The petrodollar architecture, deuterium fusion, and the case for direct-conversion compact pulsed FRC

J. Samper Finberg[a)]

Laurelin Technologies Inc., Dover, Delaware, USA

(Dated: 17 May 2026)

The case for fusion in the United States is, before it is a climate or industrial-policy case, a national-security case. This paper makes that case in four moves. First, the petrodollar arrangement, on the orthodox reading of the political-economy literature, is an architecture of substantial net cost to the United States — forward-deployed naval power, a binding ceiling on sanctions enforcement against energy exporters, and an active financing channel for the adversaries the same architecture must then suppress — while the dollar's reserve role is sustained, on the rigorous reading, by mechanisms that do not depend on oil invoicing. Second, the only durable exit is a backstop fuel substrate whose abundance dissolves the geological rent on which the architecture rests; we formalize the condition as a Hotelling rule with a public firm-dispatchable ceiling. Third, the substrate category that delivers this property is fusion and, within fusion, the deuterium–deuterium (^{2}H–^{2}H) terminal-fuel commitment, because no other candidate cycle sources its fuel outside another sovereign's strategic-inventory posture. Fourth, one combination of architectural commitments composes into a viable ^{2}H–^{2}H operating envelope: a compact pulsed field-reversed-configuration geometry, direct electromagnetic energy recovery, and a container-class packaging envelope matched to the post-ADVANCE Act regulatory framework. We argue the combination as a category, and name Laurelin Technologies as one program pursuing it. The decision is, in the strict sense, definite.

[a)]Electronic mail: info@laurelin-inc.com

I. INTRODUCTION

The argument of this paper is the following. The petrodollar arrangement, as it has been operated since the 1974 Treasury–Saudi reconstruction, is a protection-and-obedience sovereignty relation conducted in the vocabulary of financial incumbency. Its operative costs to the United States — in forward-deployed naval power, in a binding political ceiling on the enforcement of sanctions against energy exporters, in an active rent-financed channel that has funded the adversaries the same architecture is then required to suppress, and in the opportunity cost of an allied energy policy structured around a chokepoint that closes on the Iranian Majlis's calendar — are in 2026 paid in real time. The benefits the arrangement is supposed to secure for the dollar's reserve role are, on the weight of the political-economy literature since Spiro[1–4], secured by other mechanisms — by the depth and liquidity of United States financial markets, by the credibility of United States institutions, and by the Federal Reserve swap-line architecture that operates whether or not oil is priced in dollars[5,6]. The exit from the cost side of the ledger is therefore not a policy reform; it is a technology decision at the fuel-substrate level. The only category of technology that delivers the exit is a substitute that dissolves the geological rent on which the architecture rests, and the substitute category that delivers that property, on the public physics and the public operating record, is fusion with deuterium (^{2}H) as a terminal fuel, recovered electromagnetically, and packaged inside a transportable container-class envelope. Laurelin Technologies is one program pursuing this combination.

The deployment window in which the argument is operative is narrow and is set by load. United States and global electricity demand is now growing at roughly 3%–4% annually after a decade of slow growth[7]; the International Energy Agency's central projection for the data-center channel alone is approximately 1,300 TWh by 2035[8]; the United States interconnection queue stood at roughly 2,289 GW at the end of 2024 against historical completion rates of order 14%[9,10]; and the hyperscalers have acknowledged the gap as an acceleration of capacity deployment rather than a retreat[11]. The window selects for firm power sited behind the meter, on the schedule of the host facility, without a new transmission interconnect. That is the architectural shape the political diagnosis and the load arithmetic together select for, and it is the architectural shape the paper argues for in §V.

The paper is organized in eight sections. §II develops the petrodollar arrangement as an

2026-05-18 03:12 UTC - IP 216.73.216.54 - Columbus, OH, US (39.9587, -82.9987) - EVENT c944871f-2d71-4293-9e44-ed1da55d6279

architecture of net cost, including the two analytical anchors — the net-cost ledger (Eq. (1)) and the Hotelling rule with backstop ceiling (Eq. (2)) — around which the political diagnosis is organized. §III makes the affirmative case for nuclear as the category of generation matched to firm-power deployment in this decade. §IV narrows the case from nuclear in general to deuterium fusion specifically. §V is the core of the paper: it makes the affirmative case for the four architectural commitments that compose into a viable ^{2}H–^{2}H operating envelope, with the ^{2}H–^{2}H Lawson penalty (Eq. (4)), the engineering-Q inequality (Eq. (5)), and the configurational beta relation (Eq. (3)) as its analytical spine. §VI maps the fusion landscape against the six public architectural coordinates that the rest of the paper takes as the relevant unit of comparison, and locates Laurelin explicitly within it. §VII states the risks and the open problems that the architectural commitments inherit. §VIII closes with the backstop-milestone condition (Eq. (6)) that ties the political and architectural halves of the argument together. The full evidentiary surface — the empirical reading of the petrodollar literature, the primary-source engagement with the political-philosophy texts, and the full public-physics scaling arguments — is in the long-form companion paper, of which the present paper is the executive summary; the long paper is available from the corresponding author on request.

II. THE PETRODOLLAR ARCHITECTURE AS NET COST

The orthodox academic position on the petrodollar is sharper than the policy discourse and is rarely stated in the policy discourse, which is why the cost ledger has not been audited in public. Spiro's foundational reconstruction of the 1974 Treasury–Saudi negotiations established the recycling arrangement as a deliberate political construction through which Saudi surpluses were routed into United States Treasury markets outside the standard international-institutional channels[1,12]; the subsequent consensus in the political economy of money is that the construction was contingent and that the dollar's reserve role is sustained by the depth and liquidity of United States financial markets, by the credibility of United States institutions, and by the network effects of incumbency rather than by the mechanics of oil invoicing[2–4]. Helen Thompson's contrary position[13] is the strongest version of the case that the petrodollar matters and is engaged at length in the long-form companion; the synthesis we use here is the one the weight of the literature supports.

2026-05-18 03:12 UTC - IP 216.73.216.54 - Columbus, OH, US (39.9587, -82.9987) - EVENT c944871f-2d71-4293-9e44-ed1da55d6279

The architecture imposes four cost channels, three of them positive and one of them now approximately zero. The first positive channel is the protection cost. Posen's command-of- the-commons accounting[14] sets out the structural naval primacy required to keep the chokepoints open; the cost is paid in Fifth Fleet rotations from Naval Support Activity Bahrain, in nine-month carrier-strike-group combat deployments to United States Central Command, in Operation Prosperity Guardian and its successor formations against the Houthi infrastructure, in the offensive Operation Poseidon Archer strikes ongoing since January 2024, and in the mine-counter-measures and escort posture currently sustaining the Strait of Hormuz against the Iranian closure that began 10 March 2026[15–23]. Talmadge[24] concluded that the United States can reopen the Strait at bounded cost, on a timescale of weeks; the empirical test of that assessment is being passed at roughly the predicted cost. The Department of Defense's 2009 estimate of one casualty per twenty-four fuel- and-water resupply convoys in the Iraq and Afghanistan theaters[25,26] is the tactical-level arithmetic; the Posen accounting is the strategic analogue.

The second positive channel is the sanctions-enforcement ceiling. The empirical literature on sanctions effectiveness places the success rate of coercive economic sanctions at between 5% and 34% of attempted episodes[27,28]; the operationally decisive finding is that the political ceiling on enforcement against energy exporters is set by the price elasticity of supply available to the sanctioning coalition. Farrell and Newman[29] call this weaponized interdependence; the empirical calibration is the 2022–2024 European experience with the price cap on Russian crude, in which the cap was set at the level at which energy-price pass-through to coalition consumers was tolerable rather than at the level at which damage to the target was maximized[30,31]. A coalition that possesses a deployable backstop substrate with effectively infinite own-supply elasticity has the political-cost gradient at the optimum vanish and the politically feasible enforcement intensity rise to its technical maximum — a comparative-static derived in the long paper from the coalition-welfare first-order condition.

The third positive channel is the rent-financed adversary loop. The resource-curse literature has documented the dynamic with unusual rigor: Ross[32,33] establishes that petroleum wealth makes authoritarian regimes more durable, increases the frequency and severity of civil war in poor producers, and finances coercive apparatus that would otherwise depend on the productive capacity of the producer's own society; Andersen, Johannesen, Lassen, and Paltseva[34] close the identification problem by matching petroleum-price shocks to leaked offshore tax-haven deposit data and find that oil-price spikes in autocracies are followed within quarters by detectable rises in deposits, with the elasticity essentially zero in democracies; the institutional bifurcation documented in Mehlum, Moene, and Torvik[35] is the mechanism. Russia drew 45% of its 2021 federal receipts from oil and gas[36]; the same revenue stream financed the 2022 invasion of Ukraine, sustained the war effort against an EU sanctions regime that the same energy dependence had forced to be soft, and supplied Iran with the hard-currency cushion that financed the IRGC, the Houthis, and the proxy network whose Strait-of-Hormuz operations the United States Navy has been actively suppressing since 2024. The loop is closed; the loop is expensive; the loop is the architecture.

The orthodox benefit channel — the imputed dollar privilege — is the channel the political-economy literature has driven toward zero on the rigorous reading. The dollar's most recent share of allocated official foreign exchange reserves stood at 58% in early 2024, the lowest share since 1995[5,6]; the diversification has been running into nontraditional currencies and into gold; and the modern buttress of the dollar system, on the post-2008 reading, is the Federal Reserve swap-line architecture, which functions as the de facto international lender of last resort for major allied central banks and which operates whether or not oil is priced in dollars[4,5].

