A Vacuum Tube Tester For the Rest of Us

Ian Lacy & Leif Clark
6.101 - Analog Electronics Lab
May 16, 2023

A Vacuum Tube Tester For The Rest Of Us

Abstract

In order to ameliorate the undue expense imposed on vacuum tube (VT) enthusiasts by the unfortunate scarcity of tube testing devices, we have designed and constructed a mutual conductance tube analyzer with the following functionality. The analyzer places the tube in its intended operating regime and electronically compares the input and the output of the tube amplifier. The mutual conductance value of the tube is easily measured as a voltage at the output of our implemented device.

Introduction

Vacuum tubes were once the cutting edge of scientific invention. They revolutionized communication, and set the stage for the information revolution when silicon took over. Now, however, they've become obsolete for all but a few applications. Regardless, they remain a common hobby among electrical engineers, and are a critical piece of engineering history. Unfortunately, vintage vacuum tube equipment is becoming rarer and more expensive, meaning tube testing and measurement is a more and more niche process. Our project aimed to aid that by making a vacuum tube tester and characterizer to lower the barrier of entry to vacuum tube hobby projects as well as to keep this incredible piece of history alive. The fundamental goal of our project was to design and build a vacuum tube tester capable of displaying the mutual conductance at multiple discrete operating points.

Background

Vacuum tubes, as the name suggests, consist of a number of electrodes-- from one (the lightbulb) to over a dozen (a multi-stage photomultiplier tube)-- enclosed in a vacuum. The vacuum chamber is glass but potentially metal-shielded. Vacuum tubes operate via a phenomenon known as thermionic emission. When a metal is heated up in a vacuum, electrons essentially 'boil' away. In the presence of an electric field, a free electron will accelerate against the field. This results in a current! A simple two-terminal device, a *di-ode* can be made in this fashion- a heated cathode and a high voltage anode, also known as the 'plate', will allow current to flow in only one direction, as the plate isn't heated, preventing electrons from escaping. If we put an electrode with holes (a control grid) in between the anode and cathode, we can control the current. By driving the control grid to a lower voltage than the cathode, we can turn the electrons boiling off of the cathode back towards it, turning the tube off. If the voltage is at or above the cathode voltage, the tube will be fully on! Adding even more electrodes will change the characteristics of the tube further, but there's enough material there for multiple textbooks.

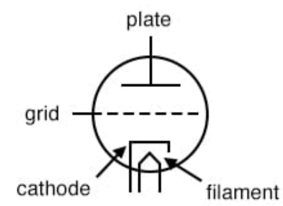

These sorts of vacuum tubes have a characteristic, known as mu, or the mutual conductance, that essentially tells you how well the tube can act as an amplifier. Much like with transistors, some uses require that tubes have a closely matched mu for best performance. Just like with transistors, we can measure this characteristic by putting the tube in a known setup! This is the core of the project.

