Wesleyan ⧫ University

PT-Symmetric Electronics

by

Joseph Caulfield Schindler

Faculty Advisor: Fred Ellis

A thesis submitted to the faculty of Wesleyan University in partial fulfillment of the requirements for the degree of Master of Arts in Physics

Middletown, Connecticut **May 2013**

Acknowledgements

First of all, it would be impossible to overstate the thanks I owe to Fred, whose mentorship for the past four years has been beyond formative. Thanks for teaching me almost everything I know!

Many thanks also go to Tsampikos, for making this entire project and thesis possible, for always providing direction and relevance to our research, and for many exciting discussions.

I owe Gideon and Nick immensely, without them I literally couldn't have finished the octoboard project. Their hard work on the Axon interfacing was amazing, and deeply appreciated.

Thanks to Hamid and Zin for making collaborating so much fun, and I hope our collaborations extend to the future. Sam, it's been a great time down in the lab, glad to share it with you. Sorry in advance if you have to make all those octoboard connectors. And to the rest of my physics peers, the physics grad students, my teachers, and all the physics faculty: thanks for making my time in the department fun and enlightening.

Thanks to Andy, Anna, Max, AK, FMR, Ty, and everyone else who spent time with me in the underworld, making the writing process more fun and considerably slower. Finally, thanks to my family for being the best, and to NF for the same reason.

Contents

List of Figures

Chapter 0

Introduction

In early 2010, a project in "parity-time symmetric electronics" was initiated in collaboration between Professor Tsampikos Kottos's theoretical physics group and the experimental physics group of Professor Fred Ellis. The basic idea of the project is to use coupled electronic LRC oscillators as a simple way to experimentally study systems with parity-time symmetries. Since 2010, the project has produced numerous experimental results and several publications. The goal of this thesis is to elucidate in detail the experimental aspects of the project, and fully document the experimental apparattii and techniques used up to this point. Besides conducting these experiments over the past several years, the major Master's research portion of this work consisted of apparatus development. Specifically, the thesis project culminated with the creation of a computer-controllable multi-unit experimental apparatus designed for future expansion of our laboratory's capability into larger-scale research in this direction. This major project is the subject of Chapters 4 and 5, as well as the Appendix A.

This thesis attempts to fullfil three major porpoises: (1) To provide an introduction to the field of PT-symmetric electronics; (2) To fully document the experiments perfomed thus far; and (3) To provide a path forward for expansion of the research to systems of greater complexity.

Chapter 1 provides an introduction to the concepts of electronics and symmetry necessary for undertaking this research and understanding the material. Chapter 2 summarizes the results of various experiments performed over several years, and discusses their experimental implementation. Chapter 3 then completes the experimental documentation by describing the apparattii and providing a broader discussion of practical experimental considerations. Chapters 4 and 5 are devoted to seeking computerized automation of the experiments via the "Octoboard", a circuit board apparatus designed for continued research in this direction. The Appendix contains full documentation of the Octoboard apparatus.

Chapter 1

The PT-Symmetric LRC Dimer

The primary objects of study in this research are coupled LRC oscillators. Using these simple and well understood devices we probe the behavior of PT-symmetric and non-conservative systems. This chapter lays the theoretical foundation for the experimental study of coupled LRC oscillator pairs.

1.1 The LRC Oscillator

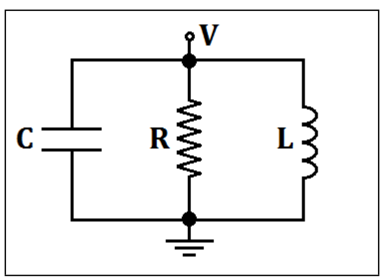

Figure 1.1: An LRC Oscillator

The *LRC* oscillator is our basic building block. Notice that although experience inclines us to regard resistance as a positive value, the results of this section extend to both positive and negative resistances. We'll soon show that both are relevant. Both L and C are considered non-negative.

1.1.1 Equation of Motion and Transient Solutions

The equation of motion for a single LRC oscillator is given by differentiating Kirchoff's current rule. Positive currents are defined downward on the page.

$$\tfrac{d}{dt}(I_C + I_R + I_L) = 0 \tag{1.1}$$

$$C\ddot{V} + \dot{V}/R + V/L = 0 \tag{1.2}$$

$$\ddot{V}(t) + \tfrac{1}{RC}\dot{V}(t) + \tfrac{1}{LC}V(t) = 0 \tag{1.3}$$

The voltage oscillates with (undamped) natural angular frequency ${\omega_0}^2 = 1/LC$. Changing variables from time to the phase $\phi = \omega_0 t$ yields, by the chain rule,

$${\omega_0}^2\ddot{V}(\phi) + \omega_0\tfrac{1}{RC}\dot{V}(\phi) + {\omega_0}^2 V(\phi) = 0 \tag{1.4}$$

$$\ddot{V} + \tfrac{\sqrt{L/C}}{R}\dot{V} + V = 0 \tag{1.5}$$

$$\ddot{V} + \gamma\dot{V} + V = 0 \tag{1.6}$$

where a dot now signifies the derivative $\frac{d}{d\phi}$. Proper scaling has allowed all but one parameter to be scaled away. Taking an ansatz $V = V_0 e^{i\omega\phi}$, the solutions are

$$V(\phi) = V_0 e^{-\frac{\gamma}{2}\phi} e^{\pm i\sqrt{1-\frac{\gamma^2}{4}}\phi} \tag{1.7}$$

which have both an oscillating part and an exponential part in the "underdamped" regime $0 < |\gamma| < 1$.

1.1.2 Impedance and Conductance

At the natural frequency ω_0 the inductor impedance and the capacitor impedance are equal in magnitude. Their common value is called the *characteristic impedance* z_0 of the circuit.

$$z_0 \equiv \sqrt{\frac{L}{C}} = |i\omega_0 L| = \left|\frac{1}{i\omega_0 C}\right| \tag{1.8}$$

On resonance the opposite phases of the capacitive and inductive impedance cause them to cancel, and the net impedance is $z_{net}(\omega_0) = R$.

It is useful to think in terms of conductance in these systems, rather than impedance. Conductance g is the reciprocal of impedance, $g = 1/z$. There are several advantages to using conductances. Parallel conductances add, unlike parallel impedances. Energy dissipated in the resistor is, for a given voltage, proportional to its conductance. Later we'll see that current through a vactrol LED is about proportional to the photoresistor conductance (A vactrol is a controllable resistance device consisting of an LED and photoresistor, see Chapter 4 for details). Moreover, in parallel circuits it is the elements with greatest conductance that dominate the behavior.

The unitless parameter $\gamma = z_0/R$ appearing in (1.6) and (1.7) is the conductance of the resistor scaled by the characteristic LC conductance. As $|\gamma|$ increases toward one, the resistor becomes very conductive and dominates the circuit's behavior, causing the rate of decay/growth to become fast relative to the rate of oscillation. The oscillation amplitude decays with time if γ is positive, and grows if γ is negative. The sign of γ is the sign of R. Throughout this work, γ will often be the manipulated variable of interest.

1.1.3 Quality Factor

For an oscillation $f(x) = Ae^{\lambda x} = Ae^{Re(\lambda)x}e^{iIm(\lambda)x}$, the quality factor Q compares the relative amounts of oscillatory versus exponential behavior, and is defined by[1]

$$Q \equiv \frac{|\mathrm{Im}(\lambda)|}{-2\mathrm{Re}(\lambda)} \tag{1.9}$$

Equivalently, under the $e^{i\omega t}$ convention this becomes $Q = \frac{|\mathrm{Re}(\omega)|}{2\mathrm{Im}(\omega)}$. Convention differences can reverse signs and should be handled carefully.

Positive (negative) Q indicates a decaying (growing) oscillation. Large $|Q|$ indicates

[1]A slightly different choice than the common definition $Q = \frac{\omega_0}{\gamma}$, which assumes more about the source of the oscillations. The two are essentially equivalent for $Q > 3.5$.

weak exponential behavior. In the passive case $\gamma = 0$, $Q = \infty$. Q is approximately the number of oscillatory cycles over which the amplitude changes by a factor of twenty.

Q also relates to the steady state solutions of a sinusoidally driven damped harmonic oscillator, and can be extended approximately to the steady state oscillations of general linear damped systems. When Q is positive it can be easily measured in this way, by the shape of the resonance curve when the system is driven at various frequencies. Modes with a high Q have sharply peaked resonances. If f_0 is the peak frequency which maximizes amplitude, and Δf is the "full width half maximum" frequency difference of the amplitude curve, then

$$Q \approx \sqrt{3}\frac{f_0}{\Delta f} \tag{1.10}$$

The Q of the LRC oscillator is given from (1.7) and (1.9) by

$$Q \approx \frac{1}{\gamma} = \frac{R}{z_0} \tag{1.11}$$

where the approximation is within 1% for $Q \gtrsim 3.5$. If $\gamma = 0$, the solutions are pure oscillations which neither grow nor decay with time. In this case, the behavior is identical whether observed forward or backward in time. Otherwise, if $\gamma \neq 0$, switching the direction of time changes the solution by turning growth into decay and vice versa.

1.2 Negative Resistance

The simplest conception of a resistor is associated with Ohm's law for a DC voltage. Maintain a constant voltage difference ΔV across a resistor, and the current that flows across it is proportional to ΔV. The current in a resistor flows from high voltage to low voltage, so we define "positive current" to be in the direction of decreasing voltage. The I-V curve is therefore a straight line through the origin with positive slope $\frac{1}{R}$.

For AC voltages, Ohm's law for a resistor is preserved and attains additional meaning. The AC impedance of an ideal resistor is real-valued and independent of frequency, meaning that I is in phase with V, and the DC Ohm's law applies at each instant.

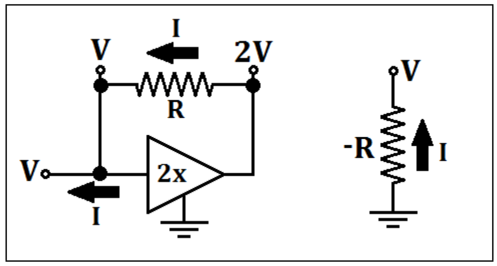

Figure 1.2: A ground-referenced negative resistance node.

So what is a "negative" resistor? It's just a regular resistor with a negative value. That is, it's a device which obey's Ohm's law with a current that flows "backwards", from low voltage to high voltage. The I-V curve is a negatively sloped line. In the case of a one-terminal ground-referenced negative resistance (the only one we'll use), this means a current $I = V/R$ flows *into* the terminal of voltage V. Compare this to the positive resistance case, where the same current would flow *out* of the terminal.

Of course, any passive linear device with this property would violate energy conservation. However, active versions can be made quite easily. Figure 1.2 demonstrates how a voltage doubling amplifier with positive feedback resistance can be used to make exactly such a device. This device will be sufficient for our purposes, but it is also possible to create a true floating two-terminal negative resistance element as shown in Figure 1.3. One can also attain an effective negative resistance with passive nonlinear elements, for example by DC biasing a tunnel diode to a region of approximately constant negative slope in the I-V curve. In this case I is still positive, so energy is conserved, but is locally decreasing as V increases, causing an effective negative resistance.

Negative resistance nodes will be drawn interchangeably as negative resistors to ground or in their op-amp form. The two forms are equivalent.

The circuit on the left of Figure 1.2, but with the resistor R removed so the feedback terminals are open, can also be referred to as a *negative impedance converter* (NIC), since any impedance placed between the V and $2V$ terminals will effectively be made negative. In our case, we are using a negative impedance converter with a resistor to

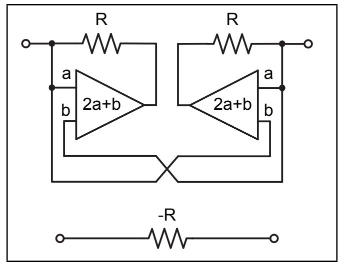

Figure 1.3: A floating negative resistor.

achieve a negative resistor.

A final note about the use of negative resistances: a negative resistance must always be used in parallel with a capacitance [4]. Networks including a two-terminal pure negative resistance structure with no parallel elements are theoretically undefined. The problem is hidden, since it only appears when a parallel capacitance is first included, then taken to the limit $C \to 0$. In that limit, any such network has an exponentially growing mode of diverging rate, which dominates the behavior, corresponding to the $e^{t/RC}$ mode of a simple RC circuit with $-R$. This mode is inconsequential for positive resistance circuits since it corresponds to an infinitely fast decay. In practice, the time constant of the hidden mode would be controlled by stray parallel capacitances. For example, a series LRC circuit with negative resistance would have an exponentially growing mode with growth rate related to the stray capacitance across $-R$. The behavior would be controlled by a circuit element we did not intentionally include. A known parallel capacitance would then need to be included to allow useful theoretical analysis.

1.3 PT Symmetry

Parity (P) reversal and time (T) reversal operations are related to spatial and temporal symmetries in a natural system. Parity-Time (PT) symmetric systems are defined by invariance under combined parity and time reversal. In English this means the following: if you look at a PT-symmetric system in the mirror while time runs backwards, you can't tell it apart from the original forward-time mirror-free case (good luck trying that out). We define the parity and time reversals in general, in terms of electronics, and in terms of a Hamiltonian formalism.

One subtlety should be illuminated from the start. A *physically* PT-symmetric system (i.e. one with PT-symmetric equations of motion) has *solutions* which are not necessarily all PT-symmetric. The symmetry of the system and of the solutions must be regarded seperately.

1.3.1 Parity Reversal

Parity reversal is defined by spatial inversion. In continuous coordinates this means P denotes a reflection $x \longrightarrow -x$ across one spatial axis[2]. Quantities depending on odd powers of x or its derivatives (e.g. p_x, a_x, etc.) are consequently inverted as well.

In the long wavelength limit of electronics, spatial considerations disappear, and spatial symmetry is reduced to a matter of circuit topology. The parity reversal is then defined by the interchange of labels corresponding to left and right sides of a circuit configuration. Parity-related symmetry is therefore easier to achieve than in a true spatial system.

In the context of the Hamiltonian dynamics, P takes $x \longrightarrow -x$ and $p \longrightarrow -p$, and H is P-symmetric if $[H, P] = 0$.

1.3.2 Time Reversal

Time reversal, as one expects, consists of inverting the flow of time. This backward propogation is acheived via the operation $t \longrightarrow -t$. Time derivatives $\frac{d^n}{dt^n}$ of odd order become inverted by this process.

In electronics time reversal effectively reverses the sign of resistive impedances while leaving reactive impedances unaltered. Applying T to Ohm's law yields

$$T[V_\omega(t) = I_\omega(t) z(\omega)]$$

[2]Actually, parity reversal can be defined by inversion of just one or all three cartesian axes in three dimensions. Such intricacies are ignored here.

$$V_\omega(-t) = -I_\omega(-t)z(\omega) \tag{1.12}$$

where the sign of I is reversed because current is the time derivative of charge. For resistors, with $V = IR$, the time-reversed statement is then

$$Ve^{i\omega(-t)} = -Ie^{i\omega(-t)}R \tag{1.13}$$

or $V = -IR$. For the reactive components, taking an additional derivative causes the negative sign to drop out. This fact is intuitively related to the dissipative nature of resistors. Energy is dissipated in a resistor in normal time, so in reverse time energy is drawn out from it.

Under Hamiltonian formalism, T takes $p \longrightarrow -p$ and $i \longrightarrow -i$. H is T-symmetric if $[H, T] = 0$. Moreover H is PT-symmetric if $[H, PT] = 0$ but $[H, P] \neq 0$ and $[H, T] \neq 0$.

1.3.3 PT Symmetry

Systems which are invariant with respect to combined PT reversal, but not under P or T individually, are PT-symmetric. Such systems are typically realized by adding balanced energy gain and energy loss mechanisms in parity-linked pairs to an otherwise symmetric system. They are non-conservative, yet the gain and loss mechanisms are equal and opposite, allowing the possibility of net energy conservation under suitable averaging. One of the strange results about these systems is that as the strength of this balanced non-conservative factor is increased, there can be a spontaneous transition from net energy conservation to net non-conservation. This corresponds to a spontaneous PT symmetry breaking of the solutions to the symmetric system. This spontaneous symmetry breaking is a hallmark of PT-symmetric systems, and has been widely studied.

The very existence of the symmetric region, the region of net conservation, is interesting in and of itself. It demonstrates that a PT-symmetric Hamiltonian can have real energy eigenvalues, despite being non-Hermitian. This contradicts the usual dogma of quantum mechanics, whose canon states the Hamiltonian of a physical system must be

Hermitian so that measured energies will be real. PT-symmetric Hamiltonians were the first non-Hermitian Hamiltonians known to have such a property. It was this discovery by Carl Bender [5] that prompted the earliest investigations of PT symmetries. The pseudo-Hermitian solutions and spontaneous symmetry breaking will be encountered in our treatment of the LRC dimer.

PT-symmetric systems have been studied theoretically in numerous fields including quantum mechanics, optics, electronics, and thermal physics, and experimentally in optics and electronics (see introduction of [4]). The fact that PT-symmetric systems show similar behavior in all these fields, despite the different physics, is an enlightening confirmation of the power of symmetry.

1.3.4 Symmetry in Electronics

As mentioned above, utilizing electronics to study spatial symmetry in this case provides a great advantage. The long wavelength electronic limit ensures that all spatial and geometric considerations are eliminated in favor of circuit connectivity considerations. For this reason PT-symmetric circuits of arbitrary complexity can be assembled according to the following simple rules [4]

(1) Reactive elements must either occur in parity-linked pairs, or connect parity-inverted nodes.

(2) Resistive elements must occur in parity-linked pairs of opposite sign.

(3) Each negative resistance element must occur in parallel with a capacitance.

1.4 The PT Dimer

The pair of coupled LRC oscillators depicted in Fig. 1.4, one with linear amplification due to the negative resistance $-R$ and the other with an equal amount of linear attenuation due to the resistance R, are the unit we call the PT dimer. Coupling is provided by

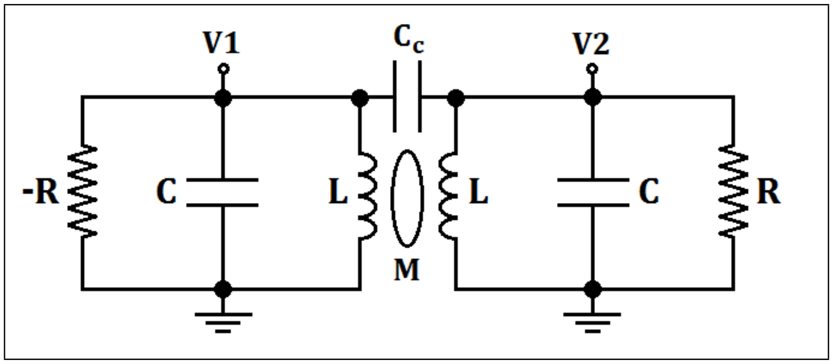

Figure 1.4: The ideal PT dimer.

a mutual inductance M at the inductors and a coupling capacitance C_c connecting the nodes. This circuit represents the simplest electronically realizable demonstration of PT symmetry with nontrivial mode behavior, and is the basic experimental apparatus used throughout this work. This section analyzes the normal mode dynamics of the dimer. Both capacitive and mutual inductive coupling are included for generality, although in practice we utilize only one or the other at a time.

The normal modes are found from Kirchoff's laws. Positive currents again run down the page. The equations of motion are given by the inductor voltage rule and Kirchoff's current law at the gain and loss sides. Note that without coupling each oscillator has natural frequency $\omega_0 = \frac{1}{\sqrt{LC}}$ and characteristic impedance $z_0 = \sqrt{\frac{L}{C}}$. These will be useful scaling factors allowing a switch to unitless parameters. The equations of motion are

$$V_1 = L\dot{I}_1^L + M\dot{I}_2^L \qquad -\frac{V_1}{R} + C\dot{V}_1 + I_1^L + C_c(\dot{V}_1 - \dot{V}_2) = 0 \tag{1.14}$$

$$V_2 = M\dot{I}_1^L + L\dot{I}_2^L \qquad \frac{V_2}{R} + C\dot{V}_2 + I_2^L - C_c(\dot{V}_1 - \dot{V}_2) = 0 \tag{1.15}$$

Assume normal mode solutions of frequency ω', where all quantities oscillate as $e^{i\omega\phi}$ if $\phi = \omega_0 t$ and $\omega = \frac{\omega'}{\omega_0}$. Eliminating I_1^L and I_2^L, and scaling the remaining parameters, the set of equations is reduced to the linear homogeneous system

$$\begin{pmatrix} \frac{1}{\omega}\frac{1}{1-\mu^2} - \omega(1+c) - i\gamma & \omega c - \frac{1}{\omega}\frac{\mu}{1-\mu^2} \\ \omega c - \frac{1}{\omega}\frac{\mu}{1-\mu^2} & \frac{1}{\omega}\frac{1}{1-\mu^2} - \omega(1+c) + i\gamma \end{pmatrix} \begin{pmatrix} V_1 \\ V_2 \end{pmatrix} = 0 \tag{1.16}$$

where $\gamma = \frac{z_0}{R}$, $c = \frac{C_c}{C}$, and $\mu = \frac{M}{L}$. Frequency ω is scaled by the natural frequency of the uncoupled oscillators, and c and μ are unitless measures of the coupling strength. Moreover, γ is the symmetric gain/loss strength parameter associated with the resistance, and is identical to the scaled conductance parameter gamma we saw for the single oscillator in Section 1.1. Notice that although the dimer has five free parameters, the scaled equations of motion depend on only three parameters. The unitless problem is independent of the values of L and C, which were scaled out and only need to be specified if a return to unit quantities is desired. Once γ, c, and μ are fixed, specifying L and C determines frequency in Hz and R in Ohms.

The PT symmetry of the equations of motion is easily confirmed. In the framework of (1.16) the PT operation swaps the indices 1 and 2, and reverses the sign of i. These operations in conjunction leave the equations unaltered. The symmetry of the equations of motion correspond to the *physical* symmetry of the system. Figure 1.5 illustrates the presence of physical symmetry.

The above matrix equation (1.16) has nontrivial solutions only if the coefficient matrix has zero determinant, allowing us to identify the four normal mode frequencies.

$$\omega = \pm \frac{\sqrt{\gamma_c^2 - \gamma^2} \pm \sqrt{\gamma_{PT}^2 - \gamma^2}}{2\sqrt{1+2c}} \tag{1.17}$$

where the PT symmetry breaking point γ_{PT} of the solutions is

$$\gamma_{PT} = \left| \sqrt{\frac{1+2c}{1+\mu}} - \frac{1}{\sqrt{1-\mu}} \right| \tag{1.18}$$

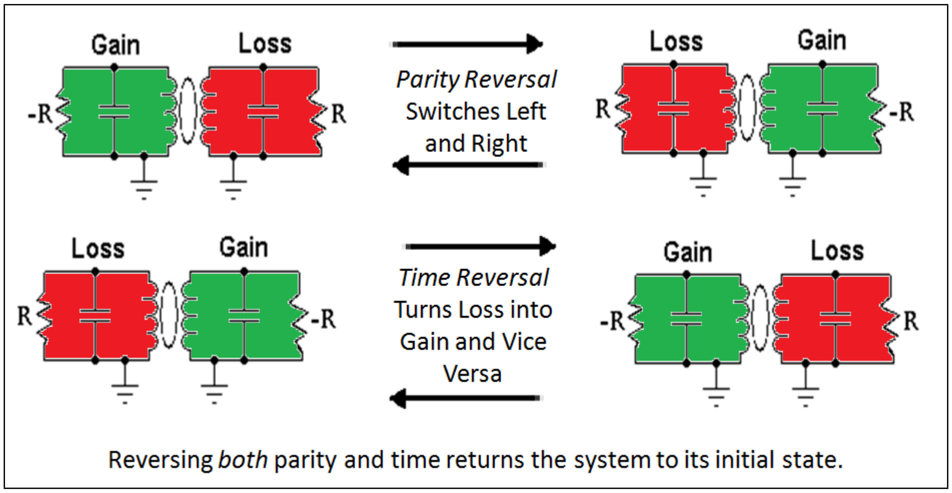

Figure 1.5: The physical symmetry of the system corresponds to its symmetric equations of motion.

and the upper critical point γ_c is

$$\gamma_c = \sqrt{\frac{1+2c}{1+\mu}} + \frac{1}{\sqrt{1-\mu}} \tag{1.19}$$

The solutions have several regimes of behavior, as depicted in Figure 1.6. There are actually four normal modes, but in the parametric regions of interest, if we seek only purely real solutions, there is redundancy, leaving just two distinct modes to be considered. In these regions the two solutions may be called $\omega_{\pm}$, where the outer $\pm$ of (1.17) has been taken positive and thrown away, and the subscript is associated with the inner $\pm$.

While $0 < \gamma < \gamma_{PT}$, the four eigenvalues are purely real, and come in two sets of positive-negative pairs. The motion in this regime is simply oscillatory, and occurs at two distinct normal mode frequencies (the positive and negative frequencies of equal magnitude are essentially identical). This interval is known as the "exact phase" because these modes are "exactly" PT-symmetric.

The symmetry of the solutions can be assessed by considering observing V_1 and V_2 normally, and then under parity-time reversal. If there's no way to distinguish the two data sets, the behavior is symmetric. Such is the case in the exact phase. Note that while any single oscillation is T-symmetric, and a pair of in-phase oscillations

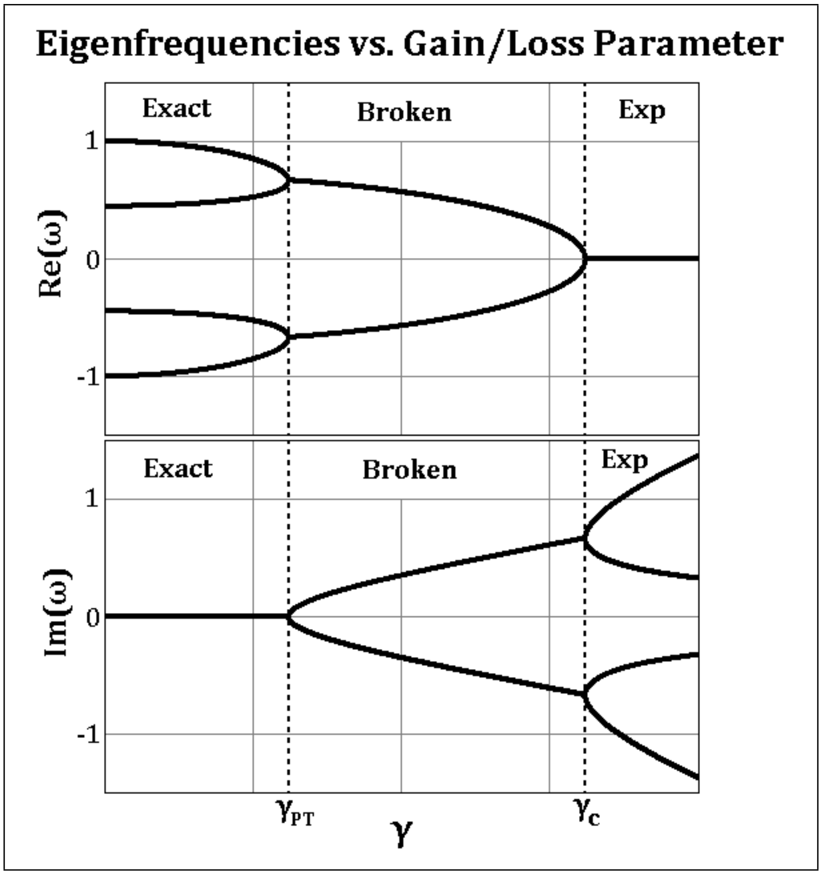

Figure 1.6: Real and imaginary parts of the normal mode frequencies versus γ for $\mu = 0$, $c = 2$. The qualitative features do not depend on the choice of μ and c. Frequencies are defined by the $e^{i\omega t}$ convention. The exact phase is characterized by four purely real oscillating frequencies. The broken phase has exponentially growing and decaying oscillations. The exponential phase has purely exponential non-oscillating solutions.

would be P-symmetric, the individual T and P symmetries of our dimer in the exact phase are broken by the phase relationship between V_1 and V_2. To check the symmetry mathematically, the P and T operations must be applied to the eigenvectors, which are shown in (1.20) below for the exact phase. In the exact phase the eigenvectors are invariant under PT reversal.

The region beyond the exact phase, where $\gamma_{PT} < \gamma < \gamma_c$, is known as the "broken phase". In this phase the PT symmetry of the normal modes is "broken"—they are no longer PT-symmetric. In the broken phase the frequencies come in two complex conjugate pairs. Again, the pairs are identical but for the sign of the oscillating frequency.

There is one exponentially decaying and one exponentially growing mode, both of which oscillate at the same frequency. The symmetry is clearly broken, since each mode grows in one time direction and decays in the other.

For large $\gamma > \gamma_c$, the modes have no oscillatory part, and simply blow up or decay away exponentially. These modes correspond to the overdamped modes of a single oscillator. This "exponential region" is of little interest to us, and is largely ignored.

Once the eigenfrequencies $\omega_\pm$ are known, the eigenvectors $\begin{pmatrix} V_1 & V_2 \end{pmatrix}^{\mathrm{T}}$ associated with each mode are determined up to a complex amplitude A by the matrix equation (1.16). The eigenvectors can easily be assessed numerically for any parameters. In the exact phase, the full solutions for the $\omega_\pm$ modes are given analytically by

$$\begin{pmatrix} V_1 \\ V_2 \end{pmatrix}_\pm = \begin{pmatrix} 1 \\ -e^{i\varphi_\pm} \end{pmatrix} A e^{i\omega_\pm \phi} \tag{1.20}$$

with the usual real part convention, where

$$\varphi_\pm = \frac{\pi}{2} - \arctan\left[\frac{1}{\gamma}\left(\frac{1}{\omega_\pm}\frac{1}{1-\mu^2} - \omega_\pm(1+c)\right)\right] \tag{1.21}$$

The exact phase exhibits left and right side voltage oscillations of equal amplitude. The phase shifts $\varphi_\pm$ between the gain and loss side oscillations are the expected $\{0, \frac{\pi}{2}\}$ when $\gamma = 0$, and as γ increases the phase shifts of each mode converge toward a coalescence at $\gamma = \gamma_{PT}$.

The exact phase solutions are precisely the real-valued, pseudo-Hermitian solutions anticipated earlier for a system with PT symmetry. We also witnessed the spontaneous symmetry breaking controlled by the gain/loss parameter γ at the transition from exact to broken phase, as predicted.

1.5 Non-Conservative Systems

PT-symmetric setups like our dimer should be understood in the greater context of non-conservative systems—that is to say, systems that exchange energy with the environment and hence lack energy conservation. In this context PT symmetry is a special case which represents a precarious balancing act. The mechanisms of energy loss must be exactly compensated by equal and opposite energy gain mechanisms.

Due to the constraints of real electronic components, the dimer is never truly balanced. In reality there is only the matter of how close we can come to the symmetric case. Moreover, some interesting phenomena require intentional deviation from this equal-and-opposite configuration. In fact, many of the strange phenomena observed with our dimer are rightly attributed to the system-environment interactions, not to the symmetry, even with the dimer in the symmetric configuration.

To put these considerations into context and understand the effect of imbalances on our experimental system, we must analyze the more general case of the imbalanced dimer, depicted in Figure 1.7.

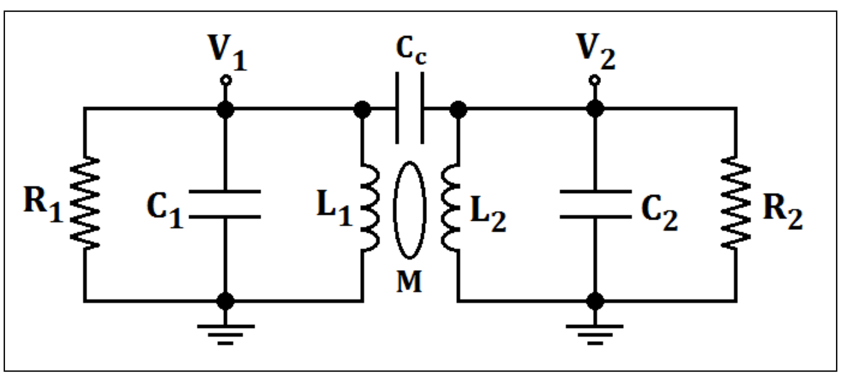

Figure 1.7: An unbalanced dimer, with arbitrary left and right component values.

For the imbalanced dimer, the inductor voltage and current rule equations become

$$V_1 = L_1 \dot{I}_1^L + M \dot{I}_2^L \qquad \frac{V_1}{R_1} + C_1 \dot{V}_1 + I_1^L + C_c(\dot{V}_1 - \dot{V}_2) = 0 \tag{1.22}$$

$$V_2 = M \dot{I}_1^L + L_2 \dot{I}_2^L \qquad \frac{V_2}{R_2} + C_2 \dot{V}_2 + I_2^L - C_c(\dot{V}_1 - \dot{V}_2) = 0 \tag{1.23}$$

Proceeding as in the PT case, with the $e^{i\omega' t}$ convention, and recasting the equations in matrix form,

$$\begin{pmatrix} \frac{1}{\omega}\frac{\lambda}{\lambda-\mu^2} - \omega(1+c) + i\gamma_1 & \omega c - \frac{1}{\omega}\frac{\mu}{\lambda-\mu^2} \\ \omega c - \frac{1}{\omega}\frac{\mu}{\lambda-\mu^2} & \frac{1}{\omega}\frac{1}{\lambda-\mu^2} - \omega(\chi+c) + i\gamma_2 \end{pmatrix} \begin{pmatrix} V_1 \\ V_2 \end{pmatrix} = 0 \tag{1.24}$$

where the parameters $\omega_1 = \frac{1}{\sqrt{L_1 C_1}}$, $z_1 = \sqrt{\frac{L_1}{C_1}}$, $\omega = \frac{\omega'}{\omega_1}$, $\gamma_1 = \frac{z_1}{R_1}$, $\gamma_2 = \frac{z_1}{R_2}$, $c = \frac{C_c}{C_1}$, and $\mu = \frac{M}{L_1}$ are scaled using component values of the oscillator on side 1, and the parameters $\lambda = \frac{L_2}{L_1}$ and $\chi = \frac{C_2}{C_1}$ relate the inductance and capacitance of the two sides.

It is easily confirmed that PT symmetry arises as a special case of the general form, being attained when $\lambda = \chi = 1$ and $\gamma_2 = -\gamma_1 = \gamma$. Under these conditions (1.24) reduces to (1.16).

At this point the eigenfrequencies can be found by setting the determinant of the coefficient matrix equal to zero, and numerically solving the characteristic polynomial. Analysis of the eigenmodes is carried out in the following section.

1.6 Dimer Eigenmode Maps

In this section we investigate the normal modes of the dimer as set up in the previous section. Since the free parameter space is now much larger, the analysis is more complicated than for the PT symmetric case. Sticking to the precedent of observing mode behaviors versus γ, we examine the unbalanced system via charts with γ_1 and γ_2 on the axes. On such charts information about the modes is represented by contour plots and/or shading.

With the γ taken as variables, choosing λ, χ, c, and μ adequately specifies the system (L_1 and C_1 scale out unless units are desired). Therefore if λ, χ, c, and μ are specified for a given γ_1/γ_2 chart, then each point on the chart represents a single point in the dimer parameter space (up to the scaling by L_1, C_1, of course). Associated with each such point of parameter space are four normal mode frequencies. But just like the

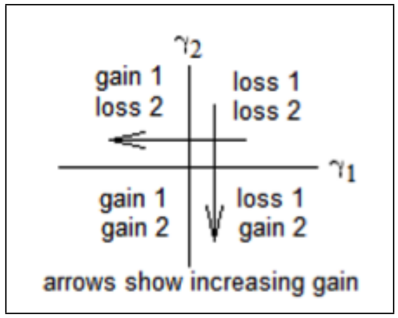

Figure 1.8: Arrows indicate directions of increasing gain on a γ_1/γ_2 chart. These charts are used to examine the functional behavior of the modes as the γ values are manipulated.

PT case, the modes in the oscillating regions (the ones of interest to us) are redundant in the form of real-conjugate pairs, and can be thought of as only two distinct modes. For large γ magnitudes there are, again, overdamped type modes with pure exponential behavior.

To summarize, after fixing λ, χ, c, and μ for a given γ_1/γ_2 chart, each point on the chart can be associated with two distinct normal modes. Each mode at the point is defined by its complex frequency (and an eigenvector). Throughout the rest of the section we use the pure capacitive coupling $c = 0.2$, $\mu = 0$, unless otherwise specified.

1.6.1 Stability

From an experimental standpoint, it is essential to know about the circuit's stability. The circuit is said to be *unstable* if it has one or more exponentially growing mode. Otherwise it is called *stable*. The stable circuit is desireable since all transient solutions are decaying. For decaying transient signals the linear equations of motion will always remain valid. Additionally, the decay of the transients allows the isolated study of steady-state phenomena. Under the $e^{i\omega t}$ convention, unstable modes have negative imaginary part.

On an unstable circuit configuration, any noise will activate the exponentially growing mode(s). Such a mode inevitably grows to the point of op-amp saturation in the negative resistor, rendering the linear equations of motion, and thereby our theoretical framework, invalid. Moreover, non-vanishing transients make it impossible to isolate steady-state solutions when the system is driven.

The stability regions of the imbalanced dimer are quite sensible. Figure 1.9a indicates universally applicable rules about the dimer stability, which apply regardless of the parameters of λ, χ, c, and μ. The yellow region shows if both sides of the dimer are lossy (no gain present), the system is always stable. The pink region shows that if there is more total gain than total loss, the system is always unstable. If gain is present on one side, but there is more total loss than gain, the stability of the system depends on the other parameters. Somewhere in the grey "transition region" lies a line seperating the stable part of the chart from the unstable. The boundary line itself has no imaginary frequency component and is neither stable nor unstable.

A specific case with PT symmetry and capacitive coupling is shown in Figure 1.9b. The spontaneous symmetry breaking of Section 1.4 is visible along the diagonal line $\gamma_1 = -\gamma_2$.

In Figure 1.9b, for large γ_1, the boundary line approaches the line $\gamma_2 = 0$. This is because as $\gamma_1 \to \infty$, V_1 essentially becomes grounded. The coupling capacitor acts simply as a capacitance to ground, and the value of γ_2 controls the gain/loss character of a mode dominated by the γ_2 side of the circuit. This type of feature should generalize to arbitrary parameter choices. With respect to stability then, the choice of parameters only controls the shape of the boundary line in the vicinity of the origin.

1.6.2 Normal Mode Contour Plots

Contour plots help study the properties of the modes. The plots shown here have blue lines of constant frequency and red lines of constant Q. In Figure 1.11 only one mode is represented at each point, which can be misleading, so interperet carefully. Different modes are shown at different points. In these figures the red and blue lines always correspond to the dominant (most unstable or least stable) mode, with most-negative imaginary frequency component. Lines of constant Q are shown only where the dimer

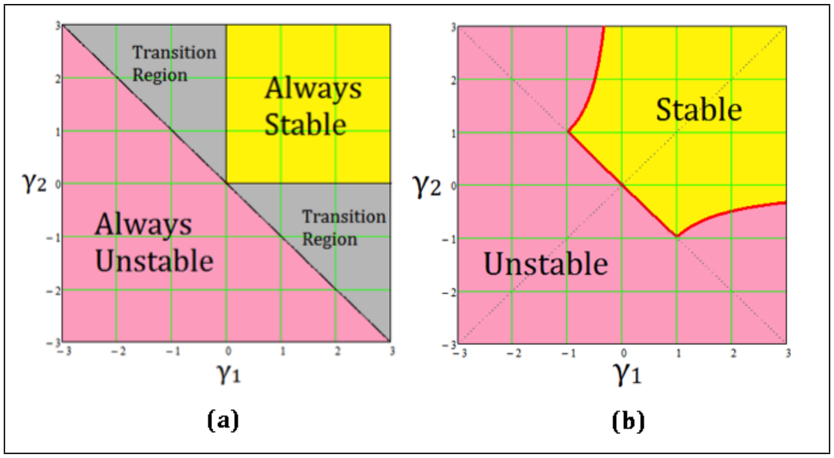

Figure 1.9: Stability regions as functions of γ_1 and γ_2. (a) For any set of dimer parameters. (b) For the balanced case $\lambda = \chi = 1$, with $c = 0.2$ and $\mu = 0$.

is stable, and the outermost red contour $Q = \infty$ is the stability boundary. The full elucidation of these charts and the modes they represent is complex beyond the scope of this section. Experimentally useful details will be the ones highlighted here.

