



FLARE

A BOLD STRATEGY TO UNLOCK FUSION POWER

2025



CONTENTS



FOREWORD	4
EXECUTIVE SUMMARY	6
• FLF's Approach – Fusion via Low-Power Assembly and Rapid Excitation (FLARE)	7
• Next Steps	8
1. FUSION TARGET DESIGN FOR LOW-POWER DRIVERS	10
• Physical Principles for Minimising Power and Energy Requirements	11
• FLF's Approach	14
2. FUEL COMPRESSION	16
• Pulsed Power-driven Isentropic Fuel Assembly	17
• Stability and Manufacturing Advantages	19
• Research Pathway	20
3. HEATING AND IGNITION	22
• Pulsed Power-driven Self-ignition	23
• Short Pulse Laser-driven Ignition	32
• Research Pathway	34
4. PULSED POWER FOR INERTIAL FUSION ENERGY	36
• Challenges in High-Power Pulsed Power for Inertial Fusion Energy	37
• Low Power Systems and Potential Benefits	38
• Assessment of Recyclable Transmission Lines	40
5. CONSIDERATIONS FOR A VIABLE REACTOR DESIGN	44
• Compatibility with Existing Infrastructure and Technology	45
• High-level Requirements	45
• The Lithium Blanket	48
• Fusion-driven Shock Loading and Chamber Survivability	50
• Management of Radioactive By-products	51
6. A PATH TOWARDS ECONOMICALLY VIABLE IFE	52
• Development Pathway	53
• Power Plant Implications	54
• Investment Case	56
APPENDIX	58
GLOSSARY	72

FOREWORD



Fusion energy represents one of the most profound scientific and engineering opportunities of our time: a virtually limitless source of clean, safe, and reliable power. Using fuel of unmatched energy density, it offers the potential to transform global energy systems and sharply reduce dependence on carbon-intensive sources. Breakthroughs in Inertial Confinement Fusion (ICF), most notably the achievement of ignition at the National Ignition Facility in 2022, have accelerated momentum towards Inertial Fusion Energy (IFE), prompting countries such as the U.S. to expand national programmes and strengthen long-term support for commercialisation. For the UK, this is a pivotal moment to increase investment and ensure it keeps pace with international progress, positioning itself as a global leader in a technology that is no longer just a theoretical prospect but a credible path towards commercially viable abundant clean energy.

Over the past ten years, First Light Fusion (FLF) has established a strong foundation in IFE research and development. We have advanced novel target designs, developed specialised manufacturing methods, and built the modelling and experimental tools needed to test and refine our approach. This foundation, combined with the modularity and adaptability of our groundbreaking FLARE concept introduced in this paper, opens the door to a wider range of partnerships across industry and academia that will be essential to turning fusion into a practical reality.

The challenge that now lies before us is not to show that fusion can work once, but to demonstrate that it can be delivered consistently, affordably, and with the gain required for practical power generation. This requires well-understood barriers to be overcome: the efficiency and cost of drivers; the precision of target manufacturing; the resilience of reactor chambers, and the sustainable production of tritium. Addressing these issues demands innovation not only in physics and engineering, but in system design and economic thinking. FLF's approach addresses this challenge with a strategy grounded in pragmatism and scalability. By transferring performance requirements from complex drivers to innovative target design, deploying efficient pulsed power technology, and utilising a lithium pool reactor concept, the pathway described here seeks to reduce

complexity, enhance robustness, and deliver a commercially viable solution. Crucially, this integrated system is conceived with power plant economics at its core.

The implications are significant. A simplified, cost-effective route to inertial fusion energy would secure substantial advantages for the United Kingdom. Moreover, the collaborative development model outlined in this review aligns with national and international strategic priorities in energy security, clean energy, and defence.

This document sets out a clear and credible roadmap toward realising fusion as a practical energy source. It reflects both the scale of the challenge and the opportunity available to those prepared to innovate and to collaborate. I commend it to policymakers, researchers, and industry leaders as a serious contribution to one of the defining energy challenges of the 21st century.



Mark Thomas
Chief Executive Officer
First Light Fusion Limited

EXECUTIVE SUMMARY



Fusion power offers the possibility of an almost limitless source of clean energy, so concentrated that a single kilogram of fuel contains as much energy potential as 10 million kilograms of coal. This compelling vision of a world powered by fusion is encouraging governments and investors to pour billions into their favoured approaches, seeking to overcome the scientific, engineering, and economic barriers to a practical energy source.

Two main approaches dominate the field of fusion energy: Magnetic Confinement Fusion (MCF) and Inertial Confinement Fusion (ICF). In ICF, a tiny capsule of fuel is rapidly compressed and heated until fusion begins and becomes self-sustaining, a stage known as ignition. Once ignited, a burn wave is propagated into additional fuel, releasing vast amounts of energy.

The challenges of building a practical IFE power plant are well understood. Traditional IFE approaches rely on extremely powerful drivers and near-perfect precision, which drive costs into the billions and stretch development timelines far beyond practicality. Enormous amounts of energy are required to compress and ignite the target, yet typically less than one percent of that energy couples to the fuel. The design of the power plant must be able to withstand the rigours of repeated high-energy pulses at the cadence required for continuous electricity generation. The reliable breeding of tritium, an essential but scarce fusion fuel, must also be achieved at scale.

At the heart of the challenge lies gain, the extent to which a fusion reaction produces more energy than is delivered to the fuel. The current record gain, achieved in May 2025 at the U.S. Department of Energy's National Ignition Facility (NIF), is 4. According to FLF's economic modelling, a gain of at least 200 is needed for fusion to be commercially competitive, while a gain of 1,000 could deliver power at exceptionally low cost. The FLARE concept outlined in this paper is designed to achieve gains of up to 1,000.

Overcoming these hurdles will not be possible with incremental improvements alone. What is needed is a step change in approach: simplifying the driver, shifting more of the performance burden onto the target, and designing a reactor that is inherently robust, scalable, and affordable.

FLF'S APPROACH - FLARE *Fusion via Low-power Assembly and Rapid Excitation*

FLF proposes a new pathway - FLARE - focused on simplicity, efficiency, and real-world power plant economics. Our concept rests on three key pillars:

1. Innovative Target Design

Rather than relying on ultra-precise, high-power lasers, FLF uses cylindrical targets with a dense, opaque pusher to compress the fuel using modest input energy. Losses are reduced, confinement is improved, and ignition is triggered by an auxiliary source such as a short-pulse laser or pulsed power system. By decoupling compression and heating, this approach lowers the power required from the driver whilst enabling higher energy gain.

2. Pulsed Power-Driven Fuel Compression

Pulsed power offers a lower-cost, higher-efficiency alternative to lasers. Our low-voltage design eliminates much of the complexity that has historically limited pulsed power-driven IFE.

3. A Lithium Pool Reactor

The fusion reaction takes place inside a liquid lithium pool dynamically structured with inert gas. This design absorbs neutrons, breeds tritium, captures heat, and protects the reactor walls without the need for complex solid structures. It extends reactor lifetime, lowers costs, and positions the UK as a global supplier of tritium, a vital resource for both inertial and magnetic fusion.

Together, these innovations create an integrated system in which the driver, target, and reactor are mutually compatible, robust, and economically attractive: an end-to-end concept designed for deployment.

The Value of This Solution

- **Scientific Advantage:** Grounded in decades of research, but assembled into a unique and novel technology, this offers innovation but with low research risk.
- **Technical Advantage:** Reduces the power burden on the driver, simplifying engineering and improving robustness.
- **Economic Advantage:** Lower capital and operating costs provide a credible path to affordable electricity.

- **Strategic Advantage:** Positions the UK as a leader in IFE, with wider benefits for defence, aerospace, and advanced materials.
- **Partnership Advantage:** A deliberately collaborative development model that shares both risk and reward, and maximises benefits for the UK supply chain.

NEXT STEPS


FLF's pathway to realising this vision is pragmatic and partnership driven. We will:

- Continue validating each component of the scheme including target physics, pulsed power coupling and reactor survivability on existing national and international experimental facilities.
- Expand collaborations with key UK partners and international laboratories to accelerate progress.
- Mature our modelling, design, and prototype targets to build the scientific and engineering case for the first demonstration reactor.
- Leverage early commercial applications of our target and pulsed power expertise to generate revenue and strengthen supply chains.

In this document, we present a rethinking of IFE, offering a novel, simpler pathway, grounded in the rich scientific heritage of the field. This approach creates the foundation for a new ecosystem of partnerships, reduces development risk, and leverages existing technology and supply chains to accelerate deployment. By starting with an innovative target concept and working outwards, we show how IFE, from development to deployment, can become an achievable goal, making clean, abundant, affordable energy a reality.

By beginning with FLF's unique target technology, we can now propose a path to delivering our vision:
Powering a World Worth Inheriting.

1. FUSION TARGET DESIGN FOR LOW-POWER DRIVERS



The generation of energy from nuclear fusion using inertial confinement is a race in which the fuel must be ignited and burned before it explodes due to extreme internal pressure. The race is one of power balance; it can only be won if the energy from fusion exceeds all the losses from the system. Under typical Inertial Confinement Fusion (ICF) conditions, losses from the compressed fuel can reach petawatt levels (1 PW = 10^{15} watts), setting a fundamental requirement for the power delivered to the target.

Conventional ICF addresses this with extremely high-power drivers. However, these systems are costly, complex, and suffer from severe operational stresses. As a result, these driver architectures are poorly matched to the needs of a practical fusion power plant. FLF is pursuing a fundamentally different approach: reducing the power burden on the driver by asking more of the target.

A low-power pulsed power driver can be shown to meet the functional and economic requirements of a viable fusion reactor. Delivering the required energy to the target at lower power necessarily results in a longer pulse duration. Consequently, the target must be capable of efficiently absorbing and utilizing energy over an extended timescale to achieve the compression and heating required for the ignition and burn-up of the fuel.

FLF is developing a high-gain target concept that satisfies these constraints, based on well-established physical principles and amenable to validation on existing High-Energy-Density (HED) facilities.

PHYSICAL PRINCIPLES FOR MINIMISING POWER AND ENERGY REQUIREMENTS

1. Decoupling Compression and Heating to Ignition

A well-studied method to reduce driver energy requirements is to decouple the compression and heating stages. In conventional hotspot ignition, such as that demonstrated at the NIF (see the Breakout Box: NIF, hotspot, FI and volume ignition), fuel compression and heating occur simultaneously, a configuration that requires precise symmetry, pulse shaping, and timing. The simultaneous heating opposes compression, increasing driver energy and power specifications. Separating these two processes relaxes requirements, improves stability, and significantly reduces the energy required for ignition.

- Fast ignition (FI) is an alternative scheme for ICF in which the fuel is first compressed to high densities without forming a central hotspot. A small region is then rapidly heated using a short, intense pulse, typically delivered by a laser or charged particle beam.
- This approach eliminates the need for ultra-precise implosion symmetry (as no highly converged hotspot needs to be formed) and removes the energy budget required to form a hotspot in pressure equilibrium with the surrounding fuel.
- FI can, in principle, achieve significantly higher gains for the same driver energy or the same gain for lower energy. The approach offers a significant potential reduction in ignition threshold energy compared to hotspot [R.B. Stevens, UCRL-JC-135800].
- The separation of the heating and compression allows compressive work of the fuel to be maximised with (ideally) no increase in entropy. This latter point is crucial: Entropy generation during compression reduces the amount of useful work available for assembling the fuel. Irreversible processes such as shocks or turbulence increase entropy and make compression less efficient. By contrast, quasi-isentropic compression minimises entropy generation, allowing more of the input energy to go directly into achieving the high fuel densities required for ignition. The efficiency benefits of isentropic fuel compression have been recognised from the beginning of ICF research [J. Nuckolls, Nature (1972)].

2. Reduced Losses Through Radiation Recycling

A strategy to reduce the energy and power needed for ignition is to minimise energy losses from the system. This energy can escape via several channels, including thermal conduction, hydrodynamic expansion, and critically in the fusion conditions under consideration, radiative losses. As the fuel heats up it emits radiation which can escape and carry away energy, unless trapped. In principle, it is possible to accomplish this by achieving extremely high fuel densities, such that the fuel is opaque to its own thermal X-ray emission.

However, this requires an accordingly large compression provided by a faster implosion and greater driver power delivered to the target. Operating at lower power, and therefore lower assembled fuel densities, necessitates an alternative mechanism for reducing radiative losses.

This challenge can be addressed by encasing the fuel in a high-opacity (typically metallic) pusher that implodes and compresses the fuel.


- The radiation field from the hot fuel establishes thermal equilibrium with the compressed pusher, recycling a portion of the emitted radiation back into the fuel.
- Suppressing radiative losses in this way significantly lowers the ignition temperature from approximately 5-10 keV (1 keV = 11.6 million K) in traditional hotspot ICF to as low as 2.5 keV, (see Figures 1 and 2).
- This “*equilibrium ignition*” configuration is well established and recognised as part of a minimum energy route to fusion ignition [S. Colgate LA-UR-88-1268; S. Colgate LA-UR-92-2599; K.S. Lackner AIP conference proceedings (1994)].
- Another significant difference to the hotspot scheme arising from this change is that the fuel is typically heated volumetrically, i.e. uniformly throughout its full mass. Volumetric burn is typical of equilibrium ignition, but it fundamentally limits the achievable target gain, since the entire fuel

volume must be heated homogeneously. While values close to the minimum viable gain required for commercial IFE are theoretically possible, real-world effects will prevent realisation. Higher gain volumetric target designs are therefore desirable to ease the margin on other reactor components.

3. Increased Fuel Burn Fraction


Once ignition occurs, maximising fusion yield from the compressed fuel is critical. The total fusion yield is limited by expansion and cooling after ignition. Expansion occurs in response to the rapid increase in internal pressure after ignition and can quench the burn prematurely, that is, before a substantial amount of the fuel is burned. A high-density pusher can tamp this expansion, keeping the fuel compressed for longer, for an increased burn fraction. As a result, the required fuel mass is also reduced, which in turn reduces stress factors on the chamber and wider reactor infrastructure. The high pusher mass, typically more than an order of magnitude greater than the fuel mass, can restrict gain potential unless further accounted for in target design (see the Breakout Box: NIF, hotspot, FI and volume ignition).

Figure 1: Fusion power balance



Contours in the $(\rho R, T)$ plane showing the conditions under which net fusion self-heating balances all losses for deuterium-tritium fuel, in both spherical and cylindrical geometries. Areas inside (above and to the right) of the curves represent a positive power balance and net fusion heating. Solid lines represent spherical configurations, while dashed lines denote cylindrical cases with aspect ratio $a=h/2R=1.0$. Red and blue curves correspond to isobaric and isochoric configurations, respectively. Black curves represent volumetric self-heating conditions assuming suppression of radiation losses through high-albedo boundaries; the “ideal” curves include no radiative losses, while the realistic ones include 15% of optically thin bremsstrahlung losses. The cylindrical configuration exhibits higher losses due to an unfavorable surface-to-volume ratio but still achieves net self-heating under reduced radiation loss conditions. [E.S. Dodd, Phys. Plasmas (2020); O.S. Jones, Phys. Rev. Lett. (2004); S. Atzeni, The physics of Inertial Fusion].

Figure 2: Extension of plot from reference [S.E. Wurzel, Phys. Plasmas (2022)].




Contours of $Q = 1$ and $Q = \infty$ plotted against fuel temperature in keV. Red dashed line indicates a reduction of radiation loss equivalent to an albedo of 0.85. [E.S. Dodd, Phys. Plasmas (2020), O.S. Jones, Phys. Rev. Lett. (2004)]. Dashed and dotted lines show the effect of reducing or removing radiation loss, respectively. While this is an oversimplification, the extension of the ignition parameter space under the assumption of reduced radiation loss is evident.

Breakout Box: NIF, hotspot, FI and volume ignition

A reactor aims to efficiently burn as much fuel as possible in each pulse, producing the maximum amount of output energy for the minimum input energy. There are several established fuel configurations, each having their own merits.

Lawrence Livermore National Laboratory's National Ignition Facility (NIF) used the 'Central Hotspot' method, described below, to achieve scientific breakeven in the laboratory, for the first time in 2022. This result, achieved using 192 lasers, has prompted a large growth in public and private investment in inertial fusion energy, (see Figure 3).

Figure 3. Three different methods for igniting fusion fuel.



In the central hotspot approach, a fuel pellet is symmetrically compressed, typically in a spherical geometry, until a small central region reaches the high temperature and pressure required for fusion

ignition. Once the hotspot ignites, a burn wave propagates outward into the surrounding colder, denser fuel, sustaining the fusion reaction. This method has been successfully demonstrated at the NIF using laser-driven, spherically imploding targets.

Fast ignition decouples the compression and ignition phases to enhance overall energy efficiency. The fuel is first compressed to high density using a conventional driver, such as long-pulse lasers or pulsed power. Then, a separate high-intensity, short-duration driver, typically a petawatt-class laser, is used to rapidly heat a small region of the compressed fuel to ignition conditions. This localised hotspot initiates a burn that propagates through the fuel. Fast ignition is designed to reduce the total energy required for ignition compared to the central hotspot method, potentially improving energy gain.

Equilibrium, or volume ignition, involves heating the entire fuel volume nearly uniformly to conditions sufficient for fusion. Ignition occurs throughout the bulk of the fuel, and the burn is sustained across the whole volume simultaneously. Enclosing the fuel within a high-density shell can reflect thermal radiation, redistribute energy losses, and provide inertial confinement, helping to maintain uniform heating and reduce the external energy required for ignition.


FLF'S APPROACH

FLF's concept builds on these established insights and prior research to provide an economical and high-TRL approach to ignition, burn, and high-gain from a low power driver, (see Figure 4).

Fuel is assembled to high densities via a low power pulsed power-driven cylindrical implosion using

a high density, high opacity pusher. The pusher reduces radiation losses and increases inertial confinement and fuel burn up fraction. Ignition is triggered in a small fuel region, driving a burn wave into an adjacent larger assembled fuel mass to provide high gain. Figures 5-7 show previous concepts aligned with this approach. Multiple options on how to produce this ignition event are under investigation, with two being presented in Section 3.

Figure 4: Simplified conceptual images showing the stages of FLF's target process.

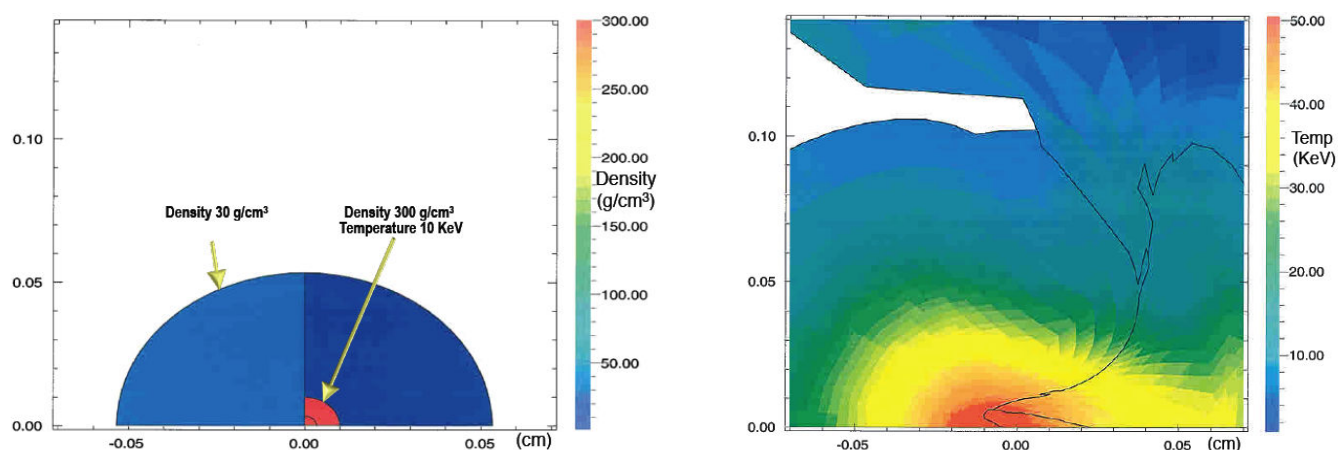


1. Isentropic compression: A low power, pulsed power driver is used to isentropically assemble fuel in a high-Z pusher. The target will use FLF's amplifier technology to create an isentropic loading profile to the fuel from a single impulsive pressure drive (amplifier structure not shown for diagrammatic simplicity).

2. Fast heating: Once compressed, a portion of the fuel is rapidly heated, in this case by a short pulse laser to the point of ignition. Many other options for heating the pre-compressed fuel are possible including target concepts based purely on a single pulsed power driver.


3. Ignition and Burn: Once ignition occurs a burn wave propagates into the rest of the fuel which is inertially confined by the high- Z pusher.

Figure 5: Modelling of a high-gain target concept.




Initial configuration for ignition across a 10 times density step and the resulting ignition profile. J H Nuckolls. Grand challenges of inertial fusion energy. 2010 J. Phys.: Conf. Ser. 244 012007. DOI 10.1088/1742-6596/244/1/012007. © IOP Publishing. Reproduced with permission. All rights reserved.

Figure 6: Conceptual design for a high gain target.



Reproduced from R.W. Paddock, et al., Potential High Gain Target Designs for IFE, IFE workshop (2022). A conceptual design for a high gain target based upon shock compression of a high mass hemisphere separated from a fast ignited wetted foam hemisphere by a high-density carbon glide plane. Electron driven fast ignition is used to ignite a small hemisphere, which in turn acts as an ion source to fast ignite a larger hemisphere. Considered to be an extreme form of cone-guided fast ignition where the cone has an angle of 180 degrees.

Figure 7: Ion driven systems Fast Ignition.



Ignition and burn of an ion beam driven cylindrical target significantly larger than those envisioned in this concept. Ion beam power 3.5 PW. Ramis R, Meyer-Ter-Vehn J. On thermonuclear burn propagation in a pre-compressed cylindrical DT target ignited by a heavy ion beam pulse. Laser and Particle Beams. 2014;32(1):41-47. doi:10.1017/S0263034613000839, reproduced with permission.

2. FUEL COMPRESSION



In fast ignition schemes the heating of the fuel is decoupled from the compression. Therefore, it is not necessary to shock-heat the fuel during the main fuel implosion. This opens the possibility for more gentle, quasi-isentropic compression to high density with low energy requirements.

- Isentropic assembly of the fuel occurs without generating significant entropy, that is, avoiding shocks, heating, and dissipative processes during compression. This leads to colder, denser fuel.
- While purely isentropic compression is not achieved in practice, quasi-isentropic compression is accessible with careful pulse shaping and/or structured target design. Some degree of shock heating may still occur due to imperfections or practical constraints. A small increase in the implosion adiabat can be beneficial in providing stabilisation.

