# Muon Catalyzed Fusion

By: Himani A. Mishra, Thomas Rimer

AEP 4840 Final Project Report

Cornell University

Spring 2026

Primary Supervisor: Professor David Hammer

## Introduction (Himani and Thomas)

When two, typically low Z-number, atomic nuclei collide with high enough kinetic energy, their nuclei fuse, forming a new higher-Z atom and releasing substantial energy in the process. This phenomenon is called “fusion,” and promises a safe and abundant source of zero-carbon renewable energy. However, fusion energy faces significant engineering and physics challenges standing between the current R&D setups and commercial viability. Most of these challenges are associated with maintaining the extremely high temperatures required to sustain fusion. However, there is an alternative fusion architecture which side steps these difficult plasma conditions; “cold fusion,” also known as “muon catalyzed fusion” (MCF) uses muons to catalyze fusion reactions at or near room temperature, substantially reducing reactor complexity, and serving as a potential path forward for commercial fusion energy.

Traditional fusion architectures (e.g. tokamak, stellerator, z-pinch, and ICF) rely on extremely high temperatures (hundreds of millions of K) to overcome the coulombic barrier between particles. These high temperatures require substantial energy to form, aggressively degrade reactor walls, and require advanced systems to maintain which are beyond current technological capabilities. For reference, deuterium-tritium (DT) fusion requires the lowest plasma temperature of any fuel combinations and yet efficient fusion does not occur until 100-150 million K. For “cleaner” reactions, like deuterium-deuterium (DD) fusion, upwards of 400 million K is required for commercially viable energy production. The containment of the plasma necessitates extremely advanced supercooled, superconducting magnet systems which adds tens to hundreds of millions of dollars to the reactors total cost and still continues to pose substantial engineering challenges.

Muon catalyzed fusion, however, promises to avoid all these traditional fusion obstacles. Muons are negatively-charged subatomic particles, 207 times the mass of an electron. When injected into a near-room-temperature fusion fuel, a muon will knock out an electron orbiting a hydrogen nuclei. This forms a “muonic atom,” whose bonding distance is correspondingly 207 times smaller. **This is because the Bohr radius is inversely proportion to mass, causing the muon to sit 207 times closer to the nucleus than an electron.** It is reduced so dramatically that, when a muonic hydrogen atom bonds with a normal hydrogen atom, the nuclei are sufficiently close to fuse. A  muon catalyzed fusion reaction still yields the same amount of energy as traditional kinetic fusion. After a muon catalyzed fusion reaction occurs, the muon is almost always ejected back out into the reaction vessel, allowing a single muon to catalyze hundreds of fusion reactions. Thus muon catalyzed fusion allows for the same energy fusion reactions while avoiding the incredibly challenging technological hurdles associated with high temperature plasma generation and containment.

However, state of the art muon sources, based on proton particle accelerators, consume significantly more energy creating muons than the muons themselves generate via fusion. While significant effort has been invested in both increasing the amount of fusion energy released per muon and lowering the amount of energy required to generate muons, net positive muon catalyzed fusion is still far from commercial viability.

We investigate the dynamics and design of a muon catalyzed fusion reactor in two parts. First, we explore how muon catalyzed fusion reaction rates are affected by various reactor conditions including initial muon energy,  temperature, pressure, and T:D ratio. All findings come from our own Geant4 (a Monte-Carlo based CERN simulation software for particle physics) simulations of a diamond anvil muon catalyzed fusion reactor, which show strong

agreement with established MCF literature. Second, we propose a design for a muon source which could be used in a commercial MCF reactor, based on the particle accelerator at the Paul Scherrer Institute in Switzerland. This design was CAD'ed from scratch in Fusion360, and all components are presented in 3D renders.

While this proposed muon catalyzed fusion reactor and vessel design will itself not be capable of achieving net-positive fusion, it will contain all necessary elements of a reactor that could one day bring MCF to the commercial stage. Therefore, as advancement in particle accelerators, muon generation, and muon reaction dynamics occur—in fusion and non-fusion contexts—our design could one day be adapted to a net-positive muon catalyzed fusion reactor.

## Section I: Muon Catalyst Fusion Reaction Dynamics with Geant4 Simulations (Himani)

### Basics of MCF (Himani)

Muons are an elementary particle similar to an electron, but with around 207 times the mass. They belong to a class of particles known as leptons, which are fundamental subatomic particles that do not partake in the strong interaction, the fundamental force that holds quarks together to form protons and neutrons and binds them to form nuclei. Muons are highly unstable, and decay after a mean lifetime of around 2.2 μs, into an electron, antineutrino, and muon antineutrino. They can penetrate deep into matter during their lifetime before decay.

Muonic atoms form when a muon is slowed in matter and captured by an atomic nucleus. Because a muon is 207 times heavier than an electron, its radius is 207 times smaller, which brings it much closer to the nucleus. Because of this close proximity, the binding energies are also much higher than **with** electrons. Thus, the energy barrier for a muon to bind two nuclei and form a muonic molecule is **narrower**.

The process of muon catalyzed fusion starts with a beam of negative muons entering a deuterium (D) and tritium (T) **target, which could be high-pressure gas or liquid.** This requires the generation of negative decay muons, **from $\pi^{-1}$ decay**, which must be slowed down the appropriate energy depending on the specific reactor cell, but typically on the order of a few MeV; Section II discusses the challenges with producing such muons. The muon is then captured by either a D or T nucleus, forming a muonic atom. If the muon binds to a D atom, it quickly transfers to a T nucleus because the binding energy is higher. Then, the μt atom collides with a $D_2$ molecule, forming a muonic molecule: dtμ. The nuclei in the dtμ then fuse and we obtain our regular DT fusion reaction: $d + t \rightarrow \alpha + \mu + n + 17.6$ MeV, with the neutron getting around 14.1 MeV of the energy. The muon then goes on to catalyze more reactions. However, there is also a chance that the muon might stick to the alpha particle, a process known as alpha sticking. In that case, the reaction is: $d + t \rightarrow \alpha\mu + n + 17.6$ MeV. There is a chance the muon can be shaken off after sticking to the alpha particle, a process known as reactivation.

**Introduction to Geant4**

Geant4 is a software package developed by CERN that simulates the interactions and energy transfer of particles as they travel through matter using Monte Carlo simulations. It is useful for a variety of applications, from high energy physics to medicine and space studies. The system allows us to define different particles, including leptons (**e.g.** electrons and muons), hadrons (**e.g.** protons, neutrons, and pions), bosons (**e.g.** photons), and ions; materials; detectors; and system geometry. We can define parameters, such as the concentration of reactants, materials used in a detector, and ambient conditions, including pressure and temperature of a reaction vessel. Then, we track primary and secondary interactions of particles (elastic and inelastic interactions and electromagnetic and hadronic interactions). The particle trajectory is broken up into discrete steps and at each step, the software uses established probability distributions to perform a prediction for the next step. Users specify Physics Lists, pre-packaged model configurations that have been theoretically and experimentally tested. In the present project, we use the QGSP_BIC physics list, which has precise cross-section data to track neutrons precisely as they scatter in the gas cell, essentially for quantifying neutron flux [1]. **This physics list is the industry standard for precise neutron and hadronic interaction modeling in high-density environments.** We also use the G4EMStandardPhysics model to accurately model the electromagnetic interactions, especially the slowing of muons from higher energies to lower energies upon entering the diamond anvil cell. Finally, we implement the G4MuonicAtomDecay and G4MuonMinusCapture modules, which allow us to model muonic atom decay and capture processes accurately.

