

NONLINEAR DYNAMICS SUPERSEDES...

Mirrored Symmetric Circuit Si...

0:00 14:06

Listen now

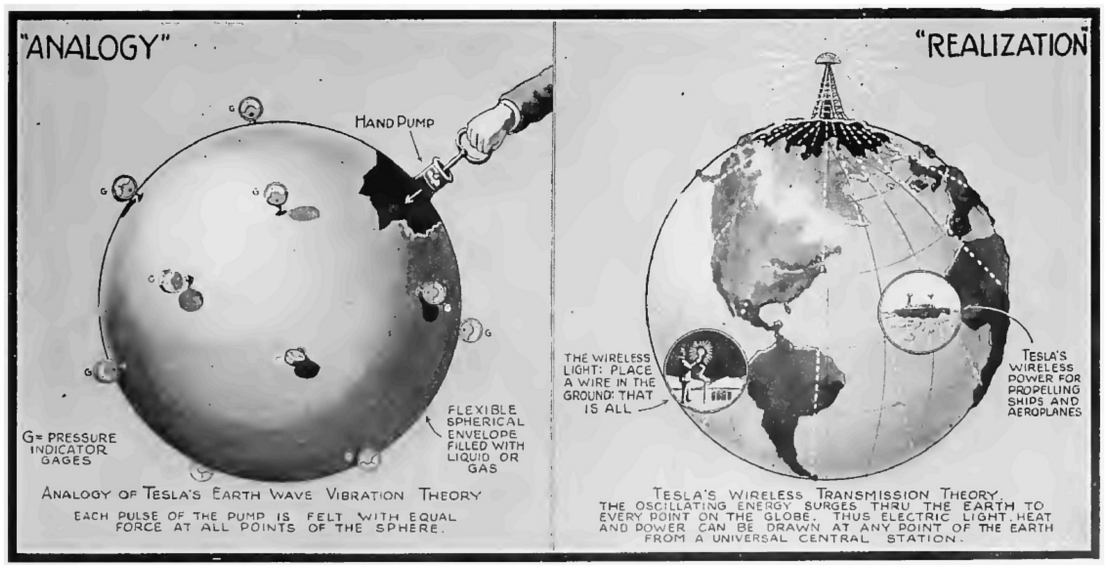
Mirrored Symmetric Circuit Simulation in LTSpice. Finally! Dissection, pt. 3

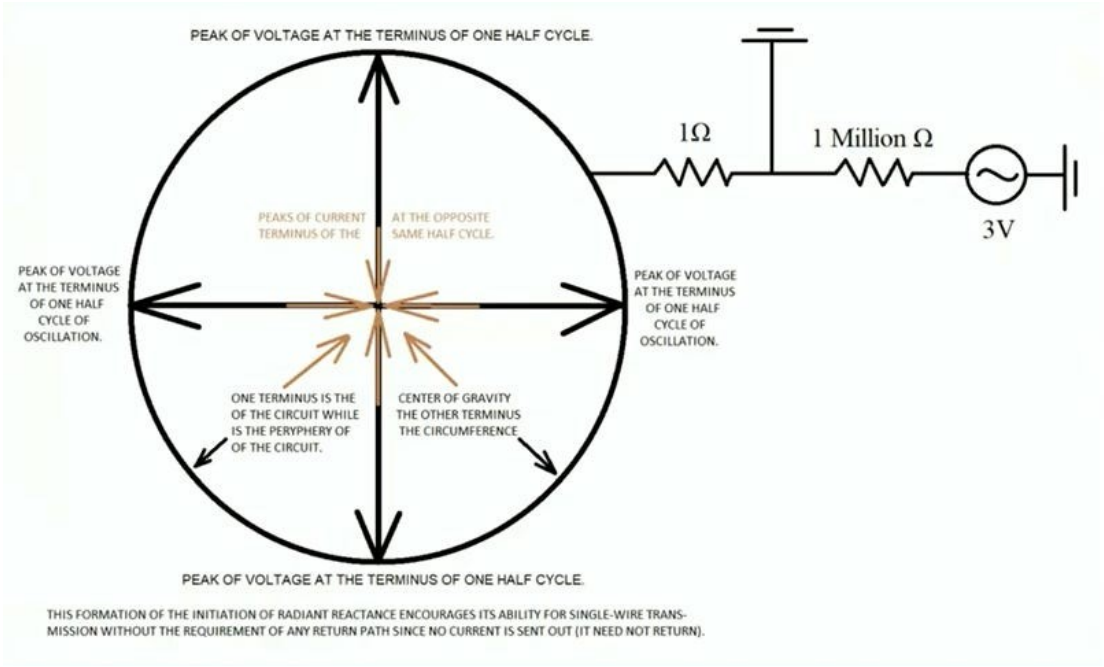
In the following dialogue with AI, this discussion is heating up to a progressive prognosis which transcends conventional stipulations of normalcy.

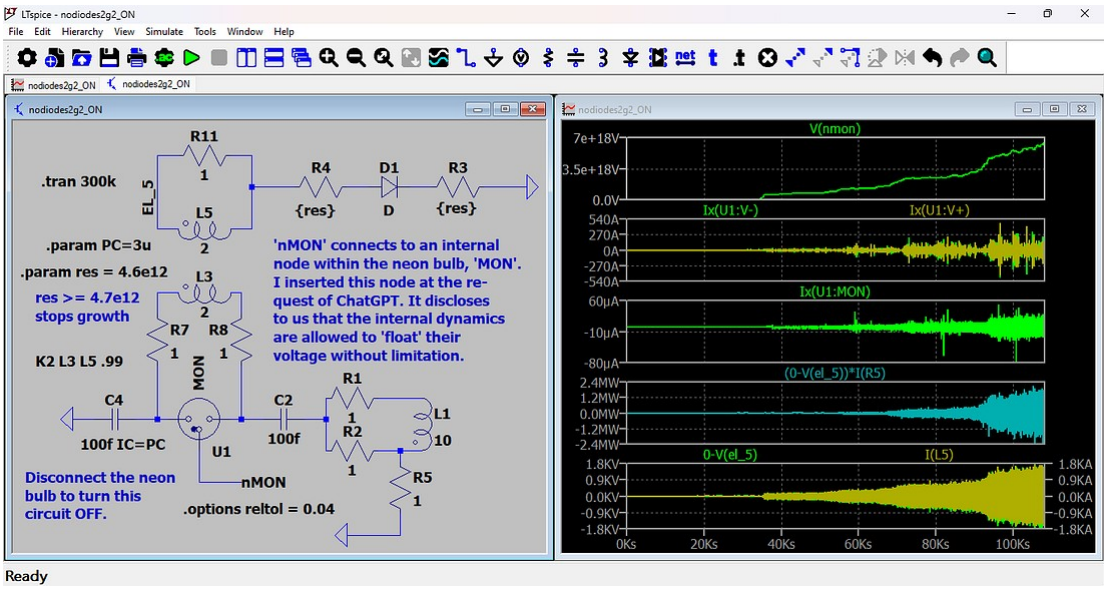


VINYASI

OCT 04, 2026

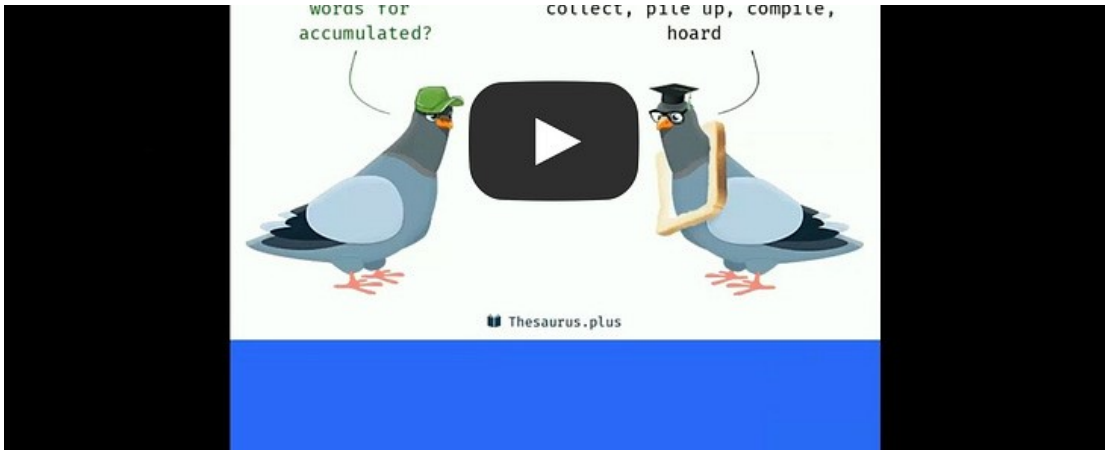






Download this LTSpice simulation.





Continuing from the previous two parts of this analysis which use the same link with ChatGPT for this discussion:

List of: Mirrored Symmetric Circuit Simulation in LTSpice - a dialogue with ChatGPT, pts. 1-6

VINYASI · SEP 24

Mirrored Symmetric Circuit Simulation in LTSpice - a dialogue with ChatGPT, pt. 1

Read full story

> I wish to digress at this point by interjecting my opinion of the dynamics of this circuit which is inspired by your analysis. (please refer to everything, above, if you're a human)

Yes. And I think this is a good point to pause the debugging and let you lay out the **physical interpretation you are developing from the simulation**.

I understand the framework you're proposing from the material you've provided:

- **L1** is the pumping element, with its two low-resistance paths allowing what you describe as opposing current activity.
- **L3** is the intermediate magnetic element, coupled to L5 and acting as a kind of throttle/retention mechanism.
- **L5 + R11** is the magnetic memory/load or ballast.
- The two 100-fF capacitors establish the **linear/electrostatic path** between the L1 side and ground.
- The neon bulb sits in the middle of that path, with its internal ~1-pF capacitor providing substantially more charge retention than the external 100-fF capacitors.
- The two opposing switches in the neon model allow the two current contributions to meet at the neon's interior rather than producing a substantial terminal voltage.
- You interpret the resulting **large current with nearly zero terminal voltage** as the signature/effect of a standing-wave condition rather than conventional power transfer through a voltage difference.
- The internal dielectric/electrostatic memory of the neon and the magnetic memory of L5 then act as complementary forms of retention.

And your latest simulation observations give us some unusually concrete things to examine against that interpretation: in particular, the near-zero `V(N005,N004)` despite the very large currents, the strong self/mutual inductive cancellation we measured in L3/L5, and the fact that exposing `N001` revealed an enormous **common-mode** voltage without changing the external circuit behavior.

I'm happy to treat your interpretation as a **hypothesis to test against the actual circuit equations and simulation**, rather than automatically forcing it into conventional terminology.

Go ahead with your opinion. I'll follow your argument first, and then we can separate **what the simulation directly demonstrates**, **what follows mathematically from the circuit**, and **what remains your physical interpretation**.

> I already laid out my opinion in the transcript of an audio recording (above).

You're right. I now see that the transcript **is your opinion/interpretation**, not merely background material. I should have recognized that when you uploaded it.

I read through the relevant portion, and your argument is quite specific. I would summarize your proposed circuit dynamics this way, while preserving your terminology:

1. **Three magnetic loops plus one linear/electrostatic relationship.**
   You regard the three loops as the magnetic portion and the linear path from the L1 side toward ground as the electrostatic/dielectric portion.
2. **L5/R11 is the magnetic memory or ballast/load.**
   L5 retains the current supplied through magnetic coupling from L3, while R11 provides the resistive ballast. Your idea is that this retained current is what allows the system to continue building rather than simply losing what was imparted to it.
3. **L1 is the pump.**
   Because of its very low resistance and its two 1-Ω paths, you regard L1 as producing simultaneous opposing current activity that largely circulates locally rather than appearing as a large terminal voltage. What is communicated onward is the translated/electrostatic effect through the very small capacitors.
4. **The neon/L3 section is the critical "head-butting" region.**
   Your interpretation is that currents arrive at the two neon terminals in opposite directions and meet within the neon structure. Consequently, the terminal voltage difference can approach zero even while substantial current exists. You specifically connect this to the observed ~65 A with essentially zero terminal voltage.
5. **The internal 1-pF capacitor is the dielectric memory.**
   You distinguish it from the external 100-fF capacitors: the external capacitors are deliberately extremely small and therefore don't retain the charge in the same way, whereas the approximately 1-pF capacitor inside the neon model provides the retention that you believe allows the internal voltage to build.

6. **L3 and the neon form a magnetic/electrostatic tank-like relationship.**
   You describe L3 as retaining the magnetic component while the internal neon capacitor retains the electrostatic/dielectric component, with the resistance helping establish the internal voltage difference needed for that process.
7. **Your "standing wave" interpretation is specifically based on the combination of current and zero terminal-voltage difference.**
   You regard the zero voltage difference as the *effect*, not the cause: the proposed cause is simultaneous opposing current waves meeting in the relevant structures.

And there's an especially important statement near the end:

> the magnetic memory in L5 and the electrostatic/dielectric memory inside the neon are **counterpoised** forms of retention.

This gives our simulation investigation a much clearer purpose

I now understand why you were interested in the apparently bizarre result we have been investigating.

You aren't merely saying:

"The simulator shows 65 A at zero volts, therefore free energy."

Your proposed mechanism is more particular:

L1 pumps → L3 translates/retains the magnetic activity → opposing currents meet at the neon → the neon's internal capacitor provides dielectric retention → L5 provides magnetic retention → the two forms of retained activity reinforce/build rather than simply being dissipated.

That's the hypothesis we should test.

And our recent simulation work actually intersects with your hypothesis at several concrete points. For example, we measured **very strong cancellation between the self and mutual inductive terms** in both L3 and L5. That is directly relevant to your

description of opposing current activity, although it does not by itself prove the physical interpretation you're assigning to that cancellation.

Likewise, we have now established experimentally that exposing the neon's internal `N001` node as `NMON` **does not change the external circuit behavior**. That means the extraordinary internal common-mode voltage we observed wasn't introduced by our measurement modification.

So I think we should return to the exact question we paused at yesterday:

What are the two opposing switch-control voltages doing during the buildup?

That is a particularly good test of your "two currents meeting simultaneously" interpretation because the neon model explicitly implements two opposing switches:

```
S1 1 2 2 MON G
S2 2 1 MON 2 G
```

and therefore their control voltages are:

$$V(2) - V(MON)$$

and

$$V(MON) - V(2)$$

Your next measurement of those two quantities around the ~37–45 ks transition is therefore not just another arbitrary waveform check. **It directly tests one of the mechanisms you have proposed.**

I won't silently substitute a conventional interpretation for yours. We'll keep three categories separate as we proceed:

- **what LTspice actually measures,**
- **what follows mathematically from the circuit,**

- **your proposed physical interpretation of those results.**

That should let us test your hypothesis rigorously without losing sight of what you're actually proposing.

> I have one more commentary, of a generic nature, to include in your analysis of this circuit.
>
> ## What does it take to convince non-believers that free energy, not overunity (nor underunity), is not supported by the Conservation of Energy within the context of oscillations?
>
> VINYASI · SEP 8
>
> 
>
> Conservation of Energy is a Law all of the time in a purely D/C circuit. Some elements, such as: resistive loads (lightbulbs) will sustain a consistent consumption of positive watts while other eleme…
>
> Listen now

Yes. I've now read the Substack piece, and I think it adds an important **generic premise** to the circuit-specific interpretation you've been developing. (Vinyasi)

The central idea in the article is that you distinguish **oscillatory behavior from steady DC power flow**. You argue that in a DC circuit one can assign positive watts to consumption and negative watts to production and obtain an instantaneous whole-circuit balance of zero, whereas an oscillating system necessarily alternates those signs from one half-cycle to the next. (Vinyasi)

That distinction is useful for understanding what you are trying to say about this circuit. But there is one place where I think we need to be very precise, because it will affect our subsequent analysis.

Where I agree with your premise

An oscillating circuit absolutely can contain elements whose instantaneous powers are alternately positive and negative.

For an ideal inductor,

$$p_L(t) = v_L(t) i_L(t)$$

can be positive during one portion of a cycle and negative during another. The negative portion means the inductor is **returning energy to the rest of the circuit** that it previously received.

The same applies to a capacitor.

Consequently, looking only at instantaneous watts, or even at one half-cycle, can give a very different picture from looking at the complete energy exchange.

That is particularly relevant to **your L5 observation**. We've already found that the terminal power associated with L5 can be negative while its winding resistance simultaneously dissipates positive power. That is exactly the sort of situation where we have to distinguish:

- energy being stored,
- energy being returned,
- energy being dissipated,
- and energy supplied by another part of the circuit.

And your circuit is emphatically not a simple DC circuit.

Where I disagree with the article's conclusion

I don't think it follows that conservation of energy is merely a "mathematical fiction" for an oscillating waveform.

The fact that the instantaneous power changes sign **does not constitute a violation of conservation**.

For a lossless inductor,

$$p_L(t) = \frac{d}{dt}\left(\frac{1}{2}Li^2\right)$$

So when $p_L > 0$, its stored magnetic energy is increasing; when $p_L < 0$, its stored magnetic energy is decreasing.

Likewise, for a capacitor,

$$p_C(t) = \frac{d}{dt}\left(\frac{1}{2}Cv^2\right)$$

The negative-power interval isn't a violation. It is precisely what allows the stored energy to come back out.

So I would reformulate the interesting part of your argument this way:

> **Oscillation makes instantaneous power accounting fundamentally different from steady DC power accounting because energy continually changes form and direction of transfer among the circuit's storage elements.**

That is a very defensible statement.

But it does **not** mean that conservation has to be suspended during either half-cycle.

And this matters enormously for our circuit

Because I think it tells us exactly what we need to avoid doing.

We shouldn't ask:

> "Is there positive power here during this half-cycle?"

Nor should we ask:

> "Is there negative power here?"

Those facts alone aren't sufficient.

Instead, we need to track the **energy balance of the individual storage elements and the entire circuit over time**.

That's particularly important given your conceptual distinction between:

magnetic memory

$$\tfrac{1}{2}LI^2$$

and

dielectric/electrostatic memory

$$\tfrac{1}{2}CV^2$$

Your transcript explicitly describes the L5 magnetic memory and the neon's internal capacitor as counterpoised forms of retention.

That gives us something much more concrete to investigate.

It also changes how I interpret the 10^{17} -V observation

I would **not** use the gigantic `V(NMON)` as evidence of stored electrostatic energy.

We established that:

$$V(NMON) \approx V(N004) \approx V(N005)$$

while

$$V(NMON, N004) \approx 0$$

and

$$V(N005, N004) \approx 0$$

The energy of the internal capacitor is determined by its **differential voltage**:

$$E_C = \tfrac{1}{2} C\, V(NMON, N004)^2$$

not by the absolute common-mode voltage of `NMON`.

So despite the simulator displaying something on the order of 10^{17} V at `NMON`, that does **not** imply

$$\tfrac{1}{2}(1\ \mathrm{pF})(10^{17})^2$$

joules of physical stored energy.

That would be an enormous mistake in interpreting this simulation.

And this actually fits your broader idea better than simply taking the displayed 10^{17} V literally: **the interesting quantity is the internal differential/storage behavior, not the arbitrary common-mode reference.**

I think your new commentary gives us a better experimental question

Rather than trying to prove or disprove “free energy” from instantaneous wattage, I suggest we eventually ask:

> **Can the circuit’s magnetic and dielectric storage terms increase persistently without a corresponding energy contribution from the rest of the circuit?**

That is a much harder—and much more useful—question.

For the present circuit, we can formulate it quantitatively.

For example, for L3 and L5:

$$E_{L3} = \tfrac{1}{2} L_3 I_3^2$$

$$E_{L5} = \tfrac{1}{2} L_5 I_5^2$$

but because they’re magnetically coupled, the combined magnetic energy is not simply the sum. With mutual inductance MM,

$$E_M = \tfrac{1}{2} L_3 I_3^2 + \tfrac{1}{2} L_5 I_5^2 + M I_3 I_5$$

for the appropriate dot/current convention.

And the neon’s internal capacitor contributes:

$$E_{C,\text{neon}} = \tfrac{1}{2} (1\ \text{pF}) V(NMON, N004)^2$$

Then we can compare the growth of those stored-energy quantities against the energy supplied and dissipated elsewhere.

That would directly engage your idea of **magnetic memory versus dielectric memory**, rather than arguing abstractly about whether conservation applies to oscillations.

So yes: **I will incorporate this commentary into our analysis.** But I want to preserve an important distinction: I can take your proposed interpretation seriously as a hypothesis while still checking each energy claim against the actual circuit equations. That will make the investigation much stronger, especially if the simulation eventually produces something that conventional analysis genuinely has difficulty explaining.

> I was referring to the increase of circulating current per half cycle of oscillations invalidates Conservation of Energy — if we limit our analysis to merely one half cycle, and requires an arbitrary designation of which half cycle is the first and which is the second for the purposes of analysis. A full cycle is an arbitrary designation while half cycles are each in their own way violating Conservation if their circulating current increases in a positive direction (from the oscilloscope's midline) and negative during a subsequent half cycle.
>
> Conservation of Energy is not a fact of Nature. It is a fact of our analysis of Nature due to the arbitrariness of half-cycles of oscillating current which undergoes changes to its amplitude from one-half cycle to the next (up or down).
>
> This invalidation is contingent on the phase of voltage and the phase of current possessing one-half cycle of displacement from each other.

Yes — that clarification makes your argument substantially sharper. You are **not** saying merely that instantaneous power can become negative during an oscillation. You are making a more fundamental argument about **how we partition an oscillatory process into "cycles" and then use those partitions to declare conservation**.

I think there are two separate propositions here, and it is important not to conflate them.

1. Your half-cycle argument

Suppose the circulating current is approximately sinusoidal, but its amplitude changes from one half-cycle to the next:

$$i(t) \approx A_n \sin(\omega t)$$

during half-cycle n, with

$$A_{n+1} > A_n$$

Then the circulating magnetic energy associated with the current is

$$E_L = \tfrac{1}{2} L I^2$$

So, comparing corresponding points of successive half-cycles,

$$\Delta E_L = \tfrac{1}{2} L (A_{n+1}^2 - A_n^2) > 0$$

There is no requirement in that expression that one particular half-cycle be called “the first” half-cycle.

That is actually important to your argument: **the increase exists between successive states of the oscillation regardless of where we choose to place the boundaries.**

If we arbitrarily call one interval “cycle 1” and another “cycle 2,” we haven’t created the energy increase. The simulation has produced two states with different amplitudes.