PULSED POWER-DRIVEN ISENTROPIC FUEL ASSEMBLY

A low-power pulsed power driver will naturally lend itself to quasi-isentropic fuel assembly and align well with the operational requirements of IFE. Pulsed power offers a cost-effective, high efficiency means of compressing fuel.

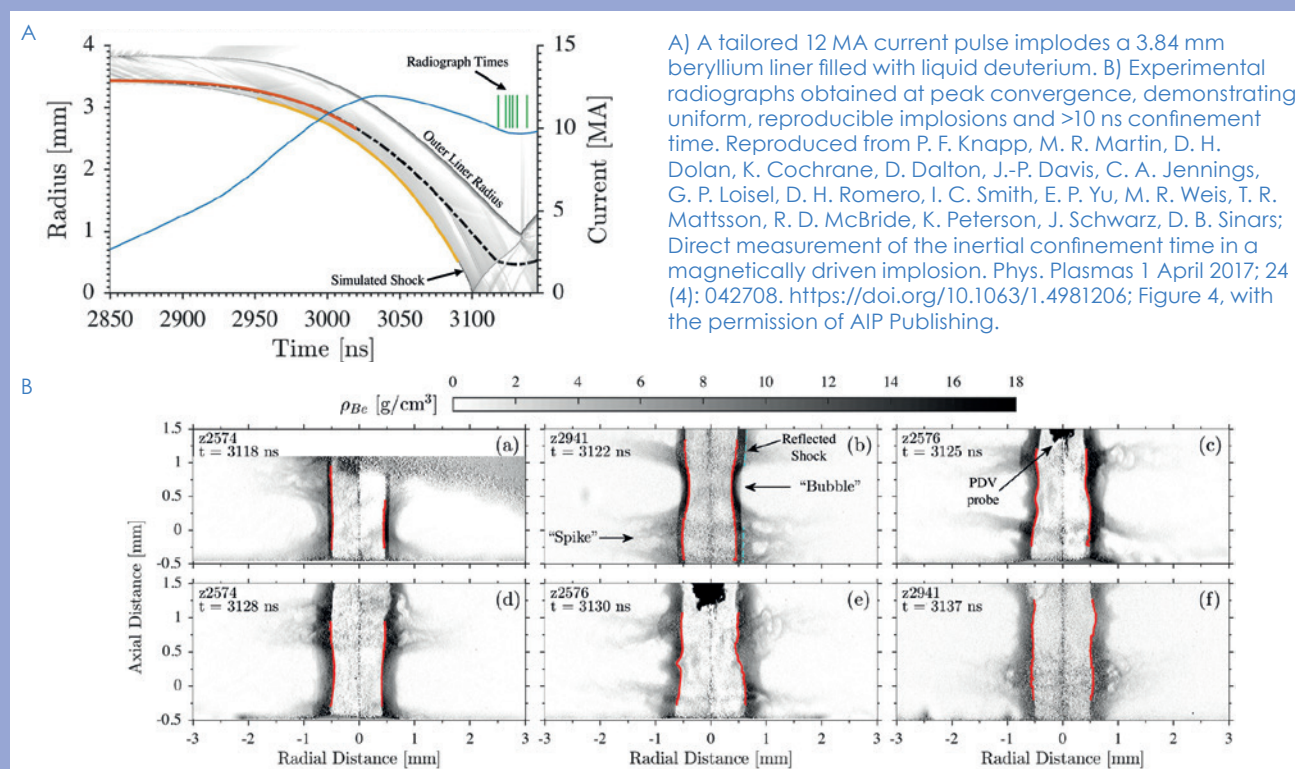
- Cylindrical liner-driven implosion of deuterium fuel has already been demonstrated at sub-ignition scales. This is enabled through controlled current waveforms and structured target configurations that suppress shock formation and reduce entropy [P. Knapp, Phys. Plasmas (2017); D. A. Yager-Elorriaga, Phys. Plasmas (2022); Lemke, J. Appl. Phys. 119, 015904 (2016); Rahman, Phys. Rev. Lett. 74, 714 (1995); Weinwurm, J. Phys.: Conf. Ser. 500, 082002 (2014)].

Breakout Box: Demonstration of pulsed power-driven isentropic compression of deuterium in cylindrical geometry: Sandia's 'Eddy' platform

Pulsed power-driven quasi-isentropic compression of liquid deuterium has been demonstrated on Sandia's Z machine. Liquid D₂ was gently squeezed to 10 g/cm³, 60 times its initial density, via the magnetically driven implosion of a

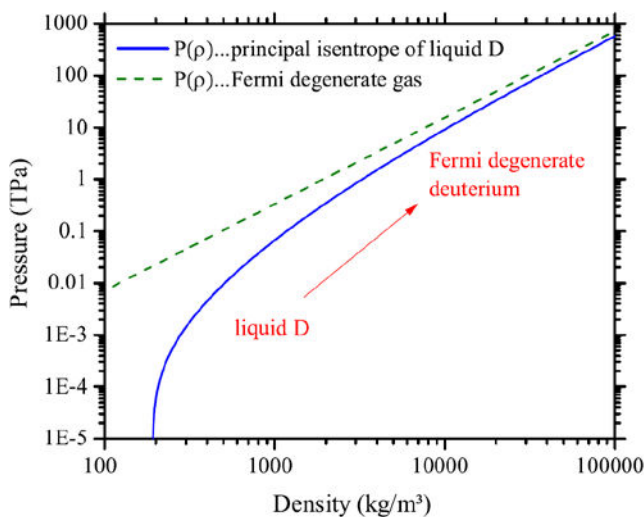
cylindrical metallic liner. The current pulse was ramped to a peak current of 12 MA, confining the fuel at a radius of ~500 μ m and a pressure of 100 Mbar, for 14 ns. X-ray radiographs obtained close to peak radial convergence demonstrate the stability and reproducibility of the system, evidenced by the uniformity and integrity of the liner inner surface, (see Figure 8).

Figure 8: Data from deuterium isentropic compression experiments on Sandia's Z machine.




- Pulsed power-driven liner implosions are more compatible with long-duration-low-power implosions than lasers, since the electrical discharges can be extended over long durations ranging from 50–1000 ns, with wall-plug efficiencies of > 20 %. In contrast, the longest pulsed laser systems require beam stacking (an inefficient usage of beam architecture with respect to the available space) to exceed 20 ns.
- The natural topology of pulsed power allows for straightforward coupling of magnetic pressure to cylindrical targets, as opposed to spherical implosions. This contrasts with laser systems, which require precise beam shaping and synchronisation of many individual beams to achieve drive symmetry, particularly in cylindrical or spherical geometries.
- The combination of decoupled heating and compression, with radiation recycling afforded by the liner, is expected to be compatible with less stringent final density requirements on the fuel. Consequently, only a modest degree of target convergence is required, even in cylindrical geometry. There is no requirement for the added complexity of spherical convergence that a laser driven system could otherwise bring, although a different set of instability mechanisms must be tackled.
- The use of a high density, low compressibility liner reduces the energy investment required to achieve pressure balance with the assembled fuel (see Figure 9).
- The areal density of the fuel refers to the product of the density and radius of a spherical target, appearing as a single term in simplified models of power balance between the fusion yield and aggregated losses. There exists a minimum areal density for ignition and burn, meaning that a lower density of fuel can, to some extent, be traded off against a larger radius/volume. In cylindrical systems this scaling does not hold precisely, although numerical studies show that the concept of areal density is still useful in predicting performance. In a cylindrical system the areal density can be tuned independently via the height of the liner, thereby partially decoupling it from the radial convergence.
- On Sandia National Laboratories (SNL) Z machine, a beryllium liner has been used to compress liquid deuterium to 10 g/cm³ and a pressure of 10 TPa, with a peak implosion velocity of ~ 25 km/s. In those experiments the liner was observed to have been compressed by a factor of eight, reaching a final density of ~15 g/cm³. (see the Breakout Box: Demonstration of pulsed power-driven isentropic compression

Figure 9: Principle isentrope of liquid deuterium in pressure-density space.



Reproduced from [M. Weinwurm, PhD Thesis. ICL (2014)]. Fermi-degenerate limit plotted as dashed green line. A fuel density of 100 g/cm³ implies a fuel pressure of 600 TPa. The high fuel pressures must be matched by the pusher. The high starting density of the proposed pusher reduces the energy investment required to achieve pressure balance with the assembled fuel.

Figure 10: Liner current required to balance isentropic fuel pressure as a function of stagnation radius.



The Fermi-degenerate deuterium gas approximation is used for $P(\rho)$. Target densities of 10 g/cm³, 30 g/cm³ and 100 g/cm³ are plotted. Corresponding areal density given on right hand axis.

of deuterium in cylindrical geometry: Sandia's 'Eddy' platform).

- For a gold liner, the required pressure could be met with three times less compression. The equivalent ram pressure to the beryllium liner would be achieved with an implosion velocity of just 14 km/s.
- Extending this simple scaling to the levels of fuel compression required for a reactor, and assuming pusher density tracks with final fuel pressure, one finds that implosion velocities of ~40 km/s are required to reach a fuel density of 30 g/cm³, and ~ 75 km/s for 100 g/cm³. Assuming an equal pusher mass for each case, a significant energy penalty is found for low density pusher materials scaling with the quadratic dependence on velocity.

Isentropic compression leads to favourable system-level driver specifications, that is, lower voltage and peak current requirements, helping to reduce both the complexity and cost of IFE power plants.

- The current waveform from a pulsed power driver is strongly coupled to the implosion dynamics of the cylindrical load. Lower implosion velocities lead to lower load voltages, which translate into reduced driver voltage requirements.
- At peak compression, the thermal pressure of the fuel balances the magnetic drive pressure from the liner current, (see Figure 10). For isentropic compression to densities above 1 g/cm³, the fuel becomes highly degenerate, and the pressure balance can be approximated as [M. Weinwurm PhD Thesis, ICL (2014)]

$$3.12 \times 10^6 \rho_D^{5/3} \approx \frac{\mu_0 I^2}{8\pi r^2}$$

- Crucially, sufficient areal density for ignition and burn when compared to [J.H. Nuckolls, J. Phys (2010)], 1.5 g/cm², can be achieved with peak currents of 30 - 40 MA. These relaxed electrical requirements enable the use of simpler, more modular pulsed power architectures that are inherently compatible with high-efficiency, high-repetition fusion power systems.

STABILITY AND MANUFACTURING ADVANTAGES

Instability control is a central challenge in IFE; slow, smooth compression improves robustness against Rayleigh–Taylor (RT) and Magneto–Rayleigh–Taylor (MRT) growth. Isentropic compression supports low-convergence, low-velocity implosions with thicker liners, key features for reliable fuel assembly and practical manufacturing.

- Early modelling by Zimmerman and Nuckolls demonstrated that with a suitable ignition source untamped fuel densities as low as 30 g/cm³ may be burned [J.H. Nuckolls, J. Phys. (2010)].
- The presence of a high-opacity, high-density pusher in FLF's proposed target concept lowers the density for ignition due to the additional inertial confinement time.
- Volume convergence is required only to raise fuel density, not temperature, reducing radial convergence ratios to 13–24 for assembled densities of 30–100 g/cm³. These are well below the requirements of hotspot ignition.
- Thicker, heavier liners with slow implosion velocities enable reduced in-flight aspect ratios and improve stability against MRT.
- Control of the acceleration and deceleration profiles of material interfaces over a given change in velocity can trade-off the instability growth rate (a function of the magnitude of the acceleration), against the time spent unstable, to improve uniformity at stagnation. In some cases, rapid deceleration is actually preferable for robustness, since stagnating before the fall line reaches the axis is correlated with improved stability [L. Welser-Sherrill, Phys. Plasmas (2008)].
- Lower pressure stagnation enhances inertial confinement and simplifies the timing of a separate ignition source.

Cylindrical geometry offers many target fabrication benefits. Unlike spherical targets, there is no need to incorporate symmetry-breaking features such as fill tubes, joint lines, or mounting structures. This simplification reduces tolerance demands and enhances scalability for mass production.

RESEARCH PATHWAY

FLF is investigating methods to achieve isentropic compression of cryogenic fusion fuel using a low-power pulsed power driver. FLF has been developing amplifiers capable of delivering ramped or isentropic loading profiles to targets using techniques such as reservoir release [J. Edwards, PRL, (2004)]. While initially explored as platforms for materials science in collaboration with external partners, this work is now being extended to FLF's large pulsed power facility, M3, as part of a broader programme. The aim is to convert a single impulse from a low-power-driven liner into effective isentropic compression of a cryogenic fuel volume. A key focus of this research is understanding the trade-off between introducing greater structural complexity in the target and the otherwise additional driver sophistication required to tailor the current pulse.



3. HEATING AND IGNITION



Once the fuel is assembled it must be heated to the ignition temperature required by the particular scheme. As shown in Figure 11, this may be conservatively estimated to be 2.5 keV (30,000,000 K) for the proposed concept. This heating must occur on a timescale shorter than the disassembly time in order to ensure enough of the fuel burns to enable energy gain.

At this temperature, the speed of sound (the characteristic speed for hydrodynamic action) in DT is around 500 km/s. For a compressed fuel region with a radius of ~50 μm , this implies a disassembly timescale of around 100 ps (0.0000000001 seconds) but this timescale will likely be increased by the added inertia provided by the high-density pusher. For the FI scheme to heat and burn substantial fuel this limit sets a useful upper bound on the rate at which energy must be delivered.

There are numerous options for igniting the pre-compressed fuel in the proposed configuration. FLF is initially exploring two candidate approaches:

- 1. Target-based power amplification and topological control**, which aims to achieve volumetric ignition via a structured pulsed power-driven target.
- 2. Short-pulse laser-driven ignition**, leveraging the high-power delivery of ultrafast lasers, similar to other fast ignition schemes. This option is aided by the enhanced confinement and reduced energy losses provided by the dense, high-opacity pusher.

These two approaches each offer distinct benefits and challenges. The laser-based pathway may provide the quickest and lowest-cost route to an integrated validation experiment but it requires line-of-sight access through the lithium blanket to the target. The pulsed power variant imposes greater demands on the driver compared to pure isentropic compression but preserves the option for nearly complete lithium blanket coverage.

PULSED POWER-DRIVEN SELF-IGNITION

Self-igniting targets demand high implosion velocities to heat the fuel and overcome energy losses during compression. Because pusher kinetic power scales as the cube of velocity ($\propto v^3$), this places increased demands on the driver.

Traditional laser-driven spherical ICF relies on high velocity and high convergence to heat the fuel to the ignition temperature. This is typically impractical for slower, pulsed power-driven cylindrical implosions with weaker volume convergence compared to spherical systems. IFE concepts such as Magnetised Liner Inertial Fusion (MagLIF) solve this by adding laser preheat and an external magnetic field, enabling the generation of fusion-relevant conditions at lower velocities and reasonable convergence. For more details, see the Breakout Box: Challenges for ignition with pulsed power-driven cylindrical implosions.

An alternative approach to relaxing ignition requirements is to employ a high-density, high-opacity pusher to increase inertial confinement and suppress radiative losses, (see the Breakout Box: Feasibility of cylindrical ignition with radiation recycling). However, this introduces a fundamental trade-off between achieving favourable ignition conditions and managing target complexity. Effectively coupling such a pusher to an inertial fusion driver motivates the development of multi-shell target designs.

Breakout Box: Challenges for ignition with pulsed power-driven cylindrical implosions

In a simplified model of traditional, laser-driven spherical ICF, an imploding pusher adiabatically compresses preheated fuel to the ignition temperature of 5-10 keV. The preheat is supplied by the strong shock that is driven into the fuel by the pusher.

Adiabatic compressional heating to a given final temperature (T_{ig}) is dependent on the radial convergence ratio, C , the geometric convergence factor ($g = 2$ for cylindrical, $g = 3$ for spherical) and the initial (preheat) fuel temperature (T_0):

$$T_{ig} = T_0 C^{\frac{2g}{3}} \quad (1)$$

It is unfeasible to reach the ignition temperature using pulsed-power driven cylindrical liner implosions for two reasons [S. Slutz, et al. Phys. Plasmas 2010 PoP]:

1. Adiabatic heating scales less favourably with radial fuel convergence for cylindrical versus spherical geometry. To reach the T_{ig} for a given T_0 would require unreasonably large radial convergence that would be highly susceptible to hydrodynamic instabilities.
2. Magnetically-driven liner implosions are slow compared to laser-driven implosions, limiting the available shock preheat. Within the constraints imposed by implosion stability considerations, cylindrical liner velocities are limited to around 100 km/s for a 100 ns drive (compared with laser-driven implosions at several hundred km/s

for a 10 ns drive). It is also challenging to reach the required radial fuel areal density for alpha stopping in cylindrical geometry. The fuel areal density scales as:

$$\rho r = \rho_0 R C^{(g-1)} \quad (2)$$

where R is the initial fuel radius. Hence, pure cylindrical target designs require higher radial convergence or much larger, higher yield targets.

The MagLIF approach

The MagLIF concept from Sandia National Laboratories in the U.S. introduces several kJ of auxiliary laser preheat prior to the magnetically driven cylindrical implosion of a liner by the 20 MA "Z" pulsed power machine. This preheat removes the requirement for a strong radial shock, permitting implosion velocities that are relatively slow (~100 km/s) compared to traditional ICF. To reduce thermal conduction losses from the preheated fuel into the liner wall during the relatively long (100 ns) implosion time, an external seed 10-30 tesla magnetic field is applied along the cylindrical axis. During the implosion, the preheated fuel is compressed quasi-adiabatically, significantly increasing its density while maintaining elevated temperatures, thereby reaching fusion-relevant conditions at a modest (~20) convergence ratio. The seed magnetic field is also compressed by the liner implosion, reaching several thousand tesla at stagnation. This can impede the radial transport of fusion alphas, thereby also easing requirements on the stagnated fuel areal density [Slutz, Phys. Plasmas (2010)].

Breakout Box: Feasibility of cylindrical ignition with radiation recycling

A simplified equilibrium ignition model for spherical geometry was developed within the multi-shell laser IFE literature. Here, we present the same model in cylindrical geometry, as a quick way to understand the cylindrical ignition landscape with substantially reduced radiation losses [D.S. Montgomery, Phys. Plasmas (2018); K. Molvig, Phys. Rev. Lett. (2016)].

A volume of DT is compressed and heated to ignition by a cylindrically convergent, high-opacity, high-density pusher. The fuel volume has

mass per unit length $M'_{DT} = \pi \rho R^2$, radial convergence, $C = R_0/R$, and density $\rho = C^2 \rho_0$. The mass required to achieve a pre-defined ignition radial areal density, $(\rho R)_{I_1}$, for a specified design convergence is:

$$M'_{DT} = \frac{\pi (\rho R)_{I_1}^2}{\rho_0 C^2} [\text{g/cm}] \quad (1)$$

The pusher heats the fuel via a strong shock followed by adiabatic compression. The post-shock temperature, T_s , is set by the implosion velocity, u_I . For an ideal DT plasma, $T_s = 4.3 \times 10^{-4} u_I^2$ keV, for u_I in cm/ μ s. This sets the initial temperature for the adiabatic compression phase, which begins at the time of shock collapse onto axis. This occurs with $C = C_s \approx 2.8$.

During adiabatic compression, the imploding pusher is decelerated by the back pressure of the fuel, and the pusher kinetic energy goes into the fuel internal energy. These considerations determine the implosion velocity required to reach the ignition temperature, T_I (keV), for a given convergence:

$$u_I = 48 \left(\frac{C_s}{C} \right)^{2/3} T_I^{1/2} [\text{cm}/\mu\text{s}] \quad (2)$$

By considering the pusher-fuel energy balance at the point of ignition, we can now find the required pusher mass for a given DT fuel mass. We retain the ignition margin parameter from the Revolver [K.Molvig, Phys. Rev. Lett. (2016)] literature, $\eta_I \sim 0.9$, which accounts for residual energy in the pusher at stagnation:

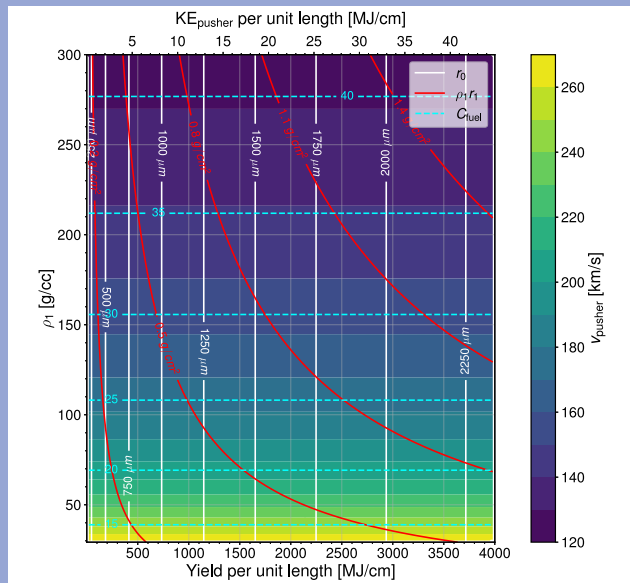
$$M_P' = \frac{M_{DT}}{[1-\eta_I^2]} \left(\frac{C}{C_s} \right)^{4/3} [\text{g}/\text{cm}] \quad (3)$$

The pusher kinetic energy requirement, $\frac{1}{2} M_P' u_I^2$, can now be expressed as a function of the fuel mass and is found to be independent of the convergence and implosion velocity:

$$E_P' = 114 \frac{M_{DT}}{1-\eta_I^2} T_I [\text{MJ}/\text{cm}] \quad (4)$$

The cylindrical ignition parameter space identified by this 0D model is plotted in Figure 11.

Figure 11: Cylindrical equilibrium ignition model results for $T_I = 2.5$ keV, $C_s = 2.8$, $\eta_I = 0.9$



Yield per unit length assumes a 40% burn fraction. Fuel radius plotted as white contours. Fuel convergence plotted as blue contours. ρ_I plotted as red contours.

The equilibrium ignition temperature makes a cylindrical system more viable by reducing the required radial convergence and implosion velocity. Sufficient fuel density can be achieved for $C \approx 25$. However, for a robust ρ_I (red contours), the tendency is for high yield targets, with direct implications for driver energy, power and chamber management.