**Reaction Rates and Probabilities**

The efficiency of the MCF cycle is determined by a series of reaction steps and pathways, some of which are competing. Each step has a specific reaction rate, and in the case of several reaction pathways, a certain probability of being taken. However, for the reaction to be viable, the sum of the reaction rates within the cycle must be lower than the muon decay rate, **$0.45 * 10^6$ / s [2]**. This is the probability per second that a muon will decay into electrons and neutrinos. Before fusion can occur, we must undergo the following steps. First, the muon must be captured by either a D or T nucleus to form a muonic atom. This

happens at a rate of $1.2 * 10^{-12}$ $cm^3/s$. Next, if the muon is captured by a D atom, it transfers to a T atom, as it is more energetically favorable for the muon to be tightly bound to the heavier T atom. The transfer from D to T is known as isotope exchange and has a rate of $0.5 * 10^{-14}$ $cm^3/s$. Next, the μt atom **combines with a D atom** to form a **muonic molecule containing D, T, and the muon**. This is the rate limiting step of our reaction and is dependent on the temperature and pressure of our reaction vessel, but averages around $0.25 * 10^{-14}$ $cm^3/s$ [2].

There are also pathways where we may lose the muon. First, when forming the muonic molecule, muon may be captured by the nucleus at a rate of $3.0 * 10^{-14}$ /ns. There is also a pathway where the muon sticks to the alpha particle around 0.7% of the time. If the muon sticks, it cannot catalyze further reactions unless it is reactivated, which has a probability of around 30% across research studies [3].

**Experimental Setup**

The experimental setup used in this experiment is inspired from the setup documented in the 2024 International Conference on High Energy Physics presentation from Acceleron researchers and collaborators at UC Davis and Fermilab [4], a deliberate choice to verify simulations. Additionally, we spoke to the Fermilab collaborator and head of Mu2e (muon-to-electron conversion), Dr. Kevin Lynch for advice on impelementing the simulation geometry and design **(K. Lynch, personal communication, March 9, 2026).**

The diamond anvil cell is composed of a front and back diamond, each 2.5 mm in diameter and 0.5 mm thick. We use a high-pressure DT gas cell, **with amorphous fluid of D and T ions at** 3 liquid hydrogen density (LHD) as the fusion fuel. 5000 muons are introduced into the system, each with an initial energy of 4.1 MeV. The muons are slowed down to 0.7 MeV by the time they reach the diamond anvil cell entry point. **It takes around 3000 MeV to deliver this μ, which means we would need around 160 reactions per muon just to get the initial input energy back. We could consider other fusion reaction, like $n + Li^6 \rightarrow He^4 + H^3 + 4.78$ MeV, where we would then only need 130 reaction per muon to recover the initial energy.**

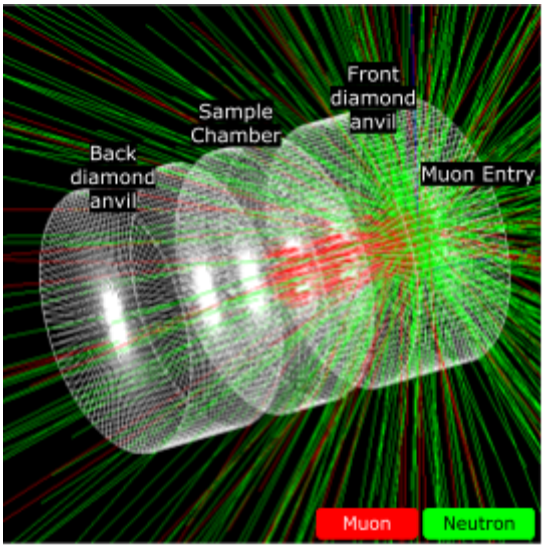


**Figure 1.** Diamond Anvil Cell Setup in Geant4 Simulation: 0.5 mm front diamond anvil, 0.25 mm sample chamber housing DT target at 3 Liquid Hydrogen Density, 0.5 mm back diamond anvil.

Considering the reaction rates and avenues of loss in our reaction from the previous section, we are interested in studying muonic molecule formation rate, number of reactions per muon, total energy generated, and neutron flux. We also consider the energy a muon must have to reach the target. Thus we run the following four experiments:

1. Effect of initial muon energy (0.3-0.8 MeV, increment by 0.1 MeV) on muon stopping time and location in diamond anvil cell.
2. Effect of temperature (7, 10, 100-1500 K - incrementing by 100 K) on number of fusion reactions per stopped muon, total number of reactions, total energy generated, muonic molecule formation rate, and total neutron flux.
3. Effect of pressure (1-10 GPa, incrementing by 1 GPa) on the same parameters as for temperature.
4. Effect of T:D ratio (0.1-0.9, incrementing by 0.1) on the same parameters as for temperature.

**Initial Muon Energy Experiment**

The first experiment is to test the relationship between initial muon energy and the resulting stopping depth and stopping time in a DT gas cell.

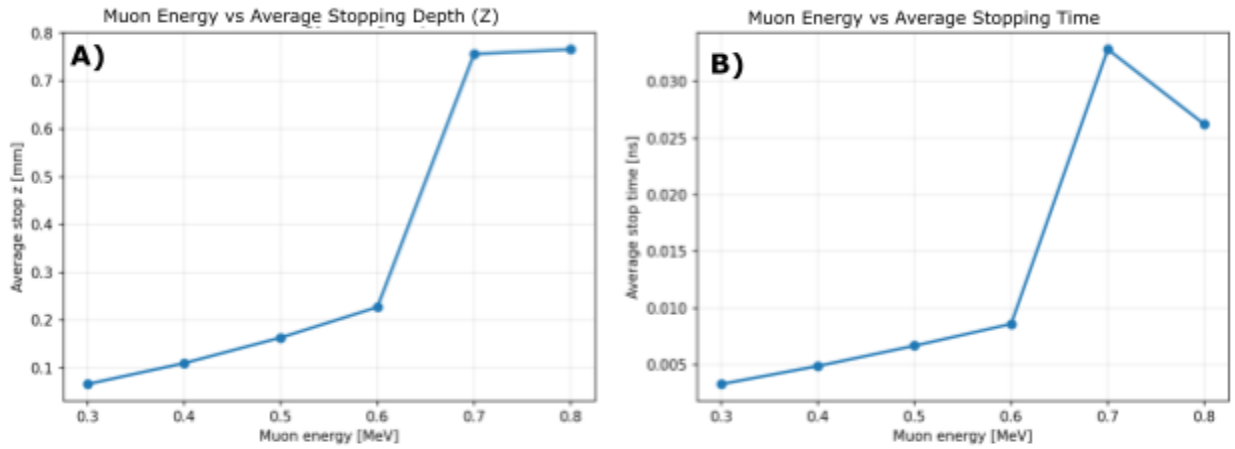


**Figure 2.** (A) Initial muon energy versus stopping depth in diamond anvil cell, where first 0.5 mm is front diamond anvil, followed by 0.25 mm sample chamber, and finall 0.5 mm for the back diamond anvil. (B) Initial muon energy versus average stopping time.

