The Pulsed High Density Experiment (PHDX): Design, Construction, Results

John Slough and Samuel P Andreason

University of Washington

2016

University of Washington

Abstract

The Pulsed High Density Experiment (PHDX):
Design, Construction, Results

John Slough
Samuel P Andreason

The purpose of this paper is to present the conclusions that can be drawn from the Field Reversed Configuration (FRC) formation experiments conducted on the Pulsed High Density experiment (PHD) at the University of Washington. The experiment is ongoing. The experimental goal for this first stage of PHD was to generate a stable, high flux (>10 mWb), high energy (>10 KJ) target FRC. Such results would be adequate as a starting point for several later experiments.

This work focuses on experimental implementation and the results of the first four month run. Difficulties were encountered due to the initial on-axis plasma ionization source. Flux trapping with this ionization source acting alone was insufficient to accomplish experimental objectives. Additional ionization methods were utilized to overcome this difficulty. A more ideal plasma source layout is suggested and will be explored during a forthcoming work.

TABLE OF CONTENTS

Page

List of Figures v
List of Tables xi
1 Introduction 1
2 Field Reversed Configuration 3
2.1 FRC Formation 4
2.2 Ionization 6
2.3 Instabilities/Transport 7
2.4 Translation 8
2.5 Compression 8
3 Experimental Layout 11
3.1 High Power Systems 12
3.2 Control Linkages 13
3.3 EMI 14
3.4 Magnetic Fields 15
3.5 Known Susceptibilities 16
3.6 Vacuum Boundary 17
3.7 Vacuum Control 19
3.8 Vacuum Pumps 21
3.9 Safety Considerations 24
3.9.1 Electrical Isolation 24
3.9.2 Redundancy 25
3.9.3 Electrocution by Coax 26
3.9.4 Capacitor Explosion 27
3.9.5 Capacitor ‘Opens’ 27
3.9.6 Capacitor Dielectric Absorption 28
3.9.7 Collapse of Axial Load-bearing Columns 28
3.9.8 Quartz Chamber Implosion 30
3.9.9 Feed-thru Rifle 31
3.10 Labview 7.0 31
3.10.1 Labview Functions 32
3.10.2 Run Modes 32
3.10.3 Memory Usage 33
3.11 Control System Operation 33
3.11.1 Single Device 33
3.11.2 Multiple Devices 35
3.11.3 For Hardware Developers 40
3.12 Matlab R14SP2 40
4 Power Systems 55
4.1 Modularized Systems 55

4.2 High Voltage Practices 57
4.3 High Voltage Design Considerations 60
4.4 Reversal Bank Units 65
4.5 HV Power Supplies 72
4.5.1 Requirements 73
4.5.2 Power Supply Basics 74
4.5.3 Converter Topologies 75
4.6 Ignitrons 79
4.7 Copper Vapor Switches 83
4.8 Capacitors 88
4.9 Voltage Monitoring 88
4.10 Bias Modules 94
4.11 Cone Magnets 95
4.12 Z-PI 98
4.13 RMF 100
4.13.1 Saturable Core RMF 101
4.13.2 Frequency 102
4.13.3 80 cm PHD RMF 103
4.14 Plasma Guns 105
4.14.1 MPD Background 109
4.14.2 MPD on Experiment 110
4.15 Fast Gas Valve (Puff) 115
4.16 Control Systems 119
4.16.1 Bank Voltage 120
4.16.2 Interlock System 122
4.17 Instrumentation Rack 124
4.18 System State 126
5 Diagnostics 129
5.1 Magnetic Loops 129
5.1.1 Driven Coil Voltage 129
5.1.2 Unshielded Flux and B Loops (PHD) 131
5.1.3 Self-Resonant Frequency 133
5.1.4 Passive/Numerical Integration 134
5.1.5 Capacitive Coupling of Driven Coil to Flux Loop 137
5.1.6 Capacitive Coupling of Driven Coil to B Loop 140
5.1.7 Shielding Effectiveness of Coax 141
5.1.8 High Voltage Resistors 143
5.1.9 Capacitors 148
5.1.10 System Model 152
5.1.11 Vacuum Reference Shots 155
5.1.12 Spatial Effects 161
5.1.13 Spatial Effects (Solved) 180

5.1.14 Standard vs. Magnetically Fitted 183
5.1.15 Summary and Commentary 189
5.2 Interferometer 190
5.2.1 Design Criteria 190
5.2.2 Acoustic Results 196
5.2.3 Electronics Results 198
5.2.4 Options 199
5.3 Visible Light Array 200
5.4 Tomography 203
5.5 Gated CCD Spectrometer (HR4000) 208
5.6 16 Channel Time Resolved Spectrometer 210
5.7 Monochrometer 213
5.8 Single Frame Color CCD Camera 213
5.9 Fast Framing Camera 216
5.10 Residual Gas Analyzer 218
5.11 Capacitance Manometer 218
5.12 Miscellaneous 220
6 Analysis 222
6.1 Lift Off Trapped Flux 222
6.2 Axial Equilibrium 224
6.2.1 80 cm Experiment 228
6.2.2 40 cm Experiment 229
6.3 Poloidal Flux Lifetime τ_φ 230
6.4 Energy 232
6.5 Energy Lifetime τ_E 233
6.6 Plasma Inventory Lifetime τ_N 236
6.7 Kinetic Parameter s 238
6.8 Collision Mean Free Path 239
6.9 Neutral Density 239
7 Results 243
7.1 Ionization 243
7.1.1 MPD Location 245
7.1.2 Basic Magnetic Timings 246
7.2 Parameter Optimization 247
7.2.1 East MPD 250
7.2.2 RMF Voltage 253
7.2.3 Secondary Gas Fill 254
7.2.4 Discussion 255
7.3 Run Parameters 256
7.3.1 LHDI Scaling 256
7.3.2 Flux Trapping 257
7.3.3 Lifetimes 260

7.4 Select Shots 266
7.5 Finale 272
End Notes 274
Bibliography 281

LIST OF FIGURES

Figure Number		Page
2.1	FRC Cross-section	3
2.1.1	RMF FRC Formation	6
3.0.1	Experimental Layout	12
3.4.1	External Magnetic Fields	16
3.6.1	Permeability of Insulators	18
3.7.1	Vacuum Control Enclosure	20
3.8.1	Single Phase Power vs. Two Phase Power	23
3.9.7.1	Conceptual Diagram of Vacuum Loading on End Cones	28
3.9.7.2	Equations for Determining Failure of Support Columns	29
3.9.7.3	Axial Support Rods Showing Deflection	30
3.11.1.1	2028 Driver Front Panel	34
3.11.1.2	2028 Driver Front Panel with Data	35
3.11.2.1	Driver Interface Control of Multiple Vis	36
3.11.2.2	Jorway 221A Driver Front Panel	37
3.11.2.3	Joerger Model TR Front Panel	38
3.11.2.4	ACQ196 Front Panel	38
3.11.2.5	Delay and Countdown Front Panel	39
3.11.2.6	Delay and Countdown Diagram	39
3.12.1	Programmed Plotting of an Integral	42
3.12.2	Improved PLG Options	43
3.12.3	Improved PLG Plotting	43
3.12.4	Windowed FFT Interface	46
3.12.5	2D Windowed FFT Interface	47
3.12.6	0D Automatically Populated Logbook	47
3.12.7	Colormap Options	48
3.12.8	Available Color Schemes	49
3.12.9	Spectra Analysis Utility	50
3.12.10	Residual Gas Analyzer Fitting Interface	51
3.12.11	Image Processing	52
4.2.1	High Voltage Probes can Fail Spectacularly	58
4.2.2	Corona Tracking Over Insulators	59
4.2.3	Electrical Discharges through Vacuum System and Across Boards	60
4.3.1	Potting Mistakes	62
4.3.2	Corona Destroys Insulators Quickly	64
4.4.1	LSX Reversal Bank Module Design	66
4.4.2	Module Evolution	67
4.4.3	LSX HV Distribution	68
4.4.4	LSX Triggers and Supplies	68

4.4.5 HV Power Supply Evolution 69
4.4.6 Reversal Bank Pair SPICE Model 70
4.4.7 Ringing of Modules Due to Start Switch Characteristics 70
4.4.8 Influence of Snubbers on Ringing of Module Outputs due to Cable Z .. 71
4.4.9 Impact of Start Switch Pair Timing Mismatch 72
4.5.2.1 Generic Power Supply Diagram 74
4.5.3.1 Simplified Flyback Converter Topology 75
4.5.3.2 Practical Flyback Converter Topology 76
4.5.3.3 Flyback Power Supply Traces Demonstrating Feedback 78
4.6.1 EEV BK488D Type D Ignitron 80
4.7.1 Copper Vapor Switch..... 83
4.7.2 Damaged Copper Vapor Switch 84
4.7.3 Jitter of Copper Vapor Switch 85
4.7.4 Jitter Aging of Copper Vapor Switch 86
4.7.5 Flux Loop Traces of Copper Vapor Switch Test Setup..... 88
4.9.1 Old Voltage to Frequency Converters 90
4.9.2 New Voltage to Frequency Converter's Linear Response Curve..... 92
4.9.3 Voltage to Frequency Converter Unit and Diagram..... 93
4.10.1 Bias Modules 94
4.11.1 40 cm Experiment Axial Field Coil Layout..... 95
4.11.2 Cone Magnets 96
4.11.3 Cone Vacuum Boundary Structural Diagram..... 97
4.12.1 Axial Discharge 99
4.12.2 Axial Discharge Asymmetries..... 100
4.13.1 RMF Magnetic Fields Vector Plots 101
4.13.1.1 Saturable Core RMF 102
4.13.2.1 End on Imaging of Axial Discharge Interaction with RMF 103
4.13.3.1 Present RMF Setup 104
4.13.3.2 RMF Frequency and Time Characteristics 105
4.14.1 Ringing Theta Induced Breakdown 107
4.14.2 MCAS on STX-HF 108
4.14.2.1 MPD Source Schematic..... 111
4.14.2.2 Installed MPD Source (+z, West End)..... 111
4.14.2.3 MPD Bank and Operational Image..... 112
4.14.2.4 MPD Diagnostic Traces..... 112
4.14.2.5 MPD Diagnostic Traces for Shot 917..... 113
4.14.2.6 Operational Images of East MPD 114
4.14.2.7 Hollow Copper Center MPD 115
4.14.2.8 Operational SXR Blocking MPD..... 115
4.15.1 Parker Hannifin Corporation (General Valve Division) Puff Valves... 117
4.15.2 Characteristic Plots of Low Voltage Puff Valve Operation..... 118
4.15.3 Inlet Gas Conditions for 80 cm Experiments..... 119

4.16.1.1 Bank Voltage Remote Control Systems 122
4.16.2.1 Interlock Systems 124
4.17.1 Screen Racks 126
5.1.2.1 B and Flux Loop Diagram 132
5.1.4.1 Passive Integrator and Response 134
5.1.4.2 Simple Passive Integrator Correction Code 135
5.1.4.3 Typical Magnetic Signals for Long Time Scales 136
5.1.4.4 Typical Flux and B_{ez} Traces 137
5.1.5.1 High frequency Ground Loop Conceptual Schematic 137
5.1.5.2 Simplified Model of Capacitive Coupling to Flux Loop 138
5.1.5.3 Frequency Response of Capacitive Noise Coupling to Flux Loop 139
5.1.6.1 Frequency Response of Capacitive Noise Coupling to B Loop 141
5.1.7.1 Transfer Impedance 142
5.1.8.1 Conductance of Resistors Exposed to Pulsed High Voltage 143
5.1.8.2 Thermal Model Results for High Voltage Integration Resistors 146
5.1.9.1 Impedance Characteristics of Through-hole Capacitor Series 149
5.1.9.2 Effect of DC Voltages on Capacitance 150
5.1.9.3 Frequency Dependence for Common Film Dielectric Types 151
5.1.9.4 Humidity Dependence for Common Film Dielectric Types 151
5.1.10.1 SPICE Model for Flux Loop 152
5.1.10.2 Frequency Response of Loop Model with Characteristics Removed 153
5.1.10.3 Frequency Response of Model for Magnetic Diagnostics 153
5.1.10.4 Zoomed Plot of Frequency Response of Magnetic Diagnostics 154
5.1.10.5 Phase Lag of Flux Loops Derived from Model 154
5.1.11.1 Measured Phase Difference Between B and Flux Diagnostics 155
5.1.11.2 Measured Relative Phase vs. Modeled Phase Response 156
5.1.11.2 Uncompensated Measured Relative Phase vs. Modeled 157
5.1.11.4 Spectral Energy Density of Mid-plane B Probe 158
5.1.11.5 Measured Flux Loop with SPICE Correction 159
5.1.11.6 Comparison of Vacuum Reference Method vs. Frequency Correction 160
5.1.11.7 Spectral Plot of Excluded Flux for Vacuum and Plasma Shots 161
5.1.12.1 Nonlinear Model for Skin Effect 162
5.1.12.2 Program for Solving Nonlinear Skin Effect Equations 163
5.1.12.3 SPICE Model of Driven Strap 163
5.1.12.4 Comparison of Model vs. Simple Inductor/Resistor 164
5.1.12.5 Normalized Current Distribution in Driven Strap 165
5.1.12.6 Normalized Current Distribution as a Function of Axial Location 166
5.1.12.7 Normalized Current Distribution as a Function of Radial Location 167
5.1.12.8 Flux Contours and B Amplitude @ High Frequency (Single Strap) 168
5.1.12.9 Flux Contours and B Amplitude @ DC (Single Strap) 169
5.1.12.10 Flux Contours and B Amplitude @ High Frequency (Three Straps) 170
5.1.12.11 Timescales Apparent in Numerical Model 171

5.1.12.12 Inductance Change as Field Penetrates Strap 172
5.1.12.13 Effective Resistance Decreases as Flux Penetrates Strap 173
5.1.12.14 Current Distribution in Single Driven Strap becomes more Diffuse 173
5.1.12.15 Ratio of B_z Measured by Diagnostic to B_z at Center 174
5.1.12.16 B_z Diagnostic is Insensitive to Small Axial Position Errors 175
5.1.12.17 B_z Diagnostic is Insensitive to Small Radial Position Errors 176
5.1.12.18 Standard Approach Magnetics for Shot #917 @ 6.1 μs 177
5.1.12.19 Standard Approach Magnetics for Shot #917 @ 8 μs 177
5.1.12.20 Standard Approach Magnetics for Shot #1166 @ 10 μs 178
5.1.12.21 Good Overlay of Corrected B Trace with Flux 180
5.1.13.1 Magnetically Fitted Magnetics for Shot #917 @ 6 μs 181
5.1.14.1 Solving for Excluded Flux and B_z Averaging Options 184
5.1.14.2 Comparison of Methods for Excluded Flux Shot #917 185
5.1.14.3 Select Magnetically Fitted Profiles for Shot #917 186
5.1.14.4 Peak Excluded Flux Comparison of Methods on Dynamic Shot 187
5.1.14.5 Residual Errors for Two Methods 187
5.1.14.6 Select Magnetically Fitted Profiles for Shot #1166 188
5.2.1.1 CAMAC Interferometer Board 192
5.2.1.2 Bragg Cell and Laser Enclosure 193
5.2.1.3 Interferometer Optics Table Diagram 194
5.2.1.4 Pictures of Interferometer Table as Implemented 195
5.2.2.1 Acoustics on Interferometer 196
5.2.2.2 Interferometer Electronics Diagram 197
5.2.2.3 Results from Interferometer Driven with 470 Hz Test Signal 198
5.2.4.1 Graph of Spectral Attenuation for Readily Available Fiber Types 199
5.2.4.2 Interferometer Strength Ratio 200
5.3.1 Thermally Stabilized VLA Board 201
5.3.2 Expected Spectral Response of VLA Detector 202
5.3.3 Spectral Transmission through Wratten Gel Filter & Fiber 203
5.4.1 PMT Tomography System Summary 204
5.4.2 Comparative Calibration Results for Tomography System 205
5.4.3 Integrating Sphere and LED Light Source 205
5.4.4 LED Spectra 206
5.4.5 Comparison of Tomography and End-on Imaging 207
5.5.1 Line Distortion of HR4000 Due to Line Source Saturation 209
5.5.2 End-on Optics Light Collection 210
5.6.1 SPEX Spectrometer Diagram (Czerny-Turner) 211
5.6.2 Side View of 16 Channel Spectrometer (SPEX) Optics Bench 211
5.6.3 SPEX PMT Characteristics from Datasheet 212
5.6.4 SPEX Spectral Response 212
5.6.5 SPEX Relay Optics 213
5.8.1 Canon EOS Digital Rebel 300D Camera Timings 214

5.8.2 Interesting Sampling of SLR Images 215
5.8.3 Not All Pixels are Captured Over Same Time 216
5.9.1 Imacon Frame with Fine Structure at Various Resolutions 217
5.11.1 Pressure Monitoring 219
5.11.2 Pulsed Vacuum Pump Response 220
5.12.1 Damaged 5 x Coaxial Attenuator (2W) 221
6.1.1 Calculation and Comparison of Lift Off Flux Methods 224
6.2.1 Rigid Rotor Dependence on X_s in Equilibrium 226
6.2.2 Density Profile as Determined by Profile Constant K 227
6.2.3 Relationship Between Line Averaged Density and Peak Density 227
6.2.1.1 Axial Driven Coil Flux Profiles 228
6.2.1.2 Axial Position of FRC During Stationary Shot 229
6.3.1 Magnitude and Lifetime of Poloidal Flux 231
6.3.2 Trapped Flux Characteristic Loss Timescale 232
6.4.1 Energy Information Available from Magnetic Diagnostics 233
6.5.1 Energy Loss Timescale Derivation 235
6.5.2 Energy Loss Timescales for 80 cm Experiment Shots 236
6.6.1 Plasma Inventory Timescales for 80 cm Experiment Shots 237
6.6.2 Peak Plasma Inventories Inside Separatrix 237
6.6.3 Total Plasma Temperature 237
6.7.1 Kinetic Parameter s 239
6.9.1 Total Neutral Density Delivered to Vacuum System from MPD 240
6.9.2 Chamber Pressure Spikes After a Plasma Shot 241
7.1.1 Flux Trapping Shows Strong Dependence on Ionization Methods 244
7.1.2 Bias Second Half Cycle Operation 244
7.1.2.1 Plug and Bias Field Timings 247
7.2.1 West MPD Timing Optimization 248
7.2.2 West RMF Timing Optimization 249
7.2.1.1 East MPD Timing Optimization 251
7.2.1.2 East RMF Timing Optimization 252
7.2.1.3 East Axial Discharge Timing Optimization 253
7.2.2.1 East RMF Voltage Optimization 254
7.2.3.1 Gas Back Fill Effects on Trapped Flux and Lifetime 255
7.3.1.1 Energy Lifetime Scaling With Lower Hybrid Drift Form 257
7.3.2.1 Time Spent with FRC Currents on the Wall Show Little Effect 258
7.3.2.2 Flux Trapping vs. FRC Inventory 259
7.3.2.3 Energy Lifetime Dependence on Total Inventory 259
7.3.2.4 Maximum Relaxed Poloidal Flux Dependence on Inventory 260
7.3.3.1 Inventory Lifetime Shows Little Correlation with Other Lifetimes 260
7.3.3.2 Energy Lifetime Dependent on Flux Lifetime 261
7.3.3.3 Energy Lifetime Dependence on Flux Trapping Efficiency 262
7.3.3.4 Energy Lifetime Correlation with Relaxed State Poloidal Flux 262

7.3.3.5 Energy Lifetime Dependence on Relaxed State Energy 262
7.3.3.6 Energy Lifetime Dependence on FRC Volume 263
7.3.3.7 Energy Lifetime and External Magnetic Field in Relaxed State 263
7.3.3.8 Energy Lifetime Dependence on X_s 263
7.3.3.9 Energy Lifetime and Total Temperature 264
7.3.3.10 Energy Lifetime Dependence on the Kinetic Parameter 264
7.3.3.11 Energy Lifetime Dependence on Radial Profile Peaking 264
7.3.3.12 Energy Lifetime and Peak Plasma Density 265
7.3.3.13 Dimensionless Parameter Operating Regime 265
7.3.3.14 Spitzer Resistivity at Separatrix Compared to Hall Term 265
7.4.1 Visible End-on Imaging Negatives 266
7.4.2 Visible Spectrum for Stationary and Dynamic Shots 267
7.4.3 Stationary FRC Separatrix Radii 269
7.4.4 Poorly Timed Dynamic Formation X_s 270
7.4.5 Well-timed Dynamic Formation 271
7.4.6 When Vacuum Shots Don't Match Plasma Shot Behavior 272

LIST OF TABLES

Table Number Page

4.18.1 80 cm Experiment Timeline....128
5.1.12.1 Characterizing Constants for Individual Magnetic Diagnostics 1 179
5.1.13.1 Characterizing Constants for Individual Magnetic Diagnostics 918+ ..182

1
Introduction

The purpose of this paper is to present the conclusions that can be drawn from the Field Reversed Configuration (FRC) formation experiments conducted on the Pulsed High Density experiment (PHD) at the University of Washington. The experiment is ongoing. The experimental goal for this first stage of PHD was to generate a stable, high flux (>10 mWb), high energy (>10 kJ) target FRC. Such results would be adequate as a starting point for several later experiments.

We have chosen to thoroughly present the current iteration on this stage of the experiment, concentrating on work done in the 80 cm diameter opaque fused silica formation chamber. Previous efforts were focused on a 40 cm diameter experiment conducted in a clear quartz tube. The basic result of the earlier experimental phase could be concisely described as a study in the instabilities associated with standard formation techniques. Including the earlier phase herein would have required a significantly altered analysis of the key diagnostics and drive system compared with the 80 cm experiment. It was judged a far more useful exercise to thoroughly describe and analyze the present experimental phase.

This document is organized into seven chapters. The second chapter introduces basic FRC concepts. This section has been kept short since the critical material is better covered by the various included references.

The third chapter concentrates on general experimental layout, factors affecting layout, safety concerns, vacuum system components, data acquisition, and control.

The fourth chapter covers the design of the various power systems connected to the experiment. This includes the modularization design choice, power supply design, and high voltage/power switching. The different magnetic systems are discussed. Also included are descriptions of ionization systems such as the pulsed rotating magnetic field (RMF) system, axial discharge (ZPI) system, and magneto plasma dynamic (MPD) plasma sources.

The fifth chapter concentrates on in-depth design and analysis of the primary magnetic diagnostics. Additional sections are included on the interferometer design, and other secondary diagnostics and commercial equipment used as diagnostics.

The sixth chapter describes the basics for the data analysis. An attempt is made to include the analyses in a form simple enough to easily apply to other experiments. This chapter concentrates on using the processed products of the various diagnostics (often in combination). Emphasis is placed on the assumptions necessary to maintain the problems' tractability.

The seventh chapter presents and discusses the results. Operational optimizations are included in sufficient detail to understand what timings worked best and their likely inter-relations. The primary energy confinement time scaling law from a lower hybrid drift instability saturation analysis is compared against experimental results. This section discusses many of the available figures, but also

includes additional figures with data applicable for understanding stability and other parameters not discussed in detail. Plots describing the available operating regime are included. A brief discussion of how the results will be improved upon is included. Last, this section discusses the gross general behavior of the FRCs generated in this experiment.

2
Field Reversed Configuration

This chapter is meant to introduce the concepts important in understanding FRCs. The first section describes the configuration in the simplest terms; neglecting important details to emphasize the overall concepts. Subsequent sections build on this base, fleshing out the necessary ideas. Given the wide range of phenomena present in any plasma and the requirement for more than magneto-hydro-dynamics (MHD) to understand FRC stability, the physics of FRCs can be very complex to understand in detail.

A field reversed configuration (FRC) refers to a magnetic configuration that is a compact toroid (i.e. simply connected). As a matter of definition, the poloidal magnetic fields (B_r, B_z) are much greater than the azimuthal magnetic fields (B_θ). A standard conception of a FRC is shown in figure 2.0.1[1]. Normally, a cylindrical coordinate system is used (r, θ, z). The magnetic fields in the r and z directions dominate. The closest simple fluid analogy is that of a smoke ring. The paths the smoke particles follow are that of the field lines. A FRC then is multiple smoke rings superimposed and stacked up together along the rotational axis. Usually, an FRC is assumed to be rotationally symmetric about the z axis, reflection symmetric about the mid-plane.

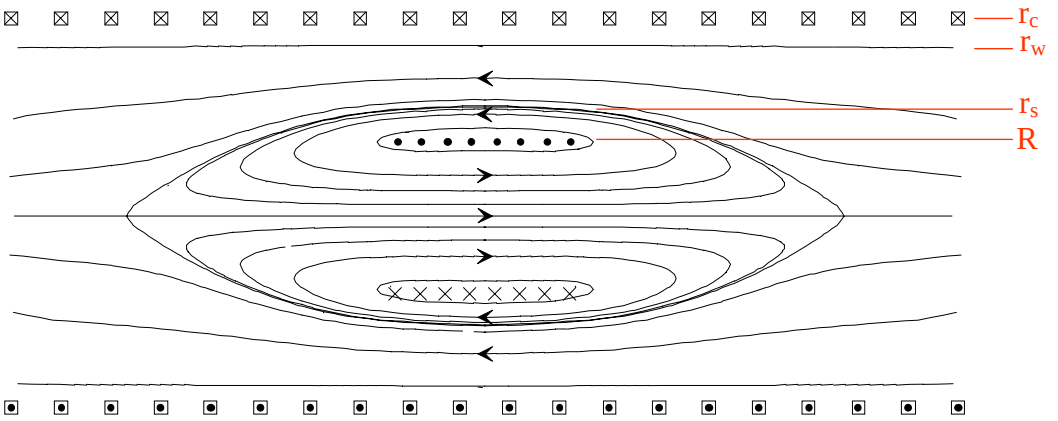

Figure 2.1: Cross-section of a relaxed FRC. The primary defining features are dominance of B_z and B_r magnetic fields over B_θ and translatability afforded by no equipment linking the hole of the torus to solid structures. The radius of the separatrix (r_s) defines the last closed magnetic field line of the plasmoid. The magnetic null (R) is typically at a radius of $r_s/2^{½}$. In experiments, most FRCs are prolate. Many analyses assume elongation $\varepsilon >> 1$ so that the magnetic field curvature at the mid-plane is negligible.

An externally applied axial magnetic field supports the FRC's radial pressure, keeping it off the vacuum boundary. The conductors that support the external currents generating these fields are at a radius r_c outside the dielectric vacuum chamber wall at r_w. The simplest FRC experiments are frequently conducted with a single turn elongated coil doubling as a flux conserver and shaper rather than the

segmented coil suggested in the illustration. The radius of the separatrix (r_s) denotes the radius of the last closed field line. The plasma density is usually assumed low here, for an internally well-confined FRC. An important non-dimensional number $x_s=r_s/r_c$ is used for scaling studies in large part due to the assumption of constant flux at all axial locations inside the external coil associated with the simplest coil constructions.

Peak plasma pressure can be found at the magnetic null ($R = \frac{r_s}{\sqrt{2}}$). When the FRC is sufficiently elongated, the magnetic curvature forces at the mid-plane can be assumed negligible. Then, the pressure from the external magnetic field (if the external coil has B parallel to the wall everywhere) is equal to the plasma pressure at the null $\frac{B_e^{\,2}}{2\mu_0} = n_0 k T_t$. Finally, plasma pressure maps as a function of flux surface $\vec{J} \times \vec{B} \approx \nabla P_e$ (from the generalized Ohm's law Hall terms).

Beta (β), as defined for FRCs, is the local plasma pressure divided by the external magnetic pressure. This definition is used so that average beta <β> is finite despite having infinite plasma beta, as normally defined, at the magnetic null. A simplified relation exists between average beta and separatrix radius for cases where the FRC is fully contained within a constant flux conserver/shaper. There are additional requirements (that are met to varying degrees) that we go into during analysis. The simplified relation is $\langle\beta\rangle = 1 - \frac{1}{2} x_s^2$.

2.1 FRC Formation

There are at least three inherently different formation techniques in general use. By 'inherent', we mean that the dominating physics aspects are different.

The most straightforward technique is the field reversed theta pinch (so-called because all the applied currents run in the theta direction). This technique does not have to require magnetic field reconnection in a region with plasma.

Utilizing the magnetic Hall term, it is possible to form FRCs with the application of magnetic fields transverse to the axially applied bias fields. One way to realize this approach has been by rotating magnetic field (RMF) current drive. Both of these first two techniques require, or rather can take advantage of, dielectric vacuum boundary conditions (no stainless steel vacuum chambers in the formation region).

The last approach takes two spheromaks (compact toroids similar to FRCs distinguished only by the significant presence of toroidal magnetic fields) with toroidal fields of opposite sign and forces them to magnetically reconnect. The toroidal fields are annihilated on some timescale (hopefully faster than the lifetime of the plasmoid) leaving an FRC. The spheromaks can be generated with electrodes and a bias magnetic field.

The programmed formation field reversed theta pinch (FRTP) technique is attractive to us due to the relative simplicity of the physics involved. The necessary

formation steps are easy to summarize and understand. Start with a uniform axial magnetic field (bias field) in the central section with magnetic cusp fields off to either end. Add a uniform plasma distribution to the inside of the formation chamber. Rapidly reverse the applied axial magnetic field in the center section so that the external magnetic field no longer forms a cusp. As the external pressure from the magnetic field initially decreases, the plasma's pressure (and trapped magnetic field) is supported by a sheath that forms on the vacuum chamber boundary. This direct contact between the high current region and the wall can result in significant ablation of the wall material, and is one of the chief drawbacks to this technique. A significant amount of flux can also be lost during this period. It is important to note that the electric field applied to change the magnetic field couples to both electrons and ions, resulting in some very small spin-up of the ions in the direction opposite the electron flow. If external magnetic fields are applied in a way that limits electron mobility, the ions can end up carrying more of the current.

When the external field is sufficiently reversed this current carrying region is pushed off the wall. This process directly results in radial dynamics. The viscosity of the plasma (and interactions with the external electrical circuit) plays a key role in damping radial oscillations. The FRC now needs to come to internal pressure equilibrium. The initially low plasma pressure is insufficient to prevent an axial contraction due to magnetic forces. This results in the axial implosion and bounce.

The second formation technique is based on rotating magnetic fields (RMF) transverse to a steady-state applied bias field. If the frequency of rotation is lower than the ion cyclotron frequency the plasma simply deforms to the new field geometry (creating a rotating bar shape). If the RMF frequency is greater than ω_{ci}, but less than ω_{ce}, then the electrons will dominate the response to the presence of the magnetic field. The RMF would induce an opposing current in the electrons. However, due to the presence of the original axial magnetic field, cross-field conduction is poor, resulting in a large j_z. This j_z, coupled with the radial component of the RMF (B_r) produces an azimuthal electric field that has a time-averaged steady component. When certain conditions are met, conditions on electron/ion relative mobility and collisionality, it is possible for this $j_z B_r$ term to overcome ηj_θ, increasing the flux contained within the FRC. The chief advantages of this approach are minimal interaction with the wall during formation, and the hope of being able to maintain a FRC in steady state.

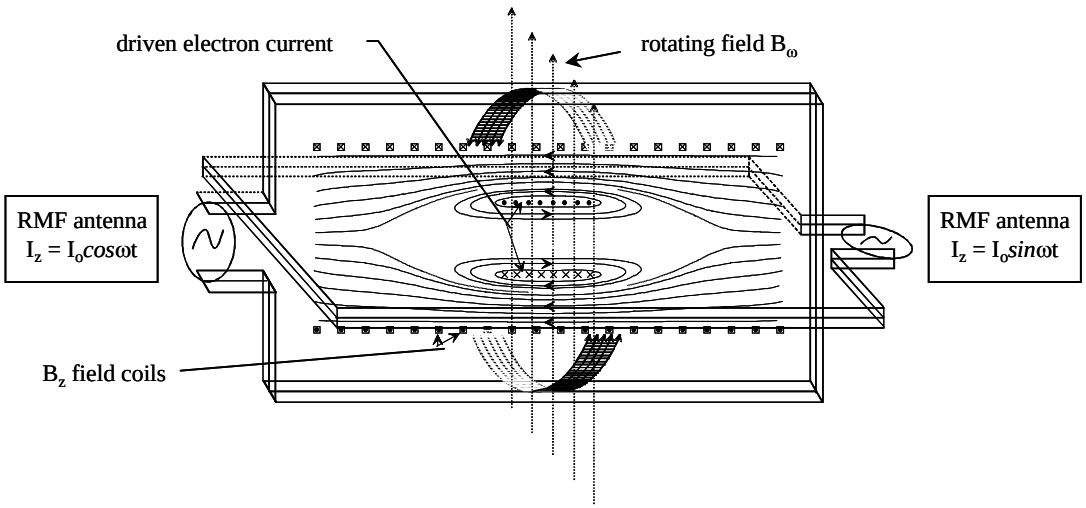

Figure 2.1.1. Basic RMF generated FRC concept. Two phases of RMF antenna are used to generate a rotating magnetic field orthogonal to the initial external bias fields.

Most people prefer to think of this effect in terms of the applied RMF dragging the magnetized electrons as it rotates, resulting in a net current. From this way of thinking it is clear that there is a maximum current that can be driven, corresponding to the electrons rigidly rotating in phase with the RMF ($j_\theta = -en_e(r)\omega r$). It is also clear that this rigid motion can weakly couple to the ions through collisions with the rotating electrons. The resistivity spins up the plasma in the same direction as electron flow until this momentum exchange is balanced by some other force (such as that due to ion loss/ionization or injected beams). While inventory, flux, and electron temperature are sustained, there is reason to believe that initial ion temperature gained during formation is not sustained. Understanding the initial process generating higher ion temperatures may be worthy of further study.

The ability of the RMF to "sweep up" and push in plasma radially, the various complex flows, the essentially 3-dimensional nature of RMF; all these effects provide a fertile ground for working on challenging aspects of plasma physics. The technology for generating the large external RMF field (on the order of the axial field) at the frequencies required (100s kHz) in a continuous mode is expensive. One question to ask is whether the required RMF powers employed must increase as the obtained plasma parameters improve. Past focus has been on impurity control under the theory that impurity radiation is limiting obtainable parameters.

The third formation approach, via merging spheromaks, is an interesting physics study in reconnection. It also differs from the first two approaches in that electrodes are usually employed in creating the spheromaks to be merged.

2.2 Ionization

Surprisingly, ionization is one of the primary sticking points for many FRC experiments. The objective is to break down a neutral gas at low pressures (<5

mTorr) in a uniform way. Some experiments use a backfill followed by an axial electrode discharge or RF antennae glow. It is often possible to pump away this initial discharge to lower pressures, permitting lower fill pressures. This method is hard to extend below 1 mTorr due to the Paschen curve and typical machine length scales involved. It is also infeasible in situations where a static backfill is unacceptable; for example experiments that make use of beams for long term sustainment may not be able to operate with significant neutral pressures outside the separatrix between radial wall and FRC.

Ionization is frequently supplemented by 'ringing' some of the axial field banks. The approaches listed so far are simple to implement. However, they frequently lead to needing additional systems to correct undesirable side effects. An example of this might be the use of high order magnetic barrier fields (B_r, B_θ) to prevent damage to the wall. They can be created by axial current feeds along the wall with alternating current directions going around the tube. They are often compared to a picket fence.

The obvious way to extend to lower pressure operation with similar technologies is to run an electrode discharge at a locally high pressure, with the discharge occurring through the gas as it enters the system. Since the discharge does not extend through the bulk of the chamber the plasma is free to cool, recombine, and interact with the walls (hopefully at some slow rate). A concern with this method is the potential magnetization of the plasma generated in this fashion (i.e. will it cross field lines). We want either a uniform plasma distribution or a hollow cylindrical shape for FRC formation. We generally need this initial plasma on field lines that are initially closed off from the axial boundaries.

For PHD to work well, we need on order 10^{20} ions (~2.5 kJ ionization energy) in a timescale that fills the formation chamber. The minimum power required in the ionization system is then simply equal to $P = \frac{Energy}{Time} = \frac{Energy * Velocity}{Length}$.

2.3 Instabilities/Transport

From a MHD perspective there are many unstable modes in FRCs. A short list would include the internal tilt, fluting (on multiple length scales), rotation, and the "inner-tube" instabilities. The first version of the PHD source section includes operating modes that demonstrate many of these instabilities, including a configuration where a Kelvin-Helmholtz instability was believed to be affecting the ends of an axially contracting FRC[2]. Obviously, minimization of the starting point from which the modes continue to grow can help only to a limited extent.

Stability for most of these modes can be improved by the gyro-kinetic effects when hot ions are present. This effect can be loosely characterized by the s number of the configuration. Crudely, s is the average number of gyro-radii from the magnetic null to the open field lines. The important implication is that if the ions are cold, the plasma is MHD unstable and has short lifetime. If the ions are hot, the number of gyro-radii across the device is small, the plasma is stable and has short

lifetime (by particle transport). Presumably there is a middle ground where the plasma is stable with reasonable lifetimes. It is this middle ground that the PHD concept is attempting to fully exploit.

Operations in this middle regime were demonstrated in previous experiments. The large s experiment (LSX) was conducted Bellevue, WA during the early 1990s[3]. One of the primary results was the demonstration of stability for s > 4. Previous experiments include the triggered reconnection experiments (TRX-1[4] TRX-2[5]) in Washington State and the field reversed experiment (FRX) series of experiments at Los Alamos[6].

Additionally, there is some speculation that the instabilities may be stabilized through sheared plasma flow[7]. Translation has the potential to introduce flow through drag and acceleration.

2.4 Translation

The simple unlinked geometry of the FRC results in a freedom of movement along its axis. There are many benefits to translation in a pulsed power system. The geometry and translation means that the high power equipment can be spread out over a larger footprint without the difficulty of power transmission lines converging in three dimensions (as opposed to two with translation). Since much of the deposited energy is unidirectional kinetic, the particle loss rate should be unaffected (except from the effects of acceleration), unlike energy added to thermal or magnetic components. This design concept may also prove more tolerant to maintenance during operation in the final embodiment. It may, for example, prove possible to deactivate one accelerator coil at a time for switch replacement.

There are also side benefits to acceleration not related to the overall system design. Some experiments have been shown incapable of generating an FRC in-situ but are still able to generate translated FRCs; this was true during at least some phases of the translation, confinement, and sustainment experiment (TCS)[8].

The final compression and heating of the FRC to fusion temperatures may be accomplished by converting the kinetic energy. Azimuthal electric fields can be generated in a translating FRC far in excess of those generated by a capacitor bank. This experiment is different from the magnetized target fusion (MTF) approach in the nature of the compression. The pulsed high density approach, and the problems that must be addressed to realize it, may have a wide range of possible applications beyond a simple power plant.

2.5 Compression

It is important to distinguish between the various ways to compress an FRC. For gas dynamic processes, fast processes generally are adiabatic and slow processes are non-adiabatic. The timescale which is important for comparisons in gas systems is the time required for the system to come into thermal equilibrium with the surroundings. For plasma, if it comes into equilibrium with its surroundings there is

no plasma. Ergo, slow is redefined to be slow on a timescale compared to the various frequencies in the plasma (most particularly the ion gyrofrequency). The slow plasma processes are adiabatic. Fast non-adiabatic magnetic compressions may be possible[9] and under some conditions should not be ignored. Also, for the slow case the expected scaling is different for wall compressions (r_w changing with constant flux) versus flux compressions (r_s changing with constant r_w).

The scaling presented is from a one dimensional adiabatic FRC compression scaling by R.L. Spencer, M. Tuszewski, and R.K. Linford[10]. For wall compressions with $\gamma=5/3$: length ~ $r_w^{2/5}$, temperature ~ $r_w^{-6/5}$, density ~ $r_w^{-12/5}$, magnetic field ~ r_w^{-2}. However, the analysis assumes that the FRC is in equilibrium for every part of the compression. Additionally, the FRC must be highly elongated. The explanation given in the Spencer et al paper for different results versus a numerical scaling analysis was "insufficient elongation for small values of x_s". A flux compression scaling requires knowledge of the internal magnetic field profile. We will repeat the diffuse case scaling here; Spencer et al. lists the scaling: length ~ $x_s^{19/10}\langle\beta\rangle^{-7/10}$, temperature ~ $x_s^{-13/5}\langle\beta\rangle^{-3/10}$, density ~ $x_s^{-39/10}\langle\beta\rangle^{-3/10}$, magnetic field ~ $x_s^{-13/4}$.

Additional considerations must be applied when considering experimental data. In particular, magnetic mirrors or anti-mirrors in the end regions can significantly impact the axial equilibrium that is assumed to take place. If the FRC would be longer in equilibrium than the machine mirrors permit then no length scaling might be observed for some range of compression.

A more in depth and complete view of FRC research up until 1988 can be found in Michel Tuszewski's review paper[11].

3
Experimental Layout

The experimental layout is the solution to balancing multiple design constraints. These constraints include utilizing pre-existing equipment as well as minimizing materials and design costs. The experiment must also be made to fit within the laboratory space without compromising safety clearances. Additionally, allowances for ease of maintenance and inspection must be made.

Mechanical supports, vacuum systems, and the requirements for electrical power and safety introduce their own constraints.

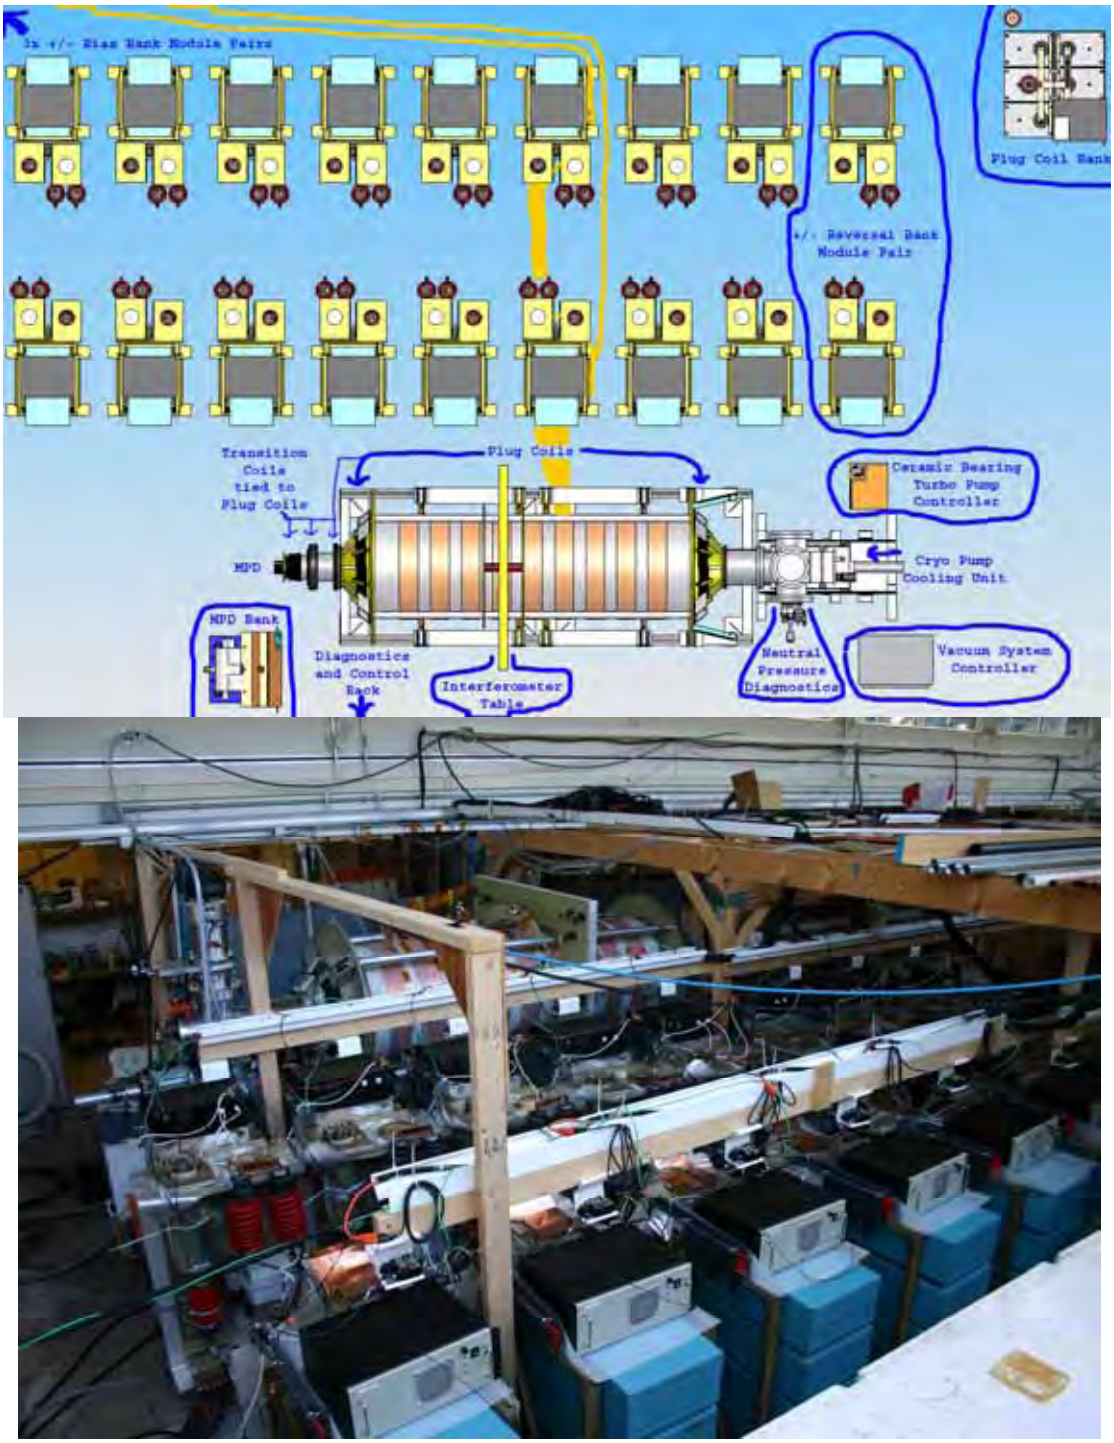

Figure 3.0.1. Experimental layout. This image was taken from a position near the plug bank. At this stage of the experiment, the standard turbo pump has been replaced with a magnetically levitated turbo. The magnetically levitated turbo is not adversely affected by the experiment due to the short duration of the experimental axial fields and the heavy metal walls of the turbo.

3.1 High Power Systems

The primary drivers in the power system layout are distance and cost. Quite simply, for power transmission, cost scales like distance squared. For a given allowable inductance and resistance in a transmission line, if the distance is increased not only does the transmission line have to become longer, it also has to lower its inductance and resistance per unit length. If the distance increases by a factor of two, the individual cables need to be twice as long, but there also need to be twice as many of them. It is often the case that the cables necessary to connect capacitor banks to driven coils are difficult to work with due to high voltage hold-off and current carrying requirements. The greater the distance involved, the greater the total cross sectional area (and hence weight) of the cables.

3.2 Control Linkages

Similarly, the distance from the diagnostic and control rack to the experiment and pulsed power banks is a function of noise on diagnostic cables (outer shield quality), and allowable attenuation from optical fiber cabling (both for diagnostics and controls). For the cheap plastic (1 mm) fiber components that we use this translates to a length range of 1-23 m or 5-30m dependent on the drive level (40 mA and 60 mA respectively) for guaranteed performance. These fiber interconnects are based on the Agilent 'versatile link' system. We typically use the HFBR-1521/HFBR-2521 (650 nm) combination. It has the advantages of easy termination, visible output, shielding, and built in amplification.

The minimum distance specification is due to saturating the receiver amplifier, causing it to take longer during the transition from lighted to no light. Under ideal conditions the 1 – 20 m drive level can be extended with no changes to 1 – 40 m, albeit at the risk of longer delays and flickering of state in the lighted state. 'Ideal conditions' means that all the fiber terminations are of high uniform quality and the temperature of the semiconductor components is 25° C.

An additional option once existed through the use of 200 µ hard clad silica (HCS) fiber for longer link distances with a lower attenuation of 0.008 dB/m (HFBR-H* & HFBR-V*). This fiber works with the same transmitters and receivers as the 1 mm POF. The standard 1 mm POF has an attenuation of 0.22 dB/m (HFBR-R*). The HCS fiber is available from other fiber manufacturers (under different part labeling schemes).

Repeatability of connection losses has been quite good with the utilized 'snap in' connectors. Tests on an individual unit (10 m factory produced cable) for thousands of disconnect/reconnect cycles indicates a standard deviation of 1.6% on the receiver end, 2.9% on the transmitter end. The distribution of results did not conform to a Gaussian profile, but rather a profile with several peaks. The standard deviation inherent in the test setup was 0.16%.

Pure triggers can be heavily overdriven (< 1 Amp) to give greater length limits or improvements in delay jitter. Obviously, this drive level should not be maintained for significant times. Dispersion (due to multi-mode fiber) and optical delay eventually become serious issues. It should also be realized that overdriving the LED results in a different spectral profile. The boards in use allow for the

possibility of drive currents up to 300 mA. However, it is preferable for the sake of interchangeability to utilize only a single level of drive current.

Some systems, such as the bank voltage monitors are much more dependent on the guaranteed specifications holding true, and are thus less capable of working properly with low signal levels due to extra long fibers. The primary concerns are false transitions increasing the apparent frequency and insufficient signal to cause any transitions. This has actually been of benefit in that power failure (if the equipment is designed appropriately) results in a short term apparent over voltage, shutting down the system (if not ignored). The cost of the transmitter, receiver, and terminated fiber for a link of 20 m comes to around $36. No expensive additional equipment is required to make custom terminations.

The closest alternative system (HFBR-1404/HFBR-2402 820 nm) utilizing a glass multimode (62.5/125 μm) fiber comes to around $78. The incremental cost for length is roughly the same for both systems at < $1/m. The more expensive semiconductors are almost 2 x the cost while the terminations are an order of magnitude more expensive. The terminations require specialized equipment to terminate. Hundreds of fibers are required for the overall system.

Longer distances from the primary rack are also possible either with the use of multiple racks (i.e. one rack per each section handling triggering, diagnostics, etc.) or via signal boosters/repeaters (for digital optical signals (everything except plasma diagnostics)). 1mm POF is quite useful for plasma diagnostic purposes whereas smaller fibers are more difficult to utilize. If HCS fiber were more widespread, distance wouldn't be a concern (10-20 x greater length limits).

Originally, it was foreseen that more rack space would be required for additional experimental sections regardless, and hence it made the most sense to go with a potentially length-restricted system. This had the distinct benefit of permitting us to use a visible wavelength, greatly simplifying troubleshooting and permitting the use of plastic fiber that can easily be terminated. It also meant that coaxial diagnostic cables would remain short and hence relatively free of contamination without special measures. The approach of multiple racks has the additional benefit of removing handling damage once an experimental section performs to required specifications.

3.3 EMI

The remaining driving layout factor is stray EMI from the pulsed power banks. When the banks are fired, high voltages are switched generating a great deal of high frequency noise. The frequencies can be expected to be those typically present in high current discharge arcs (broadband, MHz to 100s of MHz). The transmission of these frequencies via the driven antennae on the experiment will be significantly reduced by the presence of snubbers on the outputs of the banks. Additionally, the output cables are not grounded at the experiment (a significant departure from standard experimental design with critical implications in where currents flow when the banks are fired).

For our purposes, the actual transmitting structures are of small size (10 cm at most) and located on the pulsed power banks. The radiation can then be expected to

fall off with (distance/10 cm)^2 rather rapidly to relatively small values (assuming there are no focusing elements in the structure). It is known that placing a notebook computer within approximately one meter of a reversal bank can result in permanent computer failure. It is also known that the use of long unshielded maglev turbo cables can result in touchdowns.

Proper shielding techniques are not cheap to reliably implement. This makes it desirable to place the diagnostics rack (and user's computers) as far as reasonably possible from any of the pulsed power banks. The tradeoff is that increased lengths for diagnostic cables results in increased susceptibility. Hence the diagnostics rack is placed on the opposite side of the experiment from the banks and all the large banks are placed up against the far wall away from occupied spaces.

3.4 Magnetic Fields

Of course, the experiment itself has large and relatively low frequency magnetic fields. With the diagnostic rack in its chosen position, the magnetic fields (at full bank charge) are believed to be barely capable of 'saturating' the rack enclosure's magnetic material (if the field were continuous). A plot of expected variation in magnetic field intensity with distance is displayed in figure 3.4.1. Fortunately, the frequency is sufficiently high such that the current carrying capability of the shield contributes. Thus far, the witnessed 'internal' noise is small enough to not be a dominant problem for most of our diagnostics. An accurate analysis of penetration is unfortunately complicated due to the need to consider a variable permeability (as a function of field) for the steel and the current-carrying (skin) effect. The characteristic skin depth can easily range from 20 μm to 3 mm dependent on effective permeability (at the lowest frequency of interest). The enclosure thickness is approximately 2 mm. This rack, and its attenuation specifications, is discussed more thoroughly as part of the Power Systems (chapter 4).

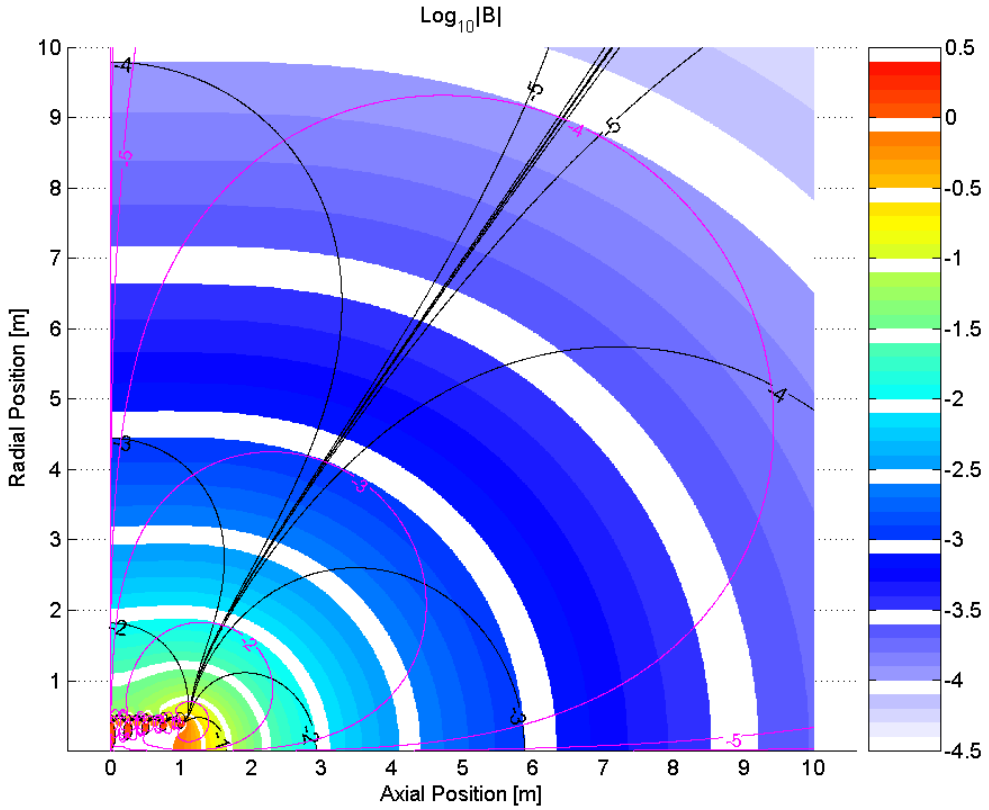

Figure 3.4.1. External Magnetic Fields. This is a base 10 logarithm plot of magnetic fields normalized to the axial magnetic field on axis at mid-plane. A factor of 10 change in magnitude is indicated by a difference of 1. The filled colors indicate absolute magnitude of magnetic field. The magenta contours indicate levels of radial magnetic field. The black contours indicate levels of axial magnetic field. This is plotted for the 80 cm experiment with all coils carrying equal current. The instrumentation rack sits approximately 3-4 m off-axis. Note that 3-4 m off-axis is better than 3-4m on-axis.

3.5 Known Susceptibilities

Obviously, our experiment has the potential to cause interference in systems where physically large ground loops naturally exist or where such systems are naturally sensitive to radiated EM.

The best example of this is common communications systems: radio and wired networks. The high powers are present for short enough duration such that systems at significant distances are generally not too adversely affected (akin to the pop heard on commercial AM radio caused by poorly-installed electric fence systems). While transformers are commonly used at connection points in Ethernet networks, the purpose of these transformers is primarily not high voltage isolation, but rather balancing. Only limited voltages across the transformer are acceptable (usually a capacitor is placed from primary to secondary with a rating near 1 kV). For the frequencies of interest Ethernet cable forms ground loops with AC power systems. We have witnessed Ethernet ports destroyed by potential differences

between the Ethernet cable and the host system. Usually the capacitor across the transformer is destroyed along with a host of other parts. For this reason, where significant distances are involved we utilize fiber optic transceivers.

Generally, this has meant using fiber to bridge between clusters of computers and the diagnostics rack (as well as the outside world). We also use surge protectors on phone lines anywhere that a powered phone is connected. These measures have been successful in preventing damage and maintaining operations. Unfortunately, interference often manifests in poorly shielded computer peripherals such as USB connected devices, keyboards, and mice. Notebook computers have the advantage of the most commonly influenced peripherals being integrated and thus effectively shielded by the chassis. The downside being that their chassis design has a weight constraint and thus has less massive shielding.

3.6 Vacuum Boundary

The majority of the vacuum vessel surface is composed of fused silica. This permits magnetic and electric fields to interact with the plasma without interference. Previous experiments[12] have shown it to be one of the few practical insulators that does not have deleterious effects on theta pinch formation techniques. The regions to either side of the formation region use stainless steel as the vacuum facing surface. In the case of the cones, a thin layer (1.52 mm) has been coated onto a FR4 structural substrate. Most of the flanges, crosses, etc., necessary for vacuum pumps and gauges are composed entirely of stainless steel (with copper-gasket seals where possible). Some O-rings persist in these sections where quick connect flanges are used. The distaste for O-rings in these sections derives more from the dangers of using greases (and creep) than from the inherent permeability of the O-ring materials.

Most of the joints in the vacuum system are sealed with small quantities of epoxy, copper gaskets, or O-rings. However, the joints between tube sections and from tubes to cones are accomplished with the addition of a Kevlar reinforced polypropylene (original plans for LSX were for ultra high density polyethylene) band that has a large exposed surface area (4” wide, 0.1” thick at thinnest point; 80 cm diameter). These bands have quartz pieces placed overlapping on the inside to present a suitable plasma boundary. Due to the flexible nature of these bands, the quartz pieces cannot be sealed to them. Permeation rates are dependent on many different environmental factors often including humidity (sometimes extremely so), temperature, pressure, solvents present, etc.

Given the dimensions and materials involved the dominant gaseous leak source is likely due to permeation through these bands and not the O-rings (at least for gasses). It should be understood that permeability is strongly dependent on temperature. Additionally, humidity can cause some materials to swell, changing the overall mechanical dimensions and the seal.

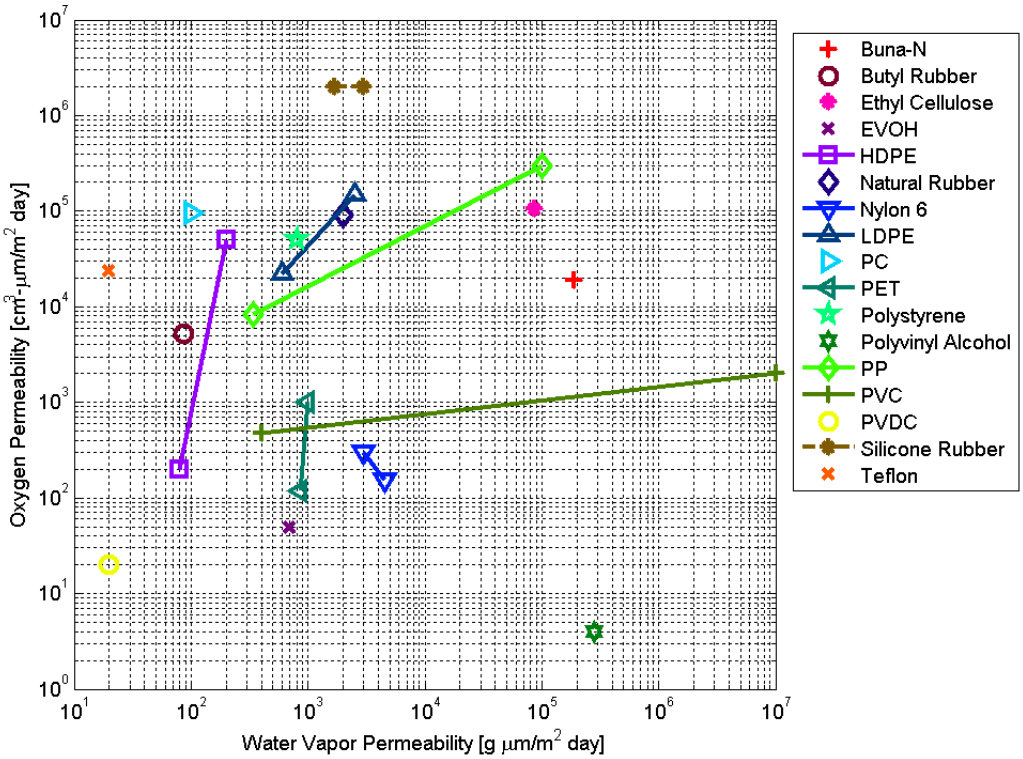

Figure 3.6.1. Permeability of possible vacuum system sealing materials. Information gathered from multiple sources[13 14 15 16 17]. The points in this figure are only approximate as a wide range of values is possible for each material (as demonstrated for some materials with multiple data points; this may be due to different thicknesses of material measured). LDPE has approximately 3 times the oxygen permeability of HDPE. Another source, with less complete tables, has listed HDPE as 100 times better than LDPE for O_2. The effect of thickness is not strictly linear as absorption and evaporation are frequently dominant drivers especially for thin films. We presume that the low permeability points on these ranges are due to tests utilizing very thin films. Note that PVDC (aka WWII era Saran wrap) can only be manufactured in thin films. Buna-N (NBR) has approximately 10-20 times the oxygen and water vapor permeability of Viton (a fluoropolymer, FKM). Oxygen ranges from 3 to 15 times the permeability of Nitrogen for many plastics. A permeation leak tends to distort the N_2/O_2 atmospheric ratio as opposed to a pinhole leak.

Common O-ring materials include Viton and Buna-N. Kalrez has higher permeability than either of these two materials. Depending on the source of the O-ring, additional processing may be required to remove volatiles from the purchased O-ring or remove solvents from the bulk of the O-ring after cleaning. These materials are commonly used to seal systems filled with various liquids rather than high vacuum systems. Hence, O-rings are commonly chosen for reactivity with processes as opposed to concerns appropriate for our systems. It should also be understood that the presence of solvents can significantly worsen the permeability rates and that they can take very long times to pump out of the bulk of the material.

It is not uncommon to use aluminum foil as a temporary dust cover for vacuum components. There is some risk associated with this practice in that the manufacturing process for aluminum foils includes the use of food safe lubricants. Presumably, most of these oils are burned off during the annealing process. However, complete burn-off is not guaranteed. The FDA requires that foils made for food contact have less than 15 µgm/in^2 of residual lubricant[18].

While vacuum greases should not be required with well-prepared surfaces, in practice their use saves significant time when not abused. We have two forms of vacuum grease commonly available in our laboratory. The most commonly used is an ultra high vacuum grease that is hydrocarbon-based, and resistant to irradiation. Apiezon® L has a vapor pressure of 7 x 10^{-11} Torr and a working range from 10 to 30° C[19]. This grease does not tend to creep (low surface tension effect that results in transport along surfaces). It can be removed with a wide variety of solvents or simply baked (@ 400°C) off. Fumes generated from its burning or interacting with hydrogen plasma are generally considered safe.

Silicone greases, while very cheap, are not appropriate for any of our systems (couple of months lifespan, high vapor pressure, creep, etc.).

Krytox® LVP, a fluorinated grease (mix of PFPE and PTFE), has many attractive properties. It has ultra low vapor pressures at room temperature (10^{-13} Torr). Its working temperature range extends from -15 to 300°C. It is very stable. Its chief downside is that with a surface tension of 17-25 dynes/cm it is likely to creep. Since it can creep, it could interact with plasma on the wall. Since it is fluorinated, the products of burning it are dangerous (e.g. Teflon flu). Since it is insoluble to common solvents, it is quite difficult to clean off thoroughly and thus potentially a serious irritant to repair processes involving high temperatures.

3.7 Vacuum Control

The first question that should be asked is whether a fully-automated vacuum control system is worth the time and effort spent. The answer to this question is very dependent on user requirements and specifications for the primary equipment. The Osaka turbo pump model we use has survived multiple catastrophic ‘up to air’ events. One event involved a 4” diameter sudden hole as an improperly secured tube was sucked in through a port. The likelihood for damage may be reduced by mounting the turbo pumps up-side down. Also, cryogenic pumps have survived gate valve openings up to atmospheric pressure while fully cooled. Obviously, such events should be avoided; however, they are not necessarily catastrophic to the pumps.

Most large gate valves derive their motive force from air pressure. Pneumatics suffer from compressor failures and air leaks or clogs. This writer has never seen a motor driven gate valve. Such an actuator would be relatively simple to design using the position switches indicating state as feedback. It would have the benefit of higher reliability than pneumatics, could be easily electrically shielded, and could even close on power loss (UPS systems) if so desired. We consider the absence

of such a control option to be very curious; there must be some reason for it. Hand crank valves do exist.

Figure 3.7.1. Vacuum control electrical enclosure. An air pressure reservoir can be seen at upper right along with the vacuum gauge controller (middle right) and air pressure actuator valves (bottom) driven by a microcontroller control board (lower right).

Our system hardware was built with the expectation of establishing automatic control at some point. The outputs from a microcontroller interface board can be used to directly control most air actuator valves. A reserve air pressure tank is included internal to the vacuum control electrical enclosure. A pressure regulator, filter, and one way air valve are externally mounted to the enclosure. The gauge controller is a Televac MM200 with cold cathode controller and thermocouple measurement modules. The pressures can either be read digitally over the serial port or monitored via semi-logarithmic analog outputs. Unfortunately, the gauge connections are not filtered or surge protected upon entering the overall enclosure. So far, this hasn't proven to be a critical oversight. External communication is accomplished via a 100BT <-> fiber converter.

The microcontroller interface board outputs 9V@1.5A to 24V@2.5A. 7 outputs are available per power supply connection. 3 supply connections are permitted, operating between 9 and 24 volts. Input current per supply connection is rated at 5A. There are thus a total of 21 switched outputs. These outputs can easily drive common SSRs (for example the Omron G3NE series (220 A inrush for 20 ms)) to switch AC voltages or fibers to switch distant (and isolated) valves and pumps.

Each output has a physical manual override that reports back (in aggregate) to the microcontroller. The current states are displayed next to the override switches via LEDs. 4 inset user pushbuttons are also supplied on the microcontroller interface board to allow for switching outputs between user defined states. The user controls require the use of a small screwdriver or pen to prevent accidental activation. The microcontroller (PIC18F8720) has 100BT Ethernet functionality on the board along with 2 standard (TIA/EIA-232-F) DB9 serial ports, a 5V 9 bit input/output or 8 channel, 10 bit, A->D converter port, and a dedicated programming port in addition to the inputs/outputs provided on the microcontroller interface board. This allows for significant future flexibility.

3.8 Vacuum Pumps

At various times several different vacuum pumps have been used. In the beginning, a 480 VAC cryogenic pump was available along with a small, ceramic bearing supported, turbo-molecular pump. Later, a high speed magnetically levitated turbo-molecular pump became available (Osaka Vacuum, LTD TG1113M (pump) / TD711/1111-C (controller)). Controllers are individually tuned for the pumps they drive. The cables are interchangeable. The Osaka pumps may not have the best specifications, but they have proven nearly indestructible in practice. We have bought all the refurbished units we could find.

The magnetically levitated turbo-molecular pump is rated for a base pressure of 7.5×10^{-9} Torr. With the proper setup it should be capable of pumping at a rate of 1100 l/s (N2). The models that we have are technically water cooled. However, there is no internal difference between the water cooled and fan cooled versions. We use an external fan to cool the housing/heat sink of the pump unit. During power outages the rotational energy of the rotors supplies electricity to the levitation system until the rotors slow down to 30% of top speed. Unfortunately, power loss qualifies as an 'error' and puts the pump into a safe mode (i.e. not spinning). This pump is designed to withstand 5 unintended touchdowns before requiring service of its protection bearings. Simple power failures and standard operation do not count towards this total as they do significantly less damage than a failure of magnetic suspension at full speed. Steady magnetic fields are acceptable so long as their magnitude is less than 3 mT.

We have had only one confirmed full speed touchdown event. This was caused by the substitution of one maglev pump for another. The refurbished pump substituted in came with a 5m long, unshielded cable (supplied by the refurbisher, not Osaka). A shot fired at 40% energy resulted in interference causing a high speed touchdown to bearings. Up to that point we had only used factory 1.8m shielded cables between controller and pump. Additionally, users should be aware that subsequent attempts to levitate (stop followed by start) that fail result in incrementing the touchdown counter. Sometimes the vanes can be re-levitated while spinning down on the touchdown bearings, sometimes not.

There is one oddity regarding the external control port that users should be aware of. From conversations with Osaka's technical support it was learned that

several of the ‘grounds’ specified for the external control connector can not be connected. These grounds correspond to different internal grounds. We originally learned of this discrepancy while checking continuity on the port. Connecting the grounds can result in certain noise induced errors. This restriction is simple to comply with via either opto-isolators or relays.

The maglev turbo was made for single phase power systems. Its nominal input voltage is 200 V @ 50 Hz and 220 V @ 60 Hz. It is rated for over voltage category II, i.e. office installation, not rated for an industrial environment. Its rated maximum power is 1.6 kVA. The maximum current draws we have actually witnessed imply a significantly lower maximum power draw during normal operation. Acceptable input voltage is from 180 to 240 V with an input current of 1.5 A under minimal load (i.e. operating with no gas loading and no backing pressure). The mean of this input range is 210 V. An excerpt from the accompanying manual (emphasis added) points out some special requirements:

> 3) Input Power
> Select an isolated source with a low noise, surge and a good voltage regulation. All models must only be connected to a power source or system that is adequately isolated (floating) from ground and **never to a mains-type (branch circuit) supply**. Connection to any other supply will render the equipment unsafe in terms of lethal electrical shock…

The implication is that this piece of equipment is not compliant for an office installation without the presence of an external isolation transformer (industrial applications are even more stringent in their requirements). We might suspect that the risk is not from normal operation but rather from power line spikes that occasionally occur (even with surge protectors). Humidity and smoke can influence acceptable clearance distances. Communicating with the manufacturer provides the further enlightenment that the insulation distances were not sufficient “to be certified as a conformable product to UL standards without” an external floating neutral and hot isolation transformer.

Note that the designers originally assumed single phase input power. I.e. one connection hot, one connection neutral (i.e. close to ground potential), and a safety/chassis ground connection. Also realize that in a conventional switching power supply the voltages present with respect to ground are lessened only by 20% (for a very relevant example) by connecting the neutral input to one phase of 120 VAC and connecting the hot to a 120° different separate phase of 120 VAC (to obtain 208 VAC single phase from 208 VAC 3 phase). The plots of voltages present with respect to ground can be seen below. If this was a suitable solution the designers would have listed it in the manual (as it costs no money). In the proper solution (as specified in the manual), the presence of the isolation transformer insures that no significant current can flow to the user if an arc to the user occurs, the same cannot be said for the 2/3rds of a three phase approach.

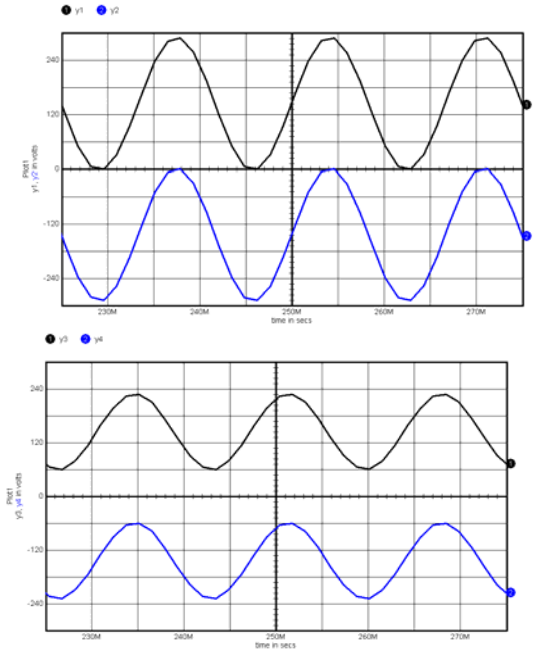

Figure 3.8.1. Comparison between single phase and 2/3 of a 3 phase line inputs. The diagrams describe the voltages present w.r.t. ground of the two terminals on the DC capacitor after rectification. No power draw is assumed for simplicity. The diagram at left is driven with a single phase 208 VAC source. The diagram at right is driven with two separate phases of three phase 120 VAC.

Due to this very specific warning / installation instruction we chose to buy an isolation transformer (<$1k) that also takes the readily available 110 VAC single phase and transforms it to 208 VAC single phase with a floating output. Electronic Specialists, Inc. makes such a unit (ISO/T-34 (14) B208) that also incorporates surge suppression and filtering. 'Alternate connection A' is the manual listed optional alteration that results in floating power outputs. Note that the earth/safety ground must be connected (3rd pin) through the transformer.

Additionally, we choose to add a standard uninterruptable power supply (Tripp Lite's OMNIVS1500XL ~$150) to the input of the isolation transformer to deal with brownout and temporary power loss scenarios. Some care is required in choosing a specific model of UPS since the acceptable deviation in input voltage for our equipment is only +/- 15%. It is not uncommon to run into models that don't boost the output voltage sufficiently during severe brownouts. The model we use is specified for –18%/+8% chiefly because the automatic voltage regulation transformer turns ratio for the severe brownout boost is not set ideally for our application. We have found that this model of UPS degrades when subjected to frequent (i.e. every hour for weeks on end) brownout conditions until it is no longer able to maintain +/- 15%. From a certain perspective, this might be considered a feature since significantly more expensive units (that utilize AC->DC-battery-DC->AC converters continuously, Tripp Lite's SU1500RTXL2Ua ~$600) would work more reliably but really shouldn't be required. All other standard UPSs that we've used have failed to maintain sufficient voltage through severe brownout conditions to avoid tripping the

maglev pump's under-voltage lockout. Alternatively to finding a better UPS, a slightly different isolation transformer might be found that specifies 110 -> 208 VAC rather than 120 -> 208 VAC. Care should be taken to insure that the maximum input voltage specification is not exceeded.

The backing pump for the turbo needs to be capable of 15 m^3/hr. This translates to roughly 400 watts if efficiency could be 100%. We use a BOC Edwards XDS 10 dry scroll pump rated for 13.4 m^3/hr @ 360 watts with an ultimate pressure rating of 40 mTorr. It does require annual maintenance and some attention to initial setup instructions. It can be overwhelmed by gas flow from the turbo (as could an appropriately overrated backing pump under certain circumstances).

3.9 Safety Considerations

The layout of an experiment directly impacts various safety concerns. The high-voltage equipment placement and work space around equipment should be first and foremost in keeping with National Electric Code specifications. This section is not intended to act as a substitute for either code or common sense. While our systems can be brought offline for servicing, it is conceivable (although we hope highly improbable) for failures to result in the presence of lethal energies when the power is off. Other safety concerns are addressed below and incidentally through the power systems chapter.

3.9.1 Electrical Isolation.

The general rule is to not allow any wire connections to pass through the experimental enclosure. Of course, this can prove to be prohibitively expensive due to the cost of high power isolation transformers. The primary concern is that any electrical connection linking inside the experimental area with outside the experimental area can carry enough energy to kill a person. A secondary concern is that the driving of current through paths external to the enclosure can result in hazardous voltages being present between separate external grounds. When the experiment is fired, mega-amps flow in some parts of the system. Only a small fraction of this current would be required to momentarily lift an enclosure connected ground (or electrical power ground) to potentially lethal voltages relative to other (user accessible) grounds.

This threat can be removed in several ways. The fool-proof approach is to provide galvanic isolation via isolation transformers and a fully enclosed conducting shield around the experiment. Clear advantages exist in that the EM generated by equipment associated with the experiment would be contained. In practice, such an approach would be very expensive due to the need for isolation transformers capable of handling large amounts of power and, more significantly, the large amounts of metal required to adequately enclose a large volume experiment with high stray magnetic fields.

With our experiment, the majority of the vacuum chamber is non-conducting. Additionally, the primary high current systems are magnetic (and individually insulated to the extent possible). Our experimental boundary is non-conducting in

nature. Ensuring that all power entering the enclosure follows the same external path (i.e. one conduit) and enters at the same location (making external experimentally directly-driven ground loops unlikely) is a feasible approach that satisfies safety concerns. Internal to the experimental enclosure this can be thought of as a star-point ground. Clearly, not allowing internal ground loops to form is a concern from an electrical noise and equipment survival standpoint. Skimping on insulation and lack of insulation redundancy (for example between flux loops and driven loops) is highly unwise. For example, the potential exists for two times the applied bank voltage to appear on the driven loops due to the inductive nature of the termination and the presence of a transmission line alone. Normally, you'd be counting on snubbers and a slow closing switch to limit this. For a small savings in material cost, labor, and simplicity one runs the risk of destroying thousands of dollars in equipment and lost time by designing systems with no derating from absolute maximum withstand parameters. This will be discussed further in the 'power systems' section (chapter 4).

In our setup, aberrations from a no-ground-loop implementation arise due to the 480 VAC electrical feed for the cryogenic pump and the 120 VAC required for the axial discharge system and vacuum gauges and controls. These systems' grounds are connected at the vacuum chamber as well as at contractor's ground. Unfortunately, the 480 VAC power feed runs along a different path from the rest of the power connections. They did form a significant internal and external loop. Fortunately, while the external loop via the 480 VAC cryogenic system could be driven, it was more directly bypassed by the small internal loop of the 120 VAC vacuum control and 120 VAC axial discharge system.

If the ground of the coaxial cable driving the axial discharge were interrupted for any reason (for example by someone putting a resistor in line with it) the preferred return would have been through the vacuum control's 120 VAC power line (that also fed the axial discharge system). Only small currents would have flowed external to the experimental enclosure. If, in addition, the vacuum chamber ground became disconnected (highly unlikely) the axial discharge would drive large currents through the 480 VAC ground back to the circuit breaker box and back along the 120 VAC ground to the axial discharge capacitors.

The 480 VAC and 120 VAC external feeds were widely separated in space, such that bridging the two by one human (without the use of extension cords) would have been impossible. To ensure that the external ground loop wasn't further driven by unanticipated arcs, the 480 VAC wiring internal to the enclosure was kept widely separated from any other conductors internal to the enclosure except at the cryogenic pump (so that no high current feeds could arc to it).

3.9.2 Redundancy

Any element of a system required to safely dissipate stored energies should be made in such a way as to be safe in the case of a single-point, and in a lab environment a two-point, failure. For capacitor banks we typically have a single point failure resistant fast dump. To add a second level of protection our voltage monitors act to dissipate the charge over long timescales, and more importantly warn

the user (both audibly and visually) that dangerous voltages still exist in the monitored system. This makes the faults detectable. The reason such redundancy in a lab setting is necessary is that the systems are experimental and assembly is often imperfect. The often-told story of a student repairing a dump resistor with a non-conducting adhesive is an example of a situation where the extra redundancy clearly proves its worth[20]. In this system, it would still be two failures away from being invisibly dangerous. The second dump resistor would have to fail, as well as the connection of the voltage monitor to the bank. The voltage monitors report back to the control system as well as positively indicating their powered state. Our dump resistors also double as charge resistors; making it impossible to charge the bank if the dump resistors are fully disconnected. This is only somewhat effective, as will be explained in the power systems chapter.

3.9.3 Electrocution by Coax

The majority of energy temporarily stored in the coaxial feeds between reversal banks and driven coils is actually due to snubbers (resistive, capacitive filters) placed on the bank outputs. It would take nearly 3 miles of coax to equal the capacitance in one of these snubbers. At 0.5 μF per snubber, and 6 snubbers per bias bank segment, a maximum of 1 kJ @ 25 kV is available. Note that charging the cable to 25 kV and not discharging would require the firing of one switch followed by a bank dump. High powered 'Tasers' use 1.8 joules. Electric fence energizers utilize anywhere from 0.1 to 36 joules per pulse. The external exposure threshold for cardiac ventricular fibrillation is widely assumed to be between 10 and 50 joules. To be under this threshold the charge voltage must decay below 2.6 kV before anyone can reach it. This implies that the RC discharge time must be at least 2.3 times quicker than the time taken for a human to enter the enclosure and reach any of the exposed coax or connected systems. Redundancy can be accomplished by multiple snubbers, and hence multiple discharge resistors, connected on the same segment. Most of our snubbers are linked to multiple segments due to the bias banks.

If we assume it takes 60 seconds for a human to gain contact, then a 50 MΩ resistor would be required across each output (worst case). Similarly, 4 minutes implies 200 MΩ and 12 seconds 10 MΩ. So why not put a small resistor in place? The resistor must not only be capable of dissipating the energy in the cable and snubber safely; it must also be able to withstand the continuous application of 25 kV in the event that an ignitron glows continuously (and thus conducts) during charging. Thus, a 10 MΩ resistor would have to be rated well above 60 watts for reliability. A 50 MΩ resistor would have to dissipate 13 watts (MOX96025005, 16 watts, $16). The MX480 resistors used in monitoring are only rated to 7.5 watts which implies a 200 MΩ resistor, 4 minutes, plus locks on the doors. Unfortunately, this means that procedure has to be followed for safety.

Alternatively, a lower value resistor can be used that is rated for the power supplies' output (plus the energy of the banks). Of course, the obvious question to ask is: do the snubbers really have to be so large? The short answer is no, not if the experiment doesn't ring (i.e. crow-barred with no ringing theta ionization system).

Having smaller snubbers has implications for transmitted noise. The snubbers presently in place have ~ 1 m coax connections implying significant inductance and inability to effectively 'snub' high frequencies. Dissipation, in some form, is needed to prevent large oscillations in current across the start switch causing large oscillations in voltage across the crowbar switch and the diagnostics. However, this dissipation can take other forms and the amount of energy that has to be dissipated can be reduced by lowering certain stray capacitances. We haven't made a significant attempt at alternate approaches due to the development time that would be required. This is discussed in more detail in the power systems section.

3.9.4 Capacitor Explosion.

The primary consideration here is how much energy a single capacitor is capable of taking before breaching the outer casing. This energy limit then defines the individual bank (module) size limit. An example of a typical failure leading to distortion of the outer case is an internal short (for example due to excessive charge voltage). The potential exists for the energy in all the capacitors connected to the shorted capacitor flowing into the short. The primary energy storage capacitors we use (14 μF 25 kV GE Cat#30F1141) have a recommended 15 kJ rupture benchmark. This is unfortunately not a guarantee that rupture is impossible at any energy level. The amount of energy on one switched bus of our reversal banks is 14 kJ @ 20 kV. The insulating oil used is combustible but not flammable. Thus, the risk of a sustained fire from a ruptured capacitor is low. The experimental enclosure is expected to act as a second layer of protection against shrapnel given the uncertainties involved (the somewhat unlikely possibility of magnetically coupled energy from other banks playing a role, and the uncertain energy required to rupture the capacitor).

3.9.5 Capacitor 'Opens'

This is the most potentially dangerous form of capacitor failure chiefly because it is usually not immediately apparent that a failure has occurred. This failure type could occur when a capacitor is run with currents exceeding absolute maximum ratings resulting in the fusing of internal wiring. We have seen several such failures in auxiliary capacitor bank systems. This kind of failure can be recognized by an excessive change in effective capacitance. The danger associated with the failure is heavily dependent on capacitor construction. Consider what might happen if a subset of the capacitors capacitance becomes disconnected while still retaining some charge. This charge might be seen leaking out to the still connected capacitor segments for extremely long timescales. The first time this failure occurred we watched the current decay over days. If a capacitor in such a state does not have it's outputs shorted this effectively allows it to charge the still connected segments, with potentially deadly consequences should human contact occur. Also, it might be possible to re-establish a 'good' connection to the segment via mechanical motion or vibration resulting in the sudden release of the remaining energy. This is a second reason why 'dead' dumps, not just resistive dumps, are advisable for when capacitors are not in use.

3.9.6 Capacitor Dielectric Absorption

In a properly-designed capacitor bank this effect should never be a concern due to the presence of automatic dump resistors and shorting sticks across the bank when not charged. However, it is easy to neglect hidden capacitance in transmission lines and snubber components as well as capacitor bank components during disassembly. When capacitors are charged, some fraction (up to 1% dependent on dielectric type) of the energy input into the capacitor does not come out upon discharge. Some of this energy is dissipated in the form of heat, while some of it is still stored in the dielectric that requires more time to access (memory). The basic effect is that the short application of a resistor (or even a short) to dump a capacitor does not result in a complete discharge of all the energy. The remaining energy is then free to partially recharge the 'fast' capacitance resulting in the improperly-trained technician getting a potentially lethal shock upon contact.

3.9.7 Collapse of Axial Load-bearing Columns

The axial compressive forces created by the vacuum are transmitted by four aluminum alloy (6061-T6 96% pure aluminum) rods from the FR4 end cones. This alloy has a modulus of elasticity of 70×10^9 N/m^2 (10^4 ksi) and an expected yield strength of 255×10^6 N/m^2 (37 ksi) in compression. Aluminum alloys' compressive yield strengths vary greatly (10-600 MPa). The force on a rod from the end cone's vacuum loading can reach 5×10^4 N (6 kip). The rods are actually cylinders with an outer diameter of 2.5 inches and an inner diameter less than 1.5 inches. In order to do the analysis properly, assumptions as to eccentric loading and end support have to be made. The most critical realization can be summarized in Figure 3.9.7.1.

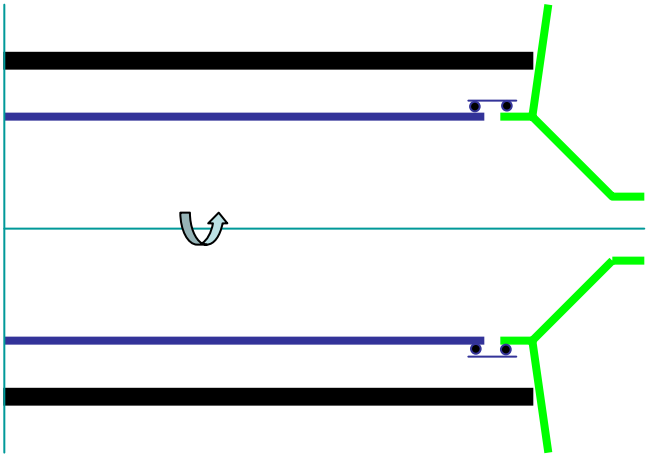

Figure 3.9.7.1. Vacuum loading conceptual diagram. Silica vacuum chamber is blue. End cone structure is green. Aluminum axial support rods are black. As the vessel loads, the end cone can be expected to distort in the manner emphasized here.

Due to how the aluminum rods are attached to the end cone, it is most closely approximated by a pinned condition (free to rotate), with a highly eccentric loading (i.e. not loaded through the neutral axis of each rod). Looking at the system under

load one can actually see the air gap between rod and end cone. The total vacuum loading on the cone is 12 kips. The appropriate equations to analyze the rods can be found from an undergraduate text on structures[21].

$$\sigma_{\max} = \frac{P}{A}\left[1 + \frac{ec}{r^2} Sec\left(\frac{L}{2r}\sqrt{\frac{P}{EA}}\right)\right] \text{ (maximum stress)}$$

$$\upsilon_{\max} = e\left[Sec\left(\frac{L}{2}\sqrt{\frac{P}{EI}}\right) - 1\right] \text{ (maximum deflection)}$$

$$I = \frac{1}{4}\pi\left(r_{out}{}^4 - r_{in}{}^4\right) \text{ (second moment of inertia)}$$

$$A = \pi\left(r_{out}{}^2 - r_{in}{}^2\right) \text{ (cross-sectional area)}$$

$$r = \sqrt{\frac{I}{A}} \text{ (radius of gyration)}$$

e (eccentricity) = 1.25”
c (neutral axis to outer structure) = 1.25”
A = 3.14 in^2
I = 1.67 in^4
r = .728“
E = 10^4 ksi
L (effective length of column) = 100”

Figure 3.9.7.2. Equations for determining failure of support columns.

The maximum loading occurs if two diagonal rods are slightly shorter than the opposite two, splitting the vacuum loading down to 6 kip per column. Note that this can also occur due to a pre-deformed cone. With the eccentricity defined for all the loading at the edge of the column (consistent with loaded observation), the maximum stress is 11.4 ksi. This is 6 times greater than the average stress, but still leaves us with a safety factor of ~ 3. The maximum deflection under this load is 0.9 “. With all the columns sharing an equal load the maximum deflection is 0.34 inches.

The use of the “secant formula” hides the relation between length and maximum stress. Euler’s formula ($P_{cr} = \frac{\pi^2 EI}{L^2}$) applies to cases with no eccentric loading. Euler’s formula defines the pressure (or effective length) at which the column will buckle. For comparison, the secant formula yields a buckling condition with the column length = 153 inches. Euler’s formula yields 165 inches. The difference is due to eccentric loading.

One last point, if we had assumed that the rods were ‘fixed’ to the cone, rather than pinned, we would have made the mistake of taking our effective length as half

the physical length. Later, we would have extended the length of the experiment with what would have appeared to be a healthy safety factor. The columns would have yielded (unless the FR4 and metal plates could sufficiently distort to guarantee even loading) until the unloaded columns picked up their share of the load (leading to some permanently-bent rods).

Figure 3.9.7.3. Axial support rods under vacuum loading. The red lines have been added as a reference for comparison. The deflection can be seen (deflection ~ 20% of diameter, 0.5" deflection) viewed at an angle likely to indicate only 70% of the total deflection. This figure is intended to emphasize how easy it is to overlook such a seemingly large deflection, not to quantify it, as the rod may have been bent to begin with due to earlier use on a longer experiment or due to original manufacturing deviances. Note that other structures that should be straight show as straight, lowering the probability of a camera lens aberration explanation.

3.9.8 Quartz Chamber Implosion

The available energy in the vacuum is approximately $E=P\ V= 0.1$ MJ. About the same amount of energy as available from one standard D cell battery. It is a serious mistake to assume that all this energy is going to end up in one small shard. How exactly is all the energy going to get coupled into a small shard? A more realistic approach would be to consider a small chunk of quartz cracking loose. The kinetic energy imparted to this piece is approximately $\Delta E = F\ d = P\ A\ d$. After the shard no longer seals the hole, the pressure on the shard can be expected to rapidly decrease. The mass of the shard is $m = \rho\ A\ d$. The kinetic energy is $\frac{1}{2}\ m\ V^2$. The velocity imparted is then $V = \sqrt{(2\ P / \rho)} \sim 9$ m/s inwards (20 mph). Not exactly a bullet. Of course, if P could be increased to say 100 atmospheres it would be quite dangerous.

3.9.9 Feed-thru Rifle

More dangerous possibilities arise from tube feed-thrus sealed by O-rings. This allows for lower average densities (ρ) with a cylindrical cross section (and thus higher velocities). Still worse is the rifle barrel situation. A piece breaks loose inside a metal tube connected to the vacuum system and is accelerated over the length of the tube. Of course, if it is going to pose any danger it then has to have sufficient energy to break through the other side. For a typical feed-thru (1/4 inch) with reasonable assumptions the maximum available energy is P*V* (diameter of feed-thru / diameter of vacuum vessel)^2 ~ 6 joules. This is significantly less energy than a firearm bullet (100s of joules) and is on the order of energy expected from air rifles and paintball guns. Standard eye protection is sufficient to protect against this threat while in the experimental enclosure. Obviously, larger feed-thrus present significantly higher risks; keep in mind that in all cases they have to be able to break through the opposite side to present significant shrapnel danger. Unfortunately (from a certain point of view), no one was around to witness the 4 inch diameter feed-thru getting sucked in and breaking a thin-walled quartz tube. The debris field did not indicate a high degree of danger.

3.10 Labview 7.0

Labview (from National Instruments) is a limited multi-threaded programming language that interprets program schematics. These schematics are largely non-text based. Individuals not trained in any of the classical programming languages may find this approach easier to understand than C or Matlab. The downside is that this schematic based approach has low information density, is indeterminate in order of execution, and thus takes more time to debug and fully understand than working in Matlab, Basic, C, Java, or the other standard programming languages (for the same tasks). A page of code in Matlab can take 10 pages to implement in Labview.

A schematic in Labview can run in multiple orders, so care must be taken when multiple elements use the same hardware. How it runs in debug mode is not how it runs normally.

The code also tends to be painful to add comments to. The reason being that free space must be made manually, which means moving the rest of the code while not breaking any links. So, commenting after the code is running becomes a painful and laborious process for any complex tasks. One way is to add “Description and Tips” pop-ups; which don’t print out well. Due to the manageability issues of an icon-based programming language, most visually complex functions are shockingly simple in execution. For the same reason the code has to be heavily chunked into sub-functions (sub-“Visual Interfaces”) to avoid massive sprawl.

Given these complaints, why do we use Labview? It works. Labview is stable with very few crashes. Its low information density is sometimes an advantage in understanding how a new program works. Using Labview ensures compatibility with National Instrument’s hardware. It is also hardware centric; it supplies all the GUI and communications interfaces necessary to make a practical system. Many

third party manufacturers make hardware with Labview drivers. We have also made Labview "drivers" for certain CAMAC vendors. Labview is limited multi-threaded.

We typically use Labview from a remote computer connected via Microsoft's 'Remote Desktop Connection".

3.10.1 Labview Functions

Another peculiarity is that all functions are global functions. If a function called in one program (VI) has a different structure (but same name) as that called in a different program, then both programs can not be run simultaneously (and should not even be loaded simultaneously to prevent corruption). Let's say we want to modify a program that downloads data from a digitizer. We have a model 2028 driver; we also need a model 2008 driver. Most of the 2028 sub-functions are hardware specific and may need to be altered. These are the steps we would have to do:

1) Use the OS to copy the 2028 driver to a new name.
2) Use the OS to copy all the sub-functions to new names.
3) Labview has no find and replace command that is useful for this purpose. We go through all the code, manually changing references to 2028 driver sub-functions into references to 2008 sub-functions.
4) Alter the code to make it work for a 2008.

3.10.2 Run Modes

The default options result in 3 running threads: One for the user interface and two for executing user code. The following settings can be accessed by right clicking on the VI connector pane icon and selecting "VI Properties". VIs can be run in one of six execution engines; this can be set in Execution: "Preferred Execution System". Most of these engines have two associated threads. All the driver VIs used here have resource semaphores that prevent VIs from accessing the same resources at the same time. Lower execution times can be achieved by setting data acquisition sub-VIs into a different execution engine from the other code sections. Of course, what would really speed things up is if each VI were allowed its own execution engine. As it stands, this is what typically happens:

1) All the VIs queue up to download data from digitizers.
2) While the first digitizer downloads, all the other VIs wanting the same interface give up their processor time due to the resource semaphores.
3) The first VI to download then queues up to store the data.
4) That VI then has to wait for it to be its turn again.
5) Most of the other VIs complete downloading data
6) The first VI gets to the point of actually storing its data…

A lot of allocated memory can be required at one time. It is also not guaranteed that best use of the processor is made. If this becomes a serious issue it can be addressed by changing the execution priority such that after downloading data, that VI has a higher priority (or gets a different thread) than the rest.

Programs (VIs) are more accurately thought of as being analogous to an ActiveX object rather than as a C function. The default settings allow sub-functions to remember previous calls made to it from different callers. For our purposes the most important setting is “Reentrant execution”. If this box is not checked then data can be unintentionally shared between multiple VIs. In theory, this should not be a problem. In practice, why take the risk? Checking the “Reentrant execution” box has the side effect of disabling the debug modes.

3.10.3 Memory Usage

There are a couple of ways that VIs use memory in hidden ways. Most problems result from using graphs. The memory a graph takes is proportional to the largest data set that has EVER been placed in it (followed by a file menu save). Attempts to clear the data, or replace it with small data sets, result in lower memory usage but not as low as its original state. Graphs are not efficient users of memory to begin with; when running multiple (large memory) digitizers this becomes a serious issue the programmer has to cope with.

3.11 Control System Operation

There are typically several prerequisites for running the included drivers and applications. The usual prerequisites are:

1) A recent version of Labview (typically version 6+)
2) An installed communications interface device and drivers (such as a GPIB interface or Ethernet)
3) DatAcqSvr.exe (an ActiveX control that handles the HDF 5 database accesses). DatAcqSvr.exe needs to be placed somewhere on the hard drive and run from a command prompt: “DataAcqsvr /RegServer”. The proper use of administrator privileges is required for Windows Vista.

Individual device drivers have varying requirements associated with hardware expectations. Note that LeCroy 8901A GPIB interfaces are NOT functionally equivalent to LeCroy 8901 GPIB interfaces (one or two bits in GPIB communications have to be different for CAMAC commands to be executed properly). One device requires that the GPIB read command (used to cause CAMAC execution) request bytes; the other requires that it request 0 bytes.

3.11.1 Single Device

Individual “drivers” can be run manually to control single digitizers, pulse generators, etc. These drivers can either be contained within a single ‘.vi’ or contained within a ‘.lib’. Within a ‘.lib’ file the appropriate ‘.vi’ file to run is the VI with the same name as the library. For example, the library named ‘2028 Driver for 8901A.llb’ must have a vi named ‘2028 Driver for 8901A.vi’ contained within.

Upon opening a single driver, it is common to be presented with a list of warnings that referenced functions were found in different locations. Certain VIs vary with Labview version and hence should not be included in library files. In general, only VIs made by the user should be included in library files. For example,

any “Semaphore” VIs should not be included in the library (or other versions of Labview will fault). After acknowledging any warnings, the user will be presented with a screen similar to the following:

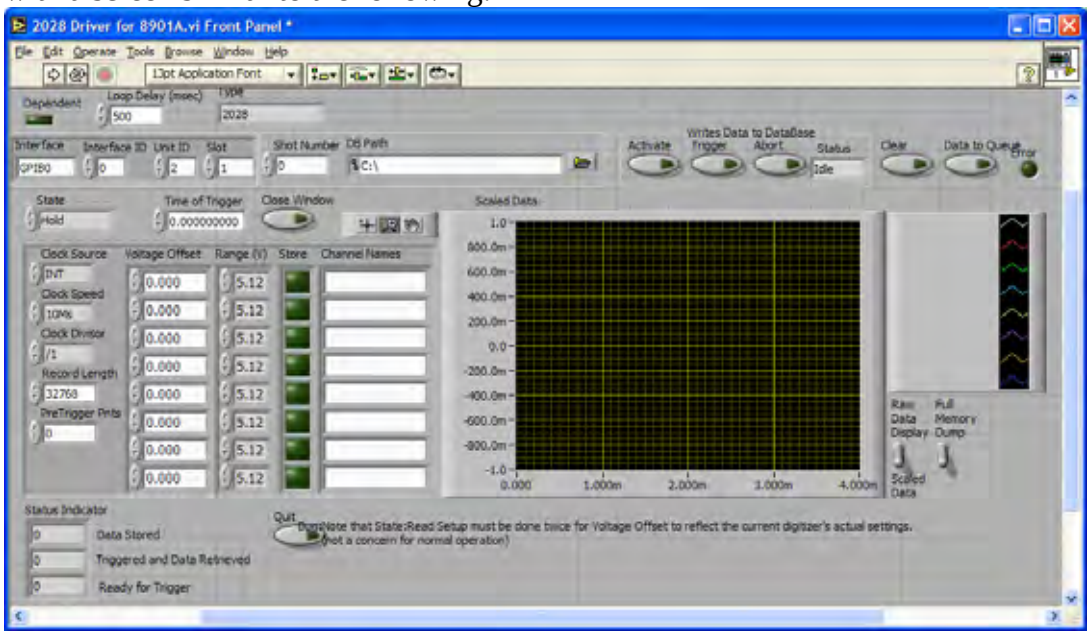

Figure 3.11.1.1. 2028 Driver for 8901A Front Panel.

Most drivers have “Pop-up” tips that help explain the meanings of the various fields and switches. Most drivers are expected to have the following names for compatibility: “Dependent”, “Type”, “Interface”, “Interface ID”, “Unit ID”, “Slot”, “Shot Number”, “DB Path”, “Activate”, “Trigger”, “Abort”, “Status”, “Error”, “Time of Trigger”, “Close Window”, “Quit”, and “Status Indicator”. Not all of these need to be present for every driver. For an example of a minimal driver see “Delay.vi”

The “Interface” field needs to be set before running the VI. This field is used to prevent conflicts between multiple drivers using the same state-full communications hardware (CAMAC, GPIB, parallel, serial). If two drivers talk to digitizers under the same crate controller, they should have the same (arbitrary) “Interface”. If two drivers talk to digitizers in different crates, but on the same GPIB bus, they should have the same (arbitrary) “Interface”. If two drivers talk to digitizers in different crates, on different buses, they can have a different “Interface” and download data at the same time.

In the case of a GPIB device, “Interface ID” refers to the unique GPIB identifier (usually zero) associated with the GPIB card.

For a CAMAC/GPIB interface, “Unit ID” refers to the GPIB address associated with the interface.

Slot number refers to the CAMAC slot the device is in.

The user needs to “run” the VI so that the driver can start responding to commands. Click the “run” button:  . Select an existing directory, such as “C:\”, for “DB Path”. To make the digitizer ready for a trigger click “Activate”. While the

"Status" reads "Awaiting Trigger", the VI polls the digitizer every ½ second (as specified in "Loop Delay (msec)") to determine whether a hardware trigger has been received and whether the digitizer is done with the acquisition. To trigger the digitizer via software, click "Trigger". The driver then downloads, displays, and stores any data. Note that "Store" must be on and "Channel Names" must have a valid string to store the data to an HDF database file.

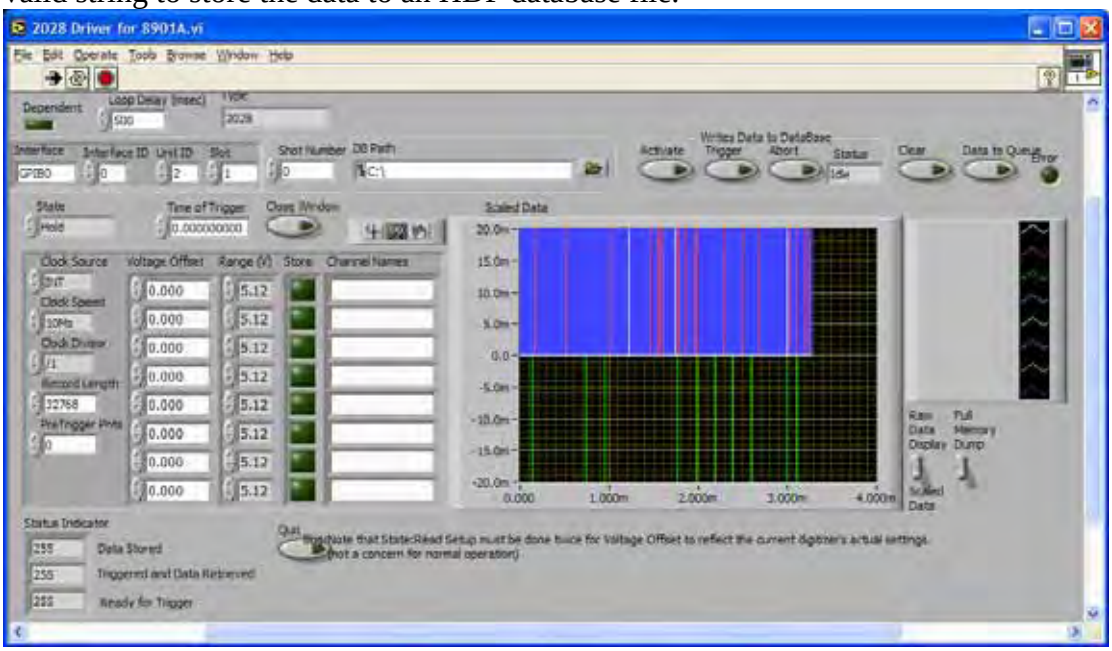

Figure 3.11.1.2. Example of 2028 Driver Front Panel after Collecting Data.

"Close Window" is used to quit or hide the driver window. Never use the Labview "Stop" icon in normal operation since this leaves semaphores un-reset.

3.11.2 Multiple Devices

Multiple device drivers can be run simultaneously using the "Single Device" methods so long as proper use of the "Interface" field is maintained. Note that multiple instances of the same driver file can not be run so simply. To overcome this limitation, and provide automation for the data acquisition process, a "controlling" VI was written.

The "Driver Interface.vi" acts as a control program over all the individual drivers. It deals with the tasks of making individual running driver copies for every digitizer, output buffer, logical unit, etc. It allows for the saving of individual settings and automatic control during shots. This information is saved as ".cfg" files, allowing for later return to previous conditions by archiving them. All front panel control settings are saved (not just the special named controls listed earlier). Most importantly, it is malleable (based on an object driver concept). All the drivers are continuously executing, allowing for functions such as bank charging control and monitoring as well as vacuum system control (with some careful attention required in hardware implementation). Adding additional hardware takes a few minutes, without

need to change any underlying code. It is also possible, with some minor modifications (of how multiple copies are made), to make an executable version of Driver Interface. It can then run without a separate Labview license.

Figure 3.11.2.1. Driver Interface Front Panel.

Initial setup:

1) Run “Driver Interface.vi” inside “Driver Interface.lib”
2) Start “Driver Interface” 
3) Add the wanted drivers to lines in the “Devices” array by clicking on the folder icon  . The number of drivers is limited only by available memory.
4) Insert the appropriate “Interface”, “Interface ID”, “Unit ID”, and “Slot” information for the device
5) Click on “Active” until the green light stays on. The “Error” light should clear
6) Insert the appropriate “DB Path”, “Shot Number”
7) Turn off “Automatic Trigger”, or set “Location” to one of the drivers. Typically set to a digital pulse generator, or a single digitizer for continuous testing.

To setup the individual drivers click on “Front Panel” and set up the driver as normally done when running a single device. “Close Window” keeps the VI running but hides it from view.

To save the driver’s settings, click on “Active” once to turn off and save the driver, then again to load and run the driver. This may take more than two mouse clicks.

Drivers respond the same as when run alone. The “Driver Interface” adds the ability to start them all together (“Start”), automatically trigger one of them

("Automatic Trigger"), automatically increment the shot number ("AutoIncrement"), and run multiple shots with no user intervention ("Continuous Fire"). When "Automatic Trigger" is off, individual drivers can be triggered by clicking in the "Trigger" column.

"STOP" closes all the drivers without saving and stops "Driver Interface" in preparation for saving the "Driver Interface" front panel settings.

Saving the "Driver Interface" settings is accomplished by "STOP" followed by "Operate: Make Current Values Default" followed by "File: Save".

Several additional "Drivers" are available beyond those displayed below:

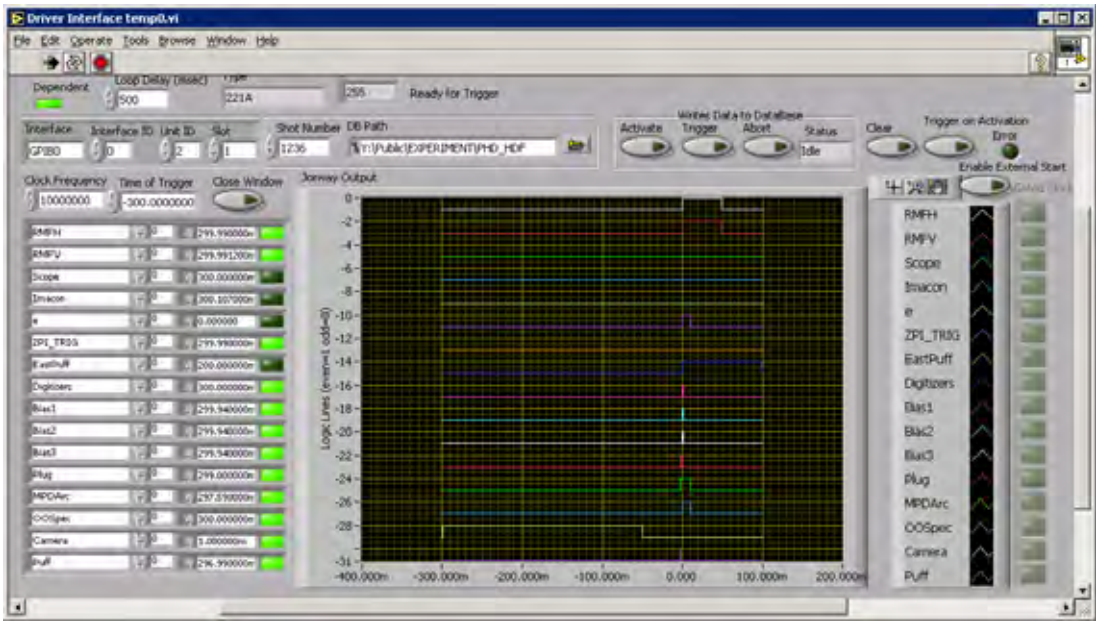

Figure 3.11.2.2. Jorway 221A Driver. 16 channel, 512K memory, ≤10 MHz, arbitrary digital waveform generator (CAMAC).

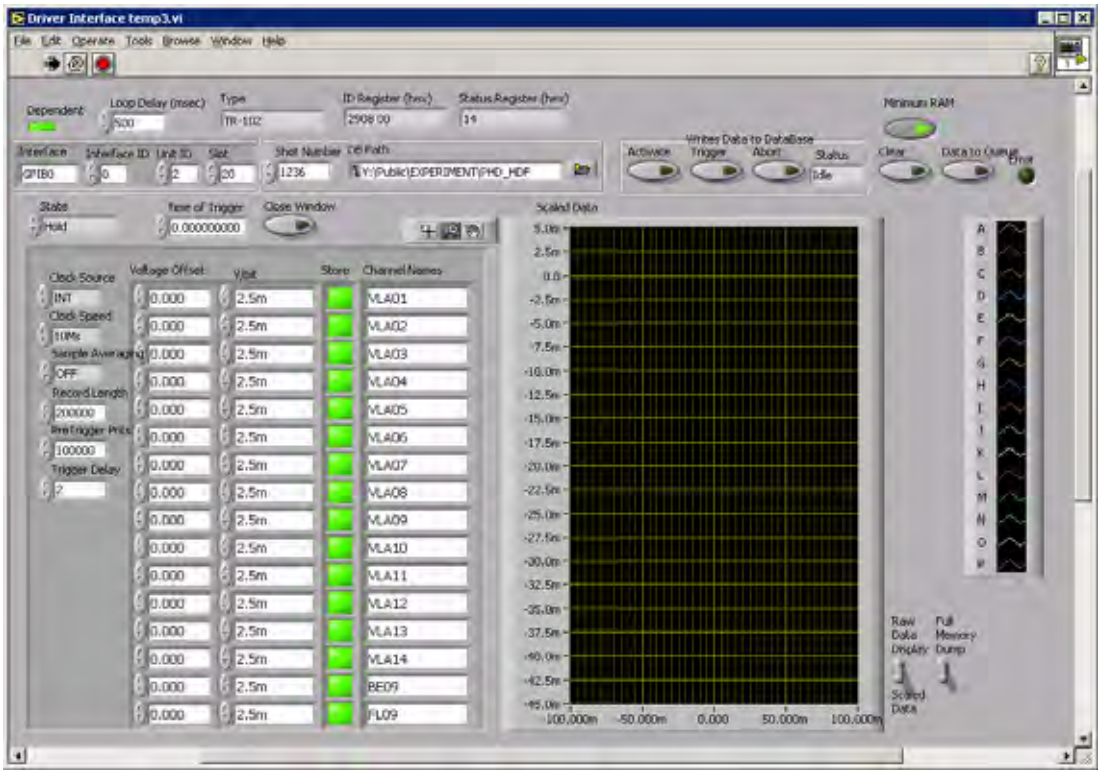

Figure 3.11.2.3. Joerger Model TR Series Driver. ≤16 channel, 12 bit, ≤40 MHz, ≤512 KSamples each channel, Transient CAMAC digitizer.