Consider a $C = 25$, $\rho_I = 0.5$ design point. These constraints fully specify the initial fuel radius, the pusher outer radius and the pusher velocity, achieving 108 g/cm³ with a pusher velocity of 176 km/s and a yield of 130 MJ/mm for a 40% burn fraction. In this example, the pusher KE is 1.2 MJ/mm and the pusher power is $P = 0.5 \times \rho u_I^3 \times 2\pi R_0 = 430$ TW/mm. It is evident that high yield, volume burn cylindrical target specifications have significant implications for driver energy and power requirements. FLF are exploring efficient, target-based power amplification as a method to minimise these requirements.

Relaxation of (ρ_I) below 0.5 could be justified for a truly volumetric ignition that does not require propagating burn. Further, the aspect ratio of the assembled fuel profile could enable favourable alpha stopping along the z-axis. An auxiliary mechanism for increasing the fuel density prior to implosion could also be employed [S. Colgate, IEEE (2006), R. Menikoff, Phys.Fluids (1990)].

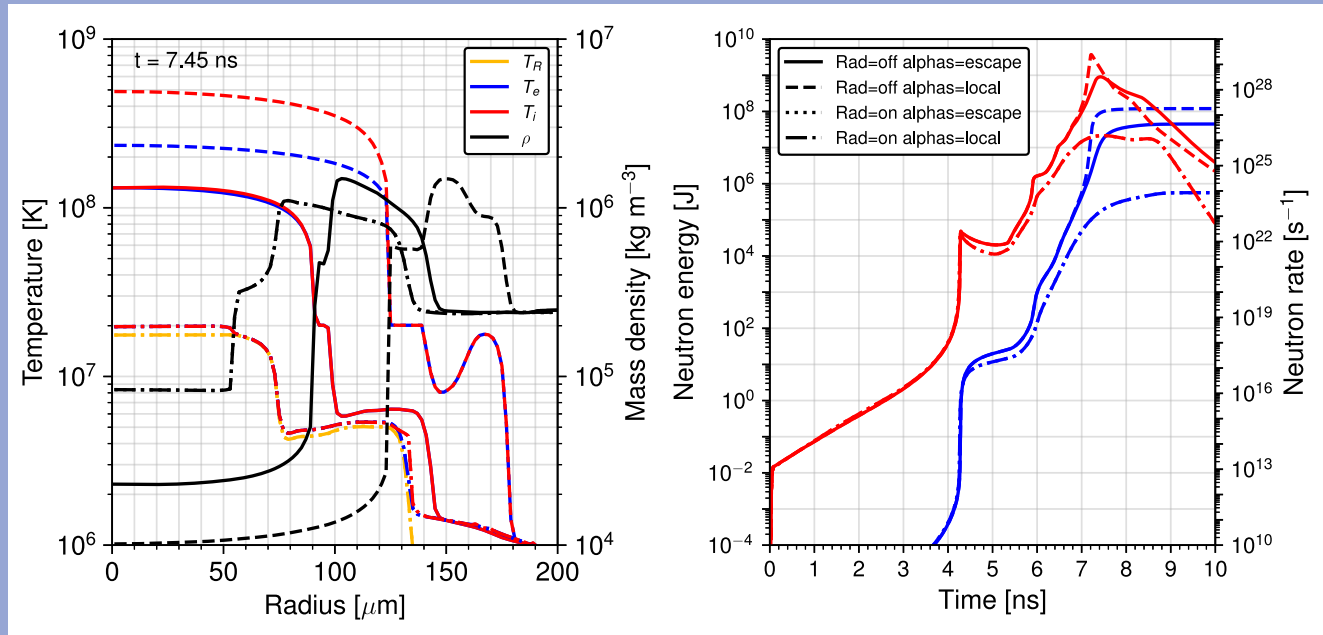
Idealised 1D hydrodynamics simulations were performed to validate the use of the model for exploring the parameter space of interest. The assumption of negligible radiative losses is inherent in the choice of ignition temperature. Therefore, it is appropriate to perform validation exercises of the above model with radiation losses switched off. We simulate an implosion with $R_0 = 0.12$ cm, a pusher velocity of 181 km/s, and pusher outer radius of 0.16 cm. The initial fuel density is 0.173 g/cm³ and the pusher is otherwise initialised at STP. We compare two alpha deposition models, such that alpha energy is either: a) deposited fully within the cell the alpha is born, or b) deposited fully or lost fully according to a local escape probability, which is based on the cell-local alpha mean free path. This provides an upper and lower bound on the amount of alpha heating, respectively yielding 120 MJ/mm and 49 MJ/mm. Both setups achieved $C \approx 16$, likely reduced from

the simplified model design convergence due to self-heating generating sufficient pressure to stagnate the pusher at larger radii.

Detailed accounting for residual radiative losses into the pusher wall should verify the efficacy of radiation recycling. Radiation-hydrodynamics simulations of the same design were performed with multi-group opacity tables covering the photon energy spectrum from $10^{-5} - 10^5$ eV. The pusher opacity is modelled as copper, scaled to account for the difference in number density for a given mass density. Figure 12 shows the radial profiles within 23 ps of bang time for all simulations. In the simulation used to identify the conditions for the spectrum, 11 group boundaries evenly distributed in log space covered the same spectral range. We define ignition to require total alpha energy deposition to exceed the initial pusher kinetic energy and bang time as the time of peak neutron production rate. The simulations performed with the optimised group structure achieve $C \approx 18$ and $C \approx 15$, and ignited with a bang-time $T_i^{bwa} = 9.70$ keV and $T_i^{bwa} = 49.7$ keV, for simulations with and without the alpha particle escape factor applied, respectively.

The performance is largely dominated by the alpha deposition model, with radiation cooling the bang-time T_i^{bwa} by <15% when the alpha escape factor is applied, and <5% when it is not. The additional convergence given by radiative cooling acts to counter the reduction in burn rate, to the extent that the overall neutron yield is increased by $\approx 3\%$ when all alphas deposit their energy in the cell they are born. When the escape factor is applied, the neutron yield is reduced by <12%. We caveat these results with the observation that from 10-50 keV, the proportion of the radiation energy density that lies at photon energies below 10^5 eV falls from 99% to 18%. Consequently, our tables will require larger photon energy ranges to accurately model very high temperatures. Performance may also improve further with more appropriate treatment of gold opacity, as radiative losses will be reduced. [K.Molvig, Phys. Rev. Lett. (2016); J. Scargle, Astrophys. J. (2013); A. Price-Whelan et al., Astrophys. J. (2022)].

Figure 12: 1D radiation-hydrodynamics simulation results for a cylindrical gold pusher configuration.



STP gold pusher velocity = 181 km/s. Liquid DT fuel, $R_0 = 0.12$ cm. (Left) Radial profiles of density (black), electron (blue), ion (red), and radiation (orange) temperatures at bang time for the cylindrical implosion described above. Simulations with radiation transport disabled with (solid) or without (dashed) an escape factor applied to reduce the alpha heating in cells where fusion is occurring are compared. Also shown are results for simulations with radiation transport modelled, with the same respective alpha deposition models given by the dotted and dot-dashed curves. (Right) Time histories of the cumulative neutron energy (blue) and the neutron production rate (red) for the same simulations.


Overview of Multi-Shell Long Acceptance Time Targets

The multi-shell class of inertial fusion target designs offers a practical framework for realising self-igniting triggers driven by long-pulse technologies. Velocity amplification is achieved via the sequential collision of shells of increasing density, enabling efficient energy transfer from the driver to the inner high-density pusher surrounding the fuel. This approach has been the focus of extensive research at major national laboratories over several decades.

FLF are continuing to develop this target technology, trading input energy for output power. Multi-shell target concepts could play a significant role in the pursuit of simple, affordable driver technology for IFE (see the Breakout Box: Velocity multiplication example in cylindrical geometry).

- Velocity multiplication via quasi-elastic impact of multiple shell collisions was first openly described in cylindrical geometry, as a method to drive a high-power, 200 km/s gold pusher using a low-power, 5 km/s aluminium liner (see Figure 13 A) [S. Colgate, LA-UR-06-3716].
- The broad class of multi-shell target concepts has evolved over several decades of research from US national laboratories, culminating in the spherical double shell designs currently being tested on the National Ignition Facility [D.S. Montgomery, Phys. Plasmas (2018)].
- From energy and momentum conversation, velocity multiplication can be accomplished with a drop in neighbouring shell mass. High-mass, low-velocity shells progress to lower-mass, higher-velocity shells.
- Cylindrical and spherical imploding systems allow for impedance matching and efficient energy transfer between different materials via shell convergence. This provides a means to grade the target from a high-density pusher to a low-density outer shell that can be efficiently coupled to the driver. This presents a lower-mass, more-efficient target approach to one based purely on mass ratios.
- The low-density outer shell can be accelerated to high velocity by the ablation pressure from a laser or the magnetic pressure from a pulsed-power machine (see Figure 13).
- Low-density buffer materials between the heavy shells act to mediate quasi-elastic behaviour by cushioning each impact, ramp-driving the subsequent shell and minimising wasteful shock heating. Modest convergence of individual shells also helps to minimise shell entropy and limit the growth of instabilities.

Figure 13: Some existing multi-shell target concepts



A) Multi-shell pulsed power-driven cylindrical target design, S. Colgate et al., *Equilibrium ignition by a cold adiabat: ICF via multi-collision velocity multiplication*, LA-UR-08-6915 (2008): Work performed by Los Alamos National Laboratory under the auspices of the U.S. Department of Energy. B) Revolver, a multi-shell, laser-driven spherical target design, C. Tyler. *Mission ignition 1663 LANL* (2019): © Copyright Triad National Security, LLC. All Rights Reserved. Unless otherwise indicated, this information has been authored by an employee or employees of the Triad National Security, LLC, operator of the Los Alamos National Laboratory with the U.S. Department of Energy. The U.S. Government has rights to use, reproduce, and distribute this information. The public may copy and use this information without charge, provided that this Notice and any statement of authorship are reproduced on all copies. Neither the Government nor Triad makes any warranty, express or implied, or assumes any liability or responsibility for the use of this information.

Breakout Box: Velocity multiplication example in cylindrical geometry

Mature design concepts presented in the multi-shell literature provide a baseline approach to inform the feasibility of achieving fusion relevant conditions with a low-power driver via target-based power amplification [D.S. Montgomery, Phys. Plasmas (2018); K. Molvig, Phys. Rev. Lett. (2016)].

From the high-density pusher geometry and velocity specifications, the geometry and velocity of additional shells with decreasing density and increasing radius can be fully specified by satisfying three conditions:

- Shell mass ratio
- Shell velocity ratio
- Kinetic pressure equilibrium ($P_{KE} = \rho u_I^2/2$) at shell impact

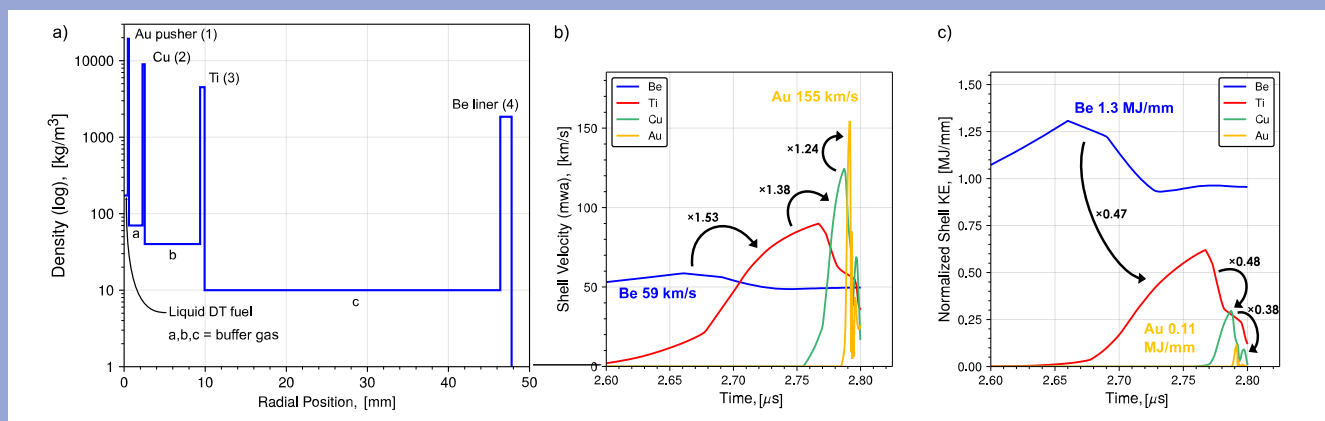
Through energy and momentum conservation, velocity multiple trades off against kinetic energy transfer efficiency. For modest shell convergence, and a carefully designed low density buffer layer, a mass ratio of 4:1 can deliver a velocity multiple of 1.4 with a transfer efficiency of ~50%.

For a thin-shell implosion in cylindrical geometry, shell density scales as $1/r$. The relative ambient shell density ratio, determined by material choice, sets the required shell convergence to satisfy (3) for a given velocity ratio.

Additional shells are specified until sufficiently low density is achieved for efficient coupling to a pulsed-power driver. The driver parameters can be matched to deliver the required outer shell velocity, given its specified geometry and impact radius. This approach allows a self-consistent, end-to-end assessment of driver scale requirements for volume ignition in pure cylindrical geometry, as a function of the fuel mass.

The principle of the process is illustrated with an example calculation and 1D MHD simulation in Figure 14. The design is specified by a very basic application of simple design rules above. No optimisation or advanced design options have been introduced.

Figure 14: Demonstration of basic multi-shell design principles for target-based power amplification with a 1D MHD simulation.



Outer liner accelerated by a coupled, 60 TW long-pulse driver. The outer liner aspect ratio and acceleration time can be reduced by increasing the driver power. The physical scale and energy requirements can be reduced by minimising the igniting fuel mass. Velocity multiplication via 3 collisions, from 59 km/s Be liner through to 155 km/s Au pusher. a) Initial geometry. b) Mass weighted average velocity of each shell as a function of time. c) Kinetic energy of each shell as a function of time.

Pulsed Power as an Attractive Driver for a Multi-Shell Target Implosion

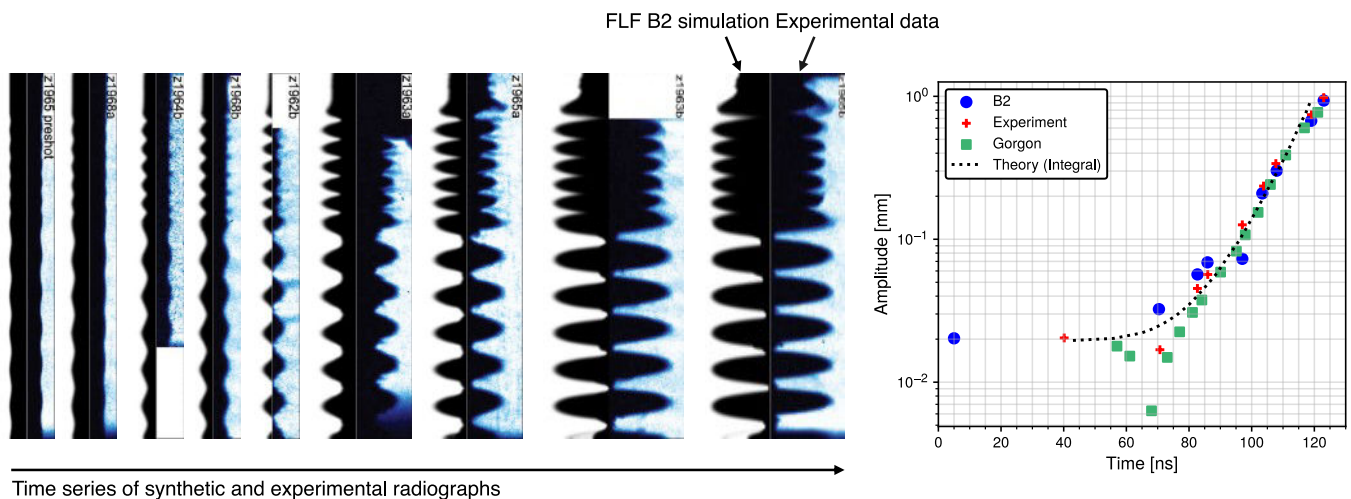
Pulsed power-driven cylindrical liner implosions offer a promising approach to multi-shell target design. Liner implosions convert electrical energy into kinetic energy as the liner converges towards axis. To enhance power amplification while maintaining implosion stability, target design must balance convergence and robustness to instabilities. This can be achieved by a graded density structure similar to that used in well studied spherical laser-driven multi-shell targets. Unlike spherical designs, cylindrical geometry reduces fabrication challenges related to symmetry-breaking features and enables better diagnostic access, supporting more robust design.

- A pulsed power-driven cylindrical liner implosion provides a convergent pressure drive for a subsequent system of nested shells.
- Liner implosions are a mechanism for pulse compression. Inherent amplification of the native driver power occurs because the current is coupled to the liner surface as it implodes. This is in contrast to laser-driven spherical implosions

that require complex dynamic zooming to provide the same function.

- Delivered power increases according to $I(t)^2 dL(t)/dt$. The load inductance, L , increases due to radial convergence of the current sheath during the implosion, as stored electrical energy is converted into the liner kinetic energy. Liner kinetic power is maximised when impact or stagnation occurs close to peak current.
- Lower-density materials can be imploded to higher velocities within the design constraints for implosion stability. This sets the requirement to decrease shell densities as we design radially out from the high-density pusher.
- A higher-convergence liner design will amplify power more effectively but also be more susceptible to disruption by the MRT instability.
- Diagnostic access in spherical geometry is limited. Surrogate laser-driven diagnostic experiments are performed in cylindrical geometry to enable additional measurements. An inherently cylindrical system offers better diagnostic opportunities [Sauppe, HEDP (2024)].

Figure 15: FLF simulations of perturbed liner implosion experiments on the Z machine.



Synthetic X-ray radiographs from FLF MHD simulations are compared to experimental data from fast, low-aspect ratio aluminium liner implosions on Z. Azimuthally correlated, low-amplitude sinusoidal perturbations of 200 μm and 400 μm wavelengths were machined into the outer surface to seed a single-mode MRT instability. B2 simulations were run in 2D with a fine 2 μm resolution and smoothing at the liner outer edge to capture the very low initial amplitudes without excessive grid imprint. The liners were driven by the load current from Sinars et al., with synthetic radiographs generated at the experimental diagnostic times; MRT amplitudes were extracted for direct comparison. The synthetic radiographs and MRT growth curves showed excellent agreement with experiment, validating the B2 model's treatment of magnetic boundary conditions, liner shock dynamics, and resistive MHD-driven instability growth, thereby increasing confidence in its ability to model single-mode MRT development in fast implosions. Experimental data from D. B. Sinars et al. *Measurements of Magneto-Rayleigh-Taylor Instability Growth during the Implosion of Initially Solid Al Tubes Driven by the 20-MA, 100-ns Z Facility*. Copyright © The Author(s), 2010. Published by Phys. Rev. Lett. 105, 185001 – Published 25 October, 2010 DOI: <https://doi.org/10.1103/PhysRevLett.105.185001> and licensed under a Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/3.0/>).

Instabilities and Mitigation Strategies for Magnetically Driven Liner Implosions

Magnetically driven liner implosions are susceptible to the MRT instability, which can disrupt symmetry and confinement. Multi-shell designs ease convergence requirements on the outer liner and improve robustness, though the primary instability sources remain under study.

- MRT often originates from the electrothermal instability (ETI), which seeds perturbations that grow rapidly after vapourisation through the electrochoric instability (ECI). In fast single-layer implosions, the combined ETI–ECI–MRT action limits viable aspect and convergence ratios to ~ 6 and ~ 20 , respectively, (see Figure 15) [K. Peterson, Phys. Plasmas 19,092701 (2012); J. Pecover, Phys. Plasmas (2015); S. Slutz, Phys. Plasmas (2010); S. Slutz, Phys. Plasmas (2016)].

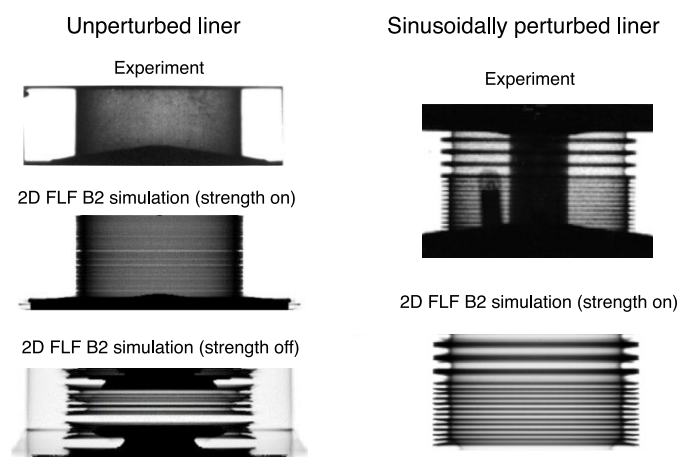
In ignition designs using target-based power amplification, MRT control is critical for the outer liner, which is expected to reach a convergence ratio of ~ 4 – 5 . Thicker liners resist MRT but reduce implosion velocity, requiring a balance between robustness and performance.

- Improved liner manufacturing uniformity can reduce seed perturbations. However, real-world

limits on surface and bulk defects restrict the potential of this approach to stabilisation. [Ryutov, Rev. Mod. Phys. (2000)].

- Dielectric coatings can tamp vapour expansion, reducing electrochoric growth and enabling higher aspect ratios without excessive MRT [Awe, Phys. Rev. Lett. (2016); Harvey Thompson, IEEE Trans. Plasma Sci. (2024)].
- Magnetic shear from a strong axial field can stabilise the $m=0$ mode, while axial velocity shear disrupts short wavelength growth [Awe, Phys. Rev. Lett. (2013), Awe, Phys. Plasmas (2014); Zhang, Phys. Plasmas (2006); Sinars, Phys. Plasmas (2011)].
- Deliberate long-wavelength seeding can suppress more damaging modes and allow targeted mitigation within inner components.
- Laser driven double-shells are vulnerable to low-mode three-dimensional asymmetries, whereas in pulsed power systems magnetic tension suppresses $r-\theta$ perturbations, leaving $r-z$ modes whose growth can be predicted and mitigated.
- In slower Z pinch systems, material strength can resist MRT in solid regions, (see Figure 16). Tailored current pulses could extend the period of solid-phase in the liner for increased stability [R.E. Reinovsky, Transactions on Plasma Science (2002)].

Figure 16: FLF simulations of liners driven by the Pegasus long-pulse machine, with and without material strength.



