A Primer for the *Primer on Pulsed Power and Linear Transformer Drivers for High Energy Density Physics Applications*

Hannah Hasson

January 14, 2022

Contents

Preamble

This little pre-intro is here to tell you what you are about to read. This document was meant to be a way to boil down some of the main concepts in Ryan McBride's "A Primer on Pulsed Power and Linear Transformer Drivers for High Energy Density Physics Applications" [4]. While this paper is a fantastic summary of the field of modern pulsed-power research and technology, it can be perhaps a bit confusing for anyone new to the content. I, a humble graduate student in Pierre Gourdain's lab at the University of Rochester, wrote the document you are now reading in order to present the essential material in Ryan's paper to an undergrad working in my lab. I thought it might be nice to have an accessible summary of pulsed-power technology that we could give to future undergrads (and possibly high school students) who join our lab. But if we're being honest here, I also wrote it to help myself clarify the important concepts in the field too. Additionally, I incorporated some resources from the "Introduction to High Energy Density Science" graduate course I took from Profs. Rip Collins and Ryan Rygg. This summary will only assume that you have a basic understanding of introductory electricity and magnetism, especially related to current and electric/magnetic forces.

Note that this isn't necessarily meant to be a stand-alone summary. It should be used to help you follow along with the original paper. To assist you in doing that, I will break things up according to the sections in that paper (though my section titles will be a little more snarky). I hope you enjoy this, and feel free to send any feedback or questions to hhasson@ur.rochester.edu.

1 Introduction: what the heck is high energy density physics (HEDP)?

The theme of this section, and frankly the field of High Energy Density Physics (HEDP) can be boiled down to one important realization: the units for energy density are the same as the units for pressure! Let's do a quick check of this, starting with Pascals and moving toward energy per volume:

$$Pa = \frac{N}{m^2} = \frac{kg}{ms^2} = \frac{J}{m^3} \tag{1}$$

Ta-da! What this means is that if we want to study the physics of very extreme environments (with huge energy densities), we need to make sure that we are maintaining high pressures. The field of HEDP is typically defined as approximately $P \geq 1$ Mbar, or equivalently $P \geq 100$ GPa, to be specific. This is about a million atmospheres, which is makes for some very hot, very dense plasma. This regime covers a lot of different types of experiments, including fusion, laboratory astrophysics, materials science, and more! Here is a useful, though slightly chaotic plot that shows where a number of physical systems sit with respect to the temperature and density threshold for HEDP.

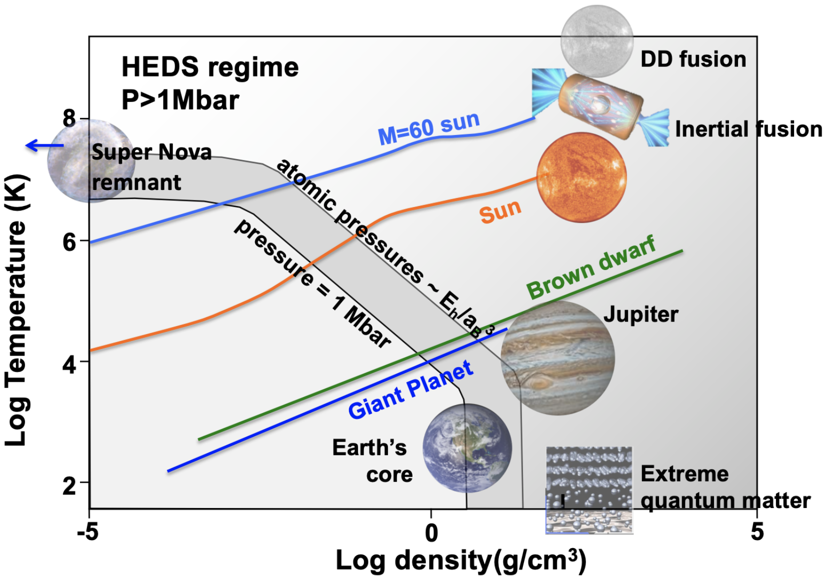

Figure 1: A diagram of some different types of plasma and where they sit relative to the high energy density physics (HEDP) threshold. You may also see it referred to as HEDS, or high energy density **science**. Everything outside of the line marked "pressure = 1 Mbar" lies in the HEDP regime. This includes the cores of gas giants, various layers of the sun, and the conditions in inertial fusion experiments. Reproduced from [1].

Figure 1 is from the "Introduction to High Energy Density Science" class that I took during Fall 2020 from Profs. Rip Collins and Ryan Rygg. All the way at the bottom right of the plot in the low-temperature, high density regime lie the ordinary solids, liquids and gases that we are familiar with. The top right are conditions you only find inside of stars.

Now let's identify the facilities that have the ability to make matter in HEDP conditions:

- The Z Machine at Sandia National Lab
- The National Ignition Facility (NIF) at Lawrence Livermore National Lab (LLNL)
- The Omega EP & Omega lasers at the University of Rochester's Laboratory for Laser Energetics (LLE)
- There are others too but I guess they weren't cool enough to make the cut for this paper (SLAC, Texas Petawatt Laser, etc...)

Note that all of these are either lasers or pulsed-power generators. Those are the kind of facilities that are capable of producing such extreme environments. This summary will cover the pulsed-power side of the coin only, though. Trust me, that's enough to bite off at once.

2 A simple picture of a pulsed-power HEDP experiment

This section is meant to take you through the most important principles of what pulsed-power can do, which is effectively generate enormous pressures to squeeze (or throw) things. Kind of like a powerful, angry baby. A pulsed-power driver does this by generating a very large voltage, which then drives a huge current over a very short amount of time. This absurd amount of current produces huge magnetic fields that exert pressure. Don't worry yet about *how* the machine produces this electrical pulse, just assume that it generates potentially millions of Amps of current over a couple hundred nanoseconds. The most basic and ubiquitous example of an experiment on a pulsed-power driver is a Z-pinch. Let's talk about what this kind of load ("load" = the hardware you stick in your machine and run the enormous current through) does.