The four channels combine into a single ledger identity. Write the present value of the architecture, on the political-economy parameters this section has reviewed, as

$$V_{\text{arch}} \;=\; \int_0^T e^{-\rho t}\Big[\,B(t) - \Pi(t) - L(t) - A(t)\Big]\,dt, \tag{1}$$

where $B(t)$ is the imputed monetary benefit, $\Pi(t)$ the United States protection cost (Carter Doctrine maintenance, Fifth Fleet rotations, Gulf basing, convoy operations), $L(t)$ the welfare loss imposed by the sanctions-enforcement ceiling, $A(t)$ the rent-financed adversary-loop cost, and ρ the social discount rate. The comparative-statics of the integrand are unambiguous in sign: $\partial V_{\text{arch}}/\partial\Pi < 0$, $\partial V_{\text{arch}}/\partial L < 0$, $\partial V_{\text{arch}}/\partial A < 0$, with $\partial V_{\text{arch}}/\partial B > 0$ acting on a term whose magnitude the literature has driven toward zero. The orthodox policy literature has, by convention, debated the sign of $B - \Pi$ in isolation; the political-economy literature has driven B toward zero; Posen establishes $\Pi(t) > 0$; Drezner–Farrell–Newman establish $L(t) > 0$; and the resource-curse literature with the Andersen identification establishes $A(t) > 0$ at quantitatively-bounded elasticity. The sign of V_{arch} is therefore the sign of the integral of four interior terms whose signs are independently estimated in the literature, three of them in the cost direction. The petrodollar architecture is, on this ledger and on the weight of the literature it draws from, a national-security liability rather than a national-security asset.

The pattern of the cost-benefit reversal was named, at the level of the form rather than the case, by Carl Schmitt in his 1929 essay on neutralizations: a sovereignty relation is repeatedly translated into the vocabulary of the latest neutral domain — theological, metaphysical, ethical, and finally economic[37] ("Age of Neutralizations and Depoliticizations," pp. 82, 89). The petrodollar is the apex of the economic neutralization; the arrangement presents itself in financial terms because presenting itself in financial terms is its mechanism of concealment. The correct question, on the diagnostic reading Strauss endorses in his 1932 review of Schmitt[38] ("Notes on Schmitt," pars. 27, 35), is not whether the arrangement is financially beneficial; it is what the protection commitment costs, whether it is still deliverable, and what would dissolve the underlying dependency. The response, on the explicitly anti-Schmittian normative move Strauss makes in *Natural Right and History*[39] (p. 162), is not to deepen the exception but to remove the dependency that makes the exception necessary. The technology that removes the dependency is the subject of the rest of this paper.

The analytical complement to the ledger identity is the Hotelling rule with a backstop ceiling. Under competitive equilibrium with constant real interest rate r, the net rental price P_t of an exhaustible resource (price minus marginal extraction cost c) must rise at rate r[40]; the introduction of a backstop technology available at constant unit cost $\bar{c}$ caps the resource-price trajectory at $\bar{c}$[41,42]:

$$\frac{\dot{P}_t}{P_t} = r, \qquad P_t \leq \bar{c}. \tag{2}$$

On the deployment date of the backstop the resource rent collapses to $\bar{c} - c$, and as $\bar{c}$ falls toward c the discounted present value of the resource owner's rent stream $\int_0^\infty e^{-\rho t}\,(P_t - c)\,dt$ collapses toward zero. The comparative-static is $\partial P_t^{\max}/\partial \bar{c} = 1$ in the binding regime and $\lim_{\bar{c}\to c^+} V^{\text{rent}} = 0$ in the limit of a costless backstop. The political consequence is direct: $A(t)$ in Eq. (1) is bounded above by the integrated Hotelling rent, which is in turn bounded above by the backstop ceiling; an architecture that delivers a deployable backstop at industrially-priced $\bar{c}$ collapses the rent that financed the loop, and with it, by the comparative-static of the long-paper sanctions argument, the ceiling on $L(t)$ and, through the dynamics of the alliance relations the loop sustains, the gradient of $\Pi(t)$. The exit is not, in the strict sense, a policy. It is a substrate.

III. THE NUCLEAR CASE

The backstop substrate of Eq. (2) is not abstract; it has to be a real technology with a real generation profile that the host market accepts as firm. The candidates that satisfy the firm-dispatchable criterion on the public record are three: advanced gas with carbon capture, long-duration storage paired with non-firm renewables, and nuclear (fission today, fusion in this decade). Of the three, nuclear is the category whose energy density and dispatchability make it the substrate against which the political diagnosis of §II actually closes. We make the case here at the category level. The narrowing from nuclear to fusion specifically, and from fusion to deuterium fusion specifically, is the work of §IV and §V.

The first argument is energy density. The energy released per unit mass of fuel separates fossil chemistry from fission by roughly six orders of magnitude and separates fission from fusion by a further factor of order four. Coal combustion releases of order 3×10^7 J per kilogram; the fission of one kilogram of uranium-235 releases of order 8×10^{13} J; the deuterium–deuterium fusion of one kilogram of deuterium releases of order 3×10^{14} J. The practical content of this separation is supply-chain autarky: a generation source with the energy density of fission or fusion can carry its working fuel inventory for a year or more inside the deployed unit, without an external resupply tail, and the marginal cost of fuel becomes a small fraction of the levelized cost of electricity for the entire technology family. The Department of Defense has spent twenty years documenting the operational and casualty cost of the fuel-resupply tail in expeditionary operations[25,26]; the convoy-casualty arithmetic is one expression of the energy- density gap, and the architectural answer is to close it. Solar, wind, hydroelectric, and combustion-class generation cannot close it. Only nuclear can.

The second argument is firm dispatchability. Solar and wind generation are non-firm by construction; their availability is governed by weather and time of day, and they cannot, without storage, supply the constant load that hyperscale data centers, industrial process heat, defense installations, and most of the modern grid require. Long-duration storage at grid scale is an unsolved problem in 2026: lithium-iron-phosphate batteries are commercially mature at four-to-eight-hour discharge but have not demonstrated the multi-day duty cycle that a firm replacement for fossil generation requires; flow batteries, compressed-air, and gravity-based storage are at substantially earlier stages of deployment; and the levelized cost of storage at the duration required to firm an all-renewable grid is, on the United

States Energy Information Administration's published projections, substantially above firm fossil and firm nuclear[43]. Hydroelectric power is firm and dispatchable but is geographically constrained at the required scale; the remaining United States expansion potential is small relative to the incremental load the next decade is delivering[7,8]. The category that supplies firm low-carbon dispatchable power at scale, deployable on schedules consistent with the load arrival of §I, is nuclear.

The third argument is the post-2024 regulatory rebase. The *Accelerating Deployment of Versatile, Advanced Nuclear for Clean Energy Act of 2024* (ADVANCE Act), signed into law on 9 July 2024[44], amended the Atomic Energy Act to add a statutory definition of "fusion machine" and directed the Nuclear Regulatory Commission to regulate near-term fusion devices under the byproduct-material framework (10 CFR Part 30) rather than under the utility-scale reactor framework (10 CFR Part 50)[45,46]. The Commission transmitted the draft proposed rule and the program- specific licensing guidance of NUREG-1556 Volume 22 as SECY-24-0085 in December 2024[47]; the Nuclear Innovation Alliance's regulatory-implementation summary of November 2025 records the NRC as having completed thirty of the thirty-six ADVANCE Act milestones tracked on its public dashboard[48]; and the statutory backstop in the 2019 Nuclear Energy Innovation and Modernization Act requires the fusion framework to be finalized no later than 31 December 2027. A fusion machine is now licensed on the timescale and under the inspection regime of a particle accelerator or industrial-irradiation facility, not a utility power plant. Licensing devolves, in much of the country, to the host state's radiation control program under Agreement-State authority — a fifty-state, rather than a single-Commission, licensing surface. The capital expense of regulatory engagement falls by roughly an order of magnitude relative to the Part 50 pathway; the schedule burden falls similarly; and public deployment can be contemplated on a timeline measured in single- digit years rather than the decades associated with utility-fission licensing.

The fourth argument distinguishes fusion from fission within the nuclear category. Advanced fission is a real candidate substrate and is well represented in the post-2024 procurement vehicles (the Department of Defense Strategic Capabilities Office's Project Pele[49], the Defense Innovation Unit's Advanced Nuclear Power for Installations[50]); we do not argue against it. We argue that fusion, when delivered, carries three properties that fission does not. First, fusion's terminal fuel is hydrogen-isotope rather than uranium-isotope, which removes the fuel cycle from the Nuclear Suppliers Group trigger list, from the Bureau of Industry and Security's safeguarded- inventory posture, and from the long-term waste-storage problem that fission inherits from the actinide chain. Second, fusion's proliferation surface is structurally smaller: a fusion machine does not produce weapons-relevant fissile material as a by-product of operation, and the Atlantic Council's April 2025 issue brief on fusion non-proliferation and export controls is the first serious public treatment of the governance design that a coordinated allied fusion bloc would require[51]. Third, fusion is the category against which the deuterium-fuel-cycle argument of §IV closes the political loop with §II; fission, even in its most advanced forms, operates on a different ratio of fuel-cost to capital-cost, and the geological-rent collapse of Eq. (2) does not bind through fission in the same way it binds through a hydrogen-isotope substrate. The case for fusion is the case for the nuclear category minus the nuclear-specific costs.

The nuclear category therefore satisfies the energy-density test, the firm-dispatchability test, the post-2024 regulatory test, and (in the fusion sub-category specifically) the fuel- cycle test that the political diagnosis of §II put in front of any candidate substrate. The narrowing to fusion specifically, and to deuterium fusion specifically, is the subject of the next section.