Block Diagram

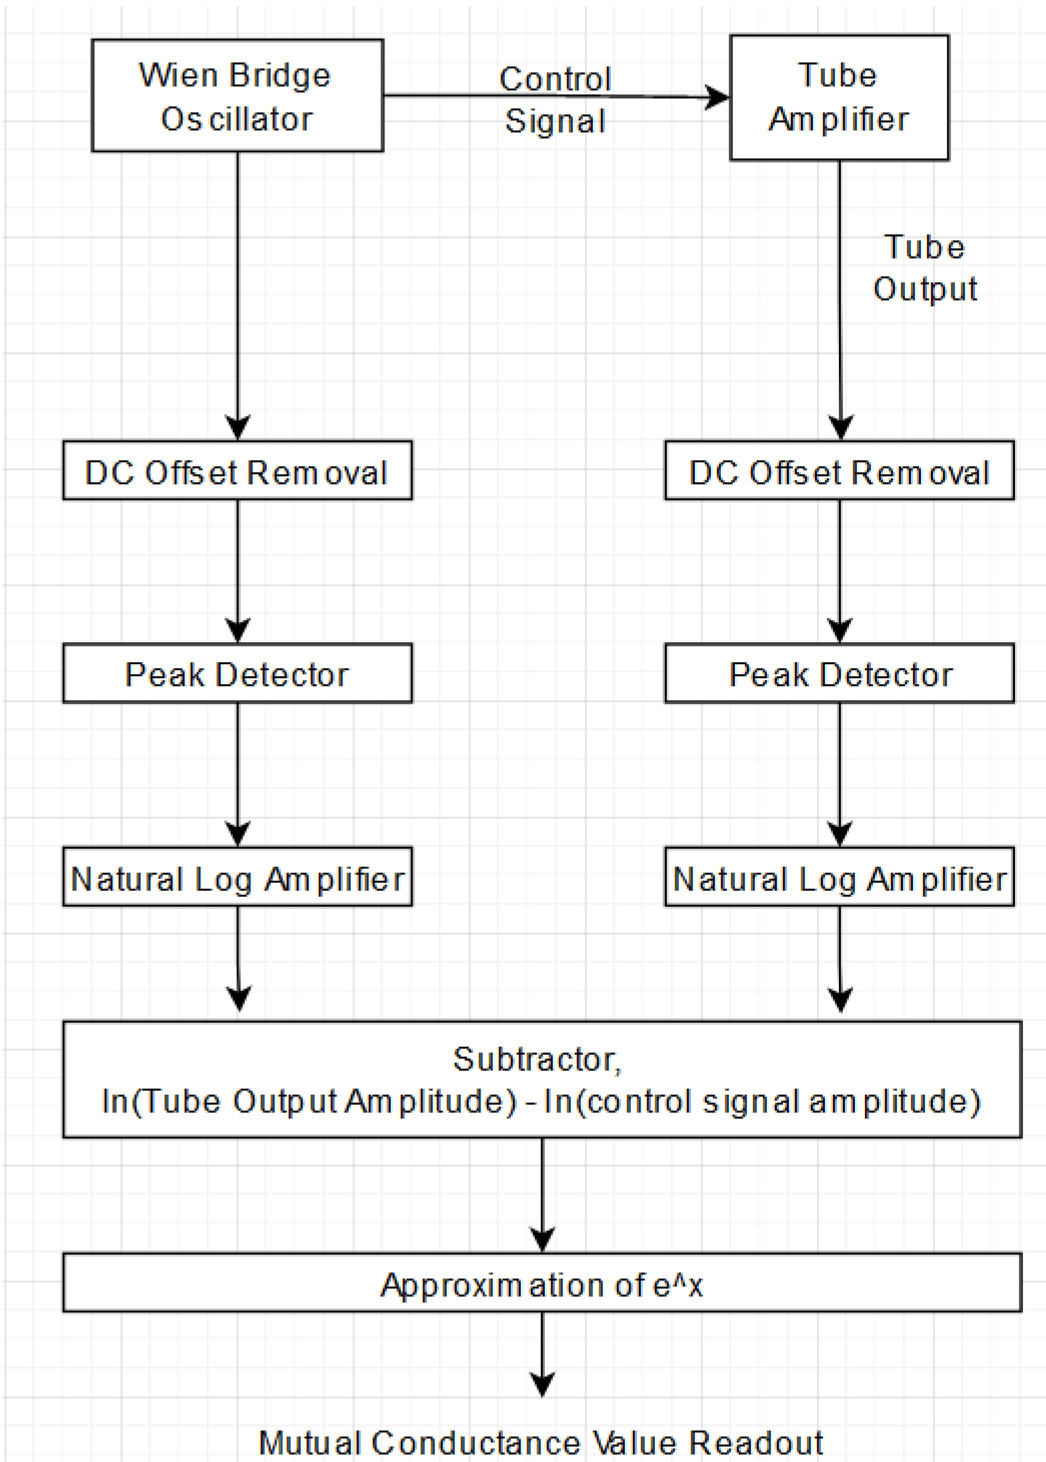

Figure 1. *Block Diagram of Implemented Tube Tester System*

The Amplifier

To measure the amplification of a signal, one first needs to generate a signal. This is done with a canonical Wien Bridge oscillator circuit, tuned to a frequency of 1kHz. Since the output swings between the rails by default, the signal is attenuated to 1 Vpp through a simple potentiometer. It is then fed into an op-amp subtractor circuit to apply an offset of -2 VDC to the signal. This signal is passed to the grid, the middle electrode, of the triode. The plate of the triode is connected to +30 VDC through a 15k resistor, and the cathode of the triode is connected to ground. The output is measured between the 15k resistor and the plate, and for this setup it is typically a 5 Vpp sine wave riding on approximately a +10 VDC offset. The exact magnitude and offset of the output depends, of course, on the characteristics of the tube, and in particular the mutual conductance, the goal of measurement. Though mu is a bit of a nebulous concept, depending on the exact setup of the amplifier, in the same setup two matched tubes will behave the same. So though the measurement is not particularly meaningful on its own, two tubes having a very close measurement in the same setup will behave very similarly in another setup, making them suitable for a differential amplifier pair, or a left/right channel pair. This is why matching mu is important. Additionally, most tubes have an expected mu for certain operating conditions, meaning that you can roughly gauge the health of the tube by measuring it. If the measured mu is drastically low, the tube may need replacement.

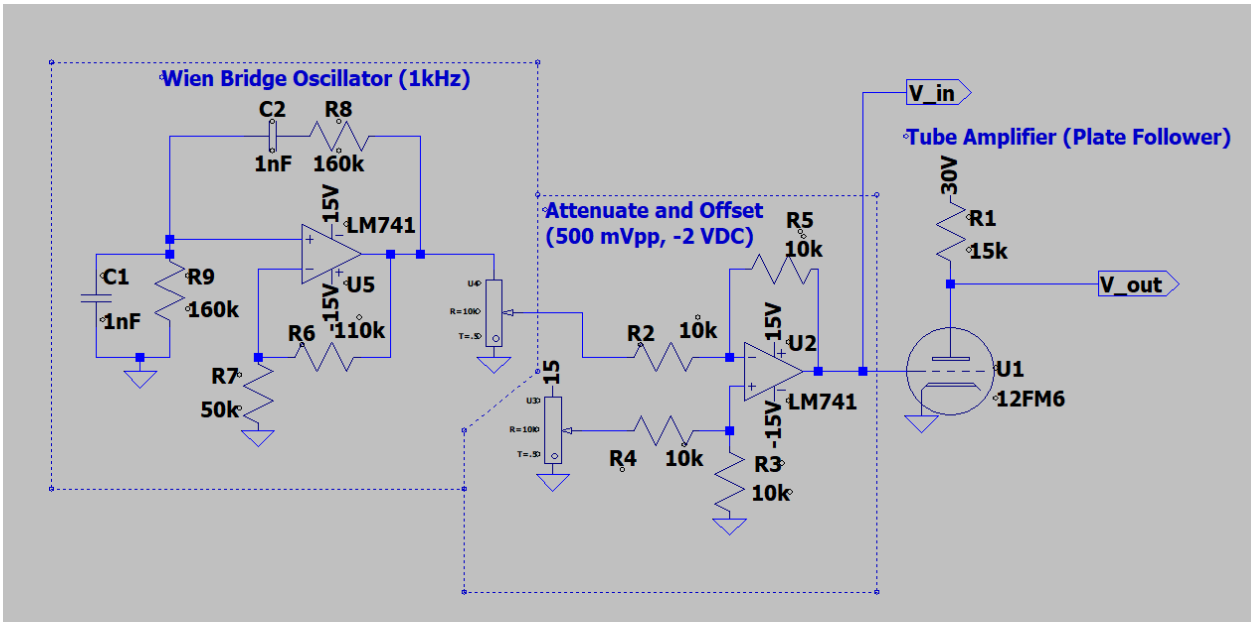

Figure 2. *The Wein Bridge Oscillator, Voltage Offset, and Plate-Follower Amplifier.*

The Division Block

Dividing the tube amplifier's output voltage by its input voltage is fundamental to determining our parameter of interest, mutual conductance. Mutual conductance is a measure of the ratio of the amplitudes of the signals going into and out of the tube amplifier. However, division poses a unique challenge in an analog system. Canonical, reliable topologies exist for addition, subtraction, integration, and derivation. There is not a simple, widely implemented division topology. Notably, there is a topology for division that is pedagogically popular, but very seldom actually implemented for reasons we encountered.

Although more complex (and possibly more accurate) solutions exist, we placed two constraints on our project. First, we wanted to conduct the mathematical manipulation of the values with an analog circuit. We imposed this constraint out of our own interest in the subject, and a desire to fulfill the spirit of a class entitled "Analog Electronics Lab." Second, we decided to limit our choice of components to those that could be considered reasonably obtainable and inexpensive. The inspiration for our project was the financial inaccessibility of tube testing equipment, and we wanted to design a device that addresses that problem. The self-imposed constraint to work with easily available and cheap components was a major driving force between our decisions to rely heavily on LM741 op-amps and avoid log amplifier topologies which required super-matched transistor pairs.

A quick internet search reveals a seemingly straightforward circuit topology for the division of two DC values:

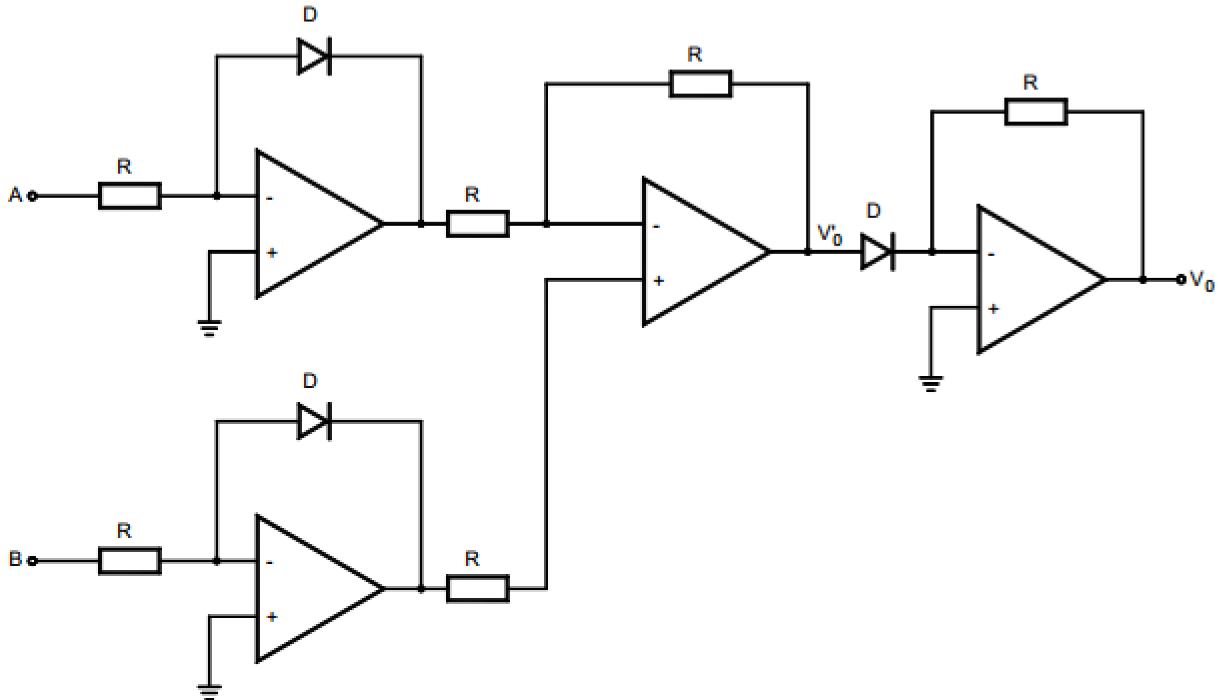

Figure 3. *A Pedagogically Popular Division Circuit*

The circuit in Figure 3 is intended to output the quotient of A and B, based off of this equation:

`A/B = e^(ln(A/B)) = e^(ln(A)-ln(B))`

The two op-amps on the left are implemented to present `ln(A)` and `ln(B)` at their output. The op-amp at the center is implemented to subtract `ln(B)` from `ln(A)`. The op-amp on the right is implemented to output e`^(V'o)`, where V'o is, in this case, equivalent to `ln(A) - ln(B)`. We used this pedagogically popular division circuit as a basis for our divider, and quickly encountered operational limitations which demonstrated why this circuit is popular as a teaching tool, but rare in actual implementation.

The natural log amplifiers shown above rely on the nonlinear I-V characteristic of the diode. However, the actual relationship between the input and the output of the log amplifier topology shown above is not `V_out = ln(V_in)`. A more detailed discussion of why this occurs and how we overcame it is offered above, but here it is sufficient to say that the hypothetically functional circuit in Figure 3 was woefully inappropriate for our purposes. Due to these issues, the actual implementation of the division block was significantly different from the hypothetically functional circuit in Figure 3.

Below is a schematic and a walkthrough of the function of the division block, as implemented.

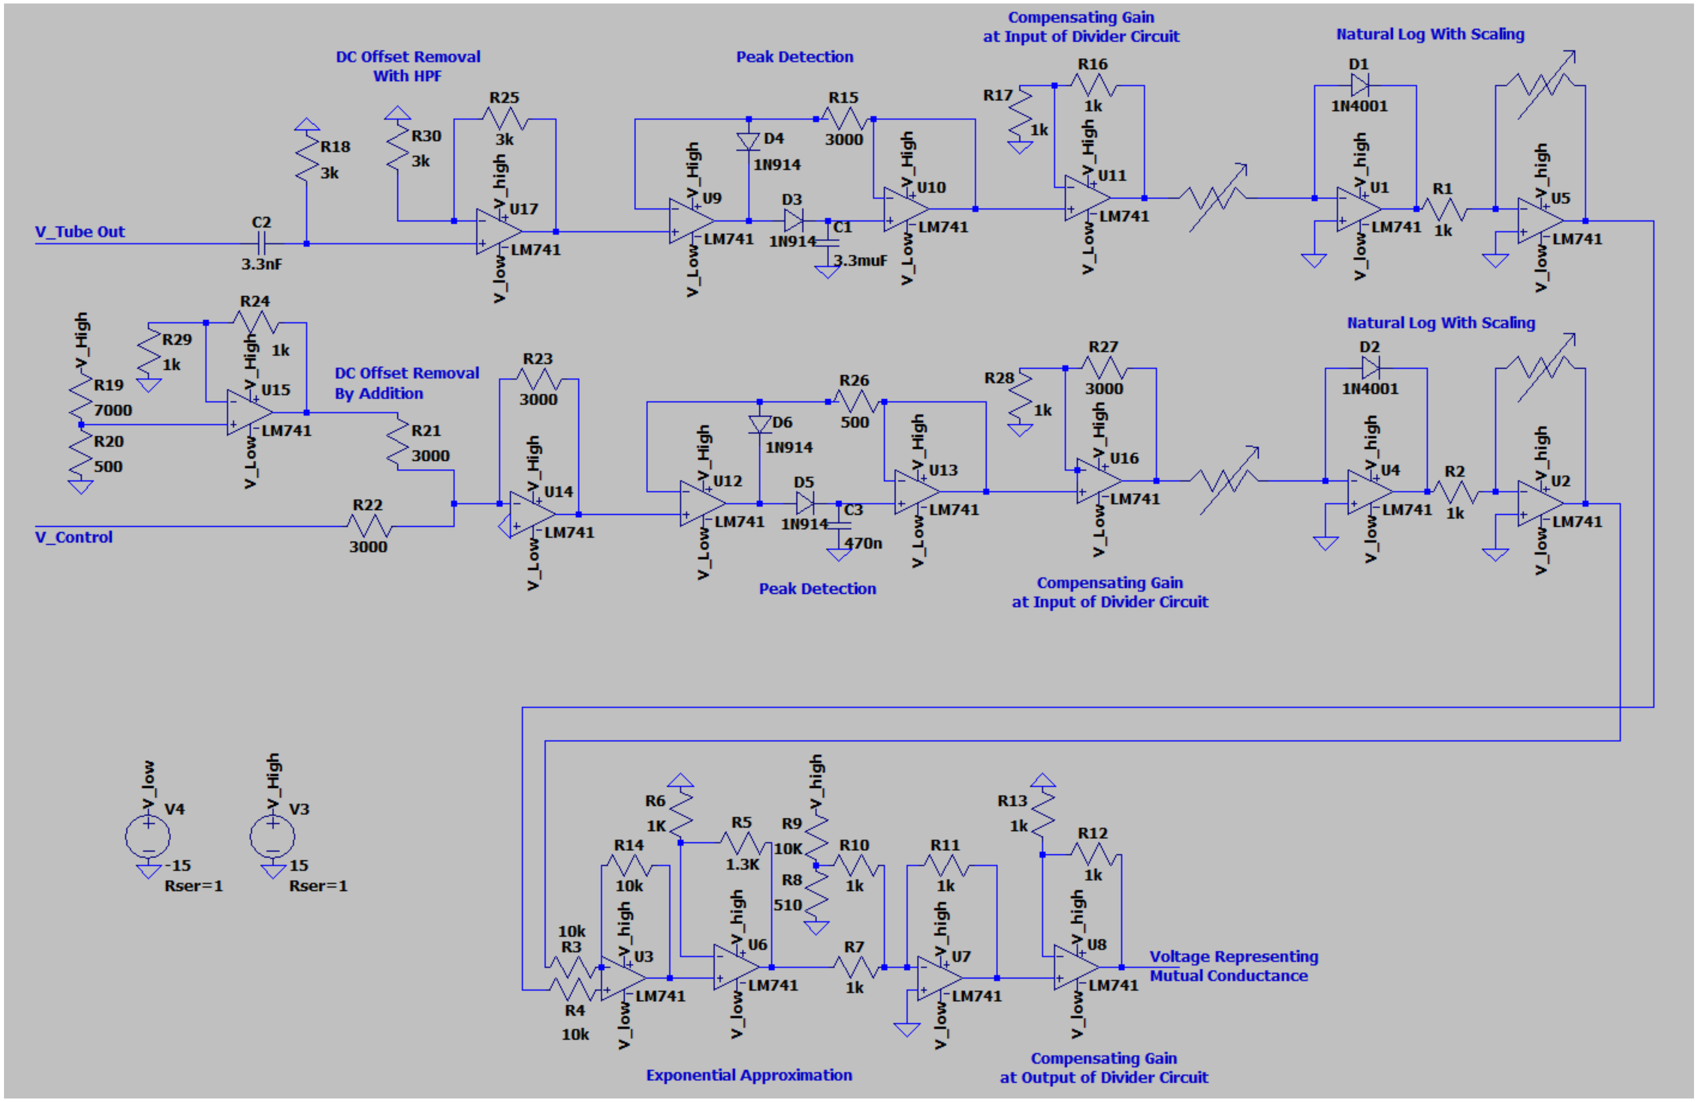

Figure 4. *Schematic for the division block. Any following device identifiers are in reference to this circuit*