2.1 The Z-pinch: squeezing a soda can into oblivion

A Z-pinch is the vanilla of the pulsed-power experiment varieties; it is the simple base on which many other flavors/experiments are formed. We will start by pretending we have a little soda can filled with and surrounded by vacuum (for the purpose of this analogy, we ignore the top and bottom of the can). What are its properties? It is just a conductive cylinder with super thin walls at some radius r from the center (Fig. 2).

It also has an inductance L, which is what slows down the *change of current* in it. L is determined by the shape and size of the load.

Figure 2: Please enjoy this mediocre drawing I made of a can with no top or bottom on it.

We stick this little can into the machine so that the top and bottom are connected to electrodes, and then apply some giant voltage (V) across the soda can to drive a current down the wall. This time let's look at the actual figure from the paper, which is arguably nicer than my crappy drawing.

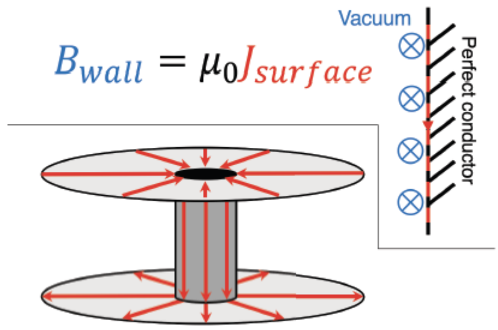

Figure 3: This figure, reproduced from [4], shows the current flowing in the top electrode, through the can, and finally into the bottom electrode. On the top right, there is a zoom-in of what the current looks like in the wall of the can, as well as the magnetic field immediately tangent to the can's surface. The magnetic field strength at the can's wall is given by B_{wall} and the linear current density (current/circumference) is $J_{surface}$.

Because of the inductance of the can, the rate of increase of current is limited to

$$\dot{I} = dI/dt = V/L, \tag{2}$$

which comes from the definition of voltage generated by an inductor. Typically the current rise-time for a low-inductance load on a pulsed-power driver is on the order of 100 ns. In these experiments, the current is assumed to flow on the surface of the conductor. This has to do with the *skin depth*, which tells us how deep past the surface of the conductor the current flows. This quantity is very small for high frequencies, meaning that you will have effectively only surface current. Since the current rise and fall in a pulsed-power driver is really short (100s of nanoseconds), it ramps up and down just like high-frequency alternating current, and thus the skin depth is small.

Once the current is running through the conductor, a magnetic field is produced in the vacuum around the can (see Fig.3). This field is tangent to the can at each point and has the magnitude $B_{wall} = \mu_0 J_{surface}$. Why is that the case? Let's do some right-hand rule shenanigans to figure it out.

If we look at the field produced at one point on the circular cross-section of the can due to several pieces of current, we get the net field at that point. Try it with these four points I illustrated for your viewing pleasure. The current from the lower three points produces a field at the point on top, which we can reduce to a net B field.

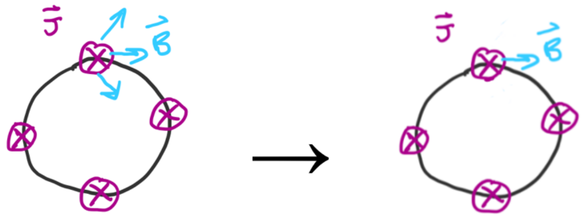

Figure 4: The magnetic field for an arbitrary position θ on the circular cross-section of the can/cylinder. The current in the can is going into the page (purple X's). The lefthand figure shows the magnetic field lines (blue arrows) at the top point due to the lower three pieces of current; and the righthand side shows the net magnetic field at the same point, with the up and down components of the field cancelling. We see that the magnetic field is tangent to the surface of the conductor, just as the previous figure showed!

You can make the density of current points as high as you want, as long as you distribute them at equal intervals along the wall's cross-section, and you'll get the same answer. So now we see why the magnetic field points tangent to the surface of the conductor. So now we can consider the field direction at each of the four current points we originally drew (Fig. 5).

But why is the magnitude $B_{wall} = \mu_0 J_{surface}$? Here, $J = I/2\pi r$, so if you plug that into B_{wall} you just get the magnetic field for a wire $B_{wire} = \mu_0 I/2\pi r\widehat{\theta}$. This makes sense, since if you applied Ampere's loop law (that loop thing you do in intro electricity & magnetism) you would draw the same loop around a tube of current as you would for a wire. Woop, there it is.

What happens next? Well, we have a magnetic field and some moving charge (the

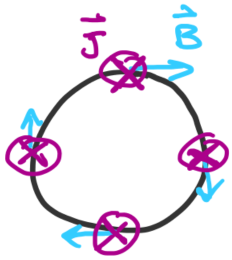

Figure 5: The net magnetic field for each point on the surface of the conductor is tangential!

current), so there's going to be a magnetic force. We'll take the magnetic part of the Lorentz force:

$$F_B = qv \times B. \tag{3}$$

But we want to deal with things in terms of *areal current density* $J = \sum_i n_i q_i v_i$ (you'll see why). This is current per area (A/m^2). Now we substitute J into the magnetic force and get

$$F_{B/A} = J \times B. \tag{4}$$

This gives the force per area because J is current per area. People like to refer to this as the "J cross B force" because I guess nobody thought of a snazzy name. If we look at the direction that this lovely cross-product gives us, we find that the force points radially inward toward the axis everywhere on the can. That means it will be crushed inward!

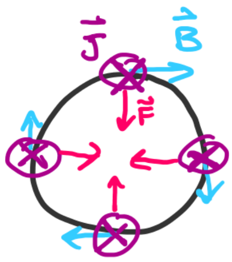

Figure 6: The "J cross B force" points radially inward along the axis, and thus the can will be crushed.

Force per area is just pressure (look at the units if you forgot)! So we can understand this as a magnetic pressure $p_{mag} = J \times B$. We like to think of this as a "vacuum pressure" that pushes in on the can, up on the top electrode, and down on the bottom electrode. Just imagine the vacuum region were filled with gas that had some pressure p_{mag}. Since we said in the beginning that pressure is the same as energy density, we can equate p_{mag} to the typical magnetic field energy density formula

$$p_{mag} = B^2/2\mu_0 \tag{5}$$

The energy density of the *electric field*, on the other hand, is insignificant in this situation because can has low voltage (and thus low E) since the it is pretty much a short-circuit.

