

Development of Electronics to Automate the Cavity Testing Process

Brad Schoenrock^{1,2}, Dave Miedlinger², Matthias Liepe²

¹*Department of Physics and Astronomy, Northern Michigan University, Marquette MI*

²*Cornell Laboratory for Accelerator based Science and Education (CLASSE), Cornell University, Ithaca NY*

5 August 2008

Testing the performance of Superconducting RF cavities is an important consideration in accelerator physics. The current set up at Cornell is somewhat old and can be arduous to generate a Q-curve with. This paper discusses some introductory mathematics necessary to the cavity testing process, the way the old setup is connected and used, the way the new setup is to be connected and used, the use of calibration software for the new equipment to give accurate results, and the use of a main program written in LabView to control and read the various devices.

I. INTRODUCTION

New accelerator facilities, such as the International Linear Collider or the Large Hadron Collider, are increasing the demand for high performance superconducting RF cavities. Ideas about new electronics and methods for measurements on the performance of Superconducting RF cavities have lead to hopes that the testing process could become automated. If this were to happen Q-curves could have more data points, which could allow for more accurate pictures of what happens to the Q values during multipacting, field emissions, ect... and may give insight into the benefits and drawbacks of various cavity preparation methods, designs, ect....

The background for testing Superconducting RF cavities is rather extensive. For a full background including cavity preparation, testing, limitations on performance, and more, I might suggest RF Superconductivity for Accelerators by Hasan Padamsee, Jens Knobloch, and Tom Hays. Another related paper I found helpful was “Basic Concepts of Measurements Made on Superconducting RF Cavities” by Jens Knobloch.

II. THE OLD SETUP & CAVITY TESTING BASICS

The whole idea of cavity testing is to measure P_e , P_i , P_r , P_t , and the frequency to obtain the intrinsic quality factor of the cavity Q_0 and the accelerating gradient E_{acc} . We obtain these quantities by a series of calculations. Equations 1 and 2 give two ways of calculating the coupling coefficient β from P_r , P_i , & P_e .

Equation 1

$$\beta = \frac{1 \pm \sqrt{\frac{P_r}{P_t}}}{1 \mp \sqrt{\frac{P_r}{P_t}}}$$

Equation 2

$$\beta = \frac{1}{2\sqrt{\frac{P_r}{P_s}} - 1}$$

With β and P_f we can calculate P_d from equation 3 which is needed to get Q_0 from equation 4.

Equation 3

$$P_d = \frac{4\beta P_f}{(1 + \beta)^2}$$

Equation 4

$$Q_0 = \frac{\omega U}{P_d}$$

E_{acc} comes naturally from a measurement of transmitted power, and the constant ($E_{acc}/\sqrt{P_t}$) found in field level calibration. (The equations come from RF Superconductivity for Accelerators by Padamsee, Knobloch, and Hays.)

To take these power measurements Cornell currently uses the setup shown in figure 1. Starting at the RF signal generator we have an RF wave that goes into an RF modulator which provides the forward power to the cavity. This wave then gets sent through a 57 dB amplifier to bring the forward wave up to power to excite the cavity. The forward power goes through a -30dB directional coupler to sample off a small fraction of the forward power for the measurement of P_f . The rest of the RF signal continues forward into a circulator that passes this wave on to the cavity. Some of this power goes into exciting the fields in the cavity, while some of it gets reflected off due to a potential barrier created by the cavity. This reflected wave gets split by another -30dB directional coupler. The part of the wave that was not taken by the directional coupler travels into the circulator and then gets dissipated by the load. The purpose of the circulator & load is to make sure that none of the reflected wave goes back into the main amplifier. The rest of the reflected wave (that was swept away by the -30 dB directional coupler) goes into two things. One is a power meter that can be connected to measure the reflected power, and the other is a crystal diode that can be employed to translate that RF wave into a voltage to be read on an oscilloscope which, based on the waveform, gives us P_r and P_e .