The γ values appearing on the axes of charts in this section are scaled by the quantity γ^1_{PT}, the PT symmetry breaking point for the equivalent PT-symmetrical dimer with L_1 and C_1, as given by (1.18) . Explicitly, $\gamma_{1,2} \leftarrow \frac{\gamma_{1,2}}{\gamma^1_{PT}}$. In the balanced case, this means the exact phase PT symmetry line is the negative sloping diagonal interval $\gamma_2 = -\gamma_1$ from $(-1, 1)$ to $(1, -1)$.

Scaled frequencies also appear, labeling blue lines on the charts. They are normalized such that 1 and -1 represent the oscillating frequencies $\omega^0_\pm$ of the two modes at the origin. Explicitly, the scaling is $f = \frac{\mathrm{Re}(\omega) - (\omega^0_+ + \omega^0_-)/2}{(\omega^0_+ - \omega^0_-)/2}$.

Figure 1.11a shows the balanced case $\lambda = \chi = 1$. Figure 1.11b shows an example with imbalanced inductance, where $\lambda = 1.01$, $\chi = 1$. In the balanced case, the stability boundary and the crossover from low-frequency-mode dominance to high-frequency-mode dominance both occur along the diagonal PT symmetry line. In the imbalanced case those three boundaries no longer coincide. Adding more imbalance exaggerates the deviation. But tailoring the LC imbalance to bring the resonant frequencies of each side back together gives imbalanced modes that look more similar to the balanced modes.

The stable region is outlined by the outermost red Q line. The stable region has two decaying modes. Reflection of the stable region across $\gamma_1 = \gamma_2$ gives the region with two growing modes. In between those two lies the region with one growing and one decaying mode.

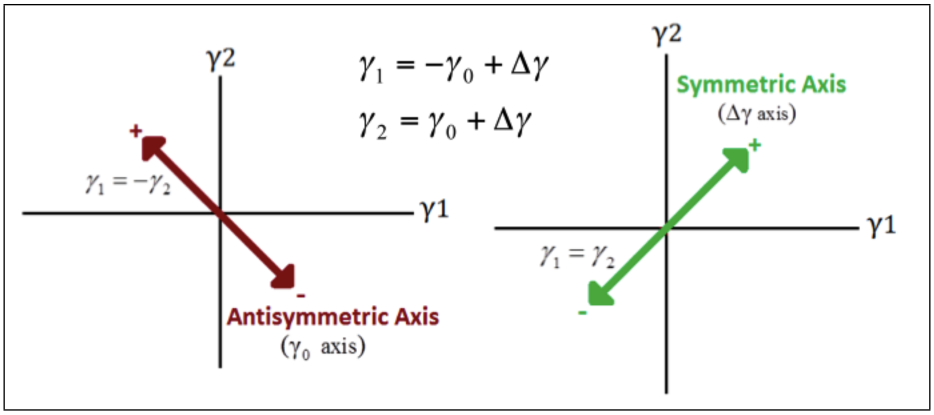

Figure 1.10: The symmetric and antisymmetric axes. The symmetric axis indicates total gain *or* loss, while the antisymmetric axis indicates balanced gain *and* loss. The positive direction of each axis moves up the chart.

To understand what's going on, it helps to look at the diagonal axes rather than the γ_1, γ_2 axes. We call the axis parallel to the PT symmetry line $\gamma_1 = -\gamma_2$ the *"γ_0 axis"* or *antisymmetric axis*. The axis perpendicular to the PT symmetry line and parallel to the line $\gamma_1 = \gamma_2$ is the *"$\Delta\gamma$ axis"* or *symmetric axis*. Figure 1.10 illustrates the axes and gives the sign conventions.

The symmetric part $\Delta\gamma$ is the average $\frac{\gamma_1+\gamma_2}{2}$. $\Delta\gamma$ dictates the *total* amount of loss or gain in the system. The antisymmetric part γ_0 is the half difference $\frac{\gamma_2-\gamma_1}{2}$. γ_0 dictates the amount of *balanced* loss and gain which is present on opposite sides of the dimer.

The mode behavior along paths parallel to these axes is relatively simple. From Figure 1.11 it is evident that moving parallel to the symmetric axis leaves frequency approximately constant while Q varies. Moving parallel to the antisymmetric axis leaves Q approximately constant while frequency varies. That is to say, lines of constant frequency point more or less in the $\Delta\gamma$ direction, while lines of constant Q point more or less in the γ_0 direction.

Figure 1.12 looks at the mode frequencies versus $\Delta\gamma$ at two values of γ_0, for parameters similar to those found in the laboratory (they were actual parameters from a PCB dimer prototype). The $\Delta\gamma$ range shown is quite large, note that it traverses the entire long diagonal of the charts in Fig. 1.11.

Evidently as $\Delta\gamma$ increases, the non-oscillating exponential component decreases linearly, while the oscillating frequency is roughly constant. When $\gamma_0 = 0$, the two modes become stable at the same point despite the imbalance. However as γ_0 is traversed outwards, the two modes' transistions to stability spread apart, and are quite seperated by the time $\gamma_0 = 0.95$.

These findings present potential experimental difficulties. For example, something we'd like to be able to do is to record frequencies near the PT symmetry line. Our standard technique for measuring normal mode frequencies is to scan through frequency with a sinusoidal drive and look for the resonant peaks. But near the PT line the imbalanced dimer is unstable, rendering that measurement impossible. Okay, fortunately the frequency won't be perturbed too much if we add symmetric loss in the $\Delta\gamma$ direction until reaching the stable region. This can work fine near $\gamma_0 = 0$. But out at $\gamma_0 = 0.95$, it's worse. First of all, a greater amount of symmetric loss is needed to stabilize the system. Second, once the system is stable, one mode has a far greater Q than the other, because of Q's reciprocality with the linearly varying $\mathrm{Im}(\omega)$. The peak with low Q will be rendered invisible. To make a successful measurement, we must add even more $\Delta\gamma$ until the Q values are of the same order. At this point one begins to wonder if the frequency really hasn't been affected. Issues like this one are not back-breaking, but care must be taken with them.

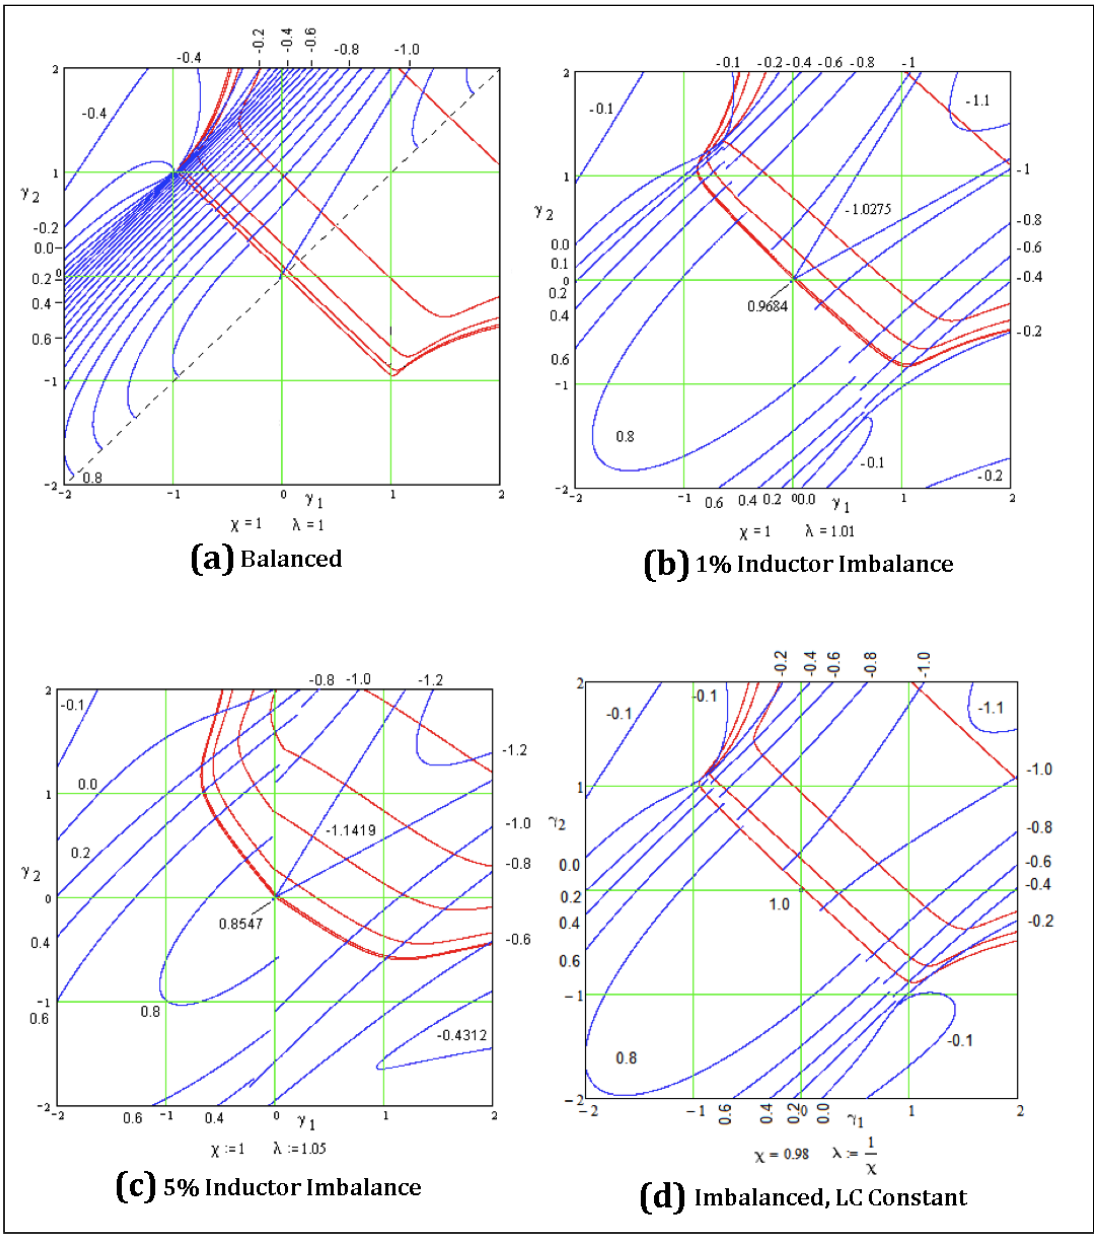

Figure 1.11: Normal mode contour plot, generated numerically. Blue lines are of constant oscillatory frequency. Red lines are of constant Q. Only one mode is represented at each point. All lines shown correspond to the mode of most-negative imaginary part at that point (most unstable or least stable mode). Frequency values shown are scaled so that 1 and -1 represent the two frequencies at the origin. Axes are scaled to γ_{PT}^{1}, the symmetry breaking value for a PT dimer with L_1C_1 on both sides. Q values are shown only in the stable region. The outermost Q line represents the stability boundary. The discontinuity seperating the sets of blue lines is the switching of which mode is dominant. Each panel has $c = 0.2$, $\mu = 0$. (a) Symmetric dimer $\lambda = \chi = 1$. (b) One percent inductor imbalance $\lambda = 1.01$, $\chi = 1$. (c) Five percent inductor imbalance $\lambda = 1.05$, $\chi = 1$. (d) About two percent each inductor and capacitor imbalance, $\chi = 0.98$ and $\lambda = (1/0.98)$ such that $L_2C_2 = L_1C_1$.

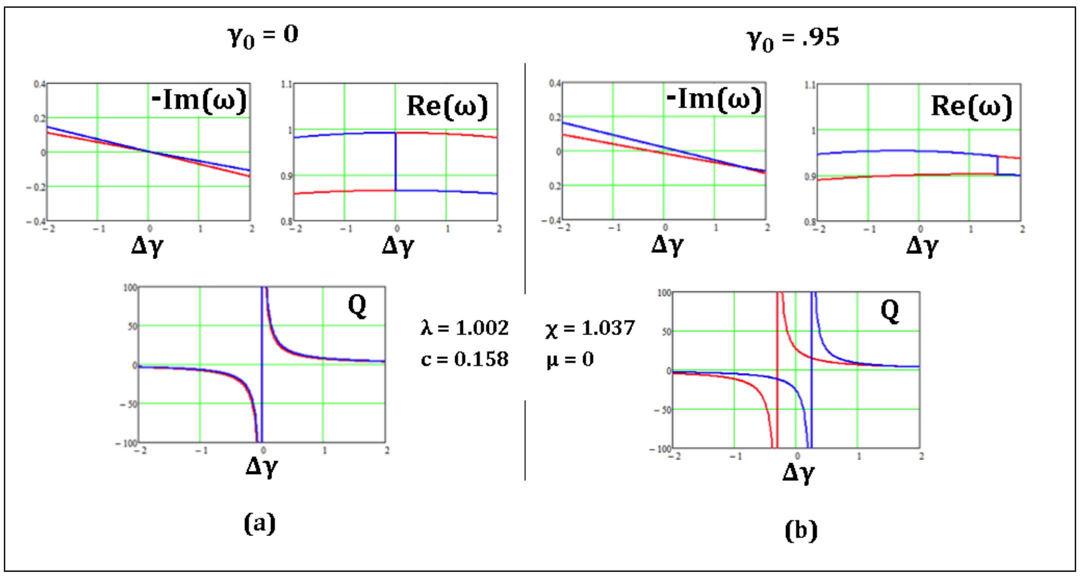

Figure 1.12: Numerical mode frequency components and Q plotted versus $\Delta\gamma$ for various γ_0. The experimentally realistic parameters $\lambda = 1.002$, $\chi = 1.037$, $c = 0.158$, $\mu = 0$ were actual parameter values present on a PCB test dimer. The $e^{i\omega\phi}$ convention is still assumed. Positive values of $(-\mathrm{Im}(\omega))$ indicate exponential growth. Positive values of Q indicate stability. $\mathrm{Re}(\omega)$ is the oscillatory frequency. Blue always represents the dominant mode. Ignore the vertical jumps at mode dominance transition points and where Q passes through ∞. The horizontal grid spacing is 1 in the scaled γ units.

1.7 Summary

We've now understood LRC oscillators and the concepts of PT symmetry. We've seen how PT symmetry can arise in electronics as a special case of coupled LRC oscillators, once negative resistance is made accessible. And we've gotten some grasp on the dynamics of the general LRC dimer throughout its vast parameter space. In Chapter 2 these findings will provide context for understanding the experiments performed with the dimer, and interpreting the experimental results.

Chapter 2

Experiments and Results

Numerous experimental results have been obtained with the PT dimer. This chapter will summarize the experiments we've conducted and divulge the experimental techniques used to acquire the data. Much more detailed analysis and discussion of the results can be found in our publications [1–4], and more general experimental details reside in Chapter 3 of this thesis. Redundancies will be avoided, and the focus here will be expanding on the experimental technicalities addressed only briefly in earlier publications.

2.1 Dimer Modes and Spontaneous Symmetry Breaking

Our first experiment with the dimer was to confirm the existence of the theoretically predicted normal mode solutions. Results for both the exact and broken phase appear in Figure 2.1. These results were initially published in [1]. The experiment was performed on the 30kHz dimer, where the trimmable parameters were the gain-side gain, and the loss-side capacitance (see subsection 3.4.1).

In the exact phase, mode frequencies were directly observed by balancing gain-loss and slightly imbalancing the capacitance, then correcting for the imbalance. Specifically, gain-loss balance $\gamma_1 = \gamma_2$ was attained by following the procedure outlined in Figure

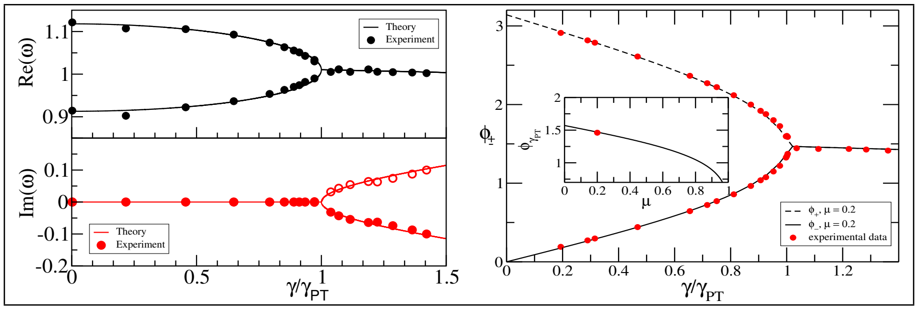

Figure 2.1: Experimentally determined normal mode frequencies for the PT-symmetric dimer with $\mu = 0.2$ and $c = 0$, as a function of γ/γ_{PT}. Lines indicate theoretical values, while symbols indicate experimental data. The left panel shows frequency data for each mode. The right side shows the phase difference between V_1 and V_2 for each mode. Frequency and phase values extracted directly from waveform captures. In the left panel, frequency is reported as the scaled value ω/ω_0. For the theoretical data $\omega_0 = 1/\sqrt{LC}$, where the dimer is symmetric. For the experimental data the capacitances are slightly imbalanced, so the scaling frequency $\omega_0 = 1/\sqrt{L\langle C\rangle}$ represents the averaged capacitance. Likewise, on the horizontal axis the value $\gamma = z_0/R$ is given by $z_0 = \sqrt{L/\langle C\rangle}$.

3.9a. This process forced the imaginary part of the frequency to be zero, and so the imaginary frequency data was automatically recorded as zero. At this point the loss side capacitance was greater than the gain side capacitance by an amount ΔC, as dictated by 3.9b. Tiny changes to the capacitance and gain-side gain from this point caused oscillations of one or the other mode. By carefully controlling the knobs, oscillation of either mode could be isolated and held below saturation. This maintaining of the oscillation at small amplitudes below saturation is also known as "marginal" oscillation. For each mode, once the dimer was brought to a state of marginal oscillation, oscilloscope waveform capture recorded $V_1(t)$ and $V_2(t)$, the voltage data at each side of the dimer. These data were were analyzed for real frequency, phase relation, and amplitude.

Data in Fig. 2.1 were reported as the scaled values $\frac{\gamma}{\gamma_{PT}}$ and $\frac{\omega}{\omega_0}$. These expressions must be defined more precisely. The experimentally measured frequency is ω. The symmetry breaking point γ_{PT} is determined by (1.18). The other two definitions take more subtlety. In the theoretical case of perfect symmetry, the definitions $\gamma = \frac{1}{R}\sqrt{\frac{L}{C}}$ and $\omega_0 = \frac{1}{\sqrt{LC}}$ are obvious. In the experimental case where C is not balanced, they break down. In that case, we use the average value of capacitance to construct the definitions. So the experimental data are, in principle, reported as $\gamma = \frac{1}{R}\sqrt{\frac{L}{\langle C\rangle}}$ and $\omega_0 = \frac{1}{\sqrt{L\langle C\rangle}}$. However, rather than basing these reported values on the nominal natural frequency, it is more accurate to report based on the experimentally determined natural frequency of 31.99kHz (see subsection 3.4.1). So the experimental data points in Fig. 2.1 actually reflect the following quantities:

$$\frac{\gamma}{\gamma_{PT}} = \frac{1}{\gamma_{PT}}\frac{1}{R}\sqrt{\frac{L}{C_0}}\frac{1}{\sqrt{1+\frac{1}{2}\frac{\Delta C}{C_0}}} \tag{2.1}$$

$$\frac{\omega}{\omega_0} = \frac{\omega_{\text{meas}}}{2\pi\, 31.99\text{kHz}}\sqrt{1+\frac{1}{2}\frac{\Delta C}{C_0}} \tag{2.2}$$

where C_0, ΔC, and R are respectively the gain-side capactitance, capactance difference, and loss-side resistance, as discussed in 3.4.1, and γ_{PT} is given by (1.18). These equations apply to both the exact and broken phase data.

A slightly different method was used to capture data in the broken phase. The capacitance dial was at the balance point, and the gain dial was left at the center of its range. One node was temporarily shorted with a wire to ground, and the short was subsequently removed by pulling the wire out from the node on the breadboard. The growing oscillation triggered an oscilloscope waveform capture of $V_1(t)$ and $V_2(t)$. Waveform data was analyzed for complex frequency, phase, and amplitude. Only the exponentially growing mode could be observed with this method, since the growing mode always dominated voltage behavior. In Fig. 2.1, data for the growing case are shown (filled dots), along with their reflections across the horizontal axis (hollow dots).

2.1.1 Wire-Pull Switching

Several experiments utilized the technique of pulling a wire from a breadboard socket as a form of switching. This was a makeshift solution to the need for a non-capacitive, instantaneous, non-bouncing switch. The method involved simply pulling a snugly embedded wire, by hand, from its breadboard hole. This method's effectiveness was tested before implementation. It is non-capacitive, and was shown to allow nearly instantaneous switching. At the time of switching from contact to non-contact, no continuous contact change was visible on the oscilloscope. It was susceptible to bouncing, but bouncing effects were found to be obviously noticeable. Since the bouncing happened randomly, the patterns were not reproduceable. In experiments, if several trials in a row proved reproduceable, the possibility of bouncing could be eliminated, and the fidelity of the switching confirmed. With good wire-pulling technique, almost all bouncing could be eliminated.

2.2 Scattering

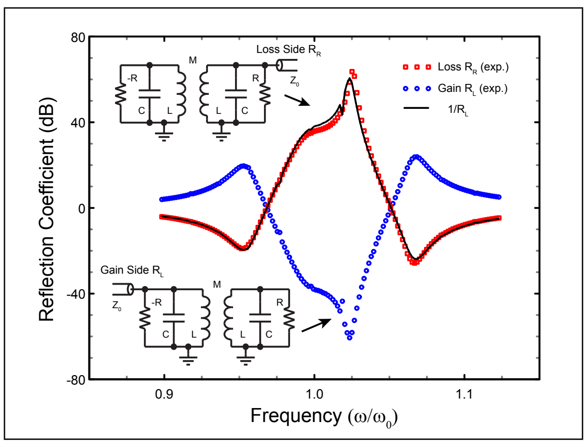

Figure 2.2: From reference [4]: Two experimental configurations associated with a simple $\mathcal{PT}$-symmetric dimer. In the lower and upper circuits, we couple a transmission line to the gain and loss sides, respectively. Preliminary experimental measurements for the corresponding reflection coefficients are shown (loss-side red, gain-side blue) along with the solid line corresponding to R_{L}^{-1} illustrate the reciprocal nature $R_{\mathrm{L}}R_{\mathrm{R}} = 1$ (see text) of the $\mathcal{PT}$-scattering. Here $\mu = 0.2$, $\gamma = 0.164$ and $Z_0 = 15.5\sqrt{L/C}$.

The next experiments investigated the scattering properties of the PT dimer [2, 4]. These experiments took several forms, but all simulated placing transmission lines at one or both sides of the dimer, and observing the scattering coefficients. We confirmed that PT systems contribute novel non-reciprocal scattering behaviors associated with their left- and right- transmissions and reflections. In all cases, input resistors simulated the presence of actual transmission lines. Forward and backward waves were extracted from the voltages and currents at either side of the TL resistors, allowing the calculation of scattering coefficients. When used, TL resistors were associated with parameter $\eta = \frac{z_0}{R_{\mathrm{TL}}}$.

These data were taken in a much more straightforward manner than those of the previous section. The dimer resistances and attached TL resistors were brought as close as possible to the perfectly balanced configuration where $\gamma_1 = \gamma_2$ and $\eta_1 = \eta_2$. Data was acquired by simply sending signals in from the signal synthesizer, measuring all relevant voltages, and extracting the scattering coefficients. The experimental data were reported directly, and compared to simulations based on the actual measured values of $\gamma_{1,2}$ and $\eta_{1,2}$.

Scattering coefficients were measured for both one-port (Fig. 2.2) and two-port configurations [2]. These scattering coefficients were determined by experiments performed on the 30kHz dimer. An additional scattering experiment aimed to show the dimer can act as a "Coherent Perfect Absorber/Amplifier" (CPA) [4]. The CPA experiment was performed on the 3MHz breadboard dimer.

In the CPA experiment, unlike usual linear scattering experiments, incident waves approached from both directions. The incoming wave amplitude on the loss side was taken as unity, with the complex amplitude of the gain-side incoming wave then being $Ae^{i\phi}$. Both incoming waves share a frequency ω. The measured quantity of interest was the "total output coefficient", defined as $\Theta = \frac{|V_{out}^L|^2+|V_{out}^R|^2}{|V_{in}^L|^2+|V_{in}^R|^2}$. The idea of this phenomenon is that at some fixed frequency, the dimer can act as either a perfect absorber, with $\Theta = 0$, or act as a perfect laser, with $\Theta = \infty$, depending on the input amplitude and phase parameters A and ϕ. The absorption happens only under extremely precise input conditions; for almost any random input amplitudes and phases, the dimer will act as a strong amplifier.

Even theoretically, the absorption point is extremely sensitive. When imbalances are introduced, the situation only gets worse. For this reason, special scanning techniques had to be used to capture the data with $\Theta < 1$ shown in Figure 2.3. Each lower data point near the absorption point in Fig. 2.3a was found by fixing frequency and scanning through values of $\phi \approx 90°$ in tightly spaced increments, then recording the minimum Θ value. Within these ϕ scans, an iterative process of measurement and resetting was used at each step, to ensure that A and ϕ were within a small tolerance level of the theoretically specified values. The portion of the bottom (green) curve in Fig. 2.3b near the minimum is an example of one of these high-precision scans.

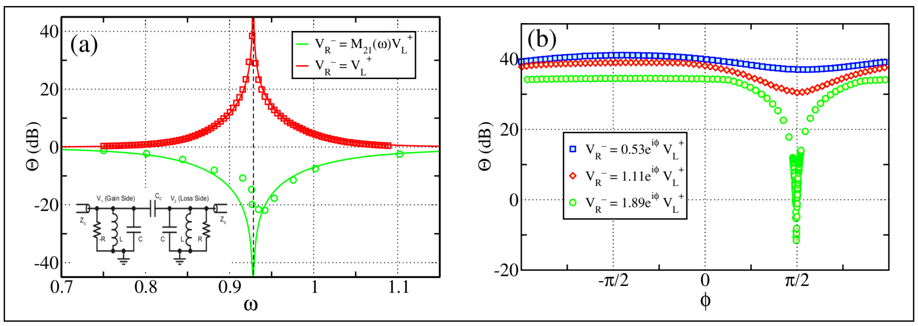

Figure 2.3: From reference [4]: (a) The overall output coefficient $\Theta(\omega)$ around the Janus amplification/attenuation frequency ω_J (vertical dashed line) for a $\mathcal{PT}$-symmetric electronic circuit coupled to two ports. The parameters used in this simulation are $\eta = 0.110$, $\gamma = 0.186$ and $c = 0.161$. The red curve corresponds to the two port coherent input excitation with $V_R^- = \mathcal{M}_{21}(\omega)V_L^+$; the blue curve correspond to a two-port input signal with $V_R^- = V_L^+$. In the former case the system acts as an perfect attenuator while in the latter as an amplifier. The dots are experimental values. (b) Plots of experimental $\Theta(\omega_J)$ as the loss side input excitation phase is changed, for several excitation amplitudes. Frequencies are scaled by ω_0.

2.3 Wave Brachistochrone

Another experimental setup investigated the dynamics of the system in the exact phase. This setup addressed a phenomenon dubbed the "wave brachistochrone" problem, demonstrating that the dimer can pass between orthogonal states faster than the limit imposed by the bandwidth theorem on conservative systems [3]. In this case the state of the system is defined by the Hamiltonian state vector $\psi = \begin{pmatrix} Q_1 & Q_2 & \dot{Q}_1 & \dot{Q}_2 \end{pmatrix}^{\mathrm{T}}$, where $Q_i = C_i V_i$. The data were taken on the 30kHz dimer.

To isolate transient dynamics of the exact phase, where $\mathrm{Im}(\omega)$ should be zero, we had to ensure that the exponential growth or decay time was many times longer than the period of oscillations. This was achieved empirically by carefully trimming the gain-side gain.

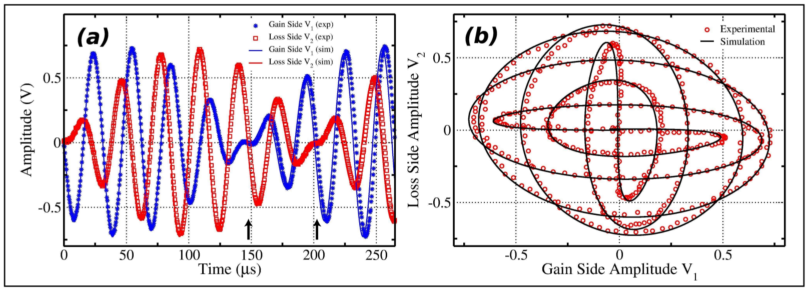

Figure 2.4: From reference [4]: (a) Gain and loss side voltages vs. time compared to the simulation. (b) Gain vs. loss side Lissajous figure for one beat period. At $t = 0$ an initial current was imposed in the gain side inductor with all other dynamical variables zero. Note that the end of the beat (indicated by the arrow near $200\mu s$) is preceded by a similar point where both voltages pass through zero (indicated by the arrow near $150\mu s$) with V_2 decreasing, and V_1 stationary. This corresponds to the complementary initial condition starting from the loss side, and illustrates an asymmetric time between the beat nodal points of oscillatory activity in the two oscillators of the dimer.

Initial conditions were imposed by connecting the +12V power supply voltage to one node through a resistance. Because of the inductor's zero DC impedance, this allowed the initial conditions $\psi = \begin{pmatrix} 0 & 0 & 1 & 0 \end{pmatrix}^{\mathrm{T}}$ or $\psi = \begin{pmatrix} 0 & 0 & 0 & 1 \end{pmatrix}^{\mathrm{T}}$ to be imposed, depending which node was connected.

Data acquisition was initiated when the system was "let go" from its initial condition by detaching the node from the power supply and resistor using the wire-pulling method of 2.1.1. The system's evolution immediately triggered a waveform capture, which was recorded and compared to a simulation. The fidelity of the initial condition was confirmed by performing successive trials. Reproduceablity indicated a "good pull", and thus successful imposition of the initial condition. The fidelity was confirmed again by close matching of the experimental waveform to the simulation. Figure 2.4 depicts an example waveform capture, along with its corresponding Lissajous figure. Performing this experiment for various values of γ, with μ adjusted at each γ to maintain a constant bandwidth $(\omega_+ - \omega_-) = \text{const}$, yielded the results depicted in Figure 2.5. These results are provided for demonstrative purposes, but must be interpreted in terms of reference [3]. Insufficient background is provided here.

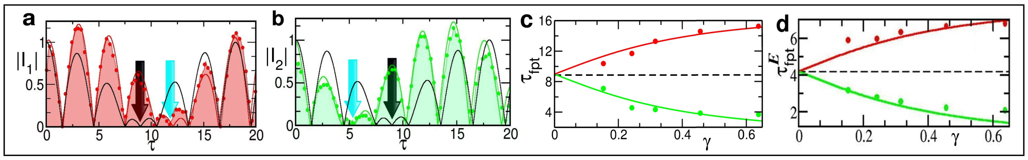

Figure 2.5: From reference [3]: A representative tachistochrone passage wave propagation for a circuit with a bandwidth constraint $\delta\omega/\omega_0 = 0.36$. (a) A typical temporal dynamics of the displacement current $I_1(\tau)$ when the initial condition corresponds to an excitation of the circuit $n = 1$ (gain side); (b) the same for current $I_2(\tau)$ when the initial condition is an excitation at circuit $n = 2$ (loss side). The black lines are numerical simulations for the passive $\gamma = 0$ circuit, while the red and green lines are numerical simulations for the active circuit with $\gamma = 0.24$ and $\mu = 0.39$. The measurements are indicated with circles. The first passage time τ_{fpt} is indicated with a cyan (black) arrow for the active (passive) dimer. (c) The extracted τ_{fpt} versus the gain/loss parameter γ. The green circles correspond to an initial condition $I_2(0) = 1$ (loss side) while the red circles to an initial condition $I_1(0) = 1$ (gain side). The solid lines indicate the theoretical result. The black dashed line indicates the beat time for a passive ($\gamma = 0$) system. (d) The first passage time τ^E_{fpt} versus the gain/loss parameter γ as it is measured from the energy exchange condition (see text). Open circles are the experimental data while lines (of similar color) correspond to the theoretical prediction $\tau^E_{\text{fpt}} = \tau_{\text{fpt}}/2$. The black dashed line denotes τ^E_{fpt} for $\gamma = 0$.

2.4 Reentrant Stability

One more interesting effect was investigated recently, taking advantage of the computer-controlled γ of the PCB prototype dimer. The concept is interesting and simple. Counterintuitively, starting from an unstable circuit configuration (the top of the vertical line in Fig. 2.6), the system can be brought stable by only *increasing* the gain (decreasing the loss, actually). Further increasing the gain then results in a return to instability. Figure 2.6 demonstrates a path through γ_1-γ_2 space corresponding to this effect. Starting at the beginning of the path, γ_1 is decreased, bringing the dimer from stable to unstable. The γ_1 side has gain, while the γ_2 side is lossy. γ_2 is then decreased to zero, bringing the circuit temporarily through a stable region.

This effect can be understood conceptually by comparison with the circuit of Fig. 2.7. In this circuit, z represents an LRC circuit with a negative resistance. Now, if

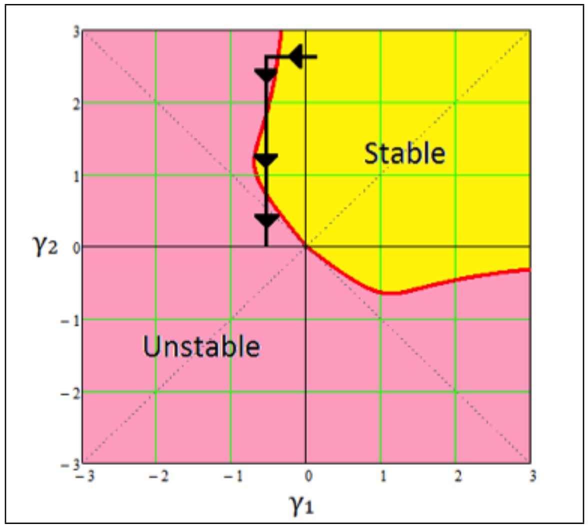

Figure 2.6: By following the depicted path through γ_1-γ_2 space, the circuit can be made to go stable→unstable→stable→unstable, all by only increasing amounts of gain. This stability diagram corresponds to actual parameters of the PCB prototype dimer, and is experimentally realistic. Compare to experimental Fig. 2.8.

$R = 0$, we simply have an additional parallel capacitance with z, which has no effect on the gain, and the system must be unstable. Likewise, if $R \to \infty$, no current passes through the right side of the circuit, and z is entirely unaffected, again forcing the system to be unstable. In between these values, current passing throught the positive resistance R causes some amount of loss, offsetting the gain of z. Somewhere along the way lies a maximum amount of net loss. If this amount happens to be greater than the gain of z, the circuit will pass through a stable configuration. In the actual dimer setup, R is replaced by a complex impedance, but the same principles apply.

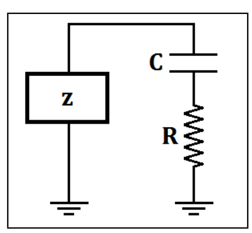

Figure 2.7: This circuit gives insight into the cause of the reentrant stability effect.

This effect was experimentally verified, as shown by the data in Figure 2.8. These data correspond to the path of Fig. 2.6. The different data sets have decreasing values of $\gamma_1 = \text{const}$ while γ_2 is varied. This is equivalent to moving the vertical line of Fig. 2.6 further and further left with each trial. The top (dark blue) data set shows that for a large enough negative γ_1, the vertical line misses the stable region entirely.

Data was acquired by varying γ_1 and γ_2 with all other parameters held fixed, and measuring V_{rms} of oscillations on the scope. When the system is stable, $V_{rms} = 0$. When the system is unstable, some nonzero value of V_{rms} is measured. The values of V_{rms} are dictated by op-amp saturation dynamics, and are unimportant here. Only the fact that they are zero or nonzero is relevant. Because increasing the gain more causes further saturation, higher up curves on this chart correspond to larger negative values of γ_1 (vertical lines further left in Fig. 2.6). The data clearly shows the expected narrowing of the stable region as the vertical line (of Fig. 2.6) is moved left, until it finally disappears entirely.

Notice that the path shown in Fig. 2.6 begins at the bottom right of Fig. 2.8, with a stable circuit. As γ_1 is increased to some unstable value, the experimental data move vertically upwards on the far right as saturation is pushed farther and farther by the gain. Once γ_1 is set at its fixed value, γ_2 is incrementally decreased. This corresponds to moving leftwards horizontally on each curve. The curves each end at $\gamma_2 = 0$, on the far left. The stable region corresponds to the gaps in the curves, where $V_{rms} = 0$ over some interval.

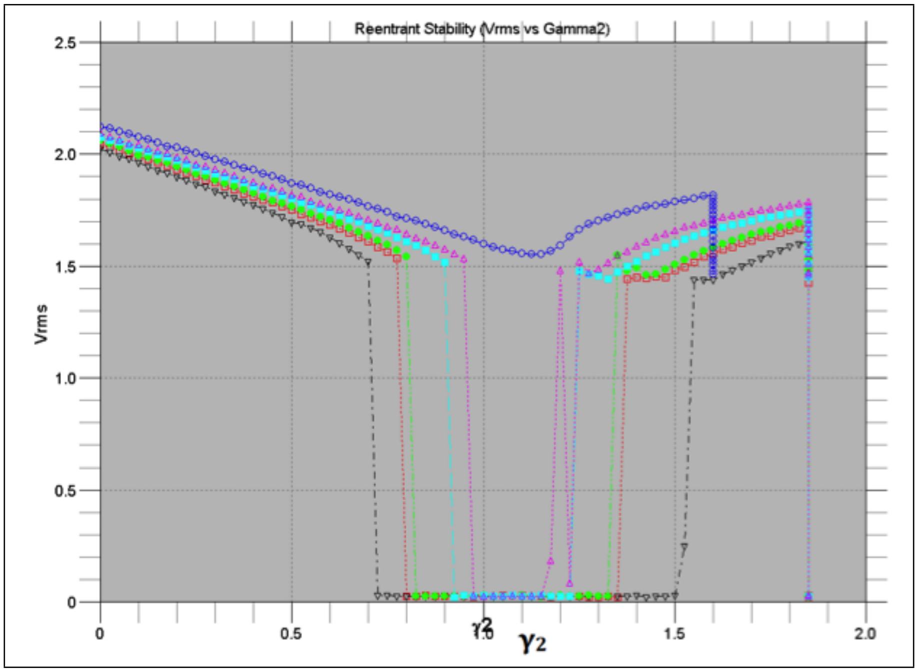

Figure 2.8: Experimental reentrant stability data. Compare to Fig. 2.6. The path shown in Fig. 2.6 begins at the bottom right, with a stable circuit. As γ_1 is increased to some unstable value, the experimental data move vertically upwards on the far right as saturation is pushed farther and farther by the gain. Once γ_1 is set at its fixed value, γ_2 is incrementally decreased. This corresponds to moving leftwards horizontally on each curve. The curves each end at $\gamma_2 = 0$, on the far left. The stable region corresponds to the gaps in the curves, where $V_{rms} = 0$ over some interval. When the system is stable, $V_{rms} = 0$. When the system is unstable, some nonzero value of V_{rms} is measured. The values of V_{rms} are dictated by op-amp saturation dynamics, and are unimportant here. Only the fact that they are zero or nonzero is relevant.

2.5 Summary

This chapter introduced the experimental results obtained with the dimers. Specific experimental techniques were addressed. In the following chapter, general experimental considerations applying to all of these experiments, and descriptions of the apparattii themselves, are discussed.

Chapter 3

Practical Considerations

The physical circuit differs substantially from the ideal circuit appearing in the dimer circuit diagrams of the previous chapters. Components are corrupted by stray impedances and by linearity limits. Attempts at balancing are impeded by component precision and thermal and temporal drift. All measurements must be interpreted with this context in mind.

This chapter details both the intentional and the unavoidable deviations from the ideal circuitry of Chapter 1. Additionally, it attempts to highlight relevant circuit design practices necessary to attain proper functioning of the dimer. The final section summarizes the various apparatii actually implemented in the laboratory, and motivates the evolution from one to the next.

The equipment regularly at use in the lab includes: A Tektronix DPO2014 Digital Oscilloscope. Two HP3325A Signal Generators. An EG&G 7265 DSP Lock-In Amplifier. Several Keithley Digital Multimeters. And several Fluke Handheld Digital Multimeters.

The oscilloscope, signal generators, and Keithley multimeters can be controlled by PC via IEEE communication with an interface in Microsoft Visual Basic. The VB control interface predates this work. The oscilloscope was a new addition, and was interfaced to allow direct measurement from and waveform capture to the PC.

