ACOUSTIC LOCALIZATION OF BREAKDOWN IN RADIO FREQUENCY ACCELERATING CAVITIES

BY

PETER LANE

Submitted in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy in Physics
in the Graduate College of the
Illinois Institute of Technology

Approved _____
Advisor

Chicago, Illinois
July 2016

ACKNOWLEDGMENT

TABLE OF CONTENTS

LIST OF TABLES

vi

LIST OF FIGURES

LIST OF SYMBOLS

Symbol	Definition
$\epsilon_\perp^N$	Normalized transverse emittance
$\beta_\perp$	Transverse Betatron function
s	Particle path length
β	Particle speed fraction of the speed of light ($\frac{v}{c}$)
E_μ	Energy of a muon
m_μ	Mass of a muon
X_0	Radiation length
$\mathbf{E}$	Electric field vector
$\mathbf{B}$	Magnetic field vector
$\mathbf{j}$	Current vector
ρ	Radial cylindrical coordinate
ϕ	Azimuthal cylindrical coordinate
z	Longitudinal cylindrical coordinate
t	Time coordinate
ω	Angular frequency
i	The imaginary unit equal to the square root of -1
c	Speed of light
∇s	Gradient of scalar function s
$\nabla \cdot \mathbf{v}$	Divergence of vector $\mathbf{v}$
$\nabla \times \mathbf{v}$	Curl of vector $\mathbf{v}$
$\nabla^2 s$	Laplacian of scalar function s

ABSTRACT

Current designs for muon accelerators require high-gradient radio frequency (RF) cavities to be placed in solenoidal magnetic fields. These fields help contain and efficiently reduce the phase space volume of source muons in order to create a usable muon beam for collider and neutrino experiments. In this context and in general, the use of RF cavities in strong magnetic fields has its challenges. It has been found that placing normal conducting RF cavities in strong magnetic fields reduces the threshold at which RF cavity breakdown occurs. To aid the effort to study RF cavity breakdown in magnetic fields, it would be helpful to have a diagnostic tool which can localize the source of breakdown sparks inside the cavity. These sparks generate thermal shocks to small regions of the inner cavity wall that can be detected and localized using microphones attached to the outer cavity surface. Details on RF cavity sound sources as well as the hardware, software, and algorithms used to localize the source of sound emitted from breakdown thermal shocks are presented. In addition, results from simulations and experiments on three RF cavities, namely the Aluminum Mock Cavity, the High-Pressure Cavity, and the Modular Cavity, are also given. These results demonstrate the validity and effectiveness of the described technique for acoustic localization of breakdown.

CHAPTER 1

INTRODUCTION

1.1 Muon Colliders

To complement current and future hadron colliders, higher energy lepton col-
liders will be needed to gain precision mass measurements of the Higgs and possible
beyond-standard-model particles [1]. Hadron colliders are excellent for use as dis-
covery machines due to the wide range of energy of their collisions. This is because
the constituent quarks that are actually involved in the collision can take on any
percentage of the total energy of their parent hadrons. Electron-positron colliders are
used for precision measurements of particle masses since electrons and positrons are
fundamental particles. Therefore, the energy of the particles involved in the collision
is much closer to the average energy of the particle bunch.

That said, electron-positron colliders are radiatively limited. In circular col-
liders, energy is lost to synchrotron radiation. Charged particles radiate when accel-
erated, and synchrotron radiation is the radiation that is induced by the acceleration
produced by the bending magnets. As (1.1) shows, the energy lost to synchrotron
radiation per turn around the accelerator ring, ΔE, is inversely proportional to the
fourth power of the mass, m. It is also inversely proportional to the bending radius,
R, and proportional to the fourth power of the current energy of the particle, E [2]:

$$\Delta E = \frac{(6.03 \times 10^{-15} \text{ MeV} \cdot \text{m})}{R} \left(\frac{E}{mc^2} \right)^4.$$

(1.1)

Hadrons are not significantly affected by synchrotron radiation loss. Conversely, the
electron with a mass approximately 1900 times smaller than the proton experiences
an energy loss thirteen orders of magnitude larger than a proton. At some point, the

accelerating elements cannot replace the energy lost to synchrotron radiation, not to mention accelerate the particles to the higher and higher energies required to produce newly discovered particles.

The energy loss due to synchrotron radiation can be mitigated to some degree by making the collider ring larger or by building a linear collider. That said, even in linear colliders the particle bunches are affected by beamstrahlung. This is energy loss due to acceleration from other particles' electric fields. At higher beam energies, this energy loss results in an unacceptable energy spread for what are supposed to be precision energy measurements [3].

Therefore, it is more efficient to achieve a desired collision energy by using a higher-mass particle rather than trying to accelerate a lower-mass particle to higher speeds. In line with this strategy, a promising alternative to hadron and electron-positron colliders is the muon collider. The muon, a fundamental particle like the electron, has a mass over 200 times that of the electron. Using muons instead of electrons and positrons reduces the radiation losses by nine orders of magnitude. A 2 TeV muon collider would offer about the same collision energy as the 7 TeV Large Hadron Collider that discovered the Higgs, but could fit in the footprint of the Tevatron at Fermilab. The difference in accelerating energy, again, is due to the fact that only a single constituent is involved in the collision for each hadron. Figure 1.1 compares the relative energies and sizes of a selection of relevant past and possible future collider machines [4].

1.2 Neutrino Factories

The standard model does not account for the masses of the neutrinos. The physics that give rise to the neutrino masses are studied through counting experiments that measure either the disappearance or appearance of particular flavors of neutrinos
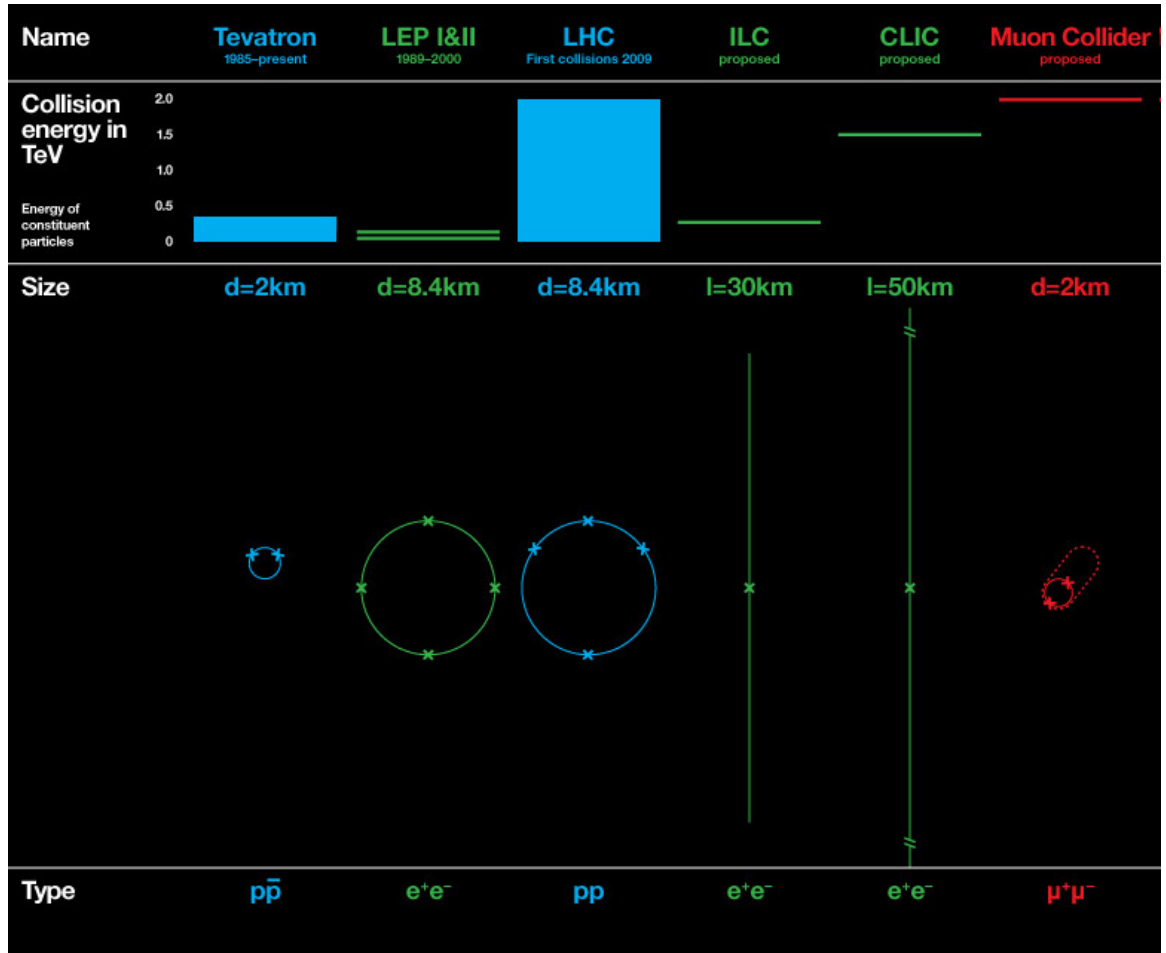
Figure 1.1. A comparison of past, present, and possible future collider machines [4]. Note that the Tevatron in fact ceased operation in 2011 [5], and the first LHC collisions were in 2010 [6].

due to the oscillation of these neutrinos between the flavors. Isotropic radiators such as the sun or nuclear reactors can serve as sources for neutrinos in these experiments [7]. For more intense and pure beams, neutrinos can be obtained from decay of focused pions. Pions can be produced by directing a proton beam at a fixed target [8, 9].

Charged pions decay 99.99% of the time via the following modes [10]:

$$\pi^+ \rightarrow \mu^+ + \nu_\mu, \tag{1.2}$$

$$\pi^- \rightarrow \mu^- + \overline{\nu}_\mu. \tag{1.3}$$

Neutrinos cannot be focused, so the size of the neutrino beam produced from pion decay is dependent on the energy of the pions and how well they are focused. The flux of neutrinos impinging on a detector of size A at a given decay angle, θ, is [11]

$$\phi_\nu = \frac{A}{4\pi z^2} \left(\frac{2\gamma}{1 + \gamma^2\theta^2} \right)^2, \tag{1.4}$$

where z is the distance of the detector from the pion decay point, and γ is the pion boost factor. Figure 1.2 plots (1.4) at several values of $\beta = v/c$, where v is the pion speed, c is the speed of light, and $\gamma = 1/\sqrt{1 - \beta^2}$. Higher energy implies a β closer to 1 and a larger flux at smaller angles to the parent pion velocity vector. Therefore, increasing the pion energy causes the decay neutrinos to be emitted in a narrower and narrower cone in the same direction as the parent pions [11]. This energy is limited by the pion production process. In addition, charged pions have a mean rest-frame lifetime of only 26 ns [10]. Typically, magnetic horns are used to provide single-pass focusing of the pions before they decay [11]. These horns are also used to remove unwanted kaons produced along with the pions, but this selection process is imperfect. Since kaons decay in significant numbers into electron neutrinos, the muon neutrino beam becomes contaminated by varying numbers of electron neutrinos [12].

A neutrino factory employing muons, on the other hand, could provide a higher-intensity, well-characterized neutrino beam. Muons decay effectively 100%
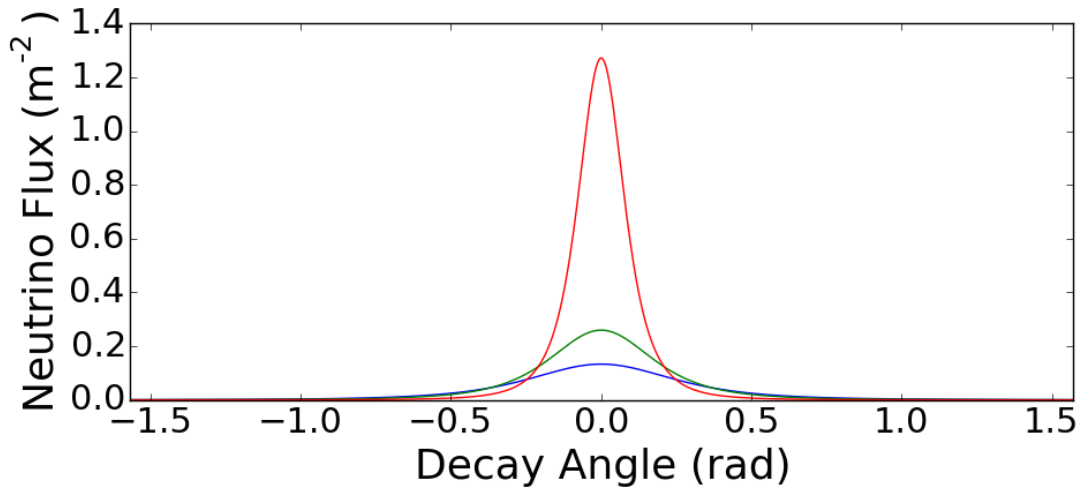
Figure 1.2. A plot of neutrino flux vs decay angle for several values of β, namely 0.9 (blue), 0.95 (green), and 0.99 (red). Higher energy implies a β closer to 1 and a larger flux at smaller angles to the parent pion velocity vector.

of the time via the following decay modes [10]:

$$\mu^+ \to e^+ + \nu_e + \overline{\nu}_\mu, \tag{1.5}$$

$$\mu^- \to e^- + \overline{\nu}_e + \nu_\mu. \tag{1.6}$$

Therefore, a muon beam would have the added benefit of producing both muon and electron-flavored neutrinos in precise proportions (50/50). Compared to pions, the longer mean rest-frame lifetime of the muon (2.2 µs) also provides the ability to accelerate and focus the neutrino-producing beam to a greater extent [10]. This would allow neutrino physicists to drastically improve their statistics and systematic uncertainties of neutrino mixing [13].

Furthermore, a muon collider and a neutrino factory based on muons could be built on the same site as part of a staged muon facility (Fig. 1.3). The neutrino factory would serve in part as a test facility for an eventual upgrade to a muon collider. The muon collider would reuse the neutrino factory's front end that supplies

the initial muon beam. Building in stages would spread the cost burden over time while guaranteeing the facility would be producing valuable results at each stage [14].
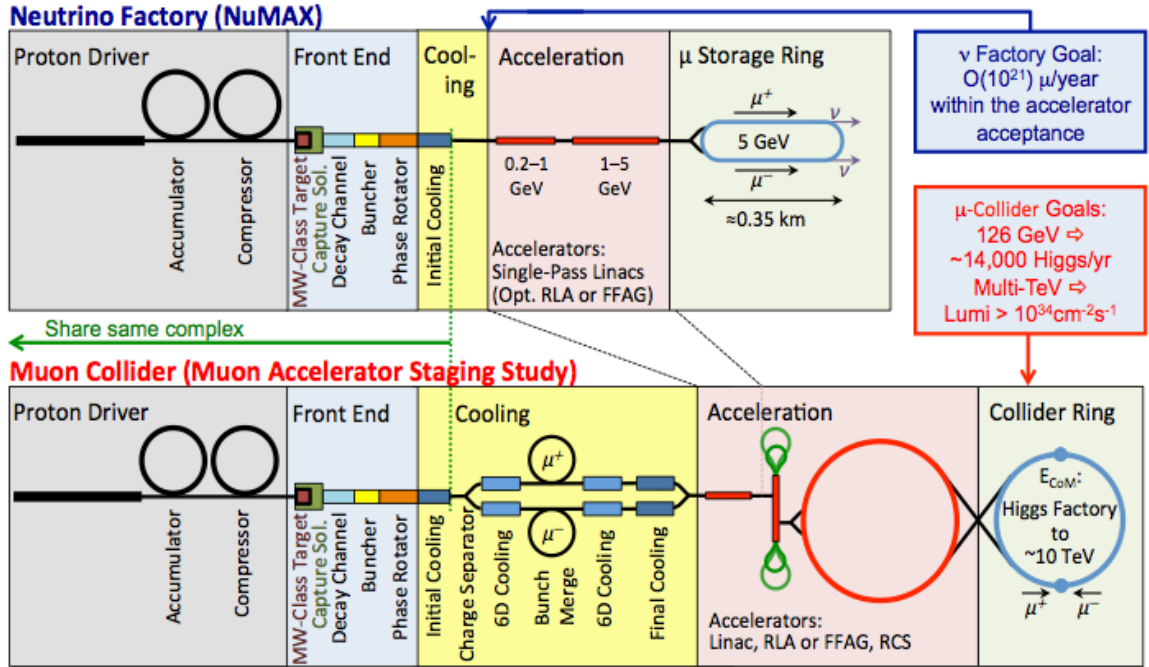
Figure 1.3. A muon collider and neutrino factory could be built at the same facility in stages using the same muon front end [4].

1.3 Ionization Cooling

The challenge in creating a muon beam for use in a collider or neutrino factory, therefore, is to collect the muons from the decay of pions and condition them into a useful beam before they also decay. This is exacerbated by the two-step production process involving a proton beam hitting a target to produce the pion beam and further decay into the beam of muons with an even larger phase space volume than the pions. To create a useful beam, the phase space volume of the muons must be reduced. This process is commonly referred to as cooling [15].

One common technique for reducing the size of a particle beam is stochastic cooling. This process uses electrical feedback loops to gradually kick the particles into

tight beams [16]. Given the short half-life of the muon, this method is too slow for muon-based neutrino factories and colliders. The only promising alternative to cool muons is ionization cooling. In ionization cooling, particles are first directed at an absorber. Muon energy is lost to ionization of the absorber atoms. The muons can also be scattered via collisions with nuclei, so care must be taken to emphasize ionization energy loss and reduce nuclear multiple scattering. Equation (1.7) describes the change in normalized transverse emittance ($\epsilon_\perp^N$) with respect to the length of the path taken through the material (s) [17]. Emittance is a measure of the size of the beam in phase space:

$$\frac{d\epsilon_\perp^N}{ds} \approx -\frac{1}{\beta^2} \left\langle \frac{dE_\mu}{ds} \right\rangle \frac{\epsilon_\perp^N}{E_\mu} + \frac{1}{\beta^3} \frac{\beta_\perp \, (0.014 \text{ GeV/c})^2}{2 \, E_\mu \, m_\mu \, X_0}. \tag{1.7}$$

The first term on the right represents the cooling of the muon beam due to ionization energy loss. X_0 is the radiation length of the absorber and Z is the atomic number of the absorber material. The value of X_0 decreases with increasing Z; thus the second term on the right, representing muon beam heating due to nuclear multiple scattering, grows with the number of protons in the absorber nuclei. Thus, it is better to use "low-Z" absorbers such as liquid hydrogen or LiH. The remaining values in the above equation are the speed (β), betatron function ($\beta_\perp$), muon mass (m_μ), and muon energy (E_μ).

For a given particle mass and energy, the only other way we can reduce beam heating is by decreasing the betatron function, $\beta_\perp$. This can be accomplished by using strong solenoid magnets to keep the muons close to the design trajectory. One consequence of this is that the RF cavities are exposed to strong magnetic fields. For this reason, only normal conducting RF cavities can be used. Superconducting RF cavities cannot be used in a cooling channel that uses solenoid magnets since the magnetic field would disrupt the mechanism that gives the cavities their superconducting properties [18]. This, in turn, requires that the accelerating cavities be operated at

close to their maximum accelerating gradient in order to guarantee a sufficiently compact cooling channel. By efficiently cooling the muon beam, the number of muons lost due to decay is kept at an acceptable level.

To reduce the longitudinal emittance, a process called emittance exchange is employed whereby higher momentum muons are forced to take a longer path through the absorber material than lower momentum muons. This can be done by bending the beam to introduce dispersion, spreading the particles out transversely in proportion to their individual momenta. The beam is then passed through an absorber that is shaped specifically to expose more material to muons which were bent the least (Fig. 1.4) [19].
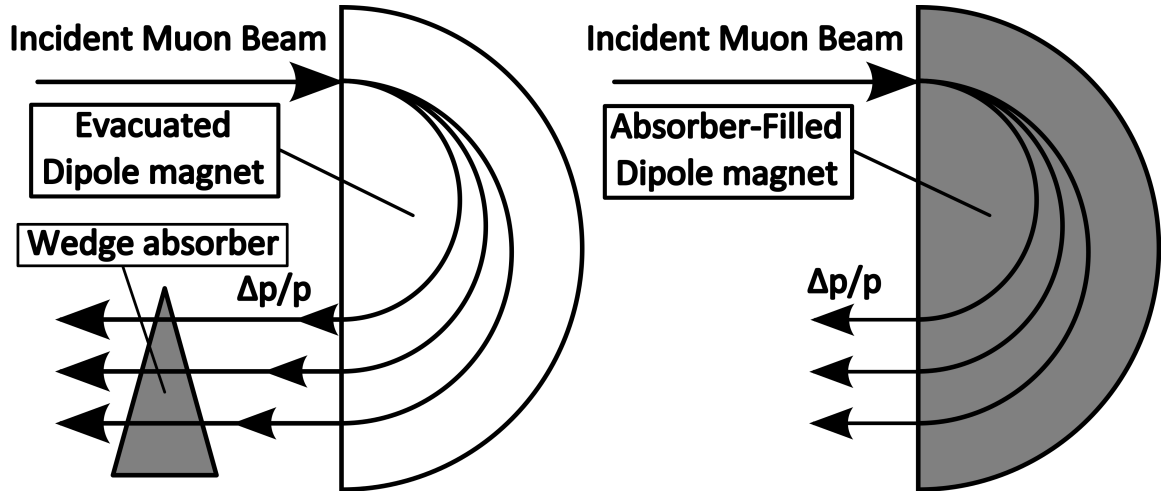
Figure 1.4. Emittance exchange [20]. After (left) or while (right) introducing dispersion to the muon beam via a bending magnet, muons pass through an absorber shaped to expose more material to muons with the highest energy (i.e., that are bent the least).

To restore the energy lost in the absorbers and maintain a longer lifetime via relativistic time dilation, the muon beam is re-accelerated with RF cavities in the longitudinal direction. The combination of transverse and longitudinal emittance

reduction by the absorbers and reacceleration produces a muon beam with the same average energy with which it started but more tightly focused (Fig. 1.5).

Figure 1.5. Ionization cooling with re-acceleration: 1) ionization energy loss (green), 2) multiple scattering off of absorber nuclei (red), 3) re-acceleration with RF cavities (blue) [21].

1.3.1 Muon Ionization Cooling Experiment. The Muon Ionization Cooling Experiment (MICE), located at Rutherford Appleton Laboratory in the United Kingdom, aims to demonstrate the feasibility of using ionization cooling to produce low-emittance muon beams [22]. The goals of the experiment as of May 2016 are to demonstrate transverse emittance reduction via absorbers by the end of 2016, followed by sustainable, transverse emittance reduction and re-acceleration.

1.4 RF Cavity Breakdown

RF cavity breakdown is the sudden collapse of the accelerating field inside a cavity due to the equalization of the accelerating gap surface charges. While there are various theories on what exactly is the cause, in general, there are four stages of breakdown: electron field emission, material ejection, plasma formation, and surface charge cancellation (Fig. 1.6) [23].

It is thought that asperities on the inner surface of the cavity are ultimately the source of breakdown. A process called Fowler-Nordheim field emission causes electrons to be emitted from the tips of the asperities via quantum tunneling due to an enhanced electric field that lowers the potential barrier at the tip surfaces [24]. At some point, cavity wall material is ejected, perhaps due to damage from the field-
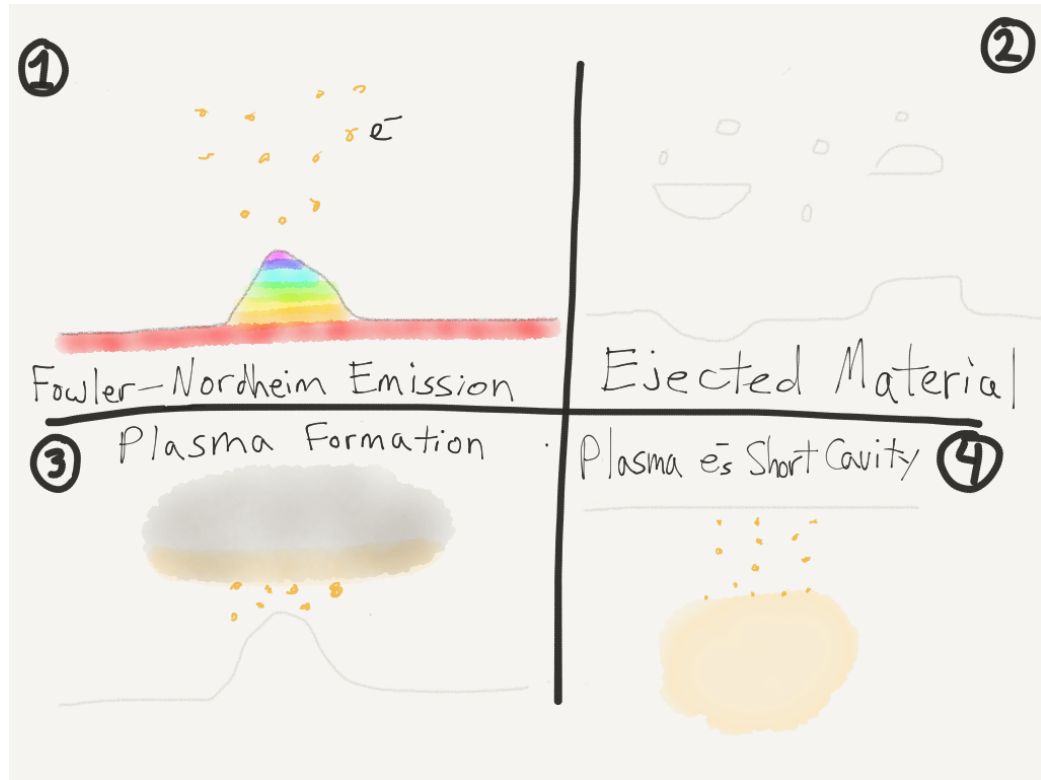
Figure 1.6. Cartoon of the four general stages of RF cavity breakdown.

emitted electrons, or some other process associated with electromagnetically induced stresses. This ejected material is then bombarded by the stream of electrons coming from the asperities, creating a plasma. The accelerating field of the cavity then transports the free plasma electrons to the opposite cavity wall, canceling the surface charges that sustain the electric field [23] (see Ch. 2 for details on electric fields and surface charges in RF cavities). This rapid discharge of electrons that collapses the accelerating gradient is referred to as a "spark".

Studies have found that the maximum accelerating gradient of a cavity, the gradient above which breakdown readily occurs, can be reduced when the cavity is exposed to a strong, solenoidal magnetic field [25]. The field-emitted electrons are focused by the magnetic field. This is thought to lower the barrier for breakdown arc formation in some way [26]. For example, the focused electrons may create more asperities or more quickly ionize ejected material.

When an RF cavity breaks down, significant damage may be done to the inner cavity surface if the amount of ejected material is large enough. Upon inspection, one can find pits as well as small patches where once molten metal has re-solidified on the surface (see Fig. 1.7). This can eventually degrade the performance of the cavity to a point where it needs to be refurbished or replaced. That said, most cavities go through a conditioning process whereby breakdown occurs frequently at some initial maximum gradient for some period of time, after which breakdown occurs much less frequently or not at all. One possible explanation is that the asperities that may be responsible for creating electric field enhancement are destroyed, resulting in a smoother surface [27]. Therefore, some breakdown appears to be healthy, while too much can be detrimental to the cavity's operation.
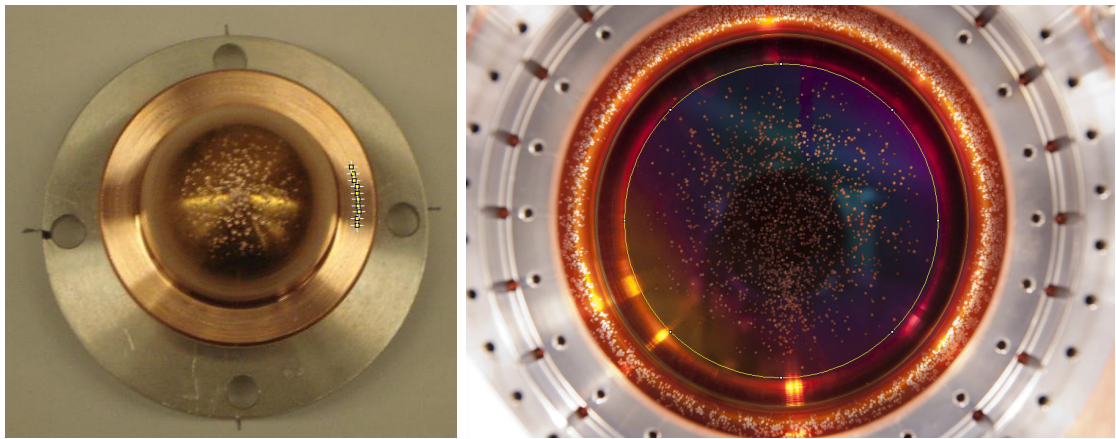
Figure 1.7. Examples of breakdown spark damage [28, 29].

1.5 MuCool Test Area

The MuCool Test Area (MTA) at Fermilab was constructed in 2003 for the purpose of testing and developing component technology for muon ionization cooling [30]. Because muon ionization cooling channel designs will likely employ high-gradient, normal conducting RF cavities, the MTA group has been engaged in a program of testing normal conducting RF cavity designs at high gradients. It is often

important in such studies to determine what the maximum gradients of certain cavity designs are both with and without strong magnetic fields.

The experiment hall has three RF stations, two for 805 MHz operation and one for 201 MHz operation (Fig. 1.8). Station 1 is collocated with a 5 T superconducting solenoid magnet that can also receive proton beam siphoned from Fermilab's H^- linac. Station 2 sits off to the side and upstream from station 1 and is used for diagnostics and testing 805 MHz cavities when no magnetic field or particle beam is needed. Station 3 allows for operation of larger cavities at 201 MHz such as the MICE prototype and production version cavities. This station is situated behind station 1 where cavities too large to fit in the solenoid can still be exposed to fringe fields from the solenoid magnet.
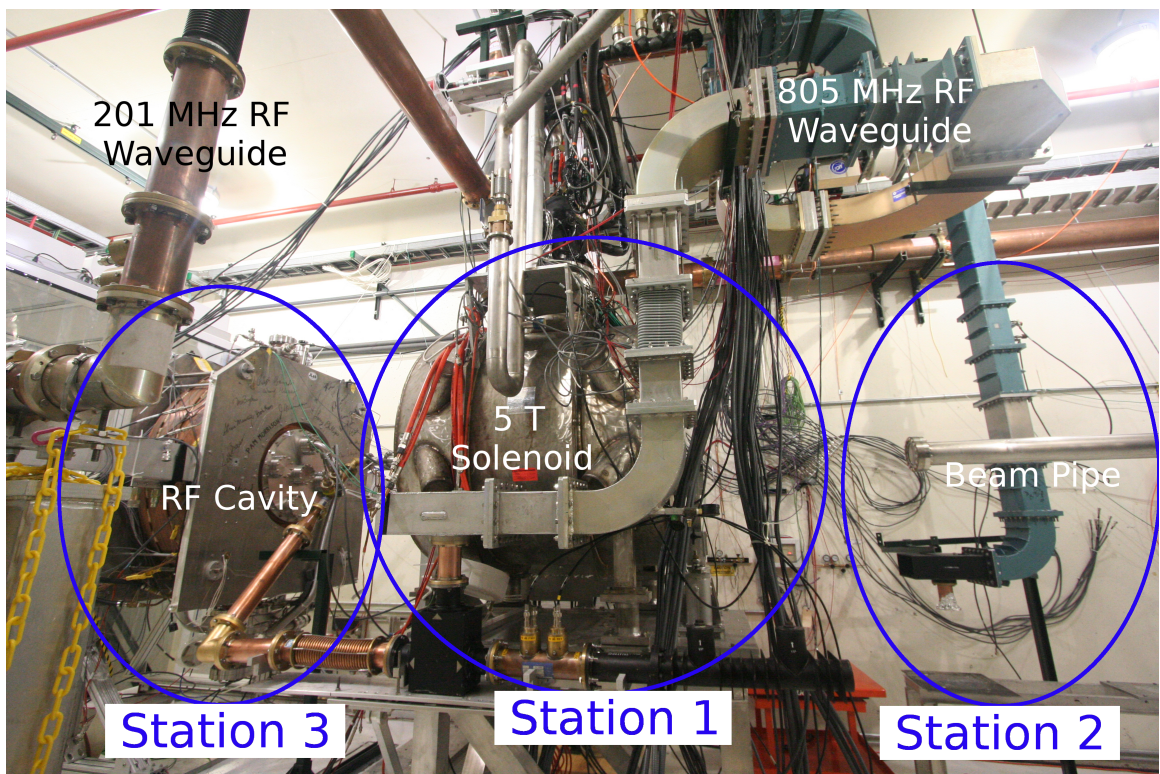
Figure 1.8. The MuCool Test Area experiment hall [29].

1.6   Cavities

Several RF cavities have been tested in the MTA experiment hall over the past five years. During this time, two of these cavities in particular, the High-Pressure Cavity and the Modular Cavity, have been used to provide results for the research described in this document. In addition, a mock cavity was machined out of aluminum specifically for this research. An introduction to these three cavities follows.

1.6.1 Aluminum Mock Cavity. Created for the purpose of testing localization algorithms on a physical cavity analog, the Aluminum Mock Cavity (AMC) has the same overall dimensions as the High-Pressure Cavity (see Sec. 1.6.2) but with no bolt, coupler port, or gas port holes except a few used for attaching handles (Fig. 1.9). It comprises two 12" diameter, 2" thick end plates, and a center ring with an outer diameter of 12", an inner diameter of 9", and a length of 3.2". As the name suggests, this cavity is made entirely out of aluminum. This makes it much easier and safer to handle compared to working directly with, say, a spare High-Pressure Cavity made of stainless steel.
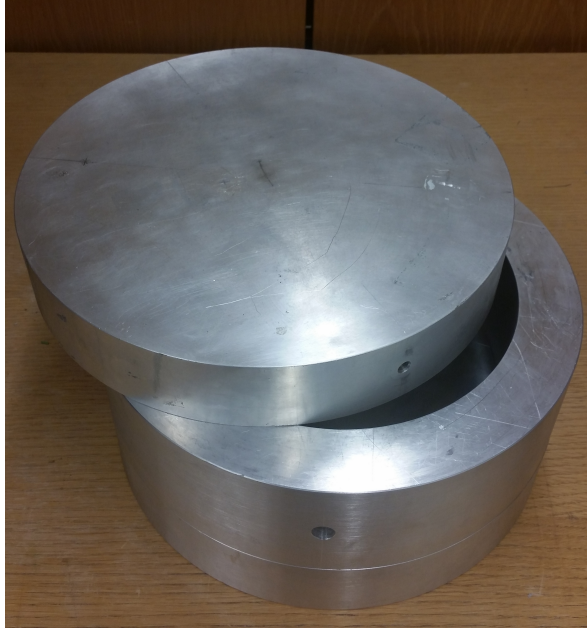
Figure 1.9. The Aluminum Mock Cavity.

1.6.2 High-Pressure Cavity.   Made of thick, copper-coated stainless steel, the High-Pressure Cavity (HC) was built for testing the concept of a normal conducting cavity that performs well in a strong magnetic field by using a high pressure fill gas to inhibit breakdown and cool the cavity [31]. It was driven by 40 μs, 805 MHz RF pulses from a spare klystron for the Fermilab H⁻ linac. Select HC engineering drawings can be found in App. A.

The HC is a modified pillbox cavity. Since the AMC was built to the same general dimensions, the HC also comprises 2" thick end plates and a 3.2" long ring. The inner and outer diameters are 9" and 12" respectively. Axial electrodes, or buttons, can be attached to the inner surface of the end plates to intensify the electric field at the axis and all but guarantee that the breakdown will originate on the buttons (Fig. 1.10, top). Adding the buttons decreases the frequency of the cavity, so the inner radius is smaller than that of an ideal pillbox (see Sec. 2.1) to bring the frequency back up within the 805 MHz klystron's bandwidth.

The HC can also be configured with dielectric inserts for dielectric loading tests instead of buttons (Fig. 1.10, bottom). The dielectric decreases the resonant frequency of the cavity. Again, the smaller radius compensates for this effect.