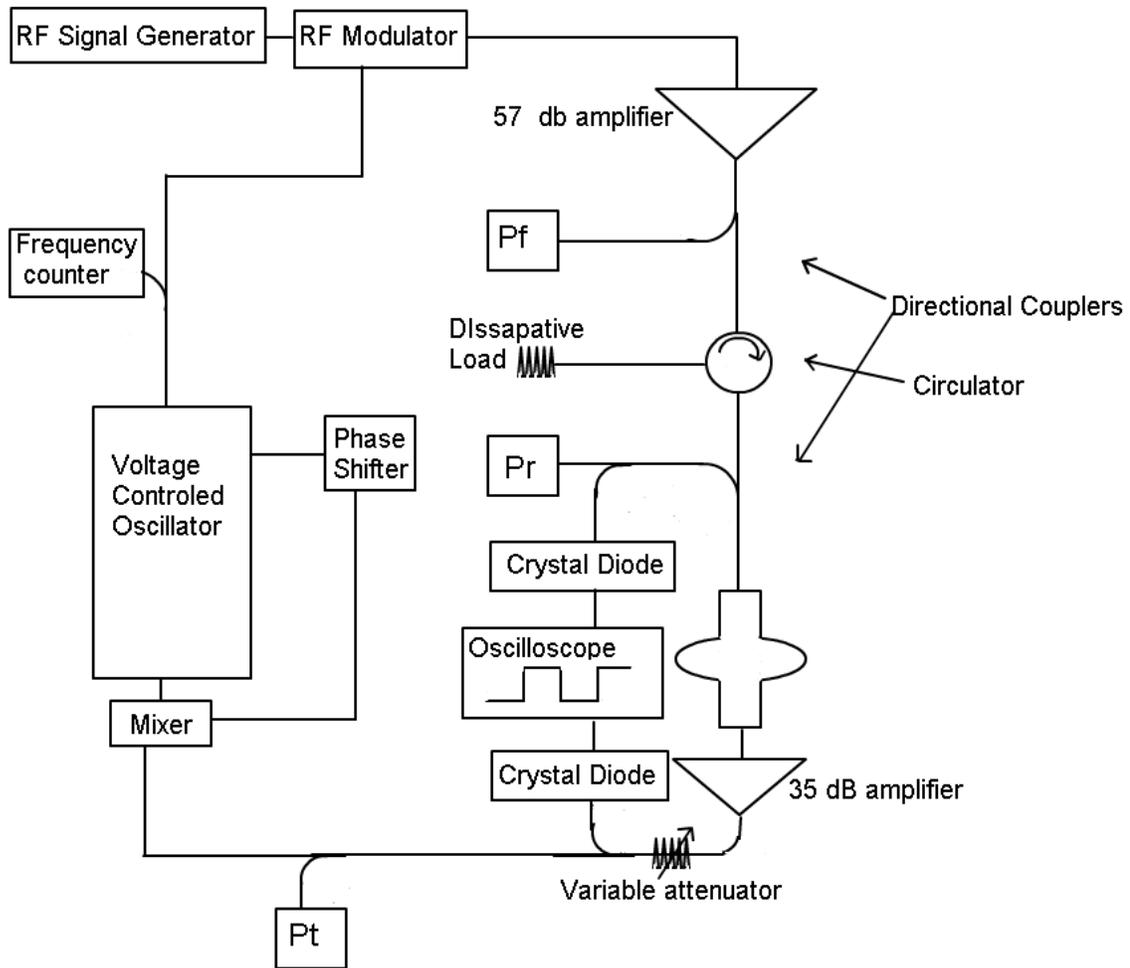


Figure 1: This block diagram is a schematic for the cavity testing circuit currently in use at Cornell University.

On the other side of the cavity there is a 35 dB amplifier followed by a variable attenuator to control the transmitted power levels. From here we employ another directional coupler to sample off the wave and run it to a crystal diode to be read on the oscilloscope which yields some more information on the powers coming from the cavity. Also from this transmitted wave a power meter is attached to measure P_t with which E_{acc} can be calculated. From here the transmitted wave enters the input of a mixer. The mixer gives a voltage out that is conveniently proportional to the product of the amplitudes of the input waves and the cosine of the phase difference between them. What this means is that the phase locked loop is dependent on the amplitude of the transmitted power, which can change. The voltage then goes in to the aptly named voltage controlled oscillator. The Voltage Controlled Oscillator outputs a RF wave whose frequency is shifted by the voltage input. That RF output goes into a trombone phase shifter (which is a miserably coarse adjustment) and back to the mixer as its second input. After the voltage controlled oscillator there is a frequency meter and then back to the RF modulator to complete the circuit.