As initial muon energy increases, from 0.3 to 0.8 MeV, the muons travel deeper into the diamond anvil cell. The first 0.5 mm are the front diamond anvil (2.5 mm diameter), followed by 0.25 mm of DT gas cell, and then a 0.5 mm back diamond anvil (2.5 mm diameter). From Figure 2a, we see that a minimum energy of 0.7 MeV is needed for the muon to penetrate the diamond anvil and actually reach the DT target. However, if the energy is even slightly higher, 0.8 MeV, the muon has too much energy and rushes past the target and hits the back diamond slab. Thus, a future reactor design would have to have precise sensors to track the muon stopping speed. Figure 2b shows us that the 0.7 MeV muons take around 0.03 ns to actually reach the target. This is important to know, as the muon decay time is 2.2 μs, so the time to reach the cell cannot exceed that. This informs the rate at which we can inject muons. At 0.03 ns, we can inject 33.3 billion muons/s, but this is limited by the target saturation and the fact that no accelerator facility can generate muons at this high rate. As discussed in section II, the highest rate achieved currently is $10^{10}$ **/ s** by the Research for Nuclear Physics at Osaka University [5], but these are **positively charged muons**, while we require negatively charged muons. No accelerator facility can generate muons at the high rate needed for the MCF setup. **Next, we look into how the temperature affects the fusion yield for MCF.**

**Temperature Experiment**

Next, we test the effect of temperature on the number of fusion reactions per stopped muon, the neutron

flux, the number of reactions, total energy generation, and muonic molecule formation rate, holding at a constant mid-range pressure of 5 GPa.

From Figure 3, as the temperature increases, the number of reactions per stopped muon, the number of fusions per stopped muon, the amount of energy produced, the neutron flux, and muonic molecule formation rate all increase. Figure 3c shows that the greatest increase in number of reactions occurs from 0 to 200 K, **followed by** considerable increase from 200 K to 800K, after which increasing the temperature does not provide considerable increase in the energy produced. Increasing temperature increases the number of reactions, as it provides greater kinetic energy to the muons and thus facilitates more reactions. The muonic molecule formation rate is highly sensitive to temperature due to the resonance required for molecular binding **between the tμ and D atom**. Below 200 K, the particles do not

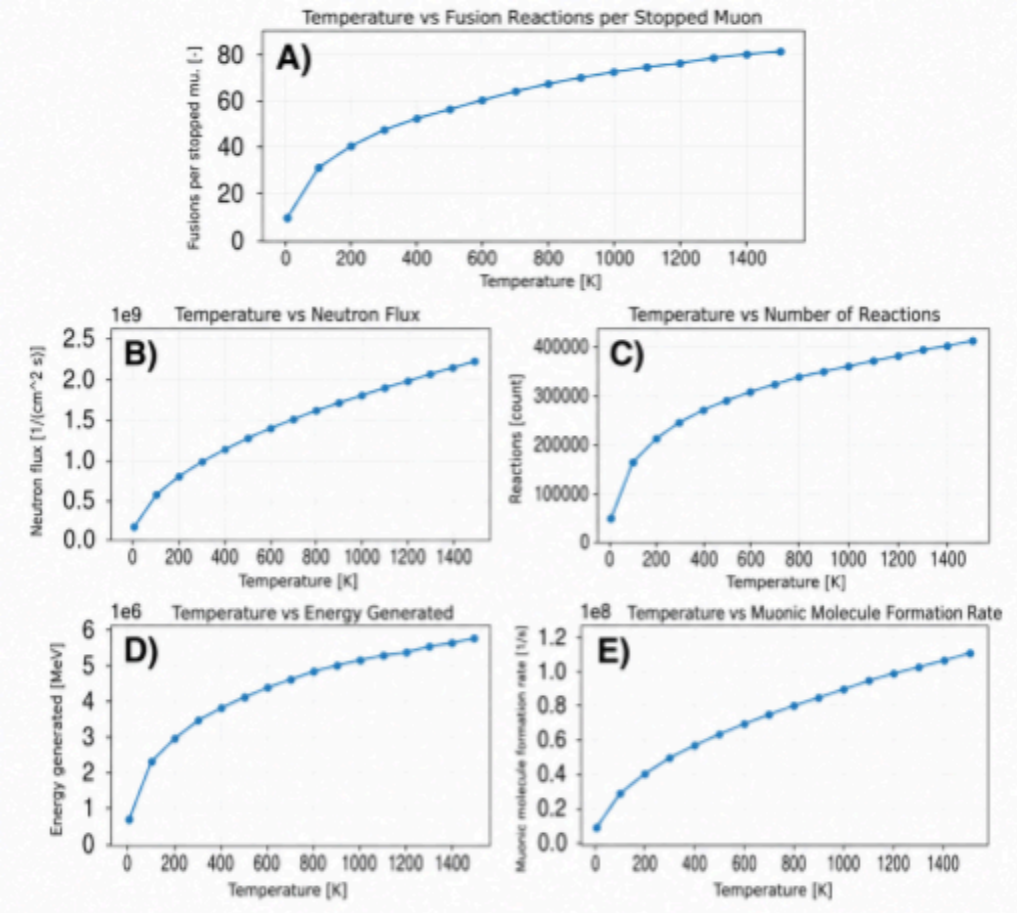


**Figure 3.** Effect of Temperature (7, 10, 100-1500K) on (A) fusions per stopped muon, (B) neutron flux, (C) number of reactions, (D) total energy generated, (E) muonic molecule formation rate

have enough kinetic energy to trigger resonant conditions for binding. The formation rate climbs sharply from 500-800 K, achieving around 1.1 * $10^8$/s. Post 1200 K, the rate increases more slowly as the thermal energy far exceeds the resonant energy.

As for the number of fusion reactions per stopped muon, under 200 K , the yield is negligible (under 10 fusion reactions). At around 800 K, we have around 375,000 total reactions for the 5000 muons. We achieve around 75 reactions per muon, which is still only 26% of the required 300 fusions per muon to achieve energy breakeven. Thus, alpha sticking and reactivation must be addressed to increase fusion yield per muon. **We now look into how the pressure affects the MCF fusion yield and how this differs from the temperature effects.**

**Pressure Experiment**

We test the effect of pressure on the number of fusion reactions per stopped muon, the neutron flux, the number of reactions, total energy generation, and muonic molecule formation rate. We maintain a constant temperature of 800 K **for the high density fluid**, the maximum temperature we saw from experiment 1 that provides significant energy yield improvement.