Figure 3.11.2.4. D-TACQ Solutions ACQ196 Series. ≤96 channel, 16 bit, ≤500 KHz, ≤ 1 GB RAM total, Transient (or continuous with driver modifications) digitizer. Ethernet interface. Note that this hardware requires significant set up accomplished outside Labview.

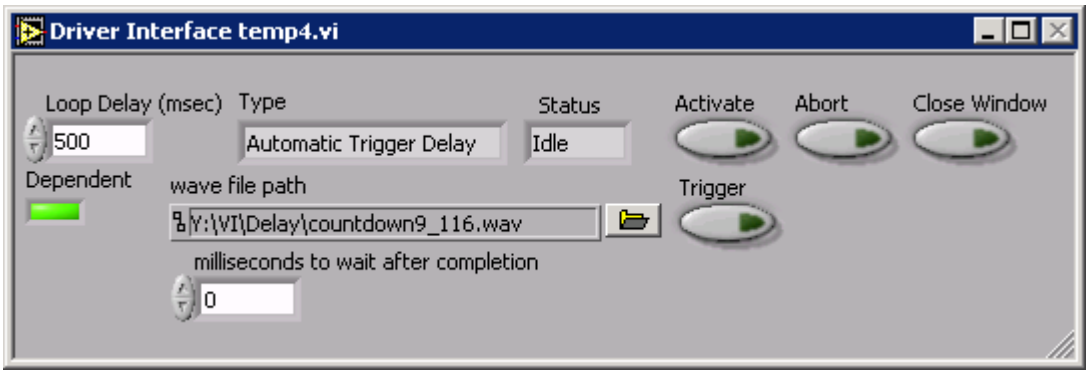

Figure 3.11.2.5. Delay and Audible Countdown Driver. Requires a system sound card.

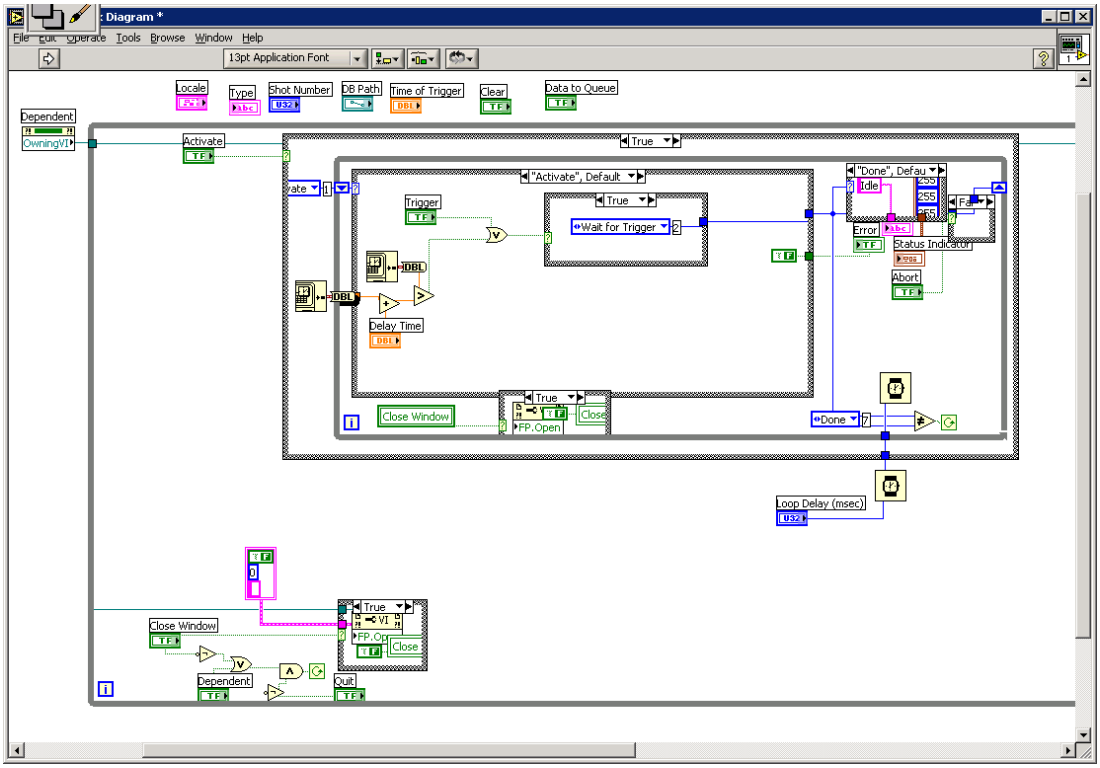

Figure 3.11.2.6. Delay and Audible Countdown Driver Diagram. Not much is required to make drivers operable with "Driver Interface".

The drivers available may not be sufficiently complete for all uses. Nevertheless, it is relatively simple to implement new drivers via the copy and hack approach (without having to understand how things work). The usual approach is to copy, run in debug, and then hack. It is easiest when copying a driver with a close modus operandi. It typically takes this author between one hour and a day to fully implement a new driver. More time is required with hardware that does not adhere to the usual conventions and approaches (e.g. when the clock frequency returned by hardware is rounded to a couple of significant figures without specifying this in the manual).

3.11.3 For Hardware Developers

As a practical matter, we have actually come to prefer CAMAC for its simplicity and easy availability. However, the ideal approach, in our opinion, is to use Ethernet as a communications interface for each digitizer. We have been developing several custom Ethernet connected microcontroller devices. Numerous books exist (with step by step instructions) on the subject. The real costs are in the first time development (not in parts).

Security is not required; in fact we discourage it in most applications since these devices can be readily placed on a secondary independent network or fire-walled. Adding security adds numerous complexities in implementation (unless using third-party libraries) and numerous headaches for users trying to get hardware to work for the first time. It is easier to fix a separate hardware firewall's security hole (used by millions) than a digitizer's (used by hundreds). It is also easier to VPN through a familiar hardware firewall than to deal with multiple manufacturers' differing security implementations. This applies to Linux digitizer implementations as well as custom. A simple 'recognition code' can be used in addition to IP address to prevent accidental interference, full security implementations are not required by us. At the least, the option to disable security should exist. Digitizers should be "plug and play". Similar complaints apply to certain oscilloscope network implementations. An example here is dueling VISA implementations. The natural result of standard wars is mud on all faces.

3.12 Matlab R14SP2

Matlab is a single threaded interpreter. The Matlab language provides a huge number of instructions that covers more than just matrix operations while also providing the ability to add more instructions through the use of external (user compiled) libraries. The result is an interpreter that is capable of approaching compiled (with full optimization) C performance but is much easier to use. Matrix operations can often be optimized with special command sets on Intel CPUs. Matlab© is capable of taking advantage of the hardware as developers add support; we gain the benefits without doing complex programming. Matlab is a simple language to program and debug with many useful GUI features. It is very well suited for data analysis and display.

The chief downside is its intentionally single-threaded nature. This means that a complex analysis and simulation can't be run while continuing to look at other data (without implementing cooperative multitasking in user code; requiring comprehensively careful programming). It also means that it is incapable of taking full advantage of multi-processor computers. Multithreading for large array matrix operations may still be possible through the external libraries interface. The second structural flaw is probably due to its Java base. Total memory allocation is limited to 2 GB which graphs (with many data points) are quickly able to consume due to inefficient memory usage. Earlier versions of Matlab could be operated via "Remote Desktop"; however, its license manager will not allow Matlab 7.0.4 to be started from

"Remote Desktop". They have claimed this is due to support issues. It does crash more frequently while being used remotely to display graphics. Matlab 7.0.4 can still be started remotely by using programs like RealVNC (http://www.realvnc.com) immediately after the remote computer restarts. More recent versions of Matlab (R2006b) can be run remotely without difficulty (according to MathWorks website). In addition to these architectural deficiencies, there are many known bugs capable of crashing Matlab©. Bug fixes come out on a regular basis (but are not freely distributed). For example, flawed behavior often ensues when reaching the memory limitation while displaying graphics. It can become impossible to close figures (to free memory) that run out of memory while displaying. Obviously, such flaws make it a poor choice for a control system, especially if the user is permitted to run new code along with the controls.

Most of the Matlab analysis utilities we have started out as a side project by Dr. Richard Milroy. They were originally intended to allow access and simple data display in Matlab from MDS/MDS+ database servers. The HDF5 (Hierarchical Data Format) interface routines (ActiveX) were intended to be used as a local data cache desirable because of the poor MDS server performance encountered at that time. Matlab 7.0.4 has its own internal HDF5 routines, but in this revision one of the critical routines crashes Matlab under normal circumstances (the fix to this came out in a later service pack, that we have not yet acquired). Ultimately, if Apple's OS needs to be supported, the internal Matlab HDF routines should be used.

Dr. Milroy's original routines allowed for HDF5 database access along with basic 2D plotting, frequency filtering and plotting, and some localized Redmond Plasma Physics Laboratory FRC analysis routines. The "plg.m" GUI plotting routine was the most widely used while many of the other command line routines were forgotten by this author. The routines have descriptions in "ExptAnalysis.chm". This illustrates that for others to make use of valuable programs, there needs to be a structure that naturally leads to their (repeated) discovery. The routines' (and signals') purposes should also be unambiguously apparent from their names rather than the shorthand abbreviations we all tend to revert to. These are hard rules to enforce (that this author tends to break when time is limited) but they can make software tools significantly more valuable.

Since the original routines were written many changes have been made. Our changes to Dr. Milroy's routines have mostly focused on features, adding:

1) (DatAcqSvr) Effective lossless compression (~ 30:1 typical for our long data sets) to the ActiveX HDF5 server.
2) (DatAcqSvr) More effective time array compression. Original routine did not take into account the effects of subtracting two very large numbers in deciding whether time is evenly spaced. The time array is of float (single precision, 4 bytes, 23 bits mantissa) type. This means that the maximum deviation for an evenly spaced array of local Δt from global Δt is < (maximum t around differenced points) * (1/2^23). This knowledge can lead to a proper fix to the time array problem at some CPU cost. It is also best case (i.e., careful time array generation from digitizers assumed). **The actually implemented fix**

works best when there are always some points near time = 0. For it to detect an uneven array (Δt(t=0) – (t(end)-t(start))/#points) / ((t(end)-t(start))/#points) > 1E-3 is the condition. This has the practical effect of treating arrays with constant dt, but large time offsets as if they were uneven (denying a factor ~2 overall compression). It was implemented this way to save on CPU time while still compressing most of our common time arrays. The logical fix would be to subtract a time offset and include this in that data file. This has not been implemented in order to maintain full backwards compatibility (without extensive testing). Instead we choose our experimental time offset so that digitizers record near t=0.

3) (DatAcqSvr) Read only capabilities to the HDF server. This allows plotting data from read-only devices such as CD and DVD. It also allows reasonable restrictions to be set on file servers for general access. This has greatly cut down on corruptions of the database.
4) (Matlab) The ability to run complex programs "on the fly" to calibrate (and display) signals through the standard DatExpr() data recall method. The complex program that is run is dependent on shot number.
5) (Matlab) The ability to run simple programs as data recall. For example, this allows plotting "mean(diff(t)).*cumsum(Power)" through "Plot Control" (or DatExpr()) to plot energy. It can also be used to easily plot charge transfer:

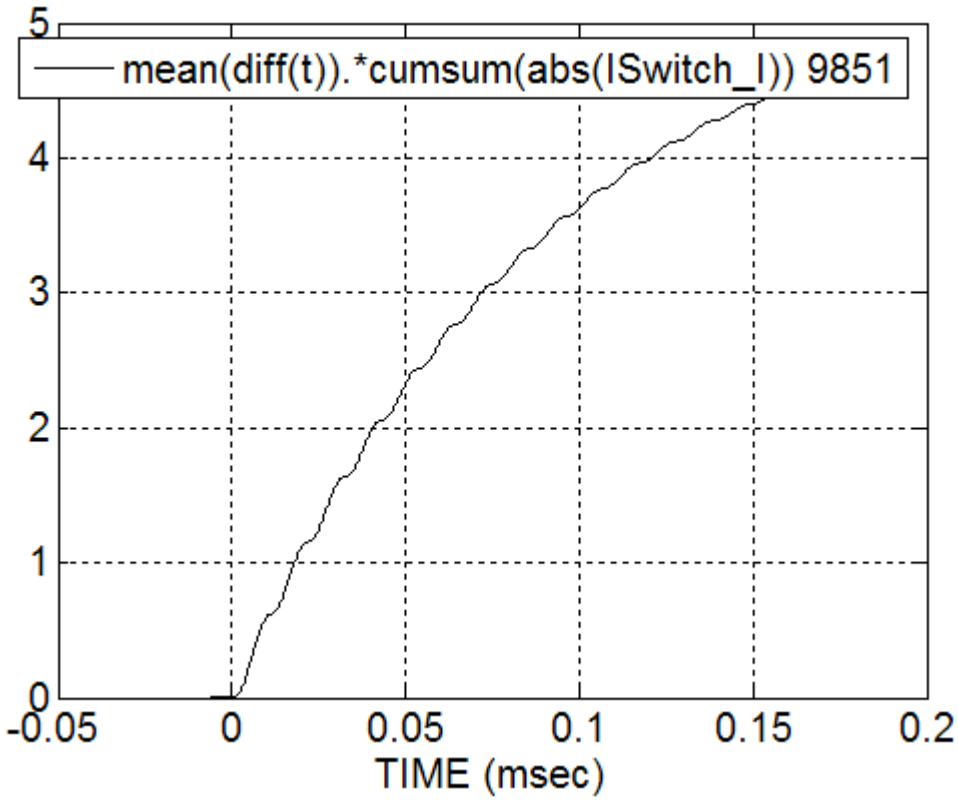

Figure 3.12.1. Simple program plotting. This displays total charge transfer through Device Under Test (DUT).

6) (Matlab) Relative timebases. Time can be plotted relative to any signal dependent on individual preference.

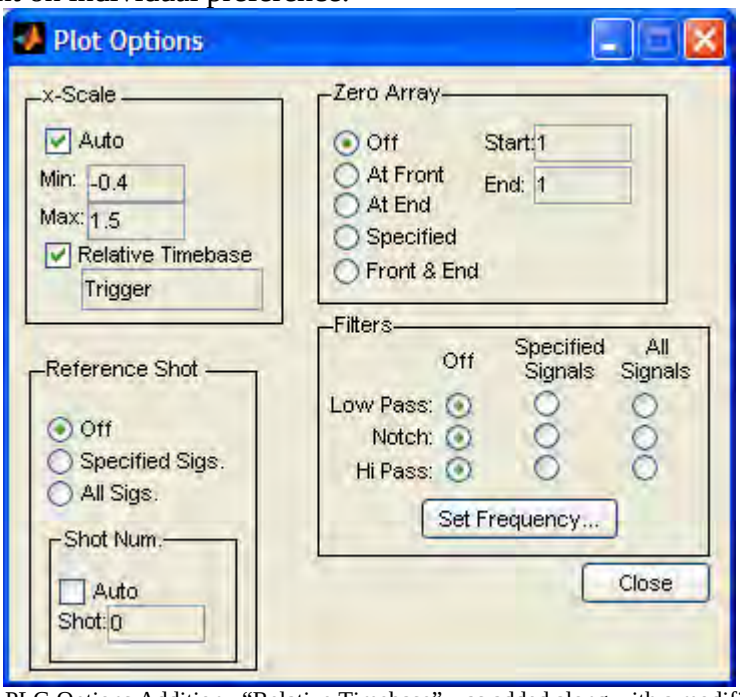

Figure 3.12.2. PLG Options Addition. "Relative Timebase" was added along with a modification to make "Specified" function in units of milliseconds rather than array index.

7) (Matlab) Signal name listing. "Items:" can be clicked to display any signals or calibrated items that can be plotted.

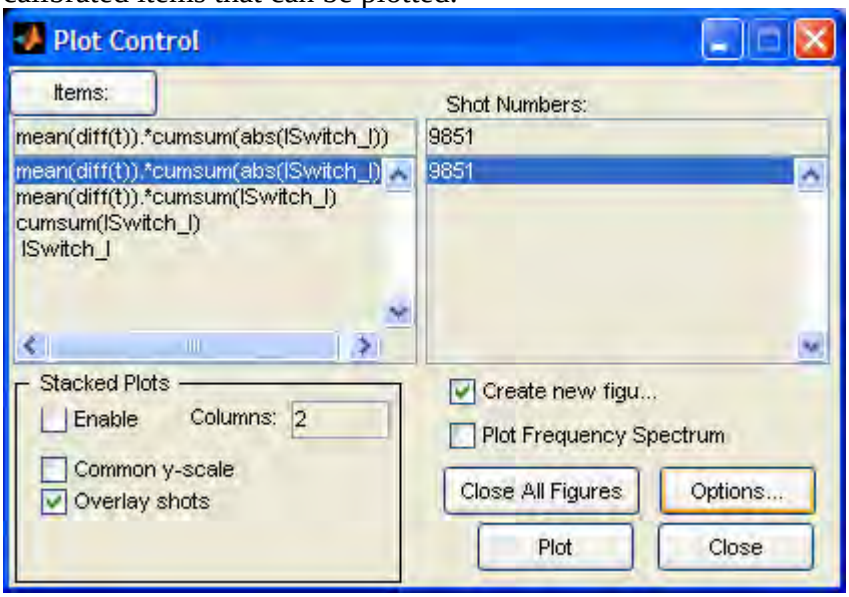

Figure 3.12.3. PLG Additions. "Items:" has been changed to display possible signal choices (rather than requiring pre-requisite knowledge of the names).

8) (Matlab) A tracing feature. When mdsVar.Tracking is defined this field reflects the calls made and the length of data returned from the database. Enabling this generates a record of all calibration and data calls made during data access. This is useful for understanding network performance issues (for example: database access over slow fully-encrypted VPN) and understanding what routines are being accessed (for dependency control) in more complex programs.
9) (Matlab) Self-installation. The initial installation (for DatAcqSvr and Matlab © routines) is the same as everyday running. The startup routine sets up all the appropriate paths plus variables and updates the DatAcqSvr ActiveX control if necessary. It also implements automatic fixes to recover from user corruption of the preferences data structures. Windows Vista has partially broken this feature; DatAcqSvr must be installed outside Matlab.
10) (Matlab) Memory resident data caching. Running complex programs as data access allows for the minimization of required hard drive space at the expense of processor time. For some very complex programs this approach makes less sense, yet the desire to use minimum hard drive space is still paramount. Higher level programs call many of the same lower level programs multiple times while accessing different related signals. For example, the excluded flux signals are dependent on magnetic field and flux loop calibrated signals. When an upper level program calls for excluded flux and external magnetic field some of the same calibration programs get run multiple times with the same results. The memory resident cache simply provides the previous run's solutions. A hash is made from the mdsVar settings (for example, active filters) so that results run with different settings are cached separately. Due to the current structure of the Matlab© implementation, it is not simple for the cache to realize when the mdsVar settings actually matter. The conservative approach has been taken to ensure accurate results. The cache treats raw data the same as processed data. Changing this could result in network VPN speed improvements as used in practice. mdsVar.CacheSize sets the cache size limit in bytes. It is common to run with very large caches (up to 500 MB) due to the lack of memory compression on the cache.

Effective compression has made it possible to record long time signals at high sampling rates with nearly no storage cost. It has brought our database down to a size where it can be backed up on one DVD. We can also afford to keep old data accessible essentially forever. Features 3 & 4 allow for segmentation and protection of base (permanent) collected data from analysis and calibrations. Any user can develop their own calibrations and analysis routines without affecting other users by changing a directory entry in the preferences structure; but this still allows easy "publication" when the routines are complete. Incidentally, a no frills Windows XP box plays the role of "server"; requiring no dedicated personnel. All computational costs are incurred by the client machines. The use of a RAM-based data cache can speed up access through VPNs (and complex processing). It was originally hoped (in

vain) that the client OS would automatically perform network disk access data caching.

Additional general plotting routines have been added. Later these routines should be integrated into the 2D plotting GUI. A generic 3D plotting option should be added as well. This list, and the lists to follow, are not complete listings.

1) WindowedFFTgui. This interface allows for plotting the time variance of the frequency components in signals (spectrum analyzer). There is currently only one implemented windowing option (Hanning) but more can easily be added. The GUI does not have all the convenience functionality of the standard 2D plotting GUI. Such features would be highly redundant when it is planned to integrate windowed fft into the standard 2D GUI.

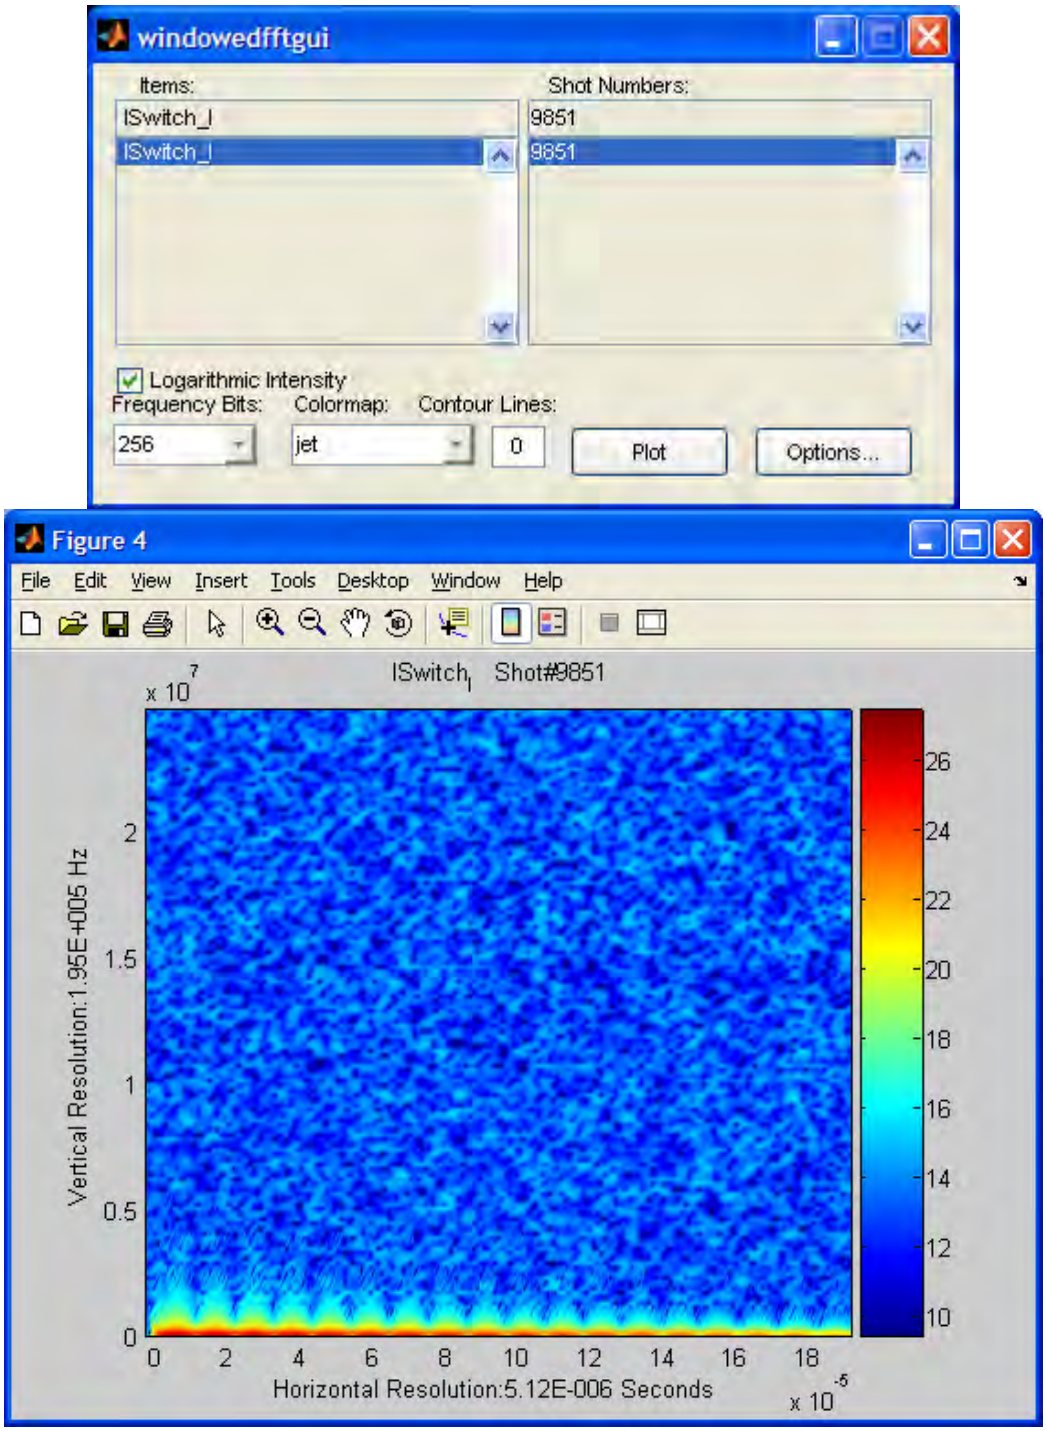

Figure 3.12.4. WindowedFFTGUI. An example of the functionality of windowedfftgui is shown using the distributed package of the routines. The speckle at high frequencies is not significant. More interesting examples can be seen in the section describing the capacitive manometer implementation. 'Bits' in this context is proportional to the number of frequency bins.

2) WindowedFFT2Dgui. This GUI allows for the plotting of spectral energy in one frequency band as a function of signal name and time. It has proven useful for plotting wave penetration information gleaned from internal probes (in other experiments).

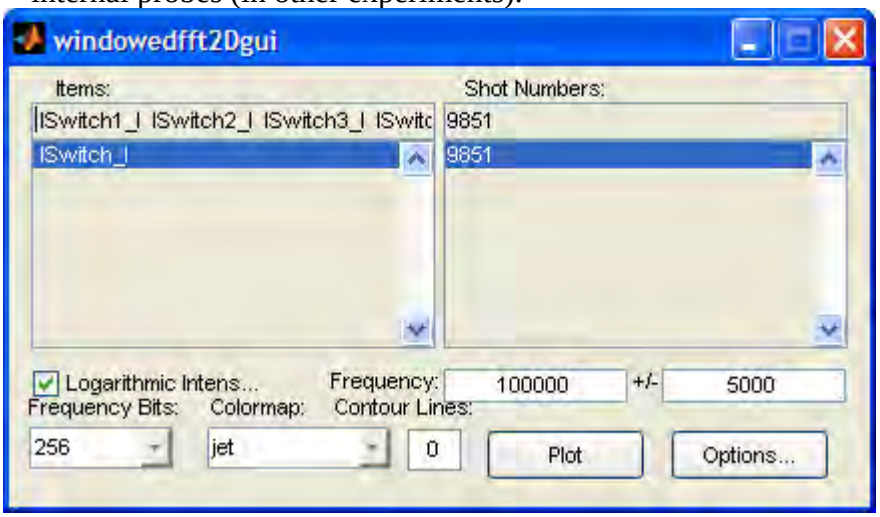

Figure 3.12.5. WindowedFFT2DGUI. As only one element is added to the GUI it should be clear why it should be simple to integrate all the preceding interfaces into one.

3) Logbook. While we keep a separate, user defined, spreadsheet as a running logbook, it has proven useful to be able to fill a "logbook" with computed values utilizing the HDF database (in addition to the Excel spreadsheet). This is in essence a 0D version of the 2D plotting GUIs. Similarly, it is capable of using complex (user defined) routines to process data. It allows for exporting the generated data into Excel or Matlab routines. It also keeps a running cache of past calculations.

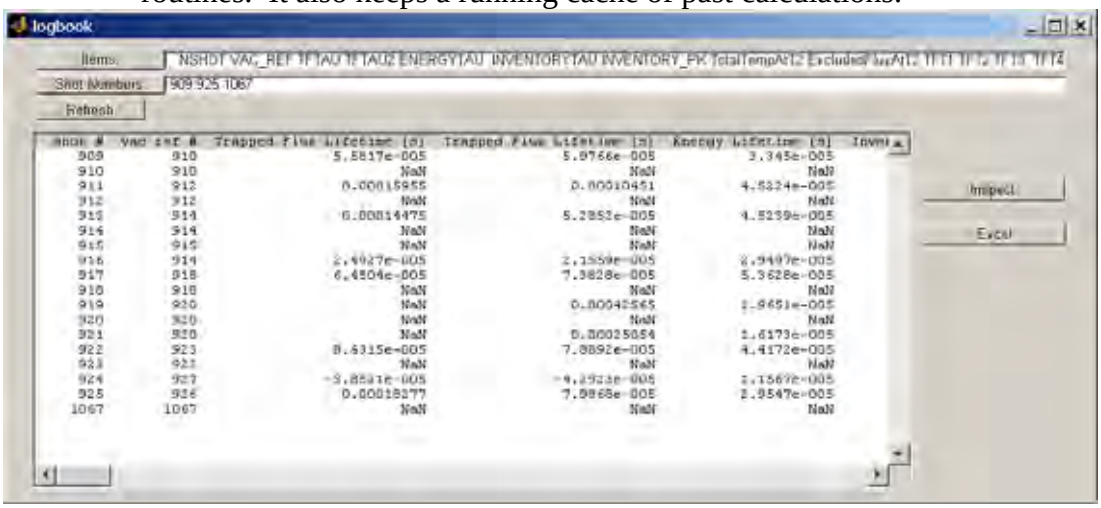

Figure 3.12.6. Logbook. Logbook is designed to provide a spreadsheet style interface to 0D data. It incorporates the functionality of Excel by allowing export of the dynamically generated data. It is possible to incorporate this functionality into the standard 2D plotting GUI. It is simple to access data from a spreadsheet logbook and display it (as demonstrated here).

4) FormatFig. FormatFig is a command line function that sets all fonts in the targeted figure to the specified font and size.
   Form: formatfig(FontName,FontSize,handle)
   Example: formatfig('Ariel',20) OR formatfig('Ariel',20,gcf)
5) SetSize. SetSize prepares a figure for copying to clipboard (or making movies) by setting the size to what is specified and positioning the figure to maximize visible area.
   Form: SetSize(width,height)
   Example: SetSize(1024,768)
6) ImagePlot. Imageplot grabs frames of figures (for making movies) without requiring that the figure be onscreen or unobstructed.
   Form: Image = imageplot(HandleFigure,… [width height])
   Example: Image = imageplot(gcf);
7) Cmap. Cmap makes it easy to apply any color maps to surfaces and figures. It is easily expanded to allow the definition of new color maps (without potential conflicts with .m files). It is also used to set text and background colors.
   Form: cmap(sChoice,sBackground)
   Example: cmap('BC','White') OR cmap('BC','Black')

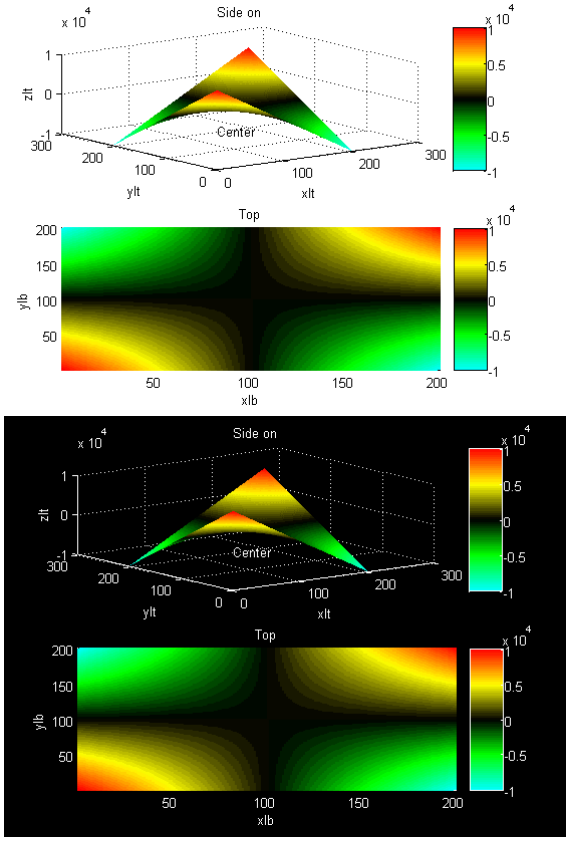

Figure 3.12.7. Cmap options for plots.

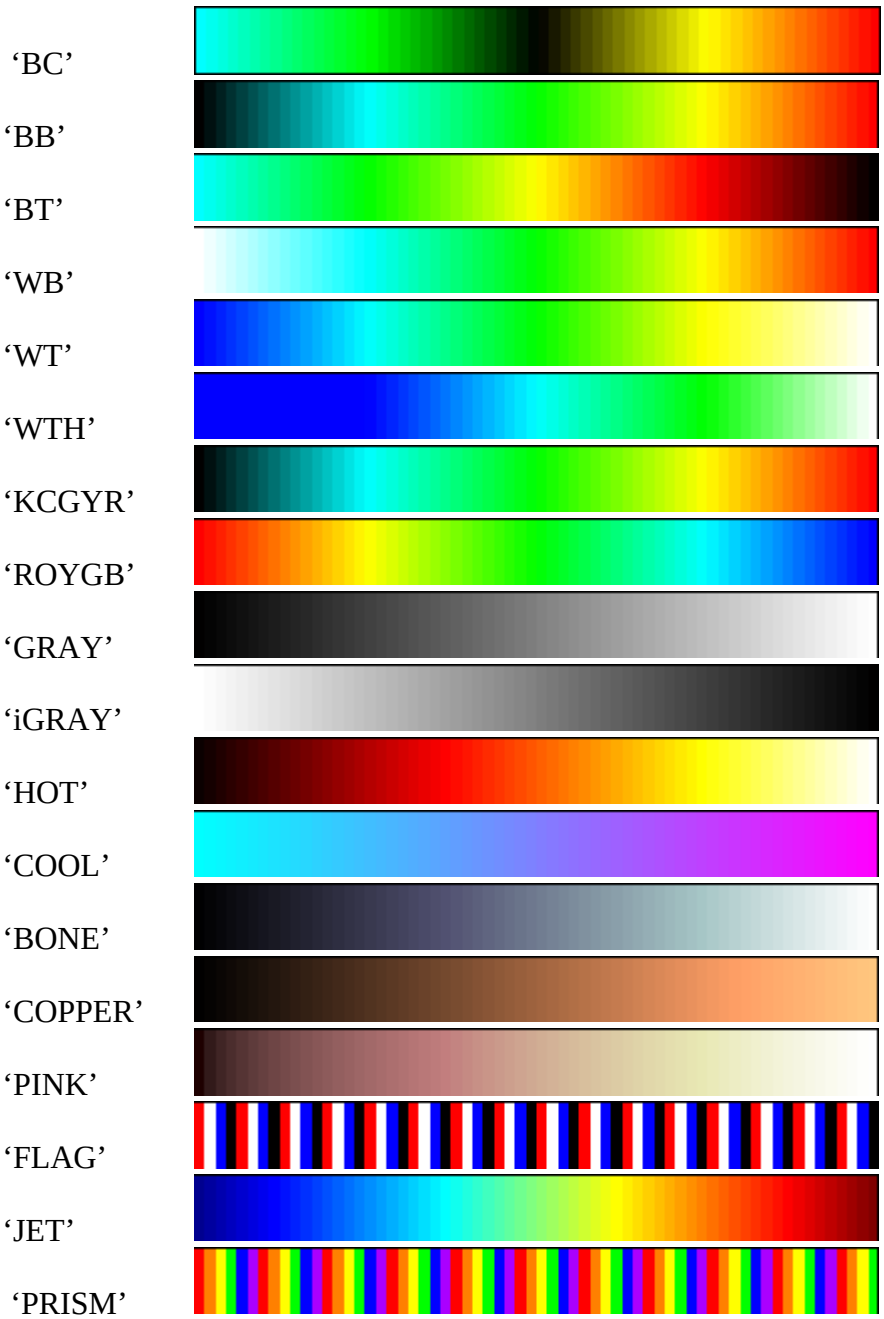

Figure 3.12.8. Cmap color options (abridged).

Specialized analysis routines:

1) SpectroID. SpectroID makes use of the NIST line database to plot and analyze spectrometer data. It includes functionality that helps suggest possible lines and identifies elements which are likely to be present when the spectrum is taken as a whole. It uses error tolerant routines so that element spectra that are missing some lines are not immediately eliminated as candidates.

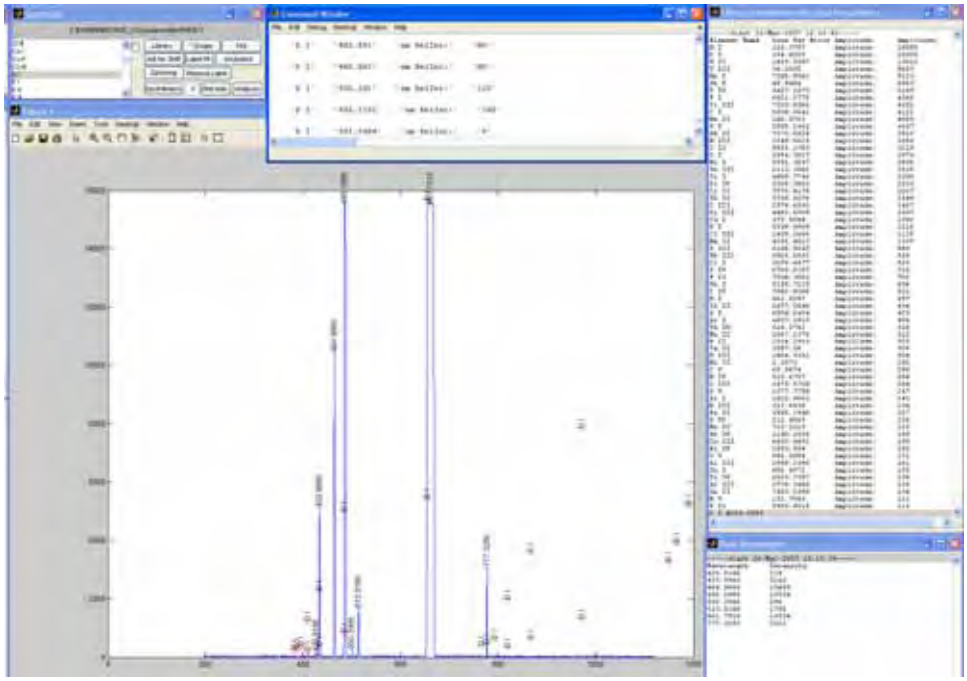

Figure 3.12.9. Plot of an experimentally observed spectra. A list of potentially present elements can be seen in the upper right. Obviously, some interpretation is required; this cannot be a fully automated routine due to the need for auxiliary information (such as the material composition of parts in the vacuum chamber). The downloaded NIST database can be edited to remove elements not present in the system.

2) RGA. RGA is a utility for analyzing mass spectrometer data. Similar to SpectroID in function, but much simpler to use due to the more limited range of recognized possibilities.

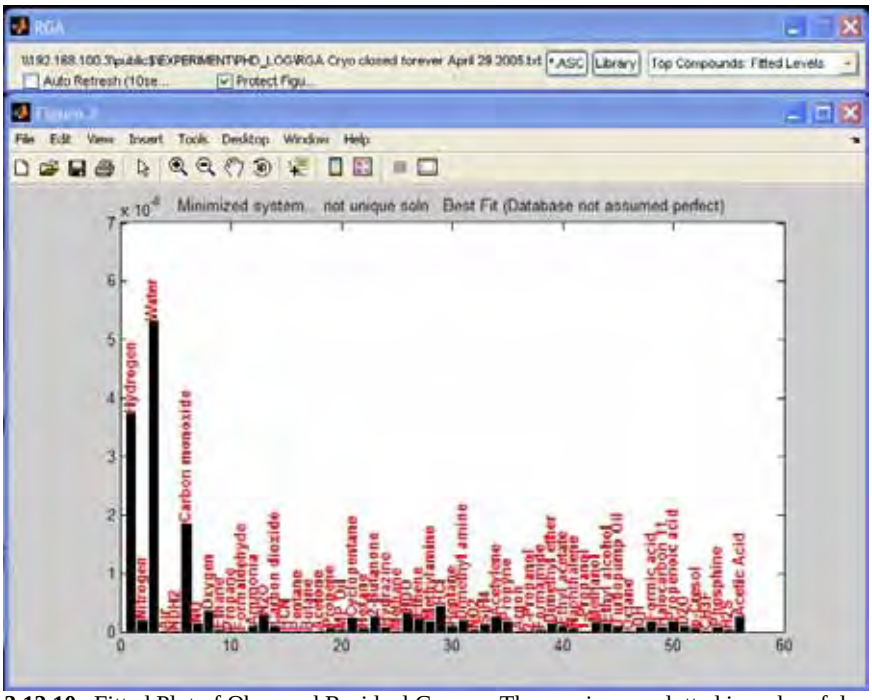

Figure 3.12.10. Fitted Plot of Observed Residual Gasses. The species are plotted in order of descending maximum possible pressure. Note that a pressure dependent mass cracking fraction has not yet been implemented.

3) ImageAnal. ImageAnal is a utility intended to aid in analyzing colored imaging data. It can be useful in tasks such as calibrating the acceptance angle of bolometers or roughly estimating parameters as a function of region.

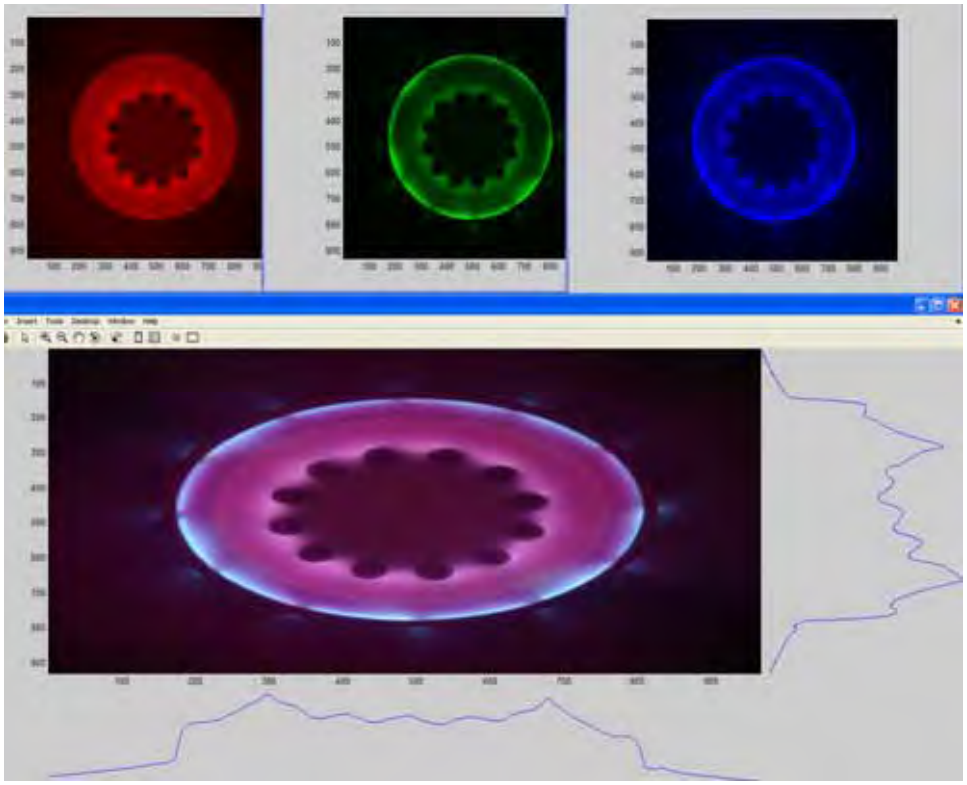

Figure 3.12.11. Decomposed images of a plasma source.

Utility routines:

1) csvdat. Comma Separated Value file to HDF database converter.
2) scopegui. Automated download to HDF (or MDS) utility for use with the Tektronics 3xxx series of oscilloscopes. Due to the Tektronics implementation, additional software drivers are required. The time that a download takes is highly dependant on the driver version used (this has resulted in serious problems running test setups since Tektronics has their own version of VISA). National Instruments' Labview uses their own VISA driver (the two do not coexist). Note that using the http:// address to access Tektronic's TDS3034B scopes requires internet access; it loads information from Tektronic's website.
3) rgb2n. Code designed to substitute for the most commonly used feature of the Matlab image processing toolbox. Used for converting RGB images to indexed colors. This is needed for some avi codecs.
4) saveall/loadall. These routines make it quick and painless to shut down and restart Matlab with the data-space intact (including figures). Used regularly it also makes it easier to recover from Matlab crashes.

Most of our non-automated settings are kept in an Excel spreadsheet that acts as a logbook for the running of the machine. This can be accessed through calibration routines for plotting. Other sources not included in the database can be (and have

been) accessed this way as well (including images and movies). This allows for a coherent way to present data (from different experiments or diagnostics) to a user. The required format should be apparent from examining the examples.

4
Power Systems

What are the characteristics of a good electrical system design? First and foremost it must be capable of accomplishing the required task. Second, the design must be reliable. Third, it must be cheap to build and design (in the case of small unit counts). The engineering effort's focus tends to progress through these criteria in order. First a prototype is built that should be capable of performing the required function. Next, flaws are found in the initial prototype that affect reliability and survivability. Once a reliable design is found efforts are eventually made to make it cheaper in production. Many factors contribute to a good system design.

The greater the foresight and experience available before making the initial (or second) prototype, the less likely it is that significant reliability problems will be encountered. The most reliable way to have difficulties is to not follow manufacturers' recommendations and specifications (this is especially true for old designs using newer components). Many components with the same part numbers are produced by different manufacturers, not always with the same 'beyond specification' characteristics. A good portion of this chapter is devoted to presenting some of the fruits of this experience.

The fact that our systems require exceptionally high voltages, high currents, and very short timescales add additional knowledge requirements that are not commonly present in electrical engineers. Some of these requirements are present in aerospace applications. As such, design guides are available and some of their key points will be summarized.

The primary non-diagnostic systems are discussed in this chapter. In many cases this includes discussion of their difficulties and failings as well as their design. Many of these systems are constantly evolving. It is important to remember that the duration of the 80 cm experiment discussed here spanned a little more than four months. The time between the last shots on the 40 cm experiment and the start of firing on the 80 cm experiment was also little more than four months. The final figure of this chapter drives home this point of constant change in most of the primary systems. The pace of change, and hopefully progress, is undeniably rapid for an understaffed, underfunded experiment. Obviously, this can not come without a price.

4.1 Modularized Systems

It is not uncommon in electronics diagnosis to encounter cascaded failures. Some part fails, and causes many other parts to fail in turn. Replacing only the first failed part found often results in another destroyed part. While this is merely a nuisance for digital electronics, it can become quite expensive in power applications. It is highly desirable to engineer systems such that relatively common fault modes do not result in cascaded damage. First and foremost this requires knowledge of which parts are most likely to fail and in what way.

It is well known that ignitron switches if improperly conditioned or maintained (or are just new) can fail in a 'glow' mode where they conduct significant currents while maintaining high voltage across them. This failure, in and of itself, is not intrinsically fatal to the switch. However, if the switch is allowed to reach too high a temperature while in this mode, out-gassing of the components may occur resulting in a higher internal pressure of the switch and a persistent degradation in voltage hold-off (not to mention a potential fire hazard).

Similarly, the ignitron can fail in a fully conducting way. In this case the power is dissipated in the charge/dump resistors.

The most reliable way to remove the possibility of collateral damage from these failures is to limit the high voltage supply's output power to levels that the resistors, switch, and capacitors can safely dissipate given their size. Note the qualifying term: 'most reliable'. Other valid approaches exist. Copper vapor pseudo-spark switches can be expected to behave similarly to ignitrons.

Solutions other than limiting power supply power do exist, such as having a computer program monitor the charge state of the banks and discontinue if anything unexpected should occur. Rather than the power dissipation of the components, the critical factor becomes the energy they can safely absorb while waiting for the computer program's response to a fault condition. This approach can become an impediment when a significant fraction of the ignitrons misbehave but are still capable of being 'burned through' and made to function.

The inherent price of a computer-monitored approach is that the complexity is higher and thus there are many ways it can fail that may not be foreseen (some of which may not be under the direct control of the end programmer). The failure modes that immediately leap to mind are viruses, other programs taking processor time, hardware errors in the setup of the monitoring system, unexpected computer hangs due to electrical noise, memory leaks degrading performance (or resulting in a swap file size increase during which the program may temporarily hang), programs conflicting with programs (for example trying to access a separate port on the same controlling hardware), and 'corner' cases. By 'corner' cases we mean cases where two parameters are at extremes (the usual connotation is the rare coincidence of two events causing problems that alone would not result in failure). Also, we must emphasize that it is never a wise practice for the developer and tester to be the same person, with no further testing, no matter how good and/or experienced they are. This presents an additional difficulty due to the limited number of highly-trained personnel available.

Most of the inherent problems can be overcome with careful programming and hardware design. The use of hardware watchdog timers, safe mode settings on power reset, and other industry recognized good practices help to minimize serious flaws. The time and expertise required to develop and comprehensively test the system should not be underestimated. This is a non-optimal solution for our design because of the time required and the existence of a relatively simple solution (limit power supply power).

Additionally, it is desirable to construct banks in a segmented way such that one component is never subjected to the energy stored in the entire bank. Isolation is required from one segment (module) to another. We typically accomplish this by having each capacitor/switch/power supply power its own individual coil. It should be realized that it is still possible for energy to couple from segment to segment via magnetic coupling (but to a degree limited by coupling and the uncoupled inductances of the system). There are also links between the segments via the shared bias banks.

This modularity has a side benefit of making it easy to replace components from spare stocks. As we have close to 30 supplies, having 3 spares is a negligible increase in cost. Modularity also greatly simplifies the task of adding new banks or running different segments at different voltages. Scaling is simple. It also makes it dangerously easy to set up new experiments based on the spare hardware. Experiments utilizing this approach have a tendency to become Thanksgiving-day turkeys if allowed to sit for too long unguarded.

The price of this highly modular approach is that the effect of individual failure rates is magnified. Looked at from a different perspective, one has ample opportunity to discover and diagnose reliability problems due to basic design flaws (rather than it remaining a 'black' and frustrating art). For example, limping along with expensive commercial supplies failing for uncertain reasons (a common situation in systems where the supply is nominally disconnected from its multiple banks before firing) has been quite frustrating in the past (especially once it was realized that some supplies had failed while their outputs were disconnected, AC powered via surge protectors, but just sitting). With the critical designs subjected to high voltage and/or high current conditions existing in-house we have been able to steadily improve reliability. We believe that most of our occasional problems (in the electronics) are now due to mechanical abuse and assembly errors. It should be added for balance that we have had some good luck with a few manufacturers of commercial high voltage supplies; the temptation to just buy something has persisted.

Modularizing systems lowers the costs associated with mistakes. It reduces the damage caused by equipment failure. It results in systems that can easily be used on other experiments. Ultimately it results in more reliable systems on an individual basis.

4.2 High Voltage Practices

If something unexpected happens in a system our advice is to turn the power off (via an insulated switch, rod, or fiber optic), then wait and think about what might have happened. Some significant time should be spent on contemplation of the possibilities. Never rush in to an enclosure. Standard HVAC systems are made to safely dissipate their stored energies in about 5 minutes. In lab settings this practice isn't always followed. Always assume everything is hot until proven otherwise as multiple (or unanticipated failures) can result in unexpected high voltages. Never allow yourself to be in contact (or near to contact) with anything metallic connected to 'inside the enclosure' parts of an operating experiment. A perfect example of this

is a monitoring coax looking at power supply voltage running outside the experiment's enclosure. Ask yourself where the currents can flow if the monitoring resistor should arc over, or the supply's 'ground' turns out to be floating (in an inductive AC sense, or resistively).

Figure 4.2.1. Commercial high voltage probes can fail. (left) The forward section of a high voltage probe (red) after an explosion occurred. (right) An example of an intact probe connected in a proper manner (so long as the power supply case is grounded). The majority of the rest of the probe ended up as a dust cloud. The probe was placed with its plastic handle too close to a high voltage plate (the plastic was slightly conductive). With the policy of not allowing any non-grounded connections to penetrate the experimental enclosure this was simply a very loud nuisance.

An abundance of caution and over-insulation are the two most critical requirements for successfully dealing with high voltage systems (frequently defined as anything over 1 kV). This author has personally created sustained arcs (seconds) several feet long through air using the module's standard 300 watt (25 kV) power supplies during failure mode testing. In this particular case an ignitron was undergoing a glow discharge with a small capacitor across the supply (~ 10 nF) followed by disconnection of the wire supplying the ignitron. This resulted in an impressively long and loud arc until power was removed. While close contact (typically within a couple inches) is usually required to start arcs directly they can also be started through the mechanical failure of components or connections and tracking over insulators for shockingly long distances (10 kV/in is not a safe assumption under certain circumstances). A similar event happened on the experiment with a disconnected high voltage lead (the plug was not pushed in to the receptacle far enough to make contact). The supply was capable of charging the module by breaking down the gap along the surface of the connector. This resulted in a rather melted and charred connector. No guarantee it was capable of discharging the remaining voltages after the shot.

The rule of thumb for tracking distance (10 kV/in) can fail in the presence of corona or high dielectric constant surfaces[22]. It is easy to cause corona by using thin high breakdown strength insulation (without punch-through) immersed in air.

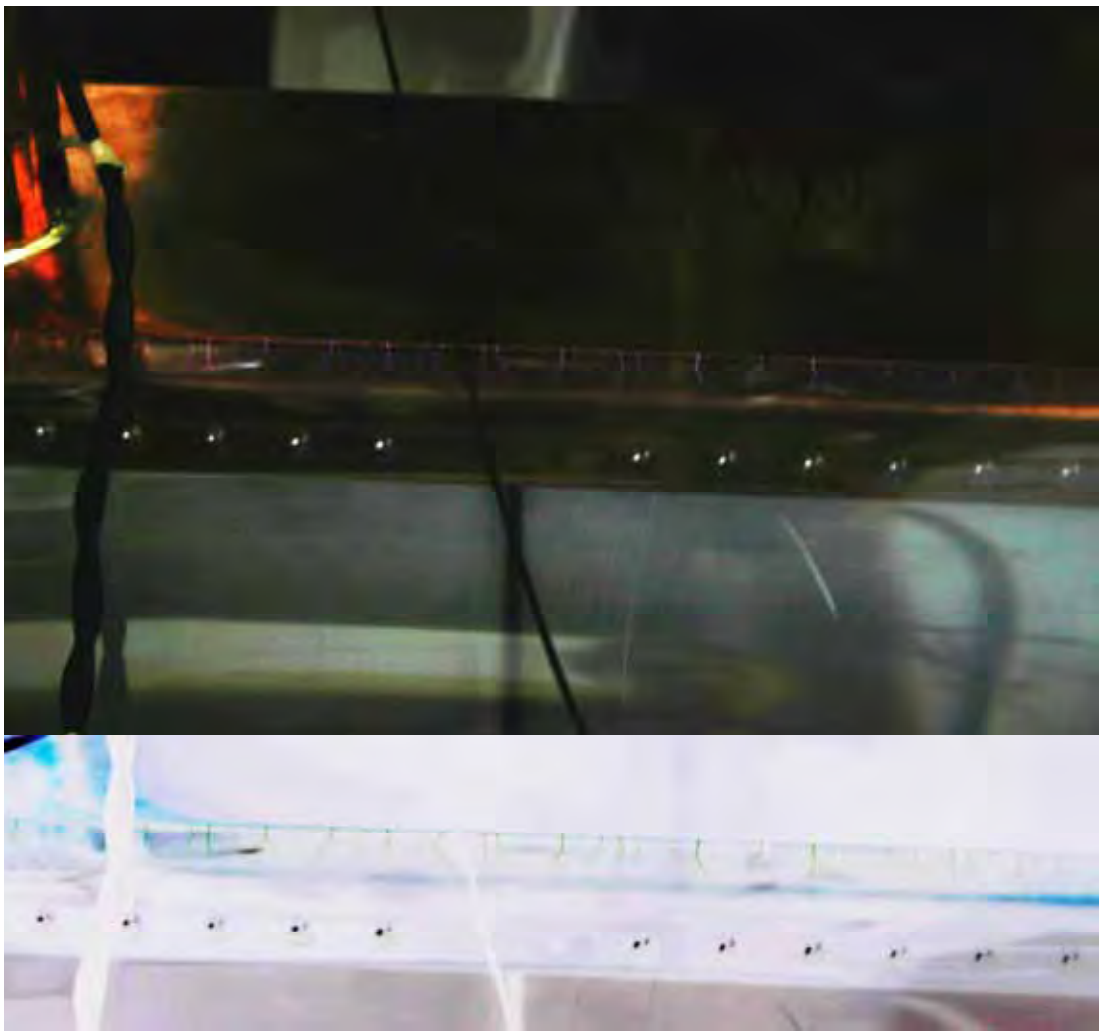

Figure 4.2.2. This is an example of corona caused by an edge and thin insulation. The capacitors are rated to 30 kV. A cropped negative is shown below the original image. The crackling of static electricity moving around could be heard leading us to take this exposure.

Some might make the mistake of assuming that the above mentioned power supply rating (300 watts, 25 kV) might limit the currents to safe levels, it does not. We relate this to help dampen down some of the feelings of over comfortableness. Our systems are designed to have multiple safety redundancies, but most of these can be circumvented. We also note that the National Electric Code §110.34 would require 5 ft of workspace even with a 25 kV system that is 'effectively guarded by insulating materials' (with no attention paid to power). The National Electric Code does not give the reasons for the rules but it isn't hard to postulate multiple reasons that are equally applicable to our experiments as to standard HVAC systems. A historical search of suggested revisions would be necessary to determine the original reasons.

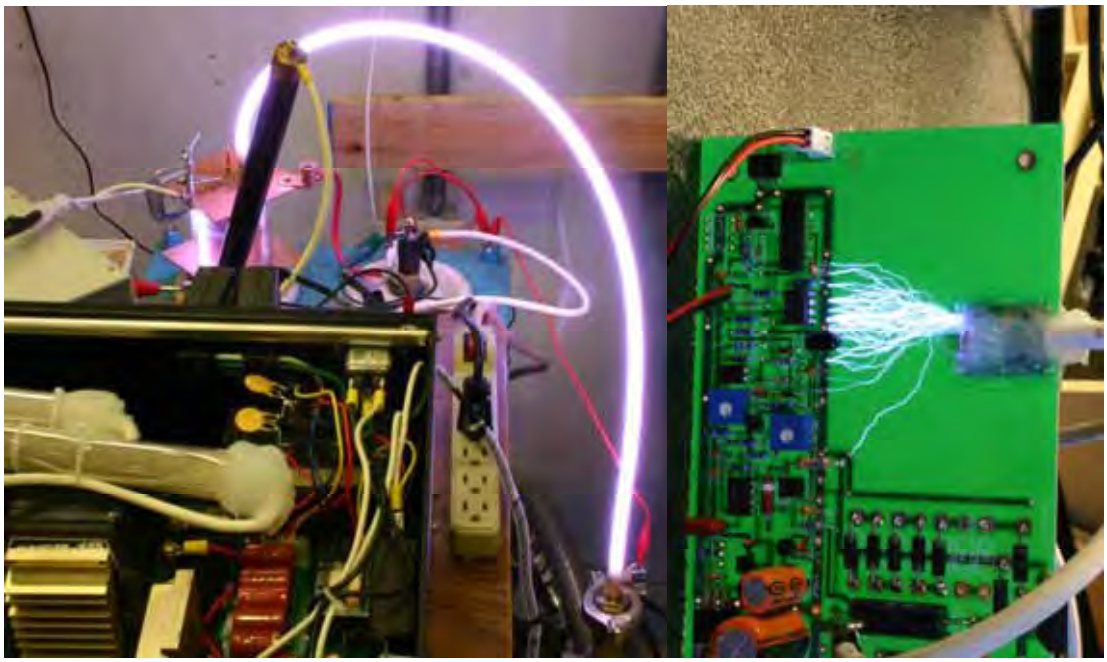

Figure 4.2.3. Discharges in uncommon places. (left) This picture was taken while testing low pressure switch designs. The triggered discharge was intended to stay within the cylindrical switch (far left glowing). The issue demonstrated has to be taken care of carefully when making use of unsealed, low pressure switches that require a vacuum pump. The implications when multiple banks are employed should be self apparent. An earlier discharge destroyed the pressure gauge. (right) This is from a very early design high voltage supply verification test of the common mode output overvoltage failure mode.

An abundance of caution and patience are the prerequisites for successfully dealing with high voltage systems. Failure modes and their effects are not always anticipated before they occur. These can present potentially dangerous situations both to people and to equipment. Foresight and planning for how to deal with failures before applying power to a system are critical.

4.3 High Voltage Design Considerations

The best accumulated summary of high voltage design and manufacturing information this author has yet found is contained within an interim report prepared under Air Force auspices[23]. It is highly recommended that anyone interested in designing reliable high voltage systems read the report in its entirety (352 pages). It is a very wide ranging report that has the basic purpose of giving designers new to the field the information necessary to not repeat the mistakes of the past:

> "For over 25 years, high voltage workshops have been held by Government and technical societies. At these meetings the same problems appear year after year, and about every 10 years, it has been observed that the same mistakes and problems reappear by the new scientists and engineers entering the field of power supply design."

While the referenced report covers the problems and issues associated with using gasses, liquids, and solids as insulators, our best experiences with permanent designs

have come from using solid encapsulants given their long term stability (and better materials compatibility). For the encapsulants we have settled upon, this is only practicable for relatively small physical dimensions. The primary component under the highest stresses has been the high-voltage step-up transformer used in our power supplies.

Originally, the power supply design used a commercial high performance ignition coil. It was later found that mineral oil (as applied without special procedures) did not provide sufficient inhibition of surface tracking, or the anticipated repair-ability that was anticipated as the chief advantages of oil. This puzzled us at first. Some of the chief problems were suspected incompatibility with the plastics used, suspected mismatch in conductivity (between the oil and plastics, resulting in plastic punch through) and the tendency for the oil to wick everywhere without fully wetting all kinds of materials. For designs with larger physical dimensions, oil (Diala AX oil) has proven more useful. We use it in insulating the high voltage switches mounted on top of single terminal capacitors (although just as messy as mineral oil). Flows are often observed in the oil filled systems (suggestive of ionic charge transfers). They also require occasional 'topping off' due to the oil lost to wicking.

We eventually found a very low viscosity epoxy encapsulant with good dielectric strength and good surface arc resistance in its filled form (EE4186/HD3561 is the filled form of Hysol RE2039/HD3561). The appropriate standards are ASTM D 149 and ASTM D 495[24]. It is important to realize that the duration of the breakdown test may be as short as 10 seconds or as long as 2 minutes dependent on the test method within ASTM D 149 chosen by the material manufacturer. The test method also defines the reported voltage as an RMS voltage (at typically between 60 and 800 Hz), i.e. peak voltage divided by $\sqrt{2}$ is reported. Arc resistance is synonymous with surface tracking damage resistance. The applied conditions get progressively tougher throughout the test (a non-linear scale).

The key factor for success with solid encapsulants was pumping out the mold and injecting the epoxy under 'roughing pump' low vacuum (~ $1/100^{th}$ atm) then reapplying atmospheric pressure well before the epoxy started to cure. Cooling (controlling the reaction rate) was necessary to prevent thermal stresses during cure that would result in interior cracking. Ironically, this successful approach was found just before discovering the Air Force design guide that neatly lays out the different approaches, techniques, and issues. For example: acrylic plastics, bare copper, rubber, waxes, fiber board, and PVC resins are all incompatible with mineral oils. Mineral oils also need to be filtered and out-gassed at low pressure and elevated temperatures for multiple hours once installed to achieve best results.

Figure 4.3.1. Potting mistakes. (left) This encapsulating block was allowed to cure at too high a temperature, resulting in an abnormally fast cure and very high thermal gradients. The epoxy interior is cracked at some interfaces with the diodes as well as internally to the epoxy. These cracks break-down, the one in the lower right appears to have developed into an arc (perhaps at higher pressure) while the more diffuse looks to be a low pressure glow discharge. Locally high electric fields on components are always a concern. (right) Another mistake is to pour the epoxy into place and then pump on it. This results in bubbles that can be trapped in odd places, and the bubbles can be seen to be larger at the top than at the bottom. A better approach is to use a tube to inject epoxy into the system while under vacuum (and components out-gassed). This results in far better gas removal, and what remains is at a lower pressure.

The remainder of this section is intended to hit some of the key points that emphasize the complexities involved without the in-depth coverage actually required to design and build a reliable (1000+ hr) high voltage module. It is advised that those serious about doing so review the original materials in order to fully understand the issues involved.

It should be well recognized that components run beneath their maximum ratings have longer working lives. Suggested applied high voltages are usually in the range of 60% of part specification for long life. For diode stacks, such as those made out of MUR1100/4100 avalanche rated diodes, our experience has suggested that a 60 kV rated stack is sufficient for operations at less than 35 kV. These diodes are similar to zener diodes in their tolerance (limited) for reversed currents. With this particular type of diode the addition of capacitors and resistors has been unnecessary (at least for the layout geometries we have chosen to use). Whether a diode is capable of being used like this is primarily a function of its breakdown characteristics. Secondary considerations are thermal dependence of leakage current and capacitance as well as the spread of values for these characteristics. While the addition of dividing resistors and capacitors may guarantee better lifetimes, they tend to be cost prohibitive. Most diodes can not be stacked without the use of additional components. Concerns do persist as to locally high electric fields due to the encapsulant choice.

Insulation choices are critical in the most intense electric field devices. Typical insulation average electric fields for effective designs are in the ranges of 50-200 V/mil (military) and 5-20 V/mil[25] (commercial). We originally chose to use a commercially available isolated ignition coil transformer (60 kV spark) in order to avoid the need to deal with high electric field issues. However, it was found during verification testing that the extent to which the secondary could be allowed to float above the primary was severely limited. For an automotive application primary to

secondary isolation is not a critical parameter. This resulted in supplies that could effectively operate in only one output polarity without reconfiguring internal components. At that time the target was 25 kV output, approximately 300 watts (although we tested to 500 watts and then de-rated). The power was limited by the internal temperature control of the transformer. It would heat up internally and the insulation breakdown strength would severely degrade. The dominant heating mechanism was believed to be due to the tape core laminations.

We still didn't want to deal with the need for coil forms and windings so we bought the components for the transformers un-potted. The secondary coil is made of small diameter wire (naturally leading to the potential for locally very high electric fields). Overall geometry, proximity of the output wire (larger diameter) connection to the coil's high potential was probably important. Locally high electric fields are so important because of what is known as the 'treeing' effect. When the local electric field causes a partial discharge breakdown, it eventually becomes conductive. This little spot where it starts generates locally higher electric fields in turn and it gradually grows throughout the insulation (looking like a little grape vine). The unfortunate effect is that the total thickness of insulation has a limited relevance in preventing eventual insulation failure. One must prevent the locally high electric field to begin with; this is often done by adding a conductive material around the small wire (effectively lowering the radius of curvature, decreasing the local electric field to a level well below the breakdown potential of the insulation).

There is a DC component to the electric field problem, addressed by conduction, but also an AC component, addressed by the choice of dielectric constants. Voids and air gaps in insulation are critical locations for corona discharges (partial discharges) to start. Most insulator solids have significant dielectric constants. For example, consider a linear geometry composed of a high potential conductor, insulator, air gap, insulator, conductor. How does the voltage initially divide across the three insulators? The answer is unevenly. One can consider each insulator as a capacitor with a virtual plate. It's simple to see that as the insulators get thinner (higher capacitance) they don't have to break down in order for the air gap to break down. This is one of the problems with using twisted pairs in high voltage applications.

Usually the purpose for using potting compounds is that the electric fields would break down air. Voids in potting compounds tend to break down easily. The Paschen curve defines the pressure distance product for breakdown. The goal is to keep any voids on the low pressure, small distance side of the curve; thus limiting the intensity of any currents. As many components operate in an oscillatory manner, no DC currents are required to result in significant damage. The practical goal is to keep voids as small as possible (as diffusion and decomposition would eventually bring any low pressure region up to atmospheric pressure over some long timescale).

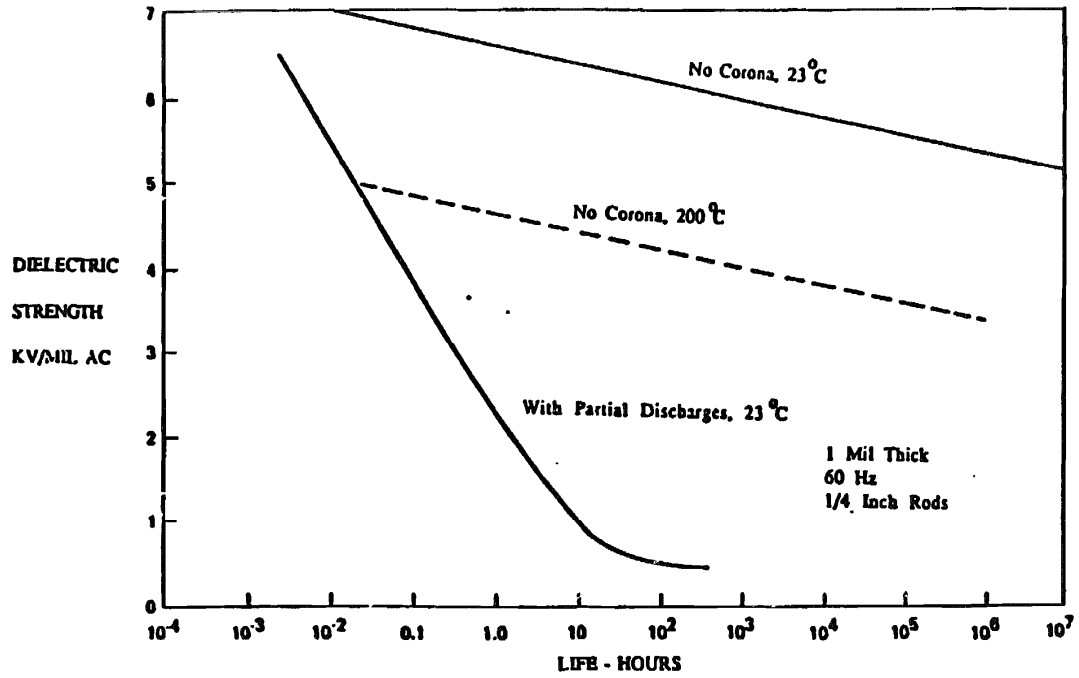

Figure 4.3.2. The effect of partial discharges (i.e. corona) on insulator lifetimes. This is copied from the aforementioned Air Force Design Guide to make the point of how lifetime is most strongly controlled by the presence of corona. A second point is that no insulator lasts forever, but moving back from high electric fields results in order of magnitude improvements in lifetime. Potting compounds are specified for a given dielectric strength for some length of time. The American Society for Testing and Materials (ASTM) publishes widely used standards for determining dielectric strength (ASTM D 149) and arc resistance (ASTM D 495). These standards are not free to copy, necessitating a trip to the library or bookstore.

It should be self apparent that there are many ways to unintentionally produce voids (and regions that act like voids): poor bonding of the encapsulant to the components, residues such as fluxes, oils, silicones, thermal expansion mismatches, and gasses generated from decomposition are but the first that come to mind. Impurities of all kinds, liquids, gasses, and solids present in the potting compounds or components can present serious issues. Some potting compounds may not be appropriate for vacuum processes due to high vapor pressures.

Eventually we hit upon a correct combination of potting compound and techniques to make a sufficiently reliable unit (with the output voltage upgraded to 40 kV) whose polarity is easily chosen and changed by the user (floating output). With the addition of an electrostatic shield between the high voltage transformer primary and secondary the rate of power supply cascaded failures dropped to zero. The transformer could internally arc due to external voltage transients without destroying the electronics. The circuits were gradually made immune to the generated EM noise without resorting to shielding; a later addition of conductive shielding around the control electronics should result in very good reliability.

Before this time it was realized how marginal the stranded 40 kV corona wire was during verification testing (and then on geometric electric field considerations). It is important to remember that the location (and geometry) of close by grounds to the HV wire can cause breakdown. Consider what the electric fields would look like if one brought a sharp grounded point up next to the high voltage wire insulation. This is a good reason to prefer coax. We have started replacing the 40 kV corona wire with 80 kV semiconductor coated stranded insulator wire (with a PVC outer jacket) in all our applications. The replacement is cheaper than the original choice at the price of a lower maximum working temperature. The original 40 kV wire was perfectly appropriate for the original 25 kV specified requirements (as one might have expected given a standard 60% de-rating).

This section has related some of the critical issues that arise in high voltage designs. In essence, it all comes down to the strength of the local electric fields and features that create them. If the electric fields are permitted to be too high anywhere the lifetime of the equipment can be expected to be short even when it doesn't result in immediate failure.

4.4 Reversal Bank Units

The basic design of the reversal and bias bank modules derives from the Large S Experiment hardware. The structure, capacitors, and basic feed plates are the same. Five 14 μF, 25 kV (90% reversal), double pole capacitors are connected in parallel for each module. The modules are designed for 100 kA each[26]. Modifications have been made to the ignitron start and crowbar portions, the output cable connections, and the module mounting arrangement (i.e. now wheeled). The frames are now hard grounded (for safety reasons) while the low voltage side of the capacitors are soft grounded through two parallel 10 kΩ resistors. The output cable shields are connected to the low voltage side of the capacitors. This results in the somewhat uncomfortable situation of having the coax shell potentially with high voltage on it with respect to ground, but only during transients. Note that this is a fundamentally inescapable result so long as the coax shell carries currents. It also enhances the probability of the transmission line acting as an antenna to a degree not limited by the cable's transfer impedance.

Note that as described later in the diagnostic section the output coil has no hard ground connection and is exceptionally well insulated. LSX had a hard grounded, center-tapped driven coil resulting in the requirement of floating the whole module chassis[27]. The hard ground made the system act as an unintentional Marx generator in the event of a single polarity (+ or -) start switch pre-fire. This would double the voltage of the opposite high side capacitor plate, potentially causing an internal breakdown between the 'high' potential terminal and its case. The LSX solution was to allow the capacitor case to charge with respect to ground. This still has the unfortunate side effect (not considered) of doubling the potential at the hot terminal, but without leading to an internal breakdown of the capacitor. Given the limited laboratory space that we have, using the LSX approach would have made the

experiment unfeasible while complying with the National Electric Code (and hence able to pass inspection, which we did manage).

Our resistors linking the case to the 'low' voltage capacitor terminal can be replaced with shorts to little ill effect. This might have benefits in terms of EMI. In our experiment, the resistors simply limit the current that can return via them rather than via the coaxial output cables. Replacing the resistors with shorts effectively results in one-terminal capacitors. The price of making such a substitution is apparent in the case where coax output shields are all disconnected (as from incomplete assembly) while the center connections remain and both modules are fired. Very high currents would flow in the ground connections with the attendant high magnetic fields defined by the cables' paths to the driven coil and the AC power line path.

Figure 4.4.1. LSX reversal bank module design. This picture shows one half (-) of a pair (+/-) required. Various states of assembly are shown. The final assembled state (not shown) has two conductive feeds (half cylinders) around the ignitrons. Suspicion has arisen that the current feeds are sufficiently asymmetrical as to preferentially form the current channel in the switches off to one side.

Figure 4.4.2. Module evolution. (left) The first modification made upon receipt of the old LSX modules was to add wheels to make it possible to move without a 'walky stacker'. This was rapidly followed by changing the ignitron stacks to singles in order to reduce stray inductance (and provide spare ignitrons). The output cables were also changed from a sliding banana plug connection to a center screwed connection (where the screw carries little to no current). The ignitron igniters were made local to the module, soon to be followed by power supplies and dump capabilities. For testing purposes, our ad-hoc pole-transformer supply was capable of supplying kilowatts. The ignitrons could be rapidly heated. (center) The ignitrons were still undesirably inductive as start switches. To alleviate this problem, two are placed in parallel with their cooling jackets used as the current returns. Note that some of the crowbars have stainless steel current feeds (used as resistors). These resistors would frequently arc to the screw heads and other close metal surfaces during firing. These arcs would sporadically destroy and disconnect the resistors from the ignitrons throughout the experimental runs. (right) The start ignitrons have been replaced by copper vapor switches. The crowbars carry too many coulombs for a similar substitution to be advisable. A 50 watt halogen spot light shining onto a blackened plate connected to the crowbar ignitron's anode can be seen; this provides a 30 C° temperature difference across the switch. Cooling of the cathode is still a potentially problematic issue given our region's summer climate. Tests on the first module (left) suggested that when the cathode reaches the region of 30 – 35 °C very little voltage hold-off capability remains. Since we need to run ~ 30 switches reliably we need to keep the cathode temperature much lower.

A basic shift in design philosophy took place between the time of LSX and our experiments. LSX centralized the triggering systems, the charging systems, and the dump systems. The design also required active events before firing to protect the power supplies. These were in essence separate stand alone systems that had to be wired in as a whole throughout the facility with long DC high voltage wires connecting capacitors to power supplies, dump relays, and charge disconnect relays. We anticipated the need to detach modules for temporary use in smaller experiments, along with a lack of available manpower to maintain and set up complex spatially integrated systems. This necessitated making the modules standalone with a higher degree of designed (as opposed to procedural) safety than before. The triggers, power supply, dump relay, service dump, and voltage monitor/indicator were integrated onto the module chassis. Remember that the modules need a switch across the capacitors providing a dead short, rather than just resistive shorts, to deal with some safety considerations. All DC high voltages are kept local to the module. Obviously the pulsed output cables can not be kept fully local. The driven output coils would also be kept as separate as possible. There is no magnet spine connected to ground in our experiment. As a consequence, we insulated coils on the experiment from day 1, avoiding the problems of arcing to flux loops that LSX had[28].

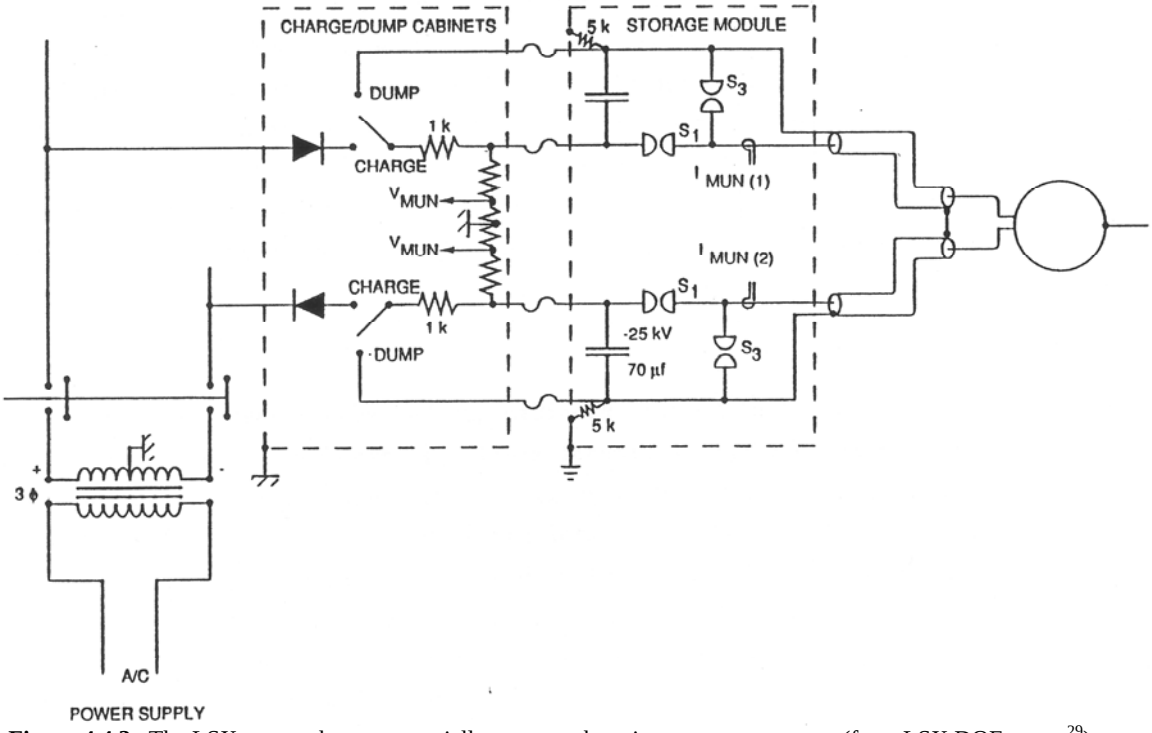

Figure 4.4.3. The LSX approach was to spatially separate the primary power systems (from LSX DOE report[29]). The power supply would power multiple storage modules. Note the use of different 'chassis / common / ground' symbols. Circuit 'ground' is not directly connected to chassis ground.

Cooling for the ignitrons is still a shared system (when it is unavoidably necessary) and has been half-heartedly attempted with limited success via piped air cooling. Heating can easily be done on an individual basis. The obvious solution to cooling is to cool the whole bank space. This removes the need for liquid heat transfer media and the risks of condensation. Simple fans can be added blowing on the ignitron cathode regions to increase the heat transfer rate if necessary.

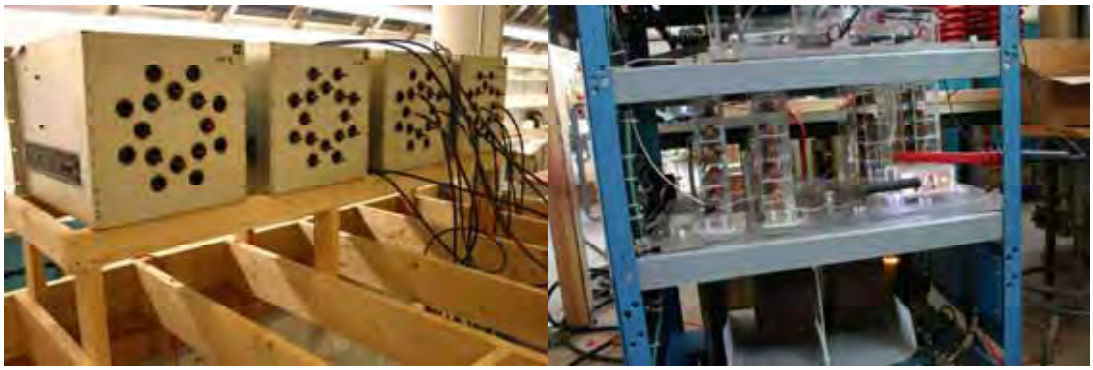

Figure 4.4.4. LSX era systems. (left) Ignitron triggers based on SCRs triggering EG&G GP85 Trigatrons. One fiber trigger was set up to drive a large number of ignitrons. Many of these old units self trigger sporadically. These have been replaced with individually triggered tube free units. (right) Old HV power supply based on a 60 Hz AC HV transformer visible at bottom fills out an entire rack. The rectification diodes fill out the center section with relays and controls on top. Upon realization that the transformer had failed this unit became scrap.

Figure 4.4.5. New HV power supplies from the prototype stage to production. Note that the final unit includes a dump relay and floating voltage outputs. The dump relay in this model is capable of reliable operation up to about 30 kV after which the generated corona causes sporadic contained arcs (non-damaging). A replacement relay has been made and tested but is not as cheaply produced as the original. The final unit has not yet been fully optimized. This picture is not representative of the final design. It represents a limited production model put in place due to time constraints.

Some aspects of the current module high current design are still not ideal. With the ignitrons installed Figure 4.4.6 represents the gross electrical circuit. The effects of coupling from coil to coil are not represented here; nor are the connections coming from the bias modules. The more subtle thing to realize from this circuit representation is that the inductance associated with the output snubber (100 nH) is not insignificant compared to that of the start ignitron & capacitors (250 nH). What was originally missing from this figure was the capacitance due to the large module feed-plates. This capacitance comes to around 10 nF at each location with a different configuration from + module to – module.

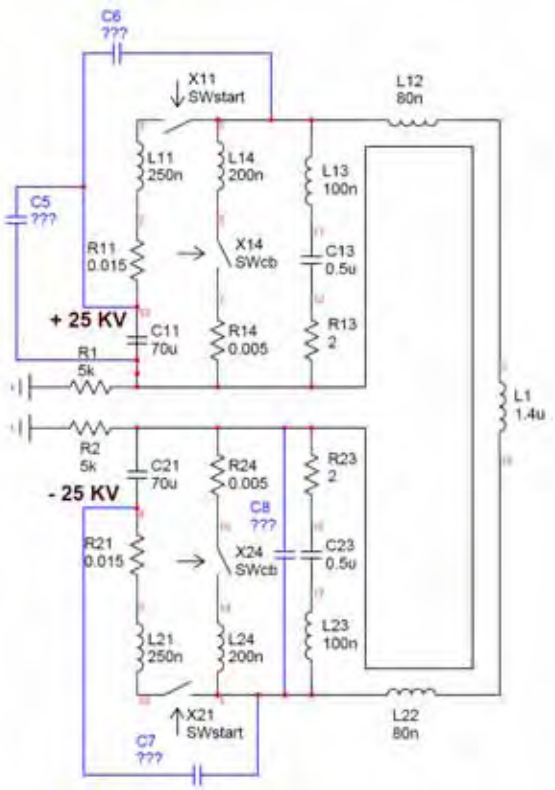

Figure 4.4.6. Effective Reversal Bank module pair. The blue indicates the locations of the inherent capacitance of the feed-plates. The capacitance is estimated at approximately 10 nF based on 800 in^2 with an 80 mil polyester dielectric. The banks are asymmetrical in their stray capacitors' locations. The snubbers are represented by the **3 components, crowbars **4. Obviously, this schematic does not break out the locations of the individual stray elements sufficient to see all the different bank resonances. It is also not appropriate for studying the effects of the transmission lines.

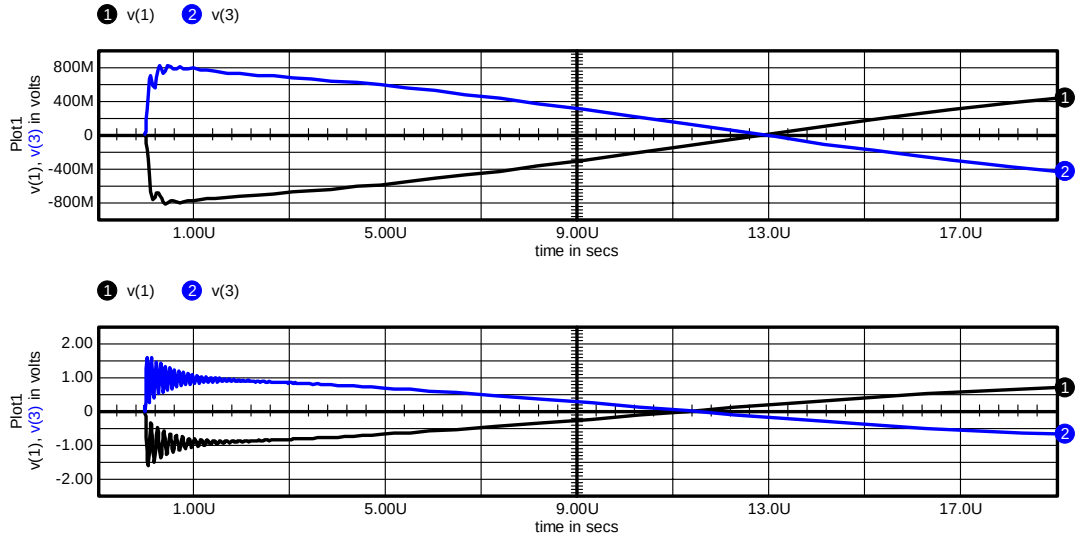

Figure 4.4.7. Normalized resultant voltage across output (and crowbar switch) of modules. (top) Ln1 = 250 nH, an appropriate value for a type D ignitron start switch. (bottom) Ln1 = 25 nH, an extreme case to show the perils of reducing only the start switch inductance too far. Ln1 = 150 nH ($R_{critical,\|} \sim 2\Omega$) is the crossover point to more

resultant voltage across the output than charged voltage. The periods are different because the total inductances are different. The two start switches are closed at the same time for these figures.

The inductive construction of the snubbers and the dissipative elements of the start switch (not properly represented here) are the major components that have to be taken into consideration. Increasing the ring frequency by either lowering the stray capacitance or start switch inductance lowers the required snubber capacitance but also requires that the stray inductance in the snubber be lower. At some high frequency the characteristics of the switch itself (how it actually closes, propensity to ring) can be expected to dominate the overall behavior.

Increasing the resistance in the start leg for high frequencies is a solution that removes the need for separate snubbers when it reaches the critically damped threshold. Finding ways to lower the stray capacitance is good. Increasing the inductance in the start leg for high frequencies (with a smooth roll off in inductance towards lower frequencies) is a solution that still requires output snubbers. An analogous approach is a saturable core, although the question of how low the saturable core inductance can go is worth asking (i.e. one might suspect that the additional space adds sufficient inductance whether saturated or not). Of course, having a well-characterized (and stable) switch with the appropriate parameters can substitute for additional components. This is the approach realized in some semiconductor applications.

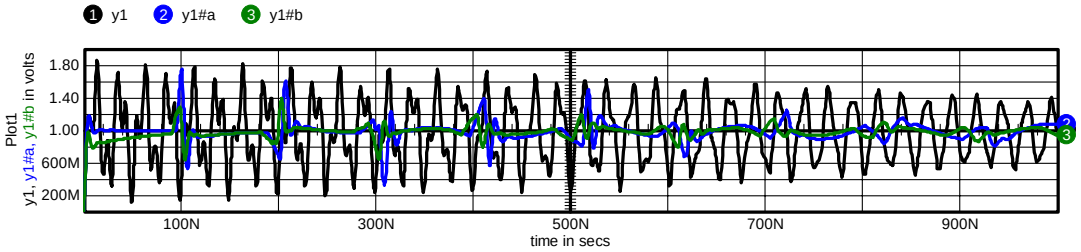

Figure 4.4.8. Normalized voltage across output of theoretical module with output cable at critical impedance and no dissipation. The stray capacitance (100 pF) and inductance (25 nH) associated with the start switch have been chosen to represent minimal values and hence high frequencies. Output voltage is displayed for two cable lengths (blue ten times longer than black). Green adds back in the original module snubber. The extreme voltage excursions due to reflected waves are clearly visible and of short duration with the original snubber in place or not. Note that these chosen values are arbitrarily chosen, not representing that of a copper vapor switch (which would have ~ 20 pF before the feed geometry is considered)

It is tempting for cases where the characteristic impedance of the output cables is comparable or less than $R_{critical,\|} = \frac{1}{2}\sqrt{(L/C)}$ to ignore the need for dissipation. In such cases one can still expect an overvoltage to appear due to the improperly terminated cable (terminated into the driven coil, i.e. high impedance). A question one might ask is: will the wave be sufficiently dissipated by 'stray' resistance (at frequencies less than 1 GHz, dissipation due to the skin effect)? The simple answer is probably not in a short length output cable. As an example, the attenuation in 20 kV RG 59 U is ~ 12 dB per 100' at 1 GHz [30], implying approximately 0.4 db per 10' at 100 Mhz.

Additional strains are put on this system due to imperfect control of the timing of start switch closure. Fortunately, these voltages are lessened by the LC filter formed by the load inductance, stray capacitances, and snubber. This results in an overvoltage of approximately 10% for every 100 ns offset in timing.

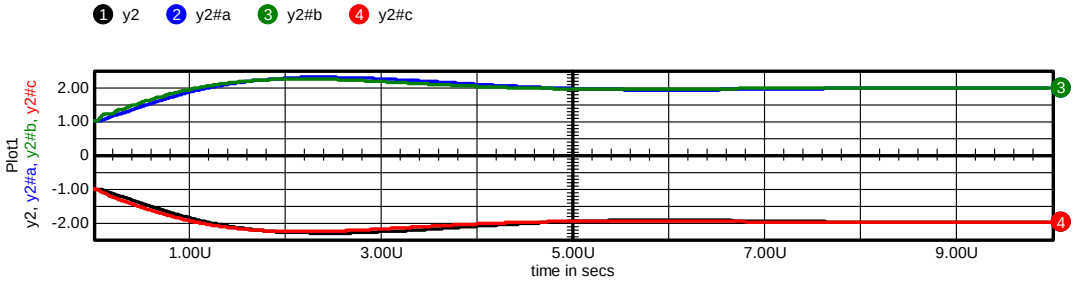

Figure 4.4.9. Voltage across open start switch with opposite switch firing. The driven coil inductance acts as an excellent high frequency filter resulting in a long rise time. (1 & 2) Trace is for the 250 nH start switch case. (3&4) The 25 nH start switch is used making an effect of capacitance asymmetry in the +/- modules apparent in the ringing on the green trace.

The remaining critical element that we have not yet discussed is coupling from adjacent coil systems. The greater the number of coils (associated with higher discretization) in a shorter experimental length the higher the coupling. This presents problems with start switch overvoltage; if one coil is fired but not crow-barred and the adjacent coil is not fired, the start switch (and associated insulation) can be over-voltaged. It also presents issues associated with current sharing. Initially, the firing of an adjacent coil results in less voltage across the remaining start switches. The same effect takes place when the crowbars are fired. This means that the triggering mechanism for the switches needs to be dominant over the effects of voltage across the switch in how the switch breaks down into full conduction. This is frequently difficult to accomplish while remaining within switch manufacturers' suggested triggering parameters. The need to have a switch that can take full voltage but be switched on at low voltage greatly limits high voltage hold-off switch options (amongst non-semiconductor switches).

As has been repeatedly noted, the simple, reliable solution to crowbar pre-fires is to choose a switch capable of withstanding twice the charge voltage (which is in fact what the LSX team effectively did by having two stacked switches). This is analogous to the high voltage power supply design review paper recommending that many components be run at 60% of rated voltage[31]. It is suggested that some components should be run with a peak as low as 25% of rated value. The simple truth is that operating equipment close to its breakdown potential is likely to result in problems from any number of sources. The most relevant unexpected example to demonstrate the point is cosmic rays' impact on high electric field semiconductor applications modeled in this 1997 reference[32] (that can be considered a reasonable starting point for further exploration). Earlier, cosmic rays were experimentally determined to be the culprit in high switch failure rates in a public transit application.

4.5 HV Power Supplies

It is worth repeating the reasons for developing a high voltage power supply in house rather than purchasing one from a commercial source. This can be distilled down to two primary rationales: costs and requirements. A third reason is that the problems encountered with HV are the same problems that need to be addressed for the other module components. I.e., if you see the same problem in different ways, it can be easier to realize what is going on. Most commercially available high voltage power supplies are not cheap, typically starting at a price of around $1 k and going up. The desired power usually only determines cost at higher power levels. But the true cost of a power supply isn't the number listed on a price tag. The true cost has to include all the attached systems and more importantly time (including down time) and maintenance. Having damaged many high voltage (and low voltage) power supplies from various vendors in the past we were all too aware of the real costs. We desired a general laboratory HV power supply that could be used in multiple situations without any additional engineering effort for each application. This generates a plethora of unusual requirements that are not mutually exclusive of each other.

4.5.1 Requirements

The supply needs to be capable of reversing output polarity. We have applications that require different polarities with respect to ground. Different capacitor bank module topologies require different native polarities usually determined by the type of high voltage switch employed.

The supply needs to be capable of operating with a floating output (with common mode voltages not anticipated to exceed the maximum output voltage). The electrodes used by an arc through a gas can be better controlled when the electrodes aren't ground referenced to anything except through the plasma. In some applications[33] these unintended currents and electric fields can result in forces that swamp those forces we are trying to measure. However, it should be kept in mind that capacitance can still play a significant dynamic role. Similarly, the high vacuum pumps can be connected to ground without drawing large currents through them while still having an electrode driven arc in the system. It needs to be able to do this safely.

The supply needs to be able to discharge capacitors as well as charge them. Dumping needs to be its native unpowered state. Having a dump switch included in the power supply removes the potential of scientists neglecting to include a dump mechanism that doesn't require their prototype to work. People have tried to dump banks (not on PHD) with broken shorting sticks, without inspecting their equipment first, and then been shocked. Also to be considered is that many commercial high voltage relays are as expensive as HV power supplies. We managed to construct a cheap, but very reliable relay that could be included in the supply and take advantage of shared systems and controls. Chris Pihl deserves credit for the mechanical design of the (< 30 kV) dump relay.

The supply needs to be tough. It needs to be capable of surviving conducted EMI as well as radiated. It needs to be able to withstand a connected capacitor bank reversing voltage during discharge. It needs to be repairable. It needs to operate

from readily available power (i.e. 120 VAC single phase). It needs to be controllable via fiber optics. It needs to be light enough that one person can easily move it.

It needs to be adjustable to low power levels (between 1 watt and 300 watts). We've actually used our supply to drive micro arc jets. It needs to be capable of the highest voltages we anticipate using (originally 25 kV, now 40 kV).

4.5.2 Power Supply Basics

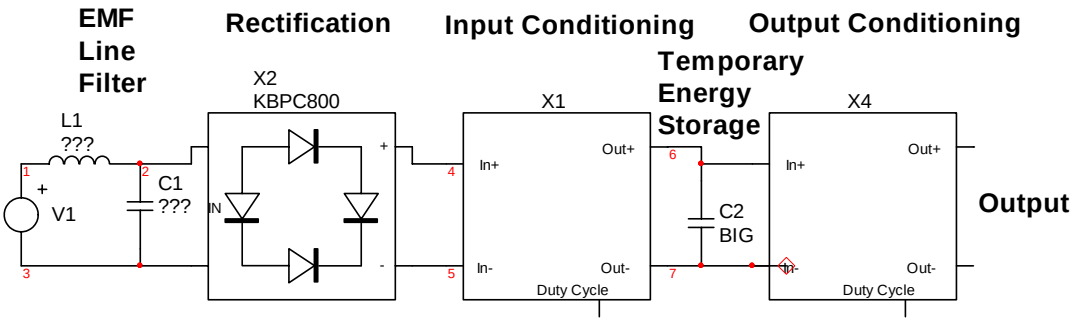

Figure 4.5.2.1. Generic power supply diagram. There may or may not be a 60 Hz AC line transformer before rectification (not shown).

Most power supplies have two primary sections. An input converter takes the available AC line power and puts it into temporary energy storage. The second section takes the power from the energy storage and converts it into the required output characteristics. This staging is necessary because continuous output power is desirable and input power = $V_0 \sin(\omega_0 t)\, I(t)$ is discontinuous (for single phase power sources). The energy lost between the utility's generator and the wall socket are roughly proportional to $R \int I^2(t)\, dt$. Using a single large storage capacitor immediately after rectification results in large current spikes near peak voltage as the diodes go into conduction, charging the capacitors. These spikes are somewhat limited due to the use of diodes designed specifically for power rectification. This is inefficient and results in significantly less available output for a given circuit breaker. Active power factor correction (PFC) circuits are used (in the input conditioning block). Their purpose is to make the current drawn from the AC line into a pure 60 Hz waveform in phase with the input voltage with minimal losses and distortion. The temporary energy storage is not at a constant DC voltage. Another converter stage (output conditioning) takes this DC voltage and supplies continuous current at the desired output voltage.

For the vast majority of our applications continuous output power is not required due to the large capacitors connected to the output (required by our pulsed experiments). This allows us to severely decrease the size of this temporary energy storage and make the output converter into the active PFC converter (significantly reducing cost and weight). The input conditioning block is replaced with an LC filter. The size of the required energy storage capacitor and LC filter is determined by the average power output and the operating frequency of the final converter stage.

Higher frequency results in smaller required components in every aspect of the power supply (except for HV insulation thickness).

4.5.3 Converter Topologies

There are many power supply converter topologies in common usage. Buck, boost, buck-boost, sepic and cuk do not provide galvanic isolation. Flyback, direct (push-pull), and some resonant mode (special class of direct) converters can provide isolation. A topology with galvanic isolation is required to permit operation of the supply directly off line voltage (with no intermediate 60 Hz power transformer). Appropriately, this is called an 'off-line' converter. The elimination of the 60 Hz power transformer saves cost but more importantly weight, size, and dissipation. However, properly trained technicians are required to service the supplies due to the internal presence of line connected components. The classic mistake is grounding of a scope probe to a point in the circuit (drawing an arc and blowing fuses if lucky) without plugging the power supply into an isolation service transformer.

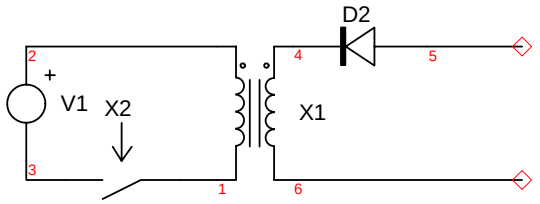

Figure 4.5.3.1. Basic flyback converter topology. Note the marks on the transformer indicating the winding ends that start at the same place and are co-wound. Node 5 will be at negative potential with respect to node 6.

Flyback topologies have features that greatly simplify certain aspects of the overall design. Their output current is inherently limited by the transformer turns ratio and the 'charging' current. For this reason they tolerate shorted outputs very gracefully. The output voltage is easily inferable from the primary voltage when the switch is open via the turns ratio, diode voltage drop, and output structure resistance (given the charging current). Only one switching component is required on the high voltage side, a simple diode, and only one on the low voltage side, a low-side switch. The number of components that have to sit at high potentials or be rated for high voltages is minimal (5 or 6 logical components for a supply that meets our requirements). This approach can be made very durable in the face of voltage spikes coming back from the load.

The chief drawback is that the power is limited by how much energy the magnetics in the transformer can store before saturating (volume and core material choices) and the frequency of operation (limited by stray inductance, losses, and winding resonances). A second drawback is the intrinsic presence of large zero-crossing voltages on the output lead. This can have a significant effect on life expectancy. This topology also requires a certain amount of finesse in making the primary circuit measurements of current and voltage. The flyback voltage spike will

be meaningless during the time it takes the stray inductance to discharge, and will be followed by a ringing on top of the voltage one would want to measure. It is also most unwise to assume that the drive voltage source is constant as it is typically unregulated, so this had better be a differential measurement. Similarly, the current during the 'charge' phase will oscillate due to transformer resonances. These transformer non-idealities are of no serious consequence for the automotive ignition circuits, but are a serious nuisance for a practical charger. Nevertheless, with the requisite amount of cleverness and attention to detail a highly successful design is very feasible for our required power levels with readily available high voltage parts. The dominant power losses are due to transformer characteristics, followed by switch losses. The costs are dominated by the high voltage components and the chassis.

The stray inductance flyback pulse can be dealt with by using feedback from the transformer to the switch. A stack of low power zener diodes can be used to partially turn the switch back on during the pulse. A fast switching diode is required in the zener diode stack. An additional zener diode is also required to limit the gate voltage. The particular turn-on characteristics of the switch (usually an IGBT) and the impedance of the gate driver determine whether this is a practical approach in any given situation. The potential exists for continuing oscillations in a poorly-designed circuit. These zener diodes also act as a fast limit (although they won't have the power rating to handle continuous absorption of the pulse). This allows the vast majority of the power dissipation from stray inductance to be dissipated in the switch; where it is easily heat sunk. Some power is wasted in this approach in that the DC power source voltage adds unnecessary heating to the switch. Also, energy is dissipated from the coupled inductance of the transformer. Standard snubber approaches require high power dissipation devices and must be tuned for a specific primary current level in order to not add more dissipation than is necessary. For the normal snubber approach, the extra dissipation can result in very low efficiencies at low output power. Zener diodes alone are expensive and difficult to find with sufficiently high power dissipation and low impedance.

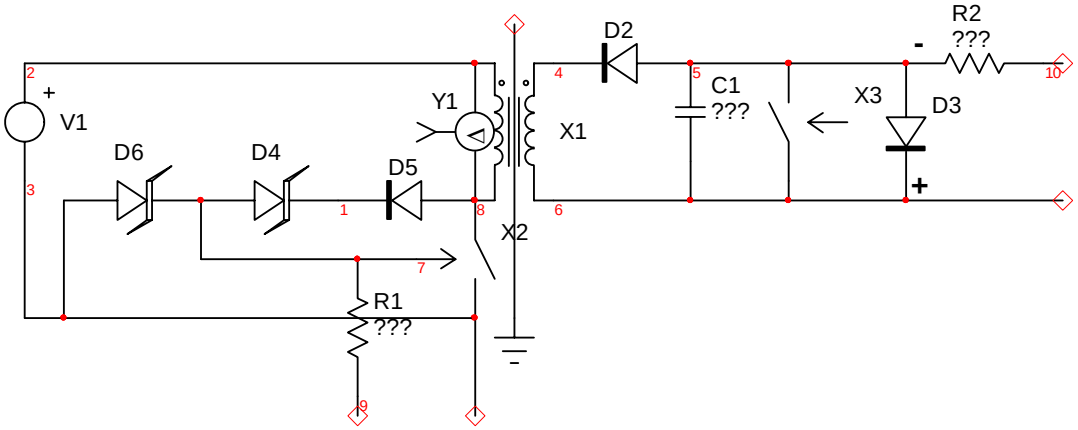

Figure 4.5.3.2. A more realistic flyback power topology. This diagram neglects to make explicit the non-ideal modeling. A strong understanding of the 'strays' is necessary for proper part choices and reliability.

The current limiting output resistor ('R2' in series with the output, external to the supply for inspection) is sized to limit peak surge currents during capacitor bank reversal (typically to less than 5 amps, although much higher currents are permissible for brief durations) and as such is not a major loss mechanism during charge. A small local HV capacitor across the output acts to slightly improve efficiency and reduce the generated noise. It also provides time for an open output circuit to be detected before over-voltage damage can result. There is a small, but significant capacitance from the secondary to chassis ground that helps limit the effects of common mode signals coming back on the power supply output leads (not shown). Obviously, some care should be exercised in preventing the common mode components from ringing to voltages beyond breakdown.

The primary HV rectification diode stack (D2) is chosen to have a much higher voltage drop than the normally non-conducting HV protection diode (D3). The two high voltage diode stacks are of different components and construction. The D2 diode stack limits efficiency at low output voltages (100 V range). Although it would be prohibitively expensive in this case, it is worth noting that one way around the diode drop problem is to make the diodes out of switching components (e.g. semiconductor synchronous rectifiers).

When the connected output bank goes through reversal, current preferentially flows through the protection diode rather than through the transformer. The transformer also has a relatively high resistance in the secondary winding. In cases where R2 is missing or bypassed (due to negligence), 'D3' eventually fails as a short during reversal (but is easily replaced). In cases where R2 = 0 and a dump command is given, the dump switch can be arc welded (easily fixed, but not a quiet event) and more serious damage could result dependent on the size of the bank. R2 is currently mounted externally (for inspection), although we would like to mount this component internally to remove user error.

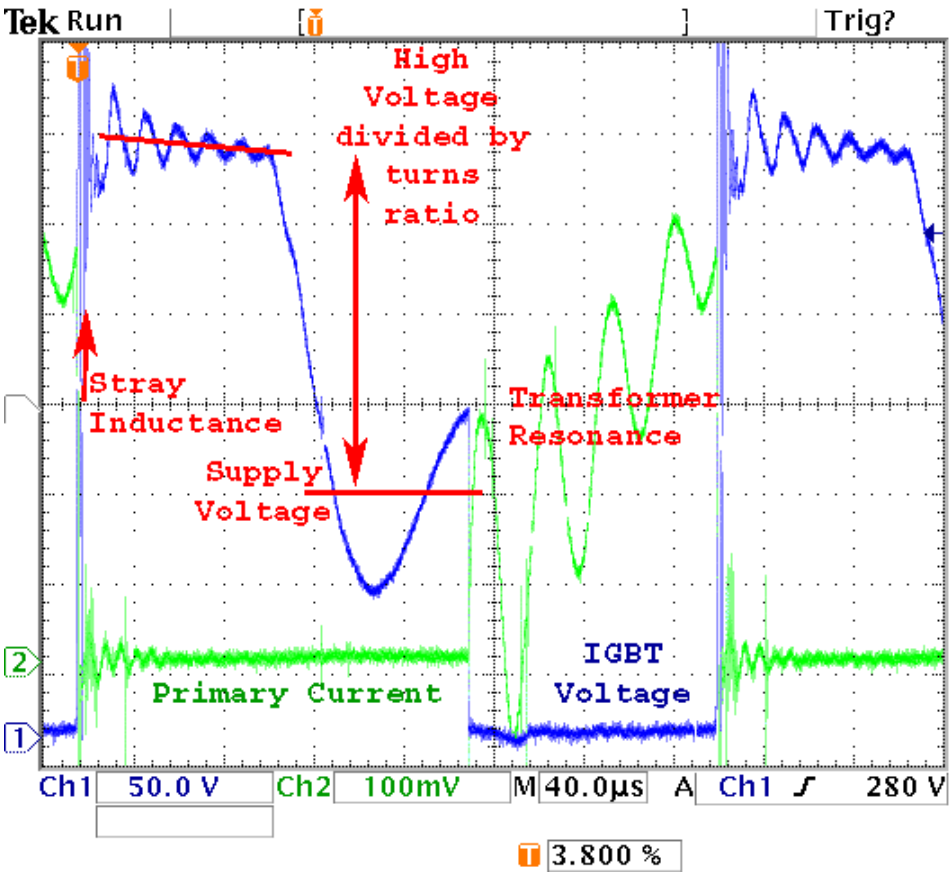

Figure 4.5.3.3. Primary traces from an operational flyback transformer. The different non-ideal transformer effects are clearly demonstrated with the transformer resonance and stray inductance highlighted. Once these effects are compensated for, and knowing that the transformer turns ratio is 1:70, it is easy to ascertain that the bank being charged is at approximately 13.5 kV. Errors due to the transformer resonance can be empirically corrected for. The determination of line supply voltage is more properly made by direct measurement (rather than inferred during the 'no-current' state as shown here). The mode of operation chosen results in an audible ascending tone (here 3.6 kHz @ 13.5 kV) amplitude modulated at 120 Hz. The supply is run such that this tone increases in frequency as the output voltage increases. As such it is yet another diagnostic to the attentive ear. This has the somewhat amusing side effect of creating a strong Pavlovian response in personnel at high voltages while testing new banks capable of arcing over. It is much preferable to run at a higher frequency due to simplification of the line voltage input structures. This is not an option due to the magnetic and resonance characteristics of our original choice in transformer. The resonance mode, at 36 kHz, should be inaudible.