Earlier I mentioned that the reflected power and the transmitted power were being sent through crystal diodes to an oscilloscope. The waveforms of these voltages can give us useful information as the forward power gets pulsed on and off. As the forward power goes from on to off, the reflected power spikes and the height of this voltage spike on the oscilloscope is proportional to the power emitted from the input coupler P_e , which then decays as the cavity energy is drawn off. When the power comes back on, the reflected power spikes again. The height of this spike is proportional to the principally reflected wave P_r . Now that we have P_e , P_r , P_f , and the frequency we have all we need to calculate Q_o .

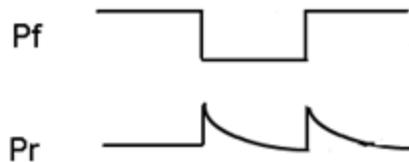


Figure 2: this diagram shows forward power and reflected power waveforms

One downside to this system is the instability of the phase locked loop caused by the mixer, which makes it hard to find resonance. Another is that all measurements have to be taken manually, which can become tedious if you are looking for a full Q-curve with many data points. The ability to get more data points in less time is another factor that would make an automated cavity test useful. Not having to take every data point by hand would mean that getting more points on a Q-curve would be fairly easy, and that could help lead to better understanding of the field levels that multipacting and field emissions occur at and the form that multipacting barriers have. These are the main reasons to develop a new system for cavity testing.

III. THE NEW SETUP

The new setup for the electronics works in basically the same way. The main differences lie in the phase locked loop, the amplitude modulator, and the log power detectors. Each of these devices is controlled/read by a computer.