1.6.3 Modular Cavity.   Intended for the study of normal conducting RF cavity breakdown, the Modular Cavity (MC) is a pillbox cavity designed to have easily swappable end plates for testing different plate materials, and side-coupled RF to reduce the electric field intensity near the coupler opening (Fig. 1.11) [32]. As it is driven by 805 MHz RF, the inner diameter is 28.2 cm. The outer diameter and length are 31.1 cm and 13.1 cm respectively. The end plates are 1.4 cm thick, and 29.8 cm in diameter. Flange rings hold the plates on the ends of the cavity, leaving an exposed diameter of about 26.4 cm to which instrumentation can be adhered. Select MC engineering drawings can be found in App. A.
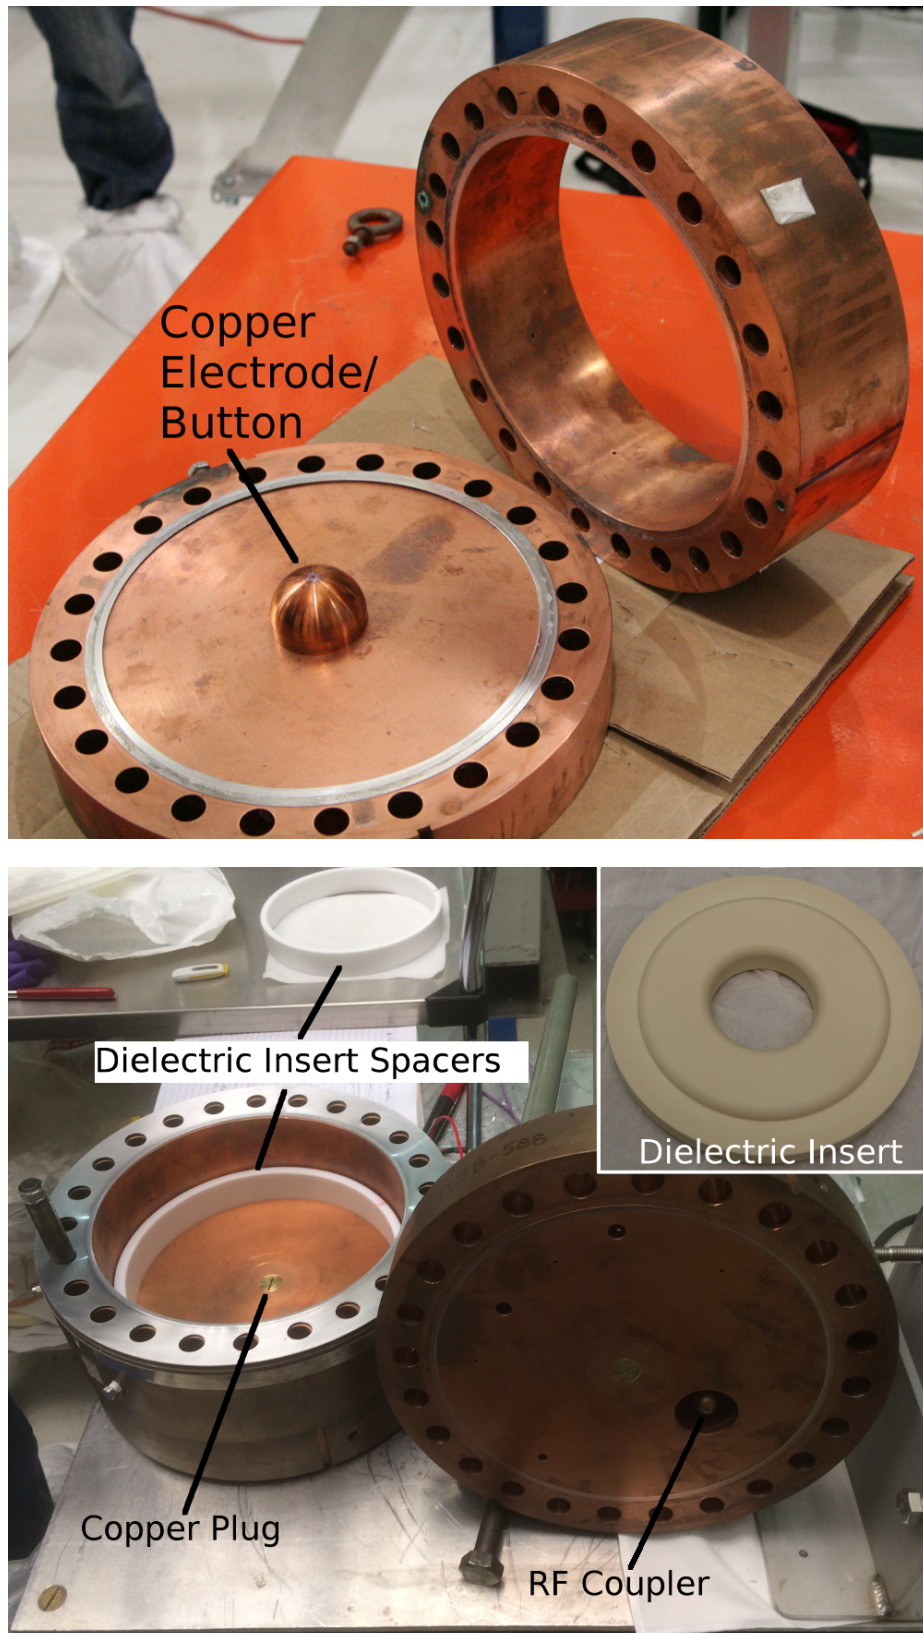
Figure 1.10. Top: HC ring and end plate with electrode [29]. Bottom: HC in dielectric insert configuration [29].
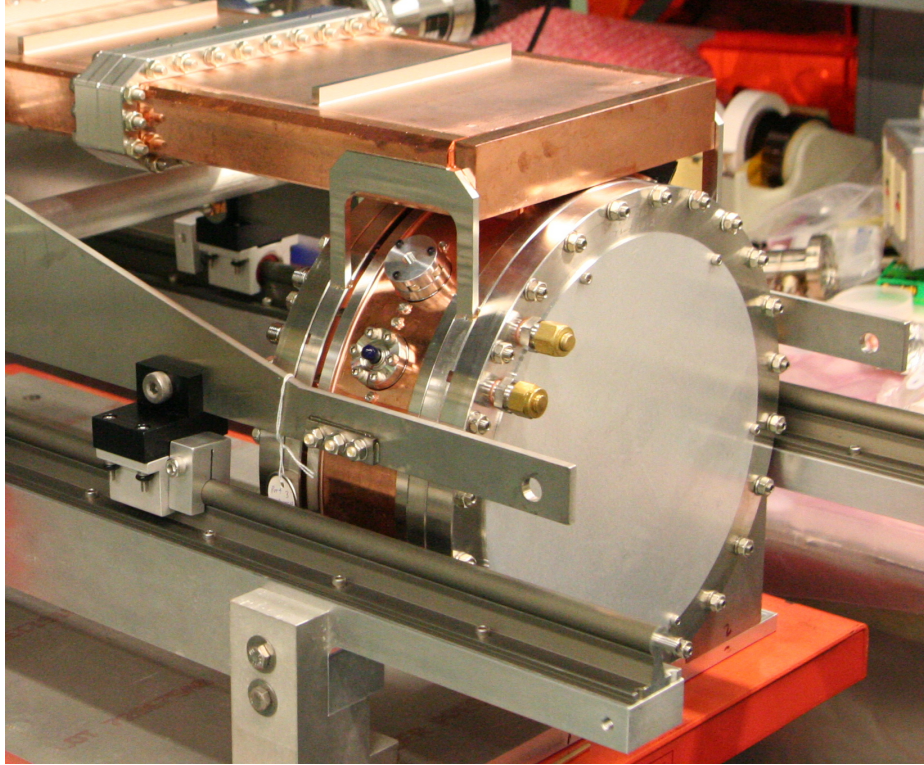
Figure 1.11. The Modular Cavity [29].

1.7 Thesis Statement

To aid the effort to study RF cavity breakdown in magnetic fields, it would be helpful to have a diagnostic tool which can localize the source of breakdown inside RF cavities. Acoustic localization of the source of thermal shock due to RF cavity breakdown using an array of microphones attached to the outside end plate surfaces of a pillbox cavity is possible with an average accuracy better than 3 cm. The average accuracy can be improved to better than 1 cm if sparks are assumed to be occurring within half the cavity radius of the beam axis. This is demonstrated through the analysis of structural mechanics simulation results and experimental data (Ch. 6).

Sound pressure estimates suggested that it ought to be possible to differentiate or separate breakdown sound from sound produced by the electromagnetic forces on the inside of a cavity during an RF pulse (Ch. 2). An appropriate data acquisition

system of hardware (Ch. 3) and LabVIEW software (Ch. 4) was built that provided and processed the experimental acoustic data. Finally, an algorithm based off of one intended for free-space localization was developed (Ch. 5) to perform this RF cavity breakdown localization demonstration.

CHAPTER 2

RF CAVITY SOUND

During either the conditioning process or normal running, it would be useful to know where breakdown sparks are forming. Though a cavity can be removed from operation, opened, and inspected for corresponding damage, it would be advantageous to be able to tell—ideally in real-time—while running whether, for example, breakdown was occurring due to normal activity or due to a faulty RF coupler. From this perspective breakdown can be categorized into that which is expected and that which is unexpected. Knowing what type of breakdown is occurring can affect how the cavity is operated or whether it needs to be decommissioned, repaired, modified, or redesigned. Furthermore, little observational data of the actual breakdown events exist due to the limitations of adding instrumentation inside the cavity. A non-invasive technique for observing additional aspects of breakdown could help refine the theories on its dynamics.

In RF cavities, sound is generated both by the sudden onset of electromagnetic forces on the inner cavity walls during an RF pulse (RF hammer) as well as by the thermal shock from the electric current during breakdown. In principle, acoustic information could be used to observe the timing of the thermal shock at a particular point on one of the inner walls of an RF cavity due to a breakdown spark current (see Sec. 2.3). By spatially separating several microphones, the physical location could be determined by comparing the relative times of arrival of the sound from the thermal shock at these different microphones.

To understand these sources of sound, a review of the electromagnetic theory of RF cavities is warranted. It is also useful to estimate the initial pressures on the

inside walls of typical cavities in order to understand the requirements for hardware and software needed to localize the source of breakdown sound (see Ch. 3 and Ch. 4).

2.1 RF Accelerating Cavities

To better understand the sources of sound in RF cavities, it is important to understand the electric and magnetic fields inside such a cavity. One of the simplest types of RF cavities is the pillbox cavity. A pillbox in this context is a hollow cylinder. To first order, all of the cavities used in this research are pillbox cavities. It is therefore appropriate to use the pillbox cavity as the basis for these electric and magnetic field derivations.

2.1.1 Cylindrical Waveguide Equations. We start with Maxwell's equations in cgs units:

$$\nabla \cdot \mathbf{E} = 4\pi\rho, \qquad \nabla \cdot \mathbf{B} = 0,$$

$$\nabla \times \mathbf{E} = -\frac{1}{c}\frac{\partial \mathbf{B}}{\partial t}, \quad \nabla \times \mathbf{B} = \frac{4\pi}{c}\mathbf{j} + \frac{1}{c}\frac{\partial \mathbf{E}}{\partial t}.$$

A cylindrical cavity is just a cylindrical waveguide with additional boundary conditions along the direction of propagation (here taken as longitudinal along the z-axis). Therefore the electric and magnetic fields in a cylindrical waveguide will be determined first. To begin, the plane-wave electric and magnetic fields in cylindrical coordinates (ρ, ϕ, z) are sought:

$$\mathbf{E}(\rho, \phi, z, t) = E_\rho\hat{\boldsymbol{\rho}} + E_\phi\hat{\boldsymbol{\phi}} + E_z\hat{\mathbf{z}} = \left[\tilde{E}_\rho(\rho, \phi)\hat{\boldsymbol{\rho}} + \tilde{E}_\phi(\rho, \phi)\hat{\boldsymbol{\phi}} + \tilde{E}_z(\rho, \phi)\hat{\mathbf{z}}\right]e^{i(k_z z - \omega t)},$$

$$\mathbf{B}(\rho, \phi, z, t) = B_\rho\hat{\boldsymbol{\rho}} + B_\phi\hat{\boldsymbol{\phi}} + B_z\hat{\mathbf{z}} = (\tilde{B}_\rho(\rho, \phi)\hat{\boldsymbol{\rho}} + \tilde{B}_\phi(\rho, \phi)\hat{\boldsymbol{\phi}} + \tilde{B}_z(\rho, \phi)\hat{\mathbf{z}})e^{i(k_z z - \omega t)},$$

with the assumption that only the real part of the complex exponential contributes to the physical fields.

Maxwell's equations in free space with the substitutions $\partial\mathbf{E}/\partial t = -i\omega\mathbf{E}$ and $\partial\mathbf{B}/\partial t = -i\omega\mathbf{B}$ are

$$\nabla \cdot \mathbf{E} = 0, \tag{2.1}$$

$$\nabla \cdot \mathbf{B} = 0, \tag{2.2}$$

$$\nabla \times \mathbf{E} = \frac{i\omega}{c}\mathbf{B}, \tag{2.3}$$

$$\nabla \times \mathbf{B} = -\frac{i\omega}{c}\mathbf{E}. \tag{2.4}$$

Taking the curl of the curl equations, one eventually gets wave equations for the electric and magnetic fields:

$$\nabla^2\mathbf{E} + \left(\frac{\omega}{c}\right)^2\mathbf{E} = 0,$$

$$\nabla^2\mathbf{B} + \left(\frac{\omega}{c}\right)^2\mathbf{B} = 0.$$

It is convenient to split these equations into transverse and longitudinal components. It will then be shown that all of the other fields can be derived from just the longitudinal component of one of the fields. To wit, the following are defined:

$$\nabla_z^2 = \frac{\partial^2}{\partial z^2}, \qquad\qquad \nabla_\perp^2 = \nabla^2 - \nabla_z^2,$$

$$\mathbf{E} = \mathbf{E}_z + \mathbf{E}_\perp, \qquad\qquad \mathbf{B} = \mathbf{B}_z + \mathbf{B}_\perp,$$

$$\mathbf{E}_z = E_z(\rho,\phi,z,t)\hat{\mathbf{z}}, \qquad\qquad \mathbf{B}_z = B_z(\rho,\phi,z,t)\hat{\mathbf{z}}.$$

$$\mathbf{E}_\perp = E_\rho(\rho,\phi,z,t)\hat{\boldsymbol{\rho}} + E_\phi(\rho,\phi,z,t)\hat{\boldsymbol{\phi}}, \qquad \mathbf{B}_\perp = B_\rho(\rho,\phi,z,t)\hat{\boldsymbol{\rho}} + B_\phi(\rho,\phi,z,t)\hat{\boldsymbol{\phi}}.$$

A longitudinal electric field is required in order to accelerate particles, so it is natural to start by addressing only the longitudinal component of the electric field

wave equation:

$$\nabla_\perp^2 E_z + \nabla_z^2 E_z + \left(\frac{\omega}{c}\right)^2 E_z = 0.$$

In cylindrical coordinates $\nabla_\perp^2$ is

$$\nabla_\perp^2 = \frac{1}{\rho}\frac{\partial}{\partial\rho}\left(\rho\frac{\partial}{\partial\rho}\right) + \frac{1}{\rho^2}\frac{\partial^2}{\partial\phi^2}.$$

Expanding $\nabla_\perp^2$ and evaluating $\nabla_z^2 E_z$,

$$\frac{1}{\rho}\frac{\partial}{\partial\rho}\left(\rho\frac{\partial E_z}{\partial\rho}\right) + \frac{1}{\rho^2}\frac{\partial^2 E_z}{\partial\phi^2} - k_z^2 E_z + \left(\frac{\omega}{c}\right)^2 E_z = 0.$$

Making the substitution

$$k_\rho^2 \equiv (\omega/c)^2 - k_z^2, \tag{2.5}$$

this becomes

$$\frac{1}{\rho}\frac{\partial}{\partial\rho}\left(\rho\frac{\partial E_z}{\partial\rho}\right) + \frac{1}{\rho^2}\frac{\partial^2 E_z}{\partial\phi^2} + k_\rho^2 E_z = 0.$$

We can solve for E_z using separation of variables. Assuming a solution of the form $E_z = X(\rho)\Phi(\phi)e^{i(k_z z - \omega t)}$ and multiplying through by ρ^2,

$$\rho\frac{\partial}{\partial\rho}\left(\rho\frac{\partial X}{\partial\rho}\right)\Phi + \frac{\partial^2\Phi}{\partial\phi^2}X + \rho^2 k_\rho^2 X\Phi = 0.$$

Dividing through by $X(\rho)\Phi(\phi)$ we can separate the ρ and ϕ terms, setting them both equal to a constant:

$$\frac{\rho}{X}\frac{\partial}{\partial\rho}\left(\rho\frac{\partial X}{\partial\rho}\right) + \rho^2 k_\rho^2 = n^2.$$

$$-\frac{1}{\Phi}\frac{\partial^2\Phi}{\partial\phi^2} = n^2.$$

Rearranging and making the substitution $x = \rho k_\rho$ yields the following two differential equations:

$$\frac{\partial}{\partial x}\left(x\frac{\partial X(x)}{\partial x}\right) + (x^2 - n^2)X(x) = 0, \tag{2.6}$$

$$\frac{\partial^2\Phi(\phi)}{\partial\phi^2} + n^2\Phi(\phi) = 0. \tag{2.7}$$

Equation (2.7) can be solved easily:

$$\Phi(\phi) = \cos(n\phi + \delta).$$

Since there is no favored azimuthal orientation, we require a periodic boundary condition ($\Phi(\phi) = \Phi(\phi + 2\pi)$) and choose our coordinate system such that $\delta \equiv 0$. The boundary condition implies that n must be an integer, therefore

$$\Phi(\phi) = \cos(n\phi), \ n = 0, 1, 2, \ldots$$

Equation (2.6) is recognized as Bessel's equation. The solutions are Bessel functions of the first and second kind:

$$X(x) = a_n J_n(x) + b_n Y_n(x).$$

Now Y_n diverges on the axis ($x = \rho k_\rho = 0$), so we reject such solutions on the grounds that they are not physical. Making the substitution $E_n = a_n$ and reverting to using ρ instead of x, we have the following z component of the electric field:

$$E_z = E_n J_n(k_\rho \rho) \cos(n\phi) e^{i(k_z z - \omega t)}, \ n = 0, 1, 2, \ldots \tag{2.8}$$

Applying the boundary condition that $E_z \to 0$ at $\rho = R$ [33], where R is the waveguide radius, k_ρ is quantized:

$$k_{\rho,nm} = z_{nm}/R, \tag{2.9}$$

where z_{nm} are the zeros of $J_n(x)$. Furthermore, from (2.5) we know that $k_z = \sqrt{(\omega/c)^2 - k_{\rho,nm}^2}$. This implies that there is a cutoff frequency,

$$\omega_{c,nm} = z_{nm}c/R, \tag{2.10}$$

below which waves exponentially decay with distance along the z-axis in the waveguide.

A similar procedure can be followed to obtain B_z. Note that there is no requirement that B_z go to zero at the waveguide radius, and therefore ω_c is different for B_z. This implies that for a particular wave in the waveguide, either E_z or B_z can be non-zero, but not both. This defines the Transverse Magnetic (TM) and Transverse Electric (TE) waveguide modes respectively. Only TM modes with non-zero E_z are of interest for particle acceleration, so from here on it will be assumed that $B_z = 0$.

By manipulating equations (2.3) and (2.4) with the assumption that $B_z = 0$, the values for the remaining components of the electric and magnetic fields in a cylindrical waveguide can be derived. To begin, the curl of the electric field in cylindrical coordinates is expanded:

$$\nabla \times \mathbf{E} = \left(\frac{1}{\rho} \frac{\partial E_z}{\partial \phi} - \frac{\partial E_\phi}{\partial z} \right) \hat{\boldsymbol{\rho}} + \left(\frac{\partial E_\rho}{\partial z} - \frac{\partial E_z}{\partial \rho} \right) \hat{\boldsymbol{\phi}} + \frac{1}{\rho} \left(\frac{\partial (\rho E_\phi)}{\partial \rho} - \frac{\partial E_\rho}{\partial \phi} \right) \hat{\boldsymbol{z}}.$$

Next, the components of $\nabla \times \mathbf{E}$ that contain E_z are equated to the matching components of $\frac{i\omega}{c} \mathbf{B}$ in accordance with equation (2.3). This gives derivatives of the electric field components with respect to z which can be evaluated:

$$\frac{\partial E_\rho}{\partial z} - \frac{\partial E_z}{\partial \rho} = \frac{i\omega}{c} B_\phi \quad \longrightarrow \quad i k_z E_\rho = \frac{\partial E_z}{\partial \rho} + \frac{i\omega}{c} B_\phi, \tag{2.11}$$

$$\frac{1}{\rho} \frac{\partial E_z}{\partial \phi} - \frac{\partial E_\phi}{\partial z} = \frac{i\omega}{c} B_\rho \quad \longrightarrow \quad i k_z E_\phi = \frac{1}{\rho} \frac{\partial E_z}{\partial \phi} - \frac{i\omega}{c} B_\rho. \tag{2.12}$$

To eliminate the magnetic field component terms, a similar procedure with equation (2.4) is performed starting with the curl of the magnetic field:

$$\nabla \times \mathbf{B} = \left(\frac{1}{\rho} \frac{\partial B_z}{\partial \phi} - \frac{\partial B_\phi}{\partial z} \right) \hat{\boldsymbol{\rho}} + \left(\frac{\partial B_\rho}{\partial z} - \frac{\partial B_z}{\partial \rho} \right) \hat{\boldsymbol{\phi}} + \frac{1}{\rho} \left(\frac{\partial (\rho B_\phi)}{\partial \rho} - \frac{\partial B_\rho}{\partial \phi} \right) \hat{\boldsymbol{z}}.$$

In this case, the B_z terms can be eliminated since $B_z = 0$:

$$\frac{\partial B_\rho}{\partial z} = -\frac{i\omega}{c} E_\phi \quad \longrightarrow \quad B_\rho = -\frac{\omega}{k_z c} E_\phi, \tag{2.13}$$

$$-\frac{\partial B_\phi}{\partial z} = -\frac{i\omega}{c} E_\rho \quad \longrightarrow \quad B_\phi = \frac{\omega}{k_z c} E_\rho. \tag{2.14}$$

Substituting equations (2.13) into (2.12) and (2.14) into (2.11) gives E_ρ and E_ϕ in terms of derivatives of E_z:

$$E_\rho = i\frac{k_z}{k_\rho^2}\frac{\partial E_z}{\partial \rho}, \tag{2.15}$$

$$E_\phi = i\frac{k_z}{k_\rho^2 \rho}\frac{\partial E_z}{\partial \phi}. \tag{2.16}$$

By substituting equations (2.15) and (2.16) back into (2.14) and (2.13), equations for B_ρ and B_ϕ in terms of derivatives of E_z are obtained:

$$B_\rho = i\frac{\omega}{k_\rho^2 \rho c}\frac{\partial E_z}{\partial \phi}, \tag{2.17}$$

$$B_\phi = i\frac{\omega}{k_\rho^2 c}\frac{\partial E_z}{\partial \rho}. \tag{2.18}$$

Finally, by evaluating the E_z derivatives in equations (2.15), (2.16), (2.17), and (2.18), the following electric and magnetic field components in a cylindrical waveguide result:

$$
\begin{aligned}
E_z &= E_n J_n(k_{\rho,nm}\rho)\cos(n\phi)e^{i(k_z z - \omega t)} \\
E_\rho &= \frac{ik_z}{k_{\rho,nm}^2}\frac{\partial E_z}{\partial \rho} = \frac{ik_z}{k_{\rho,nm}}E_n J_n'(k_{\rho,nm}\rho)\cos(n\phi)e^{i(k_z z - \omega t)} \\
E_\phi &= \frac{ik_z}{k_{\rho,nm}^2}\frac{1}{\rho}\frac{\partial E_z}{\partial \phi} = -\frac{ik_z n}{k_{\rho,nm}^2 \rho}E_n J_n(k_{\rho,nm}\rho)\sin(n\phi)e^{i(k_z z - \omega t)} \\
B_z &= 0 \\
B_\rho &= -\frac{i\omega}{k_{\rho,nm}^2 c}\frac{1}{\rho}\frac{\partial E_z}{\partial \phi} = \frac{i\omega n}{k_{\rho,nm}^2 c\rho}E_n J_n(k_{\rho,nm}\rho)\sin(n\phi)e^{i(k_z z - \omega t)} \\
B_\phi &= \frac{i\omega}{k_{\rho,nm}^2 c}\frac{\partial E_z}{\partial \rho} = \frac{i\omega}{k_{\rho,nm} c}E_n J_n'(k_{\rho,nm}\rho)\cos(n\phi)e^{i(k_z z - \omega t)}
\end{aligned}
\tag{2.19}
$$

2.1.2 Cylindrical RF Cavity Equations. A cylindrical RF cavity is created from a waveguide by simply adding flat, conducting walls at $z = 0$ and $z = L$, making L the length of the cavity. Due to reflections, independent waves traveling in both longitudinal directions at the same time must be considered. In other words, we now

have the following:

$$\mathbf{E} = \tilde{\mathbf{E}}_1(\phi, \rho)e^{-i(k_z z - \omega t)} + \tilde{\mathbf{E}}_2(\phi, \rho)e^{i(k_z z + \omega t)}. \tag{2.20}$$

$$\mathbf{B} = \tilde{\mathbf{B}}_1(\phi, \rho)e^{-i(k_z z - \omega t)} + \tilde{\mathbf{B}}_2(\phi, \rho)e^{i(k_z z + \omega t)}. \tag{2.21}$$

The boundary conditions at the end walls dictate that E_ρ and E_ϕ go to zero at $z = 0$ and $z = L$ [33]. This creates the following constraint equations:

$$\tilde{\mathbf{E}}_1 + \tilde{\mathbf{E}}_2 = 0, \tag{2.22}$$

$$\tilde{\mathbf{E}}_1 e^{-ik_z L} + \tilde{\mathbf{E}}_2 e^{ik_z L} = 0; \tag{2.23}$$

or in matrix form

$$\begin{bmatrix} 1 & 1 \\ e^{ik_z L} & e^{-ik_z L} \end{bmatrix} \begin{bmatrix} \tilde{\mathbf{E}}_1 \\ \tilde{\mathbf{E}}_2 \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \end{bmatrix}.$$

To have a non-trivial solution, the determinant must be zero. Therefore,

$$e^{ik_z L} - e^{-ik_z L} = 2i\sin(k_z L) = 0.$$

This is satisfied only if

$$k_z = \frac{l\pi}{L}, \quad l = 0, 1, 2, \dots \tag{2.24}$$

From constraint equation (2.22) we see that $\tilde{\mathbf{E}}_1 = -\tilde{\mathbf{E}}_2$. The solutions for the electric field components, therefore, are

$$E_z = E_n J_n(k_{\rho,nm}\rho)\cos(n\phi)\cos\left(\frac{l\pi z}{L}\right)e^{-i\omega t},$$

$$E_\rho = \frac{l\pi}{k_{\rho,nm}L}E_n J_n'(k_{\rho,nm}\rho)\cos(n\phi)\sin\left(\frac{l\pi z}{L}\right)e^{-i\omega t},$$

$$E_\phi = -\frac{ln\pi}{k_{\rho,nm}^2\rho L}E_n J_n(k_{\rho,nm}\rho)\sin(n\phi)\sin\left(\frac{l\pi z}{L}\right)e^{-i\omega t}.$$

Since E_ρ and E_ϕ are zero at the end walls, equations (2.13) and (2.14) dictate that $\dfrac{\partial B\rho}{\partial z}$ and $\dfrac{\partial B\phi}{\partial z}$ must also be zero at the end walls. The magnetic field

components, then, are

$$B_z = 0, \tag{2.25}$$

$$B_\rho = \frac{i\omega n}{k_{\rho,nm}^2 c\rho} E_n J_n(k_{\rho,nm}\rho) \sin(n\phi) \cos\left(\frac{l\pi z}{L}\right) e^{-i\omega t}, \tag{2.26}$$

$$B_\phi = \frac{i\omega}{k_{\rho,nm} c} E_n J_n'(k_{\rho,nm}\rho) \cos(n\phi) \cos\left(\frac{l\pi z}{L}\right) e^{-i\omega t}. \tag{2.27}$$

In general, the best RF cavity mode to operate in is the lowest TM_{nml} mode, namely TM_{010}. One of the most important reasons for this is that $J_0(k_{\rho,01}\rho)$ is at a maximum at $\rho = 0$ (the longitudinal axis). All other orders of J_n are zero at the axis. Thus, only modes with $n = 0$ provide maximum interaction with an axial particle beam. Having n and l equal to zero also has the nice effect of canceling out the radial and azimuthal components of the electric field, as well as the radial component of the magnetic field. Furthermore, since the derivative of $J_0(x)$ is $-J_1(x)$, B_ϕ is zero at the axis. Having $l = 0$ also simplifies cavity design in that the frequency is then only dependent on the inner cavity radius. Plugging (2.9) and (2.24) into (2.5) we see that

$$\omega = c\sqrt{k_{\rho,01}^2 + k_{z,0}^2} = c\sqrt{(z_{01}/R)^2 + (0\pi/L)^2} = \frac{cz_{01}}{R}. \tag{2.28}$$

From (2.28) we also see that $k_{\rho,01} = \omega/c$ and cancels the factors of ω/c in (2.26) and (2.27). In addition, $l = 0$ creates a constant accelerating field upstream to downstream.

Lastly, we take the real part of the time exponentials to get the final form of the fields for a pillbox RF cavity operated in the TM_{010} mode:

$$E_z = E_0 J_0(k_{\rho,01}\rho) \cos(\omega t), \tag{2.29}$$

$$B_\phi = -E_0 J_1(k_{\rho,01}\rho) \sin(\omega t), \tag{2.30}$$

$$E_\rho = E_\phi = B_z = B_\rho = 0. \tag{2.31}$$

Note that this reveals that the magnetic field is 90 degrees out of phase with the electric field. Figure 2.1 shows what these fields look like for select values of the RF cycle phase [34].

2.2 RF Hammer Sound Production

The sound generated by the forces on the walls of the cavity due to the electric and magnetic fields of the RF pulse injected into the cavity are referred to as the "RF hammer". This sound, though not useful for localizing breakdown sparks, is important since it provides the primary source of background noise in vacuum cavities. Furthermore, it is a well-understood phenomenon and, therefore, can also be used to validate simulation models.

To see what these forces look like, we use the equations for the electric and magnetic fields in an ideal pillbox cavity operating in TM_{010} mode, namely (2.29), (2.30), and (2.31), to determine the surface charge and current distributions. The fields then act on the surface charge and current distributions, producing pressure on the inner cavity surfaces. These pressures rise and fall very rapidly with the RF frequency (on the order of hundreds of MHz), making it impossible to detect any acoustic oscillations assuming an RF steady state. It is in fact the rapid onset and decay of the RF pulse that produces detectable sound. These pulses can be from tens to hundreds of microseconds in duration depending on the cavity.

2.2.1 Charge and Surface Current Distributions. Since E_z is zero at the inner circumference of the cavity (setting $\rho = R$ in (2.29), where R is the inner cavity radius), only the end walls will have any surface charge. Thus any forces due to the electric field will be on the end surfaces. Assuming we have an ideal conductor, we can also set $E_z = 0$ inside the metal. The following boundary condition is used to
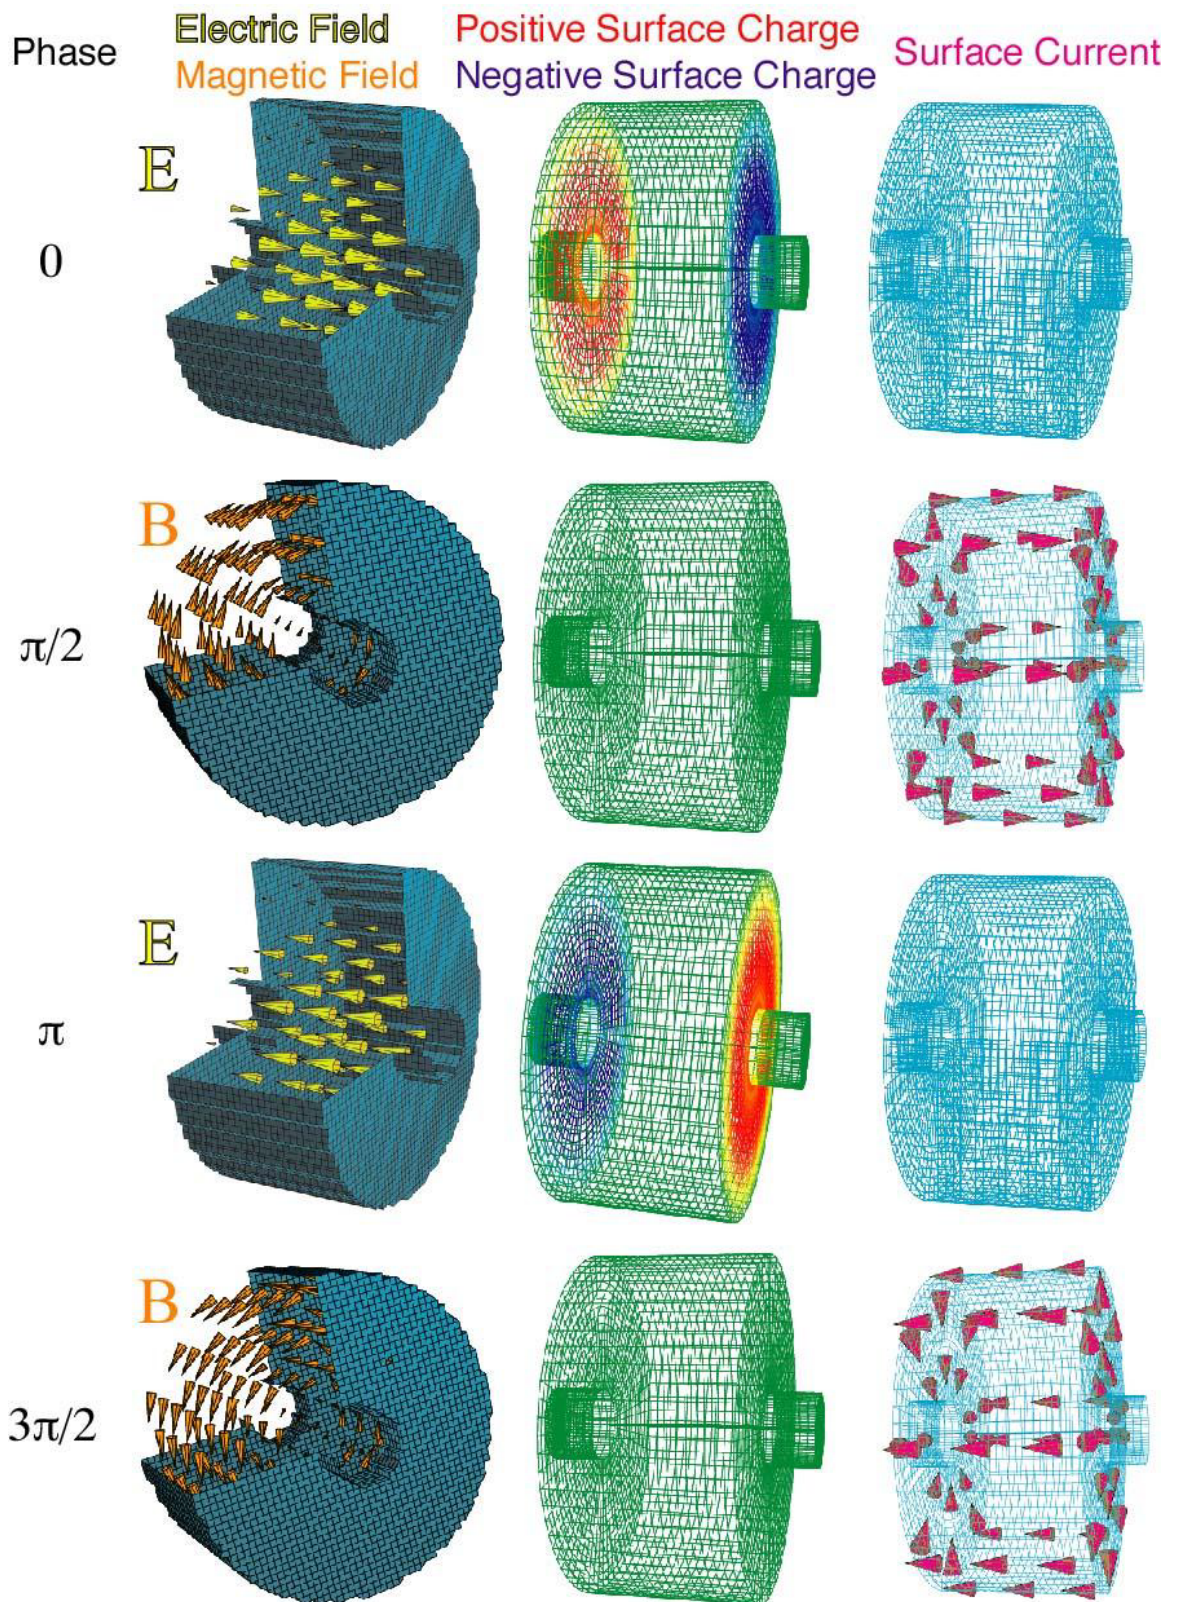
Figure 2.1. Pillbox electric and magnetic fields for select values of the RF phase, ω [34].

obtain the surface charge distributions on the end walls [33]:

$$\epsilon_1 E_1^\perp - \epsilon_2 E_2^\perp = 4\pi\sigma,$$

where ϵ is the electric permittivity, $E^\perp$ is the electric field component perpendicular to the surface, and σ is the surface charge. The numerical subscripts on ϵ and E denote the different materials on either side of the boundary. For $z = 0$, we have $\epsilon_1 = 1$ (vacuum), $E_1^\perp = E_z$, and $E_2^\perp = 0$. Plugging in (2.29) yields

$$\sigma = \frac{1}{4\pi} E_0 J_0(k_{\rho,01}\rho) \cos\omega t, \quad z = 0. \tag{2.32}$$

To maintain overall neutrality, the opposite charge must be on the $z = L$ surface, therefore

$$\sigma = -\frac{1}{4\pi} E_0 J_0(k_{\rho,01}\rho) \cos\omega t, \quad z = L. \tag{2.33}$$

To obtain the surface current distributions, we use the following boundary condition [33]:

$$\frac{1}{\mu_1}\mathbf{B}_1^\| - \frac{1}{\mu_2}\mathbf{B}_2^\| = \frac{4\pi}{c}\mathbf{K} \times \hat{\mathbf{n}},$$

where μ is the magnetic permeability, $\mathbf{B}^\|$ is the component of the magnetic field parallel to the surface, $\mathbf{K}$ is the surface current vector, and $\hat{\mathbf{n}}$ is the surface normal vector. At $z = 0$ we find that $\mu_1 = 1$, $\mathbf{B}_1^\| = B_\phi\hat{\boldsymbol{\phi}}$, $\mathbf{B}_2^\| = 0$, and $\hat{\mathbf{n}} = \hat{\mathbf{z}}$. Therefore,

$$K_\rho = \frac{c}{4\pi} E_0 J_1(k_{\rho,01}\rho) \sin\omega t, \quad z = 0; \tag{2.34}$$

where $\mathbf{K} = K_\rho\hat{\boldsymbol{\rho}} + K_\phi\hat{\boldsymbol{\phi}} + K_z\hat{\mathbf{z}}$. At $z = L$, $\hat{\mathbf{n}}$ flips and becomes $-\hat{\mathbf{z}}$, therefore,

$$K_\rho = -\frac{c}{4\pi} E_0 J_1(k_{\rho,01}\rho) \sin\omega t, \quad z = L. \tag{2.35}$$

Lastly, at $\rho = R$, $\hat{\mathbf{n}} = -\hat{\boldsymbol{\rho}}$. This yields,

$$K_z = \frac{c}{4\pi} E_0 J_1(k_{\rho,01}\rho) \sin\omega t, \quad \rho = R. \tag{2.36}$$

Together these surface currents support the time-varying build-up of charges on the end walls that was described earlier.

2.2.2 Electric and Magnetic Forces. The force per unit area (pressure) on the surface charge distribution is $\mathbf{f}_E = \frac{1}{2}\sigma\mathbf{E}$ [33]. Therefore, after plugging in (2.32), (2.33), and (2.29), the pressures due to the electric field are

$$\mathbf{f}_E = \frac{1}{8\pi}E_0^2 J_0^2(k_{\rho,01}\rho)\cos^2(\omega t)\hat{\mathbf{z}}, \quad z = 0, \tag{2.37}$$

$$\mathbf{f}_E = -\frac{1}{8\pi}E_0^2 J_0^2(k_{\rho,01}\rho)\cos^2(\omega t)\hat{\mathbf{z}}, \quad z = L. \tag{2.38}$$

For a given ρ, these pressures oscillate with time as a $\cos^2$ function. Thus the minimum magnitude is zero, and the maximum magnitude is $\frac{1}{8\pi}E_0^2 J_0^2(k_{\rho,01}\rho)$. Furthermore, the direction of the pressures is parallel to the surface normals. Therefore, this would give an average (over time) inward pressure of $\frac{1}{16\pi}E_0^2 J_0^2(k_{\rho,01}\rho)$ on the end walls.

The force per unit area on the surface current distribution is $\mathbf{f}_B = \frac{1}{2}\mathbf{K}\times\mathbf{B}$. This gives

$$\mathbf{f}_B = \frac{1}{2}(-K_z B_\phi \hat{\boldsymbol{\phi}} + K_\rho B_\phi \hat{\mathbf{z}}).$$

Plugging in (2.34), (2.35), and (2.30) we have,

$$\mathbf{f}_B = -\frac{c}{8\pi}E_0^2 J_1^2(k_{\rho,01}\rho)\sin^2(\omega t)\hat{\mathbf{z}} \quad z = 0, \tag{2.39}$$

$$\mathbf{f}_B = \frac{c}{8\pi}E_0^2 J_1^2(k_{\rho,01}\rho)\sin^2(\omega t)\hat{\mathbf{z}} \quad z = L, \tag{2.40}$$

$$\mathbf{f}_B = \frac{c}{8\pi}E_0^2 J_1^2(k_{\rho,01}\rho)\sin^2(\omega t)\hat{\boldsymbol{\rho}} \quad \rho = R. \tag{2.41}$$

For a given ρ, these pressures oscillate with time as a $\sin^2$ function. Similarly to the electric pressures, then, the minimum magnitude is zero, and the maximum magnitude is $\frac{c}{8\pi}E_0^2 J_1^2(k_{\rho,01}\rho)$. In contrast to the electric pressures, these magnetic pressures are all anti-parallel to the surface normals. This would add an average, outward pressure of $\frac{c}{16\pi}E_0^2 J_1^2(k_{\rho,01}\rho)$ on the inner cavity surfaces.