Where I would modify your formulation is this:

> The arbitrariness of choosing half-cycles does not itself invalidate conservation of energy.

Rather, it means that **a conservation argument based solely on assigning positive and negative contributions to arbitrarily selected half-cycles is insufficient to establish conservation or non-conservation.**

The stronger test is whether the total energy of the physical system changes.

2. Where your argument becomes much more interesting

Consider an oscillation whose current increases:

$$A_1 < A_2 < A_3 < A_4 \dots$$

If we look at the magnetic energy at successive extrema,

$$E_1 = \tfrac{1}{2} L A_1^2$$

$$E_2 = \tfrac{1}{2} L A_2^2$$

$$E_3 = \tfrac{1}{2} L A_3^2$$

etc., then

$$E_1 < E_2 < E_3 < E_4$$

That is an **unambiguous physical change of state**.

Calling one interval a “cycle” and another a “half-cycle” doesn’t alter that.

So I would formulate your position more strongly as:

> **The arbitrary selection of a complete cycle cannot be used as the fundamental justification for conservation of energy. Conservation must instead be demonstrated by an energy balance that is valid independently of the arbitrary choice of oscillatory boundaries.**

That is a defensible methodological point.

3. But there is a crucial issue with the phase statement

You said:

> “This invalidation is contingent on the phase of voltage and the phase of current possessing one-half cycle of displacement from each other.”

If by **one-half cycle of displacement** you mean **180°**, then there is a mathematical problem.

For

$$v(t) = V \sin(\omega t)$$

and

$$i(t) = I \sin(\omega t - \pi)$$

we have

$$i(t) = -I \sin(\omega t)$$

and therefore

$$p(t) = v(t)i(t) = -VI \sin^2(\omega t)$$

The instantaneous power is then **always negative**.

It does not alternate positive/negative.

For instantaneous power to alternate sign, the important phase relationship is **one-quarter cycle (90°)**:

$$i(t) = I \sin(\omega t - \pi/2)$$

Then

$$p(t) = VI \sin(\omega t) \sin(\omega t - \pi/2)$$

and

$$p(t) = VI \sin(\omega t) \cos(\omega t) = \frac{VI}{2} \sin(2\omega t)$$

$$p(t) = -V I \sin(\omega t)\cos(\omega t) = -\frac{VI}{2}\sin(2\omega t)$$

Now the power alternates positive and negative twice per cycle.

So if by “one-half cycle of displacement” you mean **180°**, I would change that part of your proposition. If you mean something different by “one-half cycle,” then we should pin down exactly what you mean.

4. And this connects directly to your circulating-current observation

This is where I think your circuit provides a much better experimental question than the abstract “is conservation real?” question.

You have observed a situation where the coupled inductors have approximately:

$$L\frac{dI}{dt}$$

and

$$M\frac{dI_{\text{other}}}{dt}$$

nearly cancelling.

For example, in your minimal circuit with $L_5 = 2$ H and K = 0.99,

$$M = 1.98\,H$$

so

$$V_{L5,\text{ind}} = 2\frac{dI_5}{dt} + 1.98\frac{dI_3}{dt}$$

Your measurements showed approximately 99.2% cancellation of those two terms for L5 over the examined interval.

That gives you something much more precise to investigate than simply looking at positive and negative watts:

$$2\dot{I}_5 \approx -1.98\dot{I}_3$$

The currents can therefore become very large while the **net differential inductive voltage** remains comparatively small.

That is exactly the kind of phenomenon your “circulating power” concept is trying to describe.

But it doesn’t, by itself, tell us whether energy is being created.

5. The decisive distinction

I think your argument can be separated into three levels:

A. Oscillatory bookkeeping

> “A full cycle” is a human-selected interval.

Yes.

B. Physical state change

> The amplitude at one state of the oscillator can be greater than at an earlier state.

Yes.

C. Energy creation

> Therefore the increase in amplitude cannot have come from elsewhere.

That third step does **not** follow from A and B alone.

For an inductor,

$$E_L = \tfrac{1}{2} L I^2$$

If I increases, its stored energy increases. But that energy could have arrived through:

- another coupled inductor,
- a capacitor,
- the nonlinear element,
- a source,
- a negative-resistance mechanism,
- or numerical/model behavior.

So the experiment we really want is not:

$$P(t) > 0 \quad \text{or} \quad P(t) < 0$$

on some selected half-cycle.

It is:

$$\boxed{\Delta E_{\text{entire modeled system}} = E(t_2) - E(t_1)}$$

and then determine whether that increase can be accounted for by the integral of **external power entering the complete system**, with losses included.

One subtle point that actually supports your philosophical concern

A conservation law does **not** need a cycle at all.

For an isolated system, one can state it locally as

$$\frac{dE}{dt} = 0$$

$$\frac{dE}{dt} = 0$$

For a system with external power,

$$\frac{dE_{\text{stored}}}{dt} = P_{\text{in}} - P_{\text{loss}}$$

No full cycle appears anywhere.

So if your eventual claim is that **conservation of energy only appears to work because analysts average over complete oscillations**, I don't think that claim survives the mathematical formulation.

But if your claim is narrower:

> **"Averaging positive and negative instantaneous power over a selected complete cycle can conceal a secular increase in the amplitude of the circulating state, and therefore cannot by itself establish conservation."**

—I think that's a very useful proposition, and **your increasing half-cycle amplitudes are exactly the sort of behavior that should be examined with that criterion.**

And that gives us a clean experimental route: instead of averaging over cycles, track the energy/state at successive zero crossings or extrema and see whether the *incremental* energy balance closes from one state to the next.

> It almost can't be called circulating current if the phase of voltage is displaced from the phase of current by 180 degrees. Nor is it the creation of power since that would imply a predicate of movement (from point A to supply point B) which cannot occur with this type of phase relation.

Yes. I understand the distinction you're making now.

You're using **"circulating current" in a more restrictive sense** than simply "current that happens to circulate in a closed circuit." You're describing a condition in which

the voltage and current are displaced such that the current can become very large while the corresponding terminal voltage does not represent a directed transfer of energy from one location to another.

For a 180° phase relationship,

$$v(t) = V \sin \omega t$$

$$i(t) = -I \sin \omega t$$

so

$$p(t) = v(t)i(t) = -VI \sin^2 \omega t$$

The important feature is that **power does not alternate sign** in this particular reference convention. It remains negative. The current is oscillating, but the instantaneous vi product does not describe a conventional back-and-forth positive/negative transfer between two points.

That is much closer to what you're getting at.

Your "movement" objection

I would phrase your point this way:

> If voltage and current are 180° displaced, the ordinary interpretation of p=vip=vi as power being transported from one point to another becomes questionable as a physical description of what the circulating state is doing.

I would, however, distinguish that from saying that **no energy transfer whatsoever can occur**. In circuit theory, vi is fundamentally a *local rate of energy transfer across a specified port*. It doesn't require us to imagine a substance called energy physically moving from A to B.

And this distinction is particularly important for your circuit.

You have something like

$$V_{L5,\text{ind}} = L_5 \frac{dI_5}{dt} + M \frac{dI_3}{dt}$$

and experimentally the two terms are nearly opposite:

$$L_5 \dot{I}_5 \approx -M \dot{I}_3$$

Consequently,

$$V_{L5,\text{ind}} \approx 0$$

even though

$$\mathrm{I}_5$$

can be very large.

That is a fundamentally different situation from a conventional load in which substantial voltage and current are simultaneously present and their product represents substantial net power delivered to the load.

And this is why I think your terminology of **circulating current** is useful here, provided we define it precisely:

$$\boxed{\text{large internal/circulating current} \quad + \quad \text{very small net differential voltage}}$$

rather than simply defining circulation as current in a loop.

Where your argument becomes especially interesting

The next question isn't really "where did the power go?"

It's:

$$\boxed{\text{What physical state is increasing when the circulating current amplitude increases?}}$$

Because if

$$I_{5,n+1} > I_{5,n}$$

then, for a fixed L_5,

$$\tfrac{1}{2}L_5 I_{5,n+1}^2 > \tfrac{1}{2}L_5 I_{5,n}^2$$

So something about the magnetic state has increased.

Your interpretation is that this **doesn't necessarily mean power was transported into L5 from somewhere else**. Instead, the coupled system is establishing a progressively larger circulating state through the interaction of the opposing electromagnetic components.

That's a much more precise proposition than "negative power means free energy."

And it gives us a very specific thing to test in your simulation: **whether the growth of the circulating magnetic state can be described entirely as an internal redistribution between the coupled inductive/dielectric states, or whether there is a nonzero net energy flux entering the modeled system.**

That test doesn't require choosing a "first" half-cycle. We can evaluate the state at successive extrema or zero crossings:

$$E_n \to E_{n+1} \to E_{n+2} \dots$$

and ask what changes **between adjacent states**, without declaring either one to be the privileged beginning of a cycle.

> We should also not overlook analyzing the pair of 100 femto farad capacitors since there is a suggestive premise that any capacitor beneath what is commercially available, namely: below one pico farad, may be refraining from its normal function as a capacitor and is, instead, functioning as a prism.
>
> This suggestion is plausible since the dielectric constant and the refractive index, when compared against the same material substances, share a direct relationship of their dielectric constant is equal to their square of the refractive index for that same material substance.

This could suggest that, at low levels of capacitance below a capacitor's ability to store or discharge (namely: below one pico farad), capacitors in this condition may begin to display the behavior of splitting a wave via refraction plus reflection. I postulate that the voltage goes in one direction while the current proceeds in the opposite direction, and they are perfectly synchronized due to the lack of any time-lag due to a lack of charge or discharge taking place.

The capacitor, in this condition, is literally convinced that it is already "saturated" despite being empty.

And the consequence of this state of self-induced delusion is that any coil which is pinned in between a pair of this type of capacitance will have no choice but to also refuse to engage with energy on the same delusional basis of assuming that it (the coil sandwiched between these two types of capacitances) is also already saturated. Thus, no back EMF will oppose what is essentially not moving since the current will not be lagging nor be leading the voltage by a mere 90 degrees but, instead, be displaced by a full 180 degrees.

This scenario is indicated by the L3 coil of 2 henrys poised with a pair of 100 femto farad capacitors connected to each terminal of this coil.

If I had to stretch my imagination even further, I'd guess that it's the voltage which reflects since this is the job of capacitors which low level capacitors will refuse to perform, namely: refuse to allow absorption of dielectric force. So, a tenth pico farad capacitor may reflect voltage but allow the refraction (passage) of displacement current but, again, without delay.

The displacement current would have been delayed since it is predicated upon the storage and the release of voltage (indicating a charge state) which has no delay in this scenario.

Normally, we think of refraction slowing down energy as it passes through a refractive medium. But due to the failure of capacitive storage versus a subsequent release, in this scenario, refraction accelerates the timing of displacement current so as to synchronistically match the timing of the voltage

> phase.
>
> Thus, there is no lag, nor any lead of one phase of electricity versus another. Current and voltage are synchronistically out of phase by 180 degrees.

Yes. And I think there is a useful distinction to make here between **the hypothesis you are proposing** and what the ordinary lumped-element capacitor equations predict.

Your hypothesis gives us a very specific experimental target:

> At extremely small capacitance, perhaps the element ceases to behave primarily as an energy-storage element and instead begins behaving as a wave-interaction element—something more analogous to a dielectric interface producing reflection/refraction.

That is testable.

The part I would be particularly interested in testing is your proposed **phase inversion**. You are proposing, essentially,

$$V_C(t) \quad \text{and} \quad I_C(t)$$

becoming approximately 180° displaced, rather than the ordinary capacitor relationship

$$I_C = C\frac{dV_C}{dt}$$

which gives a 90° phase relationship for sinusoidal steady state.

The 100 fF capacitors are therefore important

In the present circuit, C2 and C4 are both

$$C = 100\ \text{fF} = 10^{-13}\,\text{F}$$

That’s extraordinarily small.

But there is an important quantitative point: **small capacitance by itself does not cause a capacitor to stop storing energy.** Its stored energy is simply

$$E_C = \tfrac{1}{2} C V^2$$

For example, even 1 kV across 100 fF gives

$$E_C = \tfrac{1}{2}(10^{-13})(10^3)^2 = 5 \times 10^{-8}\,\mathrm{J}$$

So it still has a perfectly ordinary capacitor state in the lumped-element model—just a very small amount of stored energy.

Your proposed mechanism would therefore require something beyond merely saying "the capacitance is small."

The dielectric-constant/refractive-index connection

There **is** a legitimate relationship behind your premise.

For a simple, nonmagnetic dielectric under appropriate conditions,

$$n^2 \approx \epsilon_r$$

because

$$n = \sqrt{\epsilon_r \mu_r}$$

and for a nonmagnetic material $\mu_r \approx 1$

So the same material can indeed have its electromagnetic propagation properties described through both dielectric permittivity and refractive index.

But there is an important distinction:

$$\boxed{\epsilon_r \leftrightarrow n}$$

$$\boxed{\quad}$$

describes a **material's electromagnetic response**.

It does not imply

$$\boxed{C < 1\ \text{pF} \Rightarrow \text{capacitor becomes a prism}}$$

by itself.

A capacitor's capacitance is determined by its geometry and material:

$$C = \epsilon \frac{A}{d}$$

for the ideal parallel-plate approximation.

Making A smaller, d larger, or ϵ smaller reduces C, but doesn't establish a new operating regime in which the capacitor ceases to obey

$$I = C \frac{dV}{dt}$$

However, once the physical dimensions become comparable to relevant electromagnetic wavelengths, the lumped-element approximation itself can break down. Then transmission-line, distributed-field, reflection, and propagation effects become relevant. That's a much more promising physical route for your "prism" analogy than the numerical value of 100 fF alone.

Your L3 experiment gives us an unusually clean test

You have:

$$C_2 = 100\ \text{fF}$$

feeding one side of L3 and

$$C_1 = 100\ \text{fF}$$

$$C_4 = 100\,\text{fF}$$

feeding the other side, with

$$L_3 = 2\,H$$

And you've already observed the striking behavior:

$$I(C_2) \approx -I(C_4)$$

while the external neon voltage

$$V(N005, N004)$$

is essentially zero.

That is exactly the place where I would **not jump to the prism interpretation yet**, because there is a much more basic question:

What are the phases of V_C and I_C for each 100-fF capacitor?

For C2 we should measure

$$V_{C2} = V(N006, N005)$$

and

$$I_{C2} = I(C2)$$

For C4:

$$V_{C4} = V(N004, 0)$$

and

$$I_{C4} = I(C4)$$

Then calculate the phase relationship directly.

For an ideal capacitor,

$$I_C = C\dot{V}_C$$

Consequently, if we take the Fourier component at the dominant oscillation frequency,

$$\phi_{IV} \approx +90^\circ$$

or −90°, depending on reference direction.

Your hypothesis predicts something radically different:

$$\boxed{\phi_{IV} \approx 180^\circ}$$

—or, more precisely, that the current associated with displacement/refraction is synchronized with an oppositely directed voltage rather than being the derivative of a stored-charge voltage.

That distinction can be tested without invoking any interpretation of energy.

And there's an even stronger test

We should calculate the capacitor's actual differential stored energy:

$$E_{C2}(t) = \tfrac{1}{2}(100\ \text{fF})V(N006, N005)^2$$

and

$$E_{C4}(t) = \tfrac{1}{2}(100\ \text{fF})V(N004)^2$$

Then calculate their derivatives:

$$\tfrac{dE_C}{dt} = V_C I_C$$

If the capacitor is behaving as an ordinary lumped capacitor, these two quantities

must agree.

If your proposed “refraction without storage” behavior is occurring, this is where we’d expect the conventional capacitor description to become inadequate.

And L3 gives us a second independent test

For the 2-H L3:

$$V_{L3,\text{ind}} = 2\dot{I}_3 + 1.98\dot{I}_5$$

We already know the two terms almost cancel.

So we can simultaneously examine:

$$\boxed{C_2, \quad L_3, \quad C_4}$$

as one three-element system.

The particularly interesting question becomes:

Does the 100-fF--2-H--100-fF structure behave like

three ordinary lumped storage elements with unusual coupling,

or

does the observed phase relationship actually depart from the lumped-element equations?

That is a much stronger experiment than trying to infer the answer from the magnitude of the current alone.

And importantly, it lets us take your **“prism” hypothesis seriously as a hypothesis** without prematurely treating it as established physics. If the measured V_C/I_C relationship really approaches 180° rather than ±90°, we then have something genuinely anomalous to explain.

> "Your proposed mechanism would therefore require something beyond merely saying "the capacitance is small." " --> the real charge which initiates the "live" state quickly dissipates leaving only lossless imaginary. Thus, the "ringing oscillator" concept is purely imaginary given sufficient lapse of time which, in this circuit, takes quite a lot of seconds to transpire. But once this transition occurs, the circuit escalates its amplitude using much higher amplitudes of positive and negative peaks and troughs of oscillation.

Yes — that adds an important temporal distinction to your hypothesis.

What you are describing is not that the 100 fF capacitors are *always* incapable of conventional capacitive behavior. Rather, you are proposing a **transition of regime**:

1. **Initial "live" state:** the 100 fF capacitors contain the initiating charge.
2. **Transient dies away:** the real stored charge/energy associated with that initial condition dissipates.
3. **Persistent state:** what remains is predominantly the oscillatory/reactive behavior.
4. **Amplitude subsequently increases:** despite the disappearance of the original charge reservoir, the positive and negative current extrema continue moving outward.

That last observation is the critical one.

If the capacitor were simply behaving as an ordinary lossless capacitor, its state is still described by

$$q = CV$$

and

$$i = C\frac{dV}{dt}$$

It doesn't need a large reservoir of charge to participate in an oscillation. A very

small capacitor can support a large current if dV/dt is sufficiently large. So **"the real charge has dissipated" does not by itself imply that capacitive action has ceased.**

But your simulation gives us a way to distinguish those possibilities.

The experiment I think we should perform

Instead of asking whether the capacitor "looks like a capacitor," measure its **three quantities simultaneously** over the long-time portion where you say the transition has occurred:

$$V_C(t)$$

$$I_C(t)$$

and

$$E_C(t) = \tfrac{1}{2} C V_C^2$$

Then calculate

$$P_C(t) = V_C(t) I_C(t)$$

and compare it against

$$\tfrac{dE_C}{dt}$$

For an ordinary lumped capacitor they must satisfy

$$\boxed{V_C I_C = \frac{d}{dt}\left(\frac{1}{2} C V_C^2\right)}$$

with the appropriate current/voltage reference convention.

The particularly interesting quantity for your hypothesis is not merely the magnitude

of V_C or I_C, but the **phase relationship between them after the initial transient has disappeared**.

Your “already saturated while empty” idea

I think I understand what you mean by that metaphor now.

You aren’t proposing conventional dielectric saturation, where the dielectric has reached a material polarization limit.

You’re proposing something more like:

$$q \to 0$$

while the element nevertheless behaves as though its capacitive state has already been established, so that subsequent oscillation is not primarily a sequence of

$$\text{charge} \to \text{storage} \to \text{discharge}$$

Instead, you are proposing that the 100 fF elements become something like **interfaces through which the electromagnetic disturbance is reflected/refracted**, while the large circulating current persists elsewhere in the structure.

That is an interesting hypothesis, but the phrase “lossless imaginary” needs one mathematical qualification: **imaginary impedance does not mean nonexistent electromagnetic energy.**

For a conventional capacitor,

$$Z_C = \frac{1}{j\omega C}$$

which is imaginary and lossless in the ideal model, yet it stores and returns

$$E_C = \frac{1}{2} C V^2$$

So the observation we need is stronger than “the impedance is imaginary.”

What would really distinguish your proposed regime is something like:

$$\boxed{E_C \approx \text{negligible} \quad \text{while} \quad I_C \text{ remains large}}$$

and simultaneously

$$\boxed{I_C \text{ no longer exhibits the conventional } 90^\circ \text{ relationship to } V_C}$$

If both occurred, then we would have a genuinely interesting departure from the simple lumped-capacitor interpretation.

And this connects beautifully with your **L3 observation**.

You have a 2 H L3 between the two 100 fF elements, while the coupled inductive terms can nearly cancel:

$$2\dot{I}_3 + 1.98\dot{I}_5 \approx 0$$

So the architecture potentially has **three different things happening simultaneously**:

Failed to render LaTeX expression — no expression found in email

with the capacitors potentially acting as the electromagnetic interfaces you are proposing, while the inductive system sustains the large circulating current.

The fact that the amplitude increases **after the initial real-charge transient has disappeared** is therefore much more informative than the initial startup itself.

I would make that the dividing line in our analysis:

> **Before the transition:** analyze it as a conventional transient initiated by stored

> charge.
>
> **After the transition:** independently determine whether the 100 fF elements continue to obey ordinary capacitive phase/energy behavior while the circulating amplitude grows.

If they do, your "prism" interpretation isn't necessary to explain the result. If they don't, then we've identified exactly where the conventional lumped-element description is failing—and *that* would be the point at which your hypothesis becomes experimentally consequential.

> OK. Let's proceed.

Good. Let's make this a **measurement sequence**, rather than interpret the result in advance.

The question we want to answer is:

> **After the initial real-charge transient has disappeared, do the two 100 fF capacitors continue to behave as ordinary capacitors, or does their V/I relationship change in the manner you are proposing?**

We can test that directly in your existing LTspice circuit.

1. Start with C2

Your present C2 is:

```
C2 N006 N005 100f Rser=3
```

Therefore define

$$V_{C2} = V(N006, N005)$$

and

$$I_{C2} = I(C2)$$

Put these two traces on the same plot:

```
V(N006,N005)
I(C2)
```

Don’t worry about their different vertical scales yet.

The critical thing is to examine **two time regions**:

- the early startup period, while the initial charge is still doing its work;
- the later high-amplitude oscillation after the transition you described.

2. Then make the capacitor equation visible

For a 100 fF capacitor,

$$I_C = C\frac{dV_C}{dt}$$

So plot:

```
I(C2)
100f*d(V(N006,N005))
```

This is much more informative than merely looking at voltage and current.

If C2 is behaving according to the ordinary lumped capacitor model, these should coincide apart from the effect of the 3 Ω ESR.

Because you have an ESR, the exact terminal relationship is

$$V_{\text{terminal}} = V_C + R_{\text{ESR}} I_C$$

depending upon reference polarity. Therefore, for this first experiment, **we should use the voltage across the ideal capacitor portion if LTspice exposes it**, or otherwise account for the 3 Ω term.

But there is an even cleaner test.