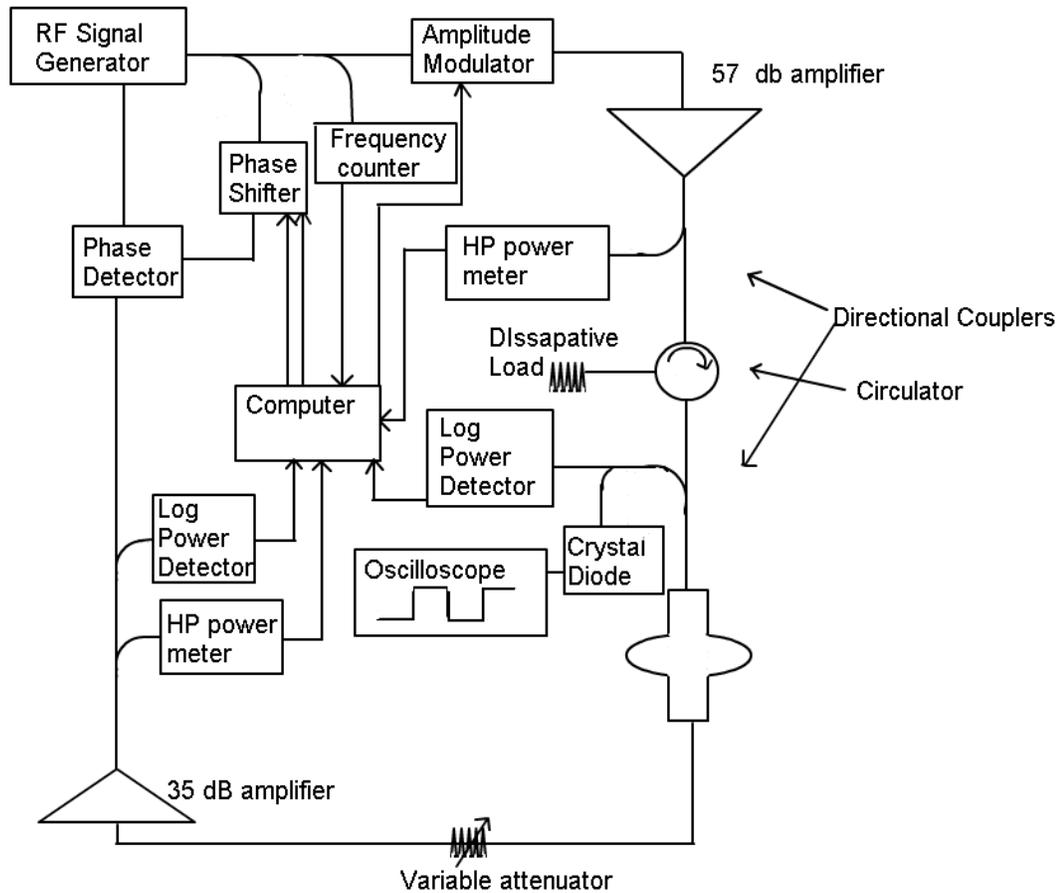


Figure 3: This shows the schematic for the new circuit that I was working on developing with the purpose of automating the cavity testing process.

The amplitude modulator is used to control the forward power to the cavity. The benefits of this are that we have the control to “pulse” the cavity. By turning off the forward power for a short time, we can measure both P_e and P_r by measuring the peak voltage on the oscilloscope corresponding to P_e , and the steady state voltage before the pulse which corresponds to P_r .

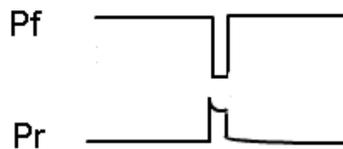


Figure 4: The forward and reflected power using the “pulsed” method

Another change is the replacement of HP power meters with log power detectors to make measurements on the various powers necessary for computing Q_0 and E_{acc} . The

faster response time of the log power detectors make them the natural choice for taking these measurements because it allows you to take rapid measurements during the pulses.

The final change is the replacement of the old phase locked loop with newer electronic devices. The phase shifter takes two voltage inputs, I&Q, and an RF reference wave, and gives an RF reference wave out with a phase difference given by $\arctan(Q/I)$ and an amplitude proportional to $\sqrt{(Q^2+I^2)}$. This gives one input to the phase detector, while the other comes from the cavities transmitted power. The nice thing about the phase detector is that the voltage output is proportional only to the phase difference between its two inputs, and is independent of amplitude of its RF input. In this way the phase locked loop remains unphased by changes in the amplitude of the transmitted power.

The main benefit of this system, however, is that all of these devices can be hooked up to a computer to read the various necessary values, and calculate the Q_0 vs. E_{acc} points more rapidly.

IV. CALIBRATION SOFTWARE

In the LabView setting several VI's were written to calibrate the devices used in the new cavity testing circuit. Four of the calibration files can be run without the cavity in the setup, while two of them are done while the cavity is cold and in place.

The first step is to hook up the RF signal generator directly to the HP power meter. Make sure the power meter is reading in mW not dBm, this is the units that all inputs should be in for calibrations. Once this is done run the VI titled "1 calibrate RF signal generator.vi" and save the output in the folder called "calibration files (spreadsheets)." This is where all calibration files are saved in this process, with each type of calibration file having its own folder within this main folder. Note that the range (-25 dBm to 5 dBm) has been chosen because of the limitations of the devices. The HP power meter has a hard time reading anything lower than -25 dBm while the signal generator is limited to only go as high as 5 dBm so as not to over work itself.

Now that this is complete you can hook up the signal generator to the log power detector to be calibrated. When you run the program "2 calibrate Power Detector.vi" you will be prompted to load the calibration of the signal generator. The calibration of the signal generator now dictates the calibrated range of the log power detector because in order to properly calibrate the detector you need to know the power of the RF input and we cannot trust the display on the front of the signal generator to be correct (sometimes being off by as much as 4 dB). The log power detectors have a listed dynamic range of -60 dBm to 0 dBm, but my calibration only supports 25 dB of that range. The good news is that once a time delay was added into the VI this made the calibration curve quite linear and this might allow one to extrapolate to lower power levels. This program does not currently do this (because the nice linearity was only recently noticed) but it should be a natural extension of the program.