A direct push-pull or resonant mode topology does not seek to store energy in the core, greatly enhancing power density. However, for best results it does require either a switch bridge configuration (4 switches) or split primary winding (2 switches) to reverse the applied polarity. Single switch configurations are possible, but at much lower duty cycles. The resonance mode version is most easily accomplished with an additional transformer winding and some capacitors. For the resonance mode version, an additional downside arises from needing components to actively limit output (or input) current. This can take the form of another converter stage requiring inductors but no transformers. It is more difficult to get the desired high voltage output as stray inductance tends to increase as turns-ratio of the transformer increases. The

capacitance of the secondary is an additional potential loss mechanism (that is more important for efficiency of low power converters).

A capacitive voltage multiplier is frequently used after this topology to obtain sufficiently high voltage (with acceptable losses) which leads to a much higher active component count on the high voltage side of the transformer. The multiplier capacitance interacts with the transformer resonance. It also means that the output voltage has to be monitored on the secondary and transmitted back in an isolated way (or not have the output be able to float). We have had good luck using this approach where space is at a premium and the output voltages are below 5 kV with low power requirements and low voltage input sources.

The high voltage power supply design effort eventually accomplished all of our technical goals. The estimated real costs (including construction labor) are far less than $1 k per unit. The downside has been that the necessary attention to detail in construction of the high voltage components prevents the assembly task from being delegated to random laboratory personnel. This results in scheduling problems as our primary purpose is not the manufacture and sale of high voltage equipment. Additionally, in-depth testing of units improperly assembled requires electronics technicians with the proper training in dealing with off-line electronics. Overall, the results have been positive.

4.6 Ignitrons

Ignitrons are low pressure mercury vapor switches. In their off state they are ideally at vacuum. The presence of gasses and electrode irregularities results in un-triggered discharges either as glows (where the switch maintains significant voltages across it) or as arcs (full conduction). The bottom electrode (nominal cathode) consists of a mercury pool. The top electrode (nominal anode) can be made of various metals or graphite[34]. Typically, the bottom section is made of a stainless steel container separated from the smaller top electrode by a glass envelope. It is desirable for any conduction to occur between the top electrode and the mercury pool (rather than attaching to the steel case)[35].

Operating properly, the vaporized mercury acts as a getter, removing gasses into a bound form. There have been many permutations on this basic design. Ignitrons are now available that are specifically intended for very high current (700 kA) or high voltage (50 kV) applications. Ignitrons are specified as capable of withstanding very high coulomb transfer (100s C). The use of ignitrons at very high current and high coulombs has resulted in disappointing results in the past[36]. Richardson Electronics, Ltd. (aka National Electronics) is our current replacement provider. Most of our current type D ignitrons are EEV BK488D left over from the LSX FRC experiments. The EEV BK488A/B peak ratings are 25 kV, 150 kA, 200 C per pulse, 5% reversal, and 4A average current. The BK488D version (the D is not a reference to type D) has three independent igniter pins. Our experience (with old ignitrons) is that whatever small improvement there may be in switching time characteristics that is obtained from firing all 3 pins with individual trigger generators

is not worth the extra effort over triggering only one pin. The extra trigger electrodes are merely spares in case one cracks or becomes significantly wetted to the mercury.

Figure 4.6.1. EEV BK488D type D ignitron with thermally insulated glass envelope. Flexible metal connections must be used (typically on anode) to prevent the build-up of mechanical stress in the device. Here, a thin metal strap is used to make the electrical connection. The painted red metal cooling jacket is the most apparent feature of this model of ignitron.

A triggering electrode (made of a carefully chosen semiconductor material) is placed in contact with the mercury pool. The mercury should not wet to this electrode. The formed meniscus presents an impedance of up to ~1 kΩ (for available units). Triggering is accomplished by applying a high positive voltage on the triggering electrode with respect to the cathode. This can either form plasma[37] or result in an increase in mercury vapor pressure. The voltage across the tube from anode to cathode acts to complete the switch breakdown and form an arc. Applying the wrong sign of voltage to the trigger can result in damage occurring between the mercury and trigger electrode even at very low currents[38].

Similarly, operating the main discharge with a reversal current is bad for the anode. EEV Lincoln suggests that graphite anodes tolerate 5-10% inverse peak current, 40% for stainless steel, 85% for high density graphite, with molybdenum being superior[39]. A second parallel tube connected anode to cathode should be considered for cases where reverse current must be carried.

As material is gettered or ablated onto the top of the mercury pool and onto the trigger electrode, the mercury can start to wet to the trigger electrode causing the trigger impedance to drop. Lower trigger impedance makes it more difficult to drive

the trigger to sufficient voltage to result in a fast breakdown. Sometimes this determines the useful life of the tube. Some trigger units are capable of increasing the impedance of the trigger pin through repeated triggering events. Unfortunately, the improvement is not persistent (as one might theoretically expect since the 'scum' is capable of moving). Ultimately, the gas released during the full switch discharge interaction with the coated mercury pool and walls becomes significant enough to overwhelm the gettering action. Pressure builds in the tube while the critical breakdown potential falls. Most of our tube failures are of this nature where the tube is no longer able to hold off significant voltages.

Ignitrons are quite useful because of their high current (and total charge) handling capabilities. They do not require very high voltages to trigger (1-10 kV will suffice for short time delays). Given appropriate triggering units, they are capable of switching from low voltages to high (~30 volts to ~ 50 kV) with no adjustments. Keep in mind that for reliability the application will need to run at less than 60% of the rated maximum voltage. Consistency of delay and effective resistance are a problem at lower applied switch voltages (anode to cathode)[40]. Thus, ignitrons can be used as crowbar switches (albeit with individual inconsistency).

Ignitrons' most glaring problems are a propensity for pre-firing and glowing even when new. In order to reach full voltage hold-off potential they need a period of electrical conditioning (presumably to smooth off the sharp electrode edges or getter the vacuum). Once an ignitron has conducted significant amounts of current the tube must not be tilted. The resultant scum on the top of the mercury pool can adhere to the glass, the anode, and the trigger. Arcing on the surface of the glass could result in the destruction of the vacuum seal (and hence the switch). At room temperature, a significant amount of mercury vapor exists in the switch. This vapor can condense out on the glass and the anode (an analogous example is a water bottle that has been lying around for several days partially filled). For this reason a temperature difference must be maintained across the switch (albeit not continuously), preventing the buildup of mercury condensate while not increasing the vapor pressure too far[41]. For ignitrons with glass envelopes, it is a good idea to insulate the glass from the ambient air so that the glass can warm up. While cooling improves voltage hold-off by lowering vapor pressure, it does result in a switch that is more difficult to trigger into full conduction.

The ignitrons we have are almost all previously used. It is likely that during transport they have been vibrated (resulting in splashed mercury). Or they may have been abused in their previous applications. In short we can't be sure of the exact cause of problems, because there are so many likely candidates.

The real cost of cooling is implementing a system that is low maintenance but still energy efficient. Summer temperature excursions exceed the maximum allowable ignitron temperature for full voltage hold-off (~ 30 °C). Additionally, discharging capacitor banks through them at a regular rate can result in significant heating (limited by our supplies to 300 watts). Early tests on our modules indicated that it was possible to heat up the ignitrons 0.8 C° per shot before the addition of resistive anode holding straps. Some form of convective cooling is often required.

The most efficient is often some form of liquid cooling conducted through tubing. Water needs to be purified to be non-conducting. In its distilled state, water is an excellent solvent capable of eroding through the ignitron cooling jackets as well as wrecking mayhem elsewhere. Other fluids exist, such as propylene glycol, which can be used. Unfortunately, many bacteria are capable of metabolizing glycols[42] generating sludge. One serious risk associated with fluid cooling is what happens if punch through from the tube to the cooling channel occurs. The tube (assuming crowbar) can fill with water and on the next shot vaporize water with mercury resulting in overpressure of the tube, followed by mercury vapor in the atmosphere[43]. All these cooling techniques run the risk of atmospheric water condensation. Perhaps the simplest solution is to cool the space, and heat the anodes with halogen lights. Full cooling down to 15 °C (or lower) may be too costly or uncomfortable for the working environment. Obviously, it is more difficult to get adequate heat transfer when the working medium is air; fans may be necessary.

We've found that a 50 watt 120VAC halogen bulb shining on a blackened plate connected to a type D anode is sufficient to increase the temperature approximately 30 C° over ambient. Some ignitrons specify a 50 C° temperature difference for 24 hrs in order to affect thermal conditioning. It is not uncommon for ignitrons to specify a maximum anode temperature (for example 55 °C). Possible explanations for such temperature limits might include the nature of the glass – steel seals and the use of epoxies. The thermal gradient may also result in significant stress at the interfaces between materials. 50 watt spotlight halogen bulbs are quite capable of charring paper. They are a potential fire hazard that must be dealt with carefully since they may have to be on continuously.

Over 120 °C (393 K), it has been suggested that sufficient atomic hydrogen can diffuse from the steel walls to destroy an ignitron within a couple of hours[44]. One would expect the diffusion rate would be governed by an Arrhenius type equation $D = D_0 \exp (-E_d / k_B T)$. However, in practice the physics are a good deal more complicated than this[45]. We are not aware of any model that accurately predicts hydrogen's out-gassing from stainless steel for a wide range of temperatures[46]. Given that caveat, $D_0 = 0.012$ cm^2 /s, $E_d = 0.56$ eV (< 700 ° C 304 stainless steel). This could be used to determine the expected lifetime of the tubes at room temperature (300-1000 x 120°C lifetime). The majority of our tubes are over one hundred thousand hours old. A more accurate experiment at elevated temperature would have to be done (i.e. precise temperature, not just 'over' 120 °C). One might expect that high temperature can damage the tube in other ways that could interfere with results. Stresses can develop at material interfaces with different coefficients of thermal expansion. Given the thermal limitations, a heat gun approach to conditioning is not always wise.

The trigger electronics must be able to repeatedly withstand the high voltages and transients generated by the switch (and its fault modes) while maintaining the capacity to deliver large amounts of power (short of destroying the trigger electrode). Traditional techniques tend to be rather expensive (in order to be reliable). Commercial units are available, one of which we gain royalties from. As such, we

aren't going to discuss the issues and solutions associated with making cheaper, solid-state, trigger units that have been tested for millions of pulses. It's easy to make something that works; it's hard to make something that works reliably. Users should be aware that different ignitrons specify different triggering requirements. Some ignitrons require significantly higher igniter currents than one might naturally expect.

4.7 Copper Vapor Switches

Our interest in alternative high voltage switches has been driven by the high replacement costs associated with ignitrons. Lower inductance switches are appealing to minimize bank voltage loss. Additionally, disposal of the large amounts of mercury involved with ignitrons is likely to become more problematic as society becomes more environmentally conscious. Pulsed Technologies Ltd produces high voltage switches that are marketed as ignitron replacements for pulse applications. They do not use mercury and are thus orientation/vibration insensitive. The trade off is that they are not capable of handling as much charge as ignitrons. Their TDI1-150k/25 (200 kA, 25 kV, 25 kJ) and TDI1-50k/45 (150 kA, 45 kV, 20 kJ) are of particular interest. They operate with a 20 – 60 Pa hydrogen fill controlled by a thermal filament reservoir.

As with the ignitrons they are triggered by creating an arc across a semiconductor surface. Unlike ignitrons, they require a negative potential pulse on their gate with respect to the cathode. They appear to have shorter delay times and less jitter (50k/45 < 40 nsec, 100 shot windowing, 120 KA, 600 joules, 6000+ shots) over their lifetimes than type D ignitrons. Adjustment of heater voltage is required on a regular basis for best performance. Our method of specifying jitter is influenced by the other characteristics of the switch going into conduction. That is to say that a change in the effective inductance or resistance of the switch (or coil) from shot to shot influences our measurement. We do not doubt that our application probably makes it very difficult to record the very low jitters or long lifetimes that are specified as each shot has a large amount of energy available to dissipate into the switch's surface.

Figure 4.7.1. Pseudospark switch pictures. The devices are permanently sealed and easily mounted on Syllac header capacitors for use in oil. The installation at right is a Diala AX oil immersed switch test / conditioning fixture.

Figure 4.7.2. Pseudospark switch tested to destruction. The unit had switched a total of approximately 15 MJ in a ringing application over its lifetime. A switch in a crowbar position lasted roughly 24 kC at just over 4 coulombs per shot. The center hole (lower left) did not originally exist. A copper electrode eventually eroded away resulting in the coating of the trigger semiconductor with copper. The hydrogen reservoir is hidden beneath the plate holding the trigger. The obvious mechanical damage was incurred while breaking open the ceramic switch. Our ignitron trigger units were readily adapted to drive these switches even to this level of copper coating. We do not believe this level of coating was accomplished in just one shot. In applications without perfectly symmetric current feeds (i.e. in a mod) we've seen similar switches come up to air via a hole punched in the electrode corner.

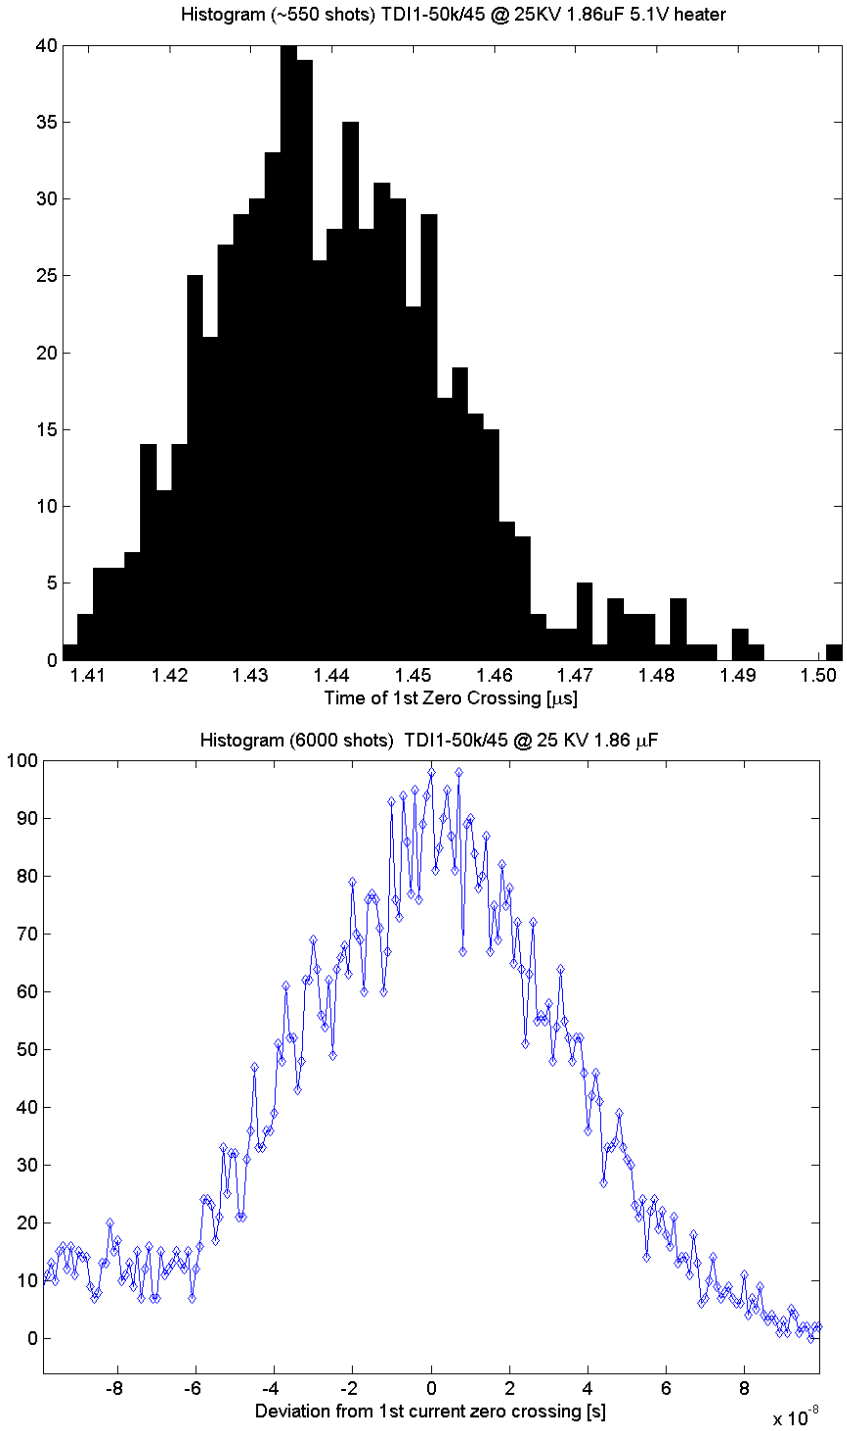

Figure 4.7.3. Apparent jitter of TDI1-50k/45. Circuit was allowed to fully ring down. Refers to time of first current zero crossing (i.e. when the capacitor has reached peak negative voltage).

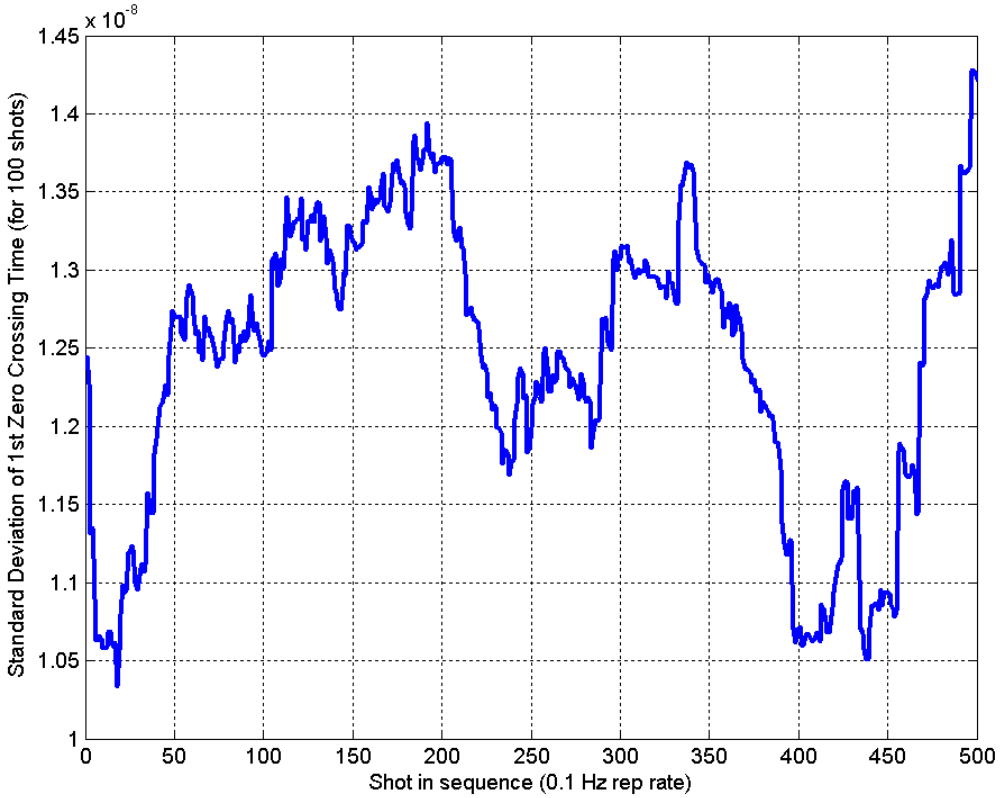

Figure 4.7.4. Early life jitter (starting from shot 200). The jitter can be seen to be less than 14 nsec early in the lifetime of this switch. Note that this is determined using a 300 MHz digital oscilloscope. These small jitter times are never recovered throughout the rest of the life of the switch.

The basic modes of bad behavior are the same as ignitrons. When the heater voltage (and hence the hydrogen pressure) is too low they do not hold off the rated voltage. Too high a heater voltage produces the same result. They tend to go into an arc mode rather than glowing. At the tuned heater voltage, switching is capable over the full range of switch voltages. They do not typically have to be cooled (unlike ignitrons), but they do have to be electrically conditioned after receipt (their required heater voltage changes very noticeably over the first couple hundred shots). It is easy for them to become too hot via switched current, resulting in higher gas pressure and breakdown.

The new complication they add to the picture is the need for a high voltage isolated power supply providing less than 20 watts from 3 to 7 volts. This is easily accomplished via a flyback switching power supply that is cheap and small. Again, this is something sold as a commercial unit. Our ignitron igniter units are easily modified to drive the gate at 6 kV. For some switch models it may be possible to realize better performance with the addition of external components. The goal is to maintain the grading grid electrode at a desired potential. This problem can be thought of as the same as that with any two stacked switches.

The mechanical arrangement of the switch lends itself very favorably to a stacked start switch / crow-bar switch arrangement on top of capacitors. This is very convenient when oil immersion becomes necessary. Various papers are available that describe the physical operation of these switches[47,48]. However, such a discussion is not necessary to appreciate their external behavior.

Additional testing of TDI1-50k/45 has taken place in FRTP applications. Testing was done on a small system consisting of 4 driven coils. Individual coils were wound as half turn in parallel units at 30 kV charge voltages (single ended). If there were no stray inductance this effectively results in 60 kV across each driven coil. Some shots were fired at 35 kV charge voltage. Each coil had two reversal capacitor / switch arrangements (one for each half turn, 40 kV & 4 µF). A bias power source (10 kV & 500 µF) and ringing theta power source (40 kV & 1.85 µF) would each power two coils. The capacitors operating over 20 kV were immersed in Diala AX oil.

The most important result from the experience was that around one hundred shots were required for the TDI1-50k/45 heater voltage and switch characteristics to stabilize sufficiently to not result in serious fault modes. The worst fault modes were of course when the crowbars conducted at the same time as starts. With 18 heaters (for 4 coils) requiring adjustment (fortunately not all, every shot) this is a lengthy process requiring that shots be fired and analyzed). Fortunately, our heaters are adjusted via multi-step dials. The conditioning phase was worrisome and time intensive. Some 300 shots were taken before disassembly and integration into another experiment.

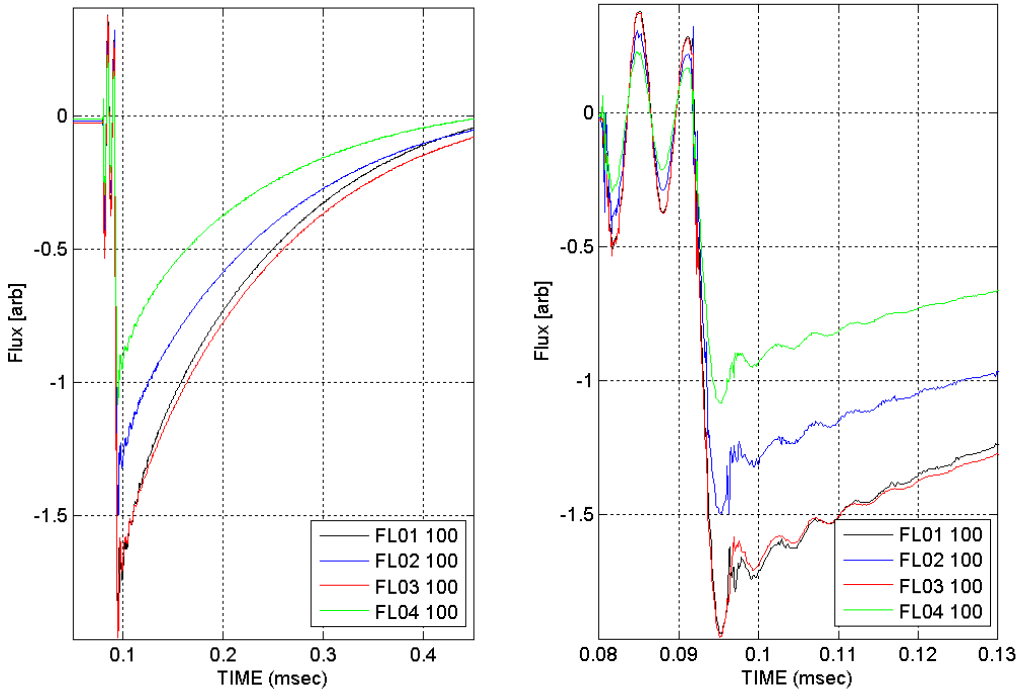

Figure 4.7.5. Copper vapor switch performing well with ringing theta @ 30 kV. Coils 1 & 3 are being fired. Coils 2 & 4 are not being fired. The decrease in field 1 & 3, then 2, then 4 is as expected. The diagnostics are 'thrown together' resulting in not the best immunity to noise sources and linearity. The ringing theta unit is fired at 0.08 ms followed by the four reversal switches at 0.092 ms and four crowbar switches at .096 ms. Peak reversal current is about 50 kA (with 30 kV charge). Approximately 5 coulombs is carried by each of the crowbar switches. The ringing induced by the ringing theta unit can be seen to continue after the firing of the crowbar units (not unexpected given their relative locations). This realization changes the interpretation ascribed to the fall from peak field. The effective resistances of the crowbars vary somewhat shot to shot. Occasionally during the conditioning phase, one crowbar breaks down too early, resulting in much higher amplitude ring with lower peak field. The bias isolation inductor had failed at an earlier date (due to cheap construction).

4.8 Capacitors

The various capacitor manufacturers, General Atomics, GE, Maxwell, etc. do a very good job of accumulating and providing the appropriate literature to select and engineer systems with their products. They include publications describing how voltage reversal affects capacitor life[49] as well as descriptions of the various non-ideal effects in practical capacitors[50]. We have already covered the safety concerns in the earlier safety section (3.9.4-3.9.6 capacitors). The high voltage review paper highlighted earlier provides most of the information required to make and understand capacitors. Suffice it to say that given the wide variety and ready availability of commercial components it makes far more sense to specify and buy what is needed than to 'do it yourself'.

4.9 Voltage Monitoring

Every capacitor module with significant energy storage needs to have a voltage monitor. Preferably, this is connected in such a position that having an individual capacitor becoming disconnected from it is highly unlikely. For our purposes the voltage monitor needs to fulfill several distinct requirements:

1) It must be continuously on.
2) It must be insulated from nearby equipment to such an extent that it is inconceivable of it being arced to
3) It must report the voltage back to the control system via fiber sufficiently fast for the control system to respond to aberrations.
4) It must give local indication (audible and visible) of the presence of dangerous voltages

LSX utilized bank voltage monitors that were effective and convenient but could not satisfy our expanded requirements. They had several very good features. They were continuously on requiring no power source other than the high voltage they were monitoring. They were isolated and compact, feeding the voltage information back via fiber. They responded to the presence of reversed voltages as well as normal polarity. However, they did have several critical flaws. They did not give any local indication of the presence of dangerous voltages. Their output indications were not linear with monitored voltage. Their response times were slow, operating at around 1 kHz at the top of their range (100 Hz more typically). In summary, they were adaptable to most of our needs but physically could not fulfill the requirement of giving sufficient local indication (simply not enough coulombs).

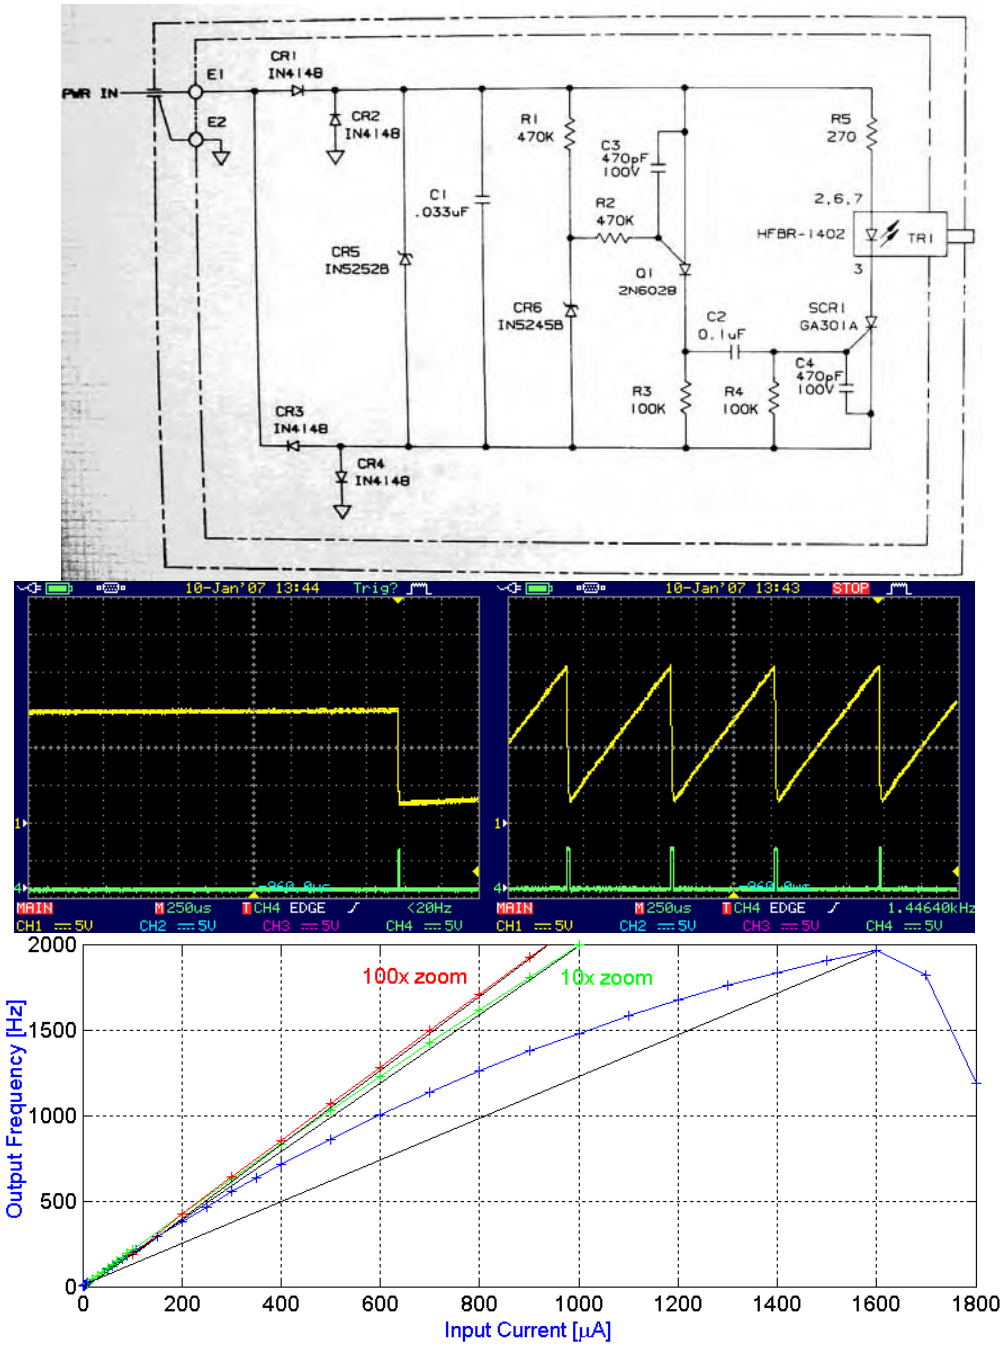

Figure 4.9.1. LSX voltage monitors. Power supply voltage monitors were also used. (top) Original schematic for voltage monitor boards. (middle) The yellow trace is the voltage across C1 while green represents a digital equivalent of light from the output. Note the difference in on time and trigger voltage. (left) Also visible is the difference in charging slopes (despite a constant current source) from beginning to end of charge. (bottom) The nonlinearity over the operational range of the device is apparent. We have chosen to control the input current rather than changing input voltage with some fixed resistor. Note that at high currents the output collapses to 0 Hz. As such, normal operation is confined to a narrow range.

The first 4 diodes act as a bridge rectifier followed by an overvoltage limiting zener diode and integration capacitor. R1 acts to bias CR6 which sets the trip point voltage for the uni-junction transistor Q1. When current starts to flow through the P-N junction of the anode of Q1 to its gate, the UJT locks into conduction from anode to cathode until C1 is mostly discharged. SCR1 acts as a buffer for driving the fiber output. This circuit's linearity is governed primarily by the non-ideal characteristics of Q1 and the variable on time of SCR1. When the input current is sufficiently high SCR1 can stay in the on condition indefinitely. The rate of rise of voltage on C1 influences the switching characteristics of Q1. This circuit is less linear with changes to input voltage monitored by a high value resistor than it is to input current. Additionally, there is a minimum input voltage for oscillation (approximately 15 to 21 volts).

This approach used on LSX was the 'traditional' programmable uni-junction transistor (PUT/UJT) oscillator. It is inherently difficult to make linear over a broad range. It is however very cheap and requires no external power source. For lower voltage operations they suffer significantly from fluctuations in input current (oscillating 20 V offset voltage). We need a device that indicates voltage presence down to ~1/1000th of peak voltage.

Such 'unpowered' approaches are limited by the very high voltages being monitored with the consequence of needing high input impedances. The power available is limited to less than the power dissipation of the monitoring resistor by the ratio of internal voltage to monitored high voltage (plus a factor of ~2 for oscillating internal voltage). With the designed capacitor size each pulse has approximately 7 µJ available for indication. Increasing capacitor size lowers frequency of indication. For comparison, our battery-powered design uses about 4 mJ in audible and visible indicators. Replacing R5 with an inductor could allow the SCR to turn off better, allowing a higher input current limit. But fundamental to this design is that the indication pulse is at a non-constant frequency. So required power = 4 mJ * frequency is a serious limiting factor. The lower limit of indication (every 4 seconds for visibility) prevents us from running extremely slowly. So, at 250 Hz (max of range) we would need at least 1 watt. The voltage of operation (20 V) to monitored voltage (20 kV) ratio at the same point in the range means that the high voltage resistor would have to be sized for 2 kW. The tricks that could be played to resolve this basic problem using magnetics rather than a resistor require a high voltage opening switch. It is simply impractical to add indication without a separate power source.

Design and construction of a battery powered unit built around an AD654 voltage to frequency chip is a straightforward affair. A constant oscillator and comparators drive the high voltage and low battery voltage indicators. A word of caution as to construction of extremely low current devices: circuit boards can leak not insignificant currents at low voltages on the surface when not properly cleaned in de-ionized water. Our design did not require a conductive shield in order to work properly. Intriguingly though, corona 'packets' could on occasion be detected bypassing the high voltage resistor. Adding a capacitive filter to the input of the low

voltage board made the problem worse as one might expect. Prior to this addition the rail diodes acted to shunt the excess charge off the input; and the RC timescale of the input was fast enough such that the output frequency was impacted for only a few periods. The obvious answer is to pot the high voltage resistors (same as high voltage probes) so that any corona packets reach ground before reaching a signal conductor. This was done after the run period described in this document.

Ideally, the voltage to frequency board is mounted with no lead length to the monitored high voltage. Because of the necessarily high input impedance (to keep the monitoring resistor's power dissipation down) coax cable is especially susceptible to currents on the shield showing up as signals; oscillatory noise is rectified by the input bridge diode. While differential approaches can be made to work using twisted pair coax (twisted pair within a coaxial shield that is chassis grounded) this roughly doubles the expense. Putting capacitive filters in place is viable once the corona 'packet' problem has been dealt with.

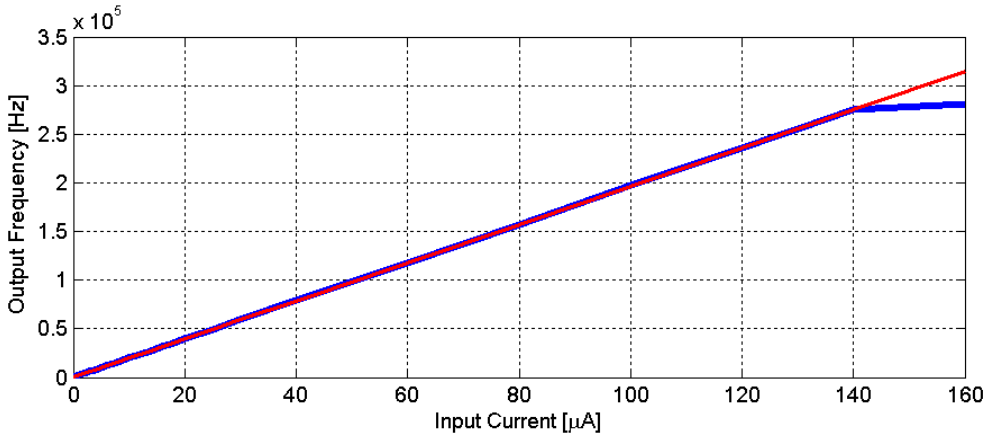

Figure 4.9.2. PHD voltage to frequency converter characteristic. Deviation from linearity is less than 0.1% over the range of 3 μA to 130 μA. Effect of exceeding input current limit is saturation, not collapse. A very low frequency output with no input serves the purpose of indicating presence and power to the control system. Input impedance is set for 40 kΩ as a matter of convenience. 5V = 250 kHz. With a 200 MΩ high voltage resistor this results in 1 KV input yielding 10 kHz.

The final design could run for 2 ½ months on 2 D cells. The primary consumer of power was the AD654. Lower power consumption is possible by sacrificing output linearity and noise immunity. Keeping the high linearity has allowed us to use this board in several diagnostic applications (such as gas pressure monitoring) without serious modifications. Solar cell battery charging could be used to extend battery life, but not under warehouse fluorescent lighting conditions. A high voltage insulated power supply is also possible without too much trouble. Super capacitors can be used to maintain power for several minutes after power removal or rechargeable batteries can be used. So far, we have elected to replace all batteries every 2 months.

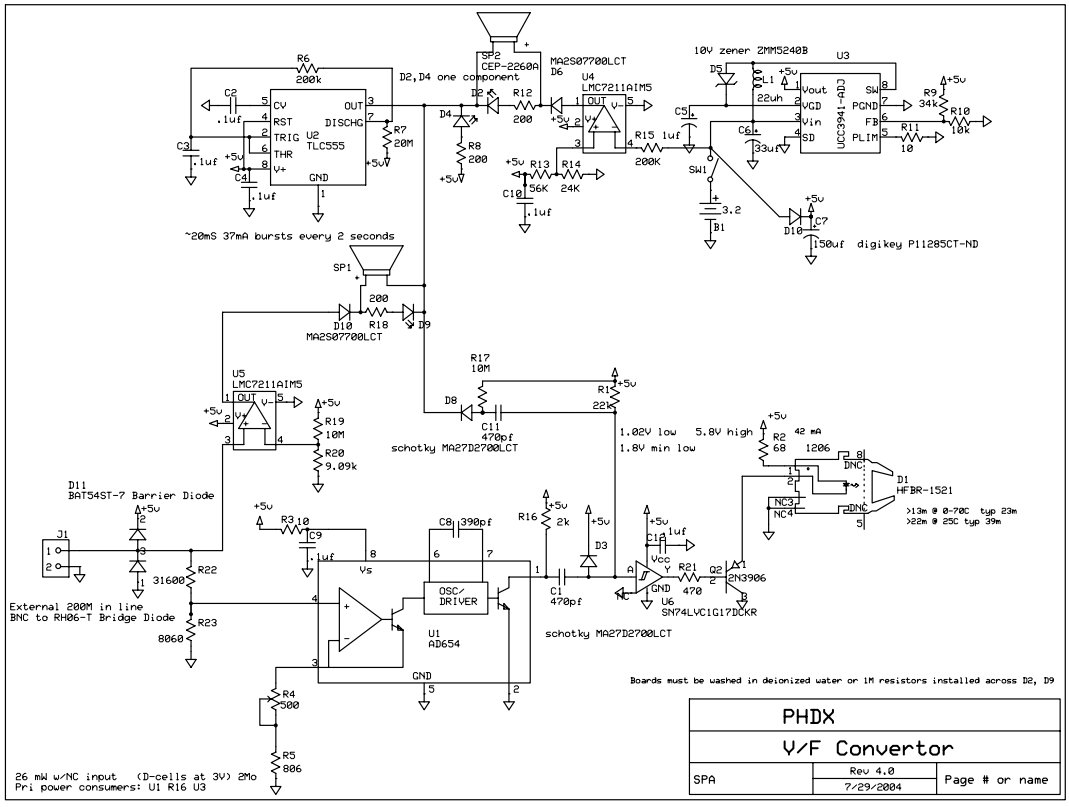

Figure 4.9.3. Voltage to frequency converter. The final design includes a metallic container and BNC input.

The 'heart beat' indicating presence to the receiving controller is not always present via the natural offset of U1 necessitating additional components. D8, R17, and C11 would accomplish this except for the occasional state of U1 where the output acts as low impedance to a negative output voltage (C1 charged to 5 V). Solutions

are removing D8, R17, and C11 by placing a high value resistor from +5 V to J1:1 (leaving the original component positions empty) or adding a low voltage drop diode in line with the output of U1:1. Note that +5 V is really set for +5.5 V (U1 has a heavily limited input voltage range dependent on supply voltage). U3 is in the lifetime buyout phase; new units will have to be redesigned.

4.10 Bias Modules

Additional modules are required in order to establish the initial 'bias' magnetic fields that determine the flux available to be trapped in the FRC during reversal by the 'reversal' modules. The primary requirement is that they have sufficiently long periods for plasma to be translated or formed near the wall for the length of the intended FRC. To accomplish this, an individual module is assembled from four 360 μF 5 kV (100% reversal) capacitors in the same style as the reversal modules. Inductance is placed inline between the bias module output cable and the reversal module outputs. This is accomplished by a 3 turn, 36 cm diameter coil mounted to each reversal module and partially shorted out by the module's frame (effectively 6.8 μH). The bias modules have no crowbars. The lack of crowbars may adversely impact the life expectancy of the ignitrons due to reverse currents.

Figure 4.10.1. Bias modules. (left) - polarity. (right) + polarity. These pictures were taken during reconstruction after the latest move. Some parts are missing.

As with all high voltage banks, the failure modes and stresses need to be considered. Due to the high capacitance and reversed polarity (bias vs reversal) there is no danger of the high voltage modules creating a significant overvoltage on the bias capacitors during failure modes. The same is true for over-currents because of the isolation provided by the inline inductors. However, the bias banks do ring and carry significant current (~ 120 kA) and charge (~ 7 coulombs for the first cycle). The high reversed currents can be expected to damage the ignitrons (that are only rated for 7.5 kA reversed current). Ignitrons in the bias bank have failed completely (only holding off ~300 volts). The ones that are still in operation can only be reliably charged to 3 kV before glow discharges and pre-fires become issues. One would expect that significant damage has been done to the graphite anodes.

An additional concern is the potential for a common mode voltage between the low voltage terminal (coax shield) and the chassis of the capacitors caused by the higher voltage reversal modules. For a fault mode to occur capable of causing this damage requires an arc from a driven coil to earth ground. Experimentally, this has not yet resulted in any failures despite the frequent occurrence of varied fault modes.

4.11 Cone Magnets

On the 40 cm experiment the end regions' magnetic field required no separate coil design. A cusp was accomplished using the last few axial field coils driven by a bias bank pair. The metallic plate at the end was beyond the region where the cusp was established.

Figure 4.11.1. 40 cm experiment axial field coil layout. (left) One bias module +/- pair was used to drive the three coils at left in series, parallel with the opposite end (smaller diameter coax). This formed a cusp with respect to the bias field in the center section (small diameter coax). Two bias module +/- pairs are used to drive the center section coils. The central coils are driven with 2 coils in series, all the sets in parallel due to the bias connections. One bias module pair drives the +z center section; the second pair drives the –z center section. One coil is left between the end coils and the central coils unpowered. The small diameter cable from the central section connects to 2 pair of bus plates (right, top) on top of the wood structure. Before shot 236 (8/22/2005) the two bias module pairs were connected directly to this buss-work. After this date they were connected through the reversal modules, as with the 80 cm experiment. (right) The buss work connecting all the central section coils was left in place. Also visible is the air cooling manifold for the ignitrons (large diameter black PVC) and the reversal bank cables (large coax bundles wrapped in polyethylene). An exhaust duct (white) for a roughing vacuum pump enclosure is visible.

The 80 cm experiment required separate coils because of the shorter aspect ratio of the device and the nearby presence of stainless steel plated end cones. Initially, surplus FR-4 potted magnets were used in series with wrapped coils (at

small diameter). The two end cone magnet sections were in parallel with each other. Upon firing the reversal banks (shielded by the stainless steel plating of the cone from the end cone magnets) the potted magnet on the pump end arced internally (shorting out), resulting in substantially increased current in the remaining small-diameter wrapped coil. The non-potted coil was wrapped on a metal form that proceeded to spark across its joint.

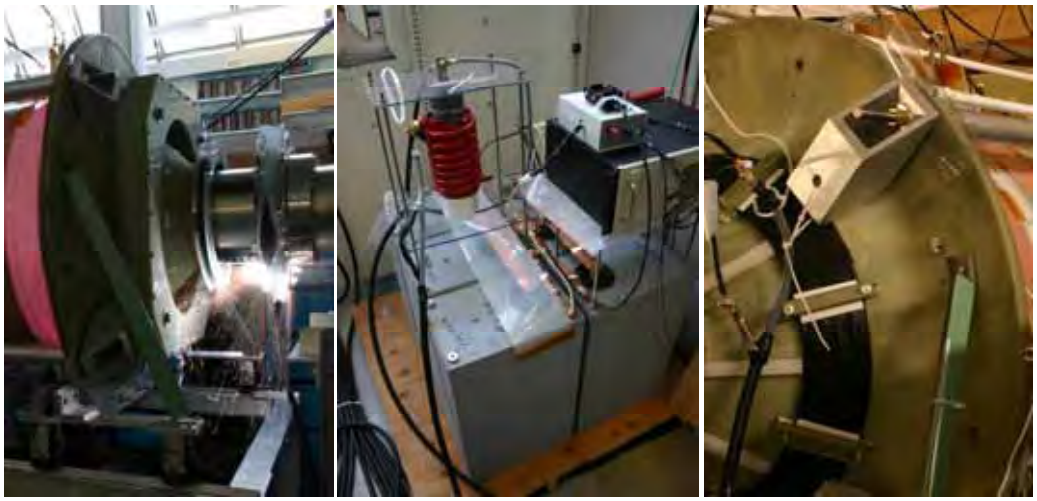

Figure 4.11.2. Original encased cone magnets replaced by 10 turn coil. (left) The original unit arced internally due to firing of reversal banks resulting in additional current through a coil connected in series. The fiberglass coated cone (45°) has a 60 mil thick layer of stainless steel along with a ½ in^2 (rotational cross section) stainless steel vacuum surface O-ring joint to the fused silica. The inner diameter is 28.3 cm. TIG welding of the cone to the sealing lip along an outer edge has resulted in significant debris being left in the gap between the two sections. (middle) Cone magnet capacitor bank. 1.35 mF (total) high energy density storage capacitors are limited in voltage by the amount of reverse voltage they experience (for a given lifespan). These capacitors are rated for 7 kV at 90% voltage reversal, 10 kV at 75%. Typically we run around 4 kV. The snubber that in this case sits across the ignitron is hidden behind the bank. The rate of current rise was sufficiently low that the snubber was required to provide a boost in ignitron ionization current. (right) The black replacement cone magnet is connected in series with the opposite end.

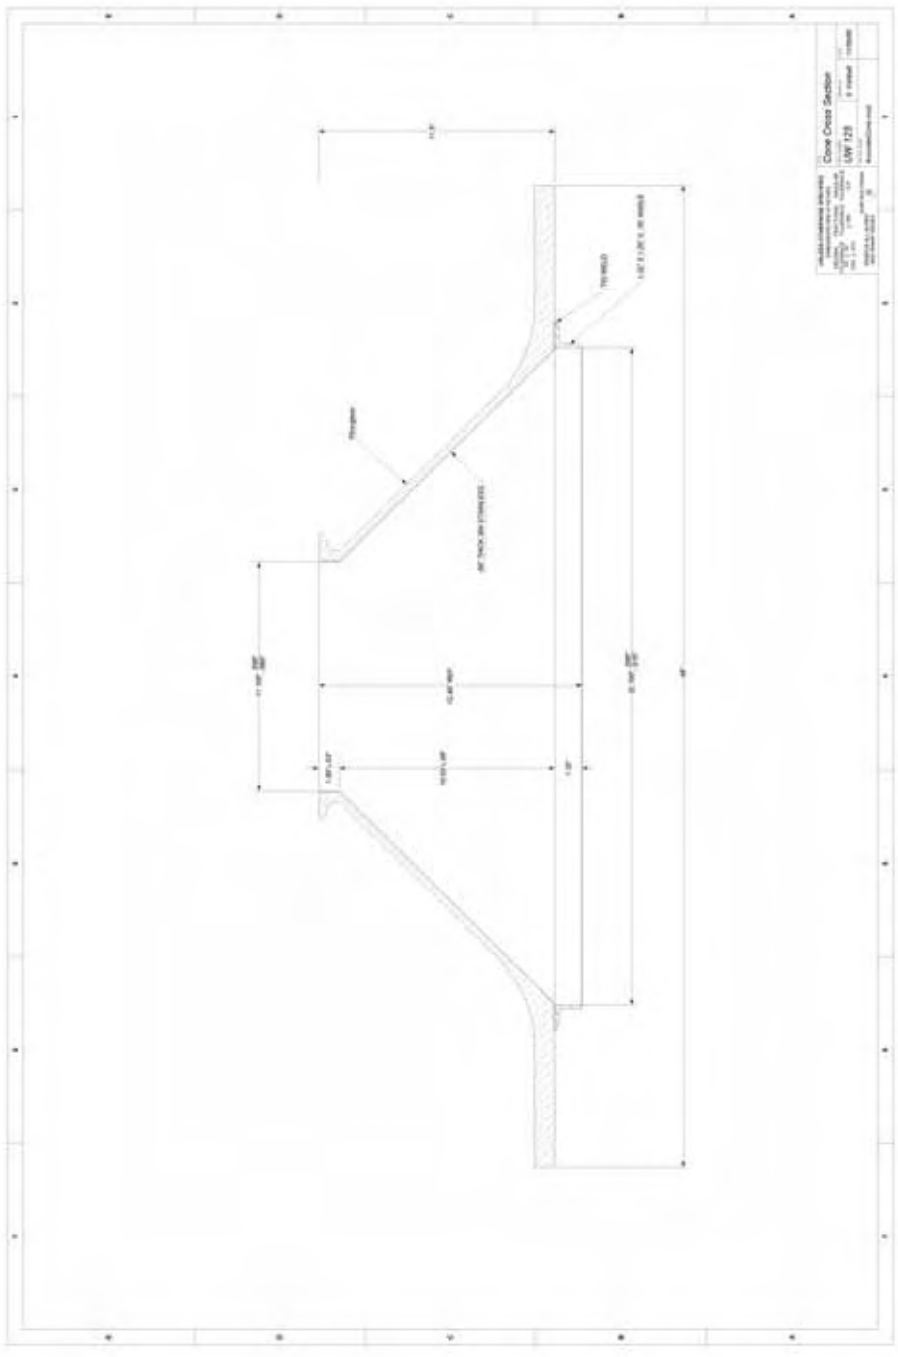

Figure 4.11.3. Diagram of original 80 cm end cones. The TCS project is the source of this hardware (diagram copied with permission of Scott Kimball).

The potted magnets were quickly replaced with 10-turn magnets wrapped on to the FR4 cones; adhered with fastenings to epoxy. The center of the new conical magnet is 40 cm from the center of the last main field coil. The average radius of the coil is about 22 cm. The axial extent is 8 cm. In the new arrangement the opposite end cone magnets are in series with each other and driven by one 1.35 mF bank. A module snubber (.5 μF, .22 ohms) is required to limit the coupled voltage (coupled from the reversal bank) and provide a boost to the initial rate of current rise in the ignitron. The rise time for the bank is 1 ms. The bank can ring or not, changing from shot to shot (sometimes the ignitron opens on zero current crossing). The second peak (negative) yields approximately half the first peak fields.

The magnetic field penetration of the cone by the plug bank was measured (with spatial precision) for the original setup with the potted magnet. The destruction of the first plug coil has made this information worthless. Since the replacement was done while operational a repeat calibration for the new coil has not been done. People objected to putting any diagnostics on this section, so all information is calculated based off the intended bank charge level (which is an approach obviously prone to error in multiple ways). Approximately 10 mWb is recorded underneath the first reversal coil when the plug bank is fired at 4 kV [724, 725] alone. At 4 kV the plug bank is capable of providing approximately 60 mWb equivalent (based on the closest flux loop, [725]) at the outer rim of the cone. This 'equivalent' is only intended to simplify some calculations as it isn't a proper way to specify the coil for spatial distribution. Simple bank calculations lead us to expect approximately 100 mWb beneath the plug coil itself.

4.12 Z-PI

Axial discharges have long been one of the simplest ways to create an initial plasma that can be further ionized during bank reversal. Their price is non-uniformity in the generated plasma and the possible introduction of axial currents that remain in the plasma. Axial discharges also generally require an even distribution of gas. The usual way to accomplish this is via a gas backfill that must be created well before the shot commences. This operational mode is undesirable for translation experiments. It isn't necessarily a show-stopper, but it certainly adds significant timing limitations.

Our use on the 40 cm experiment included two different levels of axial discharge. The first axial discharge (pre-pre-ionization, PPI) was low current (600 A, 19 mC), 19 kV, and non-ringing. The high voltage switch used was a thyratron (surplus from a previous experiment). The second axial discharge (second pre-ionization, SPI) rings with an available energy of 2.5 kJ. This unit uses an ignitron switch. The SPI was intended to strike back along the wall of the vacuum vessel to create ionization at the wall in as uniform a manner as possible. This is similar in approach to that used on TRX-1 and TRX-2[51]. Judging from the light emissions, this did happen as intended. The SPI is only fired with the axial magnetic field cusps in place. SPI and PPI use the same axial current feeds. The four over-bolts that support the structure are used for this purpose. Each axial electrical feed has a radially aligned 8 μH coil beyond the conductive endplate on the high voltage electrode end.

The ring frequency of the SPI is 17.5 kHz (which is an order of magnitude lower frequency than on TRX-1/2) [269].

Figure 4.12.1. Axial discharge arrangement on 40 cm experiment. (top-left) The radially compressed non-ringing PPI can be easily seen. (top-right) The PPI bank was inherited from LSX. (bottom-left) The SPI lights up the whole tube. (bottom-right) Highly distributed stainless steel electrodes help to result in an azimuthally even discharge.

The experiments with the 80 cm tube were not specifically designed to include the possibility of an axial discharge. The larger experiment must run at lower pressure to obtain the correct inventory of 2.5 x 10^{20} ions. At the lower pressure of 1 mTorr the gas is significantly harder to break down via axial discharge (the axial length is roughly the same). It is a simple Paschen curve P d problem. No special electrodes were placed in either of the end regions. As a consequence, axial discharges were only used via the stainless steel end cones when throwing the 'kitchen sink' at the problem of insufficient lift-off flux trapping. Predictably disappointing life-time results ensued (although more flux was indeed trapped). The

eight axial current feeds fed through holes cut into the interferometer table. This resulted in an increase in vibration of the interferometer.

Ring frequency of the axial discharge was 24 KHz. Maximum current was 25 kA @ 15 kV. A few shots were taken at 25 kV resulting in unintended breakdowns. Characteristic envelope decay timescale was 30 μs.

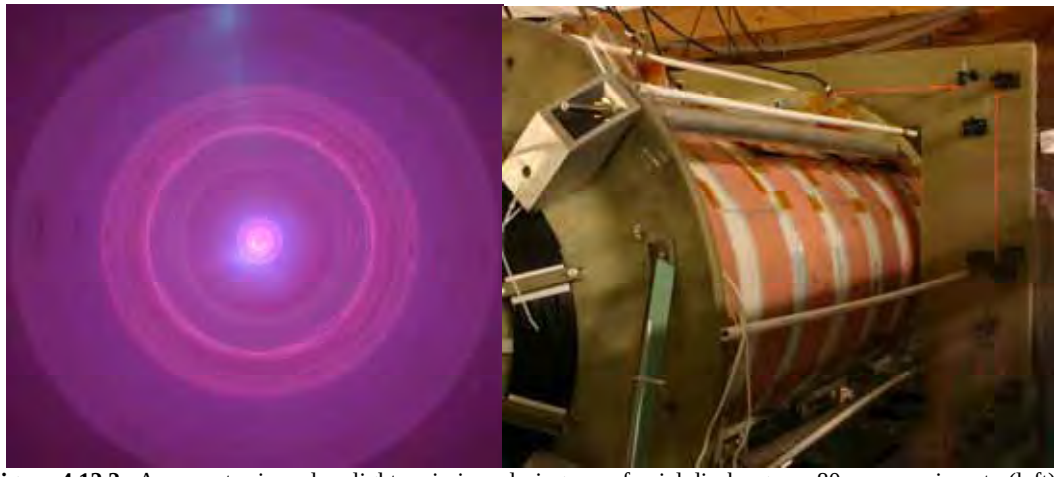

Figure 4.12.2. Asymmetry in end on light emissions during use of axial discharge on 80 cm experiment. (left) Note the bluish light at the top of the frame. This emission (at random angles, not always as concentrated) tends to be present when the axial discharge is fired. (right) The white PVC tubes contain the axial discharge current feeds.

4.13 RMF

With the components already outlined it is simple to construct any high current / high voltage decaying ringing system < 40 kV. Rotating transverse Magnetic Fields are commonly used in our experiments to enhance the initial levels of ionization pre-reversal. Additionally, their use may enhance diffusion of the initial plasma across the bias fields. In every case on PHD, the antennae are configured as two phases of Helmholtz pairs. Obviously, this form of RMF can not be considered 'pure' due to the high frequency components present due to sudden start up and the delay from the start of oscillation to rotation. There are also counter rotating Fourier components (at higher frequencies) associated with locality of the antennae. Some sources[52] suggest that a transverse oscillating magnetic field (OMF) should be sufficient to drive currents in one preferred direction without the requirement for two phases.

The original antennae (on the 40 cm quartz tube) were generally made of soldered copper pipe in an attempt to maintain a low non-coupled inductance. This had the unfortunate byproduct of increasing the feedback from initial plasma formation to relative antennae phase. The two oscillators would almost immediately lose their 90Δ degree phase difference so that trial and error was necessary to obtain the desired phase differences. A diamagnetic rotational direction seems to result in

better ionization (8% [346] vs 2% [347] +/- 2% via interferometer, both at 220 kHz, 25 kV) when combined with a low level axial discharge (19 milli-Coulombs, 19 kV, 600 amps). A consistent effect on trapped flux can be seen over a range of phases (90-155) without any strong dependence on the precise phase relationship. On average a diamagnetic rotational direction doubled the trapped flux (as compared to the paramagnetic rotational direction) with the RMF occurring 80 µS before reversal.

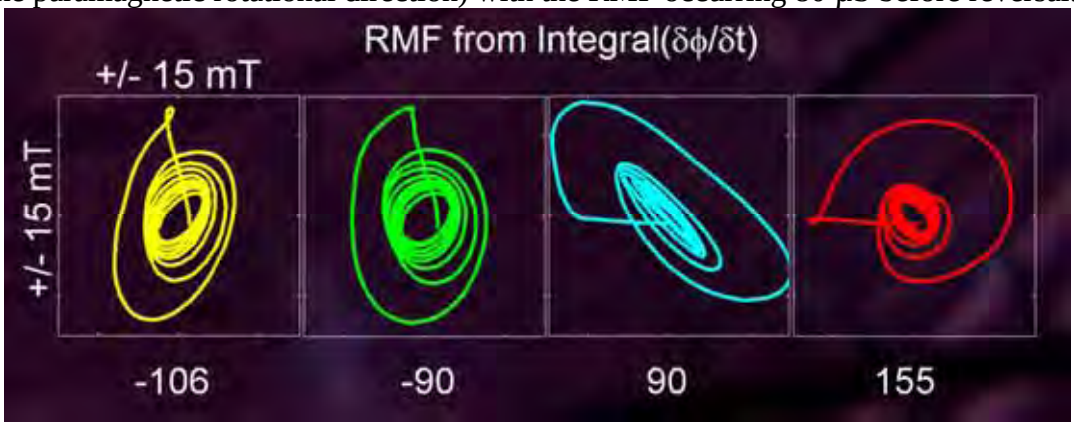

Figure 4.13.1. RMF magnetic fields on the 40 cm quartz tube. The desired 90 degree phase difference should be indicated by a circular spiral (as opposed to an elliptical one). (+) indicates the direction one would expect to be diamagnetic with respect to the bias fields (if the RMF rotates with the electrons).

Experiments were run with the RMF at different frequencies and with different drive sources. The purpose was primarily pre-ionization. It should be understood however that the temptation felt by some to use RMF for full formation has resulted in work being done beyond the scope or interests of this document. This author does not wish to get involved in any of the community conflicts.

4.13.1 Saturable Core RMF

It is possible to make a driven RMF system without the use of triggered opening switches. Such devices are two terminal non-linear components. As non-linear components they can shift energy from one frequency to another. For them to be efficient, the non-linear component must be non-dissipative. One of the most commonly encountered is any form of non-air-core inductor. At some level of current the inductor saturates and its inductance decreases. When the current reduces back below the threshold, the inductance returns to its pre-saturated state. In simple terms, it is an opening and closing switch dependent on the amount of current.

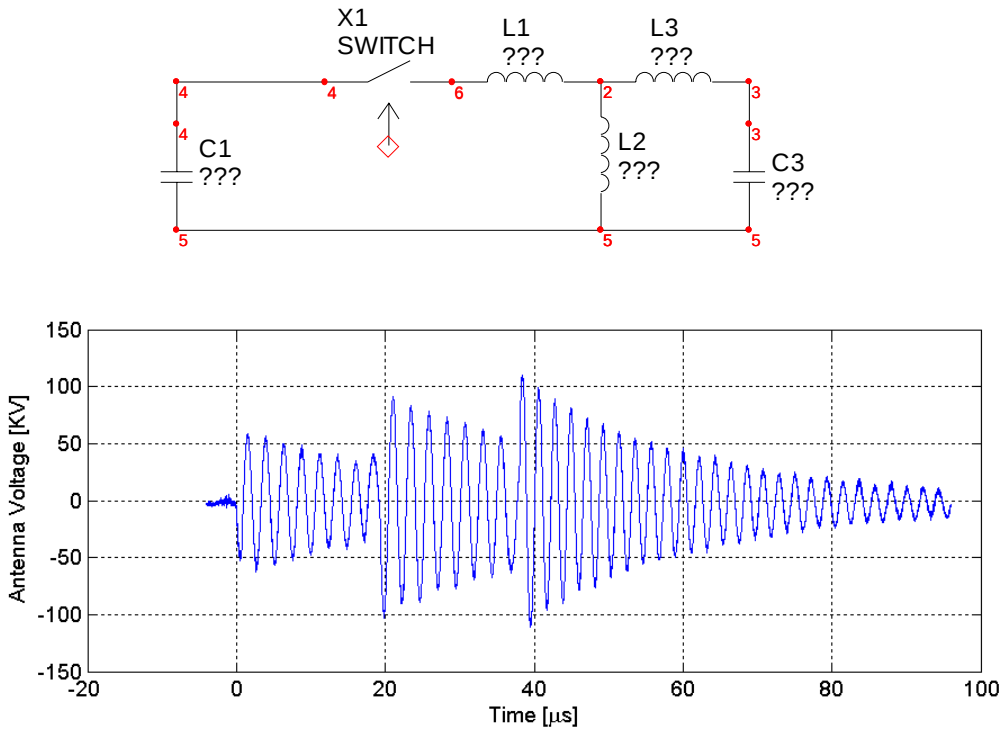

Figure 4.13.1.1. Saturable core RMF diagram. (top) The primary low frequency oscillator of C1, X1, and L1 drive L2. X1 can be an ignitron or other closing switch. L3 and C3 comprise the high frequency RMF series LC oscillator. Near zero current points in L2, the impedance of L2 increases, resulting in a short applied voltage spike to the high frequency oscillator. (bottom) The experimentally obtained antenna voltages can be quite impressive (110 kV, ~ 420 kHz from a +/- 25 kV module pair).

However, it is often omitted that such a system is highly reliant on the characteristics of the magnetic material. Remanence, hysteresis, frequency response, core losses, etc., all have critical roles to play. The magnetics can remember earlier states resulting in non-repeatable results. The ignitron used to provide the low frequency oscillation loop is also an evolving element. Coupled with the requirement for two balanced phases the idea quickly becomes impracticable for quick implementation. It is perhaps an interesting question to ask how much of the non-repeatability is due to the ignitrons versus other factors. I.e. could a more careful design take advantage of this approach? One might be tempted to try again until remembering: these tests were done without a plasma load. Also, unintended arcs were a routine occurrence.

4.13.2 Frequency

RMF was applied to the 40 cm experiment at several frequencies (25 kHz, 220 kHz, 250 kHz, 300 kHz, 344 kHz, and 425 kHz). At the lowest frequency we expected the RMF to act on ions as well as electrons. For the most part, the low frequency (high field) RMF was suspected of simply guiding the axial discharge location. The high frequency (lower field) capacitor values were far more successful.

Operating in the diamagnetic direction (with respect to bias field) with the PPI (600 amp, 19 mC axial discharge) running the RMF (220 kHz) was capable of causing 8% ionization as compared to 2% in the paramagnetic direction (neutral gas pressure: 10.5 mTorr). The ionization, % based on the interferometer, requires knowledge of the location of ionization. We assume based on the end on imaging light emissions that the majority of ionization is at the wall. This is equivalent to the stated ionization percentage being an upper limit. We must also assume azimuthal symmetry.

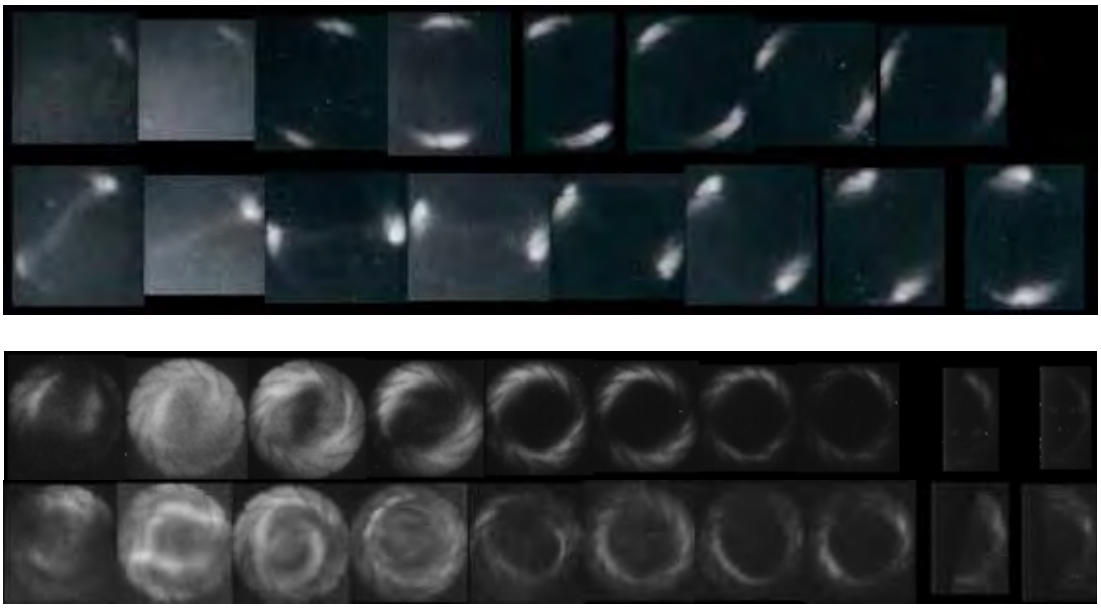

Figure 4.13.2.1. (top) 25 kHz, 30 kV. End-on emission pictures during the pre-ionization phase with an axial current present. 2 μs framing interval results in no aliasing. First set is with counter clock-wise rotation; second line uses clock-wise rotating RMF. (bottom) 220 kHz, 25 kV. These pictures are aliased. First set is with CCW paramagnetic rotation. Second set is with CW diamagnetic rotation (in bias field). The diamagnetic set's RMF phase shifts to an oscillating magnetic field (OMF) near the end.

4.13.3 80 cm PHD RMF

The antennae on the 80 cm PHD tube are made of RG-58A/U core while the feeds to the antennae are coax. The coax portions are 100" long. The increased uncoupled inductance associated with small diameter wire has made the systems' relative phase predictable. The system consists of 4 loops per side for a total of 8 loops in parallel per phase with around 450 nH of total inductance. 30 kV capacitors totaling 1.0 μF per phase are used in series with a type A ignitron to oscillate with the antenna at 235 kHz. The peak ignitron currents are less than 50 kA. This simple set up has lent itself well to operation at different frequencies (unfortunately coupled with different field strengths). The results are presented in chapter 7 due to the coupling of this system with the other sources of ionization.

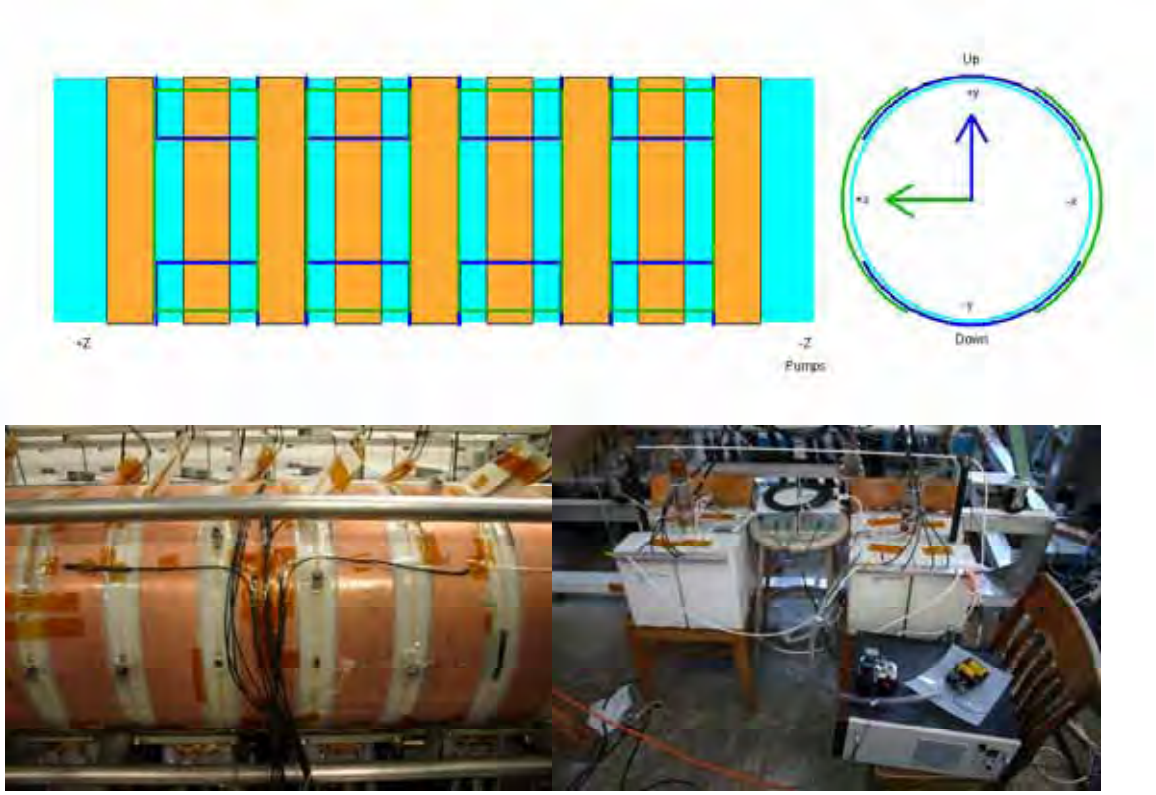

Figure 4.13.3.1. Pulsed RMF setup. (top & left) The presence of the RMF antennae is difficult to see in pictures of the experiment. The locations of the vertical antennae are indicated by the blue lines in this diagram; horizontal antennae are green. All coils for one phase are in parallel. Initial magnetic field direction is indicated by arrows. Each coil extends for 120 degrees resulting in a crudely Helmholtz arrangement. Significant local field effects can be expected at large radii and near the ends. (bottom) A simple pulsed RMF bank setup. Type A ignitrons mounted on top of 1 μF capacitor banks. The capacitance is easily changed in units of 250 nF. For test purposes ,a puff valve driver is visible connected to the vertical bank. This is being used to test the polarity of the diagnostic pickup loops. Multi-turn non-integrated (before recording) dΦ/dt loops are in place on the experiment, one per phase [RMFXLoop, RMFYLoop].

At these frequencies (and low energy levels), the condition of the ignitrons is capable of altering the behavior of the systems. The two phases are set up to charge to the same voltages (not independent). The vertical phase ignitron consistently produces a stronger field due to lower losses in the ignitron.

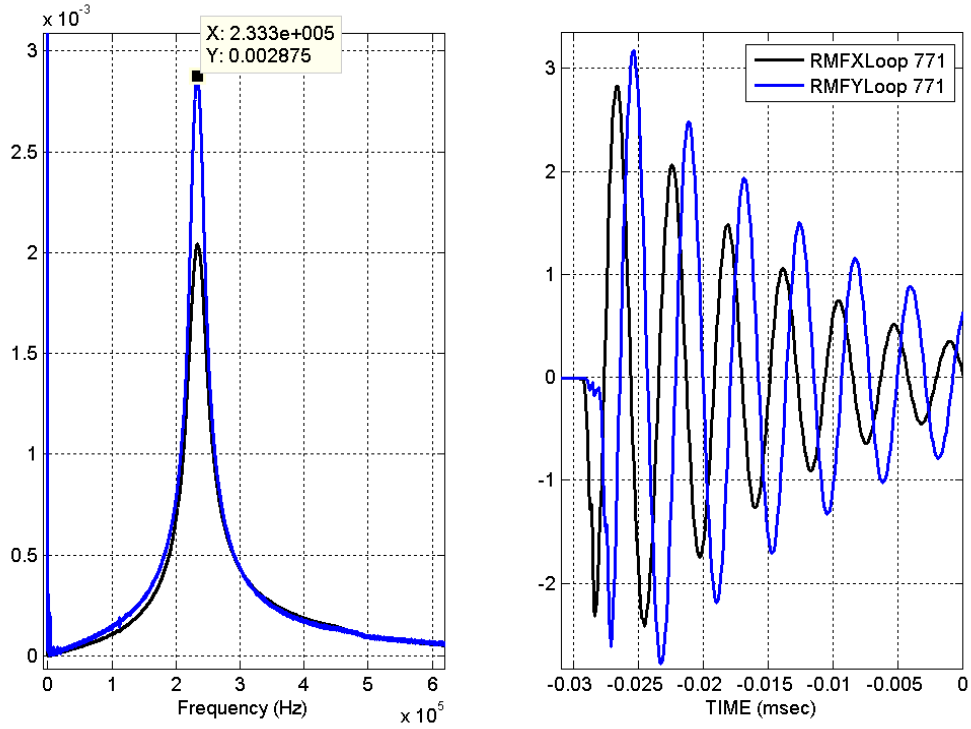

Figure 4.13.3.2. RMF characteristics as run on 80 cm experiment. (left) Fourier spectra of RMF pick up loops. (right) RMF pick up loops showing the consistency of relative phase when fired in vacuum.

4.14 Plasma Guns

Introducing a multi-pole barrier field on the walls could be expected to improve the retention of flux during the lift-off phase. Barrier fields can support the plasma pressure off the wall before reversal is fully accomplished. However, the use of a multi-pole barrier field during reversal can be expected to create its own problems. When the barrier field reaches a level sufficient to keep plasma off the wall a rotational mode is strongly induced[53].

We suspected that stability is very dependent on the uniformity of initial ionization. The logical ordering to confirm our suspicions was to improve pre-ionization via other means and then add barrier fields if absolutely necessary. RMF at sufficiently high frequency can be thought of as similar to barrier fields, but perhaps without inducing an ion rotational mode (certainly of the opposite rotation as barrier fields). Barrier fields can always be added at a later time with very little effort.

Given the suspicion that many of the major instabilities that appear after the creation of the FRC are created by ionization non-uniformities, it was deemed far more desirable to generate plasma by a highly localized arc. The expectation being that the plasma would then spread along field lines and have sufficient time to diffuse across the smaller perturbation length scales established by many small sources (without generating imbedded magnetic fields). The ultimate goal was to achieve as uniform a distribution of initial plasma as possible with a focus on providing plasma

at the wall during the start of the reversal bank firing. The desired radial distribution of plasma was uncertain. Unfortunately, pre-ionization via a gun approach would have to wait until the 80 cm experiment.

An additional incentive for the development of 'gun'-based ionization could be understood by considering the remaining available forms of ionization. The use of high frequency RF antennae is one common approach. A second is the use of a ringing-θ addition to the main axial field coils. In general these might be considered inductive modes of ionization (with a little bit of electrostatic coupling thrown in close to the antennae). The general risk with these approaches is that the breakdown voltage of most (if not all) insulators in the system is also much reduced by the high frequencies.

Ionization is not always limited to within the vacuum vessel boundary. This is especially true for long duration ionization sources that also generate UV that can pass through the quartz boundary. This can be catastrophic for the vacuum vessel if the reversal bank coils arc or drive the plasma on the outside of the tube inwards. Ringing-θ can be expected to play further havoc with the unreliable ignitrons as the frequencies are not sufficiently high as to be easily filtered out (by the snubbers). A high frequency, high voltage, and continuous RF antennae that is capable of breaking down the initial gas fill is also an excellent interference source. Such an RF approach can still be frustrated in generating ionization when not supplemented by an initial source of ionization. Using the RMF for ionization also requires the initial presence of plasma. The lower frequency of the RMF, and transverse nature of the field (limited coupling to other systems), helps to limit the unwanted side-effects present in the other options.

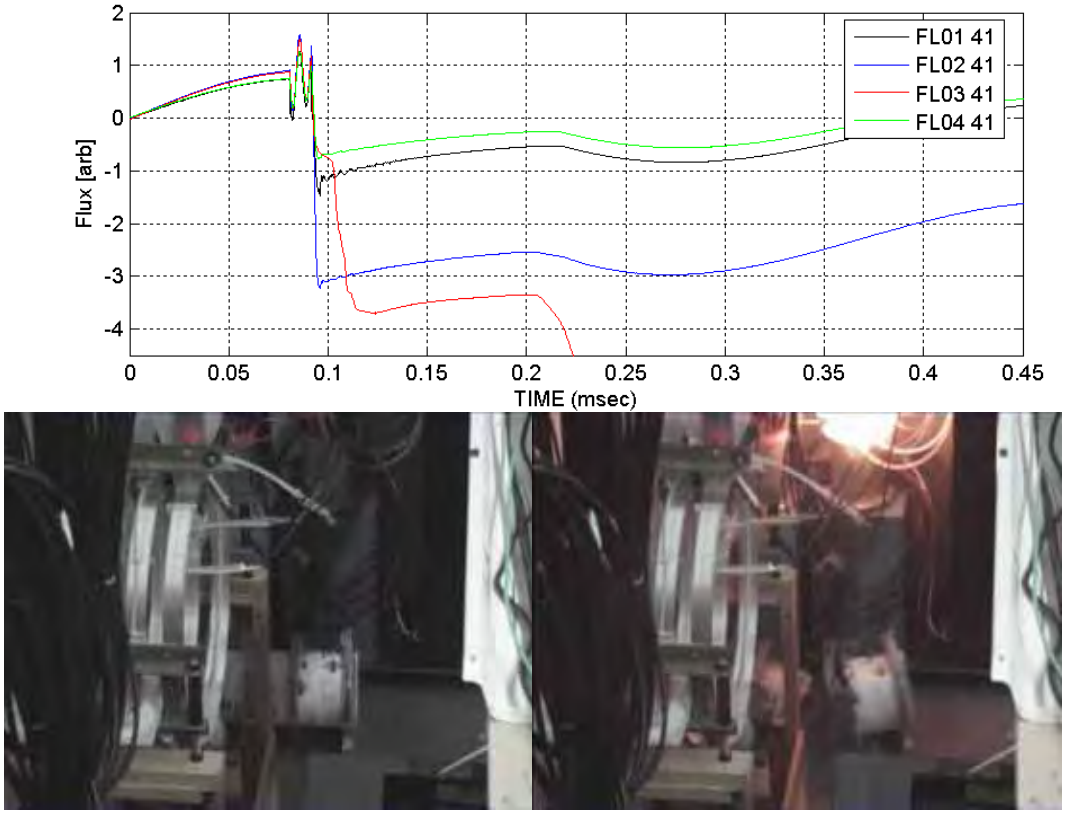

Figure 4.14.1. The price of ringing θ. Before this shot only the bias and reversal fields (with crowbar) were being fired. (top) At time zero the bias fields are fired followed by the ringing theta at 80 μs and reversal fields at 92 μs. Crowbars are attempted at 96 μs but something is already going horribly wrong. There may be corona and localized arcing occurring on the diagnostic resistors, interfering with the signals. (bottom) The driven coils can be seen (aluminum, coils 4, 3, 2, 1) at left. The power sources for coils 2 & 4 can be seen in the background along with four bias isolation inductors (for coils 2 & 4) mounted on a cardboard cylinder. The cardboard tube is of the type used for concrete forms. There were signs of an electrical breakdown beneath the coil. The interpretation was that one of the bias isolation inductors had broken down the cardboard tube it was wrapped around. Cable ties hold each inductor in place. (bottom right) They can be seen flying away from the site as the two ends separate. Thankfully, it wasn't a quartz tube under vacuum.

Additionally, an arc has occurred in one of the ringing theta power assemblies. Because the switch is immersed in Diala® AX oil a short-lived fireball erupts, consuming the oil that was vaporized by the arc. Earlier shots were done without ringing theta. The overall carnage speaks for itself. We would go on to run hundreds of shots with just ringing theta and reversal fields (coils 1 & 3) after disconnecting the bias fields. The diagnostics would continue to sporadically misbehave until potting them a short time later.

Needing a gas fill for a continuous RF antennae is undesirable from the standpoint of the overall experimental program. The potential for gas to escape ionization and move further downstream could present complications if the experimental objective should change from compression and a fast burn to sustainment. See this reference [54] for a discussion of such effects in sustainment.

Experience with 'gun' discharges was gained prior to the funding of PHD with the same basic purpose of providing highly ionized plasma from a system effectively uncoupled from the operations of the experiment. The original focus was on magnetized cascaded arc sources (MCAS) (sometimes called 'washer plasma

guns'[55]). Our interests are not in true washer guns. Washer guns make use of metallic washers as gas reservoirs. They have to be recharged on a regular basis. The guns that we have tested are very similar to a design used for injecting current into the Madison symmetric torus RFP[56]. A fast valve supplies the gas to be ionized. The cathode is located at the gas input end of the device and the anode at larger radius near the exit. The path between the two is partially baffled by metal washers with small ID and non-conducting washers at larger ID. An axial magnetic field was usually applied.

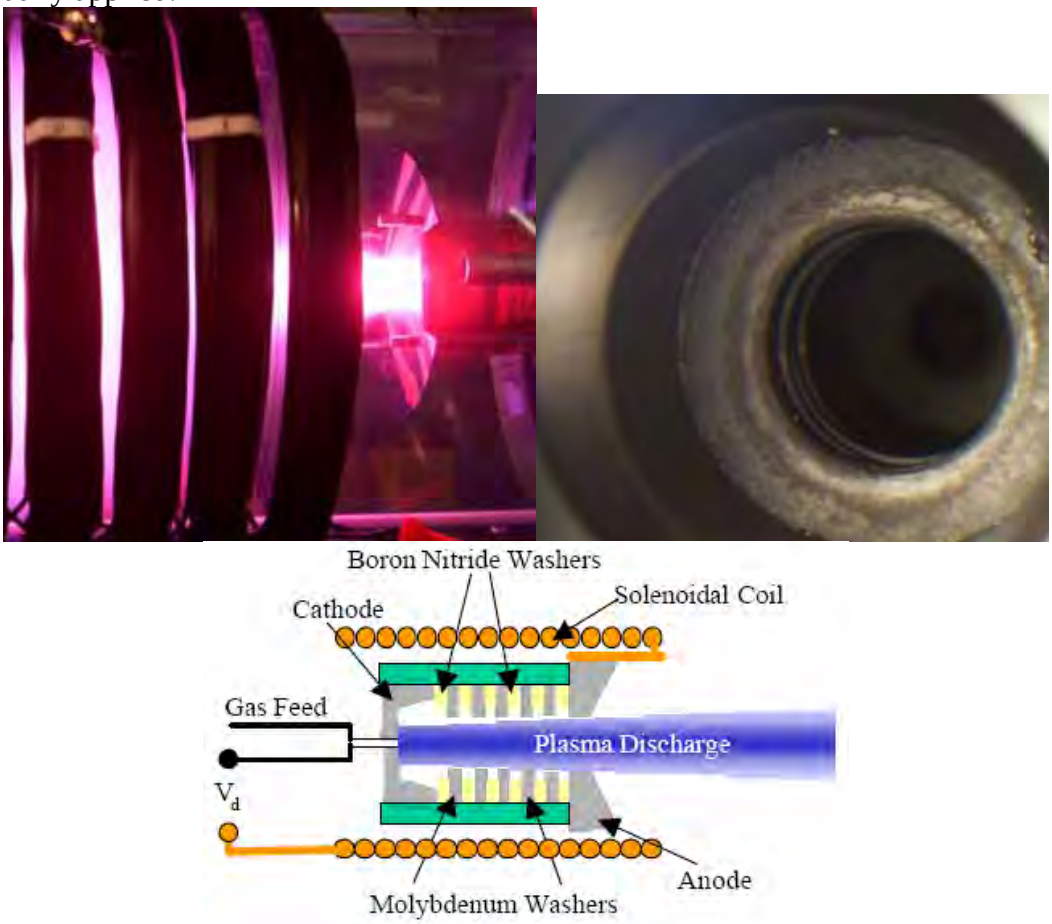

Figure 4.14.2. MCAS on STX-HF. (left) The plasma source incorporates its own high field solenoid winding. For scale, the white tube at bottom is ¼" diameter. (right) Serious erosion of the molybdenum washers was normal. A four hole gas feed can barely be seen at back. (bottom) This diagram shows a Self Applied Magnetized Cascaded Arc Source (SAMCAS). It was more typical to run the axial magnetic field solenoid from a separate bank.

Our previous experience with MCAS on STX-HF showed a ratio of ionization energy to electrical energy of ~4%. To understand why it might be so poor, consider that in addition to dissociation (4.47 eV / molecule) and ionization (13.6 eV / atom) there are also 'losses' due to heating the resultant components (electrons as well as nuclei, ~ 10 eV each) and imparting kinetic motion (several eV). It is easy to arrive at 40+ eV / ion. Losses into radiation, the boundary sheaths, and some local circuit

losses presumably make up the rest (@ ~300 eV / ion, NOT an efficient source). A summary of STX-HF can still be found here[57]. This gun was operated at 300 kW peak power, produced a plasma flow velocity of 15 km/s, and a production rate of 4-6 x 10^{21} ions/s. Given our requirement for 2.5 x 10^{20} ions, and the time taken for the first bulk plasma to reach the end of the tube (~ 250 μs), we would need a production rate approaching 10^{24} ions/s (~ 3 gm /s). This is equivalent to running an MCAS at 60 MW.

Put in slightly different terms, we need a 50 Newton plasma thruster (running in deuterium with 1500 seconds specific impulse). The required thrust for our application scales with velocity squared. A lower limit on acceptable velocity is imposed by the bias field timescale ($T_{¼}$ ~ 150 μs). Unlike a space propulsion thruster, significant recombination is unacceptable. Also, the desired design would be considered an extremely inefficient thruster due to the low output velocity, high ionization energy, and low mass of deuterium. An attempt was made to make an array version of the MCAS; it was quickly abandoned due to quality control problems and the sheer number of units required.

4.14.1 MPD Background

Magnetoplasmadynamic arcjets (MPD) come close to the required operating regime for a single source. It is important to realize that this author's tasks did not include the development of the plasma source we use. Nevertheless, it is important for operational purposes to understand the gross nature and scaling of MPDs. The single most critical parameter for MPDs is $J^2/\dot{m}$ (total current squared / mass rate). This critical parameter delineates a soft transition from the region of operation on injected gasses to the region of feeding on the solid electrodes. From the early pulsed-MPD experiments it was clear that a critical mass flow existed dependent on current and type of propellant[58]. The critical parameter tends to scale with molecular mass$^{-½}$. There is commentary warning that diatomic gasses do not necessarily follow this scaling law. Extrapolating, we expect the transition to be very crudely around 180 $(kA)^2$ / gm/s (deuterium). For thruster applications, maximizing exit velocity and minimizing plasma temperature are desirable. For our application, un-necessarily increasing the exit velocity increases the required thrust, input power, and mass flow rate making a practical system more difficult to create. More recent MPD experimental results include measured results for gasses including deuterium[59]. It is much more difficult to find appropriate results incorporating low I_{sp}, perhaps due to the low implied ionization fraction.

Some sources indicate that asymmetries are associated with initial breakdown. The duration and severity being functions of $J^2/\dot{m}$. The duration is shorter (but not zero) at higher $J^2/\dot{m}$ [60]. The tendency for erosion to increase with higher $J^2/\dot{m}$ is the primary reason for keeping $J^2/\dot{m}$ relatively low. This does not mean that there is no erosion (or its associated voltage hash) below the critical value.

The exit velocity can be expected to scale with the natural logarithm of the ratio of anode radius to cathode radius plus a constant. In all of our applications there

is an additional applied axial magnetic field. This complicates the physics further. For a comparison of models to results for the different regimes see this reference[61].

Complications arise from the expectation that ionization fraction will suffer below some critical exit velocity[61] (the velocity where the ion kinetic energy is equal to the ionization energy, 4000 s for our cases), and hence current level. To be above this threshold implies an increase in necessary thrust by a factor of ~ 10 over our previous velocity assumption, leading to a revised requirement for a 500 Newton thruster and a higher required mass flow rate (8 gm/s). The implied minimum input power is 20 MW before accounting for ionization losses, etc. This implies a current over 100 kA for reasonable discharge voltages; this is not far from the expected critical $J^2/\dot{m}$. Operation below this power level, as we will do, results in incomplete ionization, or insufficient injected inventory. The presence of an applied axial magnetic field may increase the ionization fraction at lower exit velocities (while making the physics more complicated). If the axial magnetic field is linked to the primary experimental axial fields this represents an undesirable coupling between initial bias fields and plasma gun source operations. This is a complication it would seem we can no longer avoid given the experimental results.

4.14.2 MPD on Experiment

The initial MPD designs would concentrate on large on-axis sources due to applicability for other experiments and simple implementation. Given the expected presence of asymmetries in initial breakdown phenomena, the cathode is heavily segmented with individual feeds and current limiting resistors. Initially, the plasma was required to flow past a cusp and across the bias magnetic field radially to reach the walls and trap the lift-off flux. The coil design did not allow for a connected field situation of sufficient duration, followed by cusp, and then reversal. It was rationalized that the cold plasma should be able to cross magnetic fields. Only once it was experimentally determined that the MPD plasmas were apparently not effective at crossing the transverse magnetic fields present in PHD would the focus shift to a large radii, segmented approach. The smaller radius, more numerous, and more reliable MPDs required for such an approach still have not been optimized as of this writing (and were not installed on PHD before most of this document was compiled).

The most successful design is summarized in the figures below. The primary features are a segmented (48 element) thoriated tungsten cathode and a molybdenum washer anode. The ratio of channel radii approaches unity unlike most MPD thrusters. Non-conducting (and conducting but unconnected) boundaries are added to restrict gas flow to the channel and to allow sufficient pressure for the available voltage (<5 kV) to break down the gas. Neutral gas injection is accomplished solely at the back of the device via four poppet valves. Initially, this unit was powered by a three element 1.5 mF PFN bank providing peak current of approximately 50 kA and duration of 100 μs. Voltage hash was observed at frequencies >1 MHz and amplitudes greater than 200 V at peak current.

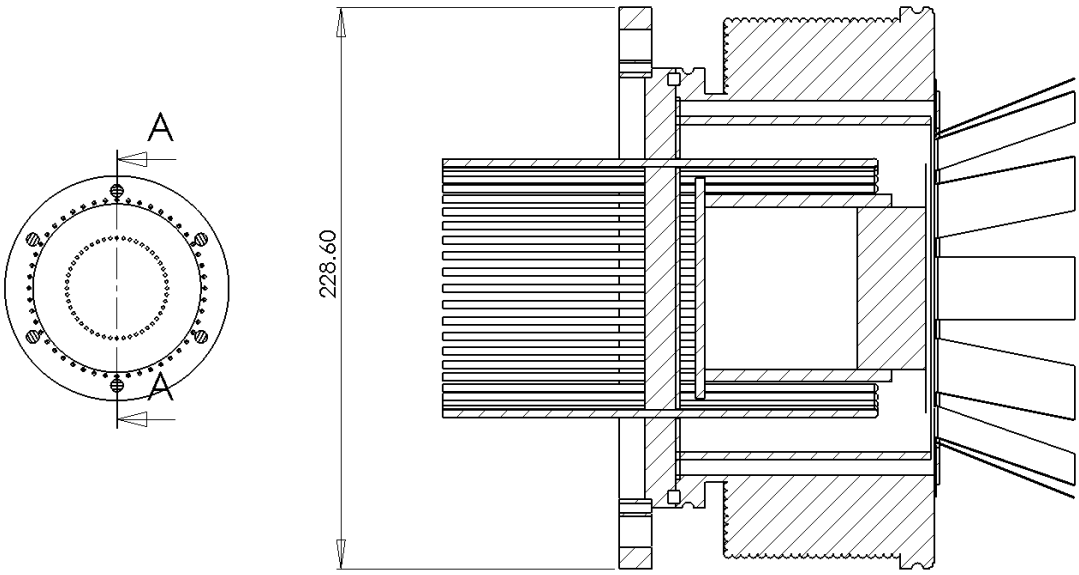

Figure 4.14.2.1. MPD source schematic. The cathode rods (1/8" diameter) are made from 2% thoriated tungsten. The OD of the final orifice is 4-7/8 inches. The center metallic section (tungsten sheet on aluminum plug) was added to help maintain the neutral pressure as an aid to initial breakdown. Inwards from the cathode rods is a glass cylinder. Outwards from the cathode rods (between cathodes and anode) is a quartz cylinder. The anode consists primarily of a molybdenum ring with copper sheets added after erosion was detected on the screw heads. Gas is injected close to axis through left flange (using four puff valves). (left) The flange includes 48 plug contacts for the anode.

Figure 4.14.2.2. MPD hardware in place on +z end of machine. The original magnetic field coils are also clearly visible. Only one cathode rod's voltage is measured through a resistor and current pick up loop. The MPD current measurement is inferred from the sum of several individual cathodes.

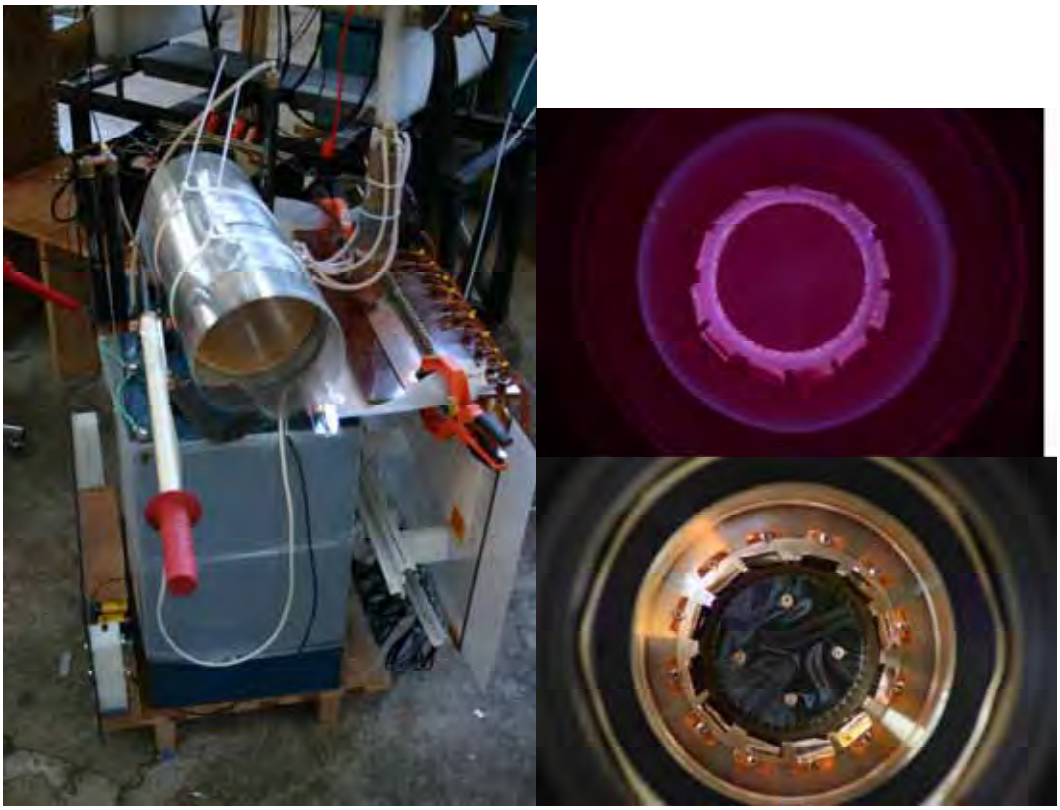

Figure 4.14.2.3. MPD PFN energy storage bank. The cylindrical tube on top is used to alter the inductance and coupling from stage to stage of the PFN. RG-58 coax is used to transmit the power to the thruster (one coax per resistor and tungsten rod). One ohm, stainless steel wires are used as individual resistors to drive each tungsten cathode separately. The use of the resistors makes the PFN much less susceptible to impedance mismatch at the price of electrical efficiency. The capacitors are not rated for high levels of reversal, or high currents. The MPD was originally used with a pulse forming network (the inductance was quickly bypassed once the stainless steel ripped away from the terminals). An ignitron crowbar would be used in place of the PFN to prevent ringing of the capacitors. A type A ignitron was initially used for the crowbar followed by a type D ignitron upon the type A's failure. (top right) The location of discharges can be ascertained from the MPD light emissions (integrated over duration of single shot). (bottom right). End on picture of MPD in relatively pristine condition.

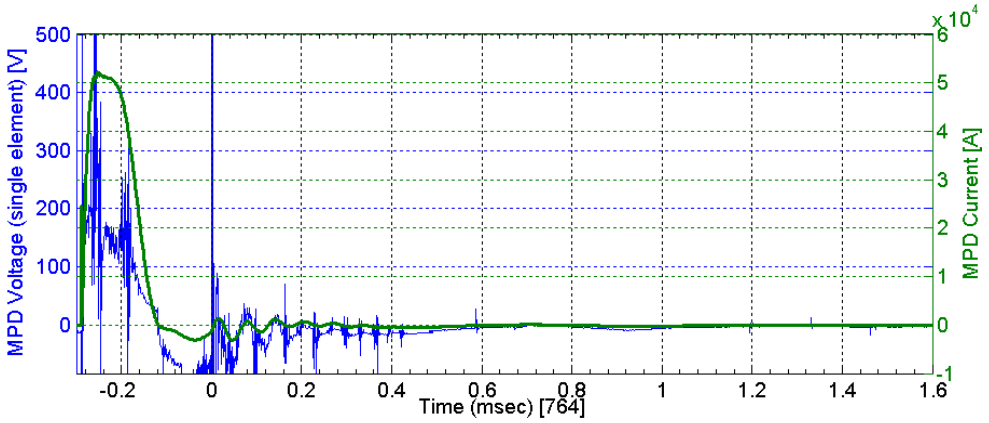

Figure 4.14.2.4. MPD driven by 1.5 mF PFN bank. The PFN was capable of driving at a power level approaching 10 MW. Reversal banks are fired around 0 μs resulting in some interesting current and voltage oscillations.

The design did accrue significant changes throughout its lifespan on PHD. The PFN bank version installed on PHD lasted less than 90 shots. The PFN bank was then changed to a standard crow-barred capacitor bank. The removal of the inductive stages was forced by their gradual tearing under high currents. However, it was clear that the MPD was not getting the plasma to where it needed to be in the formation chamber. Replacing the PFN network with a simple inductor and crowbar allowed for a longer duration of plasma output. If significant diffusion time was indeed required this provided a continual plasma source maintaining the pressure. If neutral gas escaping the MPD before discharge was a problem it was hoped that we could 'burn' through it. The inefficient resistive ballast would continue to be used.

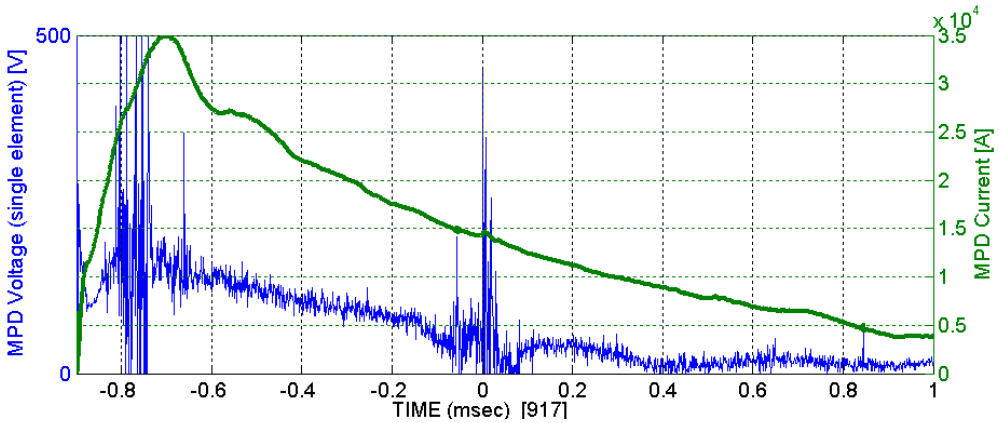

Figure 4.14.2.5. MPD current and voltage traces for shot 917. Note that we are only interested in having plasma production for less than 200 μs, removing the usefulness of the original PFN network (that at this time had already been collapsed to a single capacitor. The hash near time zero is probably due to the reversal bank firings. There were never sufficient diagnostics on this experiment to properly diagnose source operation. We are only able to see gross electrical parameters and some of the effects on FRC formation. MPD operations were to have been characterized earlier, on a separate experimental apparatus.

Erosion on the molybdenum anode and stainless screws had been detected during the development of the MPD. For this reason, copper anode 'feathers' were added before installation on PHD. This initial design showed the presence of significant sodium lines in the source. The generated plasma was interacting with the glass cylinder located inside the radius of the cathodes. Before the PFN bank was removed, this glass cylinder and the end cap (aluminum structure, tungsten foil surface) became dislodged due to gas pressure. Around this time tungsten foil was added to the surface of the glass.

Initial operations of the MPD took place with the MPD on the west end of PHD. The MPD was placed on this end because it was possible to generate the magnetic fields necessary to link the MPD and the formation chamber from the west (albeit only with long timescale magnetic fields). This was only a temporary solution as this location would ultimately be the direction the FRCs would have to translate

towards during the next phase of PHD. The east side of PHD is dominated by the presence of a 6-way cross and magnetically levitated turbo molecular pump. Externally applied magnetic fields that would be necessary to span the cross would inevitably penetrate the maglev turbo pump's shell. For this reason, when the MPD is on the east side of PHD the plasma would have to traverse the 6 way cross unguided by magnetic field.

It did not take very many shots for some of the copper 'feathers' to start bending inwards (obstructing the optical line of sight from plasma source to destination). The tungsten foil that was added to shield the glass quickly detached from the glass and contacted the segmented cathodes, no doubt resulting in an asymmetry. These obvious flaws developed while the MPD was on the east side of PHD.

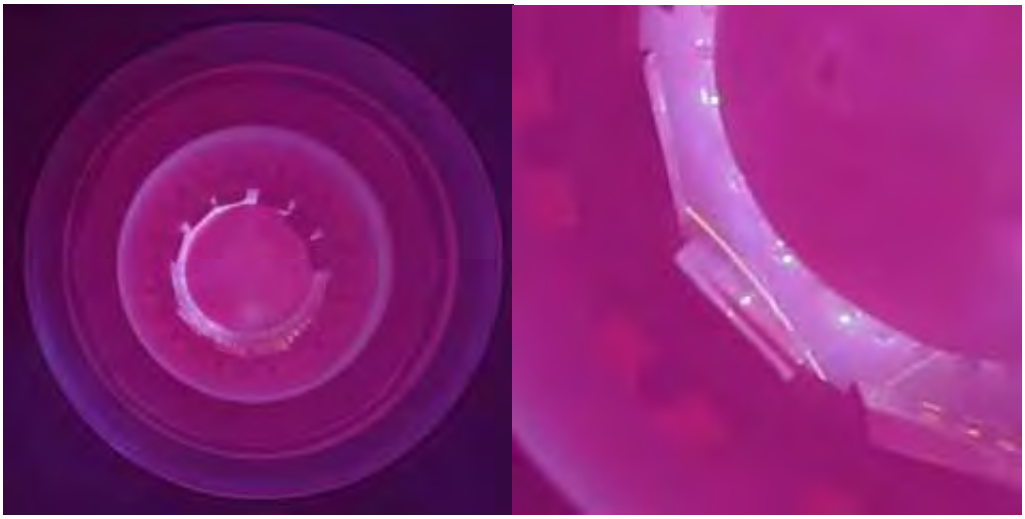

Figure 4.14.2.6. MPD on east end of experiment. (left) Shots before replacement of the MPD show the bent copper 'feathers'. (right) Molten metal can clearly be seen bouncing around the interior channel. Molten metal is visible in multiple locations both directly viewed and via reflections. This is probably the result of a loose piece of tungsten foil (that covers the inner glass cylinder) coming into contact with multiple cathodes resulting in significant current flows through poorly connected surfaces.

The MPD would be modified to remove the unreliable elements in the earlier designs. The removal of the copper feathers brings the anode design back to an earlier state. The replacement of the tungsten foiled aluminum end cap with a copper end is a change of convenience. The MPD would be operated in this mode unobstructed for only 20 shots. After this, a SXR detector would be added to the centerline of the experiment for testing in the 6-way cross.

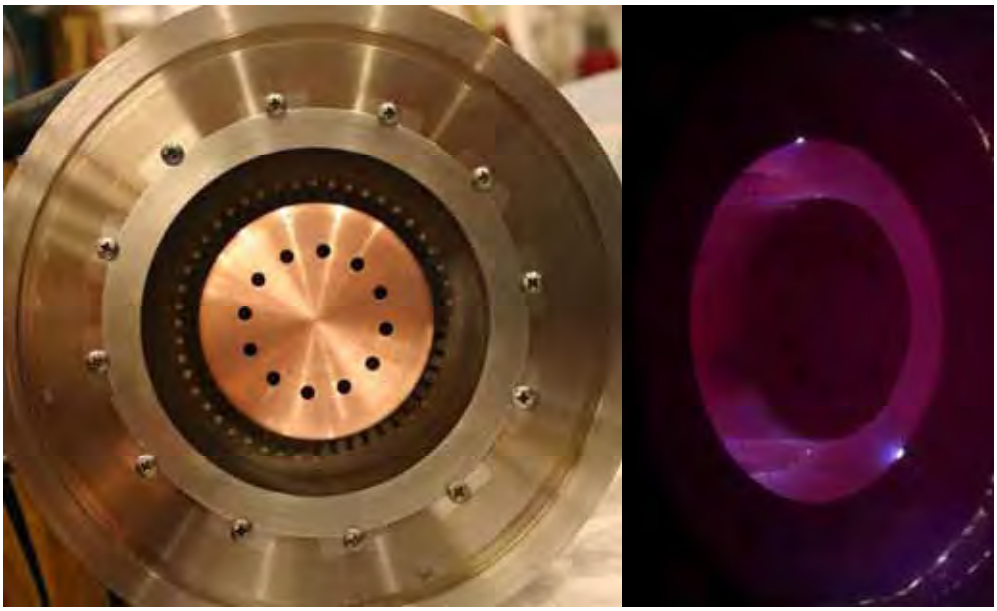

Figure 4.14.2.7. Hollow copper center MPD. (left) The primary modification to the earlier MPD designs is the removal of the copper 'feathers' and the replacement of the center tungsten lined aluminum end cap with a copper end cap. (right) Operation shows the propensity for highly localized emissions on the molybdenum ring anode.

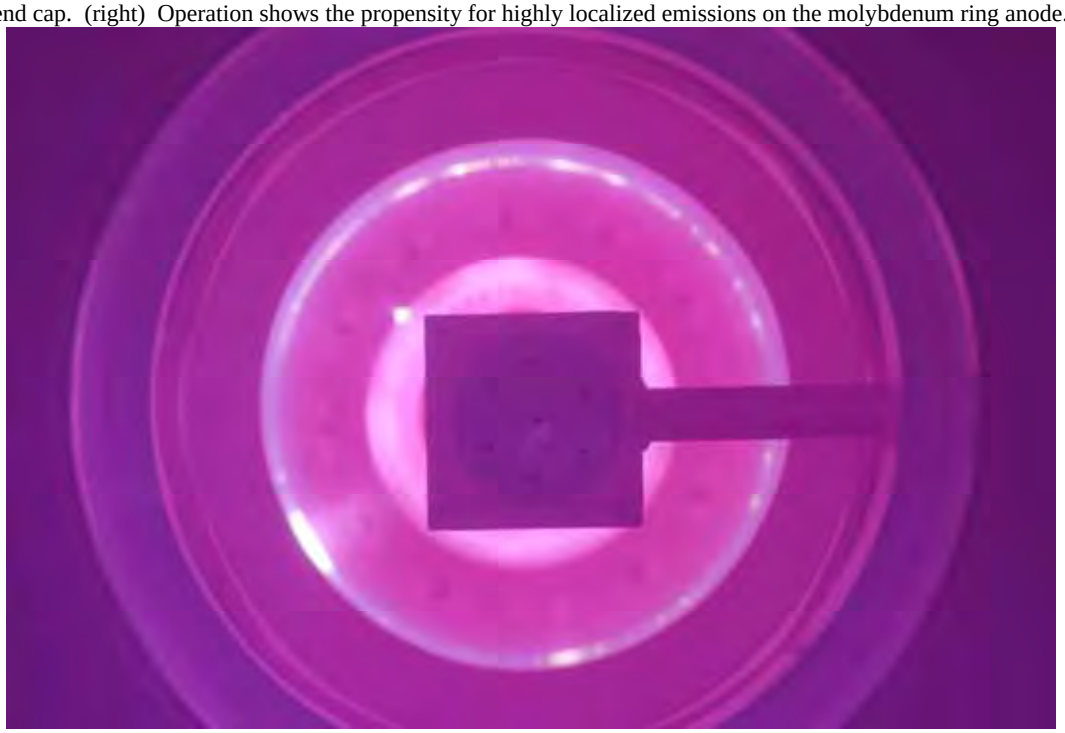

Figure 4.14.2.8. Copper capped MPD blocked by SXR diagnostic on east end of experiment. The SXR diagnostic design was looking forward to a time when the MPD sources would no longer be on axis. Unfortunately, this early test of the diagnostic had no definitive results. The ratios of the different x-ray spectra did not give self consistent results for derivation of temperature. One might speculate that this could be due to unanticipated plasma contaminants.

4.15 Fast Gas Valve (Puff)

Solenoid actuated gas valves are commonly used in various industrial control applications. They can be made capable of sufficiently high speeds for our applications. However, this is not always the intent of the manufacturer. Specifications rarely reflect their speed capability (more typically specified for millisecond actuation timescales) under increased drive levels. The Parker-Hannifin series 9 valves (part # 091-0199-900) we originally bought were capable of sub ms smooth high speed actuation while being severely overdriven. However, there are indications that their redesigned replacement valve oscillates heavily on opening, resulting in multiple gas pulses rather than a continuous stream. Examination of the valve construction shows that the armature is no longer split down the side and no doubt other changes have been made to control cost. Our understanding is that a redesign had occurred to reduce the number of unique parts required for the different valves sold.