The control signal and the tube output cannot be directly connected to the division circuit. The division circuit is designed to divide two DC values, each of which represents the amplitude of the AC signal going into or out of the tube amplifier. Because of this, we must convert these AC signals into DC values. Two circuit blocks are used in this process: DC offset removal and peak detection.

DC offset removal is accomplished with two different methods. The DC offset of the control signal is -2VDC by design, and is removed with an op-amp adder topology (U14, U15) which adds 2VDC to the signal to ensure that it has no DC offset before peak detection. DC offset removal for the tube amplifier output is accomplished differently. Because the DC offset of the tube output is not completely predictable (like the DC offset of our designed control signal), it cannot be removed with the subtraction of a constant DC value. The DC offset of the output of the tube amplifier was removed via a passive RC (R18, C2) high pass filter.

The division circuit cannot reliably divide two AC inputs because any difference in phase between the inputs creates asymptotic behavior in the output where the signals approach zero-crossings. Although applying a known DC bias to both signals does prevent asymptotic behavior near zero-crossings, the result is still periodic in time and is not representative of the ratio of the amplitudes of the two signals. In order to convert the amplitude of the signals into DC values, each signal is passed through a peak detection network (U9, U10, U12, U13) after the DC offset has been removed. Notably, the output of each peak detection network is still technically periodic in time because the capacitor (C1, C3) discharges for half of the period of the AC signal at the input. However, the total variation of the output of the peak detector networks is very small compared to the amplitude of the signal and can essentially be ignored. The output of the peak detector networks is a DC value equal in magnitude to the amplitude of the AC input signals. These values are multiplied by different constant values before entering the division circuit. The reason for this multiplication is discussed in greater detail below.

The first stage of the division circuit has an output equivalent to the natural logarithm of the DC input value. Although the output of the logarithmic amplifiers (U1, U4) is logarithmic in nature, it is a scaled version of the natural log of the input. The following amplifiers (U2, U5) up-scale the output of the natural logarithm amplifiers to account for the scaling of the actual log amplifiers. The output of the natural log amplifier and following amplifiers is calibrated to be equivalent to the natural log of the input value for the expected range of input values.

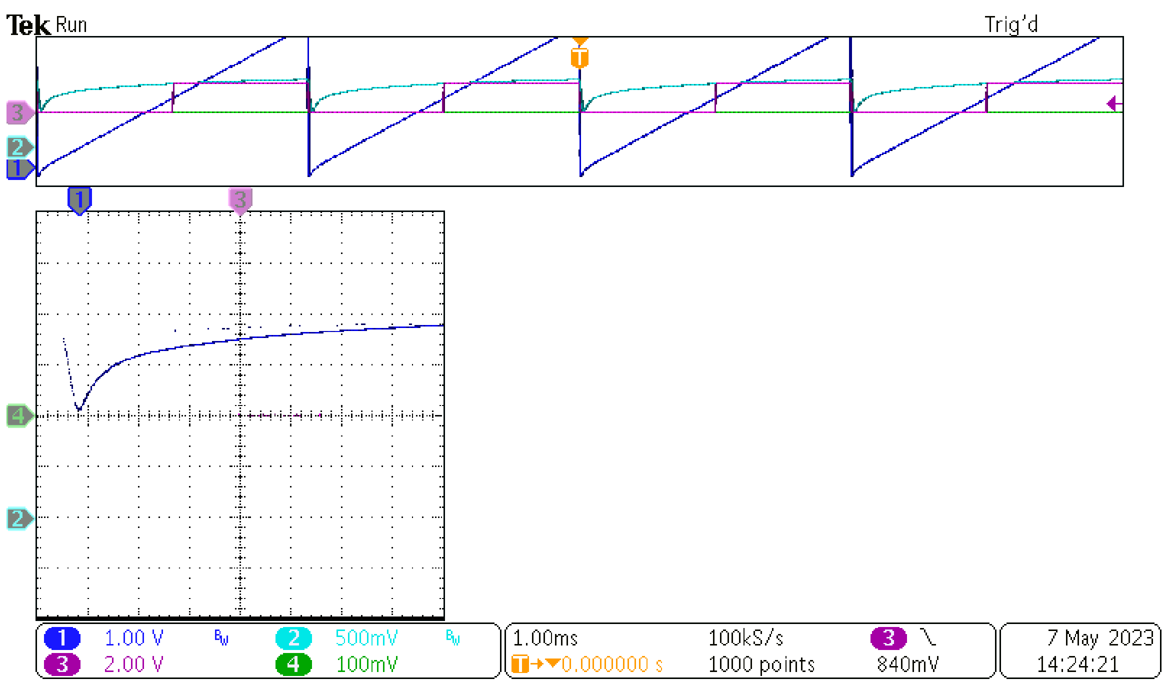

Figure 5. *Logarithm Output. The output of a log divider in XY mode. Channel 1 is a -.1 to 8 V sweep. Channel 2 is the output of the log circuit followed by inverting gain.. Channel 3 is the function generator sync pulse, used for consistent triggering.*

Device U3 is configured as a subtractor to find the difference between the outputs of the natural logarithm amplifier circuits. In mathematical terms, the mathematical constant e, raised to the value at the output of the subtractor, is equivalent to the amplitude of the tube amplifier divided by the amplitude of the control signal.

Designing the exponential amplifier posed a challenge similar to the challenges encountered in implementing the natural logarithm circuits. Although the output of the canonical exponentiation amplifier is exponential in nature, it is a scaled copy of e raised to the input value. Multiple approaches to this problem were considered. The most interesting approach was to use a Taylor series representation of e^x. The third term in the Taylor series of e^x involves x^2. In order to implement x^2 using only addition, subtraction, multiplication by constants, and natural logarithms, the following expression was designed to take the place of x^2 over our range of interest, $.7 < x < 2.5$:

$$5.8(x+.7)-5.8\ln(x+.7)-5.8$$

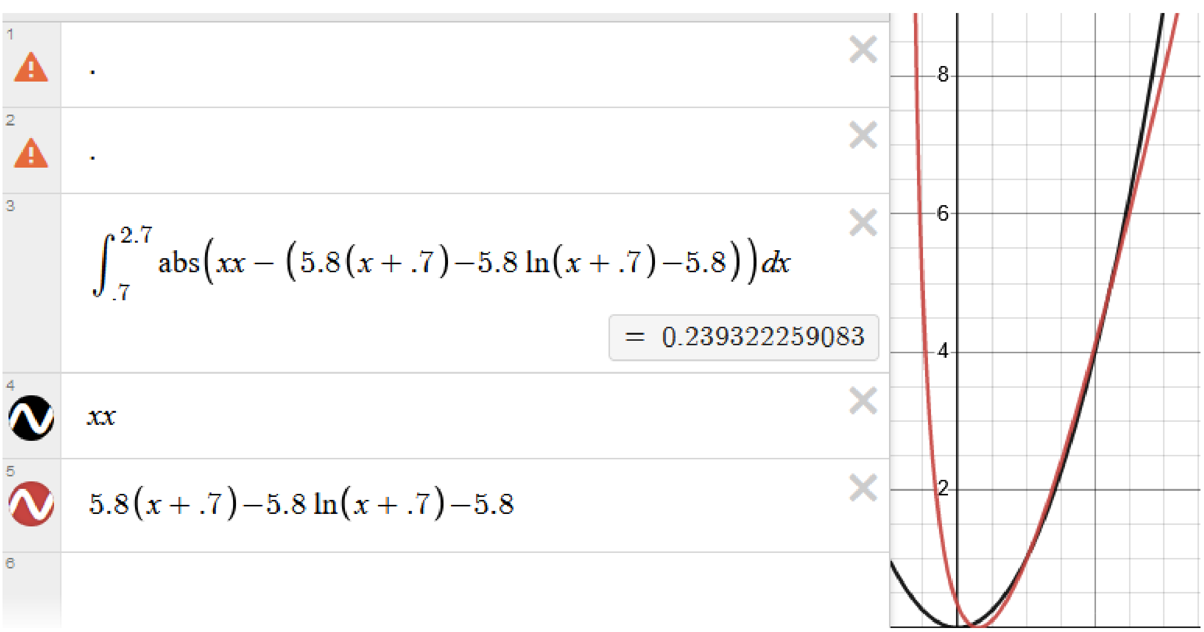

Figure 6. *Comparison of* x^2 *and approximation of* x^2

As shown in Figure 6, the approximation of `x^2` using a distorted natural logarithm is quite close on our range of interest, and enabled a similarly decent three-term taylor series approximation of `e^x` to be simplified to the following terms:

```
4.4X - 2.9ln(x+.7)
```

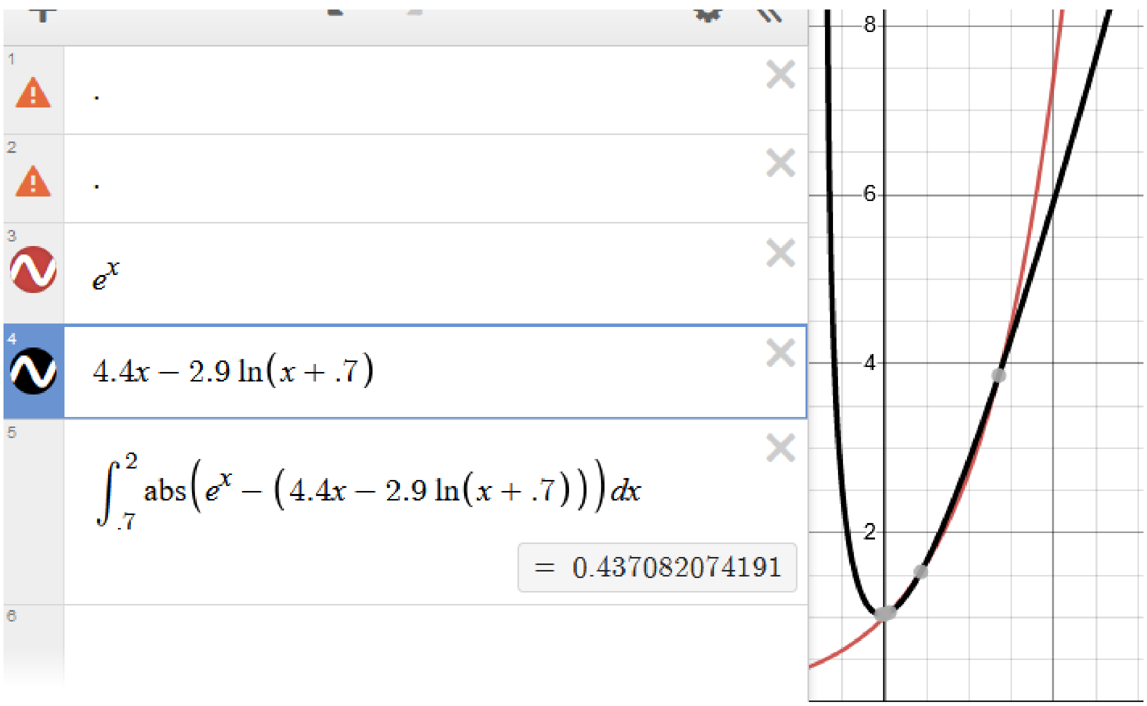

Figure 7. *Comparison of* `e^x` *and Taylor series approximation for* `e^x`

This approximation was implemented in LTspice with the following circuit:

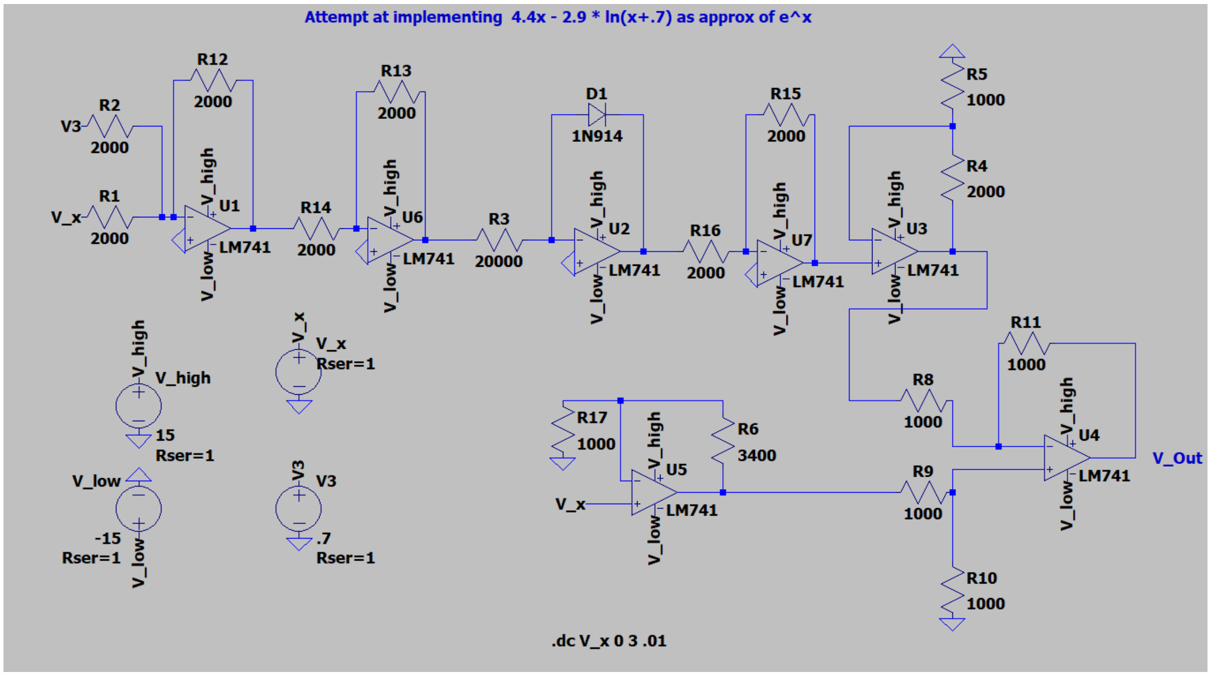

Figure 8. *LTspice simulation of Taylor series approximation of* e^x *(Device identifiers from this circuit are not referenced in the text)*

The Taylor series approximation of e^x was simulated in LTspice using a sweep of DC values between 0 and 3V at the input. The output data from the simulation was exported and graphed as individual points against e^x to graphically interpret the accuracy of the approximation. The result was underwhelming.