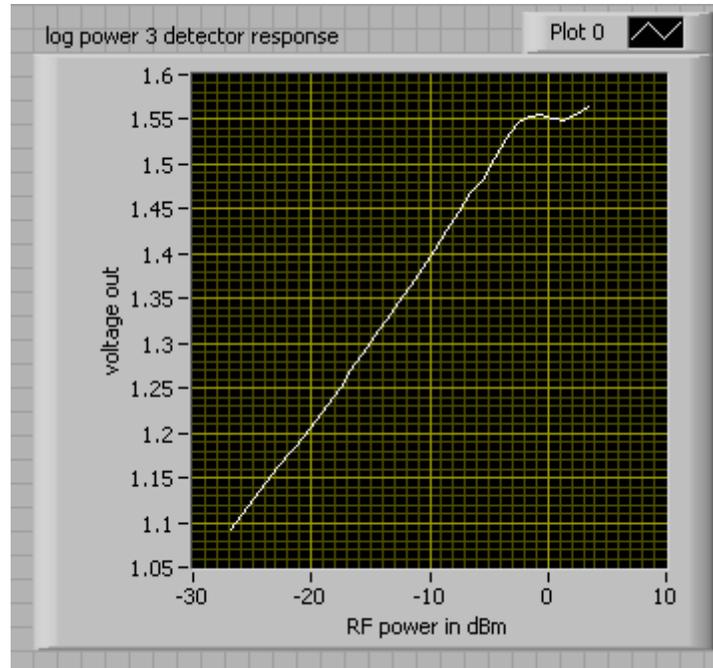


Figure 5: RF power of known level is the input of the log power detector, the voltage out is measured so that the main program may interpolate the data points and infer the power in to the device. Note the saturation around 0 dBm.

There is another VI for measuring the sensitivity of the detectors. It works on the assumption that the RF signal generator knows what it is outputting, but other than that is pretty interesting to see the lower limits of the device. This may prove useful in interpolating to lower powers.

Now that you have calibrated each of your log power detectors, the next thing to do is calibrate the amplitude modulator. Hook up the HP power meter to the output of the amplitude modulator, the RF signal generator to a -3 dB splitter (which is hooked up to the amplitude modulator and the phase shifter simultaneously) and run this program. Save its file to the appropriate folder.

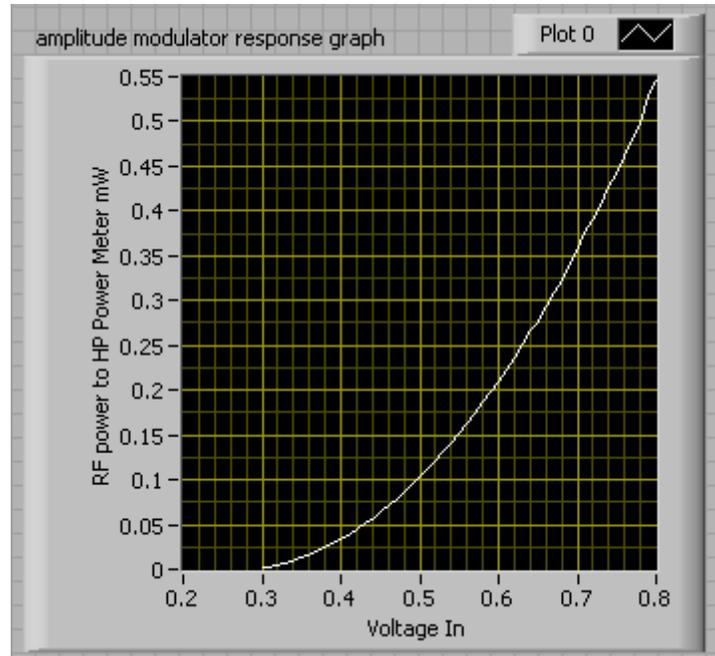


Figure 6: The amplitude modulators input I is varied through voltages that are in the “nice” range. Above about 0.8V the graph becomes very steep and difficult to control (making about .5mW its upper limit for power output). Below about 0.3V the graph turns back upward which can pose a problem for the program.

The phase shifter is the next thing to be calibrated. Open the VI “3 phase detector finding center.vi”, make sure the HP power meter is hooked up to the output of the phase shifter, run this VI, and save the output to the center of phase shifter folder. This VI simply provides corrections for the voltage offsets of the I & Q inputs to the phase shifter.