Now, we can put our magnetic field $B = \mu_0 I/2\pi r\widehat{\theta}$ from earlier into our magnetic pressure formula:

$$p_{mag} = B^2/2\mu_0 = \frac{\mu_0 I^2}{8\pi^2 r^2} \tag{6}$$

The takeaway here is that $p_{mag} \propto I^2/r^2$. So if we want to maximize the pressure and crush that soda can, we should make sure our current is very high and our soda can radius is very small. It is also important that the walls of the can need to be thin/light as well so that the force can move it easily (small moment of inertia).

2.2 An application of the Z-pinch: magnetized liner inertial fusion (MagLIF)

So what? Who cares if you can squeeze a little metal tube? I'll tell you who: any government who wants fusion energy.

The Z-pinch is applied in MagLIF, or Magnetized Liner Inertial Fusion. The main difference there is that instead of your soda can being empty, it has deuterium in it (not the kind of soda I recommend drinking). You heat up the fusion fuel inside the can (also known as the "liner") with a laser, then run the current and crush the can with the magnetic pressure we discussed. The high temperatures and densities achieved can potentially ignite fusion! MagLIF can attain pressures of 100s of Mbar, well into the HEDP regime. The big fusion laser facility NIF can only achieve up to $\sim$ 140Mbar... It sounds like a lot, but NIF has sadly been a flop when it comes to achieving fusion at this point... which is awkward because it's called the National *Ignition* Facility. But you didn't hear it from me!

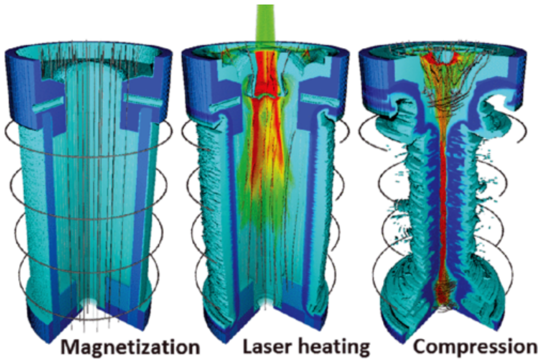

Figure 7: This is the infamous cartoon of MagLIF that everyone and their mom uses when presenting anything related to it. Gracefully stolen from [2].

2.3 Some other fun pulsed-power loads

In addition to Z-pinches (cylindrical loads on the z-axis like above), there are other cool load geometries we care about. Check out the cool diagram in Fig. 8 to get an idea of where you can put things in the current path to achieve different types of plasma geometries/dynamics.

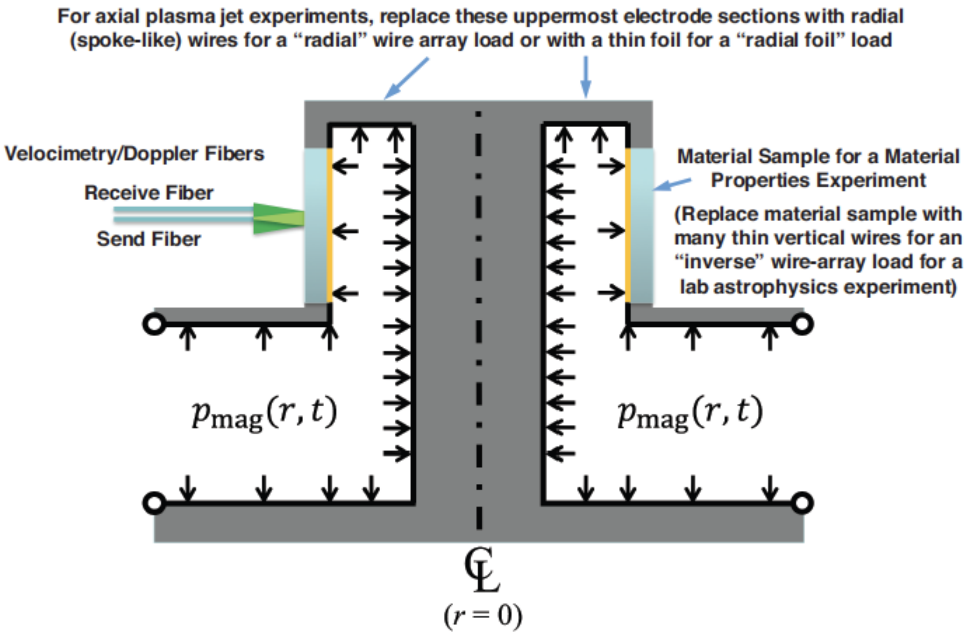

Figure 8: Once again, I just straight-up stole the diagram from [4] because I can't draw. This shows different load orientations that could be used for certain experiments. The circles at the end of the load are where it connects to the electrodes of the pulsed-power driver.

Let's go through some of the cool options and what they can do:

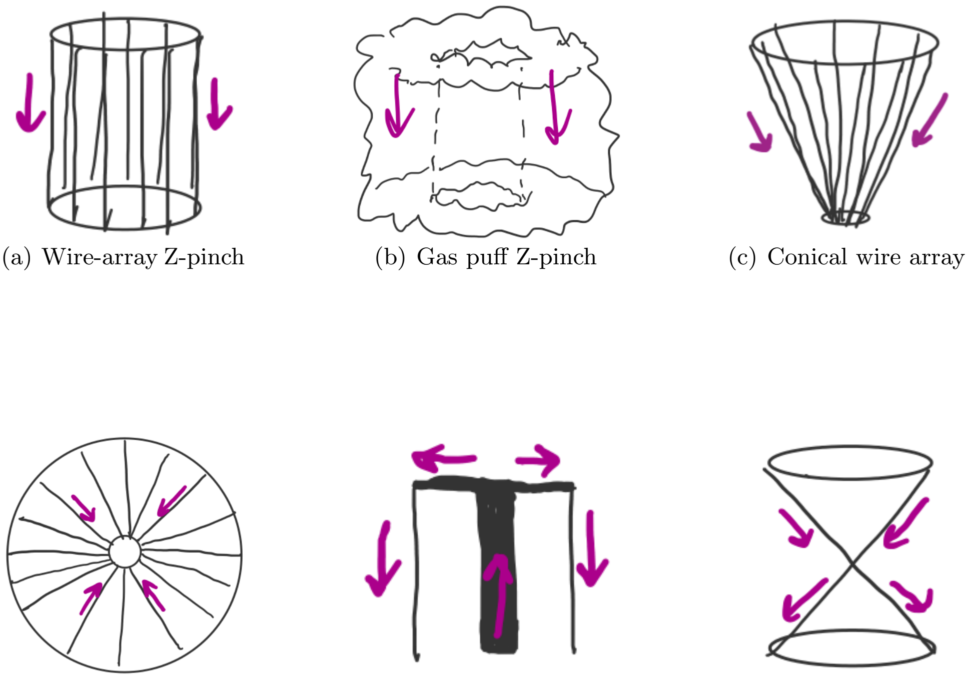

(a) Wire-array Z-pinch (b) Gas puff Z-pinch (c) Conical wire array

(d) Radial wire-array (e) Exploding wire-array (f) X-pinch

Figure 9: A pu pu platter of load hardware designs for different experimental platforms.

2.3.1 Wire-array Z-pinches

This is still a Z-pinch, but you replace the soda can with a bunch of vertical wires. This sort of thing is good for generating x-rays, which can be used for radiography (x-ray imaging). Check out the sick drawing that I made of a wire-array Z-pinch in Fig. 9(a). This implodes on axis just like a normal Z-pinch. Current follows the path in purple.

2.3.2 Gas puff Z-pinch

Ok fine this is also a Z-pinch, but you just blow a puff of gas downward in the shape of a cylinder (Fig. 9(b)) and run the current through that. These things are pretty difficult to do, and require a special nozzle that has the anode and cathode on it. You also have to pre-ionize the gas so that your giant current pulse will conduct through it, otherwise you'll break the heck out of your driver. Gas puffs can also be used to generate x-rays.

2.3.3 Conical wire arrays

Now things are getting juicy. Conical wire arrays (Fig. 9(c)) are similar to wire array Z-pinches, but here the wires ablate plasma that flows up the axis from the narrow to the wide end. This scheme is popular for generating jets to do laboratory astrophysics! You can also twist them for added fun.

2.3.4 Radial wire arrays & foils

This is the same wire array concept, but instead you don't drive the current along any vertical path. It goes straight from the outer ring (the anode) to the inner ring (the cathode). Figure 9(d) shows an above view looking down at it.

A radial foil works in a very similar way, but instead of wires you use a disk of thin foil. Many of these experiments involve a pin cathode instead of the ring used at the center of wire arrays. The current travels from the outside of the foil to the center point and then down the cathode pin. Many experiments with radial foils literally use common household aluminum foil that you can buy from the store. DIY pulsed power!

2.3.5 Inverse/exploding wire arrays

The inverse wire array (Fig. 9(e)) is the same as a conical wire array, but is not positioned in the center where the current initially runs. A thicker stalk is placed in the center, and the wires are outside of this. As the current runs through these outer wires, the magnetic pressure causes the plasma made to push outward rather than inward. See Fig. 8 at the beginning of this section for a clearer picture of where this sits in the current path.

People do lots of fun lab astro and shock physics with exploding wire arrays. You can stick a solid object in the way of the outflowing plasma and study the shocks formed! You can also put two of these next to each other and study magnetic reconnection by throwing some plasma at some more plasma.

2.3.6 X-pinches

Finally, one of the hottest (pun intended) load types in pulsed-power: the X-pinch. These are very simple yet extremely useful. They consist of two or more wires that cross at a single point (Fig. 9(f)). When the current runs through, the point where the wires cross pinches with TONS of pressure, and generates what is called a hotspot. These hotspots (there can be more than one generated) emit some spectrum of x-rays dependent on the material and wire conditions. These x-rays are a great source for doing radiography on other loads. The pinch also makes two small columns of plasma (like mini Z-pinches) just above and below where the wires cross.

3 RLC circuits on steroids: A simplified model of how pulsed-power drivers work

Alright, so we've been having fun so far. Now let's get a little bit into the nitty-gritty of things. We understand the load and what the machine puts into it: a high-amplitude pulse of current with a very short rise-time. But how does the machine generate that? What is inside these crazy sci-fi death machines anyway? What is the difference between a Linear Transformer Driver (LTD) and a Marx generator? Don't worry, these questions and more will be answered soon enough.

First, what on this earth is inside of an LTD? It's simpler than you might think. A bunch of capacitors, resistors and inductors sitting in oil, more or less! Hopefully you remember what each of these parts do: capacitors store charge/electric fields with charged plates sandwiched around dielectrics; resistors are materials that dissipate some of your energy into heat, reducing the amount of current that gets across; and inductors are components that slow down changes in current, storing the energy in magnetic fields. We can look into the layout of these components in more detail later, but now you essentially know who the man behind the curtain is.

Let's next think about the design of an LTD by identifying what we need in order to obtain high energy-density/pressure. **First: generate a ton of power.** Recall the formula for power in electronics,

$$\frac{d\epsilon}{dt} = P_{electric} = V \times I \tag{7}$$

where ϵ is energy.

Since we need high current to get high power, and $I = \frac{dQ}{dt}$ where Q is charge, then we also need to **store a lot of charge** in the capacitors. For the capacitors to do this they must have a high capacitance C, since $Q = CV$. We also want to have lots of capacitors, since the capacitance multiplies by the number of them we have ($C = n\frac{\epsilon_0 A}{d} = nC_i$, where C_i is the capacitance of an individual capacitor). You can think of the connected capacitor bank as having a ton of small capacitors next to each other whose bottom plates are all connected to each other, and all the top plates to one another. It's basically like having one giant bottom plate and one giant top plate.