The pressure equations can be easily converted to the SI system by simply substituting $f_B \to c f_B$ and $E_0^2 \to 4\pi\epsilon_0 E_0^2$, where ϵ_0 is the electric permittivity of free space:

$$\mathbf{f}_E = \frac{\epsilon_0}{2} E_0^2 J_0^2(k_{\rho,01}\rho) \cos^2(\omega t)\hat{\mathbf{z}}, \quad z = 0, \tag{2.42}$$

$$\mathbf{f}_E = -\frac{\epsilon_0}{2} E_0^2 J_0^2(k_{\rho,01}\rho) \cos^2(\omega t)\hat{\mathbf{z}}, \quad z = L, \tag{2.43}$$

$$\mathbf{f}_B = -\frac{\epsilon_0}{2} E_0^2 J_1^2(k_{\rho,01}\rho) \sin^2(\omega t)\hat{\mathbf{z}} \quad z = 0, \tag{2.44}$$

$$\mathbf{f}_B = \frac{\epsilon_0}{2} E_0^2 J_1^2(k_{\rho,01}\rho) \sin^2(\omega t)\hat{\mathbf{z}} \quad z = L, \tag{2.45}$$

$$\mathbf{f}_B = \frac{\epsilon_0}{2} E_0^2 J_1^2(k_{\rho,01}\rho) \sin^2(\omega t)\hat{\boldsymbol{\rho}} \quad \rho = R. \tag{2.46}$$

2.2.3 Electromagnetic Pressure Estimates. It is informative to have an estimate of the pressures applied to the inner cavity walls by the electromagnetic forces of the RF pulse. This can be, for example, compared with the pressure estimate from a breakdown spark to give an idea of the relative loudness of RF hammer and breakdown events. Given a moderate accelerating gradient of 10 MV/m in an 805 MHz cavity, the maximum inward pressure (at the axis where $J_0(0) = 1$) on the end walls would be

$$f_E = \frac{1}{2}(8.854 \times 10^{-12} \text{ F/m})(10 \times 10^6 \text{ V/m})^2 = 443 \text{ Pa}. \tag{2.47}$$

To calculate the maximum outward pressure (close to the inner radius) on the end walls, we can numerically optimize for the first maximum of $J_1(x)$. This yields a value of 0.5819, and we have

$$f_B = \frac{1}{2}(8.854 \times 10^{-12}\text{F/m})(10 \times 10^6 \text{ V/m})^2(0.5819)^2 = 150 \text{ Pa}. \tag{2.48}$$

Noting that $R = z_{01}/k_{\rho,01} = z_{01}c/\omega = 14.3$ cm ($z_{01} = 2.405$) for an 805 MHz pillbox cavity, the net pressure on the end walls versus radius is plotted in Fig. 2.2.

The outward pressure on the inner circumference wall can easily be calculated by plugging in the first zero of $J_0(x)$ (at the circumference, $x = k_{\rho,01}R = z_{01}$) into
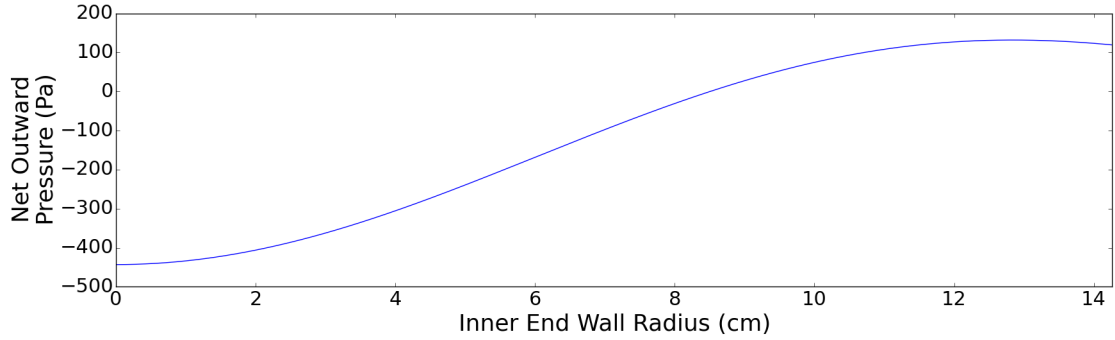
Figure 2.2. Net electromagnetic, outward pressure vs. radius on the end wall of a pillbox cavity designed for 805 MHz and operated at 10 MV/m. The center of the end wall is pulled inward while a lighter outward pressure is applied out near the inner radius.

$J_1(x)$ (yielding a value of 0.5191) and evaluating (2.46):

$$f_B = \frac{1}{2}(8.854 \times 10^{-12}\text{F/m})(10 \times 10^6 \text{ V/m})^2(0.5191)^2 = 119 \text{ Pa.} \qquad (2.49)$$

The other common size of cavity in the MTA is the 201 MHz pillbox (a.k.a. the MICE cavity). For comparison with 805 MHz cavities, an ideal 201 MHz pillbox would have a radius of 57.1 cm. At the same gradient (10 MV/m), the larger size does not change the values of the maximum electric and magnetic pressures. Figure 2.2 would be identical except for a re-scaled x-axis that runs from 0 cm to 57.1 cm. Therefore, a larger cavity would produce different RF hammer sound levels only if the cavity wall properties (e.g., thickness or material type) were different.

2.3 Breakdown Sound Production

The primary source of sound during breakdown in vacuum cavities is Joule heating as the plasma electrons are absorbed by the end wall towards which they are accelerated (see Sec. 1.4). To see this, it is informative to make an estimate of the current that is required to cancel all of the charge on one end wall of an 805 MHz pillbox cavity. We start by integrating the surface charge (2.32) over the inner end

wall surface to get the total surface charge, q_0 (in SI units):

$$q_0 = \int_{\phi=0}^{2\pi} \int_{\rho=0}^{R} \sigma(\rho) \rho \, d\rho \, d\phi = 2\pi\epsilon_0 E_0 \int_{\rho=0}^{R} \rho J_0(k_{\rho,01}\rho) d\rho = 2.4 \times 10^{-6} \text{ C}, \qquad (2.50)$$

where $R = 14.3$ cm.

Fig. 2.3 shows a typical breakdown pulse from one of the MTA's 805 MHz cavities. The maximum current from this decay, $j_{\max}$, can be estimated by first determining the time constant of the exponential decay of this pulse, τ. The time constant is the time it takes to reach $1/e$ times the maximum voltage, V_0. If we read the maximum voltage from the plot as $V_0 = 0.46$ V, then we have

$$\tau = 0.43 \text{ µs}. \qquad (2.51)$$

Since the charge on the endplates is proportional to the voltage, modeling the steepest region of the decay in Fig. 2.3 linearly yields a simple relation for finding the approximate maximum current during breakdown. Using the value of q_0 from (2.50) and τ from (2.51),

$$j_{\max} = \frac{\Delta q}{\Delta t} = \frac{q_0 - q_0/e}{\tau} = \frac{(2.4 \times 10^{-6} \text{ C})(1 - 0.3679)}{0.43 \text{ µs}} = 3.5 \text{ A}. \qquad (2.52)$$

To get a better idea of the sound generating capacity of this breakdown spark current, we can estimate the thermal expansion pressure caused by dumping the stored energy of the cavity into a small portion of the end wall of, say, a cube with sides each 1 cm long. The stored energy, U, can be found by integrating over the maximum electric field (the magnetic field is zero at this point in the RF cycle and does not contribute) [33]:

$$U = \frac{1}{2}\epsilon_0 \int |\mathbf{E}|^2 dV = \pi L \epsilon_0 E_0^2 \int_{\rho=0}^{R} \rho J_0(k_{\rho,01}\rho) d\rho = 1.0 \text{ J}, \qquad (2.53)$$

where L is taken to be 13 cm, the length of one of the 805 MHz cavities used in the MTA (see Sec. 1.6.3). We then calculate the temperature rise, ΔT, in 1 cm^3 of copper,
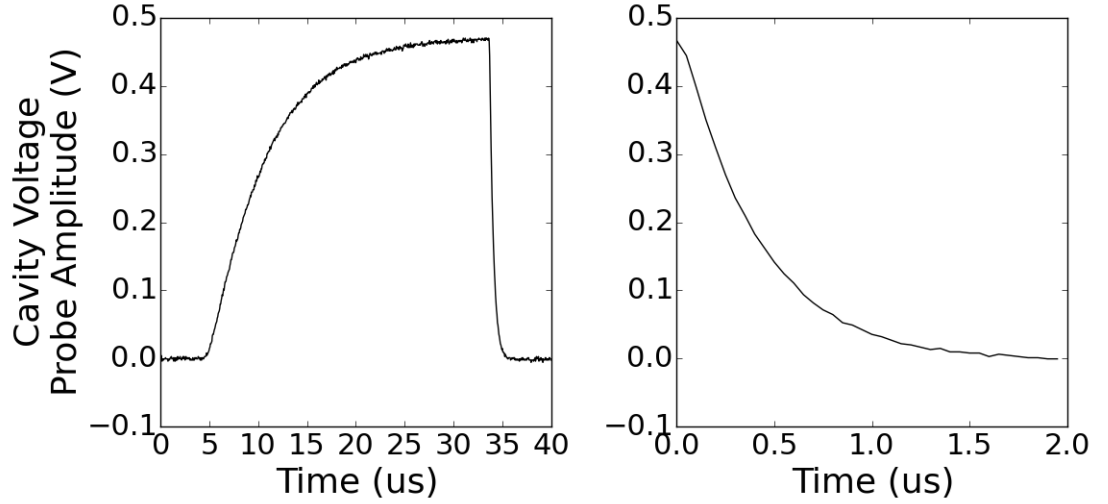
Figure 2.3. An example breakdown pulse (left) and a zoomed-in view of the decay portion of the pulse.

noting that the specific heat per mass and density of copper are $c_m = 0.385$ J/g $\cdot$ K and $D = 8.96$ g/cm^3 respectively [35]:

$$\Delta T = \frac{U}{c_m m} = \frac{U}{c_m DV} = \frac{1.0 \text{ J}}{(0.385 \text{ J/g} \cdot \text{K})(8.96 \text{ g/cm}^3)(1 \text{ cm}^3)} = 0.29 \text{ K}, \quad (2.54)$$

where V is the volume of the cube. Using the coefficient of linear expansion for copper, $\alpha = 16.5 \times 10^{-6}$ K^{-1} [35], we can calculate how much the cube would expand in one dimension, Δx, from an initial length x_0 if unconstrained by surrounding material [36]:

$$\Delta x = \alpha x_0 \Delta T. \quad (2.55)$$

The pressure, P, on the surrounding material can be found by calculating the force per unit area required to compress the expanded cube back to its original size [36]

$$P = \frac{F}{A} = \frac{AE}{Ax_0} \Delta x = \frac{E}{x_0} \alpha x_0 \Delta T = \alpha E \Delta T$$

$$= (16.5 \times 10^{-6} \text{ K}^{-1})(120 \times 10^9 \text{ Pa})(0.29 \text{ K}) = 574 \text{ kPa}, \quad (2.56)$$

where A is the cross-sectional area of the cube and $E = 120 \times 10^9$ Pa is the modulus of elasticity of copper [35]. Comparing to Fig 2.2, this pressure is three orders of

magnitude larger than the pressures that create the RF hammer sound. Even if this estimate is somewhat high, this suggests that it should be possible to detect the breakdown spark acoustic signal over the RF hammer background noise.

For the production 201 MHz MICE cavity, a similar RF trace as in Fig. 2.3 can be used to empirically determine the decay constant for this cavity: $\tau = 1.07$ µs. This cavity would have a stored energy of 11 J, and the thermal shock pressure would be 6 MPa. This is four orders of magnitude greater than the RF hammer pressure, therefore it should be even easier to separate the RF hammer background noise from the breakdown thermal shock signal in a 201 MHz cavity.

CHAPTER 3

HARDWARE

In addition to microphones, capturing sound from the surface of RF cavities for digital analysis requires several pieces of hardware. If the microphones are passive (i.e., they have no on-board amplifier), amplifiers are needed to ensure good signals. Before software can process the signals, a Data Acquisition (DAQ) device must be used to convert the analog signals into digital data. This chapter describes the hardware used for this research. A discussion of microphone mounting methods is also included.

3.1 Microphones

Two main types of piezoelectric crystal microphones were designed and built for these experiments. Originally only active microphones with built-in, fixed-gain amplifiers were used. Due to working with the MICE cavity that had additional operating constraints, a new design containing just the piezoelectric crystal was later requested. The piezos are all APC International type 850 crystals, but sizes vary among the various versions. All of these microphones were designed and built by Ilan Levine's group at Indiana University South Bend (IUSB) with design input from MTA people and some assembly assistance from the author.

3.1.1 Active Microphones. The first set of eight microphones were built with 20 mm diameter and 8.7 mm thick piezos with an on-board amplifier. The schematic can be found in App. B. The piezo was solder-pasted to a sheet of copper foil. More copper foil was used to form a cylindrical Faraday cage. Loctite five-minute epoxy was used to pot the internal space as well as coat the external surface for structural integrity.

These microphones were designed originally for the PICASSO dark matter experiment. The requirements for that experiment were to listen for the shockwave that accompanies the liquid-to-vapor transition in a superheated liquid [37]. This dictated the gain of 225 with which the on-board amplifiers were configured. Two 4.5 m sets of unshielded, twisted pair wires were used to supply power and receive the signals (Fig. 3.1, left).
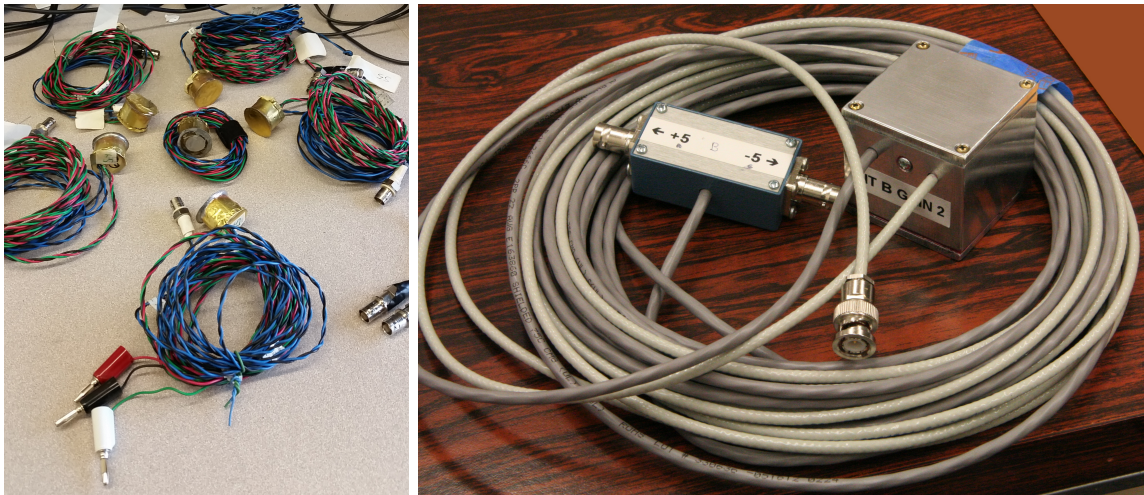
Figure 3.1. Active microphones. Left: PICASSO, gain 225 microphones used primarily on the HC. Right: one of the low gain microphones used on the ASC [29].

A second set of three active microphones were built in response to an acoustic localization experiment where the signals were severely clipping, and a test was needed to determine whether this was due to acoustic or electromagnetic noise. Two of these had gains of 2, and the third had a gain of 37. They were built with spare piezos with a diameter of 30 mm and a thickness of 8.7 mm. A 2.5 cm cubic Faraday cage was used and potted with epoxy. Two shielded twisted pair cables were used. One cable supplied power and terminated in a small distribution box that had two BNC connectors for $+V$ and $-V$. The second cable was for signal and terminated in a

single BNC connector (Fig. 3.1, right). These were too big for practical use and, after the noise test, were never employed again.

3.1.2 Passive Microphones. As mentioned in Sec. 1.5, a production cavity built for MICE was tested in the MTA [38]. It was designed to be operated inside a vacuum vessel within which no electrical power was allowed. The cavity was also operated at times in a magnetic field of up to 0.4 T. This led to the requirements that only passive microphones with non-magnetic components and low-outgassing adhesives could be used for experiments involving this cavity.
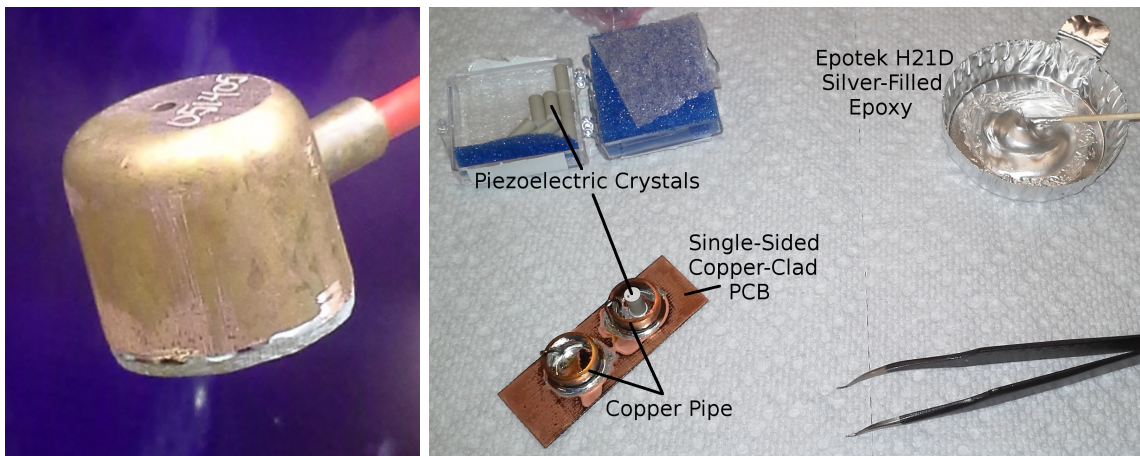
Figure 3.2. Left: a third-generation passive microphone. Right: fixing piezos in microphone bases with silver epoxy [39].

Several batches of passive microphones were built (Fig. 3.2, left). One of these microphones was subjected to a 100°C bake test with a residual gas analyzer to verify that no volatiles would be released inside the vacuum vessel. In addition, the vacuum vessel was also heated to approximately 100°C and monitored with ion gauges. There was no indication that the microphones were releasing volatiles.

The base of these microphones was a 1/4 inch length of 3/8 inch (inner diameter) copper pipe solder-pasted to a 1/16 inch thick piece of single-sided, copper-clad PC board (PCB). A 4 mm diameter, 8.7 mm long piezo was then attached to the

PCB in the center of the copper pipe section using Epotek H21D silver-filled epoxy (Fig 3.2, right).

A 50 cm length of Belden 88641 shielded twisted pair cable was used to carry the signal to appropriate connectors. Various methods were tried for attaching the positive signal lead to the exposed end of the piezo. Too much heat during soldering caused the thin, silver pad on the piezo to delaminate and lift off, making it infeasible to attach the signal lead. The technique that worked the best was to first silver-epoxy a thin strip of copper foil to the silver pad. A drop of solder was then added to the foil which was used to solder the signal lead to the foil. The foil was trimmed before continuing with assembly. To attach the negative signal lead, a wire stub was silver-epoxied to the inside of the copper pipe section. The signal lead was then soldered to the wire stub. See Fig. 3.3 for reference.
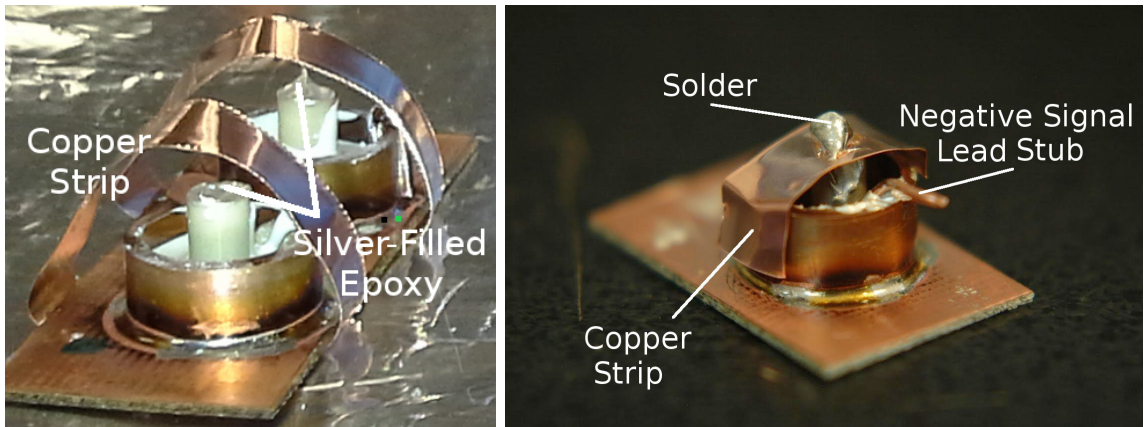
Figure 3.3. A copper strip is silver-epoxied to the piezo's silver terminal pad (left) to avoid delamination of the pad during signal lead soldering (right) [39].

With the signal leads connected, the copper pipe base was potted with Loctite Hysol 1C epoxy for structural integrity and to electrically insulate the positive signal lead from the microphone cover. A 3/8 inch copper pipe cap with a notch for fitting around the protruding signal leads was placed around the potted assembly. More silver epoxy was used to secure the base to the cap while also closing the Faraday

cage electrically. Small holes were drilled in the top of the cap to make sure no gas was trapped inside the enclosure.

In early models, the ground conductor of the signal cable was secured near the signal lead notch with silver epoxy. In later models, a 1/8 inch (inner diameter) copper tube was threaded into a hole in the copper pipe segment, acting as both a strain relief and the point of contact for the ground lead.

3.1.3 Mounting Method. Initially, double-sided tape was naïvely employed to mount the active microphones on cavities. Different types of epoxy have been used in all subsequent experiments with both the active and passive microphones. This change in mounting method was driven primarily by the requirement that the adhesive be low-outgassing for use in a vacuum vessel. NASA provides lists of low-outgassing materials for use in spacecraft [40]. Only one cavity in fact needed low-outgassing epoxy, so all of the other microphone installations that did not use double-sided tape used 5-minute epoxy. The faster cure time of 5-minute epoxy (low-outgassing epoxies had cure times on the order of 24 hours) obviated the need for clamps.

To better understand how various mounting methods may affect the received acoustic signals, a bench test was performed on the effectiveness of sound transmission using three mounting methods. A makeshift apparatus was created that allowed the dropping of a short aluminum rod with a rounded end cap repeatedly from the same height onto an aluminum surface. A microphone attached to the surface picked up the sound generated by the impact.

The drop test apparatus was built by clamping one of the Aluminum Mock Cavity's (AMC) disks on top of the AMC ring (see Sec. 1.6.1 for more information about the AMC). High tensile strength cord was tightly strung between the clamps to create a platform support. A 9/32 inch inner diameter aluminum grommet was

clamped between two square wooden dowels to form a guide platform. A die set was used to add UNC 1/4-20 threads to one end of a 13 cm length of 1/4 inch diameter aluminum rod. A 1/4-20 cap nut was screwed onto the rod to provide a smooth, rounded impact head. Figure 3.4 shows the apparatus with the rod inserted into the guide grommet.
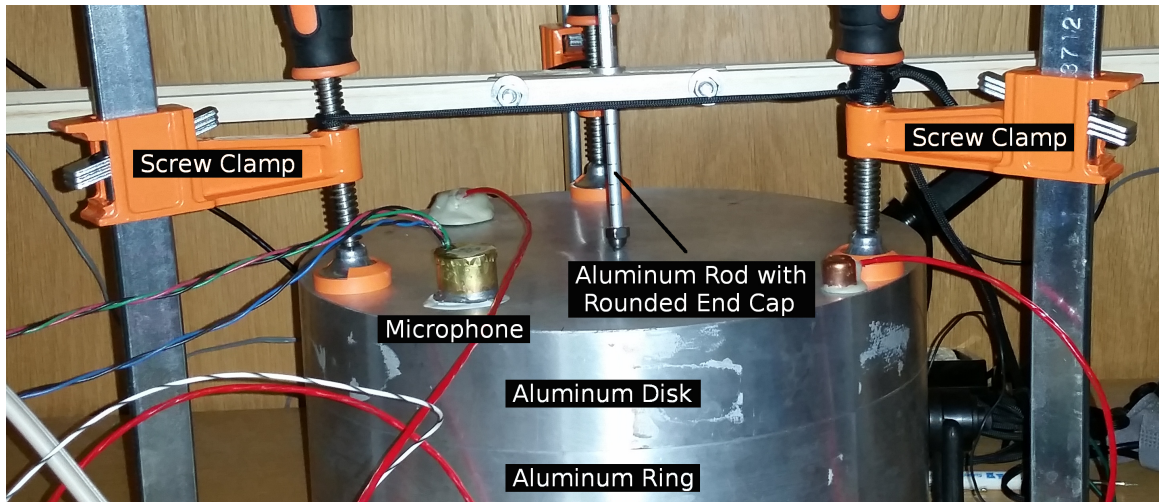
Figure 3.4. Aluminum Mock Cavity drop test setup.

The three tested mounting methods were double-sided tape, hot glue, and five-minute epoxy. Two experiments were performed whereby the rod was dropped ten times for each mounting method. The experiments differed significantly only in the cure time for the epoxy (10 minutes and 20 hours respectively). This was done to test the hypothesis that the properties of the epoxy could change as the epoxy cures. The drop test apparatus was operated by pulling the rod up through the grommet until the nut cap touched the grommet. The rod was released and caught after it bounced to create a single impact event. The peak voltage of the microphone's signal was measured using an oscilloscope. The results are quantified in Fig. 3.5.
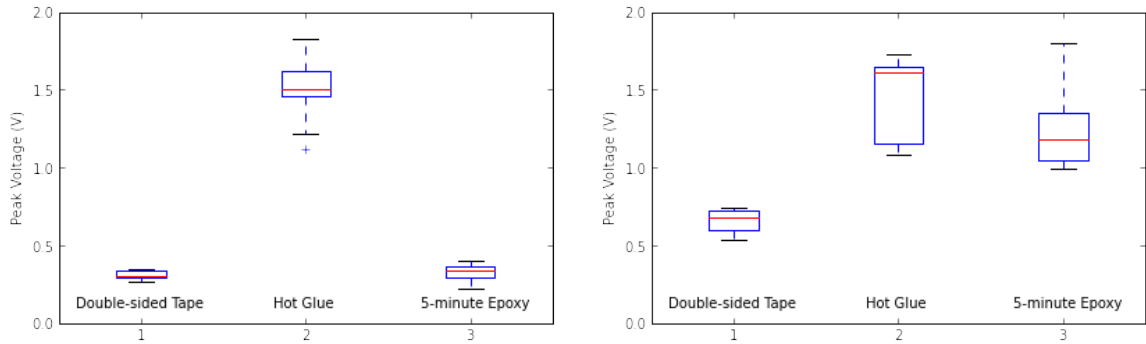
Figure 3.5. Box plots of mounting method damping effects tests. Each box (blue) represents the interquartile range (IQR) or middle 50% of the data. The bottom of the box is the first quartile (Q1), the top is the third quartile (Q3), and the red line is the second quartile (Q2) or median value. Black lines delimit the range outside of which data points are considered outliers (i.e., less than $Q1 - 1.5(IQR)$ or greater than $Q3 - 1.5(IQR)$. Epoxy cure times for the two experiments were 10 minutes (left) and 20 hours (right). Each experiment consisted of 10 trials per mounting method.

The double-sided tape method provided poor sound transmission. In addition, the transmission from one application of the tape to the other appeared to be highly variable.

The 5-minute epoxy method varies greatly with respect to cure time. After 5 minutes, the sound transmission of the epoxy was similar to that of the double-sided tape. If allowed to cure for about a day, the sound transmission of 5-minute epoxy was nearly as good as hot glue.

Finally, hot glue clearly outperformed the other methods. It had the best sound transmission and was the most consistent over time and repeated applications. Considering also its ease of use, hot glue should be considered as an alternative to epoxy in situations where outgassing is not an issue.

3.2 Data Acquisition

Two different data acquisition systems have been used. The first consisted of a National Instruments (NI) PCI-5105 card installed in a PC. This card had 8 simultaneously sampled, 12-bit analog inputs with a maximum voltage input range of ± 15 V. Although it supported sampling rates up to 60 MHz, performance and disk space issues meant that it was typically operated at only 500 kHz. This was because the card was primarily operated in continuous sampling mode during early exploration of unknown acoustic signals. The highest sampling rate that could be indefinitely sustained in this mode on was 500 kHz. This required close monitoring since the hard drive quickly filled up. Often the cavity being tested would be operated close to the maximum gradient where sparks were happening frequently. The acquisition software would be run for only a short period of time and then would be shut off while the data were analyzed.

When it was decided to instrument the production MICE cavity, new hardware was obtained to accommodate the large number of passive microphones that were planned to be placed on the cavity. This new system consisted of a National Instruments cDAQ-9188 DAQ chassis with four 9221 modules (Fig. 3.6). The modules each had 8 analog inputs with a maximum sampling rate of 800 kHz. Each card shared a single digitizer among its inputs, therefore the maximum sampling rate was only achievable when using a single channel per card. Using all eight channels brought the per-channel sampling rate down to 100 kHz. By dividing the slowest speed of sound in the cavity material, 1.8×10^5 cm/s (see Sec. 5.3.2), by the eight-channel sampling rate (100 kHz), a theoretical maximum spatial resolution of 1.8 cm could be inferred. Given that the cDAQ chassis with four analog input cards cost about as much as a single PCI-5105 card and the MICE cavity has a maximum diameter of 120 cm, this lowered resolution was judged to be an acceptable trade off in cost versus resolution.
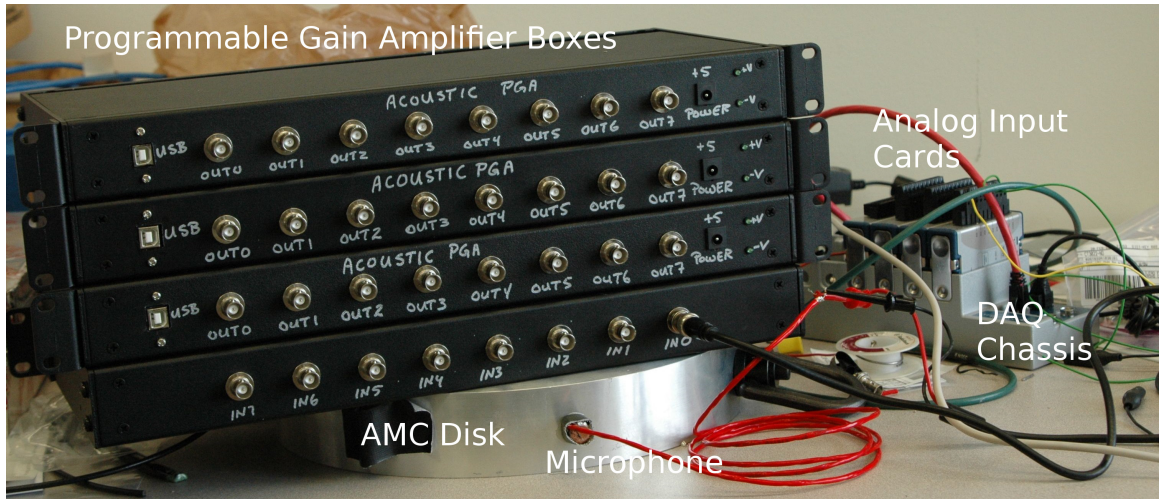
Figure 3.6. Programmable gain amplifiers and NI cDAQ-9188 DAQ chassis with four NI 9221 analog input cards.

3.2.1 Amplification. The change to passive microphones necessitated that they be augmented with external amplifiers. An 8-channel, programmable gain amplifier (PGA) board was designed for this purpose (for the schematic, see Fig. B.2). Six of these boards were fabricated along with four enclosures (Fig. 3.6). A USB connection to a microcontroller on the boards allows the configuration of the PGA chip with a gain of 1, 2, 5, or 10. The output of each PGA chip is conditioned with a 50 kHz (Nyquist frequency at 100 kHz sampling) low pass, anti-aliasing filter, and also provides a base gain of 20 that multiplies the selected gain of the PGA chip. This gives an overall set of selectable gains of 20, 40, 100, and 200. The base gain was chosen so that the maximum amplification was approximately the same as that of the active microphones.

3.2.2 Pre-Amplification. It was suspected that the cable might be too long to produce a large enough signal for the PGA. Several cavity tests with just the passive microphones and the PGA confirmed this. To understand the issue better, a test to quantify the degradation of the signal with cable length was performed. This was

done using the drop test rig that was built for the microphone mounting method tests (Sec. 3.1.3).

Several identical passive microphones with varying lengths of cable were hot glued to the same location on the cavity in turn. The rod was then dropped ten times for each cable length, and the average peak voltage was measured from the resultant acoustic signal on an oscilloscope. The blue data points and traces in Fig. 3.7 summarize the results. The following analysis explains the results in terms of charge being spread out along the cable.
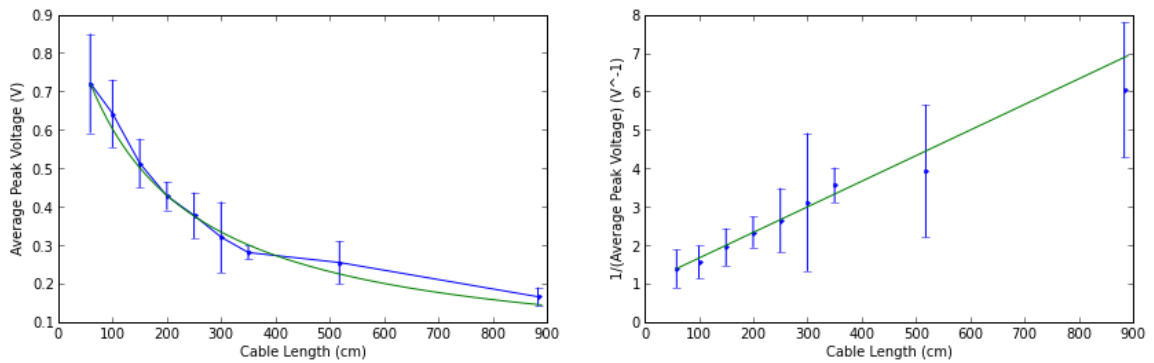
Figure 3.7. Left: Peak voltage measurements for various lengths of microphone cable (blue) and reciprocal of fit from right plot (green). Right: $1/V_{\mathrm{pp}}$ vs cable length (blue) with 1st order, weighted, least squares fit (green).

Neglecting the resistive and inductive components that become significant only at high frequencies, the piezo can be modeled as a driven capacitor. The internal microphone signal leads and Faraday cage add a capacitance C_f in parallel to the piezo capacitance. The piezo is a charge source, so the charge created from the deformation of the piezo will be mirrored at the output of the microphone. Thus with no cable attached to the internal signal leads, the microphone will present a certain charge (Q) and voltage (V_m) with $V_m = Q/C_m$, where $C_m = C_p + C_f$ and C_p is the capacitance of the piezo crystal alone.
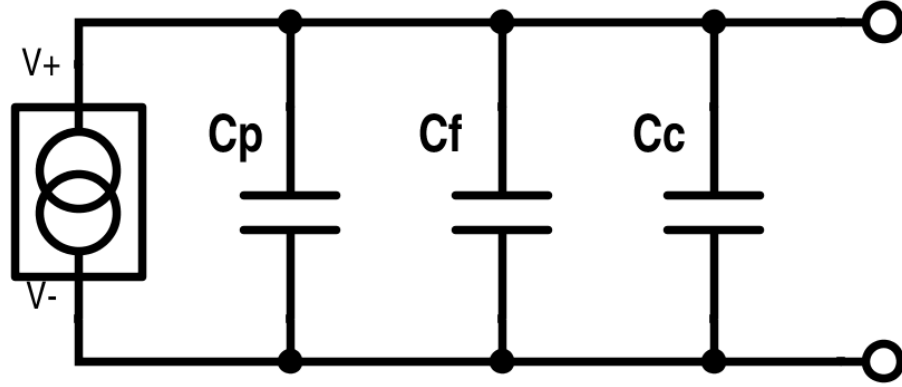
Figure 3.8. Equivalent circuit of a piezoelectric crystal, microphone Faraday cage, and signal cable.

The cable has some capacitance $C_c = C_L L$, where C_L is the capacitance per unit length, and L is the cable length. The capacitances of the microphone and the cable will also add in parallel (Fig. 3.8). Furthermore, the same charge that is seen at the output of the microphone is distributed along the signal cable. Therefore, the voltage seen at the PGA input is

$$V_{\text{in}} = \frac{Q}{C_m + C_L L}. \tag{3.1}$$

Taking the reciprocal of V_{in} then gives an equation for a line as a function of cable length:

$$y = \frac{1}{V_{\text{in}}} = \frac{C_m}{Q} + \frac{C_L}{Q} L = b + mL. \tag{3.2}$$

A first order, least squares polynomial fit weighted by the standard deviations of the average peak voltages can then be performed (see code in Appendix D) on the reciprocal of the measured peak voltages. This yields values for the y-intercept, $b = C_m/Q$, and the slope, $m = C_L/Q$, of a line that is shown in green on the right plot of Fig. 3.7. The reciprocal of this line is shown in green on the left plot of Fig. 3.7.

To use this model to predict amplifier input voltages, several capacitance values must be obtained. The microphone and cable capacitances were measured using

a digital multimeter (Extech Instruments EX330) with an accuracy stated by the manufacturer of $\pm(3.5\% + 40$ digits) [41]. The uncertainty percentage is the accuracy of the meter's measurement circuit, while the uncertainty digits value is the accuracy of the meter's analog to digital converter (ADC). The meter was tested on several capacitors and found to be accurate within the capacitor tolerances and meter uncertainties (Table 3.1).

Table 3.1. Extech Instruments EX330 Capacitance Test Measurements

Capacitor Value (pF)	Tolerance (%)	Measurement (pF)
180	10	189 ± 47
15,000	20	$12,170 \pm 830$
100	10	108 ± 44
221	5	221 ± 48
3,300	10	$3,044 \pm 140$

For a 122.5 ± 0.5 cm length, L, of Belden 88641 (the cable used with the microphones), the capacitance, C_c, was 182 ± 46 pF measured between the twisted pair conductors with one of the conductors connected to the shield. This gave a capacitance per length of

$$C_L = \frac{C_c}{L} = 1.5 \pm 0.4 \text{ pF/cm,} \tag{3.3}$$

where the uncertainty on C_c was propagated in quadrature to give the uncertainty on C_L. Belden 88641 had a nominal capacitance per length (conductor to other conductor and shield) of 1.9 pF/cm [42]. Therefore, the nominal cable capacitance is just inside of the propagated uncertainty of the measurement.

The measured microphone capacitance was $C_m = 44.0 \pm 42$ pF. APC International, Ltd. quoted the capacitance of an 8.7 mm thick, 4.0 mm diameter disk of their type 850 piezo material (see Sec.3.1.2) as having a capacitance of $C_p = 24.3$ pF [43]. The nominal piezo capacitance is well within the uncertainty of the measurement of the microphone capacitance. In other words, the microphone capacitance was indistinguishable from the piezo capacitance within the accuracy of the meter.