The next calibration is an optional addition. “4 calibrate phase detector.vi” is a program that allows you to determine if the phase detector is behaving as you would expect. If you are running the main panel for a cavity test then this is not strictly necessary as you don’t need to know the phase difference of the two waves going into the phase detector. If you would like to check that the phase detector is working, go into the examples folder and use “control phase detector.vi”. This VI uses the two calibrations noted above to measure the phase shift of a wave into one side with the phase unaltered, and the other into the phase shifter to get a specified phase shift. This may sound redundant at first glance, but that’s because it is redundant to measure the phase shift of a wave that you specified the phase shift of. Like I said it’s purely to check to see that the electronics are working.

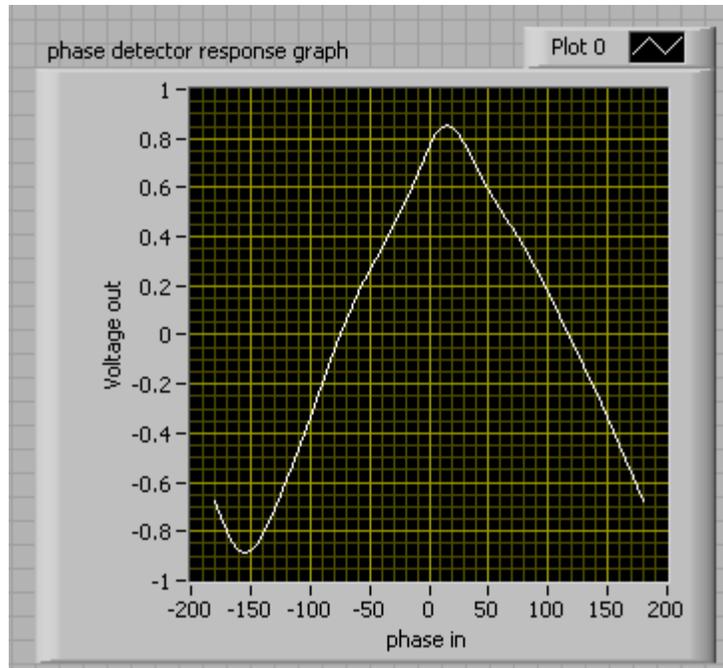


Figure 7: The phase detector is checking to make sure the phase shifter is doing its job properly. The range in use if you would like to run “control phase detector.vi” is the region between the minimum and maximum of the graph. Anything outside this range is associated to the equivalent point inside the range meaning that you only need 180 degrees of phase shift to give you the whole range for phase lock. Note that it is perfectly normal to have this graph shifted in the x-direction due to different line lengths.

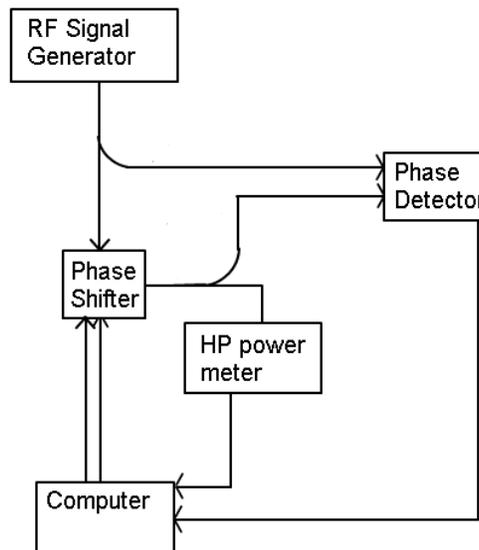


Figure 8: A diagram for how to hook up the phase detector to use the VI control phase detector.vi

The other two calibrations need to be performed when the cavity is hooked up to the system and cold. The first is “5 Calibrate Amplitude Modulator voltage sweep vs. HP power meter reading.vi” and the second is “6 calibrate amplitude modulator vs. Power Detector Voltage.vi” these allow us to convert the reading on the HP forward power meter to the voltage the log power detector puts out for an identical forward power setting on the amplitude modulator. The calibration of the log power detector can then be used to find the associated power with that voltage, and the ratio of P_r to P_f can be calculated without having to worry about line losses to/from the devices. This is one of the main advantages of this method in that line losses need not be considered between the forward power sensor and the reflected power sensor. Unfortunately I have no successful pictures of this calibration curve because the cavity that this program was tested on was too weakly coupled to get reasonable results to show from this calibration.