We also want to maximize the voltage in $Q = CV$. We can achieve that with a small inductance since $\frac{dI}{dt} = V/L$, and $\frac{dI}{dt}$ is related to the current pulse shape (which we probably won't be changing). L can be derived by doing some fun mumbo-jumbo with fluxes (see [4] for that), but I'm just going to tell you what it is:

$$L = \frac{\mu_0 h}{2\pi} ln(\frac{r_{out}}{r_{in}}) \tag{8}$$

Here, r_{out} and r_{in} are the boundaries of the vacuum region and h is the height. This inductance is the same as a perfectly conductive coaxial transmission line the height of the vacuum cavity (transmission line = pair of electrodes that carry current over long distances). Figure 10 shows a visual of what each of these parameters is referring to for this idealized case.

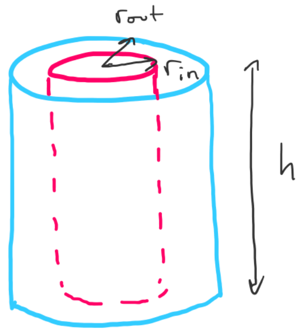

Figure 10: A drawing of an ideal coaxial transmission line. The anode is the blue outer layer, and the cathode is the pink inner layer. There is an implied layer of vacuum between them, and the anode and cathode would in practice be connected at one end by the load and connected to power supplies at the other end. You can think of this as basically a big version of a BNC cable if that helps.

This solution applies when the current is a surface current. Remember that skin-depth thing we talked about in Section 2? Let's go a little further into the conditions for it now. Remember that the skin depth tells you how deep below the surface current can run. The formula for it is

$$\delta_{skin} = \sqrt{\frac{4\rho_e \tau_r}{\pi \mu_0}} \tag{9}$$

where ρ_e is the electrical resistivity, τ_r is the rise time of the alternating current, and $\mu_0 = 4\pi \times 10^{-7}$ H/m is a constant called the vacuum permeability. If we have a rise time of $\tau_r \sim 100$ns and resistivity of $\rho_e \sim 100$nΩ we get $\delta_{skin} \sim 100\mu$m. That's pretty small compared to the typical dimensions of electrodes in these experiments, which are on the *cm* scale.

However, we are not being realistic enough. Our solution is for a straight coaxial transmission line of height h, but it's possible that the anode and cathode are curved. So now, let's make things a little spicy and evaluate

$$L = \frac{\mu_0}{2\pi} \int_{Z_1}^{Z_2} ln(\frac{r_{out}(z)}{r_{in}(z)}) dz. \tag{10}$$

Figure 11 is a nifty picture of transmission lines I stole from Ryan McBride's paper because there's literally no way I could illustrate this with my skillset.

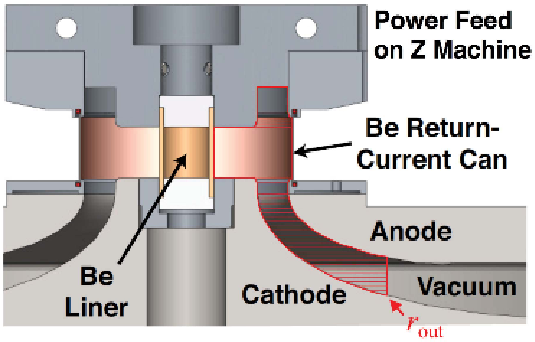

Figure 11: Reproduced from [4]. This is a slice of the inside of a cylindrically symmetric, curved transmission line. The anode and cathode are marked, and the load is a cylindrical liner shown in gold. Here it's made of Beryllium (Be), which is very very bad for humans to breathe next to.

Remember how we said we want to minimize L? One way to do this is by minimizing Δz, which is just the height of the transmission line. If you need some vertical translation, it's best to do it at a large radius, since the integral above in Eqtn. 10 evaluates to be smaller for larger r_{in}. There are some examples of optimized power feeds in [4], but I'm going to skip those and stick to the main picture.

Now it's time to find our current, $I(t)$. We will use a simple LC (inductor and capacitor) circuit to do this, since an LTD resembles that type of circuit. It's just a bunch of capacitors, inductors and resistors, remember? For now, we will assume that L=constant, and R=0 (more on this in a minute). Start with the voltage across the inductor:

$$V = L\dot{I} \tag{11}$$

Now look at the displacement current through the capacitor:

$$I = -C\dot{V} \tag{12}$$

Where did this come from? Remember that, for a capacitor, $Q = CV$. Just differentiate that bad boy and you've got a current. If you then differentiate this current with respect to time you get $\dot{I}$, which you can plug into the first equation (11). After plugging in $\dot{I}$, solve for voltage from our cute differential equation we get.

$$V = V_0 cos(\omega t). \tag{13}$$

Huh! That's just a simple harmonic oscillator (verify this yourself). We can throw this solution back into the capacitor current equation (Eqtn. 12) et viola:

$$I = I_{peak} sin(\omega t) \tag{14}$$

$$Ipeak = V_0\sqrt{C/L} = V_0/Z_0 \tag{15}$$

where Z_0 is the impedance of the circuit (Impedance is basically resistance for AC circuits).

This is all fine and dandy, BUT we neglected resistance and set inductance constant at the beginning, so the harmonic oscillating current just keeps going forever and isn't realistic. So what to do about our nonzero resistance? The resistance dissipates energy through electrons colliding with atoms or other electrons in the resistor, which generates thermal energy. This is referred to as *ohmic heating*. The ohmic heating rate is

$$P_\Omega = V_\Omega I = I^2 R. \tag{16}$$

This will contribute to damping of the harmonic oscillation.

In real HEDP experiments, we can have the power in the inductor vary with time as the amount of current changes or the current path changes shape (e.g. the soda can gets squished). The energy in the circuit that gets lost through the inductance is

$$P_L = V_L I = \dot{L}I^2 + L\dot{I}I \tag{17}$$

The last term on the right corresponds to how much energy the changing current exchanges with the magnetic field, assuming the inductor (a.k.a. the load) geometry is constant for that instant. The other term is the opposite: how much energy is exchanged with the magnetic field from a constant current as the inductor geometry changes.