We routinely actuate the valves with voltages up to 600 volts across the coil. With respect to the casing, voltages can approach 1000 volts. This clearly pushes the advisable limits for the insulation employed. A danger is present from the potential for punch through between the coil and the valve body. The danger to the user can be lessened by grounding the valve body to the circuit ground. There are then three electrical connections between valve and driving electronics. The circuitry used was based on the ignitron igniter design (and as such will not be discussed here). For slow actuation of valves (and valves that can't act as electrodes) special design considerations are not necessary. While our drivers are capable of battery operation the risk of capacitive breakdown of the insulation between valve body and coil is still present. We have not yet had a failure occur in this way (despite the valve specifications).

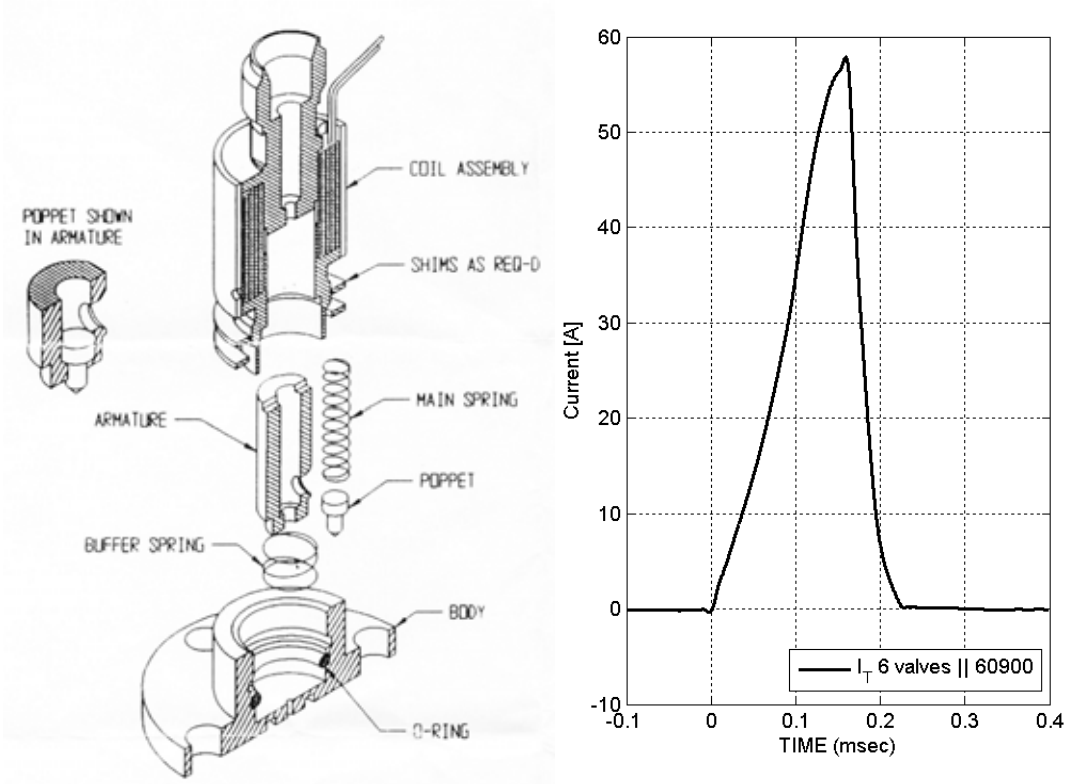

Figure 4.15.1. Parker Hannifin Corporation (General Valve Division) series 9 poppet valve. (left) This diagram was obtained from Parker Hannifin literature describing the proper shimming procedure. It is used with permission. In a standard solenoid poppet valve a coil's magnetic field acts on a paramagnetic armature to pull the poppet off the sealing surface. The timescales for penetrating the conductive surfaces can be very important. Proper use of the shims can improve the speed of the valve; when the shims are too thick the valve is held open. If the armature hits the back stops and isn't damped the potential for some oscillation clearly exists. (right) The current trace for 6 new replacement valves (in parallel) shows an increasing rate change of current with time. The implied decrease in inductance may indicate magnetic saturation of some of the components when driven with 600V. The frequency response of the hall current probe is approximately 100 kHz. Note that the duration of the open time is not limited by the drive electronics; this figure is to demonstrate the fast turn on and fast turn off of coil drive current that is possible with the electronics. The actual closing of the valve is likely to be somewhat delayed due to the need for a spring to return the poppet to the original position.

Given that these valves are magnetically actuated, it is not surprising that they require additional magnetic shielding in order to properly operate near our large axial magnetic fields. This has typically been accomplished with a steel pipe over the valve coil assembly body. We have unintentionally actuated these valves open with the magnetic fields applied to the plasma guns. An additional concern is the potential for magnetization of the various ferrous components. We have witnessed some valves (when disassembled) stick open upon actuation (but never while in actual use).

The 40 cm experiment used a back fill. The timings were not as critical and the puff valve driver used 18 volts. The valves were designed to run at voltages in this range (12 V valve). The timescale of opening is in the range that one might expect from the specifications of the poppet valve.

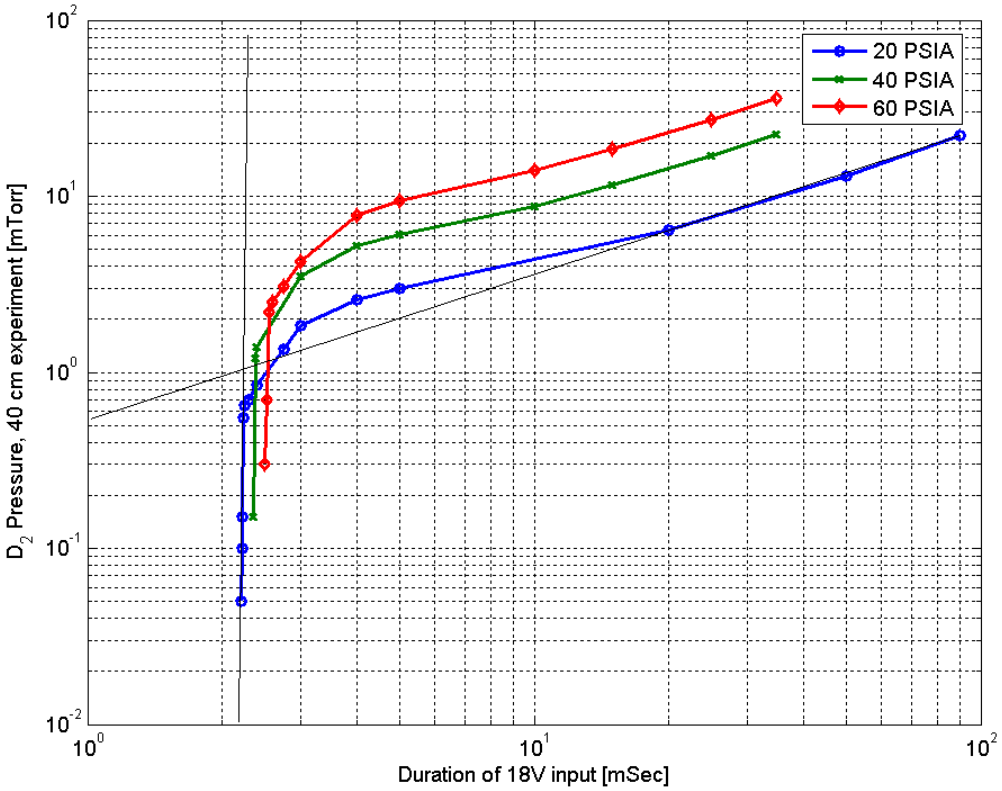

Figure 4.15.2. Characteristic behavior of softly driven 12 V poppet valve. The time to open, and hence the force applied to open, can be seen to increase with pressure. The effects of a small local gas volume on the transition from pulsed to steady state open is fairly clear from the slopes (and hence gas throughput). Short timings had a great deal of variability in end state; this was presumably due to the armature and poppet not coming back to the same original positions. For this figure the coils are driven in a crow-barred mode where the drive coil current is not rapidly zeroed.

The 80 cm experiment typically used peak voltages of 400 V or 600 V while opening and -400 V while de-energizing the solenoids. Required actuation duration was exceptionally short (<300 μs). While the primary puff valve supplying the MPD only ran with deuterium, a secondary puff valve next to the vacuum pumps allowed for supplying other gasses (helium [735-763], more typically deuterium) as a backfill.

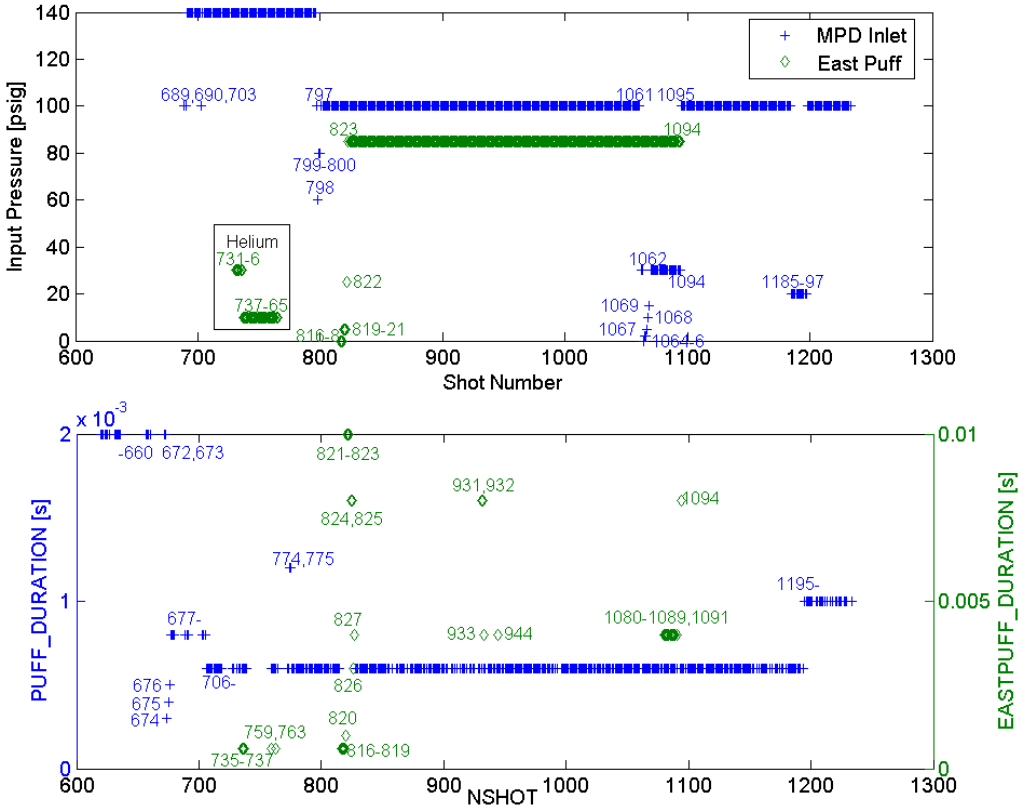

Figure 4.15.3. Inlet Gas Conditions for 80 cm Experiments. The MPD parameters are shown in blue, whilst the back fill system parameters are shown in green.

4.16 Control Systems

A large number of systems are required to control and instrument a complex experiment. Many of the systems can use hardware from earlier experiments; not all need to be developed from scratch. Early in the design process we decided to separate out systems that could be controlled manually and those for which hardware already existed. Not all the new systems would have to be fully functional in order to run the experiment and collect physics results. In general, this has meant continuing to use CAMAC (Computer Automated Measurement and Control) based systems providers for our 'pulsed' data-acquisition and control needs as well as using manual controls for slow-timescale systems. Initial development of Labview interfaces and HDF/Matlab data analysis had been completed in previous years before the funding of PHD. Many of the more modern hardware replacements that have been purchased have not had appropriate specifications to replace what we already have (for example: cost-effective, 100 ns, multi-channel data acquisition). Buying more of what we already have (such as CAMAC hardware) costs only as much as is on the price-tag; buying newly designed equipment usually costs more money because of the required integration and test time.

4.16.1 Bank Voltage

Various systems can be satisfactorily implemented with external feed back and control loops. For example, the bank-charging power supplies are designed to charge the banks to full voltage in a little under two minutes. This implies that under normal operation a significant delay is acceptable between reaching the desired voltage and instructing the supply to stop charging. By making the voltage control, feedback, and monitoring part of external circuits (to the supplies) we allow the control aspects to become common hardware for other applications (such as controlling gas pressure). For most supplies it is a matter of converting the fiber signal to an appropriate charge enable or current limit signal. The local supply set point voltage becomes the absolute maximum charge voltage. For systems that require fast timescale response it is simple to use this external control loop in addition to a local control loop to gain the desired system characteristics. So long as the timescales chosen are appropriate, it is easy for a human to take the place of the controller for testing.

One return fiber is necessary to report the voltage (as described in the V -> f section). Many multi-meters can read frequency. A second fiber controls the dump relay and instructs the supply when to charge. This control fiber has three states: steady off, 1 MHz 50% duty cycle oscillation, and steady on. Steady off corresponds to not charging and dump relays closed. 1 MHz 50% duty cycle corresponds with dump relays opened but not charging. On steady causes the dump relays to open and the supplies start charging. This approach saves some fiber and limits charging while dumped. For a digital receiving system the response time would be excellent (<1 μs); an analog approach results in longer delays but is more easily implemented (for example by duty cycle conversion to voltage).

For the demanding condition of a constant power charger $\frac{dV}{dt} = \frac{P}{cV}$ (with a low voltage on the output our supplies are not set to run with constant average power). The percentage error in charge voltage for a given delay is approximately specified by $\%Error = \frac{P}{cV^2}\Delta t \times 100\%$. With supplies running at 300 watts, a capacitance of 72 μF with a target voltage of 10 kV, and a response time of 50 ms (commensurate with an alert human) we can expect approximately 0.2% error. In practice, common frequency meters often limit the update rate to once every second. Obviously, the human can do slightly better by stopping before reaching the desired voltage and 'tapping' his way up (turning on for only an instant at a time).

It is more desirable to have bank control fully automated, but this desire must be balanced against the cost for developing a truly reliable and error safe system. "A human being is the best computer available to place in a spacecraft … It is also the only one that can be mass produced with unskilled labor." – Wernher Von Braun. Of course, humans aren't exactly error free; but neither are prototype systems. Readily available hardware allows remote control with a human in the loop approach.

A secondary consideration to voltage uncertainty is data rate. It is desirable to limit the control system bandwidth in such a way that unanticipated errors won't develop. Our way of monitoring the voltage on the various banks is through a voltage to frequency converter. The output can either be tabulated for short periods of time, i.e. (count the number of pulses for say 10 ms) times 100 = frequency, or the interval between individual pulses can be monitored. The two approaches can be classified as synchronous and asynchronous (data flow to the processor).

The synchronous approach requires longer tabulation times for high accuracy at low frequencies. If one is monitoring a 10 Hz signal, the tabulation time needs to be for at least 1 second in order to get better than 10% accuracy; 10 seconds in order to get better than 1%. On the plus side, the data rate is constant and well defined. On the other hand, for the approach to gain 1% accuracy at 1 kHz requires a tabulation time of 100 ms.

The asynchronous approach has the advantage of providing a response time limited by the frequency being monitored. That 10 Hz signal can be measured to arbitrarily high accuracy in 100 ms. The 100 Hz signal in 10 ms. The price is that the data flow isn't predictable. A second price is that the clock that counts the time interval needs to be much faster than the highest frequency to maintain accuracy. Our monitors' top of range frequency is 250 kHz; if 1% accuracy is desired a counting clock speed of 25 MHz is required. In order to accurately measure a 1 kHz signal the counter must consist of at least 15 bits. Rollover has to inform the processor or risk a false high signal. Noise pulses have an immediate effect. So the system that monitors the data has to be capable of operating at a variable data rate higher than 250 KB/s x #bytes x number of channels. The rate of data is only limited by the response frequency of the fiber optic receiver.

Here's the necessary realization to make for an easy choice: The synchronous approach can expand its tabulation time via programming when higher accuracy is desired (on individual channels) and engage in a version of the human's 'tapping' as it nears the desired level (to compensate for it's delay time). Data output for recording can be set at some reasonable rate with no fear of overwhelming the processor. For example, the base tabulation rate can be set to 1 ms with 2 bytes implying 2 KB/s times the number of channels (and a maximum input frequency of 65 MHz without rollover error).

In truth, both systems have a hidden advantage that makes them perform better than one might expect (even in a grossly simplified implementation). The desired set points are quantized, if one chooses one of these set points the accuracy becomes limited by the response time and charge rate of the supply. For example (in the synchronous system), if a frequency of 10 kHz is desired with a count time of 1 ms, the desired count is 10. If 9, continue to charge. If 10, do nothing. 9.5 kHz results in half the returns being 9 and half 10. The frequency progressively more slowly increases to 10 kHz and then sits at 10 kHz with a degree of precision limited by the charge rate. Similarly, the system can very precisely maintain 1 kHz. A side effect is that the charge rate is slowed as it approaches the target frequency. This

works the same for an asynchronous system too, so long as the clock isn't reset with every pulse.

It should be obvious why a dedicated microcontroller/microprocessor is desirable (esp. given the earlier discussions). At heart, such a system is easily and rapidly implemented. There are designs for single channel frequency meters (microcontroller based) available on the internet that could easily be converted for this purpose. In our case it is desirable to have multiple channels and an Ethernet interface to control costs (and provide recording capabilities). This is a slightly more difficult problem (less difficult now that PIC microcontrollers are integrating Ethernet functionality onboard).

Some additional features must be implemented to deal with human error modes. For example, each of the controllers should have a fiber going to the interlock hub. If the output voltage gets too high, for example from someone having switched the control fibers for two banks, the interlock fiber allows a way to remove power from the whole system. An allowance for manually controlled set points is also preferable so that a computer is not required for every application. The ubiquitousness of cheap computers makes this less a concern than it once was.

Figure 4.16.1.1. Bank voltage remote control systems. (left) The first prototype 8 channel microcontroller with Ethernet control system. The BNCs and IDCs are for manual monitoring and set points during the test/verification phases. This system can be compacted to a CAMAC slot profile once the design is verified. The system is currently set for manual input (as a bank of 8) with the set voltages being controlled by one dial (via an IDC connector at rear). Interlock connectivity is also accomplished via a rear IDC port. Ethernet hardware functionality has been tested but the final routines to implement an interrupt driven system have not completed testing for lack of time. DIP switches allow for changing the counter clock base in powers of 2. The prototypes are used to control the reversal and bias banks. At least one IDC port will be retained to allow for microcontroller programming in system and features not yet considered. The interlock fiber transmitter/receiver may be integrated into the final board. (right) Three channels of pure manual control. One unit converts the fiber frequency to BNC. The meter displays the frequency. A second unit generates the three state bank control signal. It is expected that even upon completion of the automated system these units will be used for testing/troubleshooting purposes. The three state bank control unit is also useful for test firing (because the momentary charge button is de-bounced).

4.16.2 Interlock System

Our power interlock system can be thought of organizationally as a tree. The individual sensors (emergency stops, mechanical switches, magnetic switches, phased-infrared motion detectors, laser trip-lines) can be thought of as the leaves of the tree. Each physical location has a limb that aggregates the results from the

individual sensors and sends it down to a hub. Hubs can feed into other hubs (unlatched) until reaching one trunk hub. The trunk hub allows for displaying, latching, ignoring, and resetting the detected problem (and is the same hardware as all the hubs). This root hub controls the root of the system, the primary power node, and it in turn controls the ancillary power nodes. The connecting spokes are the standard 1mm plastic fiber optic cables. A light signal indicates proper operation. This ensures electrical isolation and lack of interference. Warning lights and sensors are powered off one switched circuit. Mechanical power switches for the three phase power allow one last redundancy for safety.

Standard normally closed (series circuit) security systems are usually simple to retrofit to fiber operation. Generally, magnetic reed switches have become preferred due to their reliability and installation simplicity over mechanical switches. One permanent magnet and one normally open reed switch act as a position sensor. Similarly, phased infrared motion sensors are useful for detecting people and fires with minimal modification. Usually a local power source for the electronics is required along with the fiber transmitter and an appropriate resistor. Laser trip-lines only require a laser, a few mirrors, and something to hold the fiber (sometimes with a collimator).

Motion sensors satisfy multiple purposes. They add a useful time-delay between power up and outputting a signal. When they are powered off the same AC line as the warning lights, this enforces a reasonable time-delay between the activation of visible warnings and the power up of the experiment.

The hubs are where the interlock electronics are located. Each one includes eight interlock receiving ports. They include green LEDs to visually indicate received interlocks. They have red LEDs to indicate interlocks that have been disconnected and latched. Simple DIP switches control which interlock points are ignored (for example due to an empty port). They require a small tool to actuate. An emergency stop switch, key switch, and momentary latch reset switch can be integrated in the output. These features are usually used on the root hub. Any active interlock with no signal causes the output interlock of the hub to latch off and indicate red on the signal that caused it.

The power nodes are actuated by electronics containing Omron Electronics G3NE-220T-US DC5 solid state relays (SSRs). These SSRs integrate varistors. They have proven themselves very reliable despite the harsh electrical conditions we subject our systems to. The fiber receivers that drive them are powered by a 5 V, 4 kV Hi-pot tested, switching supply (CUI Inc. DPS050200U-P5P-DF). No false turn-ons have occurred.

The ancillary power nodes sometimes incorporate a momentary over-ride switch so that power can be applied to a load without a fiber input signal. They can also be used with a single sensor node (for example motion detector and door interlock) to create a rudimentary latching interlock system without a root hub. All power handling systems are fused.

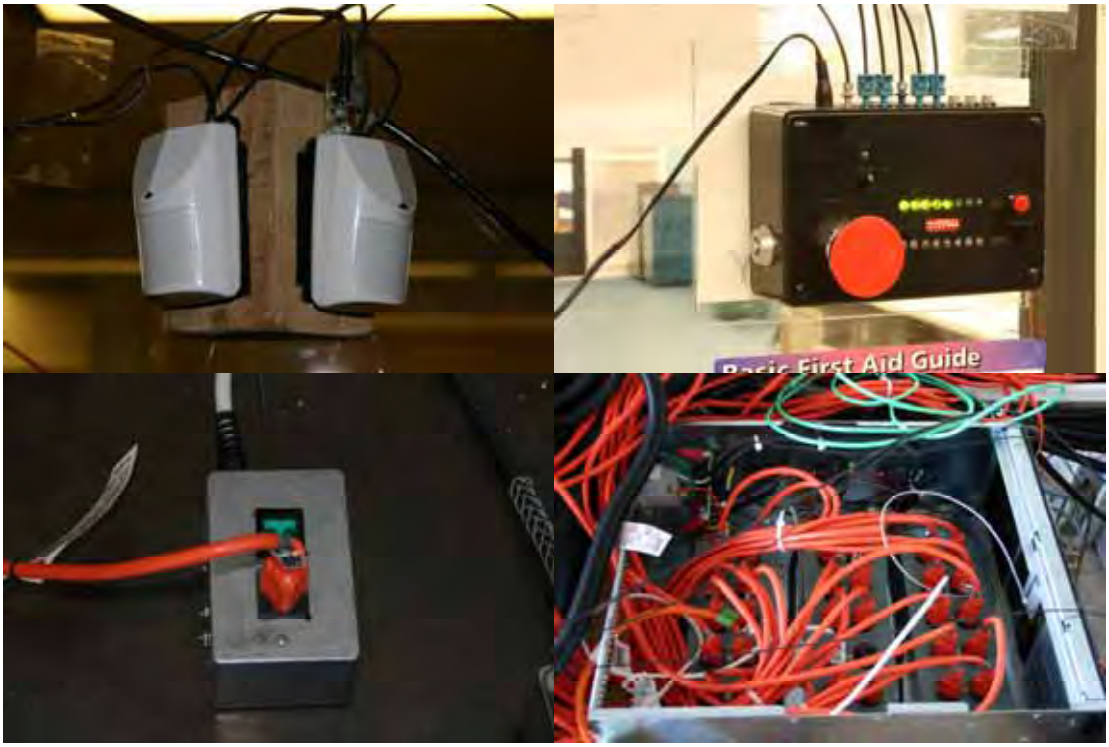

Figure 4.16.2.1. Interlock systems. (top left) phased IR motion sensors. A magnetic door switch is off frame to the left. The motion sensor at right is acting as the sensor node providing power and the output for the system. (top right) Root hub. The small red button acts to unlatch the system. Two LEDs indicate the combined state of the E-stop button / key-lock and the effective output state. (bottom left) Single fiber controlled 120 VAC power point; ancillary power node. This allows flexibility in powering some systems off separate circuits. (bottom right) Fiber controlled 3φ 120 VAC AC primary power node. This unit is scheduled for replacement with a fully enclosed unit with outlets for power strips on back.

4.17 Instrumentation Rack

Pulsed power experiments usually have some form of screen room to shield sensitive electronics and controls from the magnetic fields and EM radiation of the experiment. Proper techniques and practices can be found in various books such as Ralph Morrison's "Grounding and Shielding Techniques (4th edition)".

It is worth noting that wire mesh covered rooms are not proper screen rooms. Nor, strictly speaking, would our screen rack qualify as a proper screen room. The fan installed in the top of our rack is not properly baffled for EM radiation, nor are many of the fiber ports. There is also only a single metal wall rather than a double wall. Power conditioning and shielding is minimal. Nevertheless, our screen rack provides a large amount of electromagnetic noise attenuation and a significant amount of magnetic shielding. Given the large fields, and their large spatial extent, inherent in our experiment it is difficult to obtain excellent shielding without building a custom enclosure. We believe that external noise sources not accounted for earlier account for at most one or two digital counts on our magnetic diagnostics (i.e. mV levels).

It is difficult to isolate internal cross contamination from external sources. With the reversal banks operating at +/- 20 kV and the bias banks at +/- 4 kV a maximum excursion of 60 mV is observed on the VLA array. The VLA receiver boards sit 2 CAMAC slots away from their digitizers. Short (couple feet long) RG-174 cable is used to connect the channels. The contamination immediately responds to the firing of the reversal banks with a signal that slightly leads the reversal flux. Different channels have vastly different levels of contamination (some < 5 mV for example) despite their electrical paths being very close together. The contamination falls off with a characteristic timescale of approximately 30 μs. This last detail leads to the suspicion that the cause of the large noise signal may be related to the termination mode of the digitizer. For some length of time experimentation with the digitizer's differential mode input was attempted. Our experience was that it was incapable of acting differentially on fast time scales.

From the specifications of the rack, we can expect better than 100 dB attenuation to electric fields for frequencies less than 20 MHz. Unfortunately, the attenuation to low impedance magnetic fields is only 20-40 dB for 10 kHz – 1 MHz. These tests were carried out by the manufacturer in accordance with MIL-STD-285. This standard requires symmetric transmitter and receiver with diameters of 12" for the H-field test. The two coils are flat in a plane perpendicular to the panel under test. The edges of the coils are 12" from the shielding panel. The test is run with and without the shielding panel in place to determine attenuation.

From the calculations and specification so far, we would expect to observe on order of 5 V/m^2 for an aligned loop inside the rack. Note that the magnetic field that penetrates the rack must still penetrate the shielding of the internal cables and equipment to cause problems. Most diagnostic cables have grounded shields at the inlet to the rack and the digitizers or CAMAC equipment (the one exception being the digitizer running the differential mode test). A current would be induced in this shield. Some of this current would cause a signal on the inside of the coax in accordance with our discussion of transfer impedance (see Section 5.1.7).

Our primary rack contains the data and control server, CAMAC crate, digitizers, signal generators, oscilloscopes, fiber optic buffers, PMT power supplies, and bank control hardware. An industrial quality power line filter (Corcom 20EJT1) filters the incoming power at the rack bulkhead. This unit can be expected to provide a peak attenuation of 60 dB on the incoming power line noise. At 1 GHz, the attenuation plummets to around 10 dB. Below 500 kHz the unit is similarly reduced in effectiveness. Internal uninterruptable power is provided via a Tripp Lite OMNIVS1500XL. The UPS might be a significant source of noise during power outages. The fan is another potential source for internal electrical noise.

All coax cables are grounded while entering the rack enclosure. All attenuation or integration is also done at the bulkheads. Some effort is expended in ensuring that all signals are of similar strength inside the rack in order to reduce cross contamination and the potential for damaging electronics. Most electrical signals are grounded at two points, once upon entering the rack, and once upon entering the equipment (i.e. a digitizer). This intrinsically results in ground loops that would

allow high current signals to contaminate more sensitive diagnostics. External PMTs (such as the multichannel spectrometer) tend to generate smaller signals that are more sensitive to contamination. It is preferable to amplify and buffer these signals at the source equipment rather than in the screen rack. The visible light array detectors reside in the CAMAC crate adjacent to their digitizer (minimizing the ground loop). The same is true for the interferometer electronics.

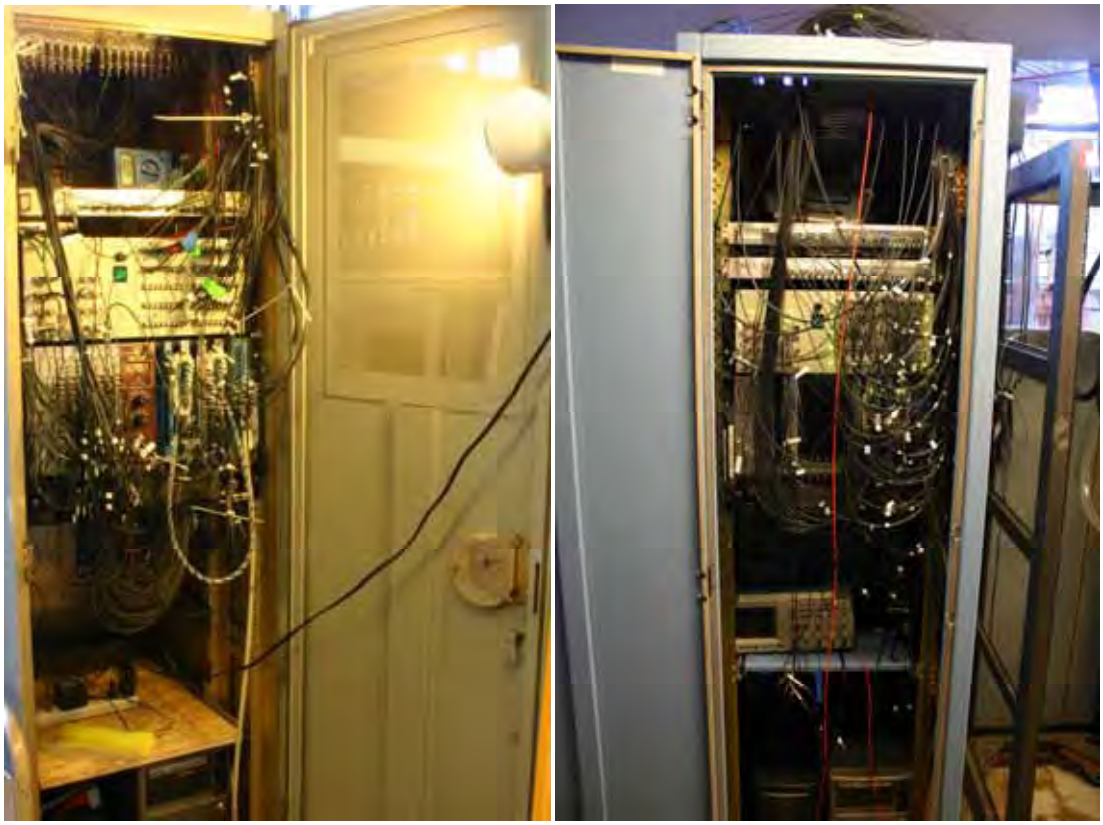

Figure 4.17.1. Screen racks. (left) A surplus rack that was formerly part of a radar installation was used for the 40 cm experiment. (right) It was replaced with a new model on the 80 cm experiment due to the presence of noise on some sensitive signals. The enclosure was manufactured by Equipto Electronics Corporation. This 'FCC' series heavy duty vertical rack is manufactured from cold rolled steel rather than a 400 series stainless steel. Racks from this manufacturer are available with slightly better attenuation than the one installed here.

4.18 System State

The constant change of systems on the experiment during operation makes it difficult to keep track of the relationships between critical systems. This is driven by the need to produce physics results at the earliest opportunity regardless of the engineering state of subsystems. The more 'proper' approach of optimizing systems off the experiment has been repeatedly neglected. The usual reason is lack of time and manpower. Premature integration leads to unreliable running of the experiment as a whole. Sometimes the experiment works well, sometimes it doesn't. Problems are always cropping up. One way to approach the organizational task is via a timeline listing the state of the different systems and the anticipated impact on overall

performance or quality of data. This aids in awareness of what information is available for analysis. For example, simply because we have a spectrometer does not mean that it is always connected to this experiment.

Many of the systems suffer limited malfunctions on a semi-regular basis. This is typical of systems that work ‘in theory’ but have not had sufficient debugging to point out deficiencies in design. Such situations are specified by the cross-hatch color patterns. A complete list of shots with failures (and their causes) is too extensive to include here. Instead we take a ‘general state of the system’ approach. Green represents proper operation with good effect. Red indicates poor operation or suspected serious impacts on the experimental operation.

Table 4.18.1. 80 cm experiment timeline. As can be seen, the majority of the systems were in a constant state of flux. Most shot sequences focused on optimizing the timing of the various systems for a given mode of operation. Significant numbers of shots were devoted purely to debugging certain subsystems. The timing of when the tungsten foil was added to the source is not definitively clear (Shot 740*); the displayed timing assumes that it happened at a recorded up to air opportunity when the glass center of the MPD channel had come loose.

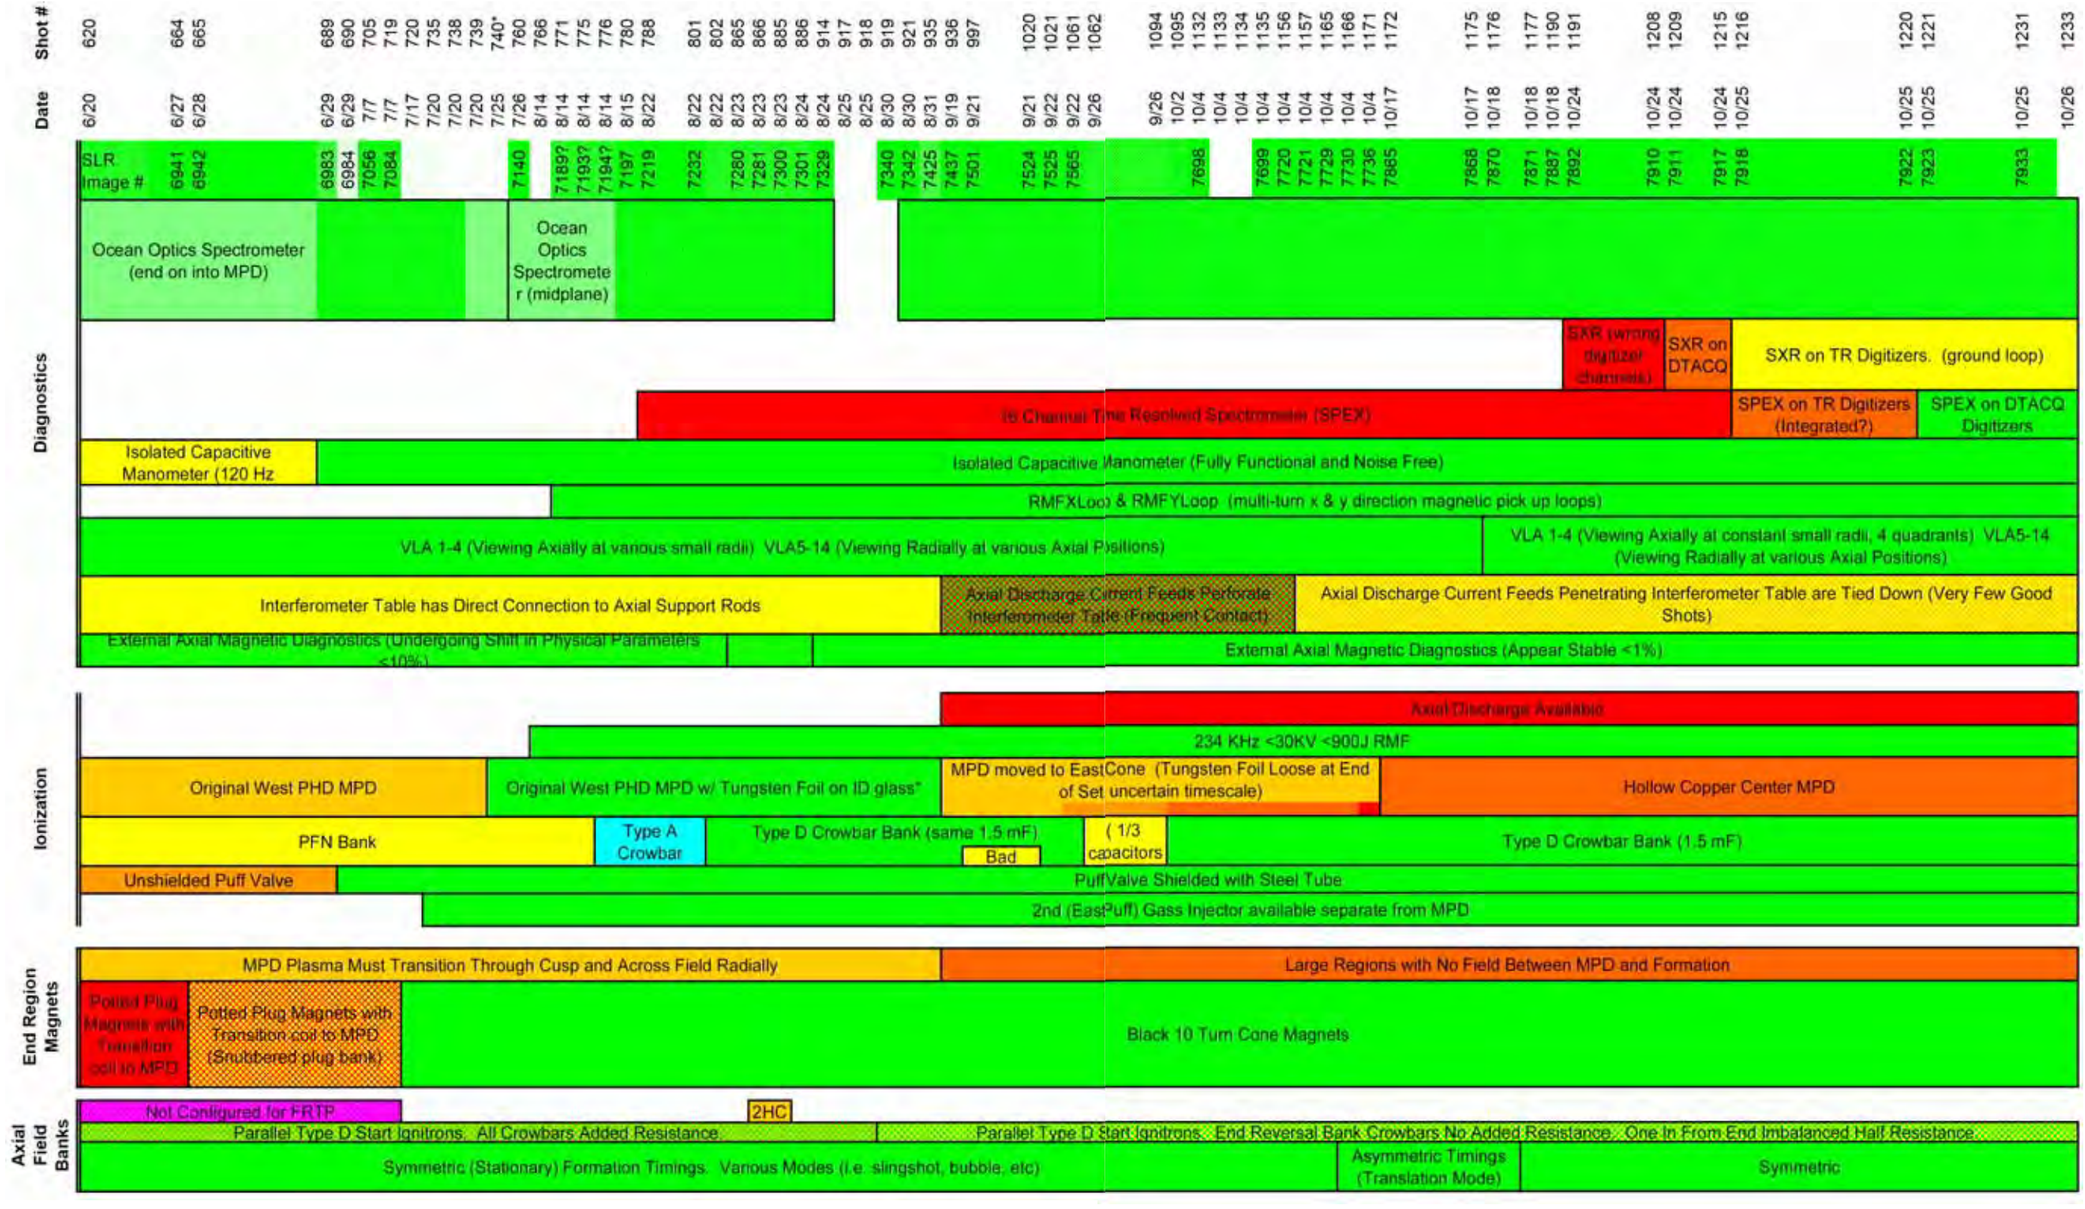

5
Diagnostics

This section describes and explains the 'simple' diagnostics in detail with the hopes of providing a thorough perspective on the complexities that are often glossed over or conveniently omitted. Conceptually simple diagnostics such as magnetic loops are considered in depth due to the general applicability of the lessons learned. The level of detail is intended to familiarize the reader with many of the critical concerns; but not always deep enough to consider every possibly relevant effect. Where readily available, manufacturer's specifications and published graphs have been used to illustrate key points. It is critically important to understand that when a parameter is not specified by the manufacturer this does not mean that the parameter will have negligible effect. Ultimately, space and time does not permit to consider every aspect of every diagnostic utilized in this dissertation.

5.1 Magnetic Loops

Perhaps the best way to explain why the axial flux and local B loops are implemented the way that they are is through a sort of design evolution. We will start from the most basic understanding of this diagnostic and gradually add elements to address various issues and uncertainties. In this process we will consider issues that also affect other diagnostics. Ultimately, the reader should understand the final design. It is also hoped that some understanding of why it is made this way and how to design similar systems will be transferred.

5.1.1 Driven Coil Voltage

Someone once asked why we do not simply monitor the voltage across the driven coil on the system. The simple answer is that you could. Consider a perfect inductor as taught in a basic electronics or physics class: $V = L\, di/dt = -\, d\Phi/dt$. If a square voltage pulse is used to drive the coil then the current in the coil ramps up for the duration of the pulse and then stays constant. Flux (in the coil) has the same behavior. Now let us add a non-ideal element, resistance of copper in the coil. Current ramps up to a level not as high as before and gradually decays exponentially on a timescale = L/R. Integrating the voltage to determine flux yields a signal that does not decay. This is clearly an error that has to be considered and possibly compensated for if this concept is going to work.

For the coils used in our system the driven straps are 6" wide x 0.11" thick x 116" long. Resistance is commonly simplified to $\rho L/A$. The resistivity of copper is 1.68×10^{-8} Ohm-m. Resistance then is approximately 100 μOhms at 0 Hz.

The inductance of the coil (with no plasma, and driven alone) is more difficult to solve for. A simple equation exists for multi turn loops with circular wire cross sections[62] that simplifies to $\mathrm{Lcircle} = R\,\mu_0\left(\mathrm{Log}\left[\frac{8R}{a}\right] - 2\right)$ for low frequencies. Our straps have highly elongated rectangular cross sections. A generic formulae for inductance is $L[\mu H] = F\, n^2\, d[in]$, where F is a function dependent on the length and

diameter of the coil, n is the number of turns, d is the diameter of the coil in inches. The function F can be solved for numerically or taken from a graphical representation of the function. It has proven difficult to obtain permission to reprint figures from reference books. Difficulties also exist in determining who currently owns the rights to some of the older sources. An example of a graphical representation of F can be found in "Reference Data for Radio Engineers" (Howard W. Sams & Co., Inc. ITT. New York. 1975. 6-1)[63]. Applying this formula yields an inductance of 1.51 μH at 0 Hz.

Inductance formulae for a wider range of geometries can be found in "Radio Instruments and Measurements" by the United States National Bureau of Standards published in 1918. It is also known as the Bureau of Standards Circular No C74. It is freely available through Google books online.

We can easily measure the single strap inductance with a meter made for this purpose (some care is required to deal with noise sources, such as 60 Hz power). The result is that @1 kHz we have 1.40 μH, 0 Ohms; @10 kHz 1.34 μH, 1 m-Ohms; @100 kHz 1.31 μH, 6 m-Ohms. We can also compare this calculation to the inductance of the total system by looking at actual data of a 36 μF bank fired into a single coil. We determine a ring frequency of 18.15 kHz. Neglecting non-ideal parameters in the system this yields a total system (coils + switches + feeds) inductance of 2.14 μH from $(2\pi f)^2 = \frac{1}{LC}$.

So the L/R timescale of the strap is ~ 13 ms (@ 0 Hz); this is conveniently much longer than the timescale of the experiment. However, there are problems with this calculation that we've hinted at. One of the problems is that the resistance and inductance are not constant for different frequencies. When field is first placed into a coil, the currents that support it are only at the surface of the conductor. Only as time passes do these currents (with magnetic field) diffuse in to the conductor. This is known as the "skin effect". It is one consideration that is important in designing high frequency (~ > 10 kHz) transformers and inductors. The skin depth[64] is $d = \sqrt{\frac{\rho}{\pi f \mu}}$, where ρ is the resistivity, and μ is the magnetic permeability. This simplifies to d = .065/sqrt(f) [m] for pure copper. Note that at sufficiently high frequencies it is only the surface (and surface condition) that matters electrically. The skin effect also reduces the inductance slightly because at high frequencies there is no magnetic field inside the conductor.

At 20 kHz (ring frequency of bank), the resistance increases by a factor of ~ 7 from local skin effect. Note that the frequencies present when a bank fires are not limited to the resonance frequency due to the sudden rise of voltage. Higher frequency signals need to be accurately recorded. There is also the tendency for the current to concentrate in the ends of the coil that is not taken into account. As to the inductance: in forming an FRC by theta pinch the plasma acts as a conductor at nearly the radius of the coil, reducing the inductance of the coil by a factor of 5-10 (80-90% coupled) during the reversal. Note that if the loop inductance is reduced (~10), the

system inductance is reduced (~2), the frequency of the system also goes up (~1.4) and resistance scales with the square root of frequency (~1.2). So this puts the strap with L/R ~ 100 μs for some length of time during reversal. Consider that we want accuracies in measuring flux better than 1%. The L/R timescale no longer looks as comfortably long and ignorable as it once did does it? We would also note that it would not be simple to compensate for errors due to the dependencies on frequency and plasma. There are additional flaws that from an engineering perspective make such a system difficult to properly implement. The additional flaws in this concept should become self apparent as we discuss noise while progressing through the design.

5.1.2 Unshielded Flux and B Loops (PHD)

Now we change our perspective to looking at the magnetic diagnostics as implemented. We have a coil driven by a +/- bank (the 'driven coil') and we have a second loop that is purely used to measure flux. They are separated from each other with two layers of polyethylene insulation (cyan). Another diagnostic loop, an azimuthally averaged B_z loop, is added beneath the flux loop.

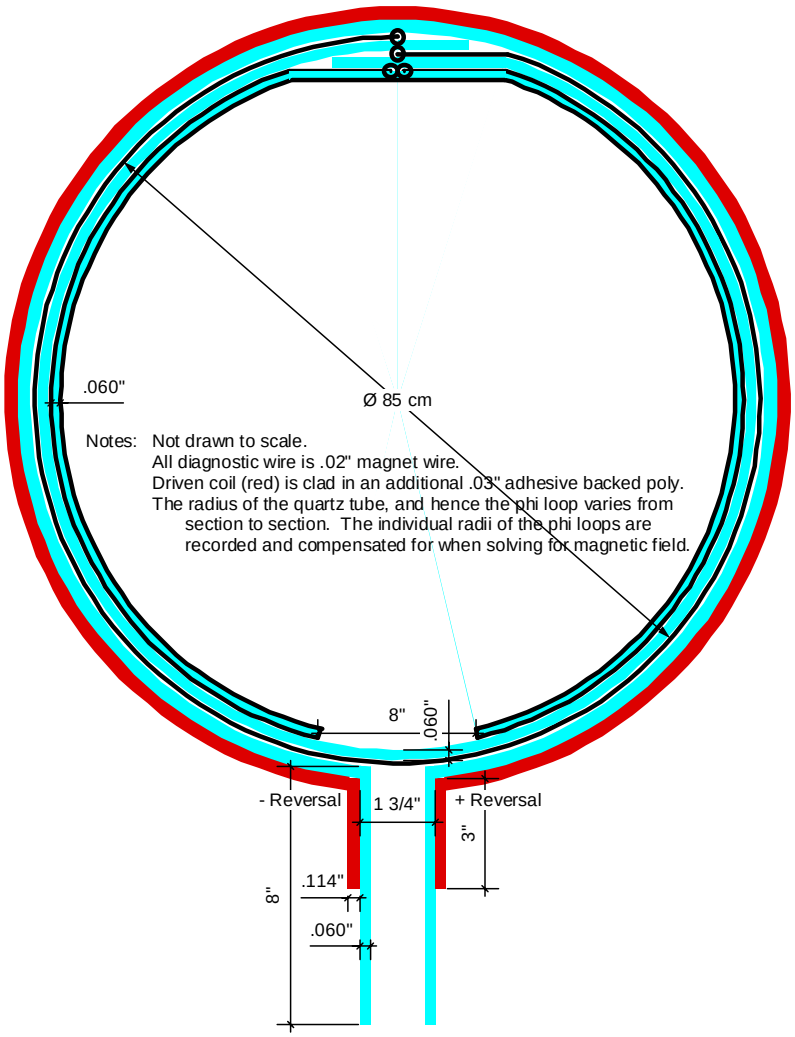

Figure 5.1.2.1. B and Phi loops as actually implemented during runs on 80 cm tube. Note that the effective area of a B loop is approximately 0.005 m^2 (the majority show effective areas 0.0042 to 0.0050), a Phi loop is ~ 0.567 m^2. A significant fraction of the cross section of the B loop is taken up by copper causing some potential problems in knowing the area. We are looking down the tube from the pump end (-Z) of the experiment. By defining excluded flux = magnetic field at edge * area of flux loop – flux, a positive magnetic field at the edge results in a positive excluded flux with an FRC inside (diamagnetic). The field introduced by the reversal banks is $-B_z$, resulting in a naturally negative excluded flux with an FRC present (diamagnetic). Various figures throughout this thesis are inverted so that excluded flux is positive. This is done because some individuals have difficulty accepting a definition that allows for 'negative' excluded flux that is not paramagnetic. For cases of a closed field line FRC, a diamagnetic situation is present when the excluded flux sign matches the external field sign; a paramagnetic situation is present when the excluded flux sign does not match the external field sign. This results in the misnomer of describing the equation's result as 'excluded flux' where the sign of excluded flux connotes the nature of the magnetism. Excluded flux in some experiments is calculated via simple electronics that are not capable of the sign inversion that would have to be added to the equation for the connotation to be correct. A more proper definition of excluded flux that allows the connotation to be true could be made: $X_f = (B_{ext}*A_\Phi - \Phi)*\text{sign(current in strap)}$ but this results in a discontinuity during formation (transitioning from the paramagnetic state to diamagnetic). It is better that the reader not associate sign of excluded flux with the nature of the magnetism.

The axial feeds for the diagnostic loops are on the opposite side from the feed points of the driven coil. This is done to limit the capacitive contamination from the coil on to the magnetic loops. One of the flux loop connections is grounded in the diagnostic rack so the precise geometry has implications for the energy stored in electric fields between the flux and driven loops. Some current still flows through this ground, from the capacitor banks, for imbalances between +/- banks. An accurate analysis taking into account details (such as resistivity of conductors and more importantly the geometry) is likely to be difficult to do. The sort of task we would rather model with a computer than try to puzzle through analytically. So instead we will make the system such that it has as little dependence on these factors as possible. We will address the potential issues one at a time.

5.1.3 Self-Resonant Frequency

It is likely that driving the coil results in oscillations in the flux loop, the important thing to realize is that the capacitance is relatively small. The most expedient way to estimate the capacitance between the driven coil and flux loop is to consider them as a transmission line. To be specific, the flux loop is a wire over the strap's plane. There are programs and formulas available on the web to give characteristic impedance and phase velocity for various transmission line geometries and materials. Knowing that $V_p = \frac{1}{\sqrt{L'C'}}$; $Z_0 = \sqrt{\frac{L'}{C'}}$, it is simple to solve for L' and C'. For the flux loop/strap system this comes to 57 pF/m; total of 150 pF. As a side note, we are rather partial to a program (free) called AppCAD[65], made by Agilent Technologies for this and other purposes.

Additionally, an important point to make is that the inductance of the flux loop is significantly larger (4 μH) and the resistance larger (.25 ohm) than that of the driven coil due to the small diameter of magnet wire. Obviously, most of this extra inductance can't be coupled to the driven strap (since they are both single turn systems) so in the sense of a 1D circuit analysis this can be considered as stray. This is one reason it is preferable to connect diagnostic loops to electronics with reasonable impedances. We don't want the L/R times getting long.

As a gross approximation, the worst case scenario is to model the flux loop as a series resonant circuit with a voltage source, followed by an inductor and resistor going into a capacitor. It is the voltage across the inductor plus resistor that is of interest as the noise signal. The stray inductance can be taken to be 2 μH. The capacitance can be found by taking the total capacitance and splitting it into the area of the two output leads so that they conduct in series (40 pF effective) closing the flux loop. Note that the resultant oscillation from a pulse is at ~ 18 MHz with a Q near 900 (not taking into account dielectric dissipation and the very important effect of the "skin effect"). This minimum frequency is beyond our range of concern (so further, more accurate analysis is not necessary). Any remaining side-effects are not model-able in an analytic way without significant effort.

Any imbalance in the drive voltages gets coupled to the flux loop's ground and high voltage connection through the 150 pF capacitance. To further the analysis we must consider the nature of the output of the flux loop.

5.1.4 Passive/Numerical Integration

For dealing with the voltage output from a flux loop ($\alpha\ \delta\varphi/\delta t$) there are two basic extreme options with a range of choices in between. The signal can be integrated in electronics to supply flux or recorded as the derivative of flux. The real versions of these extremes are never pure. However, there are very good reasons to prefer doing something closer to the electronic (accomplished by the circuit) integration. Quite simply, the timescales for which we can expect there to be high $\delta\varphi/\delta t$ are quite short (< 10 μs) yet we wish to observe the FRC for long timescales (> 1 ms). Put another way, we expect the voltage signal during reversal to be >100 times greater than that during sustainment. So, if the digitizer is scaled for 256 counts during the beginning, the signal during decay will be 2 counts (or is that 3…). The error-bars grow as time passes even for an ideal system. There are techniques that can be used (introducing an oscillation to gain sub-count resolution for example) to lower the errors but they require significant effort to do correctly. They also require that the base noise spectrum of the digitizer is significantly less than one count for low frequencies (an assumption not to be made lightly). An additional cost to doing numerical integration is the necessity of recording the entire signal from start to end.

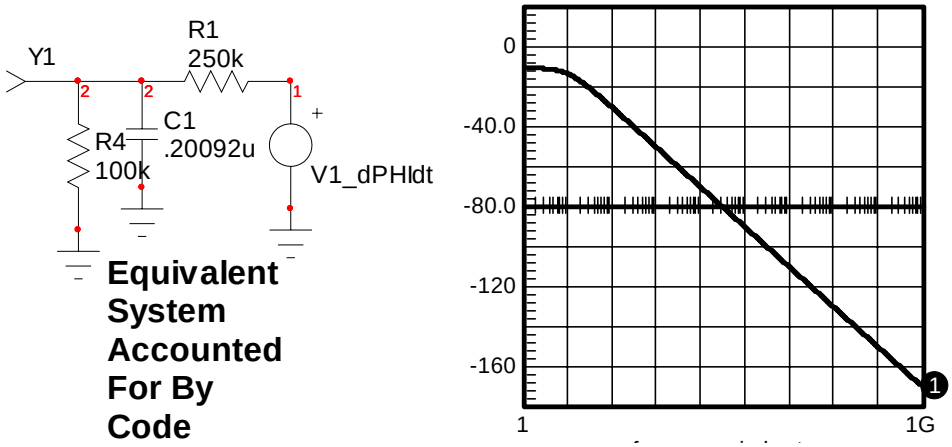

Figure 5.1.4.1. Simplest model of RC integration with the resultant frequency amplitude transform.

Most passive input structures can be simplified to the circuit equivalent shown in figure 5.1.4.1 at left with the attendant frequency response curve shown at right. A pure integrator would not have the flat frequency response at low frequency. A perfect voltage monitor would not have the constant -20 db/decade sloped section. Any approach that can be simplified to this configuration can be solved to find flux simply as $-\phi = R_1 C_1 V_{Y1} + \frac{R_1 + R_4}{R_4} \int V_{Y1}\, dt$. The integration term is solved numerically

to compensate for the low frequency response. The Matlab© code that does this correction for a flux loop looks like this:

```
function [data,time]=C00620PHI05_WB(nShot)
%GE: -phi=Rs*C*Vo+((Rs+Rl)/Rl)*Int(Vo)
%phi is flux in loop, C integration +cable capacitance, Rs resistor in
        series with loop, Rl digitizer, Vo recorded signal
sItem='FL05';
Rs=252200;      %Loop resistance
Rl=90795;       %digitizer load
C=202.6E-9;     %Integration Capacitance
a=Rs*C;         %Rs*C
sign=1;         %For reversing the sign of the output

d=datExpr(sItem,nShot,0,0);
d.y=d.y-mean(d.y(1:100));   %Zero early time signal
d.y=a.*d.y+((Rs+Rl)./Rl).*mean(diff(d.t)).*cumsum(d.y);
data=sign.*d.y;
time=d.t;
```

Figure 5.1.4.2. Simple passive integrator correction code.

The code for a flux loop is very simple and concise. We chose to keep the same integration capacitor value for both flux and B loops. So, since the B loop has less area, the integration resistor R1 (Rs) needs to be smaller (3000 ohms) to get the same signal size. These values were chosen as a result of factors (and uncertainties) that we will soon discuss indirectly. For the B loops the integration correction factor is more important. This helps demonstrate some of the difficulties in doing a numerically integrated system. For example, we have some fields that come on early in time near the ends (plug coils) with quarter cycle times longer than 1 ms. The timescale is long in order to penetrate the $1/16^{th}$ inch stainless-clad cones. The nearest flux/B loop is 40 cm axial distance. How faithfully the B signal matches the flux signal varies from shot to shot due to the numerical integration. A pair of shots are plotted below to demonstrate. Some of the loops (not shown here) are consistently off by factors ~2 for these longer timescales.

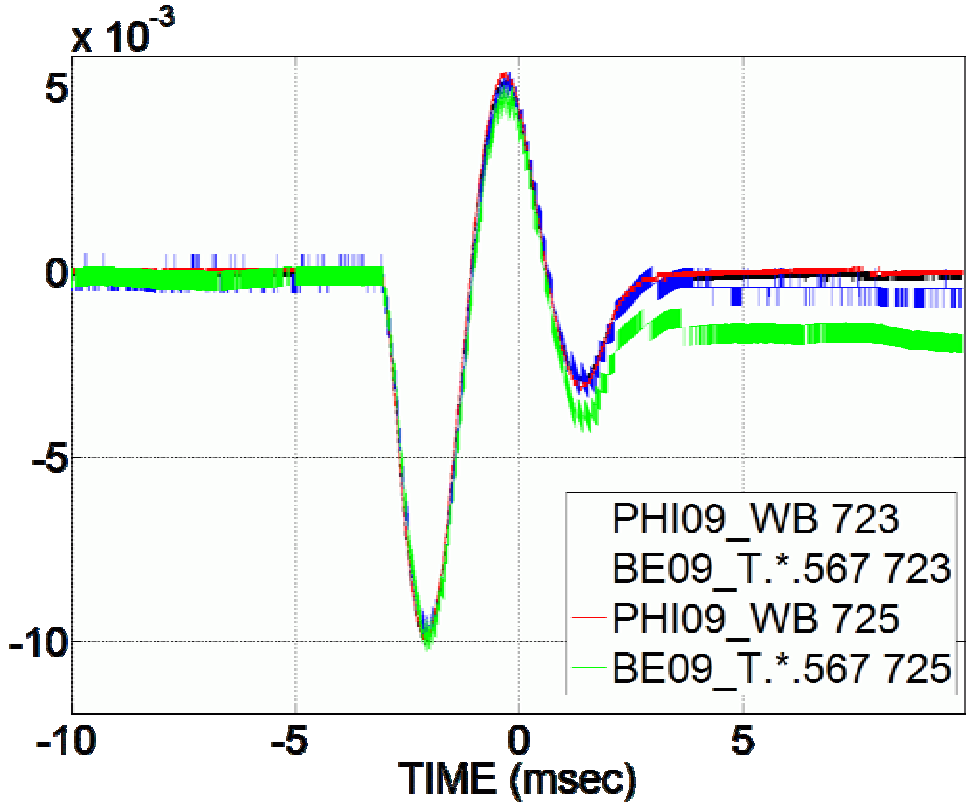

Figure 5.1.4.3. Typical magnetic signals for long time scales. Only the plug banks are firing. Note that the B signal has been normalized to the flux signal; rather than being left as a local B_z.

The flux from the plug bank at the nearest reversal coil is ~1/20th the peak flux during reversal. The flux loops see the plug fields and the numerical correction for not being a perfect integrator is ~7% at the time of peak plug field. The B loop correction is ~30%. The B loops do not consistently get sufficient signal to accurately record this field. We are using 12-bit digitizers scaled for the larger fields near time of reversal. The practical downside is that for the flux loop correction to work, we must collect data from before any fields are fired. As can be seen above this can require 10 ms (100,000 points). One of the worst B loop cases is shown below with plugs, bias, reversal banks, and crowbars firing for a loop set at the end of the machine.

There are additional complications in interpretation due to varying B_z under the flux loop; this is causing most of the error here. We actually take the flux loop signal and force the B loop signal to be consistent before any of the bias or reversal banks are fired to make a simple excluded flux signal. More complex analysis has been done using the real local B_z interpretation that can provide a significant increase in excluded flux from what is immediately apparent. However, it is a computationally taxing process. We will further consider these issues later.

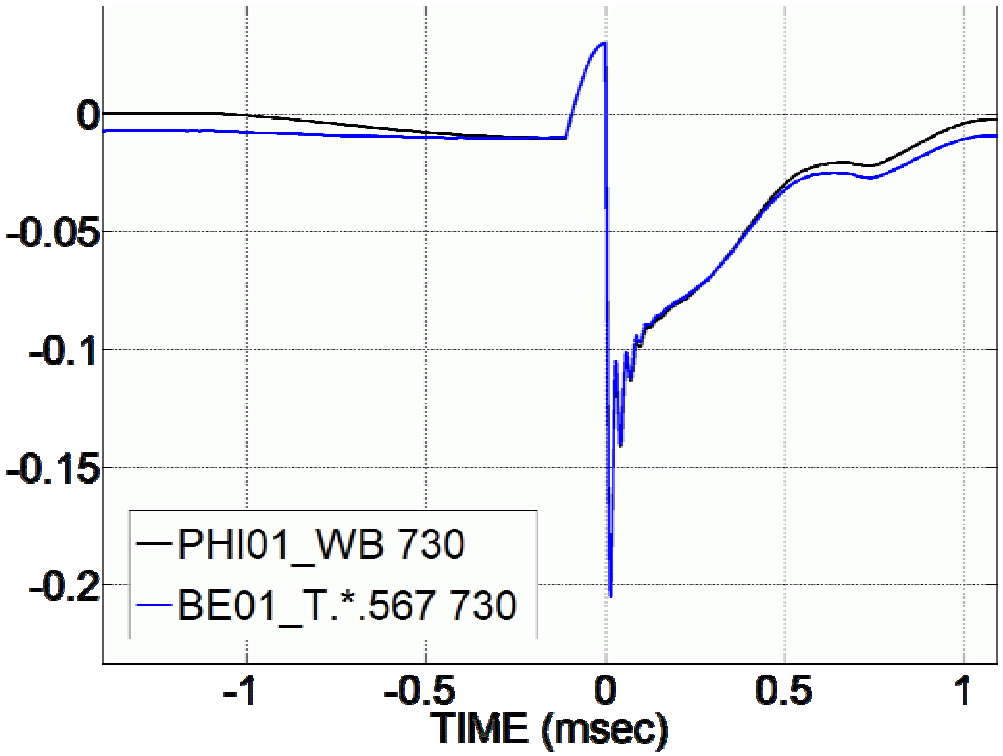

Figure 5.1.4.4. Typical Flux and B_{ez} Traces. The mismatch at long time scales is primarily due to the locality of the B_{ez} measurement.

5.1.5 Capacitive Coupling of Driven Coil to Flux Loop

Now that a basic structure for the electronics is laid out, we can further address the question: What is the impact of the 150 pF capacitance between driven coil and flux loop? Below is a diagram that describes the relevant physical connections in the system. A single driven loop is shown in red along with a basic schematic of the driving circuit. Note that there are connections through the bias modules to other straps. Similarly, the flux loops are all grounded together as they enter the screen rack (not shown).

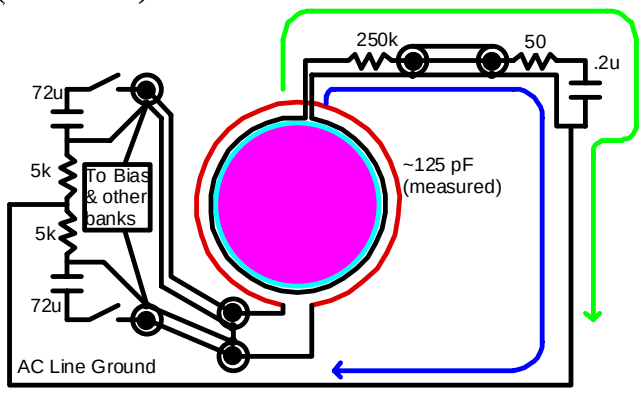

Figure 5.1.5.1. High frequency ground loop.

To gain familiarity with the system, let us consider the case where the braid connection at the flux loop is disconnected so that the dominant current path is green. To illuminate matters further let us assume that only one of the bank switches is conducting, placing 25 kV on to the coil. Any current that passes through the 125 pF capacitance must be seen across the 0.2 µF integration capacitor. So the voltage left on the output signal is $\frac{25000\ V * 125\ pF}{200\ nF} = 19\ V$. The output error signal exponentially approaches this level on a timescale dominated by 250 kΩ x 125 pF = 31 µs. A typical signal is ~ 5 V with no plasma present.

Now let us include the blue current path. The primary difficulty for transforming this layout into a circuit diagram is in estimating the inductances involved. The circuit can be approximated into a 0D form like that shown in Figure 5.1.5.2.

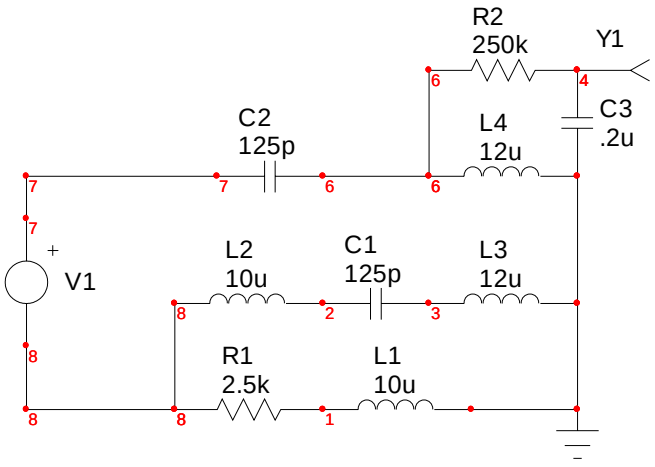

Figure 5.1.5.2. Simplified model of Capacitive Coupling to Flux Loop.

V1 represents the voltage imbalance between the banks. C2 represents the coupling capacitance from strap to flux loop. L4 represents the inductance associated with the braid of the coax going to the screen rack. R2 is the high voltage integration resistor (we are neglecting the inductance associated with the center of the coax for the obvious L/R reason). C3 is the capacitor across which we measure the output signal. L2, C1, L3 appear due to an adjacent strap driven by the same bias bank but separate reversal bank; there are more like this not considered. L2+L3+L4 has been measured (after we estimated these values) and found consistent, the difficulty is in assigning the ratio of L4 to L2+L3 for a complex geometry.

The inductance of a wire can be estimated without describing how the current into it is fed[66]. The analysis is based on defining inductance as the ratio of flux to current. Calculate the field in the space limited by the length of the wire, out to a

radius of infinity. The solutions for low frequencies and high frequencies (skin depth limited) are Llf = 0.2 (ln(2L/r) - 0.75) [µH/m], Lhf = 0.2 (ln(2L/r) - 1.00) [µH/m]. For ¼" coax braid the inductance is ~ 1 µH/m. Some additional inductance is added to L4 equivalent to the inductance of the flux loop. The last elements of the circuit, R1 and L1 are the safety ground connections going to the "low voltage" terminal on the module capacitors.

It is reasonable to note that this is a crude representation (not including various inductive couplings or multiple driven modes) that is more intended to give an understanding of the effects than to give a perfectly accurate representation. One might suspect that the biggest simplification is made in treating the coax so simply. If the coax were twisted pair and the high voltage resistor were 0 ohms, it is clear by symmetry that half the current would flow through the 'resistor'. Making the resistor non-zero only lowers this current to a level limited by the resistor. The transformer action of the cable is present with twisted pair as it is with coax. Note that the resistor is not inside the coax (i.e. this is not a fully electro-statically shielded system).

For comparison, we will also plot the case where L4 -> infinity.

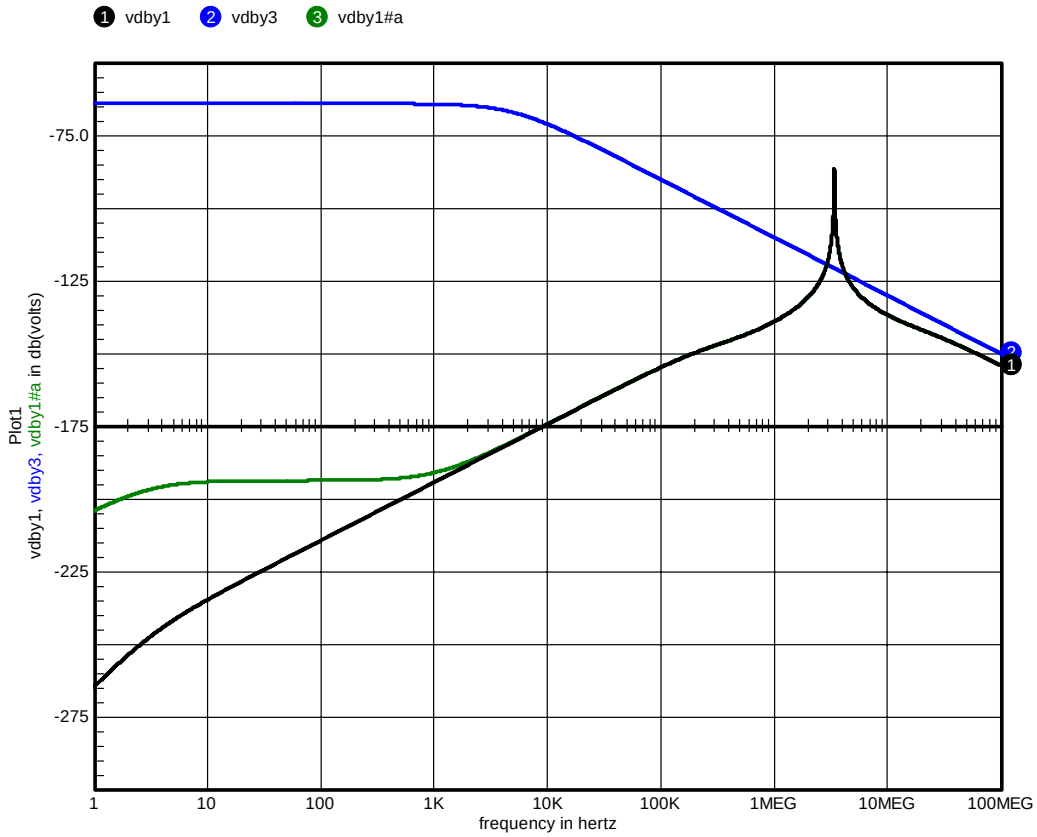

Figure 5.1.5.3. Frequency response of flux loop system to driven bank capacitive coupling. The blue trace corresponds with the flux loop having no ground connection. The black trace with the diagram just displayed. The green trace adds coax shield resistance. Mutual inductances and some additional system connections are neglected. Digitizer threshold is ~ -140 dB. Keep in mind that any impulse to this system is likely to have many distributed frequencies.

At low frequencies, the blue trace (~ -63 dB) is consistent with the 19 V output calculated earlier. The green trace adds an appropriate constant resistance (DC) to the coaxial braid's inductance of L4. The threshold for registering on the digitizer is approximately -140 dB. So even in the worst case we would not expect to see any noise via this path at any frequency below 400 kHz. The peaking of the response curve near 3.5 MHz can be understood by considering L2, C1, L3, L4, C2 as a series resonant system lacking the real resistance that would broaden and flatten the peak. Perhaps the most important lesson from this example is that keeping L4 and its resistance small is very important. There is a tradeoff between the frequency at which the system resonates and the magnitude of the current carried through the coax braid. Ideally, one would want to keep this resonance frequency high (out of the range of interest) except that this means potentially higher currents are carried. Note that adding chokes pushes the frequency lower and for the effects to be apparent the chokes have to be quite large (>20 μH). The typical snap on ferrite is ~ 1 μH.

5.1.6 Capacitive Coupling of Driven Coil to B Loop

For the B loops, this problem is approximately 40 dB worse due to the lowered integration resistance (250 kΩ -> 3 kΩ). The effective coupling capacitance to the strap is nearly the same. There is some capacitive coupling to the naturally unbalanced flux loop as well. The digitizer threshold is still -140 dB. This is still a 25 kV worst case scenario. It should be practical to reduce the effect by 20 to 40 dB (not counting the coupling to the flux loop) when +/- modules are well balanced in charge and firing timing.

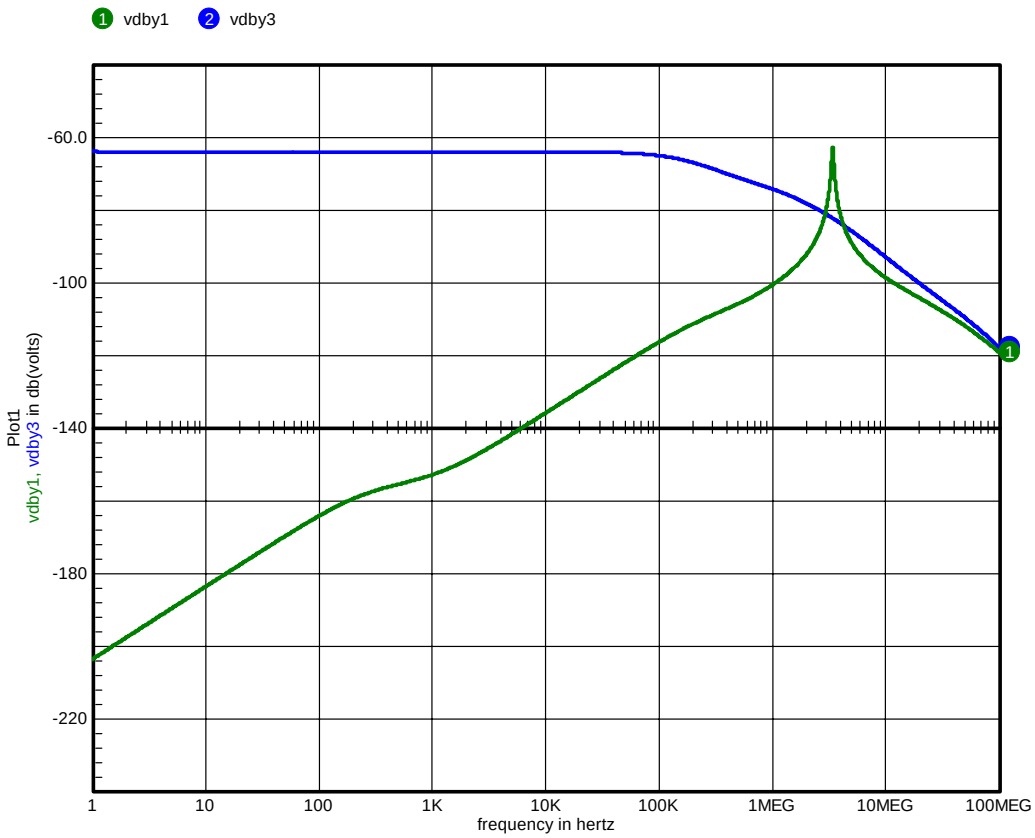

Figure 5.1.6.1. Frequency response of B_{ez} loop system to driven bank capacitive coupling. The blue trace corresponds with the B loop having no ground connection. The green trace adds coax shield resistance and the ground connection. Mutual inductances and some additional system connections are neglected. Digitizer threshold is ~ -140 dB. Keep in mind that any impulse to this system is likely to have many distributed frequencies (and hence less energy in any one frequency).

5.1.7 Shielding Effectiveness of Coax

Energy transfer occurs between the inside and outside of coaxial transmission lines. The appropriate term to look up for this effect (current on the braid inducing an electric field in the central dielectric) is cable "transfer impedance". At low frequencies the impedance is the dc resistance of the braid. If the braid was solid and the coax properly connected on both ends, the skin effect would lead to a lowering of the electric field on the inside of the cable as frequency increases (lower transfer impedance).

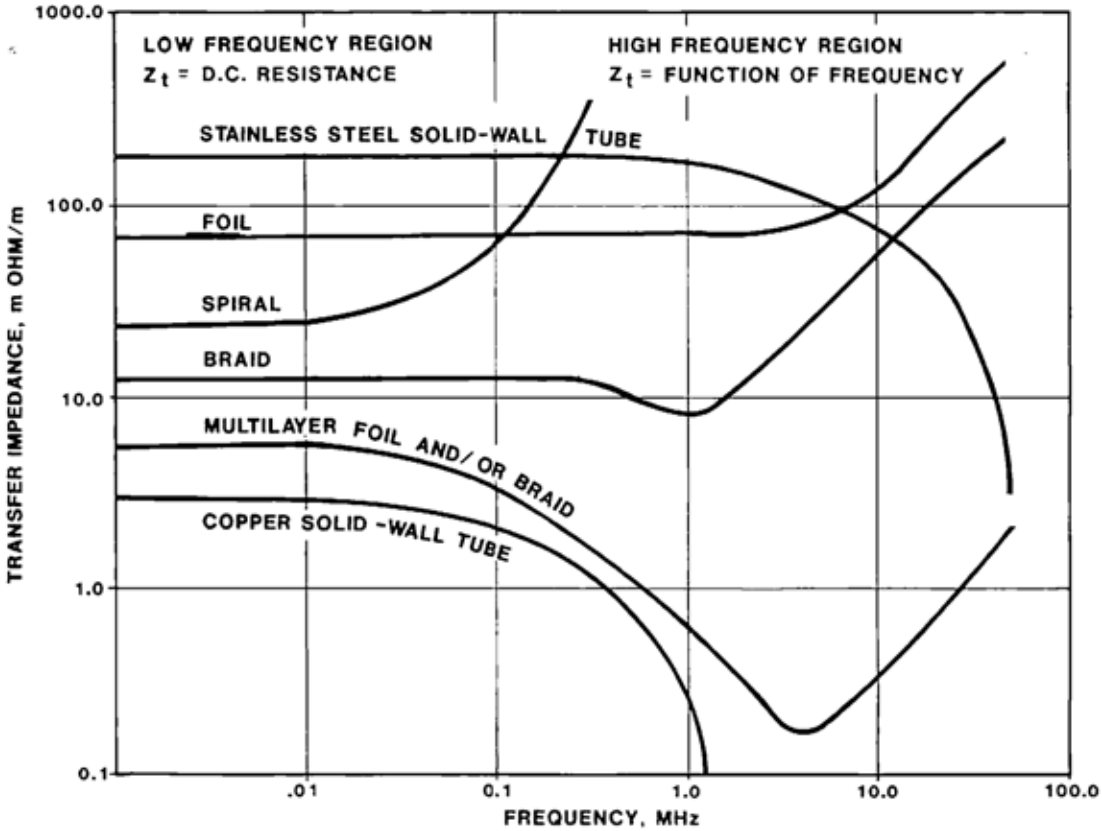

Figure 5.1.7.1. Transfer Impedance. Low transfer impedance is desirable. We commonly use RG223, a two braid coax system. Flexible coax is available with better shielding. In particular the DC transfer impedance is reduced and the minimum is extended out to 30 MHz and 100 MHz for some options (3 or 4 layers of foil or braid)[67]. This is figure 6 from "Shielded Electronic Cables for EMI Protection" by Anatoly Tsaliovich (Belden Innovators - Vol 3 No 13 - December 1982). Used with permission. See this reference for another representation of a transfer impedance plot [68].

As can be seen from the figure, RG-58 braid is not equivalent to a solid metal pipe. The weaving braid carries what would have been the outer surface currents on to the inner surface. Consider that the ~3.5 MHz ring frequency carries ~34 amps for the 25 kV worst case scenario. We use RG-223, a double braided version of RG-58 that is the same as RG-55 (RG-55 got renamed) and gives an extra 20-30 dB of shielding. We've learned this from experience, after having made the mistake of using RG-58. At this frequency with ~ 10 m of cable the transfer impedance is ~ 0.007 ohms (RG-223) leading to a worst case signal of ~ 0.23 volts that gets integrated through 50 ohms and 0.2 μF to produce ~ 1 mV for both B_{ez} and Flux diagnostics. This is just below the threshold of the digitizers.

So, some of the current that was shunted through L4 in fact ends up bypassing R2 to end up on the integration capacitor (and is quite noticeable with RG58). This makes the coupling significantly worse than the graphs shown earlier (if we were still using RG58). Caveat emptor, direct application of transfer impedance plots assumes proper impedance matching on both ends of the cable; on one end of our cable the impedance is very high, the other end is well matched at the frequencies of interest.

5.1.8 High Voltage Resistors

So, at this point we've covered the dominant forms of noise pickup and considered a simplified model for passive integration. What we've yet to discuss are the issues associated with the integrator's high voltage resistors and high frequency capacitors. They are two problems that tend to be overlooked. For example, it has often been common practice to use a string of carbon composition resistors to fill the role of a high voltage resistor. If this practice is suggested to a high voltage power supply designer, that he use carbon comps for control, he'll laugh or look at you kind of strangely. On the other hand, many audiophiles actively seek products that use them. The reason can be summed up in four words: voltage coefficient of resistance.

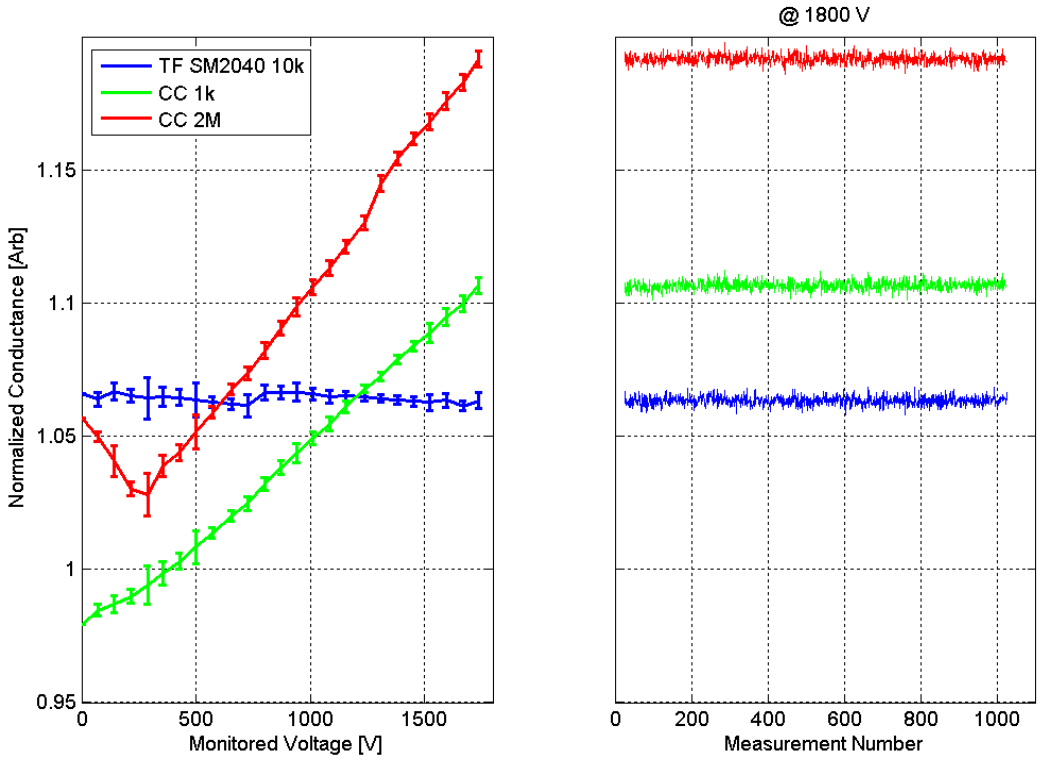

Figure 5.1.8.1. High Voltage Resistance. (left) Normalized conductance (resistance $^{-1}$) as a function of applied voltage is plotted for a 10 kΩ thick film resistor, 1 kΩ carbon composition resistor, and a 2 MΩ carbon composition resistor. These measurements are taken from numerous short (< 3 μs) pulses at voltage. This approach avoids significant heating of the resistors (and is a closer match to conditions encountered in the experiment). (right) No significant changes in the parts can bee seen for these low voltages. A plot for DC applied voltage results can be found in Horowitz and Hill[69].

The resistance of a Carbon Composition (CC) resistor typically decreases as voltage increases. This is quantified as the voltage coefficient of resistance (VCR). Note that the effect displayed in the graph is not due to temperature rise. The effect is dependent on the construction and particular composition of the resistor. From the few CC resistor manufacturers who publish the coefficient we gather it can be

guaranteed (with some effort) to a maximum that varies between -.005%/Volt to -.05%/Volt[70] dependant on value. We've taken low value CC resistors off the shelf and pulse tested them to see typical values of -.004 to -.008 %/Volt. There is an additional effect due to temperature rise that varies from -.07%/C° to -.3%/C° dependent mostly on the value of the resistor. In some ways a worse problem, carbon composition resistors gradually change their values permanently (to unpredictable levels) when exposed to sufficiently high short time overloads.

The truly devilish aspect of these resistors is the reason many audiophiles love them. The resistors add harmonics due to their high VCR that results in richer tones (perceptible to the human ear even at low voltages). Humans hear in a way that is closer conceptually to a spectrum analyzer than a digitizer[71]. The recorded signal still has the same quarter cycle time (still acts like a resistor), but the level is corrupted and new frequencies introduced. This result of a nonlinear system has the undesirable property of being difficult to unfold (as a practical matter) without a significant amount of characterization. These effects are not nearly as apparent as those from making truly obvious mistakes such as using low value wire-wound resistors.

If a designer must use carbon composition resistors for diagnostic signal applications, then it is best to use them in same value stacks (at low voltage per element) as resistive dividers (into a high impedance load) and not integration resistors. One should realize though that many of these non-ideal characteristics are not consistent from resistor to resistor in a batch; nor are they guaranteed to be consistent for resistors with original values that match.

The best class of resistors to use for high voltage diagnostic applications is the HV thick film type[72]. There are resistors in this class with VCRs of .04 ppm/Volt and temperature coefficients of 5 ppm/°C. Their active elements usually consist of a ruthenium oxide matrix, which can be modeled as a random resistor network. Care should be taken because this general class of resistor can still change value from exposure to voltages over their manufacturing conditioned maximum[73] [74] [75] [76]. It doesn't take a very long exposure (order μs) to result in a permanent shift in value[77]. This is a localized internal insulator breakdown phenomenon. These resistors generally consist of an aluminum oxide substrate with ink (consisting of ruthenium oxide and glassy compounds[78]) that is ultimately fired to produce a solid resistor (with no significant loss of mass). The implication is that on a very short timescale the joule heating is deposited into only the ink.

Most components in this class have no joule surge rating because they are intended to be used in continuous conditions, not pulsed. Some manufacturer's engineers are willing to provide such ratings in private communications. Others, including the components we are currently using, only provide the gross dimensions without the specifics of ink composition. So we must calculate and verify our own ratings; or use them well within their continuous power rating. There are two aspects of concern: melting the element and cracking the element due to thermal stresses. A third concern, the element acting like a stress gauge, is not of high concern for our setup due to the very low thermal coefficient; if this effect were unexpectedly important we'd see a tell-tale signature in the experiment from the resistance

changing as a function of time. A characteristic to watch for is a signal that does not return to zero at the proper time or rate. Some thick resistive films are used as stress gauges but the effect is very small. Highly sensitive bridge configurations are required to measure the impedance change.

The simplest model of these resistors' thermodynamics would be a thin dissipative element on top of an alumina substrate. This lends itself nicely to a 1D thermal model. The exact composition varies for different products, and this is obviously proprietary information. Fortunately, alumina, silica, and other common materials tend to have specific heat capacities in the range of 0.5-0.8 joules / (gm K). Alumina has a thermal conductivity of ~ 18 W / (m K). The dissipative layer is 12 microns thick. The alumina substrate is 600 microns.

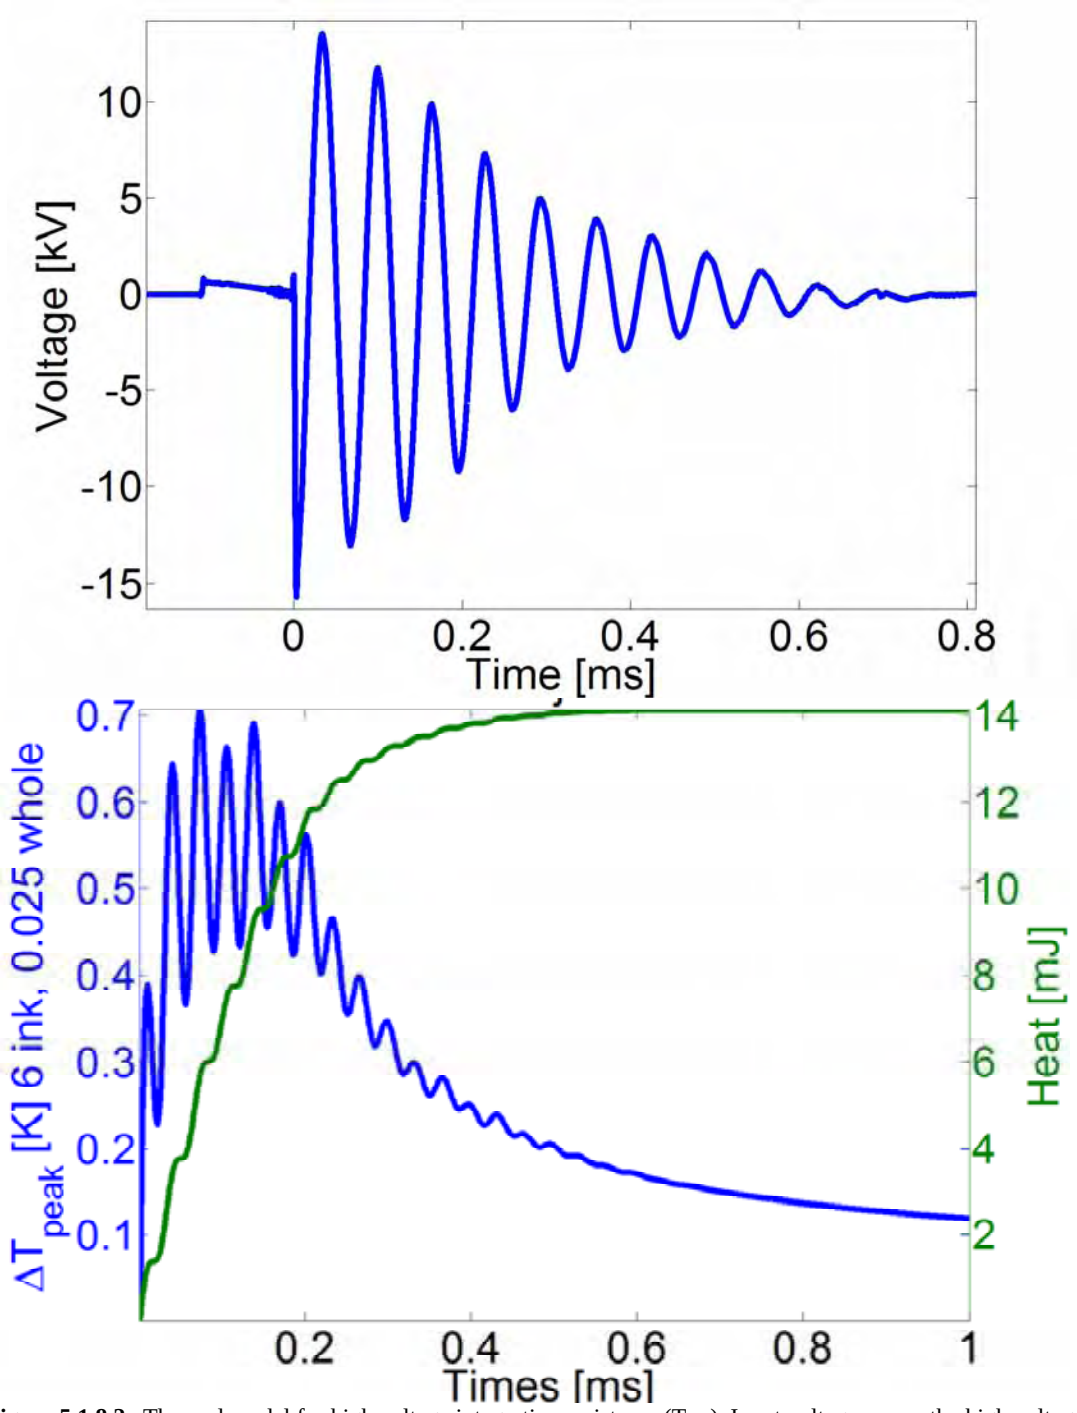

Figure 5.1.8.2. Thermal model for high voltage integration resistors. (Top) Input voltage across the high voltage resistor stack for a +/- 10 kV reversal bank shot. (Bottom) Peak temperature in a resistor and total heat stored in the resistor. If the resistor consisted of only the ink, the change in temperature would come to 6 K. Over a longer timescale the temperature increase comes to 0.025 K. Note that our resistor stacks consist of multiple SM2040s. The voltage is applied across the entire stack, not just the one element.

The results of the 1D thermal model for the system as implemented are shown. This case uses the voltages generated by a ringing theta shot as input power into the resistors. If the resistor were modeled as only consisting of the dissipative ink (0D), the peak temperature increase would have been 6 K. Similarly, if the resistor was modeled as the alumina substrate (0D), the result is only 0.025 K. Hence, a 1D model was indeed necessary (0.7 K˚). Note the change in ratio between ΔT and the heat deposited in the resistor. The (0D equivalent) heat capacity of the resistor changes as a function of frequency. If the bank were running at full power (+/-25 kV) the temperature would reach 4 K (5 series stacked 50 kΩ resistors). This suggests that in later experiments it may be safe to lower the resistance (and improve the frequency response issues to be discussed shortly). While these resistors are rated at 100˚ C, the true limit may be set by the internal thermal stresses. Thankfully, if they do crack it rapidly becomes apparent.

Unfortunately, it is very difficult to make a compact resistor capable of withstanding these high voltages. Surface tracking limits us to >= 5 inch long resistor stacks. Corona is an added concern. Because of the size of these components, the capacitance formed between the resistor and ground is also a serious concern. The inductance is not of such concern (due to the short L/R times). The high voltage resistor stacks can be modeled in the standard transmission line way.

The resistors are planar resistors with a width equal to 0.5". The 0.02" diameter ground wire runs the length of the resistor, separated by the 0.07" FR4 board, 0.06" polyethylene, and a variable air gap. Being able to calculate the effective permittivity (~3 ε_0), it is simply a matter of plugging the numbers into the previously mentioned AppCAD to find C' = 65 pF/m, L' = 0.51 uH/m. Inserting the numbers into a SPICE model is then straightforward. Unfortunately, the air gap was not well controlled. After doing the analysis detailed next; an inspection of the resistors showed that the air gaps varied considerably from one diagnostic location to another.

The wary SPICE modeler will realize that there are several different forms of transmission line equations. Some are meant to model dielectric losses at high frequencies; test the model to ensure the expected low frequency behavior. It is not unusual for the default coax models to show more DC resistance than is actually present. This is because it is common to use a resistance term to match the dielectric dissipation at some fixed (very high) frequency for DC relevant models. Caveat emptor, the standard transmission line models are effectively 1D models; energy is only allowed to propagate along the length in one dimension. The finite width of the resistor, combined with the single ground wire, may cause an error in this analysis both in frequency response and dispersion. We are uncertain of the magnitude of the error at this point but assume it is negligible compared to the primary correction. Note that SPICE transmission line models do not take into account parallel capacitance (axial capacitance across the resistive medium). This leaves us with the non-ideal solution of lumping this capacitance together, bypassing the transmission line. The effect of the bypass capacitor is larger than it would be if it were distributed, but has most of its effects at higher frequencies. For the readers' enlightenment, in resistive divider applications we would normally add a bypass capacitor of

sufficiently large size to dominate over any parasitic resistor effects. However, there is a serious price associated with this approach in terms of oscillations.

Ideally, we would be using a cylindrical resistor with a cylindrical return shell. This would make the analysis more tractable (3D -> 2D) in a truly correct fashion. Such parts are only available as custom manufactured components for such low values of resistance and hence are rather expensive with large minimum order quantities. We will return to the effects of resistor construction on frequency response shortly when we consider the full system model.

5.1.9 Capacitors

As with resistors, not all capacitor types are created equal. The class of capacitor that is closest to being "perfect" is the C0G (also called NPO) type ceramic capacitor. It is a class I dielectric (high temperature stability) commonly used in analog signal path circuits. Its chief downside is low energy density (< 0.1 uF in a 25 V 1206 case). These are the capacitors we typically use in diagnostics as they are nearly "perfect", and of sufficiently high values to be practically useful for our applications (unlike, for example, vacuum capacitors).

Most capacitor types are very far from perfect. All capacitors have a self resonant frequency determined by their geometry. This is one of the reasons why surface-mount components are so popular. There is no way to get less stray inductance than to mount (an already good device) with no leads to a transmission line. For very high frequency circuits, circuit boards can be manufactured with a distributed ceramic supply capacitor built into the layers of the board.

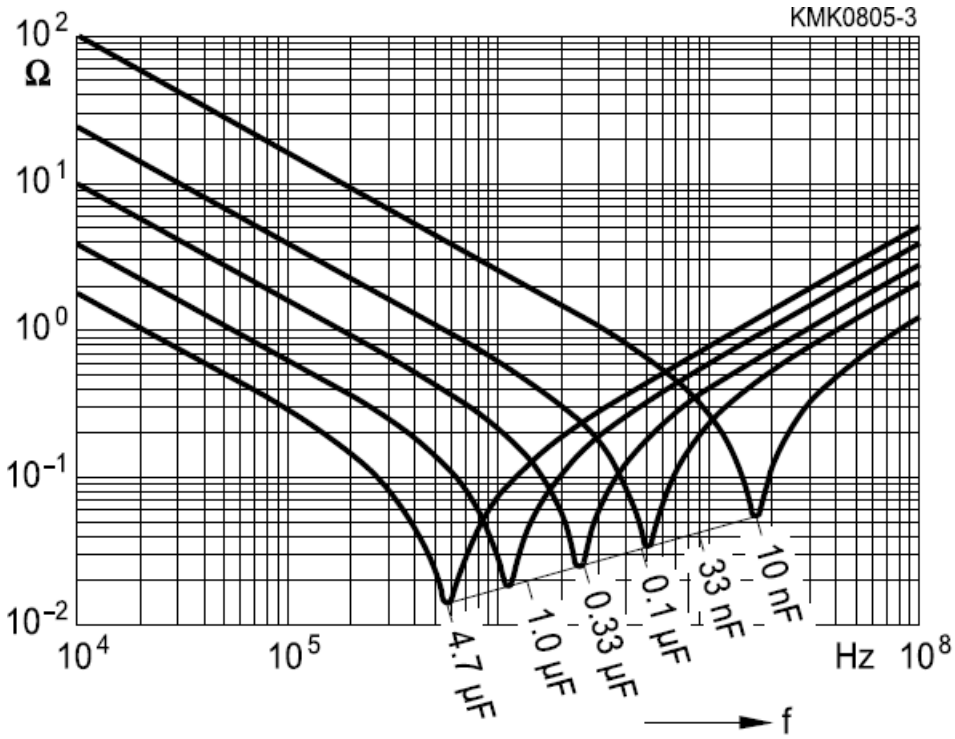

Figure 5.1.9.1. Typical resonant frequencies for low inductance, through-hole, capacitors. The 0.1 uF surface mount capacitors we use resonate near 20 MHz, the equivalent through-hole part here resonates closer to 5 MHz. Reproduced with EPCOS AG permission from datasheet[79].

The vast majority of dielectric types have significant changes as a function of temperature, frequency, humidity, electric field, and age. I've chosen to highlight a few of them here. After C0G ceramics the next best class of resistor for diagnostic use is plastic films (polypropylene, polyester, polystyrene, mica, etc). Like resistors, most of the non-idealities are not obvious from looking at datasheets. For example, a critical effect for integration capacitors can be dielectric absorption (memory effect), but this is almost never specified in component datasheets, instead it is usually lumped together with dielectric dissipation under dissipation factor[80].

Here are some typical graphs describing capacitor characteristics:

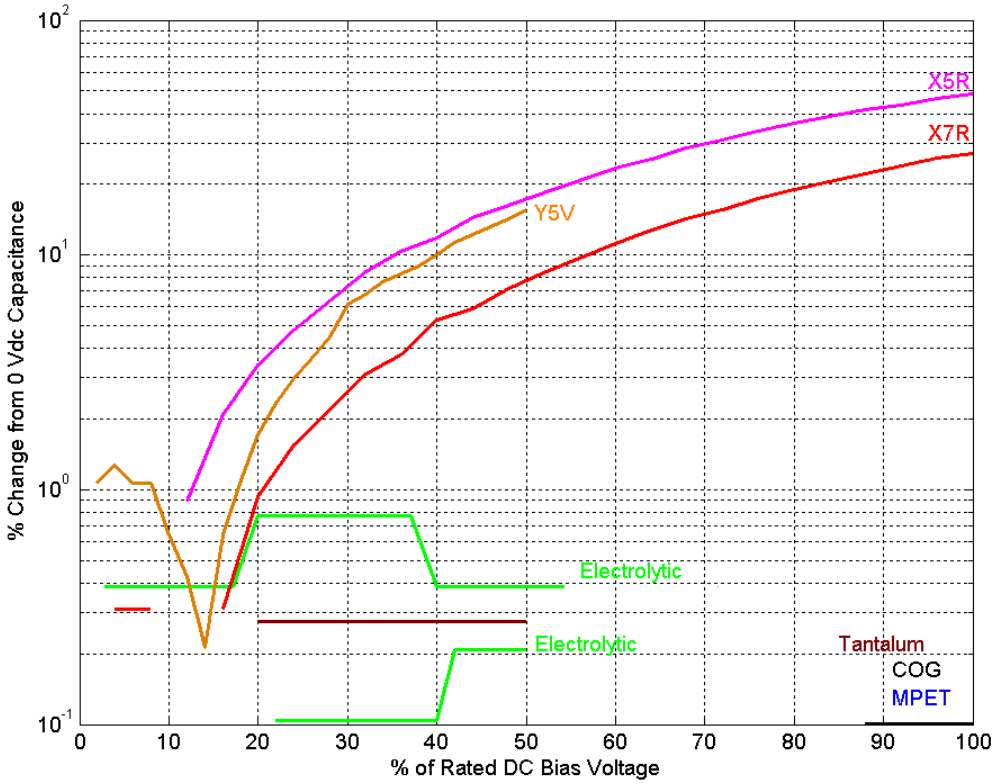

Figure 5.1.9.2. This graph describes the capacitance change from rated value for common dielectrics. The capacitor under test is charged to the DC voltage indicated, then oscillated about that voltage to determine local capacitance (c=dq/dv). These tests were performed on new parts from stock. Not all tests could be completed to full rated voltage. Some parts showed no significant deviation from initial capacitance (such as MPET, COG, and Tantalum). For parts that show significant deviations, memory effects were noticed during test. These tests were not performed in a pulsed manner. Similar plots are provided by capacitor manufacturers[81].

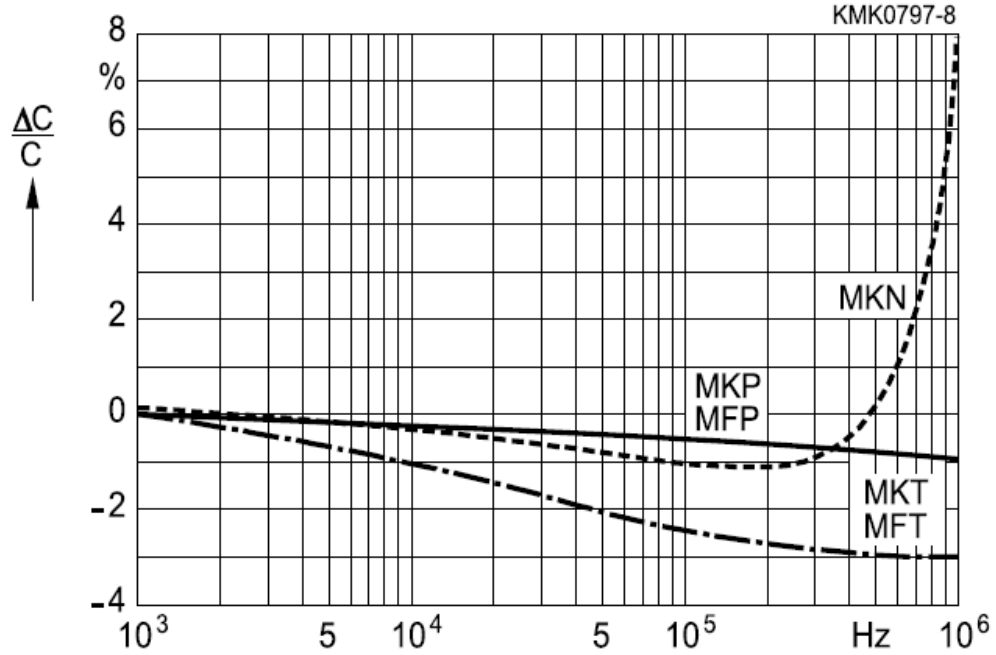

Figure 5.1.9.3. Frequency dependence for common film dielectric types. The third letter specifies the type of dielectric: N = polyethylene Naphthalate, P = Polypropylene, T = polyester. We can expect significant signal in the 100s of kHz. Used with permission of EPCOS AG[82].

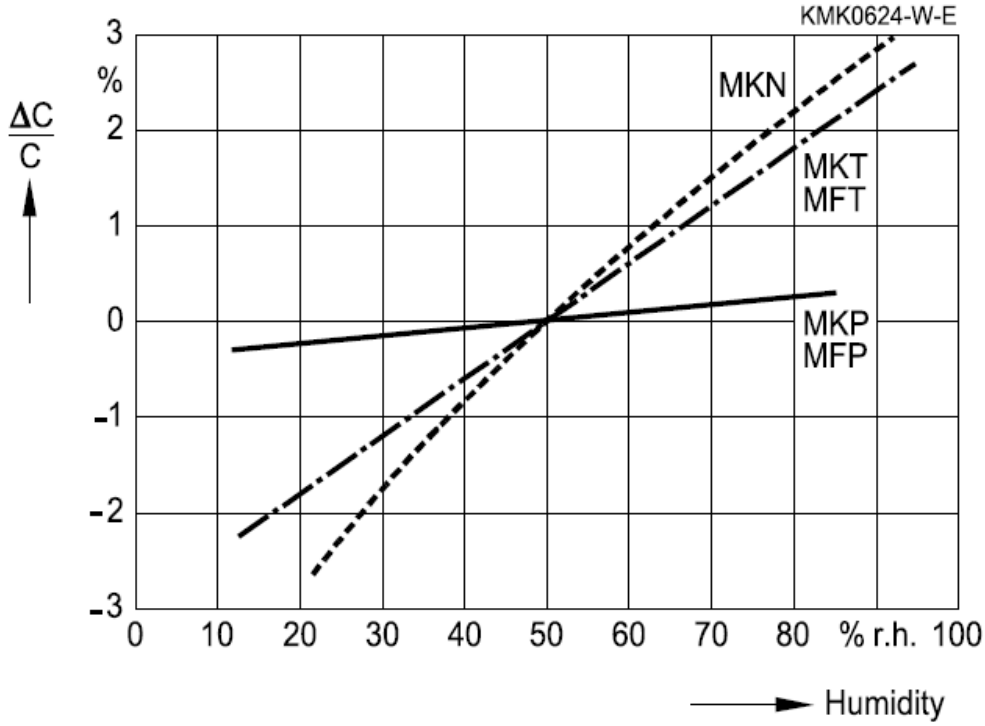

Figure 5.1.9.4. Humidity dependence for common film dielectric types. The third letter specifies the type of dielectric: N = polyethylene Naphthalate, P = Polypropylene ,and T = polyester (used with permission of EPCOS AG[83]). Given our wet and variable climate this can be a significant complication.

It should be clear to the reader why most commercial (mass produced) circuits use C0G ceramics for tasks similar to integration, where consistency and accuracy are important given the performance of the best alternatives. We've shown graphs with some of the more surprising dependencies here. If the issues are not yet clear, then pick a type and size of capacitor and ask the manufacturer for numbers describing temperature, frequency, humidity, electric field, vibration transducer, memory and aging characteristics.