Synthetic X-ray radiographs from FLF's MHD simulations are compared to experimental data from slow aluminum liner implosions performed on Pegasus (7 MA, 7 μ s), diagnosed by radiography, with shots including both unperturbed liners and liners bearing azimuthally correlated, low-amplitude sinusoidal perturbations of varying wavelengths and amplitudes to seed multi-mode or single-mode MRT instabilities. Two series are shown: one unperturbed and one with sinusoidal perturbations chosen to show both linear and non-linear evolution of the MRT instability. B2 simulations reproduced these conditions in 2D, applying the appropriate perturbations at the outer surface along with volumetric random density and temperature variations to represent surface roughness, and driving the liners with the load current from Reinovsky et al.; synthetic radiographs were generated at the experimental times. The simulations showed good qualitative agreement with the radiographs and accurately captured liner trajectories, validating the B2 model's treatment of magnetic boundary conditions, material strength, and resistive MHD-driven instability growth. Inclusion of material strength was essential. Without it, both perturbed and unperturbed liners developed excessive MRT growth, demonstrating its critical role in maintaining stability in high-aspect ratio liners until stagnation. Experimental data reproduced from Reinovsky R.E., et al. *Instability Growth in Magnetically Imploded High-Conductivity Cylindrical Liners With Material Strength* published in IEEE Transactions on Plasma Science, Volume 30, Issue 5, 2002.

Instabilities and Mitigation Strategies for Inner Shells and the Pusher–Fuel Interface

Multi-shell targets are susceptible to Rayleigh Taylor (RT) and Richtmyer Meshkov (RM) instabilities due to multiple interfaces across different densities. Multiple design features and structures can be used to mitigate these effects.

- Outer shell surfaces become RT unstable during acceleration by the low-density buffer, while inner interfaces destabilise during deceleration.
- Target imperfections such as drive asymmetries, bulk features, and microscopic defects seed these instabilities.
- Instability growth can be limited by controlling initial perturbations, density contrasts, and the timing of acceleration phases. Symmetry can be reinstated at shell collisions via multiple shock reflections in buffer layers, and shaped shock sequences can control compressibility and in-flight aspect ratio to limit instability growth. Buffer layers reduce Atwood number mismatches and high mode feedthrough, while graded-density shells, low-Z cushions, and tailored shock interactions improve symmetry at slight 1D performance cost [Sauppe HEDP (2024); Schmitt, TF workshop (2019); LA-UR-19-23624 (2019); Roycroft, Phys. Plasmas (2022); Molvig, LA-UR-16-20248 (2016); Stark, Phys. Plasmas (2024); Mikaelian Phys. Fluids (2025); Haines Phys. Plasmas (2019); Milovich Phys. Plasmas (2004); Scheiner Phys. Plasmas (2019)].
- Deliberate mode seeding and shimming allow suppression or cancellation of dominant modes using controlled thickness variations, coatings, or patterned features [K. Peterson Phys. Plasmas (2018); Stark Phys. Plasmas (2021); Haines Phys. Plasmas (2019); Mikaelian Phys. Fluids (2025)].

At the pusher–fuel interface, deceleration phase RT instability occurs as the pusher stagnates against rising fuel pressure. Non-uniform compression can degrade confinement and mix high-Z material into the fuel, absorbing energy and potentially quenching ignition. Control of the implosion trajectory can be used to mitigate these effects.

- High-Z pusher designs often employ upstream ignition, achieving burn before significant deceleration, trading efficiency for robustness. Fall line optimised designs can retain up to 70% of 1D yield in 2D compared with ~2% for yield optimised designs. While these designs resist

shock driven kinetic mix, susceptibility to diffusive and turbulent mix remains a possible concern [K. Molvig, Phys. Rev. Lett. (2016); D.S. Montgomery, Phys. Plasmas (2018); Keenan Phys. Plasmas (2020)].

Pulsed Power-Driven Self-Igniting Target Design

For IFE, a central objective is minimising the energy invested in heating the ignition fuel mass. This is essential both to reduce driver cost and to enable higher system gain while keeping the total yield within containment limits.

FLF proposes a self-igniting packet of fuel enclosed within a high opacity, high density pusher, (see Figure 17). The application of a volumetric fuel configuration significantly changes the viability of ignition in cylindrical geometry.

Figure 17: Concept for a pulsed power-driven self-ignition component




Image showing a magnetically-driven multi-shell cylindrical target concept with a low-density liner and a high-density pusher. The high-density, high-opacity pusher helps to relax requirements for ignition.

- The reduced ignition temperature due to radiation recycling improves the viability of cylindrical target systems by reducing both the required radial convergence and the implosion velocity. For convergence ratios of approximately 25, sufficient fuel density can be achieved.
- To obtain robust areal density, there is a tendency toward high-yield target designs, which in turn places stringent demands on driver energy, power delivery, and chamber management.

- Further research at FLF is directed toward enhancing volume convergence, fuel density, and adiabatic heating for a given effective cylindrical pusher convergence by means of topological control. This method requires the introduction of additional mass to redirect flow, which incurs an energy penalty. However, the increased convergence and fuel compression achieved may outweigh this additional cost.
- In parallel, FLF is developing combined design strategies that incorporate multi-shell target architectures to stabilise and increase implosion velocity, while efficiently coupling the driver to a high-Z pusher to exploit its favorable properties. When integrated with topological control, these methods aim to minimise the ignition fuel mass and, consequently, reduce overall driver requirements.

In this approach, ignition is first achieved within the small, partitioned cavity. This seed energy is used to propagate a burn in the larger, isentropically compressed main fuel body. This scheme offers a potential pathway to achieving high gain while substantially reducing the driver energy required for direct ignition of the full fuel mass.

This multifaceted concept defines a high-dimensional design space. FLF intends to exploit its expertise in artificial intelligence, machine learning, and advanced optimisation techniques to navigate this space and identify configurations that maximise gain while minimising energy investment.

SHORT PULSE LASER-DRIVEN IGNITION

The prospect of achieving fusion via decoupled compression and heating has been a mainstream proposition in the IFE community for over 30 years. The advantages of such an approach are many and varied, particularly from the perspective of the assembly of the cold fuel.

Overview of Fast Ignition

In conventional FI, a fuel capsule is first compressed using a nanosecond-scale laser, reaching high densities ($\sim 300\text{--}1000\text{ g/cm}^3$) but not high temperatures. After peak compression, a second, ultra-intense short-pulse laser (typically picosecond, petawatt-class) is used to rapidly heat a small region of the compressed fuel to achieve ignition [Tabak, Phys. Plasmas (1994)].

- The short-pulse laser produces either relativistic electrons or high-energy ions, which propagate into the dense fuel.
- The character of the heating due to charged-particle beams depends strongly on the particles. High-energy ions moving through a plasma mostly interact with free electrons via small-angle Coulomb collisions and collective excitations (plasmon generation) until they are slowed to a speed at which their velocity is similar to the background ions, whereafter large-angle Coulomb scattering becomes more important. Their stopping power curve contains a pronounced maximum, which manifests in the energy deposition as a function of penetration depth as a characteristic feature known as the Bragg peak. Conversely, the stopping power curve of electrons is generally far smoother since, being far lighter, electrons are more susceptible to collisions involving significant energy transfer at all velocities. Consequently, electron energy deposition occurs at a fairly uniform rate along their trajectories.
- In the context of heating fusion fuel, light ions (such as protons and low-Z nuclei) are able to rapidly heat a small, localised volume to fusion conditions. On the other hand, electrons tend to scatter into a diverging cone-like beam, thereby making localised heating more difficult. If the electrons can be constrained to a tighter focus, e.g. within a powerful magnetic field, then the beam draws a substantial neutralising current from the background electrons, which is more collisional by virtue of their lower mean energy. The Ohmic heating associated with the return current can then also heat the background plasma to high temperatures over very short time scales.
- If local conditions meet ignition criteria, fusion reactions initiate, and alpha heating drives a propagating burn into the surrounding cold, dense fuel.
- The system has a high fuel density requirement to help alpha particles deposit their energy more locally, improving self-heating and reducing the amount of auxiliary energy needed from the short-pulse laser.

Challenges of Laser Fast Ignition

Electron and ion fast ignition both face significant technical challenges that have limited their practical viability. These challenges relate not

only to the feasibility of delivering sufficient energy but also to the difficulty of localising that energy deposition in a small enough region to trigger ignition. The concept also demands highly accurate timing between compression and ignition, further complicating implementation.

- The original FI scheme envisaged by Tabak et al. utilised a hole boring laser to open a channel to the dense fuel core using ponderomotive acceleration of the capsule corona. A short-pulse laser was then used to create a non-thermal ('fast') electron beam via a multitude of complex laser-plasma interaction mechanisms. Later variants used a cone to help guide the beam into the cold dense fuel [Tabak, Phys. Plasmas (1994); Kemp et al., Nucl. Fusion 54, 054002 (2014)].
- The electron beam is highly divergent, making focusing into a sufficiently small volume difficult. [Green et al., Phys. Rev. Lett. 100, 015003 (2008); Pérez et al., Phys. Plasmas 17, 113106 (2010); Rusby et al., J. Plasma Phys. 81, 475810505 (2015)].
- Substantial experimental effort into realising fast-electron-driven FI was made throughout the 1990's and early 2000's but highlighted major challenges that have yet to be overcome in an integrated experiment. [Kodama et al., Nature 412, 798 (2001); Fujioka et al., Science 318, 5858 (2007); Theobald et al., Phys. Plasmas 18, 056305 (2011)].
- Whilst the divergence problem can be mitigated using resistive collimator designs, issues of fabrication requirements, magnetic field growth rate, geometric compatibility with quasi-spherical fuel assembly, and maintaining the required structural properties under extreme and rapidly varying plasma conditions, have so far precluded realisation [Robinson and Sherlock, Phys. Plasmas 14, 083105 (2007); Kar et al., Phys. Rev. Lett 102, 055001 (2009)].
- Alternatives to resistive collimators based on compression of strong seed fields to counter divergence issues have also been explored, although presently limited to simulation-based studies [Gotchev et al., Phys. Rev. Lett. 103, 215004 (2009); Wang et al., Phys. Rev. Lett. 107, 035006 (2011)].
- Given the challenges outlined, a multitude of alternative mechanisms for heating the fuel to ignition temperature have been considered. Conversion of the short-pulse laser energy into fast ions, such as protons, whilst decreasing the overall coupling efficiency, have been proposed

[Roth et al., Phys. Rev. Lett. 86, 436 (2001); Kodama et al., Nature 418, 933 (2002)].

- The chief advantage of using ions is their far shorter mean stopping range and considerably less divergent beam character. Whilst experiments have been conducted to characterise the production and beam quality of fast ion sources, no fully integrated experiments have been completed.

Hybrid Fast Ignition

FLF proposes to ignite a small portion of the fuel enclosed in a high-opacity, high-density pusher using an auxiliary heating source, such as a short-pulse laser, (see Figure 18). The application of short-pulse heating to a high-Z cylindrical liner fuel configuration (referred to here as hybrid fast ignition, HFI) significantly shifts the cost-benefit landscape compared to traditional laser-based fast ignition approaches.

Figure 18: Concept short-pulse laser-driven ignition component.




Image showing the ignition of a pulsed power assembled fuel packet via hybrid fast ignition. Cylindrical geometry provides direct-line-of-sight access to the assembled fuel reducing penetration requirements. The high-density, high-opacity pusher helps localise the energy deposition and increase confinement time compared to untamped variants.

- The use of a high-opacity pusher reduces the ignition temperature, thereby lowering the total energy required from the short pulse.
- The inertia of the high-density pusher lowers the required fuel density compared to standard laser FI.

Breakout Box: Large laser facilities for fusion physics in the UK

Since the invention of the laser in the 1960's, the UK has been at the cutting edge with successful large, high-energy laser facilities. These have been funded primarily through defence and science councils. The two central hubs have been AWE Nuclear Security Technologies (AWE) and the Central Laser Facility (CLF, based at the Rutherford Appleton Laboratory).

AWE uses lasers as a means of validation for the Weapons Assurance Program. To enable this, they built the HELEN laser in 1979 and the ORION laser in 2012. The ORION laser was the first laser to have two, petawatt class short pulse beams and ten long pulse beams that could be used simultaneously to probe new high energy density physics. The facility cost £170m (see Table 1 in Appendix A) and is also used by academics.

Figure 19: A scientist in the ORION laser interaction vacuum chamber



The chamber is 4 m in diameter. All 12 lasers will be focused on to a millimeter scale target. This photo was duplicated from the AWE website (www.awe.co.uk), Photo UK Ministry of Defence © Crown Copyright.

The CLF has led the way in developing new laser technology to enable the study of new physics. They maintain several lasers and build lasers for other countries. Their highest energy laser, Vulcan, which was the first petawatt laser in the world is currently being upgraded, at a cost of £82m. Vulcan 20-20 will be the first twenty-petawatt laser in the world, keeping Britain at the forefront of laser innovation. These lasers are used by international academic and industry groups.

A CLF led European consortium was formed in 2005 with a proposal for a gain scale inertial fusion laser facility that would leapfrog the NIF facility in U.S. The proposal was to use 'fast ignition' to enable more efficient fusion with smaller, cheaper lasers. Ultimately, further validation was requested, and this facility did not get funded beyond the design phase. This European initiative is known as HiPER (High Power Laser Energy Research).

Figure 20: A concept design of the HiPER laser facility for fusion Gain research.




Image reproduced from Danson CN, White M, Barr JRM, et al. *A history of high-power laser research and development in the United Kingdom*. High Power Laser Science and Engineering. 2021;9:e18. doi:10.1017/hpl.2021.5. Copyright © The Author(s), 2021. Published by Cambridge University Press in association with Chinese Laser Press and licensed under a Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

- Several different coupling mechanisms between the short-pulse laser and the fuel can be explored, including electrons, ions, radiation, and shock waves.
- Because the energy is partially trapped by the surrounding structure, it may be possible to exploit multiple coupling mechanisms simultaneously. The added inertia of the high-density pusher could enable confinement of the fuel long enough for energy to equilibrate and ignite.
- Some energy will be lost to the high-density pusher as a diffusive radiation wave. Longer time scales will increase the amount of material heated, adding an energy penalty for coupling mechanisms taking extended periods of time.
- The pusher aids in localising energy deposition, while the cylindrical geometry enables direct axial access to the high-density fuel, an advantage over spherical laser-driven systems. This removes the requirement for the particles to penetrate deep into the dense fuel, lowering the required particle energy, which in turn lowers the intensity requirements at the laser focus.
- Electron and ion stopping ranges can potentially be tuned to match those of fusion products, offering design simplification and improved efficiency.
- The inclusion of short pulse laser heating will require a line-of-sight to the target which is a disadvantage in comparison to purely pulsed power-driven designs.

RESEARCH PATHWAY

FLF is exploring multiple pathways to ignite fuel pre-compressed by a low-power pulsed power driver.

One promising approach focuses on target-based power amplification and topological control to ignite a small packet of fuel. To assess its potential, FLF is investigating the physical processes required for a self-igniting target driven by a long-pulse, low-power driver. Current research covers volume ignition, topology control, power amplification, liner implosions, instability mitigation, and driver-target coupling. These efforts are supported by both flagship experimental facilities and advanced multi-physics simulations, with data science methods accelerating design iteration, (see Figure 21). The team is working to consolidate these findings into a self-consistent design point to demonstrate feasibility under realistic conditions.

In parallel, FLF is exploring a HFI concept tailored to pulsed-power IFE. By combining the advantages of fast ignition with radiation recycling, improved inertial confinement, and target structures optimised for manufacturability and laser access, HFI could enable ignition at lower laser energies and intensities, reducing overall power-plant cost and complexity. Research is underway through targeted studies of fuel coupling mechanisms and integrated ignition simulations, with near-term experiments being explored at existing short-pulse laser facilities (see the Breakout Box: Large Laser Facilities for Fusion Physics in the UK).

Both approaches currently being investigated present distinct trade-offs. HFI offers the fastest, most cost-effective path to an integrated demonstration system. However, it requires line-of-sight to the target through the lithium blanket. The pure pulsed power pathway demands greater driver performance than pure isentropic compression but retains the potential for nearly complete blanket coverage and favorable power plant integration.

Figure 21: M3 load hardware.




Image 3D render of the FLF multi-shell liner hardware currently being tested on the M3 pulsed power driver.

4. PULSED POWER FOR INERTIAL FUSION ENERGY



Pulsed power systems are often seen as a lower-cost pathway to IFE, both in terms of capital and operational expenditure. This is due to higher native efficiency than laser-based approaches. For instance, pulsed power schemes for MagLIF target designs approach 10% energy efficiency, in contrast to the roughly 0.7% coupling efficiency observed at NIF. This efficiency represents 14 times saving purely on stored energy required to drive the target.

Appropriately designing fusion targets so that they can combine with modest pulsed power systems is a practical way to make IFE achievable, negating the reliance on extremely high-power drivers. This approach avoids many of the engineering challenges associated with traditional high-power architectures and allows for simpler, more reliable systems with improved reactor compatibility. By extending pulse durations and reducing voltage requirements, it becomes possible to use established components and reactor compatible materials. This direction offers a clearer path toward fusion systems that are easier to build, operate, and maintain.

The use of pulsed power introduces the requirement for a direct electrical connection to the load that is partially destroyed and replaced after every shot. While this introduces additional mass and processes to the system, this connection provides many significant benefits including: no dynamic target positioning, a single chamber entry point, high lithium coverage, no chamber clearing required, and efficient energy coupling.

The typical requirements for the transmission line and power flow architecture of a high-power driver are not compatible with rapid, low-cost replacement. Transitioning to a low-power architecture greatly reduces tolerance requirements and complexity, offering a viable path to pulsed power-driven IFE.

CHALLENGES IN HIGH-POWER PULSED POWER FOR INERTIAL FUSION ENERGY

The technical demands of scaling high power systems to achieve ignition and high gain remain substantial. Traditional pulsed power IFE concepts rely on increasing the rate of current-rise far beyond the limits demonstrated in today's most advanced facilities. For example, the proposed Z-800 driver aimed to deliver over 60 MA in approximately 100 nanoseconds, requiring ~15 MV across a six-level

vacuum insulator stack [Stygar, Phys. Rev. ST Accel. Beams (2015)].

Achieving such rapid current rise requires increased voltage, but this scaling introduces diminishing returns. As the driver voltage increases, so too must system inductance. Higher inductance limits current rise and diverts more energy into magnetic fields rather than target compression, representing a fundamental inefficiency. Therefore, much of the engineering focus shifts to minimising inductance while delivering power reliably and reproducibly to the target [Spielman, IEEE ICOPS (2021)].

Several advanced technologies have been developed to address the challenges of high-power delivery:

- **Monolithic Radial Transmission Lines (MRTLs):**
Transporting power stored in distributed capacitor banks to a central focus is challenging. Radial water-filled disk transmission lines have the attractive property of allowing the voltage to increase through impedance transformation, thereby reducing the voltage required in the storage itself. The topological change in electrode geometry between the capacitor modules and the insulator stack (typically coaxial to radial) takes place at lower voltage, and further from the machine centre where space is less constrained [Stygar, Phys. Rev. ST Accel. Beams (2007)]. However:
 - » MRTLs introduce complexity and cost into building requirements. Penetrations through the high-voltage transmission lines for structural supports are prohibited meaning that a large unsupported roof is required over the centre of the machine.
 - » The associated need for heavy radiation shielding is therefore costly and complex to incorporate into the facility.
 - » While water in the radial transmission lines is an effective material for shielding neutron radiation, it does not shield the gamma radiation emanating from both the load and neutron-activated water. Additional shielding must therefore be placed at large radius, driving up cost.
- **Magnetically Insulated Transmission Lines (MITLs):**
MITLs are a well-established technology in existing facilities like the Z Machine, (see Figure 22) [Cuneo, IEEE TDEI vol 6 (1999)]. Four-level MITL systems have been a subject of study for decades, with stringent requirements identified

for the electrode conditioning and tolerances and empirically understood power losses in the double post-hole convolute that combines the four levels. Future high-power driver concepts typically propose using six-level MITLs to reduce the system inductance, where all levels are combined by a triple post-hole convolute near the load. These designs remain unproven at full scale, with remaining uncertainties about power delivery.


- **Modular driver architectures:** Higher voltages can be achieved through modular designs such as Induction Voltage Adders (IVA), Linear Transformer Drivers (LTD) and Impedance-matched Marx Generators (IMG). Voltage differences across single components are minimised but stack up across the machine to achieve a large output voltage. This simplifies the design in areas such as switching and component spacing. However, the increased component count introduces reliability challenges. For example, the Z300 concept requires 59,400 switches. Even assuming a modest pre-fire rate of 1 in 100,000, statistical analysis suggests that over 40% of shots would be compromised by at least one early switch trigger.

Taken together, these technologies represent significant engineering and design achievements, and may be acceptable for gain demonstration platforms. However, they remain poorly aligned with the long-term needs of a commercial fusion power plant. Reactor-scale systems require robust, maintainable driver architectures and central section power flow hardware compatible with being readily mass manufactured and recycled.

LOW POWER SYSTEMS AND POTENTIAL BENEFITS

Delivering energy to the fusion target over longer timescales offers a practical and cost-effective alternative to high-power pulsed power architectures. Extended pulse lengths allow the use of higher-inductance and lower-voltage components. Accordingly, this enables the use of established, high-TRL systems such as conventional Marx generators, and more robust switch technologies operating with relaxed timing requirements. A key design threshold is reached when single-layer power feeds become viable, eliminating the need for complex stacked MITLs and post-hole convolutes. This simplification reduces manufacturing complexity and also

Figure 22: Z vacuum insulator stack, MITL and Post-hole convolute.




Schematic of the ZR 4-layer MITL and double post-hole convolute system. These technologies represent significant engineering and design achievements. While suitable for gain demonstration platforms they are not compatible with the long-term needs of a viable power plant. Images from W. A. Stygar et al. *55-TW magnetically insulated transmission-line system: Design, simulations, and performance*. Copyright © The Author(s), 2009. Published by Phys. Rev. ST Accel. Beams 12, 120401 – Published 7 December, 2009, DOI: <https://doi.org/10.1103/PhysRevSTAB.12.120401> and licensed under a Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/3.0/>).

improves reactor integration by lowering the profile of the insulator stack and transmission lines. The result is improved shielding of driver systems that is better aligned with the constraints of a commercial fusion power plant. Combining modern technology with historic low-power architectures opens a new avenue to pulsed power energy delivery compatible with reactor requirements, (see Figure 23).