For example, an imploding liner (as in a Z-pinch) would have

$$\dot{L} = -\frac{\mu_0 h}{2\pi}(\dot{r}(t)/r(t)). \tag{18}$$

Having geometry that changes with time dissipates energy in the system by converting it into both

- mechanical energy (magnetic force squishing the load)
  - $P_{mech} = \frac{1}{2}\dot{L}I^2$
- more magnetic field that fills the expanding vacuum region as the load geometry changes/compresses
  - $P_{B_{new}} = \frac{1}{2}\dot{L}I^2$

There's a proof in the full paper for why these quantities are equal, but I'm tired and I'm just going to leave that as an ∼*exercise for the reader*∼ for you to go verify that. Sorry bout it.

Let's think about some implications of these things we just learned:

What happens when things are no longer compressing? After squeezing our fusion fuel down to a certain radius in a MagLIF experiment, the fuel will start to push back and the liner will effectively bounce off of it and fly back outward. Now mechanical energy can be exchanged in the opposite direction in the P_{mech} formula, resulting in both a different $\dot{L}$ and I. This means that the current can spike post-compression! This is called **magnetic flux compression** of the driving field (and thus the driving current). The same phenomenon also causes a corresponding current *decrease* when the liner is compressing initially. This is referred to as an "inductive dip."

So we now see ways that resistance and changing inductance can pull energy out of the LC circuit. Therefore we don't see oscillating behavior in an LTD in practice. Because there is one main current pulse typically, folks will try to design their Z-pinch liners such that they will compress to their minimum radius when the current is peaking. But since there are all these weird inductance effects (i.e. the inductive dip), things become a little more complicated.

4 Marx bank generators and the behemoth that is the Z-Machine

Alright, so you want to play with the big kids now, eh? Then you've come to the right place. The Z-Machine, located at Sandia National Labs in Aluquerque, NM, is the biggest/highest-current pulsed-power driver in the world. This thing can produce up to 25 MA of current and literally causes the ground to quake a quarter mile away when it fires (speaking from personal experience). The top photo in Figure 12 shows the famous "arcs and sparks" photo from above Z as it fires. Hilariously, this sparking was not supposed to happen, and definitely doesn't happen on a normal day. The bottom photo of the figure shows Z on a normal day, without all the glamor, etc.

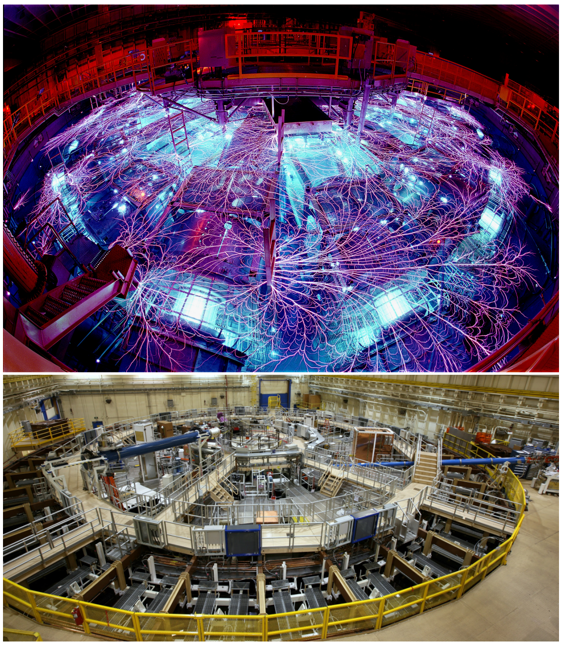

Figure 12: Ooooh pretty. On top is the classic photo of firing the Z Machine that is pasted onto every single flyer or talk ever made about it, reproduced from [6]. On bottom is Z in sweatpants with no makeup on (reproduced from [5]).

So you're probably wondering how this thing is built. At the perimeter of the machine, below the yellow handrails, are 36 capacitor banks. These are each situated in a tank of oil to protect from electrical arcing. These capacitor banks are each called a "**Marx generator**," meaning that they have a specific circuit layout (see [4]) that allows the capacitors to charge in parallel and discharge in series. This means you can charge the capacitors separately pretty quickly, and then add all of their voltages together when you discharge. Each capacitor is connected to a *spark-gap switch*, which is just a fancy jar of air whose ends have electrodes. You charge up the ends until electrical breakdown occurs, completing the circuit and releasing all that charge stored.

Here are some stats of the terrifying majesty that comes out of 36 Marx banks:

- Up to 25 MA of current (46 million 60W lightbulbs)
- Up to 23 MJ of energy (23 sticks of dynamite)
- Up to 5.4 MV of voltage (49,000 times higher than a wall outlet)
- Peak power up to 80 TW (more than all the world's power plants together)

The rise time for these Marx banks is about 1 μs due to the high inductance (remember inductance slows down change in current). But we want the same 100 ns rise time as mentioned before. So the pulse is compressed in time through some witchcraft called "pulse-forming lines." These are intermediate capacitors that progressively trigger with shorter rise times. They sit in a big ol' tank of water for insulation and because it's easier to filter junk out of water than lots of other fluids. Plus, sometimes when bubbles appear in the tank, people literally have to scuba dive in there to dispose of the bubbles. Having pockets of air can cause points where electrical breakdown occurs (bubbles=things break more often).

The next step in the machine is the transmission lines, which are separated into outer and inner by a vacuum insulator stack. This is where the water-insulated part of the machine ends and the vacuum-insulated part starts. The rest of the machine interior to this stack is shown in Fig. 13.

I know this looks like a lot and there are arrows all over the darn place. What you need to know from this is that the blue anode lines carry the current in and the red cathode lines carry the current back out, and that all the same-color lines combine at the top. These transmission lines are called magnetically-insulated transmission lines (**MITLs**, pronounced "might-uls") because the amount of current in them produces such high magnetic fields that any electrons that stray from them will get pulled back in by the magnetic force. Hence the "insulation" is provided by the magnetic field and a dielectric isn't needed between the transmission lines. The reason there are several anode and cathode lines that eventually combine is to reduce the inductance.

The final result of this is pretty similar to what you get with an LTD: about a 100 ns rise time for millions of Amps of current. So why are people switching over to LTDs?