IV. WHY DEUTERIUM IS THE TERMINAL FUEL

Fuel-cycle choice is the single architectural commitment that a fusion program cannot easily revise once hardware is in fabrication. It determines the neutron-channel engineering, the isotope-handling infrastructure, the regulatory posture, and the commodity supply chain on which the program depends. The case for deuterium (^{2}H) as the terminal fuel is, on the public physics and the public political surface, sharper than the plasma-physics question alone admits, and it is sharpest when stated against the alternatives. The relevant menu is four-wide.

The deuterium–tritium cycle, ^{2}H–^{3}H, has the highest reactivity at accessible ion temperatures and the lowest Lawson product to reach unity gain[52,53]; it is the default of the tokamak and inertial-confinement programs. It also carries the largest neutron burden of the candidate set, in the form of the 14.1 MeV channel from the dominant reaction, and it depends on tritium (^{3}H) breeding, with the tritium supply chain itself remaining a substantial unsolved problem at commercial scale[54]. The geopolitical register of the choice is sharper than the supply-chain register alone: ^{3}H is on the Nuclear Suppliers Group trigger list[55], and a commercial ^{2}H–^{3}H program is by construction a safeguarded program adjacent to the nuclear-weapons material space. The mimetic acquisition logic Thiel develops in his treatment of nuclear weapons applies to the fuel-cycle choice as well as to the warhead[56] ("The Straussian Moment," pp. 211–212): a fuel-cycle commitment that places a commercial fusion industry inside the global safeguards regime exports the architecture of suspicion to every program that elects to follow.

The deuterium–helium-3 cycle, ^{2}H–^{3}He, has an order-of-magnitude lower reactivity than ^{2}H–^{3}H[52]; its primary reaction products are charged (a proton and an alpha particle, ^{4}He), which is the basis for the common claim that the cycle is "aneutronic." That claim carries two qualifications, both well-established in the plasma-physics literature. First, ^{3}He does not exist as a commercial commodity at the scale a fusion industry would consume it. Terrestrial supplies originate almost entirely as a by-product of tritium decay in the United States nuclear-weapons stockpile, managed by the National Nuclear Security Administration and distributed through the Department of Energy Isotope Program at annual production measured in tens of kilograms[57–59]. The United States Government Accountability Office characterized the supply posture as a critical shortage as recently as 2011[59]; the Isotope Program's own description of the situation makes clear that the program is sized to meet federal scientific and national-security demands rather than commercial fusion-fuel demand[58]. Lunar regolith contains ^{3}He at concentrations of order parts per billion, but no program has demonstrated mining, extraction, processing, or return to Earth at any scale; lunar ^{3}He is not a credible near-term supply pathway. Second, a deuterium- bearing plasma will always produce neutrons from ^{2}H–^{2}H side reactions and from the ^{2}H–^{3}H burn-up of the tritium those side reactions produce; the neutron yield of a ^{2}H–^{3}He plasma is reduced relative to ^{2}H–^{3}H but is not zero, and the engineering implications of the neutron channel — shielding, activation, dose envelope, and service envelope — do not disappear at the reduced yield.

The remaining proposed pathway is to breed ^{3}He in fusion machines themselves, by running an earlier ^{2}H–^{2}H step whose neutron-producing branch yields ^{3}He, capturing that ^{3}He, and re-injecting it into a ^{2}H–^{3}He burner. Helion Energy has publicly described such an approach[60], and Helion has separately announced a commitment to deliver electricity to Microsoft from its first fusion power plant by 2028[61]. The architectural proposal is coherent. It is also, by construction, a fusion problem that no program has publicly demonstrated at the required rate: the ^{3}He yield per pulse, the capture and separation efficiency, the recycle inventory, and the energy balance of the combined ^{2}H–^{2}H and ^{2}H–^{3}He cycle are all open. A program that commits its commercial schedule to a fuel-cycle architecture which itself contains an unsolved fusion problem is, by construction, taking on the harder of the two terminal-fuel commitments while presenting it as the easier.

The proton–boron-11 cycle, p–^{11}B, sits several orders of magnitude below ^{2}H–^{2}H in reactivity at accessible temperatures, and its operating ratio of fusion to bremsstrahlung loss remains below unity for any temperature accessible to near-term thermal magnetic-confinement devices[53,62]. The cycle remains theoretically attractive for its near-zero neutron channel, but operationally it remains out of reach for a near-term compact machine.

The deuterium–deuterium cycle, ^{2}H–^{2}H, is substantially harder than ^{2}H–^{3}H at any given temperature. The Lawson-product penalty is the price of the choice and we treat it explicitly in §V. The two branches of roughly equal probability are

$$\begin{aligned} {}^2\mathrm{H} + {}^2\mathrm{H} &\to {}^3\mathrm{H}\,(1.01\,\mathrm{MeV}) + p\,(3.02\,\mathrm{MeV}) && (\approx 50\%),\\ {}^2\mathrm{H} + {}^2\mathrm{H} &\to {}^3\mathrm{He}\,(0.82\,\mathrm{MeV}) + n\,(2.45\,\mathrm{MeV}) && (\approx 50\%); \end{aligned}$$

the prompt neutron channel from the second branch carries approximately a third of the per-reaction-pair energy. If branch-product tritium is allowed to burn in-machine through $^2\mathrm{H} + {}^3\mathrm{H} \to {}^4\mathrm{He} + n(14.1\,\mathrm{MeV})$, the in-machine secondary neutron channel hardens, and a public reactor-physics analysis of ^{2}H–^{2}H operation with and without tritium removal shows that controlled removal of the bred ^{3}H reduces the fraction of fusion energy carried by neutrons in a ^{2}H–^{2}H system by approximately 25%–35% at the displacements-per-atom level[63].

The structural advantage of the ^{2}H–^{2}H commitment over the two alternatives that satisfy the firm-power deployment window admits a clean statement. ^{2}H–^{3}H operation places a fusion industry inside the inventory posture of a safeguarded nuclear material on the Nuclear Suppliers Group trigger list[55]; tritium production at commercial fusion scale is unsolved, and the existing global production base is measured in kilograms[54]. ^{2}H–^{3}He operation places the industry inside the National Nuclear Security Administration's allocation of ^{3}He, which is rationed at tens of kilograms per year[57–59]. A ^{2}H–^{2}H architecture, by contrast, sources its terminal fuel from the deuterium content of ordinary seawater — approximately one part in 6,420 by atom, in absolute quantities that exceed the energy content of all known fossil- fuel reserves by several orders of magnitude — under ordinary Bureau of Industry and Security dual-use export control[64], the deuterium having moved from Nuclear Regulatory Commission licensing into ordinary dual-use export control in October 2021. The ^{2}H–^{2}H commitment is the only candidate fuel cycle on the public technical menu whose supply chain passes through no other sovereign's strategic-inventory posture, and that property is no longer optional. It is the property required by the political surface on which a national fusion program is now operating, and it is the property that closes the loop with the Hotelling- backstop condition of Eq. (2): the Hotelling rent of the counter-Lockean substrate is, structurally, zero.

What follows is the architectural question. The Lawson-product penalty the ^{2}H–^{2}H commitment pays is real, it is roughly two orders of magnitude in the relevant ratio, and any architectural argument for ^{2}H–^{2}H has to show how the rest of the machine closes the gap. That argument is the work of §V.

V. THE ARCHITECTURAL COMMITMENTS

The deployment window of §I, the political surface of §II, the nuclear-category arguments of §III, and the fuel-cycle argument of §IV together select for a specific combination of architectural commitments. Four commitments compose, and we argue each in turn: a compact pulsed field-reversed configuration as the confinement geometry; the Lawson-gap closure that the pulsed- recovery framework makes available; direct electromagnetic conversion as the primary energy-recovery channel; and a container-class packaging envelope. The combination is not the only architecture compatible with the political diagnosis — the diagnosis selects for the category of a Lockean backstop substrate, not for any single embodiment of it — but it is one combination that the public physics, the public regulatory surface, and the public deployment record together support, and it is the one Laurelin Technologies is pursuing. We make the argument as a category, and we make it quantitatively wherever the public literature permits.

A. Compact pulsed field-reversed configuration

The field-reversed configuration (FRC) is a compact toroid without a central rod, in which the closed-flux plasma volume is sustained by its own poloidal current rather than by

2026-05-18 03:12 UTC - IP 216.73.216.54 - Columbus, OH, US (39.9587, -82.9987) - EVENT c944871f-2d71-4293-9e44-ed1da55d6279

externally imposed toroidal field. The configurational physics is well-established: Tuszewski's foundational review[65] and Steinhauer's 2011 restatement[66] set out the equilibrium pressure-balance relation against which any FRC operating point must be read. For an FRC in a flux-conserving coil, the volume-averaged plasma beta is

$$\langle\beta\rangle \;=\; 1 - \tfrac{1}{2}\left(r_s/r_c\right)^2, \tag{3}$$

where r_s is the separatrix radius and r_c is the radius of the flux-conserving coil. In the canonical operating range $r_s/r_c \approx 0.6$–0.7 documented across the LSX and TCS programs, Eq. (3) gives $\langle\beta\rangle_{\mathrm{FRC}} \approx 0.75$–$0.82$, against tokamak operating betas $\langle\beta\rangle_{\mathrm{tok}} \lesssim 0.05$–$0.10$. For a fixed magnetic pressure $B^2/(2\mu_0)$, the plasma pressure available to confine fuel scales linearly with $\langle\beta\rangle$, so a configuration in which $\langle\beta\rangle$ is an order of magnitude higher than the tokamak operating range admits either an order-of-magnitude lower B at the same plasma pressure or an order-of-magnitude smaller device volume at the same B and the same plasma content. The compactness commitment of this paper is therefore not an aspiration that the device might be made small; it is a configurational property of the field- reversed equilibrium of Eq. (3), of which the container-class envelope of §V D is the engineering expression.