- **Established high-TRL technologies:** All major components operate within the performance envelope of existing technologies. Demonstration systems already exist at the tens-of-terawatts scale, requiring no new pulsed power innovations (see the Breakout Box: UK Pulsed Power Expertise).
- **No need for MRTs:** Some machine configurations can instead use vertical transmission lines, enabling structural supports to be integrated through to the roof. This change allows for localised radiation shielding, reducing overall shielding mass and associated costs.
- **Simpler switch operation:** The design operates within safety margins of proven switch technologies, avoiding the need for innovative or untested designs.
- **Alternative to SF₆:** Operating at lower voltages and relaxed timing tolerances allows for the use of alternative gas dielectrics, reducing reliance on SF₆; a greenhouse gas 22,000 times more potent than CO₂.
- **Simplified triggering systems:** The combination of lower voltage and slower rise times allows switches to be triggered using fast high-voltage electrical pulses, eliminating the need for laser-triggering systems which are expensive and complex.
- **RTL feasibility and repeatability:** If the target can be designed to tolerate a longer input pulse, then higher system inductance and lower voltage become acceptable. This enables transmission line geometries such as single-layer RTLs, that are simpler, cheaper, and more compatible with reliable, rep-rated operation in a reactor setting (see the Breakout Box: Self-consistent transmission line design for liner implosions).

Figure 23: Cross-section of the Atlas pulsed power driver.



Simplified power flow architecture offers direct chamber access for shielding and also meets infrastructure requirements for a reactor. Harold A. Davis et al., The Atlas High-Energy Density Physics Project. Jpn. J. Appl. Phys. 40 930 (2001) DOI 10.1143/JJAP.40.930. Copyright (2001) The Japan Society of Applied Physics. Reproduced with permission. All rights reserved.

Breakout Box: UK pulsed power expertise

The UK has long been a global leader in pulsed power research, with a legacy rooted in innovation and high-impact experimentation. Central to this legacy is the AWE where pioneer J. C. Martin, widely considered the father of Pulsed Power, performed foundational work establishing many of the theoretical and practical frameworks for modern pulsed power systems. The UK's continued excellence in the field is exemplified by cutting-edge facilities and collaborations that drive progress in IFE, X-ray generation, and electromagnetic launch systems. This enduring leadership reflects the UK's commitment to scientific advancement and technical ingenuity.

Figure 24: Marx bank from the MAGPIE generator at Imperial College London.



One of the 2.4 megavolt Marx capacitor banks at the Mega Ampere Generator for Plasma Implosion Experiments (MAGPIE), Imperial College London.

ASSESSMENT OF RECYCLABLE TRANSMISSION LINES

The ZIFE program at Sandia National Laboratories evaluated the feasibility of RTLs as a key enabling technology for pulsed power-driven IFE reactors [C. Olsen, SNL Report (2006)]. RTLs act as sacrificial components that deliver multi-megampere currents from a reusable pulsed power driver to a target, while shielding upstream components from fusion debris. Studies confirmed that RTLs can be fabricated from inexpensive, low-mass materials, integrated with the target for simplified alignment, and replaced rapidly to support repetition rates of ~ 0.1 Hz. Chamber clearing and RTL geometry optimisation were shown to be compatible with debris and thermal management, establishing RTLs as a promising interface between pulsed power delivery and reactor-scale IFE operation. [ZIFE Phase I Report; SNL internal studies; Slutz et al.].

The use of RTLs in fusion power plants introduces additional demands on material handling, particularly concerning recycling and radiation protection. While all fusion concepts must manage activation products, RTL-based architectures increase the volume of material processed per shot, affecting both capital cost and the levelised cost of electricity (LCOE).

- A key mitigation strategy is to design RTLs from materials inherently compatible with the reactor environment. Where pulsed power constraints allow, using lithium significantly reduces radiological and waste-handling complexity. Lithium components can be recycled directly into the surrounding lithium pool, eliminating the need for complex separation, transport, or off-site disposal, thereby reducing operational and regulatory burdens.
- Manufacturing methods such as deep drawing, shear spinning, and flow forming are compatible with many candidate RTL materials. These techniques enable rapid, repeatable production of axisymmetric RTL components and can be adapted to include shielding features or non-line-of-sight geometries.

Transitioning from ZIFE's baseline to longer-pulse systems requires reassessment of parasitic losses such as gap opening and magnetic diffusion, both of which can increase effective system inductance and reduce energy delivered to the target. Initial feasibility assessments of single-layer RTLs for reactor-compatible power feeds indicate that low-mass RTL designs remain compatible with pulsed power and fusion environment requirements.

Assessments of a 100 MA current pulse are presented as a worst-case scenario. The system shows promise for a largely lithium-based RTL,

Breakout Box: Self-consistent transmission line design for liner implosions

The behaviour and requirements of the target underpin the design of the final Transmission Line (TL) and the rest of the pulsed power system. The desired liner implosion trajectory determines the current profile required at the load. The total current carried by the TL is constant along its length at any moment in time, when current losses are neglected. This is because the round-trip transit time for the electromagnetic wave is typically much shorter than the implosion time (even for fast ~ 100 ns implosions).

The voltage at the load end of the TL depends on the load geometry and implosion trajectory, increasing with the implosion velocity and liner height, and inversely with the liner radius. The TL voltage at the permanent machine interface is larger than the load voltage while the current is increasing, by an amount proportional to the effective TL inductance and the rate of current rise. The effective inductance depends on the actual spatial distribution of current during the pulse, which is governed not only by the profile of the gap between conductors in the TL (the geometric inductance) but also by magnetic diffusion and electron flow in the gap.

The profile of the gap between conductors needs to be carefully designed to maximise power delivery to the load. The reduction of inductance with a decreasing gap must be balanced against electron losses, ion losses, and anode plasma production that increase with a decreasing gap.

Self-consistent assessment of this trade-off is made using established vacuum power flow theory benchmarked against existing pulsed power machines, as well as detailed Particle-In-Cell (PIC) simulations of the TL. Gap closure by expanding electrode plasmas may be the determining factor on the minimum gap rather than vacuum power flow considerations, especially for long pulse durations. An iterative design process considering all the relevant physics is required to obtain a self-consistent TL design for a particular target [Ottinger, Phys. Plasmas (2006); Vandevender, Phys. Rev. ST Accel. Beams (2015)].

Fast liner implosions imply high voltages at the load and large rates of current rise, which increase the voltage across the entire TL for the duration of the experiment. Insulator surfaces at the vacuum boundary are limited to electric fields of ~ 150 kV/cm to avoid surface flashover, so interface height increases with voltage. High-power drivers must use multiple TL layers to share the current and reduce the effective inductance [Stygar, Phys. Rev. ST Accel. Beams (2015)].

Long-acceptance time targets demand a lower rate of current rise, opening the possibility of a higher inductance, single-layer TL. Preliminary modelling suggests that a single layer design can deliver load currents of ~ 100 MA in under $2\ \mu\text{s}$ (peak power < 100 TW at the vacuum interface), with a ratio of magnetic field to electric field (a figure of merit for magnetic insulation) higher than that in the Z Machine MITLs during routine operation.

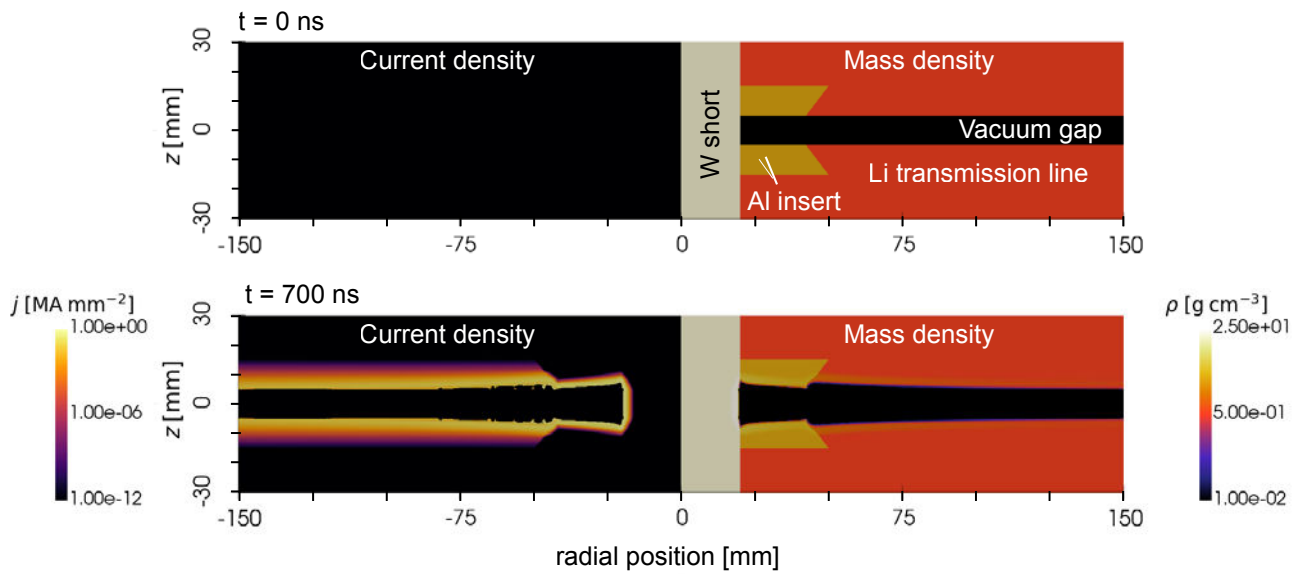
with higher-density materials likely required in critical regions such as target glide planes, where lithium's mechanical and electrical properties may be insufficient, (see Figure 25).

- Lithium's low density ($534\ \text{kg/m}^3$) compared to conventional MITL materials (e.g., aluminum at $2700\ \text{kg/m}^3$ or steel at $7800\ \text{kg/m}^3$) reduces inertial tamping, making it more susceptible to gap opening under transient magnetic pressure. Minimising gap opening during current rise is essential to avoid loss of power coupling to the load.
- MHD simulations with indicative ~ 100 MA, $\sim 1\ \mu\text{s}$ current pulses suggest lithium electrode motion has a negligible contribution to the dynamic

inductance outside of a 70 mm radius. The pulse-averaged electrode surface velocity at this radius is ~ 1 km/s. However, this analysis suggests that higher-density materials may be required at smaller radii, to be determined once detailed target geometry is available for further simulations.

- The motion of the conductor surface affects the vacuum impedance of the transmission line. However, inductance depends on the actual current distribution, which deviates from a simple surface current due to effects such as magnetic diffusion, plasma bridging, and electron flow in the vacuum gap. More detailed modelling is required to accurately quantify these contributions.

Figure 25: Recyclable Transmission Light MHD simulations.



B2 simulation examining the properties of a lithium RTL with a tungsten short circuit subjected to 94 MA in 700 ns. A static inductance of 4.03 nH is increased by 1.4 nH due to the gap opening and current diffusion. Electrode motion has a negligible contribution to the dynamic inductance outside of a 70 mm radius.

- Lithium's relatively high resistivity (90 nΩ·m at STP vs. 24 nΩ·m for aluminum) leads to enhanced magnetic diffusion. For a 1 μs rise time, the magnetic diffusion length in lithium is ~280 μm (assuming constant resistivity). [Note: nonlinear effects are expected within a certain radius.]
- Structural finite element analysis of concept RTL designs show lithium to be self-supporting over relevant dimensional scales of several meters, easing mechanical and manufacturing concerns.

Additional modelling was conducted to evaluate the direct effects of fusion radiation on RTLs made from various materials and alloys. While detailed results are beyond the scope of this document, initial neutronics simulations using OPENMC and MCNP show that a fully aluminum RTL immersed in lithium and subjected to a 5 GJ fusion yield would fall below low-level waste (LLW) criteria within three months. In contrast, steel-based RTLs would require significantly longer decay periods. Neutron-induced heating was sufficient to melt aluminum out to a ~50 cm radius, while photon heating was found to be confined to the immediate target vicinity.

5. CONSIDERATIONS FOR A VIABLE REACTOR DESIGN



COMPATIBILITY WITH EXISTING INFRASTRUCTURE AND TECHNOLOGY

The first generation of fusion reactors are likely to resemble existing fission reactors, primarily because this approach offers the fastest and highest-TRL path to deployment. The aim is for the "fusion island" (where heat is generated) to serve, as far as possible, as a replacement for the traditional "nuclear island" (including the reactor core).

- If the temperature, pressure, and neutron flux within the fusion island fall within the operational envelope of fission systems, then existing steam generators, Rankine cycles, and turbine technologies can be adopted.
- Where this compatibility is achieved, established nuclear-grade materials, such as currently approved steels, can be used, facilitating regulatory approval. This is a key advantage over approaches that require the development and qualification of novel materials.
- Compatibility with fission technology reduces investment risk: large portions of the plant can rely on standardised components, lowering overall development cost and technical uncertainty.
- In the UK, supply chains for Pressurised Water Reactors (PWRs) are well-developed and are expected to remain active and mature further over the next 20–30 years, as the current fleet of planned reactors is constructed.
- A concept reactor configuration outlined in this Section is shown in Figure 26.

HIGH-LEVEL REQUIREMENTS

The current generation of fission reactors under construction in the UK, at Hinkley Point C and Sizewell C, are PWRs based on the European Pressurised Reactor (EPR). This design is therefore used as the reference point for comparison.

The key technical requirements for fission reactors along with fusion and ICF specific requirements, such as tritium breeding and shock mitigation, can be used to assess the feasibility of IFE schemes.

The technical requirements for fission systems broadly define the characteristics of an ideal reactor concept. These are generally met by modern fission systems, though safety and radioactive waste management remain areas of

concern. These latter points are key motivations for pursuing fusion as an alternative energy source along with wider safety considerations, fuel efficiency (high energy release in the D-T reaction), and sustainability.

Fission Reactor Requirements

- **High safety standards:** Fission safety challenges vary according to reactor type but include the high radiotoxicity of spent fuel, the potential for thermal runaway, high operating pressures, and the requirement to retain fission products in the event of failure or an accident.
- **Minimise radioactive waste:** Over successive reactor generations, waste volumes have been reduced through the implementation of Best Available Techniques (BAT). However, the production of long-lived, high-activity actinides remains a central challenge.
- **Compact neutron source:** A localised neutron source facilitates shielding and radiation management. Shielding at source is a tenet of good radiation safety practice and generally simpler for smaller sources.
- **High primary energy efficiency:** Efficient capture of energy, particularly from neutrons, is critical to overall system performance.
- **Thermal compatibility with power conversion systems:** Compatibility with current standard electricity generation infrastructure suggests selection of an appropriate coolant temperature (e.g. ~600 K for PWRs).
- **Safe operating pressures:** Pressure must be controlled to assure retention of radionuclides in the event of an accident or failure (e.g. ~150 bar for PWRs).
- **Long operational lifetime:** Minimising Levelised Cost of Electricity (LCOE), particularly through enhanced reactor availability and lower operating costs demands long reactor life (e.g. ~60 years for current PWR designs).
- **Minimise maintenance requirements:** Maintenance burden must be kept low to maintain availability, reliability, and economic performance.
- **Reliable reactor control:** Control systems that are conceptually simple and operate reliably when needed, are essential for operational uptime and safety.

- **Cost considerations:** Minimising capital and operational expenditure and hence LCOE, remains a priority for all forms of nuclear power.

Fusion-Specific Requirements

As well as the general criteria above, fusion systems face additional technical requirements:

- **High tritium breeding ratio:** Self-sufficient tritium breeding is necessary to sustain deuterium–tritium fuel cycles.
- **Low-toxicity, low-activation breeder blanket:** Where possible, the blanket structure materials are to be selected to minimise activation and the production of toxic byproducts.
- **Minimise blanket cost:** The need for costly, problematic lithium enrichment or additional consumable (routine replacement) wall structures should be avoided.

ICF-Specific Requirements

As with Magnetic Confinement Fusion, there are further unique challenges specific to Inertial Confinement Fusion:

- **Shock mitigation:** Designs must address the impact of shock waves on the first wall, chamber and additional surrounding systems.
- **Debris management:** Repetitive fusion events generate significant target debris, necessitating effective protection and removal strategies.
- **Recyclable Transmission Line (RTL) exchange:** For systems involving a direct electrical connection, RTL replacement or refurbishment must be feasible at repetition rates of order 0.1 Hz.

THE LITHIUM BLANKET

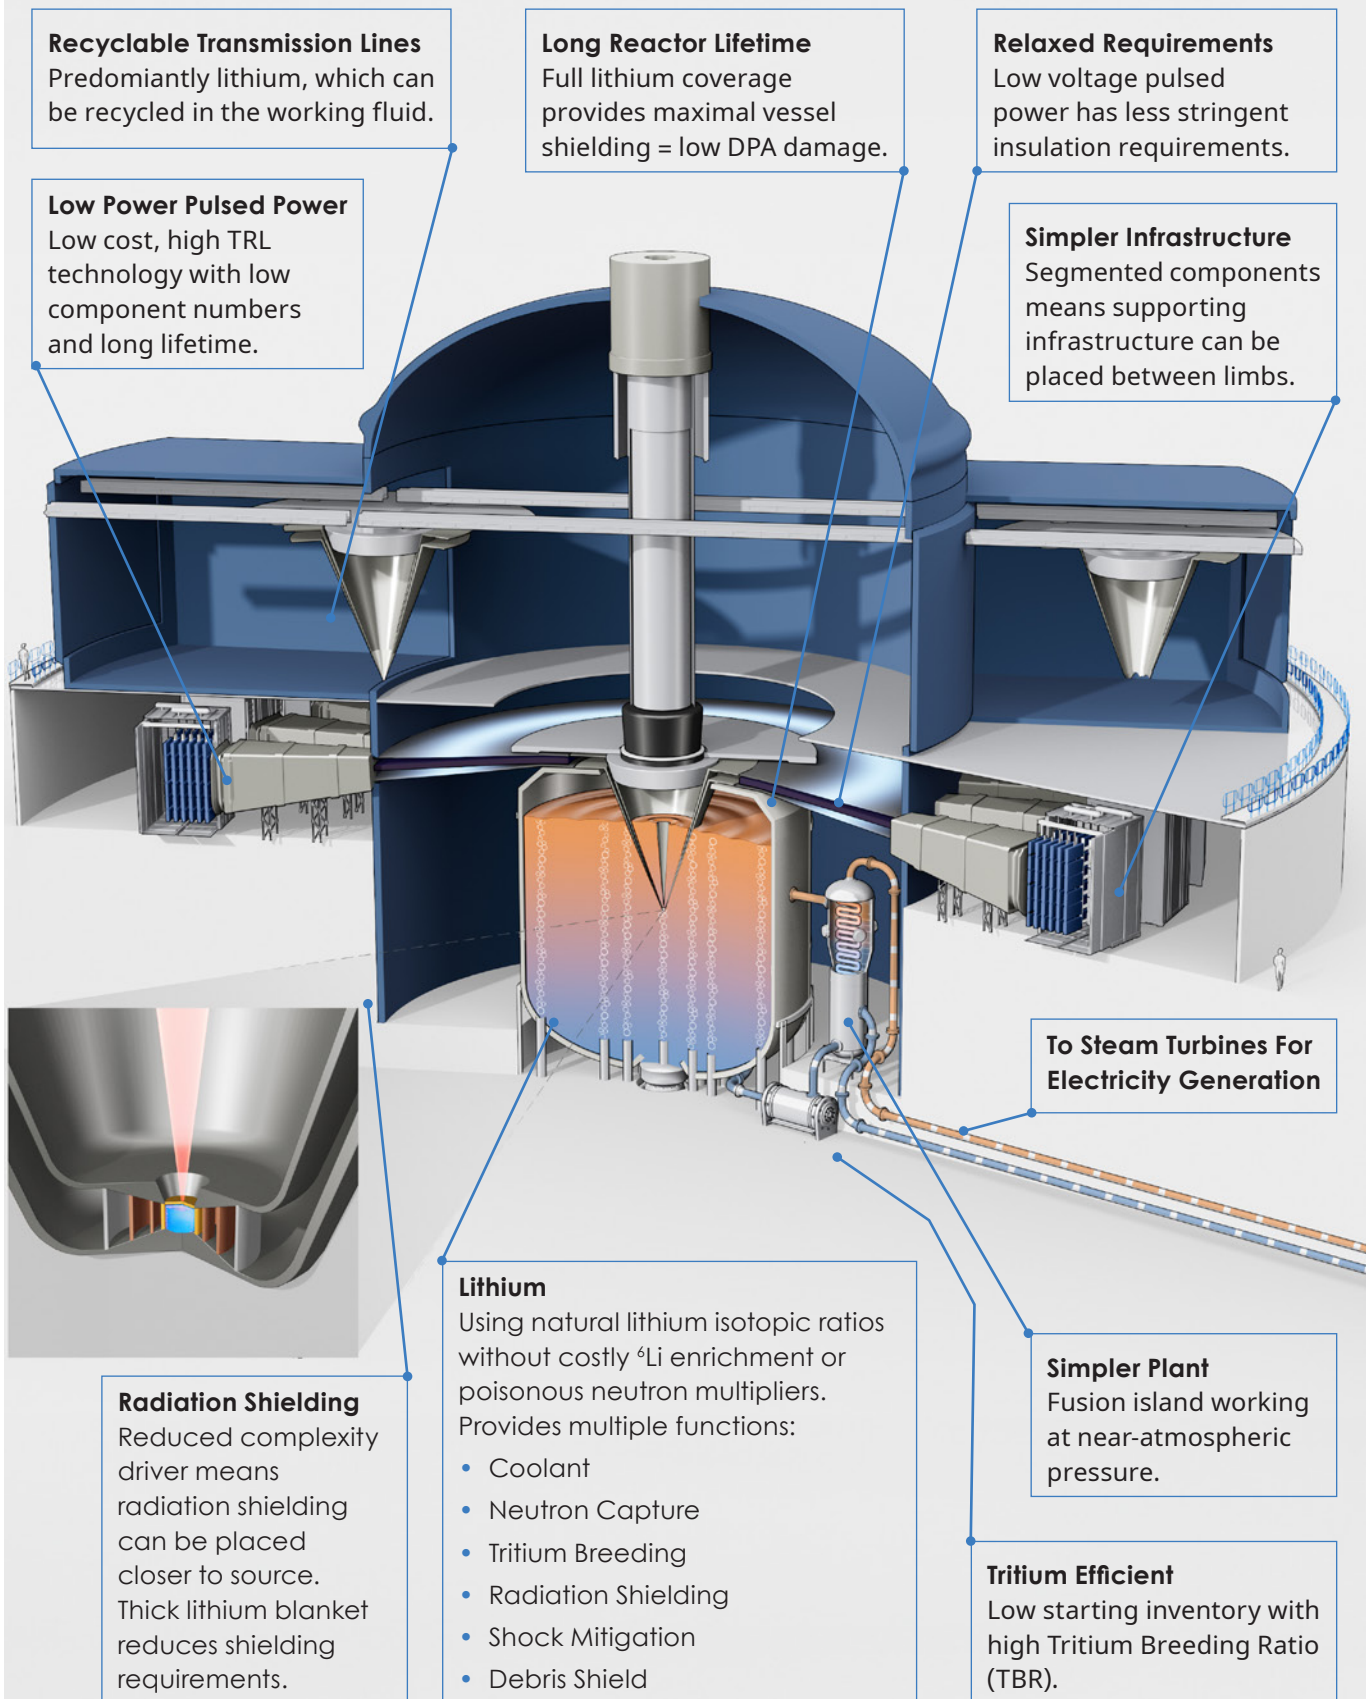
Both the fusion and IFE specific requirements can be simultaneously addressed with a correctly configured lithium blanket. The simplest implementation is a bath or pool of liquid lithium that fully surrounds the load (fusion fuel assembly). This is in comparison to more complex systems of curtains formed of hundreds of jets of liquid lithium or multi-layer wall systems that have been explored previously. These systems would have shorter lifetimes, require regular replacement, generate more radioactive waste, capture less neutron energy, breed less tritium and waste energy through pumping. The liquid bath configuration simultaneously solves the coolant, tritium breeding,

energy capture, neutron shielding, thermal protection, shock mitigation, and debris protection requirements.