3.1 The Actual Circuit

At the heart of the experimental circuit lie our friends, the simple LRC oscillators. The oscillators as depicted in Chapter 1, but with negative resistances represented in op-amp form, are the main components of the circuit. Earlier apparattii operated near 30kHz, while later apparattii operated near 3MHz. The full circuit differs from the ideal dimer in three main ways: (1) Actual components are not identical to their idealized representations, and attempts to control component values are of limited precision. (2) Parasitic conductances modify the network. And (3) additional circuitry is needed to correct for the above, to adjust parameters, and to allow measurements. The discussion of the apparatii in Section 3.4 and the full schematics, shown in Chapter 5, of the most recent apparatus will clarify the details of what additional circuitry is used.

3.1.1 Breadboards and PCBs

Both breadboarded configurations and printed circuit board configurations have been implemented in the lab. Earlier dimer models were assembled on the breadboard. Breadboarding is the inferior technique, but offers easy assembly and modification by the experimenter. Several complications are presented by the breadboard. It suffers from the lack of a true ground plane, despite being affixed to a grounded case. Inside the breadboard, the ground is provided by wires, allowing unwanted impedances to accumulate on what should ideally be a zero impedance ground. Not surprisingly, this proved problematic at 3MHz. Additionally, adjacent rows on the breadboard are linked by capacitance. The capacitance between adjacent rows was found to be approximately 2.35pF by measurement with an HP4193A Vector Impedance Meter at 50MHz. Nearby but non-adjacent rows are linked by some fraction of that capacitance. This can provide unwanted coupling between signal traces or signal and ground. Fortunately capacitance between the relevant signal nodes and ground is already included in the circuit, so stray signal-ground capacitance is only a perturbation to the desired parallel capacitance and has no effect. It is essential to ground all unused rows of the breadboard to emulate a true ground plane, prevent parasitic op-amp feedback loops, and reduce possible signal-signal interactions.

Printed circuit board (PCB) setups are superior due to having a complete ground plane, well-defined trace impedances, shorter trace lengths, lower impedance power strips, better power decoupling, and fewer hidden parasitic elements. Obviously PCBs should be used if they are attainable and if the circuit design will not need to be substantially altered once assembled. Recent dimer models are assembled on PCB. Stray capacitances remain a possible issue, but can be calculated based on trace dimensions if necessary. Again, stray signal-ground capacitances are in parallel with a large capacitor in our setup, and have no effect. Unintended feedback and signal-signal interactions should be nearly eliminated for the PCB due to the introduction of a ground plane seperating all signal points.

3.1.2 Buffering

All voltage measurements are taken at the output of a voltage following buffer implemented by op-amp with direct negative feedback. Each dimer node, then, lies on the positive input terminal of such a buffer. Buffering is necessary since oscilloscope probe capacitances would upset resonant features and cause an intrusive measurement if attached directly to the nodes.

3.1.3 Inherent Losses

With $R = 0$ in the ideal dimer configuration, and thus only L and C present, each oscillator making up the dimer does not actually have $\gamma = 0$. Indeed, as is wont to be the case in macroscopic systems, there are inherent losses. Sources of inherent dissipation occur mainly in the inductor due to resistive and magnetic core losses. Additional dissipation can come from dielectric loss in the capacitor and from resistive losses in the signal trace, but these contributions are comparitively small. The amount of dissipation inherent to the LC combination can be quantified by observing the Q factor of the LC resonance. Typical amounts of inherent loss depend on the component choices.

The inherent series inductor losses (see Section 3.2) are compensated by a negative resistance in parallel with the inductor. Since the series resistance is cancelled by a parallel negative resistance (with respect to the inductor), the cancellation is not theoretically perfect. The effect of the incomplete cancelation is not entirely known—the calculation is confounded by the frequency- and amplitude-dependent nature of coil losses. This effect is a non-trivial deviation from the theoretical dimer, and we postulate that troubles with predicting the experimental modes very close to $\gamma = \gamma_{PT}$ are related to this issue. For $\gamma \approx \gamma_{PT}$ the dimer exhibits what appears to be a chaotic interplay of the two modes, signaling the presence of some nonlinearity that may be related to this issue.

3.1.4 Balancing and Control

In order to see effects related to PT symmetry, we had better be able attain a close balance, or symmetry, between the two sides of the dimer. How close can we get to perfect symmetry? How close *must* we get? What controls allow us to adjust the dimer parameters, and how much control do they afford? These questions are addressed very generally here, and in more detail later.

Using identical nominal component values, but without carefully adjusting things, we'd expect the inductances and capacitances of the two sides of to be balanced within about 5-10%, based on the tolerances of typical components and considering potentially imbalanced parasitic elements. We can manually adjust the dimer's L and C balance by inspecting the resonances and the components carefully, scrupulously accounting for the parasitic elements, then adding small trimming capacitances and magnetic shards to trim inductances. The attainable precision of balancing essentially just depends on how much effort one is willing to put into the balancing effort. Anything more than 1%

tolerance balancing requires quite a lot of effort. Getting much closer than that would be limited by component drifts.

Balancing resistances is slightly easier, due to ease of measurement and availability of many resistor values in the lab. We can trim resistances to a somewhat higher precision, say within 1Ω. By using a precision trimming potentiometer in one of several configurations, even smaller adjustments can be made. Nonetheless, all such efforts may be somewhat vain. Imbalances will always be there, just slightly more hidden. The more important question isn't how closely we can balance components, but how much they need to be balanced. How precise we need to be depends on what effects we want to see.

When we're trying to balance resistances, we are probably trying to balance gain and loss elements in an attempt to attain constant amplitude signals with no growth or decay. This could be because the data itself is supposed to be a real frequency mode, as in section 2.1. Or it could be part of an attempt to acquire some other measurement where an observable signal is needed without pushing against the linearity limit, for example, finding the self-resonant frequency of an inductor by inducing an oscillation with negative resistance gain. But due to the system's linearity and the impossibility of exact component balancing, there will never be a truly constant amplitude behavior. Any mode, if left to itself, will eventually decay to zero or grow to the linearity limit. The best we can do is acheive behaviors which are approximately constant over lab timescales. Considering we are observing oscillations in the kHz-MHz range, times on the order of 1s are very long in the system timescale.

Parameter control has been provided in different ways at different stages of the experimental process, as needed. In earlier versions: gain and loss were trimmed by potentiometers in the negative resistance gain units, the capacitance difference between the two sides was trimmed by a variable capacitor, and mutual inductance coupling could be adjusted by coil seperation. In later versions, capacitive coupling could be adjusted by swapping components, gain and loss were still tied to potentiometers, and the capacitance difference was left fixed. Even later, vactrols allowed computer control of gain and loss. As a rule of thumb, things with knobs (i.e. potentiometer, variable capacitor) give very good relative control of parameters at the expense of absolute control.

The inductors, op-amps, and vactrols deserve more in depth analysis. Inductors and op-amps are tackled in the upcoming sections. Vactrols will be exposed at length in Chapter 4.

3.2 Inductors

Physical electronic components are never what they seem. The idealized components that make up circuit diagrams are merely useful models. A physical device is defined not by its representation on a circuit diagram, but by its physical behavior, in the form of its I-V characteristic. A general way to represent a physical two-terminal device is by its complex impedance $z(\omega, V_0)$, which defines the current $I = \frac{V}{z(\omega,V_0)}$ across the terminals for any magnitude and frequency of a voltage oscillation $V = V_0 e^{i\omega t}$.

On the other hand, the I-V characteristic of physical linear devices (resistors, capacitors, inductors) *can* be modeled well by *combinations* of the idealized linear components. For example, a resistor has a small capacitance between its leads which will become significant at high enough frequencies. A capacitor has dielectric losses that can be modeled by a series resistance. Any length of wire has both inductance and resistance between its terminals. When statements like these are represented on a circuit diagram, we are making a theoretical model of the physical system. The model is as good as the predictions it makes about $z(\omega, V_0)$.

Resistors and capacitors are actually modeled well by their theoretical counterparts. Inductors, however, have much more going on. To acheive large inductances requires the use of either a great length of wire allowing for many turns, or a magnetic core material allowing maximal flux linking between the turns. Either way, complex losses are introduced. These losses generally depend on both frequency and amplitude.

In the case of any closely turned lengths of wire, frequency-dependent losses are introduced by the skin effect and the proximity effect, which force current to pass through a smaller cross-sectional area of the wire, causing increased resistive losses with increasing frequency. Using Litz wire and special winding techniques can reduce, but not eliminate, these losses.

Using magnetic core materials reduces the need for long lengths of wire, but magnetic cores introduce losses of their own, called magnetic hysteresis losses. These are associted with the work done to move around the core's magnetic domain boundaries, and can be found by the area enclosed by the B-H magnetization curve. Power lost by hysteresis is tied to both amplitude and frequency. Higher frequencies increase the number of magnetization cycles per second, causing increased power loss. Meanwhile, large amplitudes increase the area of the magnetization curve, also causing increased power loss. Additional inductor nonlinearity is introduced by magnetic saturation of the core, when large coil currents cause the magnetic material to be fully magnetized.

So magnetic core inductors are both amplitude and frequency dependent, while air-core inductors are mainly frequency dependent. However, in most cases size constraints limit us to the use of magnetic core inductors. The effect of amplitude dependence is that if a growing current passes through a magnetic core inductor, its lossiness will increase until an amplitude equilibrium is reached. This effect can be minimized by using the core at small magnetizations. We typically model both resistive and magnetic losses as resistances in series with the inductance.

In addition to coil losses, coil capacitance is important. Each turn of the coil, in a sense, has a capacitance to the other windings, causing a distributed capacitive effect along the coil. This is typically modeled by a capacitance in parallel with the inductance.

The model of a physical inductor shown in Figure 3.1 is a logical and common choice. The dissipations can be modeled by a resistor for any particular case, but due to the amplitude and frequency dependence, the value of R depends on the frequency

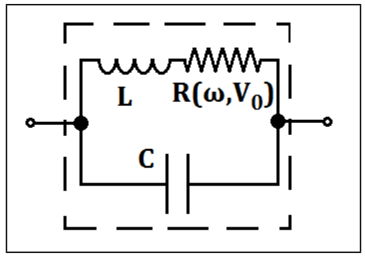

Figure 3.1: A model for a physical inductor. The series resistance is due to resistive and magnetic losses. The parallel capacitance is due to coil self-capacitance. Note that an equivalent parallel resistance $R_{\|} = \frac{z_0^2}{R}$ can be found from the series resistance R by matching the Q according to $R_{\|} = z_0 Q$.

and magnitude of current through the inductor. So long as the magnetic material is far from saturation ($I_0 \ll I_{sat}$), the resistance value should be essentially amplitude independent and depend only on frequency. Note that the model for a physical inductor is actually a damped LC resonator, and inductors really do display this self-resonant behavior.

Commercially available inductors are typically designed for use at a specific frequency. Several useful quantities besides the inductance are designated in an inductor's specifications. The "test frequency" f_{test} is the frequency where the listed Q-factor of the inductor was evaluated. This is usually also the frequency where that inductor's Q is the highest, and the frequency at which it is intended to be used. The Q is the quality factor of the resonance associated with an LC combination, which includes the inductor and an additional capacitor, that resonates at f_{test}. The self-resonant frequency f_{SR} gives an idea of the self-capacitance on the inductor. The approximate coil current which causes magnetic core saturation is I_{sat}.

For a given amount of inductance and a desired frequency, smaller devices rely more heavily on magnetic effects and thus typically have greater losses and lower Q. Also, due to the presence of less total magnetic material, smaller devices have a lower current saturation level. And, like for any device, smaller inductors have smaller associated capacitance.

Eliminating amplitude dependence of the inductor is crucial to the dimer experiment. Why? Amplitude dependence would alter the equations of motion, making them nonlinear. Now, the equations are actually already nonlinear due to the op-amp saturation. However, the op-amp saturation is discontinuous and very simple (see Section 3.3), whereas the inductor saturation is gradual. In order to have the best control over and knowledge of what is going on with our setup, we want to make the op-amp saturation the only relevant non-linearity. This amounts to making sure $R(\omega, V_0)$ doesn't change much over the amplitudes allowed by our op-amp.

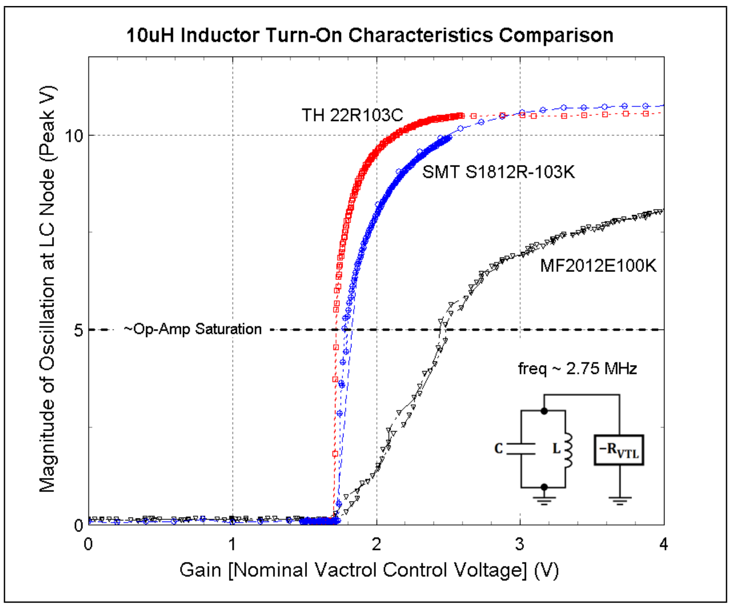

Figure 3.2: Linearity test of three 10μH inductors at ∼2.75MHz. A parallel LC configuration with C = 330pF was set up in parallel with a variable negative resistance, the value of which was controlled by a vactrol. Oscillation amplitudes at the LC node were measured by the oscilloscope through a buffer while gain was increased by reducing the negative resistance. The inset depicts the setup schematic. The vertical axis indicates oscillation amplitude at the LC node. The horizontal axis corresponds to the amount of gain; it is a voltage which controlled the vactrol's conductance. "TH" denotes through-hole style technology, whereas "SMT" denotes surface-mount. Red symbols denote the 22R103C, blue the S1812R-100K, and black the MLF2012E100K.

Figure 3.2 demonstrates a linearity test of effective resistance versus amplitude for three 10μH inductors at ∼2.75MHz. The linearity is seen by observing the slope with which the oscillation progresses from just activating (its first excursion above zero) up to the op-amp saturation point. If the slope is gradual, as it is for the MLF2012E100K, then amplitude dependent losses are causing an equilibrium, meaning there is a strong nonlinearity. The sharp slope of the other two trials tells us that those inductors are fairly linear over the amplitude range allowed by the op-amp. The Q of each coil can be estimated by the amount of gain necessary to induce some oscillation. This chart suggests that the three inductors tested have similar Q at this frequency. The tiny MLF2012E100K surface mount inductor was seen to perform poorly with respect to linearity. The somewhat larger surface-mount S1812R-103K and the much larger 22R103C through-hole inductor performed well. The surface mount inductors used on our PCB apparatii are comparable to the S1812R-103K.

3.3 Op-Amps

Operational amplifiers (op-amps) are the active devices which allow us to implement voltage following buffers and negative impedance converters. Theoretically they are three-terminal devices, where the three terminals include an inverting input V_{-in}, a non-inverting input V_{+in}, and an output V_{out}. In principle they operate according to a few simple "golden rules". (1) No current flows into or out of the two input terminals (infinite input impedance). (2) When the inputs are unequal, the output is $\pm\infty$, depending whether the higher voltage is at the inverting or non-inverting input. (3) When external feedback allows it, the amp does whatever necessary to make the inputs equal, making the output voltage finite. These are idealizations.

In practice, op-amps work by multiplying the difference $V_{+in} - V_{-in}$ by a number called the "open loop gain". Ideally the open loop gain is infinite, but in reality it is typically on the order of 10^4 or larger during proper operation. Their output is limited by the power supply rails, so when the inputs are unequal the output voltage is equal to the positive or negative power supply voltage. We'll assume the use of symmetric power rails and call the supplies $\pm V_s$. For normal operation of the amp, feedback loops are used to equalize the inputs, allowing finite output. With or without feedback, the fact remains that the amp can never output more than its power supply levels, and any time it tries to, it is said to be saturating. The input impedance of an actual op-amp is large, but not infinite. All op-amps have a frequency bandwidth over which their normal operation is defined. Outside of this bandwidth the open loop gain, which tends to drop off with frequency, is no longer large, and the slew rate (maximum $\frac{dV_{out}}{dt}$ provided by amp) can't keep up with signal changes.

Different op-amps have different strengths and weaknesses. Some have extremely high input impedance and open loop gain but operate over a small bandwidth. Some can utilize almost their entire power supply voltage range, while others saturate at some smaller fraction of their power supply levels. Some have extremely high bandwidth, but may sacrifice input impedance.

We found that "video buffers", op-amp chips intended for use with video signals, provide one of the best options for op-amp applications at any frequency. They are rated for use up to hundreds of MHz, more than ten times greater than our current experimental bandwidth. Input impedance for these devices remains large despite the bandwidth. They usually come as fixed-gain units (referring to the closed-loop gain), with preset internal feedback networks, meaning they can only be used for specific applications, but for our purposes they are just right. Figure 3.5 compares manufacturer-specified gain ratios versus frequency for various op-amps.

When an op-amp is functioning as an ideal op-amp, it has the capability to work as a perfect voltage following buffer or negative resistance element, at least over its power supply voltage range. The feedback configurations necessary for these functions are shown in Figures 3.3a and 3.3b respectively. In both cases feedback ensures $V_{+in} = V_{-in}$, which can simply be called V_{in}.

On an ideal, but voltage-limited, buffer, the output exactly matches the non-inverting input for any input within the supply range, and equals the supply voltage when the

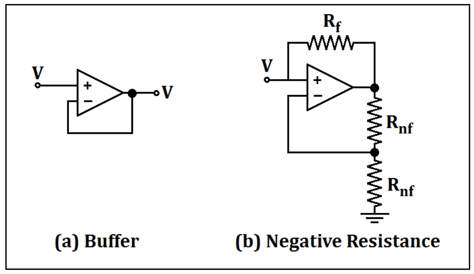

Figure 3.3: (a) A voltage-following op-amp buffer. (b) A negative resistance op-amp configuration. The negative resistance value is the positive feedback R_f.

input leaves the range:

$$V_{out} = \begin{cases} V_{in} & \text{if} \quad |V_{out}| < |V_s| \\ \pm V_s & \text{if} \quad |V_{out}| \geq |V_s| \end{cases} \tag{3.1}$$

In reality, both the input and output pins will saturate at some $|V_{sat}| < |V_s|$, where the saturation voltage $\pm V_{sat}$ depends on the supply levels and the amp's design, and is usually specified as the voltage swing. Figure 3.6 confirms that the ADA4862 video buffer operates well as a linear follower up to at least 20MHz.

In the negative resistance configuration, the amplifier is actually working as a voltage doubler, where $V_{out} = 2V_{in}$. The saturation condition is therefore $|V_{in}| > \frac{1}{2}|V_{sat}|$. Up to saturation, the device behaves as a negative resistor $-R_f$. The idealized I-V curve for the op-amp based negative feedback device with the saturation nonlinearity included is shown in Figure 3.4.

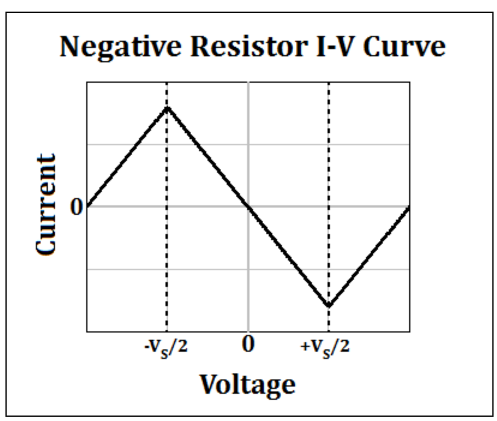

Figure 3.4: Idealized I-V curve of a negative resistor with saturation included.

Attaining an actual I-V curve like the one depicted in Figure 3.4 could potentially be trouble if the op-amp is being used is near its bandwidth limit. If gain or slew rate become an issue, the negative resistance I-V curve will be corrupted. In such a case, the I-V characteristic would likely become amplitude dependent for a given frequency. Even if the amp is operating ideally, stray impedances could play a corrupting role. To attain a good experimental model requires the inclusion of parasitic capacitances. The model shown in Figure 3.7 includes almost all of the relevant stray capacitances, for maximum generality. The theoretical expression for the impedance of the negative resistance element with the stray capacitances of Fig. 3.7a included is

$$z = \frac{-R_f}{(1 + i\omega R_f C_f)(1 + i\omega R_{nf} C_{nf}) - i\omega R_f C_{in}} \tag{3.2}$$

which has all the correct $C \to 0$ limits. It acts like a negative resistance to ground in parallel with a capacitance C_{in} to ground—which is exactly what it is. The small feedback capacitances in the NIC act to partially cancel out the parallel capacitance. So a simpler nearly equivalent model would include only the negative resistor and an effective C_{in}.

The data shows that the I-V curve observed does in fact correspond to a negative resistance parallel RC circuit, which shows that the ADA4862 is in fact operating ideally at 3MHz as a negative resistance device. Deviation from a straight line in the I-V data is due to capacitance, not imperfections of the ADA4862 chip. The stray parallel capacitance has no effect in our setup, as we have already accounted for them. The ADA4862 functions correctly as a negative resistance at 3MHz.

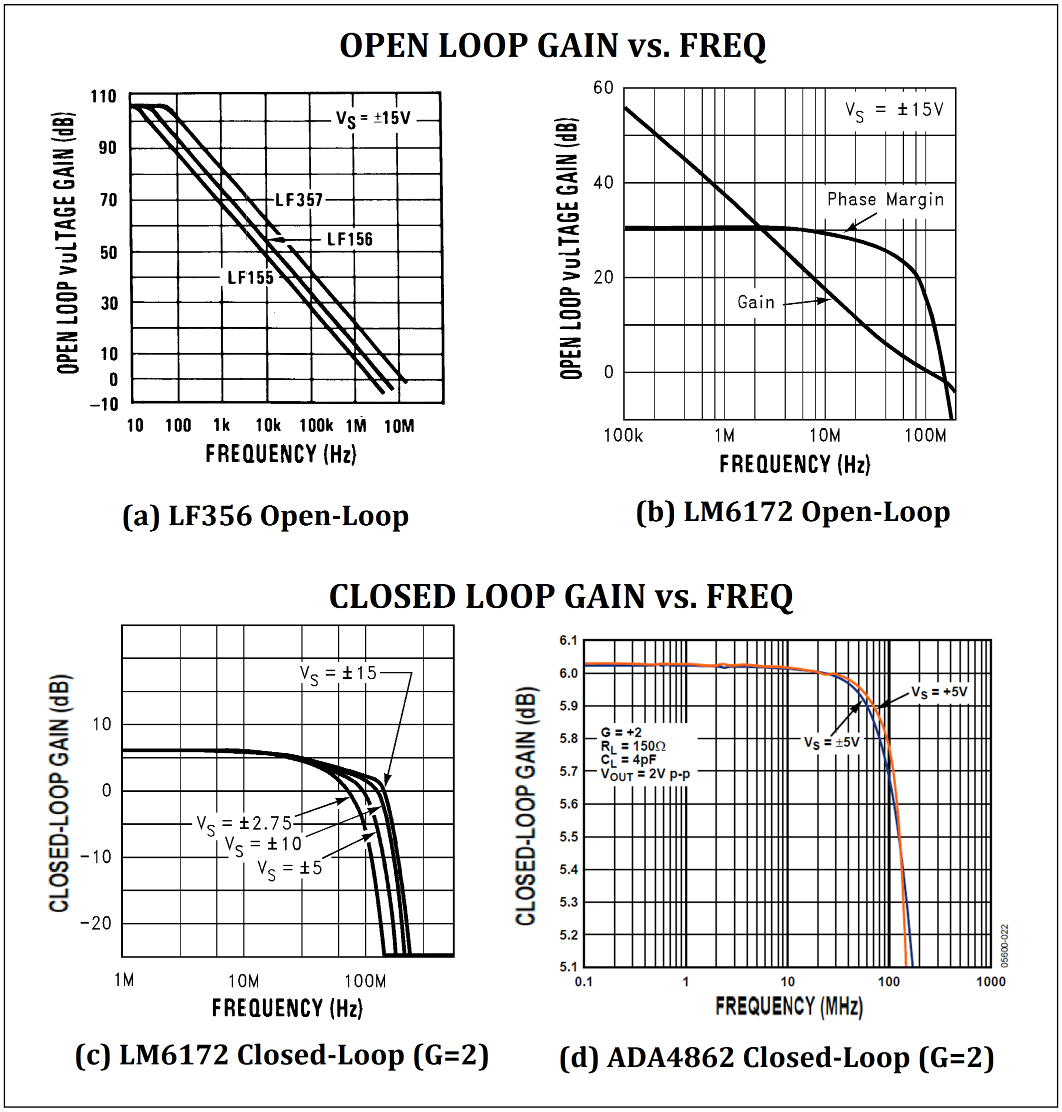

Figure 3.5: Comparison of open- and closed-loop gain ratios versus frequency for several op-amps. Each figure is taken from the manufacturer's datasheet. The top panels compare the open loop gain of (a) the LF356 Op-Amp and (b) the LM6172 Op-Amp. Bottom panels compare the closed loop gain of (c) the LM6172 Op-Amp and (d) the ADA4862 Video Buffer. Note the very different scales on the vertical axes. "G=2" refers to the closed loop gain multiplier, meaning that in the bottom panels the feedback network is configured for the amps to double the input voltage.

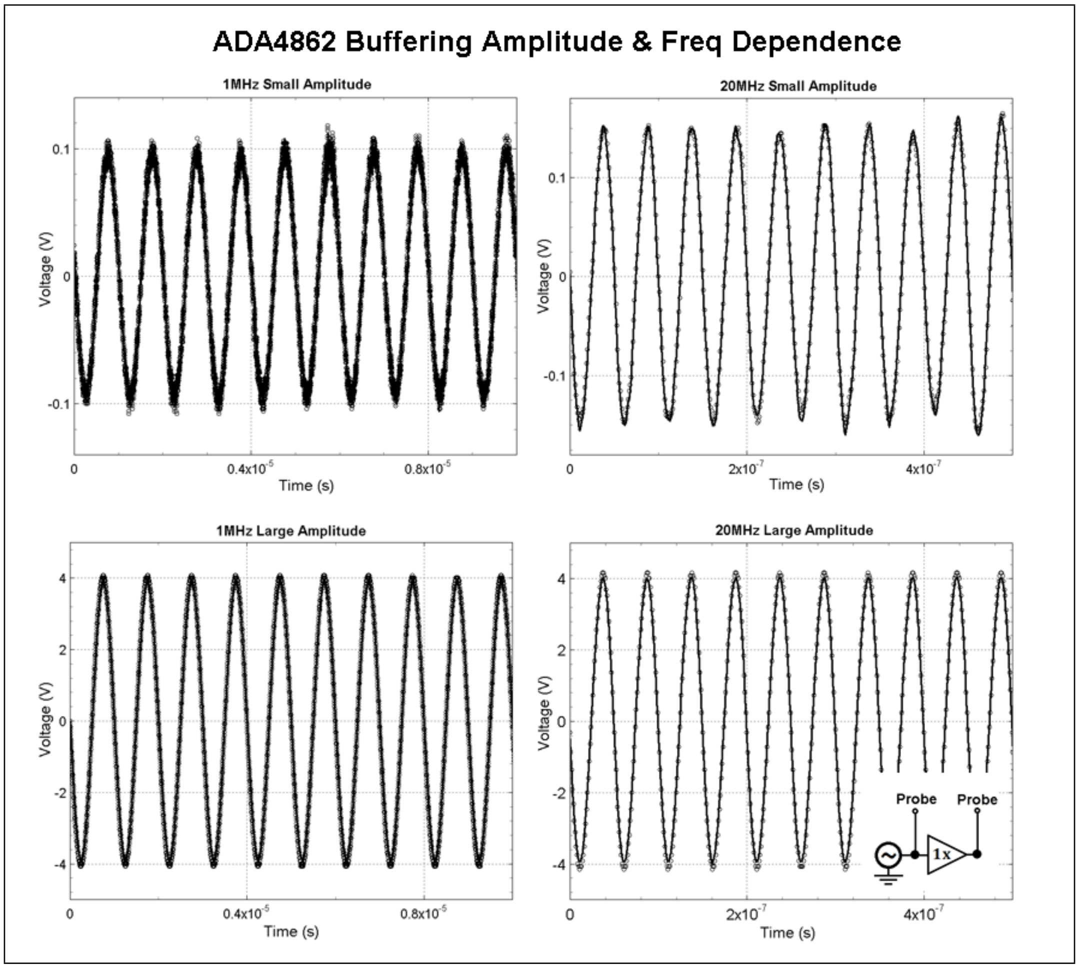

Figure 3.6: Waveform captures of the ADA4862 working as a voltage following buffer. The setup is indicated in the bottom-right inset. The input V_{in} is shown as a solid line, while the output V_{out} is shown as symbols. Responses for small amplitude (top) and large amplitude (bottom) are each shown at 1MHz (left) and 20MHz (right). The axes are voltage (volts) versus time (seconds). Notice that the ±4V swing in the large amplitude curves is pushing right up against the chip's rated voltage swing capabilities. Scope measurements were made by high impedance 10x probes. These data were taken using the PCB prototype dimer.

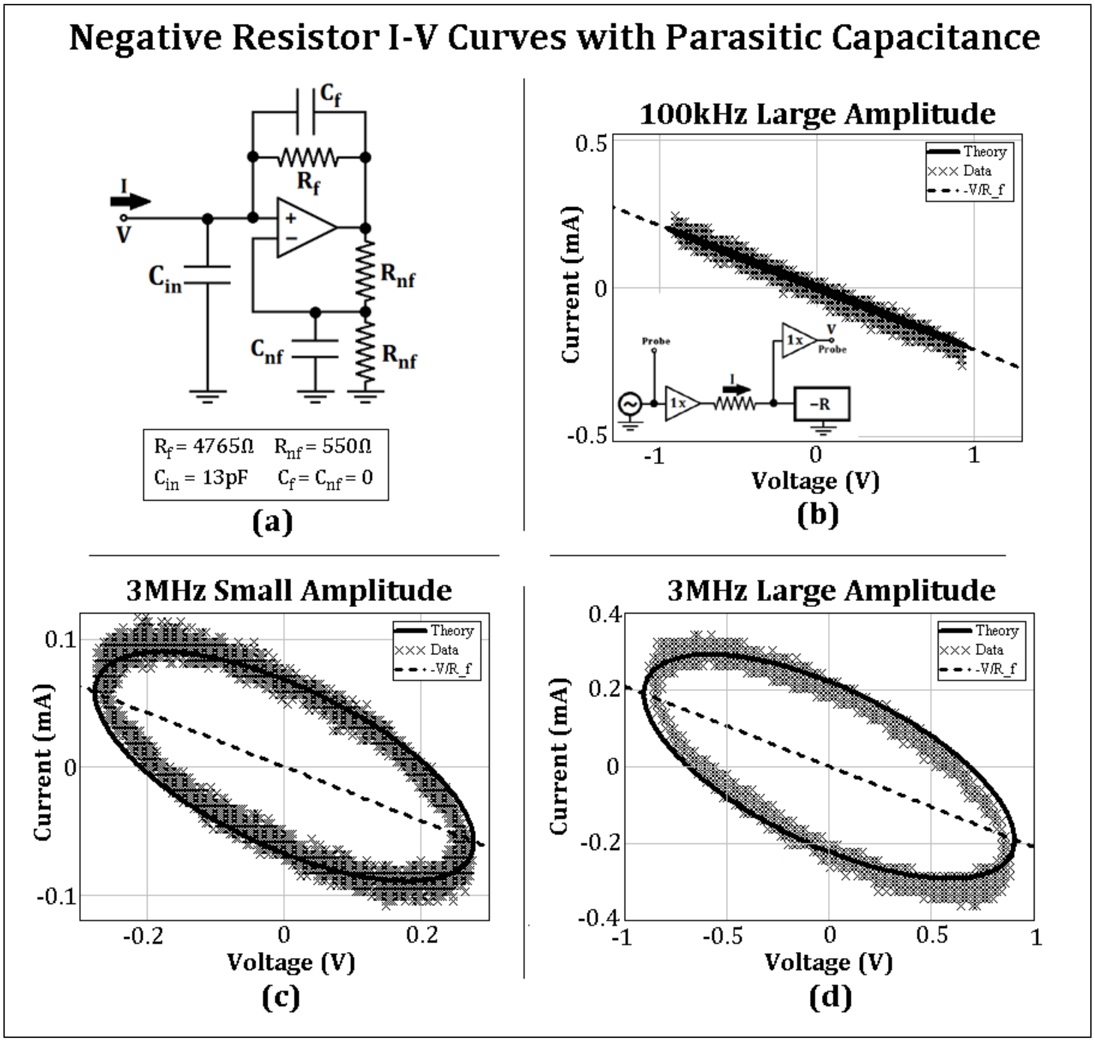

Figure 3.7: Test of the effectiveness and linearity of the ADA4862 as a negative resistor at 3MHz. The top left panel (a) shows the model of the tested negative resistance circuit. V and I depicted in the model are the voltage at and current into the negative resistance node, and correspond to the horizontal and vertical axes of (b),(c),(d). Theoretical predictions for the model are represented by solid lines, and are equivalent to equation (3.2). Experimental I-V data is represented by symbols. The I-V line for a perfectly ideal negative resistance is indicated by the dashed line $I = -\frac{V}{R_f}$. The experimental setup is shown in the inset of (b). These data were taken using the PCB prototype dimer. The experimental data fit the model well, including the amplitude linearity. This figure indicates that the ADA4862 is an effective negative impedance converter.

3.4 Apparatii

This section describes various iterations of the dimer that have been implemented in the laboratory. Each is intended to overcome shortcomings of the previous iterations. The setups are referred to by their circuit board type and the approximate natural frequency of their isolated oscillators, which is their approximate operating frequency range. Except for the Octoboard, each dimer was built so that one side, the "gain side", can be used exclusively for values $\gamma < 0$, while the "loss side" can be used exclusively for $\gamma > 0$.

3.4.1 The 30kHz Dimer

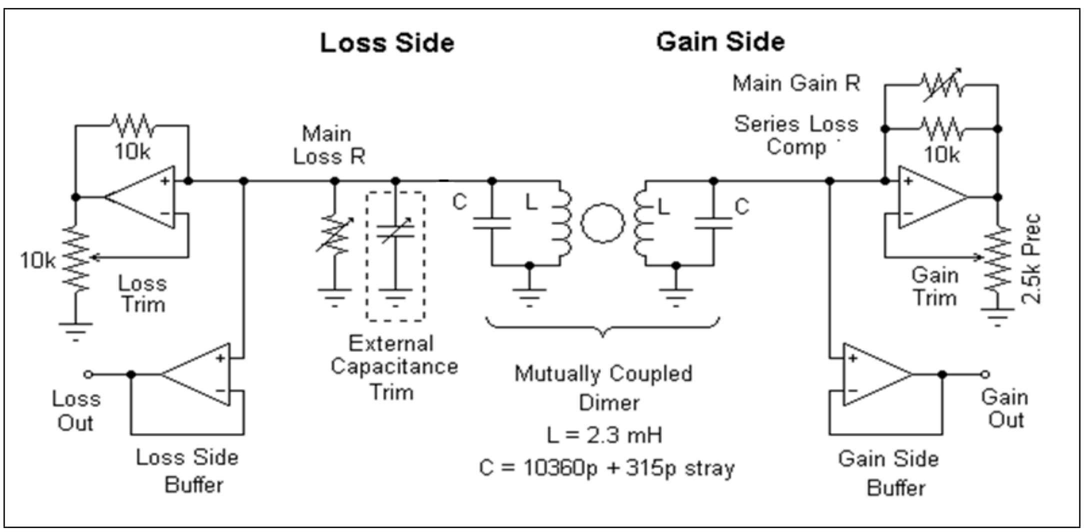

Figure 3.8: The 30kHz Dimer.

The original apparatus, equivalent to the circuit of Figure 3.8, operated at about 30kHz, and employed mutual inductance coupling as the sole coupling component. The circuit was breadboarded and mounted to a grounded external casing with regulated ±12V power supply. This apparatus was used to obtain the data of sections 2.1, 2.3, and some of 2.2.

Each inductor was wound with 75 turns of #28 copper wire on 15cm diameter cylindrical PVC forms in a 6 x 6 mm loose bundle for an inductance of $L = 2.325$mH and a coil self-capacitance of $C_{coil} = 322$pF. L and C_{self} were determined by fitting a curve to measured resonant frequencies with various parallel capacitances. The inductances were matched to within 1% by repositioning one of the turns [2, 4]. The coils were mounted coaxially with the center-to-center bundle separation used to control mutual inductance M.

During experimental trials the gain side capacitance was fixed, while the loss side capacitance was variable. There were several contributions to the gain side capacitance C_0: (1) the main oscillator capacitance provided at each side by a 10 000pF silver mica capacitor; (2) the 320pF coil self-capacitance; and (3) an additional 360pF silver mica capacitor to put the variable trim capacitance on the loss side into midrange at balance. These combine to a total gain-side capacitance of $C_0 = 10\,680$pF. The gain side therefore had nominal resonant frequency $f_0 = 1/(2\pi\sqrt{LC}) = 31.94$kHz and characteristic impedance $z_0 = \sqrt{L/C} = 467\Omega$. The measured isolated resonant frequency was $f_0 = 31.99$kHz.

The loss side capacitance differed from the gain side by (1) not including the 360pF capacitor, and (2) including a GR722-M precision condenser as the trim capacitance. The condenser's capacitance was extracted from the dial reading GR by the relations $C = (1145 - GR)$pF, and allowed capacitance control to half picofarad precision. The capacitance imbalance was defined by $\Delta C = C_{loss} - C_{gain}$. This imbalance amount was given by $\Delta C = 1145 - 360 - GR$, implying that the dimer was balanced for $GR = 785$. On balance, the nominal f_0 and z_0 are the same as for the gain side. An isolated resonant frequency of $f_0 = 31.99$kHz was measured on the loss side for $\Delta C = -267$pF. From experimental determination, both sides had $Q \approx 24$ at 31.99kHz.

Buffering and negative impedance conversion were provided by LF356 op-amps on the ±12V power rails. The linearity limit was dictated by the saturation of the voltage doubling output of the negative impedance converter. Two decoupling capacitors to ground, one 100pF and one .1μF were included adjacent to each LF356 power pin. Additional decoupling was provided on board at the power rails. Ground rows on the breadboard were wired directly to the casing.

To compensate for coil losses, the loss side included a negative resistance, adjustable via gain-control potentiometer in the negative feedback circuit, parallel to the LC combination. The coil-compensating positive feedback resistor was 10kΩ, but its effective negative resistance could range from zero to infinity by adjusting the potentiometer. With the loss side isolated and decoupled, the potentiometer was permanently adjusted to just below the threshold of oscillation, to set $\gamma = 0$ when no additional resistance is attached. After this process the potentiometer setting was near the center of its range. This setting was left alone once set.

The gain side included a twenty-turn precision potentiometer for gain-control, in the same negative feedback configuration as the loss side. Like the loss side, a 10kΩ compensating resistor was included in the positive feedback gain loop. On the gain side, additional resistances used to set γ would be placed into the same negative impedance converter in parallel with the 10kΩ resistor, a practice justified by the addition of parallel conductances. Therefore the potentiometer setting controlled all gain elements on the gain side, not just the compensating portion. The net positive feedback resistance R_f consisted of 10kΩ in parallel with an additional gain resistance chosen based on the desired γ value. A variable called G was defined as one thousand times the fractional gain potentiometer setting, and could be read directly off a register connected to the potentiometer dial, to half-unit precision. The total effective negative conductance $(R_f^*)^{-1}$ was therefore $(R_f^*)^{-1} = (R_f)^{-1}\left(\frac{G/1000}{1-G/1000}\right)$ so that $G = 1000$ corresponds to infinite gain, and $G = 0$ to no gain.

To set the factor γ in a PT symmetric way, its value on the loss side was first set by the addition of a parallel resistor to ground from the loss side node. A nominally equivalent resistance was then added to the gain side, and G was adjusted to attain very close balancing. An algorithm for "zooming in" to the gain-loss balance point by varying G and ΔC is defined by Figure 3.9. Once balance was attained, small adjustments to ΔC could be used to make either mode unstable, to allow observation.

At the balance point the gain-loss balancing was nearly exact, but the capacitances remained imbalanced by a nontrivial amount ΔC, whose value was determined by γ. Using a deft touch on the knobs, an almost arbitrarily small oscillation could be held to constant amplitude for long lengths of time in either mode, so long as γ was far from γ_{PT}.