These two programs can (and probably should at some point) be merged into one program that takes both readings at the same time. This will take nominal but not completely insignificant effort to incorporate in to the main program, the reason that was not done previously is because the other programs were already written and could be coalesced in the main program to give the information that was needed.

V. THE MAIN PANEL

The first thing you might notice when you try to run “main panel” is that you are prompted for calibration files. It prompts you for several files to load, and tells you the order you need to load them as they come up. Also important of course is to make sure every device is hooked up properly. If they aren’t hooked up properly the program will gripe at you because it can’t find the droids it’s looking for.

From a programming standpoint all my calibration files are loaded and then broken up into 1-D arrays which are passed into the main event structure where the work is done. Each event has comments in it as to how it works so if you are interested in that you will need to obtain a copy of the program.

The upper half of the front panel is dedicated to the calibration files for the devices discussed in the previous section. As you load the calibrations you can check to make sure they look like what you would expect.

There are multiple controls around the front panel to control the channels to which each of the devices are connected. This way if you have transmitted power being read by a different HP meter (GPIB0::12::INTR rather than GPIB0::13::INTR) then you can easily change them around from the comfort of your seat. It was originally intended that log power detector 2 be hooked up for transmitted power, log power detector 3 be hooked up for reflected power, HP meter with the address GPIB0::13::INTR be hooked up to the forward power, and the HP meter with the address GPIB0::12::INTR be hooked up to measure the transmitted power (which is not necessary if you are measuring the transmitted power with the log power detector, but the program is set up to do so anyway).

There are also controls for the signal generator, the phase shifter, and the amplitude modulator. The first two can be adjusted to find phase lock and the control for the amplitude modulator can be used to set the forward power to the cavity.

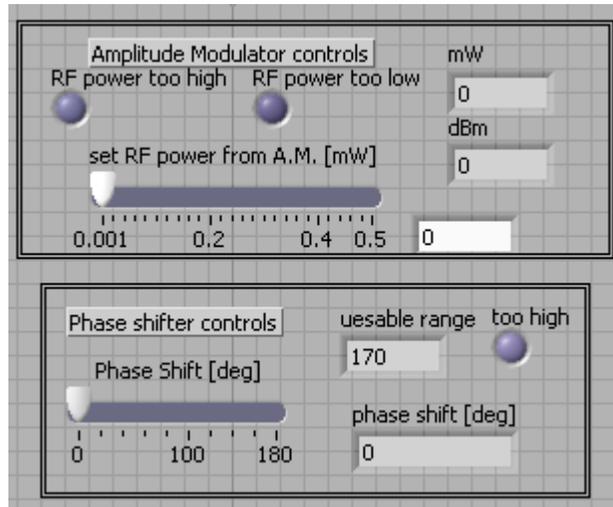


Figure 9: Controls for the amplitude modulator and the phase shifter/detector

There are readouts for single readings from each of the log power detectors. In this way you can manually take readings of reflected power and transmitted power if you so choose.

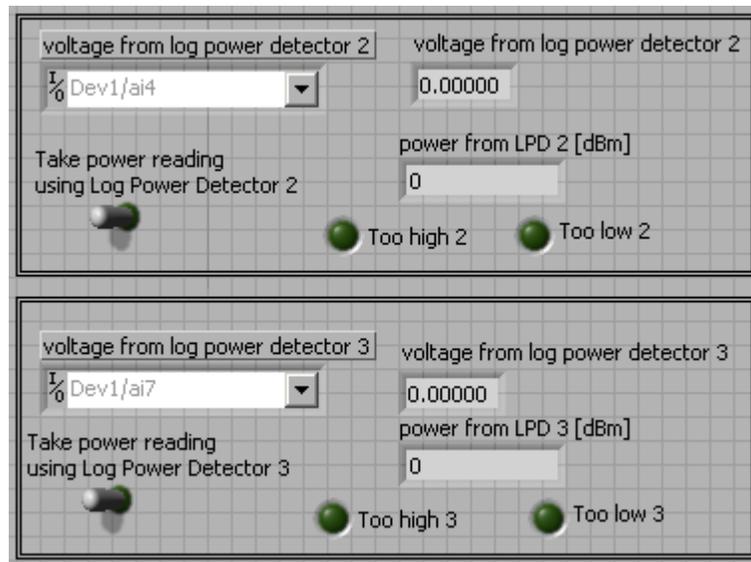


Figure 10: Readouts for the log power detectors. These are single shot readings, not part of the automated pulsing process.