- For compatibility with PWR nuclear islands, the operating temperature of the lithium is anticipated to be 600 K. This is a suitable working temperature, well above lithium's melting point (454 K), but far below its boiling point (1615 K).
- Using liquid lithium instead of water (as used in a PWR) means the coolant loop does not need to be pressurised to 150 bar to prevent boiling. The lithium loop can be operated at, or near, atmospheric pressure. Secure containment of activated material is still required but is facilitated by the lower pressure.
- Operating at 600 K, rather than the ~800 K proposed by Gen IV liquid metal reactors (defined by Generation IV International Forum (GIF)), brings several benefits. It allows the use of standard nuclear steels, keeps corrosion rates within acceptable limits, and provides higher solubility for tritium. While higher temperatures can improve thermodynamic efficiency, the combination of reduced corrosion, better tritium retention, and compatibility with existing PWR infrastructure supports the use of the lower temperature.
- High tritium solubility in lithium is important to prevent high partial pressures of tritium at the lithium surface, where it could otherwise permeate through the hot steel containment. The same applies to transfer pipes leading to tritium extraction systems and steam generators. Higher solubility also reduces tritium release in the event of an accident. However, it may increase the complexity of tritium extraction depending on the method used.

The effectiveness of a liquid lithium blanket is further enhanced when using natural, unenriched lithium comprising approximately 92% ⁷Li and 8% ⁶Li. ⁶Li is both a controlled substance and difficult and expensive to enrich, representing a major economic challenge for fusion systems requiring enrichment. The primary commercial process (Calex) is problematic on environmental grounds, due to the use of mercury. In the proposed lithium bath configuration, high-energy neutrons interact with the lithium blanket before being moderated or absorbed by structural components, allowing access to key high-energy interactions and enabling both tritium production and energy capture from the abundant ⁷Li (see the Breakout Box: Unenriched isotopic lithium advantages).

Figure 26: Conceptual reactor configuration.



Showing a possible configuration of the reactor and conventional steam and electricity generation. Liquid lithium flows through the steam generator and the steam used to generate electricity in a turbine in the usual manner.

Breakout Box: Unenriched isotopic lithium advantages

The ${}^6\text{Li} + n \rightarrow \text{T} + \alpha$ reaction is well known for its high cross-section (see Figure 27) at thermal energies and exothermic nature (releasing 4.8 MeV). However, ${}^7\text{Li}$ is often overlooked in reactor designs due to several common concerns:

- It has a relatively high reaction threshold (~5–6 MeV).
- The ${}^7\text{Li} + n \rightarrow \text{T} + n + \alpha$ reaction is endothermic (absorbing 2.47 MeV).
- Its neutron cross-section is significantly lower than that of ${}^6\text{Li}$ at low energies.

Despite these drawbacks, ${}^7\text{Li}$ offers important advantages:

- It constitutes over 90% of natural lithium, eliminating the need for enrichment. Enrichment of ${}^6\text{Li}$ is costly, environmentally problematic, and restricted due to its dual-use potential.
- Above ~5 MeV, the cross-section for ${}^7\text{Li}$ becomes comparable to or larger than that of ${}^6\text{Li}$ (see Figure 26).
- The ${}^7\text{Li}$ reaction is neutron-neutral, it emits a secondary neutron, which can then react with ${}^6\text{Li}$. This creates the potential for a single source neutron to yield two tritium atoms overall (TBR ≈ 2), while also balancing energy through sequential reactions.
- The endothermic reaction of ${}^7\text{Li}$ absorbs energy, but the follow-on ${}^6\text{Li}$ reaction is more exothermic, resulting in a net positive energy yield if the neutron is recaptured efficiently.

Additionally, several neutron interactions with lithium can result in:

- Neutron multiplication (via (n,2n) reactions), increasing breeding efficiency.
- In-situ conversion of ${}^7\text{Li}$ to ${}^6\text{Li}$, further boosting tritium yield over time.

Neutronics modelling shows net positive energy deposition of 10.3 GJ from 10 GJ of neutrons.

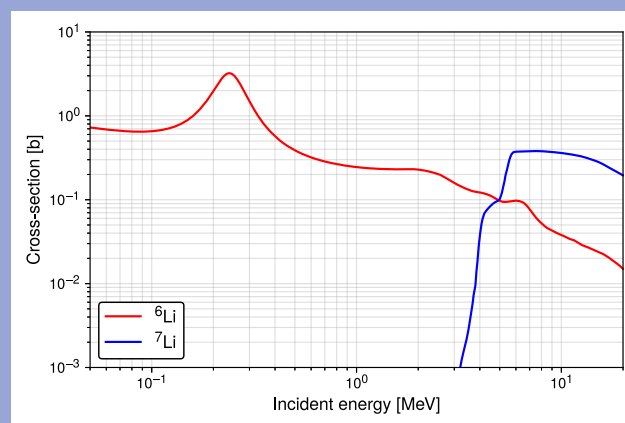
To exploit these advantages, high-energy source neutrons such as the 14.1 MeV neutron from D-T fusion, must interact with lithium before being

moderated by shielding or vessel walls.

- This requires placing lithium as close as possible to the reaction zone, ideally using it as the first wall.
- Achieving this configuration requires the chamber to avoid high vacuum, as hot lithium has a non-negligible vapour pressure that would otherwise compromise vacuum integrity.


Neutronics simulations confirm that very high Tritium Breeding Ratios (TBR > 1.8) are achievable using a 2 m thick lithium blanket. These results have been cross-validated using both FLUKA and MCNP simulation codes (see Figure 28).

Figure 27: Neutron cross-sections.



Cross-sections for the tritium producing reactions in ${}^6\text{Li}$ (MT=105) and ${}^7\text{Li}$ (MT=205) showing the higher cross-section of ${}^7\text{Li}$ above ~5 MeV reaction threshold [ENDFVII-B].

Figure 28: Cumulative tritium production through 2m thick lithium spherical blanket at 8m from a neutron source.



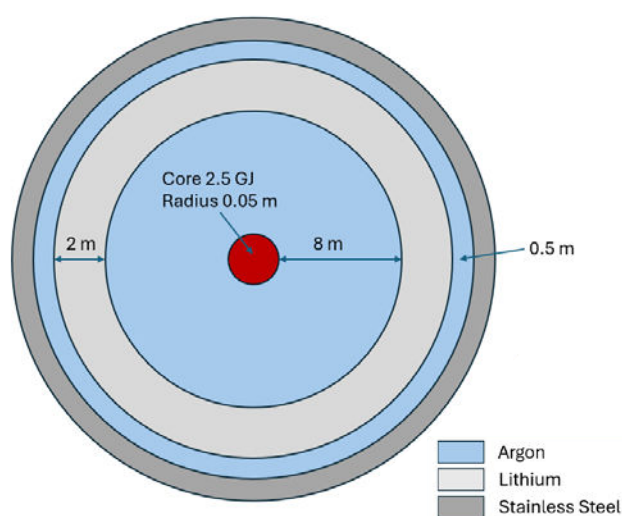
Calculated separately for ${}^6\text{Li}$ and ${}^7\text{Li}$ in MCNP where the TBR was found to be 2.06. Total tritium production per shot = 41.3 mg. Total for the year is 86.8 kg. (42.1 kg consumed in the d-t reactions). The tritium production in ${}^7\text{Li}$ plateaus as the average neutron energy drops below the threshold for this reaction. Tritium production in ${}^6\text{Li}$ begins to plateau as the available neutrons decline.

Neutron Energy Capture and Lithium Bath Effectiveness

The ability of the lithium bath to capture as much of the neutron energy as possible is a critical and primary function. Losses at this stage propagate throughout the entire system and reduce its overall efficiency. Modelling shows that a 2 m thick lithium bath allows the capture of 99.9 % of the energy produced, which also helps reduce neutron induced damage to the chamber (see Figure 29).

Figure 29: Lithium blanket configuration.

Reaction Chamber Cross-Section



Sketch of the reaction chamber arrangement showing the 5cm radius core heated with 2.5 GJ energy surrounded by 8 m of argon, 2 m of lithium another 0.5 m argon layer and finally a 5cm thick stainless-steel chamber.

- Using a 2 m thick layer of lithium, neutronics modelling has shown that only ~4% of the neutrons escape the lithium, and their average energy is reduced from 14.1 MeV to approximately 10 keV. This results in a primary neutron energy capture efficiency exceeding 99.9%.
- This compares favourably with previously studied lithium jet arrangements where neutron scattering around the jets resulted in ~25% neutron leakage.
- Modelling indicates that the lithium pool will increase in temperature by 2.4 degrees per 1000 tonnes per 10 GJ shot. This indicates that it will take approximately 123 shots to attain the 600 K operating temperature, equivalent to half an hour at 4 shots/minute.
- The heated lithium can be pumped to a heat exchanger for direct steam production, provided

that tritium loss to the water and the risk of lithium-water reactions are reduced to acceptable levels.

- The very high solubility of tritium in liquid lithium (especially when compared to lead–lithium eutectics), combined with low-permeability coatings, may mitigate tritium leakage. This could eliminate the need for a secondary coolant loop to isolate tritium from the water/steam system.
- The small number of low-energy neutrons emerging from the lithium also implies reduced neutron-induced damage to the surrounding chamber. Monte Carlo simulations (MCNP) estimate a Displacement Per Atom (DPA) of 2.3×10^{-9} per 10 GJ/shot in stainless steel. This corresponds to a projected structural lifetime exceeding 180 years of continuous operation at 4 shots per minute, well above the 60-year design life of the UK's current generation of PWRs. In terms of neutron fluence, this suggests compatibility between this fusion scheme and existing fission-grade materials.

FUSION-DRIVEN SHOCK LOADING AND CHAMBER SURVIVABILITY

Shock loading is a critical challenge in IFE reactors. As in the early development of the internal combustion engine, where pulsed energy was converted into continuous mechanical output, IFE must address the conversion of pulsed fusion events into a viable continuous power source. There are various sources of shock in the fusion event. For a working yield per shot of 10 GJ, it is found that by introducing a transient passive structure into the lithium pool, manageable chamber dimensions and lithium masses can meet shock and debris mitigation requirements.

- A fusion neutron yield of 10 GJ is energetically equivalent to approximately 2.5 tonnes of TNT, although the energy is deposited as high-energy particles rather than expanding gas. Nevertheless, this represents a substantial mechanical load.
- The primary source of energy is emitted as neutrons. However, this is not the primary source of physical shock in the system as the neutron attenuation length is on the order of 40 cm in lithium and so the energy is distributed over a large volume. This effect is further enhanced because the passive transient structuring of the lithium around the load, which pushes the lithium

blanket further out, lowering the fluence and increasing the mass of lithium heated.

- For every 10 GJ of neutrons produced there are also 2.5 GJ of alpha particles. As large, charged particles interacting through the Coulomb force, the alphas have an extremely short range on the order of micrometers.
- The energy density of the deposited alpha energy is therefore many orders of magnitude higher than that of the neutrons and the principal source of shock in the system.
- The main concern is preventing damage to the containment chamber. This is not only a matter of avoiding yield or hoop stress failure in a single shot, but of remaining below the fatigue limit of the chamber material over tens of millions of cycles.
- A large chamber is required to manage these loads. To stay within the fatigue limit of stainless steel at an operating temperature of 600 K, the peak chamber pressure should remain below 10 bar.
- Without additional mitigation other than the expansion zone around the load, hydrodynamic calculations performed using B2 show that for a spherical chamber of 10 m radius and 50 mm wall thickness, peak pressures at the chamber wall are only a few bar, well within the design specification. Recent simulations indicate that with the addition of passive transient structuring of the lithium, the chamber size may be significantly reduced.
- While large, these chamber dimensions are a manufacturable and practical starting point, especially if the vessel does not need to operate under pressure or in a vacuum.
- Although shock stresses are transient, and transient yield strength in stainless steel is typically higher than the static value (especially at high strain rates), a conservative approach using static yield limits has been adopted.

A key modification enabling the reduction of the chamber dimensions is the introduction of transient passive structuring into the lithium bath via argon gas curtains. This provides a volume for the shock to expand into before it interacts with the lithium blanket. The use of argon gas removes the need for:

- Retaining walls for the liquid lithium which would otherwise become the first wall.

- Using jets to form the blanket [ZIFE Phase I Report; SNL internal studies; Slutz et al.].
- Using a solid lithium blanket or retaining wall, both of which would have to be cast new for each shot with significant heat loss/lower efficiency.

MANAGEMENT OF RADIOACTIVE BYPRODUCTS

Although fusion reactors produce far less radioactive waste than fission reactors, and at much lower activity levels, they will still generate some radioactive materials during operation and decommissioning. Sources of waste include activated and tritiated reactor structural materials, precipitates, steel leachates and other impurities extracted from the lithium coolant loop, and slag generated from the recycling of fusion load materials.

- None of these waste streams are expected to be high-level waste (HLW), defined as having both high activity and significant heat generation. The majority will fall into the intermediate-level waste (ILW) category, and according to neutronics modelling, is expected to decay to low-level waste (LLW) within 10 years and fall out-of-scope (OoS) within 100 years.
- Materials used in the fusion load will be recycled on site.
- Neutron activation of structural components is greatly reduced by the high attenuation provided by the 2 m thick lithium blanket, which serves as the first wall. De-tritiation of steel is relatively straightforward and can be performed using thermal desorption methods.
- The lithium blanket itself, anticipated to weigh ~1000 tonnes, becomes only mildly activated. This is mainly composed of ^6He and tritium. All other radioactive nuclides together contribute less than 0.01%. ^6He is a beta emitter with associated bremsstrahlung and therefore readily shielded, decaying to ^6Li .
- A key regulatory consideration is the tritium inventory. The maximum permitted inventory in fusion reactors is likely to fall in the 0.1-1 kg range and can be minimised through efficient recovery.
- Numerous methods for extracting tritium from lithium have been proposed. In time, these may be able to directly reduce tritium content to OoS levels, allowing for lithium re-use and eliminating its classification as waste (from a tritium perspective).

6. A PATH TOWARDS ECONOMICALLY VIABLE IFE



Delivering IFE as a viable power source for the national grid requires a structured and transparent research pathway that enables stepwise progress grounded in fundamental science and supported by the broader research community.

Within this framework, the economic viability of any IFE concept is dominated by the design of the target. The target is not merely a consumable; it dictates the specification of the fusion driver, repetition rate, chamber architecture, and supporting balance-of-plant systems. Variations in achievable yield per shot directly influence the frequency of operation required to meet plant-level power output thereby impacting capital expenditure, operational costs, and overall system efficiency. Consequently, the selection of the target concept is pivotal: it defines the trajectory and cost profile of the research and development (R&D) pathway [Hawker, R. Soc. A (2020)].

Target choice determines not only the technical but also the financial contours of development; investment principles become central. A core principle of infrastructure investment is that risk and capital deployment must be aligned, (see Figure 30). Early-stage technical risk elevates the financial discount rate, limiting investability and slowing progress. Therefore, the two key investment questions for any fusion concept are:

- **Risk-Adjusted Cost Minimisation:** Does the approach lower the integrated, risk-weighted cost of delivering grid-connected fusion power?

- **Value Maximisation:** Does the resulting power plant, and the innovations generated along the development pathway, deliver long-term commercial returns?

Given the early stage of development, providing precise cost projections would imply a false level of certainty. Instead, this review emphasises the underlying physical and engineering principles that govern the drivers of cost to illustrate the potential of FLF's approach.

DEVELOPMENT PATHWAY

Leveraging Established Research

FLF's concept builds upon decades of inertial confinement research, consolidating proven scientific insights with proprietary technologies developed in-house. By adopting a component-first methodology, the system can be decomposed into modular subsystems that are amenable to rapid, low-cost testing on existing facilities. This approach mitigates risk by isolating potential failure modes early, reducing overall program exposure.

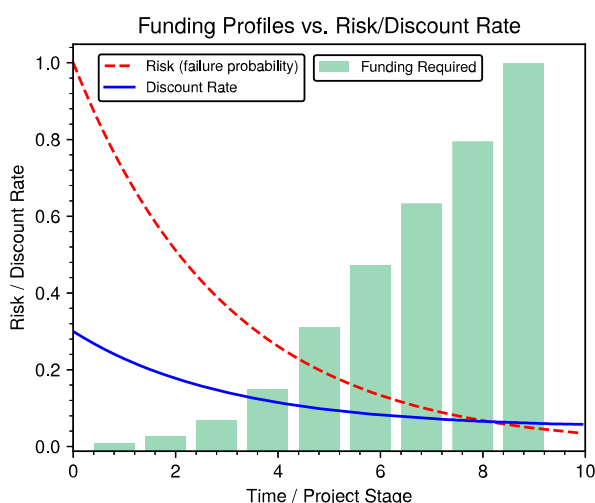
This component-first approach mirrors aspects of the United States Stockpile Stewardship and UK Warhead Assurance Programmes, where high confidence in system performance is achieved through rigorous modelling and sub-scale validation rather than full scale testing. Applying this philosophy to IFE reduces technical risk and enables progress without waiting for large demonstration facilities.

This approach also invites collaboration with academic institutions and the wider fusion community, ensuring transparency and reproducibility while reducing risk. Through partnerships and shared research programmes, a broad peer reviewed base of evidence can be assembled to validate the performance and scalability of each subsystem.

Accelerated Learning Through Rapid Iteration

The target is the dominant lever for system performance. Its design governs achievable gain and, consequently, plant economics. FLF's ability to prototype and evaluate targets at high cadence, combined with advanced simulation and data science tools, enables accelerated feedback loops.

Figure 30: Project funding idealised risk against timing.



Plot notionally showing how risk and discount rate decrease as the project matures.

Early-Stage Validation Using Existing Infrastructure

The initial phases of development can be executed on current experimental platforms, minimising capital outlay while reducing risk:

- **Proprietary Facilities:** FLF's existing experimental platforms provide unique capabilities for target compression and shock generation.
- **Institutional Collaborations:** Partnerships such as the CRADA with Sandia National Laboratories, the Z Fundamental Science Program and the Prosperity Partnership allow distributed component testing and validation.
- **Laser Facility Access:** Engagement with external laser laboratories enables investigation of ignition physics without immediate need for new high-cost infrastructure.

This approach has already demonstrated results. For example, the 2.5 TPa in quartz achieved with the Big Friendly Gun (BFG) illustrates that capabilities once requiring multi-billion-dollar national facilities can now be reached with lower-cost, modular systems.

Demonstration Facility Requirements

The transition from component validation to net energy gain will ultimately require a dedicated pulsed power driver. Scaling laws indicate that a gain-relevant pulsed power machine could be constructed for \$100M–\$200M (50MJ to 100MJ at \$2/J), representing less than 5% of the inflation-adjusted cost of the upgraded NIF (~\$5.3B), see Table 1. Critical attributes of this approach include:

- **Mature Driver technology (high TRL):** Reduces technical and supply chain risks.
- **Lower Intensity:** Enables cost-effective, robust systems with reduced component stress.
- **Integration Flexibility:** Compatibility with existing short-pulse laser facilities for hybrid component validation.
- **Staged Investment:** Capital commitments can be deferred until major technical risks are retired, improving risk-adjusted returns.

Table 1: Machine Cost Comparison, the PP drivers assume 80MJ stored energy.

	Low Intensity PP	High Intensity PP	NIF Type Laser
Estimated Cost	~\$160M	~\$480M	~\$5,300M

Additional Considerations

Other power plant challenges can be addressed more effectively through rapid derisking and the shared risk–reward dynamics of a strong partnering ecosystem. By simplifying the most difficult technical hurdles, opportunities open for technology transfer and read across from adjacent industries, accelerating progress and reducing development costs.

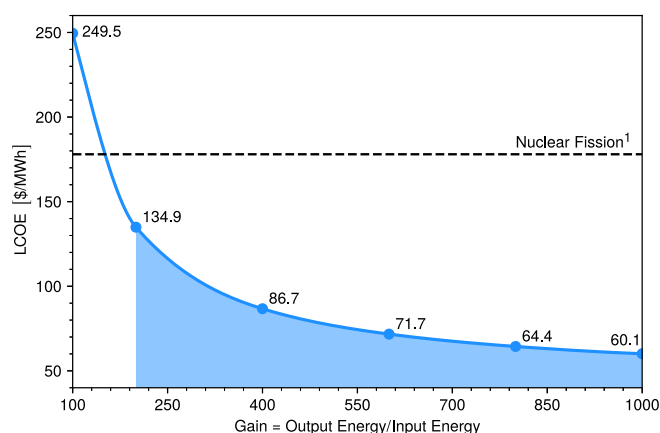
In addition, partners stand to benefit from the broader innovation journey itself. FLF has already generated a portfolio of valuable products and services from its research, and further value is likely to emerge along the path toward commercial fusion energy. These intermediate innovations not only enhance the return on investment for partners but also build momentum and confidence in the long-term goal of delivering fusion power.

POWER PLANT IMPLICATIONS

High-Gain Target Physics

High gain is the single most important determinant of IFE plant economics. Monte Carlo studies indicate that gains above 500 are associated with optimal Levelised Cost of Energy (LCOE), (see Figure 31). Even moderate gains above ~200 deliver substantial cost reductions by lowering the required shot repetition rate for a given power output [Hawker, R. Soc. A (2020)].