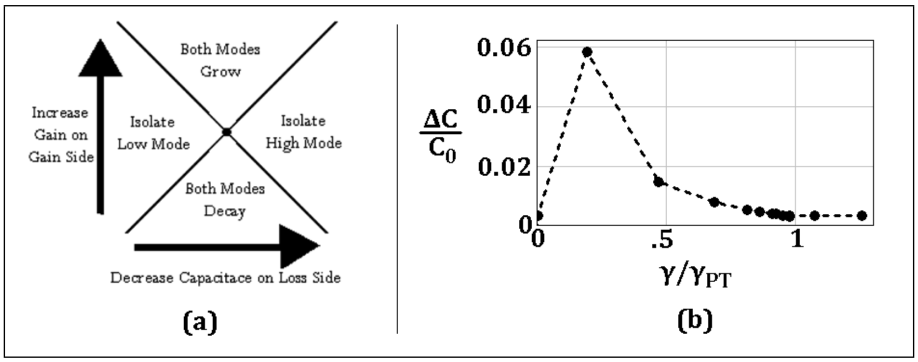

Figure 3.9: Gain-loss balancing of the 30kHz dimer. (a) The special gain-loss balancing point. The horizontal axis represents the capacitance difference $\Delta C = C^L - C^G$ between the two sides. Moving rightward on the chart indicates decreasing ΔC. Horizontal movements are controlled by the GR722M precision condenser dial. The vertical axis is the effective negative conductance $(R_f^*)^{-1}$ of the gain side, or the gain. Upward movements indicate increasing gain. Vertical movements are controlled by the gain-trimming potentiometer dial. The dot at the center of the X represents a point where the gain and loss are exactly balanced, but the capacitance is imbalanced by an amount ΔC_0. The goal is to attain a stable dimer configuration just barely below that center dot. Then the slightest dial changes can cause a marginal instability in one, the other, or both modes, allowing each mode to be observed individually in a state of gain-loss balance. Note that when viewed this way, the mode frequencies are perturbed by the presence of additional capacitance ΔC_0. An algorithm for finding this balance point by zig-zagging along one of the bottom boundary lines of the X follows. (i) Reduce gain until all modes decay. ending just inside a border of the X; (ii) Change ΔC a bit in whichever direction doesn't immediately cause instability; (iii) Increase gain until something oscillates; (iv) change ΔC enough to kill the oscillation, then a bit further; (v) Repeat (iii) and (iv) until tiny capacitance changes cause a switching from the high frequency to low frequency zone, with only a tiny "dead zone" in between. (b) The capacitance imbalance ΔC arrived at by this process, as a function of γ. Notice it approaches zero as $\gamma \to \gamma_{PT}$. Beyond γ_{PT}, ΔC was held fixed at its asymptotic value. Disregard the first point near the origin.

The 30kHz dimer was subject to several shortcomings: (1) Size. The size and clunkiness of the 2.32mH coils made duplicating the dimer impractical; (2) Low Q. The low Q-factor signifies the presence of too much coil loss, which exacerbates the issue with series-parallel resistance cancellation discussed in 3.1.3. (3) Large operating frequency range $\frac{\Delta f}{f_0}$. Since the frequency difference Δf between the low and high modes is significant compared the operating frequency f_0, frequency dependent losses in the coil may differ significantly for the two modes. Such an effect would disrupt the predicted mode structure.

To improve performance with regard to these characteristics, we built a dimer operating in the MHz range. Increasing the natural frequency would allow use of smaller inductances and capacitances, and hence smaller components, allow higher Q to be attained for the oscillators, and improve the ratio $\frac{\Delta f}{f}$. Goals included obtaining a leaner, better performing dimer, and the possibility of duplicating and chaining together multiple dimers. Coupling was switched from pure inductive to pure capacitive to allow greater control and reduce crosstalk coupling in the case of dimer chains. The result of these improvements was the 3MHz breadboarded dimer described in the following segment.

3.4.2 The 3MHz Breadboard Dimer

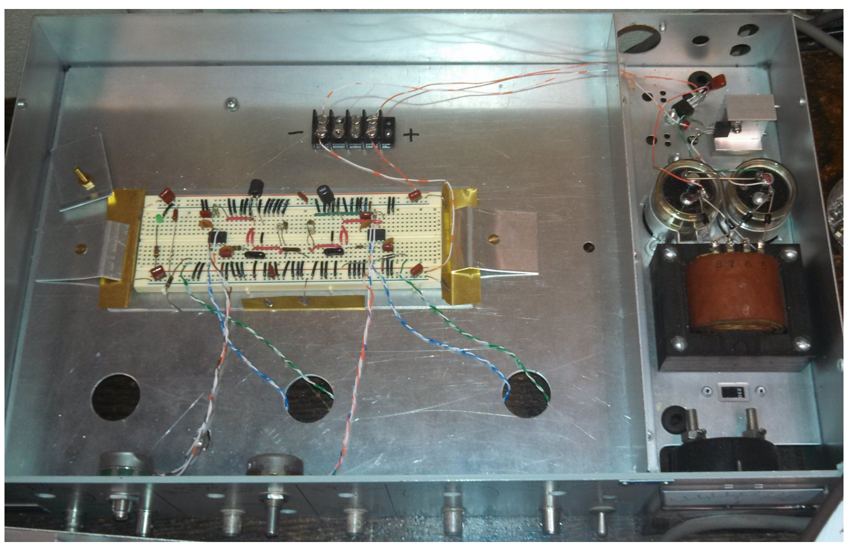

Figure 3.10: The 3MHz breadboard apparatus, mounted to casing.

This second incarnation of the apparatus operated in the range 2.5-3 MHz, and employed only capacitive coupling. It was breadboarded and mounted to a grounded external casing with regulated ±12V power supply. Breadboard ground was wired to

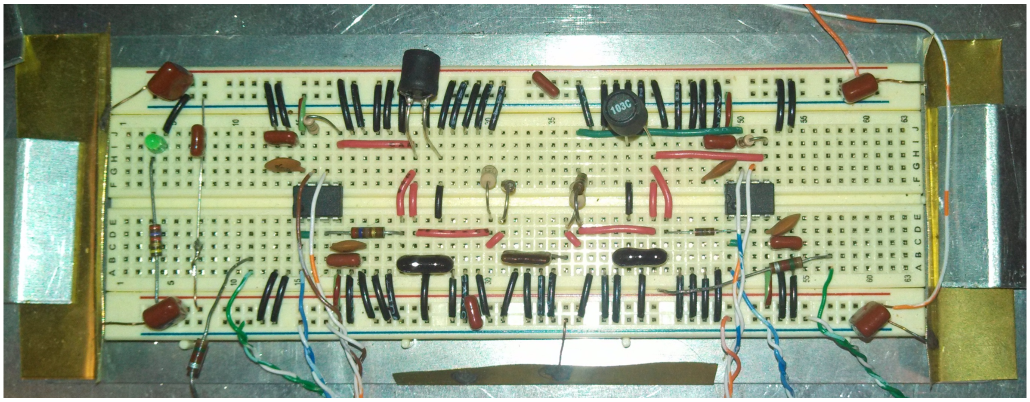

Figure 3.11: The 3MHz breadboard apparatus, close up.

the external casing. This apparatus was used to obtain the CPA data of section 2.2. If viewed from above, the left side was the loss side (denoted by superscript L) and the right side was the gain side (superscript G).

On each side a nominal inductance of 10μH was provided by a 22R103C magnetic-core radial lead inductor, while a silver mica capacitor of nominal value 330pF provided the main capacitance. The inductors were mounted with orthogonal coil configurations to eliminate unwanted mutual inductance coupling. By measuring the resonant frequency and assuming the nominal inductance values of 10μH, each side was estimated to have 25pF of additional capacitance, attributed mainly to coil self-capacitance and the 5% capacitor tolerance. Measurement with the vector impedance meter yielded $L^L = 9.88\mu$H and $L^G = 9.95\mu$H. The meter was not calibrated for absolute measurements, but gives a good idea of the matching between the two sides. The isolated loss side of the dimer was measured to have a resonant peak at $f_0^L = 2.68$MHz with $Q^L = 81.5$. The isolated gain side was observed to have $f_0^G = 2.666$MHz, with $Q^G \sim 80$. Calculations where PT symmetry was assumed used the estimated average values

$$L = 10\mu\text{H} \qquad C = 355\text{pF}$$
$$f_0 = 2.67\text{MHz} \qquad z_0 = 168\Omega$$

which are theoretically and experimentally consistent. Coupling capacitance was adjustable by simply plugging or unplugging coupling capacitors, but the experiments appearing in this work exclusively used $C_c = 56$pF.

Buffering and negative impedance conversion were provided on each side by the two channels of an LM6172 dual op-amp. The LM6172 has a higher bandwidth than the LF356, and was found to operate suitably as a negative resistance at 3MHz. However, its open loop gain decreases drastically with frequency in the MHz range, leaving 3MHz near its upper limit for safely providing negative resistance. The amps allowed a voltage swing of about ± 10V from the 12V rails. The circuit's linearity limit was dictated by saturation of the voltage doubling output of the negative impedance converters.

Each negative impedance converter was equipped with a gain-adjusting potentiometer. Coil loss was compensated on each side by placing a 15kΩ resistor into the negative impedance converter and adjusting each isolated side to just below the threshold of oscillation with the potentiometer. This left both potentiometers near the center of their range, which is consistent with $R = Qz_0$. After this point the potentiometers were not adjusted. When used, additional gain resistors were placed into the negative impedance converters parallel to the 15kΩ, meaning the gain potentiometer setting affected both the compensating gain and the main gain resistors. During the scattering experiments performed with this dimer, gain-loss balance was not carefully trimmed, but was checked by comparison to simulation.

This apparatus experienced several significant problems, all of which were ultimately eliminated. The first was discovered immediately upon the circuit's assembly, in the form of an undesired 120MHz oscillation. The oscillation was caused by some stray phase-shifted feedback loop associated with the LM6172. Such oscillations are a common ailment in high speed op amp applications. The problem was easily cured by placing a 270Ω resistor in series with each NIC input. The series resistance, combined with the LM6172's 2.5pF input capacitance to ground, formed a low-pass filter of $f_{3dB} \sim$ 200MHz, which was low enough to kill the unwanted oscillation but high enough not to affect our 3MHz signal. Subsequent apparatii employed a similar low-pass filtering resistance at the amplifier inputs as a precaution.

The other major problem was identified when experimental scattering data was found to be anomalous. Investigation revealed large oscillations occuring along the breadboard ground and power circuits when the dimer was driven by the signal generator. Connection of a capacitance from oscillating power or ground points to the external casing ground noticeably suppressed the oscillations, and was dependent on location. Evidently, some resonance of the breadboard wiring was within the operating frequency range. To suppress the oscillations, a more robust connection was made between the external casing ground and the breadboard ground rails, and power decoupling capacitances were added copiously. Rather than at one point, the breadboard ground rail was ground to the casing at five points distributed evenly across the ground rails. Small ground loops were deemed necessary to suppress the resonances. These modifications successfully stabilized the power and ground circuitry, allowing subsequent experimental trials to be fruitful.

Once its kinks were worked out, the 3MHz breadboarded dimer functioned well. The guiding principles for development of the apparatus beyond this point were to (1) develop an apparatus that could be used to chain together multiple dimers; (2) allow automated computer control of the gain-loss parameter; and (3) avoid the pitfalls of breadboarding by utilizing a printed circuit board. The octoboard represented the endgame for these advancements, but first a prototype PCB dimer was needed to ensure the Octoboard could be successful.

3.4.3 The 3MHz PCB Prototype Dimer

This dimer incarnation was designed as a prototype for expansion into the realm of printed circuit board design. The PCB prototype deviates from the preceding apparatii by acheiving automated gain control via vactrols, eliminating gain trimming potentiometers, and switching to ±5V supply rails. The board was manufactured by Pad2Pad[1] and

[1] *Pad2Pad Custom PCBs.* Mahwah, NJ. <www.pad2pad.com>.

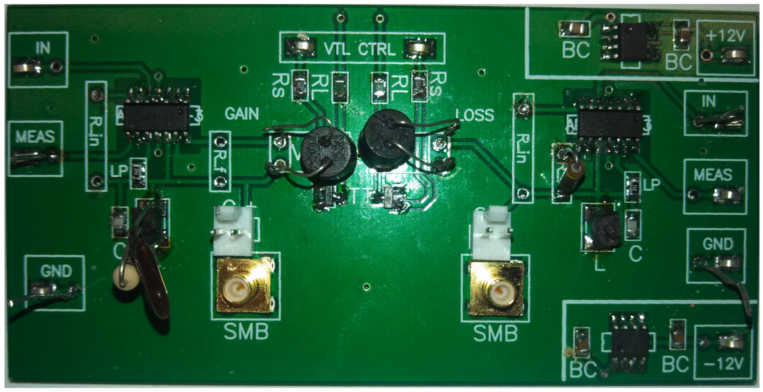

Figure 3.12: The 3MHz PCB prototype apparatus.

designed on PC using Pad2Pad's freely downloadable circuit design software. Gain and loss sides are asymmetric, and are labeled on the top silkscreen. The reentrant stability data of section 2.4 was captured with this dimer.

The two-layer board is a 2″ x 4″ rectangle of 0.062″ thick FR4 dielectric ($\epsilon_r = 4.8$). Traces and ground plane have copper weight 1oz/ft^2. Traces are of varying width. The ground plane is etched to a gap width of 0.020″ around traces. All unused top-side and bottom-side board space is covered in a copper ground plane, giving traces the "coplanar waveguide with ground" configuration, which can be used to calculate trace impedances. Thermal connections are used at grounded soldering pads. Components were ordered seperately, and the board was assembled by hand in the lab. Standard 0805 SMDs were used wherever possible, but inductors, vactrols, integrated circuits, and various other bits have different footprints.

On-board ±5V voltage regulators are powered by and grounded to an external ±12V supply. Inductance was supplied on each side by a 10μH API Delevan S1812R-103K SMT inductor, and capacitance by a 330pF SMT MLCC.

The inductors were carefully matched by attaching a small magnetic shard to the lesser with beeswax. The capacitances were then estimated by assuming one of the inductors to have the nominal value and inspecting the isolated oscillator resonances.

The resultant estimation of system parameters is:

$$\begin{aligned} L^G &= 10\mu\text{H} & L^L &= 10.015 \pm 0.05\mu\text{H} \\ C^G &= 355\text{pF} & C^L &= 368\text{pF} \\ f_0^G &= 2.67\text{MHz} & f_0^L &= 2.62\text{MHz} \\ Q^G &\sim 60 & Q^L &\sim 60 \end{aligned}$$

Buffering and negative impedance conversion were provided by the ADA4862-3 triple video buffer which was well within its bandwidth and operated ideally. A negative resitance on the gain side and a resistance on the loss side were each controllable via vactrol. Additional parallel gain or loss resistances were added on either side to offset the vactrol ranges.

Vactrol LED current was controlled through a transistor by an externally supplied voltage. The npn transistor was wired with its emitter to ground, and the vactrol LED on its collector along with a 47Ω current limiting resistor. External voltage was applied through a 220kΩ resistor to the transistor's base to control the LED current. The vactrols were typically kept at a preset resistance of ∼5kΩ during proper use.

This board was primarily used to investigate vactrol properties and test vactrol calibration procedures. The detailed results are presented in Chapter 4. To summarize: despite difficulties getting them to behave, it was confirmed that vactrol gain could be calibrated and then repeatably implemented to some level of precision. After the vactrols were calibrated, the dimer confirmed that self-oscillation of the isolated oscillators could be observed near the theoretically predicted resistance $R_{vtl} = Qz_0$. Qualitatively correct properties of the stability maps of subsection 1.6.1 and of the PT mode coalescence were also confirmed. Quantitative vactrol performance tests are discussed later. The protype was also used to take data on the ADA4862 video buffer performance.

Once satisfied with the vactrol's capabilities, and having learned many lessons from the prototype, the Octoboard project was undertaken.

3.4.4 The Octoboard

The octoboard is a printed circuit board with eight identical oscillators, each with independent computer control over γ, and additional circuitry for signal input and output. It is the subject of Chapter 5.

3.5 Summary

This chapter has taken our experiment out of the conceptual realm and into the physical world. Performing the experiment requires ways to control and measure the experimental setup, which can all become quite complicated. It's essential to make certain that everything is doing what we expect it to, and leave nothing up to assumptions.

Chapter 4

The Vactrol

To explore the γ parameter space of our dimer requires controlling resistances on either side. For the purpose of digitally controllable variable resistance, we turn to a device known as the "vactrol". The vactrol consists of an LED enclosed with a light-sensitive photoresistor. The LED light shining onto the photocell controls its resistance (Fig. 4.1).

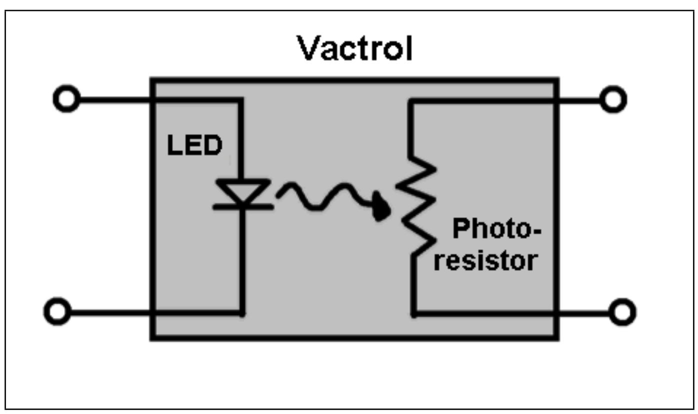

Figure 4.1: A vactrol consists of a light-emitting diode enclosed with a light-sensitive photoresistor. The light emitted by the LED controls the resistance of the photoresistor, allowing a vactrol's resistance to be controlled by a current with no coupling between the control circuit and resistor circuit.

The defining characterization of a vactrol is its photocell conductance versus LED current. These should have an approximately linear relationship, since LED light intensity is proportional to current, and photocell conductance is related to the number of carriers freed by photonic excitation, and thus to the intensity of light to which it is exposed. Photocells are resistive in the dark, and conductive in the light. So the photocell end of a vactrol conducts well when a lot of current passes through the LED, and doesn't conduct at all when no current passes throught the LED. Figure 4.2 shows photocell conductance g_{pc} versus LED current I_{LED} for the commercially available Silonex NSL-32 vactrol.

4.0.1 Why Vactrols?

Vactrols are a good choice for our application, despite a few drawbacks. Our requirements dictate the use of a variable resistance with (1) linear, frequency-independent resistance, (2) low capacitance, and (3) the possibility for non-intrusive computerized control of resistance. Other variable resistance options couldn't fullfill these needs. A traditional mechanical potentiometer, for example, both lacks computer control and has far too much stray capacitance at 3MHz. Vactrols offer typical capacitances on the order of only several picofarad, allow for resistance control with a DC signal that couples minimally to the signal side of the circuit, and linear resistance.

4.0.2 Options and Availability

Unfortunately vactrols are somewhat obsolete, having been replaced almost entirely by phototransistor optocouplers in engineering applications. At this point, electronic music hobbyists are the primary contemporary vactrol users. Only several manufacturers have available models, and, not surprisingly given the nature of their usual applications, it seems that not much effort has gone into their miniaturization or perfecting their performance, either of which we'd have appreciated. Commercial models are designed for relatively large currents (1-40mA for the NSL-32). Nonetheless, the available options have proved sufficient for our purposes so far.

The commercial options known to us include a line manufactured by Silonex, and one

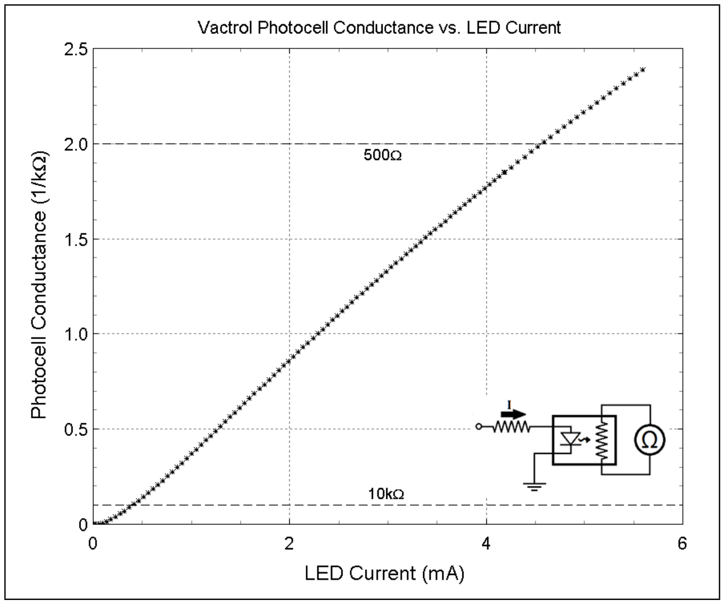

Figure 4.2: Vactrol NSL-32 photocell conductance versus LED current. Test setup shown in inset. The vactrol was aged and preset to near $g = (5\text{k}\Omega)^{-1}$ preceding data acquisition. When the vactrol is "off", with little or no current passing through the LED, conductance is negligible. Beyond a threshold of about $g = (10\text{k}\Omega)^{-1}$, conductance goes nearly linearly with current.

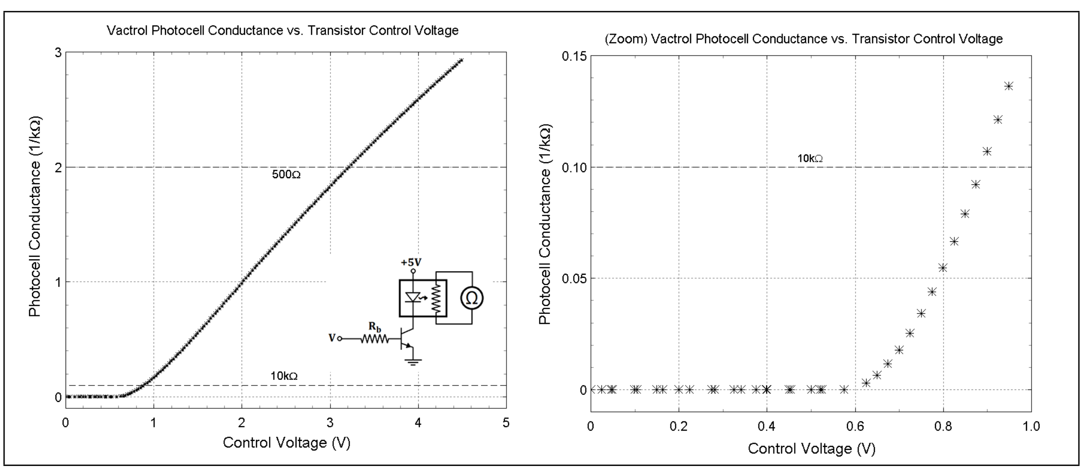

Figure 4.3: NSL-32 photocell conductance versus transistor control voltage. Similar to Figure 4.2 but with the current drawn via transistor rather than directly. This data was taken on the PCB test dimer ($R_b = 220\text{k}\Omega$); the setup is depicted in the inset. The vactrol was aged and preset to near $g = (5\text{k}\Omega)^{-1}$ preceding data acquisition. While $V < V_{be}$ of the transistor, the transistor is off and no current passes, causing the initial zero conductance region. Once the transistor is on, the curve looks like Fig. 4.2. A magnification of the turn-on region is shown on the right.

from Perkin-Elmer. The Silonex models are smaller, sleeker, and cheaper, so, choosing the model based on our desired resistance range, we are using the Silonex NSL-32 vactrol. It is also easy to build vactrols in the lab that are marginally superior to the commercial models. A nice model was made by facing a small high-efficiency LED against the face of a photocell, with a small clear plastic layer seperating them, and enclosing the unit in a bubble of black epoxy. This model was slightly smaller and less capacitive than either commercial model, performed comparably, and operated at a much lower current. However, in the interest of repeatability and student sanity, we decided to stick with the commercial model. Not to mention that the homemade models were prone to losing legs.

4.1 Characterizing the NSL-32 Vactrol

Our goal, with the vactrol, is to attain a system where a conductance value γ is entered into the computer, which then adjusts an analog voltage associated with the vactrol LED, causing the vactrol photocell to have conductance γ. Ideally, this would all happen instantaneously, and the photocell conductance would be exactly γ. Obviously, such a thing is not possible. The ability to approximate that ideal system is limited by the performance of the vactrol. As we'll see in this section, the limitation is pretty severe.

The tactics for attaining and assessing such a system in practice go as follows:

(1) Set up a control interface where some computer-controlled parameter dictates the LED current. For example, a voltage-output DAC in series with the LED and a resistor to ground.

(2) Perform a "calibration" experiment, which associates each setting of the control parameter with a measured value of photocell conductance.

(3) Fit the calibration data to a function which gives the control parameter value as a function of measured conductance. This is a "calibration" function, $V(\gamma)$.

(4) Additionally attain the inverse function, directly, by fitting the data in the reverse direction, attaining the function $\gamma(V)$. Don't try to invert the calibration function. Good matching between the forward and inverse fit will be a sign of good fitting.

(5) Test the calibration by seeking desired conductances according to the calibration function, and measuring to see how close to the desired value is attained.

(6) Determine the accuracy, precision, and repeatability conditions of the calibration, then proceed to use the calibration with its limitations in mind.

Before tackling a calibration, we need to know generally how we expect the vactrol to behave. In short: it behaves badly, unless used under very specific conditions. Its usefulness is significantly hampered by (1) slow response of the photocell to changes in light intensity, (2) strong hysteresis effects, and (3) long-term temporal drift.

These issues are tricky, because besides making implementation more difficult, it made characterizing the vactrol's behavior quite challenging, since it seemed to be changing all the time. The issues were initially identified when the first round of calibrations lacked repeatablity within the desired precision. After much experimentation, we were able to identify conditions under which suitably repeatable calibrations could be attained. At that point we didn't attempt to exactly quantify the messy hysteretic behavior, in favor of forging forward with the implementation. Our understanding of the vactrol's intricacies is therefore defined by a set of qualitative rules about the vactrol's behavior and a set of instructions allowing the vactrol to be used for our purposes.

To elucidate the complexity of this issue, consider the following. If we want to do something as simple as record the resistance versus LED current, we have many options. How long should we wait to make a measurement after changing the control voltage? First of all, it strongly depends on what size jumps we are trying to make. If we try to increase resistance by 5kΩ and wait 10s to measure, we'll get a very different result than if we wait 30min. After 30min we'd be pretty sure it's settled down close to some seemingly asymptotic value, and if we kept repeating that 30min experiment, we'd get pretty much repeatable results. After only 10s, the resistance would be far from its asymptotic value, and would certainly be changing. But if we kept repeating that 10s experiment we still might be able to get repeatable results, by picking off the curve at the right point every time. And if we used the 30min method, besides it being obviously impractical, our calibration would get "lost" by instating such a large change for such a long time, and we wouldn't be able to get back to our zero value. So doing something like the 10s method is actually better. If we go to smaller resistance steps, then something on the order of 10s will be improved in it's performance. With proper practices and small resistance steps we can get settling times to be quite fast. But that doesn't necessarily mean going to ever smaller jumps would be good. At some point the photocell would be averaging over tiny jumps and thus changing contiunously. The dynamics of that, we have no idea about.

At this point the reader is probably frustrated by the lack of quantitative statements involved in this discussion. Unfortunately, quantification depends on many factors. When was the last time the vactrol was off? What conductance was it held at in the 24 hours previous to the experiment? How old is it? Add this to the fact that every vactrol is different, and we have a bona fide nightmare.

The important point here is that all of these horrible issues can be ignored if we devise a scheme that yields useful calibrations. The best we can do with this info is use it to inform our practices. The list at the end of this section summarizes the rules and instructions for best practice that have been extracted based on our knowledge of the experiments. The captions of Figures 4.4 and 4.5 explain some experimental tests in detail, partly revealing how we've come to the conclusions stated in this section. The following section 4.2 demonstrates that successful calibrations were obtained with these practices.

Figure 4.2 shows a typical current-conductance curve for an NSL-32 under proper practices. As predicted, the relationship is about linear. There is a "dead region" at low currents, where conductance is nearly negligible. The curve becomes linear in the vicinity of $g = 0.1\text{k}\Omega^{-1}$. Note that the dead region includes several points at the origin, where the input voltage into the node shown in the inset is smaller than the LED voltage drop, meaning no current flows.

Figure 4.3 is very similar to 4.2, except for being taken with a different setup. As seen in the inset, this figure utilizes a transistor driven current. In this case the dead zone is related to the transistor's base-emitter voltage drop. Besides the different dead zones and horizontal axes, the two are basically equivalent. The right-hand panel shows a zoom of low current, low conductance response. Note that the linear region starts around V=0.9V, R=10kΩ. This fact was used to choose octoboard DAC-biasing parameters (section 5.1.4).

Figure 4.4 demonstrates the vactrol's long-term temporal drift properties, showing that over the course of a week drifting is constant but decreasing in rate, appearing to approach equilibrium. For this reason we believe new vactrols should be "aged" or broken in before calibration and use. It also demonstrates the vactrol's hysteretic thermal response, which is not particularly relevant to us.

Figure 4.5 demonstrates an example test of the vactrol's temporal response to changes in LED current settings. Two trials are shown; one traverses a large resistance range, and the other a smaller range. Many other similar trials were taken.The control current was flip-flopped back and forth between two values over half-hour intervals. The data show that the vactrol overshoots attempted changes, and that larger resistance changes cause worse overshooting. Moreover, settling times to asymptotic values are seen to exceed 30min. Thermal fluctuations at room temperature are also visible. Figs. 4.4 and 4.5 precede the proper practice procedures.

The following list summarizes the principles, rules, and tips for successfully implementing and understanding the vactrol.

Vactrol Proper Use: Rules and Tips

(1) The vactrol's resistance is always changing. Getting a good calibration is about knowing how the timescale of the experiment matches up with the vactrol drift timescale.

(2) New vactrols will be subject to long-term drifting over days-long timescales. This must be eradicated by "aging" the vactrol before using it or calibrating. Each vactrol should be set to a high conductance ($g > 1\text{k}\Omega^{-1}$) for at least a week, allowing the long term drift to at least partially equilibrate.

(3) Hysteresis effects are important. Therefore each vactrol must always be kept in a standard state when not performing a specific function, so that hysteresis effects will be be identical between trials. Specify a "preset" level near the middle of the intended conductance range. Keep the vactrol at this setting at all times when not in use. Err on the side of higher conductance.

(4) Avoid ever having the vactrol off, or using very small conductances. Putting it into a state of high resistance ($R \gtrsim 10\text{k}\Omega$) causes much worse hysteresis than higher conductance settings, and can take a long time to recover from.

(5) Settling time to some apparently asymptotic value is an increasing function of step size. Big steps on the order of ΔR=10kΩ may take hours to equilibrate. Small steps on the order of ΔR=100Ω may take less than a second. Fast and slow effects may compound each other.

(6) The vactrol will overshoot attempted conductance changes. Bigger changes to the conductance will result in worse overshooting of the desired conductance, and a greater initial rate of change of conductance following the change. (Overshoot meaning it will *change too much*, not *be too large*.)

(7) Control temperature if possible, thermal drifts are noticeable at room temperature.

(8) Calibrate directly between the conductance and the control parameter. Use a calibration routine that mimicks the intended use pattern as closely as possible.

(9) Once calibrated, perform calibration checks regularly, including checks of the temporal response to attempted changes. Calibrating and then checking gives a more precise way to figure out what the vactrol is doing than just fumbling around with parameter settings.

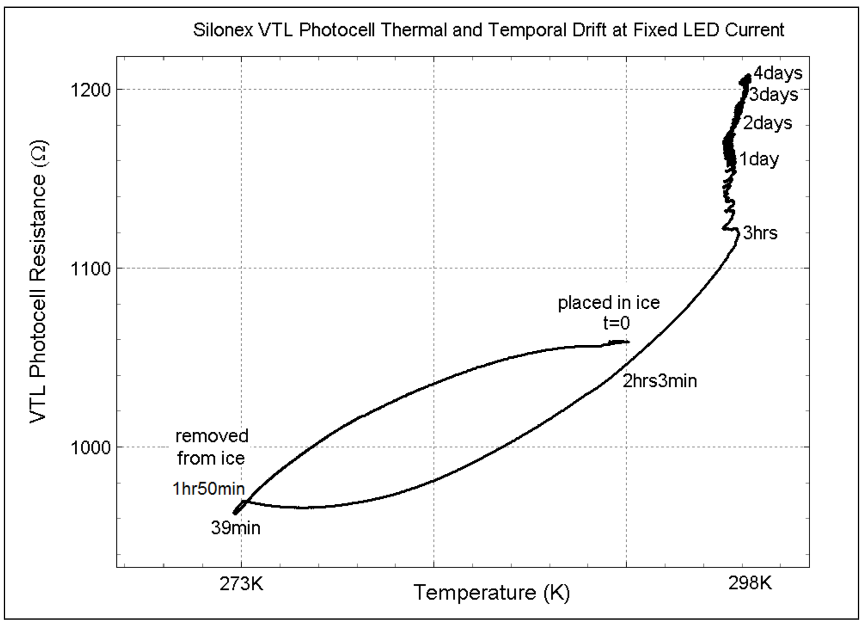

Figure 4.4: A long-term capture of NSL-32 photocell resistance versus temperature, at a (nominally) fixed LED current. This data was taken with the transistor setup on the PCB test dimer. Thermal and temporal changes could have affected both the vactrol and the transistor. Relative temperature was measured by an uncalibrated thermistor taped onto the PCB next to the vactrol, and the horizontal scale was shifted to match the known temperature of ice. The setup was initially placed into an ice bath and data acquisition was started. Equilibrium with the ice was reached after 40min. The setup was removed from ice after 1hr50min. Return to equilibrium with room temperature was reached after a total time of 3hrs. The setup was then left at room temperature for several days of data acquisition. These data demonstrate both thermal hysteresis and long-term aging effects. This data was taken with a fairly new vactrol, which had not seen many hours of use prior to the trial. The decreasing rate of temporal drift seen over the period of several days indicates aging effects which slowly equilibrate. Due to these and other data we believe the vactrols should be intentionally "aged" before calibration and use, allowing some of the temporal effects to be minimized. The hysteresis over large thermal changes is not especially relevant to our purposes.

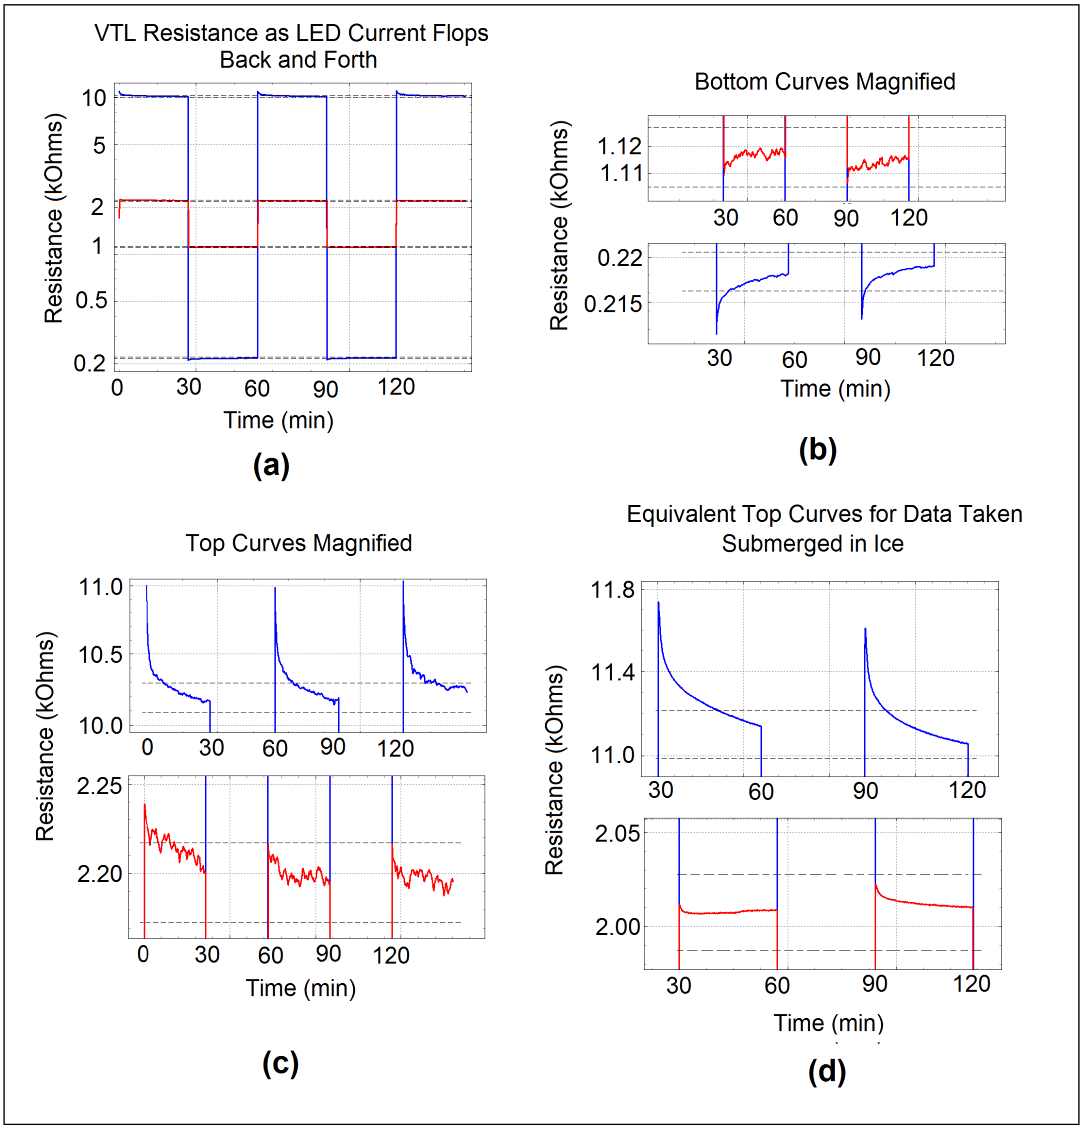

Figure 4.5: NSL-32 resistance versus time as LED current is switched back and forth over thirty minute intervals for large (blue) and small (red) changes in current. Vertical axis is logarithmic in (a), linear in the others. Pairs of dashed black lines, showing $\pm 1\%$ markers relative to an approximate average final resistance value for each current setting, provide a sense of scale. A magnified view the lower halves of each curve of (a) is shown in (b). The upper halves are magnified in (c). The data of (a), (b), (c) were taken at room temperature. The same experiment was performed with the apparatus submerged in ice; a sample of that data is shown in (d). This data was taken on the PCB test dimer. These data highlight many important features of the NSL-32 vactrol behavior. (1) They tend to overshoot attempted changes in either direction, then come back toward the equilibrium value. (2) The magnitude of overshooting, and the time taken to come within 1% of the end value, depend on the size of the attempted change in resistance. (3) Even though for the large changes in R, it takes a long time for R to equilibrate, it follows the same path each time, which can still be useful as long as time scales are paid attention to. (4) The fluctuations along the curves at room temperature are due to thermal drifting, since the temperature-controlled ice data do not show the same fluctuations.

4.2 Calibrating

Despite all the difficulties, we were able to obtain some decent vactrol calibrations. Figure 4.6 depicts a test of a resistance versus control voltage calibration attained on the PCB test dimer gain side. The control voltage in this case is the voltage into the 220kΩ resistor at the base of the transistor. The calibration was obtained by fitting resistance data to a Laurent series. This method is not recommended, it is better to fit the conductance with a polynomial. Nonetheless, useful results were obtained with this calibration. As the calcheck figure shows, measured resistance values matched the desired values to within 1% when checked on the same day. To within fairly lax precision standards, the calibration held at least over the course of two weeks. Over that time period, the calibration, and one on the loss side, allowed the stability reentry data of Section 2.4, along with some PT mode coalescence data, to be captured. The fact that those computer-controlled experiments matched well with theoretical predictions serves as a major vindication of our continued use of the vactrol.