Data were obtained using the active microphones (see Sec. 6.2) suggested that the peak voltage of a typical acoustic wavefront signal was around 2.2 V (for example, Fig. 3.9). The gain of the on-board amplifiers is 225, therefore a typical active microphone peak piezo voltage should be approximately

$$V_p = \frac{2.2 \text{ V}}{225} = 9.8 \text{ mV}. \tag{3.4}$$

The active and passive microphone piezos are made of the same material and have the same unstrained height, h. Therefore, for a given pressure, the piezos will be compressed an equal distance [35],

$$Y = \frac{Fh}{A\Delta h} = \frac{Ph}{\Delta h} \quad \longrightarrow \quad \Delta h = \frac{Ph}{Y}, \tag{3.5}$$

where F is the compressional or extensional force on the piezo, Y is the Young's modulus, A is the cylindrical cross-sectional area, and $P = F/A$ is the sound pressure. This implies that the voltage, V_p, is the same for the passive microphone piezo. To wit, for a particular piezoelectric material, the charge constant, d_{33}, is defined as [44]

$$\epsilon = d_{33}E, \tag{3.6}$$

where $\epsilon = \Delta h/h$ is the axial strain on the piezo, and E is the electric field between the two ends of the piezo. Assuming a uniform electric field, the voltage between the ends of the piezo is approximately $V_p = hE$ ($h \gg \Delta h$), therefore

$$\frac{\Delta h}{h} = d_{33}\frac{V_p}{h} \quad \longrightarrow \quad \Delta h = d_{33}V_p. \tag{3.7}$$

Plugging (3.5) into (3.7) gives the following expression for V_p that is independent of A,

$$V_p = \frac{hP}{d_{33}Y}. \tag{3.8}$$

With a 600 cm cable connecting a passive piezo to the PGA, a typical peak voltage seen at the input to the PGA would be

$$V_{\text{in}} = \frac{V_p C_p}{C_m + C_L L} = \frac{(9.8 \text{ mV})(24.3 \text{ pF})}{44.0 \text{ pF} + (1.5 \text{ pF/cm})(600 \text{ cm})} = 0.25 \text{ mV}, \tag{3.9}$$

This would produce a peak amplified signal of

$$V_{\text{out}} = V_{\text{in}} G_{\text{PGA}} = (0.25 \text{ mV})(200) = 50 \text{ mV}, \tag{3.10}$$

where G_{PGA} is the maximum PGA gain.
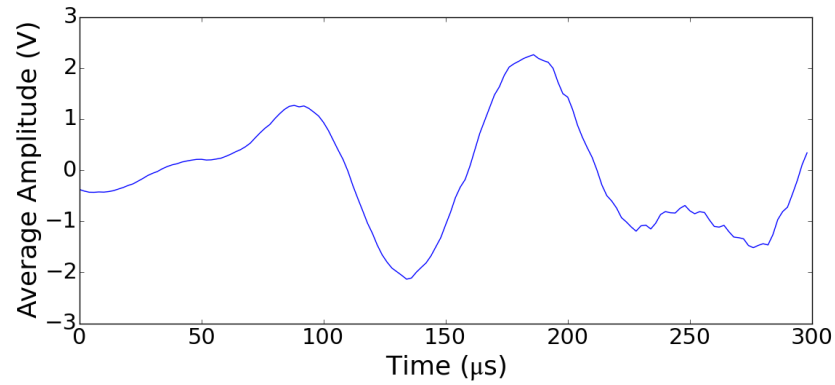
Figure 3.9. Typical HC Spark Signal Wavefront.

To maintain as good a signal-to-noise ratio as possible as well as the flexibility of the variable gain of the PGA, a voltage buffer pre-amplifier with a low gain of 10 was designed that could be used in the strong magnetic field close to the cavity (Fig. 3.10). The schematic is included in App. B. Assuming 1 m of signal cable is required in order to reach the pre-amp, the voltage at the input to the pre-amp would then be,

$$V'_{\text{in}} = \frac{V_p C_p}{C_m + C_L L} = \frac{(9.8 \text{ mV})(24.3 \text{ pF})}{44.0 \text{ pF} + (1.5 \text{ pF/cm})(100 \text{ cm})} = 1.2 \text{ mV}. \tag{3.11}$$

The voltage seen at the input to the PGA would be,

$$V_{\text{in}} = V'_{\text{in}} G_{\text{pre-amp}} = (1.2 \text{ mV})(10) = 12 \text{ mV}, \qquad (3.12)$$

where $G_{\text{pre-amp}}$ is the pre-amplifier gain. After amplification by the PGA, the 805 MHz cavity acoustic signal would then have a peak voltage of

$$V_{\text{out}} = V_{\text{in}} G_{\text{PGA}} = (12 \text{ mV})(200) = 2.4 \text{ V}. \qquad (3.13)$$
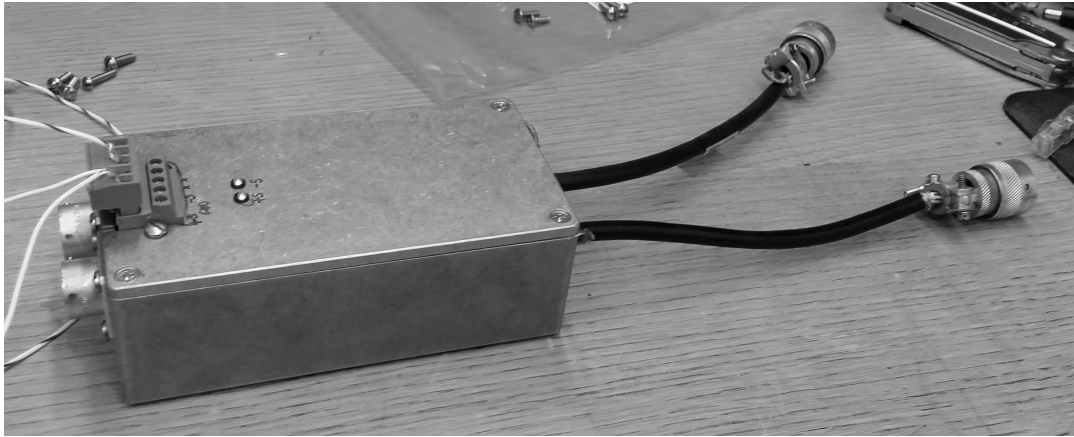
Figure 3.10. Microphone 8-channel pre-amplifier with a fixed gain of 10.

This is a 48-fold increase in V_{out} compared with the estimated value before the pre-amp was added (see (3.10)). In terms of impedance, the pre-amp converted the high output impedance of the piezo into a low input impedance at the PGA, preserving a stronger signal for amplification.

Considering that the 201 MHz MICE cavity was larger and stored more energy, it was expected to be approximately an order of magnitude louder than the 805 MHz cavity (see Sec. 2.3). Therefore the PGA ought to produce a peak output voltage of 24 V at a gain of 200. This is beyond the maximum output voltage of 5 V for the PGA board, but a gain of 20 ought to produce a peak voltage similar to what was seen using the active microphones (around 2.4 V). Therefore, this amplifier configuration should allow for comfortable amplitude ranges for passive microphone signals on both

805 MHz and 201 MHz cavities. Figure 3.11 shows a diagram of the final hardware layout.
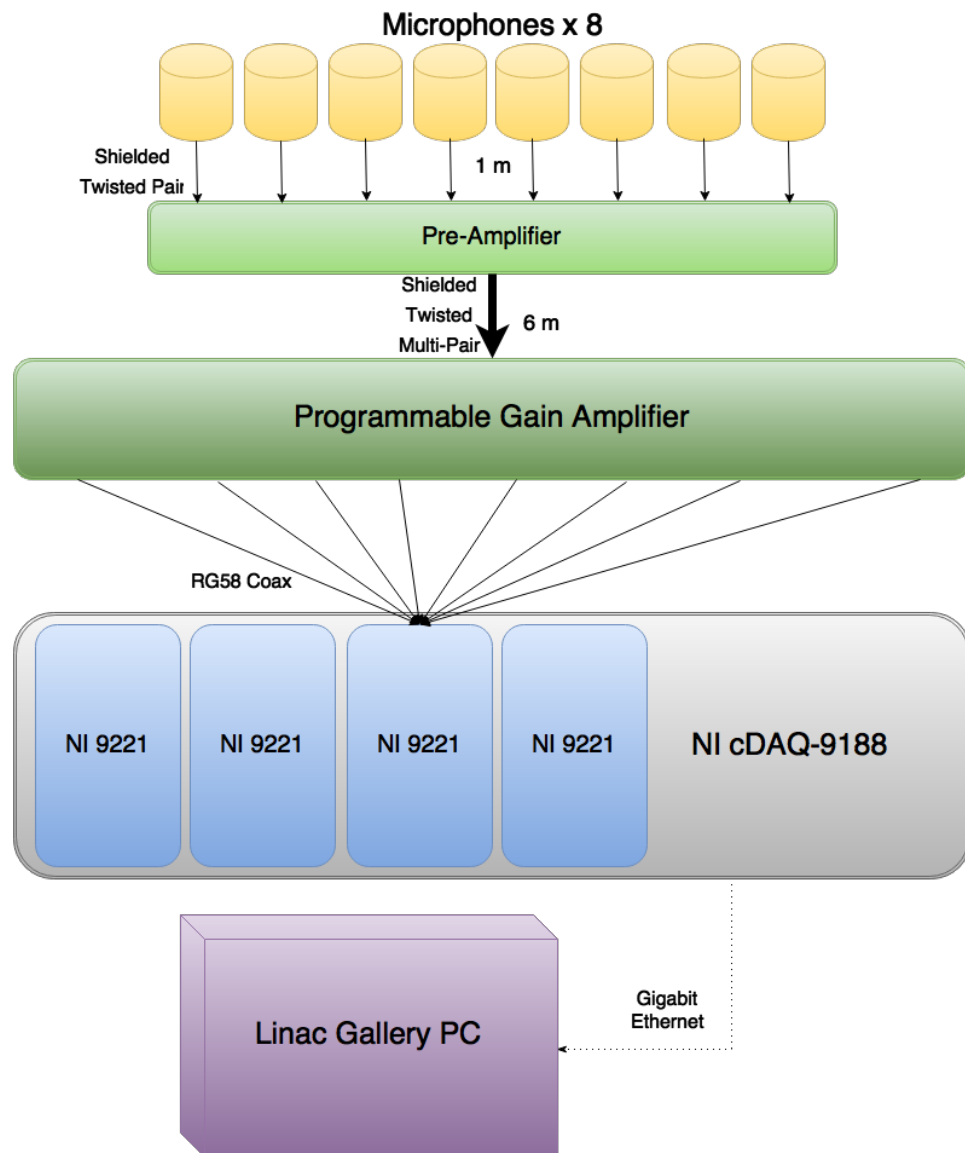
Figure 3.11. Hardware layout diagram showing the flow of acoustic signal data from the microphones at the top to the DAQ computer at the bottom. Arrow labels denote the type of cable used.

CHAPTER 4

SOFTWARE

National Instruments' LabVIEW software is a visual programming environment that is intended for use with their line of data acquisition hardware. LabVIEW was used to build several applications, the most important being the one that retrieved, archived, processed, and displayed information about the acoustic signal data. Various off-line helper applications were built to explore the data and test algorithms. Finally, though not directly related to data acquisition, several applications to support cavity run shifts were built as well. Details about some of these applications are provided here. See App. C for a list of acronyms used in the following sections and elsewhere.

4.1 Breakdown Acoustics Controls Application

The application that controls acquisition and archives, processes, and displays information about the acoustic signal data is referred to as the Breakdown Acoustics Controls (BAC). Figure 4.1 shows a screen shot of the BAC. A functional diagram (Fig. 4.2) is included as an aid for the following discussion. At the highest level, the BAC comprises several producer and consumer loops that pass information either via queues or notifications. The DAQ Producer (DPL) is triggered by the RF turn-on TTL (Transistor-Transistor Logic) signal. This signal is routed from the MTA control area to the experiment hall where it connects to one of the DAQ chassis's Programmable Function Input (PFI) ports. The DPL fetches 5120 samples (51.2 ms at 100 kHz sampling rate) of data from each of the channels and sends a notification to the Trigger Detector Loop (TDL). Depending on whether a breakdown has occurred or not, the TDL either places the data chunk on the Spark Trigger Queue (STQ) or

the RF hammer Queue (RHQ) for processing in the background by the Archive and Analysis Loop (AAL).
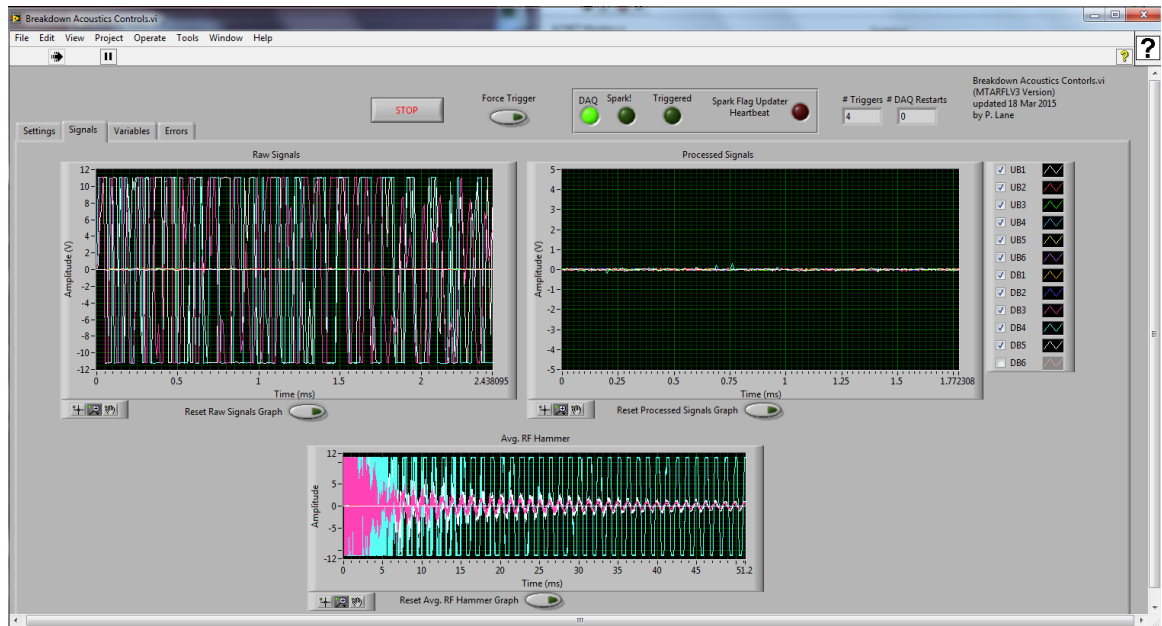
Figure 4.1. Breakdown Acoustics Controls screenshot.

Breakdown is detected by the main cavity control application written by Dave Peterson (also using LabVIEW) by analyzing light levels and voltage pickup characteristics. This state is monitored and periodically sent as Universal Datagram Packets (UDP) to the BAC via a separate application written by the author called the UDP Heartbeat Transmitter (UHT). The UDP Heartbeat Monitor (UHM) loop in the BAC processes notifications from the UHT. If breakdown has occurred, the TDL tags the data received from the DPL to indicate to the AAL that the data should be logged as coming from an actual breakdown event. This is necessary because the application user may press the "Force Trigger" button to force the AAL to treat the next set of signal data as spark data primarily for the purposes of archiving a sample of the current RF hammer signals.
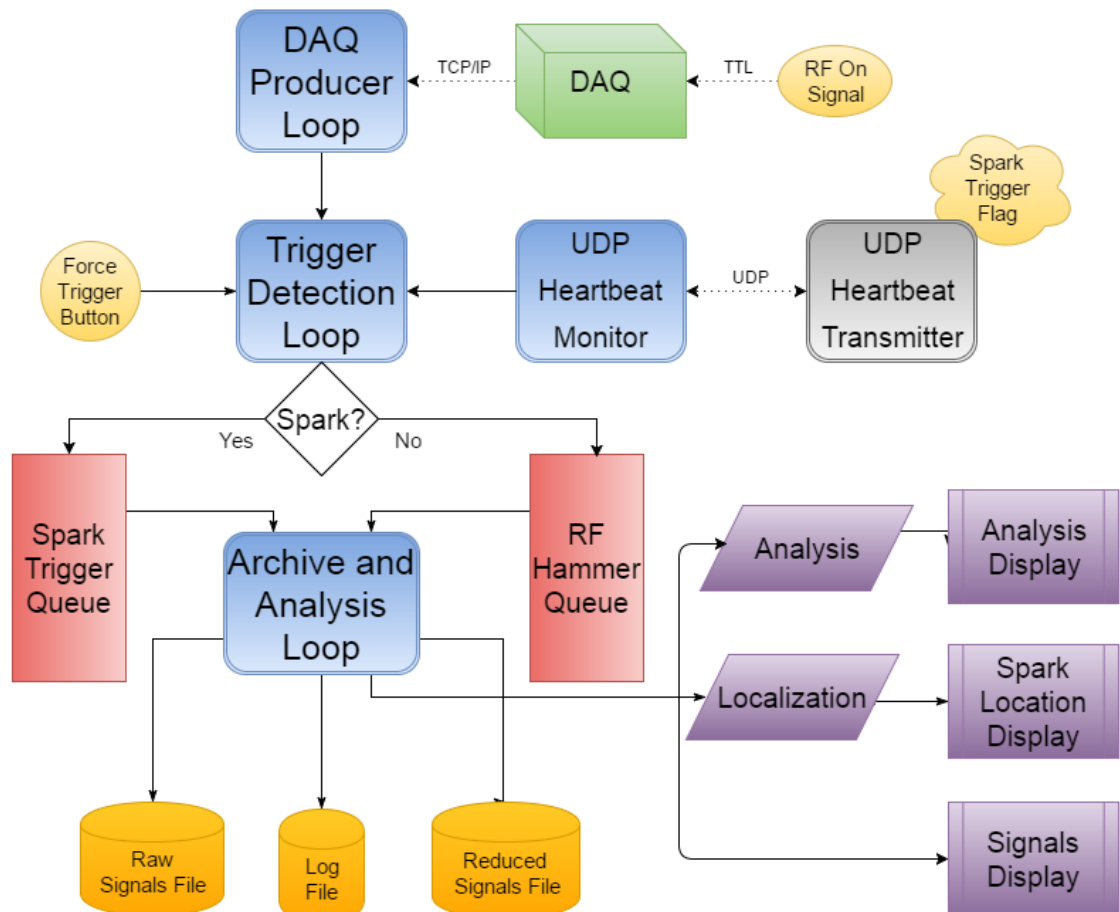
Figure 4.2. Breakdown Acoustics Controls diagram. Blue objects are the independent processing loops. Yellow objects represent external sources or sinks of data. External hardware and software are denoted by green and gray objects respectively. Signal processing queues are represented by red objects. Purple objects represent analysis computation and presentation.

The RF hammer is the most significant source of background noise. The RF hammer signals are sent to the RHQ so that they can be used to remove this noise from the spark signals (see Sec. 5.3). The RHQ is configured to hold only five RF pulses worth of signal data. If it is full, the next set of signal data that is added causes the oldest set of data on the queue to be deleted. This allows the calculation of a rolling average of the RF hammer signals on each microphone channel. The AAL does this calculation and displays it in a graph every 10 ms in the background when the STQ is empty.

If the STQ is not empty, the AAL gives priority to processing the spark signal data. The AAL removes the oldest data set, archives it, subtracts the average RF hammer from the data set, archives the reduced signal data (the data set after RF hammer subtraction), logs the event, performs localization and other analyses on the reduced data, and finally displays the results through various graphs and indicators.

The localization function implements the algorithm discussed in Sec. 5. It assumes that microphones are placed only on the upstream and downstream surfaces of the cavity. This restriction on the microphone configuration has been adhered to since early powered cavity tests indicated that signals coming from different surfaces (e.g., the upstream end cap surface and the curved ring surface) were too dissimilar from each other for the localization algorithm to work. The microphone locations are configured in a file that is specified in the main BAC window. The signals are first trimmed down to a smaller region that is thought to contain the acoustic wavefront. The offset and duration of this "correlation region" are configuration options in the main BAC window. The signals are then split into upstream and downstream sets and processed individually in turn. Thus, one of the outputs of this algorithm is a set of two predicted (x, y) coordinates for the location of the spark end points on the

inner surface of the end plates. Refer to Fig. 4.3 to see the coordinate system used by all acoustic localization software.
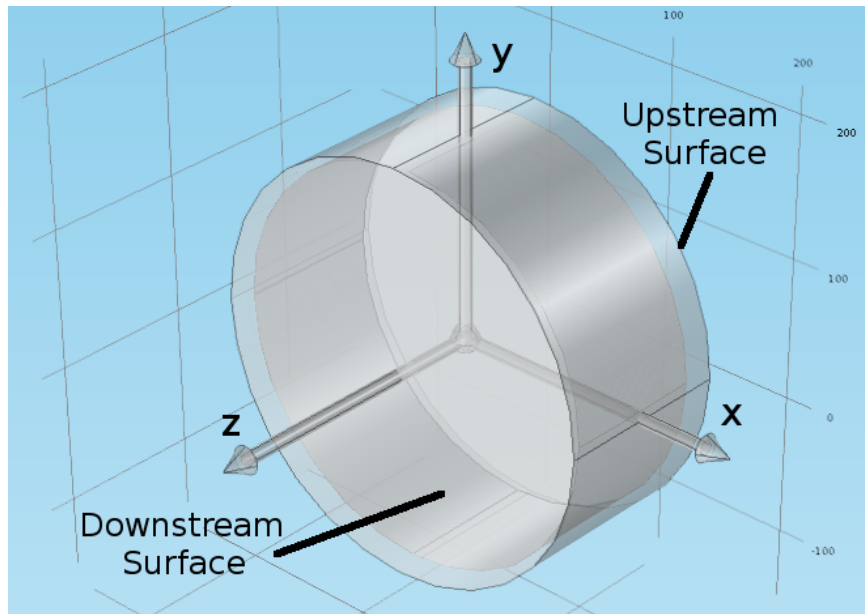
Figure 4.3. Coordinate system used by acoustic localization software.

4.2 ACNET Monitor

Fermilab has a network of computers and other devices called ACNET used to control and monitor the accelerator complex [45]. The MTA group uses this network as well to monitor its own experiments. The ACNET Monitor was built to alert the shifter when the values of the outputs of selected devices go out of range (Fig. 4.4). It also has a convenient summary of current ACNET device values.

The official configuration for the current experimental run is maintained off-site to allow remote updates. It is periodically probed by a computer on the firewall perimeter and, if it has been updated, forwarded to the ACNET Monitor via a web service running on the same computer. Configurations can be created via a custom editor for offline use. This has been useful for monitoring when cavities have started or finished run cycles. Since the MTA group does not have a policy governing the dissemination of cavity operation status, the only ways of determining the operational

status of a cavity are to either contact individuals working with the cavity or to monitor ACNET devices that change value when the operational status changes.
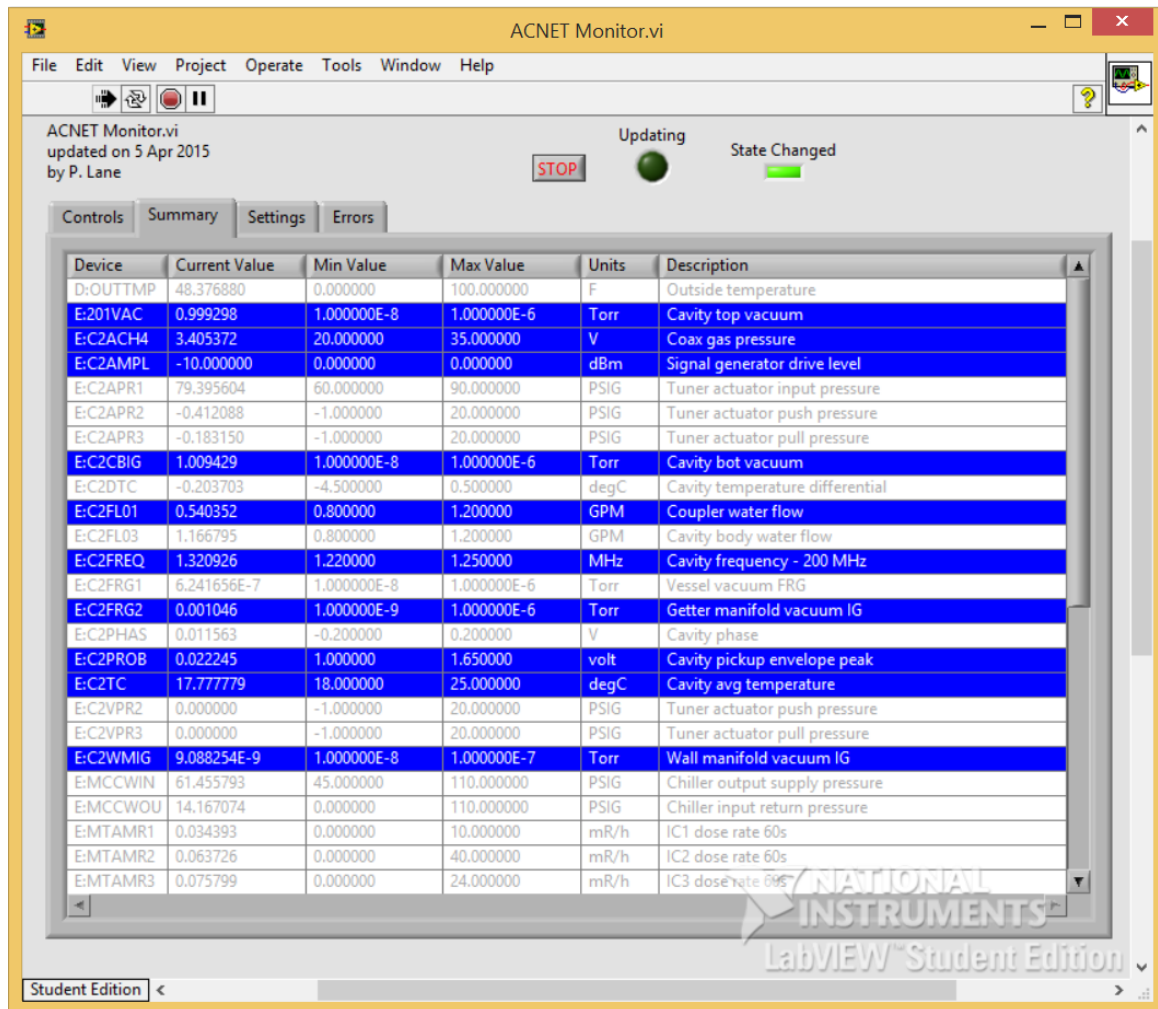
Figure 4.4. ACNET Monitor screenshot.

4.3 Controls Log Monitor

In response to desires that there be periodically logged summaries of cavity run events, the Controls Log Monitor (CLM) was created. The application runs automatically at the end of each shift. It parses the main cavity control log to identify potential breakdown events, analyzes the RF decay slope to determine whether the event was a true breakdown, pulls additional contextual information such as solenoid

magnet current from ACNET, and creates an event summary log entry in the MTA Electronic Collaboration Logbook (ECL).

The CLM can also be forced to generate an on-demand summary for past shifts with the option to create a log entry for this summary in the ECL. Obtaining summaries for cavity runs that were executed before the application was built is an example of its utility additionally as an offline tool. A screenshot of the application is shown in Fig. 4.5.
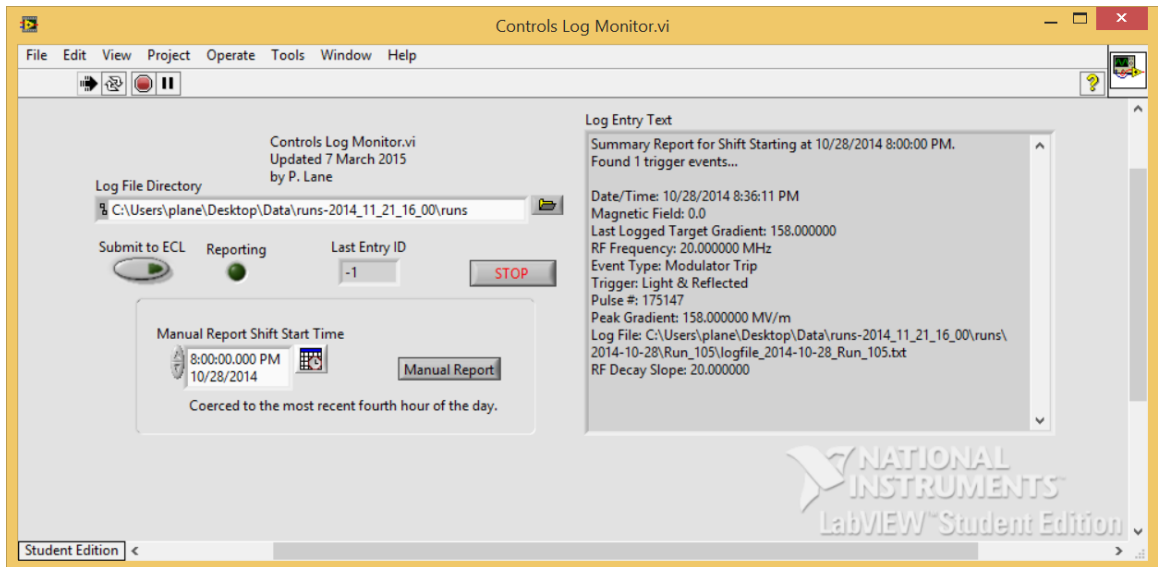
Figure 4.5. Controls Log Monitor screenshot.

4.4 Simulation Software

For purposes of developing an algorithm for localizing breakdown sparks, one would ideally like to perform a set of experiments in which after each breakdown event the cavity would be immediately powered down, opened up, and inspected. This would allow directly associating spark damage to particular acoustic signals. This is impractical, therefore simulations are relied upon to validate the localization method. It must then suffice, following cavity runs, to compare the results of the validated localization algorithm applied to actual acoustic breakdown signals with a

set of spark damage sites to see if there is a reasonable correlation in space. This section provides a brief introduction to the simulation software used for this purpose. Details on the aforementioned simulations can be found in Ch. 6.

4.4.1 Finite Element Analysis. Finite Element Analysis (FEA) uses the Finite Element Method (FEM) to find approximate solutions to physical boundary value problems that are represented by partial differential equations. FEM is a numerical technique that subdivides a problem domain into smaller pieces (finite elements), and then composes the overall solution from the solutions of the finite elements. FEA simply applies FEM to the vertices of a mesh over a geometric model. The mesh consists typically of non-overlapping triangles in 2D and tetrahedrons in 3D that share sides and vertices (see Fig. 4.6 for High-Pressure Cavity (HC) mesh examples) [46].
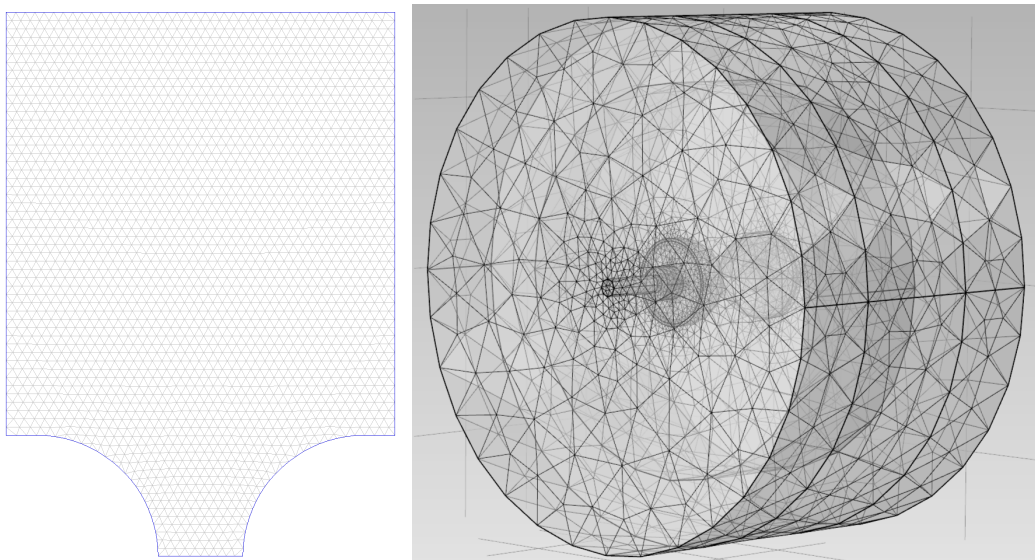
Figure 4.6. Examples of 2D and 3D meshes of the HC with buttons. The 2D mesh (left) was generated by Poisson Superfish and represents a half cross-section of the HC's resonant chamber. The curved voids at the bottom of the mesh are where the buttons protrude into the resonant chamber. The 3D mesh (right) was generated by COMSOL Multiphysics and represents a simplified version of the entire cavity.

4.4.2 Poisson Superfish. Poisson Superfish is an FEA software package for calculating, among other things, the electric and magnetic fields in cylindrically symmetric

RF cavities. It is not being actively developed anymore but is available for download with limited support from the Los Alamos Accelerator Code Group [47].

4.4.3 COMSOL Multiphysics. COMSOL Multiphysics is an FEA software package that allows one to build 2D and 3D models and then simulate one or more physical processes concurrently [48]. COMSOL comes by default with solid mechanics capabilities that can be used to simulate sound inside solids. An additional acoustics model was also acquired to simulate the fill gas in the HC and its interactions with the cavity wall.

CHAPTER 5

BREAKDOWN LOCALIZATION ALGORITHM

Successful localization of breakdown in multi-cell accelerating structures has been done before, but that work sought only to locate the individual cell where the breakdown occurred and relied on the proximity of that cell to one of the microphones [49–51]. Instead, the technique proposed here aims to localize breakdown within a larger, solitary RF cavity in at least two dimensions. In other words, the goal here is to identify the coordinates of the breakdown source in a 2D plane rather than just the closest microphone to the breakdown.

Similar to the LIGO experiment that recently detected a pair of merging black holes, two general approaches to localization algorithms were considered: high-latency, high-accuracy and low-latency, approximate [52, 53]. One method of high-latency localization of sound emission in metal plates similar to one of the LIGO methods is based on Bayesian analysis [54]. These techniques are best suited for offline reconstruction of the emission source. More low-latency techniques for localizing sound sources in metal plates include methods based on the Continuous Wavelet Transform (CWT) to isolate the initial peak of the sound wavefront [55]. Such a technique is often used for real-time applications where computational time should be kept as low as possible.

One of the goals of this research was to develop a method that could localize breakdown in real-time as a diagnostic for a running cavity. Therefore low-latency methods that sacrifice some accuracy were preferred. Future work could entail a similar system as LIGO whereby an approximate prediction is given immediately, and a more accurate prediction could be computed in the background and delivered

at a later time. Although CWT analysis on simulated signals was usually able to determine the temporal order of signal wavefronts, the predicted wavefront times were too inaccurate for calculating the spark source location. Figure 5.1 illustrates this with a simulated signal example.
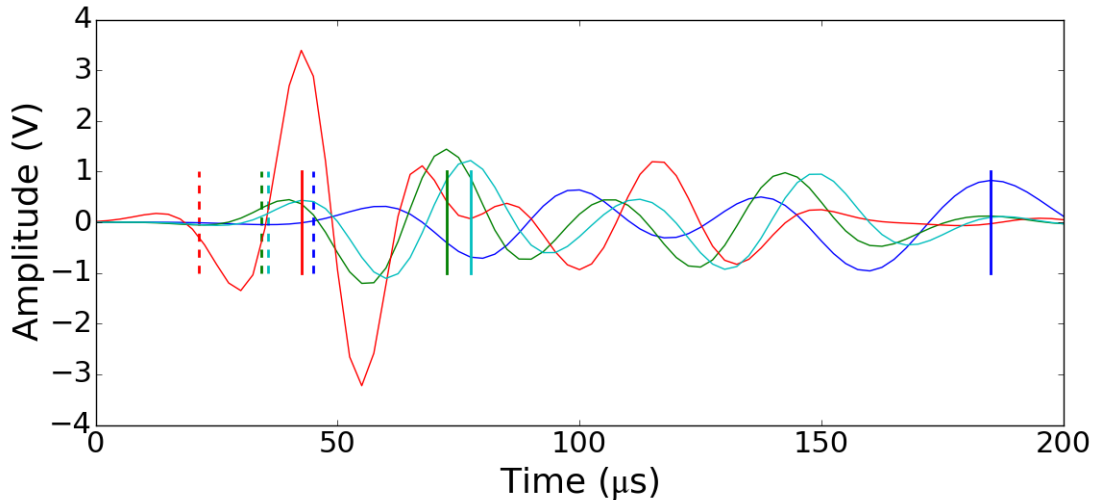
Figure 5.1. CWT analysis of simulated signals correctly determines the temporal order of the signals, but is not accurate enough to produce usable timestamps for localization. The solid vertical lines are the wavefront times predicted using CWT analysis, while the dashed lines are the actual wavefront times calculated from the known source and microphone locations.

The alternative approach to CWT analysis that showed the most promise was to adapt a free-space localization technique for use with a solid medium (i.e., copper or stainless steel) instead of air. Sound and radio waves from another location in a room, for example, have been localized using the Time Difference of Arrival technique (TDOA) [56, 57]. TDOA is an example of a low-latency but low-accuracy free-space localization method. On the high-latency but accurate end of the spectrum of free-space localization, beamforming techniques have also been used successfully to localize line-of-site noise from sound sources such as wind turbines [58]. A hybrid technique named Accumulated Correlation (AC) was chosen for the basis of the breakdown

localization algorithm presented herein. It has been shown to be just as accurate as beamforming while being more computationally efficient, like TDOA [59].

5.1 Accumulated Correlation

AC starts by computing the cross-correlation between each pair of signals:

$$R_{ij}(\tau) = \int_{t_0' - \frac{W}{2}}^{t_0' + \frac{W}{2}} x_i(t)x_j(t - \tau)dt, \tag{5.1}$$

where τ is a particular delay value, t_0' and W are the center and width of the time region in which the sound is expected to be heard, and x_i and x_j are the ith and jth microphone signals. Evaluating R_{ij} for a discrete range of τ values yields a cross-correlation vector. The value of each element in the vector gives an indication of how likely the corresponding delay, τ, equals the true delay between the two signals. The limits of integration correspond to the "correlation region" mentioned in Sec. 4.1. Currently these values are configured by the user of the acquisition software, but in a production system, this region could be determined automatically through additional signal analysis.

Fig. 5.2 shows the cross-correlation between an example signal and the same signal shifted 12.5 units later in time. The peak in the cross-correlation vector sits at -12.5, indicating that the first signal (blue) most likely precedes the second signal (green) by 12.5 time units.