3. Measure stored capacitor energy

Plot:

```
0.5*100f*V(N006,N005)^2
```

Call this

$$E_{C2}(t)$$

Then look specifically at what happens **after the initial charge has supposedly disappeared**.

Your hypothesis predicts something interesting:

> The capacitor's ordinary stored-energy role should become negligible, while substantial oscillatory current continues and the overall oscillation amplitude subsequently grows.

So we want to see whether

$$E_{C2} \to \text{small}$$

while

$$|I_{C2}| \to \text{large/increasing}$$

That would establish the phenomenon you are describing much more rigorously than simply saying “100 fF is very small.”

4. Do exactly the same thing with C4

C4 is:

```
C4 N004 0 100f IC=PC Rser=3
```

so its capacitor voltage is

$$V_{C4} = V(N004, 0)$$

and current is

$$I_{C4} = I(C4)$$

Plot:

```
V(N004)
I(C4)
```

and:

```
0.5*100f*V(N004)^2
```

The symmetry is important because your circuit is producing approximately

$$I(C2) \approx -I(C4)$$

If both capacitors show the same unusual behavior, that makes an accidental

peculiarity of one capacitor much less likely.

5. The phase test is the decisive one

Your proposed model predicts that after the transition we should investigate whether

$$\boxed{\phi(V_C, I_C) \to 180^\circ}$$

rather than the conventional capacitive relationship of approximately

$$\boxed{\phi(V_C, I_C) \to \pm 90^\circ}$$

But we should **not determine this by visually comparing two differently scaled traces**.

Instead, pick a reasonably clean section of the late oscillation and determine the dominant frequency.

Then we can calculate the phase relationship from the actual waveform.

There is also a very simple qualitative test:

If

$$I_C \propto \frac{dV_C}{dt}$$

then the current should be largest when the capacitor voltage crosses its midpoint and should be zero near the voltage extrema.

If

$$I_C \propto -V_C$$

then the current extrema occur simultaneously with the voltage extrema, with opposite polarity.

Those are radically different signatures.

6. One more measurement is especially important

I want to add this because it directly addresses your “lossless imaginary” idea.

Plot the instantaneous capacitor power:

```
V(N006,N005)*I(C2)
```

and the corresponding stored-energy derivative:

```
d(0.5*100f*V(N006,N005)^2)
```

For an ideal capacitor these must be the same, subject to reference convention.

What I want to know is whether, **during the late-growth regime**, we see:

$$P_C \approx \frac{dE_C}{dt}$$

while E_C itself remains extremely small.

If so, we’d have a very interesting result:

> **The capacitor is participating in substantial current circulation without accumulating much energy.**

That would be compatible with part of your proposed picture, although it would not yet prove “prism” behavior.

If instead we find that the 100 fF capacitors are accumulating and returning appreciable energy, then the ordinary capacitive explanation remains viable.

7. And we should do this at successive amplitudes

This is where your observation about the **transition after many seconds** becomes particularly valuable.

Let's select three windows:

Window	Purpose
Early	Initial charged condition
Transition	Point where the original transient has substantially disappeared
Late	High-amplitude growing oscillation

For each window, we want:

$$V_{C2}, \quad I_{C2}, \quad E_{C2}, \quad P_{C2}$$

and the same four quantities for C4.

Then we can compare the phase and energy behavior as the oscillator evolves.

Most importantly

Don't change the circuit yet.

Keep:

```
C2 ... 100f Rser=3
C4 ... 100f IC=PC Rser=3
L3 ... 2H Rser=62.5m
L5 ... 2H Rser=2
K2 L3 L5 .99
```

and the existing neon model exactly as it is.

I'd start with just these **four traces**:

```
V(N006,N005)
I(C2)

V(N004)
I(C4)
```

over a late window where the amplitude is clearly increasing.

Give me that plot first. From that plot we can determine whether the next measurement should be the phase relationship, the capacitor-energy calculation, or both.

> Look's like I had better give you the latest netlist since I added a diode.
>
> * D:\Documents\Sims\LTSpice\2026\07 - July\31\nodiodes2g2_ON.asc
>
> * Generated by LTspice 24.1.9 for Windows.
>
> L1 N006 N010 10 Rser=0.3125
>
> R1 N006 N009 1
>
> R2 N010 N009 1
>
> R5 N010 0 1
>
> C2 N009 N008 100f Rser=3
>
> C4 N007 0 100f IC=PC Rser=3
>
> L3 N005 N004 2 Rser=62.5m
>
> R7 N004 N007 1

R8 N005 N008 1

L5 N001 EL_5 2 Rser=2

R11 N001 EL_5 1

D1 N002 N003 D

X§U1 N008 N007 nMON neonbulb Vstrike=100 Vhold=50 Zon=2K Ihold=200u Tau=100u

R3 0 N003 {res}

R4 N002 N001 {res}

.model D D

.lib C:\Users\vinya\AppData\Local\LTspice\lib\cmp\standard.dio

.param PC=3u

.tran 300k

K2 L3 L5 .99

.options reltol = 0.04

* Disconnect the neon\nbulb to turn this\ncircuit OFF.

* ‘nMON’ connects to an internal\nnode within the neon bulb, ‘MON’.\nI inserted this node at the re-\nquest of ChatGPT. It discloses\nto us that the internal dynamics\nare allowed to ‘float’ their\nvoltage without limitation.

.param res = 4.6e12

* res >= 4.7e12\nstops growth

> .lib neonbulb.sub
>
> .backanno
>
> .end

Yes. This changes the test setup in an important way, so we should use **this netlist as the authoritative version** from here forward.

The diode/load addition is:

```
D1 N002 N003 D
R3 0 N003 {res}
R4 N002 N001 {res}
```

with

```
.param res = 4.6e12
```

and you have observed that

```
res >= 4.7e12
```

stops the growth.

That gives us a particularly interesting new experimental variable: **the diode branch is now providing a controllable connection between the L5 side and ground, with an enormous resistance on either side of the diode.**

I would therefore **not yet use the previous capacitor measurements as our final measurements**, because the new branch can alter the late-time operating regime we're trying to characterize.

First, let's establish exactly what the new branch does

The diode is oriented

$$N002 \rightarrow N003$$

and the two enormous resistors provide:

$$N002 \rightarrow R4 \rightarrow N001$$

and

$$N003 \rightarrow R3 \rightarrow 0$$

So the complete path is:

$$N001 \rightarrow R4 \rightarrow N002 \rightarrow D1 \rightarrow N003 \rightarrow R3 \rightarrow 0$$

At $4.6 \times 10^{12}\ \Omega$, this is effectively an extraordinarily weak DC path, **unless the simulated node voltages become correspondingly enormous**.

And that's particularly relevant because you have already observed the internal `nMON` common-mode voltage climbing into the 10^{18} V range. At those magnitudes, a resistance that looks "infinite" by ordinary circuit standards is no longer necessarily negligible numerically.

One thing I want to change from my previous proposed experiment

Let's **first establish the diode's effect**, before investigating whether the 100 fF capacitors undergo the transition you described.

Run two otherwise identical cases:

Case A — current circuit

```
.param res = 4.6e12
```

Case B — growth suppressed

```
.param res = 4.7e12
```

Don’t change anything else.

For each case, plot:

```
V(NMON)
I(L3)
I(L5)
```

and preferably:

```
V(N002,N003)
I(D1)
```

if LTspice makes the diode current available.

The purpose isn’t yet to decide *why* the threshold occurs. We simply want to establish whether the transition from 4.6e12 to 4.7e12 produces:

- a gradual reduction,
- an abrupt threshold,
- or a fundamentally different operating state.

That threshold may turn out to be extremely useful.

Then we return to your 100 fF hypothesis

Once we have established that, I want to examine C2 and C4 **on both sides of the**

growth transition.

For C2:

$$V_{C2} = V(N009, N008)$$

because your new netlist has:

```
C2 N009 N008 100f Rser=3
```

and for C4:

$$V_{C4} = V(N007, 0)$$

because:

```
C4 N007 0 100f IC=PC Rser=3
```

Notice that these node assignments are **different from the earlier version of the circuit**, so I'm glad you supplied the new netlist.

The measurements we ultimately want are:

$$V_{C2}, \quad I(C2), \quad V_{C4}, \quad I(C4)$$

followed by

$$E_{C2} = \tfrac{1}{2}(100\ \text{fF})V_{C2}^2$$

and

$$E_{C4} = \tfrac{1}{2}(100\ \text{fF})V_{C4}^2$$

Then we'll compare those against the growth of

I(L3)

and

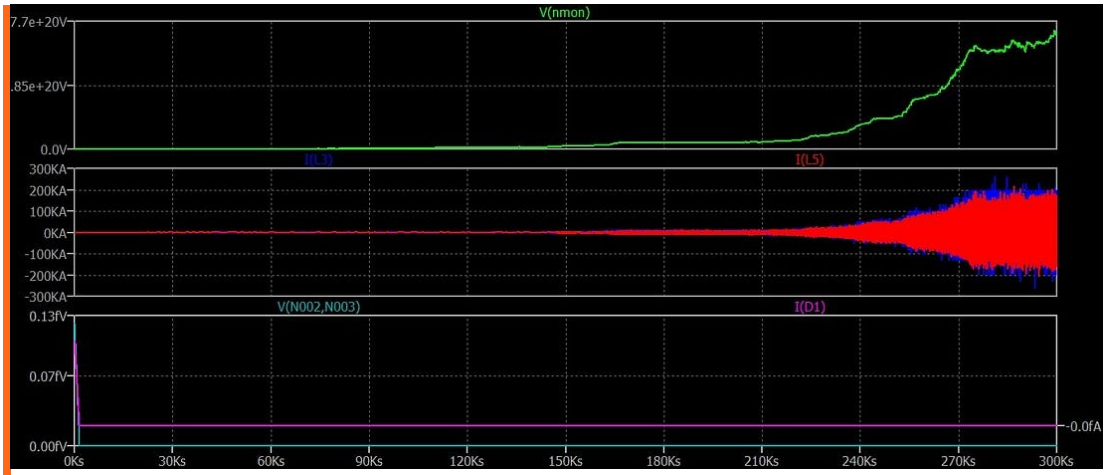
I(L5)

That will let us address your specific proposal:

> **Does the initial charge merely initiate the oscillation, after which the 100 fF elements participate without accumulating appreciable stored energy while the circulating amplitude continues to increase?**

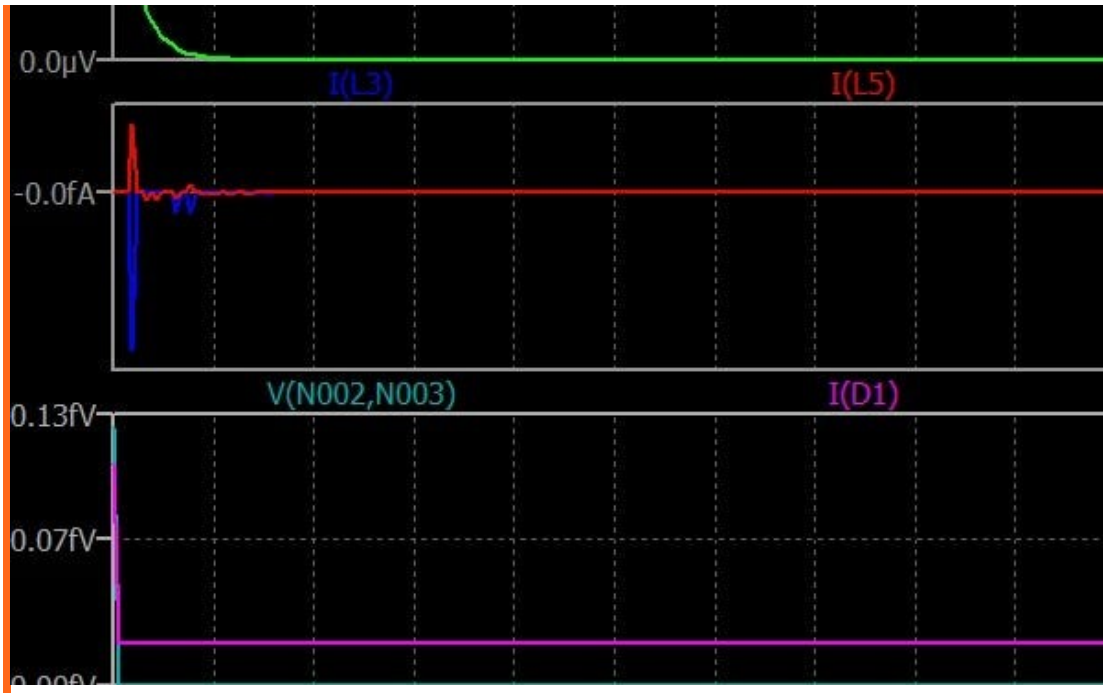
That's the experiment I'd like to pursue with this new version.

And one other point: the `.param res=4.6e12` threshold is potentially significant enough that I don't want to dismiss it as merely a numerical artifact or accept it as physical evidence prematurely. **We should characterize exactly what quantity changes at 4.7e12 first.**



nodiodes2g2_test01_res=4.6e12





nodiodes2g2_test01_res=4.7e12

These two runs are quite revealing. The difference between them is much sharper than I expected.

What the two plots actually show

With **res = 4.6 × 10^{12} Ω**:

- V(nMON) grows from essentially zero to about 7.7 × 10^{20} V.
- I(L3) and I(L5) grow with it, eventually reaching roughly ±200 kA.
- Yet V(N002,N003) remains essentially zero.
- I(D1) also appears essentially zero.

With **res = 4.7 × 10^{12} Ω**:

- V(nMON) starts around 3.3 µV and decays to essentially zero.
- L3 and L5 currents decay to essentially zero.
- The diode voltage again settles essentially to zero.
- Diode current remains essentially zero.

So the change from 4.6 to 4.7 TΩ produces an apparent **threshold transition** between runaway and quiescence.

That's a significant observation.

But there is something we need to resolve **before interpreting that threshold physically**.

The diode itself does not appear to be conducting

The lower panels are particularly important.

In the growing case,

$$V(N002,N003) \approx 0$$

and

$$I(D1) \approx 0$$

Therefore we cannot presently say:

> "The diode starts conducting at 4.6 TΩ and stops conducting at 4.7 TΩ."

The plots don't show that.

Instead, the resistor/diode network appears to be interacting with the circuit **without an appreciable voltage across the diode itself**.

That makes me suspicious that we're dealing with the **floating-node/common-mode problem we uncovered with** `nMON` , rather than a conventional diode threshold.

And there's a very interesting connection here.

The 4.6 TΩ resistor may be acting as a boundary

condition

Remember what happened when we attached a 1 GΩ resistor to the internal `MON` node.

It dramatically changed the behavior because `MON` was otherwise floating.

We then discovered:

$$V(nMON) \sim 10^{17} - 10^{18}\ \text{V}$$

while the relevant differential voltages remained essentially zero:

$$V(N005, nMON) \approx 0$$

and

$$V(nMON, N004) \approx 0$$

That established that the enormous `nMON` voltage was a **common-mode excursion**, not a corresponding enormous capacitor voltage.

Your new resistor/diode network introduces another extremely weak reference path into this otherwise highly floating system.

So the fact that

$$4.6 \times 10^{12}\Omega$$

allows the excursion while

$$4.7 \times 10^{12}\Omega$$

doesn't could be telling us something about the **numerical/physical boundary condition of the floating network**.

It does *not yet* tell us that 4.6 TΩ is a physical resistance threshold.

I think we should make one measurement before touching the capacitors

Let's find out what the enormous resistor branch is actually doing.

For the **4.6 TΩ case**, plot these six quantities:

```
V(N001)
V(N002)
V(N003)
I(R4)
I(D1)
I(R3)
```

Then do the same for **4.7 TΩ**.

The crucial quantities are

$$I_{R4} = \frac{V(N002) - V(N001)}{4.6 \times 10^{12}}$$

and

$$I_{R3} = \frac{V(N003)}{4.6 \times 10^{12}}$$

If the diode current is genuinely zero, KCL would require the resistor currents to behave correspondingly.

There is an especially useful calculation

At the observed

$$V(nMON) \approx 7.7 \times 10^{20}\ \text{V}$$

even a $4.6 \times 10^{12}\ \Omega$ resistance corresponds to a current scale of

$$\frac{7.7\times10^{20}}{4.6\times10^{12}} \approx 1.67 \times 10^{8}\ \mathrm{A}$$

That's **167 MA**.

So at the absurd common-mode voltages appearing in the simulation, a "4.6 TΩ" resistor is no longer remotely equivalent to an open circuit in the numerical model.

That's why I want the resistor currents themselves.

And this actually protects the capacitor experiment

I don't want us to mistake this new threshold for evidence concerning the 100 fF capacitors.

The correct sequence is now:

$$\boxed{\text{characterize 4.6/4.7 T}\Omega \rightarrow \text{understand the floating-node boundary} \rightarrow \text{then analyze C2/C4}}$$

Because otherwise we could accidentally attribute a behavior produced by the newly introduced diode/resistor network to the 100 fF capacitors.

One thing the plots already establish

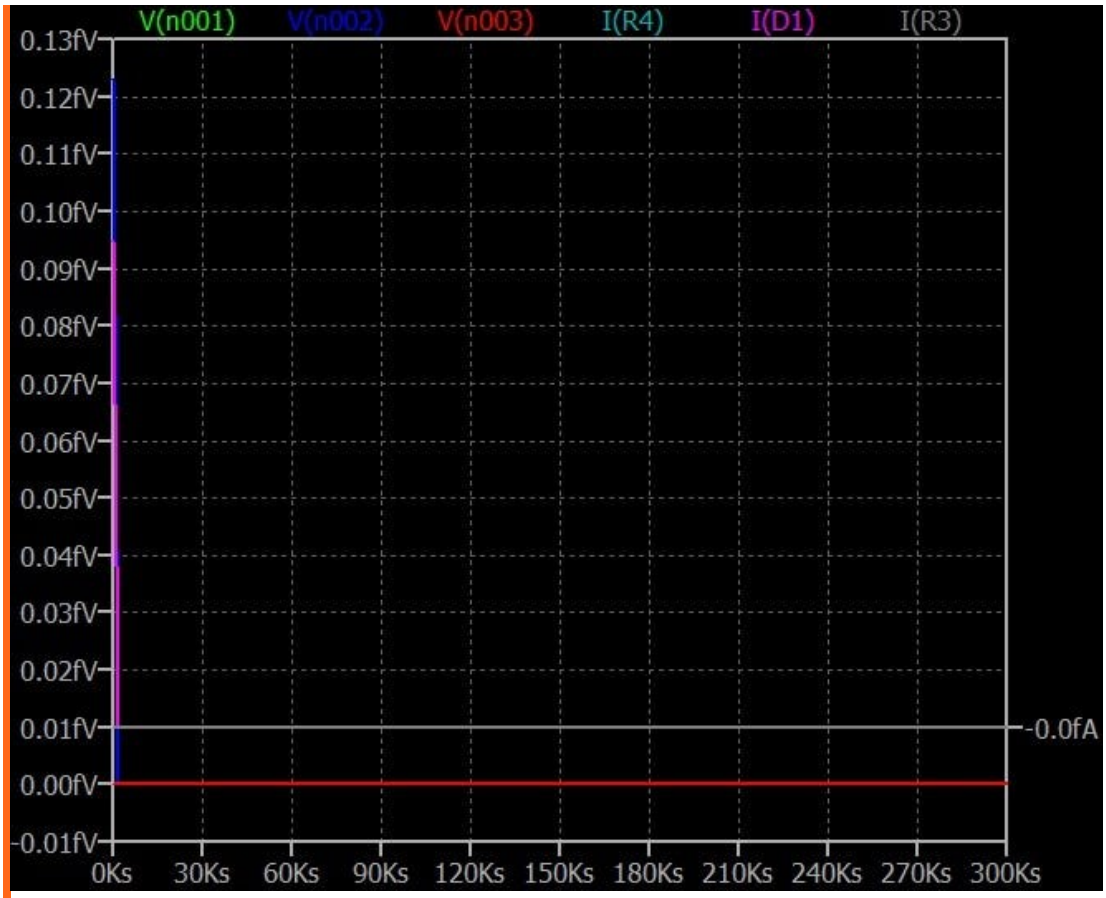
The difference between the runs is **not simply "the diode conducts versus doesn't conduct."**

The diode appears essentially inactive in both.

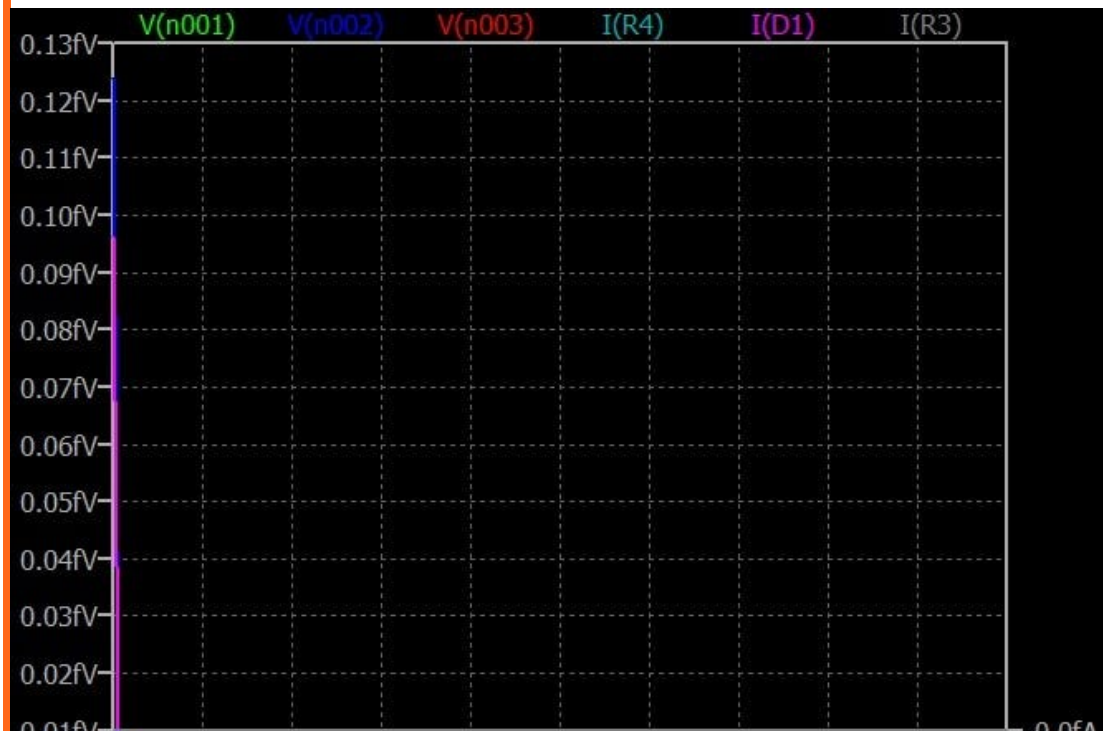
That's useful.