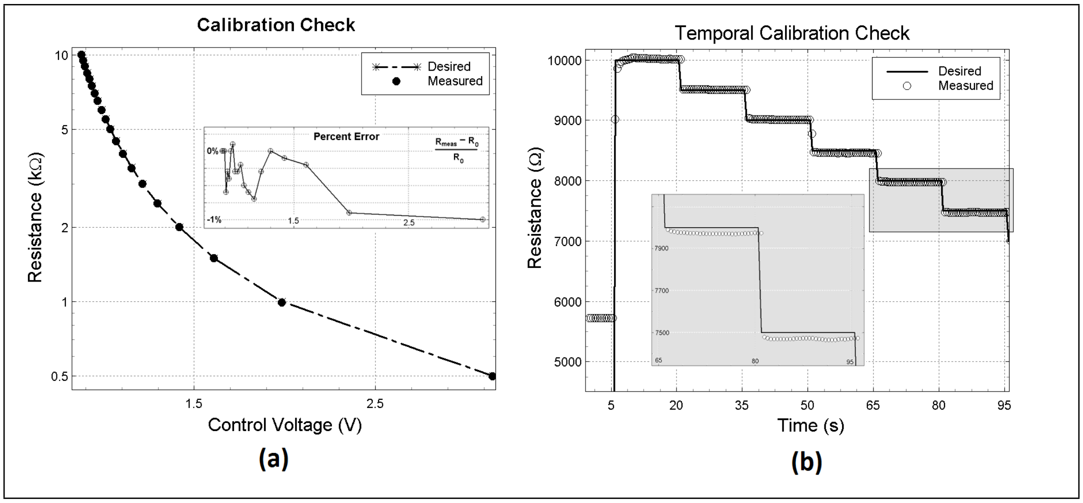

Figure 4.6: Checking the calibration of the PCB test dimer's gain side vactrol. In (a), dark circles represent measured resistance data points. The (mostly covered up) stars represent the desired resistances predicted by the calibration, and are connected by dashed lines. The inset shows percent error from the desired value, which was within 1%. On the right, (b) shows temporal data for the same test run as (a). Their relation is that the points shown in (a) are the last data points at each setting in (b) (points in (a) are at the edge of the cliff in (b)). Prior to both the calibration and the test, the vactrol was held at a preset resistance R=5.7kΩ. The original calibration data, which was fitted to provide the calibrating function, was taken in a similar data density to panel (a), and allowed a 30s settling time between steps. This calibration check allowed shorter 15s settling times, with steps ΔR=500Ω. The same calibration, as well as its associated inverse, were used regularly over the course of two weeks, and continued to yield sensible results.

4.3 Summary

The vactrol allows computerized control of the gain-loss parameter γ. We've shown that despite major issues, the vactrol can be used effectively for this purpose within certain precision and accuracy limitations. In the upcoming chapter, the octoboard exploits eight vactrols to independently control the γ of eight oscillators.

Chapter 5

The Octoboard

The Octoboard is a PCB apparatus designed for PT electronics experiments of increased complexity. The name "octoboard" comes from its eight LRC oscillator units. Using the octoboard, and even larger-scale units derived from it, we intend to conduct experiments involving multiple dimer chains, potentially in complex topological arrangements or higher connective dimensionality. Scattering effects, especially, should be very interesting for such systems. The octoboard allows maximal programmability and automation for a variety of experiments. It includes individual computerized control of each oscillator's γ via vactrol, and input and coupling connectivity are easily customizeable. Computerized vactrol control is interfaced through an external microcontroller.

The full technical specifications for the octoboard are included as Appendix A. Appendix materials will not be redundantly reproduced in this chapter, and the appendix will be referenced implicitly. Consult the appendix for circuit details. The text will focus on conceptual principles and instructions for use.

5.1 Theory of Operation

The octoboard circuitry includes:

- Eight identical LRC oscillators, each with γ adjustable for either gain or loss.

- An input port for digital signals from an external AXON microcontroller. The AXON receives USB communication from the PC, and manipulates ten digital pins on the octoboard to control vactrol levels and output selection.

- A vactrol γ-control interface. Eight DACs are controlled by serial communication with the AXON. Each DAC setting controls a vactrol LED current via transistor.

- Two multiplexed output ports intended for the oscilloscope. Each output reports the voltage at one oscillator, as dictated by AXON settings.

- Two ports for external signal input.

- On-board ±5V power regulation.

Figure 5.1 illustrates the Octoboard's computer control interface system.

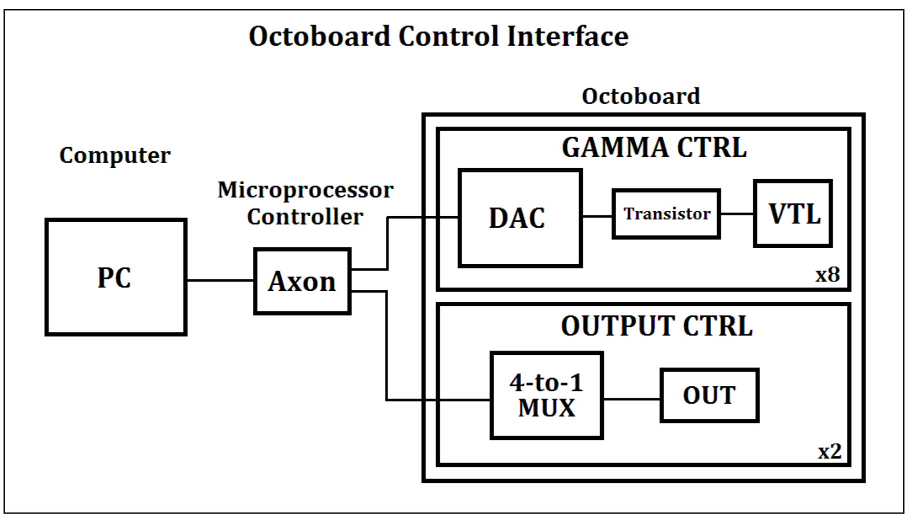

Figure 5.1: Block diagram of the Octoboard control interface structure.

5.1.1 General Use

The board is intended to be customized by coupling circuit blocks together by twisted pair at two-pin jumper headings. Twisted pairs with impedances soldered in should be used to couple various parts of the circuit together. Necessary jumper headings are provided.

For ease of testing and troubleshooting, almost all parts of the circuitry can be isolated from each other. There are copious jumper headings, most of which will be jumpered (shorted) during normal operation. Power rails include solder jumpers, which can be disconnected if necessary.

5.1.2 Oscillators

The octoboard includes eight identical oscillators. Each can be used for either gain or loss depending on the vactrol setting. Nominally, each has $L = 10\mu\text{H}$, $C = 330\text{pF}$, thus $f_0 = 2.77\text{MHz}$, and $z_0 = 174\Omega$. The LC can be isolated to allow measurement of f_0 and Q, and the vactrol photocell can be isolated to allow resistance measurement.

At each oscillator node, a triple video buffer chip provides both a buffered voltage measurement and a negative impedance converter. The third of the chip's buffers is unused and powered down. A resistor at the buffer's input, in conjunction with the input capacitance, provides low-pass filtering to suppress potential high-frequency self-oscillation. The filter doesn't affect the 3MHz signal.

The vactrol provides negative resistance, so increased vactrol conductance corresponds to increased gain. We chose to make the vactrol resistance negative so the circuit would be stable at the minimum DAC setting. The vactrol's conductance range is approximately $0.1\text{k}\Omega^{-1} < g < 3.3\text{k}\Omega^{-1}$. A 620Ω loss-biasing resistor ($g = 1.6\text{k}\Omega^{-1}$) is included to offset gain from the vactrol, so that $\gamma = 0$ will lie near the middle of the vactrol range. The vactrol range was chosen so that, for a typical dimer coupling strength, the extreme values are near $\gamma = \pm 1.5\gamma_{PT}$. If the range proves insufficient, empty soldering pads are included for both positive and negative resistances, which can be used to offset the γ range by any amount. The process behind choosing and implementing the desired vactrol range is discussed in detail in its own subsection, 5.1.4.

Once the vactrol is calibrated, γ_{vactrol} can be set by computer. The oscillator's actual

γ then has several contributions:

$$\gamma = -\gamma_{\text{vactrol}} + \gamma_{\text{bias}} + \gamma_{\text{LC}} + \gamma_{\text{offset}} \tag{5.1}$$

where γ_{bias} comes from the 620Ω biasing resistor, γ_{LC} is the equivalent parallel resistance due to inherent LC losses, and γ_{offset} is the optional offsetting contribution which could be added to the empty pads. Based on prototype dimer measurements, g_{LC} is near $0.1\text{k}\Omega^{-1}$.

5.1.3 Vactrol Control Logic

Vactrol LED current is controlled by a DAC and transistor. Working backward, the control scheme is as follows. The vactrol is on the collector of a transistor, so the vactrol current is dictated by the transistor's collector current. The collector current is in turn controlled by the base-emitter current. The base-emitter current is determined by the voltage of a DAC. There is also additional circuitry intended to offset the DAC output range. The DAC is controlled by a 3-wire serial communication scheme, interfaced through the AXON to the PC.

The serial communication scheme involves three pins on each DAC. Each has a CLOCK pin, a DATA pin, and a !SYNC (don't sync) pin. The DAC will only receive communications when its !SYNC pin is low, otherwise it will keep its setting fixed. Therefore, we provide individual control over each !SYNC pin seperately. All eight CLOCK and DATA pins are connected together, to the AXON's clock and data outputs. In order to change one DAC's setting, we simply must bring the desired DAC's !SYNC pin low, then toggle the CLOCK and DATA pins, connected to all DACs, according to the serial data transfer protocol, which can be found in the DAC datasheet. Only the DAC whose !SYNC pin was addressed will accept the communication and change its setting.

!SYNC addressing is mediated by a 3-to-8 line decoder chip, which accepts a three digit binary number, and brings one of eight output pins low accordingly. The AXON address one of the !SYNC pins by representing a three digit binary number (from 0-7) on three digital pins. These pins are wired to the input of the decoder, which brings the corresponding decoder output pin low if the decoder is enabled. An additional digital AXON pin is devoted to enabling the decoder—when the enable pin is low all decoder outputs are high; when the enable pin is high one decoder output is low. Thus to make a communication, the desired DAC is first addressed, then the enable pin is brought high. Once the enable is activated, the addressed DAC's !SYNC pin goes low, and communication can begin. Beware that the binary address 000 addresses the DAC at "OSC #1", etc.

The DAC setting D has twelve bit precision. The setting is an integer from 0 to 4095. DAC output voltage is given by

$$V_{\mathrm{DAC}} = \left(\frac{D}{4096}\right) V_s \tag{5.2}$$

To summarize the vactrol control logic circuitry:

- Relevant AXON pins are [1] Decoder Enable; [2][3][4] Decoder Address; [5] DAC Clock; [6] DAC Data.
- Every DAC's clock pin is connected, together, to [5]. Every DAC's data pin is connected, together, to [6].
- Each DAC !SYNC pin is connected to one decoder output pin.
- The three decoder address pins are connected to [2][3][4].
- The decoder enable pin is connected to [1].

And the process of changing the setting of one DAC:

(1) Set the pins [2][3][4] to the binary number corresponding to the desired oscillator.

(2) Activate the clock on pin [5].

(3) Bring pin [1] high to enable the decoder. This brings the !SYNC pin of the addressed DAC low.

(4) Toggle data pin [5], according to the serial data transfer protocol, to transfer the DAC a twelve bit binary number from 0-4095.

(5) Bring pin [1] low to disable the decoder, pulling the !SYNC pin high. Communication is terminated. The DAC voltage changes according to the new setting.

(6) Deactivate the clock.

The effect of changing the DAC setting is then:
DAC setting changed according to protocol above $\Longrightarrow$ voltage and current seen by the transistor base change accordingly $\Longrightarrow$ vactrol LED current changes $\Longrightarrow$ vactrol photocell conductance changes.

The AXON has been programmed to take inputs from Visual Basic and execute the algorithm described above.

5.1.4 Vactrol Control Analog

The analog end of the vactrol control circuitry consists primarily of the transistor which draws the LED current, and the DAC which controls current through the transistor.

Due to the base-emitter voltage drop V_{be} at the transistor, driving the base current directly with only the DAC would result in a wasted portion of the DAC range. Up to V_{be}, changes to the DAC setting would have no effect, as the transistor is off. To eliminate this waste, we can bias the base of the transistor with a voltage from the +5V supply line. The circuit shown in Fig. 5.2 is equivalent to our biasing circuit. The Thevenin equivalent of the "OUT" terminal is the equivalent circuit seen by the base of the transistor. It is with this Thevenin equivalent that we are concerned.

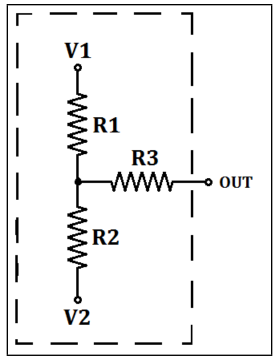

Figure 5.2: This circuit is used to bias the transistor and make the output range of the DAC more useful.

The "OUT" terminal has Thevenin equivalent

$$V_{TH} = V_1 \left(\frac{R_2}{R_1 + R_2} \right) + V_2 \left(\frac{R_1}{R_1 + R_2} \right) \quad (5.3)$$

$$R_{TH} = R_3 + \frac{R_1 R_2}{R_1 + R_2} \quad (5.4)$$

R_{TH} is independent of the voltages, and the equivalent voltage is linear in both V_1 and V_2. In our case, $V_1 = 5\text{V}$, and $V_2 = V_{\text{DAC}}$. This circuit will work perfectly for our biasing purposes. Once we choose the voltage at DAC=0, it will increase linearly with DAC setting up to its max of 5V. What remains is to decide precisely what conductance range, and therefore DAC range, is required.

We chose to aim for a conductance range of approximately $g_{vtl} \in (0.1, 3)\text{k}\Omega^{-1}$ based loosely on the following considerations:

- Small conductances should be avoided, we want to confine ourselves to the linear part of the current-conductance curve.
- The range will be offset by some resistance so that $\gamma = 0$ lies in the middle of the range. A bigger range will mean more more offsetting conductance, and thereby a greater current requirement.
- The further the vactrol is used from its resting preset state, the less accuracy it will attain.

- Having a larger range will increase our ability to perform experiments without modifying the circuit.
- For a typical inter-oscillator capacitive coupling value of 56pF, $\gamma_{PT} \Longrightarrow \sim 1\text{k}\Omega^{-1}$.

Comparison with Figure 4.3 shows that to attain our chosen conductance range with $R_{th} = 220\text{k}\Omega$, we should have $V_{th} \in (0.9, 5)\text{V}$. Deciding to err on the side of a bigger range, resistor values for the biasing circuit were chosen based on the requirements:

$$R_{th} = 200\text{k}\Omega \quad V_{th}(DAC = 0) = 0.9\text{V}$$

Since the circuit is substantially different from that used in the experiments of Fig. 4.3, we only expect our goals for the conductance range to be loosely approximated.

To summarize, the vactrol control circuitry was designed with the intention that with the DAC at its minimum setting, DAC=0, the vactrol conductance would be near $0.1\text{k}\Omega^{-1}$, and that conductance would increase linearly with DAC setting up to a maximum conductance near $3\text{k}\Omega^{-1}$. To acheive this, circuitry was added to overcome the transistor dead zone. How well those numbers were actually attained remains to be seen experimentally.

5.1.5 Input

Two input terminals are provided for external signal input. At each, a video buffer chip provides two following buffers, one for measuring the input signal, and another for coupling the input to one of the oscillators. The input is intended to be brought onto the board on a twisted pair. A 100Ω resistor to ground is included as half of a 50Ω transmission line termination. Putting a parallel 100Ω resistor at the interface of the BNC with the twisted pair will effectively spread the termination out across the connection, giving the best termination. Input coupling impedance can be soldered into the empty through-hole pads provided, or the pads can be shorted and the input coupling impedance can be incorporated into the twisted pair going to the oscillator.

5.1.6 Output and Multiplexer

Each LC node lies on the input of a voltage following buffer. Each of these buffered voltages is connected to the input of a four channel multiplexer. The multiplexer takes in the four inputs, and outputs just one. Oscillator Bank #1 is associated with MUX1. Bank #2 goes with MUX2.

The output channel of each multiplexer is controlled by the binary value of two control pins. Axon pins [7][8] and [9][10] set this value for MUX1 and MUX2 respectively. Binary zero selects OSC#1 for MUX1, and OSC#5 for MUX2. Binary one selects OSC#2 for MUX2 and OSC#6 for MUX2. The pattern continues for binary two and three. The multiplexer enable pins are hardwired so that the multiplexers always output some channel. They have no memory; their output is always determined by the binary value of their channel selection pins at that moment.

The MUX outputs each go to one additional round of buffers to ensure the output can drive the transmission line to the scope. The output can be measured by connecting a twisted pair from a BNC to the pins, or with the scope probe.

5.1.7 Power

The power circuitry includes standard ±5V regulators along with numerous bypass capacitors. Diodes prevent damage due to reverse biasing, and an LED indicates active power. Small and large bypass capacitors are provided as close as possible to the power pins of all op-amps, and additional bypass capacitors are included copiously along the power wiring on the bottom side of the board.

5.1.8 Physical Characteristics

The board was manufactured and assembled by Pad2Pad[1] and designed on PC using Pad2Pad's freely downloadable circuit design software. The two-layer board is a 8″ x 8.5″ rectangle of 0.062″ thick FR4 dielectric ($\epsilon_r = 4.8$). Traces and ground plane have

[1] *Pad2Pad Custom PCBs.* Mahwah, NJ. <www.pad2pad.com>.

copper weight 1oz/ft^2. Traces are of varying width. The ground plane is etched to a gap width of 0.020$''$ around traces. All unused top-side and bottom-side board space is covered in a copper ground plane, giving traces the "coplanar waveguide with ground" configuration, which can be used to calculate trace impedances. Thermal connections are used at grounded soldering pads. The bottom side is devoted to power circuitry, while all signals and logic occur on the top side (with the small exception of a few logic lines detouring temporarily to the bottom).

5.1.9 PC Interfacing

The octoboard logic circuits are interfaced to PC via an AXON II microcontroller. The Axon comes as an integrated digital interface, fully equipped for USB communication and digital output. Visual basic can send communications from the PC to the Axon via USB commport. The axon's customizeable operating system was programmed and loaded up for our purposes.

The Axon is in a constant loop of looking for communications of the correct syntax to trigger one of its subroutines. The two subroutines include a DAC setting routine which executes the algorithm described in 5.1.3, and an output setting routine which alters the channel selection pins [7][8][9][10]. When one of these subroutines completes, the Axon goes back to looking for more communications. Visual basic is used to trigger the subroutines by passing along communications of the correct form, which include which subroutine to trigger, and the relevant variable values.

All PC communication functions have been confirmed working. Each DAC can be set, and any output can be selected, by simply executing a line of VB code.

5.2 Calibration Instructions

The following steps should be taken to calibrate the vactrols and get the board in working order. Keep in mind that jumpers allow each vactrol to either be isolated, or incorporated into the circuit. Calibrations should be checked often, and don't need to be permanent. If different experiments call for different conditions for use, it's reasonable to use a new calibration with conditions that mimic those of the experiment. The following procedure should at least get things started. For each vactrol:

(1) Age the vactrol by running it at a high conductance, say DAC=2500, for a week. Preferably record one vactrol's conductance over this aging period to confirm that long-term effects come to equilibrium.

(2) Choose a "resting", or "preset", setting: Connect the vactrol to the oscillator, including the loss-biasing offset resistor. Determine and record the smallest DAC setting which induces self-oscillation. If everything is functioning as theoretically intended, this should be near DAC=2000. Take the setting which induced self-oscillation and subtract a few hundred. Use this setting as the preset setting for that vactrol. The idea is that we want the preset to be near the largest stable conductance value. Disconnect the vactrol again, and leave it overnight at its newly chosen preset setting.

(3) Record plots of conductance versus DAC setting under different scanning conditions. Make sure that all scans of the DAC range are terminated by returning to the preset setting. Check that for different step sizes ΔDAC, settling times dt between steps, and scanning directions, results are compatible. If things are weird at this point take note, but forge ahead anyway.

(4) Choose a calibration settling time dt_{cal}. This should be slightly slower than the speed of experimental changes. 15-30s might be a good start. Choose a calibration step size big enough that calibrating won't take hours, but small enough to provide adequate data density. Integer jumps of 100 or 200 ought to do. Keep in mind that while getting more data points sounds nice, long calibration scans risk upsetting the carefully preset hysteresis control.

(5) Perform a calibration. Scan over the DAC range with the chosen step parameters ΔDAC_{cal} and dt_{cal}, recording conductance at the end of each step. This data, DAC setting vs. conductance, will be the basis for calibration.

(6) Fit the data. Take the conductance versus DAC data and shift it so that a datapoint lies at the origin. Perform a polynomial fit to get a shifted calibration function. Shift it back to the original range to get the actual calibration function. Do the same thing for the inverse data. There should now be two polynomial functions. One, $DAC(\gamma)$, gives DAC setting as a function of desired conductance. The other, $\gamma(DAC)$, gives expected conductance as a function of DAC setting. Store the calibration parameters in a standard way so a generic calibration checking program can read out each oscillator's calibration fit functions.

(7) Perform a calibration check under similar conditions to the calibration itself. Pick a check settling time $dt \leq dt_{cal}$. For checking the calibration, it is useful to output several data. Useful plots are conductance versus DAC and conductance versus time. The first can record desired and measured conductance at each setting, and percent difference. The second should report temporal data, keeping track of what the conductance is doing in between the "official" measurements, during the stepping and settling periods. See Fig. 4.6.

(8) If the calibration check is ugly, go back and figure out what's wrong. If it's encouraging, good work. Now perform the check under several other sets of conditions to see how it holds up. Try different directions, scanning over only smaller ranges, different step sizes and settling times, or aim for random patterns of desired conductance. Taking checks is easy and changing a lot of things gives a good idea of what will effect the precision of the calibration. Definitely check the calibration under conditions that mimic the desired experiment. Continue to check the calibration regularly, probably every day or before each use, and

save these calibration checks for later reference.

Once accurate and robust vactrol calibrations are attained, anything is possible with the octoboard.

5.3 Summary and Recommended Improvements

The octoboard output switching and DAC control functions are fully functional based on initial testing. The main problem that has been identified so far is a grounding and termination issue, which causes unwanted high frequency oscillations at some of the oscillators. We believe the grounding issues can be fixed by making simple manual modifications to the board. Future circuits based on the octoboard should respect the following recommended design improvements:

- Fix known grounding issues by providing appropriate terminations and additional ground connections near the problem areas.
- "To Axon" pins should have a two-row header for use with a ribbon connector with ground. As it stands, ground is completely left out of this connector, and must be seperately connected between the axon and board.
- Mounting holes could be provided for bolting to an external grounded case for the board and axon.

Bibliography

[1] Joseph Schindler, Ang Li, Mei C Zheng, FM Ellis, and Tsampikos Kottos. Experimental study of active LRC circuits with PT symmetries. *Physical Review A*, 84(4):040101, 2011.

[2] Zin Lin, Joseph Schindler, Fred M Ellis, and Tsampikos Kottos. Experimental observation of the dual behavior of PT-symmetric scattering. *Physical Review A*, 85(5):050101, 2012.

[3] Hamidreza Ramezani, J Schindler, FM Ellis, Uwe Günther, and Tsampikos Kottos. Bypassing the bandwidth theorem with PT symmetry. *Physical Review A*, 85(6):062122, 2012.

[4] J Schindler, Z Lin, JM Lee, H Ramezani, FM Ellis, and T Kottos. PT-symmetric electronics. *Journal of Physics A: Mathematical and Theoretical*, 45(44):444029, 2012.

[5] Carl M. Bender and Stefan Boettcher. Real spectra in non-hermitian hamiltonians having PT symmetry. *Phys. Rev. Lett.*, 80:5243, 6 1998.

Appendix A

Octoboard Schematics

This appendix contains the full Octoboard technical specifications. Materials occur in the following order:

1. Circuit Diagrams and Schematics
2. Component Specifications
3. Photographic Scans
4. Layer-by-Layer PCB Drawings
5. Close-Ups of Each Functional Block

CIRCUIT DIAGRAMS AND SCHEMATICS

Generic Components

Labels not shown in circuit diagrams for these generic components:

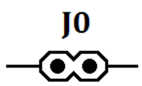

J0 = 2-pin Jumper Heading

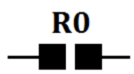

R0 = Empty Solder Pads

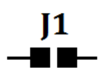

J1 = Solder Jumper

A# = Triple Video Buffer ADA4862-3

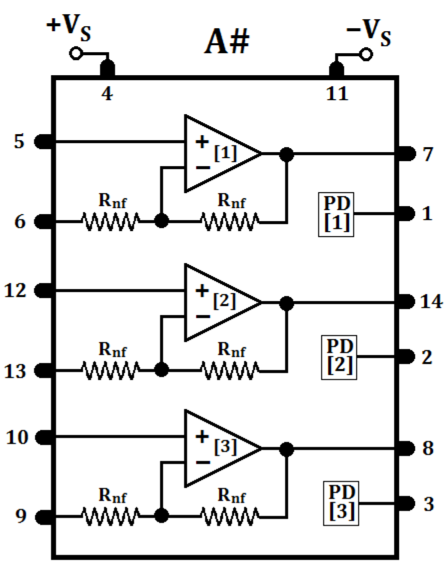

Labeled A# in these diagrams, but numbered on the board.

Inverting input pins (6,13,9) control gain. Floating yields 1x follower. Grounding yields 2x amplifier. Rnf is internal to the IC.

Power down pins (1,2,3) disable corresponding amp when held above -Vs+1V. Amplifiers enabled when PD# floating or held low.

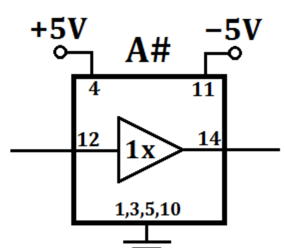

A# drawn in diagrams with internal connections suppressed or idealized, irrelevant sections suppressed, and pins with no external connection not shown. (See example to right.)

OSCILLATOR: Node

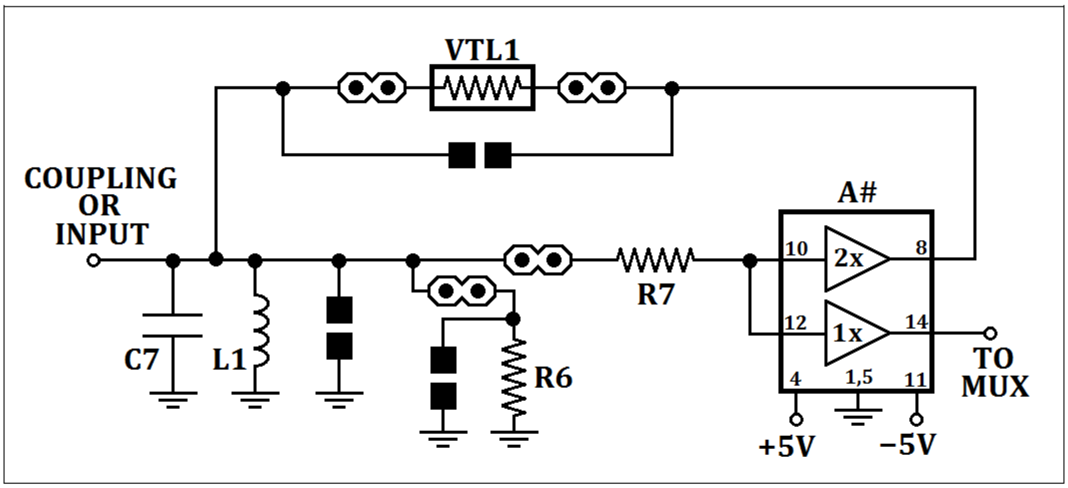

A# = Triple Video Buffer ADA4862-3

C7 = 330pF, Oscillator Capacitor

L1 = 10uH, Oscillator Inductor

R6 = 620 Ohm, Loss Biasing [gamma(DAC=0) ~ 1.6 kOhm^{-1}]

R7 = 133 Ohm, Low Pass Filtering (Self-Osc Supression)
[With 4pF capacitance of two A# input pins, f_{3dB} = 299MHz]

VTL1 = Vactrol NSL-32 Photoresistor [R_{vtl} Range ~ 10kOhms to 0.3kOhms]

Two-pin jumper headers with one side grounded provided on left to allow nodal coupling or inputs via twisted pair.

Power decoupling capacitors not shown.

OSCILLATOR: Vactrol Control

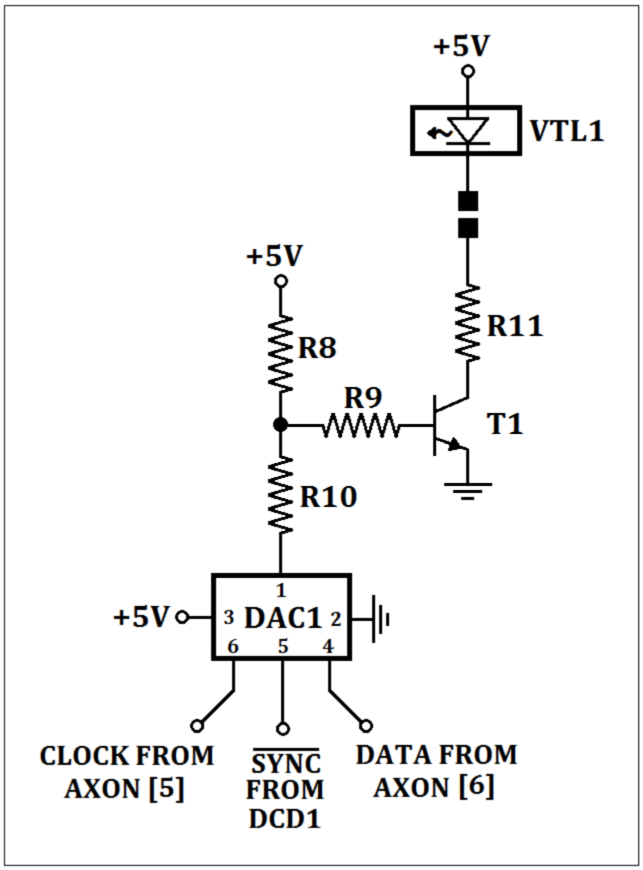

VTL1 = Vactrol NSL-32 LED, Control Gain Photoresistor

T1 = NPN Transistor (SOT-23 Package), Draw VTL Current

DAC1 = DAC AD5320, Control VTL LED Current via T1

R8 = 16.5 kOhm, DAC Range Biasing (5V Side)

R9 = 196 kOhm, Conduct T1 Base Current

R10 = 3.4 kOhm, DAC Range Biasing (DAC Side)

R11 = 47 Ohm, Limit Max VTL LED Current (to 100mA, so does nothing)

Values of R8,R9,R10 chosen so the Thevenin equivalent input into the transistor has R_{th}=200kOhms and $0.9V < V_{th} < 5V$, where V_{th} depends on DAC setting. V_{th} varies linearly with DAC setting.

"AXON [#]" refers to pin number # on the "To Axon" pin strip.

INPUT

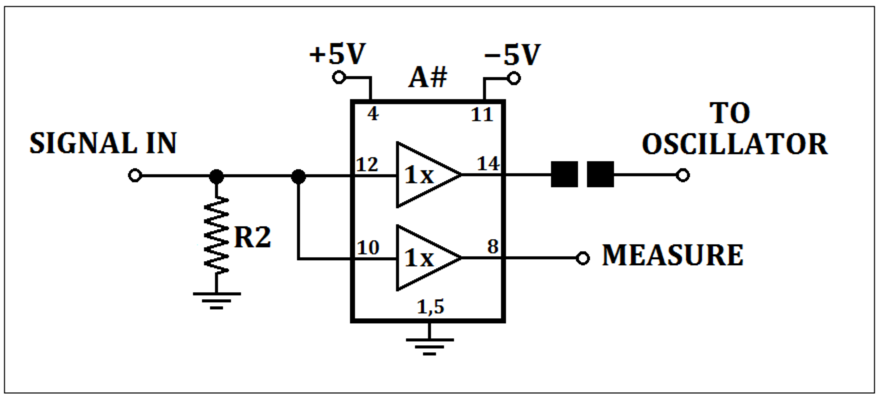

A# = Triple Video Buffer ADA4862-3

R2 = 100 Ohm, Half TL Termination

Empty pads can be used for input coupling impedance, or shorted.

Additional power decoupling capacitors not shown.

OUTPUT

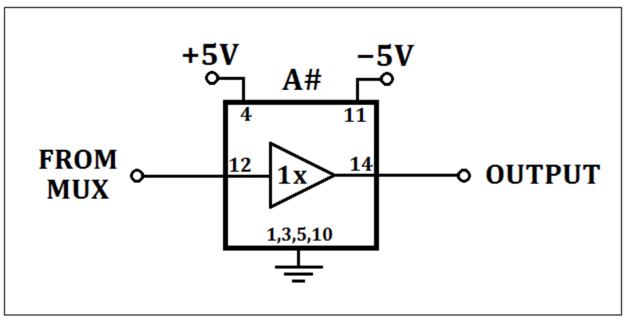

A# = Triple Video Buffer ADA4862-3

Additional power decoupling capacitors not shown.

MULTIPLEXER

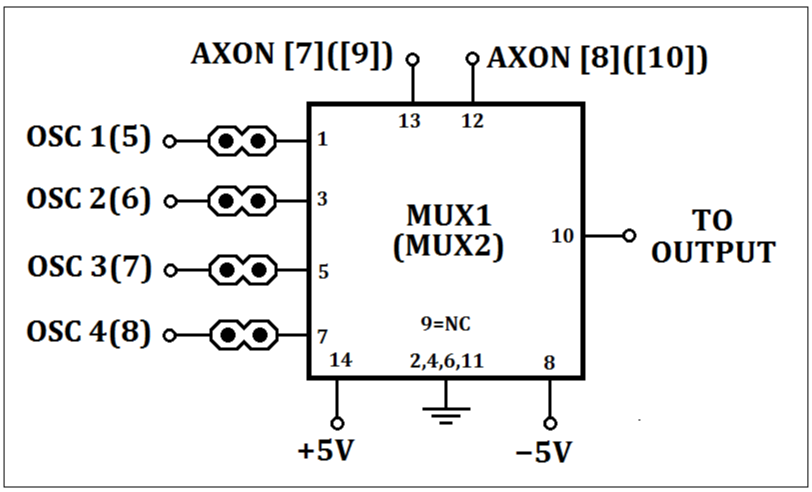

MUX# = Multiplexer AD8184, Output Channel Control

"AXON [#]" refers to pin number # on the "To Axon" pin strip.

Power decoupling capacitors not shown.

POWER

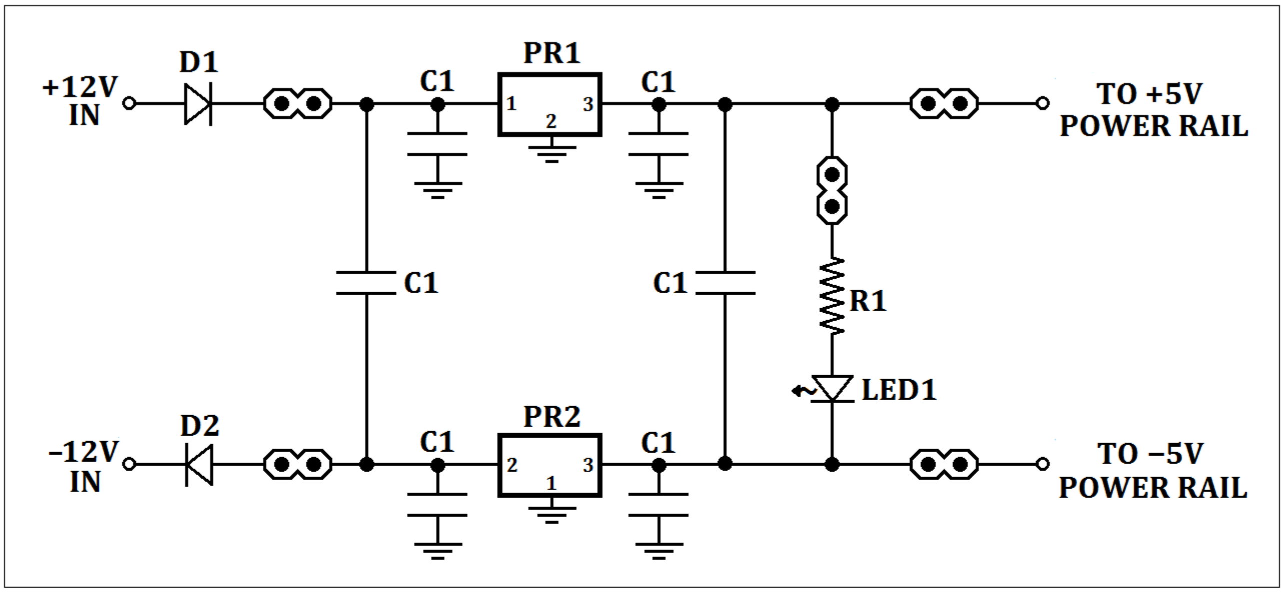


```
PR1 = +5V Regulator LM7805
PR2 = -5V Regulator LM7905

D1,D2 = Reverse Bias Protection Diodes

C1 = 1uF Bypass

LED1 = Power Indicator LED

R1 = 430 Ohm, Limit LED Current (18mA)
```

LOGIC

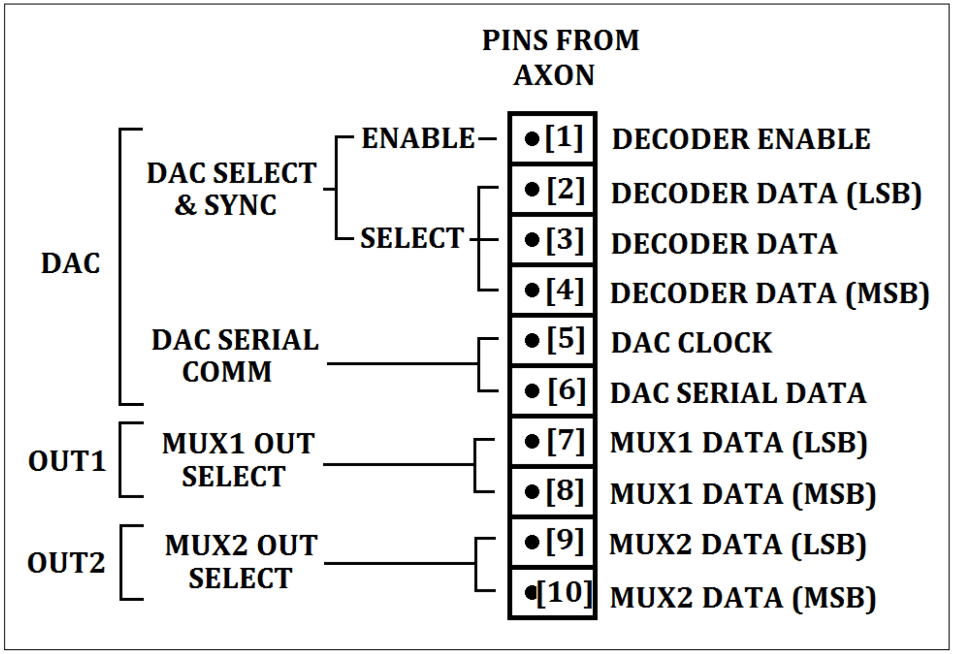

MUX Output Channel Selection:

--Each MUX is permanently enabled, so some output is always selected.

--Pins [7],[8] for MUX1 and [9],[10] for MUX2 set a binary integer from 0-3. This integer directly selects one of the four possible DAC output channels. Binary "00" corresponds to OSC#1 for Bank#1 and OSC#5 for Bank#2, etc.

DAC Communication:

--Each of the eight decoder outputs is connected to one DAC's !SYNC pin. All decoder outputs float high when disabled, so no DACs sync while decoder is disabled.

--Pins [2],[3],[4] form a three bit integer from 0-7 which selects OSC #1-#8. Binary "000" corresponds to OSC #1, and so on.

--When pin [1] goes high, decoder is enabled, bringing !SYNC low at whichever oscillator is selected.

--Serial communication occurs while !SYNC is low. Pins [5] (clock) and [6] (data) are each connected to all eight DACs, but communication only occurs at the one DAC with !SYNC low. Pin [1] low ends communication. See DAC AD5320 datasheet for serial comm data protocol.