5 Linear Transformer Drivers (LTDs): powerful donuts

Ok so we already learned in section 3 that LTDs are basically a slightly fancier version of an LC circuit. So what's with all the hoopla with people switching from Marx generators to LTD technology? The main benefit is that you don't need all these pulse-compression

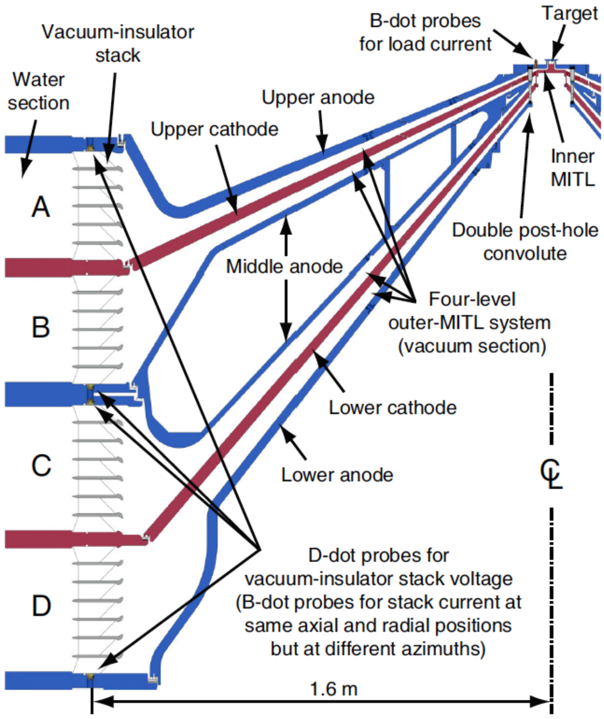

Figure 13: A cross-section of the cylindrically-symmetric geometry of the Z Machine's transmission lines. Elegantly screenshotted- *ahem* reproduced from [4].

stages. This means that you save yourself space in your machine, in addition to having less components to potentially break.

How do you make the same huge current pulse without all that extra stuff? Lower the inductance! Remember that the high inductance of the Marx banks caused the long 1 μs pulse length. So if we can make this thing super compact, we can lower L and hopefully keep the capacitance quite high too. The way people like to design these presently is in the layout in Fig. 14 which shows University of Michigan's MAIZE LTD.

The layout is the following: around the outside of the Death Donut^{TM} are a bunch of white brick things, each connected to a shiny disk. Each of these is a *pair* of two massive white capacitor bricks connected by a spark-gap switch (the shiny disk part). Same principle as discussed before with the switches, you fill them with gas and then charge them to the point of electrical breakdown. You can only see the top capacitor in each of these units, but trust me when I tell you they are each a pair. The capacitors are then combined in parallel and fed into the transmission lines, which would be connected from the inner diameter of the donut to the target vacuum chamber sitting on top at the center of the donut. The target chamber is the donut hole of the Death Donut^{TM} if

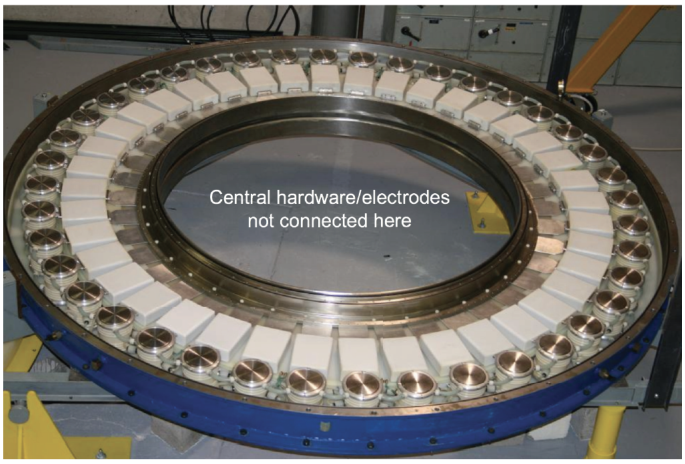

Figure 14: The MAIZE LTD module without any of the stuff in the center. Basically just a big empty electrical donut. Stolen with humility from [4].

you will.

The last component that you need to worry about is the magnetic cores, which live just interior to where the ring of capacitors are. The magnetic cores are just there to make sure the current follows the path from the capacitors to the transmission lines. Otherwise you could have the current taking a path through the metal casing of the LTD. This deters that by raising the inductance of that path.

And that's really it! This simple geometry produces a sick 900 kA of current with a rise time of 200 ns. The total inductance on this setup is about 20 nH (quite small). Another benefit of the LTD is the fact that all the components sit in dielectric oil and are enclosed in a nice metal casing. So there's not interference from electromagnetic fields outside the machine. Nice.

What's even more fun than one donut is a stack of donuts, though. You can stack these bad boys and combine them all in series along the transmission lines! This means the current runs radially to the inner wall of each donut, then travels along the transmission line at the center. That's exactly what the HADES driver at the University of Rochester (Fig. 15) does.

LTDs 1-3 are first combined in series, as are LTDs 4-6; this adds the voltage of three LTDs so that each composite module has $V_{module} = 3V_{cavity}$. These each have a separate transmission line (top/bottom anode), which are then combined in parallel in the vacuum chamber. This adds the current so that $I_{load} = 2I_{module}$. It is also very compact, and can be transported relatively easily by taking each individual LTD off and rolling them out the door (they are designed to fit!).

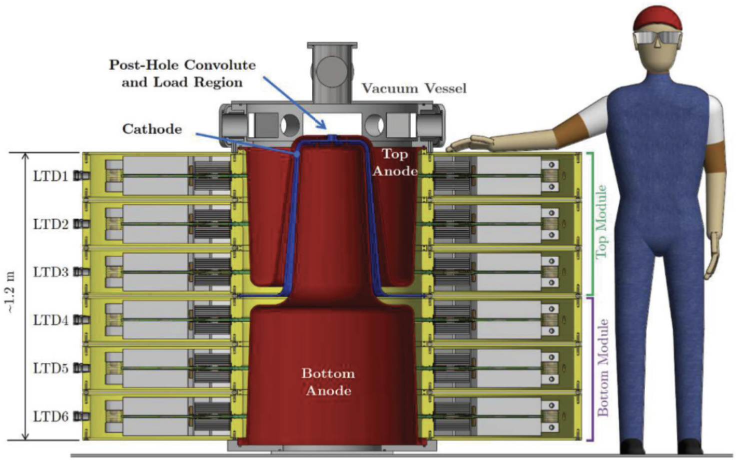

Figure 15: A cross-sectional view of the cylindrical HADES pulsed-power driver at the University of Rochester. HADES has 6 individual LTDs that are combined in groups of 3 in series, then those two groups are combined in parallel. Reproduced from [3].