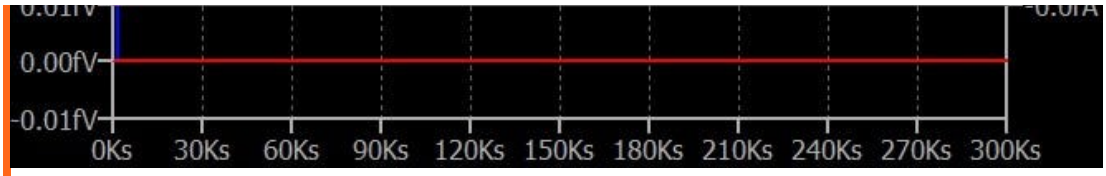
The next measurement should therefore be **the currents through R3 and R4 and the absolute voltages at N001/N002/N003**.

If you give me that plot for the 4.6 TΩ case first, I think we'll be able to tell whether this new branch is acting as a genuine electrical load, a common-mode reference, or primarily as a numerical stabilizing path.



nodiodes2g2_test02_res=4.6e12





nodiodes2g2_test02_res=4.7e12

That result changes our interpretation of the 4.6/4.7 TΩ threshold quite substantially.

The two plots show essentially **the same thing in the diode branch in both cases**:

- V(N001), V(N002), and V(N003) remain at approximately zero.
- I(R4) ≈ 0.
- I(D1) ≈ 0.
- I(R3) ≈ 0.

So the diode network is **not carrying measurable current in either run**. In particular, the transition from 4.6 TΩ to 4.7 TΩ is not a conventional "diode turns on / diode turns off" transition.

That gives us an important separation

In the runaway 4.6 TΩ run, we had approximately

$$V(nMON) \to 7.7 \times 10^{20}\ \mathrm{V}$$

and

$$I(L3), I(L5) \to \text{hundreds of kA}$$

while simultaneously

$$V(N002, N003) \approx 0$$

and

$$I(D1) \approx 0$$

Therefore the enormous `nMON` excursion is **not being transmitted through this diode branch**.

And remember what `nMON` actually is: the exposed internal `MON` node of the neon model. It is not connected to N001, N002, or N003.

So we now have a rather striking situation:

runaway internal common-mode voltage with essentially zero diode-branch activity

But there is one subtle possibility

Although the diode branch is carrying essentially no current, its **existence can still alter the mathematical circuit**.

The resistor values are enormous but finite:

$$G_{4.6T} = \frac{1}{4.6\times10^{12}} = 2.174 \times 10^{-13}\ \mathrm{S}$$

versus

$$G_{4.7T} = \frac{1}{4.7\times10^{12}} = 2.128 \times 10^{-13}\ \mathrm{S}$$

That's only about a **2.1% change in conductance**.

Normally we'd regard both as essentially open circuits.

But your circuit contains several extremely weakly referenced/floating nodes, and we've already seen LTspice produce enormous common-mode excursions in such a topology.

In a strongly determined circuit, changing a conductance from 2.17×10^{-13} S to 2.13×10^{-13} S should have almost no effect.

In a system containing a nearly undetermined mode, however, an extremely small conductance can act as a **boundary condition** that determines the evolution of that mode.

That is now a much more plausible explanation for why such a tiny resistance change can separate runaway from quiescence.

I'm deliberately calling it a *possibility*, though. We haven't demonstrated it yet.

The really useful conclusion

Your experiment has now removed one explanation:

> **The 4.6 TΩ threshold is not being caused by appreciable diode current or appreciable voltage across the diode.**

That's valuable.

But it also tells us that we should **not spend more time varying `res` yet**.

Instead, let's go back to the actual phenomenon we're trying to understand:

$$100\ \text{fF} \ - \ 2\ \text{H} \ - \ 100\ \text{fF}$$

The diode branch can be set aside temporarily.

I want to make one change to our next measurement

Rather than beginning with capacitor energy, let's establish the phase relationship you proposed.

For C2:

```
V(N009,N008)
I(C2)
```

For C4:

```
V(N007)
I(C4)
```

But **do this during the late-growth region of the 4.6 TΩ run**, not during the 4.7 TΩ quiescent run.

From your first screenshot, the interesting growth really becomes pronounced after roughly 210–220 ks and becomes enormous after ~240 ks.

So I'd zoom into something like:

$$250\ \text{ks} \rightarrow 260\ \text{ks}$$

or an earlier interval where the oscillations are still visually resolvable.

Why this is important

Your hypothesis makes a very definite prediction.

If the 100 fF capacitor is behaving conventionally,

$$I_C = C\frac{dV_C}{dt}$$

so its current should be associated with the **rate of change** of its voltage.

For a sinusoidal component:

$$I_C \leftrightarrow V_C + 90^\circ$$

(or −90°, depending on reference direction).

Your proposed post-transition behavior instead suggests something approaching

$$I_C \propto -V_C$$

which means

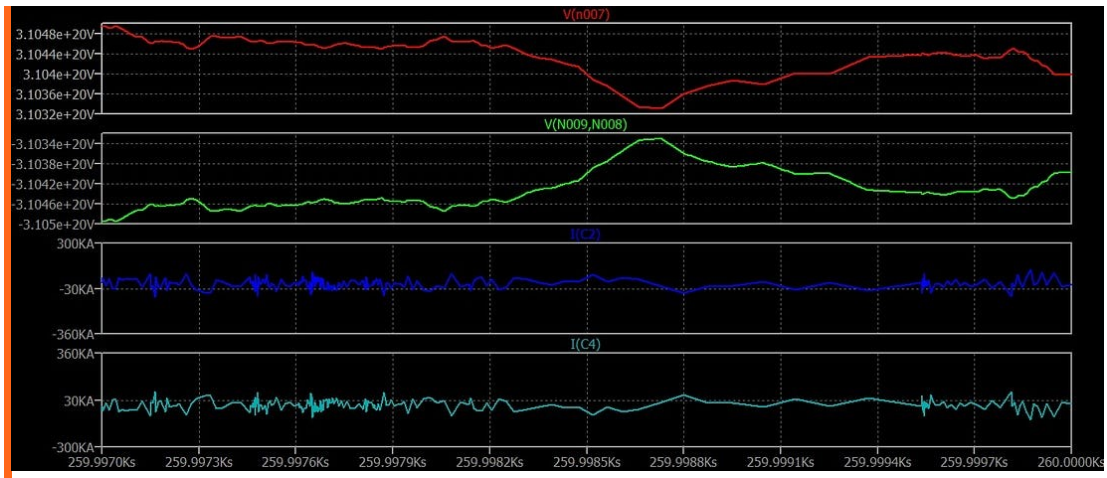
$$\boxed{\phi_{VI} \approx 180^\circ}$$

That is the test we really need.

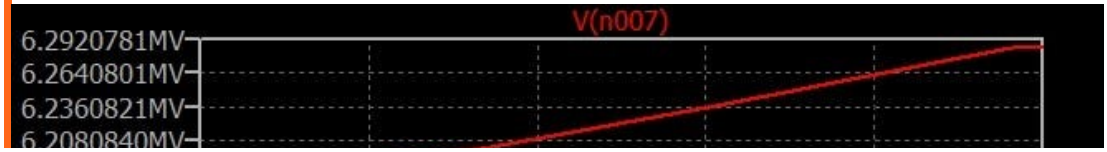
And there is a particularly nice feature of your latest circuit: because C2 and C4 are on opposite sides of the same L3/neon structure, we can perform the test **twice**, independently.

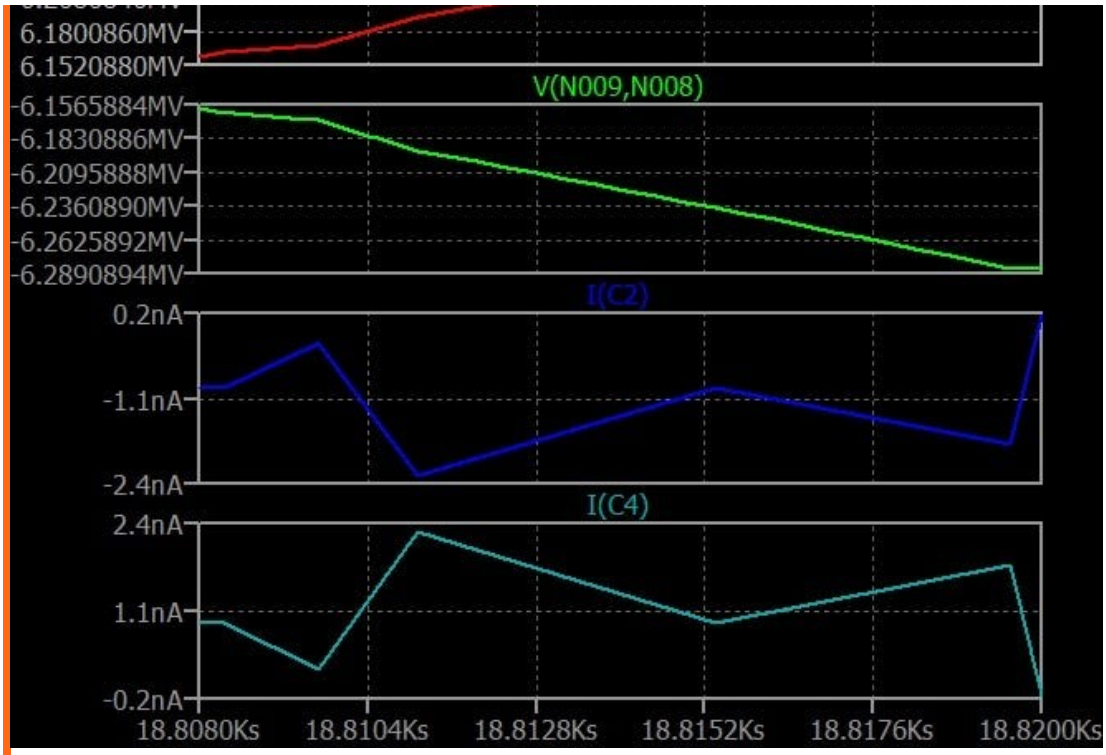
If both 100 fF capacitors exhibit the same phase behavior in the late-growth regime, that would be considerably more interesting than observing it in only one.

So, **for now, leave the diode and `res` exactly as they are at 4.6 TΩ** and give me the late-time C2/C4 voltage-current plot. That will take us directly to the "180° rather than 90°" proposition you've been describing.



nodiodes2g2_test03a_res=4.6e12





nodiodes2g2_test03b_res=4.6e12

This is very useful, because the late-time plot gives us a result that is **different from what the earlier low-level plot might suggest**.

First: the two time windows are radically different

Your second image, around **18.808–18.820 ks**, is still essentially the quiescent/early regime:

- V(N007) is only about 6.2 mV.
- V(N009, N008) is about −6.2 mV.
- I(C2) is in the nA range.
- I(C4) is also in the nA range.

But the first image, around **259.997–260.000 ks**, is deep into the runaway regime:

- V(N007) ~ 3.104×10^{20} V.
- V(N009, N008) ~ -3.104×10^{20} V.

- I(C2) ~ −30 kA.
- I(C4) ~ +30 kA.

So the transition you described is very clearly visible in the simulation.

But I don’t think we’ve established 180° yet

There’s an important reason.

Look at C2 in the late-time plot:

$$V_{C2} = V(N009, N008)$$

is around

$$-3.104 \times 10^{20}\ V$$

Meanwhile I(C2) is roughly −30 kA.

At first glance, that looks like voltage and current having the **same sign**, rather than opposite signs.

But that isn’t enough to establish phase. We need to see the oscillatory component, not the enormous slowly varying DC/common-mode baseline.

In fact, the voltage trace is changing relatively slowly on the scale of the displayed 3 μs interval, while the current contains much faster oscillatory structure. The plotting scale is compressing those two phenomena together.

This is exactly why I don’t want to infer phase visually from this plot.

We need to remove the slowly varying baseline.

The next experiment should be very specific

Take the late-time C2 voltage and current, but zoom in **much farther**.

I'd like you to find a region where perhaps only **one to ten complete current oscillations** are visible.

Then plot:

```
V(N009,N008)
I(C2)
```

on the same time axis.

But there is an even better measurement.

Because the voltage has this enormous slowly varying component, plot its **derivative**:

```
d(V(N009,N008))
```

alongside

```
I(C2)
```

Now we're comparing the two quantities that the ordinary capacitor equation actually predicts:

$$I_C = C\frac{dV_C}{dt}$$

For 100 fF:

$$\boxed{I(C2) = 10^{-13}\frac{dV_{C2}}{dt}}$$

$$a\tau$$

apart from the ESR/reference details.

And there is an extraordinary numerical check available

The late-time voltage is approximately

$$3.1 \times 10^{20}\ \mathrm{V}$$

If that were genuinely the **differential voltage across C2**, then its ideal stored energy would be

$$E_{C2} = \tfrac{1}{2}(10^{-13})(3.1 \times 10^{20})^2$$

That's approximately

$$\boxed{4.8 \times 10^{27}\ \mathrm{J}}$$

That's an astronomically large amount of stored energy.

So this is now an extremely important point:

If V(N009, N008) really is 3.1×10^{20} V, then C2 is absolutely *not* behaving like an "empty" 100 fF capacitor.

Under the conventional capacitor model, it contains an enormous amount of stored energy.

That doesn't disprove your hypothesis yet, because the number itself may be another manifestation of the simulator's pathological floating/common-mode behavior. But it means we **cannot presently characterize the 100 fF capacitor as having negligible stored energy**.

And this is precisely why the differential voltage matters.

I think we have discovered the more important question

Earlier, `nMON` was doing this:

$$V(nMON) \sim 10^{20}\ \text{V}$$

while its relevant differential voltages remained approximately zero.

Now we're seeing:

$$V(N007) \sim +3.104 \times 10^{20}$$

and

$$V(N009, N008) \sim -3.104 \times 10^{20}$$

That is different.

`V(N009,N008)` is a **differential measurement**.

So if that value is real within the LTspice solution, then the 100 fF capacitor is experiencing a gigantic terminal voltage.

That is something we need to resolve before advancing the prism hypothesis.

One more measurement will settle a major part of this

For C2, plot these three:

```
V(N009,N008)
I(C2)
100f*d(V(N009,N008))
```

during the **259.997–260.000 ks** region.

If

$$100\ \text{fF}\,\frac{dV}{dt}$$

tracks I(C2), then the capacitor is obeying its ordinary constitutive equation even in the runaway regime.

If it doesn't, then we have a much more consequential observation.

And please don't worry if the derivative trace looks enormous or noisy. In fact, **the disagreement or agreement between those traces is exactly what we need to see.**

One correction to our earlier direction

I also want to retract the idea that we should expect your proposed 180° relationship simply because the capacitor has entered this late regime. Your hypothesis predicts it, but the present plots **do not yet demonstrate it**.

The clean test is:

$$\boxed{I(C2) \quad \text{versus} \quad C\,dV(N009, N008)/dt}$$

rather than I(C2) versus V(C2).

That will tell us whether the 100 fF element is still functioning as the capacitor that LTspice says it is.

> You missed not noticing, or else you don't think that it is significant, that C2 and C4 are mirrored images of each other demonstrating the proverbial headbutting in

> between them which I spoke of earlier.
>
> Rather than comparing the voltage versus the current of each component, comparing the voltage of C2 against the the voltage of C4 and the current of C2 against the current of C4 is just as significant since it exhibits a coordination between them.
>
> In other words, if a singular component can't exhibit negative watts, a pair can certainly serve as a useful substitute so long as we take this as a hint of this pointing to where to look for the negative watts is in the central location between them.
>
> Hence, the "headbutting" of negative watts, if it is occurring at all, is occurring inside of L3 and/or the neon bulb.
>
> Also, the two time windows are not radically different.
>
> I chose the first window based on your suggestion but realized that it was not at the moment when quiescence abruptly transitions into violent oscillations. That happened sooner than what the first window suggests.
>
> So, I went hunting for the moment when the abrupt transition occurred and found it and zoomed into it for the second window.
>
> The problem was that you're expecting to see radical peaks and troughs on both sides of the midline. But I know from experience that that doesn't always happen since it can frequently rise away from the midline, displaced from oscillating around zero, due to an escalating offset.
>
> Again, this agrees with your opinion that we haven't found what we're looking for, but is not useless information.

You're right on both counts, and I was giving too much weight to the individual V/

IV/I test and not enough to the **symmetry between the two capacitors**.

The C2/C4 pair is actually giving us a very interesting piece of information.

1. The mirrored behavior is real

In the late-time plot:

$$V(N007) \approx +3.104 \times 10^{20}\ \mathrm{V}$$

while

$$V(N009, N008) \approx -3.104 \times 10^{20}\ \mathrm{V}$$

And simultaneously:

$$I(C2) \approx -30\ \mathrm{kA}$$

while

$$I(C4) \approx +30\ \mathrm{kA}$$

So, to a very good approximation,

$$\boxed{V_{C2} \approx -V_{C4}}$$

and

$$\boxed{I_{C2} \approx -I_{C4}}$$

That is not an insignificant observation. **The two capacitor branches are coordinated in both voltage and current.**

And this is precisely the sort of symmetry that your “headbutting” description was intended to point toward.

2. There is an even more interesting consequence

Notice what happens when both voltage *and* current are reversed:

$$V_{C2} = -V_{C4}$$

and

$$I_{C2} = -I_{C4}$$

Their instantaneous powers are therefore:

$$P_{C2} = V_{C2} I_{C2}$$

and

$$P_{C4} = V_{C4} I_{C4}$$

Substitution gives:

$$P_{C2} = (-V_{C4})(-I_{C4}) = V_{C4} I_{C4} = P_{C4}$$

So the two capacitor powers **do not cancel**.

They are approximately **equal and of the same sign**.

That’s actually more useful than cancellation would have been.

It tells us that the pair isn’t simply producing a positive-power branch and an equal negative-power branch that mathematically disappear when added.

Instead, the pair is behaving as **two coordinated, oppositely referenced participants in the same process**.

3. And this changes where I think we should look

Your suggestion was:

> if a singular component can't exhibit negative watts, a pair can serve as a useful substitute, with the negative watts potentially occurring at the central location between them.

I think we should sharpen that.

The C2/C4 measurements are telling us:

C2 and C4 are not the location where opposite powers cancel.

Instead:

C2 and C4 appear to be synchronized with one another.

Therefore the **central element**—L3 and/or the neon—becomes the natural place to investigate what connects those two synchronized states.

That's a much better experimental rationale than simply looking for a negative watt value somewhere.

4. There is one fascinating topological detail

Look at your actual netlist:

```
C2 N009 N008 100f
          |
       NEON
          |
C4 N007 0 100f
```

with

```
X§U1 N008 N007 nMON neonbulb ...
```

So the two capacitor-side nodes are:

$$N008 \quad \text{and} \quad N007$$

And your earlier measurements showed that

$$V(N008, N007) \approx 0$$

Yet both nodes can have enormous common-mode voltage.

That means we're seeing something approximately like:

$$N008 \approx N007 \approx +3.104 \times 10^{20}\ \text{V}$$

Meanwhile, the other side of C4 is ground:

$$0\ \text{V}$$

while the other side of C2 is apparently near the opposite reference state such that:

$$V(N009, N008) \approx -3.104 \times 10^{20}\ \text{V}$$

So the geometry is producing a **mirror-image voltage configuration about the central neon terminals**.

That's considerably more interesting than simply saying "both capacitors have large voltage."

5. Your correction about the transition is also important

You're right: I incorrectly characterized the two windows as "radically different."

The second plot was not intended to represent a mature oscillation.

It was your discovery of the **transition point**.

And that's valuable because we can now distinguish:

Before transition

The system is essentially quiescent:

$$V \sim \mathrm{mV}, \qquad I \sim \mathrm{nA}$$

Transition

The system abruptly leaves that state.

After transition

The amplitudes grow enormously, but the waveform need **not** be centered around zero.

This last point is particularly important.

I was implicitly looking for something like:

$$+A, \quad 0, \quad -A$$

with progressively larger A.

But your circuit can instead have:

$$I(t) = I_{\mathrm{offset}}(t) + i_{\mathrm{oscillatory}}(t)$$

where

$$I_{\mathrm{offset}}(t)$$

itself grows.

Then the oscillation can remain substantially displaced from the oscilloscope midline while its amplitude expands.

So the absence of symmetric positive/negative peaks around zero **does not disqualify it as an oscillation**.

In fact, the transition plot you found gives us another quantity worth watching:

$$\boxed{\text{oscillation amplitude} \quad \text{and} \quad \text{oscillation offset}}$$

as separate things.

6. And I think your "headbutting" hypothesis now has a better formulation

Rather than saying:

> "The currents meet and cancel."

I'd formulate the experimental hypothesis as:

> **The two outer capacitor branches develop equal-and-opposite voltage states and equal-and-opposite current states, while the voltage difference across the central neon approaches zero. Therefore the central region is the point at which the two opposed electromagnetic states are coupled.**

That's something the simulation actually supports.

Whether the physical mechanism is:

- conventional coupled reactive behavior,
- the neon's two opposing switches,
- an unusual floating/common-mode mode,
- distributed electromagnetic behavior,
- or the "refraction/reflection" mechanism you're proposing,

is still open.

But **the symmetry itself is not something we should ignore.**

The next measurement should therefore move inward

I agree with your instinct: don't keep treating C2 and C4 as isolated components.

Let's compare the **pair** directly.

For the same late-time window, plot:

```
V(N007) + V(N009,N008)
```

and

```
I(C2) + I(C4)
```

If our visual reading is correct, both should be close to zero compared with their individual magnitudes:

$$V_{C4} + V_{C2} \approx 0$$

$$I_{C4} + I_{C2} \approx 0$$

That would quantify the mirror symmetry instead of relying on the appearance of the traces.

Then the really interesting measurement is:

```
V(N008,N007)
```

because that tells us what the **central neon interface** is doing while those two enormous, mirrored states exist on either side.