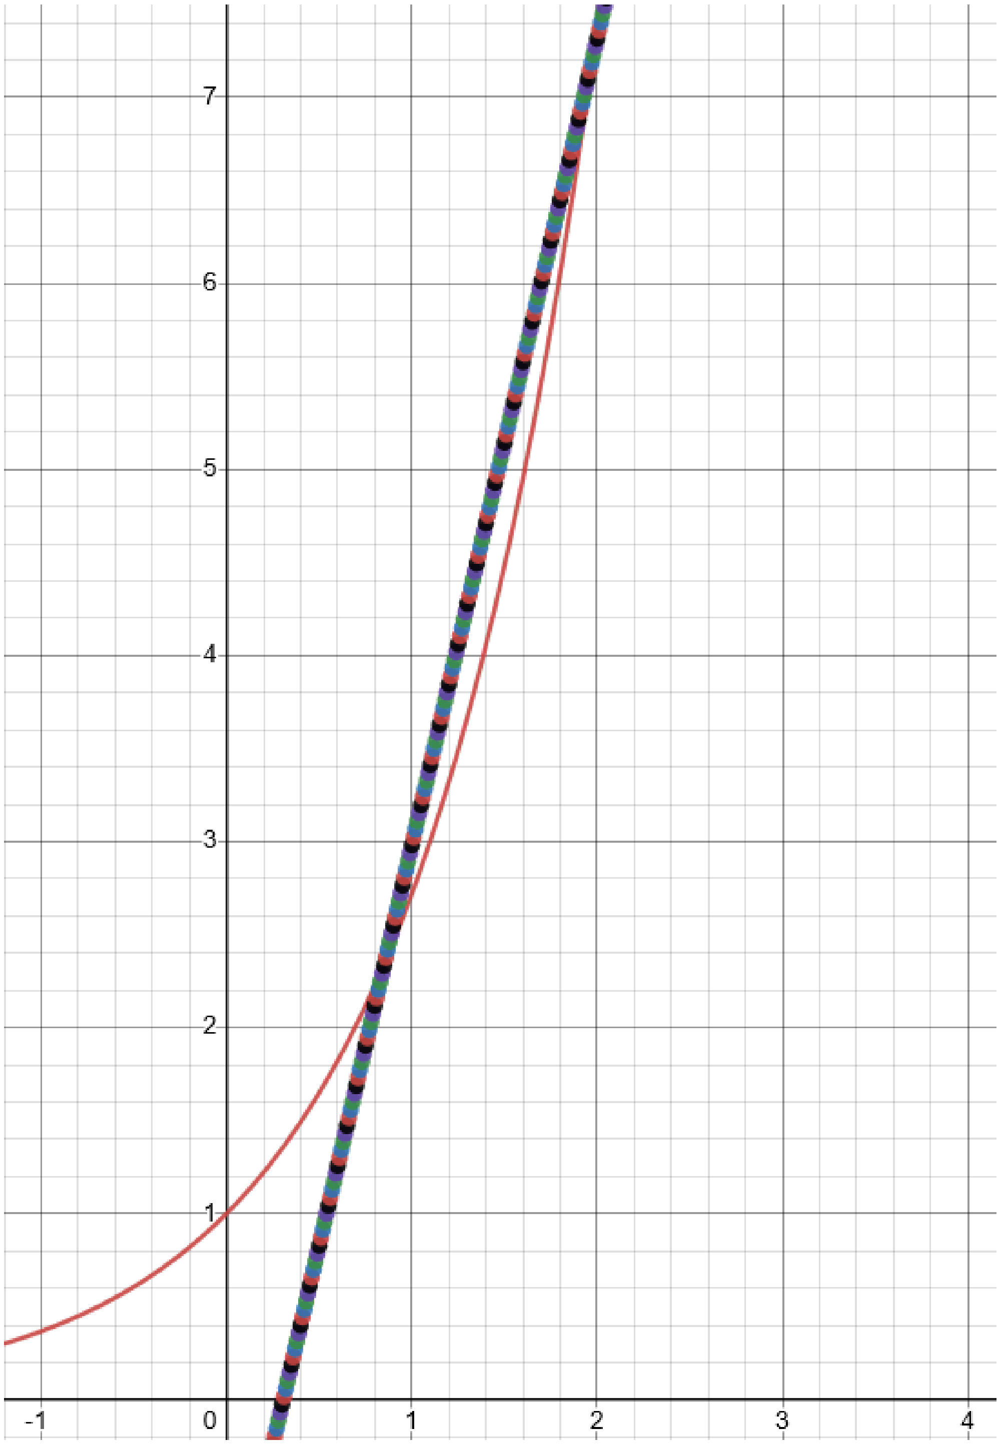

Figure 9. *Output simulated DC sweep of input values to Taylor series approximation of e^x, plotted over e^x for comparison.*

As seen in figure 9, the simulated output of the approximation is nearly linear. The only nonlinear term in the approximation is the use of a natural logarithm. Although post-logarithmic amplification was attempted to increase the influence of the logarithmic amplifier on the circuit, the results were unsatisfactory, and it was determined that implementing an exponential circuit which relied on a logarithmic circuit to function was unwise.

In place of the Taylor series, we opted to implement a linear approximation of `e^x`. The linear approximation of `e^x` is valid for a smaller range of values of x. In order for the approximation of `e^x` to be roughly accurate, the output of the subtractor must be roughly in the region where the approximation of `e^x` is most accurate. This can be accomplished by multiplying both of the inputs to the division circuit by different constants to constrain the difference in the natural logarithms of each to a smaller window. In this case, by multiplying the amplitude of the tube amplifier output by 4 and the control signal amplitude by 2, we were able to limit the expected difference in the natural logarithms of each input to a window of 1 to 1.7, where the approximation of `e^x` is most accurate. As seen below, the error of the approximation increases as the output of the subtractor circuit moves further from the intersection of `e^x` and the linear function approximating it.