Pulsed operation, in this architectural choice, is not a provisional posture pending later steady-state engineering. It is the operating mode the field-reversed equilibrium most naturally supports at compact scale, and it is the mode that makes the per-pulse energy budget — rather than the steady- state thermal power — the operating figure of merit. The public engineering literature on pulsed FRC formation, merging, and compression — from Rej's foundational FRX-C work[67] through Slough's supersonic merging-and-compression program[68] to the C-2W operational record at TAE[69,70] — is the basis on which the architectural commitment to pulsed operation is publicly supportable. The architectural payoff for accepting the pulsed mode is the analytical decomposition of §V B: the binding constraint at the operating point shifts from a steady-state diagnostic figure of merit, on which the ^{2}H–^{2}H Lawson penalty is paid, to a per-pulse figure of merit on which three independent measurable levers appear that the steady-state architecture does not have.

B. Lawson-gap closure: ^{2}H–^{2}H penalty and engineering-Q

The serious technical reader will observe that the ^{2}H–^{2}H commitment carries a quantifiable Lawson-product penalty relative to ^{2}H–^{3}H at the same temperature, and will ask how the rest of the architecture proposes to close that gap. The penalty is direct to state and easy to misread. Read off the published triple-product curves of Wurzel and Hsu[53] at the temperature of minimum requirement for each cycle ($T_i \approx 14\,\text{keV}$ for ^2H–^3H, $T_i \approx 30\,\text{keV}$ for ^2H–^2H), the ratio of minima is

$$\frac{(n_i T_i \tau_E)_{\text{min}}^{^2\text{H}^2\text{H}}}{(n_i T_i \tau_E)_{\text{min}}^{^2\text{H}^3\text{H}}} \;\sim\; \mathcal{O}(10^2), \tag{4}$$

consistent across the bremsstrahlung treatments and impurity assumptions surveyed in the same review and consistent with the standard textbook rendering of the criterion[62]. The honest answer to the gap is that it is not closed on the steady-state Q axis on which Eq. (4) is conventionally drawn. It is shifted to a different axis.

A pulsed compact machine that recovers energy electromagnetically at a protected boundary satisfies, per pulse, the net-energy condition

$$\eta_{\text{rec}}\, E_{\text{fus}}^{\text{pulse}} \;\geq\; E_{\text{drv}}^{\text{pulse}} + P_{\text{par}}/f_{\text{rep}},$$

with $E_{\text{fus}}^{\text{pulse}}$ the per-pulse fusion energy released, $E_{\text{drv}}^{\text{pulse}}$ the per- pulse driver electrical energy, η_{rec} the protected-boundary recovery ratio, f_{rep} the steady repetition rate, and P_{par} the parasitic standing load. Dividing through defines the dimensionless engineering- gain inequality

$$Q_{\text{eng}} \;\equiv\; \frac{\eta_{\text{rec}}\, E_{\text{fus}}^{\text{pulse}}}{E_{\text{drv}}^{\text{pulse}}} \;\geq\; 1 + \frac{P_{\text{par}}}{f_{\text{rep}}\, E_{\text{drv}}^{\text{pulse}}}, \tag{5}$$

which is the constraint a credible pulsed-recovery architecture must satisfy at the protected boundary, on each shot. Two structural points follow. First, the right-hand side relaxes monotonically as f_{rep} rises: $\partial Q_{\text{eng}}^{\text{min}}/\partial f_{\text{rep}} < 0$. The rep-rate lever is the lever a pulsed compact architecture holds against the steady-state Q comparison, and it is the lever a quasi-steady architecture does not. Second, the inequality is separable across three independent engineering quantities — η_{rec}, $E_{\text{fus}}^{\text{pulse}}/E_{\text{drv}}^{\text{pulse}}$, and f_{rep} — each measurable at the protected boundary per shot, rather than inferred from a steady-state plasma diagnostic chain. The steady-state Q measured against the Eq. (4) penalty is therefore not the binding constraint for this class of machine; the binding constraint is Eq. (5), on a different axis from the one on which the $\mathcal{O}(10^2)$ penalty is paid.

The architectural claim is not that the ^2H–^2H Lawson penalty is illusory; the penalty is real and is paid in plasma- performance currency. The architectural claim is that the penalty is paid on one axis and the figure of merit is delivered on another, and that the three levers of Eq. (5) compound favorably in a pulsed compact machine in ways they do not in a quasi-steady utility- scale machine. The composition of the four architectural commitments is the closure: container-class packaging shifts the relevant figure of merit from gigawatt-class plant economics to per-unit per-pulse audit economics on a deployment timescale measured in years rather than decades; pulsed merge-and- compression operation replaces the sustainment problem with a per-pulse problem; direct electromagnetic recovery makes η_{rec} the primary lever rather than an afterthought; and the ^{2}H–^{2}H fuel-cycle commitment makes the supply-chain posture independent of any other sovereign's strategic inventory. The four levers are independent; each can be improved without retiring the others; and the gain in Eq. (5) is multiplicative across them.

C. Direct electromagnetic recovery

A fusion machine that recovers its energy thermally is a heat engine. Its architectural conversation is dominated by working fluid, steam-cycle efficiency, turbine envelope, and condenser loop; the reactor itself is one block in a balance-of-plant diagram dominated by mature thermal-cycle hardware. A fusion machine that recovers its energy electromagnetically is a different kind of machine. The architectural conversation is dominated by takeoff topology, switch and converter electronics, pulsed-power protection, and the question of whether the recovered energy can be measured cleanly on each pulse. For a compact pulsed field-reversed configuration committed to deuterium as the terminal fuel, the recovery commitment that matches the architecture is to direct electromagnetic conversion as the primary channel, with thermal capture acknowledged as a necessary complement for energy not directly recoverable.

The case for the choice begins with the thermodynamic ceiling. A working-fluid thermal cycle is bounded above by the Carnot ratio $\eta_{\mathrm{th}}^{\max} = 1 - T_c/T_h$. At $T_c \approx 300\,\mathrm{K}$ and $T_h \approx 900\,\mathrm{K}$ (supercritical- CO_2 class), $\eta_{\mathrm{th}}^{\max} \approx 0.67$, with installed steam-cycle and sCO_2-cycle plant efficiencies sitting in the 0.35–0.45 band of mature commercial practice. A multi-stage electrostatic or inductive direct converter, by contrast, is bounded above by the spectrum-to-collector matching identity[71,72]: $\eta_{\mathrm{dc}}^{\max} = 1 - \langle E_{\min}\rangle/\langle E\rangle$, where $\langle E_{\min}\rangle$ is the mean energy below which charged products cannot be decelerated in the highest collector stage of the chosen topology and $\langle E\rangle$ is the mean charged- product energy at the converter interface. The direct-conversion ceiling is not Carnot-bounded, because the relation does not require a working-fluid loop and is not capped by the hot-source temperature that the materials can tolerate; it is bounded only by the spectrum-to-collector matching identity, and tightens as the converter resolves a larger fraction of the spectrum into addressable stages. The Barr–Moir tandem-mirror reactor study projected $\eta_{\mathrm{dc}} \geq 0.70$ for a sensibly-staged collector against an experimentally demonstrated single-stage value of ~ 0.48.

Three operational pressures bias a compact pulsed ^{2}H–^{2}H machine toward direct electromagnetic conversion. The first is that the per-pulse energy arrives intermittently by construction. The compression, burn, and expansion phases deposit energy at the converter interface on timescales of microseconds to milliseconds; a thermal-cycle architecture would have to absorb this in a working-fluid loop whose thermal time constant is orders of magnitude longer, with all of the buffering, instrumentation, and inertia that follow. Direct conversion couples to the per-pulse energy in its natural time domain and leaves the working fluid out of the per-pulse accounting entirely. The second pressure is auditability. A compact commercial machine deploying behind the meter at a single host site must publish energy-in and energy-out per pulse with a clean enough instrument record for a counterparty's technical reviewer to accept; a working-fluid loop adds its own losses, instrumentation, and integration time, while a direct-conversion takeoff places the measurement closer to the plasma boundary and closer in time to the pulse itself. Direct conversion is, in this sense, an evidence-architecture choice as much as a thermodynamic one. The third pressure is the packaging envelope. Container-class packaging, as argued in §V D, does not accommodate the working-fluid volume that a steam-cycle plant of comparable thermal capacity would require, and the balance-of-plant volume that a thermal loop occupies is the volume an electromagnetic-conversion takeoff does not need.

The compositional argument with §V B is the load-bearing one. The engineering-Q inequality of Eq. (5) has η_{rec} as a primary lever, and direct electromagnetic recovery is what makes η_{rec} a recoverable physical channel for a pulsed compact device, not an asymptote. A program that commits to a thermal-cycle balance-of-plant is committing to make η_{rec} a Carnot-bounded asymptote on the secondary side; a program that commits to direct conversion as the primary channel is committing to make η_{rec} a measured recovery ratio at the protected boundary, with uncertainty bounds and off-nominal behavior recorded. The two commitments produce different evidence ledgers and different deployment timelines, and the second is the one the architectural argument of this paper requires.