From Figure 4, as pressure increases, the number of fusion reactions per muon also increases, but at a considerably lower rate than temperature. Between 1 to 10 GPa, the number of reactions only increases from 65 to around 80 fusion reactions per muon. Increasing pressure also increases the number of reactions, energy generated, and muonic molecule formation reaction rate; there are more muons per unit volume **inputted**, allowing greater interaction and higher collision frequency between particles, facilitating more reactions. At 1 GPa, the system produces around 342,000 total reactions **for a total of 5000 muons**, which climbs to around 390,000 reactions at 10 GPa, a ~14% increase. The energy generated increases linearly from ~4.8 MeV to ~5.45 MeV (13.54% increase) across this pressure range as well. In contrast to the steep increase and subsequent plateau for increasing temperature, the metrics appear to scale linearly with increasing pressure.

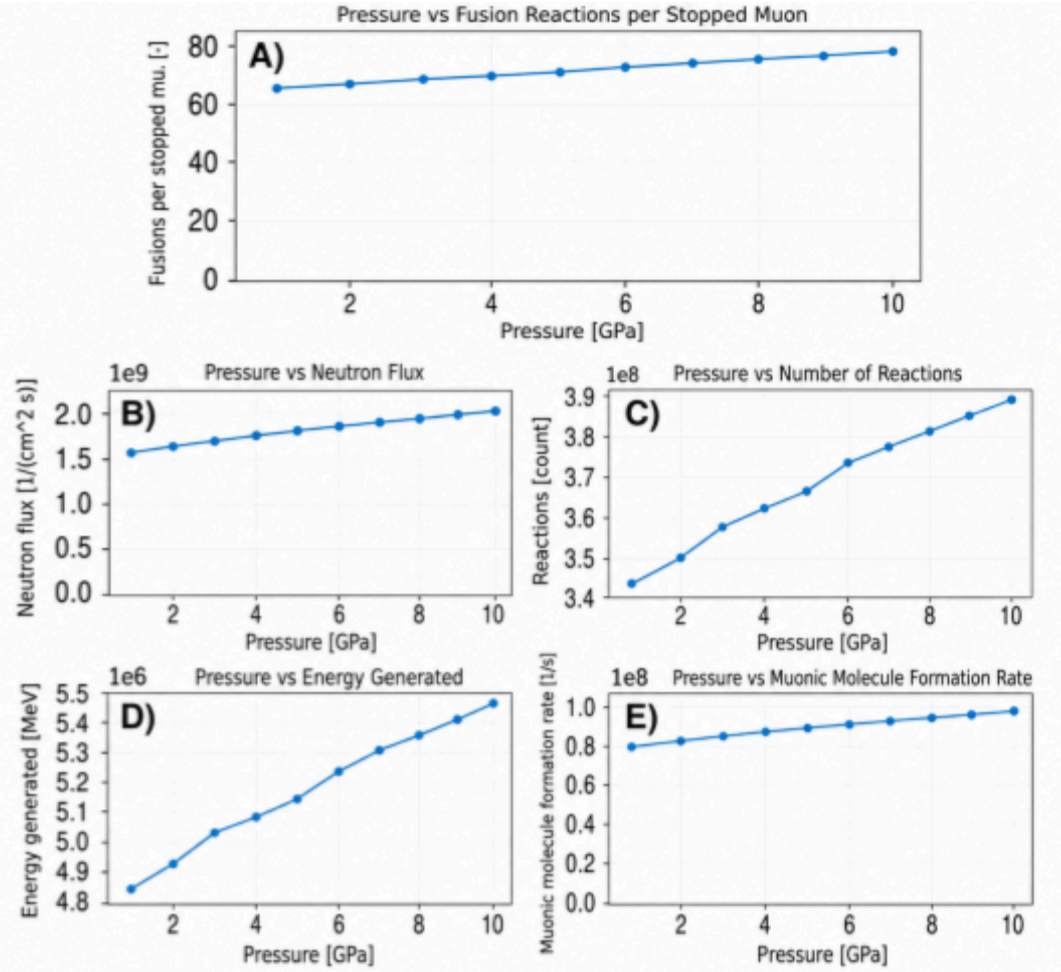


**Figure 4.** Effect of Pressure (1-10 GPa) on (A) fusions per stopped muon, (B) neutron flux, (C) number of reactions, (D) total energy generated, (E) muonic molecule formation rate

The greater the pressure, the more collisions and also the greater probability of reactivation, as the collisions are more likely to knock off the muon stuck to the alpha particle. However, we are limited to around 10 GPa, because the cracking limit of the diamond anvil is around 50 GPa. **Using this 10 GPa and 500 K, we test the effects of T:D ratio at constant temperature and pressure.**

**T:D Ratio Experiment**

Next, we test the effect of T:D ratio on the number of fusion reactions per stopped muon, the neutron flux, the number of reactions, total energy generation, and muonic molecule formation rate, held at a constant temperature of 500 K and 10 GPa, **to observe T:D ratio sensitivity in the high-growth phase of the resonance curve rather than at the plateau**. **While 800 K provides the peak yield, the T:D ratio was**

**tested at 500 K to evaluate mixture sensitivity during the formation phase.**

From Figures C and D, we see that a ratio of 0.3 produces the greatest number of reactions and energy, which matches the 0.35 T:D ratio that Acceleron Fusion is implementing in their reactor [4]. At 30% T:D, we achieve around 90 fusions per muon and around 6.3 * $10^6$ MeV produced. Beyond 50%, the performance plummets below 80 fusions per muon. Unlike the other metrics, as the tritium fraction increases, the muonic molecule formation decreases linearly. The formation rate is at its maximum (2.2 * $10^8$/s) when the T fraction is lowest. At 30%, we have already almost decreased the rate by ~31.82%. This