5.1.10 System Model

All the knowledge of the signal path characteristics that has been collected thus far can be reduced to an equivalent SPICE model:

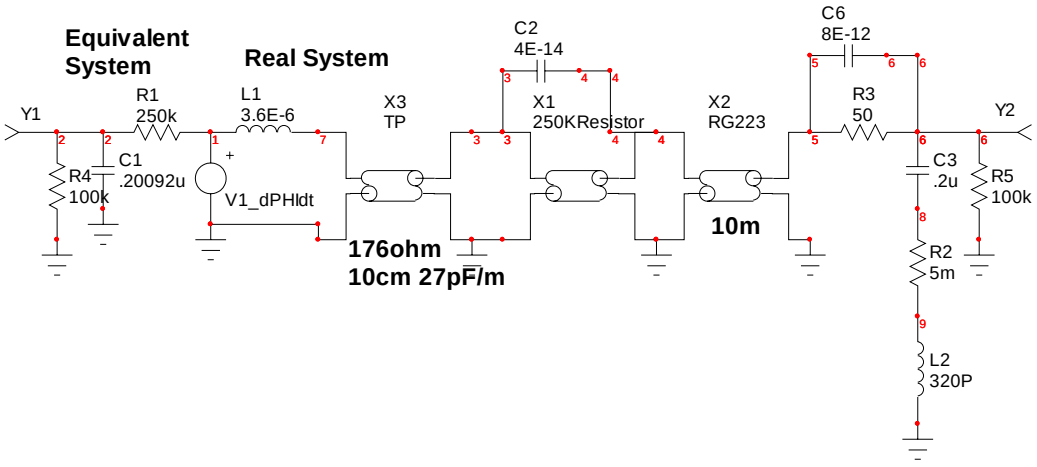

Figure 5.1.10.1. SPICE model for Flux Loop.

C2 is the bypass capacitance associated with the high voltage resistors' geometry. C3, R2, L2 are all intrinsic to the integration capacitor. L2 is extrapolated from manufacturer's specifications for different valued capacitors using the same manufacturing process; it has not been experimentally verified due to lack of proper equipment. Missing from the diagram is a 320 V sidactor at the output of the high voltage resistor used to prevent arcing in case cables get disconnected (it rolls off the 100+ MHz frequencies due to its significant junction capacitance). The simulation results are not expected to be accurate at these higher frequencies due to missing dissipative terms (skin effect in transmission lines[84], dielectric absorption, etc). Likewise, the above model can not be subjected to a transient analysis without adding more dissipative terms due to the implied restriction on time steps from the high frequency response.

Individual elements from the "real system" model can be removed one at a time to determine their effect on the system.

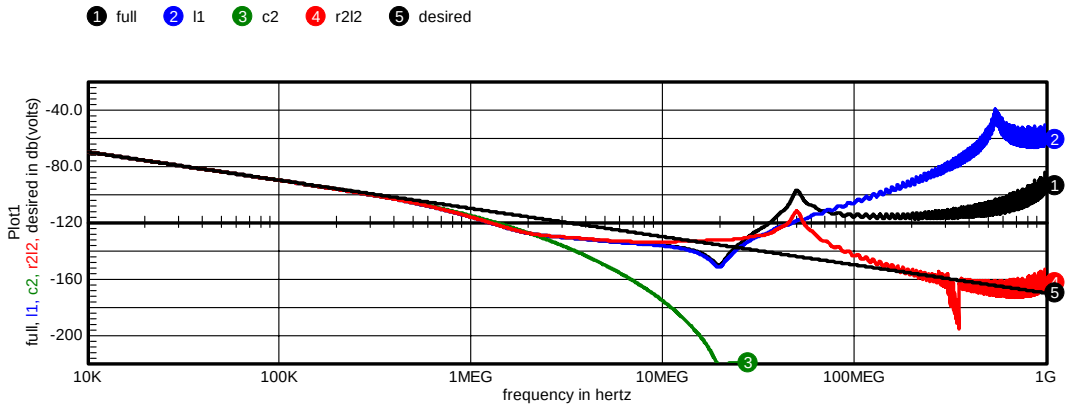

Figure 5.1.10.2. Frequency response curves due to various real elements.

This plot makes it clearer what aspects of the frequency response are caused by which stray elements. The lower response from 200 kHz onward is caused by the distributed 250 kΩ resistor's capacitance to ground. There is also stray capacitance (3 C2) that bypasses this pole, causing the frequency response above 2 MHz to level out. The 20 MHz natural resonance frequency of our integration capacitor causes the small dip in frequency response at 20 MHz (4 R2L2). The resonance at 50 MHz is caused in part by the inductance of the pick up loop (2 L2).

The "real" system model can be compared against the "model"-ed model (aspects corrected for in Matlab code). Here, we will display the frequency response for B and flux loops:

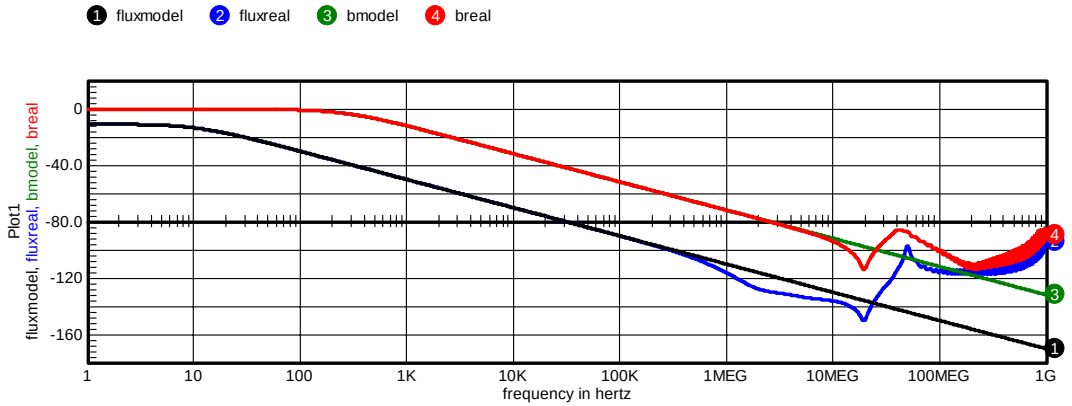

Figure 5.1.10.3. SPICE model versus simply correctable model.

The blue trace is our best estimation of frequency response of the flux diagnostics. The black trace corresponds to the frequency response our standard analysis software corrects for. The red trace corresponds to the response for a B loop diagnostic. The input signal to the B loop is expected to be ~ 46 dB less (/200) than that of the flux loop.

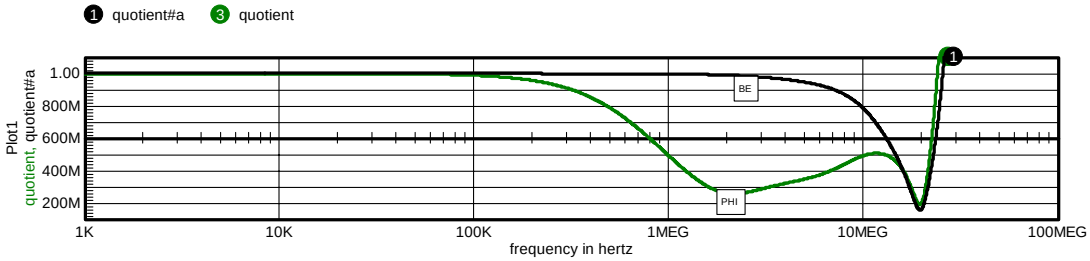

Figure 5.1.10.4. This figure shows the frequency response in the SPICE model normalized for the correctable response in Matlab. The flux loop does not have as good of a frequency response as the B loop.

The effects of the resistor transmission line (250 kΩ) are not as strongly felt by the external B loops, due to the lower 3 kΩ resistor value. This leads to a small excluded flux signal (~2% of total flux) caused by this difference (+ other more minor effects) for vacuum reference shots. It also leads to the potential for serious errors in lift-off flux if the method used is dependent on the signal while the flux is rapidly changing (immediately after reversal bank firing).

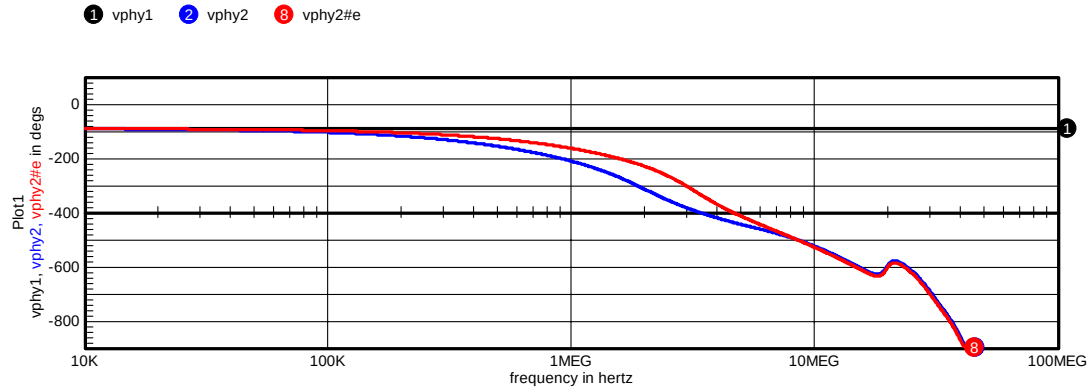

Figure 5.1.10.5. Phase Lag of Flux Loop for varying HV resistor's transmission line characteristics. X1 (2) C' = 65 pF/m expected; (3) C' = 32 pF/m

The ratio of "real" to "model" provides significant left-over error for the flux loops. Unfortunately, this cannot be directly used for corrections due to the poor manufacturing of the high voltage resistor assemblage and the uncertainty associated with the real package's bypass capacitance behavior. The dimensions are small and the tolerances are not good with the potential for a variable air gap (that would dominate capacitance). Therefore the high voltage resistor model is not guaranteed to be fully accurate. There is also the potential for noise in this frequency range (~0.1-3 MHz) as discussed earlier (see: Capacitive Coupling of Driven Coil to B Loop). Vacuum reference shots can be used to determine the relation between B and flux as well as providing a comparison against our model. It should be kept in mind though that the B signal may have significant non-random noise (limiting the utility of the comparison).

5.1.11 Vacuum Reference Shots

The plots show the phase relationship between B_{ez} and Flux. If there were no noise (and no spatial factors, as will be discussed shortly) the plots should line up not quite perfectly with a SPICE system model. So long as there are no nonlinearities in the system, the use of a vacuum shot should be sufficient to correct the system. Vacuum reference shots have low high frequency content and hence poor SNR at high frequencies. Since conditions such as charge voltages and timings change from shot to shot, only relative phase information is easily recoverable and not relative magnitudes. Grouping phase information is not as dependent on the particular type of noise (unlike magnitude). Even then, the information >500 kHz is of very poor quality as indicated in the colorbar floor of 0.8. The surrounding diagnostic pairs produce very similar results. Note that there is a noise source near 2.5 MHz of very narrow bandwidth and uncertain origin. It is present even when no shot is being fired. This could be due to an oscillatory interaction between the current mode amplifier (in the digitizer) and the integration capacitance. Not all the digitizer channels have the effect to the same degree.

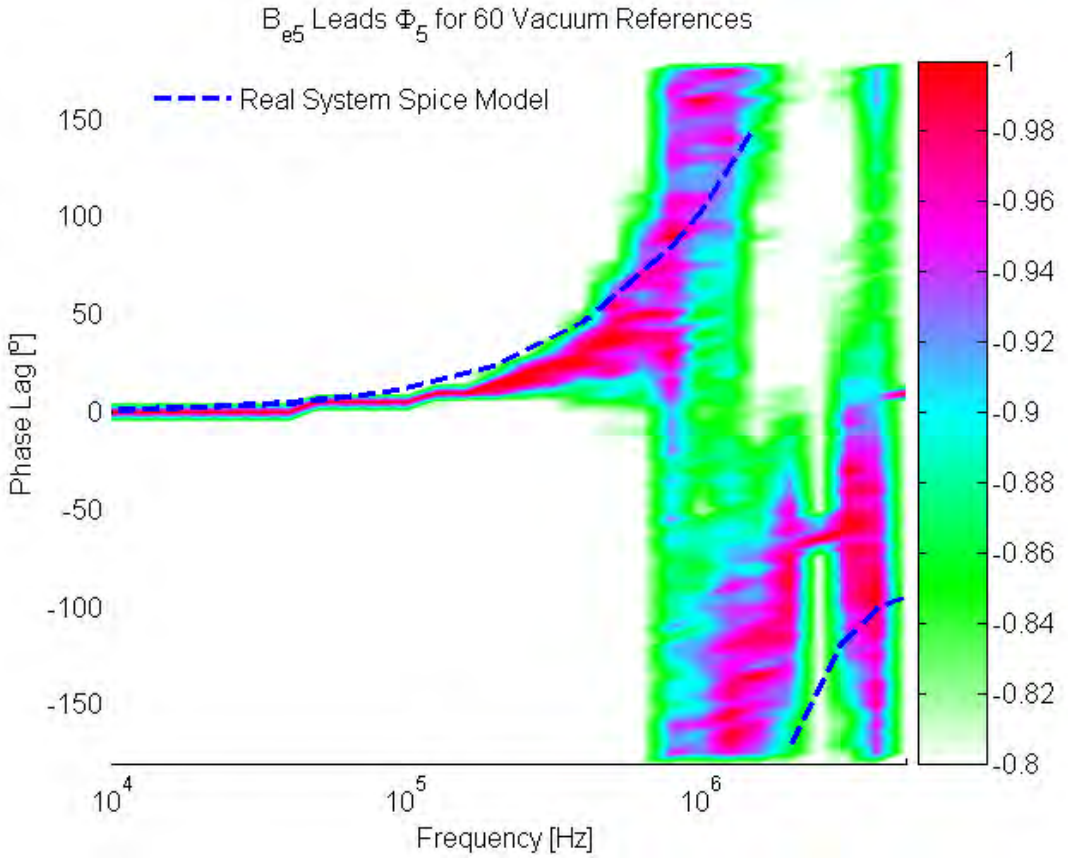

Figure 5.1.11.1. Phase difference between B_{ez} diagnostics and Flux diagnostics. The color surface can be thought of as indicating probability. The blue dashed line indicates the expected relative lag from our full SPICE model.

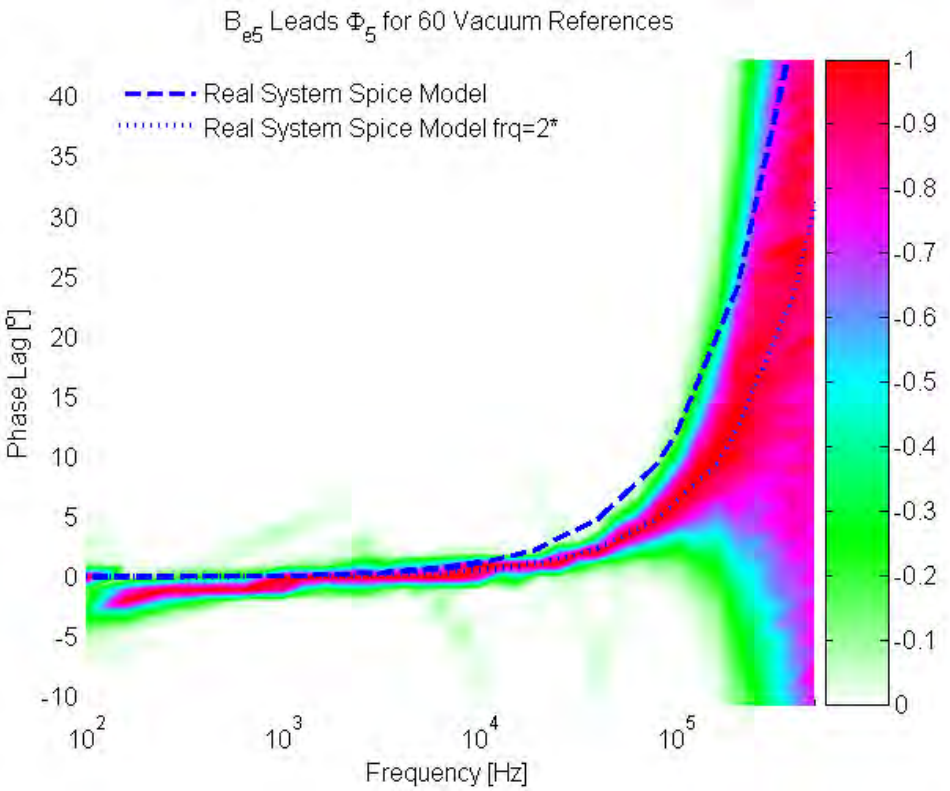

Figure 5.1.11.2. Close-up on phase difference between B_{ez} diagnostics and flux diagnostics. The dotted line corresponds to the real system model with ~ ½ the distributed capacitance. The "real system model" has been altered (for low frequencies) to assume that the integration droop has been compensated for as discussed in "Passive/Numerical Integration". The remaining lag at low frequency is likely due to spatial effects (to be discussed shortly).

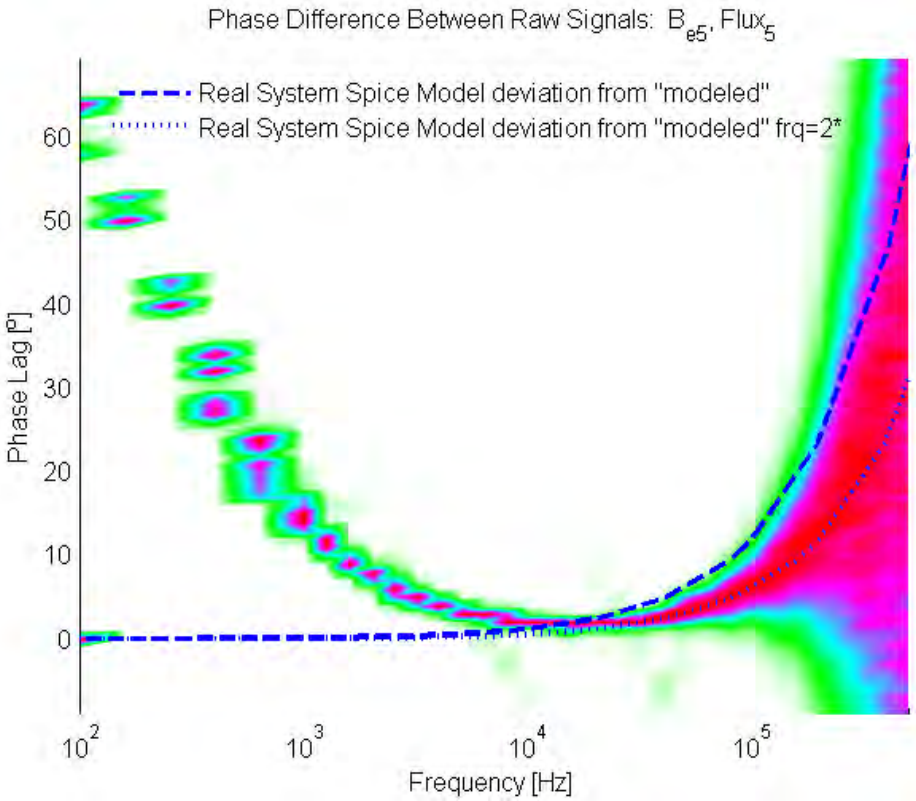

Figure 5.1.11.3. Phase difference between raw signals: B_5, F_5. The large deviation at low frequency is compensated for by the numerical integration corrections based on manually measured parameters.

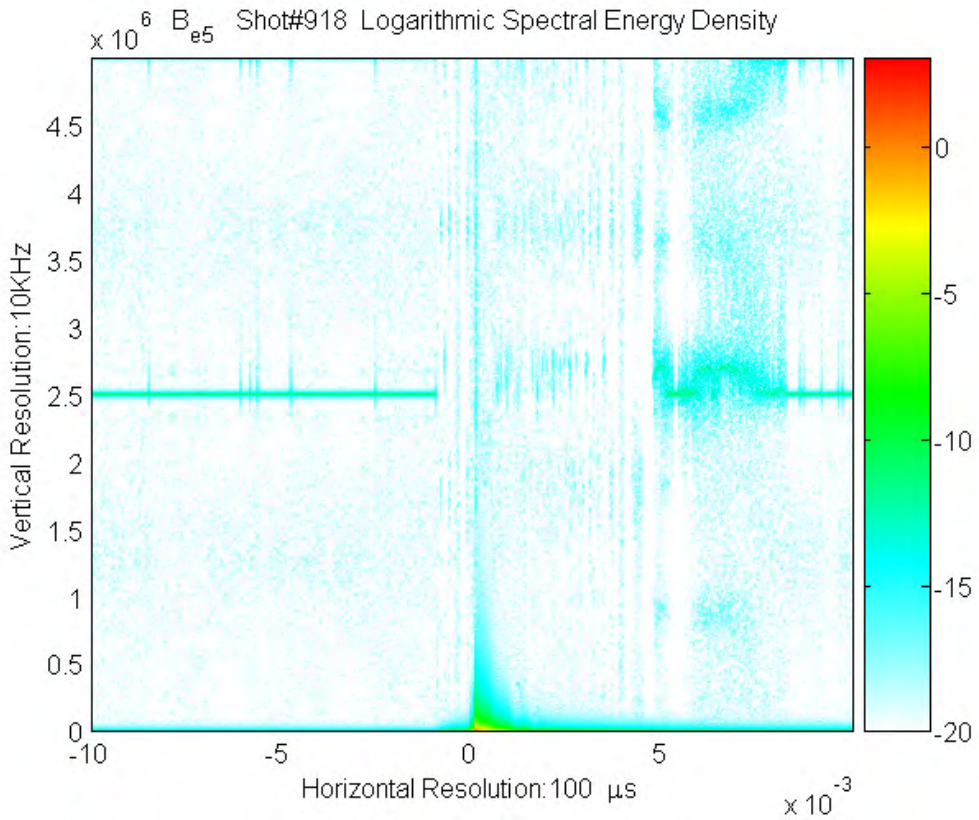

Figure 5.1.11.4. Spectral Energy Density of Mid-plane B Probe. This is taken for a vacuum reference shot. Note the 2.5 MHz line that is present before the shot begins. The previous phase diagram changes negligibly if the data is windowed to remove data before and after the shot.

We can correct the flux loops for the assumed frequency response error using the SPICE model information (interpolated). The basic effect, as can be seen below, is very similar to a simple constant time delay. Rather than running multiple simulations for different levels of assumed distributed capacitance, we use the base model result with a multiplier on the frequency axis (think of this as knowing a basic RC response curve and simply scaling it). Obviously, multiplying the frequency of the adjustment curve is not an exact solution, but for the frequency and adjustment range of interest it results in no more than a 10% error in the correction (which is indistinguishable from experimental data). More capacitance results in more delay (more phase error). Obviously, this effect is greatest when the flux is changing rapidly.

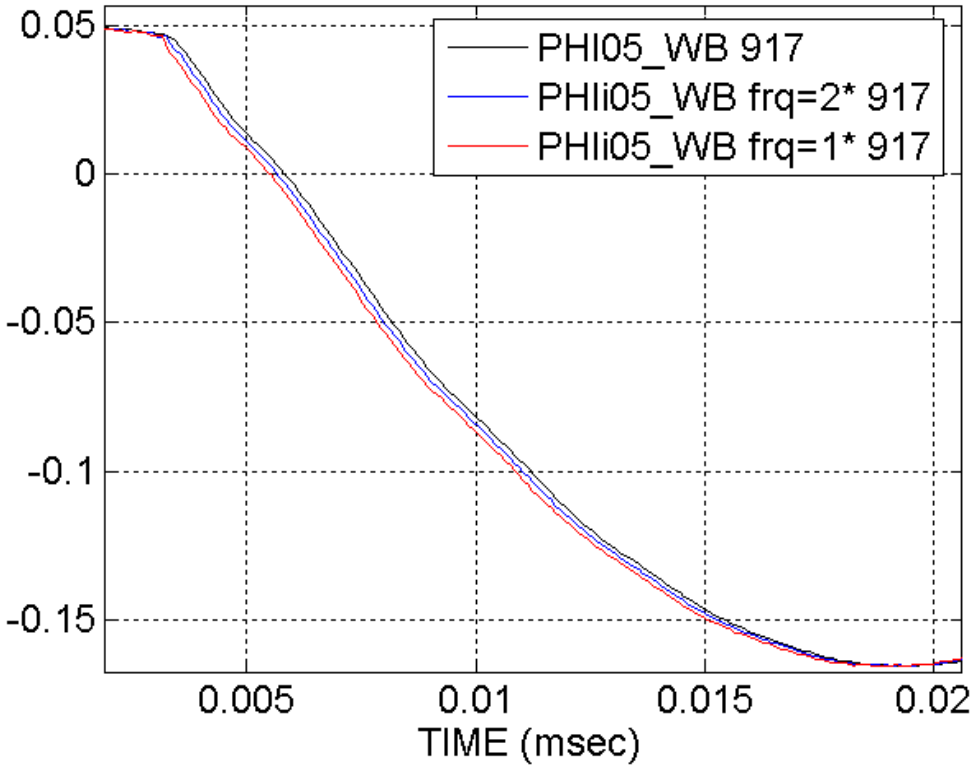

Figure 5.1.11.5. Flux loop with SPICE correction. 'i' indicates a FFT inverse transform (filter) has been performed. We correct the signal in frequency/phase space per the modeled system results. 'frq=2*' assumes ~ 32 pF/m. 'frq=1*' assumes ~ 65 pF/m. In future figures outside this section PHI0n_WB and anything dependent on it includes individualized compensations for this high voltage resistor transmission line effect.

It would be natural to assume that this correction would be very important for most of the results we display (especially excluded flux). It isn't. The diagnostic that it has the greatest initial impact on is excluded flux. Excluded flux is found by subtracting flux from scaled B_{ez}. In many cases this means subtracting two large numbers to find one small number. In calibrating the local B_{ez} loops we use corresponding vacuum shots to solve for the effective B loop area (our "standard" approach). This is done because $B_z(r,z)$ is not uniform in either r or z directions for vacuum shots. The driven straps act as local flux conservers, but their length is not large compared to their diameter. For excluded flux we usually subtract the residuals left over from calculating excluded flux for the vacuum reference shot with the same magnetic bank firing parameters. This has the useful side effect of making first order corrections for the non-uniformity of B_z as well as ameliorating the effects of frequency response due to high voltage resistor capacitance (to first order only). This can be more clearly demonstrated in the plot below.

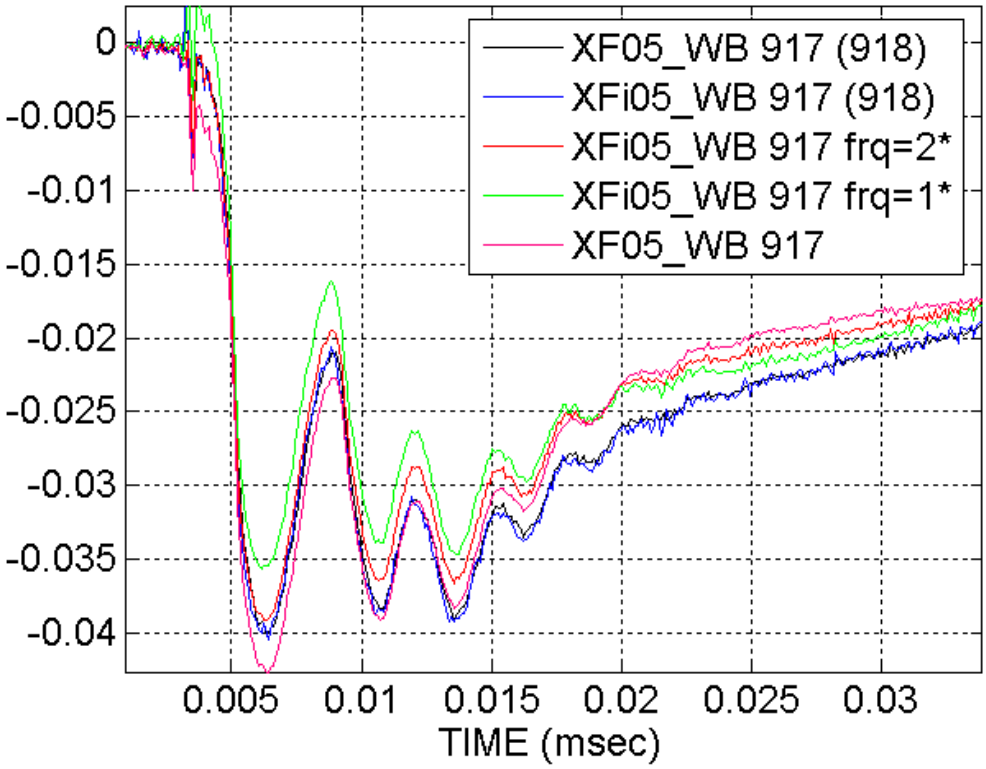

Figure 5.1.11.6. Effects of vacuum subtraction on excluded flux. The black trace is the result of calculating excluded flux without correcting for frequency response, instead subtracting a vacuum reference. The blue trace corrects for frequency response and subtracts a vacuum reference. The remaining shots show the results while not subtracting any vacuum reference. Spatial effects are not yet compensated for in this figure.

Similarly, our measure of separatrix radius is not heavily affected by the frequency corrections because it is based on excluded flux and B_{ze}. The one item that is affected is our (soon to be discussed) correction for spatial effects (i.e., not long flux conserver) due to its direct dependence on the flux loop diagnostic.

The figure below displays the frequency spectrum of excluded flux for both vacuum reference shots (without vacuum subtraction) and plasma shots.

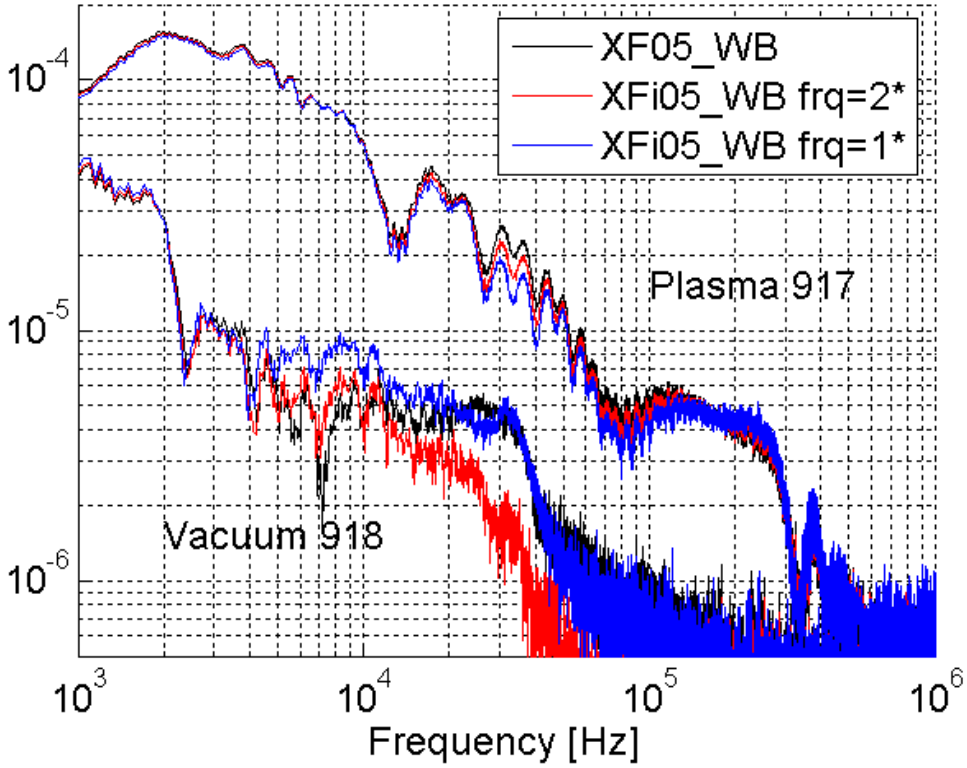

Figure 5.1.11.7. Spectral plot of excluded flux. It can be seen that the plasma signal is a factor of ~10 above the unexplained sources over the entire spectral range of interest.

5.1.12 Spatial Effects

Now that we've dealt with the basic E&M [0D] we can start to consider the effect of having an experimental setup without one long flux conserver. This is not important for making the B and flux diagnostics work; it is important for understanding what they mean. The theoretical inductance of the single driven coil in free space is ~1.324 μH @ VHF. As discussed earlier the theoretical inductance at low frequency is ~1.54 μH. The difference is due to the redistribution of currents in the driven coil from the outer surface of the conductor to the inside, and more importantly from our perspective, the faster decay of flux in the axial ends.

Experimental approach. Since $d\Phi/dt \sim \rho j$ while timescale of decay ~ $\Phi/d\Phi/dt \sim 1/j$, we would expect the high spatial frequency components to decay to the low frequency (12 ms) solution faster than the overall decay by a rate proportional to j_{local}/j_{mean} x (τ_{bulk}). Obviously, there is no one defining timescale in this redistribution since the spatial distribution is not constant (in time). This is immediately apparent with an understanding of how skin depth penetration is usually modeled for use in SPICE. It is common to model the inductive element as the parallel combination of multiple series elements consisting of inductance and resistance [Sen and Wheeler]. For example, we know the resistance and inductance for a driven strap at multiple

frequencies from experiment (as mentioned earlier). To solve for the equivalent system from what we know, we must solve the equations of Figure 5.1.12.1.

$$\sum_{n=1}^{4}\left((0 * 2 * \pi * L_n * \dot{\imath} + R_n)^{-1}\right) == \frac{1}{0 * 2 * \pi * (1.44\,u) * \dot{\imath} + (116\,u)},$$

$$\sum_{n=1}^{4}\left((1\,k * 2 * \pi * L_n * \dot{\imath} + R_n)^{-1}\right) == \frac{1}{1\,k * 2 * \pi * (1.4\,u) * \dot{\imath} + (300\,u)},$$

$$\sum_{n=1}^{4}\left((10\,k * 2 * \pi * L_n * \dot{\imath} + R_n)^{-1}\right) == \frac{1}{10\,k * 2 * \pi * (1.344\,u) * \dot{\imath} + (1000\,u)},$$

$$\sum_{n=1}^{4}\left((100\,k * 2 * \pi * L_n * \dot{\imath} + R_n)^{-1}\right) == \frac{1}{100\,k * 2 * \pi * (1.308\,u) * \dot{\imath} + (6000\,u)}$$

Figure 5.1.12.1. Nonlinear model for skin effect.

These equations can be generalized into the form $\Sigma[n=1:N](I\,\omega[w]\,L[n]+R[n])^{-1} = (I\,\omega[w]\,L[w]+R[w])^{-1}$. Note that these are nonlinear equations, and the precision with which we know the resistive parts is very poor. Nevertheless, the equations are easy to numerically solve for single solutions (just not with high accuracy).

```
xg=[1.4E-6 1E-3 1.4E-5 1E-1 1.4E-4 1E1 1.4E-3 1E3];xg0=xg;
for j=1:1000
   x=abs(fminsearch('fminStrap',xg0.*10000.^(rand(size(xg))-.5) …
      ,optimset('MaxFunEvals',1E5,'MaxIter',1E5,'TolX',1E-9)));
   if fminStrap(x)<fminStrap(xg)
      xg=x;
      (sum(xg([1 3 5 7]).^-1)).^-1
   end
end
for j=1:1000
   x=abs(fminsearch('fminStrap',[xg(1) xg(2) xg(3) xg(4) 1.4E-4*10000^(rand-.5) …
            1E1*10000^(rand-.5) 1.4E-3*10000^(rand-.5) 1E3*10000^(rand- …
            .5)],optimset('MaxFunEvals',1E5,'MaxIter',1E5,'TolX',1E-9)));
   if fminStrap(x)<fminStrap(xg)
      xg=x;
      (sum(xg([1 3 5 7]).^-1)).^-1
   end
end

fminStrap(x)                          5.0759e-005
x([1 3 5 7])./x([2 4 6 8])*1E3        12.1920   0.0830   12.1920   0.0041  [ms]
x([1 3 5 7])*1E6                      1.4349    26.5442  100.6654  48.6219 [uH]
x([2 4 6 8])                          0.0001    0.3198   0.0083    11.7875 [ohm]

function ER=fminStrap(X0)
X0=abs(X0);
f=[0 1E3 10E3 100E3];L=[1.44E-6 1.4E-6 1.344E-6 1.308E-6];R=[116E-6 300E-6 1000E-6 6000E-6];
ER=0;
for j=1:length(f)
   tmp=( f(j) * 2*pi* L(j) *i + R(j))^-1;tmp2=tmp;
   for k=1:(length(X0)/2)
      tmp=tmp-( f(j) * 2*pi* X0(k*2-1) *i + X0(k*2))^-1;
   end
   ER=ER+abs(tmp./tmp2)^2;
end
```

Figure 5.1.12.2. Program for solving nonlinear skin effect equations.

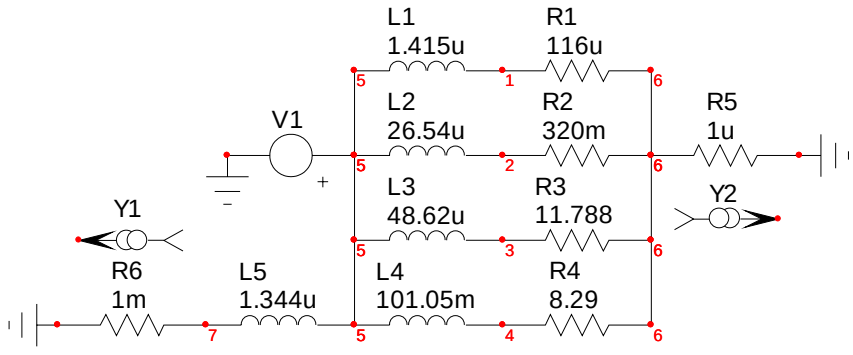

Figure 5.1.12.3. SPICE model of driven strap. R6 and L5 are chosen to match the strap at 10 kHz. L_n and R_n are the results from solving the equations with experimental data (supplemented by an analytically based estimate). R5 is of negligible value, added to make the measurement of the total current simple.

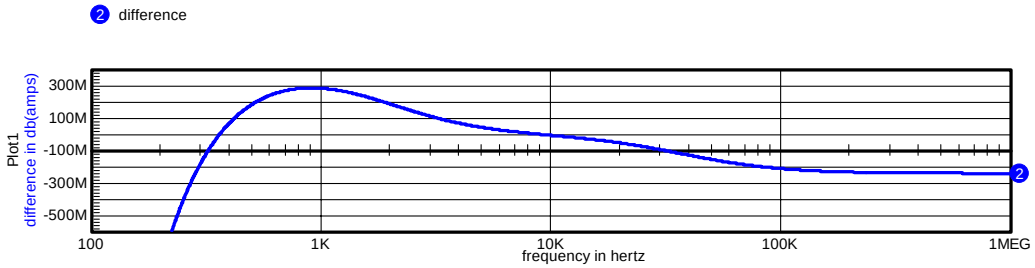

Figure 5.1.12.4. Difference in frequency response between simple strap and measured strap. The system is effectively normalized to the response at 10 kHz. 200 MΩ is approximately 2% difference between the ideal one inductance/one resistance and the modeled system.

Unfortunately, the experimental measure of resistance at each frequency is painfully imprecise (with at most only one significant figure). We present it here as an example of method, to be used in more appropriate cases or cases where other approaches are impractical.

Numerical approach. Consider the initial distribution of currents in a highly conductive coil. All the currents are being carried in the very thin skin depth layer. To visualize this condition, and understand current distributions (as well as solve for inductance) we'd like to be able to numerically model and display this system. The catch is that this current layer is inherently limited by the model's resolution. An additional limitation is imposed by the need to have our computational space include a large amount of uninteresting space (so that the measure of energy associated with the currents is correct and we don't want a close fitting shield (boundary condition) interacting with the coil). This limitation can be compensated for by available reference materials but we'd rather not depend on them.

From a reference[85], we can see that if we want the inductance calculation to be better than 1%, we need the boundary (shield) to be ~ 4 times the radius of the coil. This result is dependent on the ratio of coil radius to length and is thus not a more generally applicable rule of thumb. Hence the minimum boundary radius is 2 meters. We choose our boundaries to be 4 meters both axial and radial with 1 mm grid resolution.

The code ("sor.dll") we use allows us to specify flux through magnetic potential and iteratively step to a solution throughout free space in a cylindrical geometry. The basic approach used is a successive over-relaxation code. The axial z -> -∞ boundary is magnetic field perpendicular to boundary. At the r -> ∞ and z-> ∞ boundaries the magnetic field is restricted to be parallel to boundary. The "SOR" code we are using was originally developed and tested by Richard Milroy. We modified it to patch into Matlab for reasons that will soon become clear. It is also capable of solving for field from currents, but is not used for such purposes here. This code is how we've determined that a driven coil has ~1.324 μH @ VHF.

The spatial distributions of current and magnetic field shown below utilize an uneven mesh to accomplish higher resolution in the strap and a larger domain. The inductance solutions for a constant grid spacing are consistent to better than 1%.

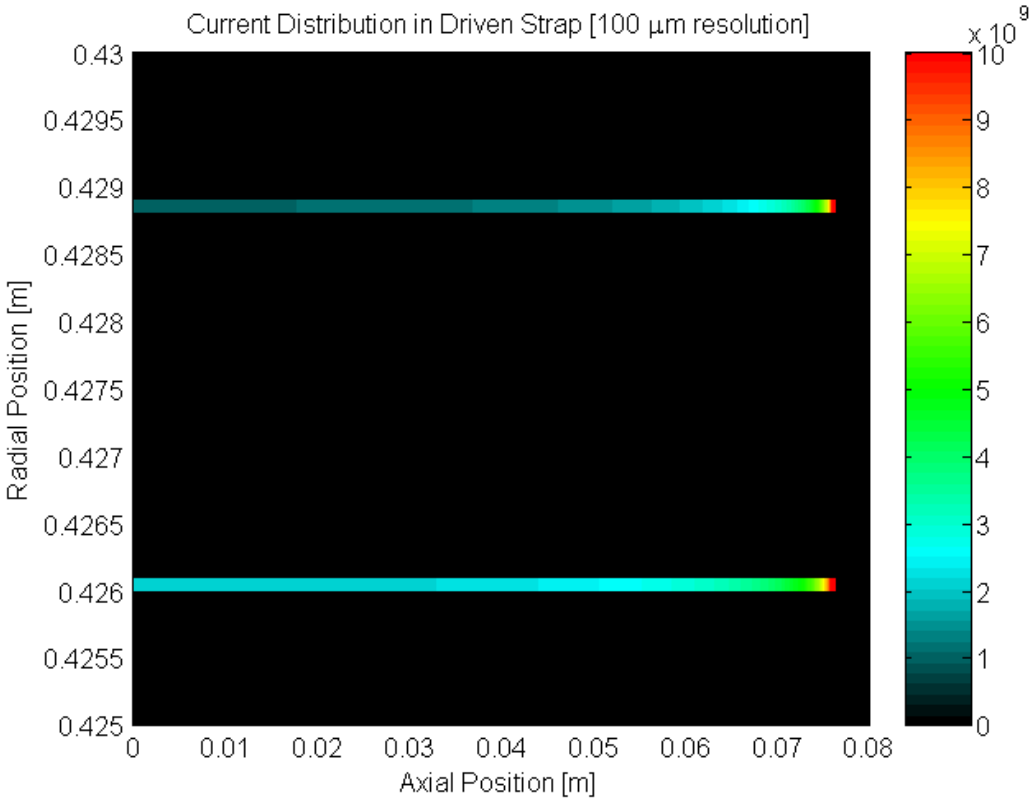

Figure 5.1.12.5. Normalized current distribution in driven strap. The initial current flow (j_θ) along the inside and outside surfaces is displayed in color.

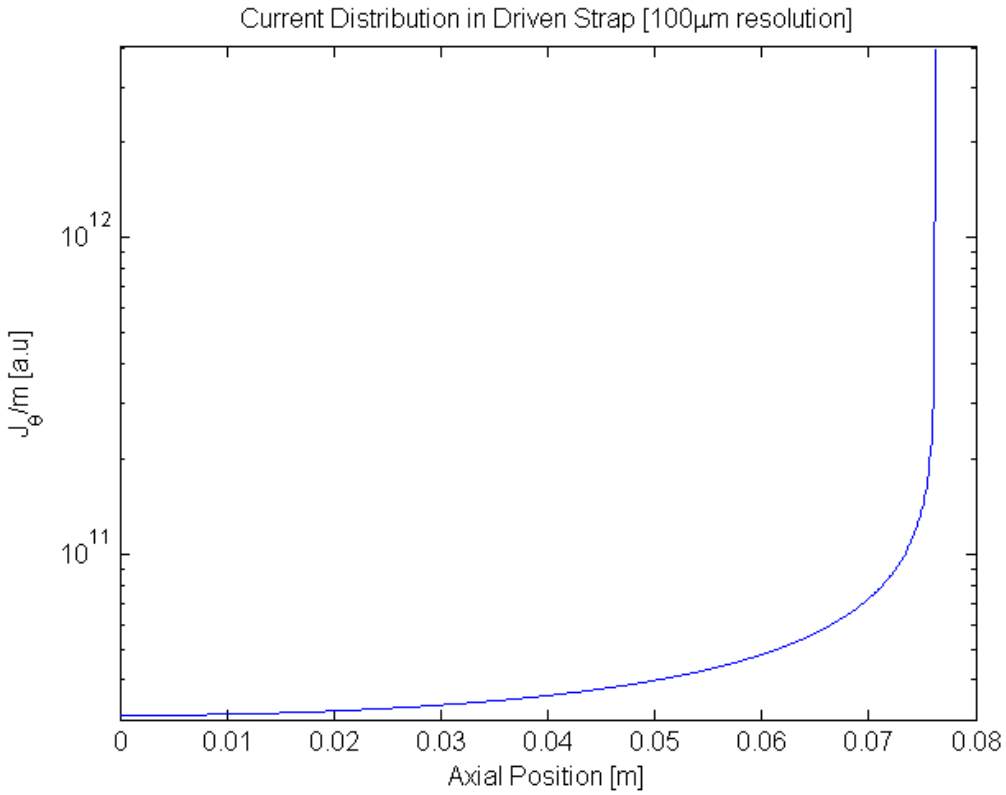

Figure 5.1.12.6. Normalized current distribution in straps as a function of axial location. For each axial location in the numerical simulation the theta currents have been summed across all radii.

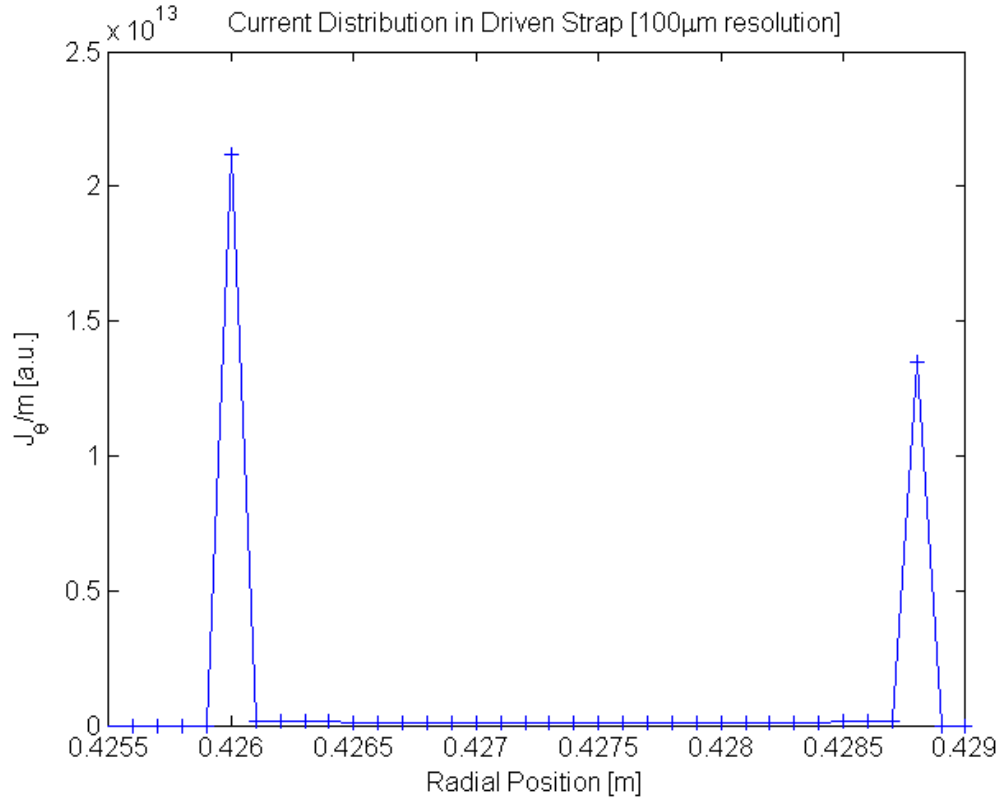

Figure 5.1.12.7. Current distribution in strap as a function of radial location. For each radial location in the numerical work the theta currents have been summed. There is a factor of two, not yet applied, for the symmetric boundary condition at z = 0. The spatial resolution has not been factored into the ordinate, hence arbitrary units.

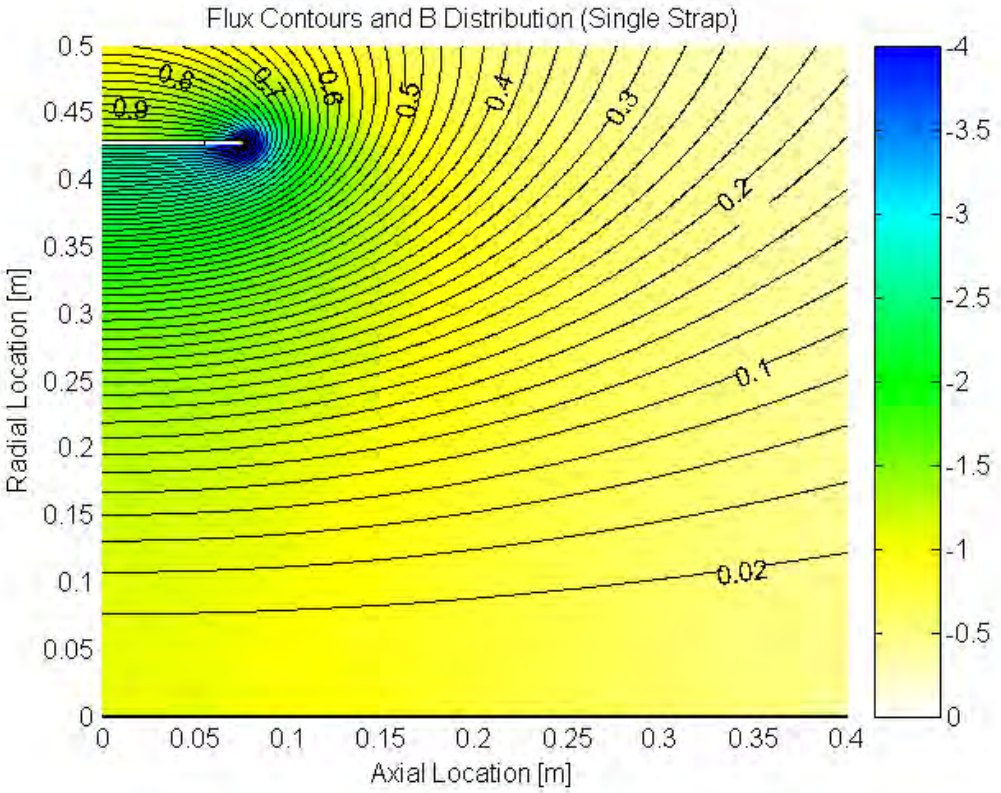

Figure 5.1.12.8. Flux contours and B amplitude @ high frequency. Note the mirror like feature at the end of the strap.

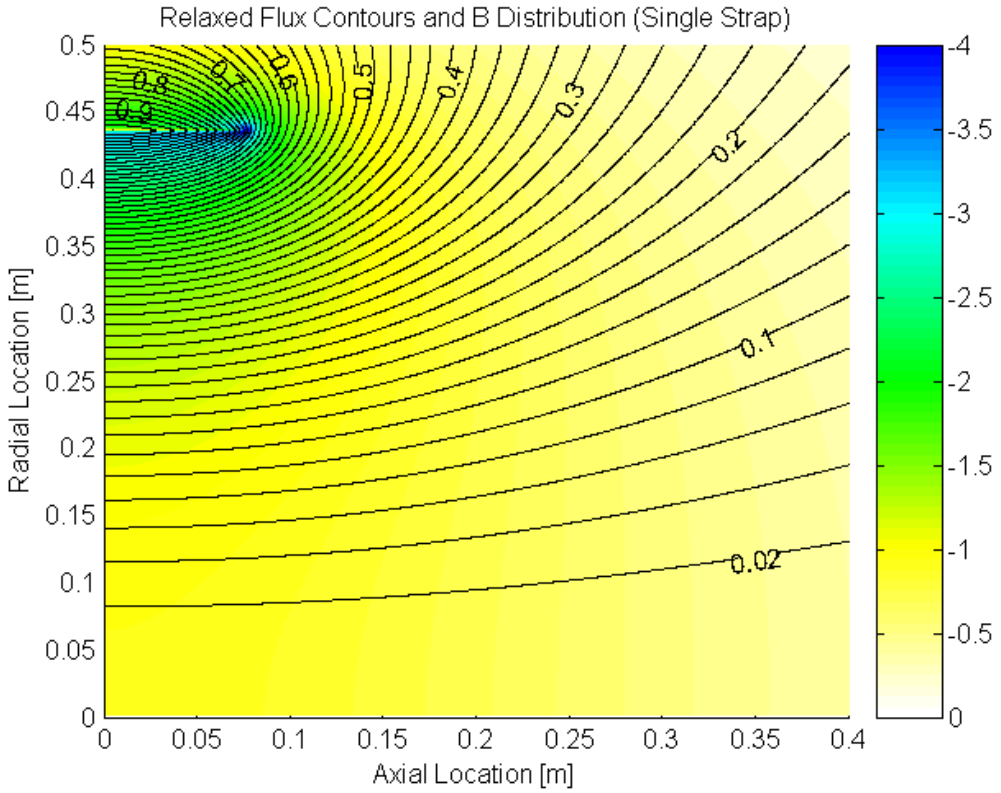

Figure 5.1.12.9. Flux contours and B amplitude @ DC.

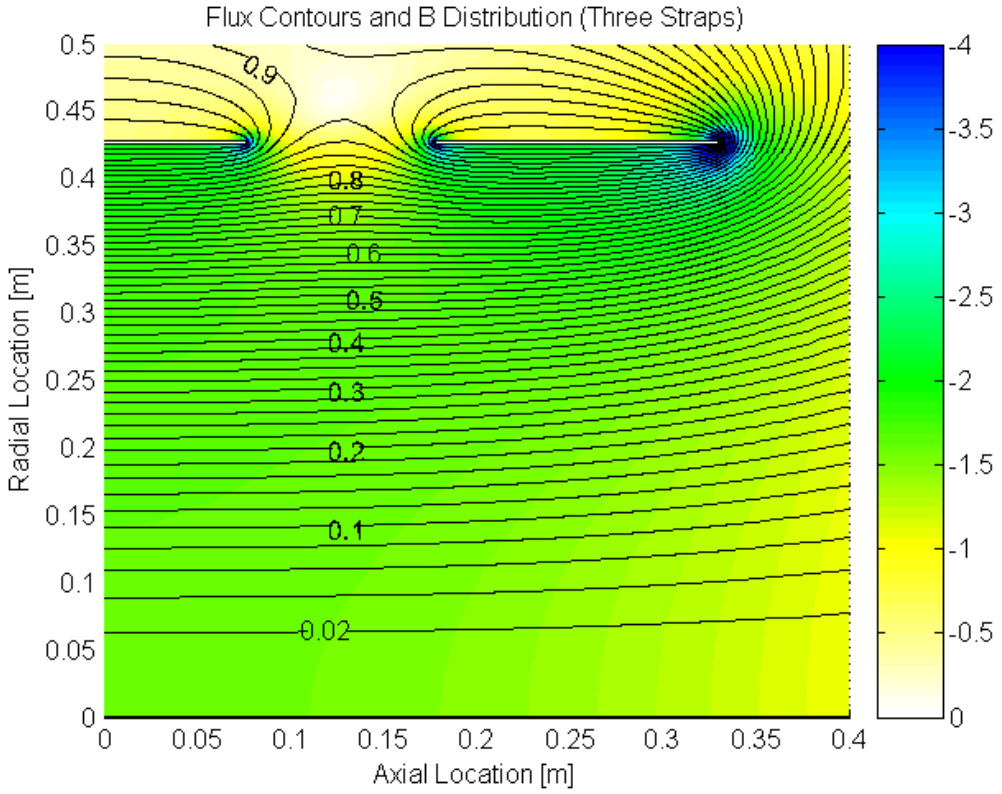

Figure 5.1.12.10. Flux contours and B amplitude @ high frequency with three straps.

When one computation over this domain takes more than 3 days to complete, we cannot afford to time-step the solution forward with this high resolution. Allowing uneven meshing allows better local resolution but not a greater number of overall grid points (it also affects the solution convergence time). Great care must be exercised due to the risk of undetected errors. We choose to drop the resolution away from the coil, then explicitly time-step forward. The SOR code takes flux as it's solved and initial quantity, so we will add a simple flux change (based on resistivity and current) to each step in the Matlab driving code. It should be realized that applying this brute force method to experimental situations would be impractical due to the long timescale bias and cone fields (i.e. our time-steps may have to start with the earliest fields).

The highly localized currents in the numerical model will "decay" towards the even current distribution solution. Our time step size is limited by numerical stability, dependent on the finest grid resolution. We have chosen two levels of resolution that do not capture all of the high frequency details; but does give a sense of the time taken to penetrate locally into the strap, as well as the axial differences (within conductors: Δr, Δz = 500 μm; Δt = 1 μs and Δr, Δz = 100 μm; Δt = 0.025 μs). Total domain size limited to 6 meters.

The timescales (based on total currents) shown in the numerical model are ~12.4 ms, 155 µs, and 19.9 µs. When 4 timescales are allowed they become 12.9 ms, 217 µs, 31.1 µs, and 9.3 µs (in descending order of amplitude).

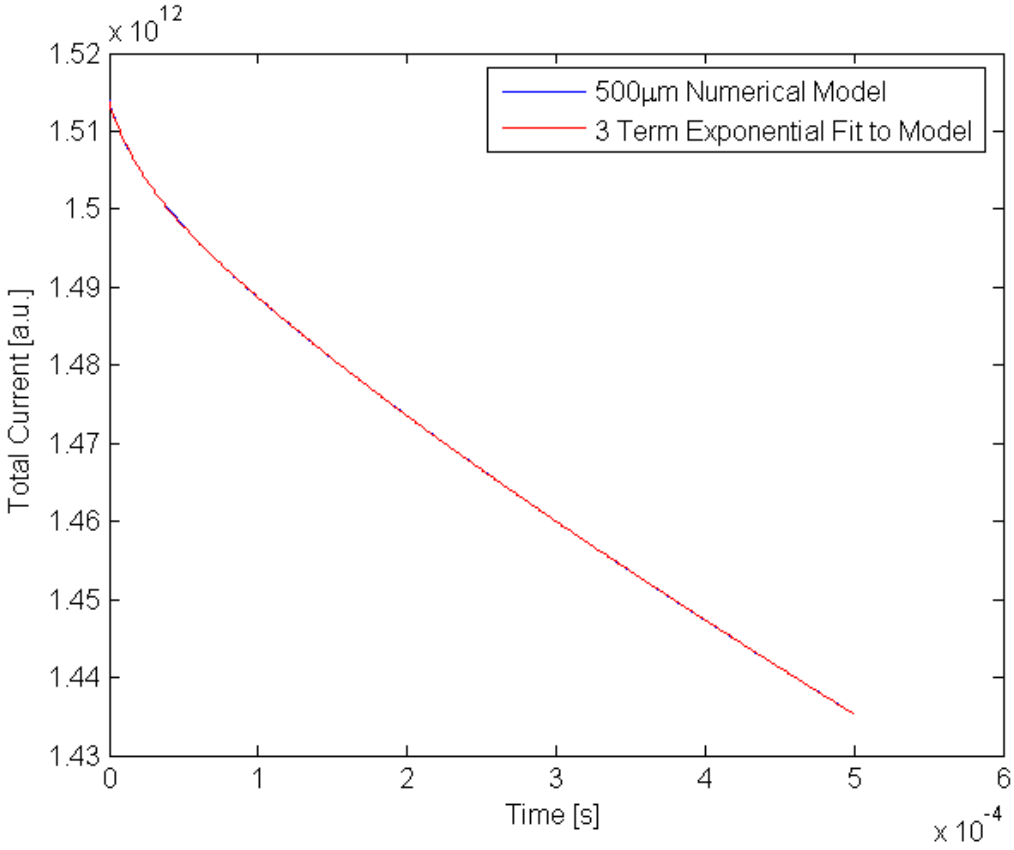

Figure 5.1.12.11. Timescales apparent in numerical model. This figure displays the total current shown in the numerical simulation along with a fit to that curve. The numerical fit is done using the same basic method as discussed for an experimental approach to the problem.

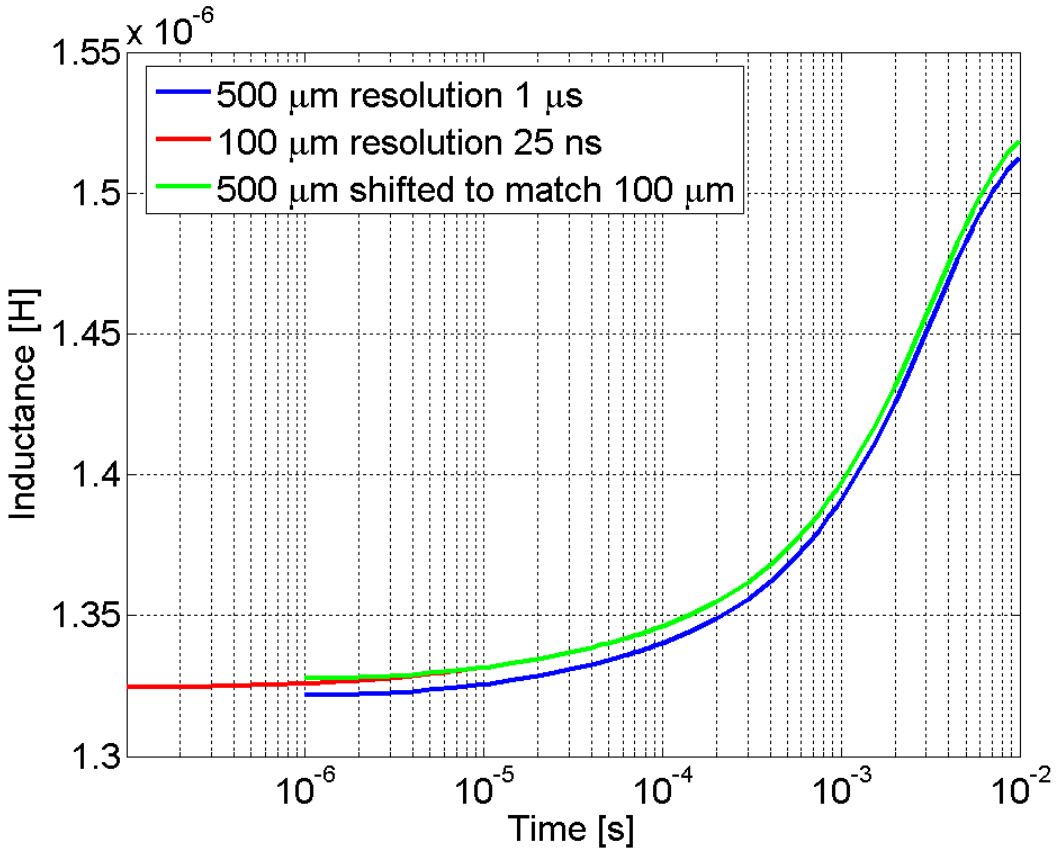

Figure 5.1.12.12. Inductance changes as field penetrates driven strap. Two different grid resolutions were used for this figure due to restraints on time step and computational time.

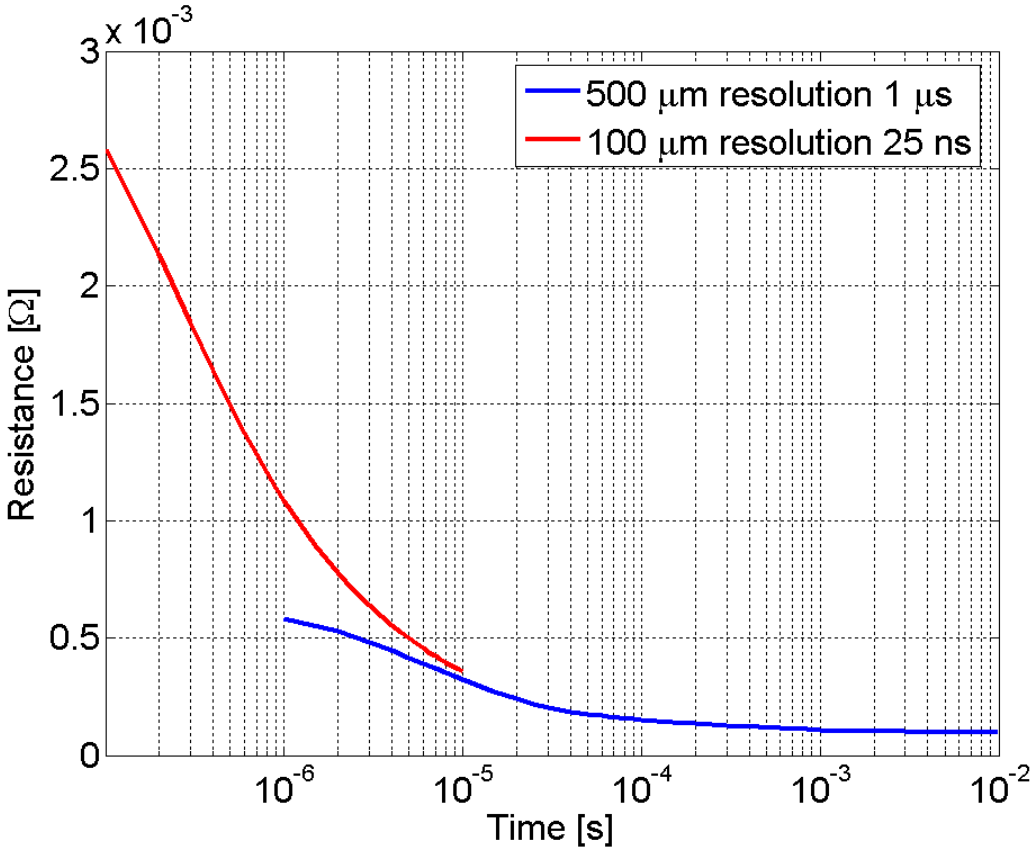

Figure 5.1.12.13. Effective resistance decreases as flux penetrates strap.

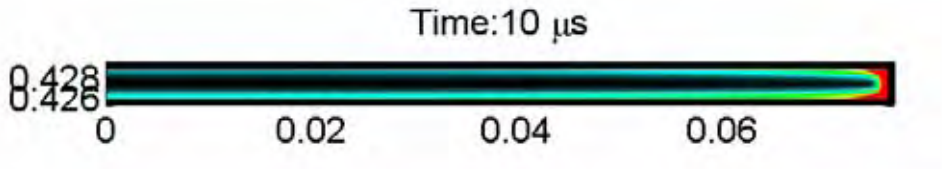

Figure 5.1.12.14. Current distribution in single driven strap after 10 μs.

As has been noted, the computational resources are too limited to provide very accurate corrections for time dependencies using high resolution (on a regular basis). This restricts us to operate in frequency regimes where detailed, time dependent, corrections are essentially unnecessary. From the following figures it can be seen that allowing a 1% error in field restricts us to ~ 20 μs, 2% ~ 65 μs, 5% ~ 270 μs, 10% ~ 700 μs. It should also be understood that these restrictions are relaxed by driving adjacent straps due to the lower concentration of currents (making this more of a worst case analysis). Note that the presence of plasma at large radii restricts the timescale (as discussed at the beginning of this chapter). It is not difficult to compensate for this in the final code, but it is not going to be a very accurate correction for reasonable run times. This correction is something we do not generally

do. To some limited extent, vacuum reference subtraction can correct some signals (such as excluded flux) for these errors. Note that vacuum reference subtraction cannot possibly correct for the additional strap currents caused by plasma, i.e. the large flux excluder limit.

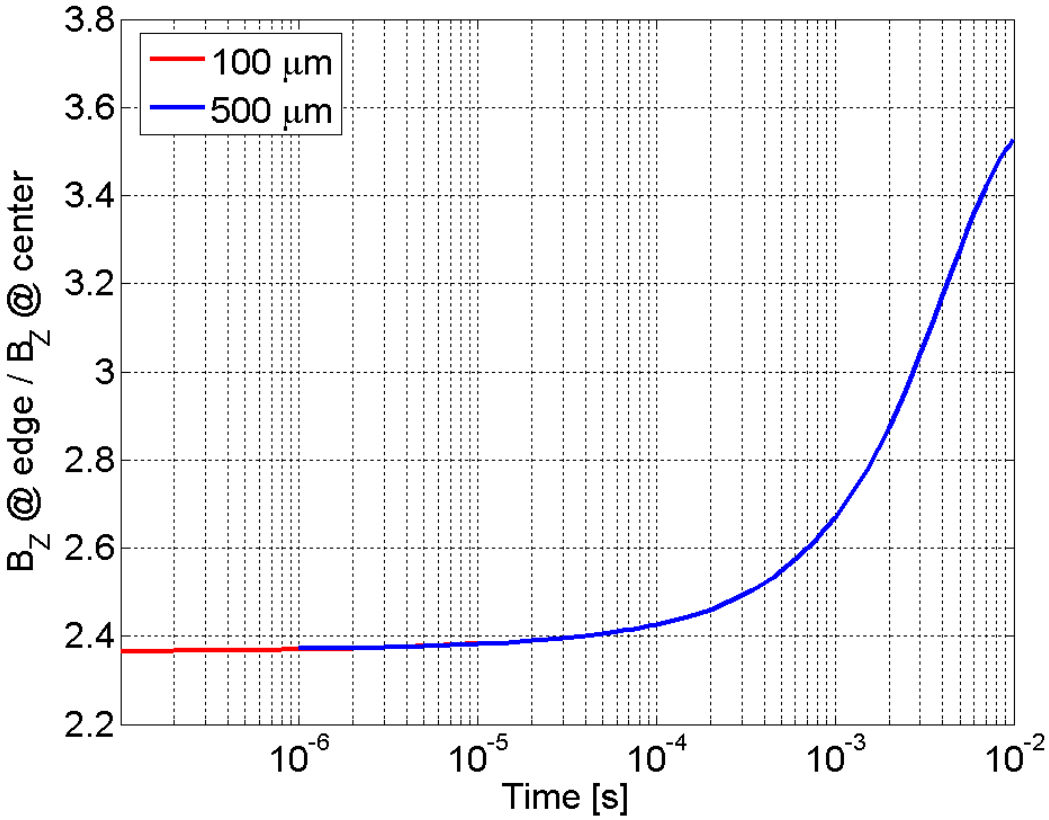

Figure 5.1.12.15. Ratio of B_z measured by diagnostic to B_z at center for a single driven strap. For this plot, measurements are made on the mid-plane defined by the driven strap.

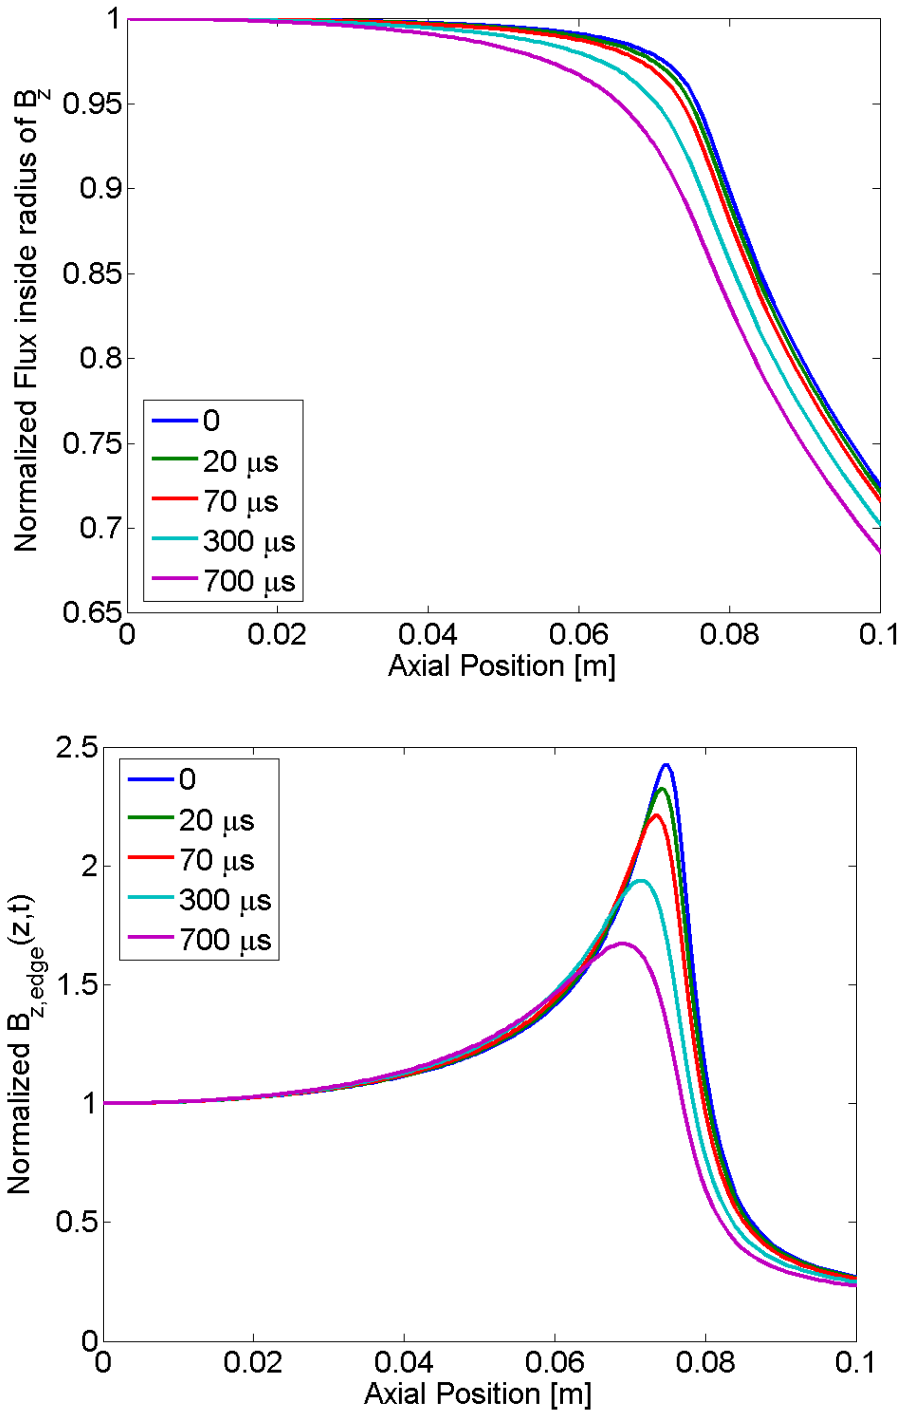

Figure 5.1.12.16. B_z diagnostic is insensitive to exact axial location. Profile for single driven strap. The magnetic field diagnostics are located on the mid-plane defined by the driven strap (i.e. z=0).

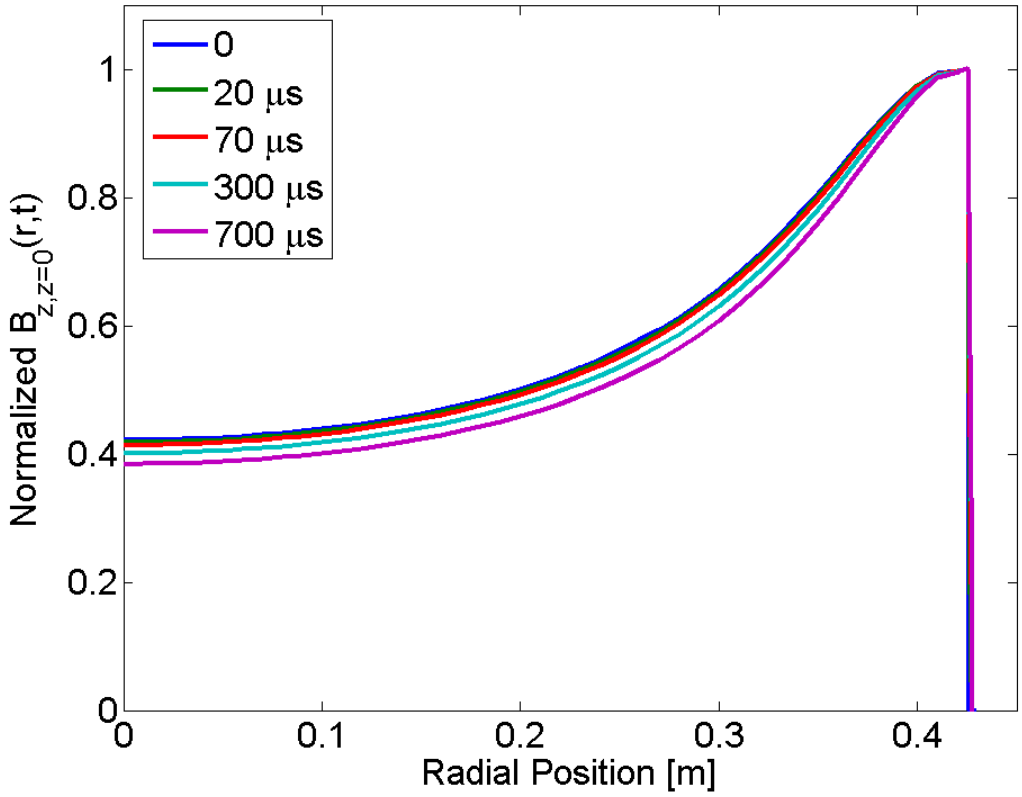

Figure 5.1.12.17. B_z Radial Profile from a single driven strap. B_z diagnostic is insensitive to exact radial location (near strap) until late times.

From the above figures it can be seen that measuring B_z at the strap is not the same as measuring B_z at the separatrix of an FRC. Note that this is displayed for only one coil being driven; when adjacent coils are driven this effect is significantly reduced. Under normal circumstances the amplitude of the B_z signal is normalized to the average B_z derived from the flux loop for a vacuum reference shot. This is done to solve for an "effective" area of the B_z loop. This has been the standard approach that most experiments hardwire into their diagnostics. This "effective area" then is assumed constant for most experiments. We do this assignment during the analysis phase, based on a specified vacuum reference shot. An alternative approach is to solve for an effective area array, allowing the area to change as a function of time. These two approaches generate nearly indistinguishable results for standard "in-situ" shots. Neither of these approaches is perfect, the presence of a flux excluding object alters the ratios for both so they must be seen as (at best) first order corrections.

Some understanding of the magnitude of the error can be obtained from solving for the magnetic profiles associated with the FRC radii generated by the standard "effective" area approach. We can plot what the actual measured B_z is as comparison points.

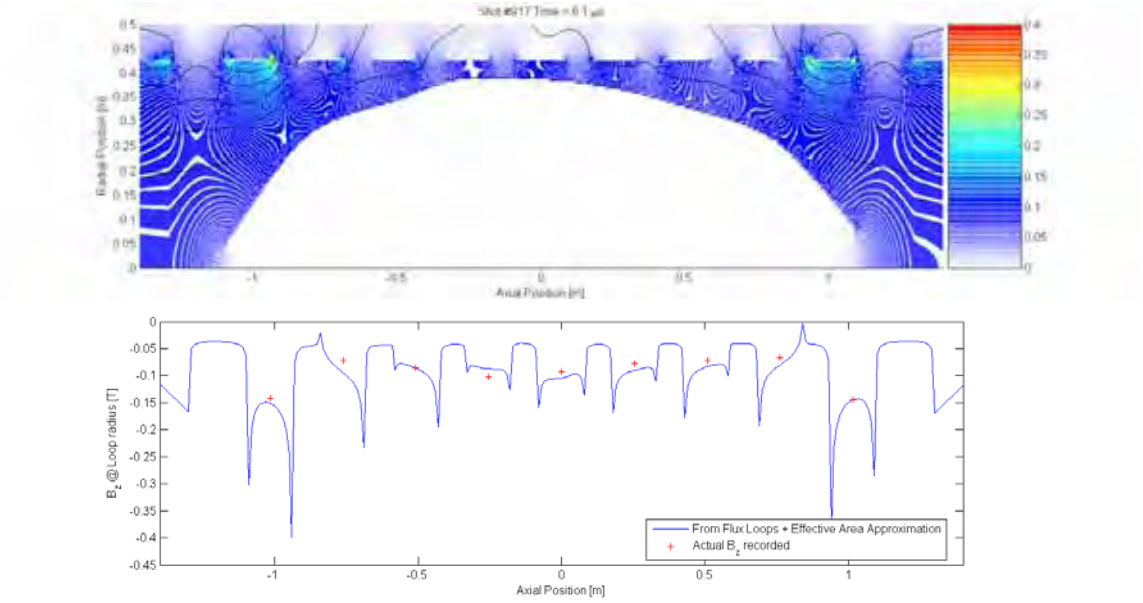

.

Figure 5.1.12.18. Standard approach magnetics for shot #917 @ 6.1 µs. This is a normal "good" in-situ shot. (Top) Black lines indicate flux contours. Color indicates magnitude of B field. This is a plot of magnetic structure derived from the FRC radius results from utilizing vacuum references for single effective B probe area and vacuum shot subtraction. R_s = ((Excluded Flux – Excluded Flux Reference) / ($B_{z,\ effective}\ \pi$))^{½} . (Bottom) Blue line displays B_z at the radii of the B probes. Red crosses indicate actual measured B_z at these locations. The actual measure of B_z is based on original calibration. The difference constitutes an error.

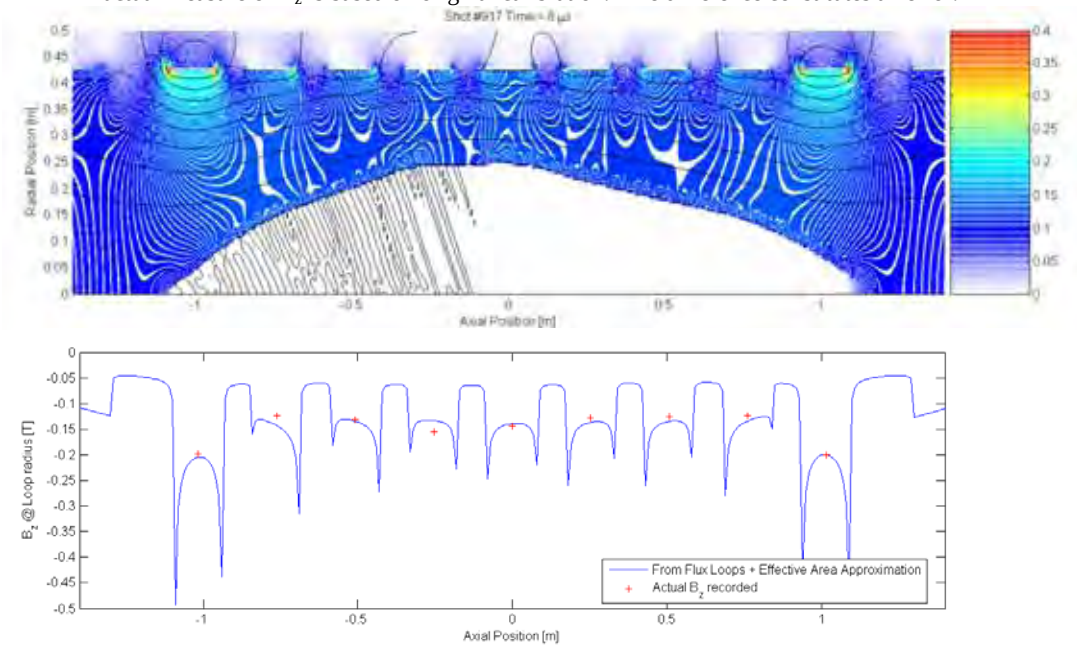

.

Figure 5.1.12.19. Standard approach magnetics for shot #917 @ 8 µs. This is a normal "good" in-situ shot. Note the continuously curved apparent shape of the FRC. (Top) Black lines indicate flux contours. Color indicates magnitude of B field. This is a plot of magnetic structure derived from the FRC radius results from utilizing vacuum references for single effective B probe area and vacuum shot subtraction. (Bottom) Blue line displays B_z

at the radii of the B probes. Red crosses indicate actual measure B_z at these locations. The actual measure of B_z is based on original calibration. The difference constitutes an error.

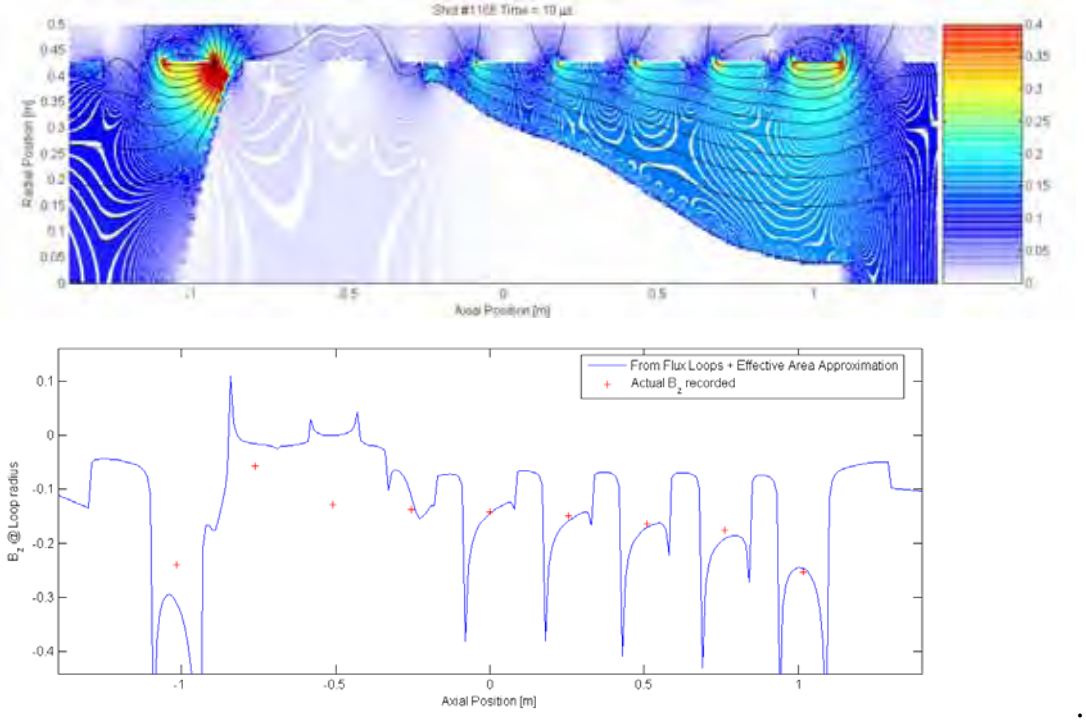

Figure 5.1.12.20. Standard approach magnetics for shot #1166 @ 10 µs. This is chosen to demonstrate the errors when large flux differences exist from one section to another. (Top) Black lines indicate flux contours. Color indicates magnitude of B field. This is a plot of magnetic structure derived from the FRC radius results from utilizing vacuum references for single effective B probe area and vacuum shot subtraction. (Bottom) Blue line displays B_z at the radii of the B probes. Red crosses indicate actual measure B_z at these locations. The actual measure of B_z is based on original calibration. The difference constitutes an error.

We have not yet discussed one critical difficulty in the diagnostic design. The actual area of the B_z loop is difficult to determine (with suitable accuracy) from mechanical constraints alone. This is likely due to the flexible nature of the diagnostic (i.e. magnet wire wrapped on polyethylene) and the small dimensions involved (0.06" polyethylene, 0.02" diameter magnet wire). Presumably, the magnet wire displaces some of the polyethylene insulation in the final configuration when it is wrapped on the vacuum vessel. Given the magnetic numerical modeling and location of the flux loop, we can expect it to record 98.4% of the peak flux in the driven strap. Similarly, for 1 Wb of flux in a single driven strap we would expect to record 2.559 T from the location (and nominal area) of the B_z loop. This yields an expected ratio of 2.6006 T/Wb between flux diagnostic and B diagnostic. Cases where one driven strap fired at a time (albeit at low +/-1 kV voltages) have been used to determine the experimental ratio T/Wb.

Table 5.1.12.1. Characterizing constants for individual magnetic diagnostics. Bz/Flux is based on nominal loop areas. Rs = series resistance, Rl = digitizer resistance. The similarity in the radius of the Quartz tube (1/2 OD) allows us to assume that inner radius of driven coils = 0.430 m, radius of flux loops = 0.4275 m, radius of B_z loops = 0.4252 m.

Loop Pair #	Axial Location [m]	~Radius of Tube [m]	Bz/Flux [T/Wb]	Bz Rs [Ω]	Bz Rl [Ω]	Bz C [nF]	Φ Rs [Ω]	Φ Rl [Ω]	Φ C [nF]	HV Φ Resistor ω Multiplier
1	-1.016	0.4246	2.4535	3057	90615	197.9	249400	91030	203.9	1.3
2	-0.762	0.4248	2.4482	3030	90850	198.4	252400	90955	202.4	2
3	-0.508	0.4251	2.3657	3053	90750	197.8	255400	90770	203.4	1.5
4	-0.254	0.4253	2.2108	3076	90860	199.6	255500	90940	203.7	1.4
5	0	0.4255	2.7077	3018	91210	198.4	252200	90795	202.6	1.8
6	0.254	0.4257	2.2265	3025	90590	198.2	250500	90695	201.1	1.7
7	0.508	0.426	2.2336	3026	90545	199.4	253900	90605	202.2	1.6
8	0.762	0.4262	2.3289	3052	90865	198	256800	90915	201.7	2
9	1.016	0.4264	2.2633	3054	90635	198.4	254400	90680	200.2	2

Presumably, the variance from 2.6006 indicates a difference in the size of the B_z loop, given that it's characteristic dimensions are the smallest (and hence most susceptible to error). Additional complications arise from loop pair #s 1 and 9's close proximity to the end cones, which are lined in stainless steel. The presence of this closed circuit changes the expected ratio by less than 1% (according to lower resolution simulations). The effect is so limited in large part due to the asymmetry as well as the extra distance between the last coils and the cones. Pair 5 has an abnormally high ratio. Pair 5 is wrapped over a joining section between quartz tube sections. With the different construction it would not be surprising if the driven strap is not applying as much pressure to the diagnostics; hence compressing them less. As can be seen from the traces below, there is nothing else abnormal about the probes' relative response.

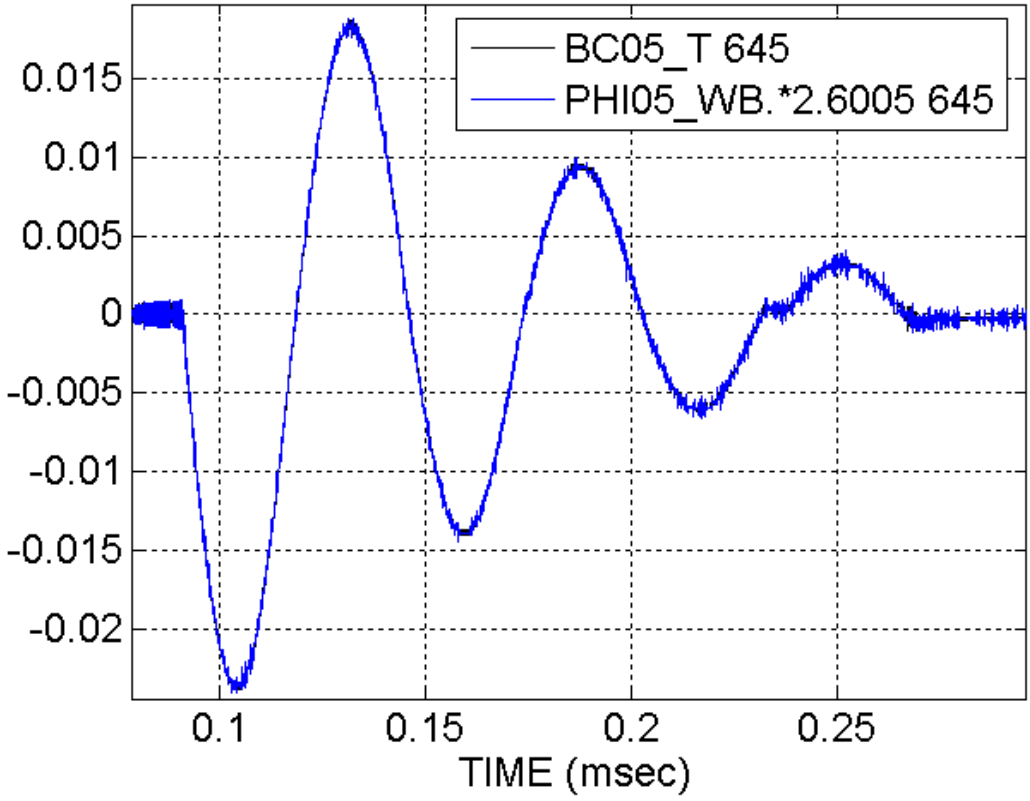

Figure 5.1.12.21. Overlay of corrected B_{z05} trace with flux. The B loop's area has been corrected for.

5.1.13 Spatial Effects (Solved)

The best (theoretically most complete) approach has been to solve for the magnetic field profile given the flux under all the loops and an assumed flux = 0 surface at the separatrix of the FRC. Feeding back the mismatch between what the B_z should be from the magnetic model versus what it is from the diagnostic, allows an iterative approach to solving for the FRC flux exclusion radius.

Certain routines, most notably fminsearch, are quite capable of taking such a complex system and attempting to find minimum error solutions. Unfortunately, this tends to be a computationally expensive approach with no assurance of the routine finding a stable (or best) solution. Indeed, under certain circumstances it fails to find better solutions while large residuals remain (without the addition of the computationally expensive random sampling routines demonstrated earlier in solving for parallel inductances of the driven strap). There also exists the possibility of multiple solutions. Quite simply, on the extrema, there is a grouped solution where the FRC appears to be a continuous mass, but there may also be fragmented solutions consisting of alternating large FRC radius and small FRC radius. The effect can be likened to the lessening of accuracy associated with doing an Abel inversion. It should be intuitively obvious that the more closely coupled the diagnostic pairs (and

the more solved radii), the more likely this technique is to have "stability" problems from noise sources. Adding more diagnostic pairs, but not increasing the number of solved radii, would help to overcome most deficiencies given a sufficiently speedy algorithm.

Up to this point, no real plasma physics have been introduced. A desire to keep the code generalized and independent of any potentially error prone physics assumptions (such as MHD and no flows) has led us to choose a separatrix boundary shape that is simply a cubic fit of the available data points. An alternative to this boundary is a flat boundary under the straps with an interpolated section below the gaps. We like to call this the "smoke-ring" approach and it is also straight forward to implement (and probably more relevant for initial formation). Note that either choice (but more so the cubic fit) may have strange implications. They allow for some locations to have negative radii and still meaningfully impact the solution between the probe locations. They will also result in uneven pressure distributions axially across the surface since the shape of the fit is arbitrary and not physics based. The currently heavy computational costs to run these corrections (with 5 mm resolution, approximately 5 minutes per frame) have dissuaded us from attempting more physics at this stage. These routines tend to be expensive, in part, because the SOR code takes nearly as long when it is given close initial values as it does when starting from scratch. Some additional physics restrictions may be added later, after the code has been optimized to run faster. Current based semi-analytic field solvers (as used to determine DC inductance) require approximately the same amount of time to run to the desired accuracy at the location of the B_z probes.

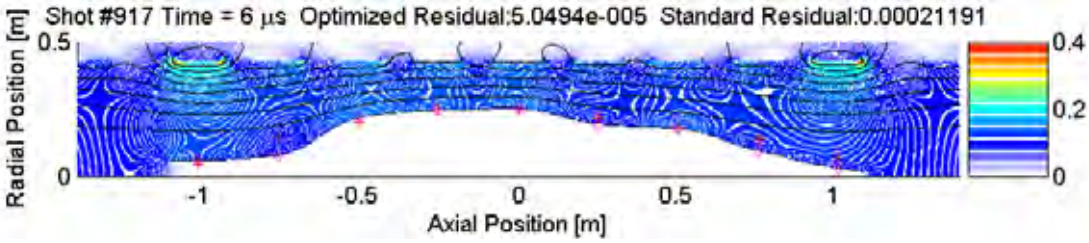

Figure 5.1.13.1. Magnetically fitted separatrix radius for shot #917 @ 6 μs. Red crosses indicate radii derived from the "standard" vacuum reference based approach. Magenta diamonds indicate magnetic simulation fitted points based on the original B probe calibration. Magnitude of B [T] is represented by the color scale with white stripes added to aid in level determination. The color scale does not accurately display the location of these stripes (aliasing problem). Black contours indicate (10 mWb Δ) flux surfaces.

One is tempted, in order to maintain simplicity and make the computational costs bearable, to assume that the area of the B probes remains a constant after installation and testing. For various reasons, this is in fact not the case; nor do the exact frequency response characteristics of the HV resistors remain constant. Some

of this may be due to polyethylene (plastic) seepage under pressure from the driven straps on very long timescales. On June 26, 2006 the magnetic diagnostics were calibrated and made operational. By August 25, 2006 the B probe areas had changed by up to +/- 10% (ironically, with pair #5 shifting negligibly). B probes 2, 3, and 4 experienced the worst shifts. It was made policy that the experimental enclosure remains locked at all times due to repeated cases of damage sustained to the experiment by individuals walking on, and through, the experimental apparatus. After this date the areas appear more or less consistent (~1%) with those tested later on October 4, 2006. The HV resistor frequency compensation was added much later (and based on an average of all available vacuum references due to the limited high frequency information). Probe resistors for loops 3 and 4 are the furthest from this average frequency response for vacuum reference 918. This difference has some (but not all) of the hallmarks of being due to capacitive coupling of a poorly synchronized +/- bank to the B probe. It is possible to fully automate the magnetically fitted vacuum reference process, but the cost of doing magnetic fitting on the plasma shots is generally too high. A lot of cases have to be run before this author is comfortable with utilizing this approach in fully automated routines.

Table 5.1.13.1. Characterizing constants for individual magnetic diagnostics (shot #918 and later). B_z/Flux is based on nominal loop areas and vacuum reference shot. Rs = series resistance, Rl = digitizer resistance. The similarity in the radius of the Quartz tube (1/2 OD) allows us to assume: Inner Radius of Driven coils = 0.430 m, Radius of Flux loops = .4275 m, Radius of B_z loops = .4252 m. B_z/Flux and HV Multiplier have been updated where significant deviations from the previous table have occurred.

Loop Pair #	Axial Location [m]	~Radius of Tube [m]	B_z/Flux [T/Wb]	B_z R_s [Ω]	B_z R_l [Ω]	B_z C [nF]	Φ R_s [Ω]	Φ R_l [Ω]	Φ C [nF]	HV Φ Resistor ω Multiplier
1	-1.016	0.4246	2.4535	3057	90615	197.9	249400	91030	203.9	1.3
2	-0.762	0.4248	2.2523	3030	90850	198.4	252400	90955	202.4	2
3	-0.508	0.4251	2.1764	3053	90750	197.8	255400	90770	203.4	1
4	-0.254	0.4253	2.4451	3076	90860	199.6	255500	90940	203.7	1
5	0	0.4255	2.7537	3018	91210	198.4	252200	90795	202.6	1.8
6	0.254	0.4257	2.1530	3025	90590	198.2	250500	90695	201.1	1.7
7	0.508	0.426	2.1353	3026	90545	199.4	253900	90605	202.2	1.6
8	0.762	0.4262	2.1962	3052	90865	198	256800	90915	201.7	2
9	1.016	0.4264	2.3765	3054	90635	198.4	254400	90680	200.2	2.5

There is an added complication due to the lack of a flux probe for the end cone coils. This has been compensated for by assuming a reasonable constant level of flux in the end cones based off several shots where the end cones were fired alone; the reversal coil flux diagnostics were able to pick up weak signals. Magnetic profile simulation allowed us to estimate the approximate flux in the cones. However, not much effort has been expended in ensuring the accuracy of the simplifying assumptions.

5.1.14 Standard vs. Magnetically Fitted

It would be convenient if we could substantiate that using the "standard" approach of vacuum references for B probe area and vacuum reference subtraction for excluded flux and R_s is sufficient for accurate results. For the excluded flux comparison, the magnetic code allows us to find the local magnetic field just outside the separatrix (with an assumption of no plasma outside the separatrix). However, a complication arises from using a rectangular grid mesh. Where the separatrix makes a step transition in radius, the magnetic field is locally perturbed (by the corner). This makes the trace of B_z @ separatrix spatially discontinuous. A choice must be made in how to average this, and the excluded flux, result. Do we want to average via a root-mean-square approach or based on a linear approach? Also, over what range do we want to average the results? Alternatively, we could take the field further radially from the discontinuities in the separatrix (effectively allowing the magnetic solver to do the averaging). We choose to assume an RMS averaging over a space between the nominal locations of the probes is sufficient. Considering that our original choice of a cubic fit interpolation for radius doesn't take into account pressure, this approach seems more natural than a linear average.

For a few select shots we will make all appropriate compensations and display the comparison of the magnetically compensated approach versus the "standard" (vacuum reference set B area with vacuum reference subtraction) approach.

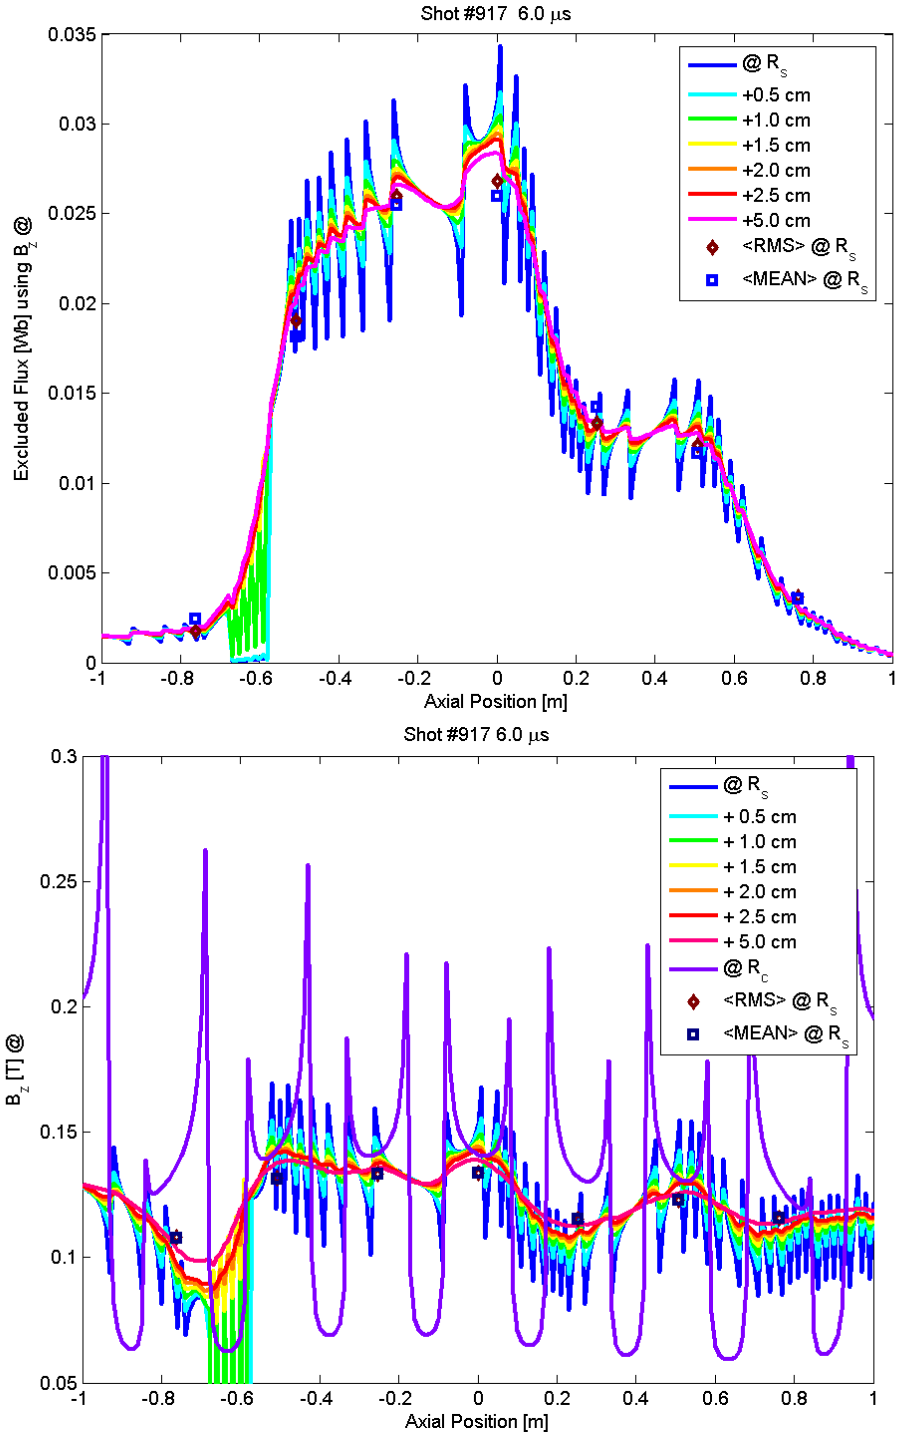

Figure 5.1.14.1. Solving for excluded Flux and B_z averaging options. (top) This plot shows the magnetic simulation results for excluded flux. (bottom) Plot of local B_z from R_s to R_c. Since the external diagnostic is intended to give a spatially averaged result (and the discontinuities are a result of the simulation) we average this result using the RMS method into regions "under" the external probes.

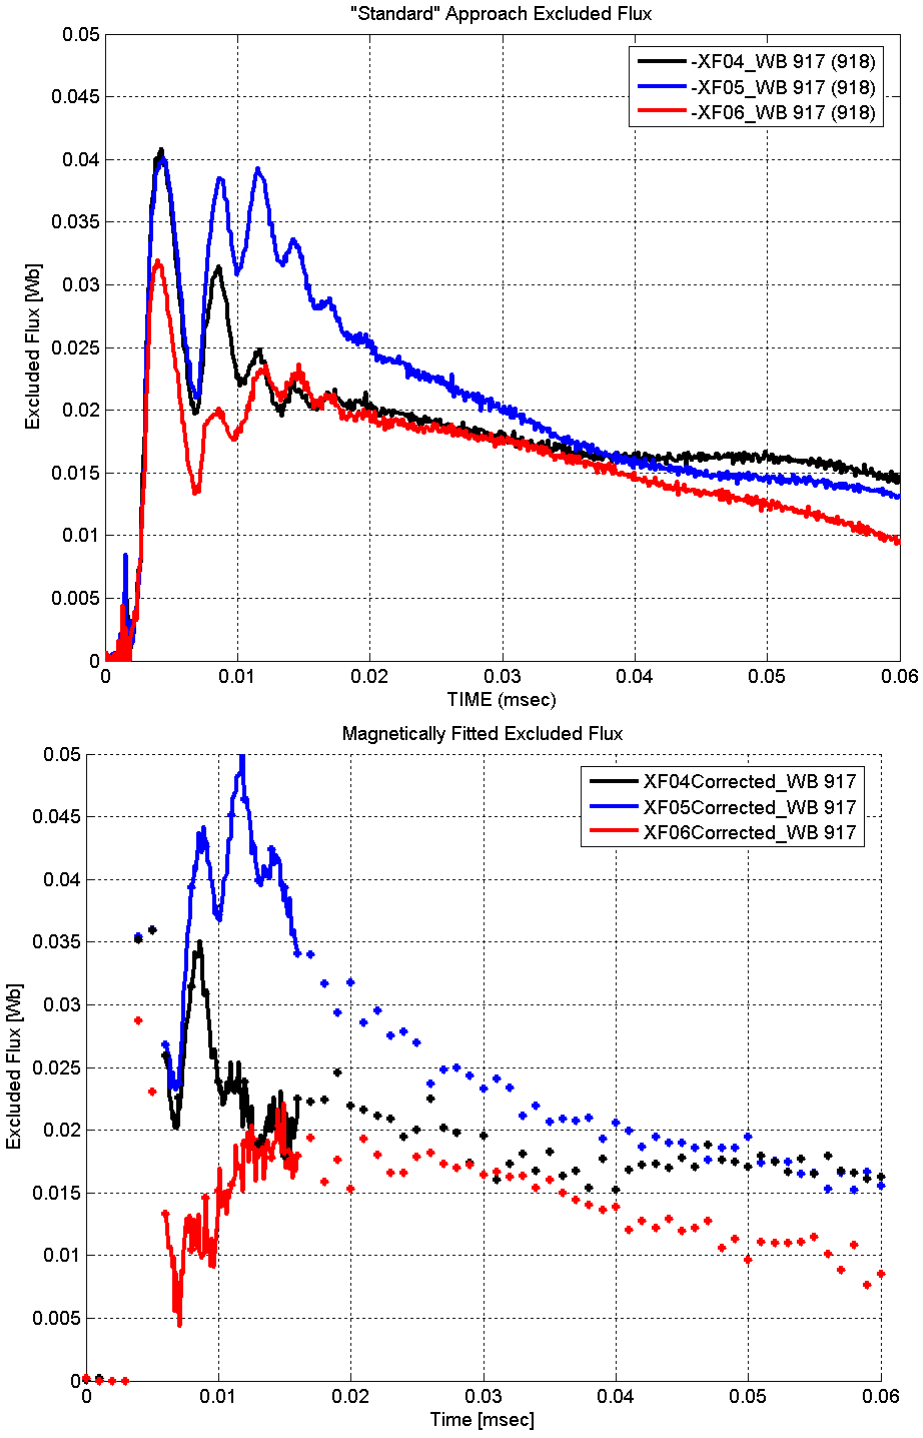

Figure 5.1.14.2. Standard and magnetically fitted excluded flux results for shot#917. The magnetically fitted approach has more noise, as expected.