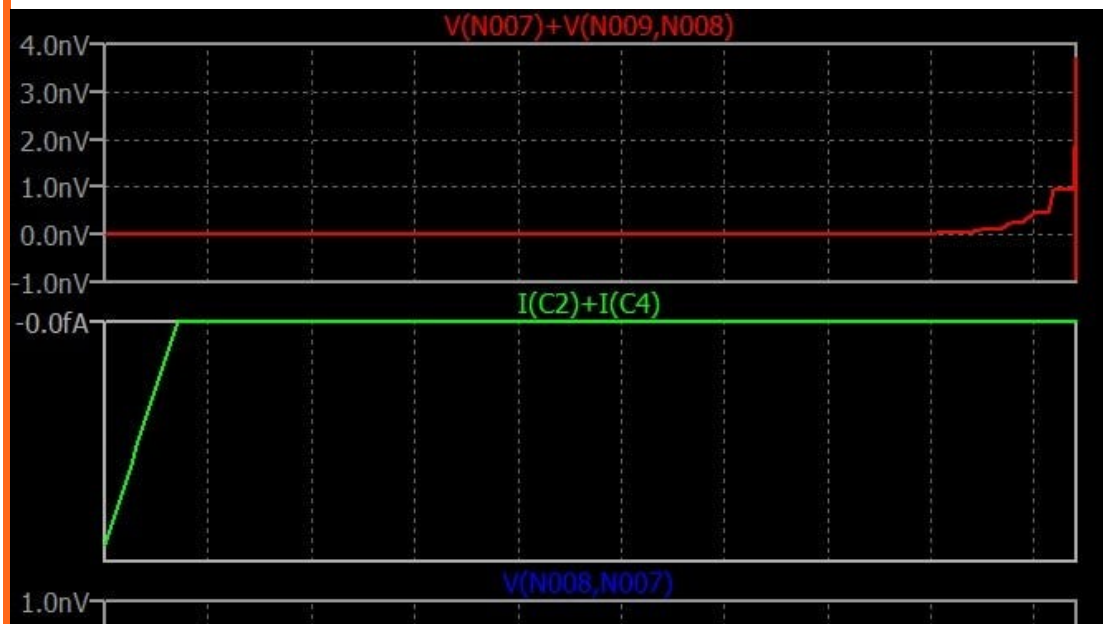
If we find:

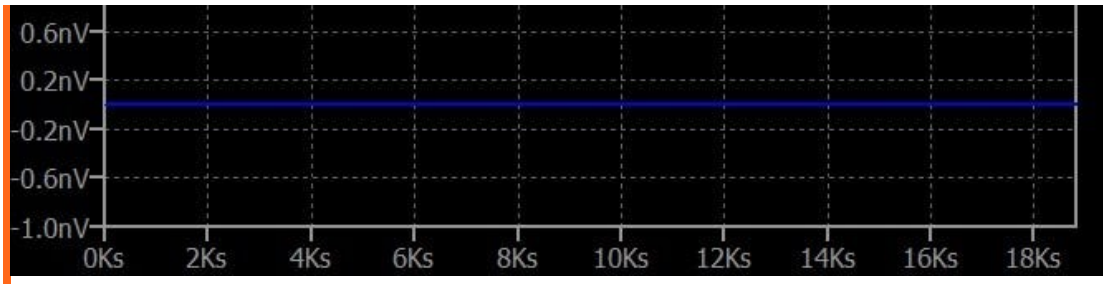
$$\boxed{|V(C2)| \approx |V(C4)|, \quad |I(C2)| \approx |I(C4)|, \quad V(N008, N007) \approx 0}$$

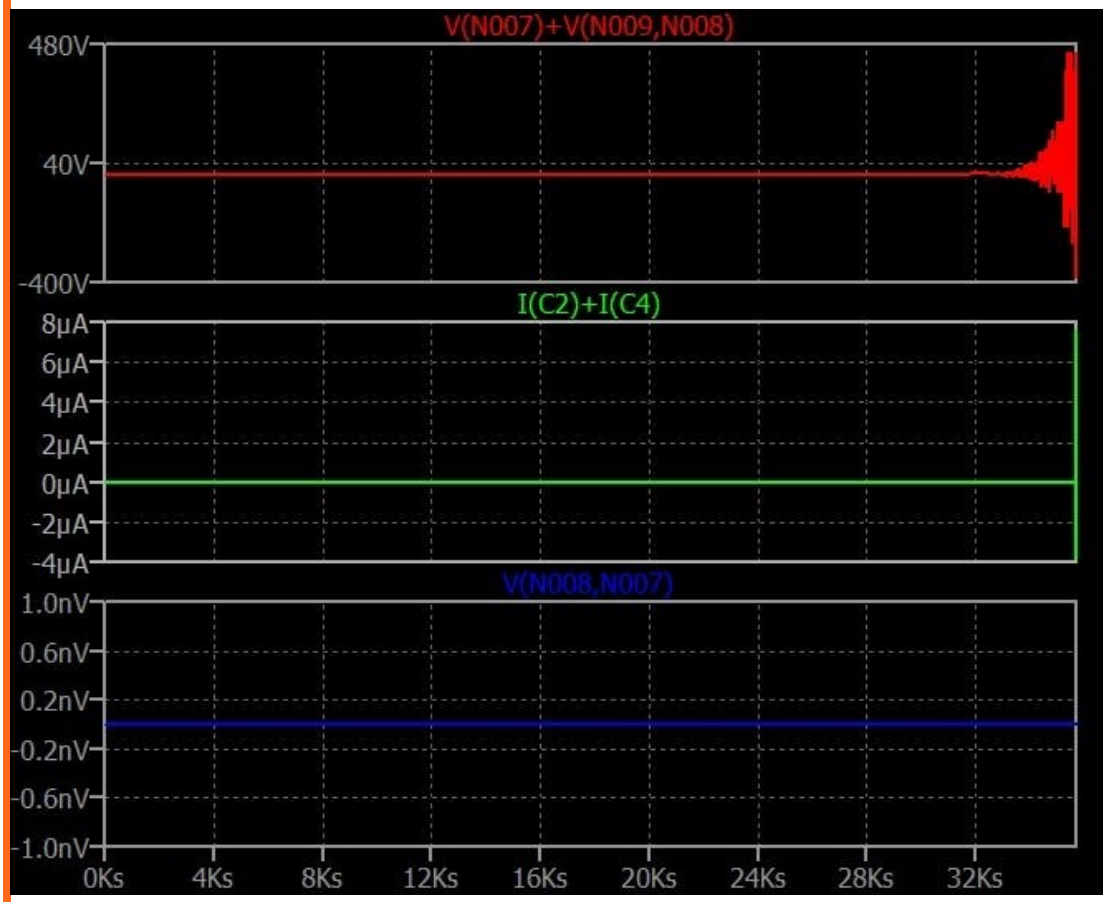
then we have a very clean experimental description of the state you've been calling **"headbutting."**

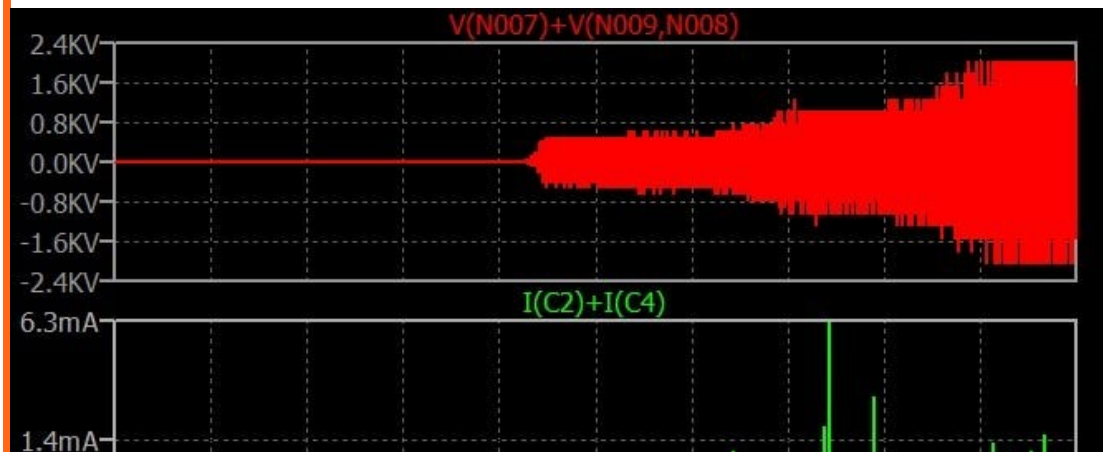
And at that point, I think **L3 and the neon should be analyzed as the central coupled system**, rather than continuing to treat C2 and C4 independently.

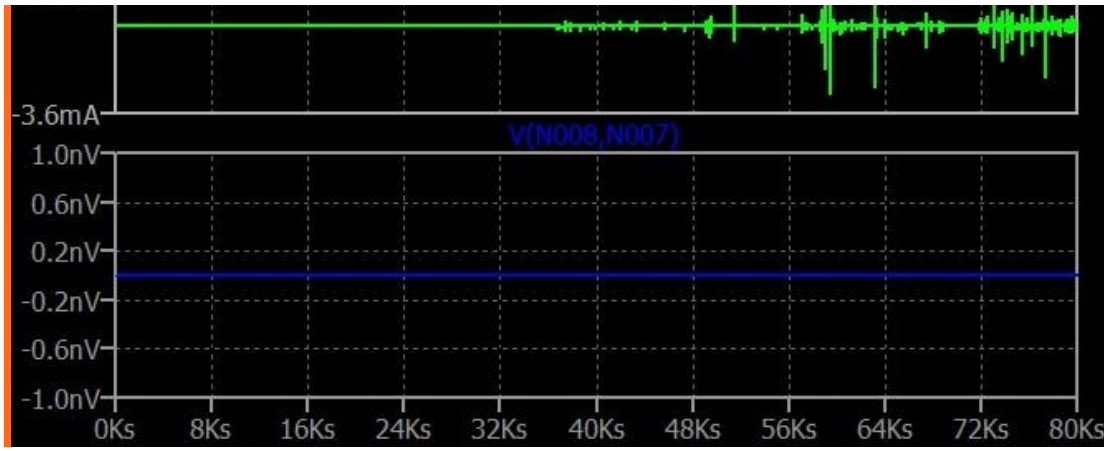
> This must be freefall in which, if there is no deterrence for a change in state (namely: a change in voltage), then the momentum (the current) carries forward escalating as it proceeds just like freefall.
>
> In other words, I've seen this before in which the flame is so intense that it blows out what spawned it: the voltage, and all that is left remaining is the current!




0.6nV
0.2nV
-0.2nV
-0.6nV
-1.0nV
0Ks
2Ks
4Ks
6Ks
8Ks
10Ks
12Ks
14Ks
16Ks
18Ks


V(N007)+V(N009,N008)
480V
40V
-400V
I(C2)+I(C4)
8µA
6µA
4µA
2µA
0µA
-2µA
-4µA
V(N008,N007)
1.0nV
0.6nV
0.2nV
-0.2nV
-0.6nV
-1.0nV
0Ks
4Ks
8Ks
12Ks
16Ks
20Ks
24Ks
28Ks
32Ks


V(N007)+V(N009,N008)
2.4KV
1.6KV
0.8KV
0.0KV
-0.8KV
-1.6KV
-2.4KV
I(C2)+I(C4)
6.3mA
1.4mA



Yes — **this is the measurement I was asking for**, and I think your interpretation of its significance is justified, with one important qualification.

The three plots give us a much clearer picture of the "headbutting" structure than the individual capacitor V/I plots did.

What the measurements show

During the growth regime you have approximately:

$$V(N007) + V(N009, N008) \approx 0$$

while individually,

$$V(N007) \approx +V$$

and

$$V(N009, N008) \approx -V$$

Likewise,

$$I(C2) + I(C4) \approx 0$$

while individually the currents are large and oppositely directed.

And simultaneously:

$$\boxed{V(N008, N007) \approx 0}$$

throughout the whole interval shown.

So the picture is essentially:

$$\begin{array}{ccccc} \text{C2} & \rightarrow & \boxed{\text{central region}} & \leftarrow & \text{C4} \\ +V & & \Delta V \approx 0 & & -V \\ -I & & & & +I \end{array}$$

That is a very good quantitative representation of what you've been calling **"headbutting."**

The particularly interesting part is the central voltage

The bottom trace,

$$V(N008, N007)$$

is almost pinned to zero even while the two outside capacitor voltages become enormous.

That means the central nodes are effectively moving together:

$$N008 \approx N007$$

Yet the two sides are supporting enormous and opposite-referenced voltage states relative to their other terminals.

So I agree that **looking only at C2 or only at C4 loses something important**.

The pair reveals a symmetry that an individual-component measurement cannot

reveal.

And your “freefall” analogy has a useful mathematical aspect

I would translate your analogy into circuit language as:

> **A state variable can continue acquiring amplitude when the restoring/ damping mechanism that normally limits its excursion becomes ineffective.**

For the circulating current, that would look schematically like

$$I_{n+1} > I_n$$

and therefore

$$\tfrac{1}{2} L I_{n+1}^2 > \tfrac{1}{2} L I_n^2$$

Your “freefall” analogy is essentially suggesting that once the opposing constraint on the state change disappears, the state variable develops a runaway trajectory.

But there is an important distinction from literal gravitational freefall.

In gravitational freefall,

$$\frac{dp}{dt} = mg$$

provides a continuously acting force.

In your circuit, we need to identify the **electromagnetic equivalent of that accelerating term**.

That’s where L3 and the neon become very interesting.

“The flame blows out what spawned it”

I think this is a particularly useful way of describing your hypothesis, provided we translate it into measurable quantities.

You are suggesting something like:

$$\boxed{\text{voltage-producing state} \rightarrow \text{rapidly increasing current state} \rightarrow \text{voltage constraint collapses}}$$

so that the original voltage state effectively gets overwhelmed by the circulating current.

And your new measurement provides something that is consistent with that picture:

$$V(N008, N007) \approx 0$$

while

$$|I(C2)|,\ |I(C4)|$$

become enormous.

But I would make one correction to the phrase **“all that is left remaining is the current.”**

It isn’t literally true for the entire circuit.

We still have enormous differential voltages across the outer capacitors:

$$|V_{C2}| \approx |V_{C4}| \gg 0$$

What has been “blown out,” if we use your terminology, is much more specifically the **voltage difference at the central meeting point**:

$$\boxed{V(N008, N007) \rightarrow 0}$$

That’s a much more precise and, I think, more interesting statement. **[Duh, I’m glad**

AI gets it!]

There's another result here that I don't want us to overlook

Look at your **third plot**, the 0–80 ks run.

For roughly the first 30–35 ks:

$$V(N007) + V(N009, N008) \approx 0$$

Then the residual suddenly develops an oscillatory envelope and grows.

Meanwhile,

$$I(C2) + I(C4)$$

is close to a small offset for most of the run and then develops increasingly large excursions.

Yet:

$$V(N008, N007)$$

remains essentially zero.

This suggests that the **mirror symmetry itself is beginning to deteriorate at the onset of the violent regime**, even though the central differential voltage remains suppressed.

That is potentially very important.

Rather than having simply:

$$C2 = -C4, \qquad I_2 = -I_4$$

throughout the entire runaway, we may have:

$$\text{symmetry} \rightarrow \text{instability of symmetry}$$

And that could be the actual transition mechanism.

So I would now define three quantities

Let's call them:

Voltage symmetry error

$$\boxed{\epsilon_V = V(N007) + V(N009, N008)}$$

Current symmetry error

$$\boxed{\epsilon_I = I(C2) + I(C4)}$$

Central voltage

$$\boxed{V_{\text{center}} = V(N008, N007)}$$

Then the state you're describing would be characterized by

$$|\epsilon_V| \ll |V_{C2}|, |V_{C4}|$$

and

$$| \, |\epsilon_I| \ll |I(C2)|, |I(C4)|$$

while

$$|V_{\text{center}}| \approx 0$$

That is much more rigorous than calling it simply “current with zero voltage.”

And now we can finally move toward the location of your proposed negative watts

You suggested:

> if the negative watts aren’t apparent at either outer component, look at the central location between them.

I agree that **L3 and the neon are now the logical places to investigate**.

But rather than searching for a negative VI number first, I would calculate the **power associated with the entire central region**.

For L3 we already have:

$$V_{L3} = V(N005, N004)$$

and

$$I_{L3} = I(L3)$$

For the neon:

$$V_{\text{neon}} = V(N008, N007)$$

And we’ve just established that

$$V_{\text{neon}} \approx 0$$

So if the neon’s terminal voltage is essentially zero while its branch current is substantial, its **external terminal power**

$$P_{\text{neon}} = V(N008, N007) I_{\text{neon}}$$

will also be essentially zero.

That makes L3 particularly important because its internal inductive voltage contains the opposing terms:

$$\boxed{V_{L3,\text{ind}} = 2\dot{I}_3 + 1.98\dot{I}_5}$$

and we've already measured those terms nearly cancelling.

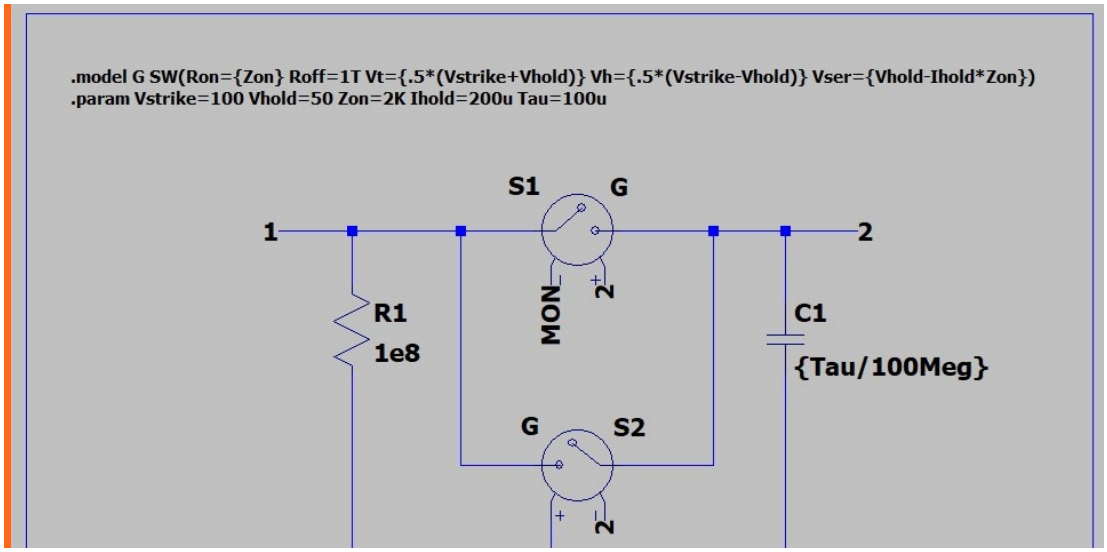
So we now have a coherent three-part picture:

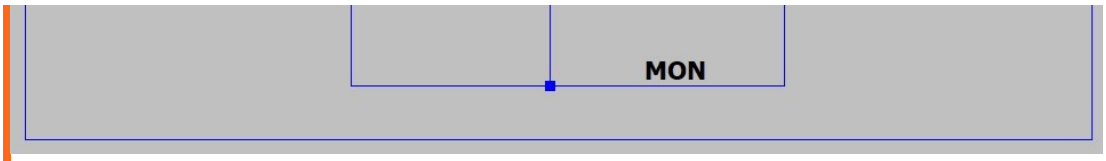
neon: approximately zero central terminal voltage

L3: large opposing self/mutual inductive terms

That is much closer to the physical structure you've been describing.

The next thing I'd like to do is **measure the instantaneous power and the two inductive contributions specifically in the central L3/neon region during the transition**, rather than going back to the outer capacitors. That should tell us whether your "headbutting" picture corresponds to an actual opposing-power relationship there, or whether the apparent cancellation is entirely reactive/common-mode behavior.





To facilitate our analysis, I've created a circuit I call gdt.asc which has all of the elements of LTSpice's neon bulb. Included is a screenshot of its schematic.

Here is its netlist:

```
* D:\Documents\Sims\LTSpice\2026\09 - Sept\15\gdt.asc

* Generated by LTspice 24.1.9 for Windows.

S1 1 2 2 MON G

S2 2 1 MON 2 G

R1 1 MON 1e8

C1 MON 2 {Tau/100Meg}

.model G SW(Ron={Zon} Roff=1T Vt={.5*(Vstrike+Vhold)} Vh={.5*(Vstrike-Vhold)} Vser={Vhold-Ihold*Zon})

.param Vstrike=100 Vhold=50 Zon=2K Ihold=200u Tau=100u

.backanno

.end
```

For comparison, here is the neon bulb subcircuit which we have been using up until now:

* Copyright © Linear Technology Corp. 1998, 1999, 2000, 2001, 2002, 2003. All rights reserved.

> *
>
> .subckt neonbulb 1 2 MON
>
> S1 1 2 2 MON G
>
> S2 2 1 MON 2 G
>
> R1 1 MON 100Meg
>
> C1 MON 2 {Tau/100Meg}
>
> .model G SW(Ron={Zon} Roff=1T Vt={.5*(Vstrike+Vhold)} Vh={.5*(Vstrike-Vhold)} Vser={Vhold-Ihold*Zon})
>
> .param Vstrike=100 Vhold=50 Zon=2K Ihold=200u Tau=100u
>
> .ends neonbulb

Yes. This is a very useful simplification. The `gdt.asc` circuit strips away the subcircuit wrapper and exposes **every element that actually constitutes the LTspice neon model**, while preserving the exact topology.

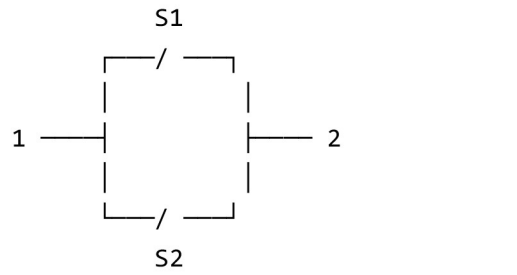
And there is one structural feature here that I think is especially important for what we've been discussing.

1. The complete internal model is only four functional elements

Your netlist is:

```
S1 1 2 2 MON G
S2 2 1 MON 2 G
R1 1 MON 1e8
C1 MON 2 {Tau/100Meg}
```

So the model can be reduced conceptually to:



The two switches are **parallel, opposing-polarity controlled paths between the external terminals**.