D. Container-class packaging

The fourth architectural commitment is to a container-class packaging envelope: reactor-core hardware fitting within a 40-foot shipping-container envelope, with balance of plant deployed as adjacent modules. Container-class hardware is the largest unit movable over ordinary roadway, rail, barge, or transport aircraft without bespoke infrastructure; sizing the reactor core to the container envelope makes the machine deployable inside the infrastructure that already exists, rather than requiring new infrastructure to be built around each deployed unit. The footprint is of order tens of meters per side, not hundreds, and the vertical envelope is set by ordinary industrial-buildings access.

The post-2024 procurement vehicles have converged on the envelope. The Department of Defense Strategic Capabilities Office's Project Pele — 1–5 MWe in four 20-foot ISO containers, 72-hour setup, three-year fuel cycle — is the configuration BWXT began manufacturing in July 2025[49]; the Defense Innovation Unit's Advanced Nuclear Power for Installations program selected eight companies in April 2025 for fixed on-site microreactors of 3–10 MW per United States Army installation[50]. The two vehicles together delineate a small, defined market — forward operating bases, fixed installations, hyperscaler campuses, remote industrial sites — in which the architectural envelope of this paper is the unit of deployment, and in which container-class fusion inherits the operational, regulatory, and supply-chain architecture that container-class fission has now publicly established. The envelope matches the Agreement-State licensing surface of §III: a forty-foot transportable neutron-producing device is, under the post-ADVANCE Act framework, licensed under 10 CFR Part 30 and inspected on the schedule of an industrial-irradiation facility, not a utility power plant. The strategic case for entering through this market is the one Thiel develops in *Zero to One*[73]: a small market one architecture can credibly dominate is the right place from which to build the operational record, the regulatory experience, and the supply chain on which subsequent expansion depends.

The four commitments compose. ^{2}H–^{2}H as the terminal fuel sets the supply-chain posture independent of any other sovereign's strategic inventory. Compact pulsed FRC operation delivers the high-β equilibrium that makes the container- class envelope a configurational property rather than an aspiration. Direct electromagnetic recovery makes the η_{rec} lever of Eq. (5) a measured channel at the protected boundary. Container-class packaging matches

2026-05-18 03:12 UTC - IP 216.73.216.54 - Columbus, OH, US (39.9587, -82.9987) - EVENT c944871f-2d71-4293-9e44-ed1da55d6279

the Agreement-State regulatory surface and the post-2024 procurement vehicles. No other combination on the public technical menu delivers all four properties simultaneously. The next section places the combination on the architectural map of the wider field.

VI. THE FUSION LANDSCAPE AND WHERE LAURELIN SITS

The architectural argument of §V is sharpest when it is placed on the map of the wider field. We treat the landscape at the level of six public coordinates: confinement geometry, fuel cycle, energy-recovery posture, packaging and deployment envelope, operating mode, and evidence maturity. The Fusion Industry Association's annual industry report[74] provides the current inventory; the grouping below is by architectural family rather than by company, because the architectural family is what sets the deployment shape and is the relevant unit of comparison. The evidence-maturity column is separate from the first five coordinates and reflects each program's public operating record at the time of writing.

The tokamak family — ITER, SPARC, Tokamak Energy — is the most mature confinement geometry, but its packaging is large by construction, its siting is utility-scale by design, and its supply-chain exposure is dominated by the tritium-fuel question of §IV. The stellarator family — Wendelstein 7-X as the operating anchor, with Type One Energy, Proxima Fusion, and Renaissance Fusion as the private programs — pays the plasma-physics advantages in coil geometry, the hardest manufacturing problem in the family. The inertial-confinement family was rebased in December 2022 by the National Ignition Facility shot; translating the result into a commercial pulse train remains an open problem (Focused Energy, Xcimer Energy, Marvel Fusion, First Light Fusion). The compact-toroid family is where the commitments of this paper sit; public programs include TAE Technologies (beam-driven FRC, p–^{11}B long-term), Helion Energy (merging-and-compressing FRC on ^{2}H–^{3}He with bred ^{3}He), Zap Energy (sheared-flow Z-pinch), and General Fusion (magnetized-target compression of FRCs in liquid metal).

Table I maps the same programs against the six public coordinates and shows the position the architectural commitments of §V occupy. The entries trace to each program's own public statements, public patents, and the Fusion Industry Association inventory[74]; the maturity column reflects the program's own public operating record at the time of writing and is given for orientation, not as an investment metric.

TABLE I: Public architectural coordinates of selected fusion programs. Each entry traces to the program's own public statements; the maturity column reflects the public operating record and is given for orientation. Laurelin is included for explicit positioning: distinct on at least three of the six coordinates from every other entry. Boldface denotes the Laurelin row, marking the architectural-positioning entry. The table is not a ranking and is not predictive of commercial outcome.

Program	**Geometry**	**Fuel**	**Recovery**	**Packaging**	**Mode**	**Public evidence maturity**
ITER	tokamak	^{2}H–^{3}H	thermal	utility-scale facility	quasi-steady	first-plasma campaign in preparation; large physics database
SPARC (CFS)	HTS tokamak	^{2}H–^{3}H	thermal	utility-scale facility	quasi-steady	full-scale TF magnet validated; SPARC under construction
Tokamak Energy	spherical tokamak (HTS)	^{2}H–^{3}H	thermal	utility-scale facility	quasi-steady	HTS coil program; spherical-tokamak demonstrators
W7-X	stellarator	^{2}H	thermal	large facility	quasi-steady	high-confinement steady-state record; mature operating record
NIF	laser ICF	^{2}H–^{3}H	thermal	national-lab facility	single-shot	ignition demonstrated 2022; rep-rate gap

2026-05-18 03:12 UTC - IP 216.73.216.54 - Columbus, OH, US (39.9587, -82.9987) - EVENT c944871f-2d71-4293-9e44-ed1da55d6279

Program	Geometry	Fuel	Recovery	Packaging	Mode	Public evidence maturity
TAE	beam-driven FRC	p–^{11}B (long-term); ^{2}H–^{2}H/^{2}H–^{3}H steps	thermal/cha product	facility-scale	quasi-steady	hot stable beam-driven FRC (C-2W); next machine in development
Helion	merging/com FRC	^{2}H–^{3}He (with ^{3}He from earlier ^{2}H–^{2}H step)	direct EM	facility-scale	pulsed	pulsed FRC merging/compression; named PPA milestones[60,61]
Zap Energy	sheared-flow Z-pinch	^{2}H–^{3}H (long-term); ^{2}H–^{2}H early	thermal	facility-scale	pulsed	FuZE-class devices; sheared-flow stabilization record
General Fusion	magnetized target (compressed FRC in liquid metal)	^{2}H–^{3}H	thermal	facility-scale	pulsed	LM26 demonstrator program
Laurelin	**compact pulsed FRC**	**^{2}H–^{2}H (terminal)**	**direct EM (primary)**	**container-class**	**pulsed**	**concept-stage; non-nuclear bench evidence to be developed**

Laurelin is distinct on at least three of the six coordinates from every other public program. The fuel-cycle coordinate places Laurelin alone in the ^{2}H–^{2}H terminal column: no other program commits to ^{2}H–^{2}H as the operating end-point, distinct from a stage on the way to another cycle. The recovery coordinate places Laurelin in the direct-EM primary column with Helion as the only other entry; the architectures diverge sharply at the fuel-cycle coordinate, where Helion's ^{2}H–^{3}He-with-bred-^{3}He commitment is the architectural decision Laurelin's ^{2}H–^{2}H commitment was made to avoid. The packaging coordinate places Laurelin alone in the container-class column, distinct from the facility-scale envelope every other program occupies; the architectural commitment to container-class is the commitment that aligns the deployment unit with the Project Pele and DIU ANPI procurement envelopes of §V D. The mode and geometry coordinates place Laurelin in the pulsed-FRC family with Helion. The evidence-maturity coordinate is the most honest entry for Laurelin: the program is concept-stage, and the non-nuclear bench evidence on which the architectural claims will ultimately stand or fall is to be developed.

The combination is the architectural claim. No other public program occupies the coordinate set {compact pulsed FRC, ^{2}H–^{2}H, direct EM, container-class}, and no other coordinate set delivers the four properties the political and architectural arguments of §§II–V together require. The combination is not historically obvious: the public-literature record on each coordinate is older than most operating programs, but the composition of the four commitments into one architectural envelope has not been publicly pursued at the program level. The Thiel test that distinguishes a new category from a relocation of the rivalry inside an existing category is the test the combination is built to pass: ^{2}H–^{2}H direct-conversion compact pulsed FRC is a category that does not yet exist publicly, not a relocation of the rivalry. The table is not a ranking. It is a positioning, and the positioning is the architectural claim.

VII. RISK REGISTER AND OPEN PROBLEMS

The architectural commitments of §V bring a particular shape of risk surface, and the case for those commitments is incomplete without an explicit treatment of it. We name the risks at the architectural level, as risks any program of this shape must confront, and we name the categories of mitigation available in the public literature. Specific embodiments inside any given machine are not the subject of this paper.

The first operational risk is long-pulse and high-duty-cycle stability. The compact pulsed

FRC operating record in the public literature is dominated by short-pulse and low-duty-cycle campaigns, from the LSX program through the C-2W operational record at TAE[65,66,69,70]. The behavior of an FRC equilibrium under sustained repetition at the engineering-Q inequality of Eq. (5) — including the question of whether the high-β equilibrium of Eq. (3) is reproducible at the rep-rate the deployment economics require — is the central open plasma- physics question for any program of this architectural shape. The categories of mitigation are well-established in the public literature on FRC stabilization (rotating magnetic field, beam sustainment, sheared-flow effects); the binding constraint is the public operating record at the required cadence, and that is the evidence artifact the field has not yet been required to produce.