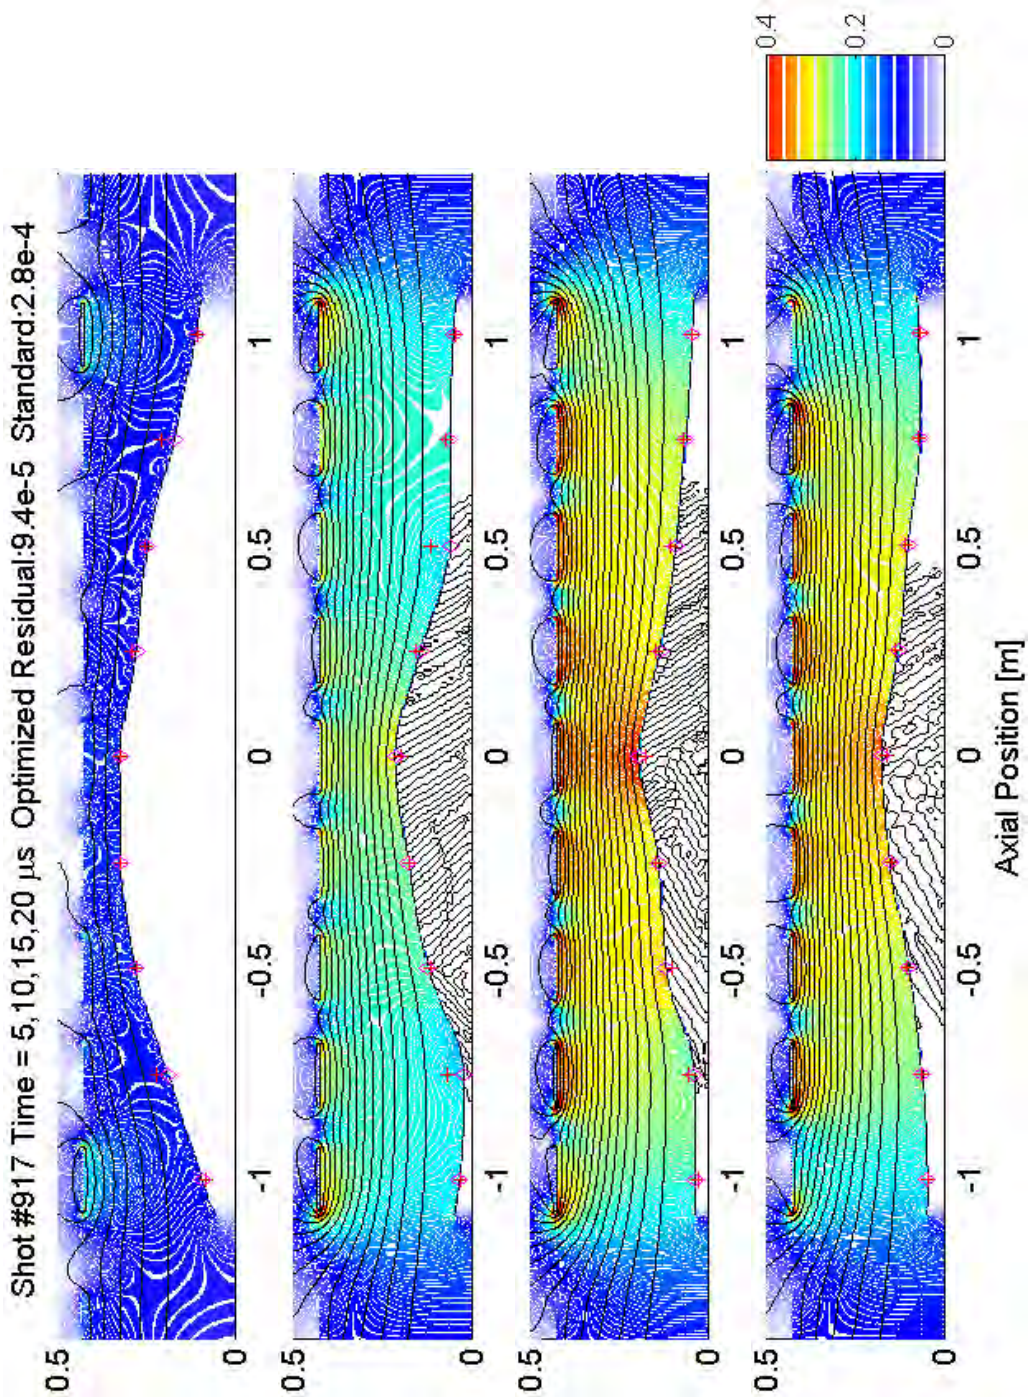

Figure 5.1.14.3. Select magnetically fitted profiles (shot #917). Red crosses indicate radii derived from the "standard" vacuum reference based approach. Magenta diamonds indicate magnetic simulation fitted points based on vacuum shot #918 B probe calibration. Magnitude of B [T] is represented by the color scale with white stripes added to aid in level determination. The color scale does not accurately display the location of these stripes (aliasing problem). Black contours indicate flux surfaces. Flux contours are shown in 10 mWb divisions.

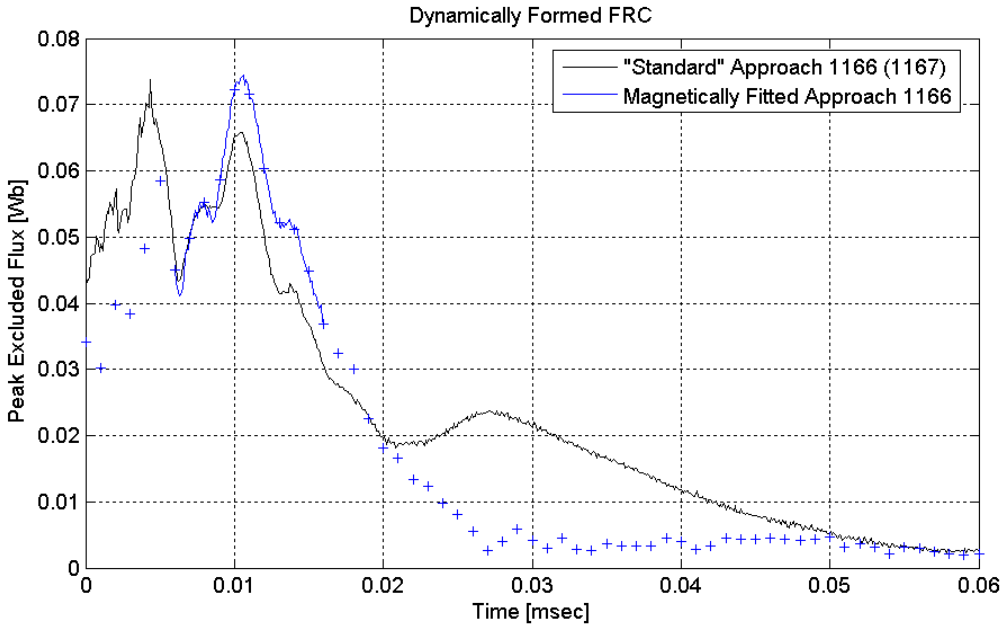

Figure 5.1.14.4. Peak excluded flux for dynamically formed FRC (shot #1166). The separatrix is closed by 5 μs. Note that even after this time the magnetic fitting can result in excluded fluxes that are above and below the standard approach. The magnetic fitting has difficulties at early times due to the interpolation shape.

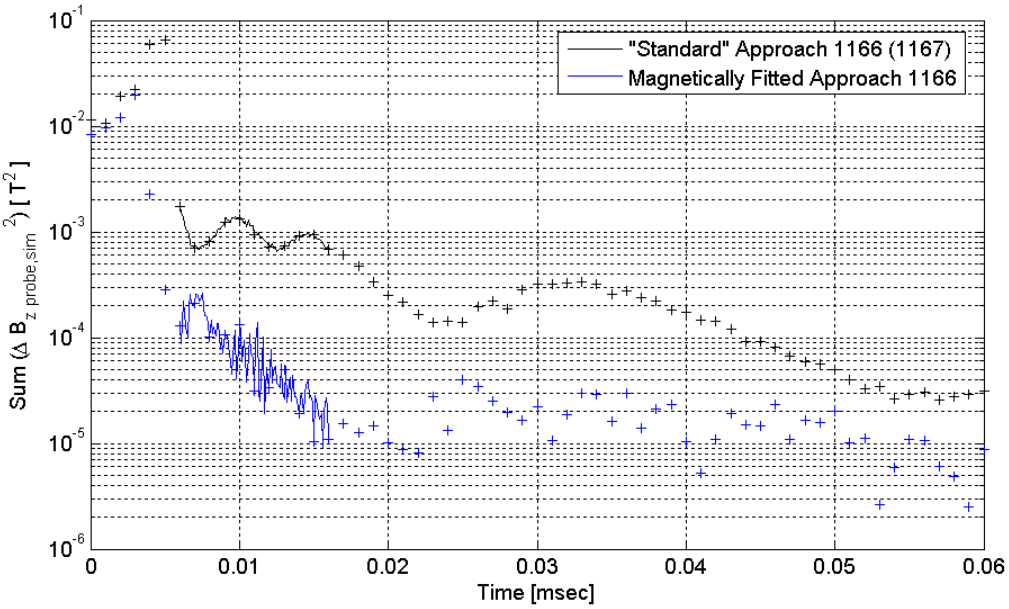

Figure 5.1.14.5. Residuals for "Standard" and "Magnetically Fitted" approaches (shot #1166). The "Magnetically Fitted" routine is not always able to find a separatrix shape that is a better match to data than the "Standard" approach. Typically it lowers the error by an order of magnitude when allowed to run for 5 minutes per time step.

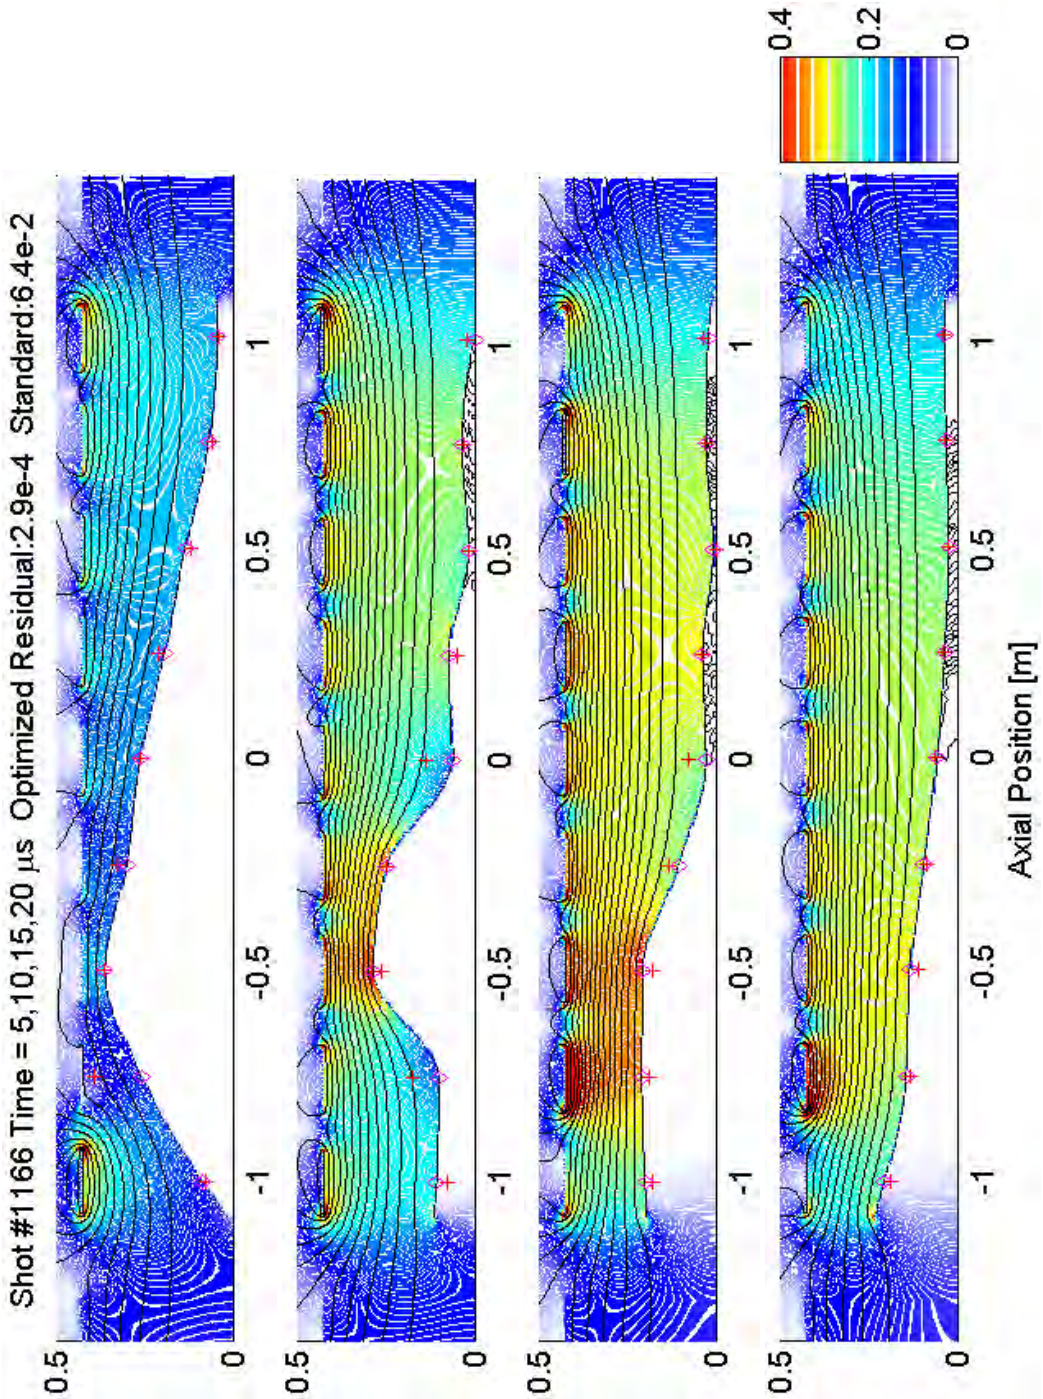

Figure 5.1.14.6. Select magnetically fitted profiles (shot #1166). Red crosses indicate radii derived from the "standard" vacuum reference based approach. Magenta diamonds indicate magnetic simulation fitted points based on vacuum shot #918 B probe calibration. Magnitude of B [T] is represented by the color scale with white stripes added to aid in level determination. The color scale does not accurately display the location of these stripes (aliasing problem). Black contours indicate flux surfaces. Flux contours are shown in 10 mWb divisions.

Intriguingly, our "standard" approach has a great deal of similarity (in the result) to the more proper approach. Unfortunately, this observation does not apply to all situations. We continue to preferentially use our "standard" approach due primarily to its low computational costs with somewhat reasonable accuracy. It is also the choice more consistent with previous precedent (and hence best for comparison to previous experiments). Note that the R_s derived from both approaches are quite comparable (within centimeters); which is consistent with experimental tests involving metal pipes inserted under the coils (in previous experiments) showing satisfactory results for the "standard" approach. Most of the error in the "standard" approach is accrued through what is effectively an inaccurate estimate of B_z at the separatrix.

We have not chosen to include here the preliminary analysis that forced us to consider a full correction. It was self-apparent from semi-analytically derived field profiles that an error should be present for either segmented coils or an FRC with $\delta R_s / \delta z \neq 0$. It was also clear that the error would in some way be a function of R_s. For long FRCs the error in excluded flux was expected to range up to +/- 20% for reasonable conditions utilizing the "standard" approach. What was not immediately clear were any analytic solutions for the magnitude of (or correction for) the error with arbitrary realistic geometries. We now have a full numerical correction available, albeit with high computational expense while also lacking a plasma physics based separatrix profile.

5.1.15 Summary and Commentary

At this point we have briefly addressed all the major issues of concern in developing a meaningful external magnetic diagnostic. We haven't addressed EM radiation pickup because of the close proximity to, and hence dominance of, the close magnetic fields. Similarly, any unintended loop areas have obvious consequences proportional to their size. These problems are well known, and well understood. In any case, proper shielding and layout is typically sufficient to limit these types of pickup. For example, the discussion of cable transfer impedance also applies to EM radiation shielding. Instead we have concentrated on analyzing the related, but less obvious, effects of coupling such as the resonances influenced by the grounding system. Note that hard connected "ground loops" are not required to have many of the problems of "ground loops". We have emphasized the importance of careful component choices and the unavoidable limitations thereof. We cannot emphasize enough that there is no such thing as an "ideal" resistor or capacitor. The design challenge is to limit these non-idealities so that they are minimized and correctable for in analysis.

Similarly, in any new design the assumptions inherent in previous designs must be understood and challenged to test if they remain valid. It has been a common unspoken assumption that the "standard" approach in solving for excluded flux is sufficiently accurate. Our opinion is that this assumption has been unchallenged for so long because of the difficulties involved in making the proper corrections. Who has the time (or alternatively computational resources)? Ironically, for most of the

work that is to follow, we shall follow this same path. At some point in the near future it is anticipated that we will have a sufficiently swift correction algorithm so as to make regular use of it. Solutions that utilize a small (and unfortunately nonlinear) matrix solution method have been implemented with some success under certain limitations. Unfortunately, the limitations correspond to the extrema cases of a FRC with a spatially constant radius and a FRC only under one strap. Since we don't currently have the time, this paper will be limited to qualitative rather than quantitative comparisons of many operating modes. It is hoped that the reader will understand our general tendency to reticence in specifying parameters that are highly sensitive to errors in excluded flux and pressure. However, it should also be noted that after the axial implosion relaxes (for in-situ shot #917) the standard approach agrees relatively well (better than 20%) with the magnetic fitted approach. At this point caution is advisable, but not disbelief.

5.2 Interferometer

The second critical diagnostic required to determine gross plasma parameters is a single-chord interferometer. Having a good interferometer combined with magnetic diagnostics allows us to determine density and total temperature. In this context, interferometer refers to a diagnostic that uses the phenomenon of coherent interference to measure the difference in the length of two optical legs: a reference beam and a scene beam. The scene beam sees a significant difference in index of refraction when free electrons are present. This allows for a measurement of average electron density. To keep the analysis simple, the frequency of light must be kept well above the cut-off frequency.

5.2.1 Design Criteria

A good interferometer also satisfies the following criteria:

1) It must be simple to maintain.
2) It must be easy to analyze the data.
3) It must be immune to noise generated from the experiment.

It must be simple to maintain. We have very limited time available to maintain the many varied systems associated with this large experiment. If a system requires frequent maintenance it will be offline for a large number of shots. If a system requires maintenance after every shot it won't provide data for more than a few shots. Some characteristics make a system easier to maintain. Using a visible (632.8 nm) He-Ne laser makes initial setup and determination of light path obstructions simpler. Using a thermally stabilized laser to prevent mode hopping (and the attendant change in location) is critical. Choosing a short laser cavity also helps to reduce mode hopping. Constructing the interferometer table from a material not prone to sag removes the need for constant fine tuning of optics.

It must be easy to analyze the data. In a simple single-detector interferometer there are two points where ambiguity exists in determining the change in path length. When the detector is at the peak, or the bottom of the trough, of a fringe it is impossible to tell whether the path length is increasing or decreasing. Assumptions

can be made as to the continuity of motion. I.e., the length was increasing so it probably still is increasing. Assuming the beginning and ending densities to be near zero also helps in limiting the ambiguity. Additionally, a practical system needs some mechanical vibration to identify the contrast and magnitude of the interference fringes. Also, all the elements in the system need to have linear responses.

The issues in analysis listed thus far can be eliminated by using a system where the effective path length (measured normalized to wavelength) of one beam is modulated sufficiently to drive the interference pattern through at least one fringe. The frequency of this driven shift must be sufficiently higher than the frequency of the path length being measured. One such system uses an acoustic-optical modulator (Bragg) cell to split the laser source into two beams of slightly different color (frequency). When they are recombined a beat is present at a frequency that is the difference between the two colors. This is known as a heterodyne interferometer. It benefits from needing only a single detector whose linearity is unimportant. The information regarding path-length is carried in the frequencies surrounding the Bragg cell modulation frequency. Nonlinearities result in harmonics of the region about the modulation frequency (i.e., creating higher frequency signals). Thus, nonlinearities in the detector or amplifiers do not significantly impact the signal.

Another reasonable solution is to use two detectors (and no modulators) in what is known as a homodyne configuration. In such a system, the information is carried in the low frequency magnitudes of the system. Having two detectors can remove the ambiguity problem mentioned earlier. However, the other original problems still persist. We choose to use the Bragg cell based heterodyne system. Rather than recording the modulated signal directly we mix it with the 40 MHz modulation signal to generate two low frequency signals. These signals combined on an X-Y plane provide phase as the angle formed with respect to origin. The distance from origin is proportional to signal strength.

The diagnostic must be immune to noise generated from the experiment. Noise takes two forms: electrical and acoustic. Dealing with electrical noise is typically easier than acoustic.

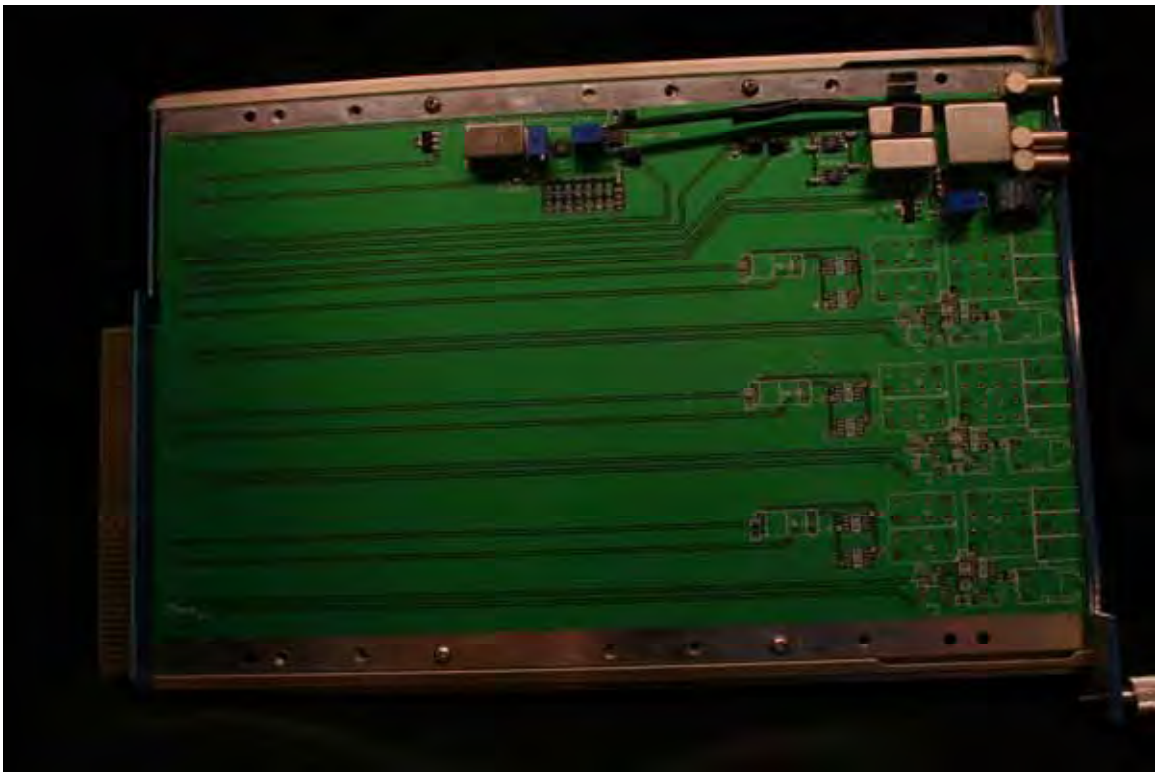

Figure 5.2.1.1. CAMAC interferometer board. All the electronics necessary to drive the Bragg cell are contained on this board. The detector, pre-amplifiers, mixers, and final signal processing are also present. Additional positions exist on the board to support multi-chord interferometry.

Keeping the electronics for the interferometer in very close proximity to the digitizers is important for reducing the coupling of ground noise into the recorded signals. The detector does not have to be near the experiment. We relay the fringe light via fiber back to the diagnostics rack. Similarly, the modulation signal for the acoustic-optical cell can be run out from the diagnostics rack on well shielded, heavy wall, coax. It is grounded as it exits the diagnostics rack. If reflections on the coax due to the Bragg cell were an issue the simplest fix is to put a heavy attenuator on the coax as it enters the Bragg cell. Our integrated Bragg cell driver has a proper 50 Ohms driving impedance with limited output power. The Bragg cell is electrically separated from the metal enclosure in which it resides to prevent large low frequency ground currents. If electrical noise coupling to the Bragg cell were an issue it could be resolved by relaying the drive signal to the laser enclosure on a fiber. The local AC power line can then be used to amplify the signal and drive the Bragg cell.

Figure 5.2.1.2. Bragg cell and laser enclosure. As can be seen, the Bragg cell (upper left corner) is kept electrically isolated from the rest of the box. The Bragg cell is an IntraAction AOM-402AF1. The laser source is a JDS Uniphase 4 mW HeNe laser. The laser is less than 16 cm long. The overall metal enclosure dimensions are 25 x 30 x 13 cm.

The AC power for the laser also provides the ground for the metal box enclosing the laser and the Bragg cell. It has not been shown that this metal box is actually necessary. In truth, it is possible for such a large metal box to do more harm than good due to the forces exerted on it by the experiment's magnetic field.

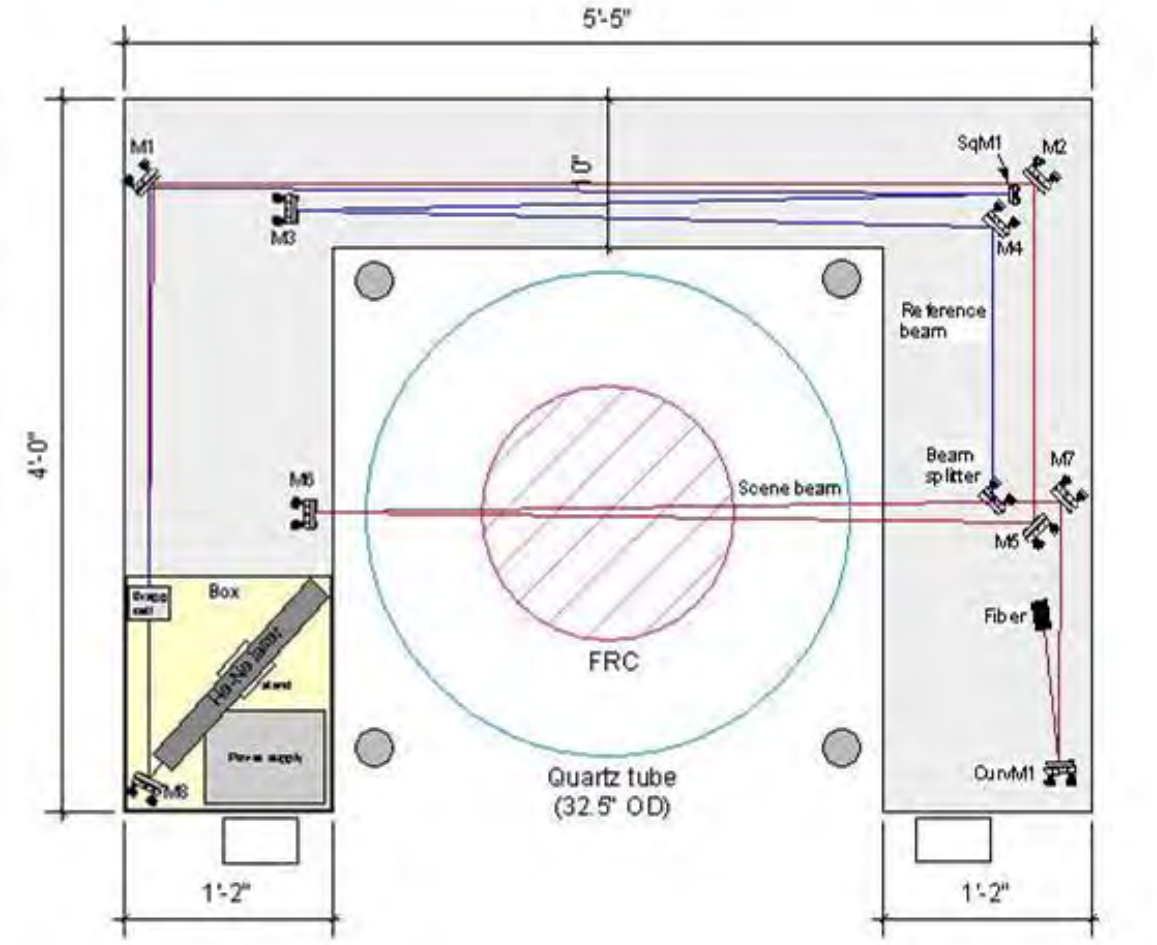

Figure 5.2.1.3. Interferometer optics table diagram. The table itself is composed mostly of FR4. As can be seen from this diagram, only half of the available light is collected. Laser intensity has not been a serious issue. The scene beam passes through two windows in the vacuum chamber.

The above figure shows the layout of the interferometer around the experiment. Like all optics tables, it is desirable to make it as structurally stiff and as massive as is feasible. Many of the structural modes associated with this U-shaped tuning fork can be expected to show up as signals. The important consideration to realize is that our experiment is not a steady state experiment. We need the interferometer signal to be stable for only ~ 400 μs after the first significant magnets are energized. Knowing that the speed of sound in air is ~ 350 m/s yields a required clearance distance of 14 cm. Similarly, many structural materials have longitudinal sound speeds approaching 6000 m/s suggesting a required distance of 2.4 meters along structures to any driven surface.

Having a delay in the time it takes for the acoustical noise to reach the interferometer is not sufficient to remove the noise that is present on a continual basis. Noise sources such as cryo, turbo, and roughing pumps are connected to the machine and introduce vibrations into the system. Airborne vibrations are also capable of

interacting with the system. The stiffer and more massive the optics system, the more energy is required to result in a given deflection. The weaker the spring constant of the supports for the optics table, the greater a driving deflection required to transmit a given amount of energy via structural materials. Of course damping is required to limit any potential resonances.

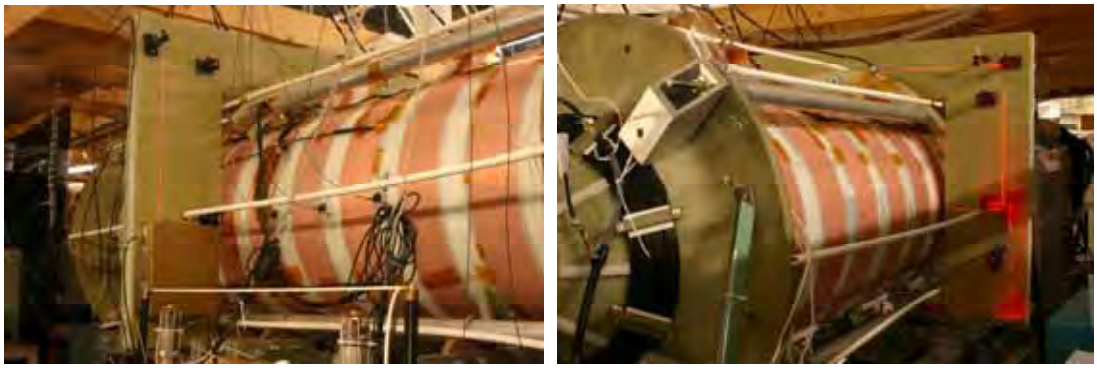

Figure 5.2.1.4. Interferometer as actually implemented. Note that the interferometer table is supported by the top two axial aluminum rods.

The design must balance conflicting design goals. Any practical interferometer system has to make some tradeoffs between the three criteria of easy maintenance, simple analysis, and immunity to noise. In addition a fourth criterion has to be taken into account: interfacing with the experiment. Earlier versions of this experiment hung the interferometer table by springs from an overhead structure. This benefited the diagnostic's immunity to structurally transmitted acoustic noise. Many milliseconds were available before the acoustic waves generated by the experiment affected the diagnostic in a major way. The downside was that it made maintenance difficult. Attempting to adjust the alignment of the optics components would inevitably set the table into oscillation about the experiment. At the time, precise alignment with respect to the quartz tube of the experiment was necessary due to striations in the clear quartz tube. In the current version, the optics table is supported by the aluminum rods that take up the axial loading in the vacuum vessel. Additionally, holes have been cut in the optics table to permit the axial discharge current returns to maintain their symmetry. The current version may go too far in making maintenance and construction easy at the expense of acoustic noise.

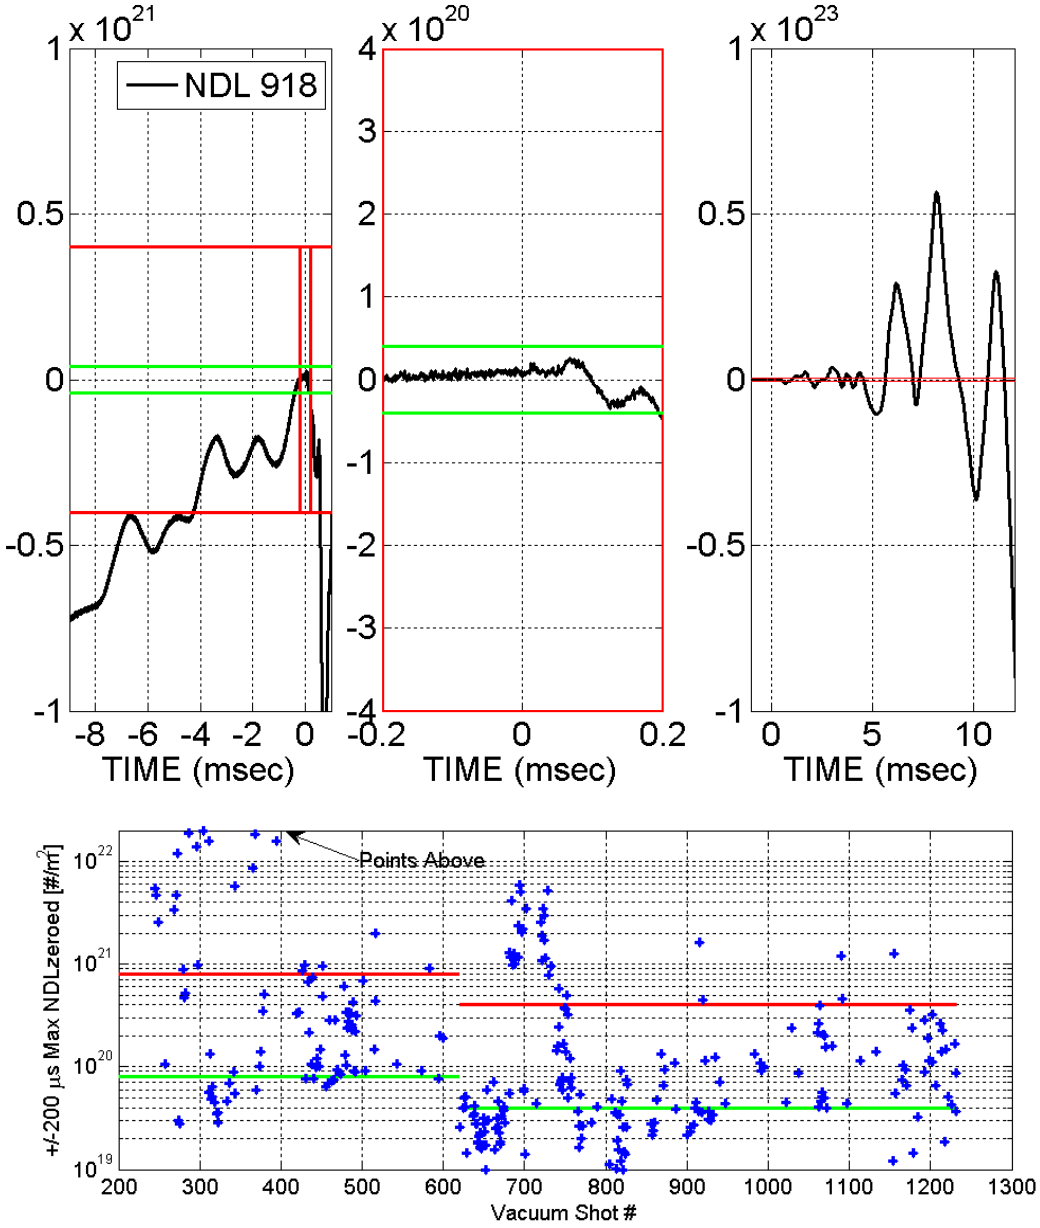

Figure 5.2.2.1. Acoustics on interferometer during vacuum reference shots. The red vertical lines bracket the time for which we need accurate data. The red horizontal lines indicate the approximate magnitude of the trace for a plasma shot. The green horizontal lines indicate 10%. (bottom) Summary of vacuum reference shot spline fit subtracted interferometer error. It can be seen that there are sequences of shots where the interferometer is reasonably well maintained and others where it is decidedly less so.

5.2.2 Acoustic Results

The current interferometer system is not perfect. The above plots show typical interferometer traces from motion before, during, and after the period of interest. Obviously, in this figure t=0 is fortuitous. The pre-shot acoustics 500 μs earlier would have exceeded 10% error. To help compensate for pre-shot acoustic

effects, the *NDL* signal between approximately -500 μs to -150 μs along with 150 μs to 500 μs is used to spline interpolate the acoustic contribution (which is subtracted) between -150 μs and 150 μs (*NDLzeroed*). The bias fields are fired at -150 μs followed by the reversal banks at 0 μs. Some lower power systems fire milliseconds before these times. The first clear sign of acoustic contamination from the experiment comes approximately 200 μs after the bias fields are fired, this is consistent with an airborne wave traveling from the driven strap 7 cm to mirror M6. The next unambiguously strong kick to the system comes at 680 μs after the bias fields. This timing would be consistent with a wave traveling at ~ 2000 m/s in the FR4 end-cone and interferometer table and ~ 6000 m/s in the aluminum support rods. The source would be the bias fields in the driven straps on the vacuum vessel. Another 150 μs after this we see a kick in the opposite direction presumably due to the reversal fields. The very strong increase in acoustics around 5 ms is most likely coupled through the air from the reversal modules (1.8 m in air, 24 m through structures). At 12 ms the interferometer has been hit hard enough to move it temporarily out of alignment (fringe contrast ->0).

Unfortunately, it is not uncommon for shorter delays in acoustic coupling to occur for many shots. These are probably caused by bridging the air gap between interferometer and tube by the careless placement of coaxial cables or other materials. Contacting items have been removed from the interferometer many times.

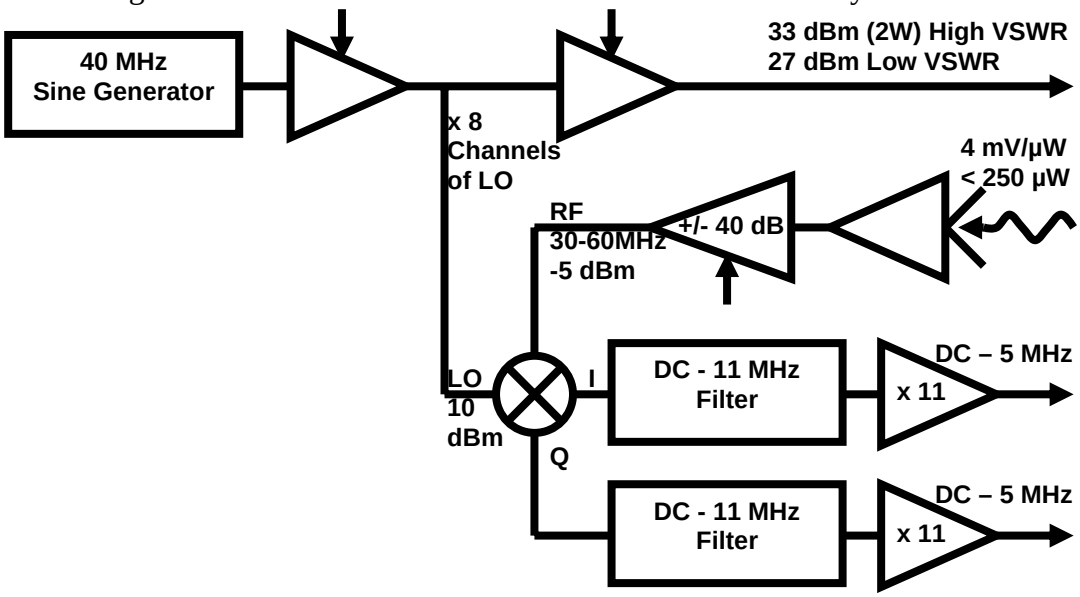

Figure 5.2.2.2. Interferometer electronics diagram. The mixer represents an I&Q demodulator (a 90° phase difference and two double balanced mixers).

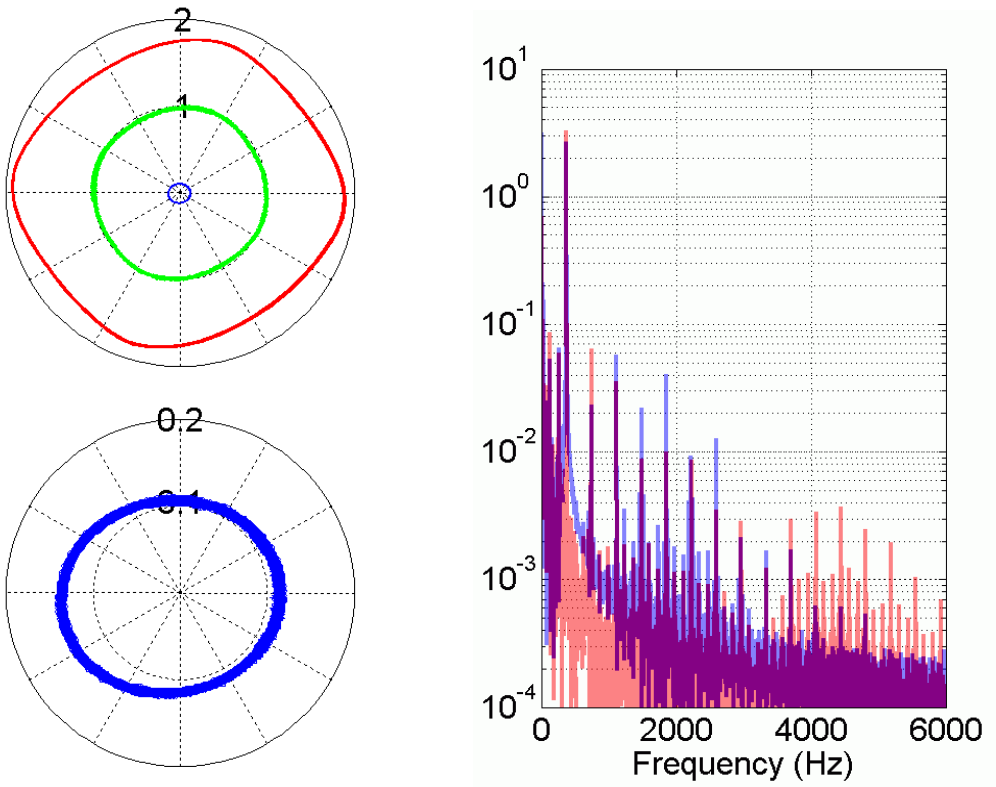

Figure 5.2.2.3. Results from interferometer driven by 370 Hz air transmitted acoustics at various levels of receiver gain. The 370 Hz test source may have unintentional harmonic components. (left) At left are plots of system outputs for three receiver gain levels. The red case corresponds with the input amplifier gain set high enough to drive the mixer just short of the 1 dB compression point. Green corresponds with a level used during normal operation. (right) The spectral plot at right shows the frequencies present in the apparent phase for the peak and minimal levels of RF mixer power. If all the power beyond 1 kHz is taken to be distortion, this still amounts to less than 1/60th of a fringe shift.

5.2.3 Electronics Results

The above figure shows the output of the system when excited by a 370 Hz airborne acoustic signal. At high RF power, the diamond shaped output is probably due to the mixer RF input being driven to the 1 dB compression point (1 dB of input RF energy is not going in to the expected output frequencies). In a mixer (or phase detector, or I&Q demodulator) the LO (local oscillator) input can be thought of as turning on and off switches (diodes in this case) in a double balanced mixer ring. The LO has to be strong enough to thoroughly turn on the diodes. The RF has to be weak enough not to change the 'ON' state of the diode significantly. This requirement is progressively more violated as the RF input power approaches the 1 dB compression point. At some point, adding RF power results in no power gain at the output.

Without the low-pass filters we would expect approximately 2 mV @ 40 MHz from LO to IF isolation. At high levels of RF input power (not obtained here), the RF can also significantly couple to the IF outputs resulting in additional 40 MHz contamination. For these reasons, additional low-pass filters are typically added after such a mixing operation.

5.2.4 Options

The phase shift recorded by an interferometer is $\Delta\varphi = (\omega / (2c\, n_c)) \int (n_e) dl = 2.818E\text{-}15\, \lambda\, NDL$. *NDL* and *NDLzeroed* do not compensate for the interferometer being a two pass instrument nor are they compensated for the radius of the plasma. These compensations are left for the user (or additional signal processing) to perform. It might be desirable to implement the interferometer with a longer wavelength laser to take advantage of the λ with NDL. Of course, phase due to path length scales with λ^{-1}. The immediate downside of a longer wavelength is the loss of visibility and difficulty in gross alignment. Commercially utilized parts are available down to 1.6 μm without any difficulties (lasers, Bragg cells, fiber, fiber receivers). See the figure below for fiber attenuation rates. Most of these parts are direct substitutions for elements in the current system. Fiber is available with useable pass-bands down to 9 μm (1000 dB/km from CorActive). Adding an infrared interferometer on the same optical path as the visible interferometer would theoretically allow for the complete removal of the acoustic effects (until the vibrations move the system out of alignment). In practice this has been difficult to achieve in past experiments. For our timescales of operation such additional complication should prove unnecessary despite the temptation of such a 'quick' fix as long as the effects from continually present acoustics can be controlled.

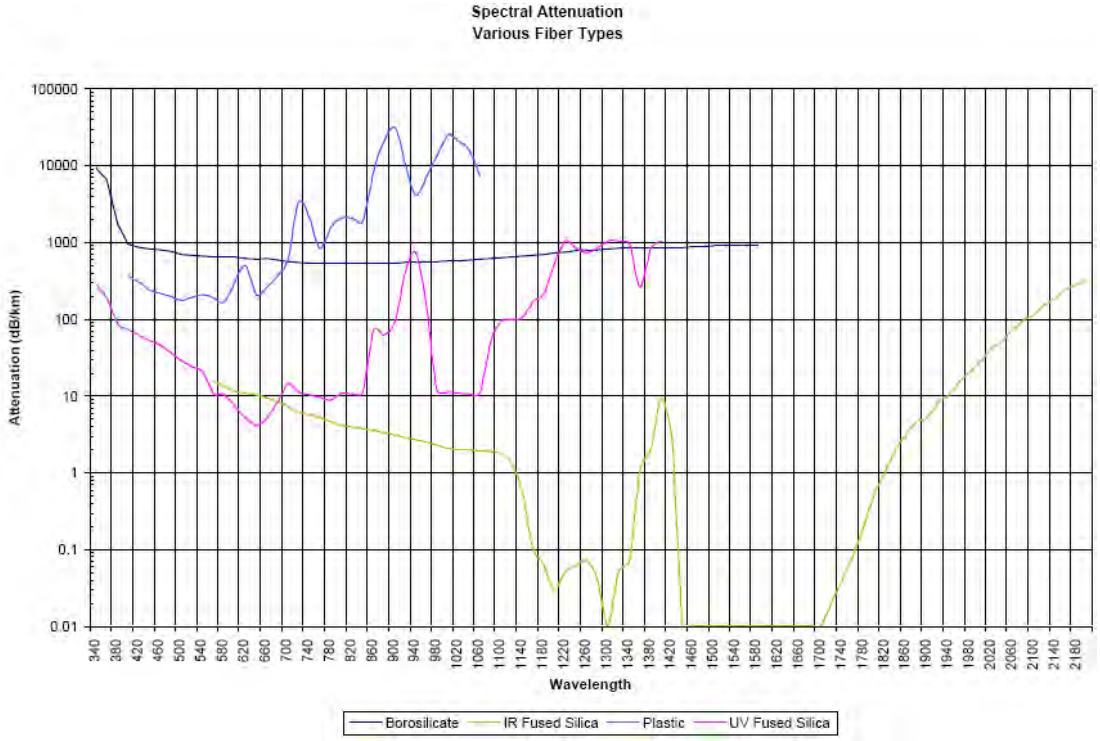

Figure 5.2.4.1. Graph of spectral attenuation for readily available fiber types. Available from Fiberoptic Systems, Inc. http://www.fiberopticsystems.com/tech2.htm. Reproduced with permission.

The simplest measure of the health of the interferometer over multiple shots can be obtained by looking at the ratio of minimum signal strength to maximum signal strength for the timeframe of interest. We choose not to look at the absolute magnitude of the signal strength because the gain of the system is variable (and can be altered without any record). This approach is simple because no distinction between plasma and vacuum shots is required. It may be unduly negative because the effects of an ellipsoidal trace and high frequency contamination would be included (that wouldn't result in a distortion of the density trace). It is also not a useful number for determining the minimum detectable density or error levels.

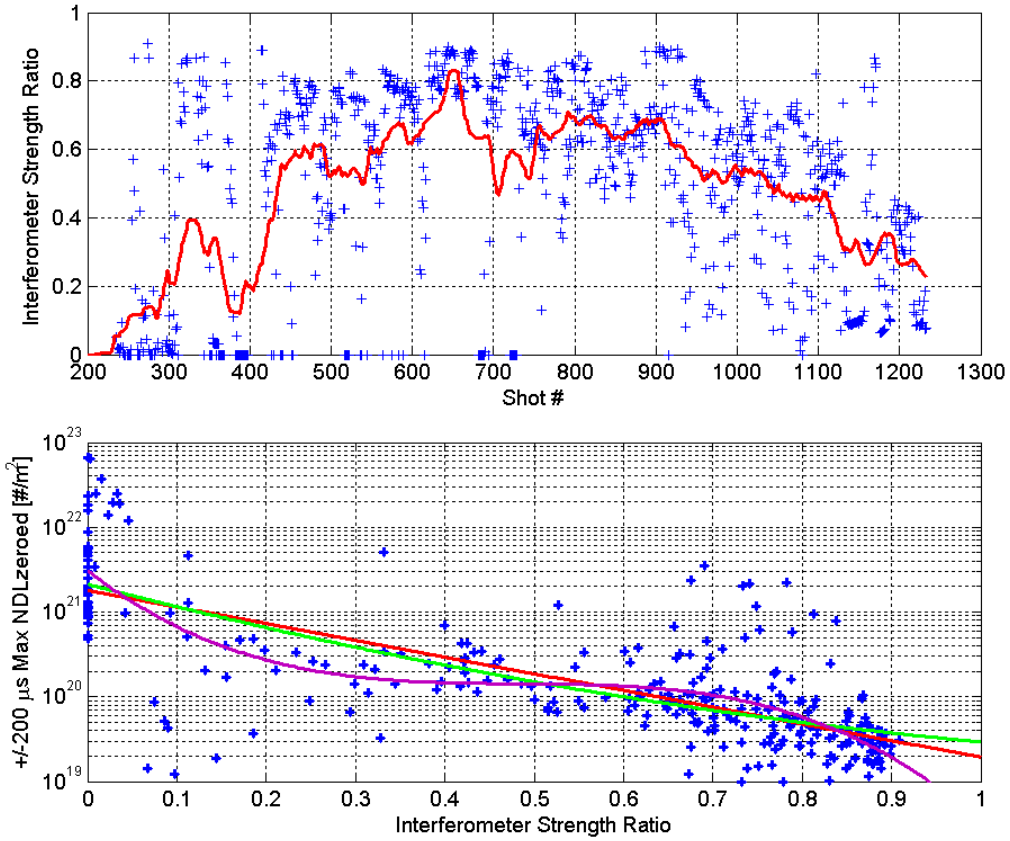

Figure 5.2.4.2. Inteferometer strength ratio. (top) The blue crosses display minimum amplitude of interferometer divided by maximum amplitude of interferometer carrier signal (for times before 200 μs). No numerical frequency filtering has been applied. The red trace is a 51 point moving average of this data. A ratio of one would be ideal. (bottom) For the vacuum reference shots the interferometer strength ratio can be compared to the error levels for *NDLzeroed*. The red trace is the result of a linear fit to the log of the data. Green the result of a quadratic; purple, cubic. A tight correlation would be useful in judging the placing of error bars on the interferometer during plasma shots.

5.3 Visible Light Array

The Visible Light Array (VLA) utilizes commercially available silicon PIN detectors with an integrated amplifier and fiber coupling. This diagnostic provides highly time-resolved visible-light intensity data through plastic fiber optics. The design has thermal control to stabilize gain and offset, although this is not believed to

be strictly necessary in our applications. Detector thermal noise is not a concern so we can do distributed heating to stabilize the temperature.

Neither the change in gain with temperature nor the spectral response of the detector has been specified by the manufacturer. We have measured a change of 0.2% per °C at 650 nm. The detector/amplifier is capable of >65 MHz frequency response. Design ΣRMS output noise is less than 1 mV. The gain of the system is approximately 4 mV/μW at 650 nm. Figure 5.3.2 shows the spectral response for a similar product from the same manufacturer. Full signal saturation occurs at approximately 1.4 V.

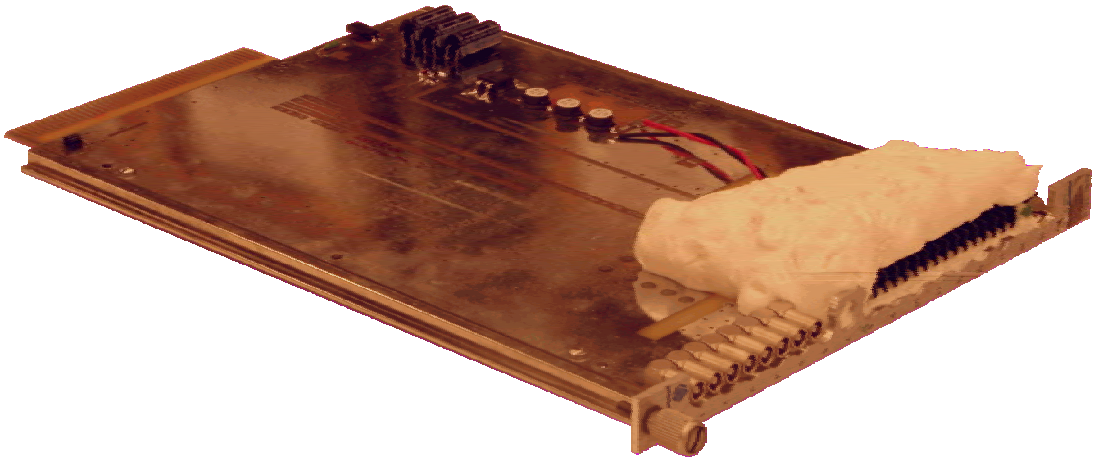

Figure 5.3.1. Thermally stabilized visible light detector array CAMAC board.

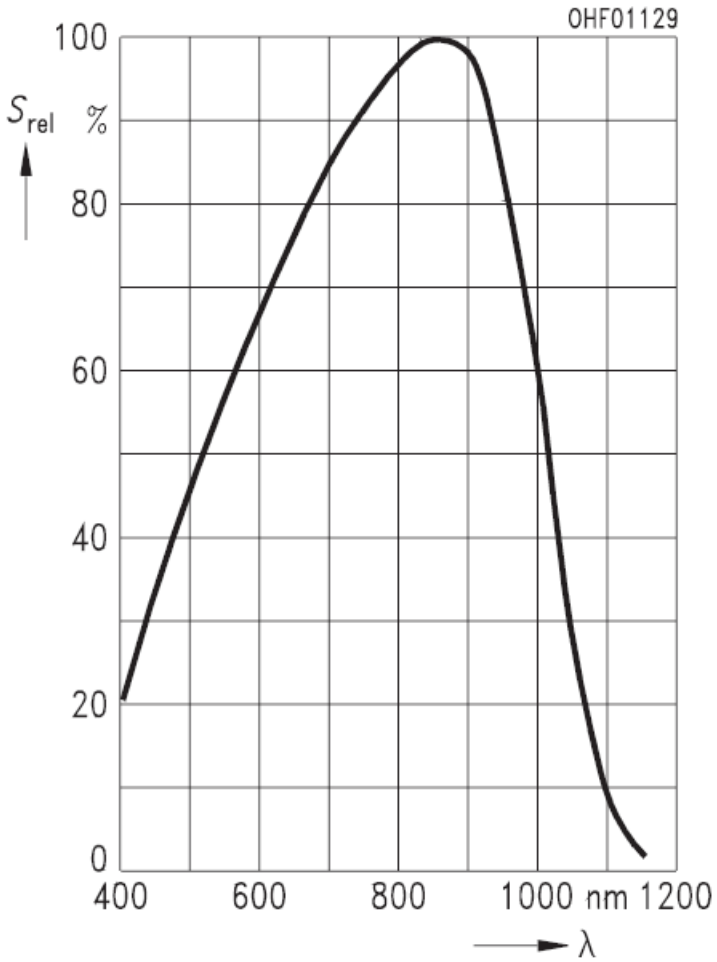

Figure 5.3.2. Example of typical spectral response for a detector with similar characteristics and purpose. Reproduced with permission from Avago Technologies (SFH250 AV01-0711EN.pdf datasheet)

This diagnostic as used does not take advantage of the AC output coupling option nor the high bandwidth; having an offset recorded with our DAQ system, or not, has no impact on the actual recorded quality. We choose to record the offset in our experiment due to the link between offset and the temperature of the detector. The collection optics is of the same design as utilized in the tomography array. Two additional Wratten Gel filters are used to window the spectra in an attempt to avoid looking exclusively at the bright deuterium ionization lines. Unfortunately, it was not realized until recently that the plastic optical fiber (POF) used has better transmission characteristics in the infrared than was previously assumed. Previous work failed to look beyond the 730 nm dip. Compare to the spectral attenuation figure in the interferometry section.

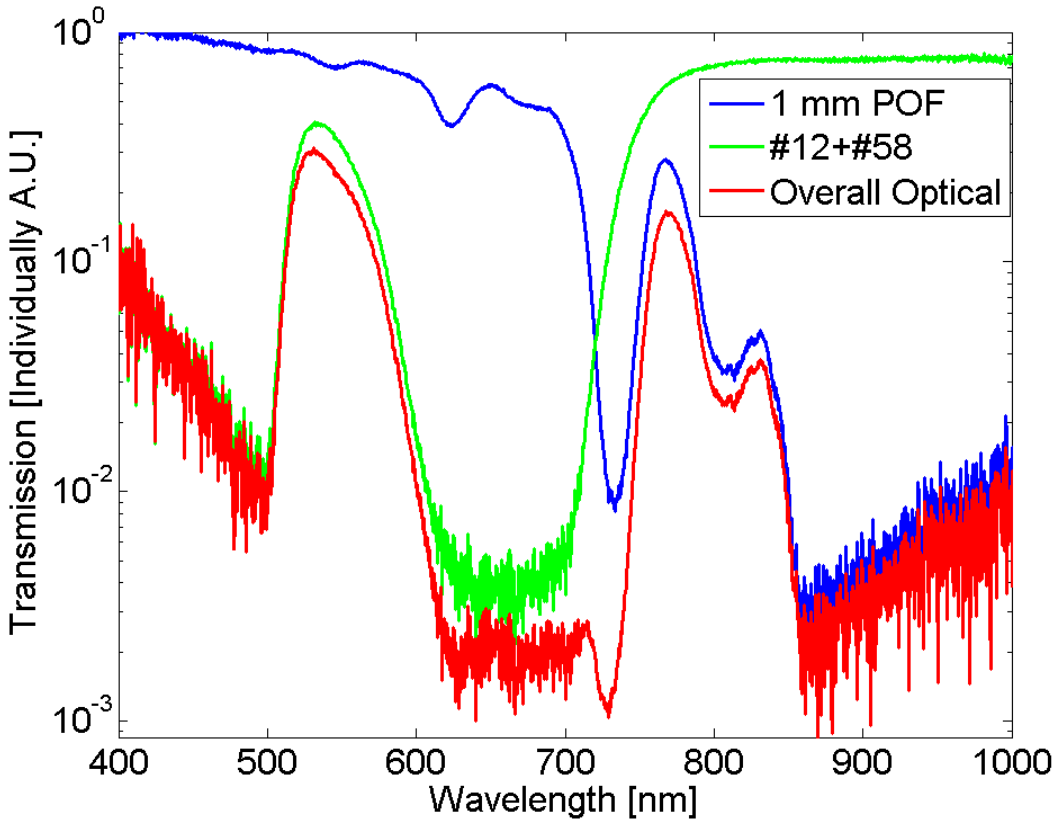

Figure 5.3.3. Spectral transmission of Wratten Gel filter combination. The transmission profiles made available by the filter reseller only describe the region between 400 & 700 nm. Initially, the NIR band was mistakenly assumed to be blocked by the plastic fiber.

5.4 Tomography

(Bremsstrahlung $\lambda = 520$ nm $\pm$ 5 nm , $\lambda = 523.5$ nm $\pm$ 0.5 nm)

The tomographic diagnostic is a left-over from the early version of this experiment (when there was a 40 cm clear quartz tube in the formation section). Those experiments did not show full agreement between the analysis based on either Cormack-Granetz, Maximum Entropy, and Minimum Fisher and the end on imaging provided by the Imacon (fast framing camera). In particular, when the Imacon showed relatively rotationally symmetric emissions (except with fine structures that VLA evidence suggested were spatially localized to the ends), the tomographic reconstructions would show definitively nonsymmetrical structures. The different analysis methods tended to yield qualitatively similar results. There are multiple possible explanations for the discrepancy. The Imacon takes an axially averaged image, so if the instability is highly localized the image would have a poor contrast ratio. Also, the Imacon was generally filtered with the same 500 - 600 nm gel wratten filter pack as the VLA. The most troublesome possible explanation is of course the possibility of setup or calibration errors in the tomography array.

This diagnostic has not been set up on the most recent experimental runs due to the lack of good optical access and manpower. Applications in shorter wavelengths should consider the potential complication of emissions from the polyethylene cladding of the fiber (fluorescence) under UV illumination. The old setup is summarized in the following figure:

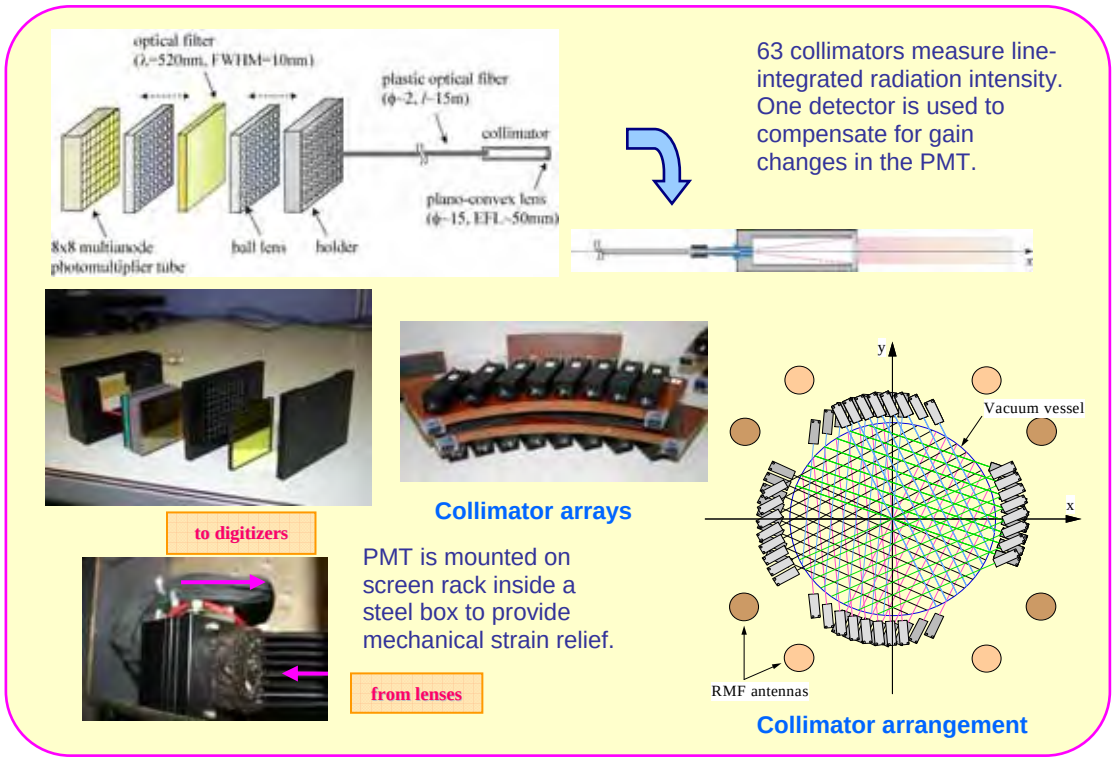

Figure 5.4.1. PMT Tomography system summary. This is from an APS poster describing the earlier version of our experiment. The PMT is a Hamamatsu H8500. The fiber used to relay the light is Mitsubishi Rayon's Super ESKA SH8001 (2.0 mm). The focusing optics around the interference filter are 5/32" diameter ball lenses. The collimator lenses are Sunex Inc's PXP105 (15 mm diameter, 50 mm EFL). The collimators typically have a viewing angle of ~1.6° (full angle). PMT is typically supplied with -900 V. Maximum allowable voltage is -1100 V.

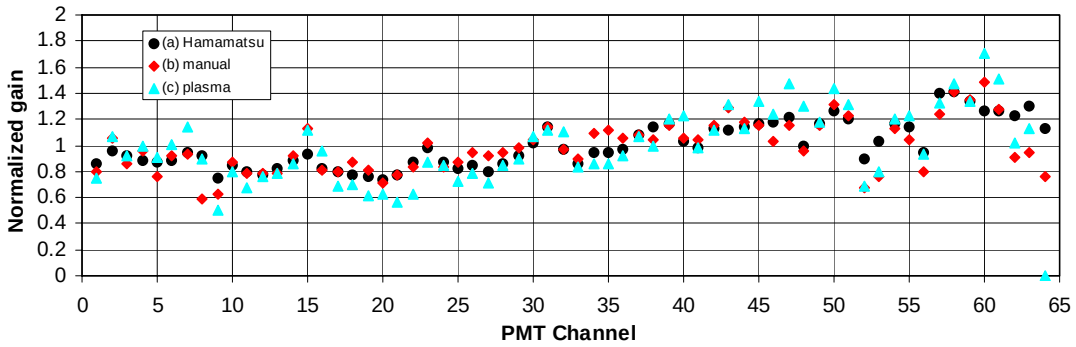

Figure 5.4.2. Comparative calibration results for tomography system. Again, this is from an APS poster describing the 40 cm version of our experiment. '(a) Hamamatsu' refers to the responsivity of the PMT channel as described by the Hamamatsu characterization of the multichannel PMT. The two remaining calibrations are full system calibrations including the effects of collimators and fiber quality. The calibration that uses an integrating sphere is referred to as '(b) manual'. The last calibration approach utilizes a theta pinch FRC during the radial implosion with all the channels viewing the r=0 axis. The required assumption is one of rotational symmetry.

The complete PMT system was calibrated for gain by use of an integrating sphere with a high power LED light source. The LEDs were chosen to provide emissions in the spectral band around the band pass filter. Obviously, it is difficult for any source to provide light intensity commensurate with that of the experiment when it is used in conjunction with an integrating sphere. Instead, the LED source provided a variable light source that could be pulsed beyond steady state max ratings and modulated to remove sources of noise from the calibration via a spectrum analyzer. Contamination from the drive signal could be ruled out by blocking the light. This made it possible to do the calibration with exceptionally weak signals (10s of mV) and still provide the exceptional relative accuracy and repeatable calibrations required for a tomographic inversion. The remaining relative mismatch between this calibration and the Hamamatsu specification is likely due to difference in fiber condition (probably end polish).

Figure 5.4.3. Integrating sphere and LED light source used for calibrating optical diagnostics response

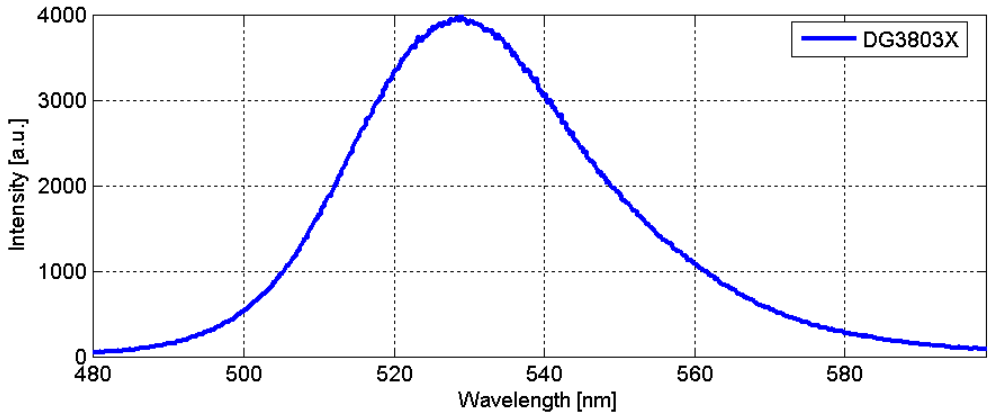

Figure 5.4.4. LED light source spectra. This is the spectra of the LEDs used with the integrating sphere.

The chosen multi-anode PMT has a very limited amount of charge available in shared internal capacitors. In addition, the internal dividing resistors include active elements in the final amplification stages for maintaining bias. A numerical method calibrating for the change in gain would require additional knowledge including starting voltage and average room-light current. The simpler solution to calibrate for gain changes due to discharge is to use a single channel with a known light source during the experiment. Since the absolute max average total anode current is 100 μA it is rather impractical to maintain a constant calibration illumination. Instead, the calibrating light source (a specially selected green fiber transmitter) is pulsed on during the shot. In our simple setup the PMT is not buffered before the analog signal is transmitted to a digitizer. The implication is that on short timescales (1 μs) proportional to the termination resistance (effectively 9.1 kΩ) and cable capacitance (30 pF/ft, ~1 m) the signal is reduced by integration.

The tomographic inversion software is based on the previous work in the Translation, Confinement, and Sustainment (TCS) experiment of George Votroubek. Hiroshi Gota holds responsibility for making the tomographic system work with our systems while George Votroubek was retained for a time to assist him. The signature result of this diagnostic (on the 40 cm version of PHD) was the identification of a possible ‘inner tube’ instability.

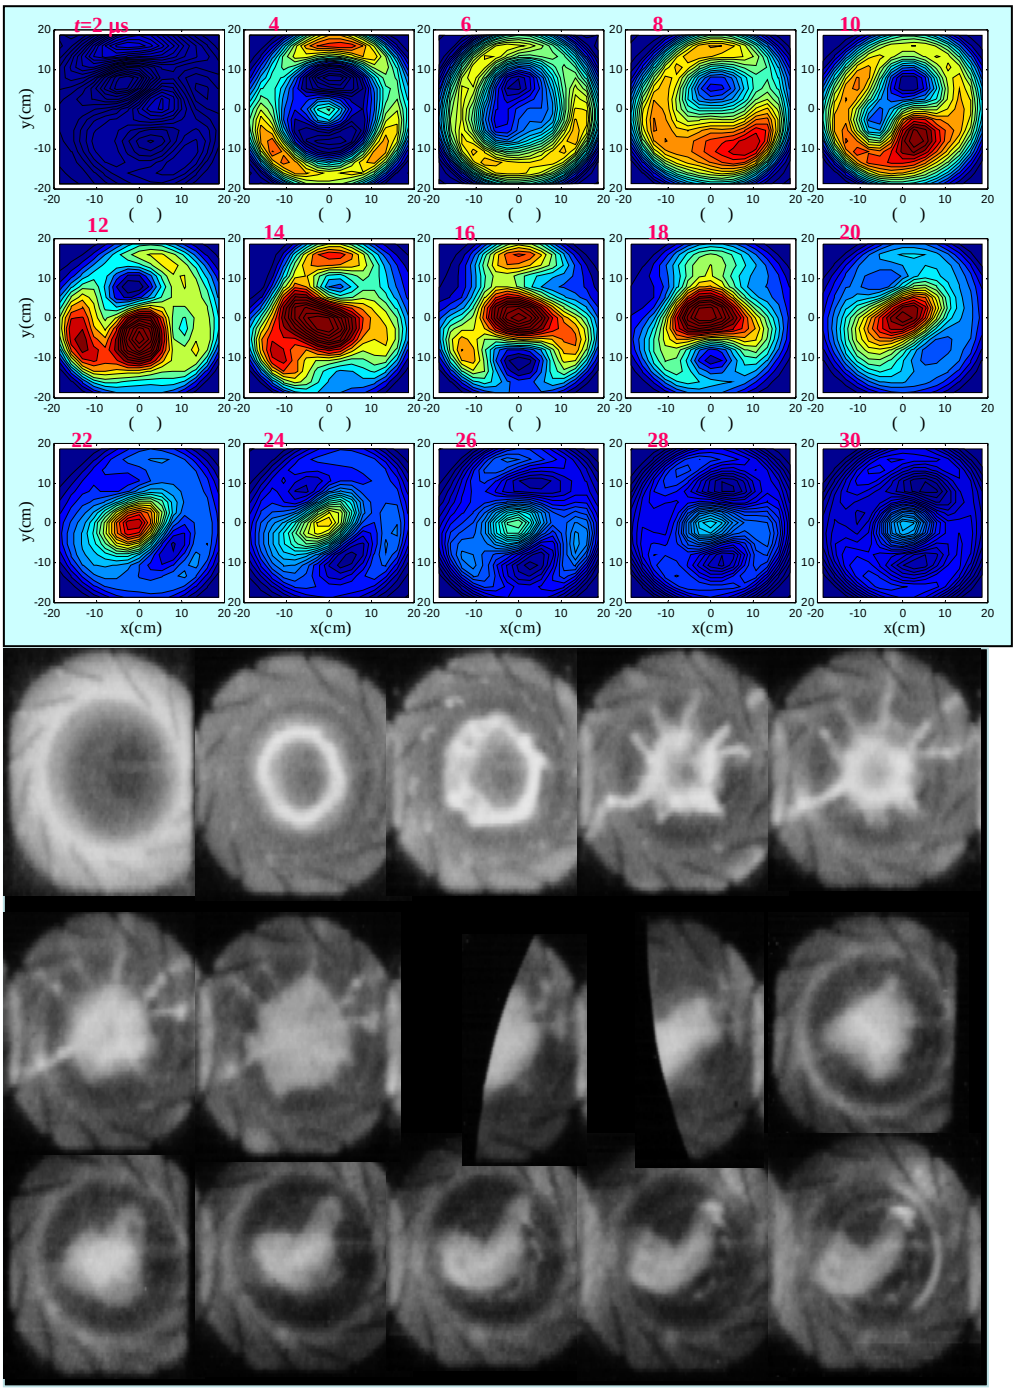

Figure 5.4.5. Comparison of mid-plane Bremsstrahlung tomography and end-on imaging. Notice the apparent 'Inner Tube' instability that develops according to the tomographic reconstruction. The radial structures in the end-on imaging are believed to lie outside the mid-plane. This example uses a Fourier fit (Cormack-Granetz) method for tomographic reconstruction.

5.5 Gated CCD Spectrometer (HR4000)

We use an Ocean Optics HR4000 (200-1100 nm, 3648 elements, 25 μm slit, f/4 symmetrical crossed Czerny-Turner, 101.6 mm focal length) as a survey spectrometer covering the UV to NIR range. An optical filter is used on the detectors to remove multi-order lines (OFLV-200-1100). We expect the resolution to be 1.1 nm FWHM for our combination of slit and grating. Signal to noise ratio (SNR) is specified as 300:1 at full scale. This instrument is expected to help us identify potential contaminants in our plasmas. It can also be used to provide candidate lines for the 16 channel time resolved spectrometer and the monochrometer. It is capable of light integration times down to 10 μs with an external trigger. The delay from trigger to the start of light detection is unspecified, but assumed to be less than 1 μs.

Figure 5.5.1 highlights this spectrometer's one serious imperfection. In this figure, a fluorescent light is used to provide a broad spectrum source with significant narrow lines. By changing the integration time we can probe for nonlinearities in the spectrometer. While no significant simple nonlinearities exist, there is a disturbing effect from saturating the instrument. The effect can clearly be seen in the higher than normal signals on the long wavelength side of the lines as they saturate. We would be curious to know whether this is due to the electrical saturation of the analog to digital converter or the Toshiba (TCD1304AP) CCD itself.

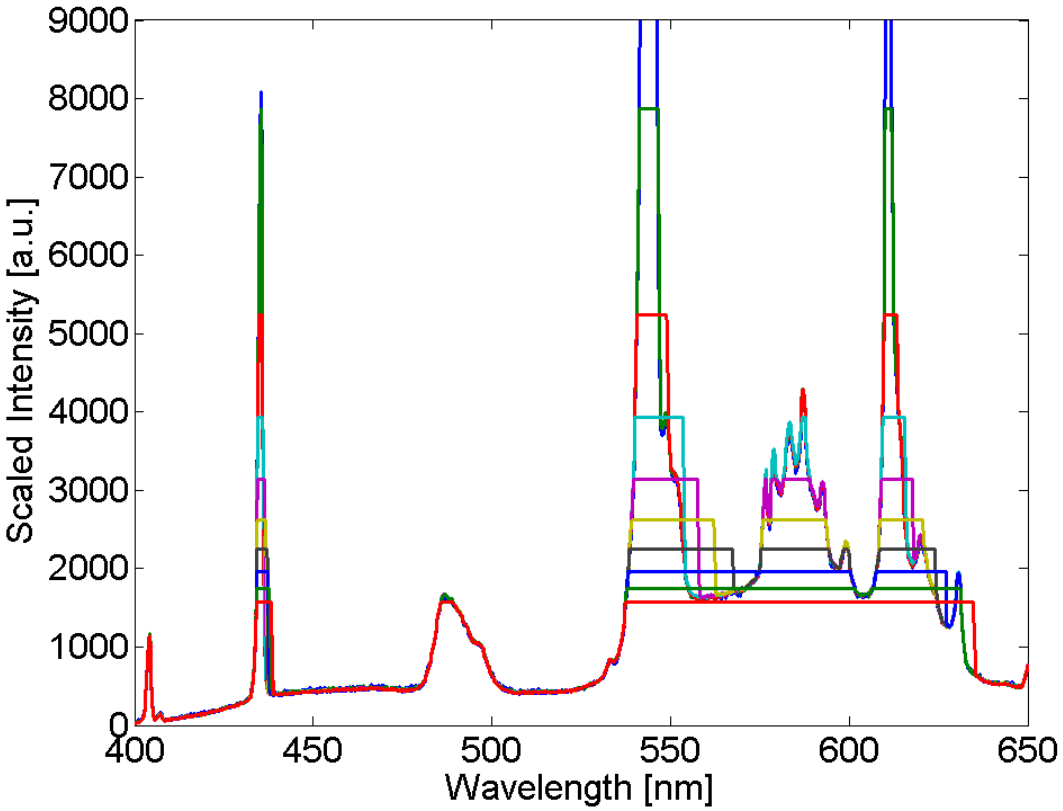

Figure 5.5.1. Line distortion occurs as diagnostic is saturated by intense lines. The recorded intensity of the light source is linearly increased by increasing the exposure time in multiples of 1/120 seconds. The recorded signal is then scaled appropriately to match original intensities.

The spectrometer has correction factors built in for sensitivity and pixel location errors. The wavelength calibration uses a third order polynomial whose constants are initially set at the factory. As with most instruments, it is recommended that it periodically be recalibrated. The polynomial coefficients are stored in the spectrometer.

The relaying optics has typically consisted of a telescope and fiber. At some times a non UV transmitting fiber was known to have been used.

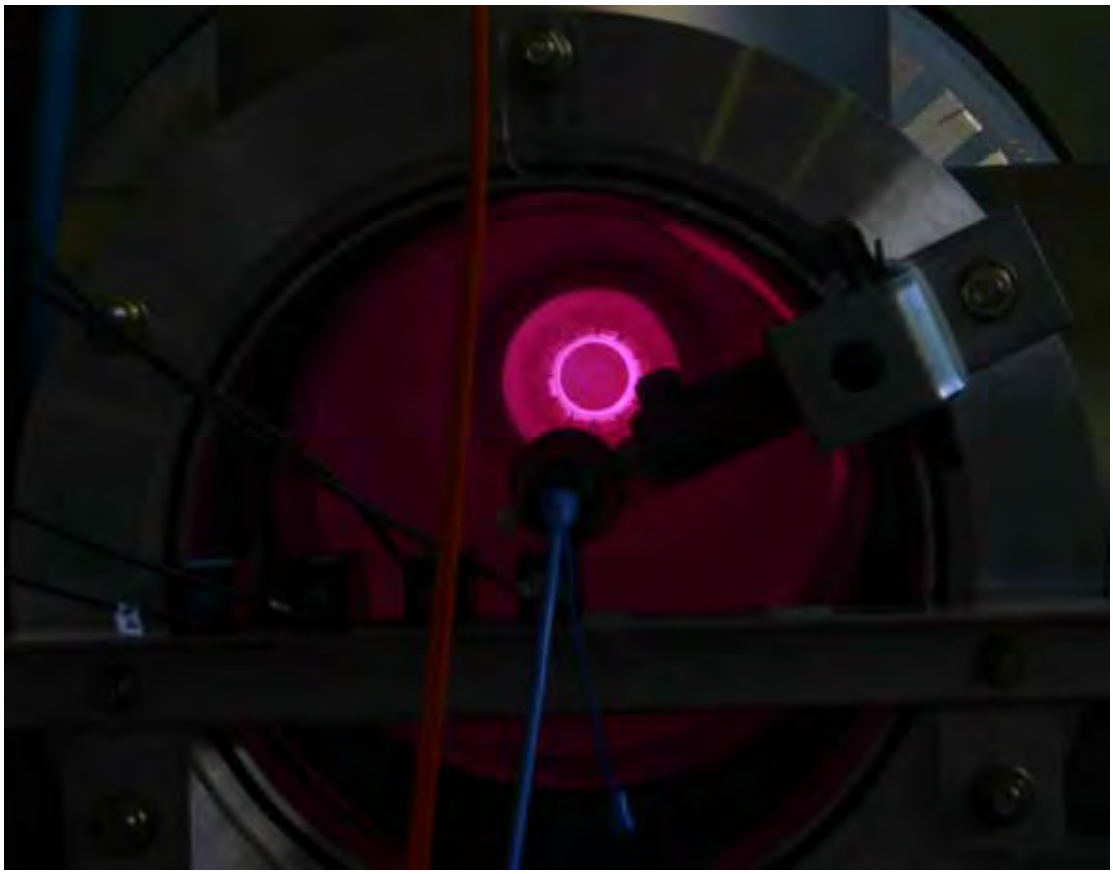

Figure 5.5.2. End on optics. The Ocean Optics spectrometer telescope and quartz (blue) fiber can be seen looking at the plasma source. Four collimators with plastic fiber can be seen mounted to look at radial locations.

5.6 16 Channel Time Resolved Spectrometer

Our high resolution diagnostic for resolving line shape is based on a SPEX Industries, Inc. 0.8m spectrometer. A 2 mm cylindrical lens is used to increase the dispersion of the instrument. A 16 channel PMT array has been fitted on the output plane to provide time resolved data. The PMT is a Hamamatsu R5900U-03-L16. Light from the experiment is coupled via lenses to the SPEX. The effective wavelength spacing of the PMT channels is 0.033 nm. Coax lines have typically been used to connect the PMT to the diagnostics rack (distances of ~6 meters) with terminations ranging from 10 kΩ to 100 kΩ. The output BNC connectors are not grounded as they exit the case. The primary objective of this diagnostic is to resolve impurity line radiation as an indicator of ion temperature and rotation. As an aside, the high frequency response and ability to significantly amplify single photons makes PMTs very useful for demonstrating the physics of improperly terminated transmission lines.

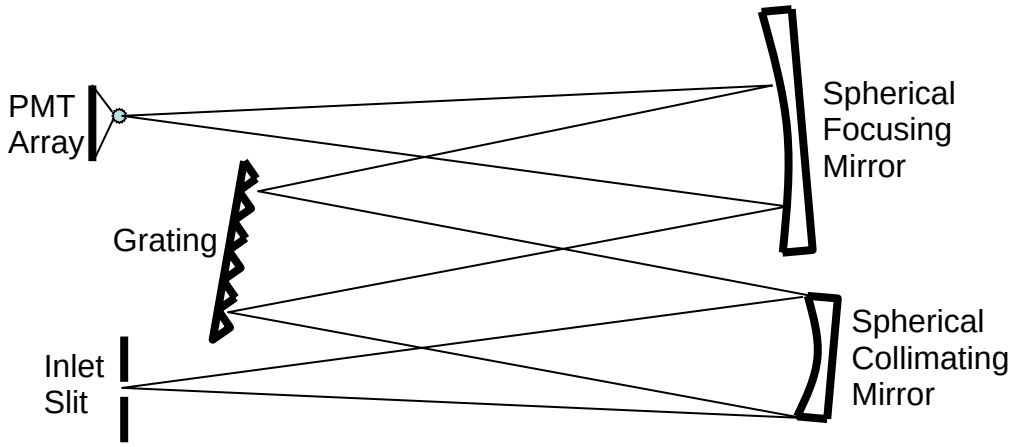

Figure 5.6.1. Czerny-Turner spectrometer diagram. Modified with a plastic cylindrical rod lens at the output.

Figure 5.6.2. Side view of 16 channel spectrometer optics bench. The basic optical instrument is configured as an 80 cm focal length Czerny-Turner spectrometer. The light inlet can be seen front left. The collimating mirror is at far right. The grating has been illuminated by the camera flash. The focusing mirror relays the lines to the 16 channel PMT (out of view behind the diffraction grating).

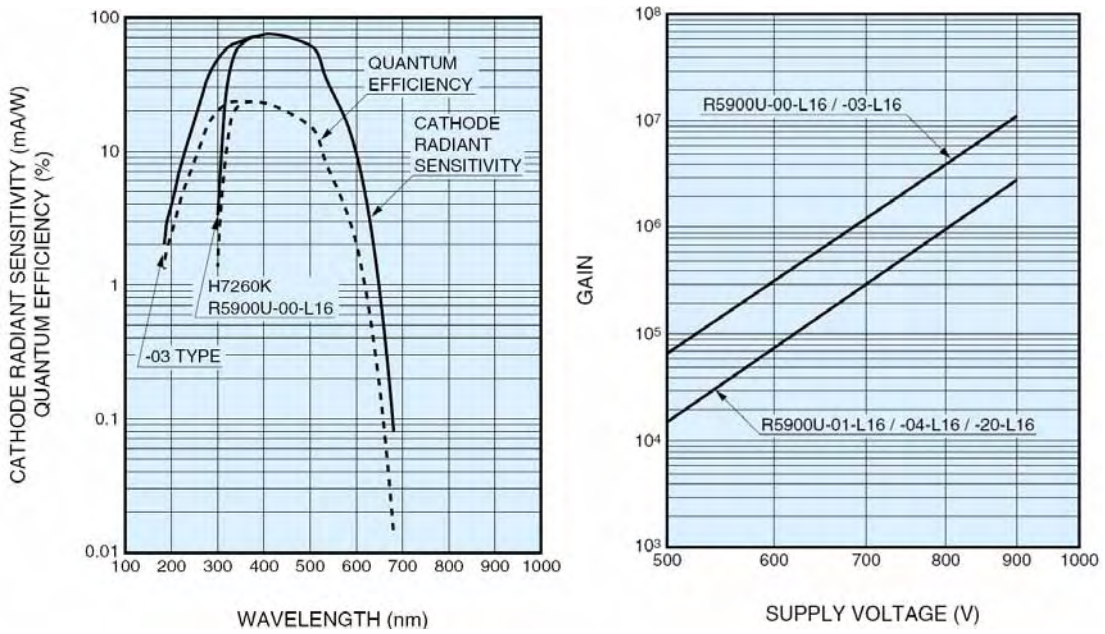

Figure 5.6.3. PMT characteristics. Reproduced from Hamamatsu R5900U-03-L16 PMT datasheet. The installed PMT has a spectral response from 185 to 650 nm. 3% optical cross-talk is expected. The PMT base has a maximum voltage rating of -900 VDC. These plots are republished with the permission of Hamamatsu Photonics Devices.

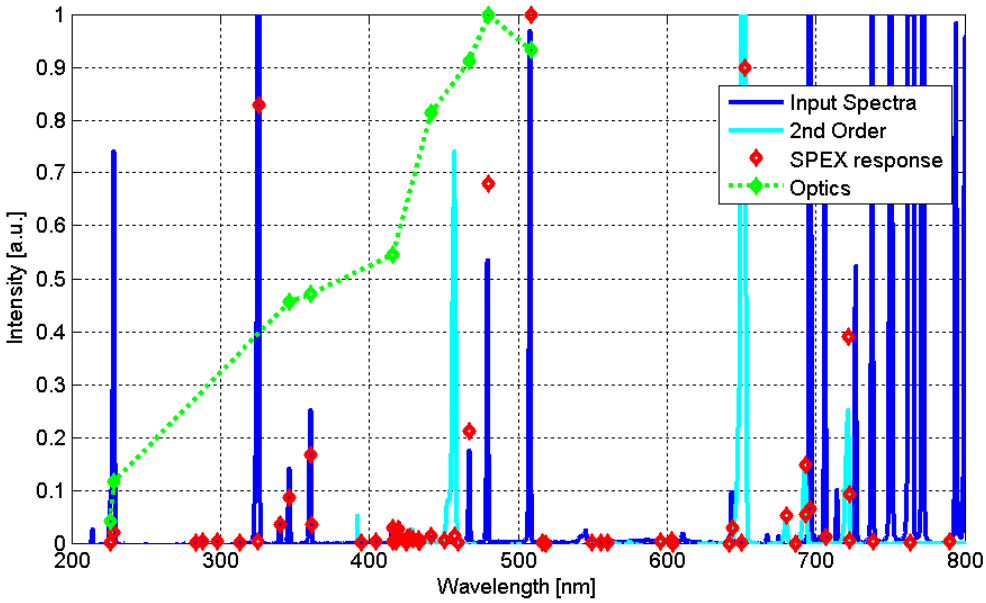

Figure 5.6.4. 16 Channel Spectrometer spectral response. The input spectra is provided by a Cd lamp as recorded by the Ocean Optics HR4000 (blue and cyan). The lines visible with the 16 channel spectrometer (-900 V) as recorded by a high impedance (10 MΩ) voltmeter (red). The response of the PMT itself can be taken out of the equation and the remaining system response (grating, mirrors, and lenses) can be found where there are unsaturated lines (green).

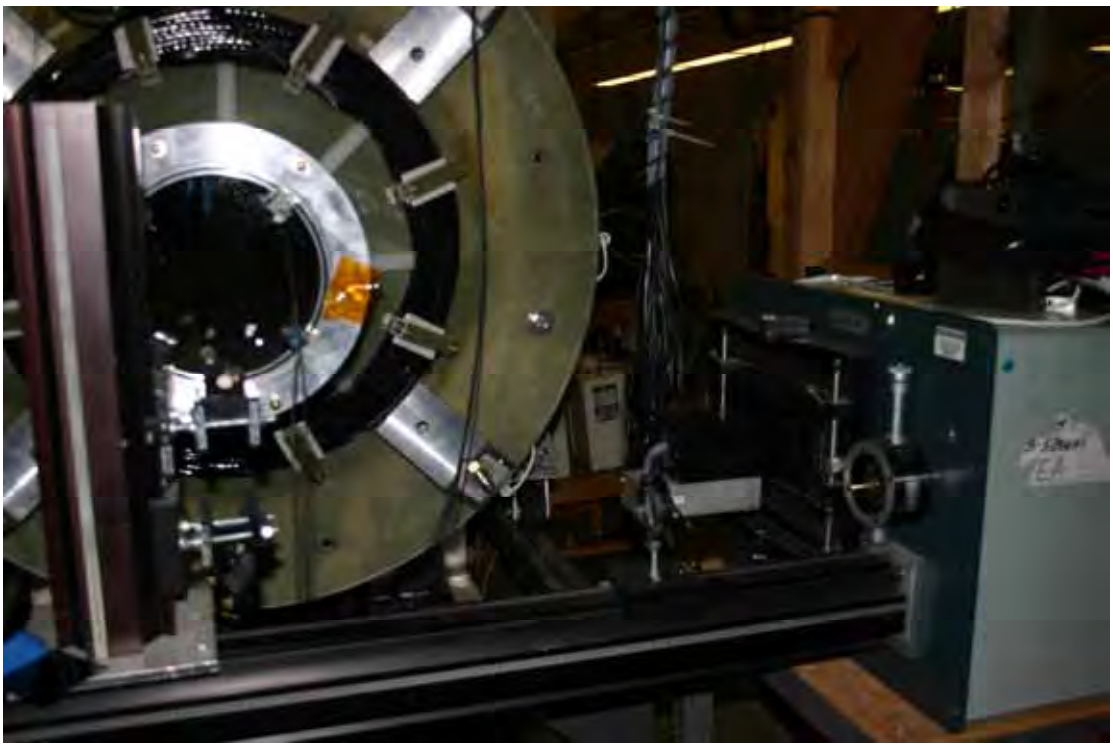

Figure 5.6.5. Spectrometer relay optics set up on experiment.

5.7 Monochrometer

A Jarrell – Ash (Division of Fisher Scientific Co. Cat #82484) 0.25 m monochrometer with PMT tube was available for observing specific line intensities with time resolution. It has a FWHM of 1 nm with a 10% off peak at approximately 1.6 nm off-center. The PMT yields ~ 300 nA of dark current at 1100 volts.

5.8 Single Frame Color CCD Camera

A 35 mm Cannon EOS Digital Rebel 300D SLR camera proved instrumental in long timescale exposures for diagnosing plasma – surface interactions as well as diagnosing aspects of machine operation. A standard 18-55 mm and 75-300 mm zoom lens were used. The camera is capable of exposures down to 250 μs at resolutions up to 3072 x 2048 with 12 bits per color (RAW format). The primary issues with this camera were timing and white balance.

The most common mistake made in using this camera as a diagnostic is leaving the camera in autofocus mode. This results in a long unpredictable delay between the command to acquire and the actual acquisition. The next most common mistake is assuming that the delay between the command and the actual start of integration is relatively constant. As can be seen from the figure below this is not the case. Keep in mind that the data used to generate these figures was for fixed orientation of the camera and acquired at relatively fixed intervals. The unfortunate result of this analysis is that the camera's flash cannot be used to establish with a high degree of accuracy when the image is acquired. I.e., we cannot set the camera to 250

μs exposures and take repeated shots to reconstruct a time history without encoding the timing in the image.

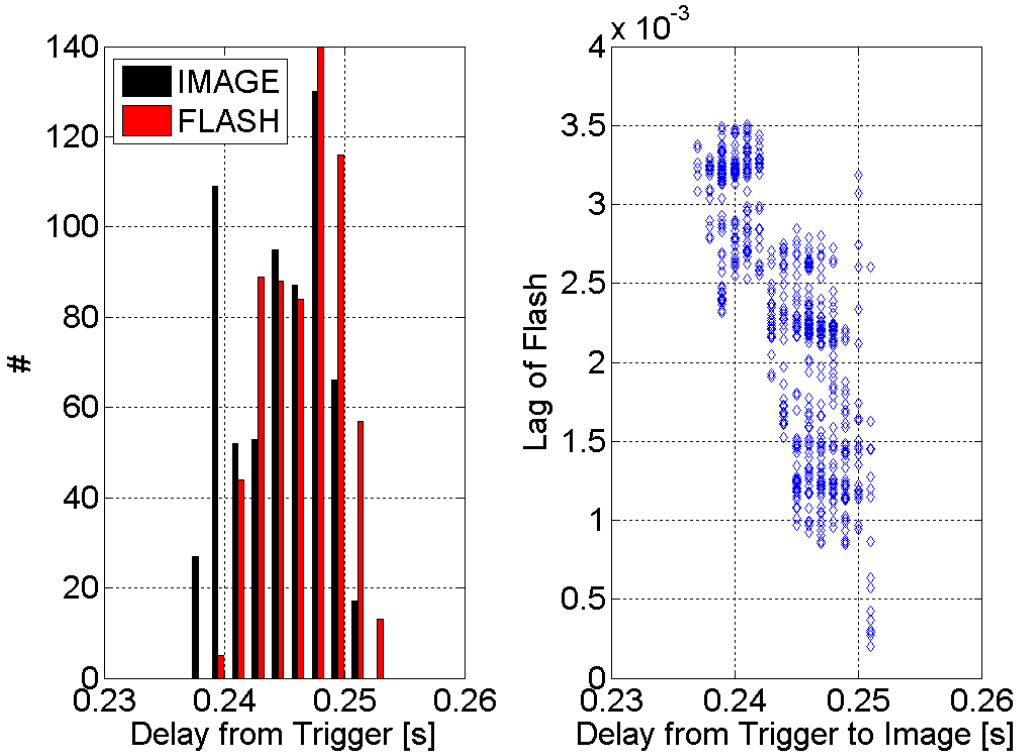

Figure 5.8.1. Canon EOS Digital Rebel 300D camera timings. The histogram at left shows the delay from triggering the camera to the actual start of image acquisition and the output from the flash adapter. The image acquisition delay tends to cluster around 244.4 ms with a standard deviation of 3.7 ms. The plot at right shows the lag of the flash signal with respect to the start of image acquisition. The delay between triggering and the start of the frame is not constant across the CCD.

Nevertheless, the utility of the camera is clear to see from the high resolution images of plasmas and arcs displayed throughout this thesis. The user should be aware that intense sources of light in one color will also eventually bleed over into the two other colors. Presumably whatever optical filter is used to resolve colors only has a sufficient extinction ratio for unsaturated images. It may be possible under carefully controlled conditions to use this effect to obtain more than 12 bits of information on one color.

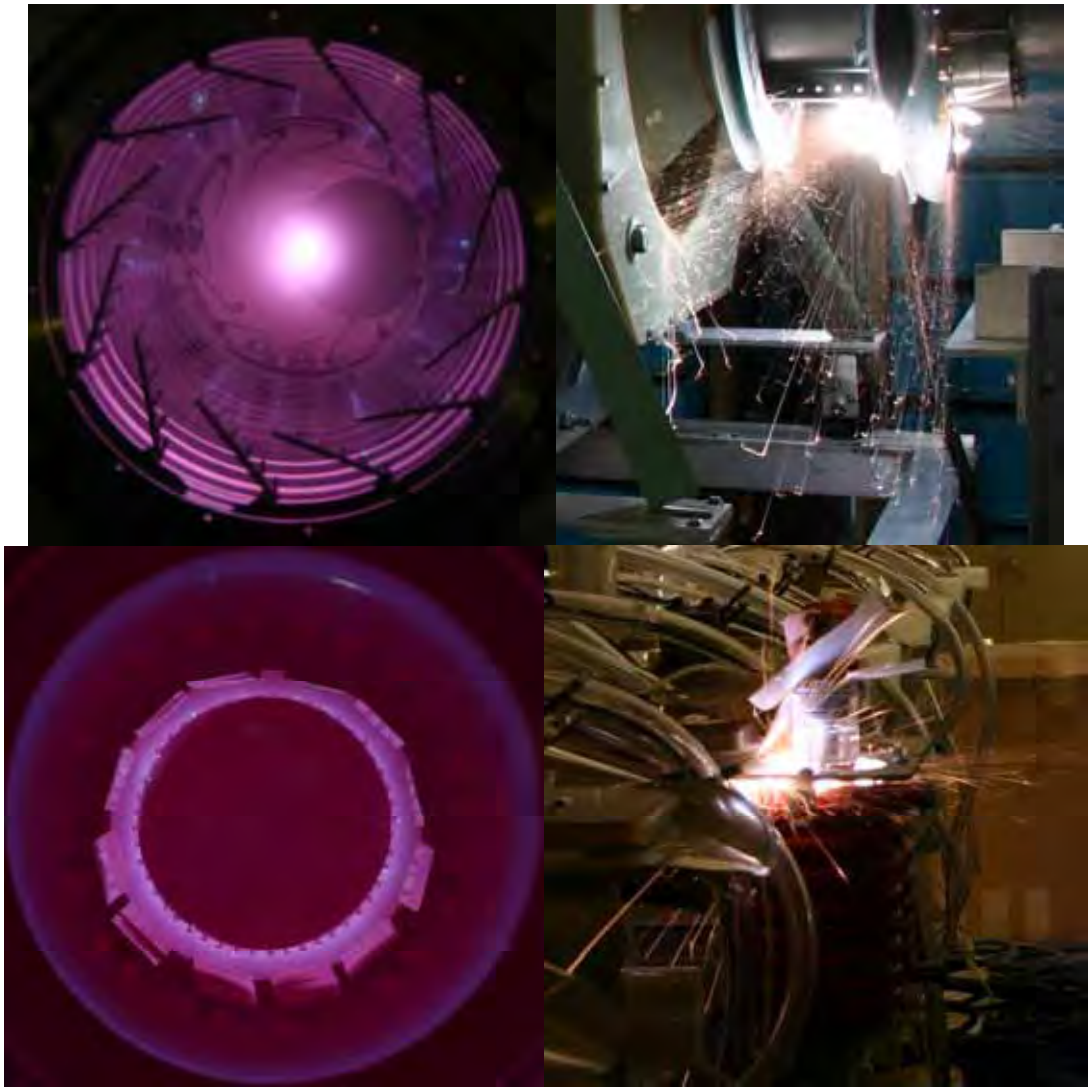

Figure 5.8.2. Sample SLR images. Upper left: Image of earlier PHD pre-axial discharge. Upper right: Magnet frame arcs due to bad electrical contact. Lower left: Plasma source behaving properly during shot (viewed with telephoto lens). Lower right: Poorly connected bank ignitron.

Figure 5.8.3. Not all pixels are acquired at the same time. The green LEDs are all on at the same times. It is presumed that near the beginning and end of capture not all the camera's pixels are acquiring, resulting in this very rare shot (exposure time: $1/60^{th}$ sec).

5.9 Fast Framing Camera

We have access to an Imacon model 790 (Hadland Photonics Limited) intensified fast framing camera. The basic operation is quite simple. Light from an experimental source impacts a 2D photocathode (behind a sapphire window) generating electrons that are focused and accelerated (by 17.5 kV) towards a phosphor screen (peak output wavelength 460 nm). The light from the phosphor screen is relayed via 5.5 micron fiber to the face that can be pressed up against film. Framing, shuttering, and focusing are accomplished with transverse electrostatic fields. Electronic shuttering is effectively accomplished by defocusing (extinction ratio $10^6 – 10^7$) after the images are captured.

From the specifications, the installed photocathode has good spectral response from 200 nm – 750 nm (10% of peak 25 mA/W @ 450 nm). The gain in total intensity is approximately 200 with variance depending on the input spectra. Much of this gain in intensity goes in to a larger output size than the photocathode however (4x area). The output has approximately 70 x 40 mm of usable area which we usually divide into 8 frames (17 x 20 mm/frame).

We most frequently use a framing module that generates 2D images arrayed two high and a variable number across generated at some time interval. A transverse staircase field is used to generate the lateral deflection. There is an aperture inside the tube with a pair of deflection electrodes on either side. When the electrostatic field

on one side is too great the electrons impact the aperture. The second pair of deflection electrodes compensates for the angle change imparted by the first pair, bringing the electrons back to their original trajectory +/- a constant deflection. The camera accomplishes these electrostatic fields by generating two high voltage damped sine waves. Initially, the two electrostatic deflectors are charged opposite each other (allowing no electrons to pass the aperture). When the trigger is received, both sides are dumped through inductors (allowing them to oscillate). Whether electrons pass through the aperture is a threshold effect (i.e., the first deflector has to have a magnitude less than a fixed value). The time it takes for the oscillator to reach zero leads to a delay from the trigger of half the inter-frame interval. The overall delay from external trigger is 70 ns + frame interval/2. The relative timing of the two oscillators is set so that on one pass near zero volts the net deflection is positive and on the next pass it is negative. This generates a frame on top and then a frame on the bottom. If the opposing deflections were linearly ramped in time the deflections would be +/- a constant during the exposure.

The system only approximates a linear ramp (from the sine wave being near zero). The magnitude of the sine wave impacts the linearity. It also determines the exposure time for each frame (currently 20% of the inter-frame interval). Each part of the frame is only exposed for ~60% of the frame exposure time. The aperture can be thought of as being swept across the image vertically. The windowing effect is caused by the aperture and the 2D nature of the electrons constituting the image. The not constant deflection is what sets the resolution at 5 – 10 line pairs/mm in the specifications (10-20 pixels/mm, 20-40 pixels/mm if fixed in position to avoid aliasing).

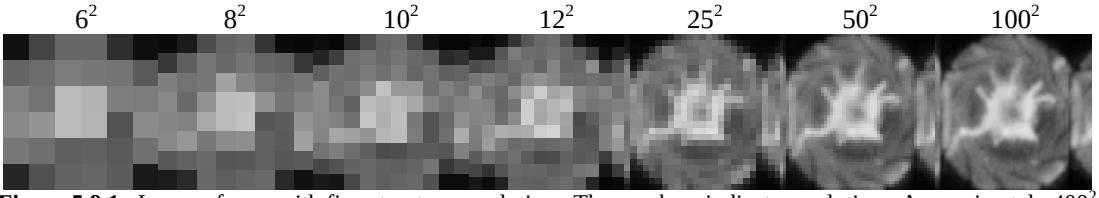

Figure 5.9.1. Imacon frame with fine structure resolution. The numbers indicate resolution. Approximately 400^2 would be the Imacon's maximum theoretical resolution.

Ideally, it would be convenient to have a fast framing system that is free of film and capable of producing as many frames as required (i.e. lifetime of the FRC). A straight forward approach to this problem could use the VLA board described earlier. From the above figure it can be seen that all the primary features are barely visible at 144 pixels with very good clarity attained at 625 pixels. With 10 MHz, 12 bit, digitizers available at $400 /channel and relatively cheap fiber, lenses and receivers for around $100 /channel the overall price is not prohibitive ($75k - $350k). At the low end of the resolution range, all the necessary electronics could fit in one CAMAC crate. While such a setup certainly would not have the high resolution of the Imacon, it would have much better temporal information. Light sensitivity issues

would also have to be examined further (particularly in regards to the sensitivity of the film currently used).

More conventional approaches would rely on CCDs, either in large chip quantities or with deep parallel shift registers. These approaches would be very dependent on what CCDs manufacturers have made available. The potential exists to be much cheaper but without more complete temporal information than that already derived from the Imacon.

It is worth noting that the tomography setup, if it were viewing end on, would see the 8^2 frame. Given this, it is not surprising that many assumptions have to be made in order to produce sharp images from tomography and that an incorrect interpretation can predominate when fine structures are present but not resolved.

5.10 Residual Gas Analyzer

We use a Stanford Research Systems 'RGA 100' for determining the composition of the base pressure in our vacuum system. It is capable of identifying partial pressures with molecular masses ranging from 1 to 100 amu. Molecules are frequently broken down during ionization resulting in a complex molecule's constituents being present in the spectra. The 'cracking fractions' are dependent on pressure (and presumably other factors) making a complete disambiguation of results difficult.

The instrument also usefully doubles as a leak checker for use with test gasses. It has been used a couple of times during the pump-down immediately after a shot to attempt a determination of what contaminants the plasma might interact with. The manufacturer does not recommend operation above 10^{-4} Torr.

5.11 Capacitance Manometer

For pressure monitoring during and immediately following shots we use a system based around the MKS Instruments© Baratron Model 622A01T. The fundamental element upon which the measurement is based is a measure of true (gas independent) pressure. The pressure being measured acts to distort a diaphragm. This diaphragm doubles as one plate in a capacitor, where the other plate is in a fixed position. The geometry change results in a small change in capacitance. The ultimate output of this commercial instrument is a 0-10V (1 Torr), < 1 mA, linearly scaled voltage. The base resolution is specified as 0.01 mTorr (0.1 mV). The time response for the pressure gauge is not specified. It is safe to assume that the response is relatively slow. We would expect it to be between 1 ms and 20 ms. To ensure safety and signal integrity we couple this signal via fiber to the data acquisition setup where it is recorded in a digital form by digitizer (inefficient but simple). This is accomplished by using a (slightly modified) voltage to frequency bank monitor to encode the pressure as a frequency. The scaling becomes 2.5 MHz/Torr. The voltage to frequency scaling is more than adequate given the resolution and time response of the device. For very small signals a vacuum offset is actually beneficial in maintaining a high frequency response.

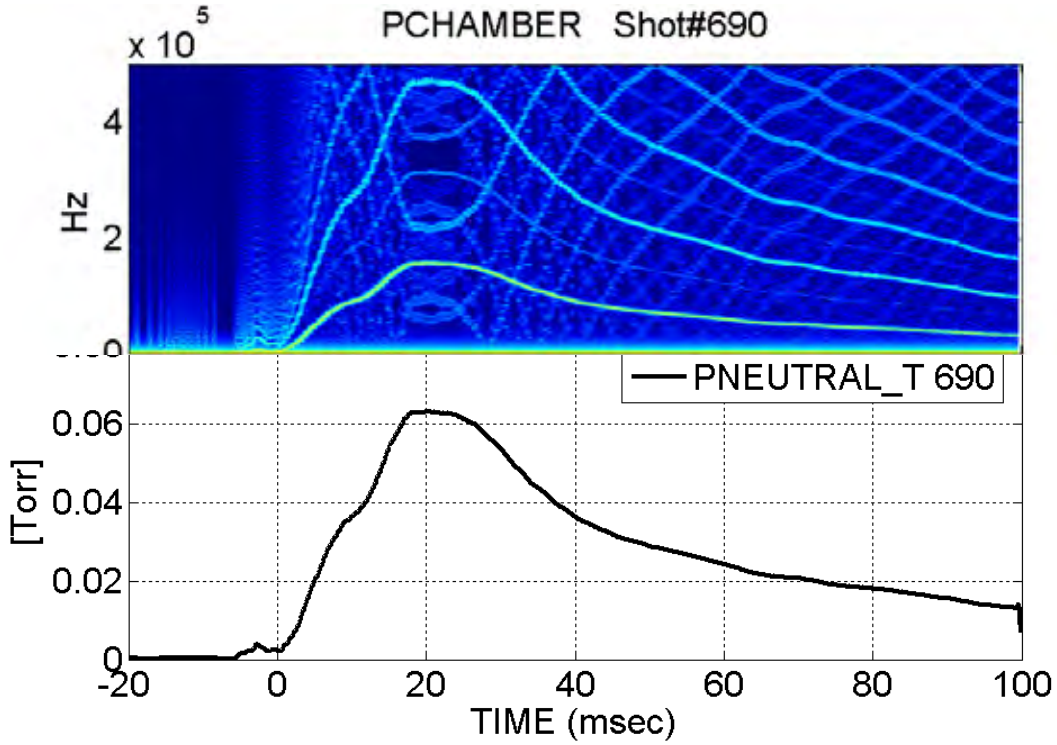

Figure 5.11.1. Typical response of pressure monitoring system to gas puff for a plasma shot. Incidentally, this figure also shows some of the utility of the *WindowedFFTGUI* routines.

This instrument has several features that should be kept in mind. The Baratron model used is not temperature stabilized. This results in a constantly shifting vacuum offset pressure. The measurement is intrinsically a strain gauge. Any coatings (for example from Ti gettering) could alter the response of the instrument. Also, the instrument is very sensitive to various forms of electrical noise. In the initial implementation the signal was contaminated with periodic spikes (20 mV) at 120 Hz and signal ramps (10 mV) at 60 Hz. The chassis of the gauge is not directly connected to the signal ground (rather they are effectively connected via 10 kΩ and 1 μF in parallel). Adding a short from signal ground to chassis ground eliminated this source of noise. The gauge and fiber optic link were being powered from a transformer isolated source that may have formed a ground loop at high frequencies. Given the long timescale required for the pressure in the vacuum chamber to stabilize, electrical noise from the experiment is not a concern. No extra shielding is installed.

The pump-down timescales for our system are 0.44 s (with cryogenic pump and turbo) and 1.8 s (with magnetically levitated turbo-molecular pump). Approximately 50 ms is required for the pressure reading to stabilize within 10% of its final injected gas value. Knowing the pump down timings permits a correction to be made in solving for the injected inventory of gas when the pumps are not closed off.