Since propagation times are proportional to distances by a factor of the sound velocity, each time delay, $\tau = t_2 - t_1$, represents a physical half-hyperboloid of sound source positions for each pair of microphones. This is because a hyperbola can be defined as the set of points where the absolute value of the difference between the distances to those points from two stationary foci is constant:

$$|d_2 - d_1| = \tau v, \tag{5.2}$$
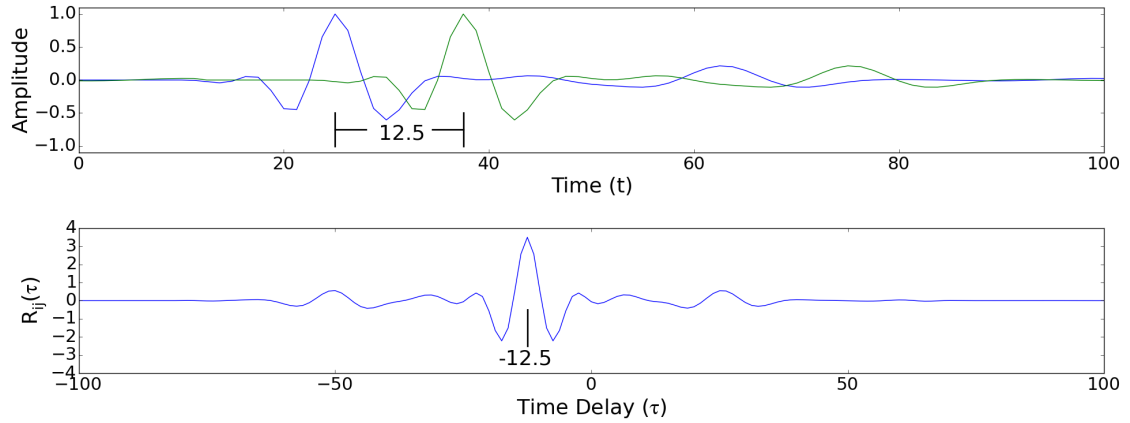
Figure 5.2. The bottom plot shows the cross-correlation between the blue and green signals in the top plot. The maximum in the cross-correlation vector indicates the most likely time delay between the two signals. Note that in this example $t_0' = 50$ and $W = 100$. Therefore the cross-correlation is twice the size of the signals.

where $d_1 = t_1 v$, $d_2 = t_2 v$, and the microphones sit at the foci (Fig. 5.3). Thus, each cross-correlation vector element indicates how likely the source is to lie on a particular half-hyperboloid in space. The maximum of the cross-correlation vector represents the most likely half-hyperboloid on which the sound source sits. That said, it is often the case that the correct half-hyperboloid is not associated with the cross-correlation element with the largest value [59].
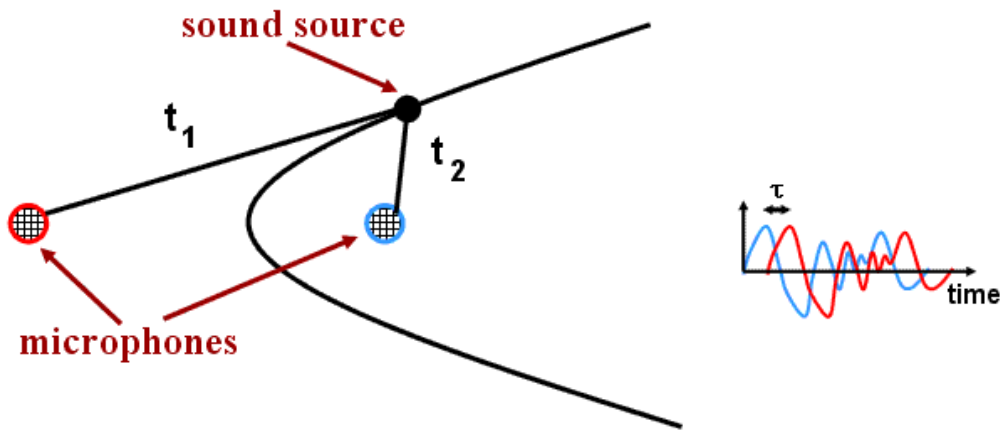
Figure 5.3. The travel delay of an acoustic signal between two microphones defines a half-hyperboloid in space with the microphones located at the foci [60].

After the cross correlation vectors are calculated, AC defines a set of candidate source locations (i.e., a regular grid of test points). Then, for each candidate location, $\mathbf{q}$, the difference of travel times is computed between each pair of transducers,

$$\tau_{ij,\mathbf{q}} = \tau_{j,\mathbf{q}} - \tau_{i,\mathbf{q}}, \tag{5.3}$$

where

$$\tau_{i,\mathbf{q}} = \frac{r_{i,\mathbf{q}}}{v}, \tag{5.4}$$

$r_{i,\mathbf{q}}$ is the distance between test point $\mathbf{q}$ and microphone i, and v is the speed of sound in the propagation medium. The cross-correlation vector elements corresponding to these time differences between signal pairs are summed to create an overall indication of how likely the sound was to have originated from the candidate source location:

$$L_{AC}(\mathbf{q}) = \sum_{i=1}^{N} \sum_{j=i+1}^{N} R_{ij}(\tau_{ij,\mathbf{q}}), \tag{5.5}$$

where the cross-correlation function, $R_{ij}(\tau)$, is that which is defined in (5.1). The candidate source location, $\mathbf{q}$, with the largest corresponding sum is taken as the source location prediction [59].

This is both similar to and different from TDOA and beamforming. TDOA also starts by computing the cross-correlation between pairs of signals. Unlike AC, TDOA uses only the maximum value of the cross-correlation vectors. It then attempts to find the intersection of the half-hyperboloids represented by these cross-correlation vector maxima. This is computationally fast, but can be inaccurate if one or more of the correct half-hyperboloids do not correspond to cross-correlation vector maxima. Beamforming, on the other hand, uses a grid of test points like AC, but instead time-shifts the signals and numerically calculates an integral to determine how likely each test point is equal to the correct source location. This can be very accurate, but is computationally intensive. By using all of the cross-correlation information, AC has

a comparable accuracy to that of beamforming while being computationally fast like TDOA.

To visually observe how AC works, a spark simulation was performed for a 28.4 cm diameter copper cavity end plate. Microphones were placed at $(x, y) = (\pm 5 \text{ cm}, \pm 6 \text{ cm})$. Figure 5.4 shows a set of plots of $R_{ij,\mathbf{q}}(\tau_{ij,\mathbf{q}})$ over the grid of test points, $\mathbf{q}$, generated while applying AC to the simulated acoustic signal data.

Each plot in Fig. 5.4 corresponds to a unique combination of i and j (i.e., a unique pair of microphones). Colors are used to represent the values of $R_{ij,\mathbf{q}}(\tau_{ij,\mathbf{q}})$, hotter (darker red) colors representing larger positive values of $R_{ij,\mathbf{q}}(\tau)$ (high correlation between signals) and cooler (darker blue) colors representing larger negative values (low correlation between signals). The location of each associated microphone pair is indicated by black crosses.

Each hyperbolic band defines a region in space where $R_{ij,\mathbf{q}}(\tau_{ij,\mathbf{q}})$ is constant. Because the cross-correlation vector is in reality discrete, groups of half-hyperboloid curves defined by unique but similar values of $\tau_{ij,\mathbf{q}}$ can map to the same cross-correlation vector value. This explains why one sees bands instead of a continuous gradation of contour lines.

Figure 5.5 then shows the result of summing the $R_{ij,\mathbf{q}}(\tau_{ij,\mathbf{q}})$ values from Fig. 5.4 together. The location prediction is taken as the dark red point which is the candidate location, $\mathbf{q}$, with the largest sum value.

The algorithm is implemented in only two dimensions for several reasons. First, as was mentioned before, comparing two signals from microphones on two different surfaces is problematic since the signals are too dissimilar. Restricting the grid of candidate locations to the end plate regions makes sense given that the strongest electric field intensities inside an ideal pillbox cavity are near the beam axis. Finally,
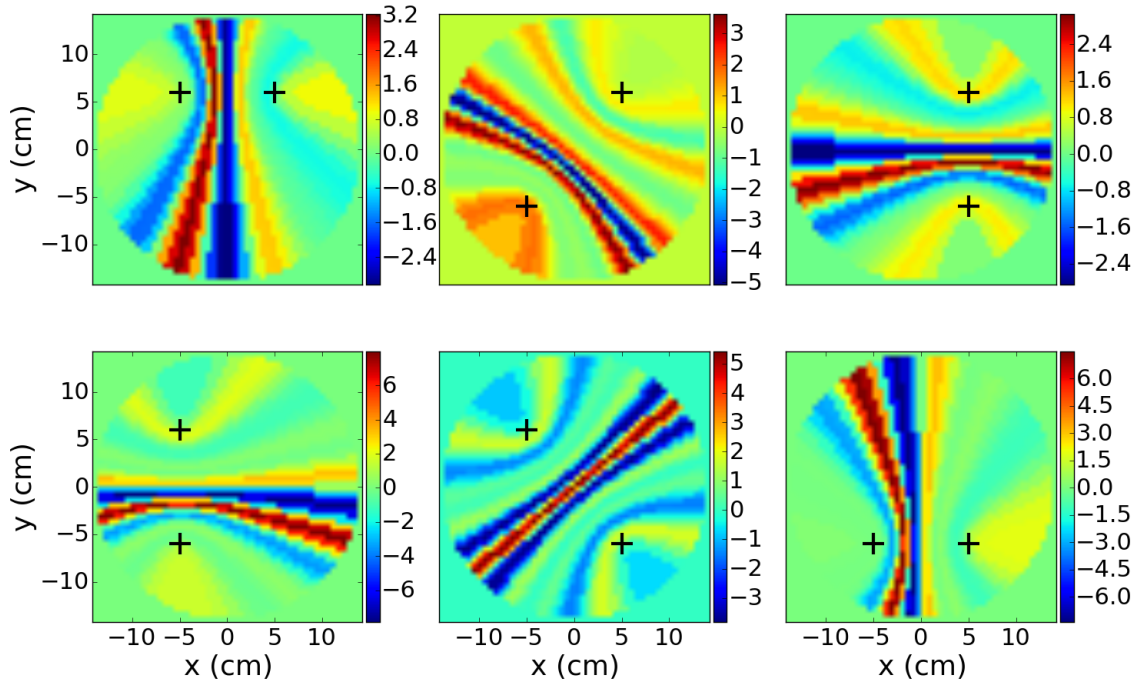
Figure 5.4. AC sum terms. Each subplot shows the values of $R_{ij,\mathbf{q}}(\tau_{ij,\mathbf{q}})$ in (5.5) for a unique pair of microphones, ij, over the grid of candidate source locations, $\mathbf{q}$. Hotter colors represent stronger positive correlation between signals, while cooler colors represent stronger negative correlation. Black crosses indicate the location of the microphones.
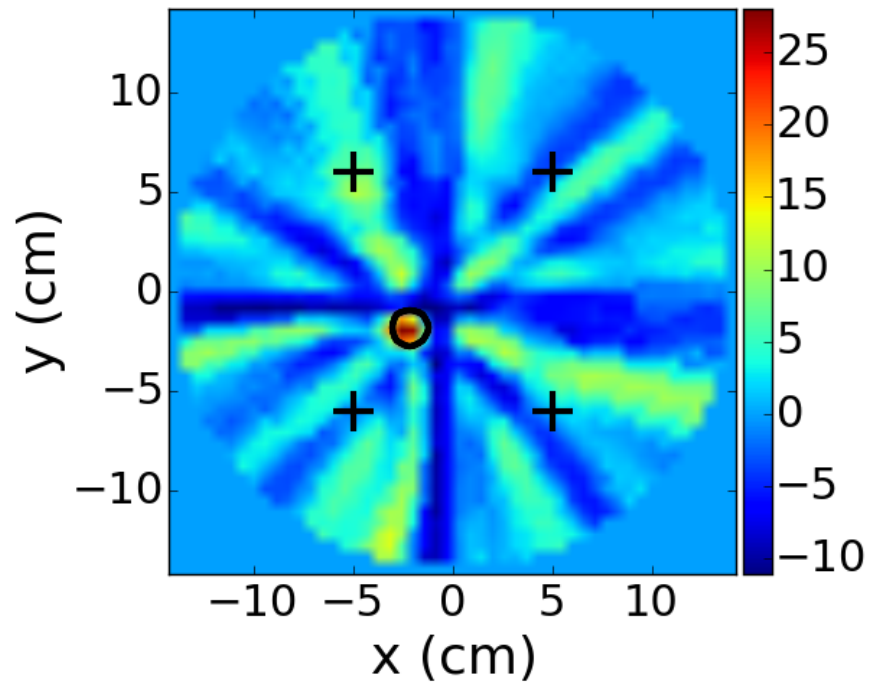
Figure 5.5. AC sums for each candidate location $\mathbf{q}$ (i.e., $L_{AC}(\mathbf{q})$). The sound source location prediction is taken as the grid position with the largest sum (dark red). The true spark source location is indicated by the black circle.

with the assumption that all sparks will occur on the inside surface of the end plate, the longitudinal position of the spark source is already known. Taking the z-axis position as a known quantity greatly improves the performance of the AC algorithm since it is searching over a flat surface instead of a volume of candidate locations.

In the event that a spark occurs near a side-coupled RF coupler, for example, the prediction ought to be coerced to a point of maximum radius in the 2D grid. This can easily be distinguished from a near-axis prediction.

Finally, it should be noted that the minimum grid size (i.e., the minimum number of points per side) for an end plate of a particular size and material is defined by the sampling resolution of the DAQ. The minimum grid size corresponds to the spatial resolution of the test grid above which the spatial resolution of the DAQ is exceeded. Trying to resolve the spark location to regions of space smaller than the DAQ can differentiate is meaningless and merely wastes computational power. Each grid side spans the diameter of the end plate, D. The maximum spatial resolution of the DAQ is equal to the slowest sound velocity in the material (see Sec. 5.3.2), v_L, divided by the sampling frequency, f_s (i.e., the smallest spatial distance that corresponds to a single DAQ sample). The minimum grid size is then simply the physical length of a grid side divided by the smallest distance that can be represented by a single DAQ sample,

$$s_{\min} = \frac{f_s D}{v_L}. \tag{5.6}$$

For example, in the case of the MC we have $f_s = 4 \times 10^5$ Hz, $v_L = 1.8 \times 10^5$ cm/s, and $D = 28.4$ cm. This gives a minimum grid size of 63 points. All MC predictions were obtained using this grid size.

5.2 Sound Wave Speed

Sound waves in solids can travel in one of several modes. Each mode has an associated propagation speed [61]. To determine an appropriate wave speed to use with the Accumulated Correlation algorithm, a velocity parameter sweep was performed using the reconstructed results of simulations as a metric. Specifically, a uniformly random set of 1,000 true spark source locations was produced and used as input to solid mechanics simulations in COMSOL (see 6.5 for more details). The results of these simulations were fed into the localization algorithm, and the predicted spark source locations were compared with the true source locations by calculating the RMS error (RMSE) of the residuals.

This reconstruction was performed for various wave speeds between 1.0×10^5 cm/s and 5.0×10^5 cm/s. The choice of these parameter sweep limits was informed largely based on the known shear and pressure wave speeds in copper (2.33×10^5 cm/s and 4.76×10^5 cm/s respectively [35]). Figure 5.6 shows the results of this parameter sweep. The velocity associated with the smallest RMS error of the residuals (denoted by a vertical red line) is assumed to be the propagation speed of the dominant mode of sound waves in the copper endplate. In this case, the sound speed matches very well with the shear wave speed in copper. This is the speed that was used to produce the plots in Sec. 5.1.

5.3 Signal Conditioning

Applying the localization algorithm directly to the raw spark acoustic signals has not been found to produce good results. Better results were obtained by
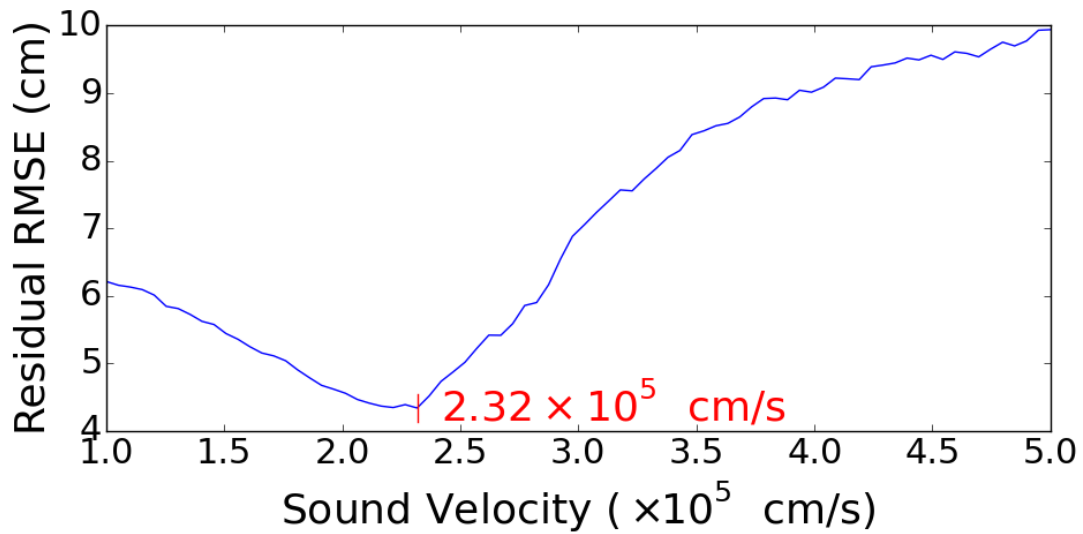
Figure 5.6. Results of a velocity parameter sweep using a model of the MC copper
  endplate. The velocity associated with the smallest RMSE of the residuals (vertical
  red line), matches well with the shear wave speed in copper.

subtracting the RF hammer background noise and taking the absolute value of the
signal.

5.3.1 RF Hammer Subtraction. It was predicted that the breakdown thermal
shock would be several orders of magnitude louder than the RF hammer (see Sec. 2.3).
In practice, the amplitude of the thermal shock varies and is often comparable with
the RF hammer amplitude (Fig. 5.7). It is therefore necessary to remove the RF
hammer component of the breakdown acoustic signal in order to expose only the
component that is from the thermal shock. For a given accelerating gradient, the RF
hammer acoustic signals do not generally differ from pulse to pulse by more than the
noise floor. Any minor differences are removed by calculating a running average of
the five most recent RF hammer signals. The average RF hammer is then subtracted
from the breakdown signals.
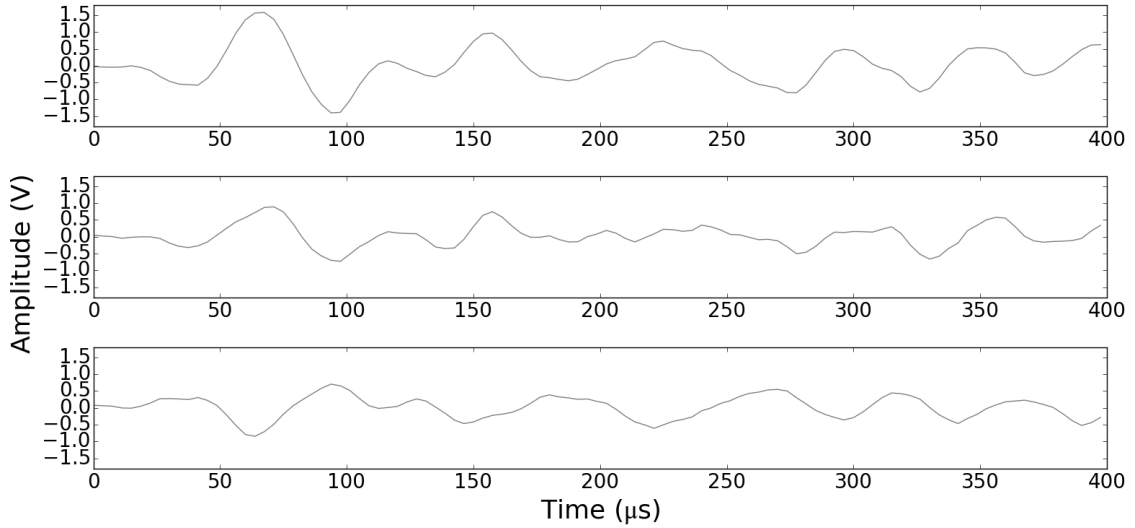
Figure 5.7. An example average RF hammer signal (top) in comparison with a relatively quiet breakdown signal at the same gradient (middle) and the difference between the two (bottom).

5.3.2 Signal Absolute Value. Although there was only one clear choice for the predicted sound source location in Fig. 5.4, AC results on simulated signals frequently showed multiple large-sum (high-probability) locations. In addition, it was not uncommon (especially for points far from the endplate center) that the true spark location was at or near a low-sum location instead of a high-sum location. In an attempt to compensate for these situations, the absolute value of the simulated signals was taken before feeding the signals to the AC algorithm. Another velocity parameter sweep similar to that in Sec. 5.2 was performed, and the results can be seen in Fig. 5.8. Again, the sound speed associated with the lowest residual RMSE is marked with a vertical red line.

The combination of taking the absolute value of the signal and using a wave speed of 1.86×10^5 cm/s significantly improved the RMSE between predicted and true spark source locations. This lowering of the wave speed may be an artifact of taking the absolute value of the signal. On the other hand, it may instead indicate that taking
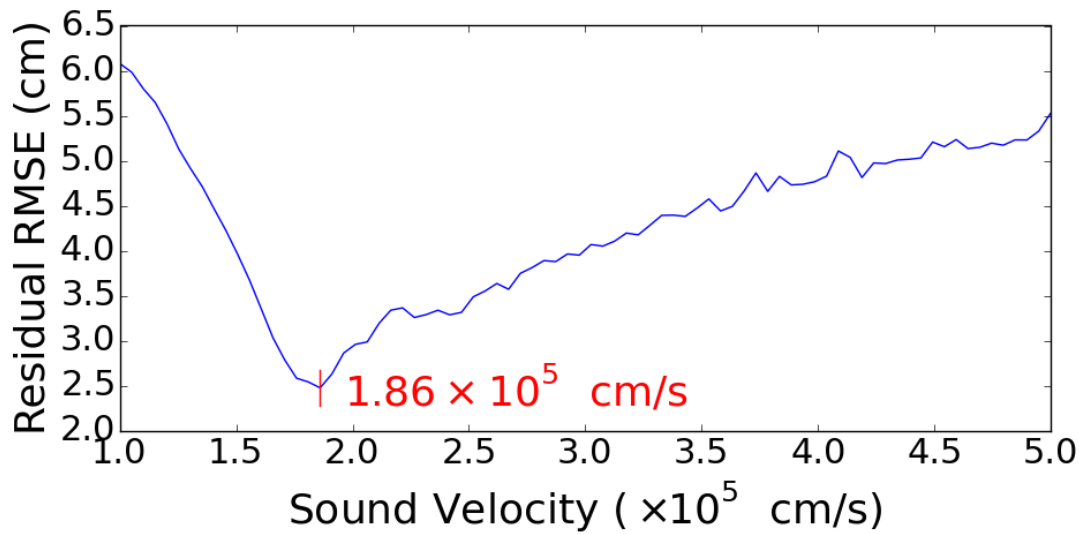
Figure 5.8. Results of a velocity parameter sweep after taking the absolute value of simulated spark sound signals. The vertical red line indicates the velocity that produced the smallest residual RMSE.

the absolute value of the signal has emphasized low-frequency, anti-symmetric Lamb waves which have similar phase velocities [62]. Either way, because of its success, this algorithm modification was used in all subsequent MC localization results.

5.4 Summary

To achieve real-time localization of spark sources inside RF cavities, a low-latency algorithm called Accumulated Correlation was chosen as the basis for a spark source localization algorithm. This algorithm is a hybrid of TDOA and beamforming, summing samples taken from the cross-correlations between each pair of microphone signals for each point in a grid of candidate locations.

To improve the accuracy of spark source location predictions, the signals were conditioned by subtracting the average RF hammer (in the case of processing live cavity data) and taking the absolute value of the signals. Before conditioning, simulations

indicated that a wave speed equal to the shear wave speed in copper (2.325×10^5 cm/s) resulted in the highest accuracy. After applying the absolute value function to the signals, a wave speed equal to 1.86×10^5 cm/s improved the accuracy nearly two-fold.

CHAPTER 6

SIMULATIONS AND EXPERIMENTS

Although it is necessary to apply the algorithm described in Sec. 5.1 to demonstrate acoustic localization of breakdown, it is not sufficient. Without corroborating evidence suggesting that the algorithm does what it was intended to, there is no reason to have confidence in its results. It is also necessary but not sufficient to include a comparison of observations of damage with predictions from acoustic data. Though not impossible, it is generally impractical to operate a cavity for the sole purpose of associating a particular instance of damage with the prediction derived from the associated captured acoustic signals. Therefore an alternative source of validation of the algorithm is required. As a substitute for direct observational evidence associating specific breakdown damage locations with thermal shock origins, results from structural mechanics simulations in COMSOL Multiphysics (see Sec. 4.4) will be presented to provide justification for trusting the results of this technique when applied to experimentally obtained acoustic data.

Section 6.1 compares the results of structural mechanics simulations with simple impact experiments on the Aluminum Mock Cavity (AMC). Section 6.2 continues this exploration by comparing the results of simulations and experiments measuring sound due to the electromagnetic forces on the inside of the High-Pressure Cavity (HC). In addition, simulations presented in Sec. 6.3 show that the sound generated by breakdown in the HC is due primarily to the shockwave produced in its fill gas. This demonstrates that the actual RF cavities are well represented by the simulation models that have been developed. Section 6.4 presents tests of the localization algorithm for self-consistency using the Dielectric Loaded High-Pressure Cavity (DLHC).

Finally, in Sec. 6.5 and Sec. 6.6, the results from localizing both simulated Modular Cavity (MC) breakdown as well as experimental acoustic data of breakdown captured during a run of the Modular Cavity in a 3 T magnetic field are presented. These results show that predictions using the algorithm described in Sec. 5.1 are accurate on average to within a few centimeters and are sufficient for differentiating between breakdown originating near the center of the cavity and that which originates further from the center near the RF coupler.

6.1 Impact Hammer

The simple geometry of the AMC made it ideal for initial simulation validation. To make sure it was understood how to make an accurate representation of a cavity and reproduce appropriate physical interactions in COMSOL, a series of impact acoustic experiments and simulations were performed that test agreement between their results. The experiments comprised striking the surface and recording signals from a microphone on an oscilloscope. This was then simulated in COMSOL, and the resulting acceleration time series data were compared with the oscilloscope signal data.

6.1.1 Impact Hammer Force Curves. A makeshift impact hammer was built by taping a spare passive microphone to a hex head driver (Fig. 6.1). A drop of hot glue was used to form a smaller impact surface directly under the microphone piezo crystal. When the AMC surface was struck by the impact hammer tip, the microphone provided an uncalibrated impact force curve (blue trace in Fig. 6.2). The following damped sine function (green trace in Fig. 6.2) was then used to approximate the force curve in simulations:

$$F = F_0 \ e^{-\gamma t/t_0} \ \sin\left(\frac{2\pi}{t_0}t\right),$$

(6.1)
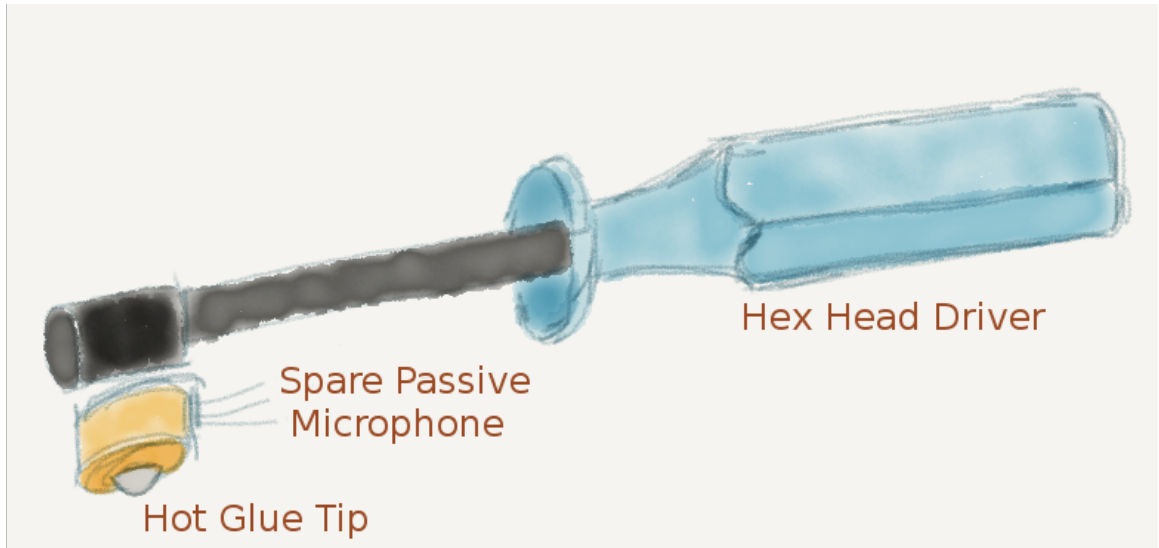
Figure 6.1. Impact hammer made of a spare passive microphone taped to a hex head driver. A rubber-like impact tip was made with a drop of hot glue.

where $F_0 = 1$ N is the force amplitude, $t_0 = 2.3 \times 10^{-3}$ s is a measure of the pulse duration, and $\gamma = 3$ controls the amount of damping.

The values of F_0, γ, and t_0 were selected to best fit the main peak of the damped sine curve (green trace) to the main peak of the captured force curve (blue trace) in Fig. 6.2. The actual value of F_0 was not important given that the voltages seen at the oscilloscope were not calibrated with a known force magnitude, and changing the value of F_0 only scaled the acceleration data amplitudes. Though the fit to the tail suggests that the value of γ is too small (too little damping), increasing this value to something that better represents the force curve tail has no discernible effect on the resultant signal. Such a change merely acts to decrease the goodness of fit to the main peak.

6.1.2 Aluminum Disk. The first experiment that was done with the AMC was to epoxy a passive microphone to the flat surface of one of the AMC end plates at an arbitrary radius of 12 cm. The surface was then struck with the impact hammer at
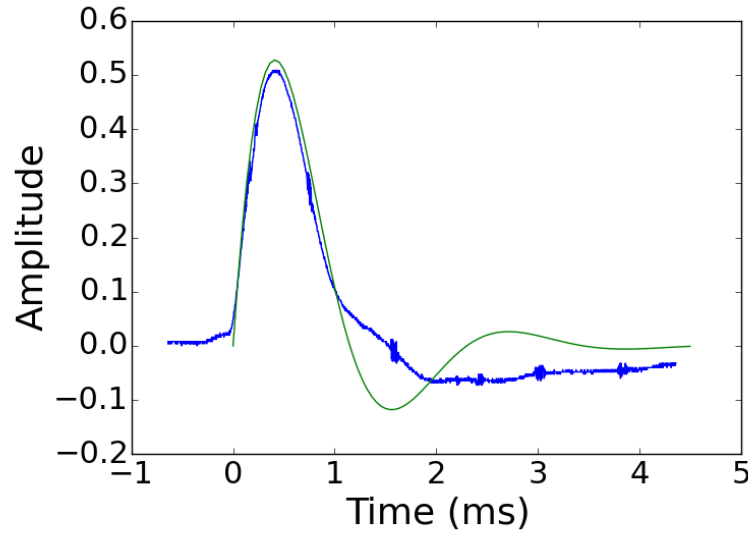
Figure 6.2. Actual impact hammer force curve as seen on an oscilloscope (blue), and an approximate version of the force curve for use in simulations (green).

the same radius, but at a position 180° from that of the microphone. The resulting oscilloscope signal can be seen in the upper left plot of Fig. 6.3.

A straightforward model of an aluminum cylinder with the same dimensions as the AMC end plate was then constructed in COMSOL (Fig. 6.4). A second cylinder 2 mm thick with a radius of 2 mm was added to represent the impact hammer tip. The damped sine function force curve (6.1) was then applied to the tip cylinder to generate mechanical waves in the aluminum disk, emanating from the point of contact between the two cylinders. The acceleration normal to the surface at the equivalent position of the pickup microphone was then sampled to produce acceleration time series data.

To simulate attenuation, Rayleigh damping was added to approximate the decay envelope of the microphone signals. In one dimension the displacement wave equation is

$$m\frac{d^2u}{dt^2} + c\frac{du}{dt} + ku = f(t),$$
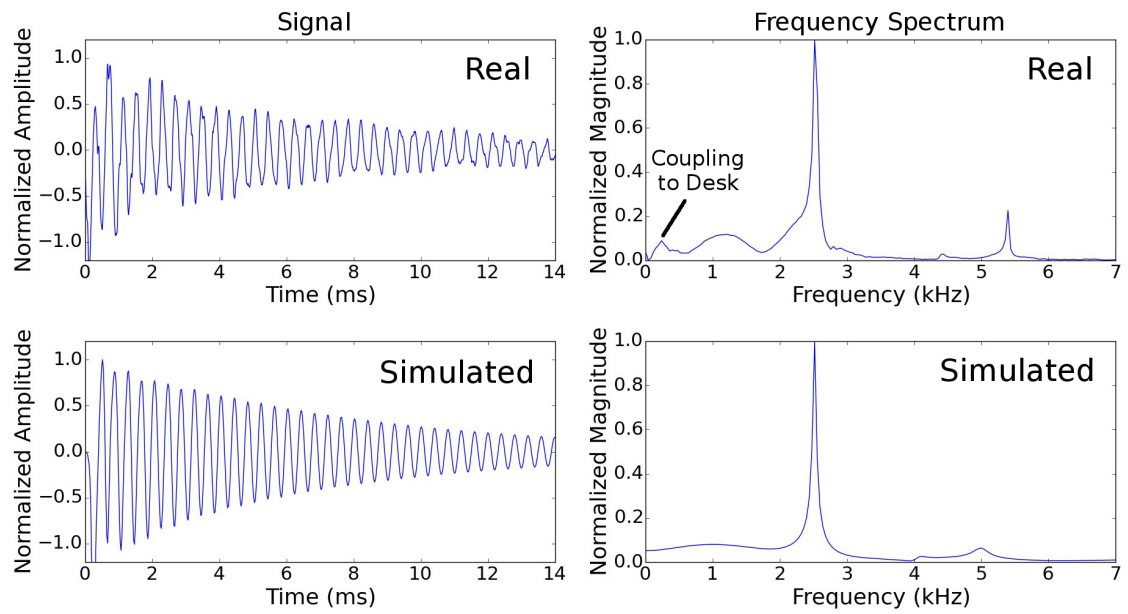
Figure 6.3. Comparison of AMC disk microphone signal data and frequency spectrum (top) with COMSOL simulated acceleration data and spectrum (bottom).
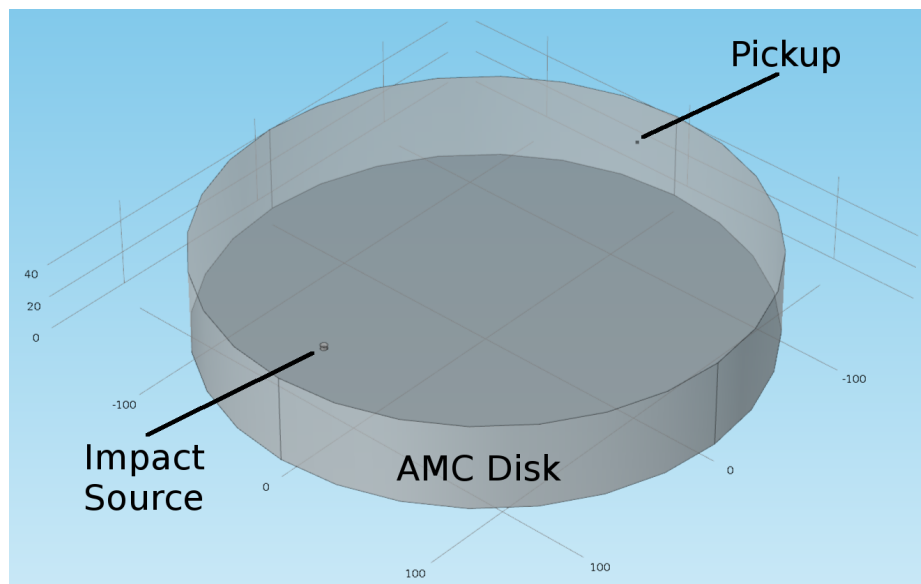
Figure 6.4. COMSOL model of the AMC disk. The tiny cylinder near the left edge acts as the impact hammer tip.

where u, m, c, k, and f are respectively the displacement, mass, damping factor, stiffness, and time-dependent forcing term. Rayleigh damping is added by setting the damping factor $c = \alpha_{dM} m + \beta_{dK} k$. In three dimensions m, c, k, and f generalize to matrices [63]. Empirically, setting $\alpha_{dM} = 0$ s^{-1} and $\beta_{dK} = 2.5 \times 10^{-6}$ s approximately reproduced the same decrease in amplitude as the real signal.

The lower left plot of Fig. 6.3 shows the normalized acceleration data from the aluminum disk simulation with tuned Rayleigh damping. The respective normalized spectrum plots of the real and simulated signals are included on the right to further illustrate the excellent agreement of the simulation. The peak frequencies taken from the spectrum plots for the microphone and simulation are 2586 Hz and 2571 Hz with an uncertainty of 40 Hz for each. The uncertainty is taken as the frequency resolution of the oscilloscope [64],

$$\Delta f = \frac{1}{N \Delta t},$$
(6.2)

where $N = 2500$ is the number of samples and $\Delta t = 1 \times 10^{-5}$ s is the sampling resolution. The simulation time step size was deliberately set equal to the oscilloscope's sampling resolution in order to simplify the analysis.

Despite the overall good agreement between real and simulated data, notice that in the real data there is a low frequency modulation of the signal. This shows up as the small peak labeled in the real data frequency spectrum. This is likely due to the coupling of the disk to the desk that it was sitting on. This frequency peak can be reproduced by adding a Spring Foundation boundary condition to the bottom surface of the large cylinder in the simulation and tuning the boundary condition's spring constants. This boundary condition simulates the coupling of the bottom surface of the large cylinder to a rigid foundation through a spring system. Thus, the motion of the bottom surface is constrained by the strength of the spring system. This boundary condition does not a perfectly reproduce the interaction of the cylinder and

desk surfaces. As a result, the low frequency peak is much stronger in simulation than in data, causing overly pronounced distortion in the signal. Therefore, this boundary condition was not used, in order to maintain the best overall signal match. In an attempt to reduce this coupling effect, subsequent experiments with the AMC were done with the cavity components mounted on small risers to lift the majority of the bottom surface off the desk.

6.1.3 Aluminum Disk and Ring.   Modeling the aluminum ring by itself showed good agreement with observed signals similar to that of the aluminum disk simulation. The aluminum disk and ring were then combined for further experimentation (Fig. 6.5). It was noticed at once that there were discrepancies between the simulated and observed signals. Modal analysis showed that there were modes representing the motion of the entire body as well as modes which represented the disk's motion nearly independent of the ring's motion and vice versa. This led to the idea that the coupling at the interface between the disk and ring was a significant factor.
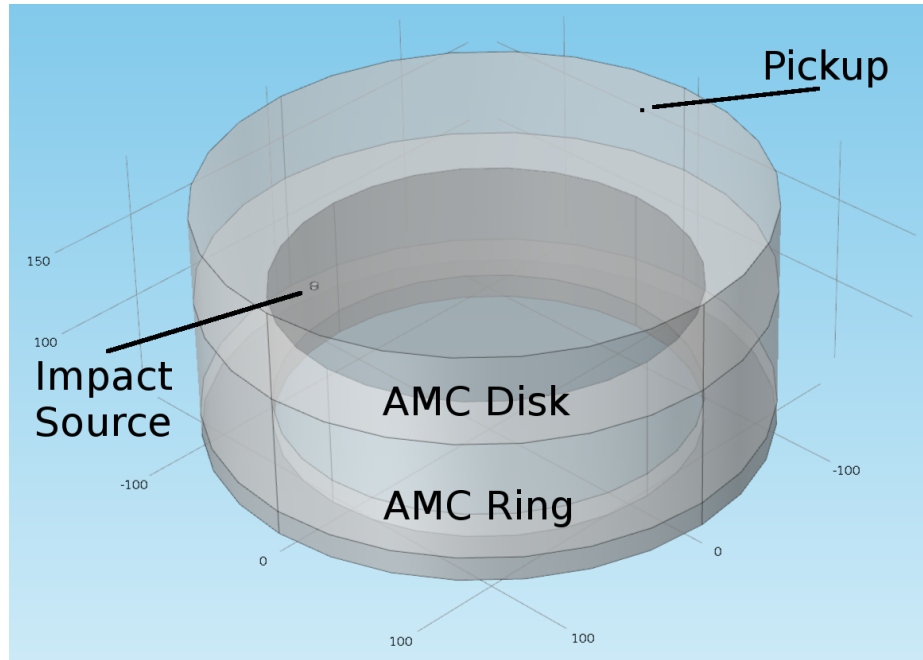
Figure 6.5. COMSOL model of the AMC disk on top of the AMC ring.