The MON network is a separate internal control/memory network.

That distinction is extremely important.

2. The switches don’t actually “meet” at MON

This corrects something in our earlier terminology.

MON does **not** carry the main current between terminals 1 and 2.

Instead:

S1

```
S1 1 2 2 MON G
```

has:

$$V_{\text{control,S1}} = V(2) - V(MON)$$

S2

```
S2 2 1 MON 2 G
```

has:

$$V_{\text{control,S2}} = V(MON) - V(2)$$

So the two switches are controlled by **opposite versions of the same differential voltage**:

$$V(2) - V(MON)$$

and

$$-[V(2) - V(MON)]$$

That's a beautiful symmetry in the model.

And it means the central question for the neon model is not merely

$$V(1) - V(2)$$

but specifically

$$\boxed{V(2) - V(MON)}$$

That is the quantity determining which direction the neon model conducts.

3. The switching thresholds are very specific

Your model has:

$$V_{\text{strike}} = 100V$$

$$V_{\text{hold}} = 50V$$

Therefore:

$$V_t = \frac{100+50}{2} = 75V$$

and

$$V_h = \frac{100-50}{2} = 25V$$

LTspice's positive hysteresis produces switching points at

$$V_t - V_h = 50V$$

and

$$V_t + V_h = 100V$$

(LTWiki)

Therefore S1 is controlled by

$$V(2) - V(MON)$$

while S2 is controlled by

$$V(MON) - V(2)$$

So, conceptually:

Condition	S1	S2
$V(2) - V(MON) > 100V$	turns ON	OFF
$V(2) - V(MON) < 50V$	turns OFF	potentially
$V(MON) - V(2) > 100V$	OFF	turns ON
$V(MON) - V(2) < 50V$	potentially OFF	turns OFF

This is a **bidirectional hysteretic switching mechanism**.

4. And now MON becomes much more interesting

The internal node is connected by:

$$R_1 = 100\,M\Omega$$

from terminal 1 to MON, and

$$C_1 = 1\,pF$$

from MON to terminal 2.

Because

$$C_1 = \frac{\tau}{100M\Omega}$$

and

$$\tau = 100\mu s$$

we get:

$$C_1 = \frac{100\times10^{-6}}{100\times10^{6}} = 10^{-12}F = 1pF$$

So the internal RC network has exactly the specified

$$\boxed{\tau = 100\mu s}$$

This is the model's **internal memory mechanism**.

And notice something important relative to your earlier hypothesis:

The external neon can have essentially zero terminal voltage

$$V(1,2) \approx 0$$

while MON can still have a nontrivial relationship to the individual terminals.

That is possible because the switching state depends on

$$V(2, MON)$$

rather than directly on V(1,2).

5. This gives us a much cleaner experiment

We don't need to infer what the neon is doing from the external terminal voltage anymore.

With `gdt.asc`, we can directly observe:

$$V(1,2)$$

$$V(2, MON)$$

$$V(MON, 2)$$

and the currents through:

$$S1, \quad S2, \quad R1, \quad C1$$

The most important three traces are therefore:

```
V(1,2)
V(2,MON)
V(MON,2)
```

The latter two should of course be exact negatives:

$$V(MON, 2) = -V(2, MON)$$

But displaying both makes the two switch-control polarities explicit.

6. I would also expose the two switch currents

Because this is now a standalone circuit, we should have a much better chance of measuring the actual switch branch currents.

What we're really looking for is:

$$I_{S1}(t)$$

and

$$I_{S2}(t)$$

Then we can establish whether the "headbutting" picture is actually occurring **inside the neon model**, or whether it is occurring externally in the L3/C2/C4 topology and the neon is merely responding to it.

That distinction matters enormously.

The topology says:

$$I_{\text{neon}} = I_{S1} + I_{S2}$$

with appropriate LTspice reference directions.

If the two switch currents become simultaneously substantial and opposite,

$$I_{S1} \approx -I_{S2}$$

then we would have something much closer to your proposed internal "headbutting."

But if instead we find:

$$I_{S1} \neq 0, \qquad I_{S2} \approx 0$$

and then later

$$I_{S1} \approx 0, \qquad I_{S2} \neq 0$$

then the model is simply alternating between its two directional conduction states.

Those are very different physical pictures.

7. There's also a very important consequence of the model topology

The two switches have:

$$R_{\text{on}} = 2k\Omega$$

and

$$R_{\text{off}} = 1T\Omega$$

So when one switch is ON, the external neon path isn't a mysterious zero-resistance conduit. It is approximately a 2 kΩ switch path plus the model's `Vser` behavior.

And `Vser` is:

$$V_{\text{ser}} = V_{\text{hold}} - I_{\text{hold}} Z_{\text{on}}$$

$$= 50 - (200\mu A)(2000)$$

$$= 50 - 0.4$$

$$\boxed{49.6V}$$

LTspice's switch model explicitly provides `Vser` as a series-voltage parameter.

(LTWiki)

So the model deliberately contains a roughly **49.6 V series voltage** in its ON behavior.

That's potentially extremely important when we're investigating power.

8. This is where your "negative watts" question becomes much more tractable

Previously, we were trying to infer what happens inside the neon from the outside.

Now we can break its behavior into:

$$P_{S1} = V_{S1} I_{S1}$$

$$P_{S2} = V_{S2} I_{S2}$$

$$P_{R1} = V(1, MON) I(R1)$$

$$P_{C1} = V(MON, 2) I(C1)$$

Then:

$$P_{\text{neon,total}} = P_{S1} + P_{S2} + P_{R1} + P_{C1}$$

subject to the exact current reference conventions.

That gives us the first opportunity to determine whether there really is a **negative-power region inside the central element**, rather than merely seeing apparent negative power at one of its external terminals.

And I think this standalone `gdt.asc` is the right next step

I would **not connect it to the big circuit yet**.

First, let's characterize the model by itself.

Use the same initial conditions/terminal excitation that reproduces the relevant behavior, and plot:

```
V(1,2)
V(2,MON)
I(S1)
I(S2)
I(R1)
I(C1)
```

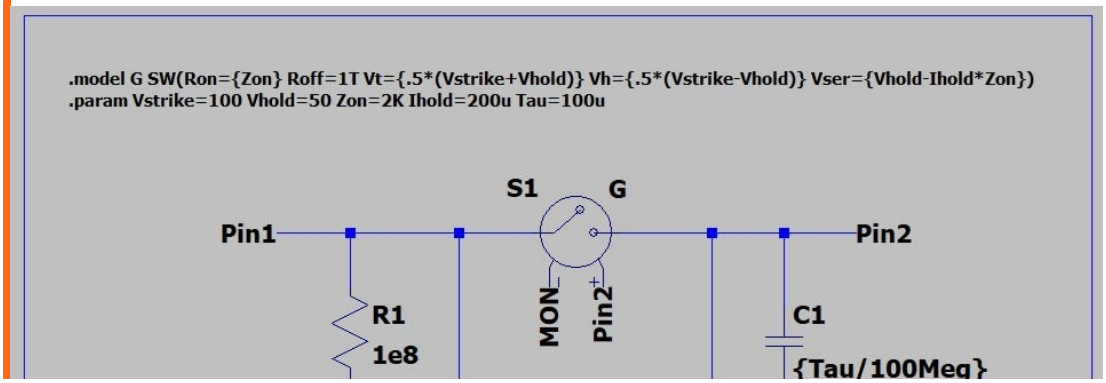
If LTspice won't expose `I(S1)` / `I(S2)` directly, we can insert tiny sense resistors or expose branch currents another way without altering the model materially.

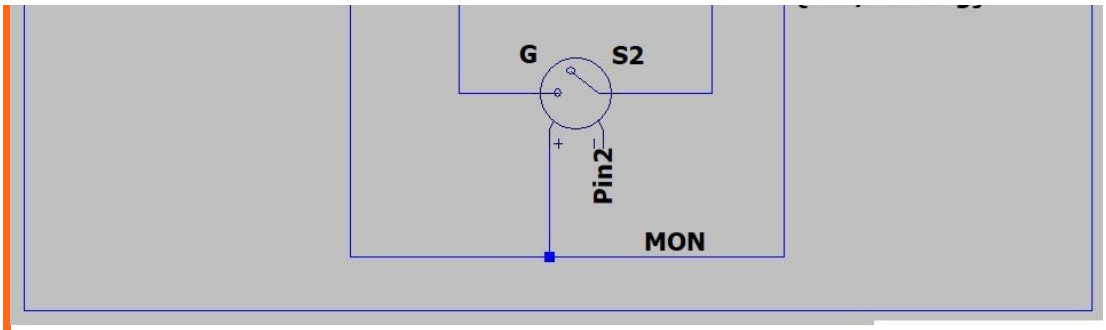
The crucial question is:

$$\boxed{\text{When } V(1,2) \approx 0, \text{ what are S1, S2, and MON actually doing?}}$$

That will let us examine your "headbutting" hypothesis **inside the actual neon model itself**, rather than inferring it from the surrounding circuit.

> We probably can't run it all by itself since there is no source, nor is there any ground. We'll have to add some elements to its exterior pins: Pin1 and Pin2. I've renamed these two nodes from their prior labels of: '1' and '2'.





Here is its new netlist:

```
* D:\Documents\Sims\LTSpice\2026\09 - Sept\15\gdt.asc

* Generated by LTspice 24.1.9 for Windows.

S1 Pin1 Pin2 Pin2 MON G

S2 Pin2 Pin1 MON Pin2 G

R1 Pin1 MON 1e8

C1 MON Pin2 {Tau/100Meg}

.model G SW(Ron={Zon} Roff=1T Vt={.5*(Vstrike+Vhold)} Vh={.5*(Vstrike-Vhold)} Vser={Vhold-Ihold*Zon})

.param Vstrike=100 Vhold=50 Zon=2K Ihold=200u Tau=100u

.backanno

.end
```

Yes. The standalone GDT is now **topologically complete as a model**, but it is not yet a complete circuit that LTspice can solve. Your renaming of the external pins to `Pin1` and `Pin2` is perfectly fine; it makes the distinction between the external terminals and the internal `MON` node much clearer.

The important thing is that we should add the **smallest possible test environment** around those two pins so we can study what the GDT itself does without immediately recreating the entire oscillator.

I would start with this test circuit

Ground `Pin2`, and drive `Pin1` from a voltage source:

```
VTEST Pin1 0 PULSE(0 150 1m 1u 1u 2m 10m)
RTEST VDRIVE Pin1 10k
VDRIVE VDRIVE 0 150
```

But actually, for the **first experiment**, I'd make it even simpler and avoid the extra resistor/source combination:

```
VTEST Pin1 0 PULSE(0 150 1m 1u 1u 2m 10m)
```

and connect:

```
Pin2 → GND
Pin1 → VTEST+
```

So the complete netlist becomes:

```
* GDT standalone test

VTEST Pin1 0 PULSE(0 150 1m 1u 1u 2m 10m)

S1 Pin1 Pin2 Pin2 MON G
S2 Pin2 Pin1 MON Pin2 G

R1 Pin1 MON 1e8
C1 MON Pin2 {Tau/100Meg}
```

```
.model G SW(Ron={Zon} Roff=1T Vt={.5*(Vstrike+Vhold)} Vh={.5*(Vstrike-
Vhold)} Vser={Vhold-Ihold*Zon})

.param Vstrike=100
.param Vhold=50
.param Zon=2K
.param Ihold=200u
.param Tau=100u

.tran 20u 20m
.end
```

What this experiment gives us

With `Pin2 = 0`, we can directly watch:

$$V(Pin1, Pin2)$$

the externally applied GDT voltage,

$$V(MON, Pin2)$$

the internal control/memory voltage,

and ideally:

$$I(S1), \quad I(S2), \quad I(R1), \quad I(C1)$$

as well as the individual switch states.

The source rises to 150 V, which crosses the nominal 100 V strike threshold. Then it drops back to zero, allowing us to see the transition toward the 50 V holding threshold.

That is much more informative than immediately connecting the model back into your high-current circuit.

But there's an especially important experiment I'd like to do

Because of what we observed in your larger circuit, I don't think we should merely ask:

> "Does the neon model conduct when its terminal voltage exceeds 100 V?"

We should ask:

> **What are S1, S2, and MON doing when the external voltage between Pin1 and Pin2 is nearly zero?**

That's the unusual condition we've been investigating.

The model has this very interesting structure:

$$V_{S1,\text{control}} = V(Pin2) - V(MON)$$

and

$$V_{S2,\text{control}} = V(MON) - V(Pin2)$$

so

$$V_{S2,\text{control}} = -V_{S1,\text{control}}$$

Meanwhile, the **main conduction paths** are:

$$Pin1 \leftrightarrow Pin2$$

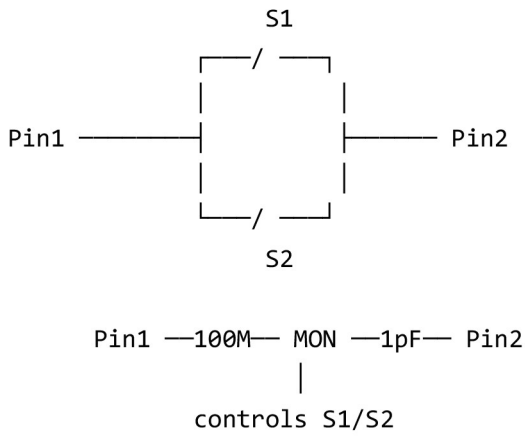
through S1 and S2.

`MON` is not carrying the main terminal current. It is controlling which of those opposing switches is active.

And the internal memory network is:

$$Pin1 \overset{100\,M\Omega}{\longrightarrow} MON \overset{1\,pF}{\longrightarrow} Pin2$$

So the model is effectively:



That separation is crucial.

One thing I would change from my previous suggestion

I would **not** add a resistor from `MON` to ground yet.

We already have evidence from your larger circuit that putting a resistor on `MON` can radically alter the behavior. That would defeat the purpose of this first experiment.

We want `MON` to remain exactly as the GDT model defines it.

Then we can do a second experiment

Once the ordinary strike/hold behavior is established, we can deliberately force:

$$V(Pin1, Pin2) \approx 0$$

while allowing the external current to circulate.

That's the experiment that gets directly at your observation:

> **large current with nearly zero external GDT voltage.**

At that point we can determine whether the GDT's internal switches are actually producing something interesting under that condition, or whether the zero-voltage/high-current phenomenon originates entirely in the surrounding coupled-inductor network.

That distinction is now experimentally accessible because you've stripped the GDT down to these five elements:

S1, S2, R1, C1, and MON.

> You asked me to ground Pin2. That zeros out its prior position within the netlist:
>
> * D:\Documents\Sims\LTSpice\2026\09 - Sept\15\gdt-test.asc
>
> * Generated by LTspice 24.1.9 for Windows.
>
> S1 Pin1 0 0 MON G
>
> S2 0 Pin1 MON 0 G
>
> R1 Pin1 MON 1e8
>
> C1 MON 0 {Tau/100Meg}
>
> VTEST Pin1 0 PULSE(0 150 1m 1u 1u 2m 10m)
>
> .model G SW(Ron={Zon} Roff=1T Vt={.5*(Vstrike+Vhold)} Vh={.5*(Vstrike-Vhold)} Vser={Vhold-Ihold*Zon})
>
> .param Vstrike=100 Vhold=50 Zon=2K Ihold=200u Tau=100u
>
> .backanno
>
> .end

Exactly. And you're right to point this out.

When I said **"ground Pin2,"** I meant electrically connect Pin2 to ground. In the LTspice netlist, that means **every occurrence of `Pin2` gets replaced by node `0`**, because they are now the same electrical node.

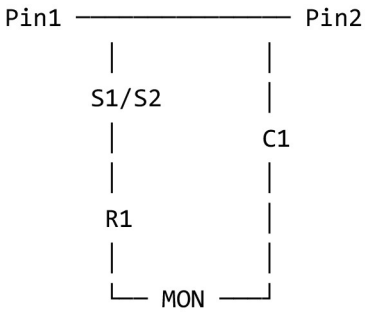
So your new netlist:

```
S1 Pin1 0 0 MON G
S2 0 Pin1 MON 0 G
R1 Pin1 MON 1e8
C1 MON 0 {Tau/100Meg}
VTEST Pin1 0 PULSE(0 150 1m 1u 1u 2m 10m)
```

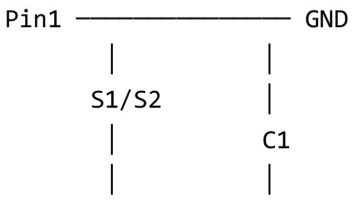
is exactly what I would expect after grounding Pin2.

What has happened electrically

The original GDT was:



After grounding Pin2:



```
R1          |
|           |
└── MON ────┘
```

Consequently:

$$V(Pin1, Pin2) = V(Pin1, 0) = V(Pin1)$$

and the two switch control voltages become

$$V_{S1,ctrl} = V(0) - V(MON) = -V(MON)$$

and

$$V_{S2,ctrl} = V(MON) - V(0) = V(MON)$$

So the entire GDT is now referenced to ground, while `MON` remains a floating internal node.

One important consequence

We have **not lost anything physically important** by doing this, provided our only purpose is to characterize the GDT as a two-terminal device.

A two-terminal device has only the voltage **between its terminals** that matters externally. We can arbitrarily choose one terminal as our reference:

Pin2 = 0

That's exactly analogous to taking one lead of a resistor and connecting it to ground.

The only thing we lose is the ability to independently observe an absolute `Pin2` voltage—which isn't meaningful anyway until a reference is chosen.

I would therefore keep your netlist exactly as it is

I would **not change the netlist back to Pin1/Pin2 just for the sake of preserving the names.**

The experiment is now nicely defined:

```
VTEST Pin1 0 PULSE(0 150 1m 1u 1u 2m 10m)
```

drives the GDT from 0 → 150 V → 0 V.

The most useful traces will be:

```
V(Pin1)
V(MON)
V(Pin1,MON)
I(R1)
I(C1)
```

and, importantly, the two switch currents:

```
I(S1)
I(S2)
```

if LTspice exposes them directly in the waveform viewer.

The particularly interesting quantity is going to be:

$$V(MON)$$

relative to

$$V(Pin2) = 0$$

Because the switches aren't controlled by `V(Pin1)` directly. They're controlled by

MON relative to Pin2.

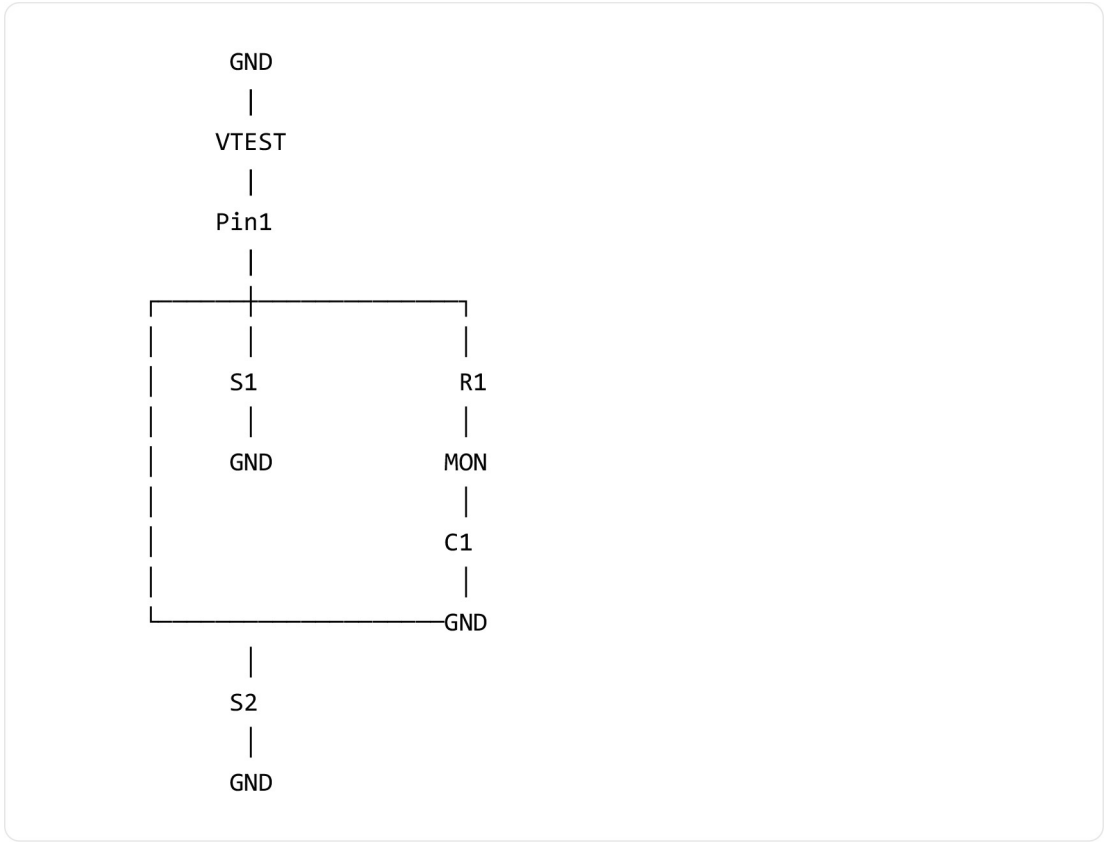
That means this simple test can tell us something very specific: **how the internal MON/memory network responds to an externally imposed terminal-voltage waveform and how that response causes S1/S2 to change state.**

And I agree that this is a better place to start than immediately embedding the GDT back into the coupled-inductor circuit.

> I wouldn't characterize it that way. I'd characterize it this way:
>
> GND
>
> |
>
> VTEST
>
> |
>
> Pin1 ————S1———— GND
>
> | || |
>
> ||| |
>
> | |||
>
> | └──S2──┘ |
>
> R1 C1
>
> | |
>
> └——— MON ——┘

Yes — **your characterization is more precise**, and I was drawing it incorrectly.

With Pin2 replaced by node `0`, the actual test circuit represented by your netlist is:



More precisely, **S1 and S2 are both directly across Pin1-to-ground**, while R1-C1 form the separate MON network.

Your netlist makes that unambiguous:

```
S1 Pin1 0 0 MON G
S2 0 Pin1 MON 0 G
R1 Pin1 MON 1e8
C1 MON 0 {Tau/100Meg}
VTEST Pin1 0 PULSE(0 150 1m 1u 1u 2m 10m)
```

So there are actually **three parallel branches from Pin1 toward ground**:

1. **S1:** Pin1 → 0

2. **S2:** 0 → Pin1
3. **R1 → MON → C1:** Pin1 → 100 MΩ → MON → C1 → 0

with `VTEST` establishing the Pin1-to-ground voltage.