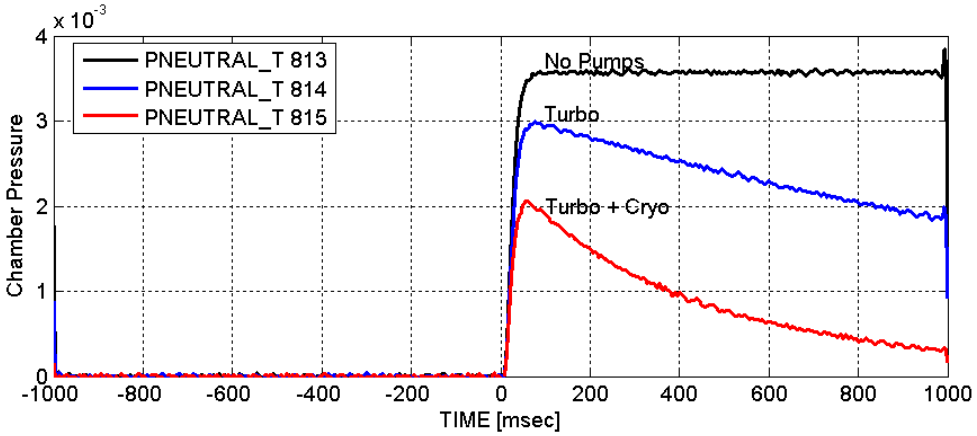

Figure 5.11.2. Effect of pumping on apparent chamber pressure. The physical location of the pressure gauge leads to a lowering of the apparent pressure due to gas flow. A simple correction for 'droop' is insufficient, a 'gain' correction is required as well. An additional 50 Hz lowpass filter has been applied (in this figure alone) to minimize noise. Here 'droop' refers to the effect of the pump pumping out the chamber (and making it difficult to determine what the peak pressure was due to the frequency response of the manometer).

Unfortunately, this diagnostic may have problems determining inventory for actual plasma shots. This instrument is a true pressure gauge. An increase in the temperature of the gas affects the pressure.

There is a less error prone way to determine total injected inventory. For this method, the gas ballast connected to the puff valve needs to be sufficiently small to register a drop in equilibrated pressure after firing. A small valve is required to prevent replenishment from the gas bottle. A pressure gauge capable of operating over a range of pressures with high accuracy is required to monitor the pressure in the ballast. While the change in pressure yields N, the system could also be used to set the pressure of the system on the fly without adjusting a mechanical regulator (a dump to the roughing pump would have to be added to decrease the pressure). The pressure could then be recorded and controlled by the bank monitoring system. Knowing injected inventory would be an almost automatic side benefit once the volume of the ballast is determined.

5.12 Miscellaneous

Many additional diagnostics exist for measuring various electrical parameters. The most reliable have been based on Stangenes Industries Inc. current probes. Fair warning, some of the probes have non-constant with changing frequency calibration factors at 100s of KHz frequencies. These probes are also very convenient for measuring pulsed voltages by the use of high power resistors. These probes are generally not noise free.

Localized multi-turn magnetic pick-up loops exist at the mid-plane to measure B_θ due to the two phases of RMF. These simple loops frequently show signals not related to the RMF. Integration, where necessary, has been accomplished numerically for these signals.

Fifty ohm attenuators and terminators are only capable of taking a limited amount of pulsed power before partial failures develop. These failures routinely result in a change in the level of attenuation rather than complete failure.

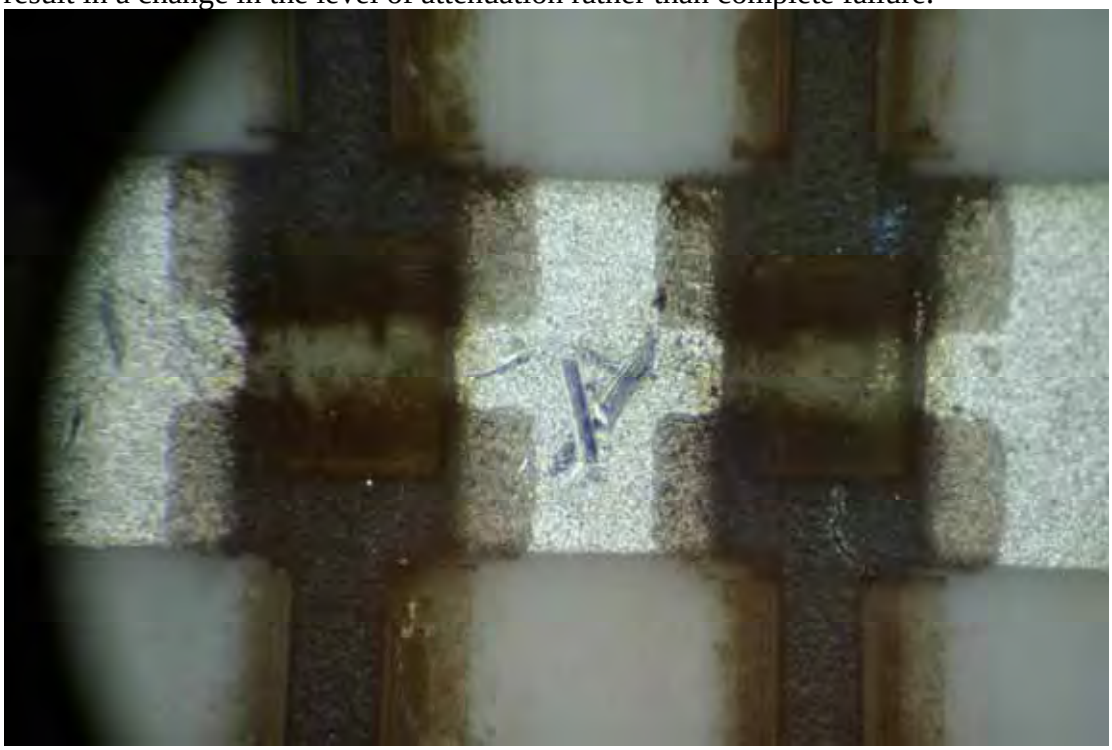

Figure 5.12.1. Damaged 5x coaxial BNC attenuator (2W). A crack due to excessive power input (100 volt levels, pulsed RMF monitoring) is visible in the lower right corner. Centers feed in on left and right. Outer ground shell attaches top and bottom.

6
Analysis

The parameters of greatest interest to us are poloidal flux φ_p, ion temperature T_i, density n, and energy in the FRC U_{FRC}. The diagnostics we have available are insufficient to the task of ascertaining these quantities without certain key assumptions. All of our diagnostics are external to the plasma. The use of no internal probes is predicated on the expectation of reaching high energy densities. Such probes would either contaminate results and/or become damaged in very short order[86]. Thus, we have no internal magnetic probes to unambiguously determine φ_p or other internal probes to determine any other localized values. The critical factor to understand is that the overall experimental objective is to attain high $n \tau T$ via translation and compression of a compact toroid. Money and time spent on diagnosing the first phase (formation) is largely wasted except where it leads to better trapped poloidal flux and avoidance of instabilities. Given the understanding gained from other experiments, this phase of experimentation is largely a matter of machine optimization not requiring overly accurate values, but instead simply whether one way of running the experiment is better than another. As such the initial diagnostic set is minimal (and the analysis correspondingly more labored).

We assume that our plasmas come to some axial pressure equilibrium when they survive for sufficient lengths of time. Similarly, we assume that the internal profiles are due to rigid rotor electron rotation. Ion temperatures are assumed to be dominant and spatially constant. Up to this point, the assumptions are expected to be of very high quality. Toroidal fields, plasma flows, and FRC end effects are assumed negligible where and when the measurements are made. Some effort is undertaken to ensure that these assumptions are in effect for displayed data (but there is no ironclad guarantee). At a much later date we may be able to translate into a proper flux conserver. Thomson scattering would also be useful in confirming some parameters directly for comparison against those obtained with the assumptions.

Our need for plasma species data is primarily governed by practical concerns. The chief amongst these is the need to avoid plating conductive materials onto normally non-conducting surfaces (or processing wall material compounds into those that conduct). As such this data does not generally get displayed (and is concentrated on effects due to the cumulative effect of the shot, not just some small length of time).

6.1 Lift Off Trapped Flux

The initial flux available for the FRC is determined by the bias field and initial distribution of plasma. A critical parameter is the amount of flux trapped within the null of the FRC when the plasma sheath lifts off the wall. This is difficult to accurately measure for several reasons. Radial pressure balance can not be blindly applied due to the inertia of the plasma.

Figure 6.1.1 displays the available relevant diagnostic signals. For this shot, we can see that early in time 50 mWb of bias flux is available within the flux loop at

the mid-plane (black trace). When the external magnetic field is decreased to zero, around 13 mWb of flux (and excluded flux) is still inside. In the standard view, the pressure of the internal field is being supported by the plasma. The plasma is in turn being supported by the physical wall, mediated by a plasma sheath layer. The vast majority of the flux has not been trapped (possibly due to a poor initial plasma distribution). At this stage, the change in excluded flux is very rapid. The crosses are located one time step apart. Any uncompensated phase error between field and flux results in a large error of the measurement. We linearly interpolate between time steps to calculate trapped flux at this stage (Lift-Off Flux from B_z=0). Each external magnetic probe pair uses the same digitizer to ensure synchronous digitization. Measuring trapped flux in this way should largely avoid the need for the corrections due to geometry discussed in the diagnostic section.

The separatrix radius appears to be pushed off the wall by approximately 4 μs. Note that if the excluded flux at this time were used to estimate trapped flux, the result would be around 20 mWb (higher than the trapped flux with no external field). We can only justify the continual decay of trapped poloidal flux, not its increase. The explanation for this error is inertia. The excluded flux can be thought of as the sum of three effects. These effects are due to internal magnetic fields, internal plasma pressure, and plasma inertia. For the separatrix to be at the wall, the plasma (near the field null) has to have been pushed in from the wall. Inertia opposes this, resulting in a higher external field than would be necessary if the plasma were massless (in turn resulting in a higher apparent excluded flux, and apparent trapped flux). Overestimation of internal poloidal flux is the result.

The reverse problem occurs as the separatrix radius reaches a minimum. Underestimation of poloidal flux occurs. Typically, the mean of these maximum and minimum values is computed to approximate the trapped poloidal flux (Lift-Off Flux from Radial Oscillations). Of course, the internal plasma pressure can distort the result (resulting in a higher apparent trapped flux than was apparent when the plasma sheath was supporting all the pressure on the wall). There is also no guarantee that the proper approach is indeed a simple average of the two measurements.

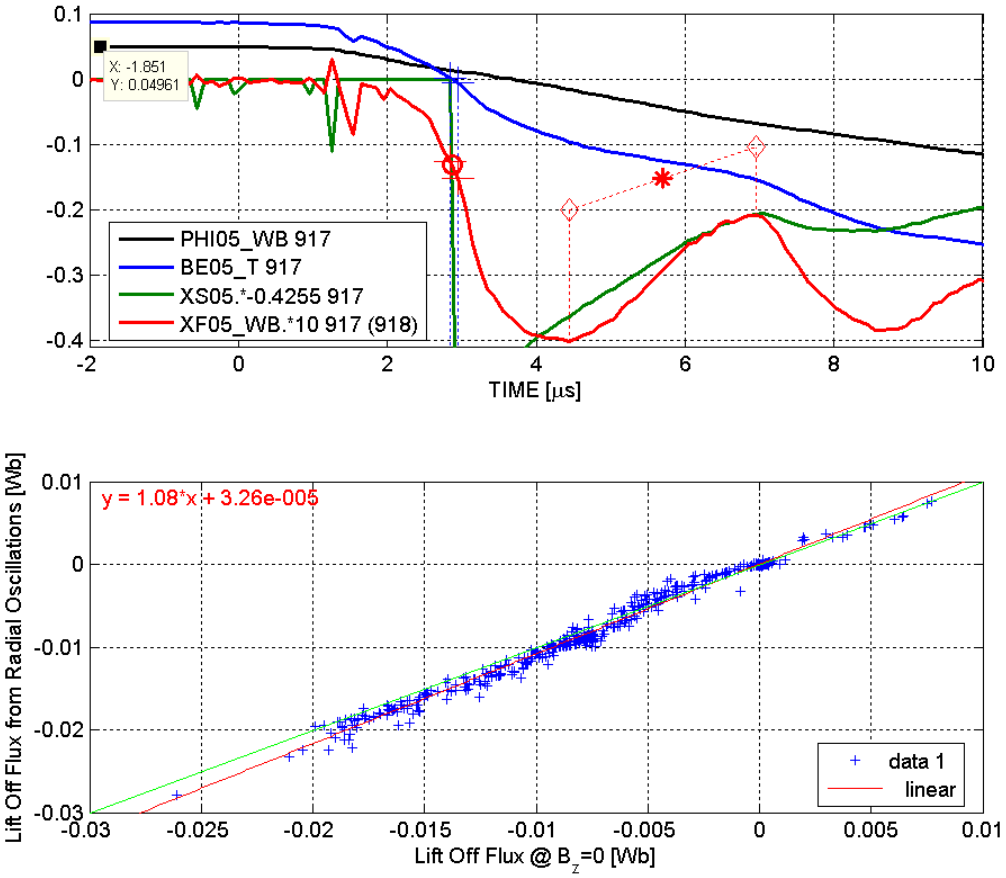

Figure 6.1.1. Calculation and comparison of Lift-Off flux determination methods. (top) The diagnostics at the mid-plane yield these traces. Axial flux within the driven strap is black. External axial azimuthally averaged magnetic field is blue. The derived separatrix radius is shown in green. Excluded flux (with a vacuum reference shot subtracted) is shown amplified by a factor of 10 as red. The cross hairs indicate the steps bracketing external magnetic field equal to zero. The red circle indicates the lift-off flux inferred from when the sheath is supporting all the pressure on the wall. The red star displays the lift of flux inferred from the radial oscillations. There are indications that the event visible in the blue and red traces at 1.4 μs (43 mWb) may be due to electrical noise; it might also be due to plasma striking the wall and momentarily forming a sheath. (bottom) A comparison of the lift off flux computed via the two methods is displayed for all 80 cm experiment shots. The shots with positive flux are second half cycle (reversal bank boosted bias field) shots. A red line is fitted to the results indicating approximately 8% higher radial oscillation inferred lift-off flux than external magnetic field equals zero inferred lift-off flux.

It should be apparent that both methods have potential for errors due to questionable assumptions. However, as a relative measure where the differences are expected to be large (say the difference between 50 mWb of trapped flux and 15 mWb), this measurement is more than adequate.

6.2 Axial Equilibrium

All of our available diagnostics are external. Our magnetic diagnostics are mounted on the vacuum boundary outside of the plasma. Light emissions and interferometer results are accumulated results that are detected outside the plasma.

They are thus a reflection of conditions for some volume of plasma. We do not have any probes that are capable of measuring the parameters of a small volume of the plasma. We do not have a Thomson scattering diagnostic. Most of the parameters we are interested in observing and comparing are internal. The net parameters of greatest interest are internal poloidal magnetic flux φ_p (including associated lifetime), and the average ratio of plasma pressure to external magnetic field $<\beta>$. The bridge between what we can measure and what we want to know has to be spanned by theory.

Theory inevitably requires assumptions. We cannot prove the validity of all the necessary assumptions. All we can do is point to similar experiments which make the same assumptions (in some cases backed up by experimental measurements). Even when similar experiments have proved their assumptions via diagnostics this still requires the assumption that our experiment is sufficiently similar for those assumptions to still hold true.

The most critical assumptions for simplified analysis relate to the profile of local plasma conditions across the configuration. If we assume that total temperature is constant inside the separatrix and that the electrons have an average motion of rotating rigidly at some fixed frequency about the geometric axis then a simple self consistent profile can be found. This rigid rotor (RR) profile is defined by $n_e = n_m \text{Sech}^2(K[(r/R)^2-1])$, $B = B_e \text{Tanh}(K[(r/R)^2-1])$, and $V_e = \Omega r$. The variables are n_e (local electron density), n_m (maximum electron density), K (profile peaking), r (local radius), R (radius of magnetic null), B (local axial magnetic field), B_e (axial magnetic field at edge), V_e (electron velocity), Ω (electron rotation rate). There is assumed to be no significant ion rotational momentum or other ion flows. It is also assumed that there is no toroidal magnetic flux consistent with the definition for FRCs. All relevant physics are assumed to be fluid-like for the purposes of the analysis. Obviously, we expect some of these assumptions to not be strictly true, but rather close enough to true to result in small to moderate errors.

Integrating the axial magnetic field profile to find flux we derive $\varphi_p = \pi r_s^2 B_e \text{Ln}(\text{Cosh}(K))/2K = XF \text{Ln}(\text{Cosh}(K))/2K$, where XF is excluded flux. Average beta ($<\beta>$) is defined as the ratio of plasma pressure to external field pressure averaged over the volume of the closed field line plasma. From the rigid rotor definitions taking the mid-plane cross-section for averaging: $<\beta> = \text{Tanh}(K)/K$. We want $\text{Ln}(\text{Cosh}(K))/2K$ as a function of X_s. An FRC axial pressure equilibrium relation of average beta (ratio of plasma pressure to external field pressure) to x_s is $<\beta> = 1- ½ x_s^2$, providing the last link to define K in terms of measured parameters. The derivation of average β's relationship to x_s depends on the existence of a long flux-conserving shell. The relation will break down for very low confinement times and very small values of the gyro-kinetic parameter s (discussed later)[87]. It is derivable from considering the magnetic force balance of the end region beyond the FRC versus the mid-plane (both where the magnetic field components are only present in the axial direction). It also requires that the relevant x_s be taken where there is no field line curvature. The immediate implication is that for the $<\beta>$ condition to hold true, the FRC must be long enough to be able to increase its internal plasma pressure

(in response to increased external field) via axial contraction. Solutions to these equations are shown in figure 6.2.1 below. A good approximation for poloidal flux is $\varphi_p = 0.31\ x_s\ XF$, where x_s is the separatrix radius normalized by coil radius and XF is the excluded flux as described in the diagnostic chapter.

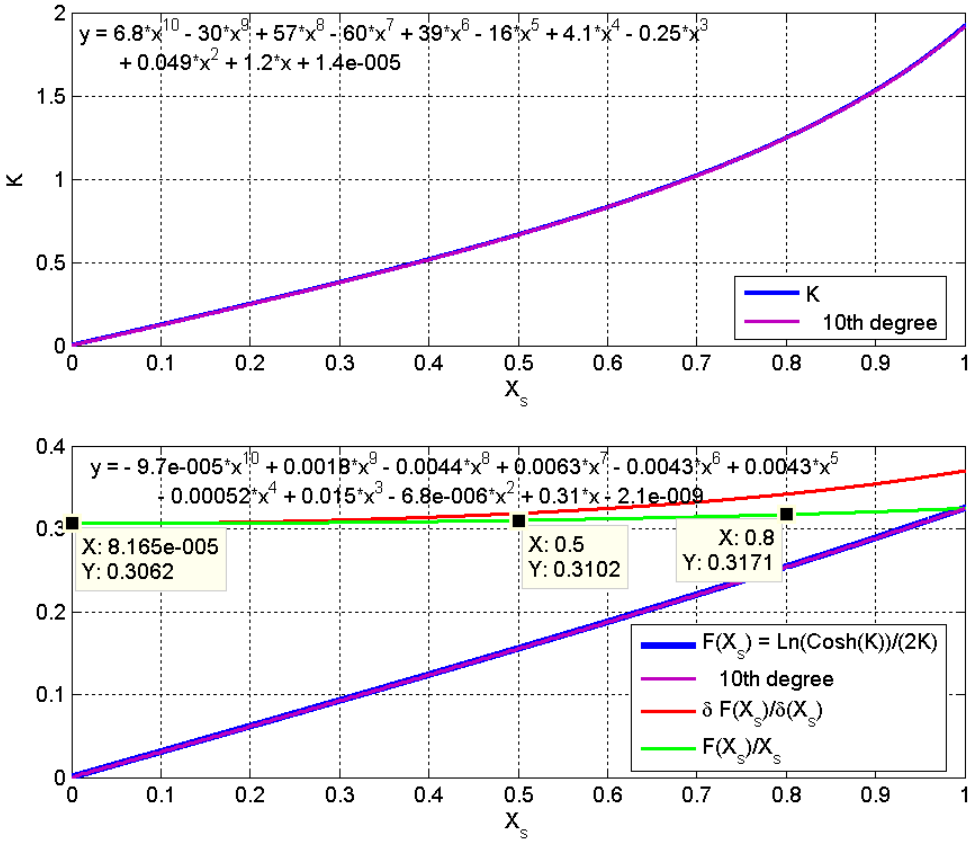

Figure 6.2.1. Rigid rotor dependence on X_s in equilibrium. (top) The profile varies significantly as a function of normalized separatrix radius. Obviously, more terms are presented than required for a good polynomial fit. (bottom) From this numerical solution we can see that $\varphi_p = 0.31\ X_s\ XF$ for a rigid rotor profile in axial equilibrium.

Similarly, given the rigid rotor assumption and the average beta relation to x_s, it is straight forward to find peak density and thus total temperature from the available diagnostics. Using the rigid rotor assumption for density we have plotted in Figure 6.2.2 the density as a function of radius (normalized to the magnetic null) and the rigid rotor profile constant K. Integrating along the sight-path of an interferometer beam passing radially through the axis of the FRC we can find the ratio of peak to average density as a function of K or x_s as below. Using radial pressure balance $k\ T_t\ n_m = B_e^2/2\ \mu_0$ it is then a simple matter to find total temperature T_t. With large error bars we may as well assume that $T_i \approx T_t$ from previous field reversed theta pinch (FRTP) formation experiments that show strongly preferential ion heating under certain conditions.

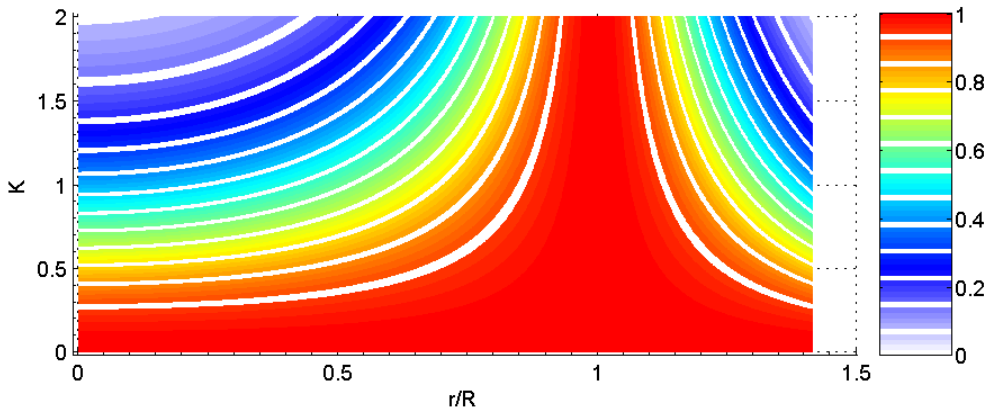

Figure 6.2.2. Density profile as determined by profile constant K. The implication of low K (and low X_s in equilibrium) is high density adjacent to the open field lines.

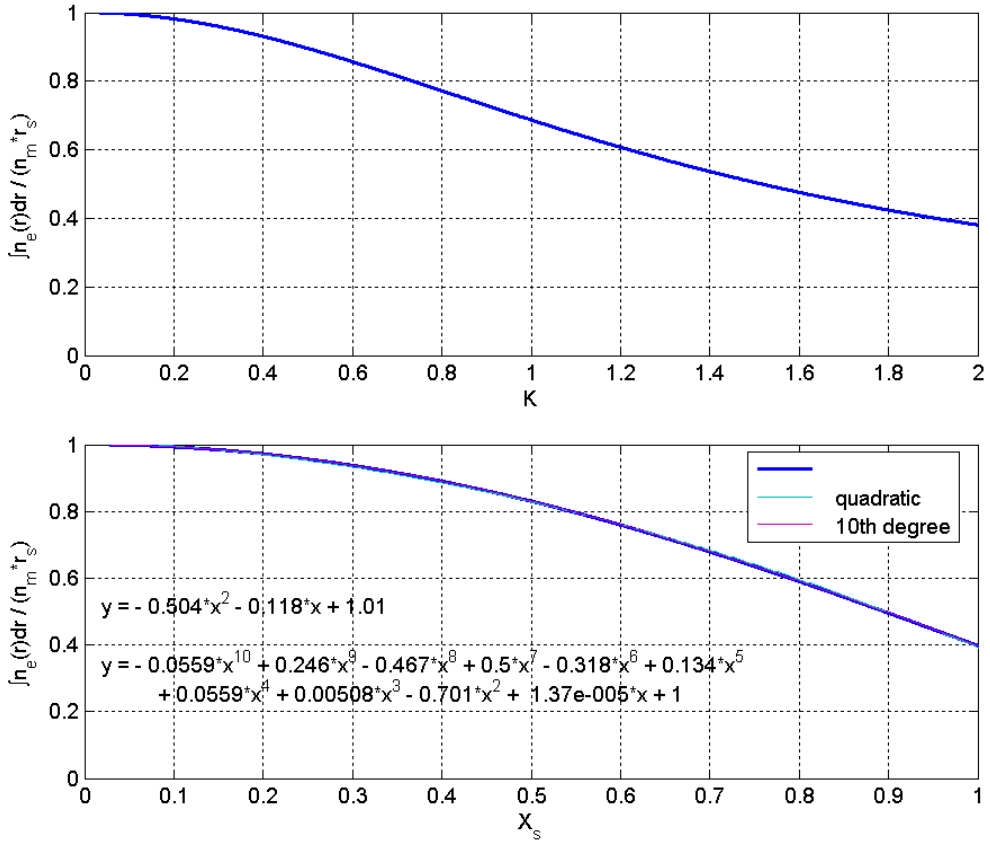

Figure 6.2.3. Relationship between average density seen by interferometer and peak density. As can be seen, the peak density varies from 1 to 2.5 times the average density (as would be determined from the interferometer) with a strong dependence on profile constant (and hence X_s in equilibrium).

We might expect that the initial distribution of plasma will not match the rigid rotor profile immediately after reversal. Also, the average beta relation to x_s will not be applicable until the FRC comes into axial equilibrium; this will be well after reversal and is expected to be governed by Alfvénic timescales. So, some time must

pass during which it is assumed that the profiles will relax into the rigid rotor profile. Once the axial contraction has been completed we assume that the simple relations will hold true. This frequently leaves little (if any) time for determining poloidal flux and its decay rate. Many of these problems occur because this experiment was designed to form and translate FRCs immediately. Axial translational stability is not required for this purpose. Additional stages for acceleration and capture were to be added sometime later upon funding availability. Funding for proper diagnostics and analysis in this more complex phase is simply unavailable. Also, running the ends at sufficient current levels to result in capture would require running the central sections at lower E_θ (longer times on the wall during reversal). Significant times spent with the plasma on the wall are expected to cause more damage at higher trapped flux levels.

6.2.1 80 cm Experiment

A more troubling problem is encountered due to the assumption in the average β relation to x_s that a constant flux conserving shell is used well beyond the ends of the FRC. For most of the stationary shots that we have run (on the 80 cm experiment) this is true to varying extents as long as the FRC isn't long enough to reach the final end coils. In general, there is no designed assurance that this should be the case. Individual axial coils are fired independently. The fields that they are able to obtain are also independent. For the end coils to have as much flux as the central coils would require much higher currents. For stationary shots, the flux on the last end coils is ~25% lower than at the center. Also, running the ends at higher fields would quickly run afoul of violating the assumed equilibrium condition by putting extra magnetic pressure on the end regions.

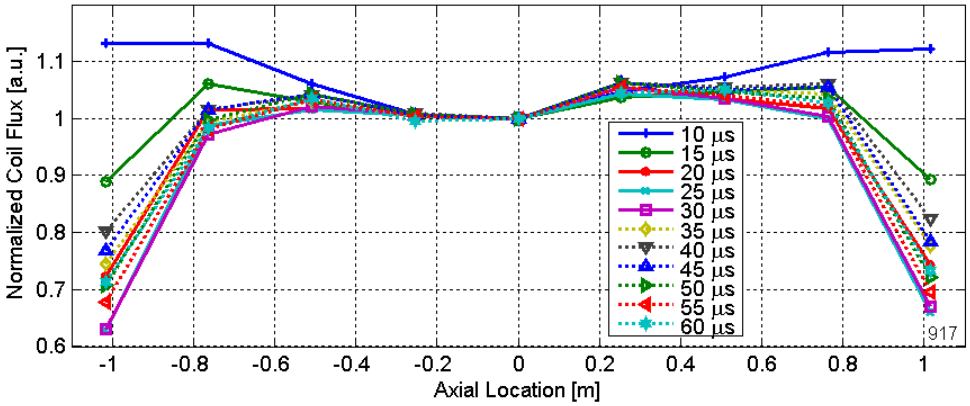

Figure 6.2.1.1. Axial driven coil flux profiles. Deviation from unity indicates deviation from a perfect axially uniform flux conserver. The extra current flowing in some coils due to the presence of the FRC probably results in some of the deviations from axial symmetry.

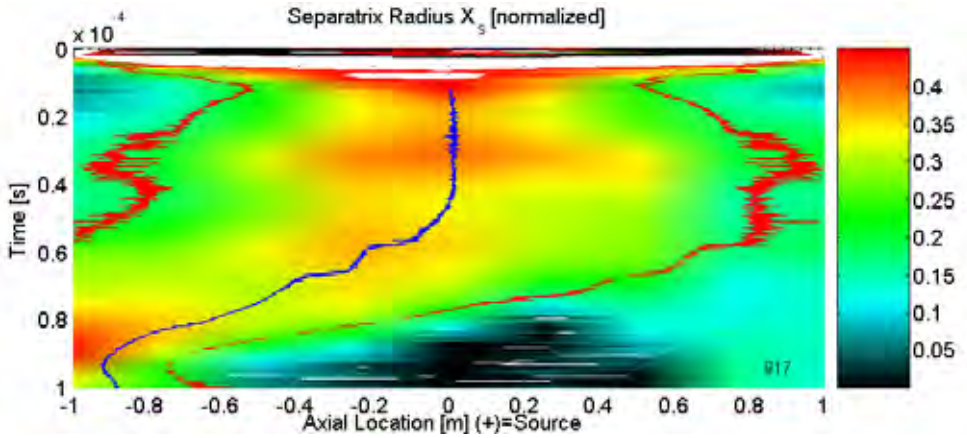

Figure 6.2.1.2. Axial Position of FRC during stationary shot. The blue line shows the axial centroid of the FRC. Red shows the location of half peak X_s. The FRC ends are perilously close to the drop off of the end coil flux throughout its equilibrated existence.

With all these complications, there can be little assurance as to the accuracy of derived quantities while using the required simplifying equilibrium assumptions. This will be obviously true for the 'dynamically formed' FRCs; and likely true for those with stationary timings as well. A more complex approach would be to feed the diagnostic information into a numerical physics model (for example MOQUI or NIMROD). This could relieve many of the analysis quandaries but the initial data to appropriately initialize the model would be difficult to obtain. Recall that there is no expectation of the axially introduced initial plasma obtaining a uniform (or even a known) distribution in the formation chamber. By necessity, an iterative technique would have to be employed varying the initial parameters (and various physics parameters such as resistivity profiles) with some reasonable limitations based on theory. Unfortunately, this is certainly a non-trivial undertaking fraught with its own perils that will have to wait.

Given the long run times for accurate simulations, it is reasonable to compare models with human-chosen conditions to experiment. However, a great deal of care may be required to prevent the unintentional introduction of observer bias (meant as it is used in the social sciences and generally) to the analysis. If the analysis could be automated, any bias could be made repeatable and thus removable (strengthening comparisons). The effects of errors in measurement could also be more easily quantized.

6.2.2 40 cm Experiment

The 40 cm experiment had the majority of the axial field coils linked together into large segments helping to assure that the flux on all segments would be similar. However, the stray inductance of the coils and interconnects was not insignificant. The 40 cm experiment was never intentionally run so as to produce a non-constant (spatially) flux coil.

Overall, the 40 cm experiment was less fruitful in studying equilibrated FRCs due to rapid destruction of the configurations by instabilities. If this phase of the

overall project were not heavily focused on obtaining reactor relevant initial formation hardware it might be worthy of further study of the particulars of the instabilities. However, as things stand, analysis of these data sets will get the short end of the stick.

6.3 Poloidal Flux Lifetime τ_φ

We expect the rate of flux loss to be proportional to the amount of flux. This naturally leads to exponential decay of poloidal flux (i.e. $\varphi(t) = \varphi_0 e^{-t/\tau}$ & $\partial\varphi(t) = (-\partial t/\tau)\varphi_0 e^{-t/\tau}$). The result is a simple relation for determining lifetime $\varphi(t) / \partial\varphi(t)/\partial t = -\tau$. Another valid approach is to take the natural log of poloidal flux as $\ln(\varphi(t)) = \ln(\varphi_0) - t/\tau$ and linearly fit to the resultant slope over some period of time. This second approach has the side-effect of frequency filtering the original $\varphi(t)$ data while the first approach (as a practical matter) requires an explicit filtering of $\varphi(t)$, not a filter of $\tau(t)$.

τ is likely capable of changing as a function of time. As noted earlier $\varphi_p = 0.31\ x_s\ XF = (0.31)\ 2\pi R^2\ x_s\ B_e$. A change in flux of the FRC must be caused by an azimuthal electric field at the FRC null from Faraday's law $\nabla \times \vec{E} = -\frac{\partial \vec{B}}{\partial t}$ transformed by $\int_S dS \bullet \nabla \times () = \oint_C dl \bullet ()$ into $\oint_{B\_null} E_r = -\int_{null\_surface} dS \frac{\partial B_z}{\partial t}$; simplifying finally to $2\pi R E_\theta = -\frac{\partial \varphi}{\partial t}$ with the assumption of rotational symmetry. The generalized Ohm's law: $\vec{E} + \vec{V} \times \vec{B} = \eta \vec{J} + \frac{1}{ne}(\vec{J} \times \vec{B} - \nabla P_e)$, which taken at the magnetic null generates $E_\theta = \eta J_\theta$ allows only for the continual dissipation of flux. Neglecting displacement currents, Ampere's law: $\nabla \times \vec{B} = \mu_0 \vec{J}$ provides $J_\theta = -\frac{1}{\mu_0}\frac{\partial B_z}{\partial r}$ assuming that there are no radial magnetic fields (consistent with being at the magnetic null). For simplicity, we can take the assumed rigid rotor profile solution to realize that $J_\theta = -\frac{2B_e K}{\mu_0 R}$. So change in flux : $\frac{\partial \varphi}{\partial t} = -2\pi R E_\theta = -2\pi R \eta J_\theta = 4\pi \frac{\eta}{\mu_0} B_e K$. Note that the flux within the null from the assumed rigid rotor profile is of negative sign relative to B_e. If we use our first approach to define lifetime we find that $\tau_\varphi = \frac{\varphi(t)}{-\partial\varphi(t)/\partial t} = \frac{(.31)2\pi R^2 x_s B_e}{4\pi \frac{\eta}{\mu_0} B_e K} = .155 \frac{x_s}{K}\frac{\mu_0}{\eta} R^2$. Lifetime will have no hope of being constant if R is not constant. As R increases, we expect lifetime to increase (if flux loss were governed by resistive losses and not instabilities).

Taking this one step further for future use, we can insert transverse Spitzer resistivity. Also, realize that x_s/K only varies from 0.82 to 0.52 (we'll approximate as ¾ consistent with $x_s = 0.5$). From this we find that classical flux lifetime is $\tau_\varphi = (0.00142)\frac{R_{[m]}^{\ 2} T_{e[eV]}^{\ 3/2}}{\ln\Lambda}_{[\sec]}$ where the remaining variables should be evaluated at the magnetic null. It is a simple matter to choose a different (less optimistic) model for resistivity to derive different scalings.

If τ_φ is allowed to vary in time, the solution of $F'(t) = -F(t)/\tau_\varphi$ is no longer valid. Instead, the solution to $F'[t] = -F(t)/\tau_\varphi(t)$ must be sought. The new solution becomes $\varphi(t) = \varphi_0 e^{\int_0^t -\frac{1}{\tau(s)} ds}$. Taken locally, $\varphi(t) / \partial\varphi(t)/\partial t = -\tau(t)$ is still true. Taking the natural logarithm of this new form yields $Ln(\varphi(t)) = Ln(\varphi_0) + \int_0^t \frac{-1}{\tau_\varphi(s)} ds$. It is obviously possible to still obtain local, or less than local, slopes via linear curve fitting.

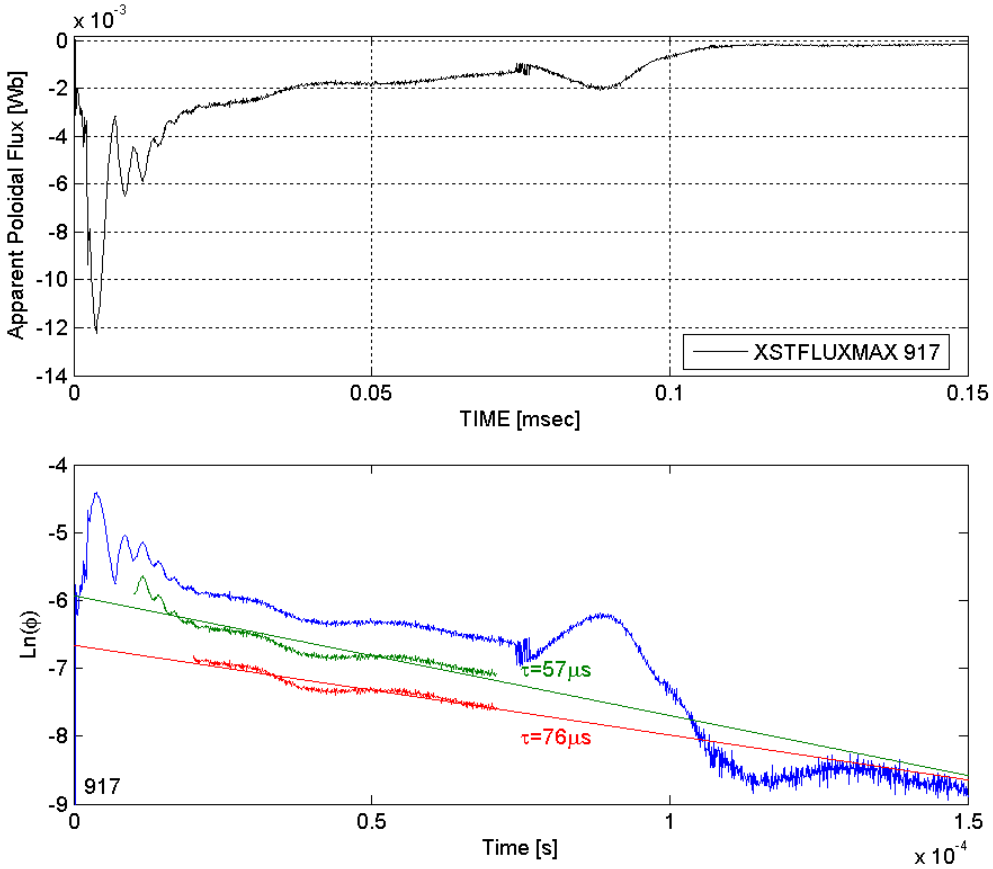

Figure 6.3.1. Magnitude and lifetime of poloidal flux. (top) The assumptions that go into deriving poloidal flux from external diagnostics are broken early in time due to a non-equilibrium state and later in time due to the FRC extending beyond the quasi flux conserving coils. A safe timeframe for analysis on this particular shot is from 20 μs to 60 μs. During this time the flux loss-rate appears to drop to zero. (bottom) For this reason it is preferable to utilize the net effect of some length of time to determine an average τ_φ. The green and red traces are plotted

shifted and for restricted times along with their linear fits and apparent τ_ϕ. The red trace includes times safely away from the axial contraction.

For comparing multiple shots, we start our 'valid' time based off the time required to reach the first peak (radial oscillation). The experiment's formation region is approximately 2.5x longer than wide. So the start of sampling is chosen to be 5x the time taken to reach radial peak so as to correspond with the axial relaxation bounce. This timeframe corresponds with the start of the red trace above. The end of sampling is determined by when the end of the FRC reaches the end of the formation section or drops below some fraction of the peak trapped flux.

At the moment, the edge detection can only happen when the FRC translates off the end. Restriction to within the last diagnosed coil (in length) removes all worthy shots from consideration. This is largely because the length of the FRC is known only to the precision available from the spacing of diamagnetic loops (0.25 m). We interpolate in a linear fashion to gain pseudo high resolution at the occasional expense of some errors in calculating FRC length.

Shots which violate equilibrium conditions (for example by continuing to oscillate radially or axially) can result in erroneous 'flyer' results. Additionally, due to the dependence on vacuum reference subtraction, it is possible for erroneous results to occur due to bank malfunctions or non-repeatable behavior.

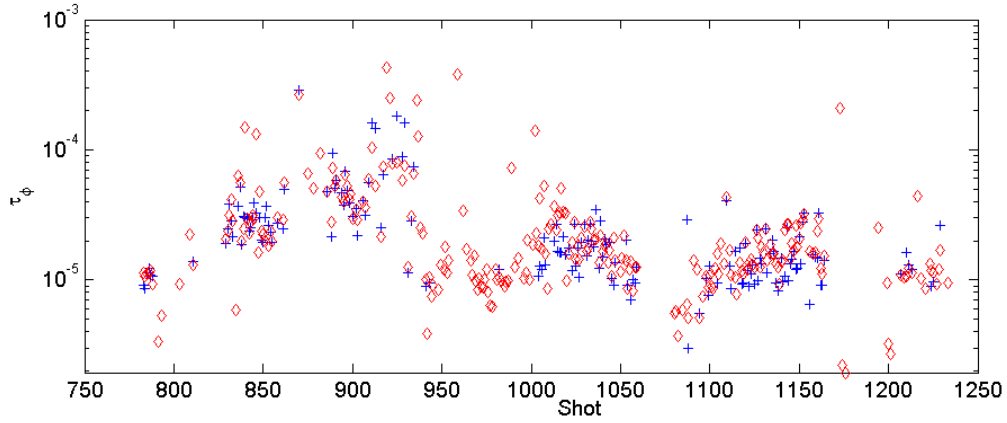

Figure 6.3.2. Trapped flux characteristic loss timescale. (blue) Analysis done with normal restrictions on location and minimum flux. (red) Analysis done with relaxed restrictions. In general, single flyers can be attributed to bank malfunctions. For example, shot 870 (blue/red 300 μs) is due to a crowbar problem on either the associated vacuum reference shot or the plasma shot. On the other hand, shot 913 is a very promising shot that eventually bounces off the end cone. After its initial formation, the restricted time analysis shows ~150 μs lifetime (from looking at only about 30 μs. The unrestricted time analysis shows closer to 50 μs is probably due to including times after the end contact. Overall, this figure shows that the shots with the best lifetimes took place while the plasma source was on the west side of the experiment (about 900) with minimal distance between the source and formation chamber. Moving the source to the east side (with vacuum pumps and no field regions separating) resulted in a significant decrease in lifetimes that could not be recovered by the addition of axial discharges or altering the source.

6.4 Energy

We choose to define the energy in the FRC as the sum of magnetic and thermal pressure energy. We neglect ionization energy, rotational and flow energy,

electric field energy, etc. Later on, we may choose to keep a separate tally of axial kinetic energy. The energy contained within the FRC is integrated over the volume enclosed by the FRC separatrix: $\int\int\int\left(\frac{B^2}{2\mu_0}+\frac{3}{2}nkT_t\right)dx\,dy\,dz$. Understanding that <β> is averaged over this same volume the total energy gives us $\int\int\int\left((1-<\beta>)\frac{B_e^2}{2\mu_0}+\frac{3}{2}<\beta>\frac{B_e^2}{2\mu_0}\right)dx\,dy\,dz$ and simplifies to $\left(1+\frac{1}{2}<\beta>\right)\frac{B_e^2}{2\mu_0}V$, where V is the volume of the FRC separatrix and B_e is taken where there is minimal field line curvature. <β> is $1 - ½\ x_s^2$ (from the axial equilibrium equations and assumptions) taken at the same location as B_e.

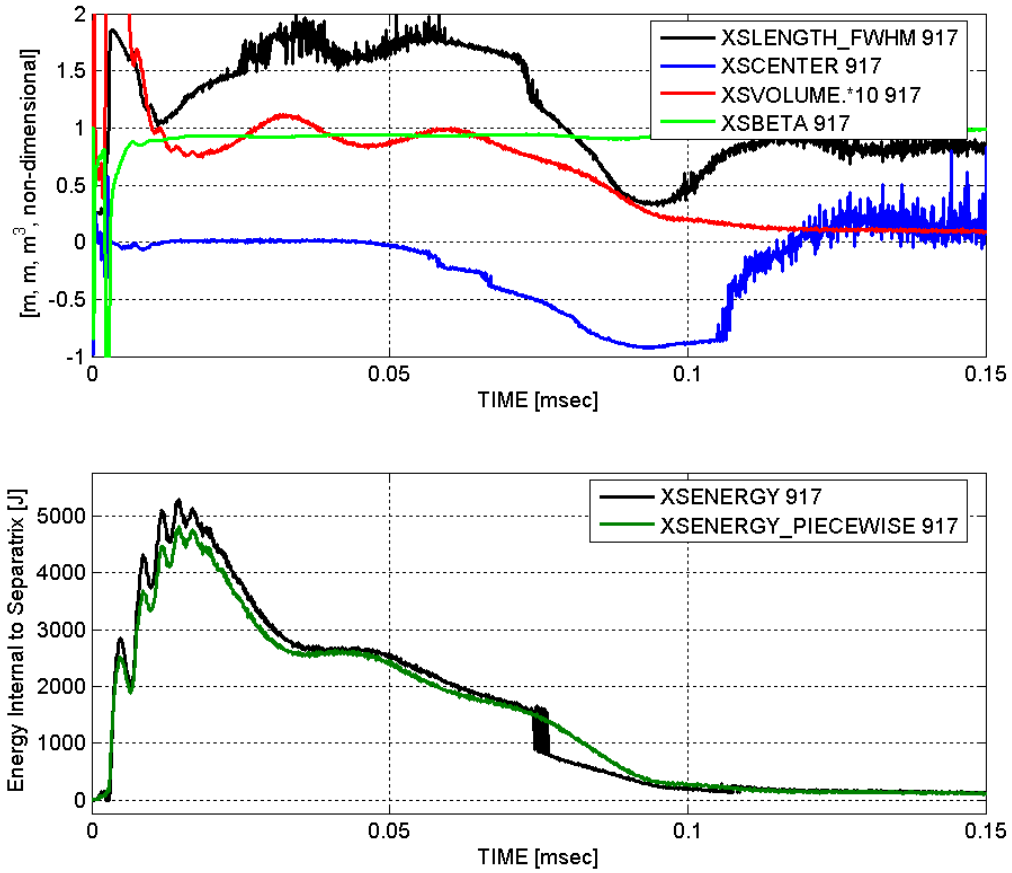

Figure 6.4.1. Information available from magnetic diagnostics. (top) (black) Full width of half maximum separatrix, i.e. FRC length. This is much more subject to noise than the other traces. An alternative formulation that would result in less noise would define length based on FRC volume and radius. (blue) Axial location of FRC separatrix radius centroid derived while ignoring parts with less than 0.7 times maximum separatrix radius. (red) FRC piecewise computed volume. (green) FRC <β> from axial equilibrium assumption. (bottom) (black) Energy in FRC computed from volume, <β>, and axial external magnetic field at peak separatrix. When the FRC passes out from under the diagnostics our measure of FRC size, length, etc becomes very poor. (green) Energy in FRC computed with <β>, localized $B_{z,e}$, and localized separatrix. This form includes energy from plasma that is not part of the FRC.

6.5 Energy Lifetime τ_E

Having the ability to solve for energy with only external magnetic diagnostic information is convenient. The next step is to monitor the change in energy with time

to define a characteristic timescale for energy loss. Energy can be lost from the FRC in a multitude of ways. Radiation is perhaps the most obvious loss mechanism followed by work done on or by the external magnetic fields on the FRC. Radiation was monitored (with a great many noise problems) by an axially located (in vacuum) bolometer during the 40 cm experiments. No measure of radiation suitable for this purpose is available on the 80 cm experiments. Work pertaining to the axial magnetic fields is nominally simple to calculate from the magnetic diagnostics. In addition to these two channels, energy can be lost with charged particles diffusing across the magnetic fields. It can also be lost with neutral atoms due to recombination or charge exchange within the FRC.

As a practical matter, it is not too onerous to display the energy loss timescale while only removing the effects of the external magnetic fields. The remaining effects are those that are due to physics rather than imperfect equipment. While this certainly complicates ascertaining scaling laws (that exclude the effects of impurity radiation for example) from displayed figures, it would entail a great deal of potentially inaccurate estimations (and wasted time) to do otherwise.

Several approximations need to be made to maintain the tractability of the problem. Evaluating work done by the external field on the FRC as the surface integral of applied pressure ($B_s^2/(2\ \mu_0)$) times the perpendicular velocity of the separatrix is not easily put in terms of the external diagnostics for all conditions. As has been investigated earlier in the diagnostic section, the magnetic field at the diagnostic locations is not the same as the magnetic field parallel to the surface of the FRC. There are two obvious variants based on external diagnostics that we might use. One takes B_z^2 and r_s under each diagnostic location as a basis. The second would use the knowledge that pressure as a function of flux surface should be approximately constant. The external parameters of interest would then just be the total volume of separatrix and the external field at the mid-plane of the FRC. Both approaches have obvious difficulties. The second approach would not take into account internal magnetic field line curvature and its effects on the pressure crossing the boundary at the ends and is thus only suitable for highly elongated FRCs. The first approach would require the diagnostic section magnetic fitting in order to be done completely properly. This is currently extremely onerous given present time constraints and the long times required to return accurate results.

We choose to use the first approach modified in the hope that the difference between B_s and $B_{z,e}$ would be partially compensated for by the non-perpendicular to surface change in r_s. As has been noted, for many conditions the standard corrected values of $B_{z,e}$ are not too different from B_s. Recall that the values of B_z we normally use are normalized to the flux loop readings (not really local values). Understand that the usual effect of taking into account work done on the FRC by external fields in our experiment is such that it makes the apparent timescale for energy loss shorter (i.e. worse confinement) during periods where the timescale would otherwise be interpreted as very long.

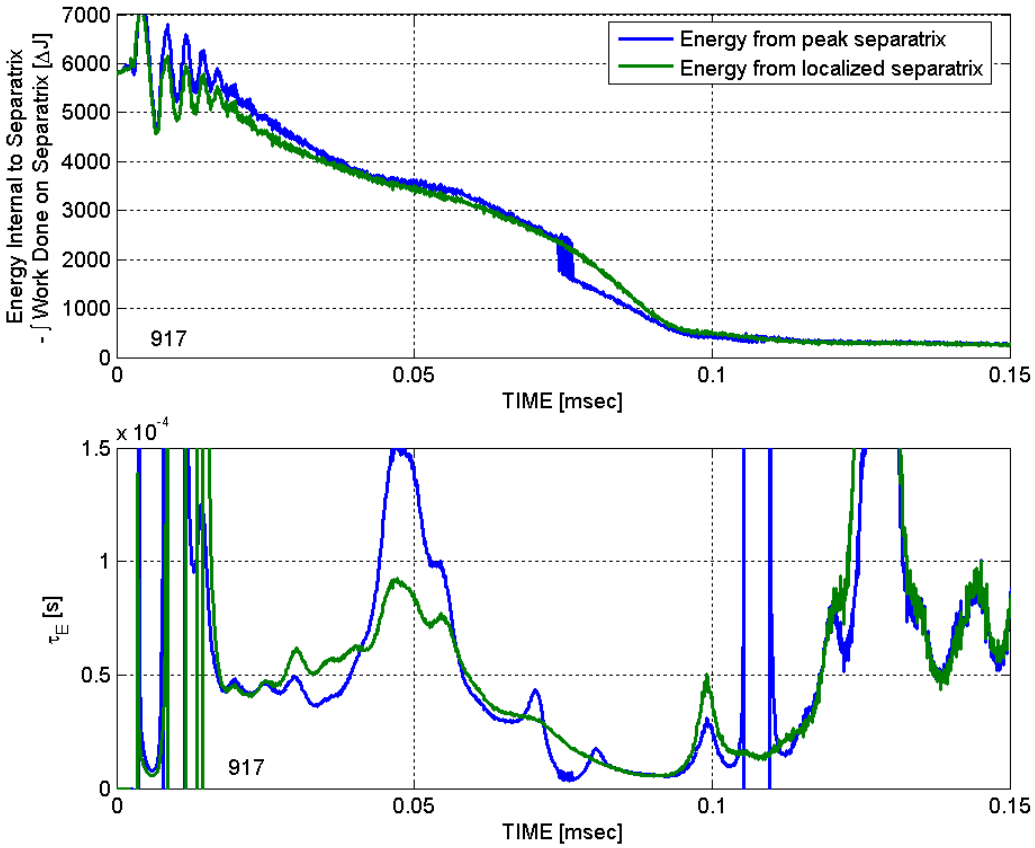

Figure 6.5.1. Energy loss timescale derivation. (top) The integral of external field work done on the FRC has been subtracted from the internal energy of the FRC. The slope corresponds to the FRC energy loss we are interested in characterizing. An arbitrarily chosen constant has been added to result in an end state near zero. (bottom) Energy lifetime using energy loss filtered at 100 KHz.

In order to display useful information, the energy traces must be filtered before taking τ_E = - Energy / ((∂ Energy / ∂t) + PdV) because of the increase in high frequencies due to the taking of the derivative. Here, the energy trace in the figure refers to the combination of FRC energy and the integral of the work done on the FRC. This is for use only in taking the derivative. The choice of the filtering timescale often determines the apparent maximum τ due to the errors in measurement and breakdown of assumptions. The need to account for work done on the FRC by external fields removes the possibility of using the slope of the logarithm of energy or energy modified by work done as the basis for determining τ_E. However, the error via such an approach could be small if the change in volume of the FRC is negligible.

The generally applicable way to reach a single τ_E for a shot is by averaging $1/\tau_E$ for the duration of the FRC. This approach does not require any frequency filtering to function. If frequency filtering is done, it tends to expand the boundaries of the time domain over which the averaging is executed. For consistency, we choose the same criteria for choosing the window as we used for determining flux lifetime (with the exception that times based on relative trapped flux are now based on relative

energy). Slight differences in the analysis frequently lead to a slightly earlier window than that used for trapped flux.

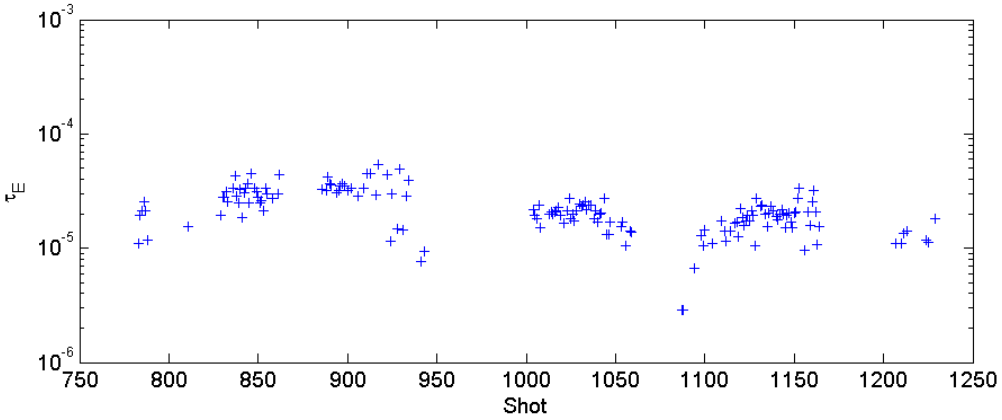

Figure 6.5.2. Energy loss timescales for 80 cm experiment shots. Only the shots for which τ_φ were available have been displayed.

6.6 Plasma Inventory Lifetime τ_N

The interferometer provides us with a measure of plasma inventory so long as the rigid rotor profile derivations (and our measure of K via x_s and axial equilibrium) holds true. The average cross section density can be found from integrating the rigid rotor density equation $\int_0^{\sqrt{2}R} 2\pi r n_m Sech\left(K\left(\left(\frac{r}{R}\right)^2 - 1\right)\right)dr = \frac{2\pi R^2 n_m Tanh(K)}{K}$ and dividing by the cross section to find that $\langle n\rangle = n_m Tanh(K)/K$. Tanh(K)/K is of course equal to $\langle\beta\rangle = 1 - \frac{1}{2} x_s^2$. So the average cross sectional density is $\langle n\rangle = n_m (1 - \frac{1}{2} x_s^2)$. We gain knowledge of n_m from the interferometer at the mid-plane combined with the relation of Figure 6.2.3 which shows the result from integrating density along a line. Our measure is only accurate so long as the FRC remains centered on the interferometer at the mid-plane. This restriction might be codified into (axial location < axial length/4) to allow some lee-way. Total inventory can then be estimated from the separatrix volume.

These constraints tend to be more restrictive than that on the other lifetime measurements, leading to more inaccurate estimates due to oscillation periods longer than sampling times. We do not have suitable diagnostics for estimation of neutral gas ionization rates (to compensate for additions to inventory by ionization).

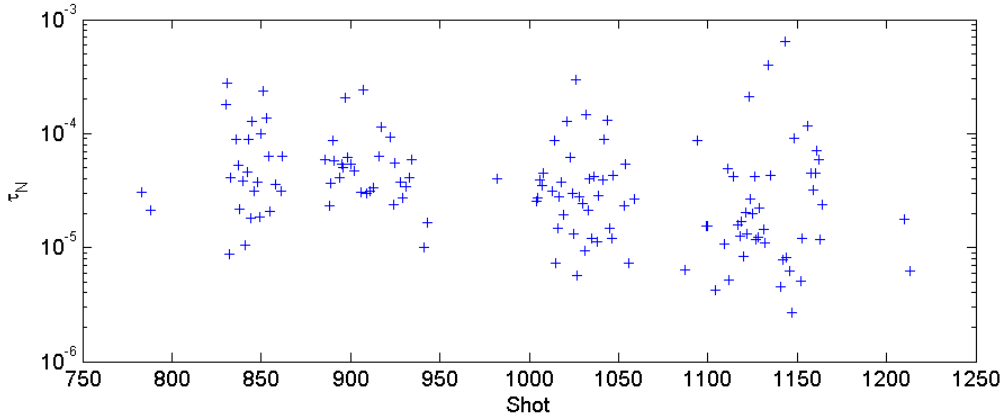

Figure 6.6.1. Plasma inventory timescales for 80 cm experiment shots. Only the shots for which τ_φ were available have been displayed.

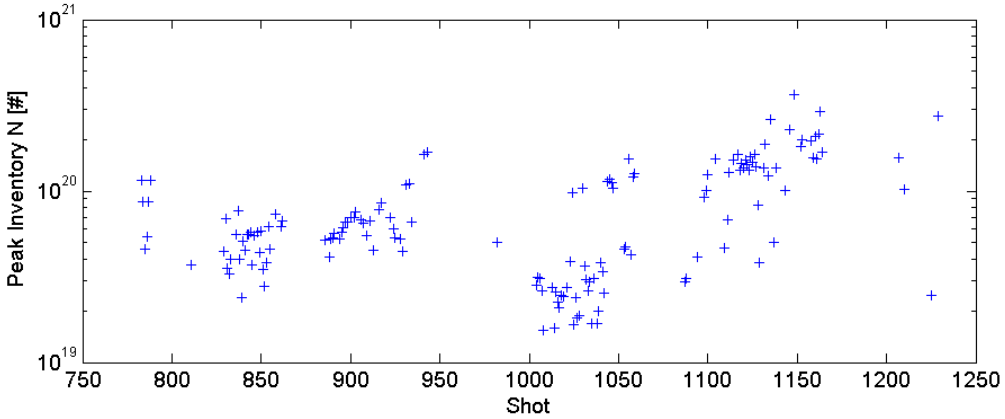

Figure 6.6.2. Peak plasma inventories inside separatrix. 9 x 10^{19} is equivalent to a uniform 1 mTorr of diatomic deuterium. Only the shots for which τ_φ were available have been displayed.

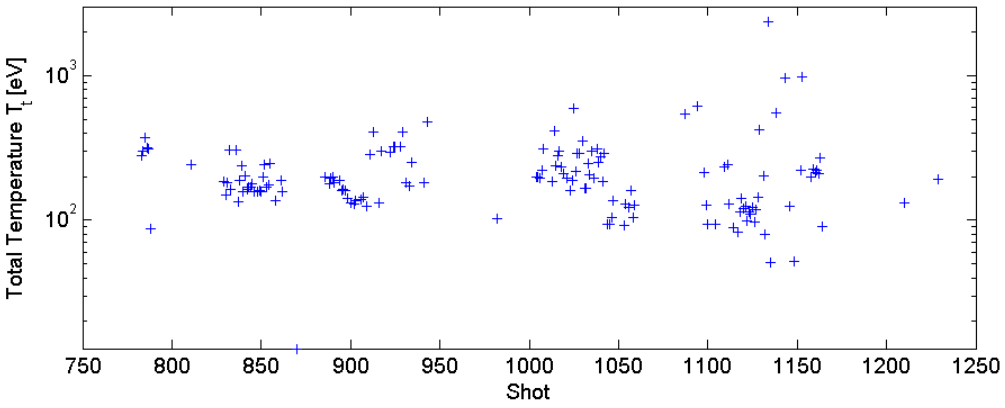

Figure 6.6.3. Total plasma temperature at first relaxed FRC time. Interferometer errors can result in large errors. Only the shots for which τ_φ were available have been displayed.

We would expect the particle loss rate to be governed by a lower hybrid drift fluctuation scaling causing particle lifetime to scale as $\tau_N \alpha s^4(1+C/s)^3$ while

experimentally scaling as R^2/ρ_i under current conditions[88]. According to the reference, C is typically 50. However, we expect any such scaling to be clouded by our imperfect accounting for inventory sources. If energy loss is dominated by particle loss (rather than by thermal conduction), we can use τ_E for τ_N, allowing us to avoid the problems of measuring inventory additions (by ionization).

6.7 Kinetic Parameter s

With ideal MHD assumptions FRCs should be unstable on a characteristic timescale derived from the Alfvén velocity. The ability of some FRCs to persist beyond the expected limits led to the conclusion that one of the assumptions must be critically wrong. This stability has mostly been explained through ion kinetic effects. One of the non-dimensional parameters used to compare theory, simulation, and experiment has been the kinetic parameter s defined as: $s = \int_R^{r_s} \frac{r}{r_s \rho_i} dr$.

The ion gyro-radius $\rho_i = \frac{v_{Ti}}{\omega_{ci}} = \sqrt{\frac{k T_i}{m_i}} \frac{m_i}{ZeB} = \frac{\sqrt{T_i}\sqrt{\mu k m_p}}{BeZ} = 1.022 * 10^{-4} \frac{\sqrt{\mu}}{Z} \frac{\sqrt{T_i[eV]}}{B[T]}\,[m]$, where μ is 2 for deuterium and Z = 1 for no impurities. With our previous assumptions associated with a rigid rotor profile, $s = \frac{Z}{2 * 1.022 * 10^{-4} * \sqrt{\mu}} \frac{r_s B_e}{\sqrt{T_i}} \left(\frac{\mathrm{Log}(\mathrm{Cosh}(K))}{2K}\right)$. The term in parenthesis can be approximated by 0.31 x_s as before. X_s can be put in terms of r_s so as to yield a relation that is dependent on excluded flux and ion temperature as variables within the experiment $s = \frac{.31}{2\pi * 1.022 * 10^{-4}} \frac{Z}{r_c \sqrt{\mu}} \frac{XF}{\sqrt{T_i}}$, where r_c is the flux conserver radius. Excluded flux is in Webers and ion temperature in electron-volts. Z is the charge state (assumed = 1) and μ is the ion mass (deuterium = 2). This is a convenient form for use in data analysis since excluded flux and ion temperature are easy to measure with appropriate diagnostics. For our purposes, we assume that $T_i \sim ½ T_t$ due to the low total temperatures associated with our current modes of operation with the justification of experimentally observed values from LSX[89]. Only at higher temperatures would we expect a transition to domination by ion temperature. If this temperature assumption is wrong, it is a trivial matter to take the results in the following section and compensate. The derivation of s is reliant on the usual average beta condition assumptions.

Of course, we must choose some time for which to display s. One approach would be to average s over the same windows that the confinement timescale are determined by. We choose to display it for the first axially relaxed time.

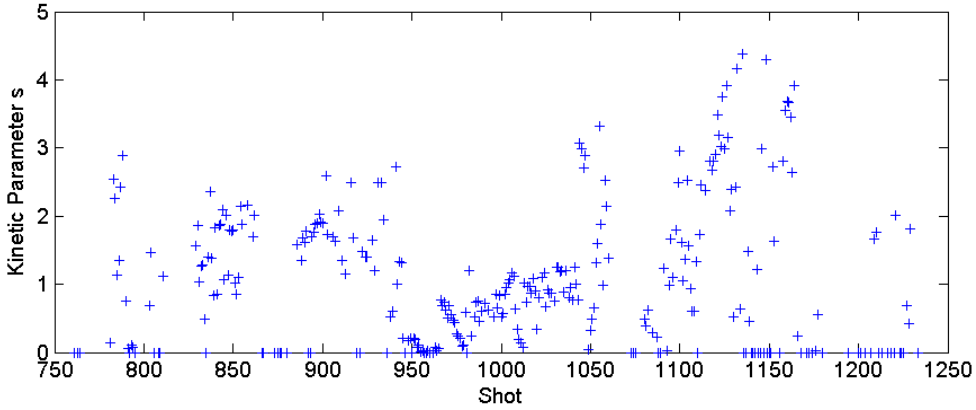

Figure 6.7.1. Kinetic parameter (little) s. Data is calculated at the first axially relaxed time. Only the shots for which vacuum references were available have been displayed.

6.8 Collision Mean Free Path

We can gain a crude definition of the ion mean free path from taking the thermal velocity divided by the ion collision rate $\lambda=v_{ti}/\upsilon_i$. From the NRL plasma formulary, ion thermal velocity is $v_{ti} = (kT_i/m_i)^{½} = 9.79 \times 10^3\ \mu^{-½}\ Ti^{½}$ [m/s] and the ion collision rate is $\upsilon_i = 4.80 \times 10^{-14}\ Z^4\ \mu^{-½}\ n_i \ln \Lambda\ T_i^{-3/2}$. The ion mean free path is then $\lambda=2.04 \times 10^{17}\ Z^{-4}\ T_i^2\ n_i^{-1}/ \ln \Lambda$. This mean free path is within a factor of 2 of other approximate solutions (another solution uses 3.4×10^{17})[90]. The Coulomb logarithm is approximately $\ln \Lambda=\ln(12\pi n(\varepsilon_0 kT_e/(n_e^2))^{3/2})$. Our interest in it arises from ion communication with the ends of the FRC.

6.9 Neutral Density

In normal FRTP FRC formation (as with the 40 cm experiments) a steady backfill gas is inserted in the chamber and ionized via axial current discharge well before the action of field reversal. In these situations the available injected neutral population is simply calculable and the action of the ionization source and FRTP formation is coupled together by a dependence on density. While attempting to provide a better form of ionization via MPD we have made it more difficult to determine total available neutral inventory. Our interest in neutral pressure independent of FRC inventory is due to a desire to understand the critical parameters determining flux loss before lift off the wall. The uncertain characteristic thickness of the sheath or profile of plasma pressure makes it difficult to use the interferometer for this purpose.

We have a low frequency measure of vacuum chamber pressure due to a capacitive bolometer; this does not simply determine the inventory available for the FRC. There are several aspects of concern. The most obvious is that our gas inventory is largely injected; not a constant pressure backfill. This means that in order to determine average gas density we must close the pump ports. Much of the gas that enters the system (and is measured) no doubt enters after it could be included

in the FRC. The total injected pressures for given tank pressure and on time also vary with the version of MPD.

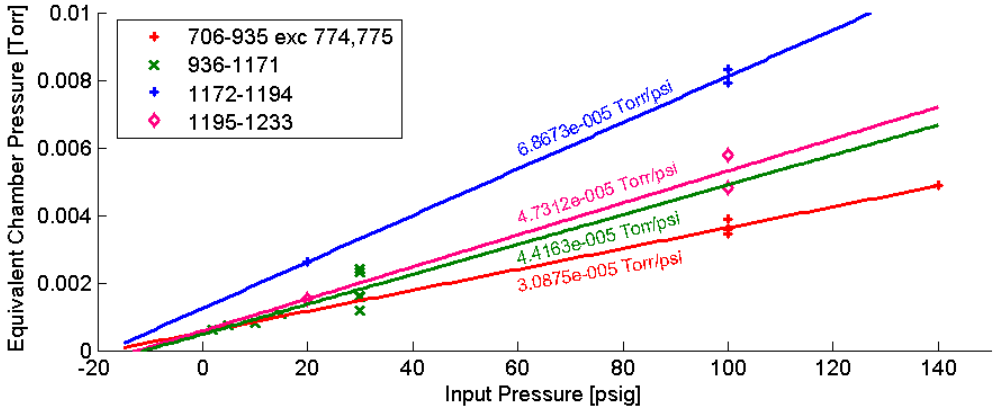

Figure 6.9.1. Total neutral density delivered to vacuum system from MPD. Input pressure refers to the input pressure to the MPD. For the shots listed here, the time for which the poppet valve is energized is always 600 µs. Earlier shots had greater variety in timing while the MPD was being optimized.

At times we have injected gas in a pseudo backfill manner via a second gas valve (*EastPuff*). The obvious price being that ionization with this approach can not be as complete as we would like. A curve fit for this independent gas puff valve's injected pressure has been found [mTorr] = (.9729+Duration[msec] x 1.5716) x (.01164+Pressure[psig] x .001435) (deuterium). This form is only expected to be accurate for long duration puffs.

A second serious complication results from the understanding that the quartz plasma chamber wall can store a significant amount of atomic deuterium on its surface (sorption). Some fraction of this inventory is no doubt liberated while the plasma sheath sits on the wall. While the active physical mechanisms could be quite complex involving multiple layers at densities commensurate with solids, a very early (1929) experiment should be applicable to our situation. In that experiment[91], the adsorption of a baked out glass boundary was measured in the presence of atomic hydrogen gas. The hydrogen gas was contained within a glass boundary at a pressure around 90 mTorr. An internal discharge was initiated by an external coil resulting in the pumping of the hydrogen on to the wall (minute timescales at these power levels). Atomization was necessary for the hydrogen to stick to the surface (this was not a cryogenic experiment). This was done with the boundary at room temperature. The pressure loss was consistent with a monolayer of atomic hydrogen at a surface density up to 6 x 10^{15} atoms/cm^2. This process saturated at the stated surface density. The pressure available from the wall in our system (if fully populated) would be equivalent to 4 mTorr of molecular deuterium fill. Our desired pressure ranges up to 10 mTorr with the most interest in levels of around 3 mTorr. Obviously, this can be the dominant source of gas.

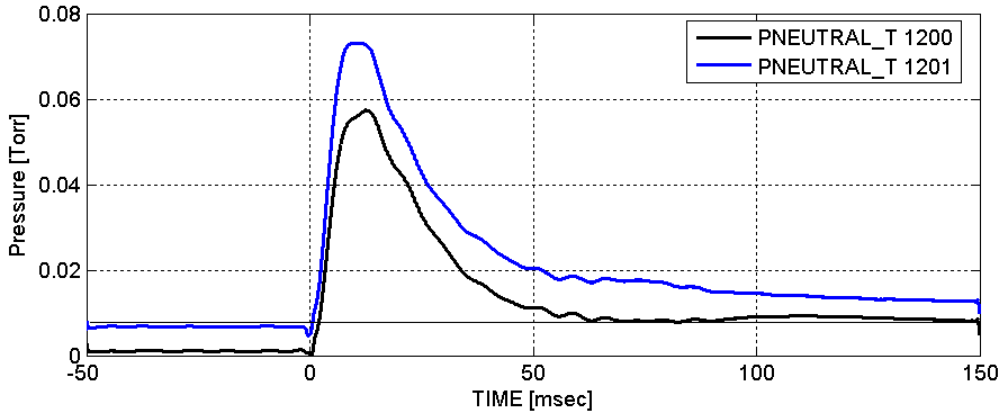

Figure 6.9.2. Chamber pressure spikes after a plasma shot. The pumps are closed off for these two shots. The pressure peak in this low power 10 kV plasma shot exceeds that of a neutral gas shot by a factor of 3. The pressure after the first shot remains relatively constant until the firing of the second shot.

If we populated the wall with some other species we would likely have impurity problems. If we stripped the wall (for example by running ringing theta pinch wall conditioning) we would probably have impurity problems under conditions with higher energy or time sitting on the wall. We would assume that firing a small number of shots in a short period would effectively strip the wall of its readily available gasses. Many of our shots have been run after ringing theta cleanup shots. Reproducibility is then likely to be dependent on the consistency of time allowed between shots versus the time required to repopulate the wall either with FRC gas or the base pressure of the system.

7
Results

This chapter is divided into four sections. The first two sections deal primarily with the difficulties encountered during operations. Chiefly, this consists of ionization / plasma distribution problems and their effects on optimizing desirable parameters. The third section presents experimental parameters inter-relationships in a form suitable for drawing scaling law conclusions. Many of the critical parameters are clearly dominated by the effects from ionization techniques. Nevertheless, within each ionization technique subset, patterns do emerge. The fourth section presents a few interesting shots that demonstrate the axial behaviors that are possible with this experiment.

Plasma shots without corresponding vacuum reference shots are generally excluded from these figures. As a result, many of the initial operation test shots are excluded so as to avoid confusing issues with inaccurate results. Dotted lines are added for illustrative purposes only; they are not numerically derived data fits.

7.1 Ionization

This section details the overall results of steps taken, and optimizations in timings made, to achieve maximum trapped flux for the varied MPD configurations.

For much of this data it helps to keep in mind the timeline from the end of chapter 4. For example, before shot 776 the MPD was operated with a PFN network that limited the duration of the output current pulse. This correlates nicely with the extremely low flux trapping shown before shot 776 in Figure 7.7.1 below. After this time the MPD bank used a crowbar to maintain current. Initially, the MPD bank crowbar was a type A ignitron. The MPD was also capable of running at slightly higher neutral gas pressures (155 PSIA vs 115 PSIA) initially, but the pressure had to be turned down due to gas line connector issues. This configuration accounts for the unusually high results from 776 to 801. Around the same time as the gas line issues, the type A ignitron failed and had to be replaced with a type D. The addition of RMF generally increased the flux trapping, sometimes by a factor of 2 over the pure MPD results.

It is worth noting that most of the original LSX results had a similar flux trapping issue. Flux trapping was governed by the time taken for the pre-ionization to diffuse out to the wall resulting in flux trapping efficiencies of 50% or less[92]. The bias fields had a longer ¼ cycle time around 200-250 μs.

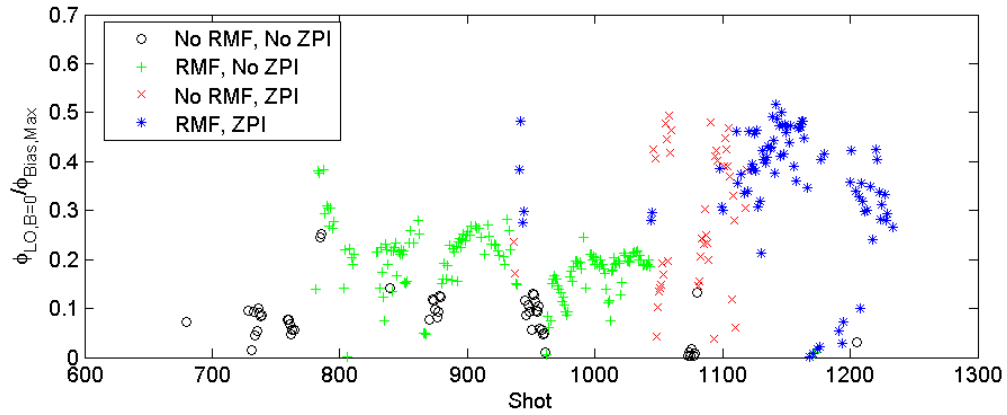

Figure 7.1.1. Flux trapping fraction shows strong dependence on ionization methods employed. The MPD source has a 6 cm radius. (black circles) The MPD, on axis and transitioning through a cusp field, generally provided a very low level of flux trapping consistent with a plasma column radius around 13 cm. (green crosses) The addition of RMF enlarged the effective plasma column radius to around 21 cm. (blue stars) Adding an axial current discharge (ZPI) from end cone to end cone improved the radius to nearly 28 cm. (red x) Axial discharges with the MPD could reach radii as large as those with additional RMF, but not as consistently.

Running with the bias banks in second half cycle (effectively bias field reversed), so as to temporarily increase bias flux with the firing of the reversal banks before oscillating to reversed field, did nothing to improve flux trapping. These shots account for the dip in trapped flux around shot 875. They are labeled 2HC in the run matrix, but do not refer to the traditional definition of second half cycle.

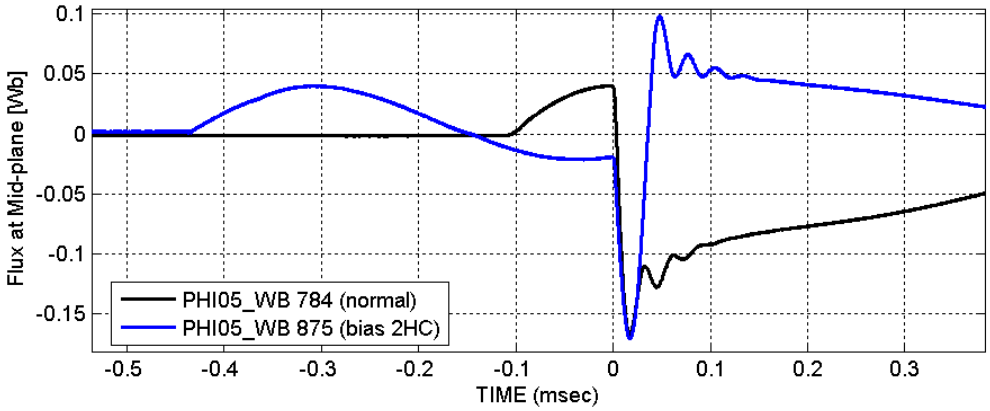

Figure 7.1.2. Bias 2HC. The plug field connections are also reversed in order to maintain the presence of a cusp field for reconnection shortly before reversal bank firings.

The repositioning of the MPD from the target (west) end to the source (east) end included the addition of the axial discharge capability. The shots [~936] corresponds to the location of the two red x-s in figure 7.1.1. With no ionization other than the MPD itself, the flux trapping after this change was of similar level as that before. This suggests that the additional distance between MPD and formation region, as well as crossing regions with no guide fields over vacuum pumps, had marginal impact on its effectiveness. However, in the presence of RMF, the east

location never managed to reach the flux trapping level seen with the MPD in its original, close in, location. A general decrease in obtained lifetimes is observed for all shots with the MPD on the east end.

While the trapped flux increased with the use of an axial discharge, the lifetimes did not. For shot 1172 onwards, a hollow copper center MPD was used resulting in no better performance. This modified source would soon have to contend with the soft x-ray detector diagnostic (shot 1191-) installed in the path between source and formation chamber. Lift off flux values never recovered.

7.1.1 MPD Location

Before moving the MPD to the east (source) end of PHD it was generally accepted that ionization needed to be done differently. It was speculated that there would be no efficient way to reconcile the metallic chamber components (i.e. vacuum pumps, cones) with the need to magnetically map the on axis source to field lines along the walls. The fields in the end cones could not be changed rapidly enough to transition from a mirrored arrangement to a cusp.

Before any plasma shots were fired in the 80 cm experiment, there were serious doubts that the MPD plasma would be able to cross the high bias fields required in the short times allotted (resulting in high electric fields seen by the plasma, if they could cross it).

The speculated solution to this problem was to build many discrete sources in an annulus near the wall, inserted through the end cones. Cold gas could certainly cross the fields, and the collision dominated conductors could certainly carry the current with no problems. The difficulties (aside from basic MPD design/physics) would be: (1) shielding the puff valves from magnetic interference, (2) ensuring that the azimuthal ripple in plasma conditions result in no instabilities, and (3) magnetically isolating the sources from the FRC during reversal. We would have the potential for two attempts at the problem. One end cone would be immediately perforated with 24 ports spaced apart from each other with roughly a maximum 50% azimuthal coverage. If this proved insufficient, the remaining end cone could be perforated to inject plasma in the gaps in the perforations of the first end cone.

If such an approach could be made to work, the upshot would be that the initial distribution of plasma density would be peaked near the wall (the location of the magnetic null). One would expect this distribution would be much better than a quasi-even distribution (as in other experiments) and the axially peaked distribution (as with our current experiments). Additionally, end on on-axis access would be preserved for diagnostics, gas flow, and FRCs.

The downside to the approach is that the dielectric quartz tube could become easily contaminated with plated metal if the MPDs are operated improperly or have manufacturing/design defects. Rushed MPD construction and testing on the operating experiment is highly unadvisable.

The last shots taken with the experiment were before November 2006. The ports for MPDs were drilled in one cone during the spring of 2007 before the summer 2007 move from our old temporary facility in Condon Hall (University of

Washington, Seattle) to a warehouse in Redmond. As of this writing, 12 / 24 of the smaller MPD sources are still under construction.

In the interim, the technique has been applied to a smaller experiment (sometimes known as IPA, inductive plasmoid accelerator) run by MSNW of Redmond, WA. We have been told the results were generally very positive.

7.1.2 Basic Magnetic Timings

Some of the key initial field timings stayed relatively constant throughout the data set. The other timings (such as ionization source timings) varied greatly. Given the similarity in optimized timing between MPD and plug field, it seems likely that the plug field represented a major barrier to the translation of the MPD plasma into the formation chamber. Optimized timings for the MPD are almost always on the leading edge of the rise in plug fields.

Two basic times were used for the bias fields. The first timing ensured that the reversal banks were fired near to the peak of the bias field. The second timing allowed approximately 40 μs more time for diffusion before firing reversal banks (at the cost of allowing bias flux to drop ~10%). It is important to keep this difference in timings in mind while considering the RMF time while the MPD was on the west end of the experiment.

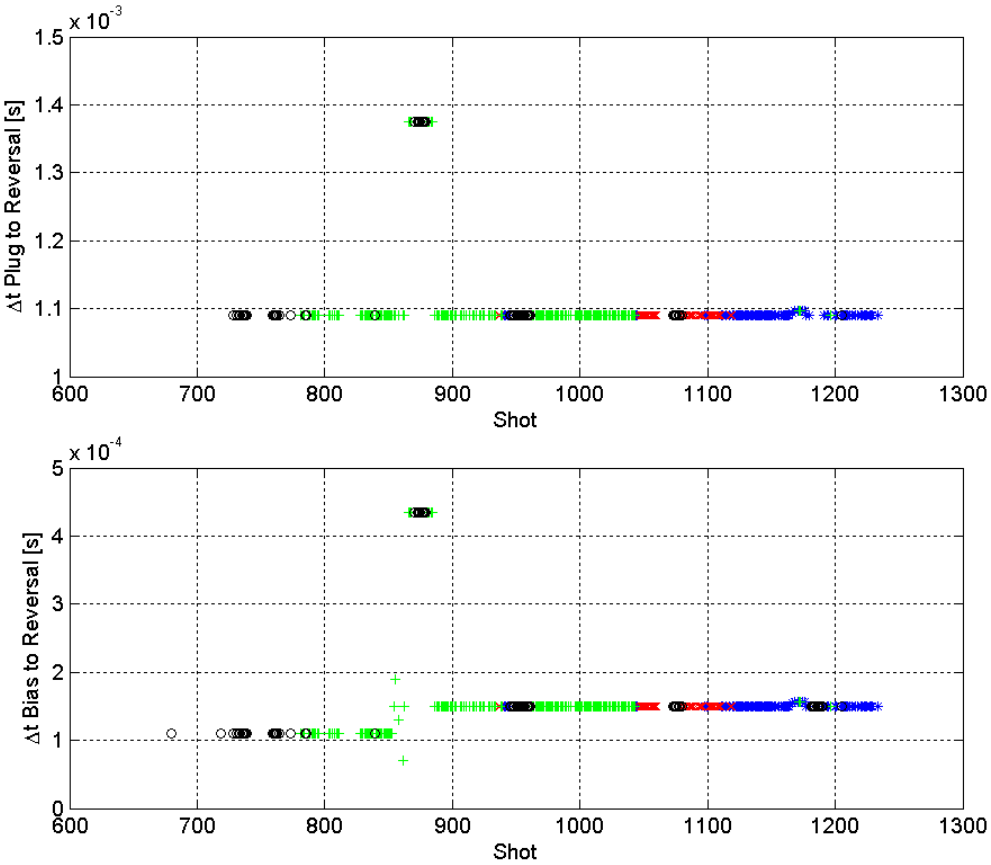

Figure 7.1.2.1. Plug and bias field timings. The one major deviation corresponds with bias 2HC timings. As might be inferred from some of the gaps, not all exploratory timings are considered or shown in the figures of this chapter.

7.2 Parameter Optimization

The initial optimization technique for the timings for MPD, bias fields, and RMF could be described as round-robin single parameter optimization. An input variable was varied until reaching peak trapped flux for that individual variable. The next input variable was then optimized for peak trapped flux. So it would go round-robin through the available control inputs. Initial input parameters would be chosen on the basis of 'back of the envelope' calculations that won't be repeated here. For the key parameters, the survey was occasionally expanded to look for 'islands' of better results. This was especially done when the rudimentary calculations were considered unreliable. Ultimately, this was a practical approach due to the need to optimize many variables including MPD input pressure, neutral gas timings, MPD voltages, MPD magnetic field, bias voltage, reversal voltage, RMF voltage, plug voltage, plug timing, secondary gas fill timing/pressure, and the list goes on. The process can be compared to the adage of blind men identifying a (silent) elephant. Although, we are more like ants crawling on the elephant… Is there a better location on the other side of this dip?

Peaking only one parameter (such as lift-off flux) does not necessarily lead to the best FRCs possible. Note that in the figures below, optimum conditions for energy lifetime are frequently achieved at conditions different than those for maximum lift-off flux. Also, some conditions are more difficult to achieve repeatable shots for. In such cases we have favored shifting towards repeatable running parameters over conditions which only occasionally yield 'super' shots. A machine with reliably repeatable operation is much more desirable for the next phases of the program.

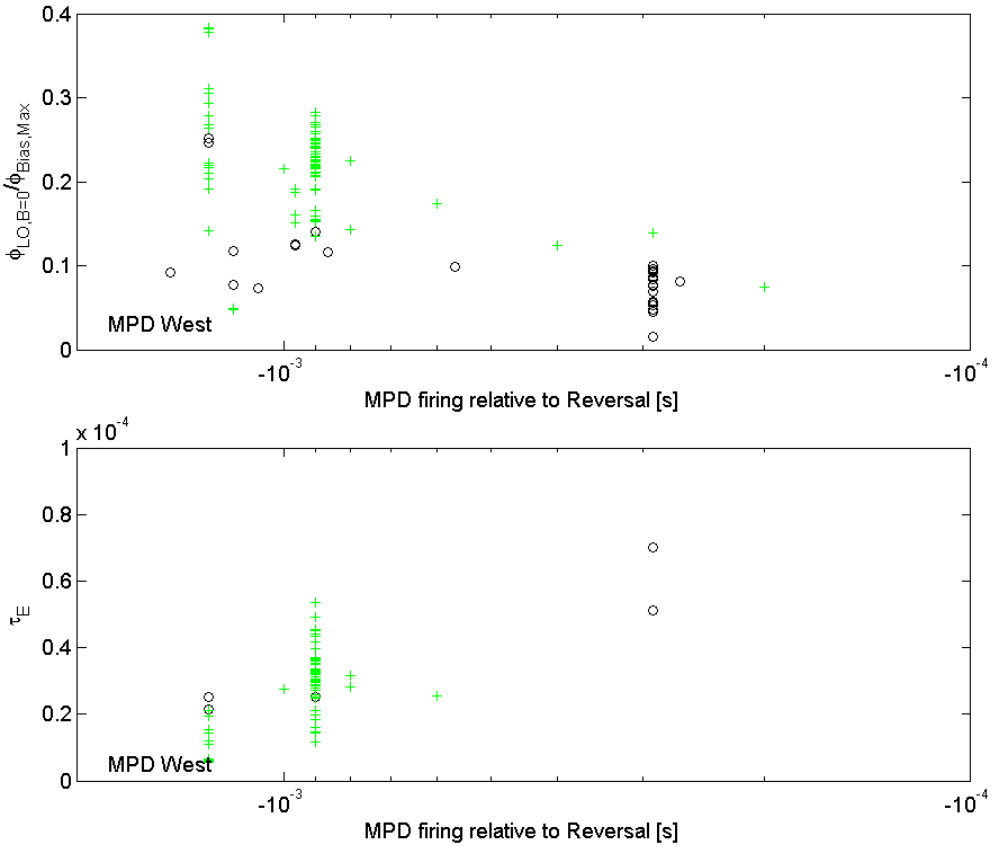

Figure 7.2.1. West MPD timing optimization. The times of the start of the MPD pulse are displayed relative to mid-plane section reversal bank firings. (top) lift-off flux trapping data is available for almost all shots. Some of the earliest optimization shots did not have vacuum references. Data without vacuum references is not displayed here. Green crosses indicate shots with RMF. (bottom) Energy lifetime is reliably available for fewer shots because the FRC has to last long enough to come into axial equilibrium. Shots with small trapped flux may be more difficult to reliably measure lifetimes for.

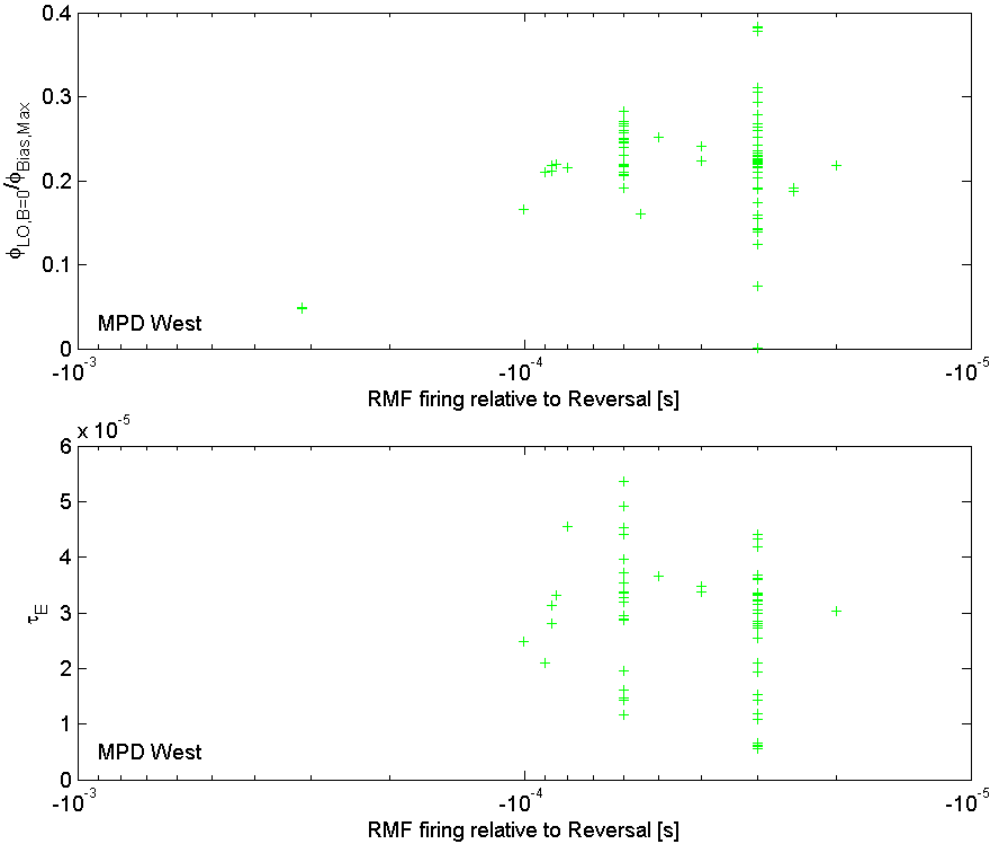

Figure 7.2.2. West RMF timing optimization. Bias banks are typically fired at -150 or -110 µs. For the '2HC' shots this timing difference was extended to -435 µs. Note how several individual shots bridge between the two primary grouped stacks. These bridging, or parameter scan, shots do not have all other parameters fully optimized. So rather than linking the tops, they link partway down. The two peak structure is not explained by displaying time relative to bias fields.

The non-grouped shots show parameter scanning operations with other parameters not always optimized. The grouped shots are generally conditions where other parameters were undergoing optimization. Capacitor bank issues limited the availability of full field reversed theta pinch formation shots; so the choice was made to optimize what we could, while we could, with the emphasis on reaching good MPD plasma production conditions rather than ensuring high quality data.

What are not shown are the optimizations of the MPD itself for inputs such as gas pressure and timing or input current. Those scans primarily took place during early shots (620-719) with no vacuum references. Those shots are not shown to avoid confusion from many pure MPD shots apparently trapping more flux than the later shots with RMF shown here. It is believed that they displayed erroneously high due to the lack of reference shots (and hence purely based on diagnostic area but no spatial corrections). A second possibility is continual evolution of the MPD. A third is the changes in the plug magnetic fields that took place when the original coils

failed. At the time, optimal conditions were mostly judged by extremely weak visible emissions results.

7.2.1 East MPD

Operations with the MPD located on the East (source) end of the experiment have more available ionization inputs. Exceptionally poor flux trapping was expected so arrangements were made to add an axial discharge bank to the main formation chamber. The results from this stage were much more striking in qualitative behavior than those from the west MPD. This may be due to the completion of the optimization of 'internal' MPD parameters and timings.

With respect to flux trapping, MPD only operation yields a clear peak in flux trapping for 1-2 ms before reversal bank firings. This timing may be dominated by the plug firings at -1.1 ms. Addition of RMF extends the breadth of the MPD timing peak to 1.5 to 6 ms. Adding axial discharge rather than RMF results in a situation where the earlier the MPD fires, the better.

The peak for lifetime does not generally match the peak for flux trapping timings. In general, the lifetimes are worse compared to those possible with the MPD on the west end. The best results were obtained with the MPD firing approximately 2 ms before reversal banks.

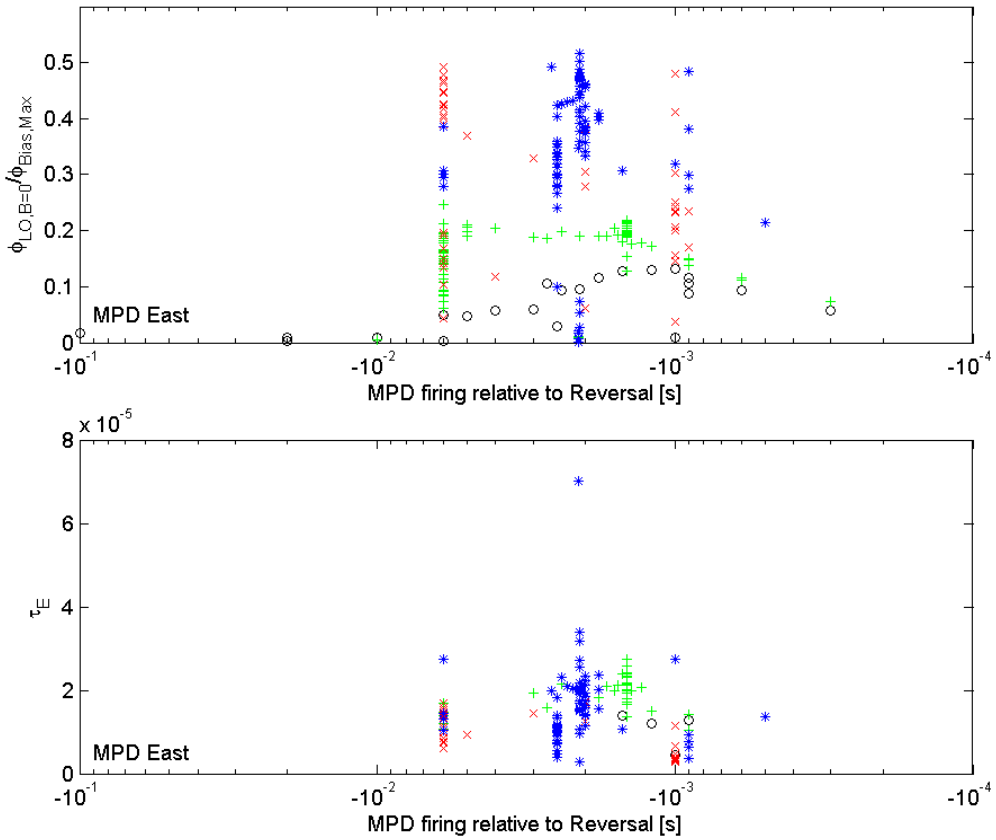

Figure 7.2.1.1. East MPD timing optimization. (top) Initial lift-off flux trapping shows different optimums for varying combinations of ionization sources. (bottom) The timings that result in better lifetimes are much more clustered around 2 ms. (black) MPD alone. The best flux trapping is achieved with the MPD starting operation approximately 1 ms before reversal bank firings. (green) MPD with RMF. A range of MPD timings from – 6 ms to -1.5 ms now give the best lift-off flux trapping. (red) MPD with axial discharge. The best flux trapping is now witnessed when the MPD is employed 6 ms before reversal bank; although it can potentially do just as well with 1 ms. (blue) MPD, axial discharge, and RMF. Best trapping is observed around 2 ms.

On the other hand, RMF timings show a clear preference for relatively early operation (~ -100 μs before reversal, shortly after the firing the bias banks). The same is true for the axial discharge during ZPI + MPD operation; as well as all three ionization systems operating. Operating the ZPI significantly closer to reversal rapidly degrades both the flux trapping and energy lifetimes. It is doubtful that the poloidal currents that the ZPI introduces remain in significant strength when the reversal banks are fired to form the compact toroid. If this were not the case, the potential for creating spheromaks would be a strong complicating factor for analysis. If there was a benefit to operating in such a mode, one would have expected our optimization technique to move the ZPI firing to a time ¼ cycle before reversal bank firing (10 μs). We do not consider that lack of such an optimization to be proof that a FRTP spheromak has to have bad lifetime. Forming a spheromak via FRTP + ZPI was never one of the goals of this project. Axial current when the ZPI is fired 100 μs before reversal decays below 1 kA at the time of reversal.

The effect of RMF operation is worthy of some speculation. The two primary possibilities are that the RMF either expands the radius of the MPD plasma column by increasing diffusion/turbulence or it creates ionization near the vacuum chamber wall. No signs of fast flux trapping (associated with plasma at the wall) have been identified independent from normally occurring noise spikes (possibly due to +/- bank firing imbalances). The RMF is generally incapable of providing ionization acting without the MPD.

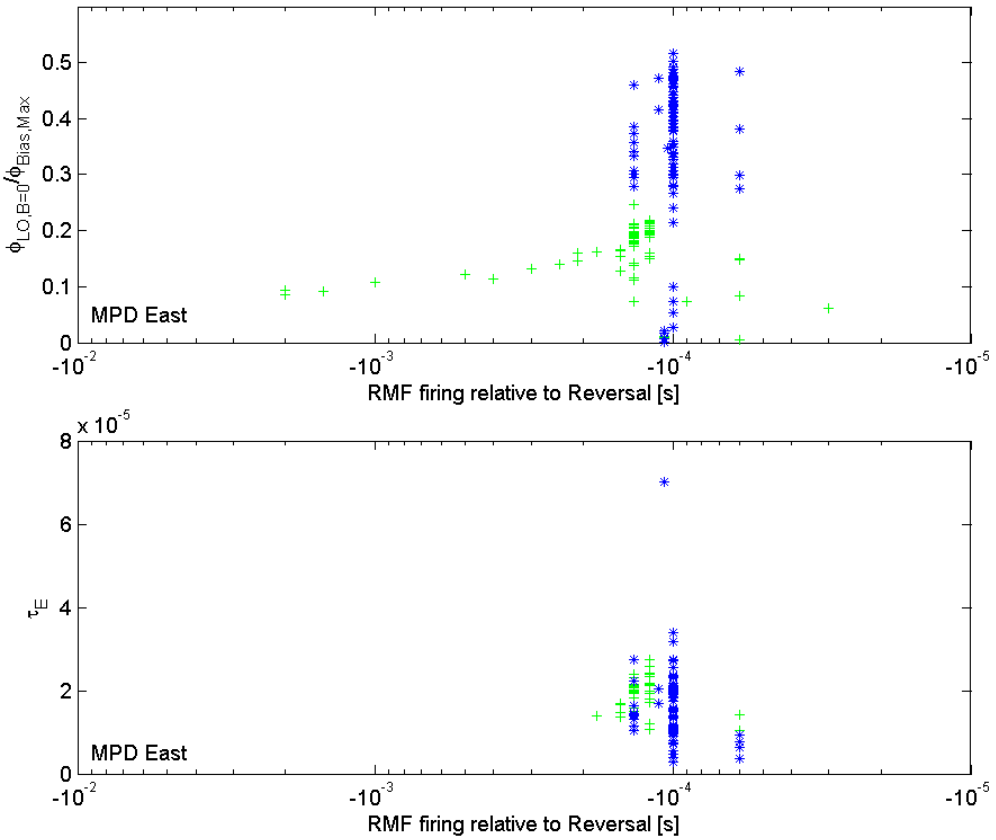

Figure 7.2.1.2. East RMF timing optimization. Note the peaks for all operating conditions and both flux trapping and lifetime for timings near -100 μs. This corresponds to a short time after the firing of the bias fields.

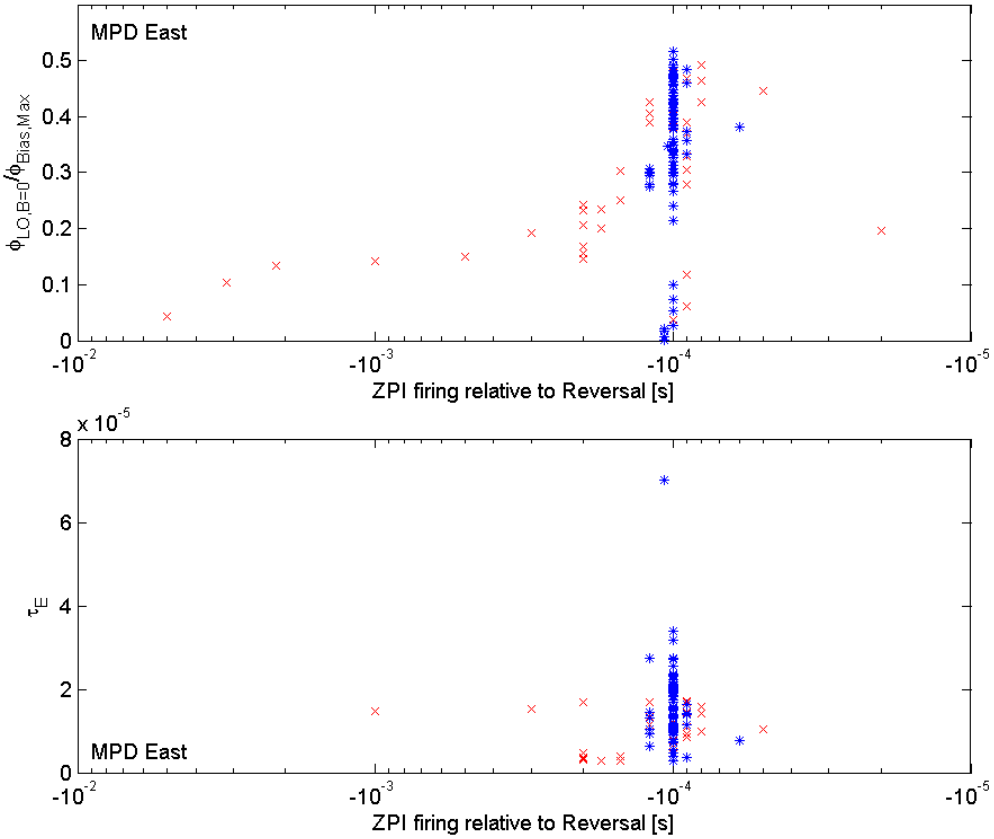

Figure 7.2.1.3. East axial discharge timing optimization. Note the peaks for operating conditions with MPD, RMF, ZPI for both flux trapping and lifetime for timings near -100 μs. Intriguingly, the MPD + ZPI operation doesn't yield a clear peak in ZPI timing for lifetime, but it does for flux trapping.

7.2.2 RMF Voltage

The ring down RMF system currently in use has a maximum charge voltage of 30 kV. Typically this is run at 25 kV to keep ignitron problems to reasonable levels. However, after the addition of the axial discharge, there was a trend towards a peak lifetime around 20 kV under some running conditions. This trend is not always clear as the addition of axial discharges tends to make repeatable shots very difficult. Just sweeping RMF charge voltage from shot to shot does not produce clear (or consistent) results.

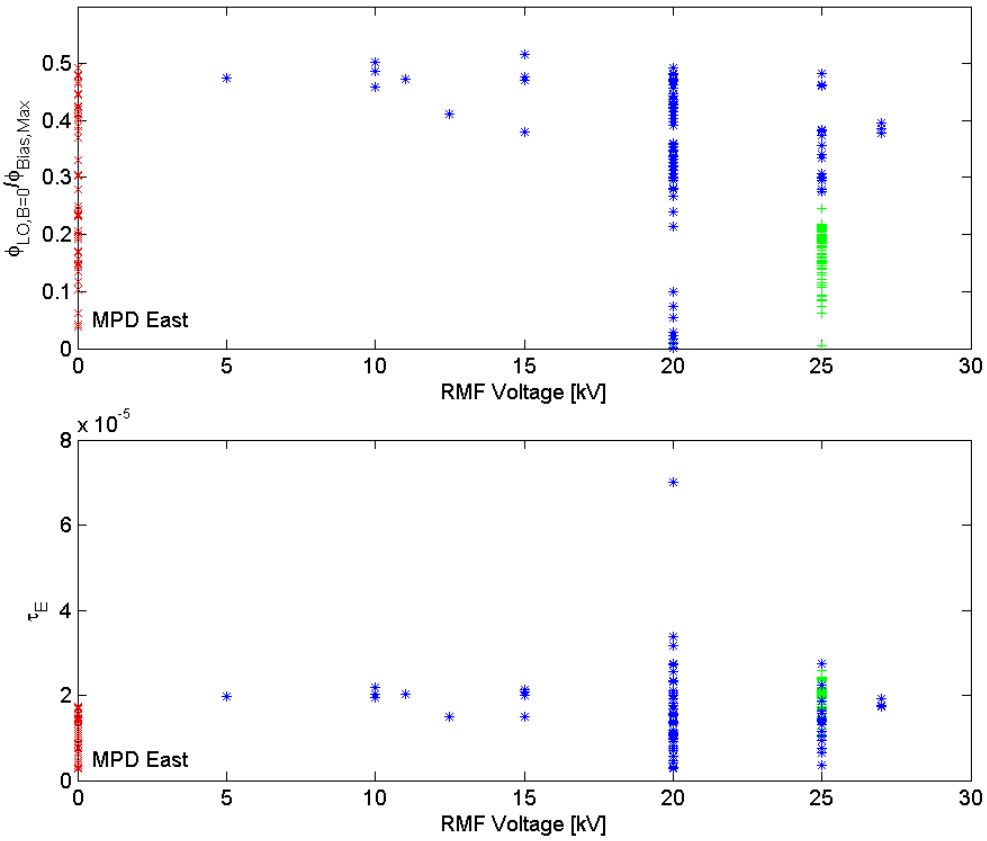

Figure 7.2.2.1. East RMF voltage optimization. Other inputs were held constant for two sweeps yielding a peak that was interpreted to be at 20 kV. The result is clearer for lifetime than flux trapping. The previous default RMF voltage was 25 kV.

7.2.3 Secondary Gas Fill

The secondary gas fill was not used extensively. This was due to both the understanding that it would not be ideal for the next phases and the fact that using it did little if anything to improve the parameters that were dwelled on the most. Primarily the system existed to provide diagnostic gasses or for cleanup (occasionally helium). Usually, significant quantities of helium persist in the vacuum system for many shots with pure deuterium.

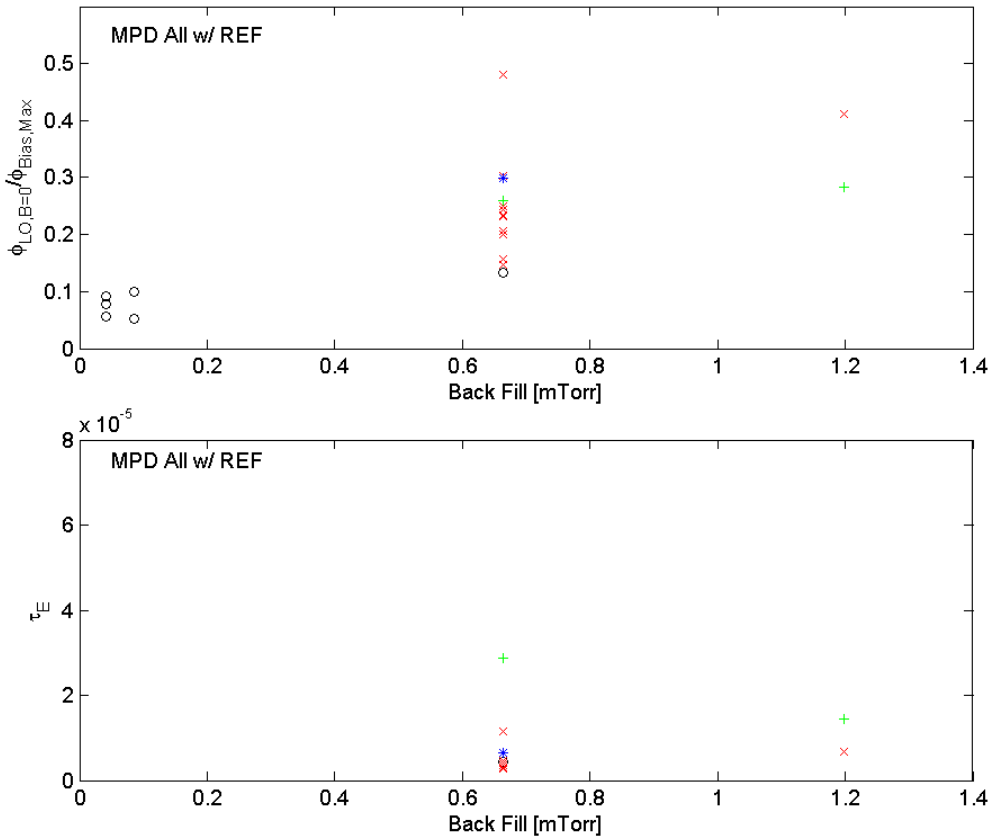

Figure 7.2.3.1. Gas back fill effects on trapped flux and lifetime. The back fill can slightly improve flux trapping under some circumstances (i.e. MPD and MPD + RMF). It probably makes energy lifetimes worse.

7.2.4 Discussion

While only a few parameter optimizations have been shown, the scope of the problem should be clear. In addition to optimizations for ionization, optimizations have taken place for reversal bank timings and for the adjustable variables associated with the MPDs. Additional objectives must be kept in mind. Simply being able to trap more flux, or achieve longer lifetimes, is insufficient. The plasma produced must also be as free from impurities as possible. The compact toroid must be made free from defects (such as excessive ion rotation) that will result in failure of later stages of the experiment.