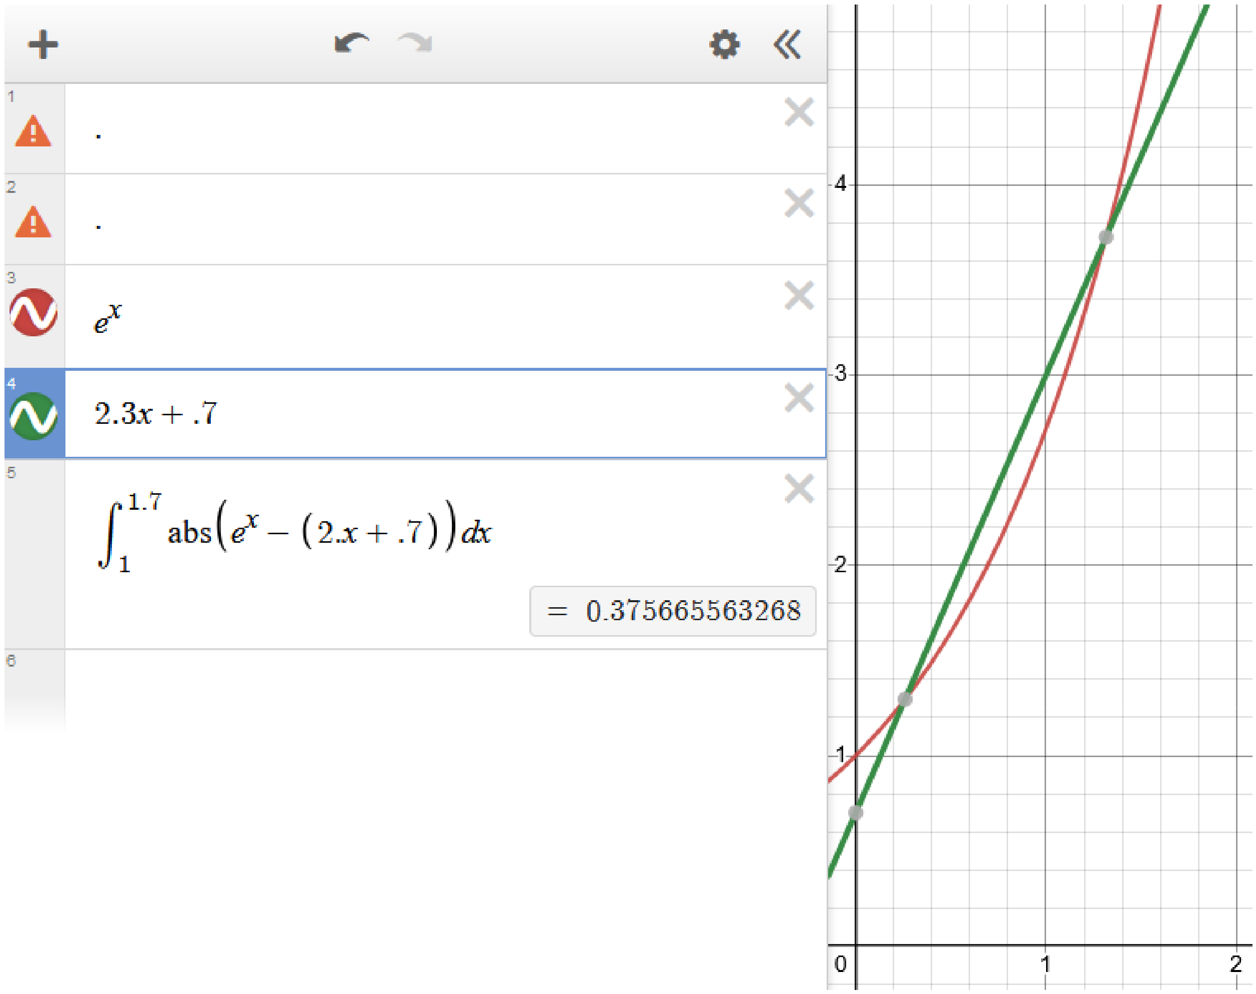

Figure 10. *Comparison of e^x and linear approximation of e^x.*

The separate gain applied to each signal before entering the division circuit must be accounted for at the output of the division circuit. Although the division circuit involves nonlinear operations, it may be considered a linear function when comparing its input and output). Accordingly, the output of the division circuit is multiplied by 2 at its output to account for the fact that the denominator of the input was subject to twice as much pre-division gain as the numerator. The voltage at the output of the final amplifier (U8) is a 1:1 representation of the mutual conductance of the tube under test.