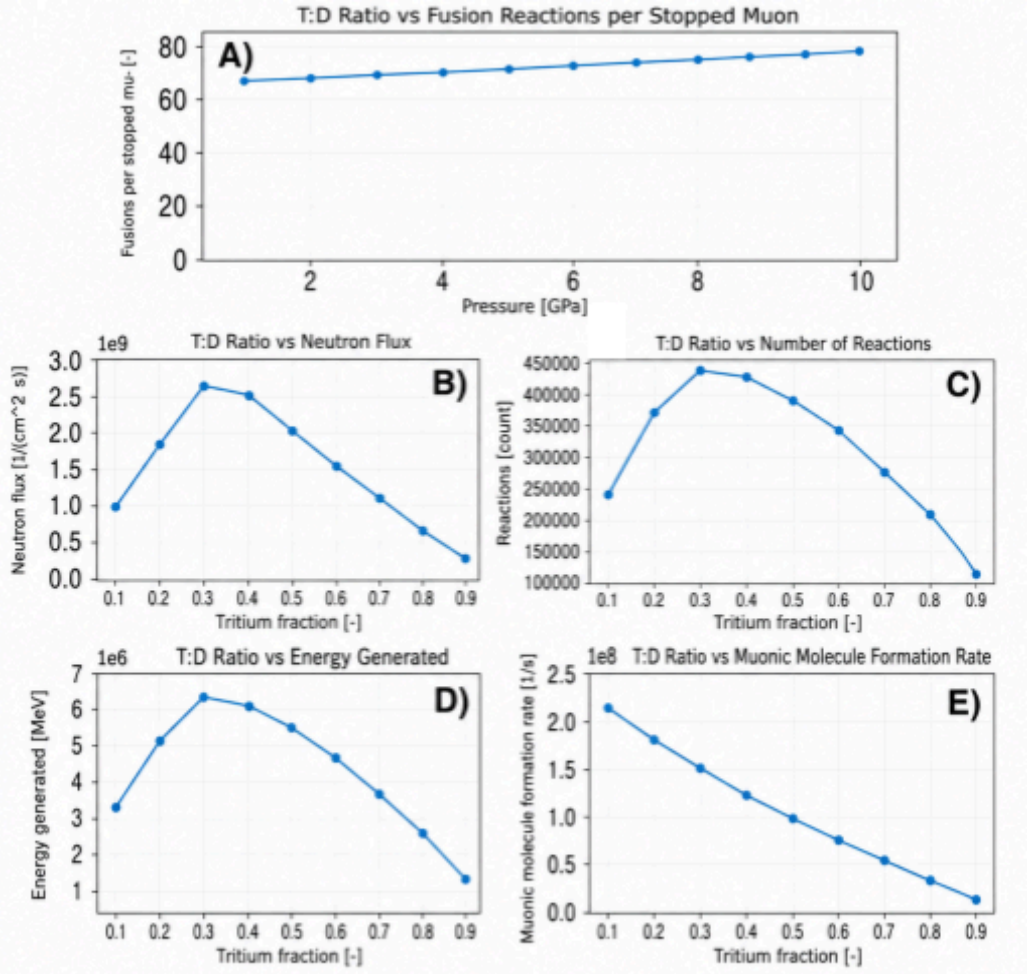


**Figure 5.** Effect of T:D Ratio (7, 10, 100-1500K) on (A) fusions per stopped muon, (B) neutron flux, (C) number of reactions, (D) total energy generated, (E) muonic molecule formation rate

trade-off is worth it however, because the probability of the molecules fusing is maximized at 30%, providing a greater net energy yield.

**Conclusion**

From the four experiments, we determine that the optimal temperature range is 500-800 K as it triggers the resonant energy required for muonic molecule formation. The ideal pressure is 10 GPa, as we want the pressure to be as high as possible to increase collision frequency and muon reactivation, but also prevent the diamond anvil from cracking. The optimal T:D ratio is 30%, as it maximizes the total reactions and energy generation, while also balancing the muonic molecule fusion rate. For the present system, the muons must enter the diamond anvil cell with at least 0.7 MeV to ensure stopping in the target and not the shielding. To produce more energy than the reactor consumes, we must reach around 300 reactions per stopped muon. With the current optimal settings of 10 GPa, 800 K, 30% T:D, we obtain around 80-90 fusions per muon, only 26-30% of the required goal. Thus, we determine that the physical parameters alone cannot bridge this gap, but solving the alpha sticking and reactivation issues can. Designing a system that can withstand higher pressures can increase reactivation, as the number of reactions scales linearly with the pressure. Additionally, decreasing the time it takes for the muon to travel to the target will allow the muon to complete more fusion cycles within its 2.2 μs lifespan.

## References

**Section I**

[1] Geant4 Collaboration. (n.d.). *QGSP_BIC*. Geant4 Physics List Guide. from https://geant4.web.cern.ch/documentation/dev/plg_html/PhysicsListGuide/reference_PL/QGSP_BIC.html

[2] A. A. Harms, K. F. Schoepf, D. R. Kingdon, and G. H. Miley, *Principles of Fusion Energy: An Introduction to Fusion Energy for Students of Science and Engineering*, Chapter 12 (World Scientific, Singapore).

[3] Stodden, C. D., Monkhorst, H. J., Szalewicz, K., & Winter, T. G. (1990). *Muon reactivation in muon-catalyzed d–t fusion from accurate p–$He^+$ stripping and excitation cross sections*. Physical Review A, 41(3), 1281–1292. https://doi.org/10.1103/PhysRevA.41.1281

[4] CERN Indico Presentation Knaian, A. N., Tripathy, S., & Lynch, K. R. (2024, July 19). *GEANT4 models for muonic atom processes, and proposed simulation package* [Conference presentation]. ICHEP 2024, Prague, Czech Republic. https://indico.cern.ch/event/1291157/contributions/5888470/attachments/2900457/5086523/GEANT4%20models%20for%20muonic%20atom%20processes,%20and%20proposed%20simulation%20package_2.pdf

[5] Seidel, M., Bi, Y., Humbel, M., Sigg, H. C., & Zhang, Y. (2017). Production of a 1.3 MW proton beam at PSI. *Physical Review Accelerators and Beams*, *20*(3), Article 030101. https://doi.org/10.1103/PhysRevAccelBeams.20.030101

## Section II: Design of Muon Generator, Collector, and Reaction Vessel (Thomas)

### Introduction to Muon Generation for MCF

A grid-scale muon catalyzed fusion reactor requires roughly $10^{19}$ muons per second to produce 1 GW of fusion power (not accounting for the energy needed to power the reactor, which would further increase the required muons by many more orders of magnitude). Thus, a reliable, large-scale source of moderate energy, negative muons is required for muon catalyzed fusion to ever become commercially viable. Presently, the most promising physical mechanism by which to generate large numbers of muons is the impingement of high energy protons—generated by a particle accelerator—on a solid target. These proton-nucelus collisions produce pions, which in turn decay into muons that can then be collected, focused, and fed into a fusion reactor cell to catalyze fusion reactions. Previously in this report, the reaction dynamics of a diamond anvil based DT reactor was detailed. In the following section, we present a design for a muon source I created in CAD, based on the accelerator at the Paul Scherrer Institute [6], but modified for use in an MCF reactor. It will begin by giving a system-wide overview of the design, and then detail each subsystem of the accelerator and target. Note that all designs in this paper were done entirely by us; no pre-made models or other third-party CAD designs were used.