An additional Thin Elastic Layer boundary condition was added to the simulation. Adjusting the boundary condition's spring constant parameters (axial and radial in 2D) brought the spectrum of the combined disk and ring model simulations into good agreement with experiment. Figure 6.6 compares these signals and spectra. The two main peaks in the simulation frequency spectrum plot agree with those found in the real spectrum to within 40 Hz (equal to the frequency sampling uncertainty).
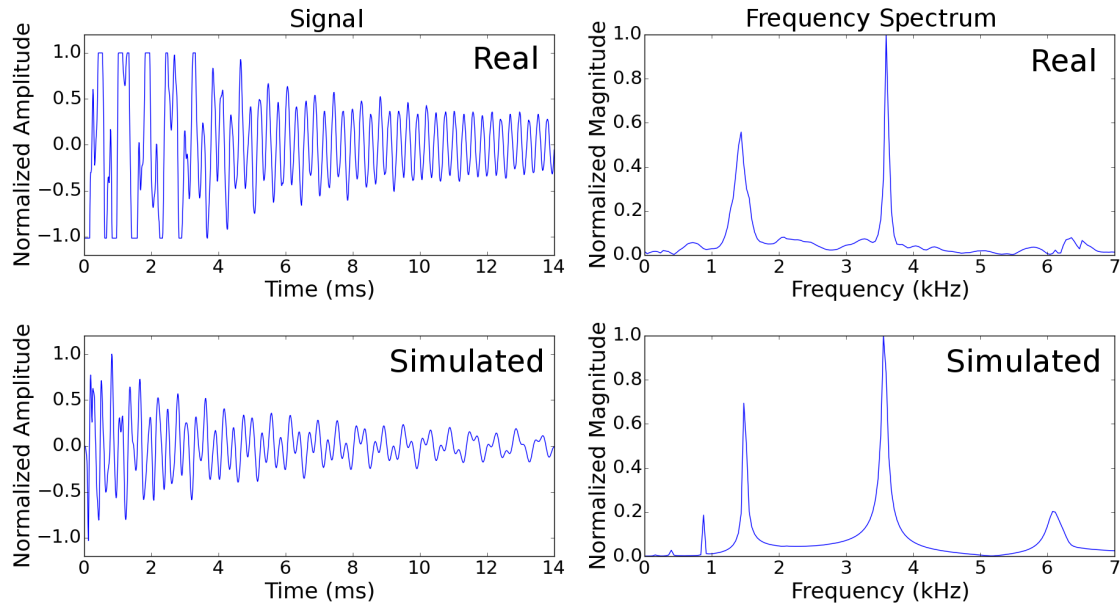
Figure 6.6. Comparison of AMC disk plus ring microphone signal data and frequency spectrum (top) with COMSOL simulated acceleration data and spectrum (bottom).

6.2 RF Hammer

Progressing to experimentation with real RF cavities, data were taken with the HC in the button configuration. Six active microphones were placed on the cavity and labeled S0 through S5. S0 and S5 were placed at a radius of 10.5 cm on the outside, flat surface of the end flanges. S1 through S4 were placed on the curved, outer circumference at 90 degree intervals (see the right image in Fig. 6.7 for reference). Each microphone signal lead was an unshielded twisted-pair terminated

in a BNC connector. Coaxial, BNC jumper cables were used to take the signals to a patch panel in the experiment hall. The patch panel connected with another in the control area. Another set of short, BNC jumpers then made the final link to the PCI-5105 card. The power leads were unshielded, twisted-triplets that terminated in two BNC connectors (one for −V and the other for ground and +V). A distribution box was built with two arrays of nine BNC connectors on opposite sides. A cable with a BNC connector at one end and banana connectors at the other end provided power from an adjustable power supply.
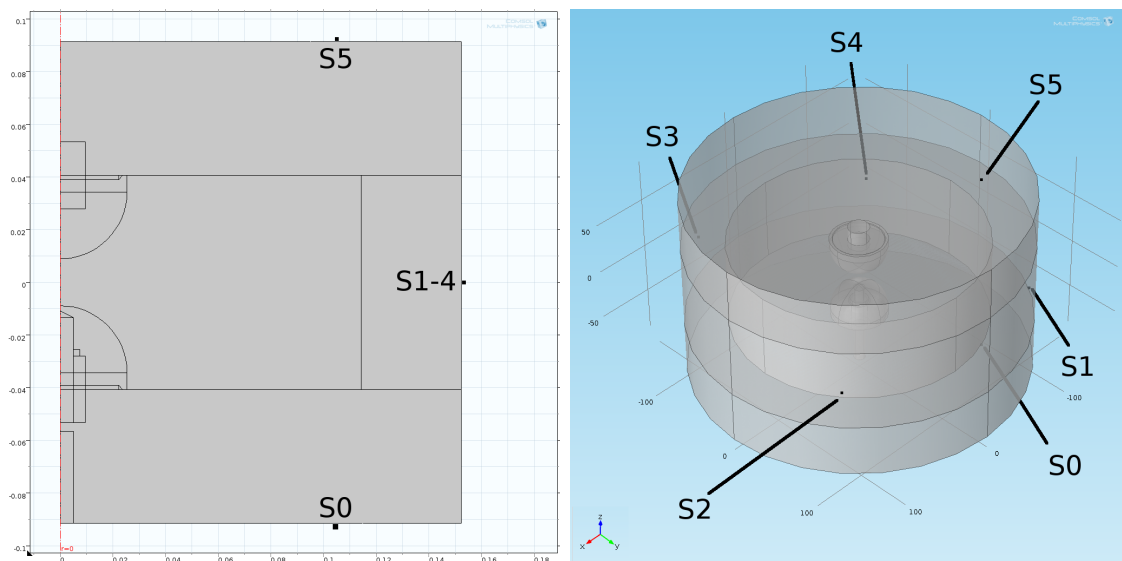
Figure 6.7. Left: 2D HC model. Right: 3D HC model. Microphones are labeled as S0 through S5. S0 and S5 are on the end plates. S1 through S4 are on the center ring.

Both 2D and 3D COMSOL models were made (Fig. 6.7). The 3D model was intended for reproducing all of the microphone signals in order to verify the localization algorithm. The final version of the 3D model did not include all details of the actual cavity such as assembly bolts and various ports. During preliminary simulations, these additions significantly increased the simulation time while not appearing to change the output of the simulations in any noticeable way. In the end, this model was primarily used to check the results of the 2D model simulations since, as ex-

plained later in this section, simulation results suggested localization of the thermal shock to the inner end plate walls of the HC would be problematic.

The 2D model is a half cross-section of the cavity about the cylindrical axis of symmetry (left image in Fig. 6.7). Running a simulation with a 2D model was much faster, so the 2D model was used extensively for testing ideas while trying to match signals and frequency spectra from real and simulated data. In this case, S1 through S4 were assumed to have produced identical signals, and only real signals from S1 were used in comparison to simulation results from this pickup.

To validate the HC model, the effect of the forces from the RF hammer on the inner cavity surfaces was simulated and compared with data collected during powered HC runs. This is a good test since the RF hammer is a well understood phenomenon. The EM fields inside the cavity can be calculated accurately using software such as Poisson Superfish. Figure 6.8 shows the visual representation of the output from Poisson Superfish.
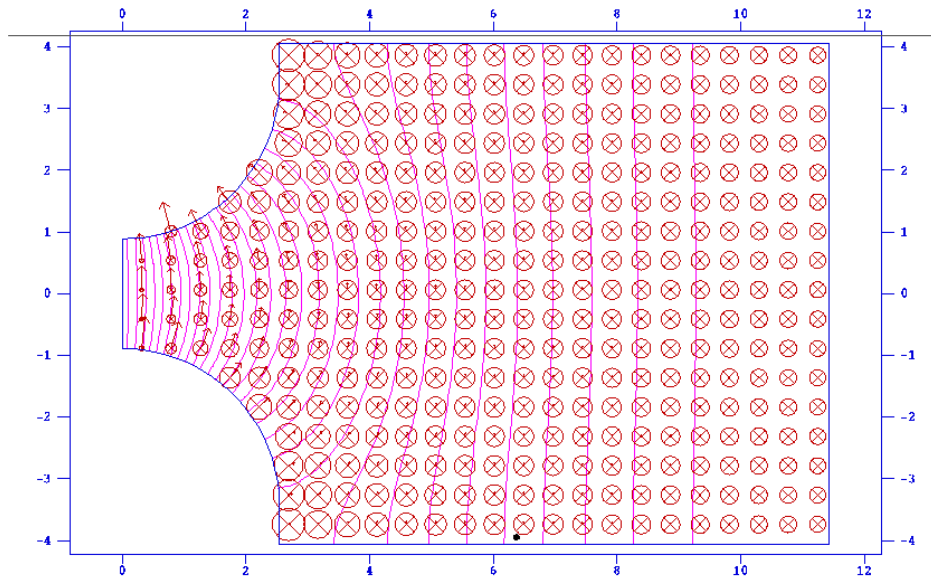
Figure 6.8. Electric and magnetic fields of the button HC calculated by Poisson Superfish. Arrows represent the electric field vectors. Circles with crosses represent the magnetic field vectors.

After exporting the field map data from Poisson Superfish, the field values at the inner surfaces of the cavity were extracted. Using the methods described in section 2.2, the forces on the inner cavity surfaces were calculated. Polynomial fits were then made and used in COMSOL. Figure 6.9 shows the force curves at a gradient of 36 MV/m together with their polynomial fits. This gradient was chosen as an empirically determined safe operating gradient just below the maximum gradient. The fits are very good, so it is difficult to distinguish between the force curves and their polynomial fits. The region between the vertical lines on the button pressure plots is not physical. It represents the gap region between the buttons, and was manipulated to get a better polynomial fit for the associated curves.
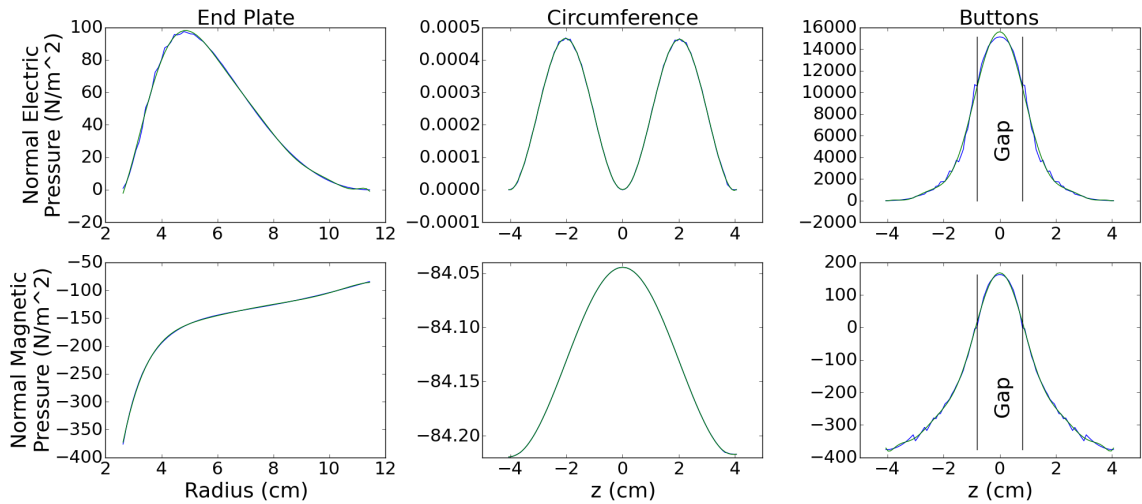
Figure 6.9. Electric and magnetic pressures on the inner cavity surfaces of the button HC (blue). Polynomial fits are in green. The region on the button pressure plots between the vertical lines is the button gap and is not physical.

If the cavity model is accurate, the signals generated from the simulation should closely match actual signals from powered cavity test data. Initial spectrum plots showed clearly that the full HC simulation did not reproduce the frequencies of the microphone signals accurately. Though the general character of the beginning of

the signals could be approximated, the main frequency of signals within the first 200 µs was larger in simulation than in the real signal data by about 25%.

Using the understanding of the coupling between surfaces gained from the AMC tests, the HC model was updated to include Thin Elastic Layer boundary conditions at the interfaces between both the end and center flanges as well as between the copper buttons and end flanges. Fine-tuning the spring constants for these boundary conditions brought the spectra of the simulated and observed signal data into better agreement by more closely aligning the three main frequency peaks from the simulation and real spectrum plots.

It was impossible, though, to shift the spectral peaks low enough to match those of the real data by adjusting the spring constants of any of the Thin Elastic Layers between cavity components. There was also a 14 µs time shift between observed and simulated signals. Further improvements to the simulation were made by adding thin pieces of polyurethane foam with thin elastic layers that represented the double-sided tape originally used to mount the microphones (refer to Sec. 3.1.3 for a discussion of microphone mounting methods). All of the spring constants were re-tuned, and after that were left alone for subsequent simulations. Figure 6.10 shows the final results.

Not only did the shape of the signals improve, but the 14 µs phase difference was eliminated. On the down side, since each individual microphone mounting spring constant had to be fine-tuned independently, this shows that similarity of the signals from different microphones is highly dependent on the coupling between the microphone and the cavity. Similarity of signals is crucial for using localization algorithms based on the cross-correlation function such as Accumulated Correlation. This is expected not to be an issue on all future experiments since epoxy was used
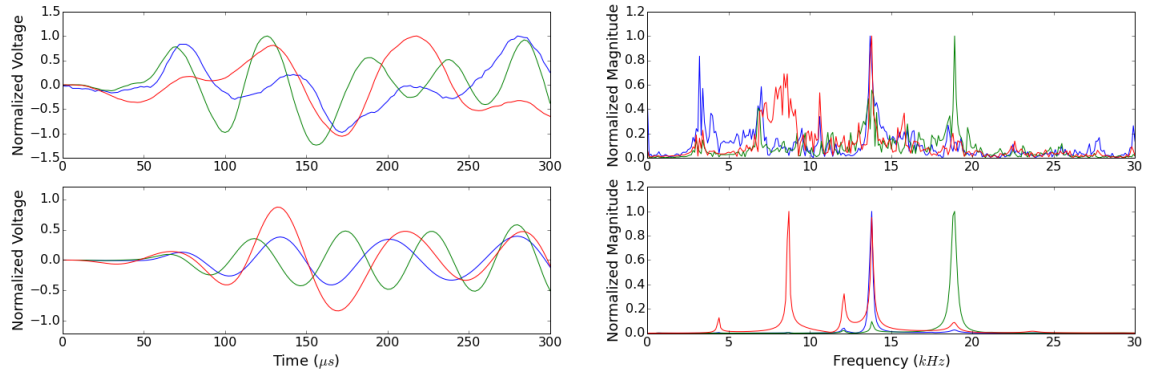
Figure 6.10. Comparison of HC RF hammer microphone signal data and frequency spectrum (top) with COMSOL simulated acceleration data and spectrum (bottom). The blue, green, and red traces correspond to the S0, S5, and S1 microphones respectively.

to mount microphones to cavities after these early experiments. Nevertheless, these results show that the HC model is a reasonable approximation of the real cavity.

6.3 Gas Shockwave

The top plot of Fig. 6.11 shows an example of breakdown signals and an equivalent 2D simulation normal acceleration plot. Inconsistent coupling of the double-sided tape microphone mounting makes localization of these HC breakdown signals inaccurate. That said, looking at the breakdown signals, they clearly exhibit an interesting feature that was not understood before simulations were performed. At around 300 μs, the amplitude suddenly increases. This amplitude jump occurs long after the wavefronts from breakdown sound generated at the button arrived at the microphones.

A simulation was performed that attempted to replicate the forces on the buttons due to a high-current electrical arc. A plot of the acceleration normal to the outer end plate surface vs time at the microphones for this simulation did not exhibit
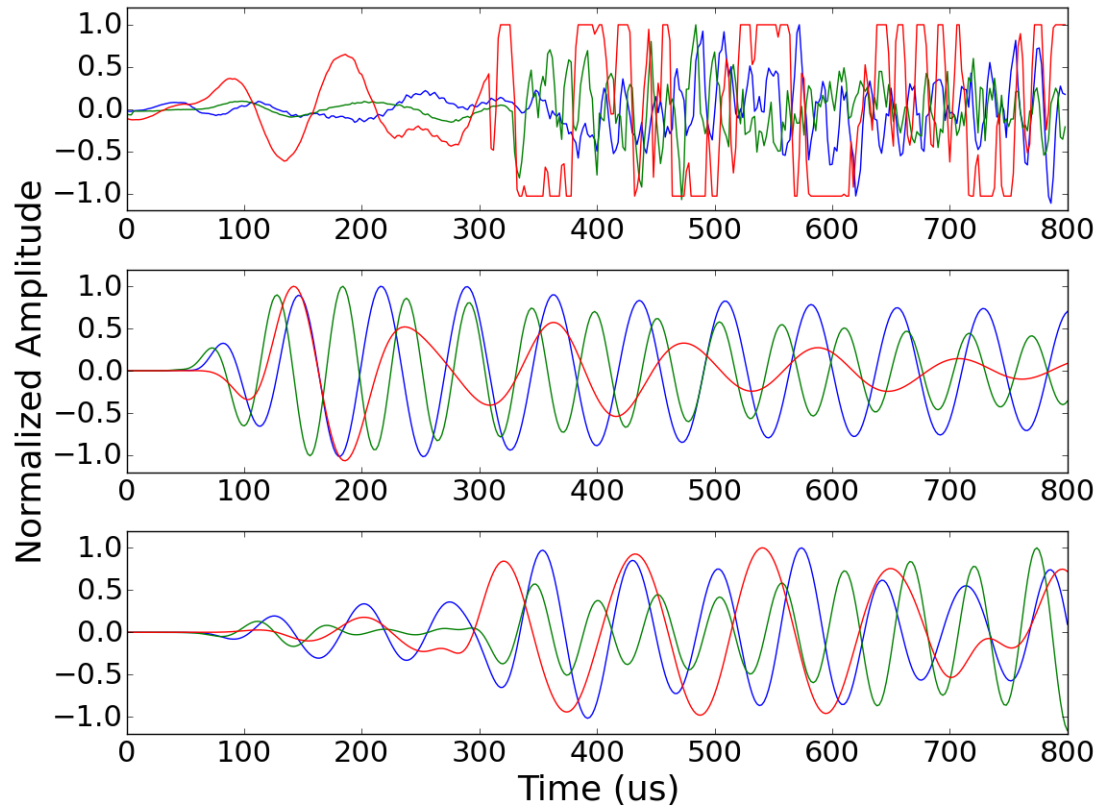
Figure 6.11. An example of an HC breakdown signal (top) with normal acceleration plots from simulations of breakdown thermal shock on the button (middle) and combined thermal shocks of the button and fill gas (bottom). Blue, red, and green traces are from microphones S0, S5, and S1 respectively.

a similar sudden rise in amplitude after several hundred microseconds (middle plot in Fig. 6.11). This led to the idea that perhaps the fill gas had a significant role.

COMSOL has an acoustics module that supports modeling sound in fluids and gases. It can also be combined with the solid mechanics module to model what are called solid acoustic interactions. The HC simulation was modified to include the fill gas and the solid acoustic interaction physics which models the sound through the gas and its transformation into solid mechanical waves at the inner cavity surfaces.

Assuming that an electrical arc between the buttons causes a fill gas shock wave, an axial, gas-flow line source was added to the HC model. This allows the specification of a gas flow that is symmetric about and emanates from the cylindrical axis. A thin Gaussian was naïvely chosen as the gas flow function:

$$Q(t) = \frac{Q_0}{\sigma\sqrt{2\pi}}e^{-(t-\mu)^2/2\sigma^2},\tag{6.3}$$

where Q_0 is the peak flow rate, μ is the peak time, and σ is the width. The intent was to approximate the initiation of a shockwave due to rapid heating of the fill gas near an axial breakdown spark. The value of Q_0 is arbitrary and only serves to scale the acceleration plot amplitudes. The value of μ was set to 35 µs and represents the time at which breakdown occurred for the real signal. The time constant from Sec. 2.3 was used for σ (i.e. the temporal width of the Gaussian). The real data were taken from an HC run where the cavity was filled with 300 PSI of nitrogen. The bottom plot of Fig. 6.11 shows the normal acceleration at the microphones that resulted from this gas shockwave simulation.

The speed of sound in 300 PSI of nitrogen at room temperature is $v_{N_2} = 352$ m/s. A naïve estimate of the distance that the nitrogen pressure wave would travel in $\Delta t = 300 \pm 10$ µs (the approximate time at which the amplitude jump

occurs) is

$$\Delta r = v_{N_2}\Delta t = (352 \text{ m/s})(300 \text{ µs}) = 10.6 \pm 0.4 \text{ cm}, \quad (6.4)$$

where the uncertainty on Δt is a judgment based on inspection of the signal plot and is propagated to determine the uncertainty on Δr.

The only inner cavity dimension that is similar to Δr is the inner radius of the cavity which is 11.4 cm (within two error intervals of Δr). This suggests that the sudden rise in amplitude in the RF hammer signals is due to the gas shockwave impacting the inner cavity circumference.

In fact, plotting the 2D pressure field of the gas is also very suggestive of this hypothesis. The plots in Fig. 6.12 show the gas shock wave before and after this impact which appears to take place at about 325 µs. This is a slightly more reasonable value for the propagation time as it yields the inner cavity radius:

$$\Delta r = (352 \text{ m/s})(325 \pm 10 \text{ µs}) = 11.4 \text{ cm} \pm 0.4 \text{ cm}. \quad (6.5)$$

There is also the propagation time through the ring wall which must be accounted for:

$$\Delta t_{\text{ring}} = \frac{3.8 \text{ cm}}{v_{\text{SS}}} = 6.6 \text{ µs}, \quad (6.6)$$

where $v_{\text{SS}} = 5.79$ km/s is the speed of pressure waves in stainless steel. The uncertainty on the arrival time (± 10 µs) can also accommodate Δt_{ring}, although admittedly, as stated above, this uncertainty was chosen somewhat arbitrarily.

Although it was likely not possible to localize breakdown in the HC via the cavity wall thermal shock, it may have been possible to do so using the gas shockwave if the microphones were mounted with epoxy or hot glue to make the signals more consistent. Also, additional modifications of the localization algorithm would have likely been needed given that this scenario ceases to be one where a single sound point source is being localized. Instead, information about the shockwave origin
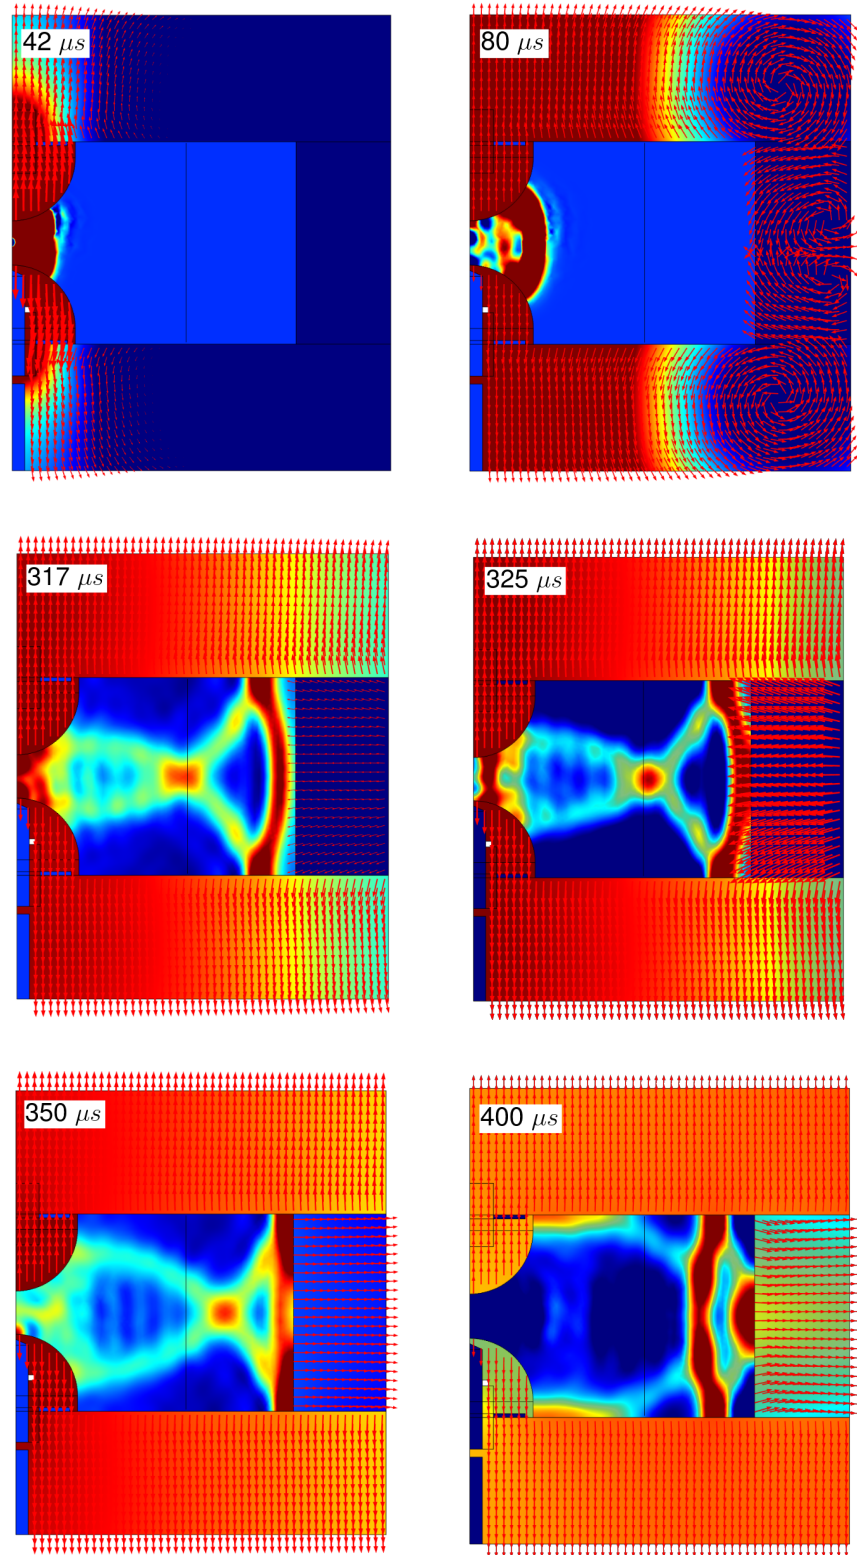
Figure 6.12. 2D plots of HC fill gas pressure (central region without arrows) as well as the solid displacement field of the cavity walls (arrows) before and after the gas shockwave impacts the inner cavity circumference. Pressure is depicted as a cool-to-warm color scale, with the highest pressure being dark red.

was transferred to the ring walls by a continuous set of impacts on the entire inner circumference. Given these considerations and that this was a special case involving a gas-filled rather than an evacuated cavity, efforts continued on other vacuum cavities to localize the more general case of thermal shock to the inner cavity walls.

6.4 Algorithm Self-Consistency

Recent dielectric loading experiments with the HC required relatively frequent breaks in running in order to change fill gas pressures and mixtures, and to change the ceramic, dielectric inserts (referred to as "donuts"). This yielded an opportunity to perform several acoustic measurements that required multiple microphone position changes. The purpose of these measurements was to verify the self-consistency of the localization algorithm.

For these tests on the dielectric-loaded HC (DLHC), three PICASSO microphones were epoxied to the downstream surface of the HC at a radius of 11 cm and distributed at 120 degree intervals. Signal and power leads were routed to patch panels that consisted of five-pin Phoenix screw terminal connectors. Microphone power was supplied by a ± 12 V wall adapter. Signals were routed directly to one of the NI-9221 DAQ modules. After running with the microphones in their initial positions, the cavity was subsequently run with the microphones rotated around the cavity axis by 120 degrees and then by 90 degrees with respect to the original positions (Fig. 6.13).

Since the three microphones are distributed evenly by angle, the 120 degree rotation could be used to directly verify whether the same signals were being reproduced by the microphones by comparing the average RF hammer from the two configurations. Figure 6.14 shows a comparison of the three microphone RF hammer signals between the original (top row) and 120-degree-rotated (bottom row) microphone positions. Each column shows the average RF hammer signal from the same
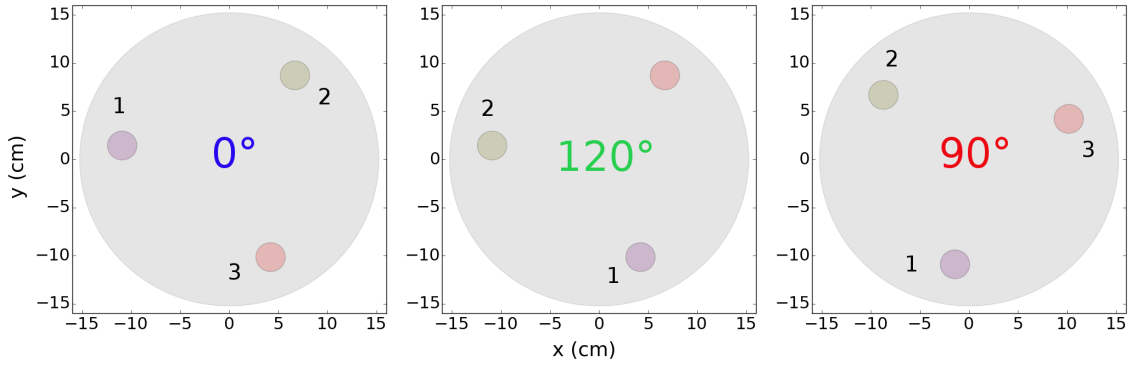
Figure 6.13. Positions of the microphones for the algorithm self-consistency tests with the DLHC.

position. The good agreement in shape between different microphones attached to the same position on the cavity suggests that sound at a particular position will be picked up similarly regardless of the microphone.
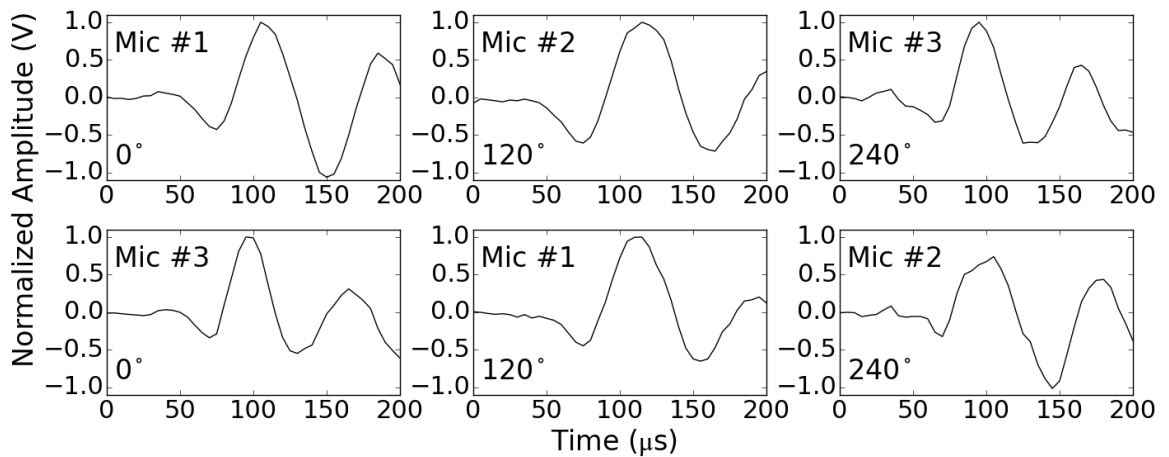
Figure 6.14. A direct comparison of the average RF hammer between the original and 120 degree-rotated locations of the DLHC microphones.

The next step in verifying the self-consistency of the algorithm was to run the localization algorithm on some breakdown signals. The goal was not necessarily to correctly localize the signal, but rather to observe whether the results of localizations from different microphone positions agreed. The 90 degree rotation was added as

an additional test of the algorithm. The 0 degree and 120 degree configurations ought to agree on the position of identical events since the signals being received are also identical. Positioning the microphones at unique positions demands that the algorithm produce a similar result from signals that are dissimilar to the signals from the previous two configurations.

Now, there is no guarantee that different breakdown events will originate from the same cavity wall location. That said, photographic evidence of breakdown damage to the ceramic donut used for these experiments suggests that there were only four general regions where breakdown was occurring. The approximate centers of these regions are marked with crosses on the plots in Fig. 6.15.
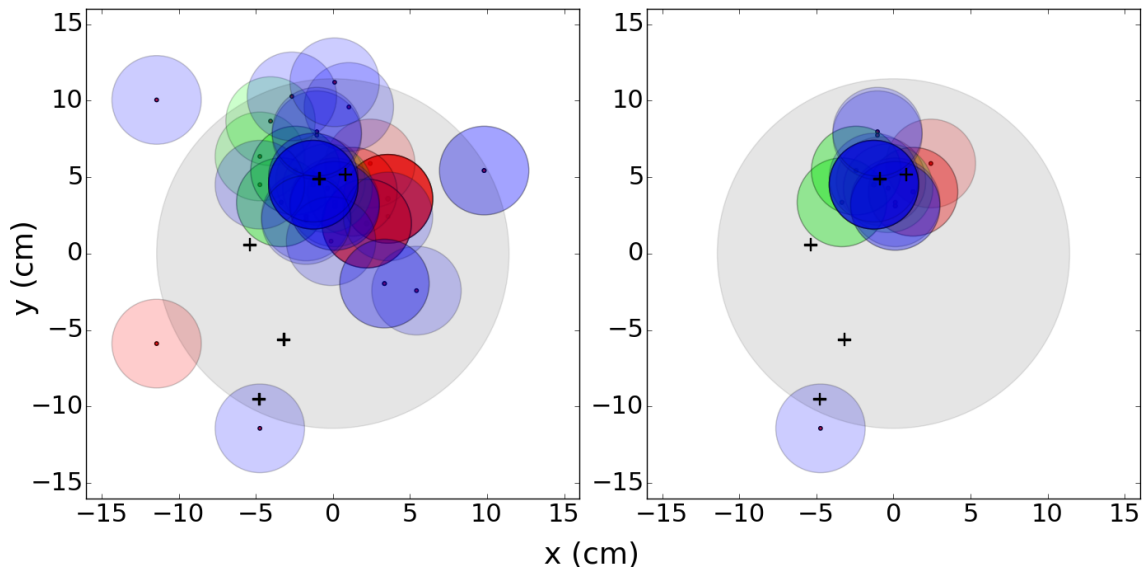
Figure 6.15. Left: all predicted locations from the DLHC microphone rotation experiment. Right: select predicted locations that are within the DAQ resolution of observed damage. The colored circles represent the position uncertainty of predicted locations from the 0 degree (red), 120 degree (green), and 90 degree (blue) rotations. The large, gray circle represents the inner surface of the end plate. The crosses are locations of observed donut damage.

Fig 6.15 shows a summary of the results from this experiment. The plot on the left shows all of the spark source location predictions. The plot on the right shows

unique spark source locations that are within the 2.9 cm DAQ spatial resolution of an observed damage spot on the donut (assuming a wave speed of 5.79×10^5 cm/s, the pressure wave speed in stainless steel, with a DAQ sampling rate of 200 kHz).

These results could imply detection of events that produced the damage on the donut above the x-axis. That said, the RF coupler is centered at 6.3 cm on the positive y-axis. It is also possible that RF hammer sound from the RF coax is not being adequately canceled out and is biasing the results. Another possibility is that much of the breakdown arcing is being produced by the RF coupler itself. It has been recommended that the RF coupler be inspected the next time the cavity is opened. Regardless, there seems to be a degree of agreement among the three different position trials.

6.5 Simulation Reconstruction

Given the simple geometry and the fact that the end plates are easily removed for inspections, the MC is one of the more ideal cavities for testing acoustic localization of breakdown. A simple 3D simulation of an MC end plate was used to collect statistics on the accuracy of the localization algorithm. The MC end plate has the advantage of being an almost perfect disk. This makes full 3D simulations relatively fast.

A set of 1,000 sparks with uniformly random source locations were simulated. Time series data were then extracted from the simulation output for the acceleration normal to the outer cavity surface at each of the virtual microphone locations. The localization algorithm described in Sec. 5 was then applied to each data set. Figure 6.16 shows the accuracy of the source location reconstructions using all of the predictions (top plot) and omitting predictions that have true source locations greater than 7.1 cm or half the endplate radius (bottom plot).

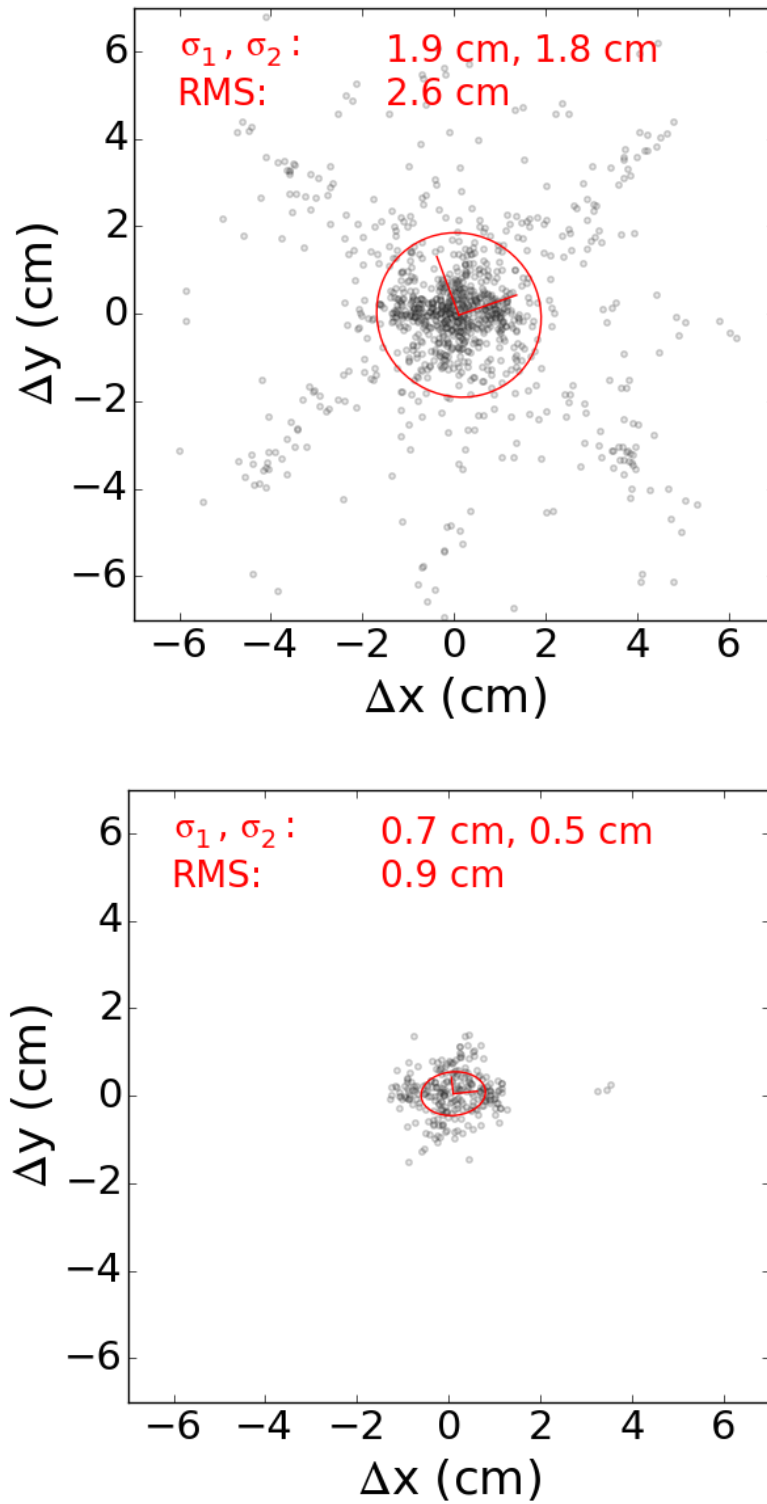Figure 6.16. Accuracy of reconstruction of the positions of 1,000 uniformly random simulated sparks. Dots represent the residuals between the predicted and true source locations. The red ellipses represent the boundaries containing 39% of the residuals, and have semi-major and semi-minor axes equal to the standard deviations of the residual distributions' principal components (σ_1 and σ_2).