The second operational risk is plasma-facing materials under the ^{2}H–^{2}H neutron spectrum. The 2.45 MeV prompt neutron channel and the secondary 14.1 MeV channel from in-machine burnup of $^2\mathrm{H}(^2\mathrm{H}, p)^3\mathrm{H}$-produced tritium together set the displacements-per-atom budget of the plasma-facing surface, the activation budget at the balance-of-plant boundary, and the service envelope on which the container-class architecture depends. The methodology is dense in the tokamak and DEMO neutronics literature[75–80] and transfers cleanly to a compact field-reversed configuration with the differences localized in geometry, spectrum, and dose-budget allocation rather than in the transport calculation. The risk surface is therefore characterizable in advance, and the evidence artifact for closure is a materials-qualification record on the operating spectrum, not the claim that the spectrum has been designed away.

The third operational risk is direct-conversion as a measured per-pulse channel. The engineering-Q inequality of Eq. (5) has η_{rec} as a primary lever, and the architectural case of §V C turns on η_{rec} being a measured recovery ratio at a protected boundary, with uncertainty bounds and off-nominal behavior recorded, rather than an asserted efficiency. The public-literature direct-conversion record is dominated by the Barr–Moir tandem-mirror tradition and the Japanese TWDEC experimental program[71,72,81–84]; the integration of direct conversion into a pulsed-FRC machine at engineering- relevant scale is the gap the field has not yet closed. The evidence artifact for closure is metered, protected, repeatable recovery in the pulse's native time domain. A recovery trace that destroys the recovery path is not a reactor-engineering result; a plasma-energy term inferred from a model is not the same object as recovered electrical work. The risk is that the integration work is harder than the family-level literature implies, and the evidence the field can demand of any program making the claim is the protected per-pulse measurement.

The political-surface open problem is the institutional architecture a coordinated allied fusion bloc requires. The post-ADVANCE Act regulatory rebase of §III is the first half; the international export-control, supply-chain- coordination, and alliance-architecture half is the open problem[51,85–88]. The closest comparison is the semiconductor industry of the past twenty-four months: the technology is necessary but not sufficient, and the strategic outcome is determined by the institutional architecture in which the technology is deployed. No program of this architectural shape can carry the institutional half by itself, and the technology choices of §V are necessary but not sufficient for the political outcome §II argues for.

The architectural thesis is empirically dispositive. The falsification conditions are direct to state. If the ^{2}H–^{2}H Lawson penalty of Eq. (4) cannot be shifted to the engineering-Q axis of Eq. (5) at engineering-relevant scale — if the three independent levers of the inequality do not compound in a pulsed compact machine the way the architectural argument requires — the architecture has failed and a different combination of commitments is the answer to the deployment window of §I. If the high-β equilibrium of Eq. (3) cannot be reproduced at the rep-rate the deployment economics require, the compact-FRC architectural family is the wrong family. If direct-conversion η_{rec} cannot be made measurable as a per-pulse channel at engineering relevance, the recovery commitment has failed and the architecture defaults to thermal capture with the Carnot ceiling reattached. If the container-class regulatory architecture is not written — if the Agreement-State licensing surface is not built out at the operational scale the post-ADVANCE Act framework now permits — the packaging commitment loses its deployment vehicle. Each of these conditions is testable on the public operating record. The architectural thesis stands or falls on that record, and the record is what the rest of the program is built to produce.

VIII. CONCLUSION

The argument of this paper closes as it opened. The petrodollar arrangement, on the orthodox reading of the political-economy literature, is an architecture of substantial net cost to the United States in the form of forward-deployed naval power, a binding ceiling on sanctions enforcement, and an active rent- financed channel that finances the adversaries the same architecture is then required to suppress; the benefits the arrangement is supposed to secure for the dollar's reserve role are, on the rigorous reading, secured by mechanisms that do not depend on it. The exit is not a policy reform; it is a substrate. The only category of substrate that delivers the property the political diagnosis requires is a Lockean backstop fuel, and the substrate that delivers the property on the public physics and the public regulatory surface is deuterium fusion with a ^{2}H– ^{2}H terminal-fuel commitment, recovered electromagnetically, and deployed in a container-class envelope inside a compact pulsed field-reversed configuration. Laurelin Technologies is one program pursuing this combination.

The analytical statement that ties the political and architectural halves of the paper together is direct. The Hotelling rent collapse of Eq. (2) binds the petrodollar architecture once a backstop substrate is available at constant unit cost $\bar{c}$; the political consequence is the dissolution of the rent-financed adversary loop $A(t)$ in the ledger of Eq. (1) and, through the political dynamics of the sanctions-enforcement ceiling and the alliance-protection commitment, the dissolution of $L(t)$ and the gradient of $\Pi(t)$. The condition under which the architectural commitments of §V satisfy the binding test is the backstop-milestone inequality

$$\bar{c} \;\leq\; \bar{c}^{\text{firm}} \;\equiv\; \min_{k\in\mathcal{F}} \text{LCOE}_k, \tag{6}$$

with $\mathcal{F}$ the set of publicly-reported firm dispatchable alternatives in the host market (advanced gas with carbon capture, advanced fission, advanced geothermal, and long- duration storage paired with non-firm renewables) and LCOE_k the levelized cost of electricity of alternative k as published in the Energy Information Administration's annual outlook[43]. The architectural claim of this paper is not that any particular program occupies any particular point inside the bracket of Eq. (6); it is that the category of generation defined by the commitments of §V is the category against which Eq. (6) is the binding test, and that satisfying the inequality is the necessary and sufficient public-record condition for the rent collapse of Eq. (2) to bind on the architecture of §II.

The decision the paper urges is, in the strict sense of the verb, definite. The architectural commitments are specifiable. The deployment timeline is compressible. The regulatory surface is built. The procurement vehicles are in place. The fuel-cycle choice closes the supply-chain posture independent of any other sovereign's strategic inventory. The recovery commitment makes the measurement architecture auditable per pulse. The packaging envelope matches the demand surface that is already chartered and already procuring. The four commitments are not options among many; they are the architecture the political and deployment surfaces together select for, and no other architecture on the public technical menu delivers all four properties simultaneously. The remainder of the work is to build the machine.

The first-named author is affiliated with Laurelin Technologies Inc. The program's public-facing material is available at `https://laurelin-inc.com`. The full-length companion paper, which carries the evidentiary surface, the primary-source engagement with the political-philosophy texts, the full public-physics scaling arguments, and the complete ten-section treatment, is available from the corresponding author. Qualified counterparties may request data-room access through the company website.

ACKNOWLEDGMENTS

The author thanks colleagues in the field-reversed configuration and pulsed-power communities for their reading of earlier drafts of the long-form paper from which this paper is condensed. Any errors of fact remain the author's.

REFERENCES

[1] D. E. Spiro, *The Hidden Hand of American Hegemony: Petrodollar Recycling and International Markets* (Cornell University Press, Ithaca, NY, 1999).

[2] B. Eichengreen, *Exorbitant Privilege: The Rise and Fall of the Dollar and the Future of the International Monetary System* (Oxford University Press, Oxford, 2011).

[3] E. Helleiner, *States and the Reemergence of Global Finance: From Bretton Woods to the 1990s* (Cornell University Press, Ithaca, NY, 1994).

[4] A. Tooze, *Crashed: How a Decade of Financial Crises Changed the World* (Viking, New York, 2018).

[5] S. Arslanalp, B. Eichengreen, and C. Simpson-Bell, "Dollar dominance in the international reserve system: An update," IMF Blog (2024), u.S. dollar share of allocated foreign exchange reserves is 58 % in 2024-Q1, the lowest since 1995. Accessed 2026-05-14.

[6] International Monetary Fund, Statistics Department, "Currency composition of official foreign exchange

reserves (COFER)," Tech. Rep. (International Monetary Fund, 2024) quarterly survey of central-bank foreign exchange reserve currency composition. Accessed 2026-05-14.

[7] International Energy Agency, "Electricity 2024 — analysis and forecast to 2026," Tech. Rep. (International Energy Agency, Paris, 2024) accessed 2026-05-14.

[8] International Energy Agency, "Energy and AI," Tech. Rep. (International Energy Agency, Paris, 2025) base-case data-centre electricity demand projection 460 TWh (2024) → 1 000 TWh (2030) → 1 300 TWh (2035). Accessed 2026-05-14.

[9] J. Rand, M. Bolinger, R. Wiser, S. Jeong, and W. Gorman, "Queued up: 2025 edition — characteristics of power plants seeking transmission interconnection as of the end of 2024," Tech. Rep. (Lawrence Berkeley National Laboratory, 2025) active interconnection-queue capacity totalled 2,289 GW at the end of 2024. Accessed 2026-05-14.

[10] Federal Energy Regulatory Commission, Office of Energy Policy and Innovation, "2024 state of the markets," Tech. Rep. (Federal Energy Regulatory Commission, Washington, DC, 2025) accessed 2026-05-14.

[11] Microsoft Corporation, "2024 environmental sustainability report," Tech. Rep. (Microsoft Corporation, 2024) scope 1+2+3 emissions up 29.1 % from 2020 baseline, primarily from datacentre construction and operation. Accessed 2026-05-14.

[12] A. Wong, "The untold story behind Saudi Arabia's 41-year U.S. debt secret," Bloomberg News (2016), accessed 2026-05-14.

[13] H. Thompson, *Disorder: Hard Times in the 21st Century* (Oxford University Press, Oxford, 2022).

[14] B. R. Posen, International Security **28**, 5 (2003).