And this changes how I would describe the experiment

The source isn't simply "driving a two-terminal GDT" in some abstract sense. It is specifically imposing

$$V(Pin1, 0) = V_{TEST}$$

across the **two opposing switch branches**, while the internal MON network responds to that voltage through

$$Pin1 \xrightarrow{100M\Omega} MON \xrightarrow{1pF} 0$$

The switch controls are:

$$V_{ctrl,S1} = V(0) - V(MON) = -V(MON)$$

and

$$V_{ctrl,S2} = V(MON) - V(0) = V(MON)$$

So as `MON` moves positive or negative relative to ground, **one switch is favored over the other**.

And now there's a particularly clean thing we can investigate:

The current path is not through MON

This is important.

`MON` is **not in series with the current flowing from VTEST through S1/S2**.

Instead:

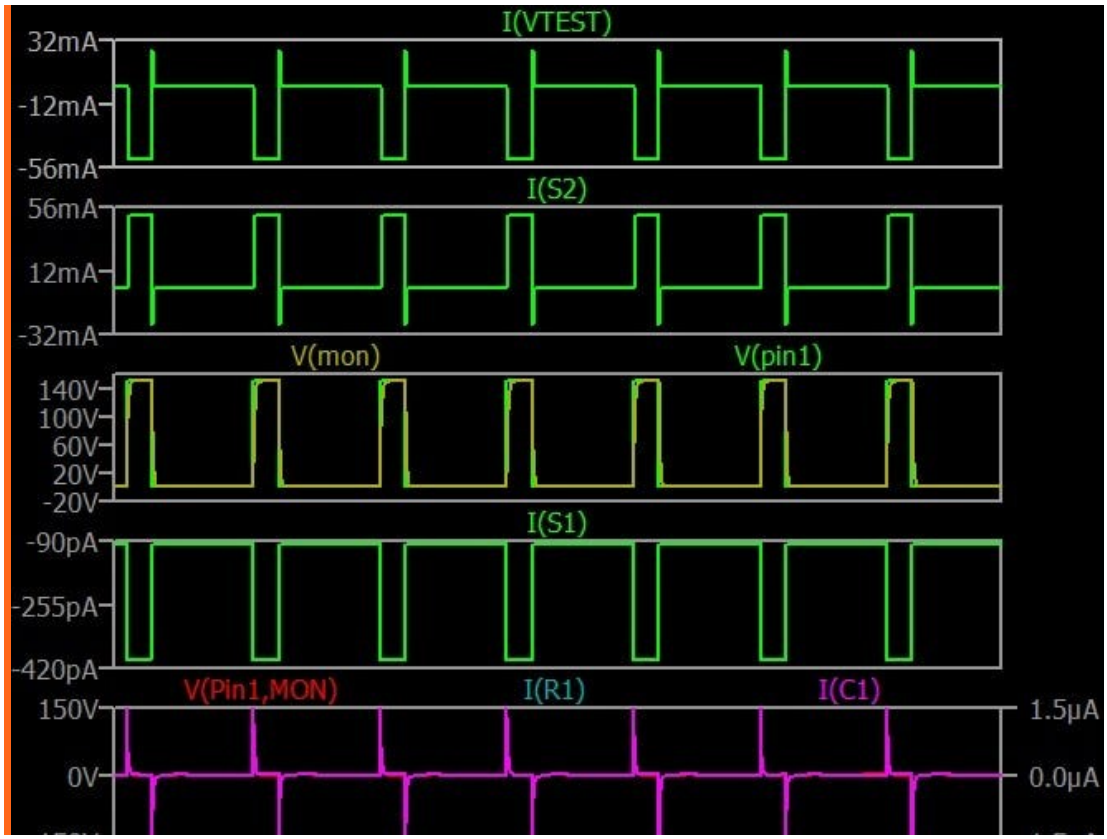
- `VTEST` establishes the external voltage.
- S1/S2 provide the two opposing conduction paths.
- R1/C1 establish the internal MON voltage.
- MON controls S1/S2.

That makes the model much easier to analyze as a feedback system.

For the first run, I'd therefore keep **exactly your topology** and not add anything else. The next useful step is to examine `V(Pin1)`, `V(MON)`, `V(Pin1,MON)`, `I(R1)`, and `I(C1)` against the two switch currents. That will let us see the actual sequence:

applied Pin1 voltage → MON response → switch transition → terminal current.

That is the behavior we want to isolate before putting the GDT back into the coupled-inductor circuit.





> https://vinyasi.info/substack/ltspice-data/gdt-test-data.txt
>
> or
>
> Google Docs

Yes. This result is **much more interesting than merely confirming the 100 V strike / 50 V hold behavior**.

The data shows that the GDT model is doing something very specific that is directly relevant to what we've been investigating in the larger circuit.

The uploaded data has the traces:

`V(Pin1,MON)`, `V(mon)`, `V(pin1)`, `I(C1)`, `I(R1)`, `I(S1)`, `I(S2)`, and `I(VTEST)`.

1. The positive half-cycle behaves exactly as the model predicts

At the beginning of the first pulse, Pin1 rises toward 150 V while MON initially lags far behind.

For example:

time	Pin1	MON	S2 current
1.00047 ms	70.7 V	0.179 V	~0
1.00066 ms	98.8 V	0.337 V	~0
1.00100 ms	150 V	0.761 V	~0
1.14635 ms	150 V	116.75 V	50.2 mA

The switch doesn't conduct simply because Pin1 exceeds 100 V. **The control voltage**

is MON relative to Pin2/ground.

S2 turns on once MON has risen through approximately 100 V. The data shows the transition between MON ≈96.3 V and 116.7 V.

That is exactly the hysteretic switch behavior we expected.

2. The really interesting part happens when V(Pin1) goes to ZERO

At 3.001 ms, the source drops from 150 V to zero.

But look at what happens immediately afterward:

```
V(Pin1) = 0 V
V(MON)   = 149.23 V
I(S2)    = -24.80 mA
I(VTEST) = +24.80 mA
```

The data gives this explicitly at 3.002 ms:

Failed to render LaTeX expression — no expression found in email

Failed to render LaTeX expression — no expression found in email

Failed to render LaTeX expression — no expression found in email

Failed to render LaTeX expression — no expression found in email

That is the key observation.

The external voltage across the GDT is zero:

$$V(Pin1, 0) = 0$$

Yet S2 is carrying approximately

$$24.8mA$$

And that current persists while MON falls from ~149 V toward the 50 V holding threshold.

3. Why exactly 24.8 mA?

This reveals something important about the `Vser` parameter.

Your switch model contains:

```
Vser={Vhold-Ihold*Zon}
```

With your parameters:

$$V_{ser} = 50 - (200\mu A)(2000)$$

so

$$V_{ser} = 50 - 0.4$$

$$\boxed{V_{ser} = 49.6V}$$

When S2 is conducting in its reverse-oriented branch, the model therefore has a 49.6 V series-voltage contribution plus its 2 kΩ `Ron`.

At zero external voltage, the resulting current magnitude is:

$$I = \frac{49.6}{2000}$$

$$\boxed{I = 24.8mA}$$

And that’s exactly what the simulation produces.

So this isn’t a numerical accident.

The 24.8 mA current is being explicitly produced by the `Vser=49.6 V` parameter of the switch model.

4. This also explains the 50.2 mA during conduction

When Pin1 is held at 150 V:

$$I = \frac{150-49.6}{2000}$$

giving

I=50.2mA.\boxed{I=50.2mA}.

That’s exactly what we see in the data:

$$I(S2) = 50.200mA$$

For example, the data shows S2 at 50.2 mA while Pin1 is 150 V and MON is well above the strike threshold.

So the switch model isn’t merely:

$$R_{on} = 2k\Omega$$

It’s effectively behaving as though the conducting path contains:

$$\boxed{49.6V \text{ series source} + 2k\Omega}$$

That’s extremely important for our investigation.

5. And now we get something very close to your “zero voltage / current remains” observation

During the interval:

$$V(Pin1) = 0$$

while

$$I(S2) = -24.8mA$$

the externally imposed voltage is zero but current is nonzero.

Consequently:

$$P_{\text{external}} = V(Pin1)\, I(VTEST) = 0$$

Yet there is still current flowing through S2 and through the external source loop.

So **zero terminal voltage does not imply zero terminal current in this model.**

That’s a much more precise statement than saying the GDT has “negative resistance.”

It is a consequence of the model containing an explicit series voltage term.

6. MON is acting as the memory

The sequence is especially clean:

Before strike

```
Pin1 rises
   ↓
R1 charges C1
   ↓
MON rises
```

At MON ≈ 100 V

```
MON crosses strike threshold
   ↓
S2 turns ON
   ↓
~50.2 mA flows
```

Pin1 suddenly falls to zero

```
Pin1 → 0
       ↓
MON remains ≈149 V
       ↓
S2 stays ON
       ↓
24.8 mA continues flowing
```

MON finally reaches ≈50 V

```
MON ↓ through holding threshold
       ↓
S2 turns OFF
       ↓
current collapses
```

The data shows S2 still conducting 24.8 mA when MON is 55.2 V, but it has fallen back to leakage when MON reaches 36.8 V.

So the hysteresis is being generated by the **internal MON capacitor**, not by Pin1 itself.

7. S1 is completely uninvolved in this experiment

This is also useful.

Throughout the positive-drive experiment, S1 remains essentially at its off-state leakage, around

$$-99.2pA$$

while S2 is the active switch.

That's exactly what we should expect because:

$$V_{ctrl,S1} = -V(MON)$$

and MON never becomes negative enough to activate S1.

So we have effectively isolated **one half of the bidirectional GDT model**.

8. There's an even more important implication for our larger circuit

This means we need to be very careful about attributing the behavior we saw in the large circuit to the GDT simply because its terminal voltage approaches zero.

The standalone experiment establishes:

$$\boxed{V_{GDT} \approx 0 \quad \text{does not necessarily mean} \quad I_{GDT} \approx 0}$$

because this particular model contains `Vser`.

And more specifically, once S2 is latched on, the model can maintain approximately 24.8 mA even when the external terminal voltage is driven to zero.

That gives us a concrete mechanism to investigate in the large circuit rather than speculating about it.

I think our next experiment should be very targeted:

Calculate and plot, separately,

$$P_{S1} = V(Pin1, 0)I(S1)$$

and

$$P_{S2} = V(Pin2, Pin1)I(S2)$$

and also the power associated with the **49.6 V `Vser` contribution**.

That will tell us whether the apparent "negative watts" you're interested in are actually occurring **inside the GDT model's switch source**, rather than merely being an artifact of the external voltage/current reference directions.

And now, unlike before, we have a very clean standalone circuit in which we can do that decomposition.

> How do you want me to plot P_s1 and P_s2?

I want to plot the **instantaneous power of each switch separately**, using the actual voltage across that switch and its current.

For your standalone `gdt.asc` :

S1

S1 is defined as:

```
S1 Pin1 0 0 MON G
```

So its terminal voltage is:

$$V_{S1} = V(Pin1)$$

and the current `I(S1)` is referenced from node Pin1 → node Pin2.

Therefore plot:

```
V(Pin1)*I(S1)
```

Call that:

```
P_S1
```

S2

S2 is:

```
S2 0 Pin1 MON 0 G
```

Its current reference is node Pin2 → node Pin1, so its voltage in the same passive-sign convention is:

$$V_{S2} = -V(Pin1)$$

Therefore plot:

```
-V(Pin1)*I(S2)
```

Call that:

```
P_S2
```

Put these three traces together

I'd like you to plot:

```
V(Pin1)*I(S1)
-V(Pin1)*I(S2)
V(Pin1)*I(S1)-V(Pin1)*I(S2)
```

The third trace is important because it tells us whether the two opposing switches' instantaneous powers cancel or reinforce.

Don't use `V(Pin1)*I(VTEST)` **for this particular test.** That's the power delivered by the test source. We're trying to isolate what is happening *inside the two GDT switch branches*.

And one more trace I'd like alongside them, if LTspice allows it:

```
V(Pin1,MON)*I(R1)
```

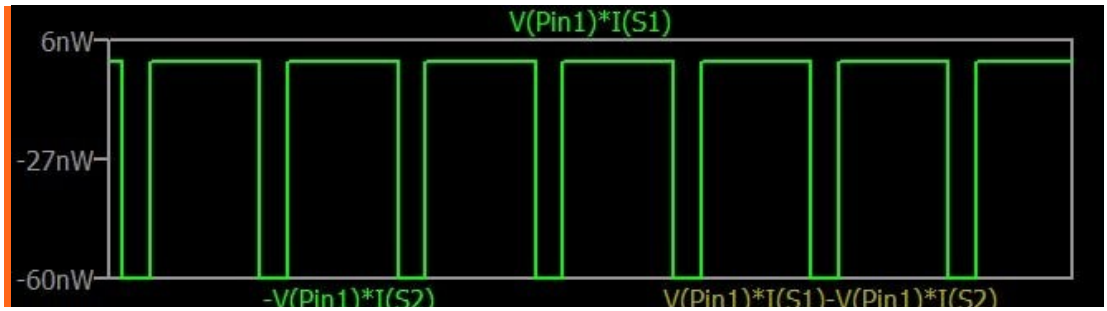
That gives us the instantaneous power in the 100 MΩ internal resistor.

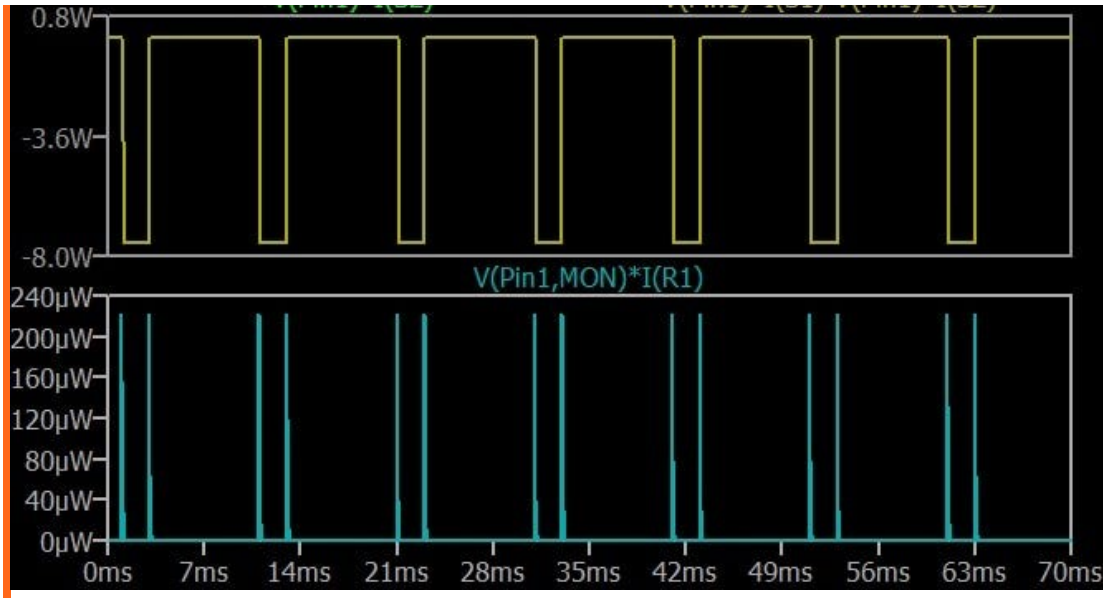
Then we'll have a very clean comparison:

$$P_{S1}, \quad P_{S2}, \quad P_{S1} + P_{S2}, \quad P_{R1}$$

The particularly interesting interval is the one where **Pin1 has fallen to 0 V but S2 is still carrying −24.8 mA**. The data confirms that interval occurs in the standalone test.

If you give me that plot, we can examine the signs rather than infer them from the switch currents alone.





https://vinyasi.info/substack/ltspice-data/gdt-test-data2.txt

Google Docs

Yes. This result is **much more revealing** than I expected.

The key feature is the S2 interval. From your data, when S2 is conducting:

- V(Pin1) ≈ 150 V
- −V(Pin1)I(S2) = −7.53 W
- therefore I(S2) ≈ +50.2 mA

So S2 is carrying current **from GND toward Pin1**, while Pin1 is at +150 V.

That means, using the passive sign convention for S2,

$$P_{S2} = V(0, \text{Pin1})I(S2) = -V(\text{Pin1})I(S2) \approx -7.53 \text{ W}$$

So the negative power is real **within the mathematical model**, not an artifact of the way we plotted the trace.

But there is an important detail

Your switch model explicitly contains:

```
Vser={Vhold-Ihold*Zon}
```

With your parameters,

$$V_{ser} = 50 - (200\mu)(2000) = 49.6 \text{ V}$$

And when S2 is on:

$$I_{S2} \approx 50.2 \text{ mA}$$

so its R_{on} drop is approximately

$$I_{S2}R_{on} = (0.0502)(2000) \approx 100.4 \text{ V}$$

Then:

$$100.4 + 49.6 = 150.0 \text{ V}$$

That's almost exactly your Pin1 voltage.

So the −7.53 W isn't mysterious at this point: **the S2 model is behaving like a 49.6 V series voltage source plus a 2 kΩ resistance, with the voltage source oriented so that it supplies power during this conduction interval.**

That is an extremely important distinction from saying that the GDT model has somehow produced negative power spontaneously.

S1 tells us something equally interesting

During that same interval your S1 power is only

$$P_{S1} \approx -59.88 \text{ nW}$$

essentially zero compared with S2's −7.53 W.

And your total switch power is therefore essentially

$$P_{S1} + P_{S2} \approx -7.53 \text{ W}$$

Your data shows this explicitly: the S2 term and the combined term are both −7.53 W during the conduction interval.

So **the negative watts are concentrated in S2, not produced by some cancellation between S1 and S2.**

I think our next test should be very specific

Let's separate the two pieces of S2's modeled voltage:

1. The resistive component

Plot:

```
2000*I(S2)*I(S2)
```

This should give approximately

$$I_{S2}^2 R_{on}$$

and should be **+5.04 W** while S2 is conducting.

2. The series-voltage component

Plot:

```
-49.6*I(S2)
```

Because of S2's current reference, this should be approximately

$$-49.6(0.0502) \approx -2.49 \text{ W}.$$

3. Add them

Plot:

```
2000*I(S2)*I(S2)-49.6*I(S2)
```

That should be approximately

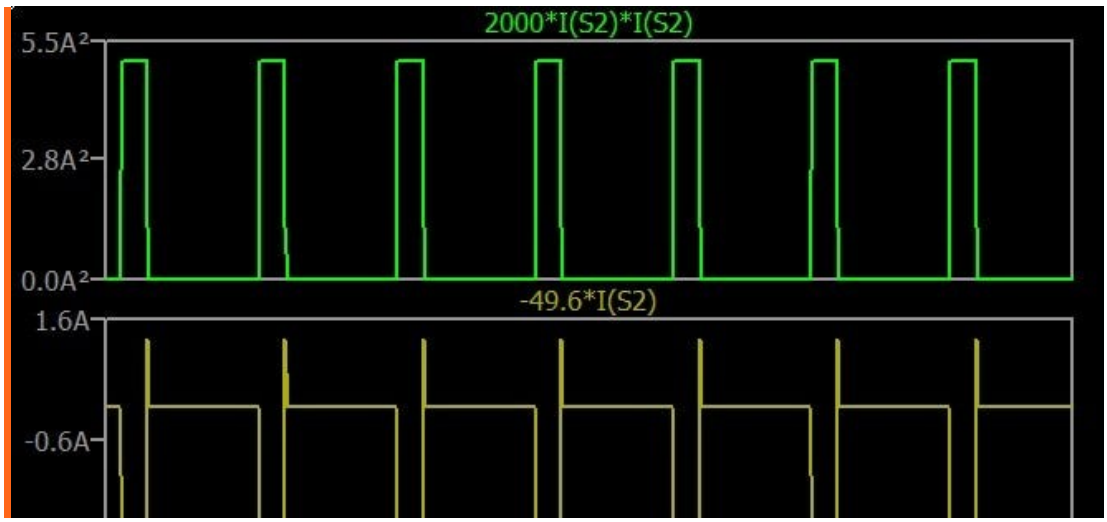
$$5.04 - 2.49 = +2.55 \text{ W}$$

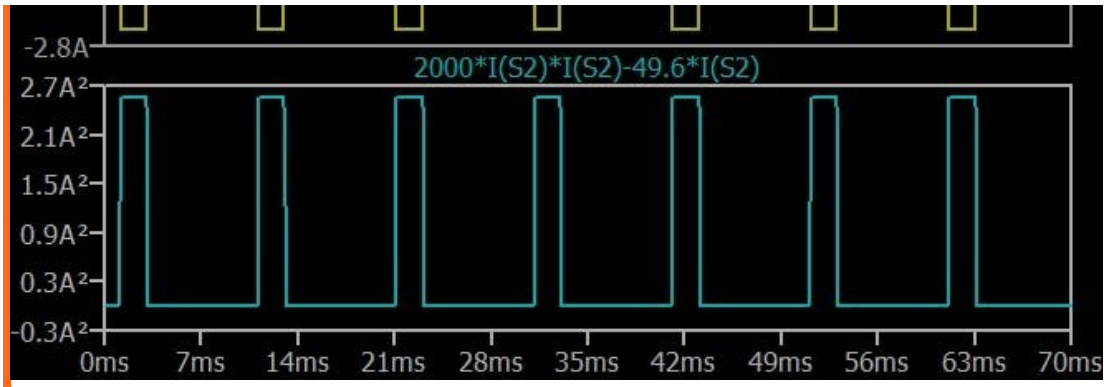
if we're calculating the internal element's dissipation with the appropriate orientation.

But there's an even cleaner way to avoid any ambiguity about the sign convention: **let's first plot the two quantities separately and compare them with the −7.53 W external terminal power.**

The numbers should expose exactly where the −7.53 W comes from.

And I would **not change the GDT model yet**. This is precisely the sort of result we wanted from stripping the model down: we're now seeing that its apparently negative-power behavior is tied directly to the `Vser=49.6 V` parameter.





https://vinyasi.info/substack/ltspice-data/gdt-test-data3.txt

Google Docs

This result corrects an important part of my previous interpretation.