The points on the plots of Fig. 6.16 indicate the residuals between predicted and true source locations. These distributions were then modeled as bivariate Gaussians, and the standard deviations of the principal components were calculated (σ_1 and σ_2 in the plots of Fig. 6.16) [65]. The red ellipses were drawn with the semi-major and semi-minor axes equal to the principal component standard deviations, and represent the boundary that contains 39% of the residuals. These standard deviations were 1.9 cm and 1.8 cm for the top plot, and 0.7 cm and 0.5 cm for the bottom plot. The RMS was then calculated by simply taking the quadratic sum of the standard deviations. These were 2.6 cm for the top plot and 0.9 cm for the bottom plot.

There are two reasons for providing the second plot. First, reconstructions of sound source locations that are close to the endplate radius are expected to be less accurate compared with source locations that are closer to the beam axis. This is due to interference from wave reflections at the radial boundary. Second, due to the electric field being most intense at the beam axis, sparks are more likely to occur near the center of the endplates. Thus, in practice, the accuracy ought to be closer to that seen in the second plot. Overall, though, the accuracy shown in the first plot ought to be sufficient for differentiating between sparks that occur near the high-field region at the beam axis and those that occur near the RF coupler on the circumference.

6.6 Experimental Reconstruction

Eight microphones were installed on the Modular Cavity (MC). Four were attached in an axially-centered, rectangular arrangement on each of the two end plates (refer to Fig. 5.5). The horizontal and vertical microphone separations are 10 cm and 12 cm respectively. The signal leads for each set of four microphones on the upstream and downstream surfaces were combined into a single connector that interfaces with a pre-amplifier (see Sec. 3.2.2). Two twenty-foot long shielded, two-pair signal cables connected the pre-amplifier output to the input of a PGA. The eight PGA outputs

were then distributed evenly across the four NI-9221 analog input cards. This allowed the sampling rate to be increased to 400 kHz since only two channels per card were used at a time. See Sec. 3.2 for more details on the DAQ hardware.

Data were taken during an MC run with a 3 T magnetic field ($B = 3$ T) in April, 2016. During this run, 87 breakdown events were recorded. After the end of the magnetic field run, the cavity was taken out of the solenoid bore and inspected in a clean room. A cluster of pits was observed and mapped out [66]. Curiously, the number of pits was nearly three times the number of recorded breakdown events. This discrepancy has not yet been accounted for by the team running the MC experiments. This further complicates the problem of directly associating particular spark damage with specific acoustic events. A statistical analysis of the results is therefore warranted.

The acoustic signals from the 87 events were collected and fed to the localization algorithm described in Ch. 5. The results of applying the localization algorithm to the four upstream and four downstream acoustic signals for the 87 breakdown events are shown in Fig. 6.17. For each plot, the red dots represent the location of observed damage pits. Black dots are the predictions produced by the algorithm. The red ellipses are centered on the damage centroid (2D mean) and mark the boundary of the one and three standard deviation regions of the damage locations distribution. Similarly, the black ellipses are centered on the prediction centroid and denote the one and three standard deviation regions of the predicted locations distribution. The three standard deviation ellipse was drawn to denote the region that contains most (94.2% to be exact) of the particles.

The damage pits have a one-to-one mirror symmetry between the upstream and downstream end plates. The centroids of the damage sites are both at (0.6 cm, −1.6 cm). The standard deviations of the principal components of the
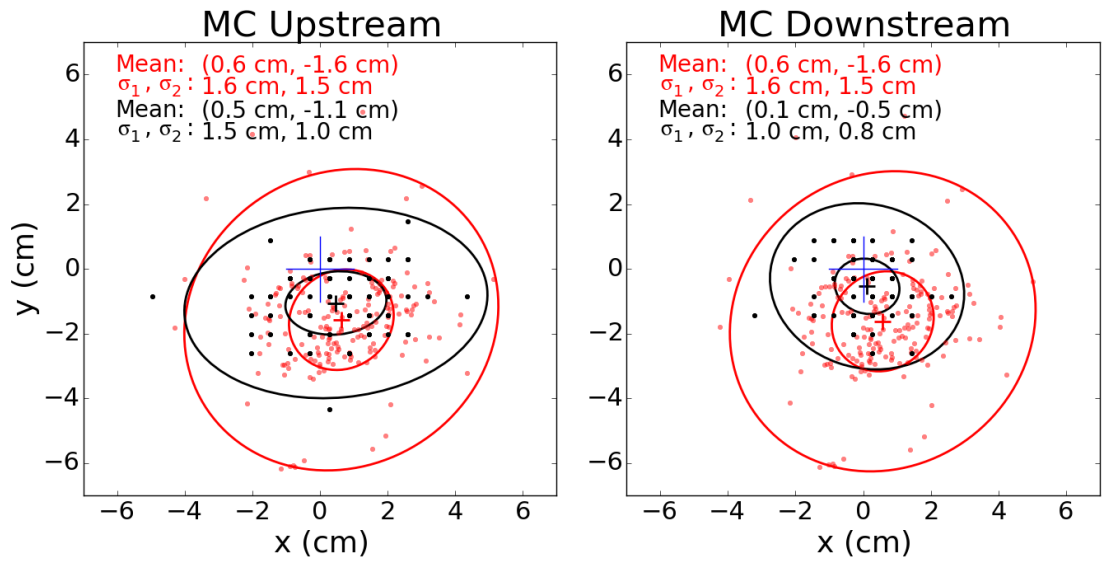
Figure 6.17. 195 breakdown spark damage site locations at $B = 3$ T (red dots) and 87 spark endpoint predictions (black dots) on the MC endplates. Also plotted are the boundaries of the one and three standard deviation regions for the damage and prediction location distributions (red and black ellipses respectively). Furthermore, the red and black crosses denote the damage and prediction location centroids. Lastly, the blue cross indicates the endplate center.

pit damage location distribution are 1.6 cm and 1.5 cm. As for the predictions, the centroids are at (0.5 cm, −1.1 cm) and (0.1 cm, −0.5 cm) while the prediction distribution standard deviations are 1.5 cm and 1.0 cm upstream, and 1.0 cm and 0.8 cm downstream. The separations between the damage and prediction centroids for the upstream and downstream end plates are 0.5 cm and 1.1 cm. This puts the damage and prediction centroids for the upstream endplate inside each other's standard deviation regions, and nearly so for the downstream endplate.

This analysis indicates that, when taken together, the location predictions agree with the locations of observed damage to a similar degree as the simulation reconstructions using the accuracy determined for points within half the endplate radius. That is, the centroids differ by 0.8 cm averaged over the upstream and downstream endplates (again, 0.5 cm and 1.1 cm respectively). Quoted here again, the RMS of the simulation reconstruction residuals for points that are located within half the endplate radius is 0.9 cm (see Sec. 6.5).

6.7 Conclusions

The initial solid mechanics simulations in COMSOL and their accompanying experiments with the AMC were useful in better understanding the importance of boundaries between cavity components (Sec. 6.1). The main peaks in the frequency spectrum of the disk-and-ring model could be matched to the real spectrum by adding a Thin Elastic Layer boundary condition and tuning its spring constants.

Additional evidence for the importance of component boundary layers and consistent mounting of microphones was found through simulating the RF hammer on the HC models (Sec. 6.2). Frequency spectra from simulated signals could not be reconciled with those from real HC signals until Thin Elastic Layer boundaries were added and their spring constants fine tuned. This is especially true with respect to

adding virtual polyurethane foam to represent the double-sided tape of the earliest data taking runs.

Although coupling inconsistencies due to the use of double-sided tape made it impossible to apply the localization algorithm to HC breakdown signals, these signals revealed interesting physics happening in the cavity. The signals acquired during HC breakdown are largely due to the shockwave in the fill gas (Sec. 6.3). An obvious feature in the signals can be explained by the effect of this shockwave impacting the inner circumference of the cavity. This might yet make it possible to localize breakdown in gas filled cavities, but this was deemed out of scope at the current stage of this research.

Tests on the DLHC demonstrated self-consistency of signals coming from particular positions on the cavity as well as the localization algorithm itself (Sec. 6.4). Due to the general location of the predictions near the RF coupler, it has been suggested that the cavity's coupler region be inspected for damage.

Simulated MC sparks were reconstructed giving some basic statistics with which the localization algorithm could be judged in terms of its accuracy (Sec. 6.5). From this evidence it is reasonable to assume that sparks can be localized on the MC to within 2.6 cm RMS on average. If desired, localization can be restricted to within half the cavity radius (7.1 cm) of the beam axis in which case the accuracy improves to 0.9 cm RMS.

Reconstruction of breakdown signals captured during a 3 T magnetic field run of the MC was successful (Sec. 6.6). Though associating individual pits with specific breakdown events is not feasible, the centroids of observed pit damage and source location predictions differ by only 0.8 cm on average between the two end plates and are within each other's standard deviation regions on the upstream endplate.

CHAPTER 7

CONCLUSION

An algorithm called Accumulated Correlation was chosen as the basis for producing RF cavity breakdown spark source location predictions from signals produced by microphones attached to the outer endplate surfaces of RF cavities. To improve accuracy, the signals are first conditioned by subtracting the RF hammer background noise and then taking the absolute value of the signals. Although the shear wave speed in copper (2.325×10^5 cm/s) was the best choice for the wave speed used by the algorithm before signal conditioning in simulated reconstructions, a wave speed of 1.86×10^5 cm/s produced more accurate results after taking the absolute value of the signal. This lower wave speed is compatible with the speed of low-frequency Lamb waves.

Acoustic localization of RF cavity breakdown has been demonstrated on the Modular Cavity (MC) during operation with a 3 T magnetic field in the MuCool Test Area at Fermilab. While experiments would need to be done to match specific breakdown damage to acoustic breakdown data in order to be certain, simulations and analysis of acoustic data as well as observations during inspection of the MC suggest that the method for acoustic localization of breakdown presented here is accurate on average to within 2.6 cm RMS and within 0.9 cm RMS using a radial cutoff of half the cavity radius (adequate to describe most of the sparks) on cavities similar to the MC. This should be sufficient for differentiating between breakdown near the beam axis and that which is due to a faulty or poorly designed RF coupler. Applying the technique to additional test cavities would help bolster this claim.

It is plausible that the technique could be improved by installing additional microphones. It would also be worth the effort to try a statistical localization algorithm such as one based on Bayesian analysis to see if the performance cost is acceptable for a real-time diagnostic tool.

The High-Pressure Cavity is an interesting special case in that the sound generated during breakdown is largely due to the fill gas shockwave and not the thermal shock on its button electrodes. If production gas-filled cavities were to be developed, it might be worth the effort to attempt localization using the time of arrival of the gas wavefront with microphones located on the outer circumference instead of on the outer end plate surfaces.

7.1 Lessons Learned

Most fundamentally, it is essential for there to be a good, consistent coupling between the microphones and the surface to which they are applied. Epoxy has good, consistent acoustic transmission; and there is evidence that hot glue could be as good or better than epoxy if outgassing is not an issue (i.e., the cavity is not in a vacuum vessel). Epoxy is a little more forgiving when trying to position the microphone in hard to reach places since it takes minutes to harden. On the other hand, hot glue would not require the installer to stand in uncomfortable positions while holding the microphones in place. While off the shelf or custom clamps are always preferred, it is not always feasible to clamp microphones to certain surfaces. Hot glue has the disadvantage that it is sometimes difficult to remove the microphones. Attempts to reheat the thin layer of glue with the glue gun tip were not successful. Microphones attached using epoxy are relatively easy to remove using a flat head screwdriver as a chisel on the hardened epoxy layer.

Active microphones are very good for this technique. That said, they have disadvantages such as increased sensor size due to the on-board amplifier, and additional cabling for power. Passive microphones can be made quite small, but their signal leads can be no longer than about a meter without adding pre-amplifiers. This is probably sufficient for most cases, but it somewhat defeats the purpose of having a long signal cable in order to keep electronics away from the experiment. Future tests of the MICE cavity could employ pre-amplifiers as was done on the MC. This would be a good demonstration of the feasibility of using passive microphones on enclosed cavity systems.

On the subject of simulations, it is important to consider the interfaces between cavity components. As was shown in Sec. 6.1 and Sec. 6.2, adding Thin Elastic Layer boundary conditions (or the equivalent in whatever simulation package is being used) and adjusting their spring constants is crucial for getting a good match to the frequency spectrum of the cavity.

7.2 Alternative Applications

If localization based on the fill gas shockwave were successful, this technique could likely be applied to other scenarios outside of accelerator physics where an estimate is desired for the source of some sound inside a container. For example, a failing component in a large gear box might be generating noise. Currently the failing component could be located by using a remote array of microphones and advanced beamforming techniques [58], but this would require opening or removing the gear box enclosure. In private communication with colleagues from Illinois Institute of Technology's MMAE department, unpublished results have indicated that localizing the source of sound within an enclosure using such current techniques does not appear to be possible. For safety or other practical reasons, opening a machine's protective enclosure might not be an option. By leaving the protective enclosure in place and

attaching microphones directly on the outside of the enclosure, an estimate of the location could potentially be made without having to expose a technician to the dangers of working near the moving machinery inside.

APPENDIX A

CAVITY ENGINEERING DRAWINGS
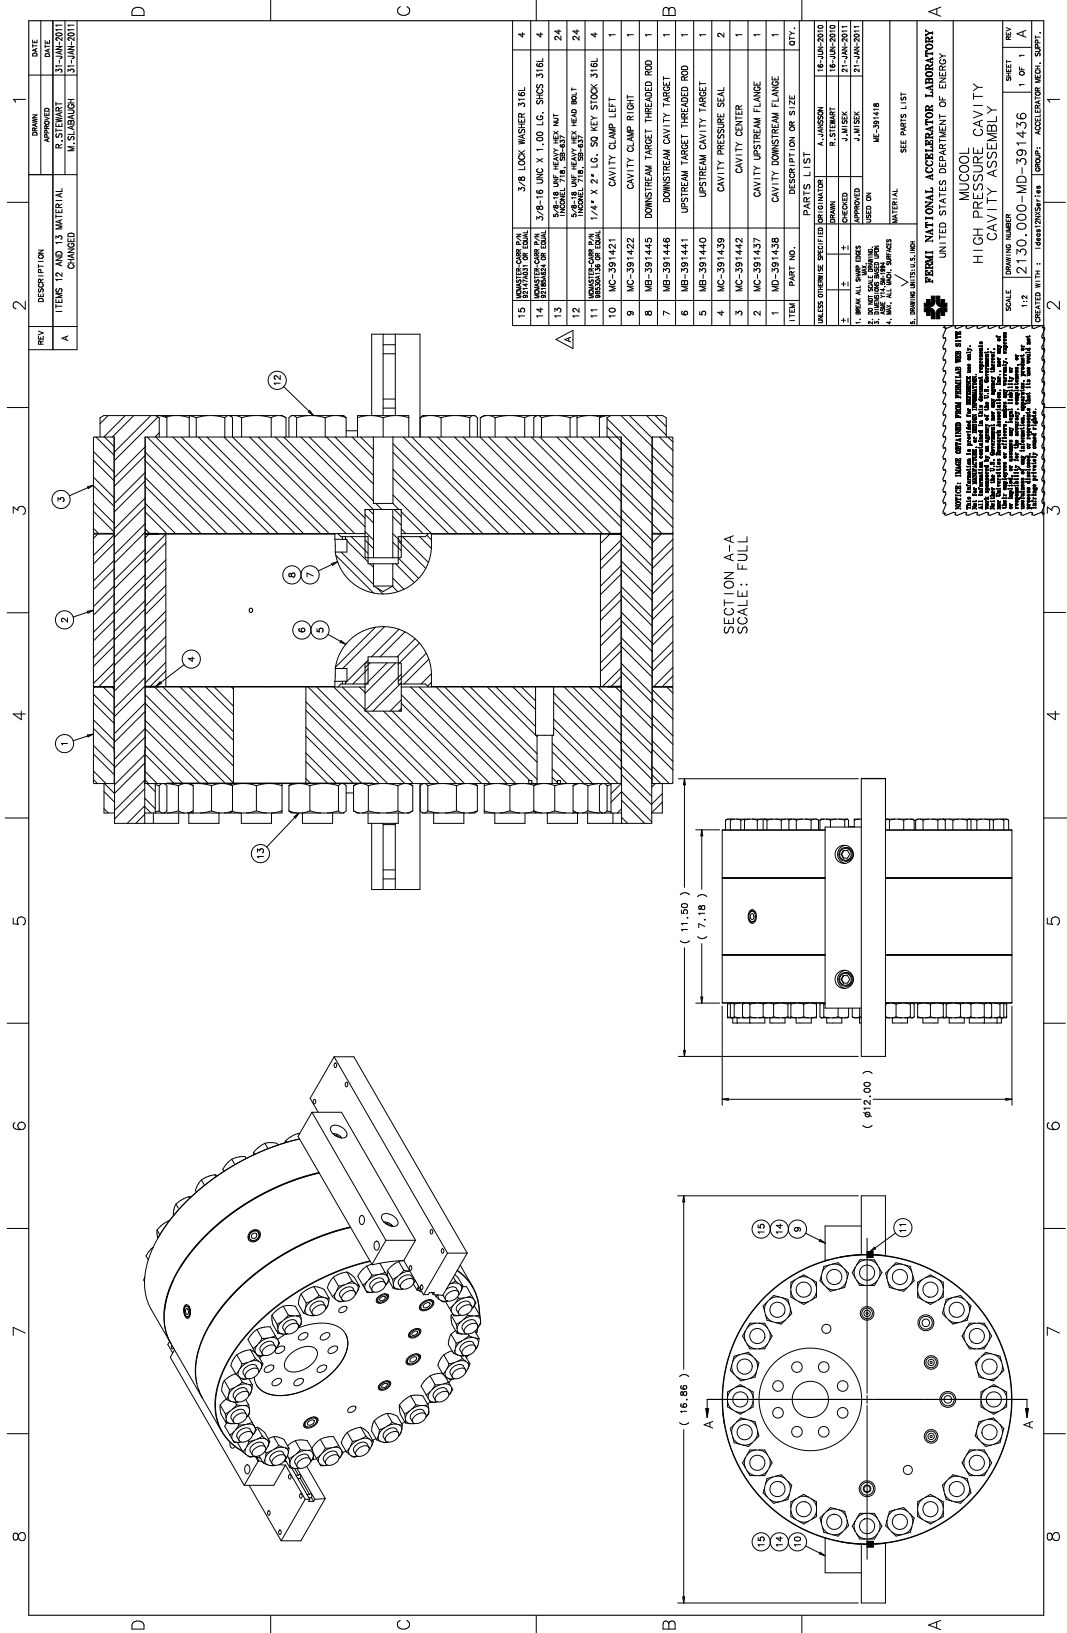
Figure A.1. HC assembly drawing [29].
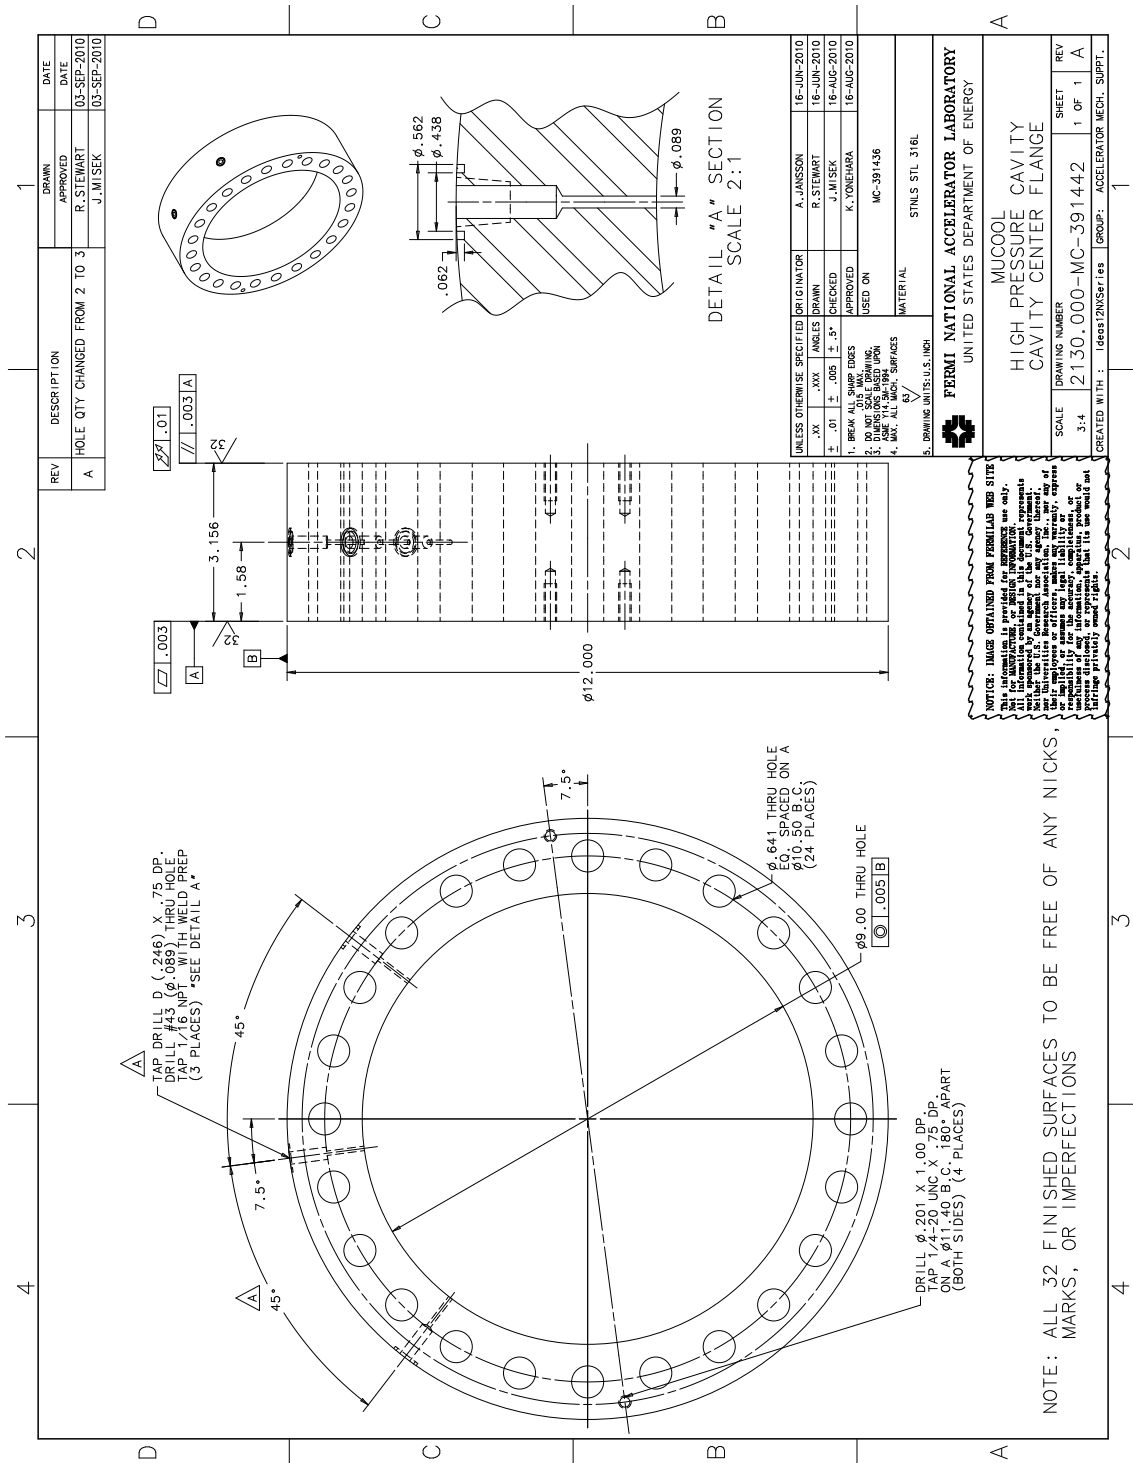
Figure A.2. HC ring drawing [29].

Figure A.3. HC downstream end plate drawing [29].
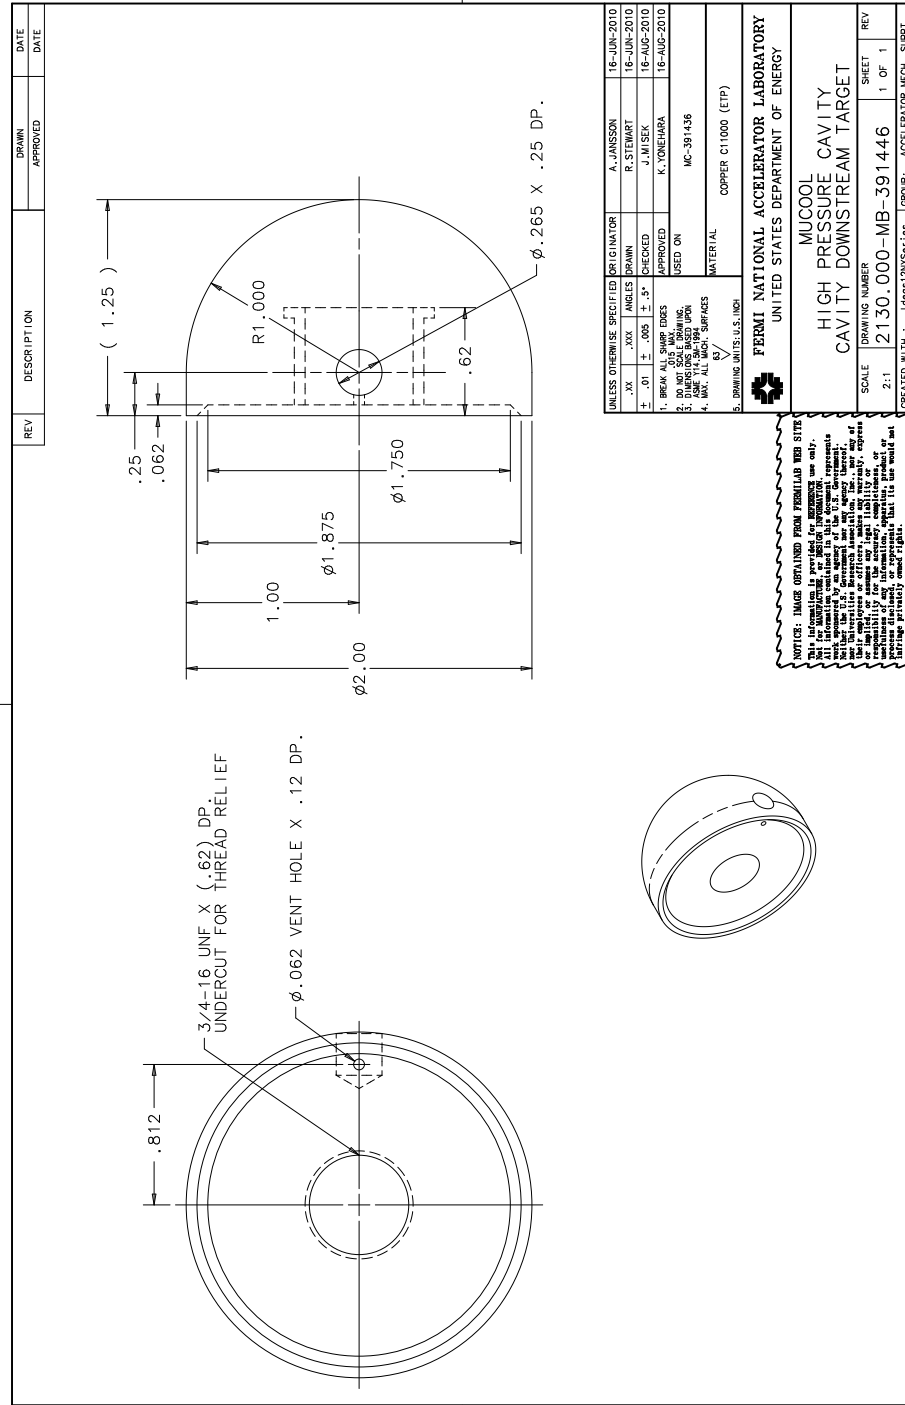
Figure A.4. HC button/electrode drawing [29].
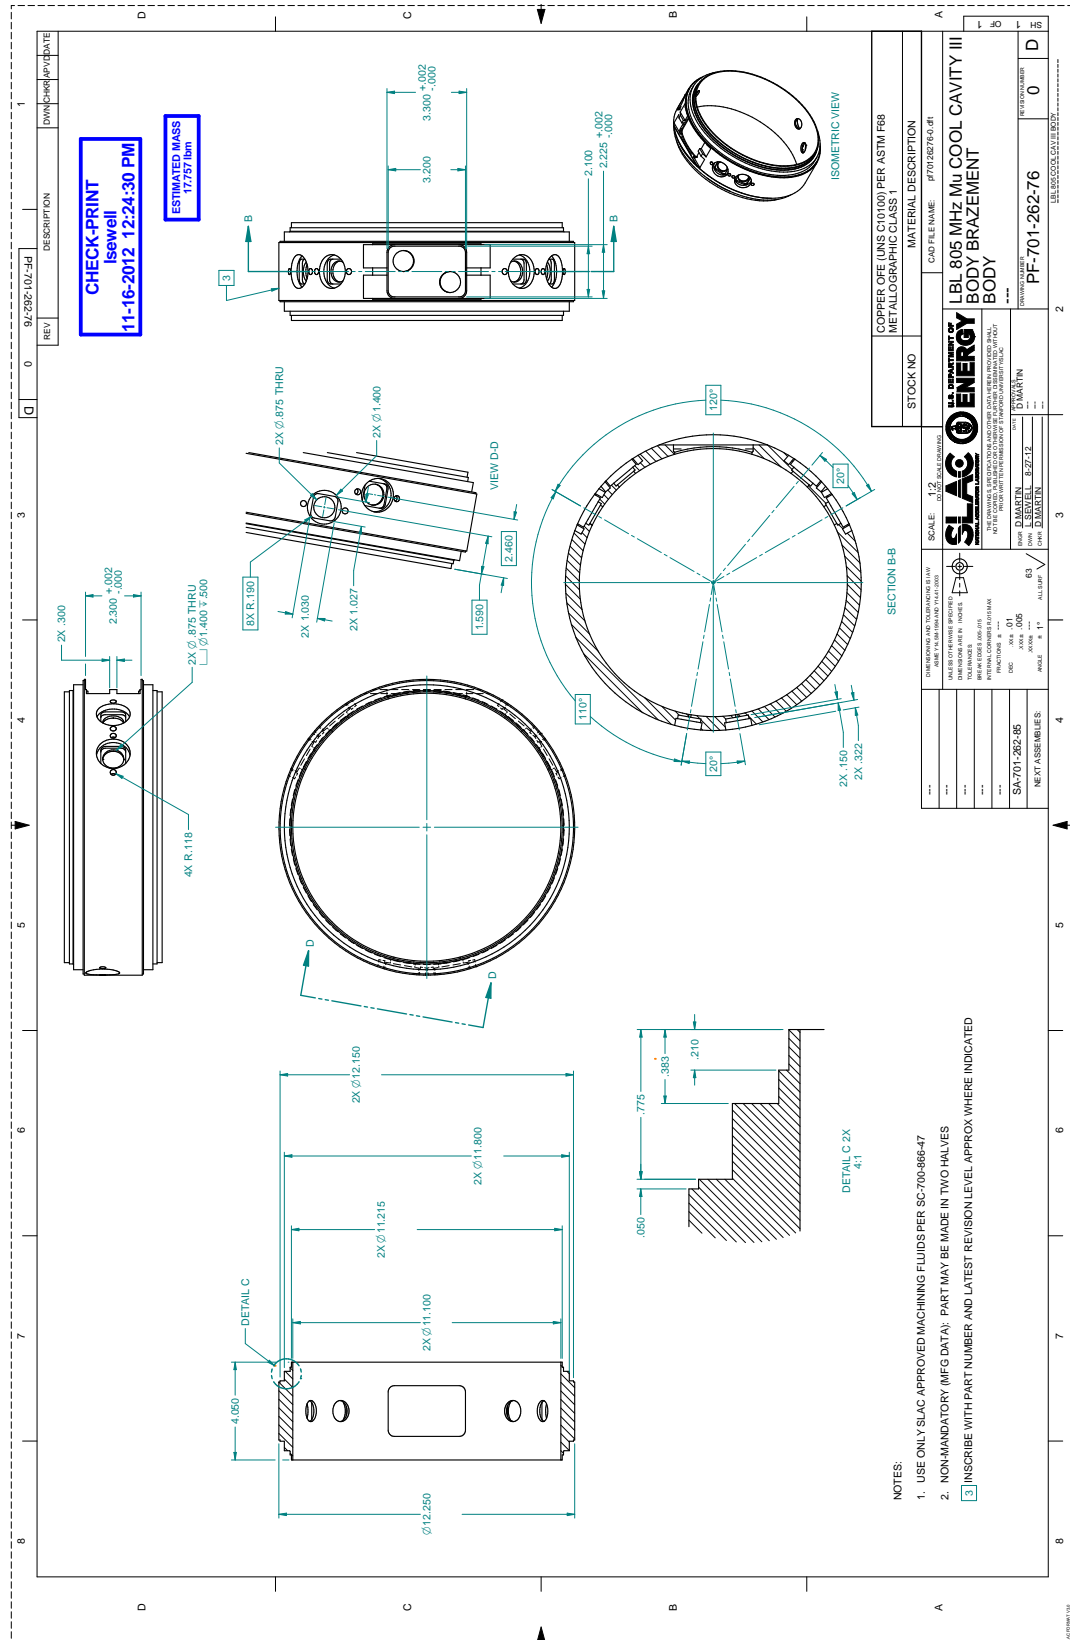
Figure A.5. MC body drawing [29].
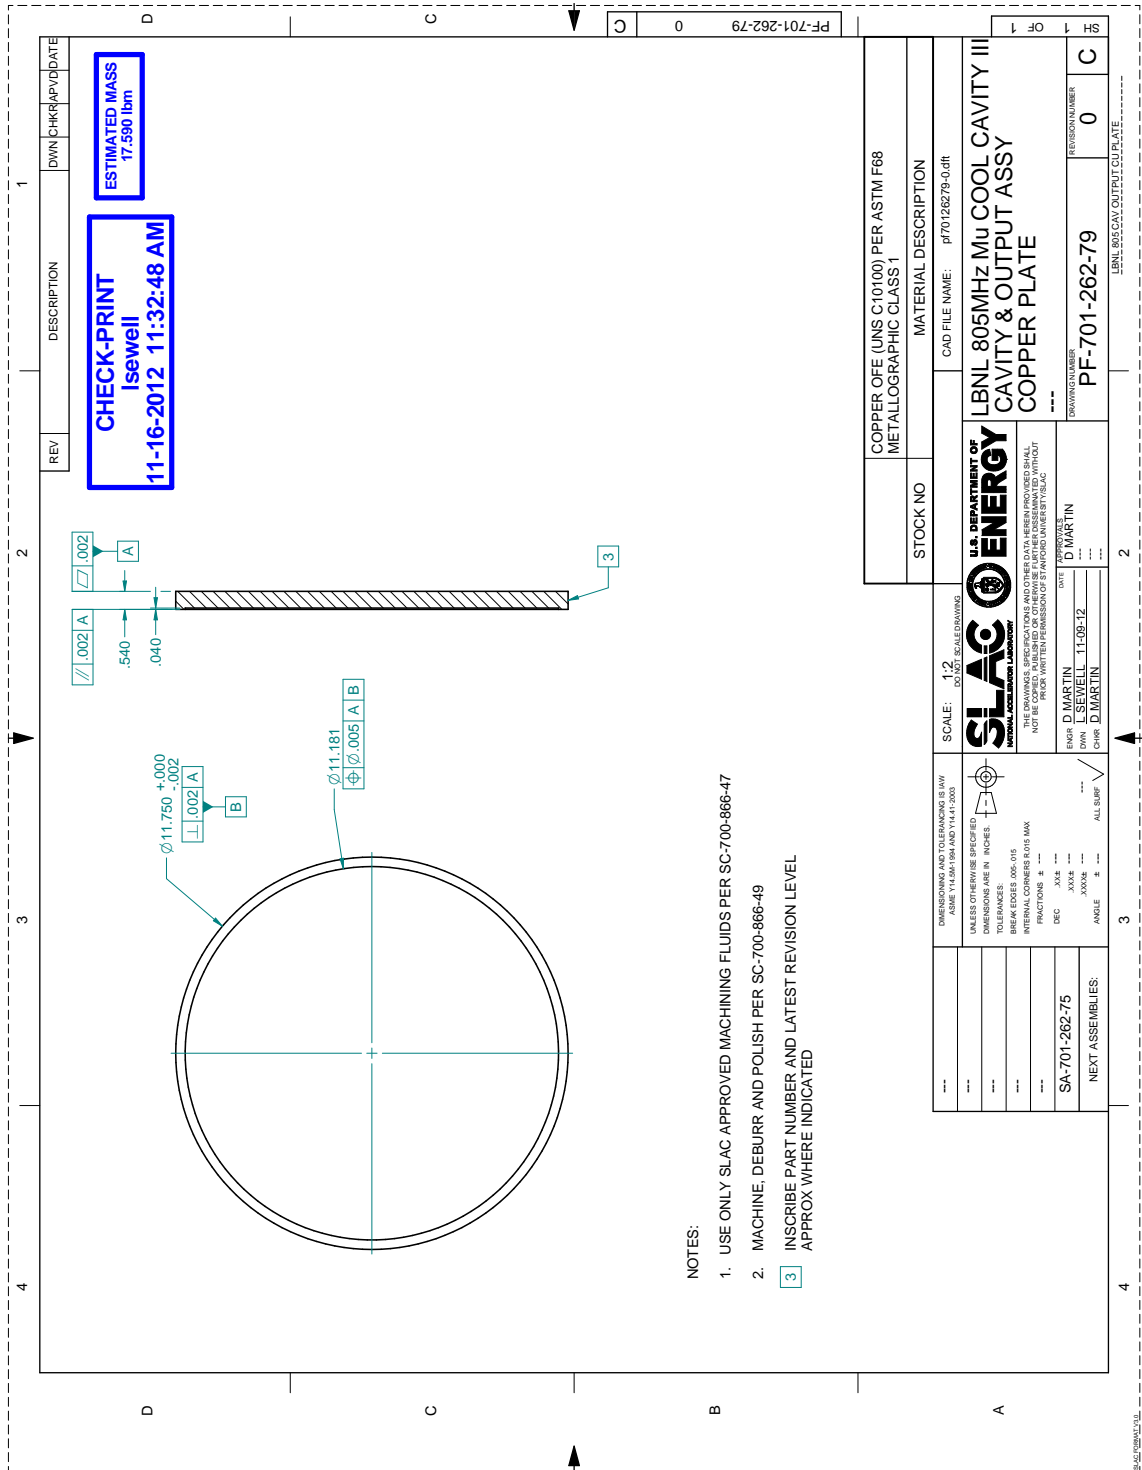
Figure A.6. MC end plate drawing [29].