Figure 31: High gain lowers LCOE



Plot showing the effect of gain on LCOE (\$/MWh) for a given system based on a simple economic model for fusion energy, Hawker, R. Soc. A (2020). Fission data from Lazard, v18.0 (2025).

FLF's concept enables high gain through two mechanisms:

- Variable fuel loading, allowing gains to be "dialled up" by increasing target fuel mass with minimal additional fabrication cost.
- Low pressure argon chamber with liquid lithium working fluid, capable of managing output energies up to ~10 GJ (gain ~1000), while enhancing heat transfer and providing intrinsic neutron shielding.

These features and low repetition rate (modelled as low as 0.25Hz) reduce the number of targets and Recyclable Transmission Lines (RTLs) required per gigawatt-hour, extend system lifetime by reducing wear, and simplify turnaround between pulses.

Low-Power Pulsed Power Drivers

Driver cost dominates IFE plant capital expenditure. FLF's approach leverages low-intensity pulsed power drivers that offer:

- High efficiency system: ~20% of wall-plug energy reaches the load.
- Favourable cost scaling at ~\$2/J compared to:
 - » High-voltage Z-pinch: \$6/J (ZR refurbishment project, total Z cost >\$13/J).
 - » Laser based systems: \$13/J (NIF including upgrade).
- Reduced technical complexity, lowering failure rates and operational costs.
- High readiness levels, given reliance on mature pulsed power technologies.
- Relaxed tolerance requirements for RTL fabrication, further lowering costs and easing supply chain constraints.

System-Level Cost Advantages

The proposed low-intensity Pulsed Power architecture provides ancillary benefits that extend beyond the target and driver:

- Simplified plant infrastructure, e.g., no large single-span machine halls).
- Simpler plant engineering, e.g., Fusion Island working at near-atmospheric pressure.
- Reduced large-volume materials costs, e.g. natural lithium isotopes serving as an affordable, multifunctional working fluid (approx. £70M,

significantly below equivalent fission materials and other proposed fusion materials (with current cost >10X this).

- Recyclability of RTLs, lowering consumables expenditure and the waste management burden.
- Improved plant longevity and significantly decreased waste stream, due to effective neutron shielding by liquid lithium.
- Naturally high Tritium Breeding Ratio (TBR) and low start up inventory.
- Potentially lower development costs for core components due to synergies with adjacent industries, including the fission sector (balance-of-plant technologies) and automotive manufacturing (recyclable components).

The low intensity pulsed power fuel assembly, coupled to the reduced ignition requirements laser, would result in a 10x lower driver cost compared to a purely laser-based system.

Table 2: Costs in \$Millions.

	HFI: PP + SP Laser	Laser FI (Low Est)	Laser FI (Mid Est)
Fuel Assembly	150	1,500	2,400
Ignition	180	1,800	2,500
Full System Total	330	3,300	4,900

Note – Based on "Focused Energy, A New Approach Towards Inertial Fusion Energy", HFI with low intensity Pulsed Power fuel assembly and Short Pulse Laser ignition, using scaling from NIF (low) and Laser Megajoule". These are not production cost estimates, which would be different.

Additional Considerations

As subsystems mature, there is also a need to validate the behaviour of reactor relevant technologies that fall outside the target and driver, such as recyclable transmission lines, neutron shielding, tritium handling, and liquid lithium systems. These components often reach higher technology readiness levels more quickly than the fusion core but must still be integrated and qualified. In this regard, collaboration with United Kingdom Atomic Energy Authority (UKAEA) and the broader civil nuclear sector will be essential.

There are strong parallels with the wider nuclear industry, particularly in materials science, high-integrity engineering, and coolant loop design.

Fusion systems will expose materials to extreme thermal and radiative environments, and the nuclear industry's experience in qualifying components, managing activation and corrosion, as well as developing safe remote handling systems will directly support IFE development.

Flexibility and Grid Value

Modern electricity markets, increasingly dominated by variable renewable energy, place high value on dispatchable and flexible generation. A plant based on FLF's FLARE concept could modulate electrical output on hourly timescales by varying pulse repetition rates. This capability allows plants to act as flexible baseload providers, increasing their commercial value. Moreover, the relatively low capital intensity reduces the financial pressure for continuous full-power operation, supporting participation in this use case.

Funding most fusion approaches requires substantial costs in early phases of the research. This presents a challenging investment case, with some investors believing that, for grounded technologies with a strong scientific pedigree, only the most highly funded projects can succeed. The development and power plant benefits of FLF's FLARE concept challenge this assumption, creating an asymmetric investment case. The downside risk is bounded, yet this offers the enormous upside reward of creating a power plant technology that can be efficient, high value and globally deployable.

INVESTMENT CASE

This pathway creates a compelling, asymmetric investment opportunity:

- **Reduced downside risk:** Grounded in established science, the early program phases require modest funding and limited capital investment, enabling rapid de-risking using existing infrastructure.
- **Scalable upside:** This technology has the potential to scale rapidly due to the high TRL supply chain, partnering ecosystem, and compelling price point. It could offer power solutions from low hundreds of megawatts to multi-gigawatts, addressing different market needs and capital appetite. It offers high gain i.e., more out for a given input, which is the most significant lever for achieving competitive LCOE. Unlocking the potential to address the multi-trillion-pound global market for clean, affordable, dispatchable energy.
- **Diversified returns:** Non-fusion products such as advanced simulation tools, data-science methodologies, and recyclable transmission line technology, provide near-term commercial value aligned with the fusion program.
- **Partnership leverage:** Collaboration with industrial and national research partners ensures access to facilities, expertise, and supply chains without needing bespoke infrastructure. Risk and rewards are shared.

APPENDIX



REVIEW OF FIRST LIGHT FUSION'S HISTORICAL STRATEGY AND CAPABILITIES

The FLARE concept builds on more than a decade of research and development at First Light Fusion. Research and development completed on uniaxially driven fusion systems provided a vital focus for developing the capabilities, scientific understanding, engineering and manufacturing skills that underpin our progress today. Combined with the design of our M4 demonstration system, which offered crucial insights into target physics, driver technology, and reactor integration, these foundations have enabled us to define a new end-to-end concept for delivering fusion energy, grounded in evidence and progressing along a credible path.

FLF was founded with the aim of reducing the complexity and cost of fusion power plants through innovative target design. Asymmetrically driven target concepts were identified as a promising route to achieve this goal. The increased design complexity required to realise such concepts became tractable when established theoretical target behaviors were combined with state-of-the-art numerical modelling, data science, and machine learning techniques. This research pathway led to a focus on impact fusion, in which the energy is delivered to the target via a hyper-velocity projectile.

- A critical component of a fusion reactor is a lithium blanket to capture neutron energy and breed tritium fuel for subsequent reactions.
- A uniaxially driven system allows for fewer penetrations into the chamber improving the integrity and performance of a liquid lithium first wall, offering key design advantages in performance and cost.
- Such design improvements include:
 - » Enhanced tritium breeding and energy capture.
 - » Better thermal load management.
 - » Reduction of material activation and neutron damage.
- Impact fusion via a hyper-velocity projectile offers specific engineering advantages, including enabling the fusion event to occur at a standoff distance from sensitive or structural components, contributing to improved system economics via:
 - » Reduce neutron-induced damage.

- » Extended component lifetimes.
- » Lower waste generation from replaceable driver components.
- Projectile-driven inertial fusion demands high precision in both energy delivery and target design.
- Unlike spherically driven capsules, uniaxial targets require larger and more intricate target designs to wrap an initially planar shockwave around a spherical fuel capsule with the correct intensity and implosion symmetry.
- The inherent and unavoidable losses due to the shaping requirement increase driver energy requirements.
- Small misalignments, such as projectile tilt or target positioning errors, can introduce significant asymmetries that impair performance.
- These technical requirements informed the engineering strategy and experimental design.
- The need for precision in target fabrication and energy delivery guided research priorities and technology pathways.

Validation Facilities

A key early challenge in FLF's development was access to suitable experimental platforms. At the time of the company's founding, there were no access pathways to external HED research facilities for validating target concepts. Initial experiments studying the underlying physics therefore had to be conducted using limited, non-specialised resources. These challenges led to the design and commissioning of multiple in-house facilities on a limited budget and the eventual design of an ignition scale driver to study uniaxial driven fusion concepts.

- One of the first platforms used was a gas gun at the University of Oxford, originally intended for solid mechanics experiments. Although not designed for fusion research, the system enabled initial tests of projectile-target interactions and provided a means to explore key aspects of the proposed approach under controlled laboratory conditions.
- Working with a limited budget, particularly in contrast with the high cost of conventional fusion research infrastructure, available options for delivering high-energy impacts were limited. Natural candidates for low-cost validation facilities were two-stage light gas guns (2SLGGs)

and pulsed electromagnetic (EM) launch systems.

- This led to the design and commissioning of custom-built platforms, including our flagship facilities, the BFG (2SLGG) and M3 (EM). Both remain in active use and have formed the foundation of FLF's in-house research capability (see the Breakout Box: M3 and BFG – High Energy Density research facilities at FLF).

- To support these experimental facilities, a suite of high-resolution diagnostics was developed to capture spatial and temporal detail across transient, HED phenomena. Additionally, diagnostic variants of targets were fabricated to isolate specific physical mechanisms, allowing for model refinement and targeted validation (see the Breakout Box: Experimental diagnostics capabilities at FLF).

Breakout Box: M3 and BFG – High Energy Density research facilities at FLF

BFG

The 'BFG' is a 38 mm bore two-stage light-gas gun. The hypervelocity projectile is principally used to drive amplifier experiments but has also been used for micro-impact experiments. It was commissioned in 2021 and has fired over 300 (311) shots to date. Designed and commissioned within 11 months.

- Total energy in the projectile 2 MJ
- Impact velocity of 6.5 km/s
- Impact planarity less than 0.5 degrees from normal
- 3 kg of gun powder per shot
- Approximately 25 m in length including impact chamber

M3

Machine 3 (M3) is a low impedance, bi-polar, capacitor discharge pulsed power generator. Reliable plate flyer, electric gun and liner platforms have been developed for study on M3. It was commissioned in 2018 and has fired over 250 (255) shots to date. Designed and commissioned within 18 months at a cost of £3.6 million.

- Total stored energy on 2.5 MJ
- +/- 100 kV charge voltage resulting in 200 kV total
- 10 MA current delivered in 2 μ s
- 192 capacitors of 2.6 μ F each
- 10 m diameter




Figure 32: The BFG facility at FLF in Oxford.



Figure 33: The M3 pulsed power facility at FLF in Oxford.

Breakout Box: Experimental diagnostics capabilities at FLF

FLF has an extensive, evolving, diagnostics suite centred around optical, X-ray, neutron, electromagnetic field measurements. Analysis is being carried out using in-house tool suites and advanced data-science, AI, and techniques.

Diagnostics:

- 2 Shimadzu HPV-X2 high-speed cameras
 - » Frame spacing as low as 100 ns (10 million frames per second)
- SIMX 16 frame ultra-high-speed camera
 - » Frame spacing as low as 1 ns (1 billion frames per second)
 - » Exposure per frame as low as 3 ns
- 2 Sydor Ross 2000 optical streak cameras
 - » Sweep windows as short as 10 ns
 - » Imaging line VISAR, fibre-based point VISAR
 - » Imaging VISAR utilises a 532 nm, multi kilowatt pulsed laser with durations up to 5 microseconds
 - » Line VISAR is captured by streak cameras and is capable of approximately 10% fringe resolution with less than 50 ps timing resolution
 - » Point VISAR utilises a 532 nm, 6 W continuous wave laser, with 5% fringe resolution and 150 ps timing resolution
- Photon Doppler Velocimetry
 - » 8 channels operating at 1550 nm wavelength
 - » Tuneable reference lasers allowing for velocities up to 70 km/s to be measured
- Fibre optic current sensing
 - » Routinely measures up to 14 MA to assess the machine current on M3
 - » 1550 nm wavelength
 - » Time-resolved optical spectroscopy
 - » PI Acton SP 2360 spectrometer coupled to a streak camera
 - » Gratings of 150 l/mm, 600 l/mm, and 1200 l/mm
- Time-resolved X-ray photodiodes
 - » Silicon PIN diodes with rise time less than 1 ns

- » Broadband spectroscopy can be achieved using an array of diodes and filters of varying materials
- X-ray point projection radiography
 - » 10-20 keV broadband source produced using an external X-pinch driver (140 kA, 350 ns)
- Neutron Detection:
 - » High efficiency, large format (1000x500x50 mm) scintillators coupled to photomultiplier tubes
 - » ^4He detectors for discriminating between gamma rays and neutrons
 - » Neutron bubble detectors for measurement of higher neutron flux than scintillator-photomultiplier tube detectors



Figure 34: Optical diagnostics at FLF in Oxford.




Figure 35: Radiographs of a pulsed power driven flyer plate experiment.

Flyer plate driven by 1 μs 2 MA pulse on CEPAGE generator at FLF. Auxiliary generator provides point source of > 10 keV X-rays from an x-pinch implosion [J. Read, Rev. Sci. Instrum. (2024)].

- Compared to large HED facilities typically used in ICF research, these internally developed drivers offer a favourable cost-to-energy ratio. Their relative simplicity also contributes to lower maintenance requirements and improved operational reliability. These characteristics enabled the team to conduct a sustained programme of experimental research without the prohibitive costs that commonly hinder the private fusion sector.
- Building on this research programme, FLF developed a conceptual design for Machine 4 (M4), a pulsed power facility aimed at demonstrating ignition-scale conditions using magnetically launched flyer plates. The M4 concept was influenced by next-generation pulsed power systems under study at Sandia National Laboratories (e.g., Z300, Z800, NGPP), which were designed to deliver peak powers in excess of 300 TW and currents above 30 MA (see the Breakout Box: M4 design outcomes).

Breakout Box: M4 design outcomes

The conceptual design of Machine 4 was completed in January 2024, using a custom-built systems code to explore the complex coupling between machine, load, and target performance. The model enabled wide-ranging parameter scans and exposed non-obvious trade-offs, challenging conventional assumptions.

By factoring in engineering risk, machine and building costs, and component readiness levels, the optimisation favoured longer pulse lengths, between 300 and 450 ns to peak current. This kept the design within a regime similar to the Z machine, leveraging known architectures and technologies where confidence in performance is high.

This design exercise highlighted the value of integrated system modelling in fusion, allowing physics, engineering, and cost to be balanced in a coherent way, and helping identify viable design spaces that are both ambitious and grounded in reality.

- Stored energy 66 MJ
- 87.5 kV capacitor charge voltage
- 5.25 MV Marx voltage
- 6650 capacitors 2.6 μF each
- 3 layer design
- 48 m diameter



Figure 36: Conceptual design point for the M4 pulsed power facility.

- Through target-based power amplification, the M4 system requirements for machine voltage and power could be reduced relative to these benchmark systems. This allowed scaling of existing, mature pulsed power technologies in comparison to newer higher voltages architectures being explored.
- The final M4 concept involved approximately 180 TW delivered to the vacuum section, with around 3 MJ of energy transferred to the flyer plate over a ~300 ns pulse. The system was predicted to attain velocities in the range 50–60 km/s.
- While this output represented more than twice the peak power of the existing Z Machine, the reduction in required system voltage, enabled by the long-pulse acceptance and internal amplification features of the targets, resulted in a significantly lower projected cost, (see Table 3).

TARGET PHYSICS FOR HIGH-ENERGY LOW-POWER SYSTEMS

- While FLF's in-house facilities are energetically comparable to other major HED and ICF platforms (see Table 1), there is an important distinction in how this energy is delivered to the target. Conventional fusion facilities typically operate with very short pulse durations, resulting in significantly higher power and drive intensity. These high-intensity conditions are important not only for overcoming radiative and conductive losses, but also for limiting the available growth time for hydrodynamic instabilities that can disrupt the compression of the fuel to the hot, dense states required for ignition. This led to the design of targets capable of amplifying the intensity from the projectile to the fuel whilst simultaneously shaping the shockwave from planar to spherical.
- Rather than attempting to match these high powers and intensities, which would require projectile velocities on the order of hundreds to thousands of km/s, FLF focused on a lower peak drive intensity combined with internal amplification mechanisms within the target.
- These mechanisms include geometric convergence, shock reflection, pulse shaping, and ramp compression, each of which contributes to concentrating energy around the fuel within the target structure.
- Because compression and heating to ignition conditions is not viable in a purely planar geometry, the design approach incorporated shock shaping and topological control strategies. These were intended to achieve the benefits of increased volume convergence associated with spherical geometries, while still using a one-sided driver.
- Together, these functions effectively increase the pressure and apparent drive velocity at the fuel interface, meeting a desired loading profile without requiring high intensity at the initial coupling point between the driver and the target.
- This approach enables the use of lower-power, more technologically mature driver systems, while still aiming to reach the extreme pressures and densities necessary for fusion.
- The resulting design strategy informed FLF's experimental programme and served as a basis for evaluating cost-performance trade-offs in inertial fusion system architecture.

Table 3: Comparison of FLF platforms to leading national and international HEDP/IFE ICF facilities.

Platform	Type	Stored Energy (MJ)	Drive Energy (MJ)	Peak Deliverable Power (TW)	Pulse Length (ns)	Drive Intensity (W/cm ²)	Approximate facility cost (M\$)	Estimated Inflation adjusted Facility cost (M\$)
BFG	2SLGG	13	2	2	1000	10 ¹²	1	2
M3	Pulsed EM	1.2	0.1	0.7	2000	10 ¹¹	4	6
M4	Pulsed EM	80	>3	180	300	10 ¹⁷	300	311
Orion	Laser	8	0.005	200	1	10 ¹⁶	200	320
Omega	Laser	39	0.04	30	0.2-5	10 ¹⁹	100	262
Z	Pulsed EM	22	3	80	100-1000	10 ¹⁷	200	341
LMJ	Laser	450	1-2	400	1	10 ¹⁵	3000	5115
NIF	Laser	400	2	500	1	10 ¹⁵	3500	6888

Numbers are representative and can range (in some cases widely) depending on several factors.

NUMERICAL MODELLING, DATA SCIENCE AND OPTIMISATION

- As FLF's understanding of target physics and reactor integration progressed, so too did the complexity of the systems required for simulation and design. Achieving a spherical implosion from a uniaxial driver over extended pulse durations, while maintaining symmetry and achieving amplification, introduced significant design challenges. These problems could not be addressed through experiment alone and required high-fidelity computational models capable of resolving coupled, physical phenomena with sufficient accuracy and robustness. Similar to the issues experienced with access to experimental facilities, available numerical models were not suited to the challenge, leading to the development of an in-house suite of capabilities (see the Breakout Box: Multiphysics hydrodynamics modelling at FLF).
- At the time of founding, the existing computational toolsets for ICF and HED physics were limited in scope. Many off-the-shelf codes lacked the flexibility, numerical performance, verification infrastructure, or core coupled capabilities needed to support non-standard drive geometries or extended drive conditions.
- In response, FLF developed a bespoke suite of multi-physics simulation tools, including magneto-hydrodynamics solvers, systems-level reactor models, and flexible simulation frameworks, tailored to the company's specific design space. These tools are built around modern verification and validation practices and incorporate material models, such as advanced equations of state and strength models, necessary for simulating structured targets under extreme pressure and temperature conditions (see the Breakout Box: Multiphysics hydrodynamics modelling at FLF).
- Despite these capabilities, the design problem remained computationally intensive and under-constrained. The presence of large parameter spaces, model uncertainty, sparse experimental data, and mixed-fidelity simulations rendered traditional design workflows insufficient. To address this, FLF incorporated methods from statistical learning and Bayesian inference to improve calibration, parameter sensitivity studies, and optimisation processes (see the Breakout Box: Integrated simulation design environment using machine learning and probabilistic artificial intelligence).

- This data-centric approach led to the development of an integrated simulation design environment combining efficient sampling algorithms, surrogate modelling, uncertainty quantification, and simulation-based inference. The resulting framework enables structured exploration of design trade-offs and supports a hybrid optimisation methodology adapted specifically to the high-dimensional, multi-objective nature of inertial fusion target design.
- These computational and statistical tools form the core of FLF's current target design workflow, supporting rapid prototyping, design optimisation, and performance evaluation across a wide range of operational regimes.

THE STAND-OFF PROBLEM AND ALTERNATIVE DRIVER TECHNOLOGIES

Translating target physics results into a practical fusion powerplant introduces several engineering and physical challenges. A primary issue for impact fusion is the stand-off requirement: maintaining projectile integrity and alignment over meter-scale distances between the driver and target. It became evident that meeting these requirements for projectile-based systems would not be viable. As a result, a strategic review was conducted to assess the application of FLF's target concepts to alternative driver technologies, with the aim of evaluating the physical viability, economic feasibility and TRL of each option.

This assessment led to a re-evaluation of target-driver integration pathways, with pulsed power systems featuring a direct electrical connection to the target emerging as the most promising candidate. This led directly to the development of the FLARE concept - a high-TRL, scalable approach to IFE based on low-power pulsed power.

Collaborations and partnerships built over a decade of scientific research, position the company to make rapid progress across multiple disciplines. Recent innovations and existing IP are also unlocking opportunities for early revenue generation in adjacent markets.

Breakout Box: Multiphysics hydrodynamics modelling at FLF

High fidelity multiphysics simulations underpin our target design process. The software we have developed in house to do this must simultaneously solve several integrodifferential equations, covering an enormous range of space and time scales, and states of matter.

There are two principal complexities of developing software of this type. Firstly, reducing the model complexity to be tractable to HPC without losing accuracy. Secondly, the high degrees of mathematical nonlinearity throughout the equations themselves, and their closures, require a high degree of robustness in the algorithms both to accuracy and execution. This is particularly important when the software is used in ML and data science applications where thousands of simulations are required to provide the data basis for these tools.

The result of these complexities is that the “know-how” of developers solving these problems can introduce subjectivity and sensitivity of the numerical solution. Even if the underlying equations of the model are agreed, the solution strategy and complexity of this computational system can introduce unexpected consequences. We have therefore developed a range of methods to solve given problems in the codes. This has been essential to facilitating numerical and model understanding and qualification of correct model implementation, and ultimately their suitability for reproducing real world observations.