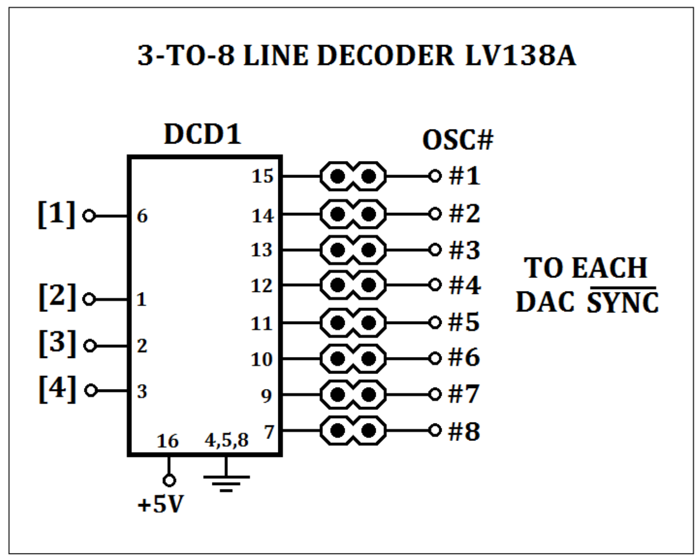

COMPONENT SPECIFICATIONS

General Info 1

Label	Type	Value	Purpose
General			
J0	Jumper	---	2pin .1"pitch Jumper Heading
A#	Video Buffer	ADA4862-3	Triple Voltage Following/Doubling Buffer
J3	Pin	---	Single Pin, for Grounding
Power			
D1, D2	Diode	---	Reverse Bias Protection
C1	Capacitor	1uF	Bypass
LED1	LED	---	Power Indicator
R1	Resistor	430 Ohm	Limit LED Current [(10V-2.2V)/430Ohm=18mA]
PR1	5V Regulator	LM7805	+5V Regulator
PR2	-5V Regulator	LM7905	-5V Regulator
J0	(see J0 above)		---
J2	Jumper	---	3pin .1" Jumper Heading, Power In
Power Line Decoupling			
C8	Capacitor	.1uF	High Fq Decoupling MLCC Low ESR
C9	Capacitor	22uF	Low Fq Decoupling Tantalum
C10	Capacitor	1uF	General Decoupling MLCC
J1	Jumper	---	Solder Jumper
Input			
R2	Resistor	100 Ohm	Half of Transmission Line Termination
R3	Resistor	z_in	[Empty Pads] Input Coupling Impedance
A1	(see A# above)		Two Voltage Following Buffers
Output			
J0	(see J0 above)		---
A4	(see A# above)		Voltage Following Buffer, Voltage Doubling Buffer

Package and Description 1

Label	Type	Package	Description
General			
J0	Jumper	---	.1" pitch, 2pin, 5.5mm pin height
A#	Video Buffer	SOIC-14	High Bandwidth Triple Video Buffer OPAMP CF TRPL LP 115mA(max out) 14SOIC
J3	Pin	---	1 pin, 8.08mm height
Power			
D1, D2	Diode	DO-214AC, SMA	1000V(V_r), 1A(I_f), 1.1V@1A(I_f), SMA
C1	Capacitor	0805	CAP CER MLCC 1UF 25V 10% X7R 0805
LED1	LED	----	LED Green, Diffuse, 2.2V drop (typ), 20mA, 30mA (max steady)
R1	Resistor	0805	RES Thick Film 430 OHM .4W 1% 0805
PR1	5V Regulator	TO-220-3	IC REG LDO 5V 1A TO-220
PR2	-5V Regulator	TO-220-3	IC REG LDO -5V 1A TO220-3
J0	(see J0 above)	---	---
J2	Jumper	---	.1" pitch, 3pin, 5.5mm pin height
Power Line Decoupling			
C8	Capacitor	0805	CAP CER 0.1UF 10V 10% X7R 0805 MLCC
C9	Capacitor	1206	Pointy-Side Positive, CAP TANT 22UF 16V 20% 1206 Tantalum ESR=1.4 Ohm
C10	Capacitor	0805	CAP CER 1UF 25V Y5V 0805 MLCC
J1	Jumper	---	---
Input			
R2	Resistor	0805	RES Thin Film 100 OHM 1/4W 1% 0805 SMD
R3	Resistor	---	---
A1	(see A# above)	---	---
Output			
J0	(see J0 above)	---	---
A4	(see A# above)	---	---

Part Numbers 1

Label	**Mfr.**	**Mfr. Part #**	**Digikey #**
General			
J0	3M	961102-6404-AR	3M9447-ND
A#	Analog Devices	ADA4862-3YRZ	ADA4862-3YRZ-ND
J3	TE Connectivity	87224-1	A26541-ND
Power			
D1, D2	ON Semi	MRA4007T3G	MRA4007T3GOSCT-ND
C1	TDK	CGA4J3X7R1E105K125AB	445-5687-1-ND
LED1	Lumex	SSF-LXH305SGD-TR	67-1742-1-ND
R1	Rohm Semi	ESR10EZPF4300	RHM430AECT-ND
PR1	Fairchild Semi	LM7805CT	LM7805CT-ND
PR2	Fairchild Semi	LM7905CT	LM7905CTFS-ND
J0	---	---	---
J2	3M	961103-6404-AR	3M9448-ND
Power Line Decoupling			
C8	Kemet	C0805C104K8RACTU	399-7999-1-ND
C9	Nichicon	F951C226MAAAQ2	493-2946-1-ND
C10	TDK Corp	C2012Y5V1E105Z/0.85	445-1590-1-ND
J1	--	--	--
Input			
R2	Stackpole Electronics	RNCP0805FTD100R	RNCP0805FTD100RCT-ND
R3	---	---	---
A1	---	---	---
Output			
J0	---	---	---
A4	---	---	---

General Info 2

Label	**Type**	**Value**	**Purpose**
Oscillator			
J0	(see J0 above)		---
(osc node)			
R0	Empty Pads	---	[Empty Pads] In case of later use.
L1	Inductor	10uH	Oscillator Inductor
C7	Capacitor	330pF	Oscillator Capacitor
R6	Resistor	620 Ohm	Oscillator Loss Biasing (gamma0=1.6 $kOhm^{-1}$ when DAC=0)
R7	Resistor	133 Ohm	Lo-Pass for Self-Osc Suppression. With 2x(2pF)=4pF capacitance from ADA4862-3 input pins, f_3dB=299MHz.
A3	(see A# above)		Voltage Following Buffer, Voltage Doubling Buffer
VTL1	Vactrol	NSL-32	Vactrol Voltage-Controlled Resistance to Control Gain
(vtl ctrl)			
J1	Jumper	---	Solder Jumper
R11	Resistor	47 Ohm	Limit Vactrol Current (only to ~100mA, so it doesn't really come into play)
T1	Transistor	NPN	Draw Vactrol LED Current
R8	Resistor	16.5 kOhm	Vtl Ctrl Biasing (5V side) (V(DAC=0)=.9V, Rth=200kOhms)
R9	Resistor	196 kOhm	Vtl Ctrl Voltage Input to Transistor Base (Rth=200kOhms)
R10	Resistor	3.4 kOhm	Vtl Ctrl Biasing (DAC side) (V(DAC=0)=.9V, Rth=200kOhms)
DAC1	DAC	AD5320	Control Vactrol LED Current via Transistor
Logic			
DCD1	Line Decoder	LV138A	DAC SYNC Enable
MUX#	Multiplexer	AD8184	Output Multiplexer
TO AXON (J4)	Jumper	---	10pin .1" Jumper Heading, AXON In

Package and Description 2

Label	**Type**	**Package**	**Description**
Oscillator			
J0	(see J0 above)	---	---
(osc node)			
R0	Empty Pads	---	---
L1	Inductor	1812	10uH, 5%, Q=50@2.52MHz, Shielded, 250mA, 1.6 Ohms (DC max), SelfRes=20MHz (min)
C7	Capacitor	0805	330PF 50V 1% NP0 0805 MLCC
R6	Resistor	0805	620 OHM 1/8W 1% 0805 Thick Film
R7	Resistor	0805	133 OHM 1/8W 1% 0805 Thick Film
A3	(see A# above)	---	---
VTL1	Vactrol	Radial Lead	2V@16mA, 40mA(max), 500-500k (typ), Dot=Cathode(Neg)
(vtl ctrl)			
J1	Jumper	---	---
R11	Resistor	0805	49.9 OHM .4W 1% 0805 Thick Film
T1	Transistor	SOT-23	NPN 45V 100MA SOT-23, Vce(sat)=600mV@Ic=100mA
R8	Resistor	0805	16.5K OHM 1/8W .1% 0805 Thin Film
R9	Resistor	0805	196K OHM 1/8W 1% 0805 Thick Film
R10	Resistor	0805	3.4K OHM 1/8W .1% 0805 Thin Film
DAC1	DAC	SOT-23-6	IC DAC 12BIT R-R W/BUFF SOT23-6 Serial Input
Logic			
DCD1	Line Decoder	SOIC-16-N	IC 3 to 8 LINE DECODER 16-SOIC
MUX	Multiplexer	SOIC-14-N	IC VIDEO MULTIPLEXER 4X1 14SOIC
TO AXON (J4)	Jumper	---	.1" pitch, 10pin, 7.5mm pin height

Part Numbers 2

Label	Mfr.	Mfr. Part #	Digikey #
Oscillator			
J0	---	---	---
(osc node)			
R0	---	---	---
L1	TDK	NL453232T-100J-PF	445-6409-1-ND
C7	AVX	08055A331FAT2A	478-6045-1-ND
R6	Panasonic	ERJ-6ENF6200V	P620CCT-ND
R7	Panasonic	ERJ-6ENF1330V	P133CCT-ND
A3	---	---	---
VTL1	Silonex	NSL-32	(Allied Elec #) 70136794
(vtl ctrl)			
J1	---	---	---
R11	Rohm Semi	ESR10EZPF49R9	RHM49.9AECT-ND
T1	ON Semi	BC850BLT1G	BC850BLT1GOSCT-ND
R8	Susumu	RG2012P-1652-B-T5	RG20P16.5KBCT-ND
R9	Panasonic	ERJ-6ENF1963V	P196KCCT-ND
R10	Panasonic	ERA-6AEB3401V	P3.4KDACT-ND
DAC1	Analog Devices	AD5320BRTZ-REEL7	AD5320BRTZREEL7CT-ND
Logic			
DCD1	TI	SN74LV138ADR	296-1676-1-ND
MUX#	Analog Devices	AD8184AR	AD8184AR-ND
TO AXON (J4)	Molex	22102101	WM2730-ND

Quantity on Board of Each Unique Component

#		Name	Digikey #	QTY
1		J0	3M9447-ND	78
2		AA1	ADA4862-3YRZ-ND	12
3		J3	A26541-ND	8
4		DD1	MRA4007T3GOSCT-ND	2
5		C1	445-5687-1-ND	6
6		LED1	67-1742-1-ND	1
7		R1	RHM430AECT-ND	1
8		PR1	LM7805CT-ND	1
9		PR2	LM7905CTFS-ND	1
10		J2	3M9448-ND	1
11		C8	399-7999-1-ND	30
12		C9	493-2946-1-ND	16
13		C10	445-1590-1-ND	11
14		R2	RNCP0805FTD100RCT-ND	2
15		L1	445-6409-1-ND	8
16		C7	478-6045-1-ND	8
17		R6	P620CCT-ND	8
18		R7	P133CCT-ND	8
19		R11	RHM49.9AECT-ND	8
20		T1	BC850BLT1GOSCT-ND	8
21		R8	RG20P16.5KBCT-ND	8
22		R9	P196KCCT-ND	8
23		R10	P3.4KDACT-ND	8
24		DAC1	AD5320BRTZREEL7CT-ND	8
25		DCD1	296-1676-1-ND	1
26		MUX1	AD8184AR-ND	2
27		To Axon (J4)	WM2730-ND	1
			Allied #	
28		VTL1	70136794	8

PHOTOGRAPHIC SCANS

TOP

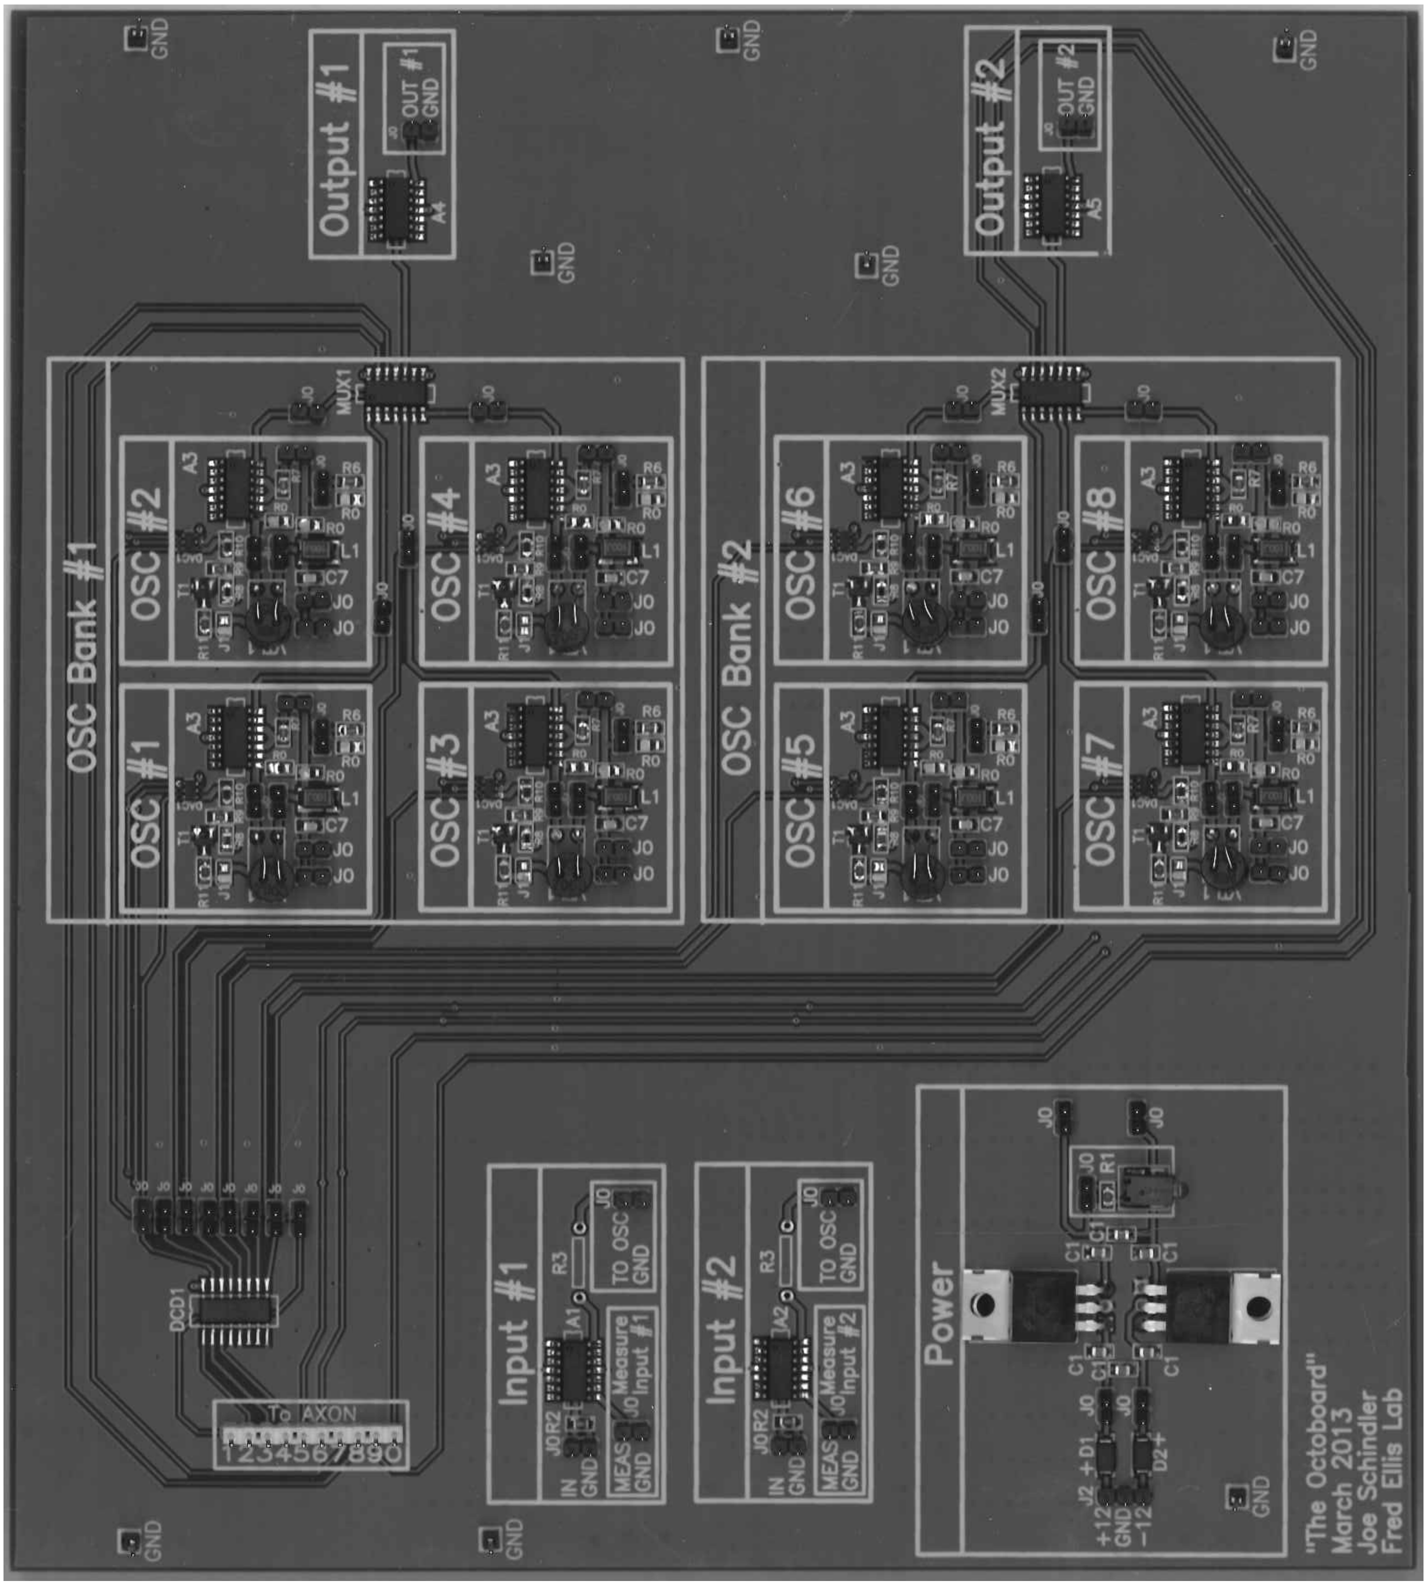
Output #1
Output #2
OSC Bank #1
OSC Bank #2
OSC #1
OSC #2
OSC #3
OSC #4
OSC #5
OSC #6
OSC #7
OSC #8
MUX1
MUX2
DCD1
To AXON
1234567890
Input #1
Input #2
Power
"The Octoboard"
March 2013
Joe Schindler
Fred Ellis Lab