APPENDIX B

CUSTOM HARDWARE SCHEMATICS
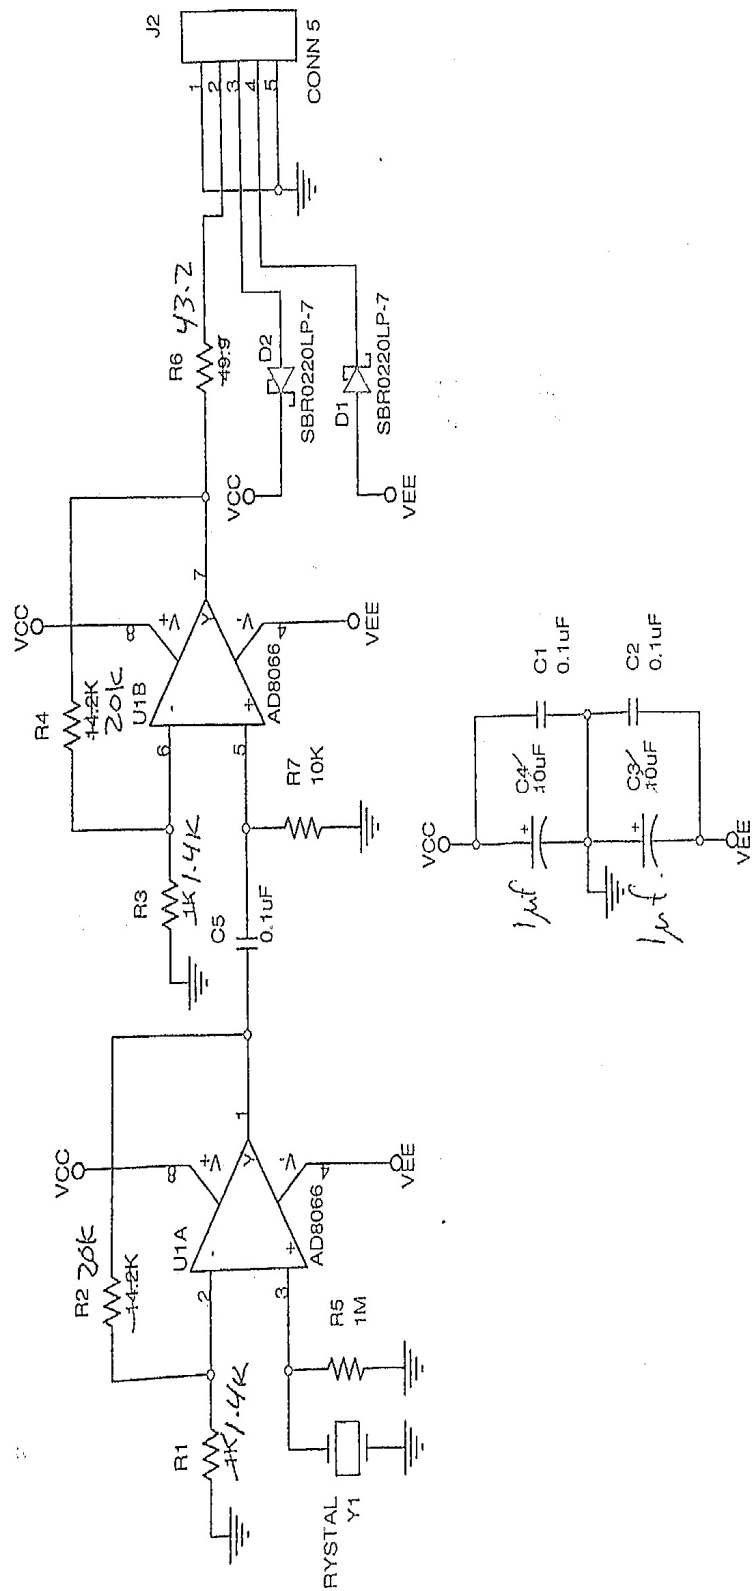
Figure B.1. Schematic of the PICASSO, active microphones [39].
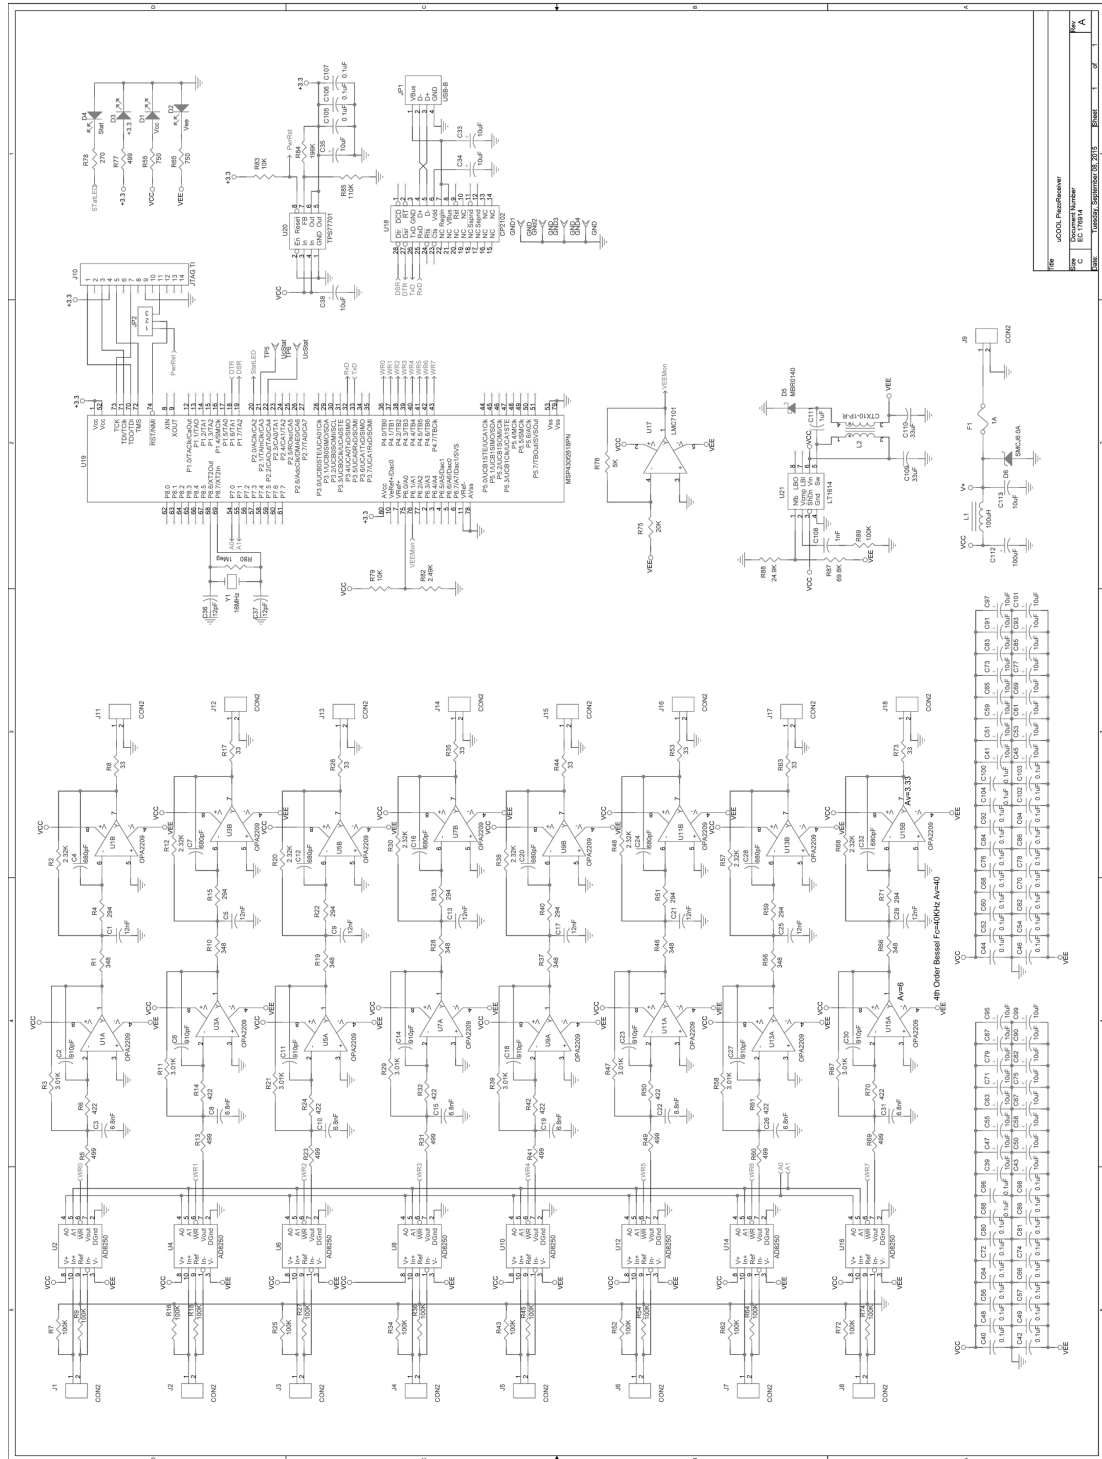
Figure B.2. Schematic of the programmable gain amplifier boards used with the passive microphones [67].
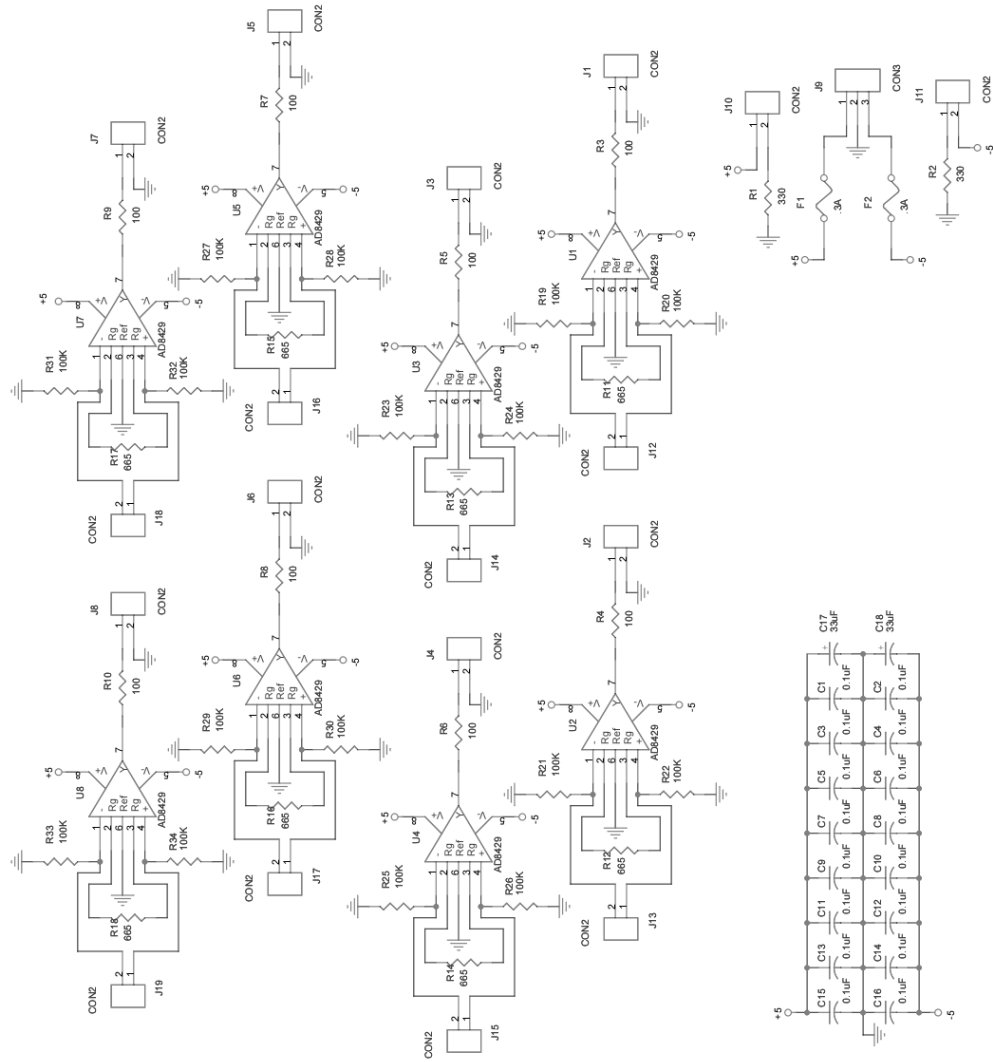
Figure B.3. Schematic of the pre-amplifier boards used with the passive microphones to drive the signal cables [67].

APPENDIX C

ACRONYMS

Table C.1. Acronyms

Acronym	Expansion
AC	Accumulated Correlation
AMC	Aluminum Mock Cavity
ASC	All-Seasons Cavity
AAL	Archive and Analysis Loop
BAC	Breakdown Acoustics Controls
DPL	DAQ Producer Loop
HC	High-Pressure Cavity
MC	Modular Cavity
MICE	Muon Ionization Cooling Experiment
MICEC	MICE Cavity
NI	National Instruments
PGA	Programmable Gain Amplifier
PFI	Programmable Function Input
RF	Radio Frequency
RHQ	RF Hammer Queue

SUS	Stochastic Universal Sampling
SSD	Sum of the Squared Differences
STQ	Spark Trigger Queue
TDL	Trigger Detection Loop
TDOA	Time Difference of Arrival
UDP	Universal Datagram Packets
UHM	UDP Heartbeat Monitor
UHT	UDP Heartbeat Transmitter

APPENDIX D

PYTHON IMPLEMENTATION OF LEAST SQUARES POLYNOMIAL FITTING

D.1 Fit Polynomial

D.1.1 Description.
X are parameters and Y are corresponding measurements with estimators $f(x_i) = \sum\limits_{j=1}^{m}(b_j h_j(x_i))$, where h_j are m linearly independent functions. In this case $h_j(x_i) = x_i^j$. To find the coefficients, b_j, we must solve the linear equation $Y = HB$. The least squares solution is the Moore-Penrose Pseudo Inverse of H times Y: $B = (H^T H)^{-1} H^T Y = DY$. To accommodate measurement error, the covariance matrix can be included so that $B = (H^T V^{-1} H)^{-1} H^T V^{-1} Y$. In this implementation, the errors are passed to the function as a list the same size as Y. These then make up the diagonal elements of V.

Table D.1. Fit Polynomial Function Parameters

Parameter	Type	Description
X	numpy.ndarray	list of parameter/independent variable values
Y	numpy.ndarray	list of measurement/dependent variable values
order	scalar,optional	The order of the polynomial to be fit (default 2)
yerr	scalar,optional	list of measurement errors (default 1, not weighted)
Returns	numpy.ndarray	list of polynomial term coefficients

D.1.2 Code Listing.

```python
def fit_polynomial(X, Y, order=2, yerr=1):
    H = np.empty((X.size, order+1))
    for j in range(order+1):
        H[:,j] = map((lambda x: x**j), X)
    D = np.empty((0,0))
    if type(yerr) == int:
        yerr = np.ones(len(X))
```

```python
        D = linalg.pinv(H)
    else:
        V = yerr*yerr*np.identity(len(yerr))
        Vi = linalg.inv(V)
        Ht = np.transpose(H)
        D = np.dot(linalg.inv(np.dot(Ht,np.dot(Vi,H))),np.dot(Ht,
        ↪ Vi))
    B = np.dot(D, Y)
    return B
```

APPENDIX E

PYTHON IMPLEMENTATION OF THE MODULAR CAVITY LOCALIZATION

ALGORITHM

E.1 Localize Spark

E.1.1 Description. This is the main spark localization function. It takes a set of signals, orders the microphones by when the associated signals arrived at their respective microphone, determines from this the octant of the closest microphone to the source, and then runs the Accumulated Correlation algorithm constrained to this quadrant.

Table E.1. Localize Spark Function Parameters

Parameter	Type	Description
times signals	numpy.ndarray	An array of time-domain signals
v	scalar	The sound wave speed in the cavity material
grid_size	int	The length of the sides of the test point grid in units of the number of points
Returns	list	A two-element list containing the x and y coordinates

E.1.2 Code Listing.

```python
import math
import numpy as np
import scipy.signal as sig


def localize_spark(times, signals, v, grid_size):
    dt = times[1] - times[0]
    mic_coordinates = numpy.array(\
        zip([5, -5, -5, 5], [6, 6, -6, -6]))
```

```
radius = 14.22   # cm
thickness = 1.37   # cm
octant = −11
signals = numpy.abs(signals)
return pgl.signal.accumulated_correlation(
    signals, dt, mic_coordinates, radius, thickness,
    [v,], grid_size=grid_size, octant=octant)
```

E.2 Accumulated Correlation

E.2.1 Description. The Accumulated Correlation algorithm estimates the source location of a signal that arrives at different times at multiple sensors. It starts by calculating the cross-correlation for each pair of microphones signals. For each test grid point, the expected time delay is calculated for each microphone. Then for each unique signal pair the difference in the expected time delay is used as an index into the cross-correlation vectors. The value in the cross-correlation vector is added to a running sum for the current test grid point. Finally, the test grid point with the largest sum is taken as the most likely source location of the signal.

Table E.2. Accumulated Correlation Function Parameters

Parameter	Type	Description
signals	numpy.ndarray	An array of time-domain signals
dt	scalar	The amount of time between signal samples
mic_coordinates	numpy.ndarray	An array of microphone coordinates (2D array with dimension N x 2)
radius	scalar	The inner radius of the cavity

thickness	scaler	The thickness of the end plate of the cavity
thetas	list	Parameterized sound velocity parameters
grid_size	int, optional	The number of vertical and horizontal test points (default 10)
octant	int, optional	The octant that the grid is to be constrained to (default −1, i.e. no constraint)
Returns	list	A two-element list containing the x and y coordinates

E.2.2 Code Listing.

```python
def accumulated_correlation(signals, dt, mic_coordinates,
                            radius, thickness, v_s, v_p,
                            grid_size=10, octant=-1):
    # - Calculate the lag matrix (skip auto correlations since
    #   ↪ they aren't used)
    lag_matrix = np.zeros((len(signals), len(signals), len(signals
        ↪ [0])*2-1))
    for i, signal_i in enumerate(signals):
        for j, signal_j in enumerate(signals[i+1:]):
            lag_matrix[i, j+i+1] = sig.correlate(signal_i,
                ↪ signal_j)
            lag_matrix[j+i+1, i] = lag_matrix[i, j+i+1]

    quadrant = -1
    if octant >=0:
```

```python
        quadrant = int(octant / 2)
# - Create a zero matrix the size of the test point grid (sum
    ↪ matrix)
sums = np.zeros((grid_size, grid_size))
if quadrant >= 0:
    if (quadrant == 0) or (quadrant == 3):
        xs = np.linspace(0, radius, num=grid_size)
    else:
        xs = np.linspace(0, -radius, num=grid_size)
    if (quadrant == 0) or (quadrant == 1):
        ys = np.linspace(0, radius, num=grid_size)
    else:
        ys = np.linspace(0, -radius, num=grid_size)
else:
    xs = np.linspace(-radius, radius, num=grid_size)
    ys = np.linspace(-radius, radius, num=grid_size)


n0 = len(signals[0])
ijs = []
for i, signal_i in enumerate(signals):
    for j, signal_j in enumerate(signals[i+1:]):
        # Note: j -> j+i+1 because of the loop optimization
        ijs.append([i, j+i+1])
        ijs = np.array(ijs)


if np.any(octant == np.array([0,1,4,5])):
    # octants 0,1,4,5
```

```python
        constraint_slope = float(mic_coordinates[0,0]) /
            ↪ mic_coordinates[0,1]
    else:
        # octants 2,3,6,7
        constraint_slope = float(mic_coordinates[1,0]) /
            ↪ mic_coordinates[1,1]
# - For each test point...
for a,x in enumerate(xs):
    if (quadrant >= 0):
        max_y = math.sqrt(radius**2 - x**2)
        dy = radius / (grid_size-1)
        max_b = int(round(max_y / dy))
    else:
        min_b = 0
        max_b = len(ys)
    for b,y in enumerate(ys[:max_b]):
    # - For each pair of microphones...
        for index,ij in enumerate(ijs):
            contrib_index = -1

            # Only contribute to the sum if the test point is
            # within the correct octant
            if (x**2 + y**2) <= (radius**2) and\
                ((octant == 0 and y <= constraint_slope*x and x
                    ↪ >= y/constraint_slope) or\
                (octant == 1 and y >= constraint_slope*x and x
                    ↪ <= y/constraint_slope) or\
```

```
        (octant == 2 and y >= constraint_slope*x and x
           ↪  >= y/constraint_slope) or\
        (octant == 3 and y <= constraint_slope*x and x
           ↪  <= y/constraint_slope) or\
        (octant == 4 and y >= constraint_slope*x and x
           ↪  <= y/constraint_slope) or\
        (octant == 5 and y <= constraint_slope*x and x
           ↪  >= y/constraint_slope) or\
        (octant == 6 and y <= constraint_slope*x and x
           ↪  <= y/constraint_slope) or\
        (octant == 7 and y >= constraint_slope*x and x
           ↪  >= y/constraint_slope) or\
        (octant < 0)):
            contrib_index = ac_contrib_index(
                mic_coordinates, [x, y],
                thickness, ij[0], ij[1],
                n0, thetas, dt)
        if contrib_index >= 0 and contrib_index <
           ↪  lag_matrix.shape[2]:
            sums[a,b] += lag_matrix[ij[0], ij[1],
               ↪  contrib_index]
# - Use the max sum matrix element to calculate the most
   ↪  likely source point
max_indices = np.unravel_index([np.argmax(sums)], np.shape(
   ↪  sums))
coordinates = [xs[max_indices[0][0]], ys[max_indices[1][0]]]
if coordinates[0]**2 + coordinates[1]**2 > radius**2:
    coordinates = [0.0,0.0]
```

```
    return coordinates
```

E.3 AC Contribution Index

E.3.1 Description. Calculate the index into the lag matrix of the element that will be added to the current test point's sum in the Accumulated Correlation algorithm.

Table E.3. AC Contribution Index Function Parameters

Parameter	Type	Description
mic_coordinates	numpy.ndarray	An array of microphone coordinates
test_coordinates	scalar	An array of test point coordinates
thickness	numpy.ndarray	End plate thickness in cm
i	scalar	Microphone pair first index
j	scaler	Microphone pair second index
n0	scalar	Number of samples in a signal
thetas	list	Parameterized sound velocity parameters
dt	int, optional	Signal sample width in seconds
Returns	list	A two-element list containing lag matrix indices

E.3.2 Code Listing.

```
def ac_contrib_index(mic_coordinates, test_coordinates,
                 thickness, i, j, n0, v1, v2, dt,
                     ↪ settling_time):
    # − Calculate the expected difference in TOA (TDOA)
    xi = mic_coordinates[i,0]
```

```python
        yi = mic_coordinates[i,1]
        dxi = xi-test_coordinates[0]
        dyi = yi-test_coordinates[1]
        di = math.sqrt(dxi*dxi+dyi*dyi+thickness*thickness)
        for order,theta in enumerate(thetas):
            v_i += theta*di**order
        ti = di / v_i

        xj = mic_coordinates[j,0]
        yj = mic_coordinates[j,1]
        dxj = xj-test_coordinates[0]
        dyj = yj-test_coordinates[1]
        dj = math.sqrt(dxj*dxj+dyj*dyj+thickness*thickness)
        for order,theta in enumerate(thetas):
            v_j += theta*dj**order
        tj = dj / v_j

        # Covert the TDOA into a lag matrix index
        tij = tj - ti
        return int(round(n0 - tij / dt - 1))
```

APPENDIX F

LAUNCHPAD CODE AND DATA REPOSITORY

All of the code and data needed to reproduce the results from previous sections have been uploaded to a public Launchpad [68] repository. The top-level URL to this repository is `http://bazaar.launchpad.net/~plane1/+junk/dissertation/files`. In addition, all of the plots can be reproduced using the `plots.py` Python script in the `src` directory. This file was derived from a copy of the actual Jupyter Notebook file used to generate the plots (`plots.ipynb`).

The `src` directory contains all of the relevant code developed for this work as well as exports of the COMSOL simulation files as Java code. These Java files are in the `COMSOL` subdirectory. Python scripts used for making plots and other analysis are located in the `Analysis` subdirectory. The `data` directory contains all of the data used to generate the plots as well as the field map data files generated by Poisson Superfish that were used to create RF hammer force curves for some of the COMSOL simulations. In `bin` one can find two scripts that were used under the COMSOL 4.4 environment to compile the COMSOL Java files and run them. Finally, the `etc` directory contains configuration files used by `DataExtractor3D.java`.

Under `src` are a number of subdirectories. The `LabVIEW` directory contains all of the National Instruments LabVIEW visual interface files that were used during data acquisition and MTA shifts. This is where the Breakdown Acoustics Controls, ACNET Monitor, and Controls Log Monitor can be found. `Library` contains a number of Python modules commonly used by the plotting and other analysis scripts. Code from App. E was derived from these modules. The `Superfish` directory contains all of the source files for Poisson Superfish simulations of the cavities that are used to generate RF force curves for some the COMSOL simulations.

The subdirectories in `Analysis` are named with the RF cavity acronym to which they are associated, namely `AMC` (Aluminum Mock Cavity), `HC` (High-Pressure Cavity), and `MC` (Modular Cavity).

The `data` directory also has subdirectories for each of the cavities used in this research. These data are a mixture of those obtained through the RF cavity acoustics DAQ (.npz files), an oscilloscope (.CSV files), and simulations (.csv and .npy files).

BIBLIOGRAPHY

[1] C. Ankenbrandt et al. Status of Muon Collider Research and Development and Future Plans. *Phys. Rev. ST Accel. Beams*, 2:081001, August 1999. `doi:10.1103/PhysRevSTAB.2.081001`.

[2] D. Neuffer. Multi-TeV Muon Colliders. In *Proceedings of Symposium on Advanced Accelerator Concepts*, volume 156, pages 201–208, 1987.

[3] S. Geer. Muon Colliders and Neutrino Factories. *Annual Review of Nuclear and Particle Science*, 59:347–365, 2009.

[4] Muon Accelerator Program, Accessed March 27, 2016. URL: `http://map.fnal.gov/`.

[5] The Shutdown Process, Accessed June 23, 2016. URL: `http://www.fnal.gov/pub/tevatron/shutdown-process.html`.

[6] CERN. LHC First Physics : First collisions in the LHC's experiments, afternoon of March the 30th, 2010, Accessed June 23, 2016. URL: `https://cds.cern.ch/record/1256279?ln=en`.

[7] V Barger, D Marfatia, and K Whisnant. *The Physics of Neutrinos*. Princeton University Press, 2012, Accessed July 13, 2016. URL: `https://books.google.com/books?id=qso8NEr3XY8C`.

[8] The T2K Collaboration. Neutrino Sources, Accessed January 19, 2016. URL: `http://t2k-experiment.org/neutrinos/sources-and-experiments/`.

[9] F. Vissani. Neutrino Sources and Properties, December 29 2014, Accessed January 19, 2016. URL: `http://arxiv.org/abs/1412.8386`.

[10] K. Olive et al. (Particle Data Group). Review of Particle Physics. *Chin. Phys. C*, 38(9):30–34, 2014.

[11] S. Kopp. Accelerator Neutrino Beams. *Physics Reports*, 439(3):101 – 159, 2007, Accessed July 13, 2016. URL: `http://www.sciencedirect.com/science/article/pii/S0370157306004431`, `doi:http://dx.doi.org/10.1016/j.physrep.2006.11.004`.

[12] R. Smith. Measurement of kaon and pion contributions to the t2k neutrino beam., 2013, Accessed July 13, 2016. URL: `http://www.t2k.org/docs/thesis/056/rjsthesis`.

[13] S. Geer. Neutrino Beams from Muon Storage Rings: Characteristics and Physics Potential. *Phys. Rev. D*, 57:6989–6997, June 1998. `doi:10.1103/PhysRevD.57.6989`.

[14] J-P. Delahaye et al. A Staged Muon Accelerator Facility for Neutrino and Collider Physics. In *Proceedings of 5th International Particle Accelerator Conference (IPAC 2014)*, page WEZA02, 2014. `arXiv:1502.01647`.

[15] R. Palmer. Muon Collider: Introduction and Status, 1998. Accessed April 20, 2016. URL: `http://www.hep.princeton.edu/mumu/palmer/palmer_9802005.pdf`.

[16] D. Möhl, G. Petrucci, L. Thorndahl, and S. van der Meer. Physics and Technique of Stochastic Cooling. *Physics Reports*, 58(2):73 – 102, 1980. `doi:10.1016/0370-1573(80)90140-4`.

[17] D. Neuffer. Introduction to muon cooling. In *Proceedings of COOL'03*, pages 26–31, 2004. `doi:10.1016/j.nima.2004.06.051`.

[18] S. Posen, N. Valles, and M. Liepe. Radio Frequency Magnetic Field Limits of Nb and Nb$_3$Sn. *Phys. Rev. Lett.*, 115:047001, July 2015. `doi:10.1103/PhysRevLett.115.047001`.

[19] D. Stratakis, R. C. Fernow, J. Berg, and R. Palmer. Tapered Channel for Six-Dimensional Muon Cooling Towards Micron-Scale Emittances. *Phys. Rev. ST Accel. Beams*, 16:091001, September 2013. `doi:10.1103/PhysRevSTAB.16.091001`.

[20] R. Johnson. Introduction: Low Emittance Muon Collider, New 6D Cooling Ideas, and MANX, August 8 2007.

[21] Muon Ionization Cooling Experiment. Repository of images and public information, Accessed April 14, 2016. URL: `http://micewww.pp.rl.ac.uk/projects/mice/wiki/_Repository_of_Images`.

[22] The MICE Collaboration. An International Muon Ionization Cooling Experiment (MICE). Technical report, January 10 2003, Accessed January 19, 2016. Proposal to the Rutherford Appleton Laboratory. URL: `http://mice.iit.edu/mnp/MICE0021.pdf`.

[23] Z. Insepov, J. Norem, T. Proslier, D. Huang, S. Mahalingam, and S. Veitzer. Modeling RF Breakdown Arcs, March 8 2010, Accessed January 14 2016. URL: `http://arxiv.org/abs/1003.1736`.

[24] R. Fowler and L. Nordheim. Electron emission in intense electric fields. *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, 119(781):173–181, March 1928, Accessed February 11, 2016. URL: `http://rspa.royalsocietypublishing.org/content/119/781/173.full.pdf`, `doi:10.1098/rspa.1928.0091`.

[25] J. Norem, V. Wu, A. Moretti, Z. Popovic, Z. Qian, L. Ducas, and N. Solomey. Dark Current, Breakdown, and Magnetic Field Effects in a Multicell, 805 MHz Cavity. *Phys. Rev. ST Accel. Beams*, 6(7):072001, July 30 2003.

[26] R. Palmer, R. Fernow, J. C. Gallardo, D. Stratakis, and D. Li. RF Breakdown with External Magnetic Fields in 201 and 805 MHz Cavities. *Phys. Rev. ST Accel. Beams*, 12:031002, September 8 2008.

[27] A. Hassanein et al. Effects of Surface Damage on RF Cavity Operation. *Phys. Rev. ST Accel. Beams*, 9:062001, June 2006. `doi:10.1103/PhysRevSTAB.9.062001`.

[28] M. Jana et al. Investigation of Breakdown Induced Surface Damage on 805 MHz Pillbox Cavity Interior Surfaces. In *Proceedings of PAC2013*, pages 1007–1009, 2013, Accessed March 17, 2016. WEPMA12. URL: `http://accelconf.web.cern.ch/accelconf/pac2013/papers/wepma12.pdf`.

[29] Y. Torun. MuCool Test Area Photo Gallery, Accessed April 14, 2016. URL: `http://mice.iit.edu/mta/jpg/photos.html`.

[30] Y. Torun et al. The MuCool Test Area and RF Program. In *Proceedings of IPAC'10*, pages 3780–3782, 2010.

[31] R. Johnson et al. HPRF Cavities for Muon Cooling. In *Proceedings of LINAC 2004*, 2004.

[32] D. Bowring, M. Palmer, K. Yonehara, and Y. Torun. RF Design and Operation of a Modular Cavity for Muon Ionization Cooling R&D. In *Proceedings of IPAC2014*, pages 42–44, 2014.

[33] D. Griffiths. *Introduction to Electrodynamics*. Pearson Education, Inc., fourth edition, 2013.

[34] Q. Wu, S. Belomestnykh, and W. Xu. Cavities, April 21 2011, Accessed March 30, 2016. URL: `http://www.c-ad.bnl.gov/ardd/docs/rf4me_v2.pdf`.

[35] D. Lide, editor. *CRC Handbook of Chemistry and Physics*. CRC Press, 84th edition, 2003.

[36] D Roylance. *Mechanical Properties of Materialis*. 2008, Accessed July 13, 2016. URL: `http://web.mit.edu/course/3/3.225/book.pdf`.

[37] V. Zacek, S. Archambault, E. Behnke, J. Behnke, M. Das, A. Davour, and T. Xie. Dark Matter Search with PICASSO. *Journal of Physics: Conference Series*, 375(1):012023, September 7 2012.

[38] R. Rimmer, N. Hartman, A. Ladran, D. Li, and A. Moretti. Closed-Cell 201.25 MHz RF Structures for a Muon Ionization Cooling Experiment. In *Proceedings of EPAC 2002*, pages 2190–2192, 2002.

[39] E. Behnke. Photos and schematics, 2014.

[40] National Aeronautics and Space Administration. Outgassing Data for Selecting Spacecraft Materials Online, Accessed March 21, 2016. URL: `https://outgassing.nasa.gov/`.

[41] Extech Instruments. Model EX330 User Manual, April 2015, Accessed July 13, 2016. URL: `http://www.manualslib.com/manual/577521/Extech-Instruments-Ex330.html`.

[42] Belden Inc. 88641 Technical Data Sheet, Accessed January 26, 2016. URL: `http://www.belden.com/techdatas/english/88641.pdf`.

[43] APC International, Ltd. APC Piezo Calc, Accessed January 26, 2016. URL: `https://www.americanpiezo.com/knowledge-center/apc-piezo-calc.html`.

[44] PI Ceramic GmbH. Fundamentals of Piezo Technology, Accessed July 11, 2016. URL: `http://www.piceramic.com/piezo-technology/fundamentals.html`.

[45] J. Patrick. The Fermilab Accelerator Control System. In *Proceedings of ICAP 2006*, pages 246–249, 2006.

[46] J. Brauer, editor. *What Every Engineer Should Know About Finite Element Analysis*. Marcel Dekker, Inc., second edition, 1993.

[47] Los Alamos Accelerator Code Group. Poisson Superfish, Accessed January 14, 2016. URL: `http://laacg.lanl.gov/laacg/services/download_sf.phtml`.

[48] COMSOL Inc. COMSOL Multiphysics, Accessed January 14, 2016. URL: `http://www.comsol.com/`.

[49] G. Gollin. Investigation of Acoustic Localization of RF Cavity Breakdown, 2003. Accessed March 31, 2016. URL: `http://www.hep.uiuc.edu/home/g-gollin/talks/cornell_acoustic_sensors_talk.pdf`.

[50] F. Le Pimpec et al. An acoustic sensor system for localizing rf breakdown in warm copper accelerationg structures. 2008. SLAC-PUB-10883.

[51] J. Klavins et al. Breakdown Localization Studies on the SwissFEL C-Band Test Structures. *Proceedings of IPAC 2013*, 2013. WEPFI058.

[52] B. Abbott et al. Observation of Gravitational Waves from a Binary Black Hole Merger. *Phys. Rev. Lett.*, 116:061102, Feb 2016. `doi:10.1103/PhysRevLett.116.061102`.

[53] T. Sidery et al. Reconstructing the sky location of gravitational-wave detected compact binary systems: Methodology for testing and comparison. *Phys. Rev. D*, 89:084060, 2014.

[54] G. Yan and J. Tang. A Bayesian approach for acoustic emission source location in plate-like structure. In *In proceedings of SINCE2013*, 2013.

[55] F. Ciampa and M. Meo. Acoustic emission source localization and velocity determination of the fundamental mode a_0 using wavelet analysis and a newton-based optimization technique, Accessed April 4 2016.

[56] J. Valin, F. Michaud, J. Rouat, and D. Létourneau. Robust Sound Source Localization Using a Microphone Array on a Mobile Robot. In *Proceedings of the 2003 IEEE/RSJ International Conference on Intelligent Robots and Systems*, volume 2, pages 1228–1233, 2004.

[57] T. Callaghan et al. Correlation-based radio localization in an indoor environment. *EURASIP Journal on Wireless Communications and Networking*, 2011(1):1–15, 2011, Accessed July 13, 2016. URL: `http://dx.doi.org/10.1186/1687-1499-2011-135`, `doi:10.1186/1687-1499-2011-135`.

[58] R. Ramachandran, H. Patel, G. Raman, and R. Dougherty. Localization of Wind Turbine Noise Sources Using a Compact Microphone Array with Advanced Beamforming Algorithms. *In Proceedings of BeBeC-2012-25*, February 2012.

[59] S. T. Birchfield. A Unifying Framework for Acoustic Localization. In *Proceedings of the 12th European Signal Processing Conference*, pages 1127–1130, September 2004.

[60] S. Birchfield and D. Gillmor. Acoustic Localization by Accumulated Correlation, Accessed April 3, 2016. URL: `http://www.ces.clemson.edu/~stb/research/acousticloc/`.

[61] K. Graff. *Wave Motion in Elastic Solids*. Dover Publications, 1991.

[62] D. Royer, D. Clorennec, and C. Prada. Lamb Mode Spectra Versus the Poisson Ratio in a Free Isotropic Elastic Plate. *J. Acoust. Soc. Am.*, 125:3683, 2009.

[63] COMSOL. *Structural Mechanics Module User's Guide*, 2012.

[64] National Instruments. FFT Fundamentals (Sound and Vibration Measurement Suite), Accessed June 22, 2016. URL: `http://zone.ni.com/reference/en-XX/help/372416B-01/svtconcepts/fft_funda/`.

[65] K. Pearson. On Lines and Planes of Closest Fit to Systems of Points in Space. *Philosophical Magazine*, 2(11):559–572, 1902, Accessed July 10, 2016. URL: `http://stat.smmu.edu.cn/history/pearson1901.pdf`, `doi:doi:10.1080/14786440109462720`.

[66] A. Kochemirovskiy, 2016. Unpublished Modular Cavity Inspection Results.

[67] S. Hansen. Schematics, 2013-2015.

[68] Canonical Ltd. Launchpad, Accessed January 14, 2016. URL: `http://launchpad.net/`.