Using diverse team expertise in modelling and simulation for fusion, HEDP, plasma physics, astrophysics, aerothermodynamics, chemical, electrical and mechanical engineering we have built a comprehensively verified and validated suite of models covering:

- Compressible Hydrodynamics
 - » Eulerian and Lagrangian, Finite-volume (FV) discretisation
 - » Multi-order gradient reconstruction schemes
 - » Several flux calculation schemes:
 - Exact Riemann solvers, robust to the complexities of highly nonlinear EoS and tabulated models
 - Numerous approximate Riemann solvers

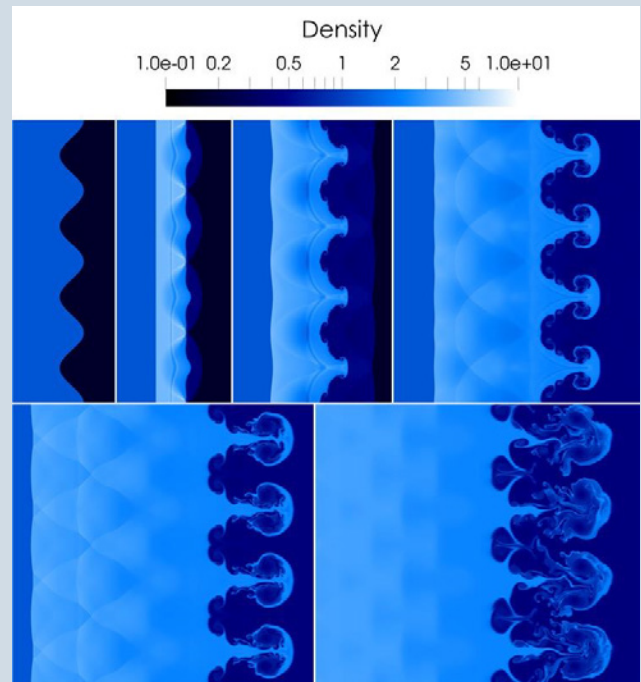


Figure 37: FLF's next generation multi-physics code for IFE and HEDP.

Maps of density at selected times from a simulation of a 2 km/s flow impacting a sinusoidal air-helium interface (1 m wavelength, 0.2 m amplitude), exhibiting Richtmyer-Meshkov instabilities. The mesh was refined adaptively around density gradients, with a resolution of 0.977 mm at the finest level, and ~7 million cells in the final snapshot at 10 ms. The simulation domain was 32 m long, but the snapshots in the Figure have been cropped to show only the region around the interface.

- » Thermodynamic non-equilibrium
 - Electron energy conserving scheme
 - Adiabatic electron approximation (electron entropy conserving scheme)
- Grid topologies
 - » Cartesian, axisymmetric and spherical symmetries
 - » 1D, 2D, 3D
 - » Material interface transport with several options for flux evaluation
 - » Eulerian material interface-capturing: Volume of Fluid (VoF)
 - » Lagrangian material interface tracking: Front Tracking (FT)
- Magnetohydrodynamics
 - » Ideal and resistive MHD
 - » Displacement current inclusive
 - » Magnetic field conservative staggered grid approach [Stone et al., Astrophys. J (2008)]
 - » Coupled circuit models

- Dissipative transport
 - » Flux-limited thermal conduction
 - » Newtonian viscous stress tensor
 - » Material strength and solid mechanics – Johnson-Cook Model
- Radiation transport
 - » Multigroup flux-limited diffusion approximation
 - » Multigroup flux-limited P1 moment closure
 - » P1 correction for vacuum light speed propagation (P1/3)
 - » Multigroup Monte Carlo solution of the RTE
 - » Laser-matter interaction
 - Ray-traced (geometric optics) with plasma based refractive index
 - Opacity based attenuation
- Nonlocal transport
 - » SNB electron heat flux [Schurtz et al., Phys. Plasmas 7, 4238 (2000)]
 - » RKM ion/electron multispecies electrothermal coefficients [Mitchell et al., Plasma Phys. Control. Fusion 66, 075005 (2024)]
 - » Explicit Monte Carlo algorithm coupled to tabulated/in-line stopping power models
- Fusion reaction and energy deposition
 - » Single group alpha diffusion [Atzeni and Caruso, Nuovo Cimento, 64B, 383 (1981)]
 - » Local deposition
 - » Various electron-ion energy partition schemes, e.g., Fraley [Fraley et al. Phys. Fluids 17, 424 (1974)]
 - » Bosch-Hale & Nayak fits to thermal reactivity [Bosch and Hale, Nucl. Fusion 32(4) (1992); B. Nayak, Annals of Nuclear Energy 60:73-77 (2013)]
 - » Nonlocal reactivity correction [Albright et al. Physics of Plasmas 20:122705 (2013); Molvig et al. Phys. Rev. Lett. 109:095001 (2012)]
- Equation of state
 - » QEOS based FEOS and SpK, both with material-specific data supplementation and calibration [Faik et al., Comput. Phys. Commun. 227, 117 (2018); Fraser et al., HEDP 53, 101136 (2024)]
 - » Helmholtz free energy-based decomposition for electrons and ions:
 - Cowan's semi-empirical blended model for the ions
 - Semi-classical Thomas-Fermi average atom approach for electrons (FEOS) and degeneracy-corrected Saha-Boltzmann population solver (SpK)
- » Bonding corrected for improved thermodynamics at reference state
- Charged-particle microphysics closures
 - » Ion transport (ion thermal conductivity, dynamic viscosity and diffusivity) from effective potential solution to linearised Boltzmann equation [Stanton and Murillo, Phys. Rev. E 93, 043203 (2016)]
 - » Electron transport (electron thermal conductivity, electrical resistivity) from generalised relaxation time approximation, integrated with material-specific data supplementation for improved agreement at ambient conditions [Lee and More, Phys. Fluids 27, 1273 (1984); Apfelbaum, Phys. Rev. E 84, 066403 (2011)]
 - » Reduced Fermi-golden rule approximation to volumetric electron-ion energy exchange, extended with finite-wavelength screening and state-of-the-art local field corrections [Hazak et al., Phys. Rev. E 64, 066411 (2001); Chapman et al., Nat. Commun. 6, 6839 (2015); Dornheim et al., Phys. Rev. Lett. 125, 235001 (2020)]
- Plasma and atomic opacity
 - » Detailed configuration accounting
 - » Bound-Bound and Bound-Free transitions
 - » Free-Free (Bremsstrahlung and Compton scattering)
 - » Busquet's simplified effective ionisation temperature model for non-equilibrium (RADEON)
 - » Various ionisation potential depression (continuum lowering) models. [Crilly et al. HEDP 48, 101053 (2023)]

Breakout Box: Integrated simulation design environment using machine learning and probabilistic artificial intelligence

The integrated simulation design environment enables the use of large simulation ensembles and intelligent exploration of models using Machine Learned (ML) and probabilistic Artificial Intelligence (AI) surrogates. The environment is primarily used with hydrocodes for target design but is agnostic to the simulation code used. Target geometries are parameterized using a programmatic computer-aided design package (<https://doi.org/10.5281/zenodo.3955118>). Epistemic uncertainties in physics closures and aleatoric experimental uncertainties are also parameterised. Probes and synthetic diagnostics are added to post-process quantities of interest automatically. A lightweight simulation marshalling system allows massively concurrent simulations scans. A variety of back ends are available to allow distribution on high performance and heterogeneous compute resources.

Many scan types can be used to launch ensembles of simulations. Optimised space-filling experimental designs allow for efficient initial parameter exploration and model analysis. Gaussian Process, potentially with deep learned kernels, and Neural Network surrogates are fit to ensemble data. Sensitivity analysis using these surrogates allows identification of key parameters of interest and regions of parameter space for further exploration using active learning. Bayesian optimisation and evolution strategies enable final optimisation of designs.

Finally, the simulation-experiment loop is closed using simulation-based Bayesian inference to update uncertainties in the model. The environment's holistic set of tools and workflows enables rapid iteration and update of IFE target and other simulation models.

[Statistica Sinica 18(2008); 171-186, "Gaussian processes for machine learning", MIT press (2006); "Proceedings of the 19th International Conference on Artificial Intelligence and Statistics", PMLR 51:370-378 (2016); "Deep Learning", MIT press (2016); "Proceedings of the 26th International Conference on Neural Information Processing Systems 2" (2012); 2951 – 2959, "Evolutionary Computation 9" (2001); Kyle Cranmer et al. 159 - 195, Proc Natl Acad Sci. (2020), 117(48):30055-30062].

Computational modelling and simulation are key tools in the development of IFE targets and systems. Hydrocodes are computationally expensive, with run-times for a single simulation measured in hours to days. The models feature high-dimensional parameter spaces and many model uncertainties. To mitigate this, machine-learned and probabilistic AI surrogates are trained on simulated data. Surrogates support rapid optimisation and uncertainty quantification to accelerate the design process. Simulated/surrogated outputs from final designs are compared to experimental data. Surrogate accelerated inference is used to refine estimates for parameters and uncertainties in the model, ensuring predictive capability in future experiments.

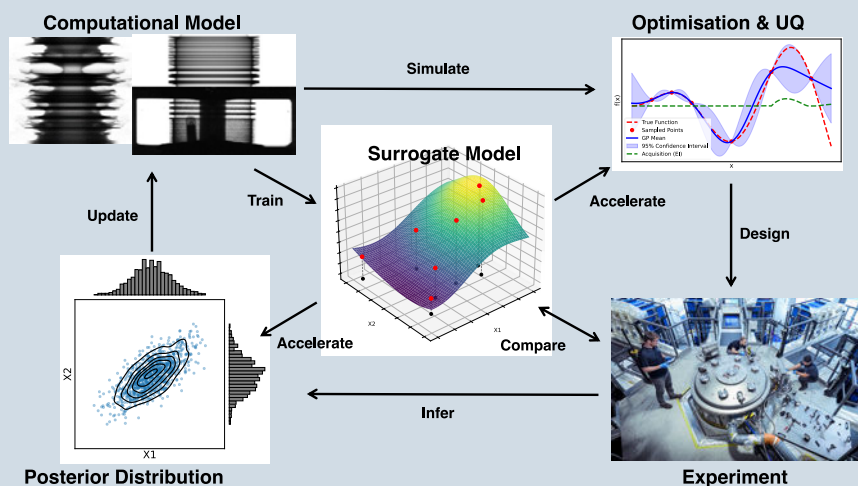


Figure 38: Conceptual diagram illustrating the role of ML/AI surrogates in accelerating simulation-based design and updating models with experimental data.

Breakout Box: High Performance Computing

Many of the physics-based simulations we perform require significant compute resources. We make use of highly parallel codes, where the work is distributed to multiple CPU cores across multiple nodes, to ensure that solutions are found within workable timescales. The workload can vary between running many small simulations in parallel, to explore various parameters which are then processed by a Machine Learning AI system, or running fewer large simulations for looking at a specific design in detail.

To enable High Performance Computing (HPC) at this scale, a 'cluster' of servers (nodes) with a fast interconnect is required. At First Light we run and maintain our own HPC resource to ensure constant availability for the computations that underpin our approach. Our main production cluster has 174 nodes, each with 48 cores (across 2 processors) and 256 GiB of RAM. We also have 600 TiB of fast SSD storage for processing and storage of results. Last year, simulations consumed around 70 million hours of compute time.




Figure 39: FLF HPC “Deep Thought”. Front of HPC racks (left) and rear of the racks (right).

PLATFORM DEVELOPMENT AND VALIDATION

To further validate and refine the underlying principles of FLF's target concepts, a multi-year research and development campaign is in progress. This work was closely integrated with computational design workflows and utilised FLF's in-house experimental facilities, as well as larger external facilities through growing collaborations and partnerships based on mutually beneficial research [see Breakout: Collaboration with European Synchrotron Radiation Facility (ESRF) and Oxford University].

- One of the key outcomes was the demonstration of shock pressures exceeding 2.5 TPa in quartz experiments conducted on the BFG facility. This result supported both the feasibility of target-based amplification techniques and the predictive capabilities of the accompanying simulation tools, (see Figure 41).
- The work substantially lowered the barrier of access to multi-terapascal pressures by extending the capabilities of two-stage light-gas guns,

allowing FLF to study HED states of matter in-house.

- Building on these findings, experiments were designed for external platforms to further evaluate key target function under more extreme conditions or with more advanced diagnostic capabilities.
- In parallel, the machine learning-based optimisation methods developed for target design were adapted to improve aspects of shock planarity and pulse shaping for Equation-of-State (EoS) measurements, facilitating broader applications in dynamic material properties (DMP) research.
- The formation of a collaboration with a team at Sandia National Laboratories led to a successful proposal to the Z Fundamental Science (ZFS) program, enabling testing of target concepts on the Z machine for DMP research. This collaboration provided validation of target performance at conditions not attainable in-house, further constraining the simulation models, (see Figure 42).

Breakout Box: Collaboration with European Synchrotron Radiation Facility (ESRF) and Oxford University

FLF invests heavily in academic collaboration, primarily through PhD and post-doctorate funding at universities such as Oxford, Imperial, York and Warwick. These collaborations enable access to expertise and equipment acting as a multiplier on FLF's capability.

One case study of this can be seen where FLF and Oxford University have collaboratively carried

out several experiments at ESRF. The synchrotron provides multiple short pulse X-rays which allow high resolution X-ray 'movies' to be generated of target collapse. These allow direct comparison to simulation for code validation.

The shock wave overcomes the material strength and causes the plastic to behave like a fluid, forming a 'jet' in the cavity. The experimental X-ray ultrafast image data taken on the ESRF is shown in the top half of each frame and the simulation for comparison is shown in the bottom of each frame. This data was all captured on a single shot.





Figure 40: Time series of a block of plastic with a hollow sphere being hit by a shock wave travelling from left to right at 4 km/s.

Figure 41: Achievable pressure in planar quartz samples using different drivers as a function of year.

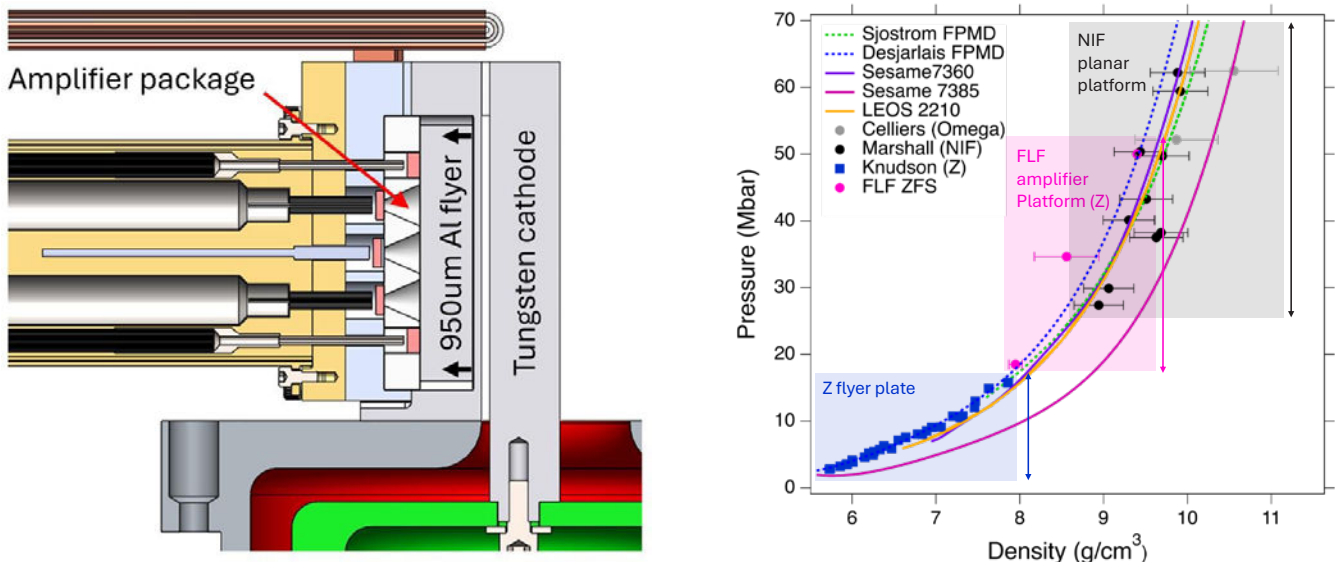


Two stage light gas guns have been used since the middle of the 20th century, but peak pressures are limited to ~ 200 GPa. Underground nuclear tests were also used to study materials at much higher pressures (~2000 GPa) between the 1960s and 1990s, before the Comprehensive Nuclear Test Ban Treaty (CBNT) came into effect. The CBNT led to the development of pulsed power drivers such as the Z machine and high-power laser drivers such as Omega and the NIF in the 2000s and 2010s.

Using hypervelocity flyer plates, Z has reached 1600 GPa in planar quartz samples, while the laser platforms have reached considerably higher pressures beyond 6000 GPa. FLF's Pressure Amplifiers have rapidly increased the accessibility of HEDP conditions over the course of four years (2021-2025).

FLF has increased the pressures achievable on two-stage light gas gun by over 12x to achieve pressures of 2500 GPa and increased the pressures achievable on Z by over 2x to 3700 GPa to enable cross platform comparison with Omega and NIF.

Figure 42: FLF Z shot (Z4110) hardware and results from a collaboration with Sandia National Laboratories via the Z Fundamental Science Program.



An aluminum flyer plate is accelerated to 45 km/s by the magnetic pressure from the 20 MA Z Machine at SNL. The flyer impacts a set of FLF pressure amplifiers. Output pressure, uniformity, and temporal duration is monitored with an array of shock velocity probes. Output pressure is 37 Mbar, allowing access to regimes previously only accessible at the National Ignition Facility at Lawrence Livermore National Laboratory.



GLOSSARY



2SLGG	Two-Stage Light Gas Gun
AWE	Atomic Weapons Establishment
B2	FLF's multi-physics fixed-grid Eulerian code
BFG	Big Friendly Gun, FLF's two-stage gas gun
CLF	Central Laser Facility
CPU	Central Processing Unit
CR	Convergence Ratio
D2	Deuterium-Deuterium
DMP	Dynamic Material Properties
DT	Deuterium-Tritium
ECI	Electrochoric Instability
EM	Electromagnetic
EoS	Equation of State
EPR	European Pressurised Reactor
ESRF	European Synchrotron Radiation Facility
ETI	Electrothermal Instability
eV	Electron-Volt, equivalent to $1.602\text{e-}19$ J or $\sim 11,600$ K
FEOS	Framework for equation of state and classical density functional theory. See Faik et al., Comput. Phys. Commun. 227, 117 (2018)
FI	Fast Ignition, a fusion scheme where separate laser pulses are used to compress and ignite the fuel
FLARE	Fusion via Low-Power Assembly and Rapid Excitation
FLF	First Light Fusion
HED	High-Energy-Density
HEDP	High-Energy-Density Physics
HFI	Hybrid Fast Ignition
HPC	High-Performance Computing
ICF	Inertial Confinement Fusion
IFE	Inertial Fusion Energy
IGNIS	FLF's next-generation multi-physics code, with adaptive mesh refinement
LANL	Los Alamos National Laboratory
LCOE	Levelised Cost of Electricity
LLNL	Lawrence Livermore National Laboratory
M3	Machine 3, FLF's pulsed power facility (the largest in Europe)
M4	Machine 4, FLF's conceptual pulsed power facility for demonstrating ignition conditions
MagLIF	Magnetised Liner Inertial Fusion
MHD	Magnetohydrodynamic
MITL	Magnetically Insulated Transmission Lines
ML	Machine Learning
MRT	Magneto-Rayleigh-Taylor Instability
MRTL	Monolithic Radial Transmission Line
MT	Reaction Type code
NIF	National Ignition Facility
PIC	Particle-In-Cell
PWR	Pressurised Water Reactor
QEOS	Quotidian Equation of State
RAM	Random Access Memory

RKM	Reduced Kinetic Model
RM	Richtmyer-Meshkov instability
RT	Rayleigh-Taylor instability
RTE	Radiation Transport Equation
RTL	Recyclable Transmission Line
SNB	Schurtz-Nicolai-Busquet model. See Schurtz et al., Phys. Plasmas 7, 4238 (2000)
SNL	Sandia National Laboratories
SpK	Atomic and microphysics code. See Crilly et al. HEDP,101053 (2023), Fraser et al., HEDP 53, 101136 (2024)
SSD	Solid State Drive
STP	Standard Temperature and Pressure
TBR	Tritium Breeding Ratio
TL	Transmission Line
TRL	Technology Readiness Level
VISAR	Velocity Interferometer System for Any Reflector (velocity measurement system)
VoF	Volume of Fluid
ZIFE	Z-Inertial Fusion Energy
Z-pinch	Plasma confinement method where a strong z-flowing current generates an azimuthal magnetic field and radial implosion
UKAEA	United Kingdom Atomic Energy Authority

ABBREVIATIONS FOR SCIENTIFIC JOURNALS CITED IN TEXT

AIP Conference Proceedings	American Institute of Physics Conference Proceedings
Astrophys. J	The Astrophysical Journal
Comput. Phys. Commun.	Computer Physics Communications
HEDP	High Energy Density Physics
ICOPS	IEEE International Conference on Plasma Science
IEEE	Institute of Electrical and Electronics Engineers
IEEE Trans. Plasma Sci.	IEEE Transactions on Plasma Science
IEEE TDEI	IEEE Transactions on Dielectrics and Electrical Insulation
Jpn. J. Appl. Phys.	Japanese Journal of Applied Physics
J. Appl. Phys.	Journal of Applied Physics
J. Phys.	Journal of Physics
J. Phys.: Conf. Ser.	Journal of Physics: Conference Series
Phys. Fluids	Physics of Fluids
Phys. Plasmas	Physics of Plasmas
Phys. Rev. Lett	Physical Review Letters
Phys. Rev. Accel. Beams	Physical Review Accelerators and Beams
Proc Natl Acad Sci.	Proceedings of the National Academy of Sciences of the United States of America
Rev. Mod. Phys.	Review of Modern Physics

Copyright Statement:

The concepts described in this paper are proprietary to First Light Fusion Limited and may be subject to registered or pending intellectual property rights. No part of this publication may be reproduced without the prior written permission of First Light Fusion Limited.

Disclaimer:

This paper includes statements about future plans, ambitions and expectations. These forward-looking statements reflect our current thinking and intentions, but they are not guarantees of future outcomes. They are based on information available at the time of writing and are subject to risks, uncertainties, and changes in circumstances that may be beyond our control. Actual results or developments may differ from those expressed or implied in these statements. Nothing in this document should be taken as a promise or commitment to take any specific action or deliver any particular outcome. We do not undertake to update or revise any forward-looking statements, except as may be required by law.



CONTACT

Unit 10 Oxford Pioneer Park,
Mead Road, Yarnton,
Oxfordshire, OX5 1QU

T: +44 (0)1865 807670
E: enquiries@firstlightfusion.com
W: firstlightfusion.com