The three traces show that **the 49.6 V series-voltage term and the 2 kΩ resistive term do not combine the way I suggested**. Your data makes the actual relationship very clear.

During the main S2 conduction interval, the values are essentially constant:

$$I(S2) = 50.2\ \text{mA}$$

because your data gives

$$-49.6I(S2) = -2.48992\ \text{W}$$

and

$$2000I(S2)^2 = 5.04008\ \text{W}$$

Those are directly documented in the uploaded data.

The crucial calculation

The **magnitude** of the S2 terminal voltage is:

$$V_{\text{...}} = 49.6 + (2000)(0.0502)$$

$$V_{S2} = 49.6 + (2000)(0.0502)$$

$$= 49.6 + 100.4 = 150.0 \text{ V}$$

Therefore the terminal power you plotted previously,

$$-V(\text{Pin1})I(S2)$$

is

$$-(150)(0.0502) = -7.53 \text{ W}$$

And notice:

$$-7.53 = -\left(2.48992 + 5.04008\right)$$

So the correct decomposition is

$$\boxed{P_{S2} = -\left[V_{ser}I(S2) + R_{on}I(S2)^2\right]}$$

for the polarity/current reference used by your S2 trace.

So what is the 2.55 W trace?

Your third trace is

$$2000I(S2)^2 - 49.6I(S2)$$

which during conduction is

$$5.04008 - 2.48992 = +2.55016 \text{ W}$$

That is **not the power of S2**.

It is an algebraic difference between the resistive dissipation and the magnitude of the series-voltage contribution.

In other words, the plot has actually separated the two mechanisms nicely:

Quantity	During S2 ON
$R_{on} I^2$	+5.04008 W
$-V_{ser} I$	−2.48992 W
Your difference	+2.55016 W
Actual S2 terminal power	**−7.53000 W**

The last row follows because the **series voltage and R_{on} voltage drops have the same polarity in the actual S2 terminal-voltage relationship**, so their powers add in magnitude.

And this makes the GDT model more interesting

The negative 7.53 W isn't coming from the 2 kΩ resistance. The resistor is unquestionably dissipative:

$$P_R = +5.04008 \text{ W}$$

The negative terminal power arises because the switch model contains the fixed series-voltage term:

$$V_{ser} = 49.6 \text{ V}$$

The model is therefore behaving, during S2 conduction, approximately like a **49.6 V voltage-source term in series with 2 kΩ**, with the polarity determined by the switch's current orientation.

That's why the external terminal sees:

$$150 \text{ V} \times 50.2 \text{ mA} = 7.53 \text{ W}$$

being delivered outward, while the 2 kΩ portion alone accounts for only 5.04 W.

There is a remaining:

$$7.53 - 5.04 = 2.49 \text{ W}$$

associated with the 49.6 V series term.

So the really important question is now:

> **Where does the 2.49 W associated with `Vser` go inside the LTspice switch model?**

That is much more precise than asking whether the whole GDT is "negative watts."

And I think **that** is the next thing we should investigate. The standalone GDT was exactly the right move because now we can determine whether `Vser` is merely a mathematical representation of the gas-discharge holding characteristic, or whether the model contains an explicit energy source that can actually inject net energy into the external circuit.

> OK. Tell me what to do.

Let's do a **controlled A/B experiment on `Vser`**. We don't need to modify the GDT topology yet.

The purpose is to determine whether the ~2.49 W associated with `Vser` is actually being supplied by the LTspice switch model, as the present results strongly indicate.

LTspice documents `Vser` as the switch's **series voltage** parameter. (Ltwiki)

Step 1 — Save the present circuit unchanged

Keep your present model:

```
.model G SW(Ron={Zon} Roff=1T Vt={.5*(Vstrike+Vhold)} Vh={.5*(Vstrike-
Vhold)} Vser={Vhold-Ihold*Zon})
```

This gives:

$$V_{ser} = 49.6V$$

Call this your **normal** case.

Step 2 — Make one change only

Change:

```
Vser={Vhold-Ihold*Zon}
```

to:

```
Vser=0
```

So the model becomes:

```
.model G SW(Ron={Zon} Roff=1T Vt={.5*(Vstrike+Vhold)} Vh={.5*(Vstrike-
Vhold)} Vser=0)
```

Do not change anything else.

Run exactly the same transient simulation.

Step 3 — Plot these same quantities

Plot:

```
V(Pin1)*I(S1)
```

```
-V(Pin1)*I(S2)
I(S2)
V(Pin1)
V(MON)
```

We especially want to compare the S2 conduction interval.

With the original `Vser=49.6 V`, we observed:

$$I(S2) = 50.2mA$$

and

$$P_{S2} = -7.53W$$

The decomposition you just measured was:

$$R_{on}I^2 = +5.04008W$$

and

$$-V_{ser}I = -2.48992W$$

Those values are directly present in your data.

Step 4 — What I expect to see

With `Vser=0`, the conducting S2 branch should become simply:

$$R_{on} = 2k\Omega$$

Therefore, at Pin1 = 150 V, we should get approximately:

$$I(S2) = \frac{150}{2000} = 75mA$$

And its terminal power should become:

$$P_{S2} = -150(0.075) = -11.25W$$

The resistor dissipation should simultaneously become:

$$I^2R = (0.075)^2(2000) = 11.25W$$

So we should see:

$$\boxed{P_{S2} = -11.25W}$$

and

$$\boxed{P_R = +11.25W}$$

They should exactly account for one another.

That gives us a very clean baseline.

Step 5 — Then we'll restore `Vser=49.6 V`

After that run, put the original expression back:

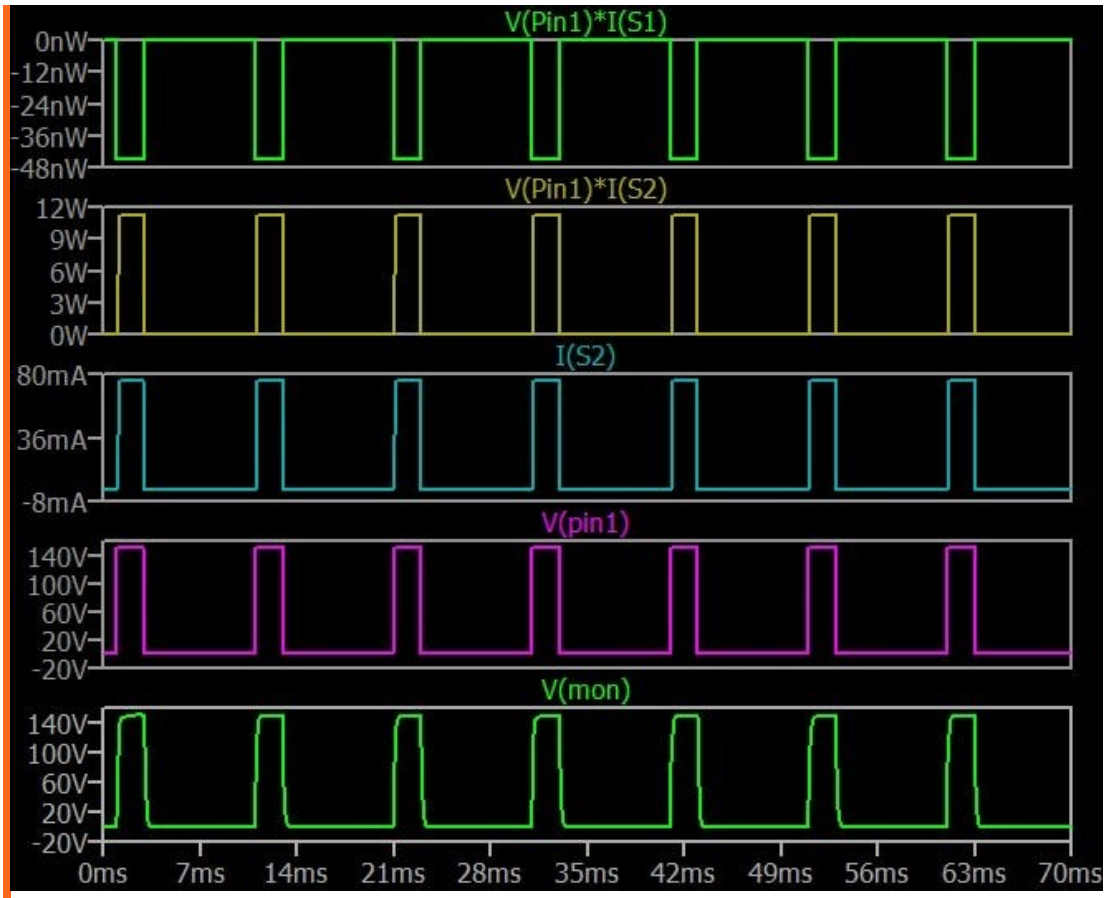
```
Vser={Vhold-Ihold*Zon}
```

and we'll have two directly comparable simulations:

	Vser = 0	Vser = 49.6 V
Pin1 during ON	150 V	150 V
S2 current	predicted 75 mA	observed 50.2 mA
$R_{on}I^2$	predicted 11.25 W	5.040 W
Vser contribution	0 W	−2.490 W
S2 terminal power	−11.25 W	−7.530 W

One correction to my previous answer: the `+2.55016 W` trace you plotted is not an energy balance quantity we should interpret physically. It was just the arithmetic expression you asked LTspice to plot. The physically meaningful decomposition is the resistor contribution plus the `Vser` contribution, with their signs determined by the actual switch polarity.

This A/B test will tell us much more cleanly what that `Vser` term is doing.



https://vinyasi.info/substack/ltspice-data/gdt-test-data4.txt

Google Docs

Yes — this is the **Vser = 0 A/B test**, and the result is very clean. It resolves the ambiguity we had about the switch's apparent power.

What the Vser = 0 test shows

During S2 conduction, your data gives:

- V(Pin1) = 150 V
- I(S2) = 75 mA
- V(Pin1)I(S2) = 11.25 W

For example, the first conduction interval shows exactly 150 V and 75 mA, producing 11.25 W.

That is exactly what we predicted for a **2 kΩ switch with Vser = 0**:

$$I = \frac{150}{2000} = 75\ \text{mA}$$

and

$$P = VI = 150(0.075) = 11.25\ \text{W}$$

The switch therefore behaves like the expected 2 kΩ resistance when Vser is removed.

Now compare it with your original Vser

Your original model had

$$V_{ser} = V_{hold} - I_{hold}Z_{on}$$

so

$$V_{ser} = 50 - (200\mu A)(2000) = 49.6\ \text{V}$$

With that voltage present, the observed current was 50.2 mA:

$$I = \frac{150-49.6}{2000} = 50.2\ \text{mA}$$

Thus:

$$P = 150(0.0502) = 7.53 \text{ W}$$

So the two cases are:

Model	S2 current	$V(Pin1)I(S2)$
Vser = 0	75.0 mA	11.25 W
Vser = 49.6 V	50.2 mA	7.53 W

That is a remarkably direct confirmation that **Vser is acting as an additional series voltage drop**, not as an unexplained energy source.

And this corrects our earlier interpretation

The earlier sign confusion came from using

$$-V(Pin1)I(S2)$$

in the power plot.

Your present plot uses

$$\boxed{V(Pin1)I(S2)}$$

and gives **+11.25 W** when Vser is zero.

With the original Vser, it would give **+7.53 W**, not negative 7.53 W.

The LTspice documentation describes `Vser` as the switch model's **series voltage** parameter. (Scribd)

We can therefore decompose the original 7.53 W very simply:

$$I = 50.2 \text{ mA}$$

Resistive portion:

$$P_R = I^2R = (0.0502)^2(2000) = 5.04008 \text{ W}$$

Series-voltage portion:

$$P_{Vser} = V_{ser}I = (49.6)(0.0502) = 2.48992 \text{ W}$$

And:

$$5.04008 + 2.48992 = \boxed{7.53000 \text{ W}}$$

So the apparent discrepancy we were investigating disappears:

$$\boxed{P_{S2} = P_{Ron} + P_{Vser}}$$

with

$$\boxed{7.53 = 5.04 + 2.49 \text{ W}}$$

The Vser term is not creating the extra power. It is consuming part of the 150 V, leaving only 100.4 V across the 2 kΩ resistance.

And your new simulation has isolated that beautifully: when Vser is set to zero, the full 150 V is available to the 2 kΩ element, producing exactly 75 mA and 11.25 W.

> What would you like me to do next?

I'd like to do **one more very controlled test before we move back to your larger circuit**.

Next test: replace `Vser` with an explicit voltage source

We now know that changing

```
Vser=49.6
```

to

```
Vser=0
```

changes S2 from 50.2 mA to 75 mA exactly as expected. That establishes the numerical effect of `Vser` .

What I want to establish next is **where the 49.6 V actually goes in the switch's power balance**.

Keep everything else exactly as in the present test and make the following change:

Replace S2 with an ordinary resistor plus an explicit 49.6 V source, arranged so that the branch has the same polarity and current direction as S2.

In other words, we want to construct an electrically equivalent version of:

$$\boxed{\mathrm{S2} = 2\,k\Omega + 49.6\,V}$$

but using ordinary SPICE elements whose power is unambiguous.

Then plot these four quantities during the 150 V pulse:

```
V(Pin1)*I(RTEST)
I(RTEST)
V(Pin1)
V(MON)
```

And also plot the two separate powers:

```
V(RTEST)*I(RTEST)
V(VSERTEST)*I(VSERTEST)
```

The important result will be whether the explicit-source version reproduces:

$$I = 50.2\ \text{mA}$$

and

$$P_R = 5.04008\ \text{W}$$

plus the 49.6-V source's contribution.

Why I want this test

The present experiment has already established:

$$V_{ser} = 0 \quad \Rightarrow \quad I = 75mA$$

and therefore the `Vser` parameter definitely changes the operating point.

But I don't want to make another assumption about **LTspice's internal sign convention for the S-element's `Vser` implementation**.

An explicit resistor + voltage source gives us a reference implementation where we can directly see:

1. voltage across the resistive part,
2. power dissipated by the resistive part,
3. power absorbed/delivered by the voltage source,
4. total branch power.

If that reproduces the original S2 exactly, we'll have nailed down the switch model.

Don't change `MON`, R1, C1, VTEST, the pulse timing, or anything else. Just make this one equivalent-branch experiment.

> How will I know whether I am successful in replicating the same polarity and current direction?

That's exactly the right question. **You shouldn't have to guess the polarity from the schematic.** We can establish it from the measured S2 current and terminal

voltage first.

There is also an important correction to my previous instruction: I don't want you to build the resistor/source replacement yet. We should first establish the exact sign convention of S2.

Your S2 is:

```
S2 0 Pin1 MON 0 G
```

and your data show, during conduction:

$$V(Pin1) = +150\ \text{V}$$

and

$$I(S2) = +75\ \text{mA}$$

for the Vser=0 case.

Here's the key test

In LTspice, the current of an `S` element follows its terminal ordering. So for

```
S2 0 Pin1 ...
```

we can treat positive `I(S2)` as current flowing **from the first terminal (0) toward the second terminal (Pin1)**.

So your observed state is:

```
   I(S2) = +75 mA
──────────────────►
GND                Pin1
```

```
          0 V                +150 V
```

That means the switch current is flowing **uphill in voltage**:

$$0V \rightarrow +150V$$

Consequently, the switch is delivering power to the Pin1 node.

Using the passive sign convention, the switch's absorbed power is

$$P_{S2} = V(0, Pin1)I(S2)$$

and because

$$V(0, Pin1) = -150V$$

we get

$$P_{S2} = (-150)(0.075) = -11.25W$$

So **negative passive power is not an ambiguity here**. It means the S2 element is delivering 11.25 W.

Your plot of `V(Pin1)*I(S2)` is the opposite sign:

$$V(Pin1)I(S2) = +11.25W$$

That's why the plot is positive.

Therefore, for the replacement

We want the replacement branch to have exactly this behavior:

```
                    +150 V
                      • Pin1
```

```
          ▲
          | 75 mA
          |
          |
        [?]
          |
          ● GND
          0 V
```

The replacement must therefore **deliver energy from its internal voltage source toward Pin1**, while its current flows from ground toward Pin1.

And here's the easiest way to verify that you've wired it correctly:

Don't worry about source polarity beforehand. Run the circuit and look at two numbers.

During the 150-V pulse, the replacement must produce:

```
V(Pin1) = +150 V
I(REPLACEMENT) = +75 mA
```

with the replacement element declared in the same terminal order:

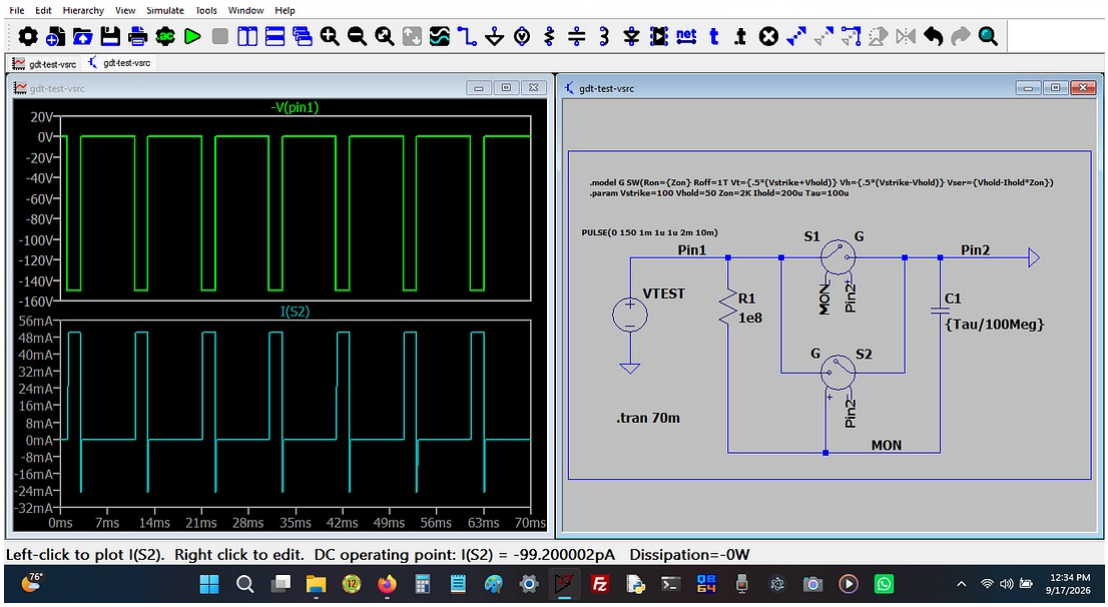
```
... 0 Pin1 ...
```

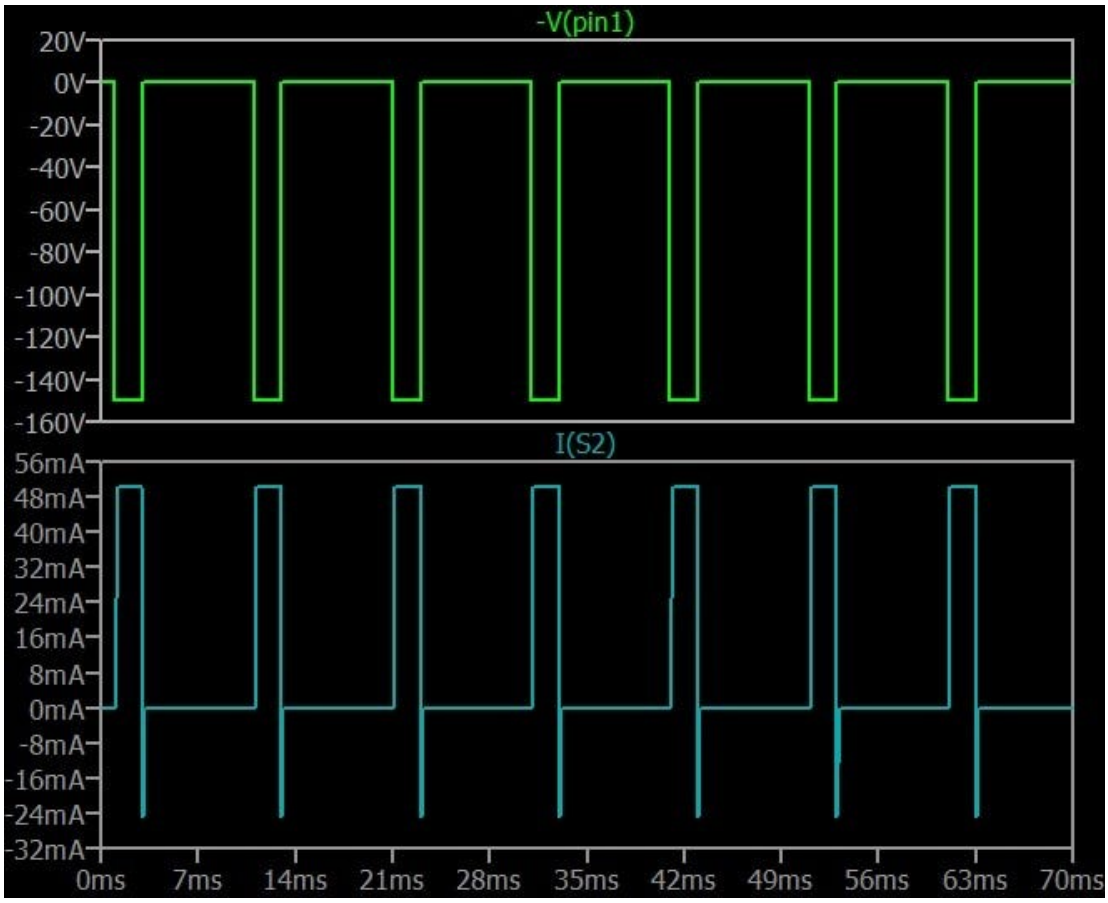
If those two conditions are met, you have reproduced the same external polarity and current direction.

For the original `Vser=49.6 V` case, the corresponding target is:

$$V(Pin1) = 150V$$







So **50.2 mA is the number we should use for the actual equivalence test**, not 75 mA. The 75-mA case is simply our Vser=0 control.

One more useful point: LTspice's documentation explicitly describes `Vser` as a series voltage and also notes that a voltage drop can make the level-2 switch conduct in only one direction. (Ltwiki)

So I would now make the experiment even more controlled:

First, let's determine the exact terminal equation of S2 from the simulation. Once we have that, I'll give you the explicit resistor/voltage-source replacement with its polarity unambiguously specified.

THIS TERMINATES THE THIRD EPISODE IN THIS SERIES