[15] J. Saballa and J. Detsch, "U.S. and Israel launch strikes on Iran," Foreign Policy (2026), operation Epic Fury: coordinated U.S.–Israeli air campaign 28 February 2026, 48 hours after the third round of Omani-mediated nuclear talks in Geneva. Supreme Leader Ali Khamenei killed in opening salvo. Accessed 2026-05-14.

[16] Congressional Research Service, "U.S. strikes on nuclear sites in Iran," CRS In Focus IN12571 (2025).

[17] CNN, "Iran begins laying mines in strait of Hormuz," CNN (2026), accessed 2026-05-14.

[18] Vortexa, "Hormuz strait special briefing, June 2025," Vortexa report (2025), strait flows for the week ending 22 June 2025 ran 4 % above the 2024 average despite the active war; transit collapse to a single vessel on 3 March 2026 documented separately. Accessed 2026-05-14.

[19] U.S. Energy Information Administration, "Short-term energy outlook, May 2026," EIA (2026), april 2026 production shut-ins of 10.5 Mbbl/d across Iraq, Saudi Arabia, Kuwait, UAE, Qatar, and Bahrain attributable to the Strait of Hormuz closure. Brent forecast central case assumes constrained throughput

through late May; full restoration not expected until late 2026. Accessed 2026-05-14.

[20]World Bank, "Strait of Hormuz disruption sends oil prices surging," World Bank Blogs / Commodity Markets Outlook (2026), april 2026 *Commodity Markets Outlook* characterizes the disruption as "the largest oil market shock in history." Accessed 2026-05-14.

[21]U.S. Energy Information Administration, "The strait of Hormuz remains the world's most important oil chokepoint," EIA Today in Energy (2025), 2024 average flow through Hormuz: 20 Mbbl/d, ∼20 % of global petroleum-liquids consumption. Accessed 2026-05-14.

[22]Iran International, "Iran's parliament approves resolution to close strait of Hormuz; decision lies with supreme national security council," Iran International (2025), the Supreme National Security Council declined to execute the closure during the June 2025 conflict. Accessed 2026-05-14.

[23]U.S. Department of Defense, "Historically successful strike on Iranian nuclear site was 15 years in the making," DoD News (2025), operation Midnight Hammer: 14 GBU-57A/B Massive Ordnance Penetrators delivered by B-2 Spirit bombers against Fordow, Natanz, and Isfahan; supplemented by Tomahawk strikes. Accessed 2026-05-14.

[24]C. Talmadge, International Security **33**, 82 (2008), canonical operational assessment: Iran can temporarily close the Strait through mine warfare and anti-ship missiles, but the United States and allies can reopen it within weeks at bounded cost.

[25]D. S. Eady, S. B. Siegel, R. S. Bell, and S. H. Dicke, "Sustain the mission project: Casualty factors for fuel and water resupply convoys — final technical report," Tech. Rep. (U.S. Army Environmental Policy Institute, Arlington, VA, 2009) casualty-rate estimates for fuel- and water-resupply convoys in Iraq and Afghanistan. Accessed 2026-05-14.

[26]U.S. Department of Defense, "Energy for the warfighter: Operational energy strategy," Tech. Rep. (Office of the Assistant Secretary of Defense for Operational Energy Plans and Programs, 2011) accessed 2026-05-14.

[27]R. A. Pape, International Security **22**, 90 (1997).

[28]D. W. Drezner, *The Sanctions Paradox: Economic Statecraft and International Relations* (Cambridge University Press, Cambridge, 1999).

[29]H. Farrell and A. L. Newman, International Security **44**, 42 (2019).

[30]B. Hilgenstock, E. Ribakova, T. Babina, O. Itskhoki, and M. Mironov, "Russian oil exports under international sanctions," KSE Institute / SSRN 4430053 (2023).

[31]A. Demarais, "The limits of the effectiveness of EU sanctions on Russia," Bruegel commentary (2023), accessed 2026-05-14.

[32] M. L. Ross, *The Oil Curse: How Petroleum Wealth Shapes the Development of Nations* (Princeton University Press, Princeton, NJ, 2012).

[33] M. L. Ross, Annual Review of Political Science **18**, 239 (2015).

[34] J. J. Andersen, N. Johannesen, D. D. Lassen, and E. Paltseva, Journal of the European Economic Association **15**, 818 (2017).

[35] H. Mehlum, K. Moene, and R. Torvik, Economic Journal **116**, 1 (2006).

[36] International Energy Agency, "Energy fact sheet: Why does Russian oil and gas matter?" IEA commentary (2022), oil and gas revenues accounted for 45 % of Russia's federal budget in 2021. Accessed 2026-05-14.

[37] C. Schmitt, *The Concept of the Political: Expanded Edition* (University of Chicago Press, Chicago, 2007) original German edition 1932; expanded English edition includes "The Age of Neutralizations and Depoliticizations" and Leo Strauss's notes.

[38] L. Strauss, in *The Concept of the Political: Expanded Edition*, edited by C. Schmitt (University of Chicago Press, Chicago, 2007) pp. 97–122, original German essay published 1932.

[39] L. Strauss, *Natural Right and History* (University of Chicago Press, Chicago, 1953).

[40] H. Hotelling, Journal of Political Economy **39**, 137 (1931).

[41] W. D. Nordhaus, Brookings Papers on Economic Activity **4**, 529 (1973).

[42] P. Dasgupta and G. Heal, *Economic Theory and Exhaustible Resources* (Cambridge University Press, Cambridge, 1979).

[43] U.S. Energy Information Administration, "Annual energy outlook 2024 — levelized costs of new generation resources," Tech. Rep. (U.S. Energy Information Administration, Washington, DC, 2024) tables of levelised cost of electricity (LCOE) and levelised cost of storage (LCOS) for new generation resources entering service in 2026, 2028, and 2035. Accessed 2026-05-15.

[44] United States Congress, "Accelerating deployment of versatile, advanced nuclear for clean energy (ADVANCE) act of 2024," Public Law 118-67; signed 9 July 2024. Section 205 amends the Atomic Energy Act to add a statutory definition of "fusion machine" and direct the NRC to regulate near-term fusion devices under the byproduct-material framework rather than utility-reactor authority. (2024), accessed 2026-05-14.

[45] U.S. Nuclear Regulatory Commission, "Fusion and the ADVANCE act — Section 205 implementation," NRC Office of Nuclear Material Safety and Safeguards (2024), accessed 2026-05-14.

[46] U.S. Nuclear Regulatory Commission, "Licensing and regulating byproduct material for fusion energy systems — SECY-23-0001," NRC Commission paper and 2024 implementation guidance (2023), NRC 2023 decision to regulate near-term fusion devices under 10 CFR Part 30 (byproduct-material framework,

Agreement-State eligible) rather than 10 CFR Part 50 (utility-scale reactor framework). Accessed 2026-05-14.

[47]U.S. Nuclear Regulatory Commission, “Secy-24-0085 — proposed rule: Regulatory framework for fusion machines,” Transmittal of draft proposed rule to the Commission, with NUREG-1556 Vol. 22 program-specific licensing guidance (2024).

[48]Nuclear Innovation Alliance, “Regulatory implementation summary: NRC progress under the ADVANCE act,” (2025).

[49]U.S. Department of Defense, Strategic Capabilities Office, “Project Pele Site Investigation Report 3 (SIR-3): Idaho National Laboratory Critical Infrastructure Test Range Complex Pad a,” Transportable nuclear microreactor demonstration: 1–5 MWe, four 20-foot ISO containers, 72-hour setup, ≥3-year fuel cycle; BWXT core manufacturing began July 2025; first electricity targeted 2028 (2025).

[50]U.S. Defense Innovation Unit, “Advanced Nuclear Power for Installations (ANPI) Program: eight eligible companies selected,” DIU/Army Office of the Assistant Secretary for Installations, Energy & Environment Commercial Solutions Opening for fixed on-site microreactors at CONUS Army installations; 3–10 MW; NRC-licensed; full life-cycle owned and operated by industry; first operation before end of 2030; eight vendors selected 10 April 2025 (2025).

[51]R. Desai, S. Hua, J. Roma, B. Bufford, J. Siebens, and N. Proffitt, “Building a path toward global deployment of fusion: Nonproliferation and export considerations,” Atlantic Council issue brief (2025).

[52]H.-S. Bosch and G. M. Hale, Nuclear Fusion **32**, 611 (1992).

[53]S. E. Wurzel and S. C. Hsu, Physics of Plasmas **29**, 062103 (2022), arXiv:2105.10954.

[54]R. J. Pearson, A. B. Antoniazzi, and W. J. Nuttall, Fusion Engineering and Design **136**, 1140 (2018).

[55]Nuclear Suppliers Group, “Guidelines for nuclear transfers (NSG part 1),” INFCIRC/254/Rev.14/Part 1 (2024), tritium is listed in the trigger list as a nuclear material relevant to weapons usability. Accessed 2026-05-14.

[56]P. Thiel, in *Politics and Apocalypse*, Studies in Violence, Mimesis, and Culture, edited by R. Hamerton-Kelly (Michigan State University Press, East Lansing, MI, 2007) pp. 189–218.

[57]D. A. Shea and D. Morgan, “The helium-3 shortage: Supply, demand, and options for congress,” Tech. Rep. R41419 (Congressional Research Service, Washington, DC, 2010) accessed 2026-05-14.

[58]National Isotope Development Center, U.S. Department of Energy, “Supply and demand of helium-3 (He-3),” Tech. Rep. (U.S. Department of Energy, Isotope Program, 2024) accessed 2026-05-14.

[59]U.S. Government Accountability Office, “Managing critical isotopes: Weaknesses in DOE’s management

of helium-3 delayed the federal response to a critical supply shortage," Tech. Rep. GAO-11-472 (U.S. Government Accountability Office, Washington, DC, 2011).