People in the field are trying to apply this technology to the future giant machines in the field. Namely, the future of the Z-Machine: the proposed Z-300 and Z-800 models [7]. Their schematics are in Fig. 16 They are designed to essentially be the Godzilla version of HADES. Instead of one stack of six donuts, Z-300 would have thirty-three stacks of **ninety** donuts. *That's 2,970 donuts (hope you're hungry).* And each one of those is an LTD like MAIZE. This design is theorized to produce 48 MA of current. The Z-800 design has 5,400 LTDs in stacks of 90 for a total current output of 60 MA. Yeesh.

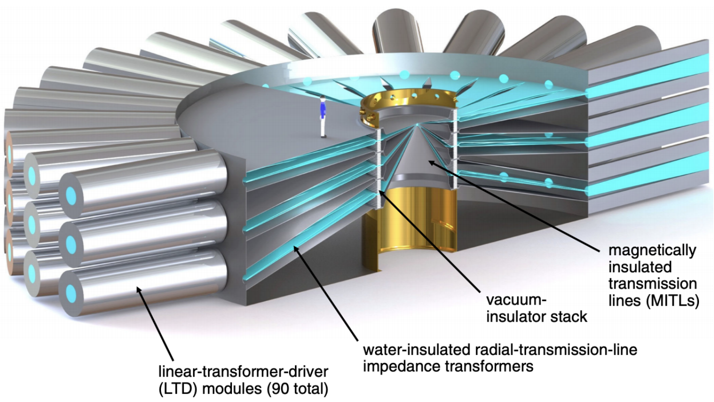

Figure 16: A cross-sectional view of the proposed Z-300 driver from [7]. This thing is truly massive at a total of 2,970 LTDs in the model. The driver is estimated to produce 48 MA of current.

6 LTD-Driven HEDP Research from Around the World + outro

I'm not just saying this because I'm lazy: I don't think it's necessary to summarize this. Go make some friends in the field and ask them about what they work on! Or read the section in the original paper :)

And that's it for this summary! Hopefully some of this content was intelligible and you learned a thing or two about this exciting technology. I hope that it helps your confidence to know that I came into graduate school knowing almost nothing about this field and these machines. We all have to start from nothing at some point in the learning process, so don't be afraid to ask questions! It all starts to make more sense with time.

Bibliography

[1] G. Collins and J. R. Rygg. Introduction, 2020.

[2] M. R. Gomez, S. A. Slutz, A. B. Sefkow, D. B. Sinars, K. D. Hahn, S. B. Hansen, E. C. Harding, P. F. Knapp, P. F. Schmit, C. A. Jennings, T. J. Awe, M. Geissel, D. C. Rovang, G. A. Chandler, G. W. Cooper, M. E. Cuneo, A. J. Harvey-Thompson, M. C. Herrmann, M. H. Hess, O. Johns, D. C. Lamppa, M. R. Martin, R. D. McBride, K. J. Peterson, J. L. Porter, G. K. Robertson, G. A. Rochau, C. L. Ruiz, M. E. Savage, I. C. Smith, W. A. Stygar, and R. A. Vesey. Experimental demonstration of fusion-relevant conditions in magnetized liner inertial fusion. *Phys. Rev. Lett.*, 113:155003, Oct 2014.

[3] P. A. Gourdain, M. Evans, B. Foy, D. Mager, R. McBride, and R. Spielman. Hades: a high amperage driver for extreme states, 2017.

[4] R. D. McBride, W. Stygar, M. Cuneo, D. Sinars, M. Mazarakis, J. Leckbee, M. Savage, B. Hutsel, J. Douglass, M. Kiefer, et al. A primer on pulsed power and linear transformer drivers for high energy density physics applications. *IEEE Transactions on Plasma Science*, 46(11):3928–3967, 2018.

[5] R. Mitnick. Sandia z machine, Oct 2020.

[6] R. Montoya. The sandia z-machine, Oct 2011.

[7] W. A. Stygar, T. J. Awe, J. E. Bailey, N. L. Bennett, E. W. Breden, E. M. Campbell, R. E. Clark, R. A. Cooper, M. E. Cuneo, J. B. Ennis, D. L. Fehl, T. C. Genoni, M. R. Gomez, G. W. Greiser, F. R. Gruner, M. C. Herrmann, B. T. Hutsel, C. A. Jennings, D. O. Jobe, B. M. Jones, M. C. Jones, P. A. Jones, P. F. Knapp, J. S. Lash, K. R. LeChien, J. J. Leckbee, R. J. Leeper, S. A. Lewis, F. W. Long, D. J. Lucero, E. A. Madrid, M. R. Martin, M. K. Matzen, M. G. Mazarakis, R. D. McBride, G. R. McKee, C. L. Miller, J. K. Moore, C. B. Mostrom, T. D. Mulville, K. J. Peterson, J. L. Porter, D. B. Reisman, G. A. Rochau, G. E. Rochau, D. V. Rose, D. C. Rovang, M. E. Savage, M. E. Sceiford, P. F. Schmit, R. F. Schneider, J. Schwarz, A. B. Sefkow, D. B. Sinars, S. A. Slutz, R. B. Spielman, B. S. Stoltzfus, C. Thoma, R. A. Vesey, P. E. Wakeland, D. R. Welch, M. L. Wisher, and J. R. Woodworth. Conceptual designs of two petawatt-class pulsed-power accelerators for high-energy-density-physics experiments. *Phys. Rev. ST Accel. Beams*, 18:110401, Nov 2015.