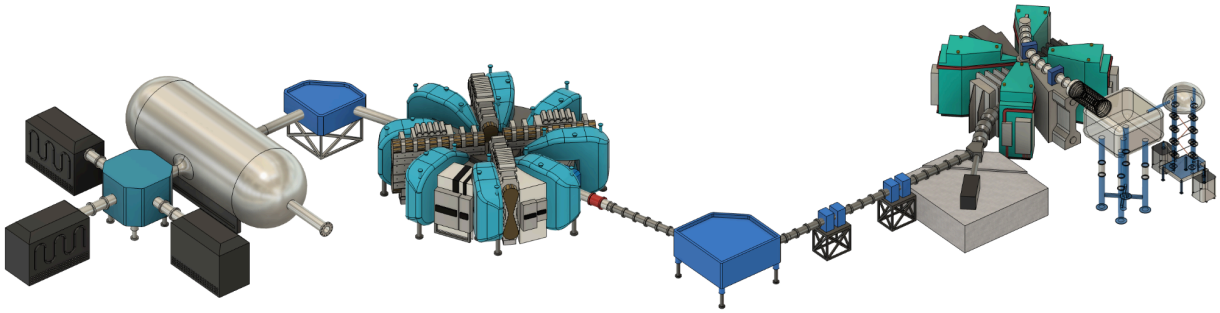

**Figure 6.** Complete CAD of the proposed muon source detailed in the following section. The cockcroft walton is on the far right. The green-teal assembly next to it is the injector. The long tube extending from the injector with blue objects along its length are the transfer beamline optics. The large, circular blue-teal assembly is the primary accelerator. Finally, the cylinder and black boxes on the far left are the target and associated beam dumps.

**Overview of the Accelerator**

The muon source and target presented in this paper are based on the accelerators at the Paul Scherrer Institute (PSI) in Würenlingen, Switzerland, the world's leading muon research facility. PSI has several sub-facilities within it that conduct a broad range of research in physical and engineering sciences, including the Swiss Light Source (SLS) (bright and tunable x-ray source for crystallography), SwissFEL (free-electron x-ray laser for observation of atomic dynamics), Swiss Spallation Neutron Source (SINQ) (neutron source for studying solid, liquid, and magnetic material properties), and the Swiss Research Infrastructure for Particle Physics (CHRISP) (using ultracold neutrons to test the standard model). However, the facility of relevance to this project is the Swiss Muon Source (SμS), powered by the High Intensity Proton Accelerator (HIPA) (which somewhat-confusingly also powers the SINQ and CHRISP facilities, but not the SLS and

SwissFEL facilities) [7]. Together, HIPA powering SμS holds the title of world's most intense continuous muon source, generating 108 muons per second for use in muonic material studies and investigations of fundamental muonic physics. HIPA and SμS therefore provide a good starting point for the design of a particle-accelerator-based muon source. However, the muons generated by SμS are all surface muons, meaning they are positively charged and cannot be used for MCF (discussed in more detail in the target section). Thus my design, specifically of the target, deviates from the original design of SμS and HIPA, though the rest of the design closely resembles that of PSI.

The proposed muon source begins with a small, microwave-based proton source inside a Cockcroft-Walton accelerator. This pre-accelerator stage not only generates the protons which will ultimately go on to collide with the target, but it uses the Cockcroft-Walton voltage source to accelerate the protons up to sufficient energy to move into the next stage; the injector. This second "injector" stage is a small, isochronous separated sector cyclotron which accelerates the low energy protons generated in the Cockcroft-Walton accelerator up to a moderate energy. These protons are then transferred via a proton beamline, focused by intermittent pairs of quadrupole magnets and steered by dipole magnets, to the final stage of acceleration: the primary accelerator. This large, isochronous separated sector cyclotron (same architecture as the injector) provides the protons with their remaining energy, bringing them to a total energy of ~590 MeV. These high energy protons are at last focused on a target, where their collisions generate pions which rapidly decay into muons. These muons are then fed to a DT reactor like the diamond anvil cell discussed in the previous section. Power from the still-energetic proton beam is recovered after the target via multiple large beam dumps. See figure 6 for an overview of the system.

**Cockcroft Walton Pre-Accelerator**

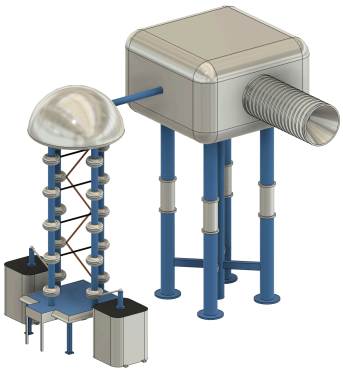

**Figure 7.** Closeup of the Cockcroft-Walton pre-accelerator.

The muon source begins in the pre-acceleration stage, where a Cockcroft-Walton generator brings protons generated in a microwave cavity up to a modest energy of 870 keV [8]. See figure 7 for a subsystem-wide render of the cockcroft walton pre-accelerator. Inside the large, silver, rounded-rectangular box is the proton source itself; radiation from a 1100 W, 2.45 GHz microwave generator (in fact the exact same power and frequency as the kind found in standard microwave kitchen appliances) is sent via waveguide to a ceramic plasma chamber filled with hydrogen gas [9]. Of the 1100 watts of wall-plug power, roughly ~350 watts makes it as microwave radiation to the plasma. This chamber, which is roughly the size of a tea cup, is constantly resupplied with ultrapure hydrogen and serves as the source of all downstream protons. Electromagnetic coils allow the magnetic field generated by permanent magnets to be tuned to exactly match the electron cyclotron resonance of the electrons in the chamber ensuring efficient coupling of microwave power to proton heating [10]. Protons liberated by this source are ejected from the tiny cavity into the larger interior of the silver metal box, with an average kinetic energy of ~60 keV [11].

This metal box is held at an extremely high (~800 kV) potential relative to ground (and the microwave cavity), which has the effect of accelerating the ejected protons up to a kinetic energy of ~870 keV. To actually generate this high potential, a Cockcroft-Walton generator (hence the name Cockcroft-Walton accelerator) is used. The generator can be seen in the foreground of figure 7. The two large boxes sitting on either side of the three-pillar tower are high voltage AC waveform generators. Using wall-plug power, they generate an AC waveform which is then fed to the tower via the insulated connectors seen on their respective tops. Bridging between the pillars of the towers are diodes and capacitors, which rectify the AC waveform, converting it from high voltage AC to extremely high voltage DC. At the top of the tower is a dome held at ~810 kV (relative to ground); this potential is transferred to the box via the blue conductor shown bridging both towers.

The net result of the pre-accelerator is a ~10 to 12 mA beam of protons (corresponding to ~10 trillion particles) accelerated to 870 keV [12]. This beam exits out of the side of the large metal box for further acceleration via the corrugated cylinder. These corrugations are insulating gaps which safely step down the high voltage of the metal box to the ground potential of the room and remainder of the muon source. Protons exit the Cockcroft-Walton preaccelerator at 4% the speed of light [12].