[60] Helion Energy, "Our technology," Helion Energy website (2024), public description of Helion's deuterium–helium-3 fuel cycle with in-machine ^{3}He production via D–D. Accessed 2026-05-14.

[61] Helion Energy, "Helion announces world's first fusion energy purchase agreement with Microsoft," Helion press release (2023), "New facility aims to deliver at least 50 MW and begin producing electricity by 2028." Accessed 2026-05-14.

[62] S. Atzeni and J. Meyer-ter Vehn, *The Physics of Inertial Fusion: Beam-Plasma Interaction, Hydrodynamics, Hot Dense Matter* (Oxford University Press, 2004) standard reference on fusion cross-sections and reactivity for D–D, D–T, D–^{3}He, and advanced fuels.

[63] G. Schnabel, D. L. Aldama, T. Bohm, U. Fischer, S. Kunieda, A. Trkov, *et al.*, Nuclear Data Sheets **193**, 1 (2024), FENDL data releases: `https://www-nds.iaea.org/fendl/` (accessed 2026-04-23).

[64] U.S. Department of Commerce, Bureau of Industry and Security, "Control of deuterium that is intended for use other than in a nuclear reactor under the export administration regulations," Federal Register 86 FR 55492 (2021), NRC licensing of non-reactor-end-use deuterium exports transferred to BIS. Accessed 2026-05-14.

[65] M. Tuszewski, Nuclear Fusion **28**, 2033 (1988).

[66] L. C. Steinhauer, Physics of Plasmas **18**, 070501 (2011).

[67] D. J. Rej, "Design and construction details of the FRX-C/T device: a compact toroid plasma translation experiment," Tech. Rep. LA-10108-MS (Los Alamos National Laboratory, 1984).

[68] J. Slough, G. Votroubek, and C. Pihl, Nuclear Fusion **51**, 053008 (2011).

[69] H. Gota, M. Binderbauer, T. Tajima, S. Putvinski, M. Tuszewski, B. Deng, S. Dettrick, D. Gupta, S. Korepanov, R. Magee, T. Roche, J. Romero, A. Smirnov, V. Sokolov, Y. Song, L. Steinhauer, M. Thompson, E. Trask, A. Van Drie, X. Yang, P. Yushmanov, K. Zhai, I. Allfrey, R. Andow, E. Barraza, M. Beall, N. Bolte, E. Bomgardner, F. Ceccherini, A. Chirumamilla, R. Clary, T. DeHaas, J. Douglass, A. DuBois, A. Dunaevsky, D. Fallah, P. Feng, C. Finucane, D. Fulton, L. Galeotti, K. Galvin, E. Granstedt, M. Griswold, U. Guerrero, S. Gupta, K. Hubbard, I. Isakov, J. Kinley, A. Korepanov, S. Krause, C. Lau, H. Leinweber, J. Leuenberger, D. Lieurance, M. Madrid, D. Madura, T. Matsumoto, V. Matvienko, M. Meekins, R. Mendoza, R. Michel, Y. Mok, M. Morehouse, M. Nations, A. Necas, M. Onofri, D. Osin, A. Ottaviano, E. Parke, T. Schindler, J. Schroeder, L. Sevier, D. Sheftman, A. Sibley, M. Signorelli, R. Smith, M. Slepchenkov, G. Snitchler, J. Titus, J. Ufnal, T. Valentine, W. Waggoner, J. Walters, C. Weixel, M. Wollenberg, S. Ziaei,

L. Schmitz, Z. Lin, A. Ivanov, T. Asai, E. Baltz, J. Platt, and the TAE Team, Nuclear Fusion **59**, 112009 (2019).

[70]H. Gota, M. Binderbauer, T. Tajima, A. Smirnov, S. Putvinski, M. Tuszewski, S. Dettrick, D. Gupta, S. Korepanov, R. Magee, J. Park, T. Roche, J. Romero, E. Trask, X. Yang, P. Yushmanov, K. Zhai, T. DeHaas, M. Griswold, S. Gupta, S. Abramov, A. Alexander, I. Allfrey, R. Andow, B. Barnett, M. Beall, N. Bolte, E. Bomgardner, A. Bondarenko, F. Ceccherini, L. Chao, R. Clary, A. Cooper, C. Deng, A. Dunaevsky, P. Feng, C. Finucane, D. Fluegge, L. Galeotti, S. Galkin, K. Galvin, E. Granstedt, K. Hubbard, I. Isakov, M. Kaur, J. Kinley, A. Korepanov, S. Krause, C. Lau, A. Lednev, H. Leinweber, J. Leuenberger, D. Lieurance, D. Madura, J. Margo, D. Marshall, R. Marshall, T. Matsumoto, V. Matvienko, M. Meekins, W. Melian, R. Mendoza, R. Michel, Y. Mok, M. Morehouse, R. Morris, L. Morton, M. Nations, A. Necas, S. Nicks, G. Nwoke, M. Onofri, A. Ottaviano, R. Page, E. Parke, K. Phung, G. Player, I. Sato, T. Schindler, J. Schroeder, D. Sheftman, A. Sibley, A. Siddiq, M. Signorelli, M. Slepchenkov, R. Smith, G. Snitchler, V. Sokolov, Y. Song, L. Steinhauer, V. Stylianou, J. Sweeney, J. Titus, A. Tkachev, M. Tobin, J. Ufnal, T. Valentine, A. Van Drie, J. Ward, C. Weixel, C. White, M. Wollenberg, S. Ziaei, the TAE Team, L. Schmitz, Z. Lin, A. Ivanov, T. Asai, E. Baltz, M. Dikovsky, W. Heavlin, S. Geraedts, I. Langmore, P. Norgaard, R. Von Behren, T. Madams, A. Kast, and J. Platt, Nuclear Fusion **61**, 106039 (2021).

[71]W. L. Barr and R. W. Moir, in *Proceedings of the Second ANS Topical Meeting on the Technology of Controlled Thermonuclear Fusion* (Richland, WA, USA, 1976) lawrence Livermore Laboratory preprint UCRL-78204 (OSTI record: https://www.osti.gov/biblio/7341986).

[72]H. Takeno, K. Ichimura, S. Nakamoto, Y. Nakashima, H. Matsuura, J. Miyazawa, T. Goto, Y. Furuyama, and A. Taniike, Plasma and Fusion Research **14**, 2405013 (2019).

[73]P. Thiel and B. Masters, *Zero to One: Notes on Startups, or How to Build the Future* (Crown Business, New York, 2014).

[74]Fusion Industry Association, “The global fusion industry in 2024,” Tech. Rep. (Fusion Industry Association, 2024) cumulative private investment in fusion companies above $7.1B; 45 declared private fusion programs. Accessed 2026-05-14.

[75]J. T. Kriese and D. Steiner, “Magnet shield design for fusion reactors,” Tech. Rep. ORNL-TM-4256 (Oak Ridge National Laboratory, 1973).

[76]V. D. Lee and Y. Gohar, in *6th Topical Meeting on the Technology of Fusion Energy* (San Francisco, CA, USA, 1985) cONF-850310–92.

[77]S. Yang and Y. Gohar, in *6th Topical Meeting on the Technology of Fusion Energy* (San Francisco, CA,

2026-05-18 03:12 UTC - IP 216.73.216.54 - Columbus, OH, US (39.9587, -82.9987) - EVENT c944871f-2d71-4293-9e44-ed1da55d6279

USA, 1985) cONF-850310–64.

[78] J. W. Fletcher, E. E. Peterson, J. R. Trelewicz, and L. L. Snead, Fusion Science and Technology **82**, 901 (2025), also available via MIT DSpace: https://hdl.handle.net/1721.1/163370.

[79] K. P. Griffin, M. T. Walsh, S. Cohen, R. Feder, J. Klabacha, and Q. Zhou, "Effects of neutron radiation and shielding recommendations for the PFRC4," Tech. Rep. (Princeton Plasma Physics Laboratory, 2014) compiled with Attila transport modeling; PDF metadata creation date 2014-11-11.

[80] International Atomic Energy Agency, "Investigations of materials under high repetition and intense fusion pulses: Report of a coordinated research project 2011–2016," Tech. Rep. IAEA-TECDOC-1829 (International Atomic Energy Agency, Vienna, 2017).

[81] H. Momota, Y. Tomita, M. Ishikawa, and Y. Yasaka, Fusion Technology **35**, 60 (1999).

[82] H. Takeno, Y. Ikeda, T. Yamada, K. Noda, and Y. Yasaka, Japanese Journal of Applied Physics **39**, 5287 (2000).

[83] K. Sato and H. Katayama, Fusion Science and Technology **43**, 299 (2003).

[84] H. Takeno, K. Shibata, S. Nakamoto, T. Furukawa, and Y. Nakashima, Plasma and Fusion Research **18**, 2405053 (2023).

[85] U.S. Department of Energy, "U.S. fusion energy strategy," Tech. Rep. (U.S. Department of Energy, Washington, DC, 2024) federal strategy document; follows the March 2022 White House Bold Decadal Vision for Commercial Fusion Energy. Accessed 2026-05-14.

[86] The White House, Office of Science and Technology Policy, "Developing a bold decadal vision for commercial fusion energy," White House Summit on Commercial Fusion (2022), accessed 2026-05-14.

[87] U.S. Advanced Research Projects Agency–Energy, "Innovation network for fusion energy (INFUSE)," ARPA-E programme page (2024), accessed 2026-05-14.

[88] C. Miller, *Chip War: The Fight for the World's Most Critical Technology* (Scribner, New York, 2022).