BOTTOM

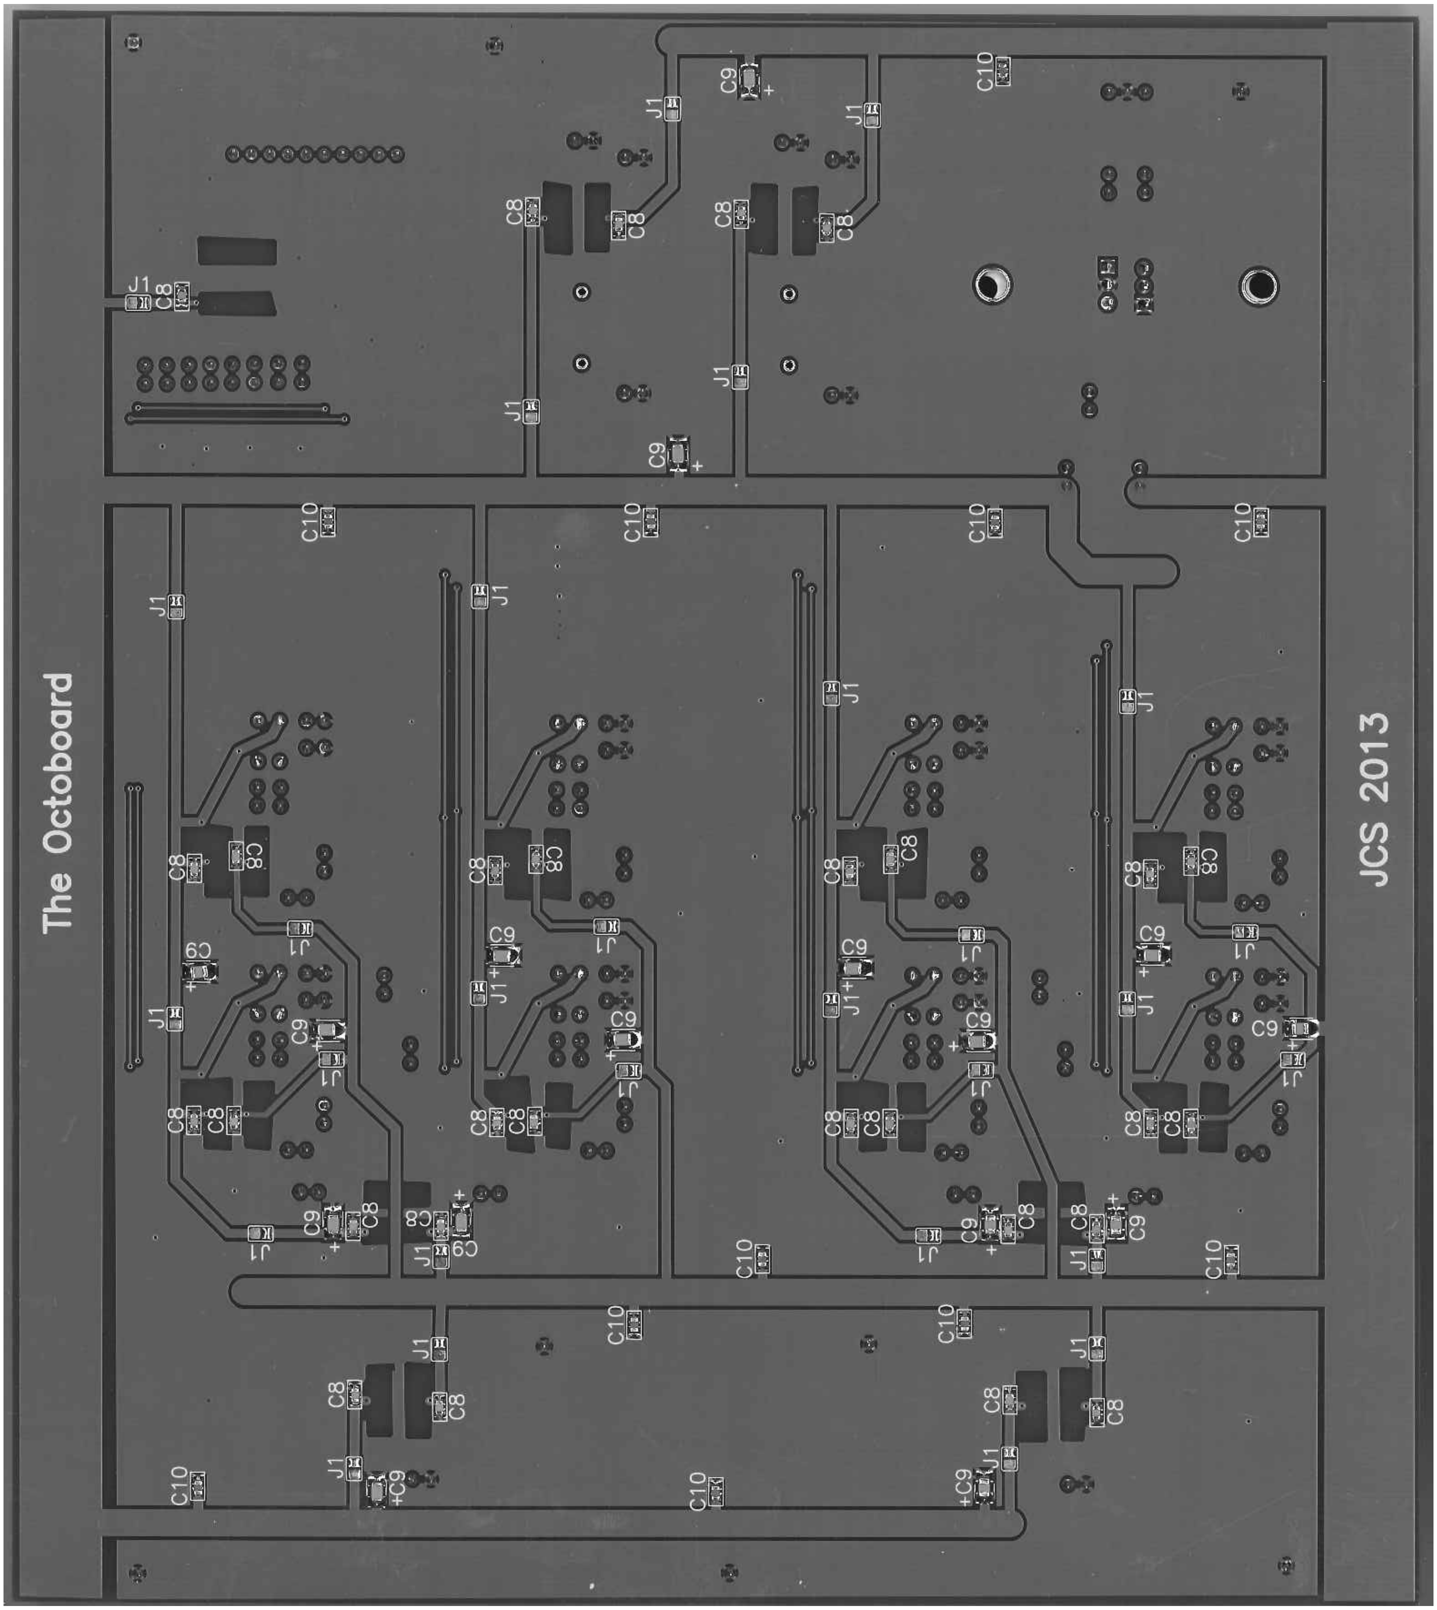

PCB LAYER DRAWINGS

TOP

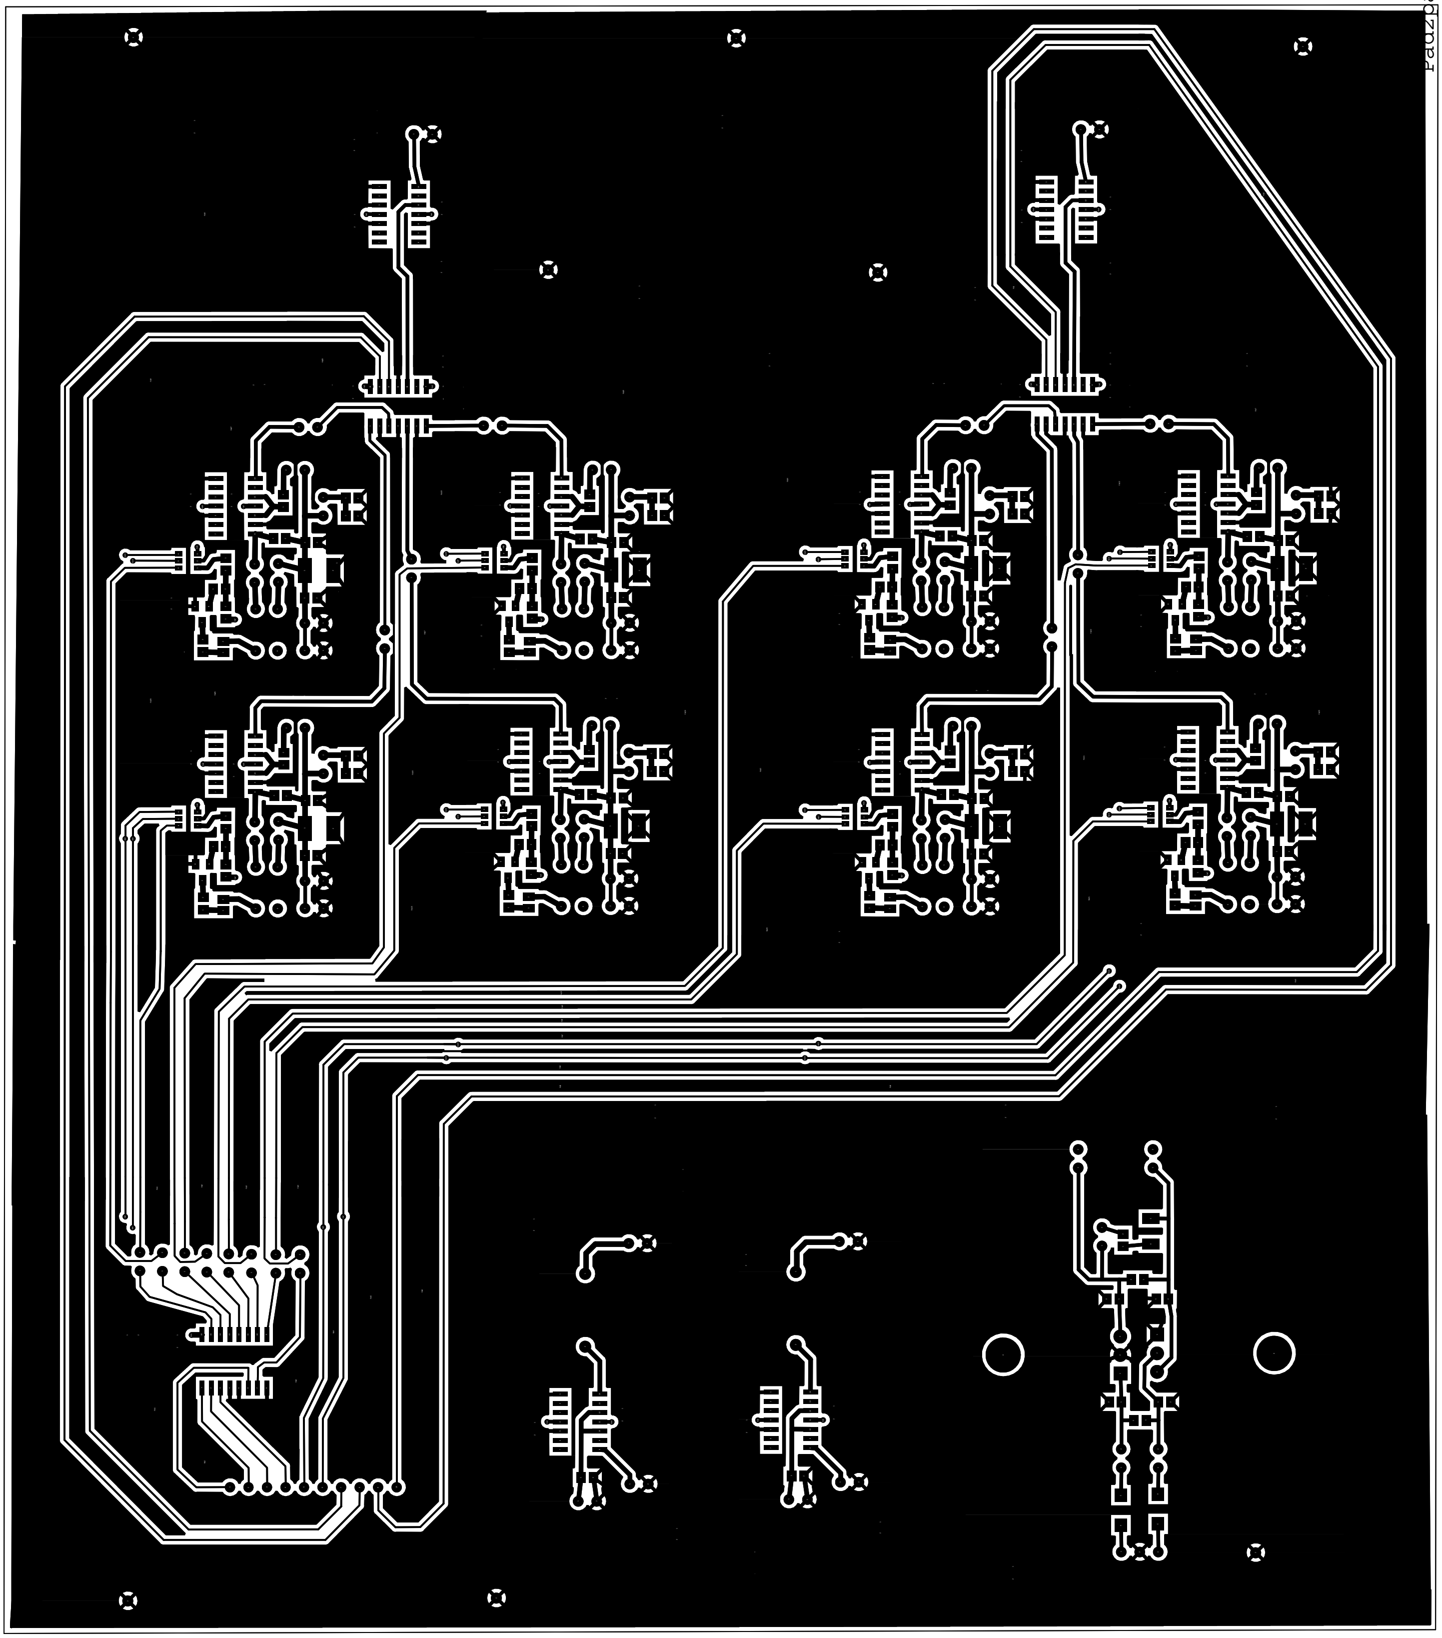

TOP SILKSCREEN

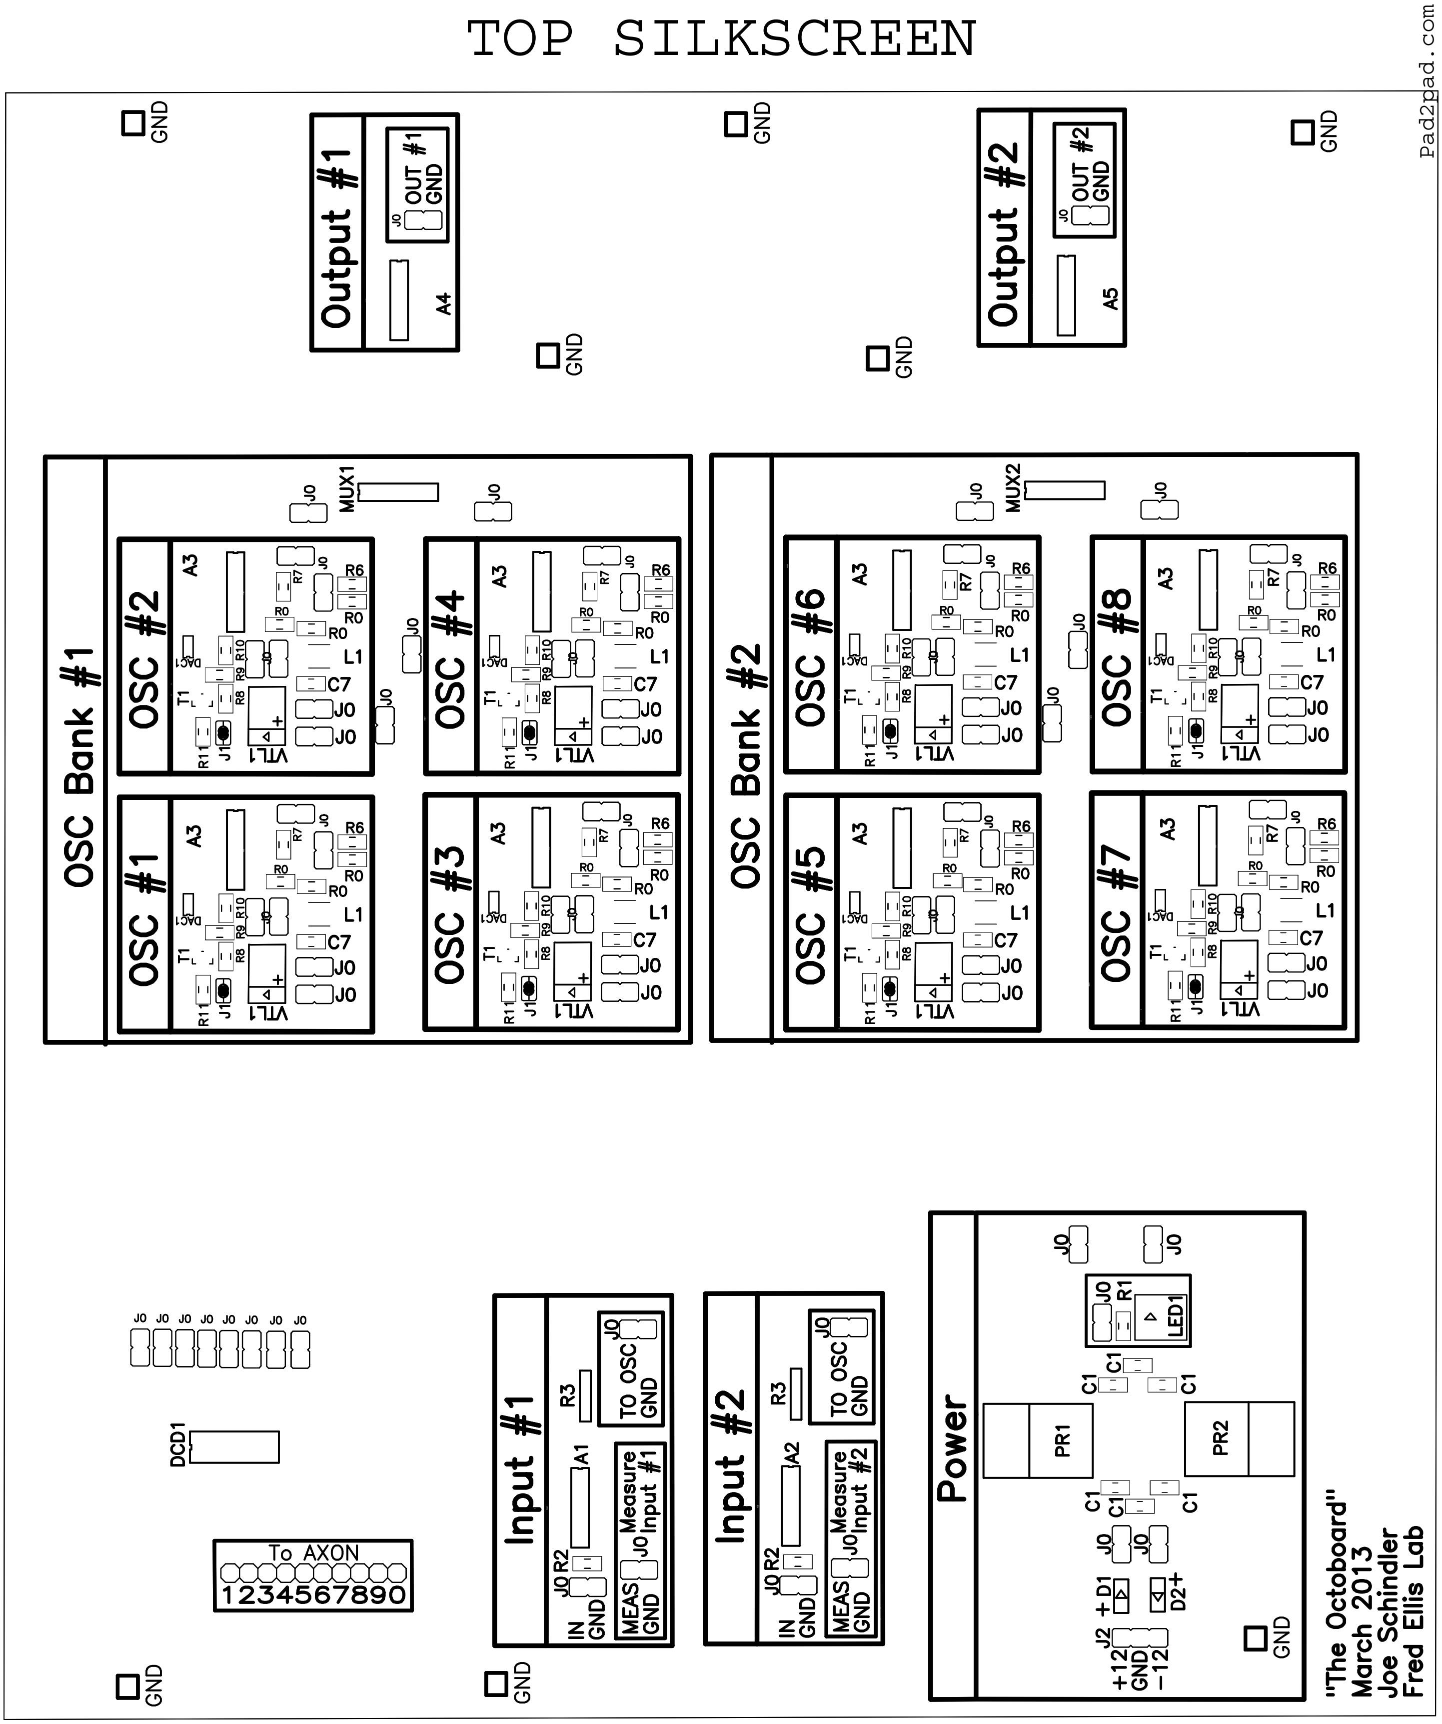

TOP UNMASK

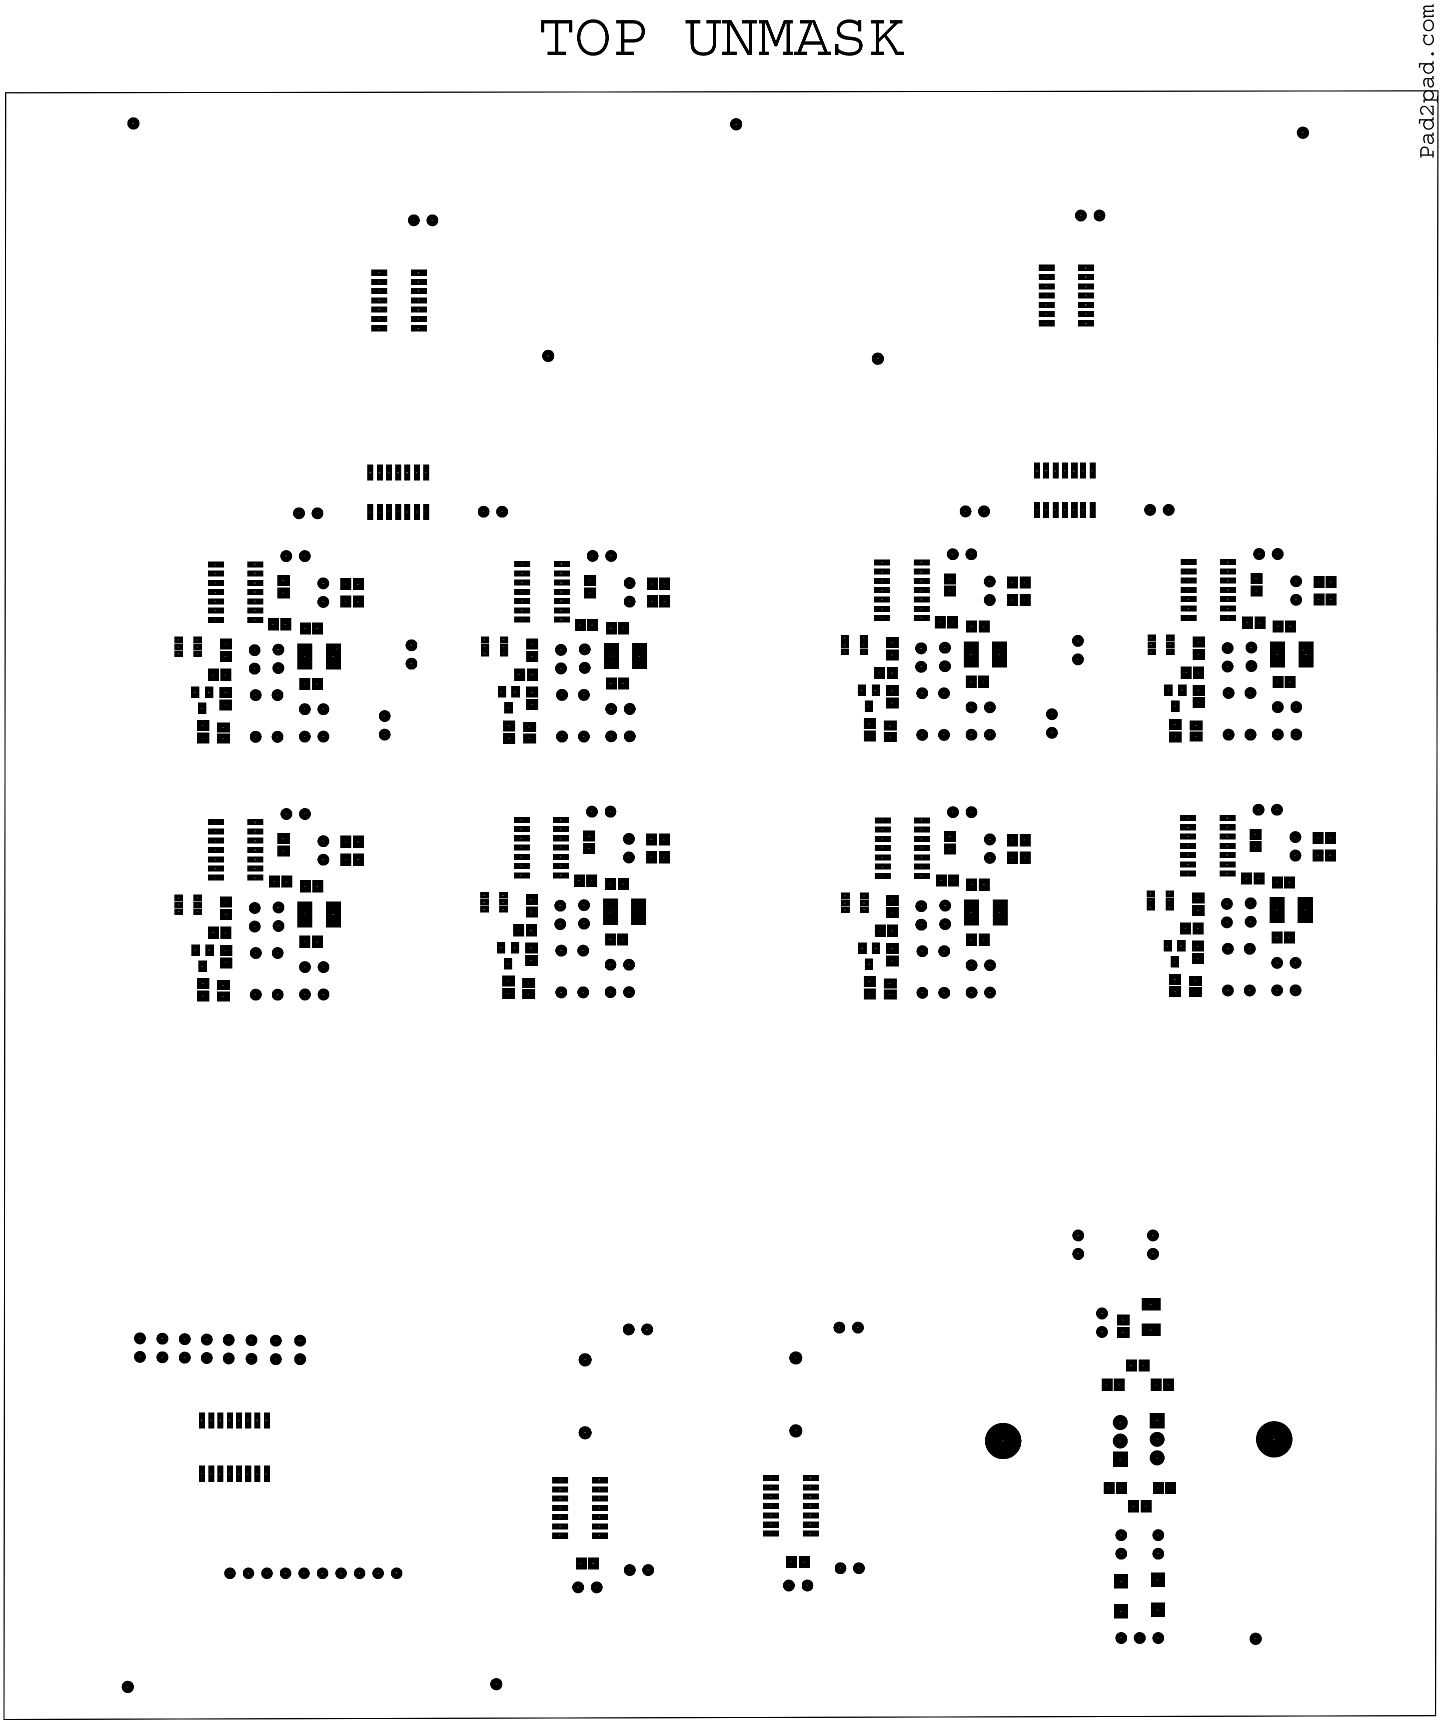

HOLES

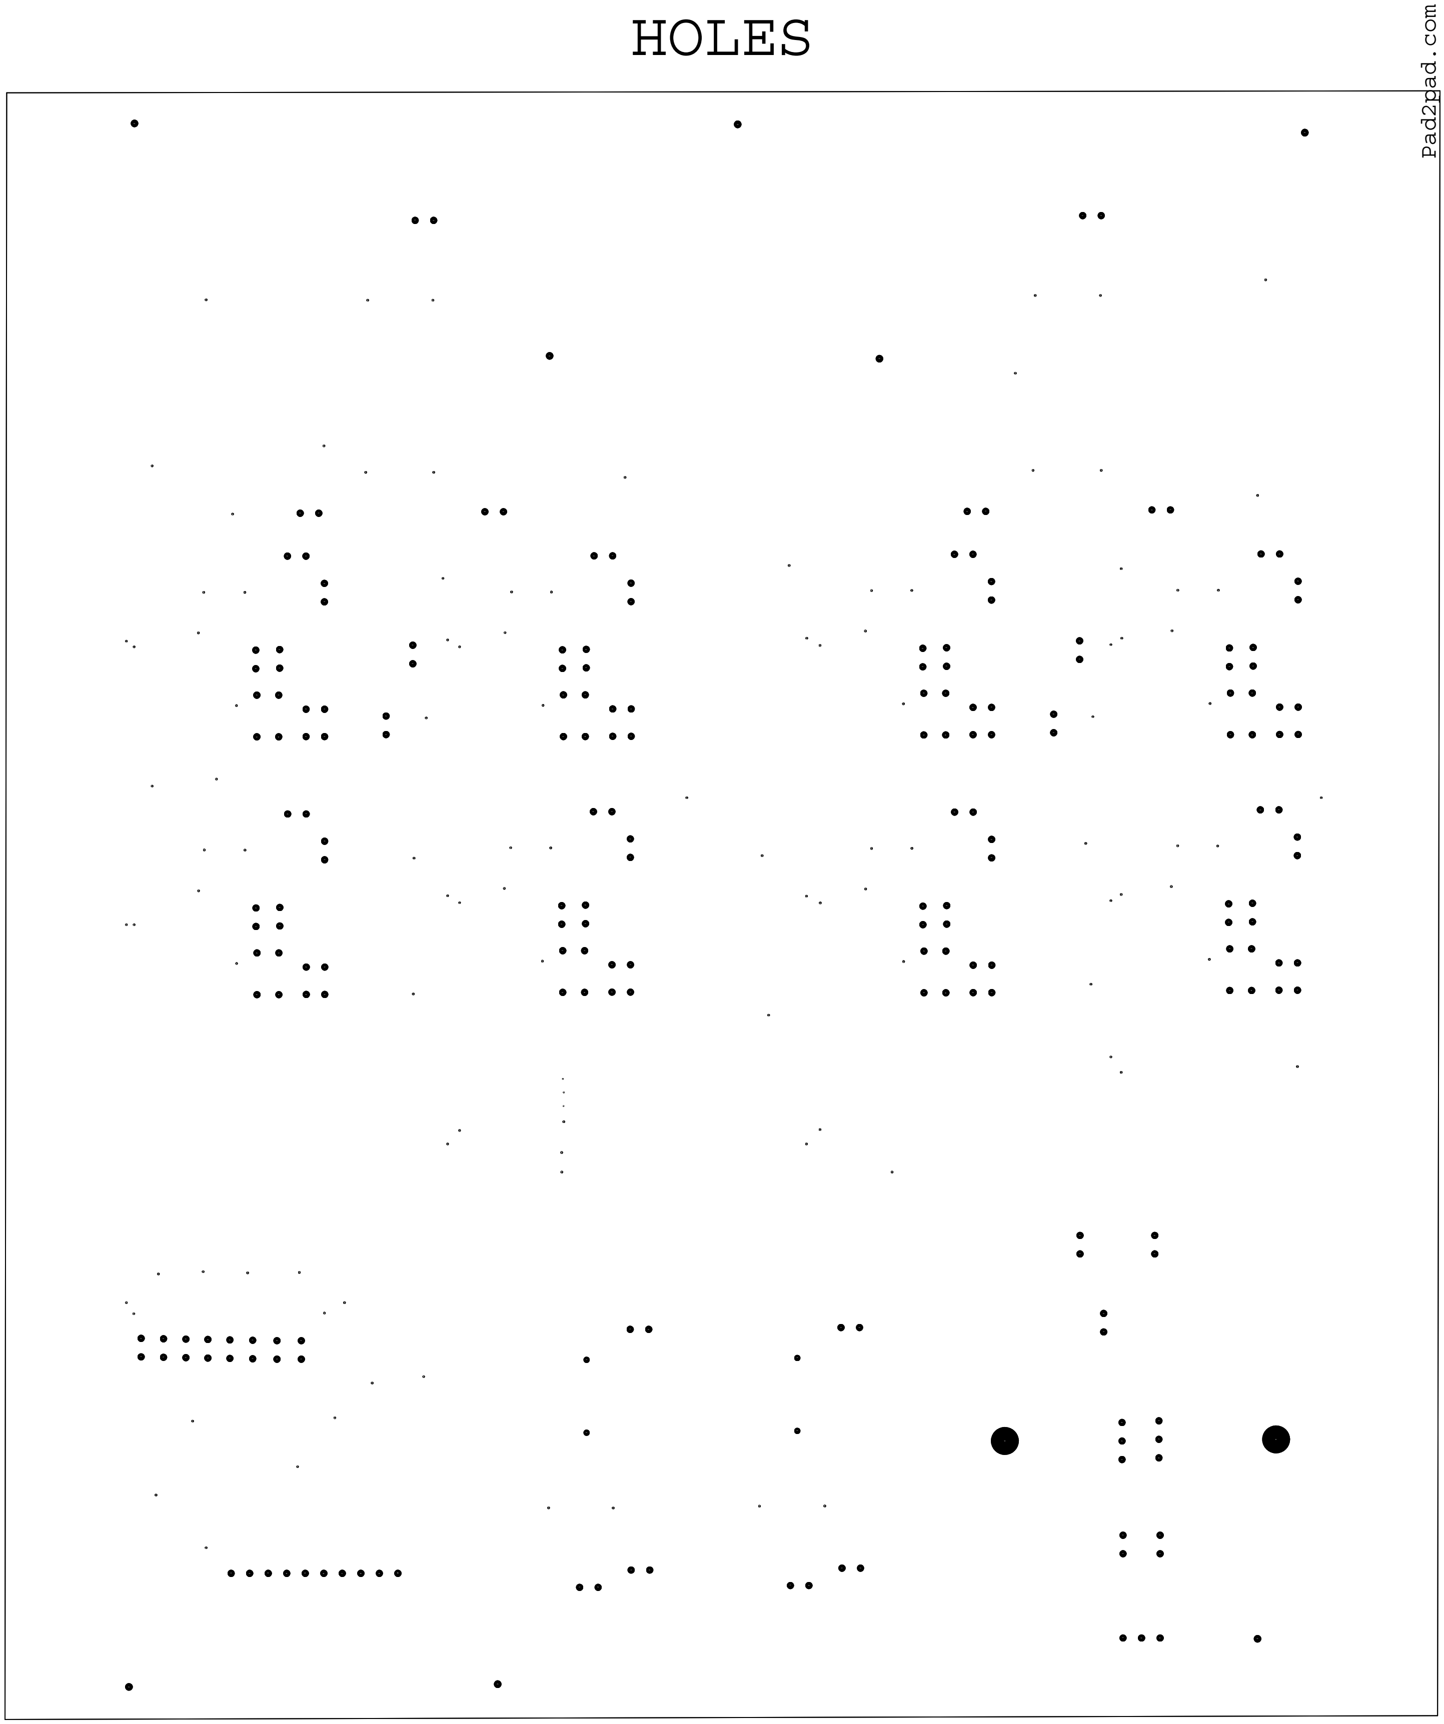

BOTTOM

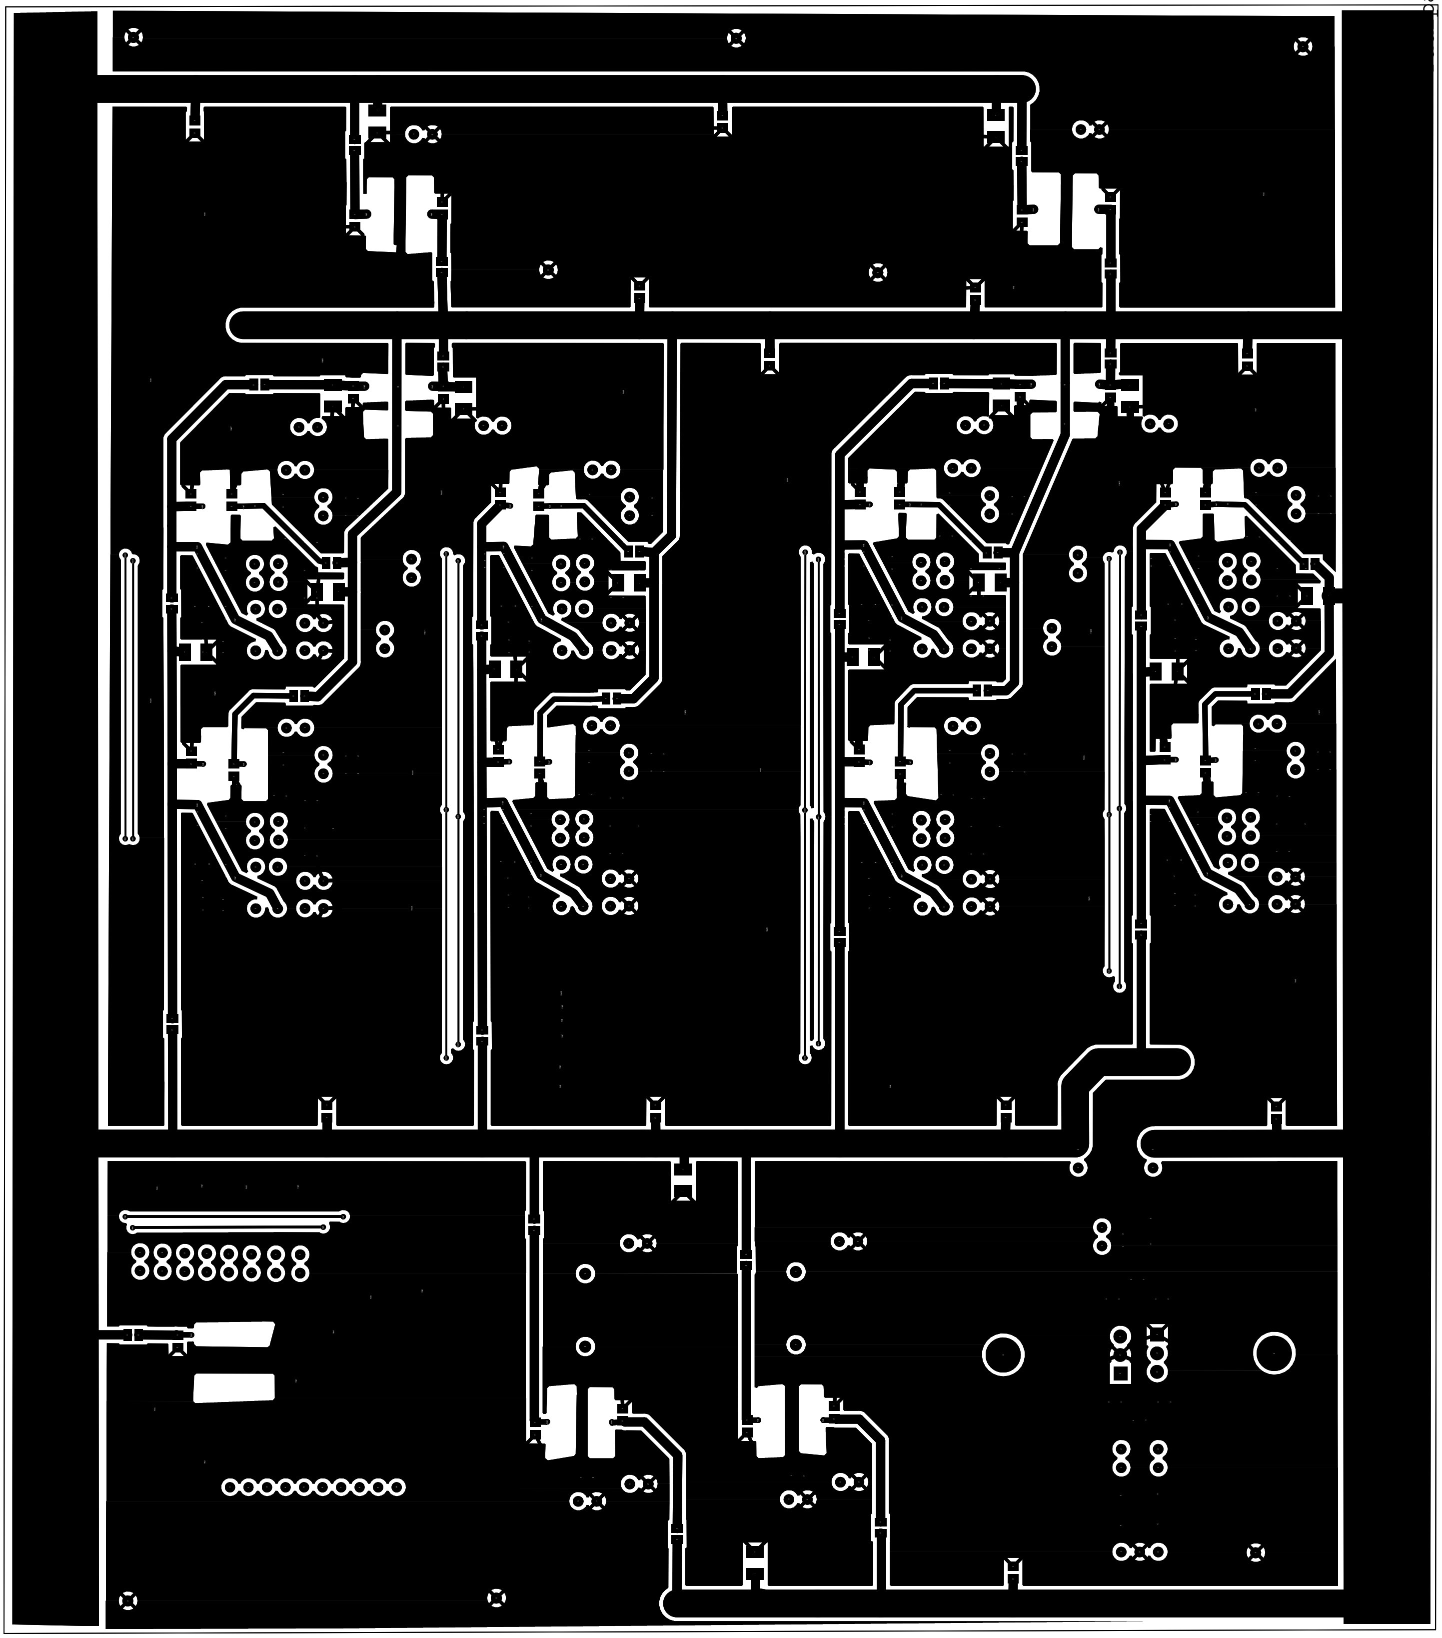

BOTTOM SILKSCREEN

Pad2pad.com

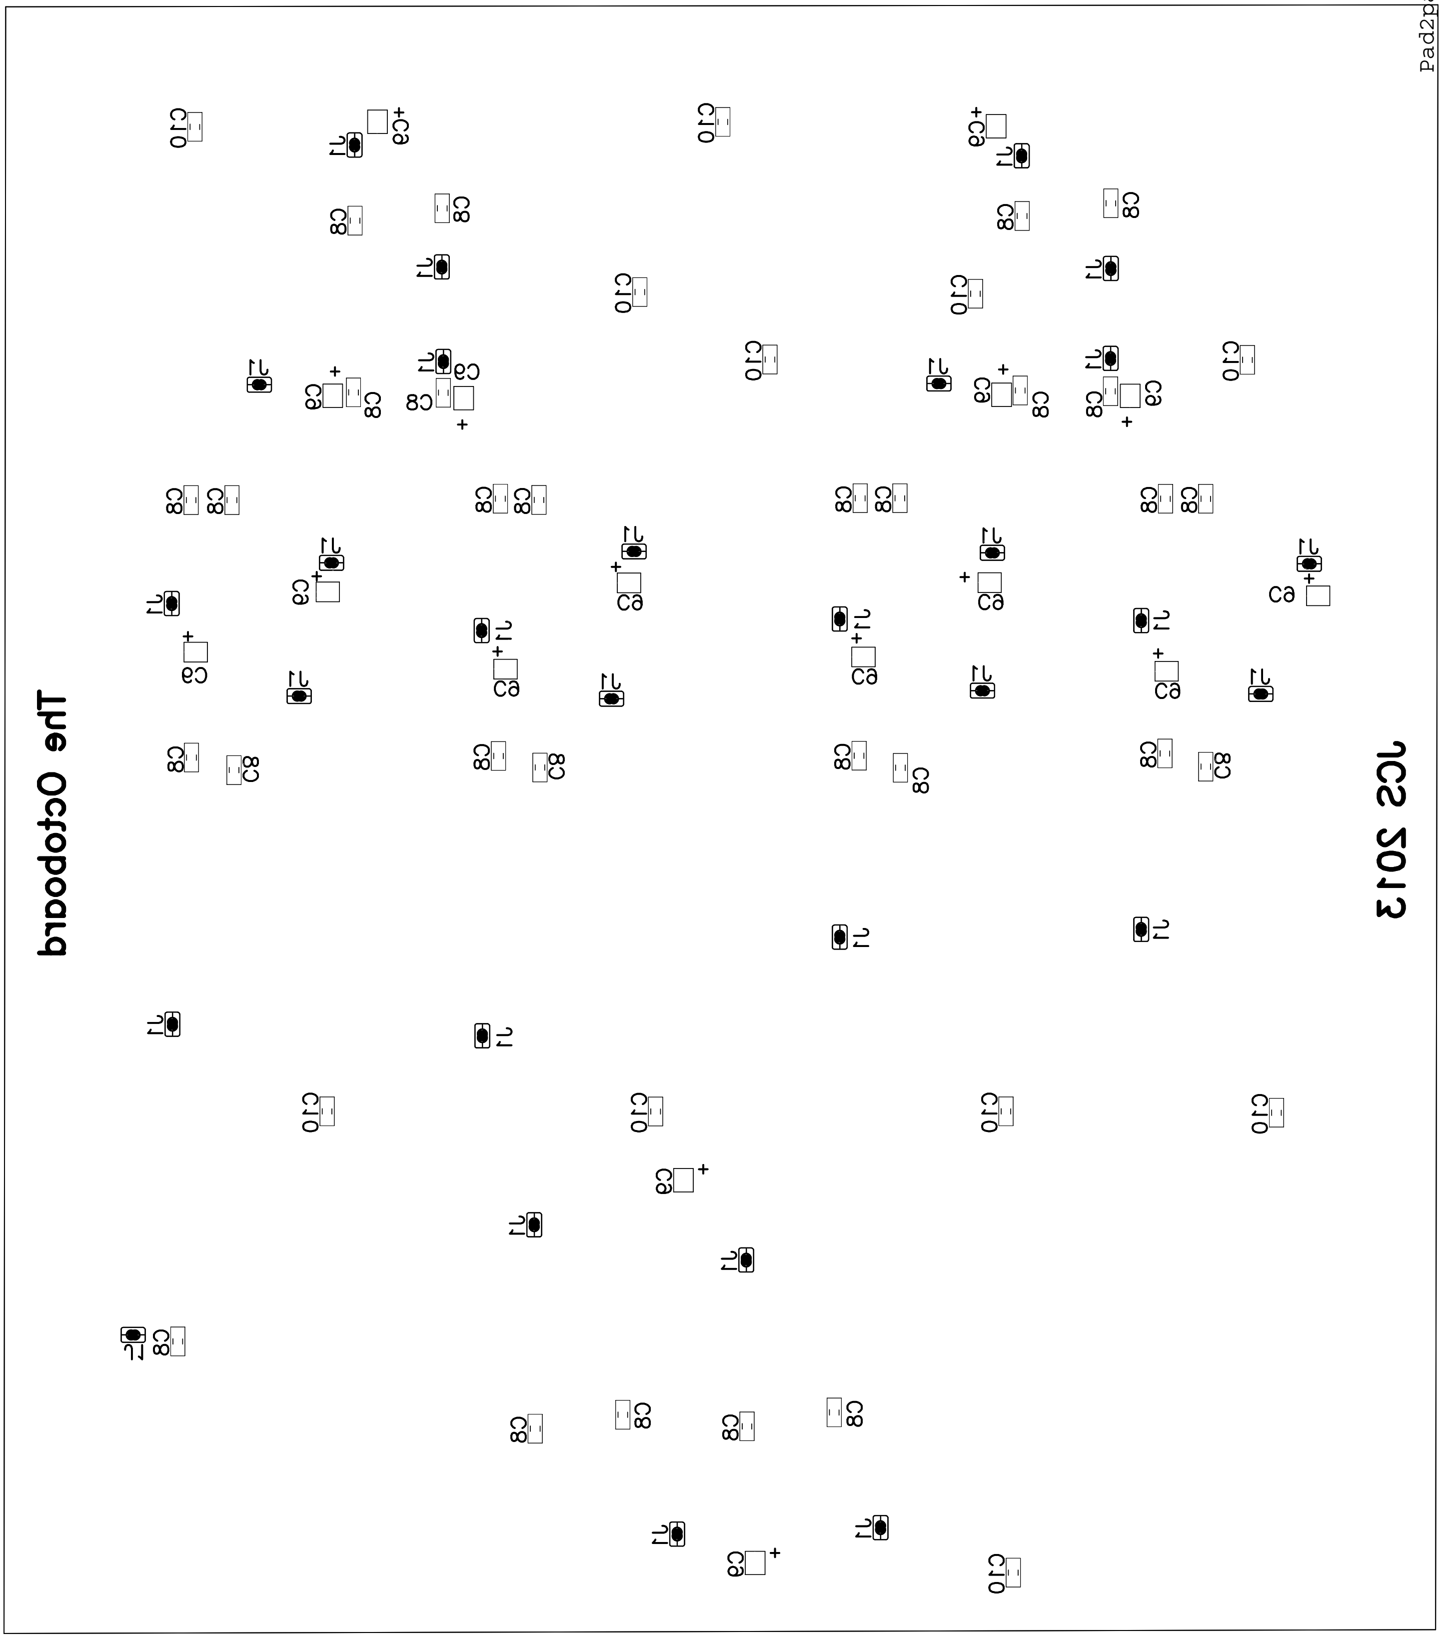

BOTTOM UNMASK

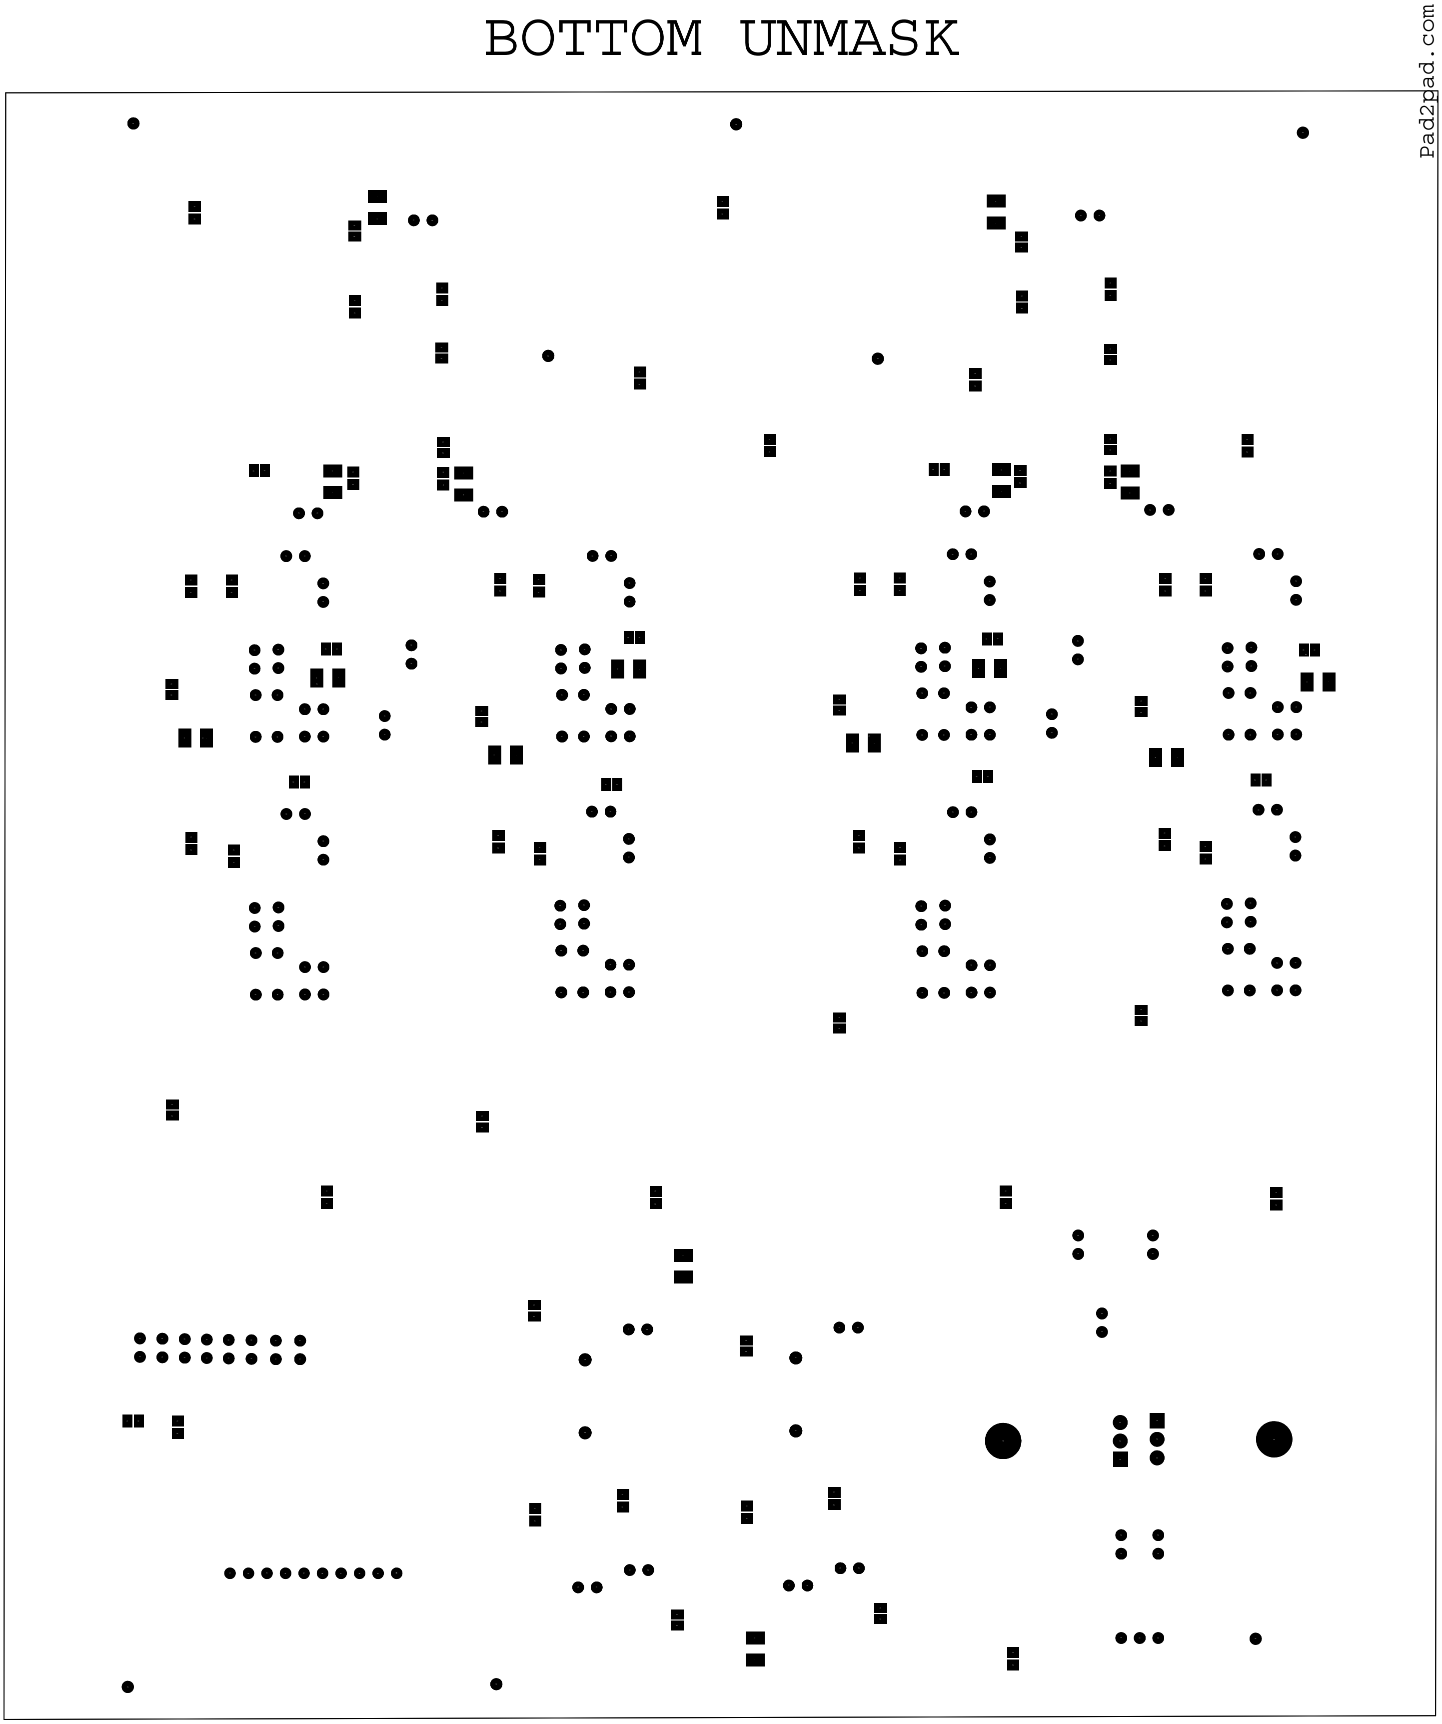

TOP + TOP SILKSCREEN

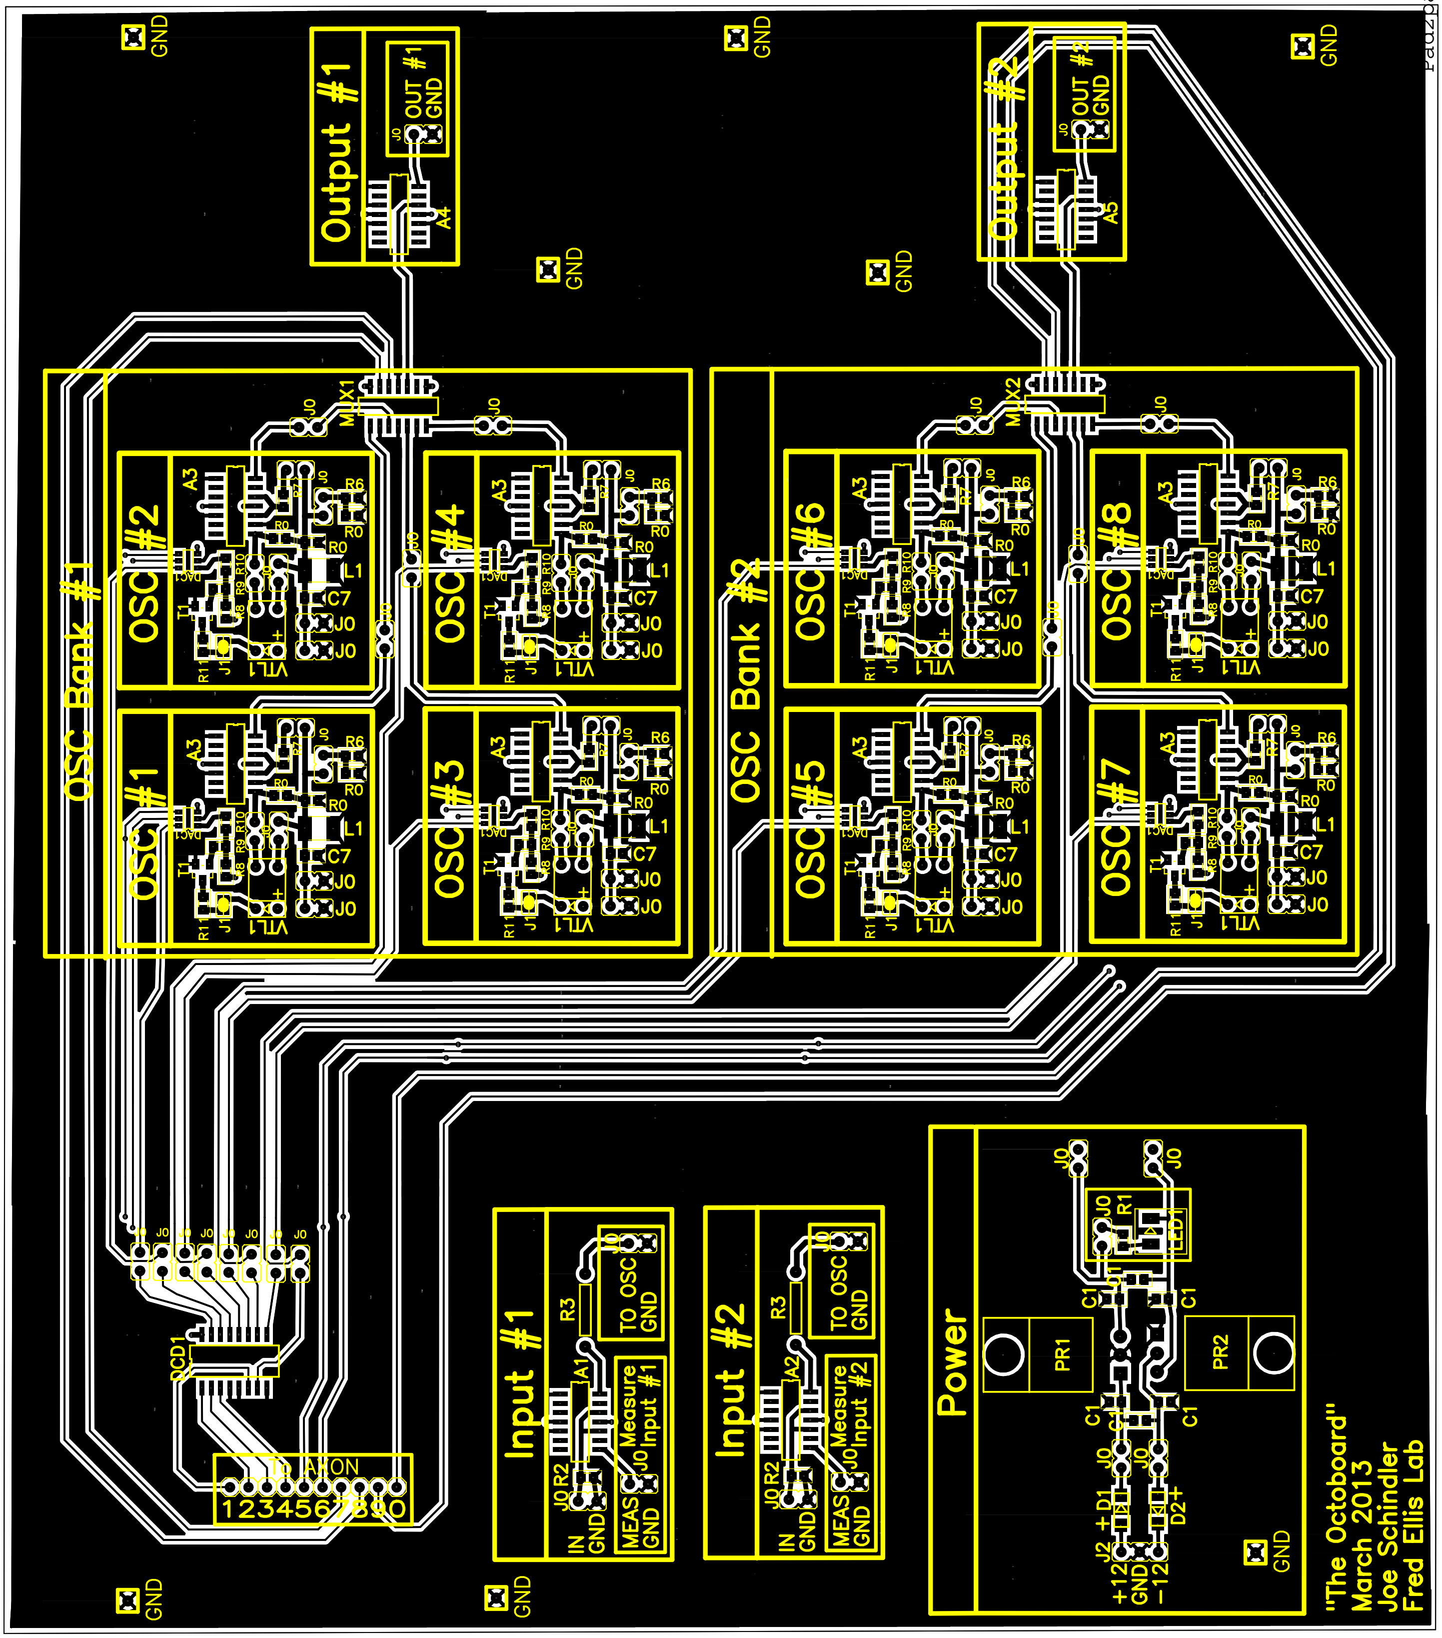

BOTTOM + TOP SILKSCREEN

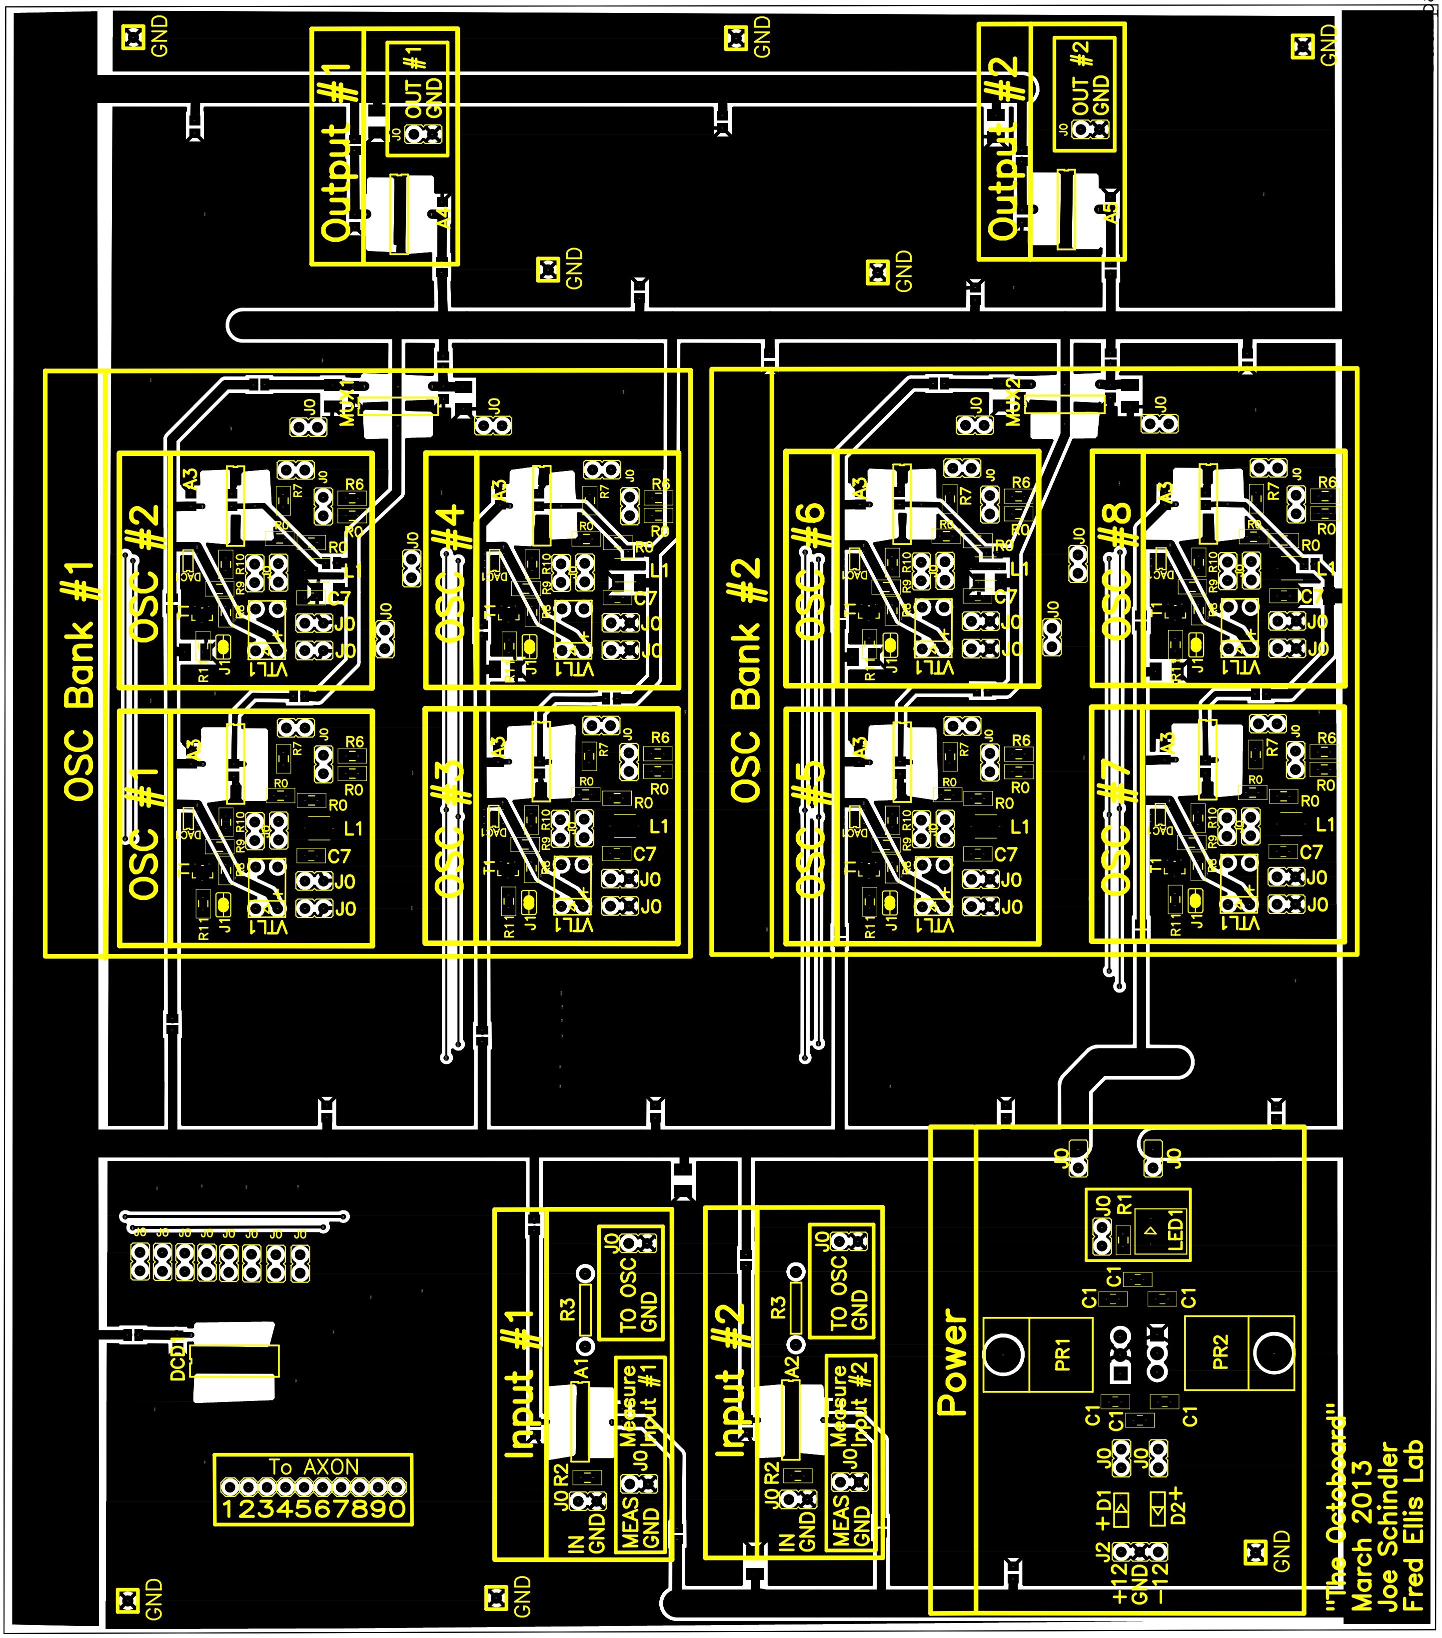

BOTTOM + BOTTOM SILKSCREEN

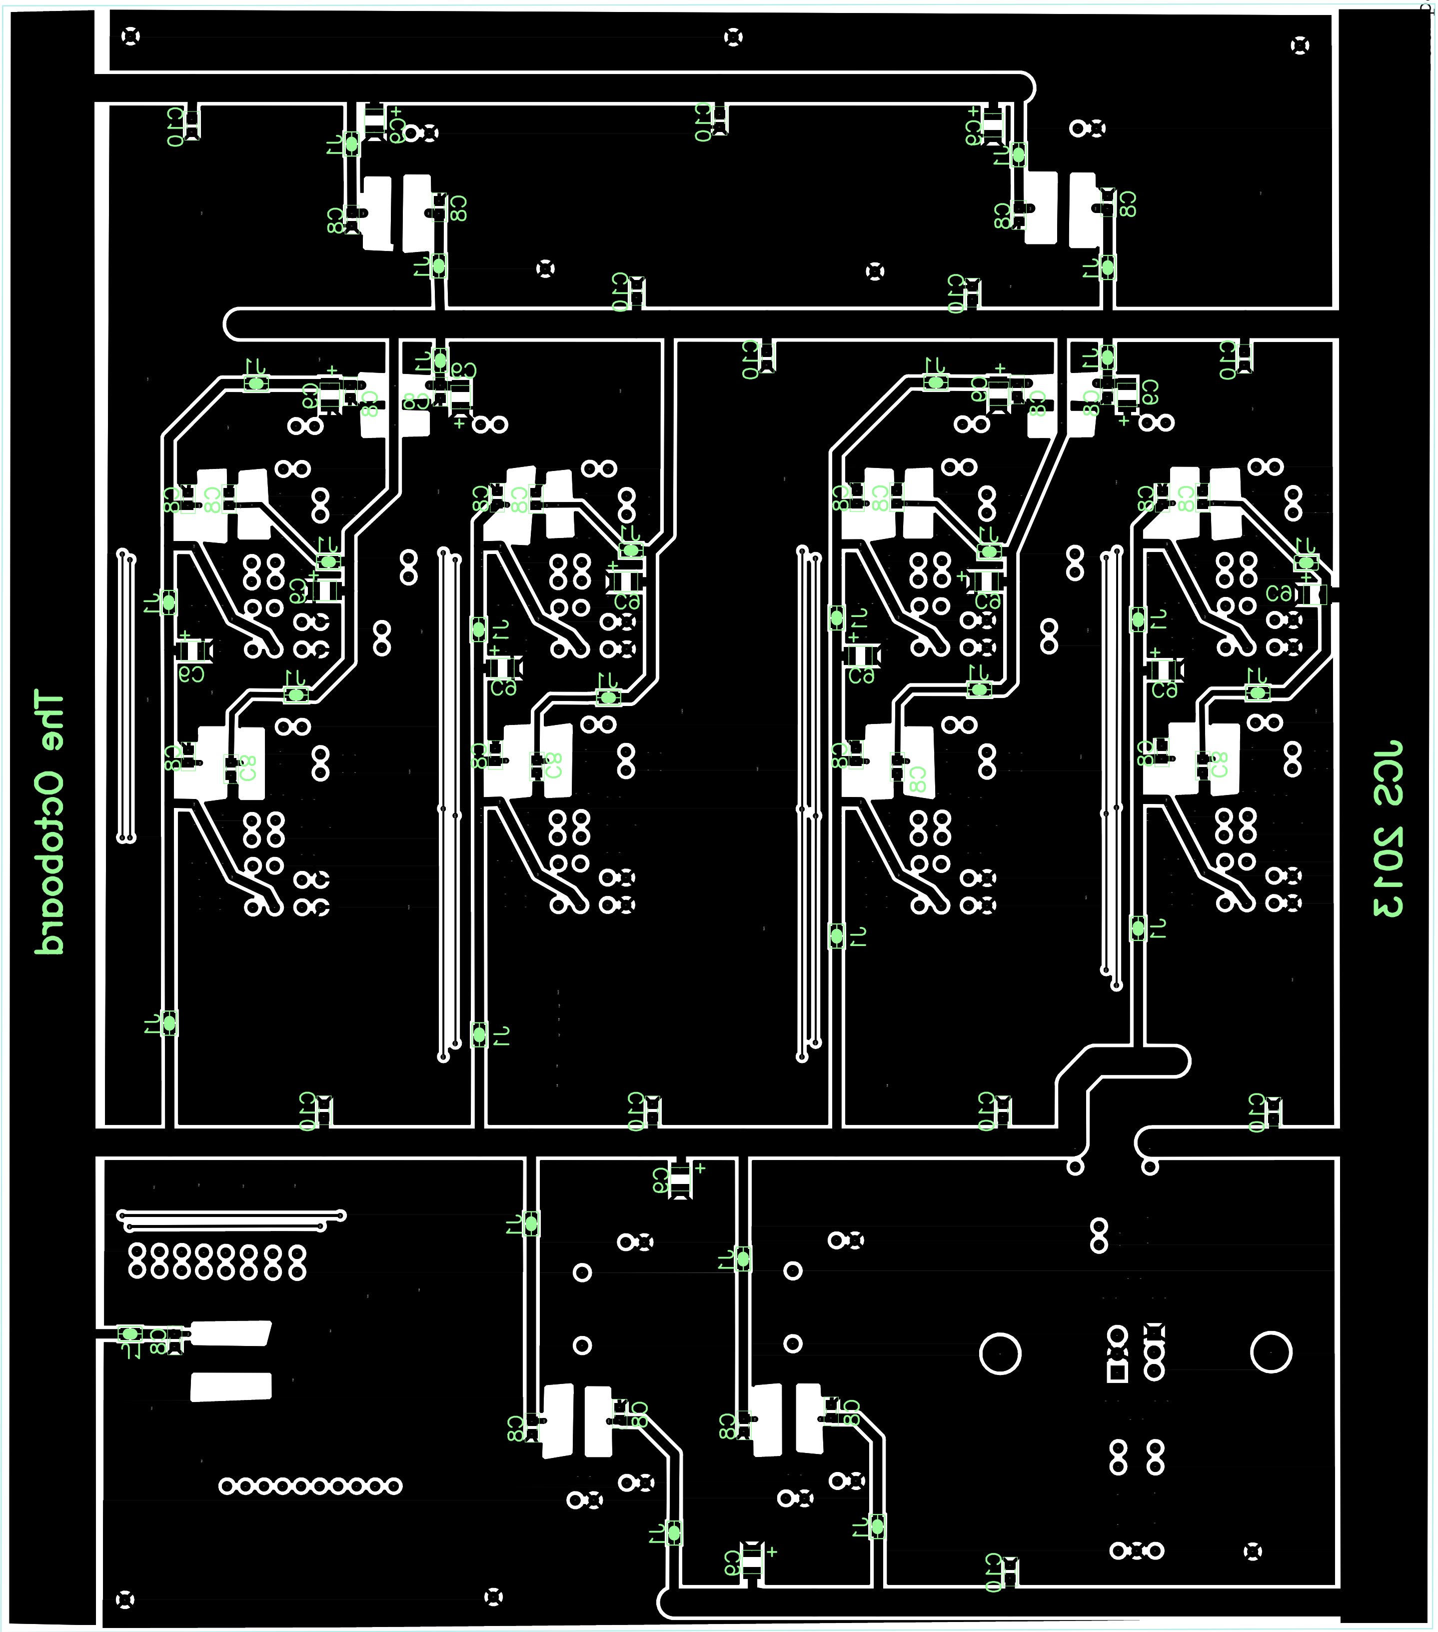