**Injector**

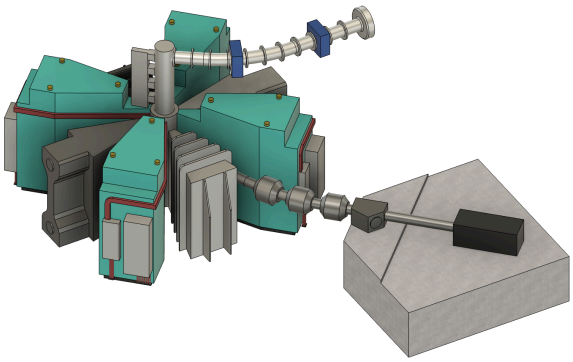

**Figure 8.** Closeup of the injector.

The injector gradually steps up the protons' energy from the 870 keV of the Cockcroft-Walton accelerator to the ~70 MeV required for protons' injection into the primary accelerator, hence the name injector. See figure 8 for a subsystem-wide render of the injector. Without the injector, the protons generated in the Cockcroft-Walton accelerator would not have sufficient energy to be fed into the primary accelerator [13], [14].

Protons from the cockcroft walton pre-accelerator enter the injector via the curved, silver pipe in the top of figure 8. Notice the blue steering magnets which redirect the beam through two 90° bends to ultimately reach the cyclotron's midplane. The injector is a separated sector cyclotron, with 4 magnets, 2 50 MHz RF resonators, and 2 secondary 150 MHz flat top resonators for beam corrections. The four magnets (easily seen as the four teal rounded objects) are equally spaced around the center of the cyclotron. Both sets of RF resonators can be seen as silver wedge-shaped apparatuses also equally spaced around the center. When protons enter the injector, the field generated by the magnets causes the protons to orbit the center of the cyclotron via the Lorentz force. Each time the protons pass through an RF resonator, they are accelerated via the electrostatic force, and gain additional kinetic energy. As protons gain energy, their orbit

increases. The nonhomogeneous magnetic field formed by the magnets ensures that despite this orbital radius increase, the protons still take the same amount of time to orbit, thus earning the cyclotron the “isochronos” title [15]. After 80 orbits, the protons have sufficient energy and are extracted via an electrostatic diverter out of the side of the cyclotron and transported via the silver tube on the lower right side of figure 8.

In standard operation, the beam is diverted from the silver exit tube via the dark grey steering magnet roughly half way along the silver tube's length. However, during startup of the accelerator, protons with insufficient energy are generated by the injector which risk damaging the primary accelerator; thus the steering magnet can be turned off and the proton beam can be allowed to impinge on the beam dump at the end of the silver exit tube, shown in black, sitting on a large concrete pad, in figure 8 [16].

**Transfer Optics**

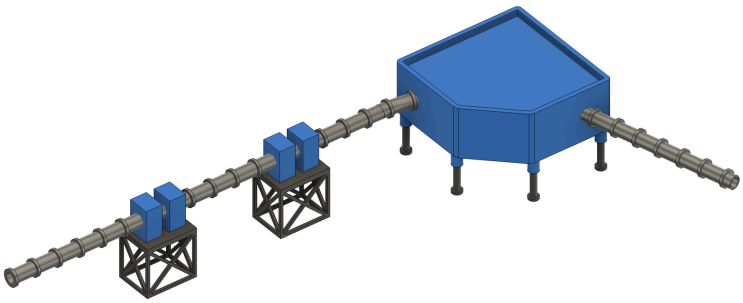

**Figure 9.** Closeup of the transfer beamline optics.

Given the injector’s role is purely to increase protons’ energy enough to allow them to enter the primary accelerator, it would seem logical that the injector feeds directly into the primary accelerator. However, imperfections in the proton beams focus and the physical layout of the

facility means that an intermediate beamline, with associated beamline optics, must be used to transfer the beam a nontrivial distance from the injector to the primary accelerator. See figure 9 for a subsystem-wide render of the example transfer beamline used in my design.

To address the issue of proton beam focus, pairs of quadrupole magnets, depicted as pairs of small blue rectangles on black support structures, compress the beam horizontally and vertically, thus refocusing it as it travels down the tube. These quadrupole magnet pairs would be placed along the tube with appropriate spacing to ensure the proton beam within the tube always remains well contained [16].

To address the issue of site layout, the large blue object is a dipole magnet used to steer the beamline around corners. In a perfect world, the injector would sit directly next to the primary accelerator allowing one to immediately feed the other. However, given the immense complexity of a cyclotron, both the injector and primary accelerator have substantial supporting hardware, not shown in the render, which must sit close to the cyclotrons. Furthermore, for safety, both cyclotrons require shielding to prevent radiation from harming people working in and around the accelerator. This shielding is achieved with thick concrete blocks, further forcing the two cyclotrons apart. Thus it is physically not possible, or at least an overwhelming challenge, to have both cyclotrons directly adjacent. Hence, a beamline capable of transferring the beam around the supporting hardware and shielding is required. It's also worth highlighting that the specific layout of a site may also mandate the use of transfer beamline optics; while a muon catalyzed fusion reactor would ideally inhabit a piece of land large enough to accommodate any desired layout of systems, this is not always the case, and buildings and other facilities may further prevent both cyclotrons from sitting adjacent to one another.

**Primary Accelerator**

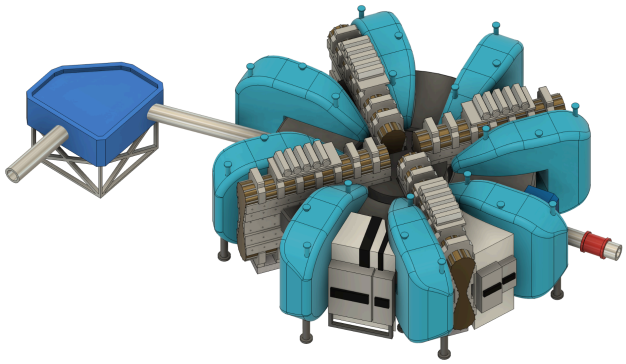

**Figure 10.** Closeup of the primary accelerator.

The primary accelerator is the final stage of proton acceleration before the beam can finally be used to generate muons. See figure 10 for a subsystem-wide render of the primary accelerator. While the primary accelerator uses the same separated sector cyclotron architecture as the injector, it's equipped with 8 sector magnets and 4 primary RF resonators (also 50 MHz). The primary accelerator is also substantially larger than the injector, with a maximum orbital radius of 4.4 m and a total exterior diameter of 15.5 m. To extract the beam, an electrostatic septum is used between proton orbitals, diverting the outermost protons to output beamline transfer optics. The final beam produced by the primary accelerator is 1.4 MW, 5-7 mm in diameter, with an output current of 2.2 mA (and the aforementioned 590 MeV proton energies: ~79% the speed of light). Another large steering dipole magnet is used after the primary accelerator to direct the protons towards the target, for the final stage of muon generation [17], [18], [19], [20].

**Target**

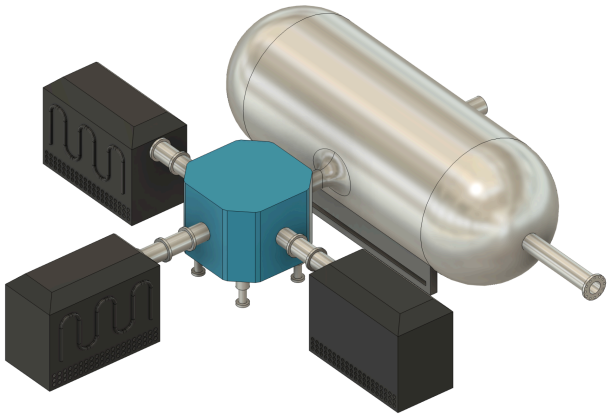

**Figure 11.** Closeup of the target

After pre-acceleration, injection, transfer, and primary acceleration, the protons finally have sufficient energy to impinge on a target and produce muons. See figure 11 for a subsystem-wide render of the target. Protons from the primary accelerator enter the large muon decay chamber and collide with a pair 2 m diameter, 5 cm thick isotropic graphite targets [21]. These targets spin at several thousand RPM to evenly distribute the beam's deposited energy over its surface and avoid local damage. A large liquid cooling system on the opposite side of the spinning wheels extracts deposited heat from the targets and returns it to the generators to be fed back into the system, decreasing losses. These targets are included in the CAD but obscured but the vacuum chamber.

When the proton beam impinges on the targets, individual protons collide with carbon nuclei, generating pions. These pions have a range of energies, some of which are low enough energy to be stopped within the target. If these stopped pions are positive, they remain in the outer orbital of the carbon nuclei, decaying after 26 ns into a positive muon which subsequently escapes into the vessel. If these stopped pions are negative, they cascade down to the nucleus of the carbon atom, shattering and preventing muon formation. Thus these monoenergetic "surface

muons" are almost exclusively positive, which cannot be used for muon catalyzed fusion. Thus surface muons cannot be used for muon catalyzed fusion. However, if the pions generated from proton impingement have sufficient energy to make it outside the target, they will decay into positive and negative "decay muons" with high energy and crucially, negative charge compatible with MCF [22], [23]. These muons would ultimately be slowed to the single-MeV scale muons via "injection windows" to match the energies described in the simulations of the previous section.

Generated negative muons are collected with large, superconducting solenoids within the target chamber and redirected to the port in the center right of the diagram. The flange of this port is exactly 8 meters from the graphite target, corresponding to the mean distance required for pions to decay into muons [24]. Thus, a DT reaction chamber specially designed for MCF, such as the diamond anvil discussed in Section I, would be mounted onto this port for the ultimate fusion reactions.

However, despite intense muon generation, the proton beam would still carry with it significant energy. Thus, to increase efficiency of the reactor, the beam would be aimed at a beam dump capable of absorbing the residual beam's energy, and feeding the power back into the system. Within the large blue vessel in the middle of figure 11 are a set of steering dipole magnets which split the beam into three smaller beams, each of which are aimed into a large tungsten beam dump. These beam dumps are connected to the same power generation system as the DT reaction vessel itself. Proton beam energy is split amongst three as even sputtering resistant materials like Tungsten cannot withstand the full power of the attenuated proton beam.

**Conclusion**

The presented particle accelerator and associated target, inspired by the design of PSI and similar muon research facilities, serve as a starting point for the eventual development of a muon source for commercial MCF. While the proposed design would almost certainly never be capable of achieving net positive MCF, through incremental upgrades, redesigns, and improvements in efficiency could one day get there. Nonetheless, the complexity of designing, fabricating, and maintaining a high efficiency, high energy particle accelerator for a continuous power plant is an incredibly challenging task which current technology is unprepared to take on. Muon catalyzed fusion, while proven, is still considerably far from commercial viability, and arguably one of the least feasible fusion architectures presently being investigated by the larger scientific community.

## References

**Section II**

[6] Paul Scherrer Institut. (2022). Impact conceptual design review.

[7] ICTP. (n.d.). Event 8728 presentation slides.
https://indico.ictp.it/event/8728/session/2/contribution/2/material/slides/0.pdf

[8] Paul Scherrer Institut. (n.d.). Cockcroft-Walton. https://www.psi.ch/en/cas/cw

[9] Paul Scherrer Institut. (n.d.). ECR ion source. https://www.psi.ch/en/cas/ecr

[10] Morton, A. C., et al. (2011). A compact electron cyclotron resonance proton source. Review of Scientific Instruments, 82(5), Article 053304.

https://pubs.aip.org/aip/rsi/article/82/5/053304/852552/A-compact-electron-cyclotron-resonance-proton

[11] Paul Scherrer Institut. (n.d.). Injector 2: A pre-accelerator for protons. https://www.psi.ch/en/news/psi-stories/injector-2-a-pre-accelerator-for-protons

[12] Paul Scherrer Institut. (n.d.). A reliable type from the 1980s. https://www.psi.ch/en/news/psi-stories/a-reliable-type-from-the-1980s

[13] Paul Scherrer Institut. (n.d.). Injector 2. https://www.psi.ch/en/cas/injector2

[14] Schryber, U., et al. (1992). The commissioning of PSI Injector 2 for high intensity, high quality. In Proceedings of the 13th International Conference on Cyclotrons and Their Applications (Paper I-05). https://proceedings.jacow.org/c92/papers/i-05.pdf

[15] Haj Tahar, M., et al. (2025). First experimental demonstration of machine learning-based tuning on the PSI Injector 2 cyclotron. arXiv. https://arxiv.org/pdf/2512.03829

[16] Presentation MOD03. (2016). In Proceedings of Cyclotrons 2016. https://proceedings.jacow.org/cyclotrons2016/talks/mod03_talk.pdf

[17] Paul Scherrer Institut. (n.d.). HIPA. https://www.psi.ch/en/cas/hipa

[18] Paul Scherrer Institut. (n.d.). Ring cyclotron. https://www.psi.ch/en/cas/ring

[19] Paul Scherrer Institut. (2024). Protons and other particles: The HIPA facility turns 50. https://www.psi.ch/en/news/psi-stories/protons-and-other-particles-the-hipa-facility-turns-50

[20] Seidel, M., et al. (2010). Production of a 1.3 MW proton beam at PSI. In Proceedings of the 1st International Particle Accelerator Conference (IPAC'10) (Paper TUYRA03). Kyoto, Japan. https://proceedings.jacow.org/IPAC10/papers/tuyra03.pdf

[21] Paul Scherrer Institut. (n.d.). Cyclotrons at PSI [Presentation slides]. https://indico.psi.ch/event/3484/attachments/5948/7494/Cyclotrons_PSI.pdf#page=22.98

[22] Chen, et al. (2026). Muon beams towards muonium physics: Progress and prospects.

[23] Nagamine, K. (2009). What are muons? What is muon science?. In Introductory muon science (Chap. 1).

[24] Nakahara, K., et al. (2008). Design of the large acceptance muon beamline at J-PARC.