Results

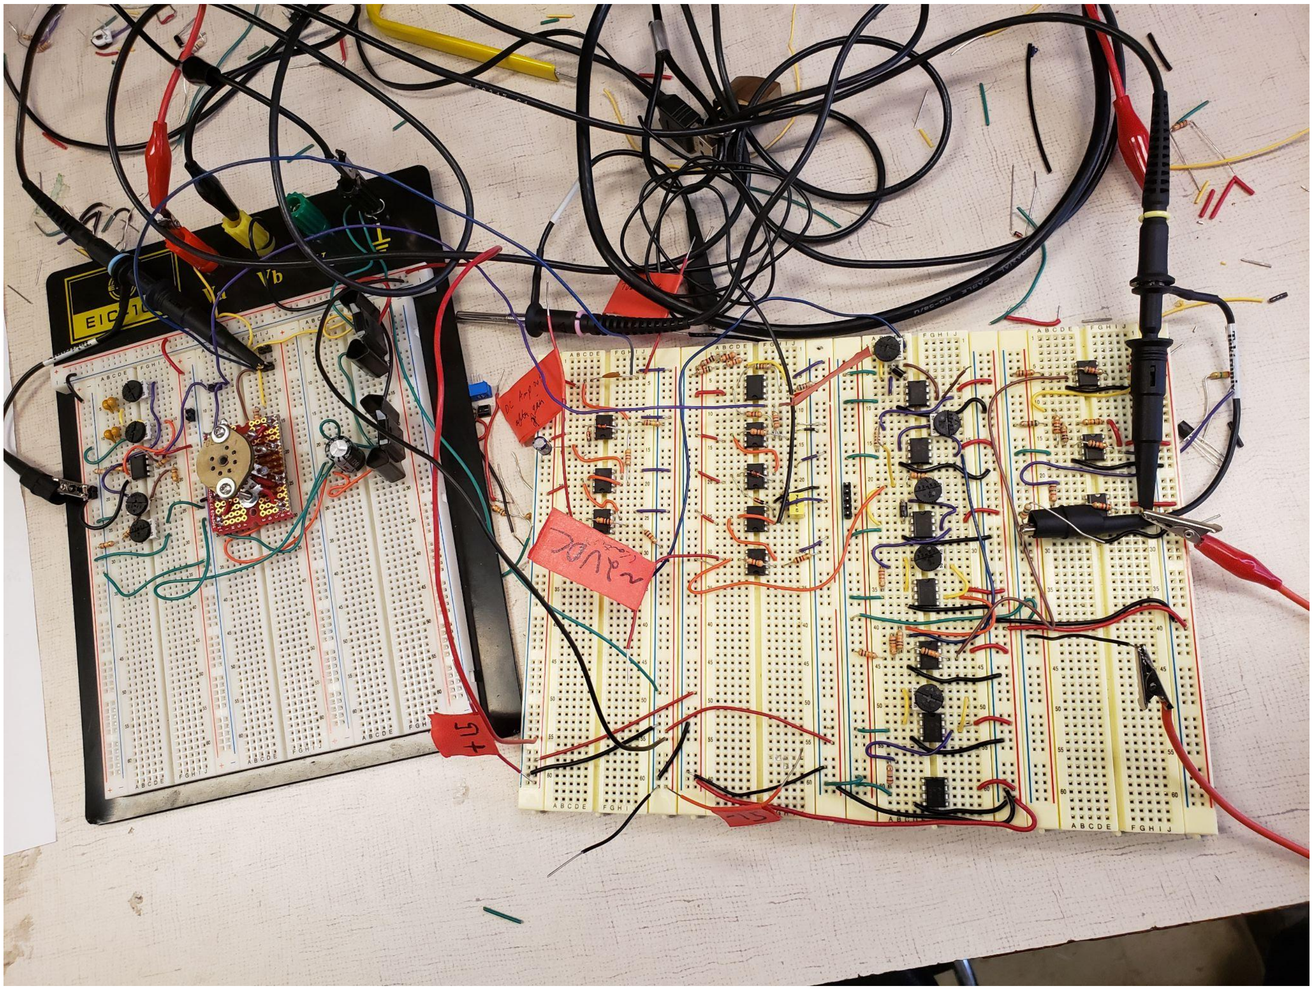

Figure 11. *The Tube Testing Circuit As Assembled*

Tube	Brand	Measured Gain (unloaded)	Project Measured Gain	Input Signal (Vrms)	Output Signal (Vrms)	Measured Gain (Loaded by project)	Measured/Project
1	GE	5.16	2.76	0.539	2.57	4.768089054	1.727568498
2	GE	5.38	2.79	0.539	2.683	4.977736549	1.784134964
3	GE	5.52	2.99	0.538	2.77	5.148698885	1.721972871
4	GE	5.22	2.67	0.538	2.59	4.814126394	1.803043593
5	GE	5.19	2.71	0.538	2.621	4.871747212	1.797692698
6	Mytron	4.8	2.55	0.538	2.52	4.68401487	1.836868576
7	Sylvania	4.7	2.46	0.538	2.39	4.442379182	1.805845196
1	GE	5.16	5.73	0.539	2.57	4.768089054	0.8321272345
2	GE	5.38	5.76	0.539	2.683	4.977736549	0.8641903731
3	GE	5.52	5.81	0.538	2.77	5.148698885	0.8861788098
4	GE	5.22	5.73	0.538	2.59	4.814126394	0.8401616744
5	GE	5.19	5.75	0.538	2.621	4.871747212	0.8472603847
6	Mytron	4.8	5.7	0.538	2.52	4.68401487	0.8217569947
7	Sylvania	4.7	5.68	0.538	2.39	4.442379182	0.7793647688

Table 1. *Tube Gain Measurements*

Table 1 shows measurements taken of the tube gain during our project. The unloaded gain was calculated directly from the unloaded amplifier. The project measured gain was the value given by the output of our project. The input and output signal amplitudes were measured through the AC-coupled scope. The measured gain is the ratio of output over input, measured while the tube was hooked into the project. The final column is the directly measured gain divided by the gain measured by the project-- note how though the ratio isn't 1, the ratio is fairly consistent between tubes, meaning that the error is consistent as well, and with further calibration accuracy could likely be greatly improved. The bottom half of the table is measurements taken after recalibrating the output of the log circuits to match predicted values, and the output was much closer to the measured gain.

Sources of Error

In order to gain a better understanding of the source of error in our device, we compared measured voltages from points within the circuit to the expected output of the previous block, and calculated the percent error occurring in each sub-circuit:

Circuit Block	Expected value [Vp, V]	Observed value [Vp, V]	% error
Tube Out -> HPF -> Peak Detection -> 2X gain	6.78	6.2	-8.55457
control signal -> offset removal -> peak detection -> 4x gain	3.00	2.12	-29.3333
Natural log, with scaling, of tube output amplitude	1.82	1.86	2.197802
Natural log, with scaling, of control signal amplitude	0.75	0.75	0
subtractor	1.07	1.02	-4.6729
exponential approximation, compared to designed approximation	3.04	2.85	-6.25
exponential approximation, compared to exponential function	2.77	2.85	2.888087
overall gain measurement*	4.52	5.69	25.88496
Expected value calculated from measured input/output of tube while operating in the test circuit			

Table 2. *Comparison of Error in Each Subcircuit for V_control = .75 Vp and V_Tube_Output = 3.99 Vp. Data Presented for the Sylviana Tube.*

The error in each subcircuit in Table 2 is not cumulative. It is specific to that subcircuit. In reality, the error propagates through the circuit, and has the potential to self correct if the relative direction of error between subcircuits is opposite at a particular operating point or if the error at a particular operating point influences the output of the subtractor in a manner that shifts the final output. Although the notion of self-correcting error may seem haphazard to note, it is important to consider that the sign of the error introduced by the linear approximation of `e^x` changes depending on whether the output of the subtractor is greater or less than 1.3 Volts (the intersection of `e^x` and the linear approximation of `e^x`).

In addition to the potential for error that either accumulates or self-corrects throughout the circuit, the error can also be amplified. The use of compensating gain prior to the logarithmic amps and prior to the mutual conductance readout amplifies the magnitude (but not the percentage) of error which is present prior to this amplification. This can be best seen at the output, where approximately a tenth of a volt deviation from the expected value becomes two tenths of a volt of error after the final 2x compensating gain.

Table 2 indicates that the largest source of error is in offset removal and peak detection. Both the tube output signal and the control signal were processed through the same peak detector topology, but were subject to different implementations of DC offset removal. Unsurprisingly, the HPF topology was more accurate than offset removal by addition of a fixed DC value. Interestingly, at this operating point, the linear approximation output was closer to e^x than it was to the approximation that the circuit was designed around.

Conclusion

Leif:

I spent the early phases of the project simulating in LTspice. This provided two important lessons for me. First, spending a significant of time working in LTspice forced me to develop a greater degree of familiarity and comfort with the software. I wish I had done so years ago. Second, I learned a hard lesson about over-reliance on simulation. Early in the project, I was under the impression that I had designed a working division circuit. I later realized that my simulated circuit incorporated complimentary errors which had somewhat miraculously canceled each other out at the particular test values I was using. Additionally, my designs behaved radically differently when implemented on a proto-board with real components (and their associated real tolerances and parasitics). This was an exemplary demonstration of an old saying - “in theory, there is no difference between theory and practice, but in practice…”

Ian and I chose to do analog division precisely because it was not the recommended approach. Although I don’t regret this decision, I do recommend that a future team requiring functional division either conduct it in the digital realm, or consider exploring an approach based on voltage controlled resistors.

This project’s greatest value to me is specific and tangible. In the last few months, I still haven’t crufted a tube tester, and my luthier friend back home still needs one. Although our design is not exactly what I would recommend for the upcoming project, it offers a strong starting point and it allowed me to further develop technical skills for my next attempt at implementing a functional tube tester.

Ian:

Primarily, I handled the design of the tube input circuit, and the tube amplifier circuit itself. I picked out a low-voltage tube, figured out an operating point, and made an adapter to fit the tube in the breadboard. After running the tube with the function generator, I measured the gain of the tubes directly by measuring the input and output with the oscilloscope, and dividing the two. I then passed these values on to Leif so he could verify that his designs would function in that area. After this I designed and built the input oscillator and offset circuit, then started building the division circuit based off of Leif’s schematics. After running into trouble with the diode-wired transistor, Leif and I researched some more and tried different diodes in the circuit, tweaking values until we got something to work. Setting the scope to XY mode and sweeping the input and measuring the output made it really satisfying to see the log curve finally show up. Though the analog division circuit was challenging and not actually all that practical, I sincerely enjoyed working on it and learning more about the specifics of the diode- it's all too easy to gloss over the detail with simplifications, such as assuming that the voltage is entirely constant across the diode, without considering whether there are uses outside of the cases where we can get away with so much simplification. Knowing when to simplify and when not to is an important skill to have as an engineer, and I'm glad that this project was able to help me practice that ability.

Citations

Pedagogically popular op-amp division circuit: (Figure 3)

https://www.researchgate.net/profile/Miguel-Martinez-Ledesma/publication/338983194/figure/fig7/AS:853928484798474@1580603824412/Analog-division-scheme-using-operational-amplifiers.ppm