MAGNIFIED DRAWINGS BY FUNCTIONAL BLOCK

OSCILLATOR

TOP:

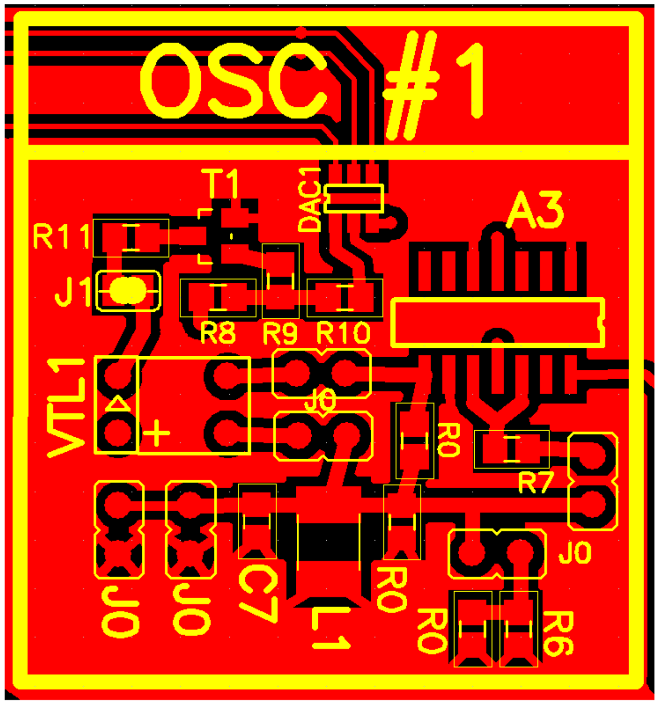

BOTTOM:

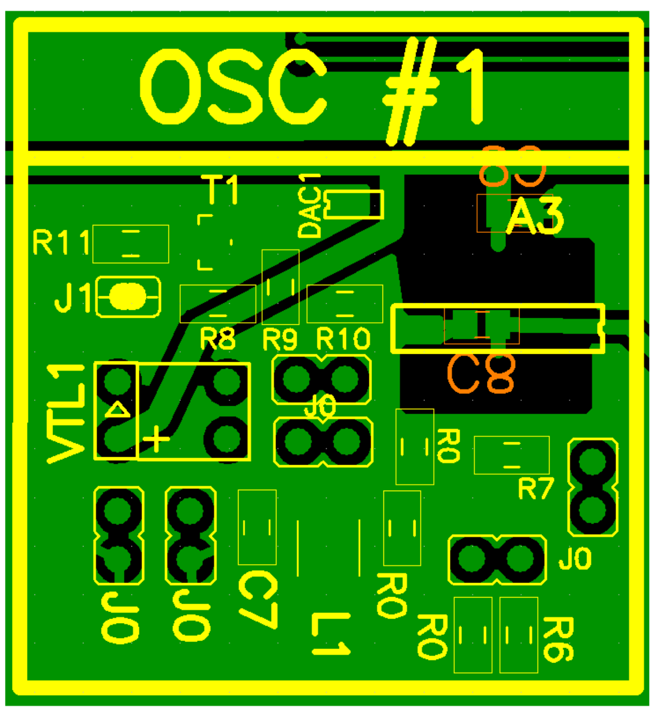

INPUT

TOP:

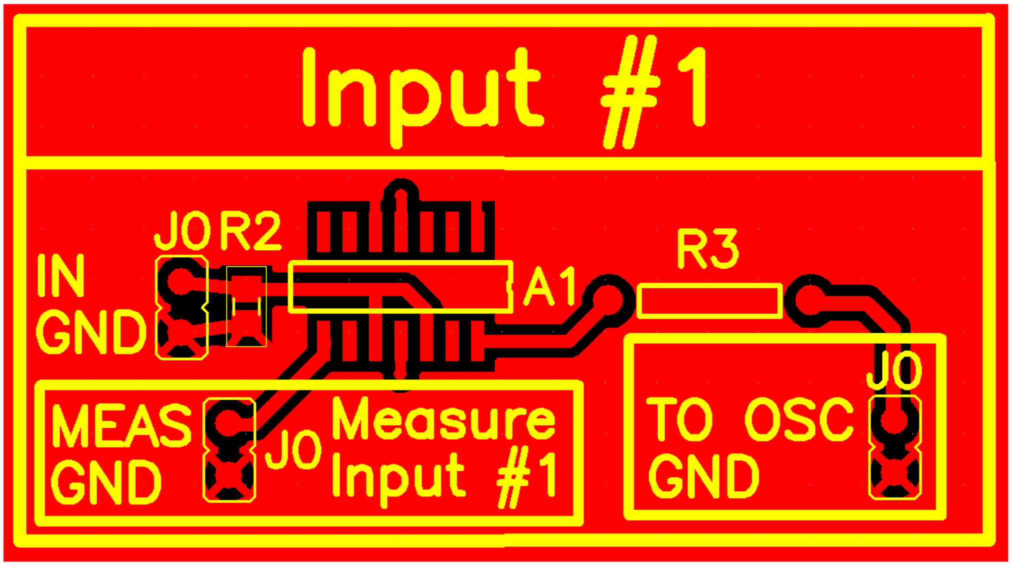

BOTTOM:

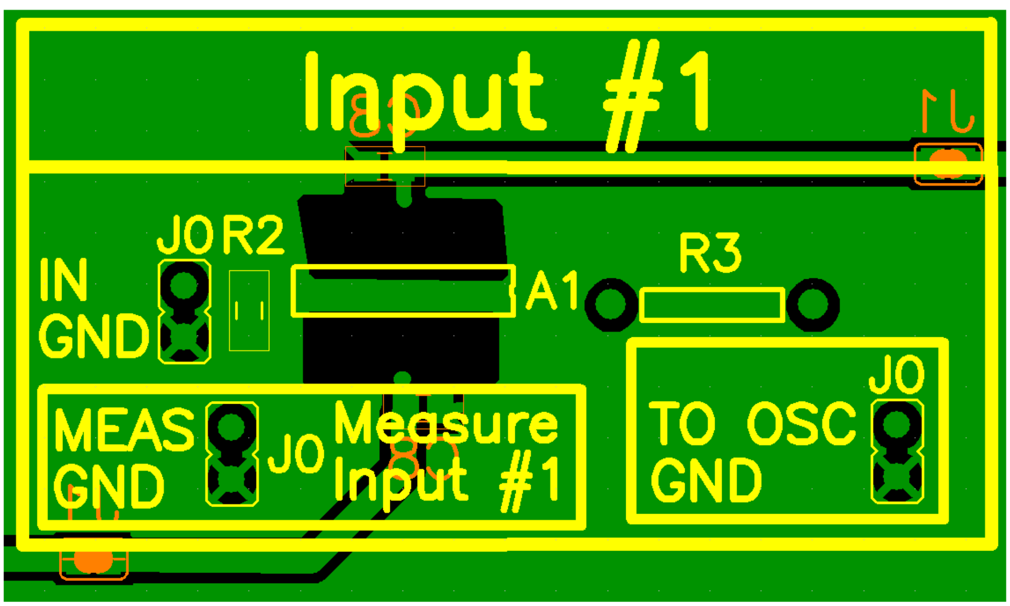

MULTIPLEXER

TOP:

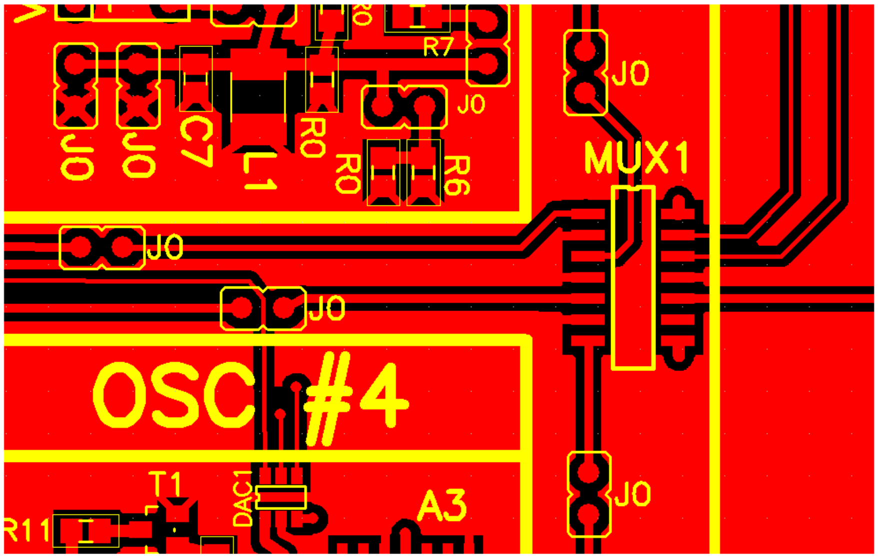

BOTTOM:

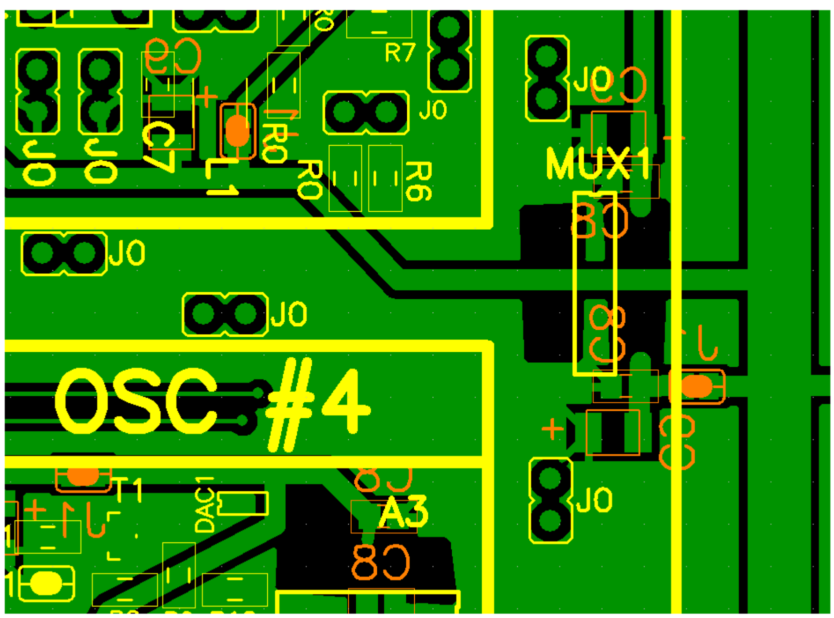

OUTPUT

TOP:

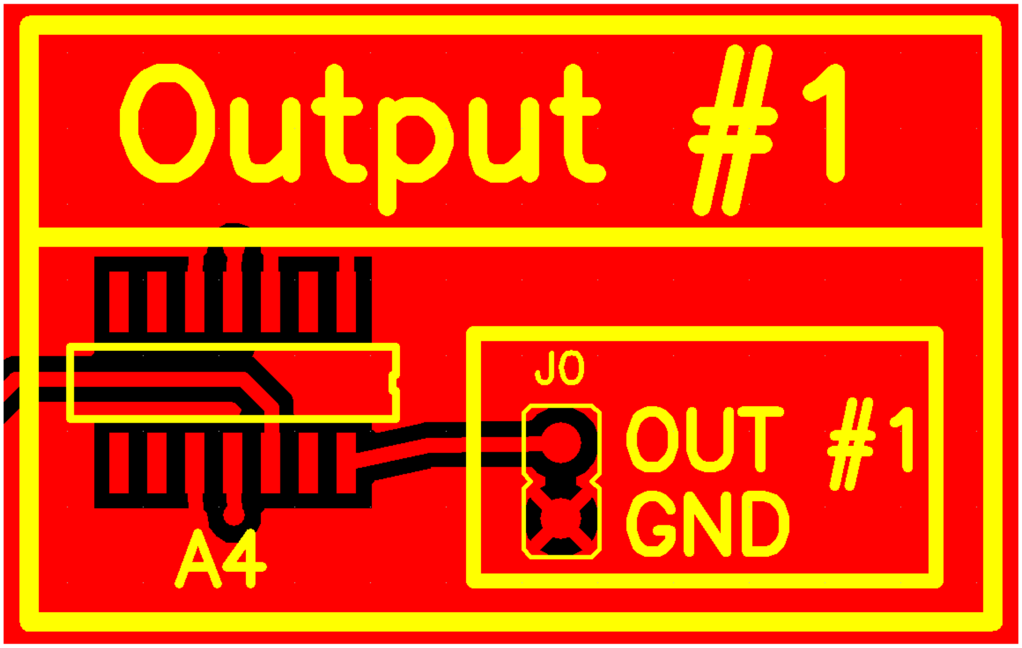

BOTTOM:

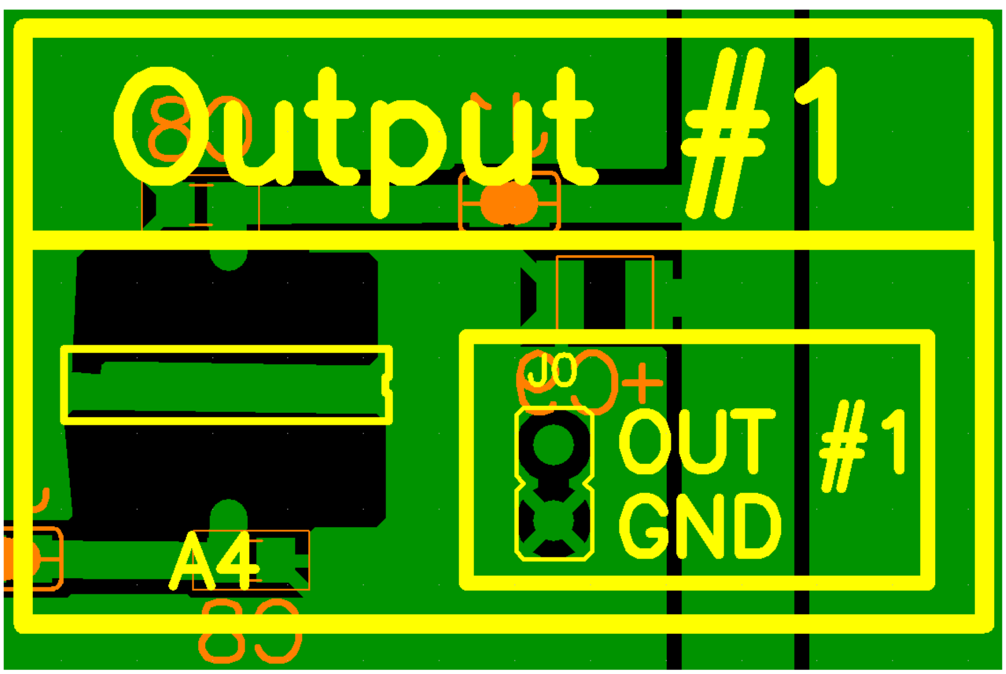

POWER

TOP:

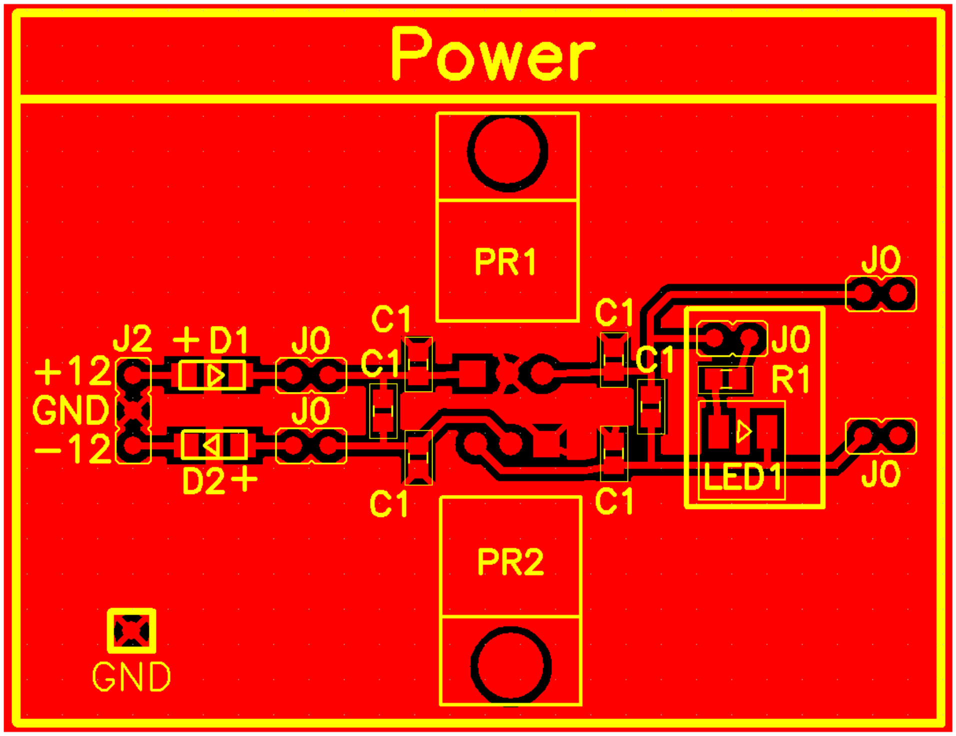

BOTTOM:

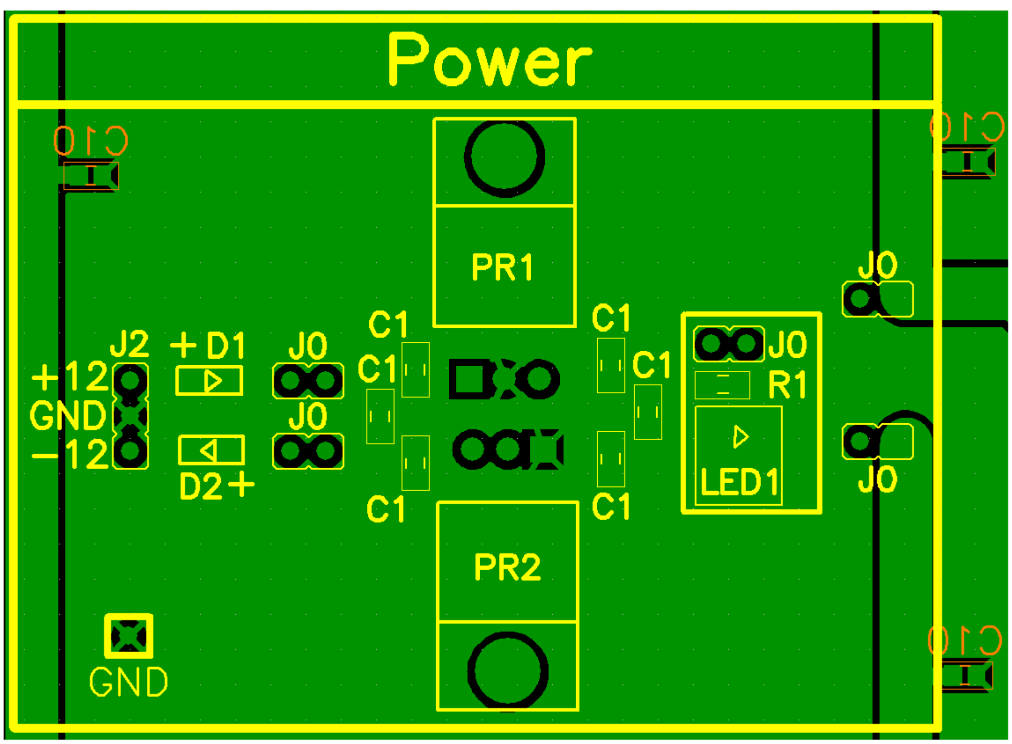

FROM AXON

TOP:

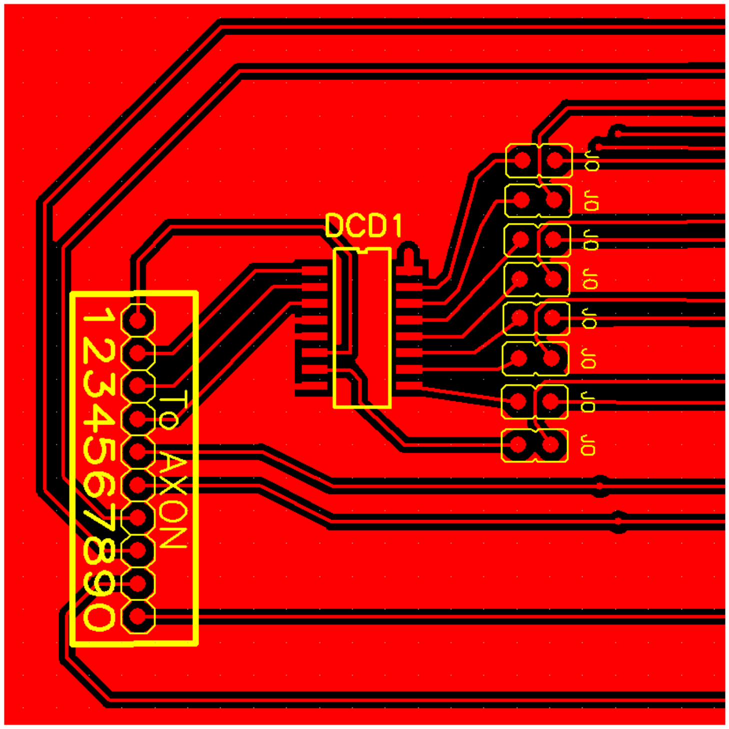

BOTTOM: