on whether the products can be used directly as exhaust. Waiting for ignition before studying exhaust handling risks discovering too late that the source and nozzle must be co-designed.

Ignition is the long pole. Achieving p-^{11}B ignition is extremely difficult with current methods. This paper does not forecast a specific ignition date. It treats ignition as an upstream source program whose timeline depends on which bottlenecks are physical and which are methodological, logistical, or coordination-limited. Autonomous research may compress the latter substantially, but the power-balance and hardware-cycle constraints still have to close.

The architecture is partially separable. Exhaust handling can be studied before ignition, but it is not independent of the source. The one-plug nozzle requires alpha production near the plug and rapid extraction into an open aft field. This makes early source-nozzle co-design essential.

The application is time-sensitive. Infrastructure decisions in the next decade (lunar bases, Martian settlements, asteroid mining) will be shaped by assumptions about available propulsion. If directed p-^{11}B exhaust is feasible, it changes the optimal architecture for interplanetary logistics. That information is needed before the infrastructure is built.

Reusability changes the economics. As detailed in Section 4.2, p-^{11}B shifts the dominant energy carrier from neutrons to charged particles. That does not make radiation vanish, but it changes the system design problem from neutron-dominated heat capture and activation management to charged-product transport, neutralization,

direct conversion, and magnetic structure. The choice of fuel determines not just propulsion performance but whether the resulting vehicle is closer to a reusable transport asset or a periodically replaced nuclear engine.

9. Conclusion

Proton-boron fusion products have a phase space structure qualitatively more favorable for direct magnetic collimation than thermal plasmas. Their entropy is concentrated in angular distribution rather than momentum magnitude. Magnetic nozzle geometries can collimate the exhaust by trading angular spread for spatial spread, consistent with Liouville's theorem, placing the conserved phase space volume in a degree of freedom irrelevant to thrust.

Adiabatic analysis establishes that collimation is physically effective at aft mirror ratios up to ~10, bounded by the Larmor radius of 2.9 MeV alphas in the diverging field region. Beyond this ratio, particles undergo non-adiabatic detachment with gyrophase-dependent and velocity-dependent smearing of the exhaust angle. The one-plug topology places the alpha source just aft of a high-field magnetic plug, reflects most forward-born products aft, and converts the remaining narrow forward loss cone into a bounded direct-conversion problem.

Using total propulsive fraction ψ rather than nozzle efficiency alone, the reference 10-day Earth-Jupiter case remains in an operationally interesting regime across a wide band: $\psi = 0.7$ gives ~30% propellant, $\psi = 0.4$ gives ~37%, and $\psi = 0.3$ gives ~42%. This robustness does not prove the architecture works. It shows that the concept is not balanced on a single idealized 90% collimation number; the next question is whether particle tracing and PIC simulation can produce ψ in this range.

The hardest unsolved problem remains p-^{11}B ignition or alpha production at useful power density. Rider's bremsstrahlung objection, non-equilibrium maintenance power, repetition-rate limits, source lifetime, and recirculating driver power are not solved by the nozzle. They are the subject of follow-on source papers. The contribution here is narrower: if nonthermal p-^{11}B alpha products are available, their phase-space structure is unusually favorable for direct magnetic exhaust.

If validated, this concept would enable propulsion performance approaching the regime usually reserved for much more exotic systems, using fuel components that are abundant and non-radioactive before use. Compared with D-T, p-^{11}B avoids tritium breeding and greatly reduces neutron-dominated system burden. The price is severe: p-^{11}B source physics, bremsstrahlung management, plug/nozzle integration, quasi-neutral exhaust transport, terawatt heat rejection, and large-bore high-field magnets.

The question this paper poses to the community is whether the phase space argument developed here is quantitatively sufficient to justify a dedicated simulation and design effort. We have tried to state the uncertainties clearly enough that readers can form their own judgment.

Acknowledgments

Iterative development and computational analysis of arguments in this paper were conducted with assistance from Claude Opus 4.6 (Anthropic), Gemini 3.1 Pro (Google DeepMind), and GPT-5.4 (OpenAI). The adiabatic limit analysis in Section 2.4 and the detachment physics in Section 2.5 were refined through adversarial dialogue with Gemini and Opus 4.5, which provided the magnetic rigidity

calculations and identified the velocity-dependent smearing of the detachment surface. Claude Sonnet 4.6 (Anthropic) provided independent adversarial review that identified the Alfvén current limit constraint, the underived Δv, and the need to strengthen the D-T reusability comparison. Subsequent adversarial review identified the two-mirror trapped-population problem, motivating the one-plug topology in Section 3.3. All substantive claims, interpretive judgments, and errors are the sole responsibility of the author(s).

References

1. Rostoker, N., Binderbauer, M.W., and Monkhorst, H.J. (1997). "Colliding Beam Fusion Reactor." Science, 278(5342), 1419–1422.

2. Putvinski, S.V., Ryutov, D.D., and Yushmanov, P.N. (2019). "Fusion reactivity of the pB11 plasma revisited." Nuclear Fusion, 59(7), 076018.

3. Nevins, W.M. and Swain, R. (2000). "The thermonuclear fusion rate coefficient for p-^{11}B reactions." Nuclear Fusion, 40(4), 865.

4. Rider, T.H. (1997). "Fundamental limitations on plasma fusion systems not in thermodynamic equilibrium." Physics of Plasmas, 4(4), 1039–1046. See also Rider, T.H. (1995). *Fundamental limitations on plasma fusion systems not in thermodynamic equilibrium*. Ph.D. thesis, Massachusetts Institute of Technology.

5. Liu, S. et al. (2025). "Feasibility of proton–boron fusion under non-thermonuclear steady-state conditions: Rider's constraint revisited." Physics of Plasmas, 32(1), 012101.

6. Ochs, I.E. and Fisch, N.J. (2024). "Lowering the reactor breakeven requirements for proton–boron 11 fusion." Physics of Plasmas, 31(1), 012503.

7. Ahlstrom, H.G. (1982). "Magnetic Nozzle Design for Plasma Propulsion." AIAA Journal, 20(7), 971–977.

8. Santarius, J.F. (1998). "Magnetic Fusion for Space Propulsion." Fusion Technology, 33(1), 28–39.

9. Atzeni, S. and Meyer-ter-Vehn, J. (2004). The Physics of Inertial Fusion. Oxford University Press.

10. Borowski, S.K. (1995). "Comparison of Fusion/Antiproton Propulsion Systems for Interplanetary Travel." NASA TM-107030.

11. Kammash, T. (1995). Fusion Energy in Space Propulsion. AIAA Progress in Astronautics and Aeronautics, Vol. 167.

12. Tarditi, A.G., Scott, J.H., and Miley, G.H. (2012). "Aneutronic Fusion Spacecraft Architecture." NASA NIAC Phase I Final Report.

13. Hyde, R.A. et al. (1991). "VISTA: A Vehicle for Interplanetary Space Transport Application Powered by Inertial Confinement Fusion." LLNL Report UCRL-ID-107066.

14. Hora, H. et al. (2017). “Fusion energy using avalanche increased boron reactions for block-ignition by ultrahigh power picosecond laser pulses.” Laser and Particle Beams, 35(4), 730–740.

15. Eliezer, S. and Martinez-Val, J.M. (1998). “Proton-boron-11 fusion reactions induced by heat-detonation burning waves.” Laser and Particle Beams, 16(4), 581–598.

16. Sikora, M.H. and Woo, H.J. (2016). “A new evaluation of the p+^{11}B reaction rates.” Nuclear Physics A, 945, 1–13.

17. Stave, S. et al. (2011). “Understanding the ^{11}B(p,α)αα reaction at the 0.675 MeV resonance.” Physics Letters B, 696(1–2), 26–29.

18. Binderbauer, M.W. et al. (2015). “A high performance field-reversed configuration.” Physics of Plasmas, 22(5), 056110.

19. Chapman, J.J. (2011). “Advanced Fusion Reactors for Space Propulsion and Power Systems.” NASA Technical Report.

20. Goldston, R.J. and Rutherford, P.H. (1995). Introduction to Plasma Physics. CRC Press.

21. van der Linden, S. et al. (2023). “First measurements of p^{11}B fusion in a magnetically confined plasma.” Nature Communications, 14, 1077.

22. MatterBeam (2019). “The Expanse’s Epstein Drive.” ToughSF Blog. Available at: http://toughsf.blogspot.com/2019/10/the-expanses-epstein-drive.html

23. Ahedo, E. and Merino, M. (2010). “Two-dimensional supersonic plasma acceleration in a magnetic nozzle.” Physics of Plasmas, 17(7), 073501.

24. Merino, M. and Ahedo, E. (2016). “Magnetic nozzle plasma exhaust simulation for the VASIMR thruster.” Proceedings of the 52nd AIAA Joint Propulsion Conference.

25. Fisch, N.J. and Rax, J.M. (1992). “Interaction of energetic alpha particles with intense lower hybrid waves.” Physical Review Letters, 69(4), 612–615.

Appendix A: Derivation of Key Relations

A.1 Alpha Particle Velocity

$v = c\sqrt{(1 - (mc^2/(T + mc^2))^2)}$ (A1)

For m_ac^2 = 3727.4 MeV, T = 2.9 MeV: v ≈ 0.039c.

A.2 Mass-Energy Conversion

$f = Q / (m\_reactants \times c^2)$ (A2)

p-^{11}B: 8.7/11,190.8 ≈ 0.078%. D-T: 17.6/4,684.5 ≈ 0.376%.

A.3 Magnetic Rigidity and Larmor Radius

For He^{2+} at 2.9 MeV: m = 6.64×10^{-27} kg, q = 3.22×10^{-19} C, v ≈ 1.18×10^{7} m/s.

r_L = mv/(qB) ≈ 0.24/B meters (A3)

At 10 T: 2.4 cm. At 1 T: 24 cm. At 0.1 T: 2.4 m.

A.4 Adiabaticity Parameter

ε = r_L/L_B (A4)

Adiabatic invariant conserved when ε ≪ 1. For L_B = 2 m: ε crosses unity near B ≈ 0.1 T, establishing the effective mirror ratio at R ≈ 10.

A.5 Specific Impulse

$v_e = v_a\sqrt{\psi}$ (A5)

I_sp = v_e/g_0 (A6)

For ψ = 0.7, v_e ≈ 0.033c and I_sp ≈ 1.0×10^{6} s. For ψ = 0.4, v_e ≈ 0.025c and I_sp ≈ 7.5×10^{5} s.

A.6 Plug Loss-Cone Fraction

For plug ratio R_p = B_plug/B_source, the loss-cone half-angle is:

$\sin^2\theta_c = 1/R_p$ (A7)

The forward-leaking fraction of an isotropic distribution is the forward loss-cone solid angle divided by 4π:

$f_{leak} = 2\pi(1 - \cos\theta_c)/4\pi = (1 - \cos\theta_c)/2$ (A8)

R_p = 10 gives f_leak ≈ 2.6%; R_p = 20 gives f_leak ≈ 1.3%; R_p = 50 gives f_leak ≈ 0.5%.