There is also a section devoted to pulsing the cavity. In this section you can dictate how much time to keep the power on for before pulsing the cavity (to make sure the

fields have reached a steady state value before pulsing), and dictate how long the pulse length will be. Once the power is switched the program takes its readings which are displayed as a grouping of indicators on the various power levels.

If the values for all power readings are acceptable (and appropriate calibration coefficients are in place) then push the Perform Field Level Calibration button to do the calculations. This will update U/Pt and allow you to calculate Qo values. The diagram that explains what the coefficients stand for is to the left of the main view.

Inside the field level calibration section there is a box which can override the readings from the power detectors. The main purpose of this is to convince yourself that the calculations work. If this is still not acceptable, then you can override the field level calibration altogether and input your own numbers for U/Pt. If you do not input any numbers, then these fields serve only as indicators of the calculations performed. Whatever values are in those fields are the ones used to calculate Qo values in the next section.

Calculate Qo point does about what the name suggests. If you flip the switch it grabs the values from the pulse cavity section and calculates a Qo point. The array "Qo data" contains information such as the coupling coefficient β , Qo ect.... Note that this does not pulse the cavity to update the numbers. As it stands the pulse cavity section and calculate Qo point function separately. If you scroll to the right in the program you will see another set of Qo data. The only difference is that this is calculated from the HP power meter reading of transmitted power, whereas the array that is with the rest of the program is calculated using transmitted power read from the log power detector.

By far and away the most important part of the program is the switch with 4 lights around it. This switch was originally used to tell if the program had died or not, but now remains due to its contribution of style points.

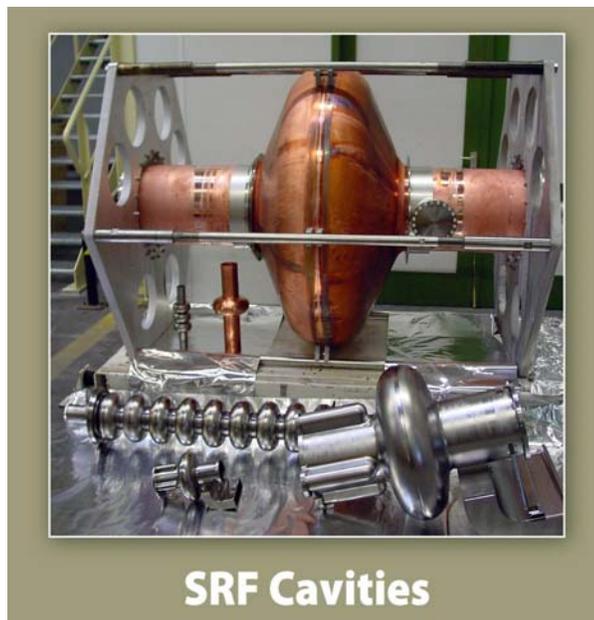
VI. RESULTS

A cavity test was performed (on 8/3/09) with the new electronics setup. In this cavity test the new phase locked loop with the electronic phase detector proved its stability. This stability made it easier to find and maintain phase lock. Also demonstrated was that the log power detectors can give repeatable and accurate results, matching up with the readings from the HP power meters nearly perfectly. Unfortunately no Q-curve could be generated due to severe undercoupling that prohibited the measurement of a decay curve for field level calibrations. Also interesting to note is that the pulse method worked well, despite the coupling problems that were encountered. All these results lend themselves to optimism that the cavity testing process can be automated to provide further understanding of Superconducting RF technology.

VII. ACKNOWLEDGEMENTS

I am truly thankful for the help from my advisors Dave Meidlinger and Matthias Liepe, other faculty Zack Conway and Georg Hoffstaetter, The entire SRF group, the REU project coordinators Rich Galik, Ernie Fontes, Lora Hine, and Monica Wesley, and every faculty who took on a student this summer. Without these hard working people at Cornell this Research Experience for Undergraduates would not happen. This has truly been a rewarding experience that I hope will continue for many years to come.

This work was supported by the National Science Foundation REU grant PHY-08498865.



The pursuit of your dreams is its own reward.