When any part of the system operates inconsistently, optimization studies become truly difficult due to the large numbers of operational variables involved. When done with the addition of non-repeatable ionization, such as that provided by the axial discharge, all optimizations become very time consuming indeed. If gas generation from metallic surfaces is involved, it is likely that the time taken between shots will matter. Instead of attempting fully convincing optimizations, multiple theories of operation are formed and progress proceeds on their basis so as to require as few shots as possible.

The concurrent optimizations could very easily result in a slanted view of the dependencies. Plotted are 'see through' slices of a high dimension function flattened into two dimensions (or three as we separate out active ionization sources). Sometimes this approach provides clearer impressions as opposed to the local optimizations that take place with as few shots as possible. The price is that not every part of this multidimensional function is evenly sampled. For example, it is possible to get the wrong impressions due to the variance of parameters that tended to get varied together. Some effort has been expended on the preceding plots to insure that this is not the case. However, this warning is most definitely relevant for the next section.

7.3 Run Parameters

Other reports tend to focus effort on comparing experimental results versus semi-empirical scaling laws. While the same is done here, additional figures are presented in this section intended to give a complete view of the operating regime achievable by the experiment in its current configuration. As such, most of the interesting parameters are included extensively graphed to expose potential relationships. Without some associated discussion, this makes for rather dry and uninteresting plots. So instead, we approach the presentation by considering the various questions regarding operation and displaying appropriate graphs. Many alternative explanations will be omitted, which is part of what makes the display of extra data interesting.

The operating regime, as the ionization situation currently stands, is somewhat limited by the need to produce FRCs of sufficient lifetime to be studied. Some of the boundaries are due to operational limits (deficiencies in ionization for example) and not some underlying physics. Quantities which are dependent on the interferometer (such as temperature) frequently have a large amount of spread in the data. Many of the shots (especially after the addition of the axial discharge) would show negative temperatures due to apparent negative densities. Unfortunately, there is no known reliable way to process the interferometer data to remove the high error data points in a bias-free way.

7.3.1 LHDI Scaling

Hoffman and Slough[93] used a combination of classical and LHD (lower hybrid drift) resistivity to show a particle life-time scaling with $(R^2/\rho_{i0})^n$ where ρ_{i0} refers to the ion gyroradius in the external field and n refers to some power, not necessarily an integer. Additional dependencies were speculated on. Later, they applied this same approach to a combination of some of the LSX, TRX, FRX, and LSM data sets[94]. A similar scaling parameter is shown in Figure 7.3.1.1. The abscissa is the general scaling law for energy lifetime, in modified form, from the LSX scaling paper. Powers have been rounded to ¼. It has been put in a form that is dependent on s as defined earlier rather than ion gyroradius in the external field. Our definition of s incorporates a dependence on x_s (due to the rigid rotor profile constant K: $s \sim x_s r_s/\rho_{i0}$).

Examination of the apparent theoretical basis[95] [Tuszewski and Linford] for the Hoffman and Slough scaling shows that the significant scaling quantity is indeed ρ_{i0}. The scaling comes about from equating radial particle transport inside the separatrix to axial transport on external open field lines. An assumption that the lower hybrid drift instability is saturated by the "wave energy bound" yields a lifetime scaling consistent with experiments. Assuming this to be the case, with our definition of s, energy lifetime should scale as $\tau_E \sim s\, r_s\, x_s^{-0.25}$ (a different form of the scaling listed in Hoffman and Slough).

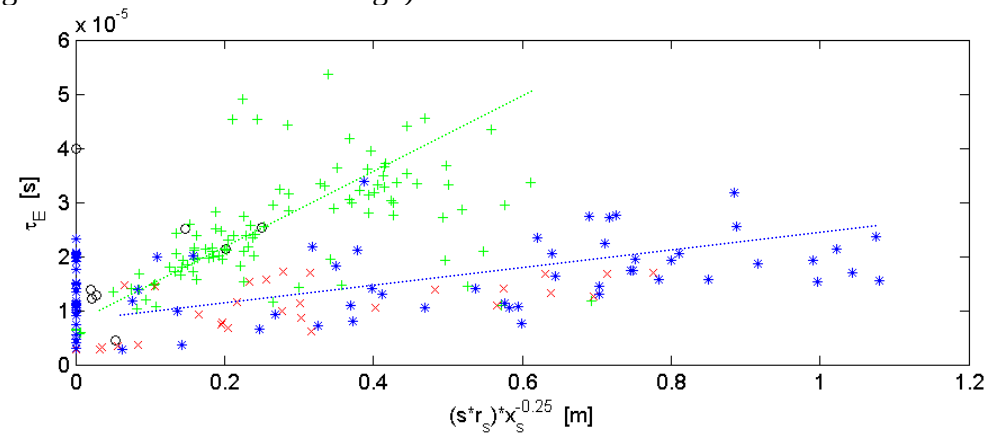

Figure 7.3.1.1. Energy lifetime scaling with lower hybrid drift. Note the dependence on method of ionization. As usual, o is MPD; + is MPD & RMF; x is MPD & ZPI; * is MPD & RMF & ZPI. Note that the majority of the pure MPD shots scale with MPD & RMF. Many of the shots on the ordinate axis are due to erroneous negative apparent densities from the interferometer.

7.3.2 Flux Trapping

One of the concerns associated with field-reversed theta-pinch operation is loss of trapped flux due to time spent with the vacuum chamber wall supporting the magnetic null currents. The positing of this question already assumes that the flux is effectively trapped by a sheath on the wall rather than assuming that all plasma to reach the wall recombines. One way to approach this question is to display flux trapped for varying levels of bias and reversal charge voltages. Another way is to plot flux trapped versus the time taken to reach lift off, and thus a measure of time spent on the wall. The expectation is that the less time spent on the wall the better. As can be seen from the figure below, there is no strong correlation between flux trapping and the time taken to go through reversal. However, there is a slight correlation between lifetime and reversal time for the cases with the MPD on the east end using axial discharge and RMF.

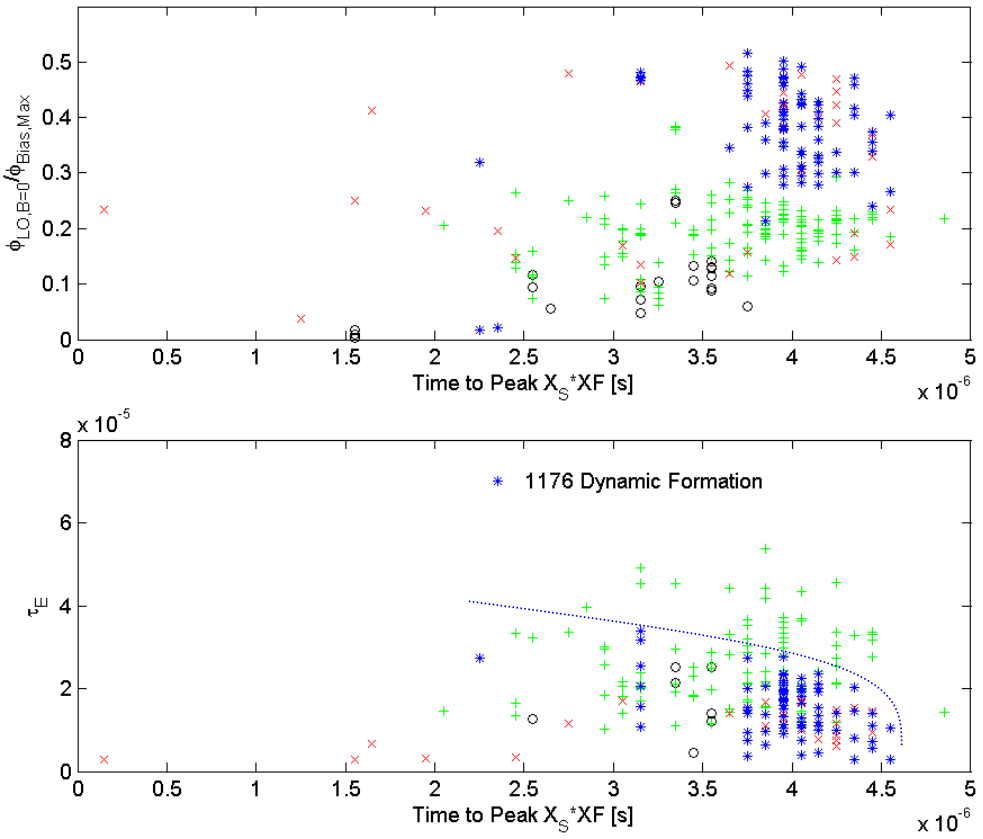

Figure 7.3.2.1. Time spent with FRC currents on the wall show little effect. (top) It is possible to trap as much flux on a long timescale as on a short timescale over the parameter range explored. (bottom) The lifetime shows what might be a dependence for the cases with axial discharges + RMF. The lifetime appears to degrade with more time spent on the wall. The dynamic formation shot should be taken with a grain of salt (as with any conditions that have only one shot). The problem will be shown later.

Another approach to the question might be attempted through the interferometer and a measure of inventory when the FRC relaxes. Unfortunately, this is reliant on an often faulty interferometer reading. Remember that the interferometer is expected to be more sporadic for shots with axial discharge due to the threading of current returns through the interferometer table. Nevertheless, there is apparently a pattern. Higher inventories tend to go with more lift-off flux trapping. One might argue that this is the result of additional ionization and gas introduction due to the discharge. Over some range of inventory with just RMF the flux trapping fraction does not appear to change but then starts to nose upwards at higher inventory.

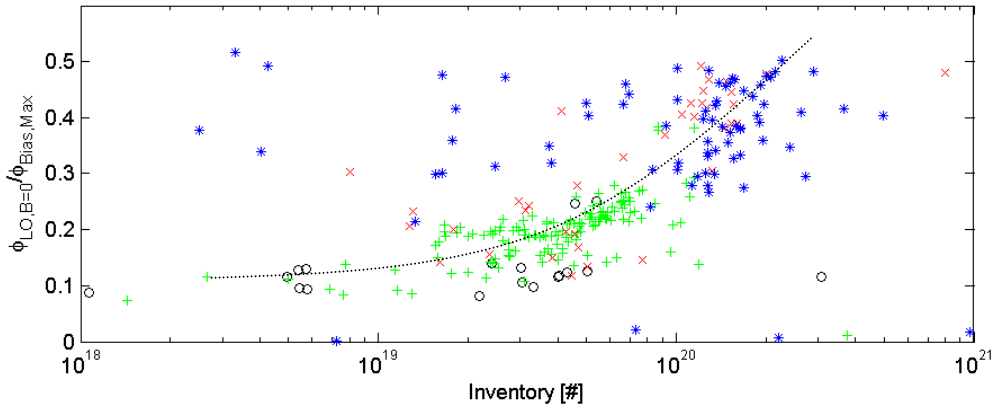

Figure 7.3.2.2. Flux trapping versus FRC inventory. The inventory listed is from the first equilibrated FRC time. Notice the trend following from the bulk of the green MPD + RMF dots through the blue MPD + RMF + ZPI. It will be interesting to see if the low inventory space is accessible with higher flux trapping once the annular segmented MPD is installed.

Lifetimes for the higher flux trapping cases using ZPI are generally poor compared to those using RMF with no ZPI. There are not many cases with minimal flux trapping using ZPI primarily due to the optimization process. Early optimizations were done with MPD & RMF. ZPI was almost always added once an initial optimum point had been found. Only for the MPD & ZPI case was a timing scan conducted resulting in cases with lower flux trapping and lower inventories.

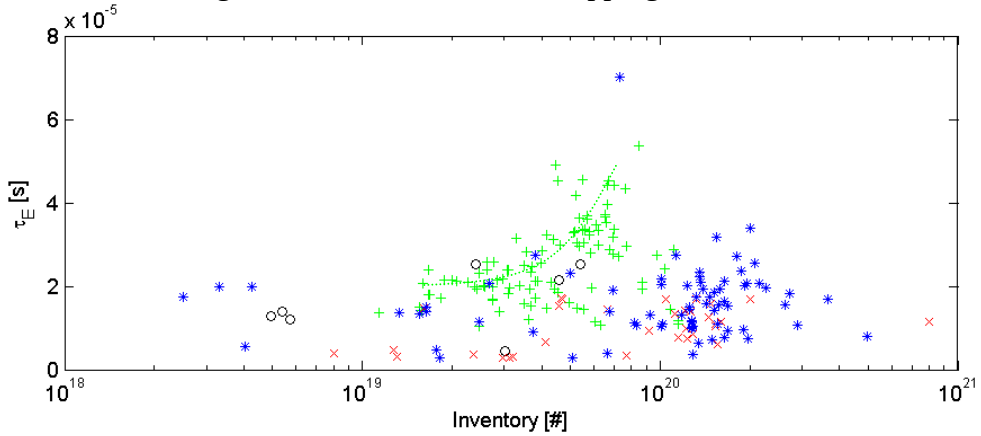

Figure 7.3.2.3. Energy lifetime dependence on total inventory. Note that maximum lifetime tends to improve with greater inventories for the MPD + RMF cases. The blue MPD + RMF + ZPI cases generally have poor energy lifetimes with no clear scaling with inventory.

Axially relaxed poloidal flux show a similar dependence on inventory. In this case the curve acts as a soft limit on maximum poloidal flux rather than defining a band trend.

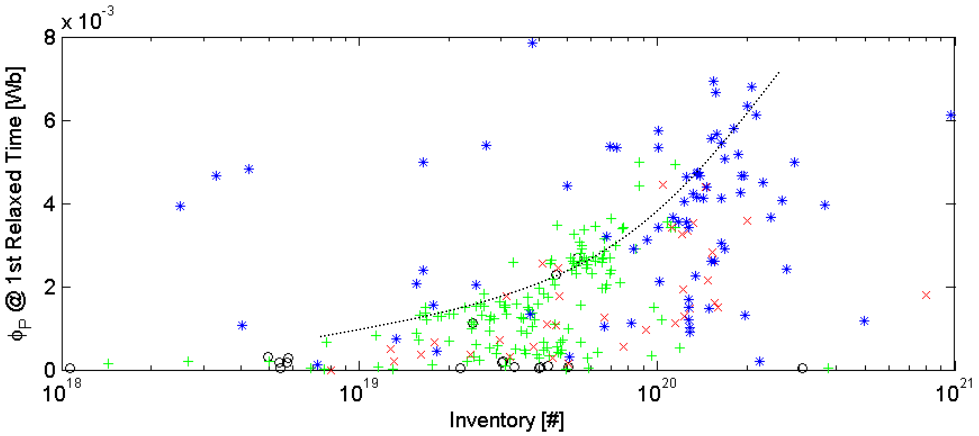

Figure 7.3.2.4. Maximum relaxed poloidal flux dependence on inventory. Many of the points showing low inventories could be due to interferometer errors.

7.3.3 Lifetimes

Inventory lifetimes show poor correlations with other measured parameters. This could be due to interferometer errors (i.e. vibration), unaccounted for ionization, or both. The quickest way to demonstrate this issue is by plotting against the other lifetimes for energy and flux.

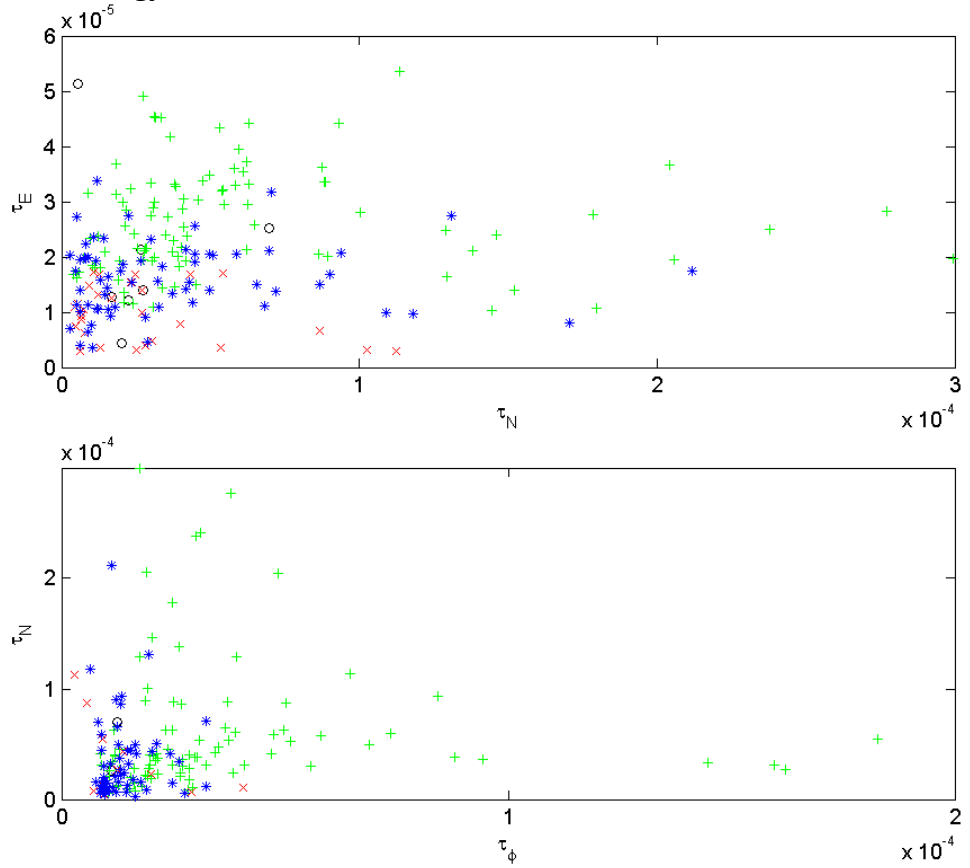

Figure 7.3.3.1. Inventory lifetime shows little correlation with other lifetimes.

We preferentially plot energy lifetimes. Energy is a well measured quantity via the magnetic diagnostic arrays. It is relatively straightforward to compensate for work done by the external magnetic fields. Poloidal flux lifetimes are also easily calculated. The lack of external inputs into poloidal flux makes it an attractive measure. However, poloidal flux is more dependent on the axial equilibrium relation and in particular the structure of the magnetic fields at the FRC ends. Poloidal flux (φ_p= 0.31 x_s XF) is dependent on the axial equilibrium average beta (<β>) for the 0.31 x_s term. Energy is dependent on the axial equilibrium average beta in a form that may be more weak: $\left(1+\frac{1}{2}<\beta>\right)\frac{B_e^2}{2\mu_0}V$. Regardless, the relation between energy lifetime and flux lifetime is fairly strong in this experiment.

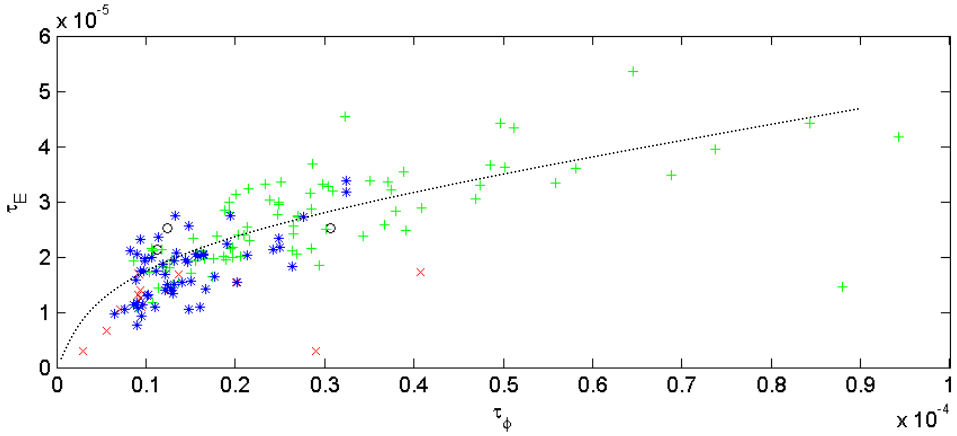

Figure 7.3.3.2. Energy lifetime dependent on flux lifetime. Note the rather strong correlation. From the scaling relations discussed earlier we would expect an essentially linear relation (with a weak $x_s^{¼}$, $r_s^{0.14}$, and $\rho_{io}^{0.07}$ difference)

Knowing that the correlation between energy and flux lifetimes is strong allows us to plot most parameters against energy lifetime alone. This is done for a wide range of parameters in the figures below. An independent power regression analysis has not yet been done, primarily due to difficulties in compensating for the errors in determining density and temperature. The purpose of these figures is to flesh out our understanding of the parameter space accessible by the current experiment, not to show any particular scaling. The dashed lines are added by hand to emphasize prominent features in the figures.

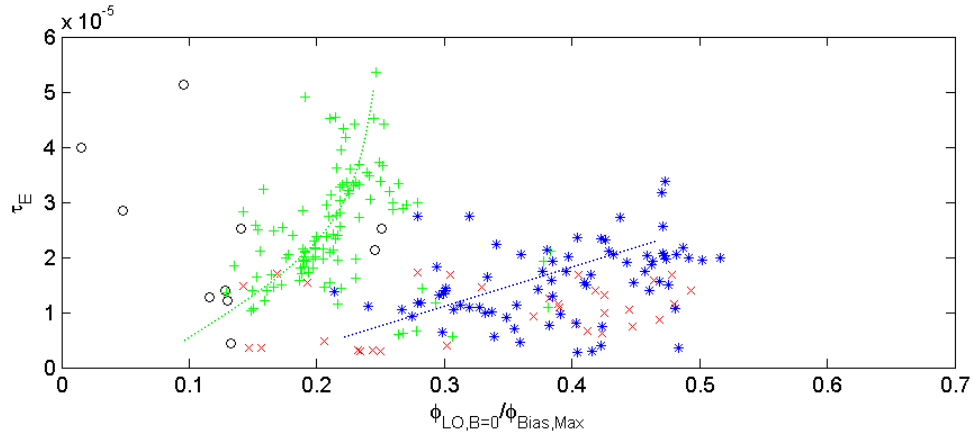

Figure 7.3.3.3. Energy lifetime dependence on flux trapping efficiency.

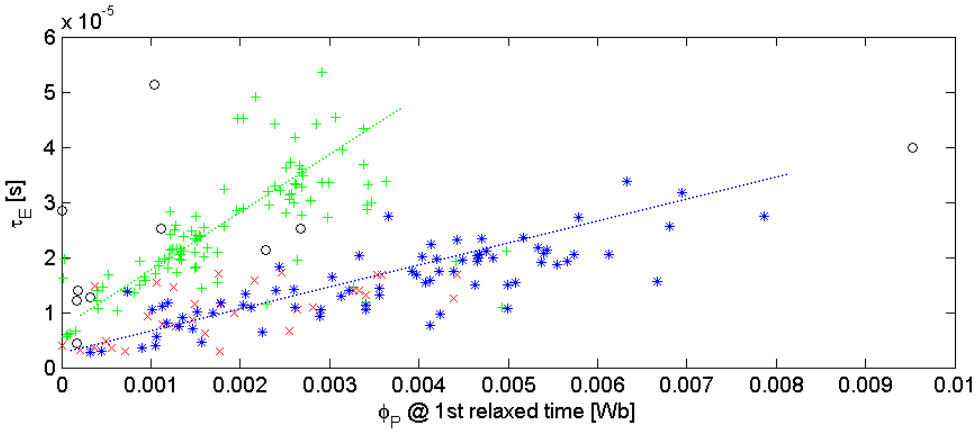

Figure 7.3.3.4. Energy lifetime correlation with relaxed state poloidal flux.

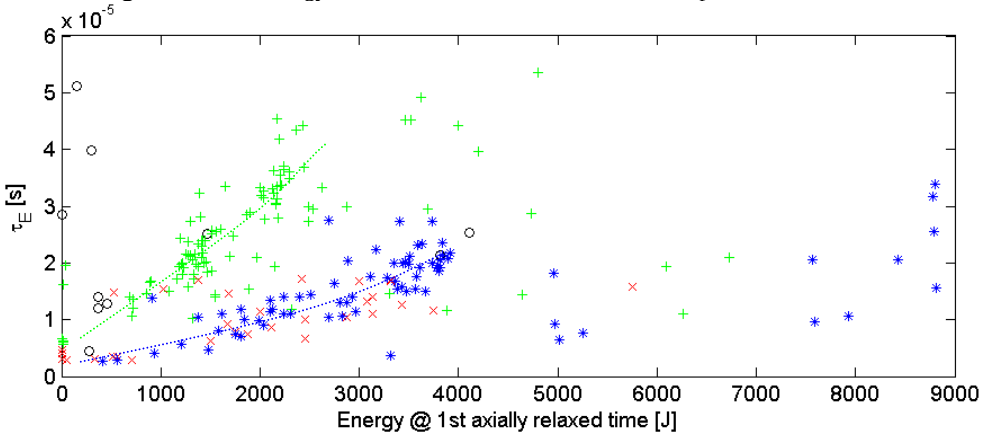

Figure 7.3.3.5. Energy lifetime dependence on relaxed state energy. The energy shown is the energy within the separatrix of the FRC.

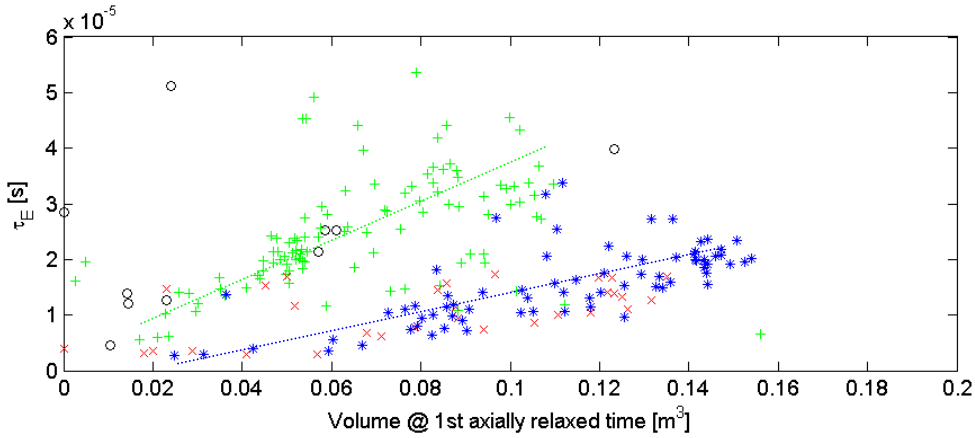

Figure 7.3.3.6. Energy lifetime dependence on FRC volume.

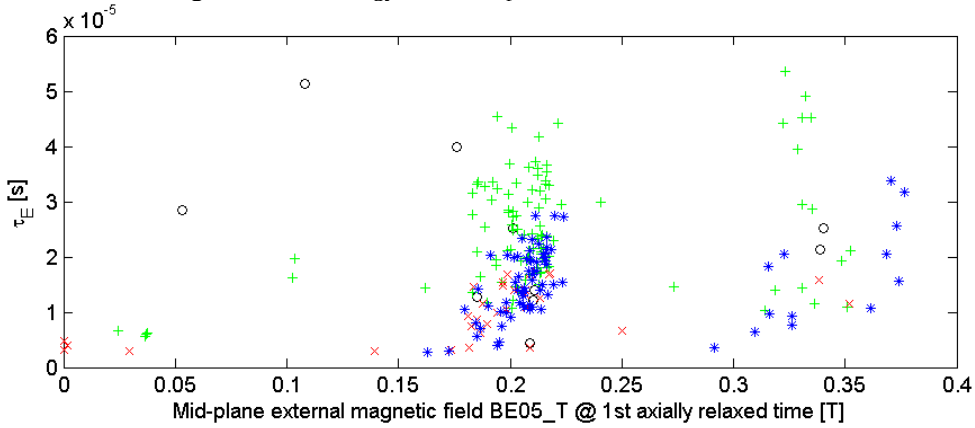

Figure 7.3.3.7. Energy lifetime and external magnetic field in relaxed state.

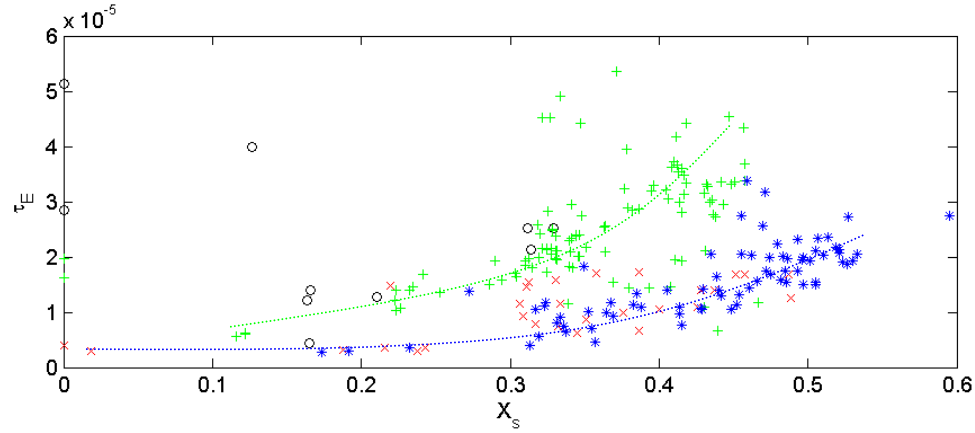

Figure 7.3.3.8. Energy lifetime dependence on x_s.

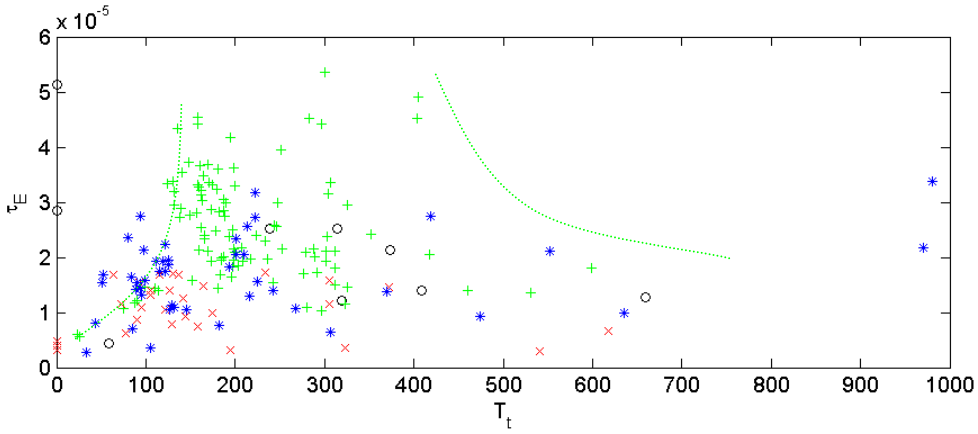

Figure 7.3.3.9. Energy lifetime and total temperature.

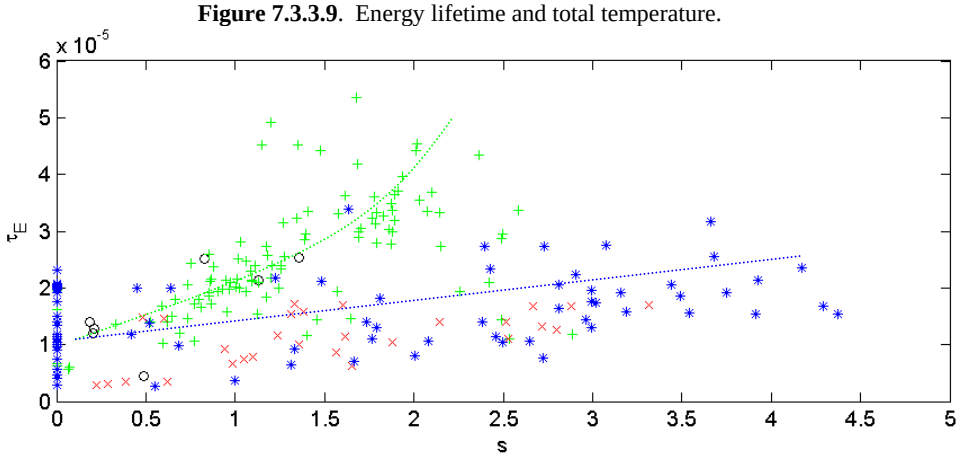

Figure 7.3.3.10. Energy lifetime dependence on the kinetic parameter.

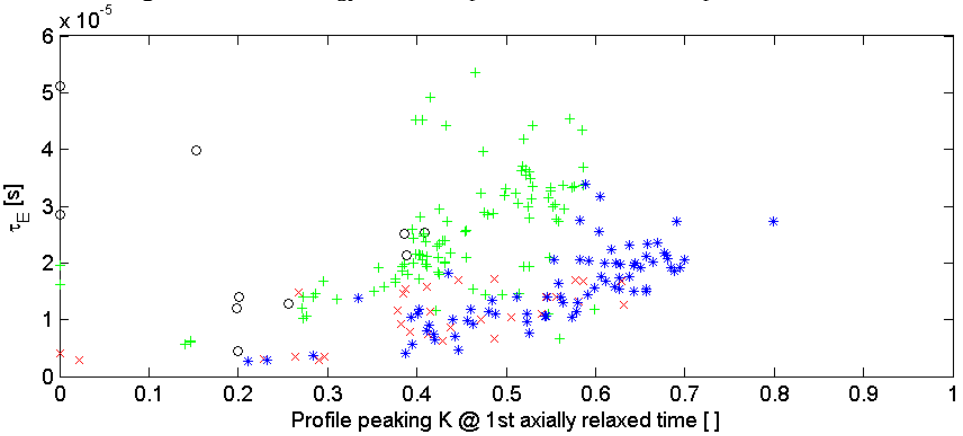

Figure 7.3.3.11. Energy lifetime dependence on radial profile peaking.

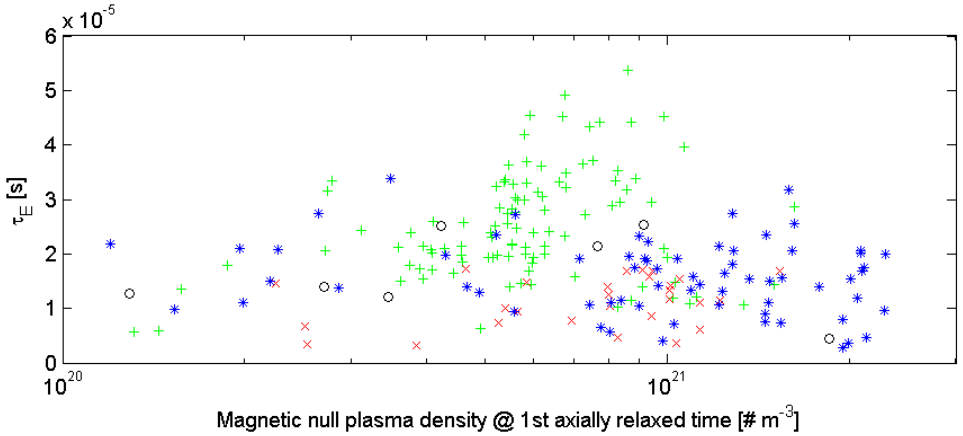

Figure 7.3.3.12. Energy lifetime and peak plasma density.

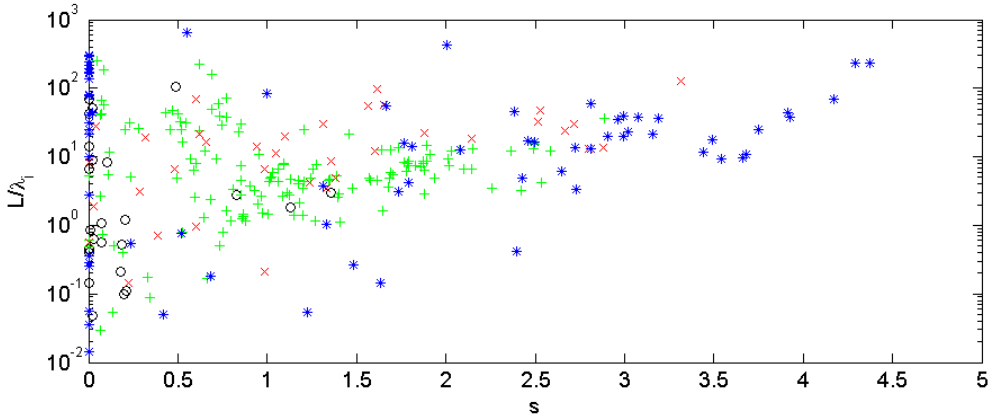

Figure 7.3.3.13. Dimensionless parameter operating regime. Length of FRC normalized by ion mean free path (taken at magnetic null) is plotted against the kinetic parameter s.

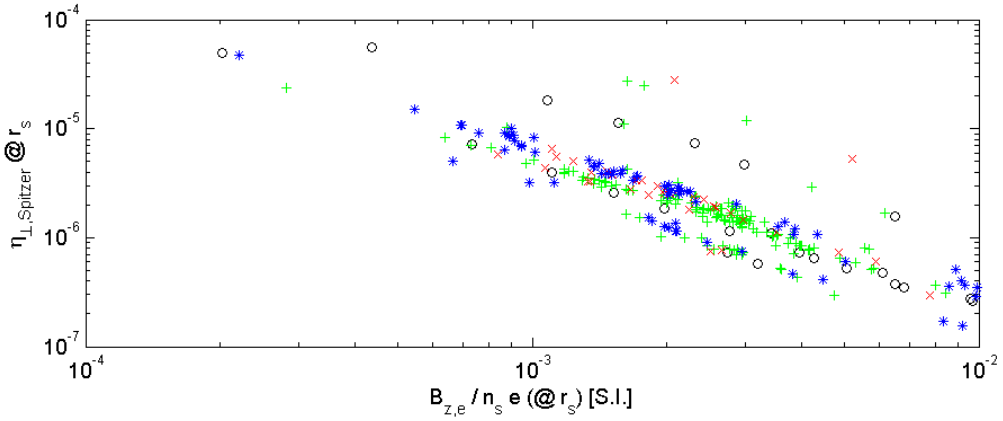

Figure 7.3.3.14. Spitzer resistivity at separatrix compared to Hall term.

Unfortunately, further analysis will have to wait for more time to become available. Despite this lack of further commentary on the plots, we hope that this dissertation is richer for having included them rather than leaving them out.

7.4 Select Shots

Several shots are worthy of a closer look as exemplars of the dataset. The two general ways of operating the reversal coils can be characterized as 'stationary' timings and 'dynamic' timings. The majority of shots are subcategories of stationary timings. In particular these shots include different end region timings designed to either speed the axial contraction or delay it. The more repeatable shots tend to have timings designed to slightly speed the axial contraction by firing the end region reversal fields slightly before the center sections. Our primary criterion is in reliably forming FRCs. This set of timings is probably not ideal for forming the best possible FRCs. Stationary timings are symmetric. Dynamic timings fire each coil in sequence. The difference in times can vary, but too large a Δt results in no FRC.

Unfortunately, except for the magnetic diagnostics, most of the current diagnostics do not provide that much useful and reliable information. This has forced us to analyze general parameters without being able to explain much of what is going on as far as internal physics are concerned. Imaging, of various forms, played a much bigger role in understanding the operation of the 40 cm experiment than the 80 cm version. This was due to the high visibility of instabilities. The optical access was much cleaner with a clear 40 cm radius quartz vacuum boundary.

The end-on viewing problem is demonstrated in the figure below. For the case of interest from a programmatic perspective very little is distinguishable. However, for the clean-up shot running with helium much is visible but the compact toroid's parameters are not so good.

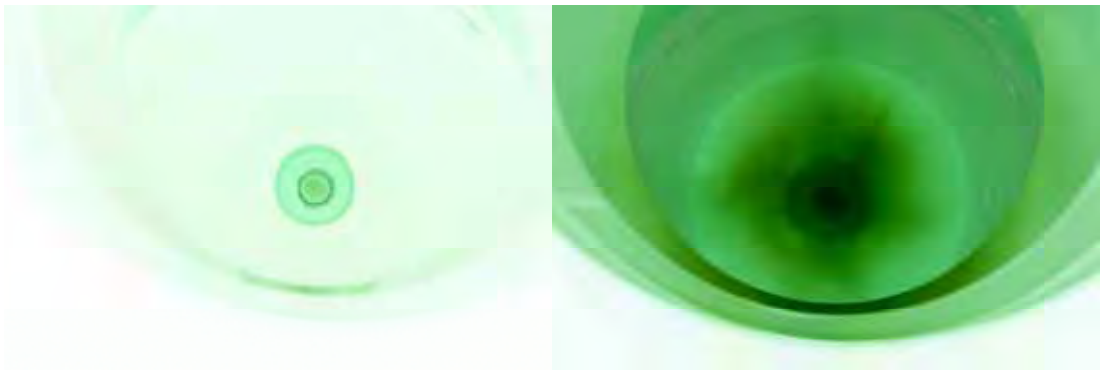

Figure 7.4.1. Visible end on imaging negatives. (left) shot 913. As a rule, most shots look like this with no distinguishable separatrix from the background light. Many reflections of the MPD are present. Enhancement of the image yields nothing useful. (right) shot 760. This shot was a special stationary shot with a helium backfill and only the MPD ionization source. The reversal bank was allowed to ring without crowbar to clean up the walls of the system. In general, the VLA detectors have the same basic problem. More light is produced during ionization processes than is produced in the bulk of the plasma. Hopefully, this will no longer be the case during compression.

Information from a survey spectrometer was available for many sporadic shots (but not #917, the preferred 'stationary' shot). The information that was gained requires interpretation; many of the conclusions are ambiguous. The most important was the lack of a recognizable carbon V line, indicating that electron temperatures are

significantly less than 300 eV. The total temperature for shot 929 was 400 eV. Since electrons are not mono-energetic, time resolved line ratios would be required to obtain a more definitive result. For the dynamic shot, the survey spectrometer yields a C IV line at 580.3 nm and a C II line at 589.0 nm. The ratio is 2400 to 6200 counts for shot 1168 and 3955 to 3108 for shot 1166. These were two dynamic shots. The carbon IV line requires 65 eV. There is a good chance these lines are only present after the translating FRC impacts an end cone.

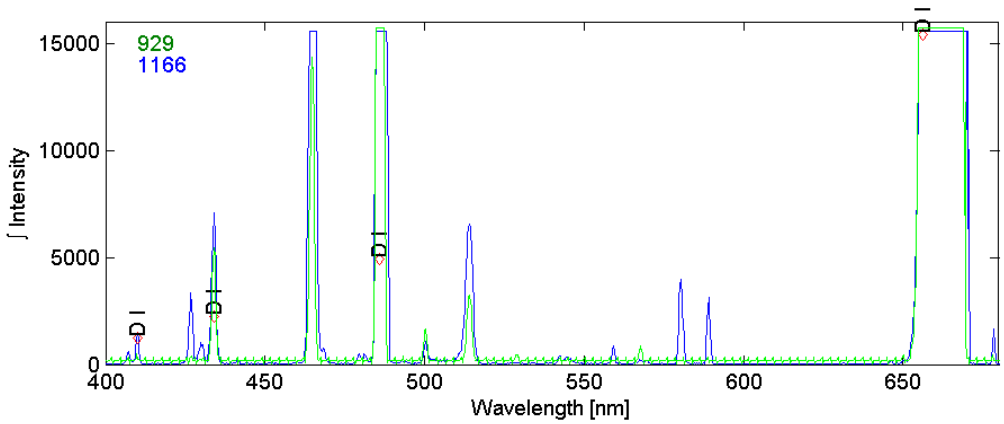

Figure 7.4.2. Visible spectrum for 'stationary' timing shot and 'dynamic' timing shot. Most of the lines can be explained by the usual suspects CI CII CIII CIV OI OII OIII. Several lines require more exotic explanations such as a line hiding at 670.9nm within the limb of a DI line. The short list of somewhat less unlikely candidates are Ti III, Li I, Cm I, Hg III, La I. Another line at 500.3nm has somewhat more likely candidates: Ti I, N II, F II, K II. Additional weak lines show Ti, K, and N as candidates. Potassium reappears as candidates in many of the weaker lines, often with no other reasonable candidates. Potassium carbonate is commonly used in the making of chemistry and tempered glasses. No elements, other than deuterium, show their full expected spectra. CIV, CIII, OIII, CI, and CII are, in order from closest to furthest, closest to having full spectra. There is no CV line at 494.4 nm. The lines are identified using a downloaded NIST database[96].

An indirect measure of ion temperature, from the shape of an impurity line, was never successfully performed on the 80 cm experiment. The only usable data on the 16 channel spectrometer took place on the last day of running, with the experiment operating under the poorest conditions. Sufficient signal strength on that day was obtained after the FRC impacted an end cone. As such, this data is not of high interest.

The results from the magnetic diagnostics are generally much more reliable. It is quite fortuitous that FRC equilibriums are highly constrained by external magnetic field parameters or the previous parameter space plots would not have been possible. In general, we are unable to efficiently generate axially stable (in the sense of not translating) FRCs. Inevitably, an FRC comes in contact with the axial magnetic field gradient associated with our end coils and accelerates away from center

The magnetic diagnostics themselves are somewhat unreliable in their readings near the ends. The spatial gradient becomes too large for the vacuum reference techniques to accurately compensate for. Additionally, the vacuum reference technique is wholly dependent on having accurate vacuum references that

perform similarly to the plasma shots. With a system that frequently has bank issues this is highly problematic. This explains the appeal of the system demonstrated (and employed on two shots) in the diagnostics section. One shot, one solution rather than two similar shots yielding one solution. However, this technique takes too long to complete (with the accuracy set to a level commensurate with digitizer accuracy) for general usage (~ 8 hrs for a sparse 100 time points per shot). The addition of accounting for resistance in the driven coils makes the analysis take even longer. The same basic problem is anticipated for using plasma numerical solutions when some of the inputs (read ionization distribution) are not available.

As can be seen in Figure 7.4.3, most of the formed FRCs are of small x_s and varying levels of axial and radial oscillation damping. It is likely that the FRCs interact with the banks quite significantly as part of their radial oscillations. In at least one case, the FRC manages to ultimately come back from the cone region. It is uncertain why this should happen. There is an outside possibility that it may be a ghost from the vacuum reference process.

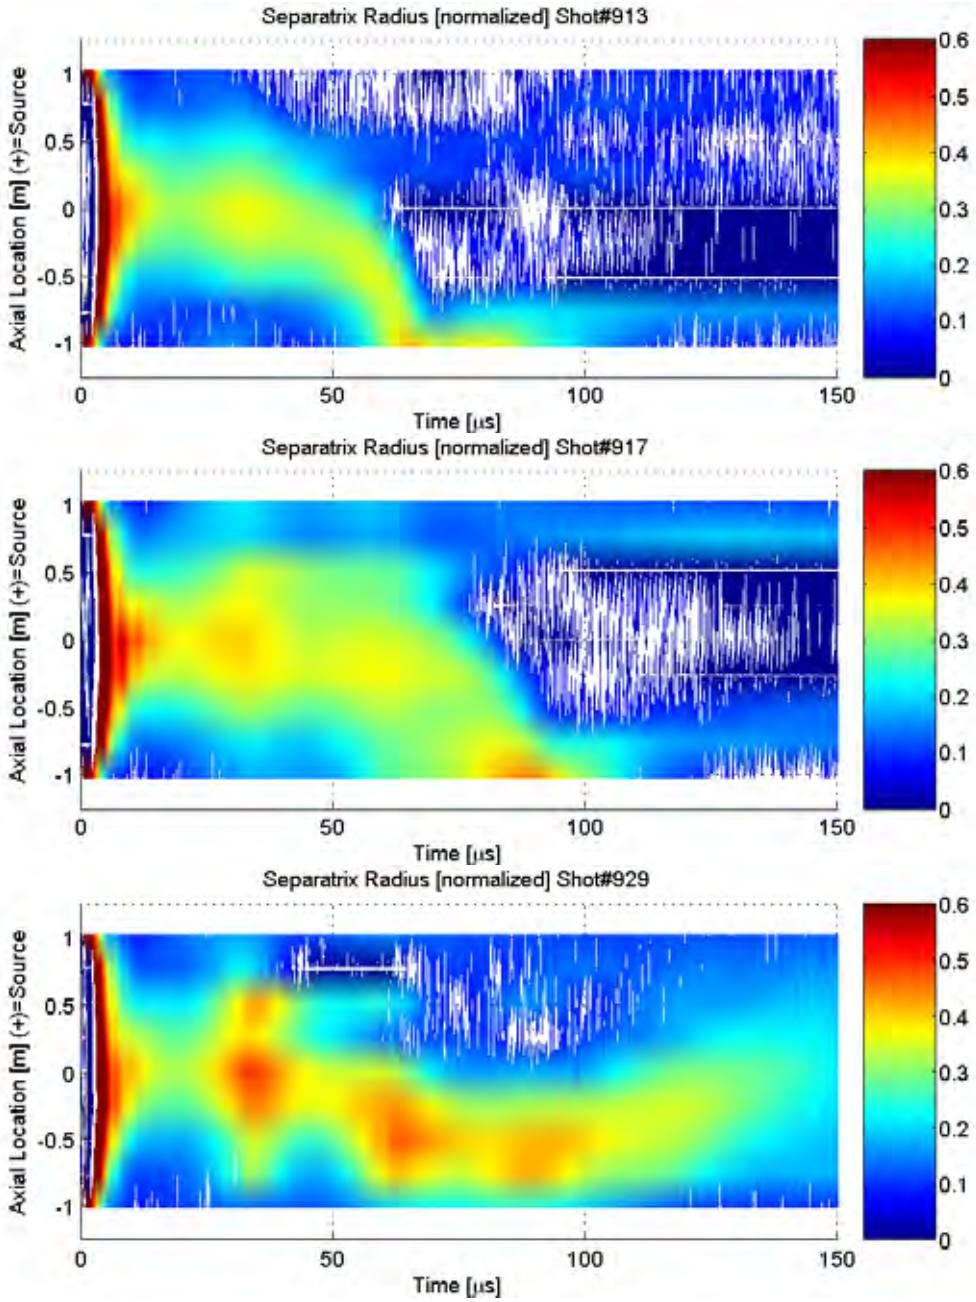

Figure 7.4.3. Stationary FRC separatrix radii. Color is indicative of x_s; hotter colors indicate larger radii. Note the varying levels of oscillation damping. Much of the oscillation is due to varying performance of bank crowbars. Shot 917 is also described in the magnetic diagnostics section.

The dynamic formation shots were done as a small sequence of 11 shots during the last stages of experimentation. In part this was done to see if we could get any shots to bounce off the end cone and stay intact despite the transition to a short, and lower field, region. It was also done because the stationary results with the

copper center MPD version were very disappointing. There was never any real danger of the translating FRC reaching the MPD.

The primary parameter to choose is Δt between reversal coil firings. In shot 1168 below the timing is too long and the results are disappointing. Shot 1166 does not represent an optimized condition, merely one of the better conditions before efforts were discontinued. In it, the separation of the FRC intact from formation is clear and something appears to persist in the end regions for some time.

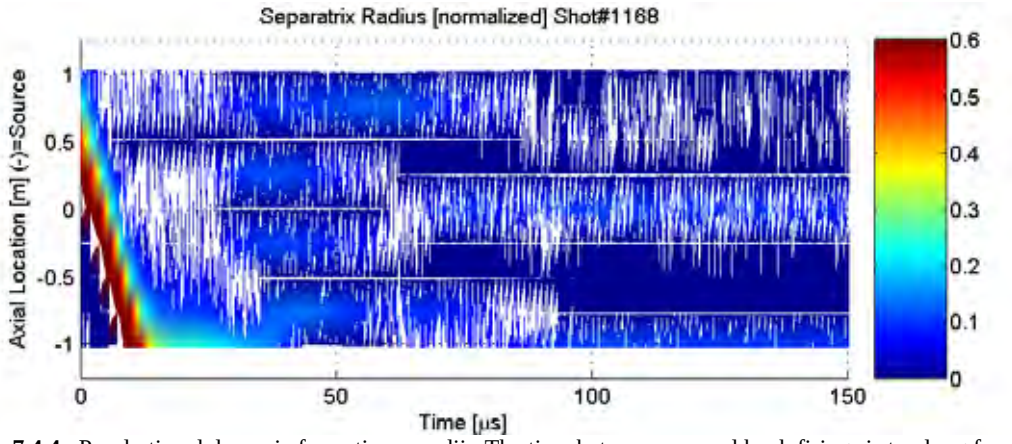

Figure 7.4.4. Poorly timed dynamic formation x_s radii. The time between reversal bank firings is too long for proper FRC formation.

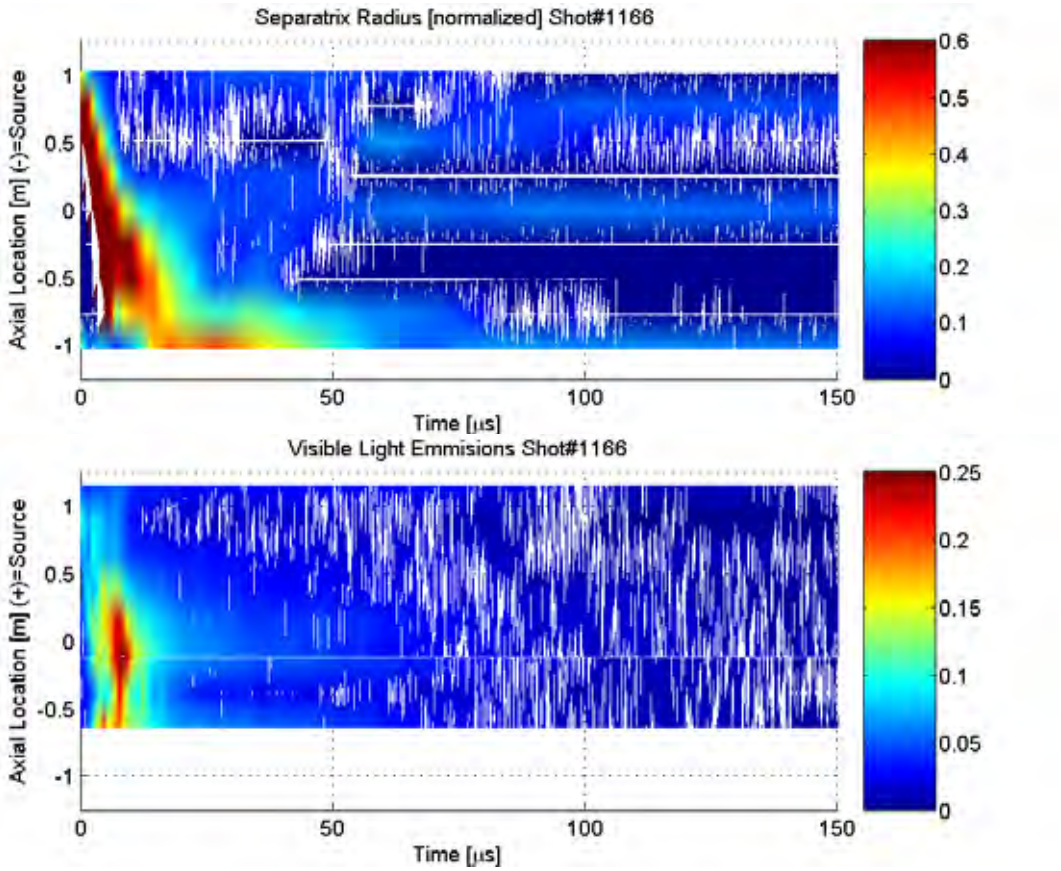

Figure 7.4.5. Well timed dynamic formation. (top) x_s radii can be seen to persist after detachment from the drive coils. (bottom) The VLA array is capable of monitoring the ionization caused by dynamic formation. The end cone on the source side (east) is not well instrumented so as to allow for extra VLA sensors arrayed azimuthally (looking radially) near the mid-plane (where the discontinuity is). Axially looking VLAs were arrayed at this time to monitor a fixed radius at 4 azimuthal locations.

We have chosen one last shot to demonstrate the hazards of relying on vacuum reference subtraction. This is the '1176 Dynamic Formation' shot clearly marked in one of the earlier figures. In the figure below, the separatrix appears to oscillate in the end region. In reality, the explanation is much simpler; the crowbar for one drive coil failed to come fully into conduction in a timely manner. This is included to demonstrate the hazards of relying on vacuum references. Fortunately, it is only a single shot that misbehaved; easily identifiable in all the previous plots by it's standing alone in parameter space.

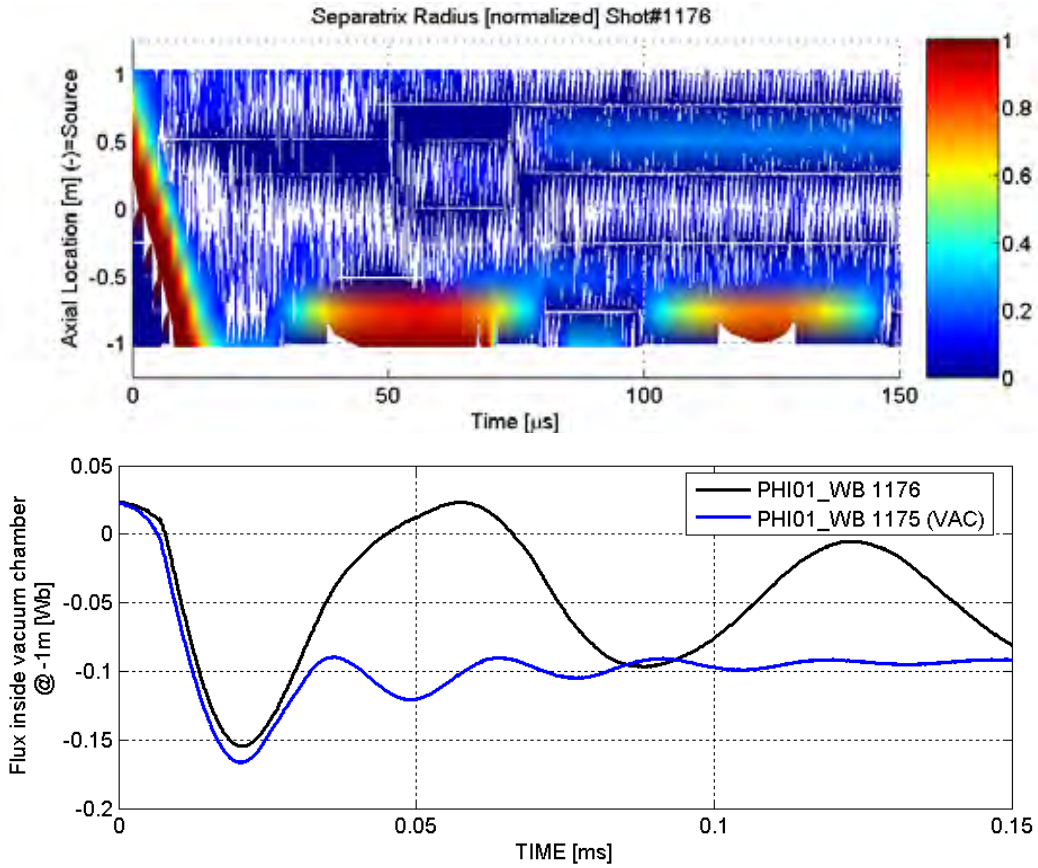

Figure 7.4.6. When vacuum shots don't match plasma shot behavior. (top) The very large separatrix radius oscillations apparent on an end coil are due to a crowbar failing to fire for the plasma shot. (bottom) The difference in flux on the coils for plasma versus vacuum is made clear here.

7.5 Finale

This document has discussed most of the design, analysis, and operational details of the experiments conducted in the LSX (80 cm) sized tubes at the University of Washington to date. Special emphasis has been paid to including critical necessary details of design. For example, homo-polymer polypropylene with fiber reinforcement was not the first attempt made by our predecessors at designing an appropriate joining section between fused silica tubes. Other material choices suffered from a host of issues.

Many of the details are no longer directly applicable (i.e. many specific parts are no longer available on the market). However, it is hoped that the discussion has been thorough enough for adequate replacements to be chosen or new designs made with awareness for the various issues. The discussions are intended as a starting point and are not fully comprehensive.

Emphasis has been placed on various safety issues in a manner that is not ordinarily done in thesis work. This is due to the tendency for safety to be last in order of program priorities. It is possible on smaller experiments to have very infrequent contact with high voltage equipment. On an experiment as large as this

one, maintenance and inspection is a regular requirement. Also, a great many things have (and can be expected to) go wrong. Hence there are a great many more opportunities for contact with potentially dangerous equipment. For this reason, additional levels of redundancy have been added and additional emphasis placed on understanding how things can go wrong.

The take-home message: In general, the national electric code is not something to be ignored; as much as it would simplify design choices to do so. Design changes without thinking through all the ramifications can be dangerous; as can jury-rigging systems and tampering with safety interlocks. The same is most emphatically true for operating half complete systems. However, as with so many other things, this thesis presents many of the issues and possible solutions but is not intended to be the only education for someone working with high voltage systems. Basically, don't try this at home (and don't sue us, ever).

Most of the emphasis in this document has been placed on the design and construction of the experiment and its associated diagnostics. This is because, done properly, these are time-consuming (and alas, never-ending) tasks. All of the diagnostics operate with various caveats attached. The external magnetic diagnostics are subject in their interpretation to magnetic field spatial variation constraints. The spatial variation constraints can be relaxed with the use of time consuming analysis techniques. Further analysis is contingent upon the applicability of rigid rotor current profiles, fluid-like plasma behavior, and a host of assumptions regarding the remaining plasma parameters.

The interferometer has problems due to continuous acoustic couplings that result in large error bars. The results are difficult to evaluate except in statistical terms (hence the reliance on all the shots rather than focusing on individual shots). The interferometer and remaining systems require a great deal of maintenance and active attention and are thus sporadically available. Most of the operator's time was in fact consumed in keeping the power systems (especially the ignitrons) functional.

Despite all these issues, useful results certainly were obtained and are presented. The apparent scaling of the energy loss timescale with the lower hybrid drift scaling is encouraging. The dependence on ionization technique is even more so. This result holds out the hope that improving ionization by providing an annulus of plasma at the peak of the bias field via segmented MPD guns will significantly improve results both in terms of flux trapping and potentially lifetime. Enough scaling graphs have been provided to keep someone occupied with analysis (and intrigued by it) for a long time to come. Examples of a few operating conditions have been presented to make the general mode of operations clear.

It is hoped that this dissertation will provide a useful reference for many researchers in the future. Such is all that can be hoped for of any thesis.

END NOTES

Note: Many of these references are better identified in the bibliography section. A web search should be sufficient to bring up most of the sources with the keywords listed herein. Many sources are not widely distributed and/or published. Some require access to DOE reports.

[1] Peter, M. *PHD Thesis.* University of Washington, Aeronautics and Astronautics. 2003.
[2] Gota, H., et al. Visible Bremsstrahlung tomographic diagnostic for the pulsed high density field-reversed configuration experiment. *Rev. Sci. Instrum*. **77**, 10F319 (2006).
[3] Hoffman, A.L., Slough, J.T. Field Reversed Configuration Lifetime Scaling Based on Measurements from the Large s Experiment. *Nucl. Fusion* Vol. 33, No. 1, (1993).
[4] Mathematical Sciences NorthWest, Inc. *Compact toroid generation, lifetime, and stability studies in linear reversed-field theta pinch geometries (TRX-1): Second annual and final report*. U.S.DOE DE-AC06-80ER53096, C.12.144.01 (1983).
[5] Spectra Technology, Inc. *Compact toroid generation, lifetime, and stability studies in linear reversed-field theta pinch geometries (TRX-2): final report*. U.S.DOE C.12.144.07 (1987).
[6] Tuszewski, M., Barnes, D.C., Rej, D.J., Taggart, D.P. Stability and compressional heating of large field-reversed configurations in the FRX-C/LSM device. *Phys. Fluids B* 3, 2205 (1991).
[7] Ono, Y., Matsuyama, T., Umeda, K., Kawamori, E. Spontaneous and artificial generation of sheared-flow in oblate FRCs in TS-3 and 4 FRC Experiments. *Nucl. Fusion* 43 (2003) 648-653.
[8] Nothing was ever published on the attempts made to create an in-situ FRC in the formation chamber of TCS (LSX/Mod) by standard theta pinch reversal. The attempts failed; therefore no attempts were made at publication. Nevertheless, translating FRCs were readily generated.
[9] Persson, Hans, Ågren, Olov. Vlasov solution describing nonadiabatic plasma heating by magnetic compression. *Phys. Fluids B* 5 (7), July 1993. 2187.
[10] Spencer, R.L., Tuszewski, M., Linford, R.K. Adiabatic compression of elongated field-reversed configurations. *Phys. Fluids* 26 (6), June 1983. 1564.
[11] Tuszewski, M. Review Paper: Field Reversed Configurations. *Nuclear Fusion,* Vol.28, No.11 (1988). 2033-2092.
[12] Spectra Technology, Inc. *LSX quarterly technical progress report No.3.* U.S.DOE (1987). Appendix 3A: Plasma Vessel Tests in Components RDAC.
[13] Massey, Liesl K. *Permeability Properties of Plastics and Elastomers: A Guide to Packaging and Barrier Materials 2nd Ed.* William Andrew Inc., 2003. ISBN 1884207170.

[14] Composite Agency - Analysis & Simulation Expertise. -vapour- gas water diffusion permeability polymers WVTR moisture. http://www.composite-agency.com/vapour_diffusion.htm. Accessed 2007.
[15] Composite Agency - Analysis & Simulation Expertise. vapour –gas- water diffusion permeability polymers WVTR moisture. http://www.composite-agency.com/gas_diffusion.htm. Accessed 2007.
[16] Professional Engineering Computations (PEC, Inc.). Outgassing and Permeation, http://lpc1.clpccd.cc.ca.us/lpc/tswain/lect8.pdf. CA. 2003.
[17] Permeability and Plastic Packaging. Accessed 2008. http://www.personal.psu.edu/irh1/PDF/Perm%20and%20Pack.pdf pg 19
[18] FDA. Surface lubricants used in the manufacture of metallic articles. *CFR Title 21*, Volume 3, section 178.3910 21CFR178.3910. Apr 2007. Available on web: www.accessdata.fda.gov.
[19] Apiezon. Apiezon L datasheet. 2007.
[20] From conversations with the UW Aerospace department's electrical engineer.
[21] Hibbeler, R.C. *Mechanics of Materials (3rd edition).*
[22] W. G. Dunbar. AD-A204 504 AFWAL-TR-88-4143 Volume II. *Design Guide: Designing and Building High Voltage Power Supplies.* Interim Report August 1988, now available through NTIS. Figure 39, pg 91.
[23] W. G. Dunbar. AD-A204 504 AFWAL-TR-88-4143 Volume II. *Design Guide: Designing and Building High Voltage Power Supplies*. Interim Report August 1988, now available through NTIS.
[24] *2008 Annual Book of ASTM Standards*, ASTM D 495-99.
[25] W. G. Dunbar. AD-A204 504 AFWAL-TR-88-4143 Volume II. *Design Guide: Designing and Building High Voltage Power Supplies*. Interim Report August 1988. now available through NTIS. Pg 4
[26] *LSX Construction Final Report*. November 1990.
[27] *LSX Construction Final Report*. November 1990. pg 48, 53.
[28] *LSX Quarterly Technical Progress Report No 16.* May 1990- July 1990. pg 12. Arcing between experiment ground and flux loop grounds.
[29] *LSX Construction Final Report*. November 1990. Figure 4-18, pg 56.
[30] Kerns, Q., Kirsten, F., Winningstad, C. *Pulse Response of Coaxial Cables.* Rev. February 12, 1964.
[31] W. G. Dunbar. ADA204503 AFWAL-TR-88-4143 Volume I. *Design Guide: Designing and Building High Voltage Power Supplies*. Interim Report August 1988, now available through NTIS. pg 43.
[32] Zeller, H.R. Cosmic Ray Induced Failures in High Power Semiconductor Devices. *Microelectron. Reliab*. Vol 37, No 10/11 pp 1711-1718. 1997.
[33] Not sure if the chamber grounding issue of Microthruster ever got published as such. Historically, there was a problem with the first electric thrusters due to the lack of a neutralizing electron emitter. Lab tests did not need an emitter to show thrust.

[34] Cummings, D. B. "Chapter 9. Pulsed Power Ignitron Switches", LLNL. After 1963. This reference is useful for its historical perspective and breadth. Also has good diagrams.
[35] Gruber, J.E. *Properties of Ignitrons for 20 KV Capacitor Discharge Circuits*. August 1967. Institut Fur PlasmaPhysik.
[36] Kristiansen, M., Honig, E. *Workshop on High Current, High Coulomb Ignitrons*. Texas Tech & USAF Armament (LANL). April 21, 1987.
[37] Cummings, D. B., Ignitron Discharge Growth During High-Current Pulses, 1962 *IEEE*. (The copy we have is unclear, possibly preprint).
[38] General Electric Company, Tube Products Department., *Ignitrons: Application Notes, Capacitor Discharge and Crowbar Service*, November 1974.
[39] EEV Lincoln, *Ignitrons for Capacitor Discharge and Crowbar Applications*, (believed to be 1980s).
[40] Stanway, M., Seddon, R. *Ignitron Firing Time and Jitter*.
[41] Seddon, R. *The Use of Ignitrons at High Voltage*. July 1987.
[42] Klotzbücher, Thimo, et al. Biodegradability and groundwater pollutant potential of organic anti-freeze liquids used in borehole heat exchangers. *Geothermics* 36 (2007) 348-361.
[43] Cummings, D. B. "Chapter 9. Pulsed Power Ignitron Switches". LLNL. After 1963. pg 22
[44] Cummings, D. B. "Chapter 9. Pulsed Power Ignitron Switches". LLNL. After 1963. pg 7
[45] Kleppa, O. J. *The Solid State: Diffusion in Metals and Alloys*. The University of Chicago, Chicago, Illinois. This reference can be considered a general solid state diffusion reference (almost a review paper).
[46] Zajec, Bojan, Nemanič, Vencenc. *Hydrogen bulk states in stainless-steel related to hydrogen release kinetics and associated redistribution phenomena*. Institute of Surface Engineering and Optoelectronics, Ljubljana, Slovenia.
[47] Bochkov, V.D., Dyagilev, V.M., Nicolin, S.V., Ushich, V.G., Seeteekh, D.S., Vlasov, P.N. The Pseudospark Switch Crowbar Unit – High Reliability, Low Cost System. *IEEE* 1999. Again, possibly preprint.
[48] Bochkov, V.D., Dyagilev, V.M., Nicolin, S.V., Ushich, V.G., Korolev, Yu. D., Shemyakin, I.A., Frants, O.B., Frank, K. Sealed-off pseudospark switches (current status and prospects). *1st International Congress on Radiation Physics, High Current Electronics, and Modification of Materials*, Tomsk, Russia, 24-29 September 2000.
[49] General Atomics Energy Products, *Engineering Bulletin 96-004. The Effect of Reversal on Capacitor Life*, Nov 2003. Sorrento Electronics, Inc. www.gaep.com/tech-bulletins/voltage-reversal.pdf
[50] General Atomics Energy Products, *Engineering Bulletin Capacitors. 2007*. www.gaep.com/tech-bulletins/capacitor-engineering-bulletins.pdf
[51] Hoffman, Alan L. *Unsolicited Technical Proposal Appendices (Volume III). FRC High Flux Formation and Large s Physics Studies TRX/LSX*, 1985. A3-1.

[52] Brotherton-Ratcliffe, D., Storer, R.G. Transverse Oscillating Field Current Drive in Spherical Plasmas. *Plasma Physics and Controlled Fusion*, Vol 31. No 4, pp 615-628, 1989. Right handed helicon or whistler wave therefore diamagnetic.
[53] Slough, J.T., Hoffman, A.L. Experimental study of the formation of field-reversed configurations employing high-order multipole fields. *Phys. Fluids B* 2 (4), April 1990.
[54] Inomoto, Michiaki, et al. Effects of Neutral Particle Density on Equilibrium of Field-Reversed Configuration Sustained by Rotating Magnetic Field. *Plasma and Fusion Research: Regular Articles.* Volume 3, 004 (2008).
[55] Jain, KK., et al. Gas injected washer plasma gun. *J. Phys. E: Sci. Instrum*. Vol. 13, 1980. Great Britain.
[56] Fiksel, G. et al. High current plasma electron emitter. *Plasma Sources Sci. Technol*. 5 (1996) 78-83. Printed in the UK.
[57] Slough, J. Application of a Rotating Magnetic Field for Confinement Enhancement and Control of MHD Instabilities in the FRC. *7th Workshop on Active Control of MHD Stability*. Columbia University, NY. Nov 20, 2002. http://www.apam.columbia.edu/fusion/mhdcontrol2002/Tuesday%20Nov%2019/Slough.pdf
[58] Malliaris, A.C., et al. *Quasi-Steady MPD Propulsion at High Power. Final Technical Report*. AVCO Corporation. Feb 1971. N71-21717. NASACR – 111872.
[59] Choueiri, E.Y., Ziemer, J.K. Quasi-Steady Manetoplasmadynamic Thruster Measured Performance Database. EPPDyL, Princeton University. Presented at the *34th AIAA Joint Propulsion Conference*. July 1998.
[60] Hoskins, W A, et al. Asymmetric discharge patterns in the MPD thruster. *AIAA, DGLR, JSASS, International Electric Propulsion Conference, 21st, Orlando, Florida*, July 1990.
[61] Gilland, James. MPD Thruster Performance Models for System and Mission Analysis. *39th AIAA/ASME/SAE/ASEE Joint Propulsion Conference*. Huntsville, Al. 2003.
[62] http://www.emclab.umr.edu/new-induct/ 2007.
[63] Howard W. Sams & Co., Inc. *Reference Data for Radio Engineers*, ITT. New York. 1975. 6-1.
[64] http://en.wikipedia.org/wiki/Skin_effect 2007.
[65] http://www.hp.woodshot.com/ 2007.
[66] http://www.ee.scu.edu/eefac/healy/indwire.html based on E.B. Rosa, "The Self and Mutual Inductances of Linear Conductors", *Bulletin of the Bureau of Standards*, Vol.4, No.2, 1908, Page 301ff.
Alternative formula http://www.consultrsr.com/resources/eis/induct5.htm#ref1
[67] Tsaliovich, Anatoly. Shielded Electronic Cables for EMI Protection. *Belden Innovators*, Dec 1982, Vol 3 No 13.
[68] Morrison, Ralph. *Grounding and Shielding Techniques, Fourth Edition*. John Wiley & Sons, Inc. 1998. pg 176, figure 10.2.
[69] Horowitz, Paul, Hill, Winfield, *The Art of Electronics*. 2nd ed, 1989, pg 373.

[70] See RC ¼,1/2 from www.megastar.com, CRC-0405 series from Tyco Electronics, Carbon Composition Resistors from ibselectronics for examples of carbon comp specifications that include specified VCR.
[71] Howard W. Sams & Co., Inc., *Reference Data for Radio Engineers*, ITT. New York. 1975. 37-24.
[72] Williams, Barry Wayne. *Power Electronics Devices, Drivers, Applications, and Passive Components*. www.eee.strath.ac.uk Table 18.1.
[73] Grimaldi, C., Maeder, T., Ryser, P., Strassler, S. A random resistor network model of voltage trimming. *J. Phys. D: Appl. Phys.* 37 (2004) 2170-2174.
[74] Stevens, E.H., Gilbert, D.A., Ringo, J.A. High-Voltage Damage and Low-Frequency Noise in Thick-Film Resistors. *IEEE Transactions on parts, hybrids, and packaging*. Vol PHP-12, No. 4, Dec 1976.
[75] Stanimirovic, I., Jevtic, M.M., Stanimirovic, Z. Multiple high-voltage pulse stressing of conventional thick-film resistors. *Microelectronics Reliability* 47 (2007) 2242-2248.
[76] Tobita, T., Takasago, H. New Trimming Technology of a Thick Film Resistor by the Pulse Voltage Method. *IEEE Transactions on parts, hybrids, and manufacturing technology*, Vol 14, No. 3, Sep 1991.
[77] Personal communication with Caddock engineer.
[78] Iles, G. S. and Casale, M. E. A. Ruthenium Oxide Glaze Resistors – New Screen Printing Preparations for Thick Film Circuitry. *Platinum Metals Rev*. 1967, 11, (4), 126-129.
[79] EPCOS AG. Figure 19 Typical impedance characteristics of film capacitors. KMK0805-3. pg 20. http://www.epcos.com/web/generator/Web/Sections/ProductCatalog/Capacitors/FilmCapacitors/PDF/PDF__GeneralTechnicalInformation,property=Data__en.pdf;/PDF_GeneralTechnicalInformation.pdf. Accessed Oct 4, 2008.
[80] Guinta, Steve. Capacitance and Capacitors 1. Understanding the Parasitic Effects in Capacitors. *Analog Devices*, 2008. http://www.analog.com/library/analogDialogue/Anniversary/21.html
[81] Murata Manufacturing Co., Ltd. *Chip Monolithic Ceramic Capacitors C02E-12.* 2007 pg 49.
[82] EPCOS AG. Figure 12 Relative capacitance change dC/C vs. frequency f (typical example). KMK0797-8. pg 12. http://www.epcos.com/web/generator/Web/Sections/ProductCatalog/Capacitors/FilmCapacitors/PDF/PDF__GeneralTechnicalInformation,property=Data__en.pdf;/PDF_GeneralTechnicalInformation.pdf. Accessed Oct 4, 2008.
[83] EPCOS AG. Figure 11 Relative capacitance change dC/C vs. relative humidity (typical values) KMK0624-W-E. pg 11. http://www.epcos.com/web/generator/Web/Sections/ProductCatalog/Capacitors/FilmCapacitors/PDF/PDF__GeneralTechnicalInformation,property=Data__en.pdf;/PDF_GeneralTechnicalInformation.pdf. Accessed Oct 4, 2008.

[84] Sen, Bidyut K. (Sun Microsystems Inc.) and Wheeler, Richard L. (Wheeler Enterprises Inc.) *Skin Effects models for Transmission Line Structures using Generic SPICE Circuit Simulators.* http://www.wheeler.com/technology/technicalpaper2/technicalpaper2.pdf.
[85] Howard W. Sams & Co., Inc. *Reference Data for Radio Engineers*, ITT. New York. 1975. 6-3.
[86] Internal Redmond Plasma Physics Laboratory results. FRCs translating over probes composed of various materials (insulators and conductors). References to this knowledge are present in APS posters (some with incorrect references). An Aeronautics and Astronautics thesis hints at portions of this knowledge but is incomplete in what is presented.
[87] Tuszewski, M., Linford, R.K. Particle transport in field-reversed configurations. *Phys. Fluids* 25 (5), May 1982. 765. (Appendix)
[88] Hoffman, A.L., Milroy, R.D. Particle lifetime scaling in field-reversed configurations based on lower-hybrid-drift resistivity. *Phys. Fluids* 26 (11), November 1983.
[89] Hoffman, A.L., Slough, J.T. Field Reversed Configuration Lifetime Scaling Based on Measurements from the Large s Experiment. *Nuclear Fusion*, Vol.33, No.1 (1993).
[90] Chen, Francis F. *Introduction to Plasma Physics and Controlled Fusion 2nd Ed* 1984.
[91] Johnson, M.C. The Adsorption of Hydrogen on the Surface of an Electrodeless Discharge Tube. *Proceedings of the Royal Society of London. Series A, Containing Papers of a Mathematical and Physical Character.* Vol. 123, No 792, (Apr. 6, 1929), pp 603-613.
[92] Hoffman, A.L., Slough, J.T. Field Reversed Configuration Lifetime Scaling Based on Measurements from the Large s Experiment. *Nucl. Fusion* Vol. 33, No. 1, (1993) 27.
[93] Hoffman, A.L., Slough, J.T. Flux, Energy, and Particle Lifetime Measurements for Well Formed Field Reversed Configurations. *Nucl. Fusion* 26 (1986) 1693.
[94] Hoffman, A.L., Slough, J.T., Field Reversed Configuration Lifetime Scaling Based on Measurements from the Large s Experiment. *Nucl. Fusion* Vol. 33, No. 1, (1993) 27.
[95] Tuszewski, M., Linford, R.K. Particle transport in field-reversed configurations. *Phys. Fluids* 25 (5), May 1982. 765.
[96] NIST ASD Team. http://physics.nist.gov/PhysRefData/ASD/index.html (2005).

Bibliography

Adler, R.J., Smith, I.D., Noggle, R.C., Kiuttu, G.F. *Pulse Power Formulary*. U.S.AF (1991).

Allfrey, S.J., Bottino, A., Sauter, O., Villard, L. The role of radial electric fields in linear and nonlinear gyrokinetic ful radius simulations. *New Journal of Physics* 4, 29.1-29.19, (2002).

Andreason, S.P. *Langmuir and radial electric field probe results for TCS*. University of Washington, Aeronautics & Astronautics thesis. (2004).

Andreason, S.P., Slough, J.T. Internal probe array for the measurement of radial electric field. *Rev. Sci. Instrum*. Vol 75, 10, Oct (2004).

Apiezon, *Apiezon L datasheet*. 2007.

Arakoni, R.A., Ewing, J.J., Kushner, M.J. Microdischarges for use as microthrusters: modeling and scaling. *J. Phys. D: Appl. Phys*. 41 (2008) 105208.

Armstrong, W.T., Crawford, E.A., Hoffman, A.L. Flux-trapping during the formation of field-reversed configurations. *Phys. Fluids* 25, 2121 (1982).

Armstrong, W.T., Linford, R.K., Lipson, J., Platts, D.A., Sherwood, E.G. Field-reversed experiments (FRX) on compact toroids. *Phys Fluids* 24, 2068 (1981).

ASTM International. *Annual Book of ASTM Standards.* (2008).

Avago Technologies. datasheets for fiber optics products.

Back, R., Bengtson, R.D. A Langmuire/Mach probe array for edge plasma turbulence and flow. *Rev. Sci. Instrum*. 68 (1), (1997).

Barnes, D.C., Schwarzmeier, J.L., Lewis, H.R., Seyler, C.E. Kinetic tilting stability of field-reversed configurations. *Phys. Fluids* 29, 2616 (1986).

Barnes, D.C. Stability of long field-reversed configurations. *Phys. Plasmas* 9, 560 (2002).

Bentham, J. *TCP/IP Lean: Web Servers for Embedded Systems, 2nd ed*. CMP Books, Lawrence, Kansas, ISBN 1-57820-108-X, (2002).

Blevin, H.A., Thonemann, P.C. Plasma confinement using an alternating magnetic field. *Nucl. Fusion Suppl*. 55 (1962).

Bochkov, V.D., Dyagilev, V.M., Nicolin, S.V., Ushich, V.G., Seeteekh, D.S., Vlasov, P.N. The pseudospark switch crowbar unit – high reliability, low cost system. 0-7803-5498-2/99 IEEE (1999).

Bochkov, V.D., Djagilev, V.M., Ushich, V.G., Korolev, Yu.D., Shemyakin, I.A., Frants, O.B., Frank, K. Sealed-off pseudospark switches (current status and prospects). *International Congress on Radiation Physics, High Current Electronics, and Modification of Materials*. (Sep 2000).

Brotherton-Ratcliffe, D., Storer, R.G. Transverse oscillating field current drive in spherical plasmas. *Plasma Physics and Controlled Fusion*. Vol. 31, No. 4, pp 615-628 (1989).

Chen, F.F. *Introduction to plasma physics and controlled fusion volume 1: plasma physics 2nd ed*. Plenum Press, New York. ISBN 0-306-41332-9, (1984).

Chen, S., Sekiguchi, T. Instantaneous Direct-Display System of Plasma Parameters by Means of Triple Probe. *Journal of Applied Physics* 36 (8), (1965).

Chiang, P., Hsiao, M. Electric field in the edge layer of field-reversed configurations. *Phys. Fluids B* 4, 3226 (1992).

Choueiri, E.Y., Ziemer, J.K. Quasi-steady manetoplasmadynamic thruster measured performance database. AIAA-98-3472. *34th AIAA Joint Propulsion Conference*. (July 1998).

Clemente, R.A. On Current Drive in Field-Reversed Configurations. *J. Phys. Soc. Japan* 67, 3450 (1998).

Composite Agency - Analysis & Simulation Expertise. vapour –gas- water diffusion permeability polymers WVTR moisture. http://www.composite-agency.com/gas_diffusion.htm. Accessed 2007.

Composite Agency - Analysis & Simulation Expertise. -vapour- gas water diffusion permeability polymers WVTR moisture. http://www.composite-agency.com/vapour_diffusion.htm, Accessed 2007.

Cummings, D.B. “Chapter 9. Pulsed Power Ignitron Switches”. Preprint from book.

Cummings, D.B. “Ignitron discharge growth during high-current pulses”, (1962). *IEEE*. (UCRL-6635).

Dolan, T.J. *Fusion Research: Principles, Experiments and Technology*. ISBN 0-08-025565-5, (2000).

Donne, A.J.H. Diagnostics for current density and radial electric field measurements: overview and recent trends. *Plasma Phys. Control. Fusion* 44, B137-158, (2002).

Dunbar, W.G. *Design Guide: Designing and building high voltage power supplies Vol I*. AFWAL-TR-88-4143, ADA204503, (1988).

Dunbar, W.G. *Design Guide: Designing and building high voltage power supplies Vol II*. AFWAL-TR-88-4143, ADA204504, (1988).

Earley, M.W., Sargent, J.S., Sheehan, J.V., Caloggero, J.M. *NEC 2005 Handbook*. NFPA, Inc., Quincy, MA, ISBN 0-87765-625-8, (2005).

EEV Lincoln. *Ignitrons for Capacitor Discharge and Crowbar Applications*. (believed to be 1980s). (incomplete copy).

Enicks, D., Oleszek, G. Vacuum science considerations for rapid reactor recovery with extremely low oxygen in low temperature low pressure chemical vapor deposition of $Si_{1-x}Ge_x$ and $Si_{1-x-y}Ge_xC_y$ films. *J. Vac. Sci. Technol. A* 24(3), May/Jun (2006).

EPCOS AG. *General Technical Information: Film Capacitors*. 2007.

Farengo, R., Brooks, R.D. Current limit in ohmically heated high-beta plasmas: Application to field reversed configurations. *Phys. Fluids B* 3 130 (1991).

FDA. Surface lubricants used in the manufacture of metallic articles. *CFR Title 21*. Volume 3, section 178.3910 21CFR178.3910. Apr 2007. Available on web: www.accessdata.fda.gov.

Fiberoptic Systems, Inc. http://www.fiberopticsystems.com/tech2.htm, Accessed 2008.

Fiksel, G., Almagri, A.F., Craig, D., Lida, M., Prager, S.C., Sarff, J.S. High current plasma electron emitter. *Plasma Sources Sci. Technol*. 5 (1996) 78-83.

Fisch, N.J., Watanabe, T. Field reversal by rotating waves. *Nucl. Fusion* 22, 423 (1982).

Fusion Energy Sciences Advisory Committee. *Summary of Opportunities in the Fusion Energy Sciences Program*. U.S.DOE. (Nov 1999).

Gabuzda, D.C. The charge densities in a current-carrying wire. *Am. J. Phys*. 61 (4), (1993).

General Atomics Energy Products. *Engineering Bulletin 96-004. The Effect of Reversal on Capacitor Life*. Nov 2003. Sorrento Electronics, Inc. www.gaep.com/tech-bulletins/voltage-reversal.pdf

General Atomics Energy Products. *Engineering Bulletin Capacitors*. 2007. www.gaep.com/tech-bulletins/capacitor-engineering-bulletins.pdf

General Electric Company, Tube Products Department. *Ignitrons: Application Notes, Capacitor Discharge and Crowbar Service*. November 1974, M-1256.

Giersch, L., Andreason, S., Slough, J. Experimental apparatus for studying rotating magnetic field current drive in plasmas. *Rev. Sci. Intrum*. 76, 093507 (2005).

Giersch, L., Slough, J.T., Winglee, R. Demonstration of current drive by a rotating magnetic dipole field. *J. Plasma Physics* (2007), vol 73, pt 2, pp. 167-177.

Gilland, J. MPD thruster performance models for system and mission analysis. *39th AIAA Joint Propulsion Conference* (July 2003).

Goldston, R.J., Rutherford, P.H. *Introduction to Plasma Physics*. Institute of Physics Publishing, Philadelphia, ISBN 0-7503-0183-X, (1997).

Gombosi, T.I. The Plasma Environment of Comets. *Rev. Geophysics*. 91RG00752, 976-984, (1991).

Gota, H., Andreason, S.P., Votroubek, G.R., Pihl, C.J., Slough, J.T. Visible Bremsstrahlung tomographic diagnostic for the pulsed high density field-reversed configuration experiment. *Rev. Sci. Instrum*. 77, 10F319 (2006).

Gota, H., Andreason, S., Votroubek, G., Pihl, C., Slough, J. Temperature diagnostics for field-reversed configuration plasmas on the pulsed high density (PHD) experiment. *Plasma and Fusion Research: Regular Articles* Vol 2, S1050 (2007).

Gota, H., Akiyama, T., Fujimoto, K., Okhuma, Y., Takahashi, T., Nogi, Y. Separatrix shape measurement on field-reversed configuration plasmas. *Rev. Sci. Instrum*. 74, 2318 (2003).

Gota, H., Fujimoto, K., Ohkuma, Y., Takahashi, T., Nogi, Y. Separatrix shapes and internal structures of a field-reversed configuration plasma. *Phys. Plasmas* 10, 4763 (2003).

Griffiths, D.J. *Introduction to Electrodynamics, 3rd ed*. Prentice-Hall, New Jersey, ISBN 0-13-805326-X, (1999).

Grimaldi, C., Maeder, T., Ryser, P., Strassler, S. A random resistor network model of voltage trimming. *J. Phys. D: Appl. Phys*. 37 (2004) 2170-2174.

Grossnickle, J.A. *Deep Fueling of Large Tokamaks by Field-Reversed Configuration Injection*. University of Washington, Aeronautics and Astronautics, (2001).

Gruber, J.E. *Properties of ignitrons for 20 kv capacitor discharge circuits*. Institut fur Plasmaphysik Garching Bei Munchen. (1967).

Guinta, Steve. *Capacitance and Capacitors 1. Understanding the Parasitic Effects in Capacitors*. Analog Devices, 2008. http://www.analog.com/library/analogDialogue/Anniversary/21.html

Guo, H.Y., et al. Improved confinement and current drive of high temperature field reversed configurations in the new translation, confinement, and sustainment upgrade device. *Physics of Plasmas* 15, 056101 (2008).

Hakim, A., Shumlak, U. Two-fluid physics and field-reversed configurations. *Phys. Plasmas* 14, 055911 (2007).

Hamamatsu Photonics Devices. *R5900U-03-L16 PMT datasheet*.

Harned, D.S. Rotational instabilities in the field-reversed configuration: Results of hybrid simulations. *Phys Fluids* 26, 1320 (1983).

Hassam, A.B., Kulsrud, R.M., Goldston, R.J., Ji, H., Yamada, M. Steady State Thermoelectric Field-Reversed Configurations. *Phys. Rev. Lett*. 83, 2969 (1999).

Hibbeler, R.C. *Mechanics of Materials (3rd edition)*. Prentice Hall. (1997). ISBN 0132569833, 9780132569835.

Hershkowitz, N. Sheaths: More complicated than you think. *Phys. Plasmas* 12, 055502 (2005).

Hoffman, A.L. et al. *The large s FRC experiment (LSX).* STI Optronics. (>=1991).

Hoffman, A.L. *Field reversed configurations and the LSX facility*. (1991).

Hoffman, A.L. *Field Reversed Configurations (and Rotating Magnetic Field Current Drive)*. (2006).

Hoffman, A.L., Guo, H.Y., Milroy, R.D., Pietrzyk, Z.A. Resistivity scaling of rotating magnetic field current drive in FRCs. *Nucl. Fusion* 43 (2004) 1091-1100.

Hoffman, A.L., Matsen, S. *Unsolicited technical proposal appendices (Volume III).* A.12.310.01 U.S.DOE (1985).

Hoffman, A.L., Milroy, R.D. Particle lifetime scaling in field-reversed configurations based on lower-hybrid-drift resistivity. *Phys. Fluids* 26, 3170 (1983).

Hoffman, A.L., Slough, J.T., Crawford, E.A. *The large s FRC experiment: Final Report*. Co. No. DE-AC06-86ER53235 U.S.DOE, (1992).

Hoffman, A.L., Slough, J.T. Field Reversed Configuration Lifetime Scaling Based on Measurements from the Large s Experiment. *Nucl. Fusion* Vol. 33, No. 1, (1993).

Hoffman, A.L., Slough, J.T. Flux, energy and particle lifetime measurements for well formed field reversed configurations. *Nucl. Fusion* 26, 1693 (1986).

Hoffman, A.L., Slough, J.T. *Next step in FRC development*. UW.

Hoffman, Alan L. *Unsolicited Technical Proposal Appendices (Volume III). FRC High Flux Formation and Large s Physics Studies TRX/LSX*. 1985.

Hoskins, W A, et al. Asymmetric discharge patterns in the MPD thruster. *AIAA, DGLR, JSASS, International Electric Propulsion Conference, 21st*, Orlando, Florida, July 1990.

Horowitz, P., Hill, W. *The Art of Electronics*. 2nd ed. Cambridge University Press. ISBN 0521370957, 9780521370950. 1989.

Howard W. Sams & Co., Inc. *Reference Data for Radio Engineers*. ITT. New York. 1975.

Huba, J.D. *NRL Plasma Formulary*. Naval Research Laboratory. 2002.

Hugrass, W.N., Okada, T., Ohnishi, M. Plasma-circuit interactions in rotating magnetic field current drive. *Plasma Phys. Control. Fusion* 50, 055008 (2008).

Hutchinson, I.H. *Principles of plasma diagnostics 2nd ed*. Cambridge University Press, Cambridge, ISBN 0-521-80389-6, (2002).

Iles, G. S. and Casale, M. E. A. Ruthenium Oxide Glaze Resistors – New Screen Printing Preparations for Thick Film Circuitry. *Platinum Metals Rev*., 1967, 11, (4), 126-129.

Inomoto, M., Kitano, K., Okada, S. Effects of Neutral Particle Density on Equilibrium of Filed-Reversed Configuration Sustained by Rotating Magnetic Field. *Plasma and Fusion Research*: Vol. 3, 004 (2008).

Jain, K.K., John, P.I., Punithavelu, A.M., Rao, P.P. Gas injected washer plasma gun. J. *Phys. E: Sci. Instrum.* Vol. 13, (1980).

Jarboe, T.R. Review Article: Review of spheromak research. *Plasma Phys. Control. Fusion* 36, 945-990, (1994).

Johnson, M.C. The adsorption of Hydrogen on the surface of an electrodeless discharge tube. *Proceeding of the Royal Society of London, Series A, Containing Papers of a Mathematical and Physical Character*. Vol. 123, No. 792, (Apr 1929) 603-613.

Jones, I.R. A review of rotating magnetic field current drive and the operation of the rotamak as a field-reversed configuration (Rotomak-FRC) and a spherical tokamak (Rotamak-ST). *Phys. Plasmas* 6, 1950 (1999).

Kerns, Q., Kirsten, F. *Counting note: Pulse response of coaxial cables*. University of California, Berkeley (Feb 1964).

Kerns, Q., Kirsten, F., Winningstad, C. Pulse Response of Coaxial Cables. *Rev.* February 12, 1964.

Kleppa, O.J. The solid state: Diffusion in metals and alloys. *Annu. Rev. Phys. Chem*. 1955.6:119-140.

Klotzbucher, T., Kappler, A., Straub, K., Haderlein, S. Biodegradability and groundwater pollutant potential of organic anti-freeze liquids used in borehole heat exchangers. *Geothermics* 36 (2007) 348-361.

Kristiansen, M., Honig, E. *Workshop on high current, high coulomb ignitrons*. Texas Tech. U.S.AF. (1987).

Kushner, A.M. Characteristics & parameters of coaxial transmission lines. *Electronics World*, (Oct 1968) pg 54-56.

Larry, John R., Rosenberg, Richard M., Uhler, Roger O. Thick-Film Technology: An Introduction to the Materials. *IEEE Transactions on components, hybrids, and manufacturing technology*, Vol CHMT-3, No. 2, June 1980.

LaPointe, M. Numerical simulation of geometric scale effects in cylindrical self-field MPD thrusters. 1N-20, 117613, P_28, *28th AIAA Joint Propulsion Conference* (July 1992). NASA-CR-189224.

Litvak, A.G. (editor). *High-frequency plasma heating*. American Institute of Physics, New York, ISBN 0-88318-765-5, (1992).

MacLatchy, C.S., Boucher, C., Poirier, D.A., Gunn, J. Gundestrup: A Langmuire/Mach probe array for measuring flows in the scrape-off layer of TdeV. *Rev. Sci. Instrum*. 63 (8), (1992).

Malliaris, A.C., John, R.R., Garrison, R.L., Libby, D.R. Quasi-steady MPD propulsion at high power. *Final Technical Report.* AVCO Corporation, AVSD-0146-71-RR, NASA CR-111872, N71-21717. (Feb 1971).

Manheimer, W.M., Lashmore-Davies, C.N. *MHD and microinstabilities in confined plasma*. Adam Hilger, New York, ISBN 0-85274-282-7, (1989).

Massey, Liesl K. *Permeability Properties of Plastics and Elastomers: A Guide to Packaging and Barrier Materials 2nd Ed*. William Andrew Inc., 2003. ISBN 1884207170.

Maqueda, R.J., Wurden, G.A., Crawford, E.A. Wideband “silicon bolometers” on the LSX field reversed configuration experiment. *Rev. Sci. Instrum*. 63 (10), (1992).

Mathematical Sciences NorthWest, Inc. *Compact toroid generation, lifetime, and stability studies in linear reversed-field theta pinch geometries (TRX-1): Second annual and final report*. U.S.DOE DE-AC06-80ER53096, C.12.144.01 (1983).

Matthews, G.F. Review Article: Tokamak plasma diagnosis by electrical probes. *Plasma Phys. Control. Fusion* 36, 1595-1628, (1994).

McCracken & Stangeby. Plasma boundary phenomena in tokamaks. *Nuclear Fusion* 30 (7), (1990).

Menown, H. Gaseous switches: The past and present state of the art. *International Pulsed Power Conference*. English Electric Valve Co LTD., (1976).

Mikhailovskii, A.B., Laing, E.W. *Electromagnetic Instabilities in an Inhomogeneous Plasma*. CRC Press, ISBN 0750301821, 9780750301824, (1992).

Miller, K.E. *The Star Thrust Experiment, Rotating Magnetic Field Current Drive in the Field Reversed Configuration*. University of Washington, Aeronautics and Astronautics, (2001).

Milroy, R.D., Guo, H.Y. Rotating magnetic quadrupole current drive for field-reversed configurations. *Physics of Plasmas* 12, 072503 (2005).

Milroy, R.D., Steinhauer, L.C. Numerical modeling of an end-plugged theta pinch. *Phys. Fluids* 24, 339 (1981).

Morrison, R. *Grounding and Shielding Techniques, 4th ed*. John Wiley & Sons, Inc., New York, ISBN 0-471-24518-6, (1998).

Nicholson, D.R. *Introduction to plasma theory*. Wiley, New York, ISBN 0-471-09045-X, (1983).

NIST ASD Team. http://physics.nist.gov/PhysRefData/ASD/index.html (2005).

NREL. *A look back at the U.S. Department of Energy’s Aquatic species program* (1998).

Ohkuni, K., Toi, K., Ohdachi, S., Takagi, S., Torus experiment group (Japan). Langmuir probe array for edge plasma study on the compact helical system heliotron/torsatron. *Rev. Sci. Instrum*. 70 (1), (1999).

Ono, Y., Matsuyama, T., Umeda, K., Kawamori, E. Spontaneous and artificial generation of sheared-flow in oblate FRCs in TS-3 and 4 FRC Experiments. *Nucl. Fusion* 43 (2003) 648-653.

Oost, G.V., Adamek, J., et al. Turbulent transport reduction by E x B velocity shear during edge plasma biasing: recent experimental results. *Plasma Phys. Control. Fusion* 45, 621-643, (2003).

Osaka Vacuum, LTD. *Operation Manual: Magnetic Suspended Compound Molecular Pump / Controller…* 1998.

Owens, T.L. A pulse-compression-ring circuit for high-efficiency electric propulsion. *Rev. Sci. Intsrum*. 79, 034701 (2008).

Parker Hannifin Corporation (General Valve Division) series 9 poppet valve schematics.

Permeability and Plastic Packaging. Accessed 2008, http://www.personal.psu.edu/irh1/PDF/Perm%20and%20Pack.pdf

Persson, H., Agren, O. Vlasov solution describing nonadiabatic plasma heating by magnetic compression. *Phys. Fluids B* 5 (7), July (1993).

Peter, A.M. *Paramagnetic Spin-up of a Field Reversed Configuration with Rotating Magnetic Field Current Drive*. University of Washington, Aeronautics and Astronautics, (2003).

Pike, G.E., Seager, C.H. Electrical properties and conduction mechanisms of Ru-based thick-film (cermet) resistors. *J. Appl. Phys* 48(12), Dec 1977.

Professional Engineering Computations (PEC, Inc.). *Outgassing and Permeation*. http://lpc1.clpccd.cc.ca.us/lpc/tswain/lect8.pdf. CA. 2003.

Proof-of-Principle Proposals in Fusion Energy Science, Consensus technical summary of the members of the review panel for, June 8-11, 1998, *Magnetized Target Fusion: A Proof-of-Principle Research Proposal*

Rade, L., Westergren, B. *Mathematics Handbook for Science and Engineering 4th ed*. Springer, New York, ISBN 3-540-65569-7, (1988).

Rej, D.J., Taggart, D.P., Baron, M.H., Chrien, R.E., Gribble, R.J., Tuszewski, M., Waganaar, W.J., Wright, B.L. High-power magnetic-compression heating of field-reversed configurations. *Phys. Fluids B* 4, 1909 (1992).

Riccardi, C., Longoni, G., Chiodini, G., Fontanesi, M. Comparison between fast-sweep Langmuir probe and triple probe for fluctuations measurements. *Rev. Sci. Instrum*. 72 (1), (2001).

Rodriguez, A., Gatrell, J., Karas, J., Peschke, R. *IBM: TCP/IP Tutorial and Technical Overview, 7th ed*. Prentice Hall Ptr, ISBN 0-13-067610-1, (2002).

Rudakov, D.L., Shats, M.G., Boswell, R.W., Charles, C., Howard, J. Overview of probe diagnostics on the H-1 heliac. *Rev. Sci. Intrum*. 70 (1), (1999).

Santarius, J.F. Fusion Space Propulsion – A shorter time frame than you think. UWFDM-1287, *JANNAF* (2005).

Sarfaty, M., Harper, M., Hershkowitz, N. A novel electro-optical probe to diagnose plasma uniformity. *Rev. Sci. Instrum*. Vol 69, No 9, (1998).

Seddon, R. *The use of ignitrons at high voltage*. (July 1987).

Sen, Bidyut K. (Sun Microsystems Inc.) and Wheeler, Richard L. (Wheeler Enterprises Inc.). *Skin Effects models for Transmission Line Structures using Generic SPICE Circuit Simulators*. http://www.wheeler.com/technology/technicalpaper2/technicalpaper2.pdf.

Slough, J., Andreason, S., Gota, H., Pihl, C., Votroubek, G. The pulsed high density experiment: Concept, design, and initial results. *Journal of Fusion Energy*, Vol. 26, 1/2, June (2007).

Slough, J.T., Andreason, S.P., Milroy, R.D. Plasma resistivity profile measurement from an external radio-frequency magnetic coil. *Rev. Sci. Instr*. Vol 75, 10 (Oct 2004).

Slough, J., Andreason, S., Ziemba, T., Ewing, J.J. *Micro-discharge micro-thruster*. U.S.AF ADA437433 (2005).

Slough, J. Application of a rotating magnetic field for confinement enhancement and control of MHD instabilities in the FRC. *7th Workshop on Active Control of MHD Stability*. Columbia University, NY, (Nov 2002).

Slough, J.T., Andreason, S.P. Rotation and Confinement of a Plasma with Finite Ion Pressure sustained by a Rotating Magnetic Field.

Slough, J.T., Hoffman, A.L. Experimental study of the formation of field-reversed configurations employing high-order multipole fields. *Phys. Fluids B* 2, 797 (1990).

Slough, J.T., Hoffman, A.L., Milroy, R.D., Crawford, E.A., Cecik, M., Maqueda, R., Wurden, G.A., Ito, Y., Shiokawa, A. Confinement and stability of plasmas in a field-reversed configuration. *Phys. Rev. Lett.* 69, 2212 (1992).

Sorrento Electronics, Inc. *The effect of reversal on capacitor life*. (Nov 2003).

Spectra Technology, Inc. *Compact toroid generation, lifetime, and stability studies in linear reversed-field theta pinch geometries (TRX-2): final report*. U.S.DOE C.12.144.07 (1987).

Spectra Technology, Inc. *LSX construction: final report*. U.S.DOE (1990).

Spectra Technology, Inc. *LSX operations quarterly progress report No.1*. U.S.DOE (1990).

Spectra Technology, Inc. *LSX operations quarterly progress report No.2*. U.S.DOE (1991).

Spectra Technology, Inc. *LSX operations quarterly progress report No.3*. U.S.DOE (1991).

Spectra Technology, Inc. *LSX quarterly final report*. U.S.DOE (1990).

Spectra Technology, Inc. *LSX quarterly technical progress report No.2*. U.S.DOE (1987).

Spectra Technology, Inc. *LSX quarterly technical progress report No.3*. U.S.DOE (1987).

Spectra Technology, Inc. *LSX quarterly technical progress report No.16*. U.S.DOE (1990).

Spencer, R.L., Tuszewski, M., Linford, R.K. Adiabatic compression of elongated field-reversed configurations. *Phys. Fluids* 26, 1564 (1983).

Spencer, W.E. *Memo: Electrolytic Resistors*. (Sep 24, 1970).

Stanimirovic, I., Jevtic, M.M., Stanimirovic, Z. Multiple high-voltage pulse stressing of conventional thick-film resistors. *Microelectronics Reliability* 47 (2007) 2242-2248.

Stanimirovic, Z., Jevtic, M.M., Stanimirovic, I. Computer Simulation of Thick-Film Resistors Based on 3D Planar RRN Model. *Eurocon 2005*, Serbia & Montenegro, Belgrade, Nov 22-24, 2005.

Stanimirovic, Z., Jevtic, M.M., Stanimirovic, I. Simultaneous mechanical and electrical straining of conventional thick-film resistors. *Microelectronics Reliability* 48 (2008) 59-67.

Stanway, M., Seddon, R. *Ignitron firing time and jitter*.

Steinhauer, L.C. (editor), et al. *FRC 2001: A White Paper on FRC Development in the Next Five Years*. http://www-ferp.ucsd.edu/PUBLIC/AC-PANEL/REC-DOCS/W-PAPERS/frc-2001.html, (1996).

Steinhauer, L.C. Improved analytic equilibrium for a field-reversed configuration. *Phys. Fluids B* 2, 3081 (1990).

Steinhauer, L.C., Ishida, A., Kanno, R. Ideal stability of a toroidal confinement system without a toroidal magnetic field. *Phys. Plasmas* 1, 1523 (1994).

Steinhauer, L.C. Plasma heating in field-reversed theta pinches. *Phys. Fluids* 26, 254 (1983).

Stevens, E.H., Gilbert, D.A., Ringo, J.A. High-Voltage Damage and Low-Frequency Noise in Thick-Film Resistors. *IEEE Transactions on parts, hybrids, and packaging*. Vol PHP-12, No. 4, Dec 1976.

STI Optronics. *Tokomak refueling by accelerated plasmoids (TRAP) experimental development proposal*. (1991).

Sutterlin, Robert C., Dayton, Gordon O., Biggers, James V. Thick-Film Resistor/Dielectric Interactions in a Low Temperature Co-fired Ceramic

Package. *IEEE Transactions on components, packaging, and manufacturing technology, Part B*, Vol 18, No. 2, May 1995.

Tipler, P.A., Llewellyn, R.A. *Modern Physics, 3rd ed*. W.H. Freeman and Company, New York, ISBN 1-57259-164-1, (1999).

Tobita, T., Takasago, H. New Trimming Technology of a Thick Film Resistor by the Pulse Voltage Method. *IEEE Transactions on parts, hybrids, and manufacturing technology*, Vol 14, No. 3, Sep 1991.

Totten, G.E. *Handbook of Lubrication and Tribology: Volume 1 Application and Maintenance, 2nd ed*. CRC Press, ISBN 084932095X, (2006).

Tsaliovich, Anatoly. Electromagnetic Compatibility Performance of Electronic Cables. *Belden Innovators*, Nov 1983.

Tsaliovich, Anatoly. Shielded Electronic Cables for EMI Protection. *Belden Innovators*, Dec 1982, Vol 3 No 13.

Tuszewski, M., Barnes, D.C., Rej, D.J., Taggart, D.P. Stability and compressional heating of large field-reversed configurations in the FRX-C/LSM device. *Phys. Fluids B* 3, 2205 (1991).

Tuszewski, M, Barnes, G.A., Chrien, R.E., Hugrass, W.N., Rej, D.J., Siemon, R.E., Wright, B. The origin of the rotation in field-reversed configurations. *Phys. Fluids* 31, 946 (1988).

Tuszewski, M. Experimental study of the equilibrium of field-reversed configurations. *Plasma Physics and Controlled Fusion*, Vol. 26, No. 8, pp 991-1005, (1984).

Tuszewski, M. Field Reversed Configurations. *Nucl. Fusion* 28, 2033 (1988).

Tuszewski, M., Linford, R.K. Particle transport in field-reversed configurations. *Phys. Fluids* 25, 765 (1982).

Tuszewski, M. Review Paper: Field Reversed Configurations. *Nuclear Fusion*, Vol.28, No.11 (1988). 2033-2092.

Tuszewski, M., Taggart, D.P., Chrien, R.E., Rej, D.J., Siemon, R.E., Wright, B.L. Axial dynamics in field-reversed theta pinches. II: Stability. *Phys. Fluids B* 3, 2856 (1991).

Uribarri, L. *Onset Phenomena in MPD Thrusters*. http://www.princeton.edu/~uribarri/ressearch/Onset/, (2006).

Vlases, G.C. *TCS edge studies final technical report*. U.S.DOE DE-FG03-01ER-546537 (Feb 2005).

Votroubek, G.R. *Tomographically Aided Study of Field Reversed Configuration Plasma Rotation and Stability*. University of Washington, Aeronautics and Astronautics, (2005).

Votroubek, G., Slough, J., Andreason, S., Pihl, C. Formation of a stable field reversed configuration through merging. *J Fusion Energ* (2008) 27:123-127.

Wang, A.K., Song, X.M., Qiu, X.M. Fluctuation-induced current and energetic electrons in reversed field pinch plasmas. *Plasma Phys. Control. Fusion* 37, 647-655, (1995).

Williams, Barry Wayne. *Power Electronics Devices, Drivers, Applications, and Passive Components*. www.eee.strath.ac.uk. Accessed 2007.

Woodruff, S., Macnab, A.I.D., Mattor, N. Adiabatic compression of a doublet field reversed configuration (FRC). *J Fusion Energ* (2008) 27:128-133.

Woodruff, S. Technical survey of simply connected compact tori (CTs): spheromaks, FRCs and compression schemes. *J Fusion Energ* (2008) 27:134-148.

Zajec, B., Nemanic, V. Hydrogen bulk states in stainless-steel related to hydrogen release kinetics and associated redistribution phenomena. *Vacuum* 61 (2001) 447-452.

Zeller, H.R. Cosmic ray induced failures in high power semiconductor devices. *Microelectron. Reliab.* Vol 37, No 10/11, pp. 1711-1718. (1997).