

Practical Design Guidelines for Fugitive Gas Detection from Unmanned Aerial Vehicles

by

William D. Tandy Jr.

B.S., University of Texas at Austin, 2004

M.S., University of Texas at Austin, 2006

A thesis submitted to the
Faculty of the Graduate School of the
University of Colorado in partial fulfillment
of the requirement for the degree of
Doctor of Science
Department of Aerospace Engineering

2017

This thesis entitled:

Practical Design Guidelines for Fugitive Gas Detection from Unmanned Aerial Vehicles

written by William David Tandy Jr.
has been approved for the Department of Aerospace Engineering

William Emery

Jarett Bartholomew

Date_____

The final copy of this thesis has been examined by the signatories, and we find that both the content and the form meet acceptable presentation standards of scholarly work in the above mentioned discipline.

Tandy Jr., William (Ph.D., Aerospace Engineering)

Practical Design Guidelines for Fugitive Gas Detection from Unmanned Aerial Vehicles

Thesis directed by Professor William J. Emery

Abstract

Simulation, design, and analysis are combined in this effort to realize a UAV-scale instrument for fugitive gas detection. The contributing material to the industry begins by extending and correlating an integrated Gaussian plume model useful for instrument predictions and trade studies, regardless of the instrument type or molecule of interest. A variety of generally applicable plots are produced from this foundation, including receiver operator curves for leak rate detectability vs. wind speed, beam diameter vs. leak rate detectability, and plots for required scan densities. The atmospheric and instrument parameter trade studies are followed by hardware-specific analyses applicable to differential absorption lidar (DIAL) instruments. A synopsis of the lessons learned from hands-on experiences in the lab further define the design space for DIAL sensors. The dissertation culminates in the detailed design and analysis of two DIAL instrument concepts. The conclusion is that a DIAL instrument capable of reliably detecting a 50 SCFH plume in winds speeds up to 7 mph is on the threshold of being achievable on a quadcopter platform. Of special note is that the effort was funded by a Pipeline and Hazardous Materials Safety Administration grant and performed in collaboration with Ball Aerospace & Technologies.

Dedication

If you think you might even tangentially be someone I have leaned on while pursuing this dream, I
promise you that you were even more important than you realize. Thank you.

Acknowledgments

Funding for this dissertation comes from the Pipeline and Hazardous Materials Safety Administration

Award DTPH5615HCAP11, Proposal 0415.07.0998B - Small-Scale DIAL for Methane Detection.

Table of Contents

Chapter 1: Introduction	1
1.1 Motivation.....	1
1.2 Methane Overview	9
1.3 A Brief History of Gas Detection	11
1.3.1 In Situ Devices	11
1.3.2 Lasers and Methane Sensing.....	13
1.4 Dissertation Focus and Outline	15
1.5 PHMSA Grant and Ball Collaboration	17
Chapter 2: Technical Foundations	19
2.1 DIAL Technical Details	19
2.2 UAV Platforms	22
2.2.1 Fixed-Wing UAVs	22
2.2.2 Quad-Copter UAVs.....	24
2.3 Choosing a Laser Wavelength	25
2.4 Pulsed vs. Continuous Wave.....	28
2.5 Laser Technology.....	30
Chapter 3: Gaussian Plume Simulations	31
3.1 The Gaussian Plume Model	31
3.2 Extending the Gaussian plume model.....	33
3.3 Model Correlation with Field Measurement	35

Chapter 4: Analysis Trade Studies.....	42
4.1 Signal to Noise Ratio Calculation	42
4.2 Choosing a Beam Diameter	46
4.3 Impact of Error on Plume Area.....	51
4.4 Quality of Measurement Sets.....	53
4.5 Scan Pattern Considerations.....	57
4.6 Influence of Beam Angle on Effective Beam Diameter	60
4.7 Speckle Considerations.....	62
Chapter 5: Wavelength-Based Reflectivity Variations	69
5.1 Introduction.....	69
5.2 Data.....	72
5.3 Methodology	76
5.4 Analysis Results.....	79
5.5 Example Applications	90
5.6 Summary	93
5.7 Application to Methane at 1.645 μm	94
Chapter 6: Hardware Planning and Considerations	97
6.1 Two Instrument Concepts	98
6.2 Spectral Separation: Volume Bragg Grating.....	99
6.3 Single Mode vs. Multimode Fiber	101
6.4 Beam Pickoff	104

6.5 Minimizing Optic Count	110
6.6 Beam Steering	110
6.7 Detectors	114
6.8 Signal Measurement Via Integration	120
6.9 Analog to Digital Conversion, FPGA, and Storage	127
6.10 Lessons Learned from Ball Hardware	128
Chapter 7: Baseline Design and Calculation.....	130
7.1 Goals and Requirements	130
7.2 Design Process	130
Chapter 8: Final Hardware Design.....	133
8.1 Design Concept 1: Quad-Copter “Traditional” Instrument	133
8.2 Design Concept 2: Fixed-Wing “Traditional” Instrument with Scanning	141
8.3 Design Concept 3: VBG Filtering.....	142
8.4 SNR Analysis.....	157
8.5 Leak Detection Performance.....	161
8.6 Approximate Cost and Weight.....	164
Chapter 9: Summary, Conclusions, and Next Steps	168
9.1 Summary	168
9.2 Conclusions.....	173
9.3 Next Steps	174
Bibliography	176

Appendix A: Cost and Weight Tables	184
--	-----

List of Tables

Table 1: Listing of commercially available mid-sized UAVs. Data comes from vendor websites.	24
Table 2: Sample Quadcopter values.....	25
Table 3: I_y and I_z are the lateral and vertical intensity coefficients. They are chosen based on empirical values and are simplified.	31
Table 4: Parameters chosen to create Figure 15	34
Table 5: Goodness of Fit R-Squared statistics for three plumes with Gaussian fits.	39
Table 6: Type and class listing of the reflectance datasets. For compactness, the meteorite and minerals classes are summarized.	73
Table 7: List of individual and cumulative losses in the optical system. In cases where polarization differentially affects losses, the minimum value of the two is displayed.	138
Table 8: List of individual and cumulative losses in the optical system. In cases where polarization differentially affects losses, the minimum value of the two is displayed.	154
Table 9: Input parameters for the SNR equation. Values are for the “traditional” design.....	158
Table 10: Output from model for input from Table 9.....	159
Table 11: Input parameters for the SNR equation. Values are for the “VBG Filtering” design.....	160
Table 12: Output from model for input from Table 11	161
Table 13: Summaries of the cost and weight of the two concepts. The Traditional concept is from the quadcopter iteration.	165
Table 14: List of major system components with unknown costs and weights	166
Table 15: List of approximate component costs.	184
Table 16: List of component weights.....	188

List of Figures

Figure 1: Greater than 50% of the 300,000+ miles of transmission pipeline was laid prior to 1970. More than 1 million miles of distribution pipeline was laid prior to 1980.	2
Figure 2: The number of serious accidents has held relatively steady. However, loose reporting laws and increasing pipeline miles obscure the true trends. [PHMSA, 2016]	3
Figure 3: Directly measured methane concentrations have increased since dedicated measurements began in the 1980's. [Steele, 2005, Hawaii, 2017, Australia, 2017]	4
Figure 4: The Environmental Protection Agency tracks methane sources within the United States [EPA, 2015]	6
Figure 5: The property damage per year shows an increasing trend [PHMSA, 2016]	8
Figure 6: Methane has one carbon atom (black) with four hydrogen atoms (white) in a tetrahedral arrangement [18]	9
Figure 7: The PNNL database for methane provides atmospheric compensated features. The subset pictures demonstrate progressive zooms into the features around 1,645 nm.	10
Figure 8: The Davy Lamp burned brighter in the presence of methane and dimmer in the presence of CO ₂ . [Davy,1816].....	11
Figure 9: Different lasers interact with the gas in different ways. Measuring those differences provides insight into the quantity and, sometimes, the distribution of the gas.	19
Figure 10: Plot comparing the values of available payload power, maximum payload weight, and the designed cruise velocity.....	23
Figure 11: Sample quadcopter values for payload weight and endurance.....	25
Figure 12: Example reflectivity plots demonstrating that reflectivity tends to be higher at lower wavelengths.....	27
Figure 13: A simple example of a Gaussian plume model predicting the mass density of a gas at a given height and location.....	32
Figure 14: The beam from a DIAL instrument (red circles) integrates the total volume and so overlap with a gas plume (blue triangle) is driven by the percentage of beam and plume overlap.	34
Figure 15: Plotted examples of the plumes parameterized in Table 4.	35
Figure 16: Example plume of fugitive methane captured with Ball's methane detecting instrument on May 16 th , 2017.....	36
Figure 17: Graphical illustration of the plume binning described in the text	37
Figure 18: Example cross-sections of plumes measured with Ball's methane detecting instrument.....	38

Figure 19: Goodness of Fit “R-Squared” plot for the plume in Figure 15. The red arrows in the first image point out the low R-Squared value example plots in the next two images.	38
Figure 20: Example simulated plume interacting with a 1 meter diameter laser beam. The image on the left is the perfect result while the image on the right assumes Poisson shot noise.	40
Figure 21 Along-axis plots with points representing Gaussian fit peaks. The first plots’ source is the plume in Figure 16, the second is from a ~800 SCFH plume.	40
Figure 22: Actual plume profiles vs. simulated plume with matched peaks.	41
Figure 23: Plots demonstrating the relationship between the system SNR, the quantity of methane in the environment, and the ability to correctly resolve the methane. See the text for how the average methane concentration was converted to SCFH.	45
Figure 24: A larger diameter beam on a plume measures more of the surrounding air than a smaller beam.	46
Figure 25: 300 SCFH plume as measured by different diameter beams. Scales relative to measurement.	47
Figure 26: 300 SCFH plume as measured by different diameter beams. Scales relative to the 0.5-meter beam.	48
Figure 27: Simulation of beam diameters on a 50 SCFH flow simulation with 50 ppm noise floor. The color scale is fixed to that of the 0.5-meter case.	49
Figure 28: A variety of curves representing the plume sizes visible to an instrument with beams of different diameters. For this plot, all other variables are held fixed.	50
Figure 29: Example receiver operator curves for a set beam diameter and a range of wind speeds. Blue circles indicate detectable plumes. Red X’s indicate non-detectable plumes.	51
Figure 30: Detectable plume area vs. SCFH leak rate for three minimum observable methane concentration above a background concentration of 1.8 ppm. See text for assumptions.	52
Figure 31: Plume dimensions as a function of noise floor and leak rate. The y-axis is the half-width of the plume.	53
Figure 32: Plot of Quality ratios versus the number of random points measured within a 100 SCFH plume with an instrument noise floor of 0.18 ppm.	55
Figure 33: Plots showing how fixed scan resolutions map to measurement quality. The jagged nature of the curves is related to “modes” of scan resolution on fixed lengths.	56
Figure 34: Overview of beam diameters on the ground (red) from an aircraft flying transverse to the plume (blue), demonstrating that beams that are too small may miss plumes.	57
Figure 35: Sample circular scan pattern with red dots indicating pulse locations on the ground, following the scan pattern displayed as the dashed line. Scan rate sets the along-track point density while pulse rate sets the cross-track point density.	58

Figure 36: Exaggerated representation of subtended arcs created by rays reflecting from the ground.	60
Figure 37: The plot on the right shows the percent of energy reflected back to the instrument relative to the peak value. The plot on the right convolves this result with a Gaussian profile.	61
Figure 38: Figure on the left investigates how much the effective diameter is reduced as a function of percentage while the right image looks at the result in terms of resulting angles.	62
Figure 39: Simulated speckle patterns with equal total energies. The images were created by introducing phase noise to stacked Zernike functions to simulate multimode fiber.	63
Figure 40: Ratios of energy from the speckle patterns. Values are absolute differences from unity. The first plot has a value of zero, demonstrating equal total energy.	65
Figure 41: Reflectivity vs. wavelength plot for two materials commonly seen around oil and gas infrastructure. Data from JPL ASTER spectral library [Baldrige, 2009].	66
Figure 42: Percent error contribution due to reflectivity variations of concrete and glossy paint for the laser speckle examples.	67
Figure 43: Wide varieties of materials demonstrate precipitous reflection increases until around 0.7 μm and decreases around 2.7 μm	74
Figure 44: Spline interpolation incorrectly added reflectivity at peaks and shelves. Matlab's "PCHIP" implementation more closely followed the underlying data and was used in the analysis	75
Figure 45: The materials in the database that are relevant to the analysis vary with wavelength.	76
Figure 46: The distribution of reflectance (%) with wavelength. Many materials, such as marble, have multiple wavelengths over a wide range with similar reflectance. The red highlights show wavelengths that would produce less than 5% error in the overall DIAL result for a given reflectance.	79
Figure 47: Complete comparison of all possible offline and online wavelength combinations reflectance that are in the database	80
Figure 48: Contour plot of the percentage of materials within the JPL library that introduce less than 5% error in a DIAL signal. Contours are in 10% increments, starting with 10+% at the outermost contour and ending at 100% in the inner most regions. Insets demonstrate some of the fine-grained structure in the dense, central region.	81
Figure 49: All of the reflectance materials in the database analyzed from 400 nm to 2,500 nm with an offline range of ± 6 nm.	82
Figure 50: All of the reflectance materials in the database analyzed from 2.5 μm to 4.0 μm with an offline range of ± 6 nm.	83
Figure 51: Absolute value of the offline wavelength separation for various thresholds for all materials from 0.4 μm to 2.5 μm with 2,287 sample sets.	83

Figure 52: Absolute value of the offline wavelength separation for various thresholds for all materials from 2.5 μm to 4.0 μm with 2,287 sample sets.....	84
Figure 53: Absolute value of the offline wavelength separation for various thresholds for the manmade materials, with 56 sample sets.	85
Figure 54: Absolute value of the offline wavelength separation for various thresholds for the manmade materials, with 56 sample sets.	85
Figure 55: Absolute value of the offline wavelength separation for various thresholds for the rock type, with 473 sample sets.	86
Figure 56: Absolute value of the offline wavelength separation for various thresholds for the rock type, with 473 sample sets.	86
Figure 57: Absolute value of the offline wavelength separation for various thresholds for the minerals type, with 1,616 sample sets.	87
Figure 58: Absolute value of the offline wavelength separation for various thresholds for the minerals type, with 1,616 sample sets.	87
Figure 59: Absolute value of the offline wavelength separation for various thresholds for the soils type, with 52 sample sets.	88
Figure 60: Absolute value of the offline wavelength separation for various thresholds for the soils type, with 52 sample sets.	88
Figure 61: Absolute value of the offline wavelength separation for various thresholds for the water types, with 9 sample sets. The analysis only went to 10 nm of separation, cutting off some of the result around 400 nm.	89
Figure 62: Absolute value of the offline wavelength separation for various thresholds for the water types, with 9 sample sets.	89
Figure 63: Sensitivity plots of the allowable percent error contributed by reflectivity variations with each line representing the ratio of offline to online received power. The nominal case used throughout the paper is 5% error with a ratio of 1.10.	90
Figure 64: The center of the three plots represents the peak values from the three main clusters of absorptive online wavelengths for methane below 4 μm . The vertical bars span regions where 90%, 95%, and 99% of all of the materials in the database contribute 5% error or less.	91
Figure 65: Plot of the CO ₂ and H ₂ O absorption lines for the first wavelength region considered for the UAV instrument.....	92
Figure 66: Plot of the CO ₂ and H ₂ O absorption lines for the second wavelength region considered for the UAV instrument.....	93

Figure 67: Plot of sorted percent error for all of the materials in the ASTER library for the online and offline wavelengths of 1645.55 and 1645.37 nm respectively, assuming that the offline to online received power has a ratio of 1.1.....	95
Figure 68: Reflectivity vs Wavelength for Tourmaline. The inset image is for roughly the 1.61 to 1.67 μm range.....	96
Figure 69: Simple schematic of a typical system that separate the online and offline pulses in time.	98
Figure 70: Simplified schematic of a system designed to process a single pulse with multiple wavelengths	99
Figure 71: Simple diagram of a transmissive VBG redirecting a desired wavelength while other wavelengths pass through. See [Ciapurin, 2006] for more information.	100
Figure 72: Plot of effective mode field diameters of single mode fibers sold by ThorLabs as of June, 2017. Curves represent max and min values in product properties. [ThorLabs, 2017]	102
Figure 73: Plot of practical damage thresholds for ThorLabs’ single mode fibers carrying continuous laser energy. Curves represent max and min values in product properties. [ThorLabs, 2017]	102
Figure 74: Plot of practical damage thresholds for ThorLabs’ single mode fibers carrying 10 ns pulses. Curves represent max and min values in product properties. [ThorLabs, 2017]	103
Figure 75: Plot of practical damage limit for ThorLabs’ multi-mode fiber carrying 10 ns pulses. [ThorLabs 2017]	103
Figure 76: Simple schematic demonstrating how DIAL, which has two laser sources, can be combined to create a single laser output path.	105
Figure 77: Example plot of how even “non-polarizing” beam splitters still have some polarization dependence. [ThorLabs, 2017].....	106
Figure 78: Plot of the error coefficient that would multiply through the rest of equation 11 as a function of optical efficiencies.	107
Figure 79: Plot of the error coefficient for ThorLabs’ 50:50 beam splitting optic at a wavelength of 1.645 μm . The range of values maps to their as-measured light splitting variations due to polarity.	108
Figure 80: Plot of ThorLabs protected gold coating for the region around 1645 nm.	109
Figure 81: Plot of the reduction in energy/photon counts per optical interface given a fixed loss value per interface.	110
Figure 82: Simple schematic of how a rotating mirror with a wedge profile can trace out a path on the ground. Angles are exaggerated to illustrate the concept.	112
Figure 83: Plot of the minimum and maximum relative angles the outside diameter of the disk can be while keeping the laser footprint within the receiver footprint.	113

Figure 84: A square wave with 50% duty cycle and 20 ns period in both the time and frequency domains.	117
Figure 85: Each curve represents the signal created by combining the modes up to its value. For instance, the long periodic form has one mode, the next curve has the first two modes, and so on.	117
Figure 86: Plots of rise time demonstrating that adding additional modes (bandwidth) provides limiting returns in recreating the square wave pulse.	118
Figure 87: The plot on the left is a time series representation of a half sine wave. The plot on the right is the Fourier transform of the somewhat discretized sine wave.	119
Figure 88: Plots of rise time demonstrating that increasing the mode count for a sine wave Fourier series reconstruction has limited utility.....	119
Figure 89: Simulated pulse shapes with an idealized return pulse on the left and a multiple reflecting pulse on the right.	121
Figure 90: Overview of the boxcar integration technique where pulse energy is cumulatively summed within the integration period. The top plot is two Gaussian shaped plumes with simulated Poisson shot noise. The second plot shows external trigger pulses being sent to begin the process of integration. The third plot shows that the integration gates begin summing energy a fixed time after the trigger is sent. The bottom plot shows the results of the summing operation with RC roll-off. The ADC reads the energy from the top of the integrated pulse before the gate ends. Triggering the start of the integration and using an appropriate integration period are two key parts of the approach.....	122
Figure 91: Example integration result for a series of centered integration gate widths. The plot on the left is a Gaussian pulse with a 150 ns FWHM duration. The middle plot shows the integrated energy as a function of the number of FWHM time periods. The plot on the right shows the percent error relative to the true energy within the pulse.	123
Figure 92: Plots results for a study investigating the sensitivity to gate location relative to the center of the pulse. The plot on the left is an ideal Gaussian pulse with FWHM time of 150 ns. The middle graph plots the integrated energy for four gate widths that slide left and right of the center. The x-axis value of zero means the pulse is centered in the gate width period. The plot on the right provides the percent error for the same parameters as the middle plot.	124
Figure 93: Example 150 ns, 50 μ J peak pulses with peak SNR values of 10, 100, and 1000. Poisson shot noise has been applied to the signal.	125
Figure 94: Percent errors as a function of the center offset and gate width for the 50 ns pulses and peak SNR ratios of 10, 100, and 1000.....	126
Figure 95: Percent errors as a function of the center offset and gate width for the 150 ns pulses and peak SNR ratios of 10, 100, and 1000.....	126
Figure 96: General overview of the traditional design concept. The letters in yellow boxes indicate optical paths discussed in this section.....	133

Figure 97: Trade study example of displacing optics in OpticStudio and calculating the resulting focus spot diameter. In this case, the lens is translated by 5 mm and the focus spot radius is 0.67 mm.	135
Figure 98: Zemax Optics Studio plots of the focal spot for the 125 mm Edmund Optics 26-506 convex-plano optic with a 375 mm focal length (left) and the 100 mm Asphericon ALL100-100-S-U aspheric optic with a 100 mm focal length (right). The best case spot radii are 416 and 35 μm respectively.....	136
Figure 99: View of the optical layout for the quad-copter design. The base plate is 11"x8". The assembly is 7" tall. The receiver telescope optic is ~4" (100 mm) in diameter.....	139
Figure 100: View of the beam splitting assemblies for the transmit, gas cell, and power paths. The supporting structures are set to semi-transparent.	140
Figure 101: The receive telescope assembly. The lens mounts to a ring with three tabs. The tabs fit into the lower left tube and are epoxied in place after alignment. The longer tube slides over this sub-assembly. A bracket holds the detector in place and slides into the tube where it is also epoxied in place. Fasteners may be a better attachment option to facilitate modifications.	140
Figure 102: An outline of how the fixed-wing instrument could be assembled. The dimensions are 9.5" wide, 19.5" long, and 9.5" tall at the tallest end (5.5" at the shorter end). The transmitted beam (blue) travels under the received beam (green). The transmitted beam strikes the outer edge of the scan mirror which has subtle angles machined into it to reflect the transmitted beam into the receive telescope's footprint. The mirror rotates with a larger angle to reflect the light in a circular pattern on the ground. .	142
Figure 103: General overview of the VBG filtering design concept. The letters in yellow boxes indicate optical paths discussed in this section.	143
Figure 104: OpticStudio rendering of the launch optics. A single mode fiber with a NA of 0.14 is collimated with an aspheric lens before interacting with a gold knife-edge prism (not shown). A plano-concave lens then diverges the beam until a plano-convex optic with a long focal length collimates the expanded beam. A perfectly aligned system creates a beam size on the ground, accounting for polarization effects, of about 14 mm in diameter at a range of 150 m (~500 feet). Minor variations of two tenths of a millimeter from the ideal alignment can expand the beam to 88 mm in diameter at 50 meters (~150 feet).....	145
Figure 105: View of the light splitting and monitor box. A collimated beam is split by the gold prism on the left with the majority of the light being sent to a beam expanding optic on the left. A small percentage is sent through a neutral density filter before being split by the small prism on the right. One half of the path goes into a GRIN optic for focusing onto a single mode fiber. The other half goes through a gas cell before being captured by another GRIN/fiber combination.	146
Figure 106: The fiber from the receive telescope is held in place by a three-point mount. Spherical washers are meant to be used between the mount and support. The real outer support would not have the three large holes in order to increase stiffness and block stray light.	148
Figure 107: OpticStudio rendering of launching a 0.26 NA fiber into a GRIN collimator then into an offset aspheric lens and finally into another 0.26 NA fiber. The fiber diameter on both sides is 62.5 μm . ThorLabs "C" coatings are assumed on all surfaces and polarization effects have been accounted for...	149

Figure 108: View of the path combining assembly. Three GRIN collimating lenses feed into an aspheric lens that focuses onto a multi-mode fiber. The notch near the fiber (and a hidden hole under the fiber) are meant for aligning the fiber before being staked in place with epoxy.	150
Figure 109: View of the VBG wavelength splitting assembly. Multiple wavelength light is introduced from the left by a collimating GRIN assembly. The VBG splits the wavelength of interest off at an angle of 10° into a GRIN/fiber assembly. The remaining light is captured by a GRIN/fiber assembly for further filtering downstream. The assembly is mounted on a TEC for temperature control of the VBG. The TEC itself is mounted on a copper heat spreader. Titanium or Ultem spacers are used between the copper and optics bench to reduce heat transfer.	151
Figure 110: Light enters the detector from two GRIN/fiber assemblies. One path is the “online” wavelength, one is the “offline” wavelength. An aspheric lens focuses the light from both optics onto the detector. The GRIN/fiber assemblies are mounted on a 1” diameter part that threads into an optic tube mounted directly to the detector for maximum stiffness.	152
Figure 111: Overhead view of the concept instrument. The platform is 13” wide and 10” tall. The structure is 4” (100mm) in diameter and 9.8” tall.	155
Figure 112: Side view of the concept instrument. The center optical structure is a tube for maximum stiffness and light isolation. It has a height of about 9.8” (248 mm)	156
Figure 113: Opposite side view.	156
Figure 114: Traditional instrument designs’ ROC plot for a quadcopter scanning with somewhat random tip/tilt. Green means “good” probability of detecting and mapping the plume, yellow means “marginal”, and red means “unlikely”.	162
Figure 115: VBG Filtering ROC plot for a quadcopter scanning platform.	162
Figure 116: ROC plot of the traditional instrument design with 1 meter resolution scanning. Green dots represent plumes with sufficient numbers of point to reliably detect and map them. Red dots represent plumes with insufficient points.	163
Figure 117: VBG Filtering ROC plot for a scanning system with 1 meter resolution.	164

Chapter 1: Introduction

Section 1.1 begins the technical narrative, outlining the motivation for advancements in the detection of gases. Section 1.2 continues the discussion with a brief history of gas detection and where the industry has been heading. Section 1.3 more fully introduces laser based systems (the focus of this dissertation) covering both their benefits and challenges. Section 1.4 outlines the goals of this work and gives a chapter-by-chapter summary. The chapter concludes with section 1.5, covering acknowledgments in more detail than what can be inserted into the acknowledgements section.

1.1 Motivation

Natural gas is about 97% methane, with the remaining percentage being ethane (for thermogenic sources) and odorants so that people can smell it (pure methane has no scent), and company-specific chemicals that aid in keeping the gas flowing through the pipes. The Pipeline and Hazardous Materials Safety Administration (PHMSA) is charged with overseeing the transportation of materials, including natural gas, in pipelines. And this is no easy task: as of 2012, there were about 300,000 miles (~483,000 km) of natural gas transmission pipeline and 2.1 million miles [3.38 million km] of distribution pipeline [PHMSA, 2016]. Transmission pipelines refers to large volume pipe that carries natural gas from the source to primary distribution centers, while distribution pipeline carries natural gas to residential and commercial customers.

The majority of all the pipeline is greater than 30 years old and, in the case of the transmission pipeline, the majority is greater than 40 years old. Perhaps surprising to readers will be the fact that about 20% of the pipeline currently in use in the United States was built in the 1930's and 1940's. Figure 1 provides an overview of the percentage of the existing pipeline installed by decade.

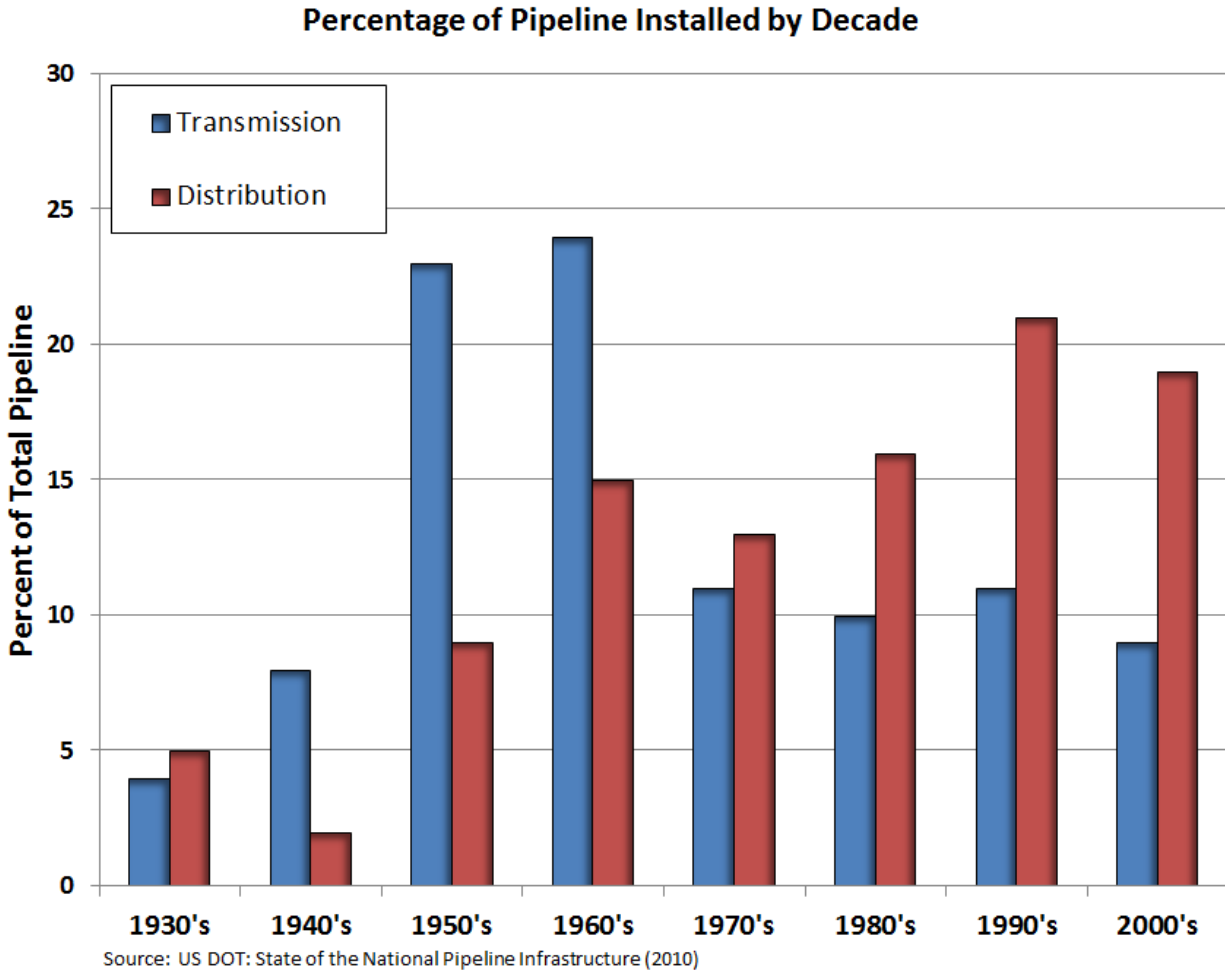


Figure 1: Greater than 50% of the 300,000+ miles of transmission pipeline was laid prior to 1970. More than 1 million miles of distribution pipeline was laid prior to 1980.

As may be expected from an aging infrastructure, failures frequently occur. Figure 2 plots the rates of “Significant Incidents” collected by PHMSA. A “Significant Incident” is defined as:

“(1) fatality or injury requiring in-patient hospitalization; (2) \$50,000 or more in total costs, measured in 1984 dollars; (3) highly volatile liquid releases of 5 barrels or more or other liquid releases of 50 barrels or more; or (4) liquid releases resulting in an unintentional fire or explosion.” [PHMSA, 2016]

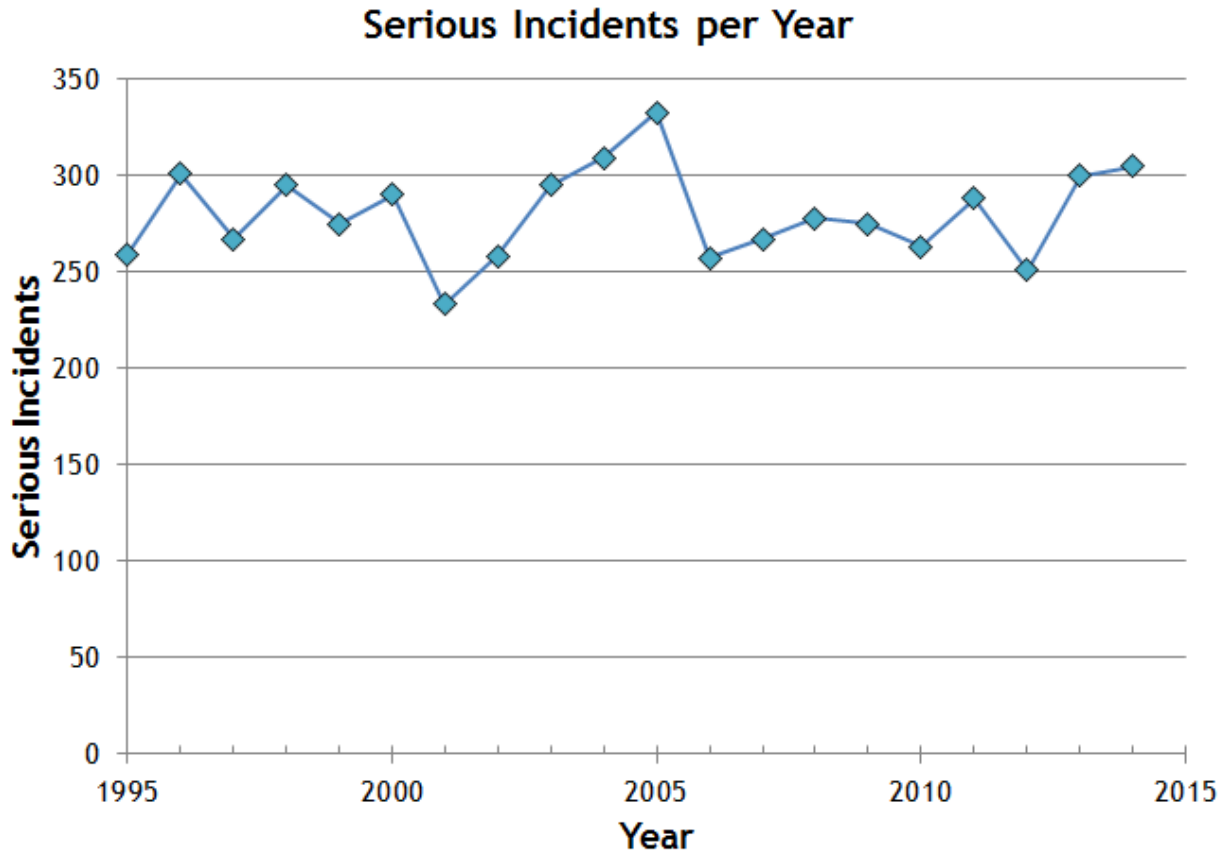


Figure 2: The number of serious accidents has held relatively steady. However, loose reporting laws and increasing pipeline miles obscure the true trends. [PHMSA, 2016]

Note that the total amount of pipeline has also increased so that the percentage of accidents is not a constant. Nevertheless, the industry widely feels that the overall trend is towards greater concern for the well-being of the infrastructure [PHMSA, 2016]. One unfortunate data point is that PHMSA listed 19 people being killed by gas pipeline accidents in 2014, of which 17 were public citizens and 2 were industry employees.

In addition to the cost of human life and maintenance is the impact that methane can have in the environment. Methane is the third most important greenhouse gas, after water vapor and carbon dioxide [IPCC,2007, IPCC,2014]. Water vapor is largely a naturally occurring part of the environment so, although understanding its role in global climate change is considered critical, there is little that can be

done to mitigate its sources. Carbon dioxide and methane also have natural, mostly uncontrollable sources such as decaying biomass and the respiration and digestion cycles of animal life. However, anthropogenic sources of these molecules are widely acknowledged to contribute significant additional quantities to the environment [IPCC,2007, IPCC,2014]. To that point, carbon dioxide has increased from about 280 ppm in 1750 to approximately 395 ppm, a 41% increase, as of 2013 largely due to fossil fuel burning [IPCC,2014]. Meanwhile, methane has increased from about 722 ppb to about 1,800 ppb over the same time period, a nearly 150% increase [EPA, 2015]. Figure 3 plots the increase in methane concentrations measured at three sites since the 1980's.

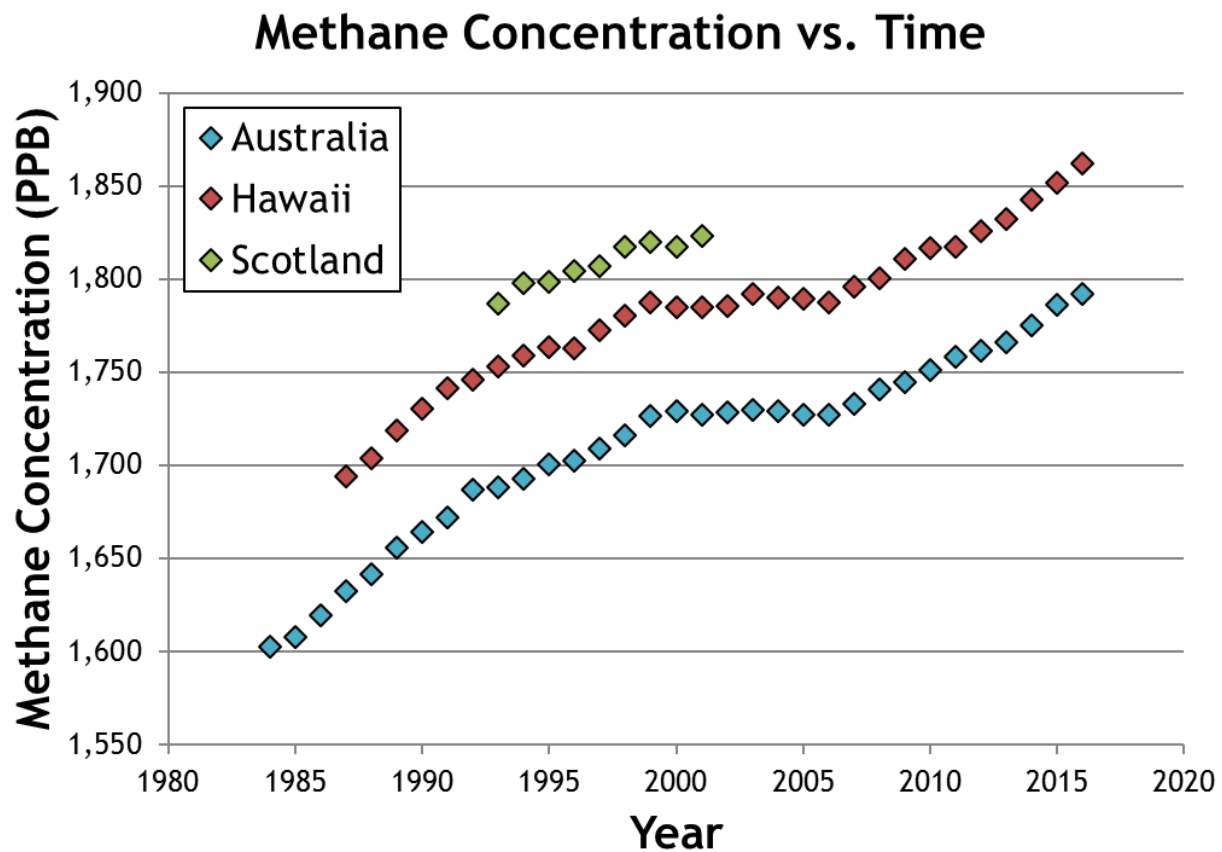


Figure 3: Directly measured methane concentrations have increased since dedicated measurements began in the 1980's. [Steele, 2005, Hawaii, 2017, Australia, 2017]

Carbon dioxide receives more attention than methane primarily due to its abundance. For instance, from 2010 to 2014, roughly 35 to 40 gigatonnes of carbon dioxide were released each year [IPCC, 2014, Le Quéré, 2013]. Meanwhile, methane release rates are about 0.5 gigatonnes per year [IPCC, 2014].

However, although the scales tip to carbon dioxide in terms of mass, methane has higher radiative forcing. According to the Intergovernmental Panel on Climate Change (IPCC),

“Radiative forcing is a measure of the influence a factor has in altering the balance of incoming and outgoing energy in the Earth-atmosphere system and is an index of the importance of the factor as a potential climate change mechanism” [IPCC, 2014].

The same references estimate that, over a 100-year time span, methane has between 21 and 25 times more effect on global warming per molecule than carbon dioxide. The IPCC thus uses an “equivalent carbon dioxide” conversion for methane emissions to try to compare the molecules more evenly. Applying their factor of 25 gives a carbon dioxide equivalent methane release of 12.5 gigatonnes per year. Therefore, methane has about 25% of the impact of carbon dioxide. The summation of all greenhouse gases, including molecules such as Nitric Oxide, results in methane providing about 16% of the total.

In the United States, the Environmental Protection Agency tracks the anthropogenic sources of methane release [EPA, 2015]. Figure 4 breaks down the primary sources. The largest contributor of all of the known sources is the oil and gas industry at 29%. Within this category are sources such as the oil industry venting methane as a byproduct of the drilling process and, key to this dissertation, transportation losses due to leaks [EPA, 2015].

Continuing with discussion of the U.S. methane emissions, the next largest contributor is enteric fermentation, which refers to the digestive process that breaks down carbohydrates in some animals to produce methane. In the U.S., the primary contributor in this category is the cattle industry [EPA, 2015]. The “Other” category includes items such as wastewater treatment at 2.3%, rice cultivation at 1.3%, and abandoned underground coal mines at 1%.

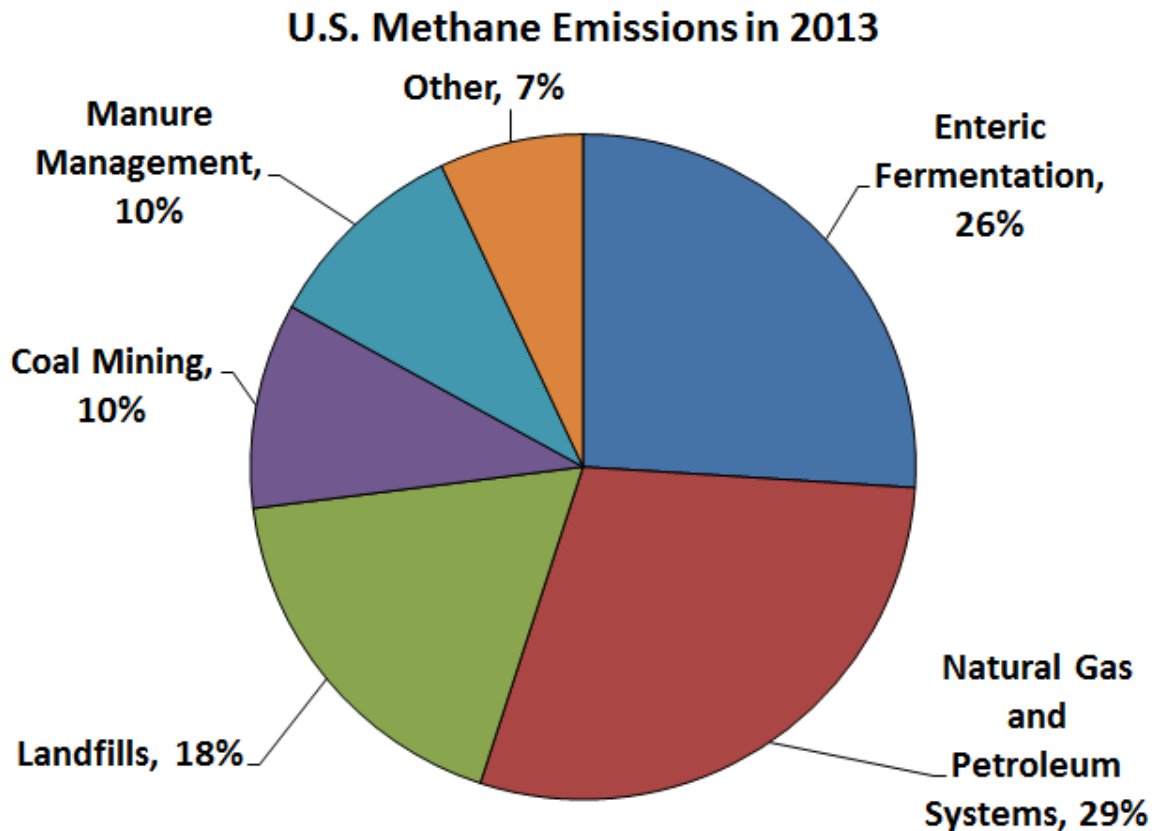


Figure 4: The Environmental Protection Agency tracks methane sources within the United States [EPA, 2015]

Of the many categories tracked by the EPA, many are anthropogenic without easy solutions. For instance, 18% of the methane comes from landfills, but short of forcing less waste the government has little they can pursue. Likewise, if more people chose to reduce their meat consumption then perhaps the enteric fermentation and manure management categories could be reduced, but this is not something the government is likely to try to enforce. Indeed, one of the few major categories that can be improved is the way natural gas and petroleum systems are operated.

This tall pole had not gone unnoticed by regulators. The impact that methane has in global climate change, the aging infrastructure, and the number of significant incidents has resulted in regulations and proposed regulations in the natural gas industry. In 2014, Colorado became the first U.S. state to pass

laws regarding mandatory annual inspections of wells and transmission pipelines [Colorado, 2014]. In June 2016, the Environmental Protection Agency drafted rules after several years of community feedback [EPA Press Release, 2016]. (Note that the new rules are currently on hold per direction of the Trump administration.) Among other requirements, they propose that all new natural gas well sites be inspected for leaks within 30 days and then twice annually thereafter.

These inspections are not cheap. According to the EPA, a well site inspection can cost \$25,000 and per-mile costs of pipeline inspections can be \$250 or higher [PG&E, 2015]. With over 6,000 currently operating gas wells in Colorado alone, this represents up to \$150 million in inspection costs [CO OGCC, 2015, Denver Post, 2015]. Most of these inspections are done manually with union labor and hand held instruments [PG&E, 2015].

However, although it may be expected that the industry would find the frequent inspections onerous, some of their concern is mitigated by the substantial costs that they pay each year to cover property damage caused by infrastructure failures. Figure 5 below plots the reported cost burden for U.S. providers of natural gas [PHMSA, 2016]. It is seen that damage has exceeded greater than 1 billion dollars twice in ten years and that the overall trend has been increasing over time. Note also that these costs do not cover lost production time and miscellaneous overhead costs associated with managing and repairing the damage.

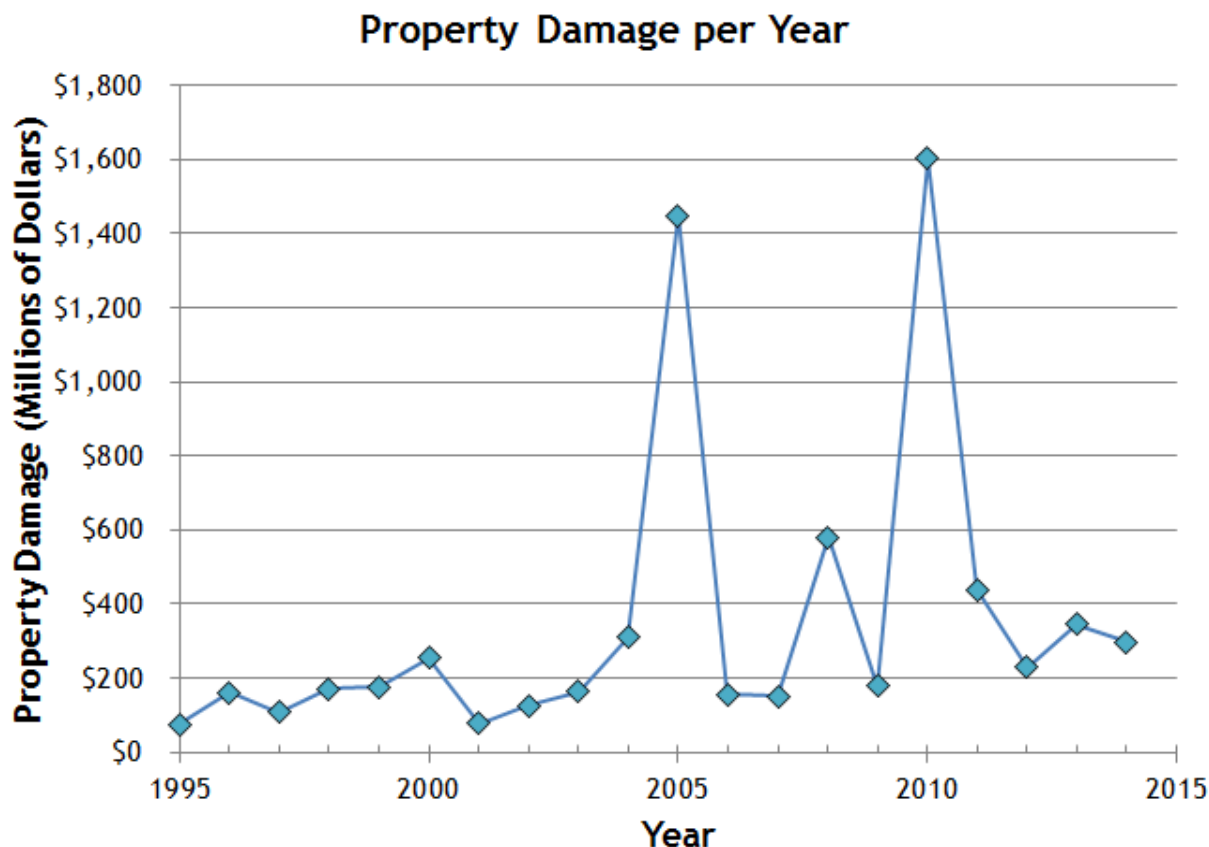


Figure 5: The property damage per year shows an increasing trend [PHMSA, 2016]

In summary, the confluence of improved understanding of methane's impact on the environment and the regulations introduced to combat aging infrastructure and increasing emissions provides an opportunity to make an impact through rigorous engineering. And, not incidentally, methane is just one of many molecules where having an economical, rapid, and technically sound data collection approach may help make a difference. Other gases, such as CO₂, NO₂, and hydrocarbons impact the environment and behave similarly to methane such that finding a solution for one may create a solution for others. The potential to push technology forward in a way useful to addressing fugitive emissions of all types motivates this dissertation.

1.2 Methane Overview

In order to understand how to detect natural gas, it is first helpful to understand some of its properties. To start, so called “natural gas”, can have anywhere between 75% and 98% methane, depending on the processing stage of the gas [naturalgas.org, 2017]. In turn, methane is the common name for the molecule CH_4 . Figure 6 shows a view of the structure. A carbon atom occupies the central role and four hydrogen atoms take up tetrahedral positions on the outside. Each hydrogen atom shares its electron with the carbon atom, bringing the total electron count around the carbon atom to 10, which is a stable, preferred result. The generally accepted mean diameter is about 0.4 nm. In chemical terms, methane is a saturated hydrocarbon and is classified as an alkane. [Mao, 2001] An interesting historical side note is that it was first suggested worth studying in 1776 by Benjamin Franklin who referred to it as “flammable air”. This suggestion led to an Italian effort that isolated and characterized it in 1778 [Volta, 1777].

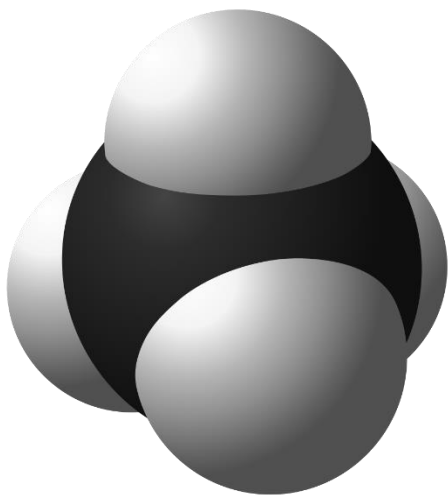


Figure 6: Methane has one carbon atom (black) with four hydrogen atoms (white) in a tetrahedral arrangement [18]

The bonds between the carbon and hydrogen atoms lead to a number of absorption features based on translational and rotational modes. These modes can be theoretically analyzed, but the theorist must take into account pressure and temperature to compensate for atmospheric broadening of the modes. Fortunately, the Pacific Northwest National Library has gone to great lengths to characterize many different gases with in situ testing. Figure 7 shows their results for methane across a broad range of wavelengths [PNNL, 2015]. The subset figures progressively zoom on the 1,645 nm region.

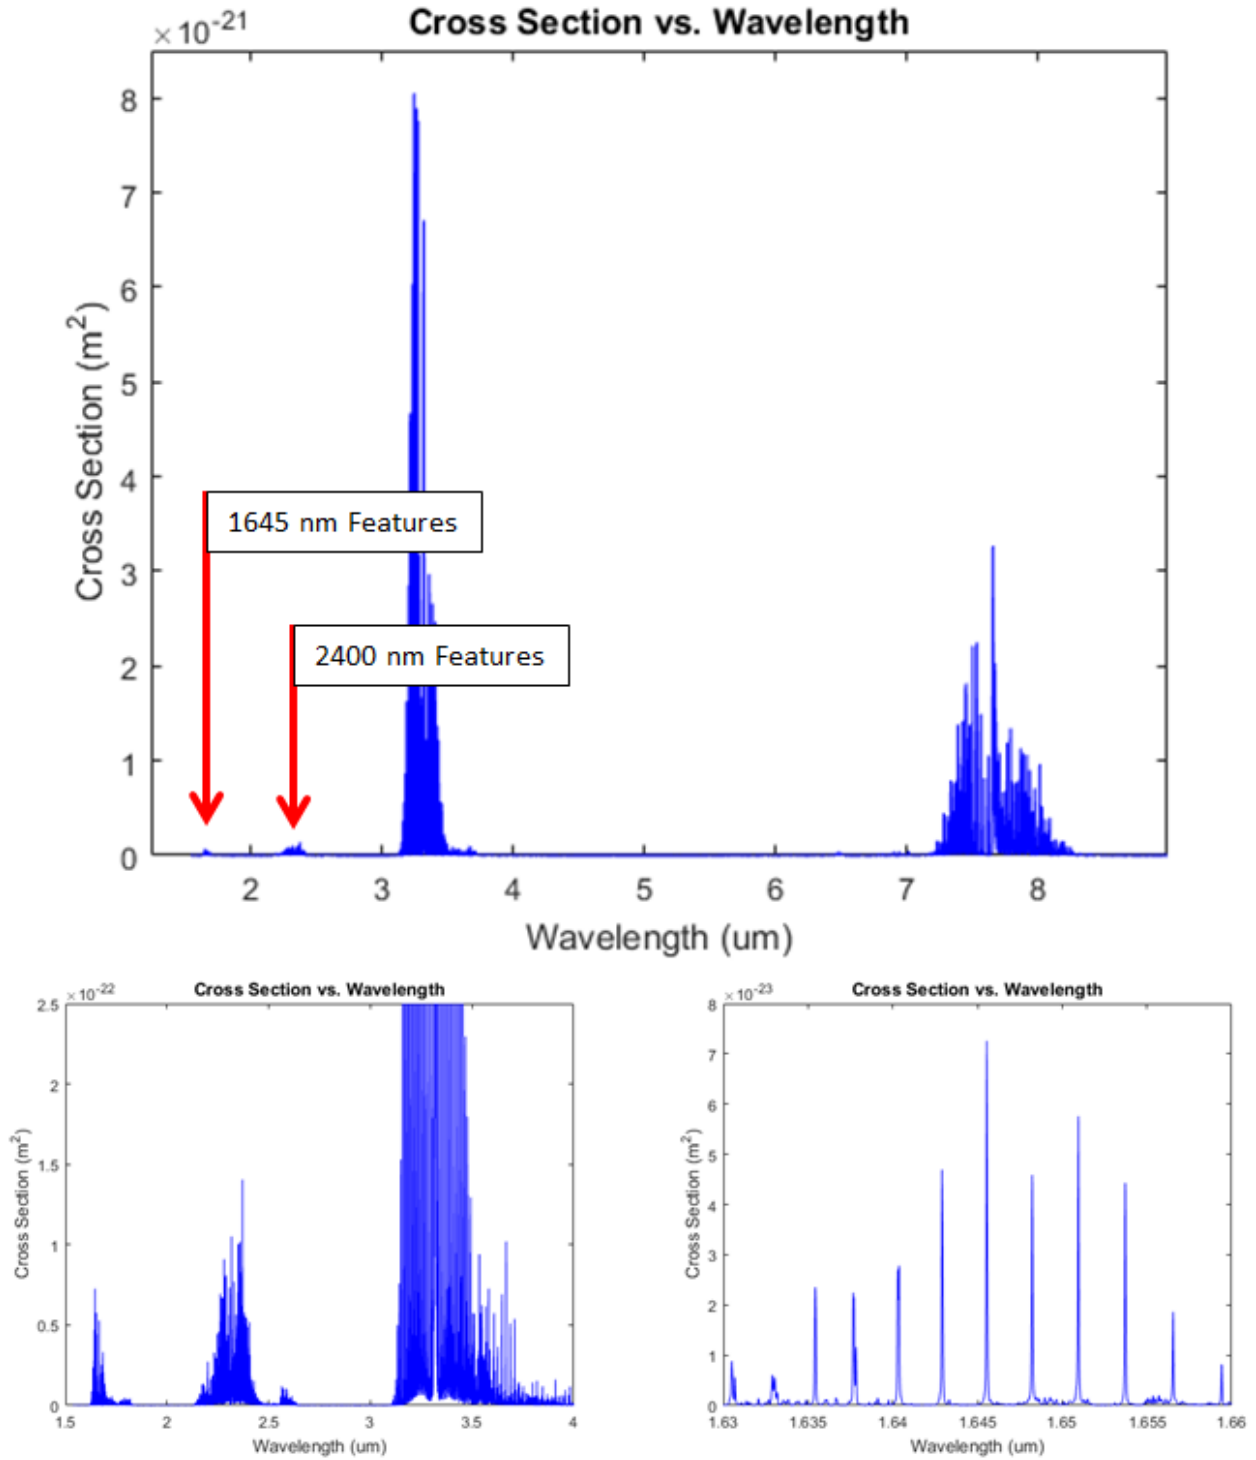


Figure 7: The PNNL database for methane provides atmospheric compensated features. The subset pictures demonstrate progressive zooms into the features around 1,645 nm.

1.3 A Brief History of Gas Detection

Sensing the presence of methane has, historically, been as much a practical undertaking as a financial one: coal mines are noted sources of methane and can asphyxiate miners as well as create explosion hazards. The challenge is that methane is both colorless and odorless. (The “smell” that most people associate with natural gas is caused by sulfur based molecules such as tert-butylthiol or mercaptan that are intentionally added to aid detection by humans.)

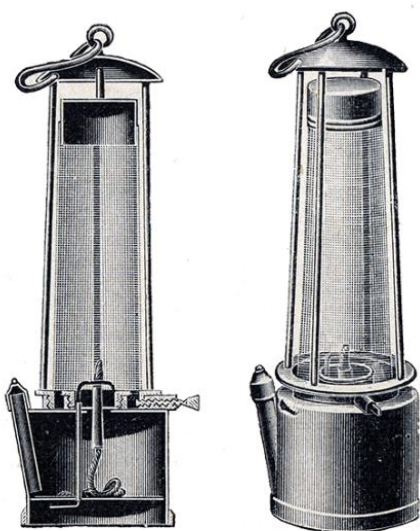


Fig. 192. Davysche Sicherheitslampe

Figure 8: The Davy Lamp burned brighter in the presence of methane and dimmer in the presence of CO₂. [Davy, 1816]

1.3.1 In Situ Devices

The first engineering solution was the Davy Lamp, invented by Sir Humphry Davy in 1815 and shown in Figure 8 [Davy, 1816]. Before entering a mine, the miner would light the lamp and observe its brightness with a gauge installed on the side. The miner would then periodically check on the flame while working underground. If the flame was brighter it suggested the presence of methane. If the flame was dimmer, it suggested carbon dioxide. In either case, the miner was theoretically alerted to a problem. Unfortunately, flames

are fickle things with many variables that control its brightness. The result was, unfortunately and unexpectedly, an increase in mine accidents as miners over trusted its effectiveness so worked in mines that would otherwise have been abandoned [Lawrence, 1990].

The solution in the early 1900's was the literal “canary in the coal mine” in which a canary was observed to react to a reduction in oxygen faster than the miners. When miners saw the canary's energy falter they knew it was time to evacuate the area [Boycott, 1908]. Although effective at both alerting miners to potential danger and as a lasting idiom, the canaries could not differentiate between methane, carbon dioxide, carbon monoxide, or any other oxygen displacing molecule. The unreliable Davy lamps and

uncertain canaries thus lead to increasing attention in the scientific community. As a result, the 1920's saw a replacement of the canaries and flame based "sensors" with two inventions that are the basis of most of the methane sensors in operation today: the catalytic combustion sensor and the gas detecting interferometer.

The catalytic combustion sensor was developed by Dr. Oliver Johnson in 1926-1927 [Kohl, 2014]. At the time, storage tanks and ocean going tankers were exploding frequently enough that Dr. Johnson was hired to develop a more reliable instrument. His solution was based on two principles. The first was that a catalyst oxidizes flammable gases at lower levels than they do in air and the second was that a hot wire has higher electrical resistance than a cold wire. Dr. Johnson used a Wheatstone bridge with two of the resistors made with the catalyst. Of the two, only one was exposed to the gas and thus would change temperature and resistance with changes in gas concentrations. The change in its resistance was shown on an analog gauge. Since then, the same principles of a catalyst and changing resistance have been used in various similar concepts such as pellistors and catalytic beads.

Nearly in parallel with the development of the catalytic sensor has been the development of the gas sensing interferometer. Interferometry was developed in the late 1800's. Its application to gas detection started in Japan in the late 1920's to deal with the same exploding ocean-going tankers [Riken Keiki, 2016]. The physical concept is that methane and other flammable vapors diffract light and create fringe lines. By using a long enough path length, the diffraction is sufficient to put the fringes on a card readable by the user to measure the gas concentration. Since then, the card has been replaced by electronic means to detect and calculate the fringe lines.

Examples of modern instruments that can claim heritage with these ancestors are the Flame Ionization Device, ring down cavity spectrometers, and nearly every "system on a chip" device that must be within the methane to sense it.

1.3.2 Lasers and Methane Sensing

The path to using laser technology to detect methane began in 1964 when R.M. Schotland introduced the first Differential Absorption Lidar (DIAL) approach to measure atmospheric water vapor [Schotland, 1964]. This came just four years after the first laser was developed by T.H. Maiman in 1960 [Maiman, 1960]. The concept is straightforward in principle: Send two laser pulses, where one laser is absorbed by methane and the other is not. If everything else is equal, including surface reflectivity, non-methane atmospheric absorption, and geometry, then the scientist can simply ratio the returned pulse energy to determine how much methane was in the laser's' path. A more thorough treatment of DIAL's technical aspects is provided in the technical section.

From Schotland's work came various systems measuring atmospheric gases, but the next major milestone relevant to this dissertation was the first airborne DIAL system. It was developed by R.T. Menzies and M.S. Shumate for flight campaigns over California and the Pacific Ocean in 1977 [Menzies, 1977]. In their work, they were able to reduce the size of the system such that it fit on a twin engine Beechcraft Queen Air. Although a twin engine aircraft, it was only capable of carrying between 2,000 and 2,4000 pounds so fitting operators and instruments in the airplane represents a significant engineering achievement [Beechcraft, 2016]. Further, they were the first to use the earth's surface as a reflecting medium by flying 2 km above it and pointing nadir.

Specific methane laser sensing technology first came into consideration with K. White's brief 1972 paper on the coincidence of the laser emission line of Erbium doped Yttrium, Aluminum, Garnet (Er:YAG) crystals and the known methane absorption line at 1645.1 nm [White, 1972]. Three years later, K. White continued with his 1975 paper on directly suggesting the use of Er:YAG lasers to detect atmospheric methane [White, 1975]. In parallel, other researchers were pursuing additional work in the mid-infrared regions near 3.39 micrometers [Barger, 1973, Hubbert, 1977, and Murray, 1978].

Using a DIAL approach for methane evolved throughout the 1970's. Indeed, there were two dozen research laboratories, including 12 in the United States, actively pursuing DIAL [Murray, 1978]. As described in this reference, the difficulty faced then was the lack of laser power to reliably bounce signals off of topographic targets instead of carefully controlled laboratory reflective surfaces.

Also relevant was the division of camps when it came to choosing a wavelength to focus on. Specifically, some researchers preferred the eye-safe near-infrared region for its safety and applicability in populated areas. This wavelength range comes at the cost of reduced absorption, as shown previously in Figure 7, so some researchers preferred the mid-infrared regions. Although this meant the laser power could be lower, achieving this wavelength has traditionally relied on optical parametric oscillator technology, which has a reputation for being complex and finicky [Bartholomew, 2015]. This debate continues today as modern systems utilize both regions.

Modern efforts in laser-based spectroscopy also continue with partnerships such as NASA's Langley and Exelis on their CO₂ detecting systems. Dobler and Browell, in particular, have been prolific with their papers covering a wide range of platforms and upgrades for their approach [i.e. Dobler, 2013]. The European Space Agency also funded the Advanced Space Carbon and Climate Observation (A-SCOPE) mission in preliminary work that involved an airborne DIAL system for CO₂ measurement [Amediek, 2009]. Although not specifically applicable to this dissertation, their contributions helped shape an understanding of the challenges and solutions involved in a modern methane DIAL instrument. The Amediek paper on reflectance, in particular, has useful data from flight campaigns that are applicable to methane's near infrared absorption wavelengths.

Methane is being targeted on both sides of the Atlantic with major efforts. France and Germany's MERLIN (translated in English as Methane Remote Sensing Lidar Mission) is a space-based platform targeted for launch in 2021. High quality papers from this effort are used within this dissertation, including the excellent discussion of instrument sensitivity by Kiemle [Kiemle, 2011]. On the western side of the Atlantic, ITT Manufacturing supported a long effort in an airborne methane DIAL system

[Steven, 2011]. Their Airborne Natural Gas Emission Lidar (ANGEL) system used an OPO laser in the 3 μm region and participated in government trials [Murdock, 2008]. That result failed to garner commercial success, but produced a new generation of researchers. Jarett Bartholomew, in particular, has established a new foothold with commercial methane detection at Ball Aerospace that targets methane in the 1.645 μm region [Bartholomew, 2015]. Another company currently operating commercially is Lasen, a helicopter-based instrument targeting the 3.4 μm region [Geiger, 1993]. They have not published extensively so details are scant.

Application of DIAL from an unmanned vehicle has also been investigated with a thorough treatment by Tamer Refaat in 2013 with their paper on ground, aircraft and UAV platforms [Refaat, 2013]. For their purposes, aircraft and UAV are analyzed at altitude of 15.24 km and 8 km respectively. Although this UAV altitude is much higher than what might be considered for local analysis, the equation heavy paper provides excellent material used throughout this dissertation.

This dissertation acknowledges the many contributions from a half-century of researchers. Their efforts have established DIAL technology as an industry-accepted way of finding real-world leaks to an extent that allowed for this government supported work (See section 1.5 for details). To wit, discussions with industry and government representatives can be summarized as “Methane DIAL technology is really great. When can we put it on a UAV to help us find leaks?” [i.e. Wood, 2015; PG&E, 2015; PHMSA, 2015].

1.4 Dissertation Focus and Outline

The focus of this dissertation, then, is to investigate and lay the foundation for a practical instrument optimized for use from an unmanned vehicle flying within the FAA’s 500-foot ceiling limit. This involves developing an end-to-end simulation approach to predict overall performance, developing hardware concepts, and analyzing these concepts. The final result is size, weight, and cost estimates for

the transmitter and receiver parts of the instrument as well as the expected leak rate detection performance.

The effort begins with Chapter 2: Technical Foundations. This chapter more fully introduces DIAL's technical underpinnings. It then looks at available UAV platforms to determine payload size, weight, and power (SWaP) targets. The remainder of the chapter discusses laser technology and lays out justification for choosing the instrument's baseline properties.

Chapter 3: Gaussian Plume Simulations develops a core contribution to the analysis effort. The equation often found in literature is extended via integration and unit conversion to provide concentrations as would be measured by an airborne DIAL instrument reporting values for the industry. Plume simulations are then correlated with in situ data collected by Ball Aerospace's DIAL instrument to show applicability to instrument performance prediction work.

With that platform in place, Chapter 4: Analysis Trade Studies mines the extended Gaussian plume equation to develop plots and guidelines useful for estimating the required instrument performance. For instance, it is shown that beam diameter strongly drives the minimum visible leak rates. The outcome of the various analyses is a Receiver Operator Curve (ROC) prediction for plume detectability as function of leak rate and wind speed.

Chapter 5: Wavelength-Based Reflectivity Variations is a modified version of a peer-reviewed publication: "Analysis of the impact of wavelength separation on reflectivity error for differential absorption lidar using the ASTER spectral library" published in SPIE's Journal of Remote Sensing [Tandy, 2017]. The paper investigated the oft-made assumption that reflectivity differences between the online and offline lasers is negligible as long as they are "close" to each other. The conclusion is that "close" is a function of the center wavelength and can be anywhere from tenths to several nanometers. Example applications for the plots are provided. An extension of the work in the paper is added with an analysis of applicability for the wavelength chosen for the UAV instrument.

With the simulation analyses developed, attention next turns to hardware in Chapter 6: Hardware Planning and Considerations. This chapter begins by introducing two instrument concepts. It then looks at sundry aspects of hardware design. Much of these analyses came about from hands-on development work with Ball Aerospace's methane instrument. A number of observations and resolutions drove research into fundamentals in order to both solve the challenges and to gain a deeper understanding of the underlying issues. These analyses may be useful to others.

Chapter 7: Baseline Design and Calculation is a short summary of the various work to that point and the requirement goals for detection. It then outlines how the pieces work together to drive an instrument concept.

Chapter 8: Final Hardware Design is a detailed optical, cost, and weight analysis for the transmit and receive concepts outlined in Chapter 6. OpticStudio's Zemax and AutoDesk's Inventor are used to run trade studies on manufacturability and performance. As the emphasis is on a practical design, components are chosen from current vendor offerings so that a bench-top system could be fairly quickly developed. The efficiencies are rolled into a signal-to-noise ratio estimate, which is then rolled into the end-to-end analyses developed previously in the dissertation. The result is a performance, price, and weight comparison between the instrument concepts.

Chapter 9: Summary, Conclusions, and Next Steps wraps up the dissertation with a discussion of the summaries and conclusions of each chapter. Discussion then covers the overall lessons learned. This work ends with an outline of the next steps that could be taken to further this work.

1.5 PHMSA Grant and Ball Collaboration

This section briefly covers the PHMSA grant and Ball Aerospace contributions in more detail than what is available in the acknowledgment section.

Funding for this dissertation came from the Pipeline and Hazardous Materials Safety Administration (PHMSA). A proposal was submitted to their request for leak detection concepts. Our group was notified

that we had won Award DTPH5615HCAP11, for Proposal 0415.07.0998B - Small-Scale DIAL for Methane Detection. The two-year, \$300,000 grant funded two PhD students and included a small stipend for hardware.

The award also cemented the topic of this dissertation: designing a UAV-sized instrument intended for pipeline leak detection. The funding was not sufficient to purchase a laser, but components could be purchased to test individually. There was also an expectation that our team would dig deep into technical details to frame future efforts.

Our team is grateful to Jim Merritt, the PHMSA manager overseeing our project. He also managed an award won by Ball Aerospace. His ability to link the two organizations in a collaborative way helped give the students a foundation to build on while introducing us to wider industry challenges.

Our team is also grateful to Ball Aerospace for their mentorship and access to hardware. Quite simply, this dissertation would be much poorer without the ability to test concepts in their labs or ask questions of the technical team. Jarett Bartholomew, in particular, went above and beyond in answering many naïve questions in the beginning and becoming a collaborative partner as our understanding grew.

There was no obligation on their part (i.e. no money) to motivate the extent of their aid. Ball Aerospace has designed and is currently operating a full-sized methane DIAL system [Bartholomew, 2017]. While some information gained from their system is withheld for competitive reasons, they largely shared information that may one day help create a competing instrument. Additionally, the overlap of some concepts enabled access to their laser, detectors, software, laboratory space, and many different types of test equipment. In return, we offered them the technical insights gained as we developed our ideas. Some of them, such as the reflectivity study, were used in their own work.

Chapter 2: Technical Foundations

The previous chapter outlined some of the history and impetus of methane detection. This chapter lays the technical groundwork for the trade studies, design concepts, and analysis of the rest of the dissertation.

First, the DIAL equations and discussion are offered. Focus then moves to UAV platforms and general design targets for both fixed-wing and quadcopter vehicles. The third sub-section discusses wavelength choices. The final sub-section looks at continuous wave vs. pulsed systems.

2.1 DIAL Technical Details

As mentioned, the DIAL approach uses two different laser wavelengths and measures how much energy they lose as they interact with the atmosphere. One laser is designed so that it strongly interacts with methane and is called the “on line” or, more generally, “online” laser. “Strongly” loosely means it is absorbed significantly more by methane than any other atmospheric constituent at that wavelength. How strong it needs to be is part of the design trade study process. Part of that process is choosing the second wavelength, which is the “off line” or, again more generally, “offline”. This second laser is chosen so that it weakly interacts with methane while interacting with the other gases along the path nearly identically as the online wavelength. Figure 9 graphically demonstrates the concept. A common approach (not shown) is to add a reflective target in the path so that the receive and transmit optics are coincident.

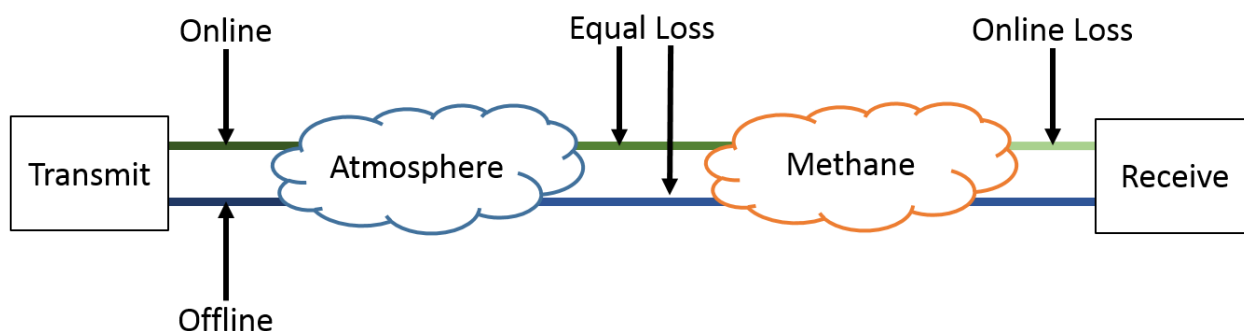


Figure 9: Different lasers interact with the gas in different ways. Measuring those differences provides insight into the quantity and, sometimes, the distribution of the gas.

Before diving into design considerations, it may be useful to take a look at the math behind the science. In full detail, a single laser (either online or offline) has behavior that can be modeled with [Chu, 2017]:

$$N_S(\lambda, R) - N_B(\lambda) = N_L(\lambda) \cdot \frac{A}{R^2} \cdot \eta_d(\lambda) \cdot \eta_r(\lambda, R) \cdot \rho(\lambda) \cdot T(\lambda, R) \quad \text{Eqn. 1}$$

N_S is the number of signal photons

N_B is the number of background photons

N_L is the number of transmitted photons

λ is the Wavelength

R is the altitude above the ground

A is the receiver area

η_d is the quantum efficiency of the detector

η_r is the optical efficiency (including geometrical overlap)

ρ is the reflectivity of the ground (or target)

T is the transmission, related to optical depth

The transmission itself is an equation. For airborne operations, it is double-path integration [Chu, 2017]:

$$T(\lambda, R) = \exp \left(-2 \int_0^R \sigma_{gas}(\lambda, r) \cdot n_{gas} dr - 2 \int_0^R \sum_{i=1} \sigma_i(\lambda, r) \cdot n_i dr \right) \quad \text{Eqn. 2}$$

σ_{gas} is the cross section of the desired gas

σ_i is the cross section of each of the other gases

n_{gas} is the number density of the gas

n_i is the number density of each of the other gases

The variable we seek to find is n_{gas} , the number density of the target gas. Two things complicate this.

First, it is deeply buried in the equation set. The second challenge is that we also need to know n_i , the number density of all of the other gases, which is nontrivial. This latter challenge is met by introducing a

second signal. Dividing two signals to create a ratio can cancel out the unknown number densities if the signal loss due to those constituents for the two lasers is the same (or very nearly the same). The equation for this looks like [Chu, 2017]:

$$\frac{N_S(\lambda_{off}, R) - N_B(\lambda_{off}, R)}{N_S(\lambda_{on}, R) - N_B(\lambda_{on}, R)} = \frac{N_L(\lambda_{off})}{N_L(\lambda_{on})} \cdot \frac{\eta_d(\lambda_{off})}{\eta_d(\lambda_{on})} \cdot \frac{\eta_r(\lambda_{off}, R)}{\eta_r(\lambda_{on}, R)} \cdot \frac{\rho(\lambda_{off})}{\rho(\lambda_{on})} \cdot \frac{T(\lambda_{off}, R)}{T(\lambda_{on}, R)} \quad \text{Eqn. 3}$$

Where λ_{off} and λ_{on} represent the offline and online wavelengths respectively. Notice that the A/R^2 term has canceled. The assumption is that the same receiver is being used for both signals, therefore the range and optical receiver area should be identical. In cases where this may not be true, the term can be appended to the equation.

From this point, the equation must be manipulated to bring out the number density of the target gas. The integrals somewhat complicate the situation, as does the exponentials. The final form of these operations looks like [Chu, 2017]:

$$n_c = \frac{1}{2\Delta\sigma_C(R)} \frac{d}{dr} \ln \left[\frac{N_S(\lambda_{off}, R) - N_B(\lambda_{off}, R)}{N_S(\lambda_{on}, R) - N_B(\lambda_{on}, R)} \right] - \frac{1}{2\Delta\sigma_C(R)} \frac{d}{dr} \ln \left[\frac{\eta_d(\lambda_{off}) \cdot \eta_r(\lambda_{off}, R)}{\eta_d(\lambda_{on}) \cdot \eta_r(\lambda_{on}, R)} \right] - \frac{1}{2\Delta\sigma_C(R)} \frac{d}{dr} \ln \left[\frac{\rho(\lambda_{off})}{\rho(\lambda_{on})} \right] - \frac{1}{\Delta\sigma_C(R)} [\alpha(\lambda_{on}, R) - \alpha(\lambda_{off}, R)] \quad \text{Eqn. 4}$$

Here, $\alpha(\lambda, R)$ is the extinction coefficient and contains both the scatter and absorption terms. The $\Delta\sigma_C$ is the difference in absorptive cross-section for the gas constituent.

The next step is to look for ways to reduce this equation further. The natural logarithm terms reduce to zero if their argument is one, which implies that the system should be designed to make the online and offline values for the arguments be equal wherever practical. Ideally, all of the ratios in lines 2 through 4 in the equation above are equal. This is plausible for things such as reflection strength and optical efficiency if the wavelengths are “very” close together, although one of the topics tackled in this dissertation is how true the reflectance assumption is as a function of wavelength.

A final simplification is to take a look at the differential aspect of the equation. The equation is taken as a function of distance. For long distances, such as from the ground to the ionosphere, the atmosphere can change dramatically. However, for distances on the order of a couple of hundred feet, the atmosphere is likely to be constant. For the purposes of this stage of the design, no assumptions are made or taken. Part of the proposed research is to understand what the impacts of the assumptions are and determine when simplifying decisions can be made.

2.2 UAV Platforms

The PHMSA grant supporting this effort specifically calls for the design of a UAV instrument. Interviews with gas industry representatives have also showed an enthusiasm for the idea [PG&E, 2015]. As part of the discussion of foundational work, UAV capabilities are discussed here.

The vehicle choice affects everything from scanning options to laser power because limits on the available payload weight and power touch nearly all aspects of performance. Another relevant factor for civil, non-military, UAV instrument developers is that, as of 2017, unmanned vehicles are not allowed to fly higher than 500 feet (152 m) above the ground. [FAA, 2015]. While this rule is in effect, an instrument developer targeting this platform will likely want to maximize performance for around this altitude. With these thoughts in mind, this section looks at two UAV categories: fixed-wing and quadcopters.

2.2.1 Fixed-Wing UAVs

For the fixed-wing aircraft, a search of manufacturers provided the data for Figure 10. The raw data is provided in Table 1. These vehicles are classified as small and medium class. The large class, such as the Predator military aircraft capable of carrying a couple hundred kilograms of weaponry, are not investigated here as their capabilities are likely to exceed the needs of gas monitoring instruments.

A couple of things are worth noting about the data. First, the manufacturers have primarily tailored their designs for the military. Insitu, Textron, and Arcturus have the most capable UAVs and are also the ones

who have received government contracts from militaries around the world. Second, these three have built mid-sized platforms with similar performance, with the exception of the Textron Shadow V2, which is nearing large-UAV capabilities. Finally, although not shown in the plot, all of the aircraft promise flight times of at least 10 hours, with many going to 24 hours, although there are disclaimers that more payload weight reduces the endurance. The large numbers of similarities indicate that UAV's in this class has been as optimized as possible with today's technology.

Based on the data, a nominal fixed-wing UAV platform performance target is considered to be:

1. A payload power of 350 W
2. A cruise velocity of 55 knots (28.3 m/s)
3. A payload weight of 16 kg (35 lbs.)

These are the values that will be used in the trade space throughout this report.

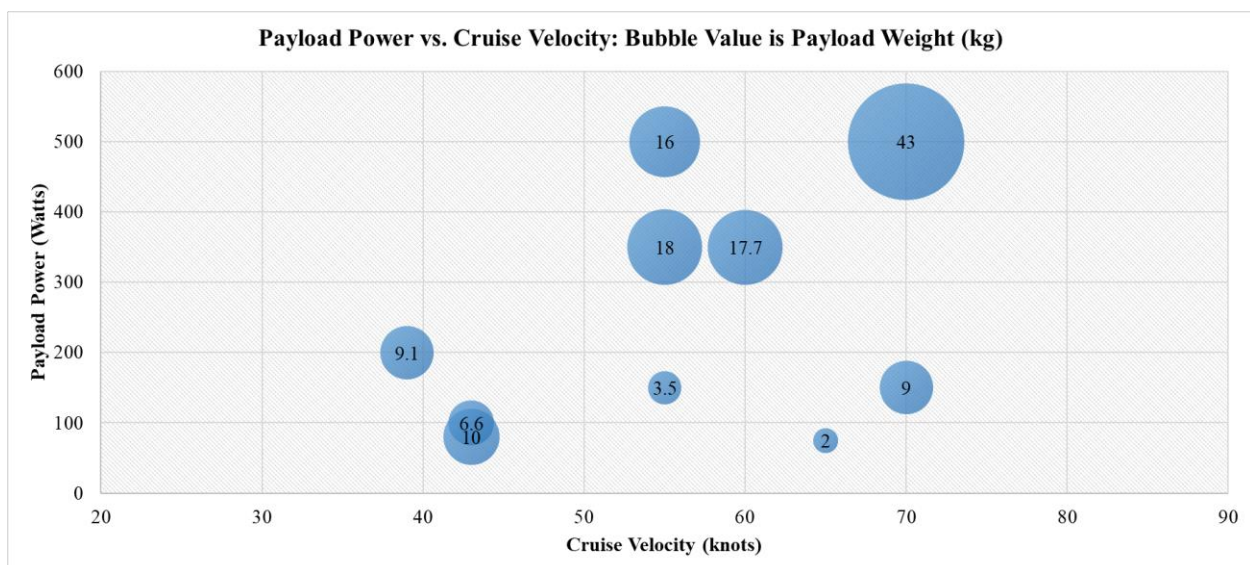


Figure 10: Plot comparing the values of available payload power, maximum payload weight, and the designed cruise velocity.

Table 1: Listing of commercially available mid-sized UAVs. Data comes from vendor websites.

Manufacturer	UAV Name	Flight Speed (knots)	Endurance (hours)	Payload	
				Weight (kg)	Power (Watts)
Insitu	ScanEagle	55	24	3.5	150
Insitu	Integrator	55	24	18	350
Insitu	RQ-21A Blackjack	60	16	17.7	350
UAV Factory	Penguin B	43	20	10	80
UAV Factory	Penguin B-Electric	43	1.5	6.6	100
UAV Solutions	Phoenix FW 240		6	11.4	
Martin UAV	V-Bat	90	8	3.1	
Martin UAV	Super Bat DA-50	65	10	2	75
Martin UAV	Bat 4	70	12	9	150
Arcturus	T-20	55	16	16	500
UASUSA	Tempest	50	1.5	4	
UASUSA	Recon	35	1	1	
Textron	Aerosonde-G	39	14	9.1	200
Textron	Shadow V2	70	9	43	500

2.2.2 Quad-Copter UAVs

For the quadcopters, the field is still in its nascent stage. Most vehicles are designed for hobbyists who may carry a lightweight GoPro or other small camera. There are quadcopters that can carry heavier payloads. However, this comes at the expense of flight endurance. All of the quadcopters found that featured meaningful capability promised no more than 20 minutes of flight time. Additionally, none of the vehicles provide payload power, meaning that part of the payload capacity must be reserved for sufficient battery power to run the instrument for the duration of the flight.

The graphical display of the data is provided in Figure 11 and the raw data in Table 2. Choosing an average is more challenging with few heavy lift options. Instead, the maximum payload option is chosen based on it having the same advertised endurance time as the other options. It is also expected that the

need to carry batteries for the instrument is going to drive a need for the most payload capacity available. Based on these data points, the instrument platform analysis will use a payload capacity of 20 kg.

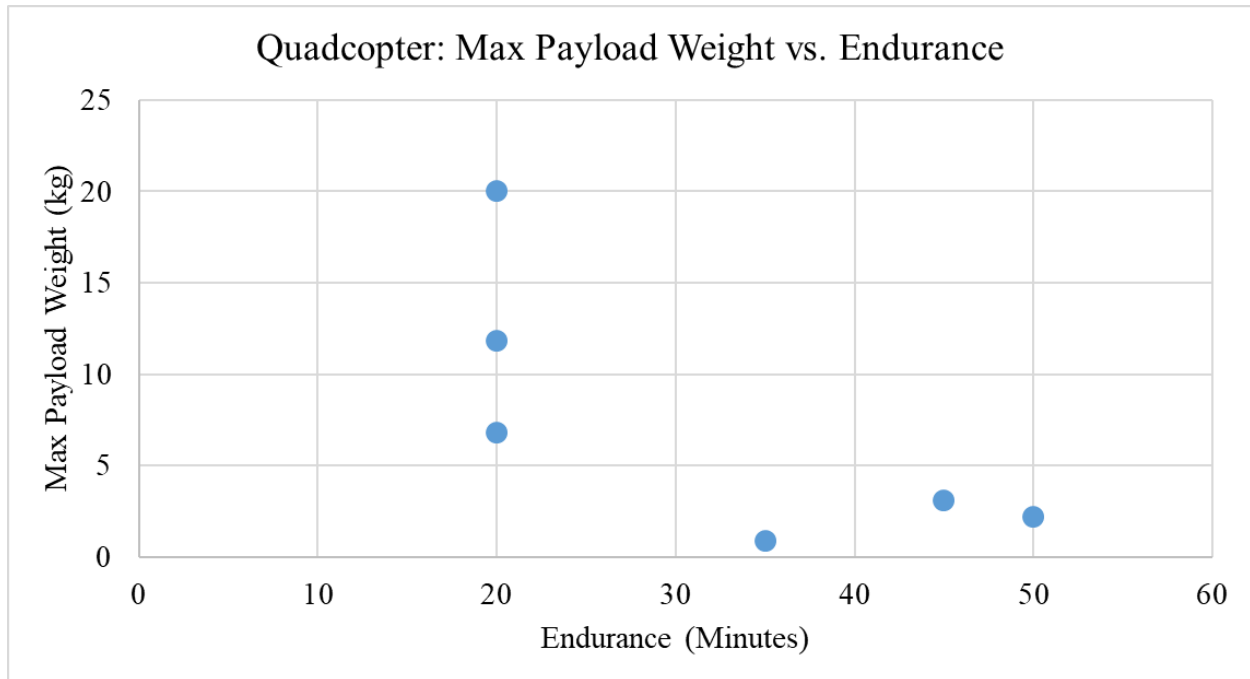


Figure 11: Sample quadcopter values for payload weight and endurance.

Table 2: Sample Quadcopter values.

Manufacturer	UAV Name	Endurance (Minutes)	Payload Weight (kg)
UAV Solutions	Phoenix 30	35	0.9
UAV America	Eagle XF	45	3.1
Lockheed Martin	Indago	50	2.2
Freefly Systems	Alta	20	6.8
AZ Creative	Green Bee 1200	20	20
Versadrones	Heavy Lift Octocopter	20	11.8

2.3 Choosing a Laser Wavelength

There are four primary groupings of methane absorption features: around 1.6 μm , 2.4 μm , 3.4 μm , and 7.5 μm . The absorption cross-sections are strongest in the 3.4 μm region, as shown in Figure 7. Choosing

which wavelength region to target is a fundamental decision with many variables at play. For instance, the commercial company Lasen targets the strongly absorbing $3.4\text{ }\mu\text{m}$ region [Geiger, 1993]. The cross section is two orders of magnitude stronger than the $1.6\text{ }\mu\text{m}$ region, leading to higher ratios in equation 4, which can mean better signal to noise ratio. On the other hand, there are several aspects to consider:

1. The reflectance of solid materials tends to be much lower in the $3\text{ }\mu\text{m}$ region. A chapter on reflectance in this dissertation investigates wavelength-based reflectivity. It is seen in that effort that reflectivity in the $1.6\text{ }\mu\text{m}$ region can often be 50% or greater while reflectance in the $3\text{ }\mu\text{m}$ region can be 5% or less [Baldrige, 2009]. This means nearly all of a transmitted pulse can be absorbed by the ground, leaving not much for analysis. Figure 12 plots the reflectivity of a few common materials as a function of wavelength to demonstrate the effect.
2. The absorption is so strong that, coupled with reflectivity problems, the online signal can be completely lost in a plume of methane. One solution is to use a higher power laser. The other is to fly lower. This latter approach is taken by Lasen as they fly on a helicopter, often not more than a few hundred feet off the ground. For a UAV system constrained to flying lower than 500 feet, this is not a problem. In fact, until UAVs can legally fly higher than 500 feet, the extra absorption is attractive.
3. Laser systems in the $3\text{ }\mu\text{m}$ region often depend on Optical Parametric Oscillator (OPO) technology [Geiger, 1993]. Although it is effective at producing the required wavelengths, it generates waste photons as part of the process, requires multiple stages, and the non-linear crystal used to generate the final wavelength can be susceptible to temperature and vibration effects [Shuman, 2013]. The complexity also introduces initial costs and can create ongoing labor costs to keep the system well-tuned.

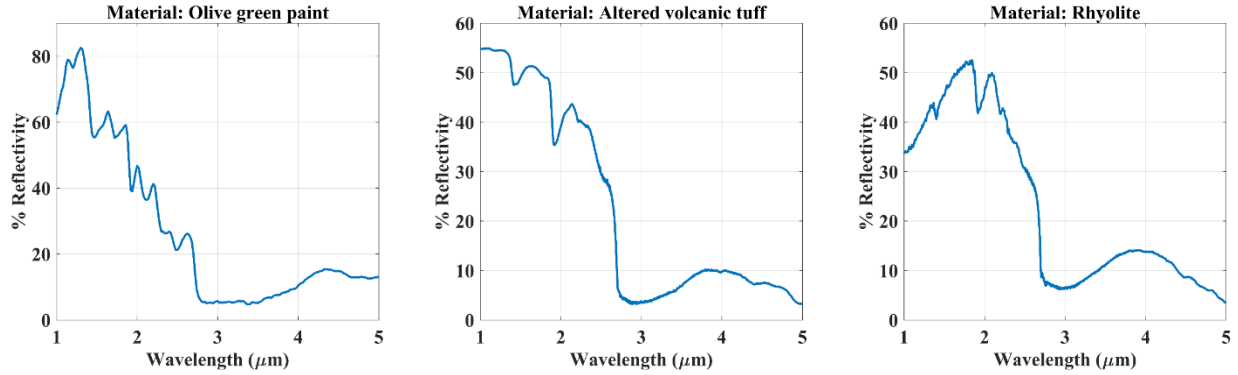


Figure 12: Example reflectivity plots demonstrating that reflectivity tends to be higher at lower wavelengths.

The 2.4 μm region has about the same reflectivity challenge as the 3 μm region. Unfortunately, where greater absorptivity is available in the 3.4 μm region so that a low-flying aircraft can still get a decent signal, the cross-section at 2.4 μm is roughly two orders of magnitude lower. This means a longer path length is required to generate the same online losses. However, long path lengths also mean high intensity losses which, when coupled with low reflectivity, create detection challenges.

Extending the list of issues is that there are not many options for generating single wavelength pulses in this region. One somewhat nearby option is Holmium:YAG lasers which can be found at 2.1 μm from commercial suppliers [Scientific Materials, 2017; Olympus Medical, 2017]. However, there is no clean path to increasing the wavelength by 0.3 μm. Another option is a CrZnSe laser at 2.5 μm, but this misses the methane features and, again, it's not easy to shift down by 100 nm [Sorokin, 2010]. The CrZnSe lasers do provide a path to 2.4 μm via a frequency comb approach where a femto-second pulse creates a broad spectrum of wavelengths that encompasses all of the methane features [Sorokin, 2010]. However, while frequency combs are a promising technology for remote spectroscopy, they are not the focus of this effort. So, even if the poor reflectance and absorption could be overcome, the lack of good laser technology in this region makes application difficult.

The 1.6 μm region has the least absorption cross-section at $\sim 7.5 \times 10^{-23} \text{ m}^2$, making methane detection over small path lengths challenging. On the other hand, as shown in Figure 12, the reflectivity can be much

higher. Additionally, the Er:YAG laser technology is easily able to hit the wavelength region with a simple diode-based, solid state laser pumping a crystal [Young, 2004]. Additionally, the challenge of low absorption cross-section becomes a strength for long path lengths, as evidenced by the effort to put this wavelength in space on the Methane Remote Lidar Mission (MERLIN) platform from Germany [Kiemle, 2011]. This wavelength is also being used by Ball Aerospace for their methane instrument [Bartholomew, 2017].

For this dissertation, the 1.6 μm region has been chosen even with path length challenges because:

1. A diode laser pumping a crystal is about as simple a laser technology as can be achieved. The UAV will need simple, straightforward solutions to minimize the costs of development and maintenance.
2. The FAA has a 500-foot altitude limit, which is low, but the double-path length is 1,000 feet. Analysis in a later section shows that absorption is sufficient to find most leaks and that detection of very small leaks is limited by other factors.
3. Additionally, if the ceiling is expanded in the future then a 1.6 μm instrument will be extendable to higher altitude whereas the 3 μm region would reach saturation limits.
4. The university's partner, Ball Aerospace, has allowed access to their laser for testing theories and some hardware. Although not required for the bulk of the work presented here, the hands on experience and ability to test theories is a valuable asset.

2.4 Pulsed vs. Continuous Wave

Laser systems can use continuous waves (CW) or emit bursts of energy in a pulse. Briefly, a continuous wave system creates an online laser modulated at one frequency and an offline beam modulated at a slightly higher or lower frequency [Pruitt, 2003]. The combination of these waves creates a modulated beam that interacts with an absorption feature. The returned signal can be detected with a lock-in

hardware scheme. The pulsed version sends energy in time scales of around 50 ns, although an exact value depends on the laser design. The pulses are measured with simple, direct detection hardware.

According to a 2012 paper by Xiaoli Sun, there are three significant benefits to choosing pulsed over continuous wave lasers [Sun, 2012]:

1. The hardware to detect pulses is more simple than the hardware for deconvolving the combined continuous waveforms.
2. The SNR of a pulsed system is roughly 4x higher than that of a continuous wave system for equal powers.
3. For equal SNR, a pulsed system requires about $1/6^{\text{th}}$ of the power of a continuous wave system.

Additionally, even optimistic CW instrument developers tend to report observation rates on the order of 100 Hz at best or even one measurement per 10 seconds [Pruitt, 2003; Dobler, 2013]. Extensive research by Exelis, Inc. and NASA Langley has also shown that backscatter from aerosols and clouds can create waveform interference. Solving the problem resulted in a patent for “tone hopping” that subtly changes the modulation frequencies [Dobler, 2013].

This is not to say that pulsed systems are without challenges. Whereas CW systems, such as the Exelis laser, are completely fiber-based, pulsed systems tend to rely (for now) on free-space optics for pumping. Proponents of CW systems argue that this can mean the transmitter is more complicated with alignment and damage thresholds [Dobler, 2013].

On the balance, though, pulsed systems seem more suited for UAV concepts. Pulsed systems may require more careful alignment and wavelength control, but the higher SNR for equal power, the more direct signal detection, and the faster measurement rates guide their selection here.

2.5 Laser Technology

As discussed briefly in section 1.6.2, Er:YAG lasers are the technology of choice for providing 1.645 μm wavelengths. White's brief 1972 and 1975 papers first investigated the technology and numerous researchers have expanded on it. Setzler's 2005 overview paper on the state of current technology provided an excellent overview of the concepts and foreshadowed coming advancements [Setzler, 2005]. Since his publishing, dozens of papers have been published that incrementally advance the technology. In broad strokes, two technologies have received the most attention:

1. Diode pumping lasers: Er:YAG crystals can be pumped at $\sim 975\text{ nm}$ $\sim 1,470\text{ nm}$ or 1532 nm (or combinations of them) to emit $1,645\text{ nm}$ wavelengths. The development of solid state diode lasers capable of resonant pumping at these wavelengths opened the doors to compact, stable systems [Setzler, 2005; Kim, 2009; Fritsche, 2014].
2. Volume Bragg Grating (VBG) filters: VBG transmissive or reflective glass filters clean up the sometimes-noisy spectrums of diode lasers. Using VBGs before, in, or after laser cavities can trim FWHM spectrums to 0.1 nm or less [Lumeau, 2010; Tang, 2014; Fritsche, 2014].

Although the budget did not allow for purchasing a laser system, research into the state of the art showed that the system would likely be a diode pump laser at $1,532\text{ nm}$ with an Er:YAG crystal doped at about 2% and a VBG filter somewhere in the line. Wall efficiencies for such a system were found to be between 20% and 40% [Young, 2003; Fritsche, 2014].

As stated previously, this dissertation benefited from having access to Ball Aerospace's methane system. Their lasers are also built on this same technology concept [Bartholomew, 2017]. It is worth saying that significant effort was made to differentiate from their efforts in order to explore new ground. However, in the case of laser technology, compact, high-power, narrow wavelength pulses are hard to improve on.

Chapter 3: Gaussian Plume Simulations

This chapter considers the target of interest: plumes of methane. Modeling and characterizing plumes was a major thrust of the research. The Gaussian plume model is presented here along with modifications for utility in DIAL simulation. With the help of data from Ball Aerospace, it is also possible to demonstrate the quality of the modeling.

3.1 The Gaussian Plume Model

It is necessary to have a plume model to create the spatial variations in density to test the instrument against. A common model with decades of history is the Gaussian plume model [Roberts, 1923, Sutton, 1947]. The underpinning equation in this model was first sketched by Roberts in 1923, expanded by Sutton in 1947, and subsequently extended by generations of researchers [Stokic, 2011]. At its heart is the assumption that a plume disperses laterally and vertically with a Gaussian distribution influenced by three factors: leak rate, wind speed, and stability coefficients. The stability coefficients are decimal values that approximate the influence of the thermal stability of the atmosphere and are provided in Table 3.

$$C(x, y, z) = \frac{Q}{U} \frac{1}{2\pi I_y I_z x^2} e^{\frac{-y^2}{2(I_y x)^2}} \left(e^{\frac{-(z-H)^2}{2(I_z x)^2}} + e^{\frac{-(z+H)^2}{2(I_z x)^2}} \right) \text{Eqn. 5}$$

x, y, z are distances in meters from the source

I_y , I_z are measures of atmospheric stability

H is the height of the source above the ground, in meters

Q is the leak rate in kg/s

U is the wind speed in m/s

Table 3: I_y and I_z are the lateral and vertical intensity coefficients. They are chosen based on empirical values and are simplified.

Thermal Stratification	Lateral Intensity (I_y)	Vertical Intensity (I_z)
Extremely unstable	0.40 - 0.55	0.15 - 0.55
Moderately unstable	0.25 - 0.40	0.10 - 0.15
Near Neutral	0.10 - 0.25	0.05 - 0.08
Moderately stable	0.08 - 0.25	0.03 - 0.07
Extremely stable	0.03 - 0.25	0 - 0.03

In this form, the concentration of a gas in kg/m^3 can be found for any coordinate within a plume. For instance, Figure 13 below is created via the following parameters:

1. A grid of sample points is distributed about a 20-meter wide sample by 20-meter long sample area
2. A query height of 0.5 meters is set
3. A leak rate of 1 kg/s is used with a 1 m/s wind speed
4. I_y and I_z are set to 0.15 and 0.13 respectively to model a near neutral lateral intensity and a moderately unstable vertical intensity, as might be expected on a warm, relatively calm day.

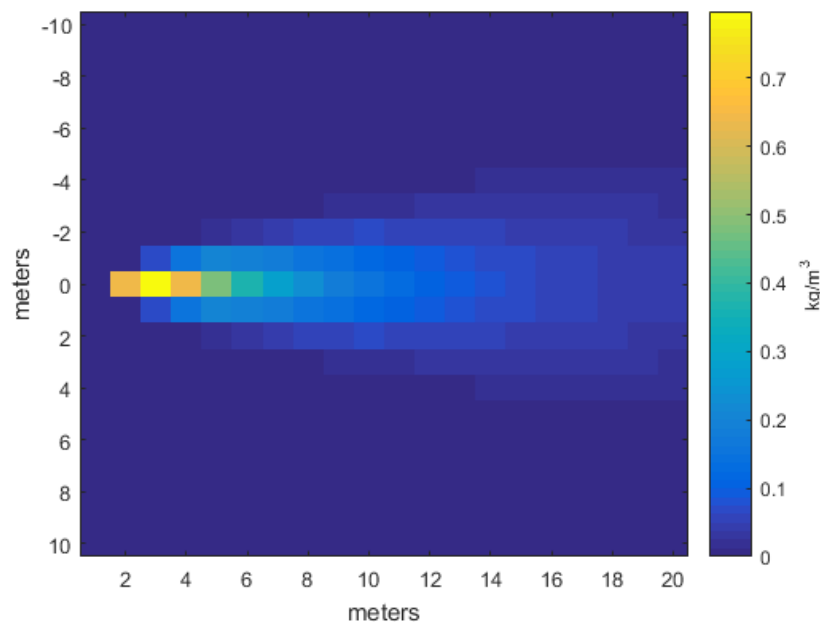


Figure 13: A simple example of a Gaussian plume model predicting the mass density of a gas at a given height and location.

This type of model is useful for in situ sensors, such as what a gas line worker might hold or what might be mounted on a truck. In fact, many researchers using in situ sensors rely on this form of the model to predict plume source and leak rate parameters [Rao, 1991, Thoma, 2012, von Fisher, 2017]. This is true even though terms like buoyancy are not considered. The reason for its popularity is that it is relatively

straightforward to forward solve or invert this plume model. Additionally, researchers make the argument that while gases, such as methane, are more buoyant (.016 kg/mol for methane vs. 0.029 kg/mol for air), in a well-mixed atmosphere and over short distances, the buoyancy is a small factor.

3.2 Extending the Gaussian plume model

The Gaussian model was built for mass density predictions. It can be extended to a total column measurement by integrating the equations in a volume defined by the laser diameter and path length. With parameters such as the local air density, it can return the concentration path length in parts per million per meter (ppmm) or a path averaged parts per million (ppm). For plume simulation, these values are above the background levels. The result of the integration and ratio math is a model useful for trading instrument design parameters such as the beam divergence angle, path length, and spatial sampling resolution requirements.

One more way to tweak the model is to convert the leak rate from kilograms per second to “Standard Cubic Feet per Hour” (SCFH). This unit of measurement is what is used in the gas transport industry and represents the number of cubic feet of gas at standard temperature and pressure (STP) released per hour. A value of 100 SCFH means that 100 cubic feet of methane at 1 atmosphere of pressure is released over one hour. To convert the Gaussian plume model, a factor is derived by assuming a methane density at STP of 0.668 kg/m³. Then, the conversion looks like:

$$1 \frac{kg}{s} = 1 \frac{kg}{s} \times \frac{1 m^3}{0.668 kg} \times \frac{1 ft^3}{0.028317 m^3} \times \frac{3600s}{1 hr} = 190,320 SCFH \quad \text{Eqn. 6}$$

So it is seen that the 1kg/s leak rate used in Figure 13 above is an extreme amount of methane! To wit, a 1,000 SCFH leak rate is considered by the industry to be significant to the point that it warrants immediate attention.

The final modification is to take into account how the beam overlaps with the plume. If a beam falls directly on the source of a leak, half will be in the downwind plume and half will be upwind and measure

no additional methane. The integration of this scenario will result in lower methane than might be expected. The peak methane is actually measured slightly downwind when the maximum overlap of the laser and the plume is reached. From that point, the measured methane will decrease exponentially with the Gaussian plume equation. Figure 14 illustrates the concept of beam locations and plume overlap. The solution is to only integrate the part of the beam in the plume and setting the rest to zero.

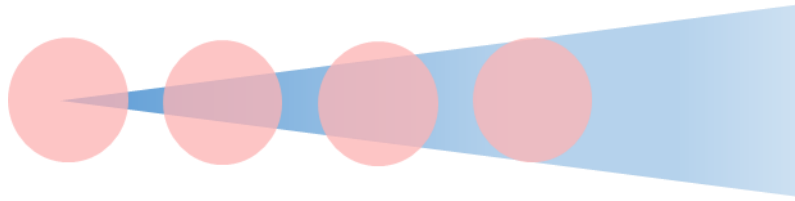


Figure 14: The beam from a DIAL instrument (red circles) integrates the total volume and so overlap with a gas plume (blue triangle) is driven by the percentage of beam and plume overlap.

The code for this approach was written in MATLAB and relied on a parameterized *integral3* function. As an example of the output, Table 4 provides two sets of parameters and Figure 15 shows the two plumes in units of ppm. The relatively low measurement at the leak source followed by a sharp peak is seen, illustrating a central expectation of real-world plumes measured with an airborne DIAL system.

Table 4: Parameters chosen to create Figure 15

Parameter	Configuration 1	Configuration 2
Leak Rate (SCFH) (kg/s)	100 (5.25E-4)	1,000 (5.25E-3)
Wind Speed (m/s)	1	2
I _y	0.10	0.40
I _z	0.05	0.15
Beam Diameter (m)	1	1
Laser Flight Path Length (m)	100	1000
Spatial Sampling Resolution (m)	0.5	0.5
Gas Release Height (m)	0	0

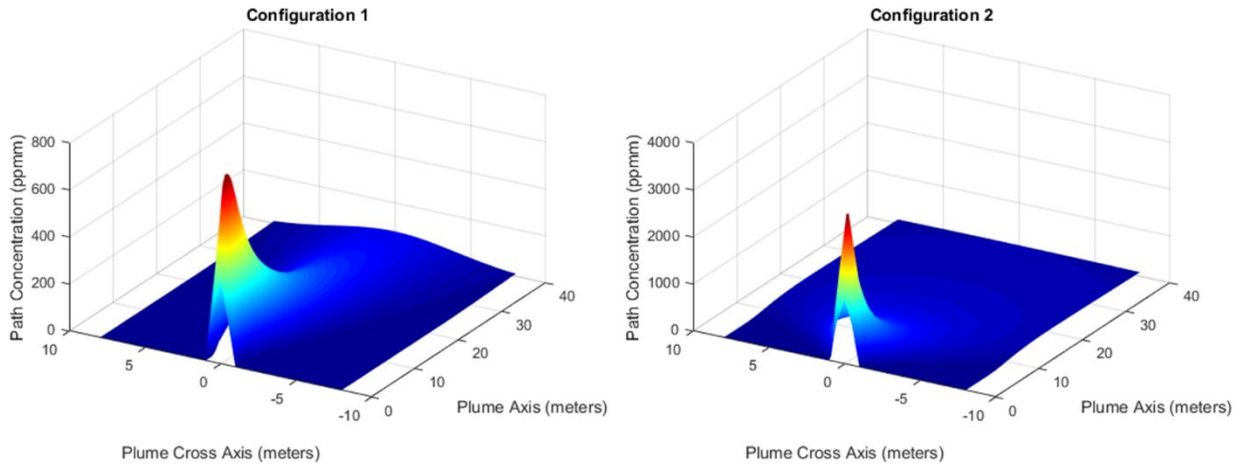


Figure 15: Plotted examples of the plumes parameterized in Table 4.

3.3 Model Correlation with Field Measurement

Empirical analysis of data helped build confidence in the integrated Gaussian plume model. The data for the analysis came from Ball Aerospace & Technologies' previously discussed methane instrument. The latest version of Ball's instrument samples at 10,000 Hz, giving point density sufficient to demonstrate model correlation. An example plume imaged in field operations in the gas fields of northeast Colorado is shown below in Figure 16. In this example, the blue areas are background measurements and the values range from 0-1,000 ppmm, although the reddest parts of the plume may be above 1,000 ppmm.

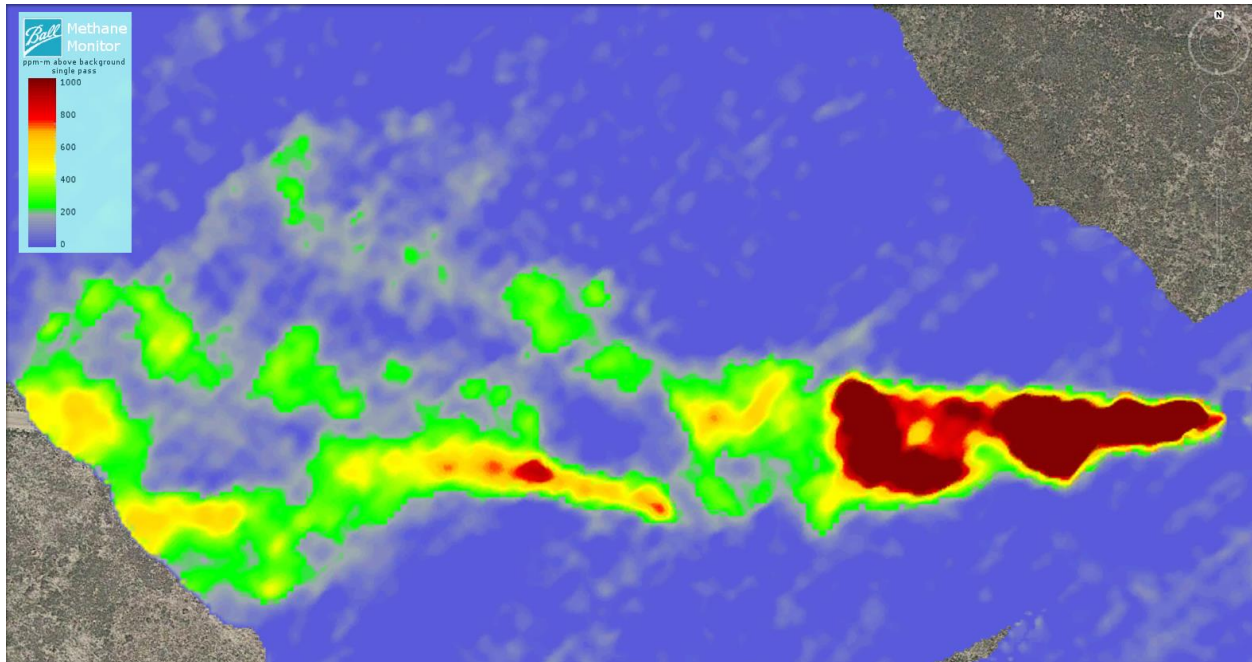


Figure 16: Example plume of fugitive methane captured with Ball's methane detecting instrument on May 16th, 2017.

A couple of things are observed with simple visual analysis of the plumes. First, the plume is not smoothly distributed along its axis. In other words, it appears lumpy and the model does not cover that behavior. The Gaussian model is geared towards time-averaged solutions rather than instantaneous snapshots, though. However, a second observation is that the boundaries of the plume roughly follow the shape of a triangle, which is what the Gaussian model predicts.

Numerical analysis of the data followed a series of steps:

1. **Isolate just the plume:** Calculate the Eigenvector of the plume by isolating the points in the data above the mean value of the points in the area. Some noise was carried over into the reduced set. This was addressed by using a weighted matrix in the Eigenvector calculation where the weights were the concentration values. In this way, points near the noise floor contributed less than points well above the floor.

2. **Orient the plume:** Rotate and translate the plume to align with the primary Eigenvector as the “X-Axis”. Although not necessary, it makes for easier visual confirmation of results and more consistent analysis.
3. **Convert coordinates:** Convert the latitude and longitude of the georectified data to X and Y coordinates in meters.
4. **Bin data:** Slice the plume up into 5 meter bins along the X-axis to look at local effects.
5. **Curve fit:** Look at the behavior of points within each bin. Find the peak value by fitting a Gaussian profile to the points.
6. **Plot results:** Create a plot with the peak values plotted at the center of each bin. Plot Gaussian fit statistics. Plot curves against data.

These steps are illustrated in Figure 17. There, a simulated plume has been aligned with the X-Axis. The colored bins in the second plot are in 5 meter lengths and each bin has a Gaussian curve fit to it as seen in the third plot. The peaks of the curve fit are used in the final step.

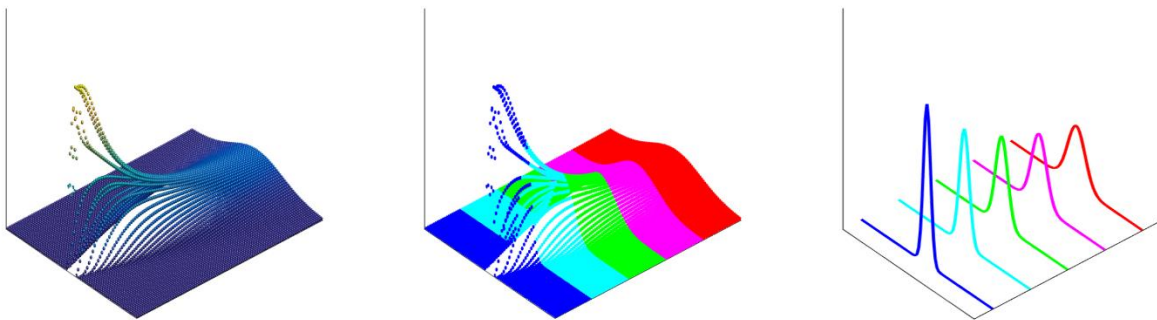


Figure 17: Graphical illustration of the plume binning described in the text

The above process was implemented in MATLAB for plumes provided by Ball for this exercise. The first is the one in Figure 16 and has an unknown, but very large, release rate. The second and third are from plume release rates of about 800 and 400 SCFH respectively from Ball testing.

Three example cross-section plots of bins of real data are shown in Figure 18. It is seen that the Gaussian fit works perhaps even better than one might expect from real-world data.

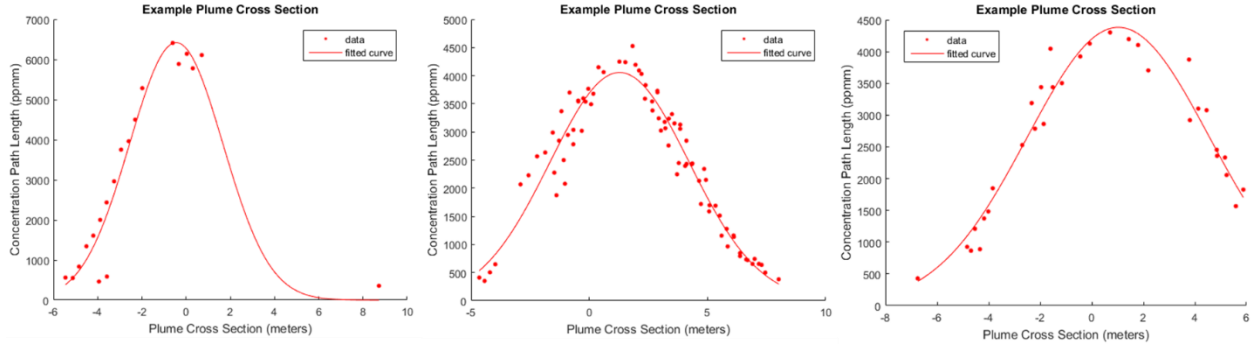


Figure 18: Example cross-sections of plumes measured with Ball's methane detecting instrument.

The results of Figure 18 look good graphically. Numerically, the answer is more nuanced. Figure 19 shows three plots: the first is the R-Squared value for every bin along the length of the plume in Figure 16. It is seen that there are many points above 0.75, indicating that the fit works well. However, there are also bins with low R-Squared values. Two such points are plotted in the second and third plots of Figure 19. It is clear that the fits miss most of the points. However, although the fit is “bad”, the intent of the Gaussian distribution is met in the region of increasing concentration.

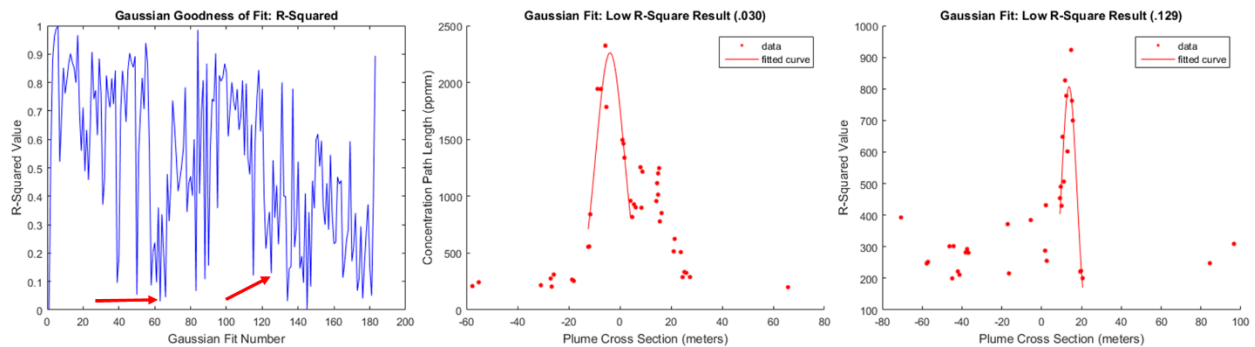


Figure 19: Goodness of Fit “R-Squared” plot for the plume in Figure 15. The red arrows in the first image point out the low R-Squared value example plots in the next two images.

Further analysis shows that the quality of the fit trends downwards along the length of the plume. In other words, the Gaussian assumption works best near the source of the leak. Table 5 provides the mean and standard deviations of the R-Squared values for the three plumes as the overall result, and then the leading and trailing portions of the plume. The leading edge is defined as starting from the source of the plume, up to the peak, and then down to the change in the knee-of-the-curve, as seen in Figures 15 and 17.

Table 5: Goodness of Fit R-Squared statistics for three plumes with Gaussian fits.

Plume	Release Rate	Overall		Leading		Trailing	
	SCFH	Mean	STD	Mean	STD	Mean	STD
1	Unknown	0.524	0.27	0.807	0.15	0.493	0.264
2	~800	0.442	0.31	0.862	0.16	0.331	0.237
3	~400	0.580	0.36	0.733	0.33	0.537	0.360

From investigating the low R-Squared values, it seems that some data cleaning may improve the result. For instance, the noisy floor values around the peaks could be cleaned through simple thresholding or tossing values with significant peak-to-valley differences. However, this is not pursued further because the intent of demonstrating that a Gaussian model can simulate plume contours is sufficiently met, at least in the cross-axis direction near the leak source. Attention next turns to the along track axis where the lumpiness of the plume may have a stronger influence.

The first expected behavior is the unique profile of rapidly rising methane concentration followed by an exponential decrease. Refining expectations, shot-noise was introduced into the simulation as a Poisson distribution. Shot-noise is the noise-per-measurement that occurs due to myriad influences such as electrical noise. The result is seen in the second image of Figure 20. The overall shape is retained while local values may be noticeably higher or lower than the noise-free example.

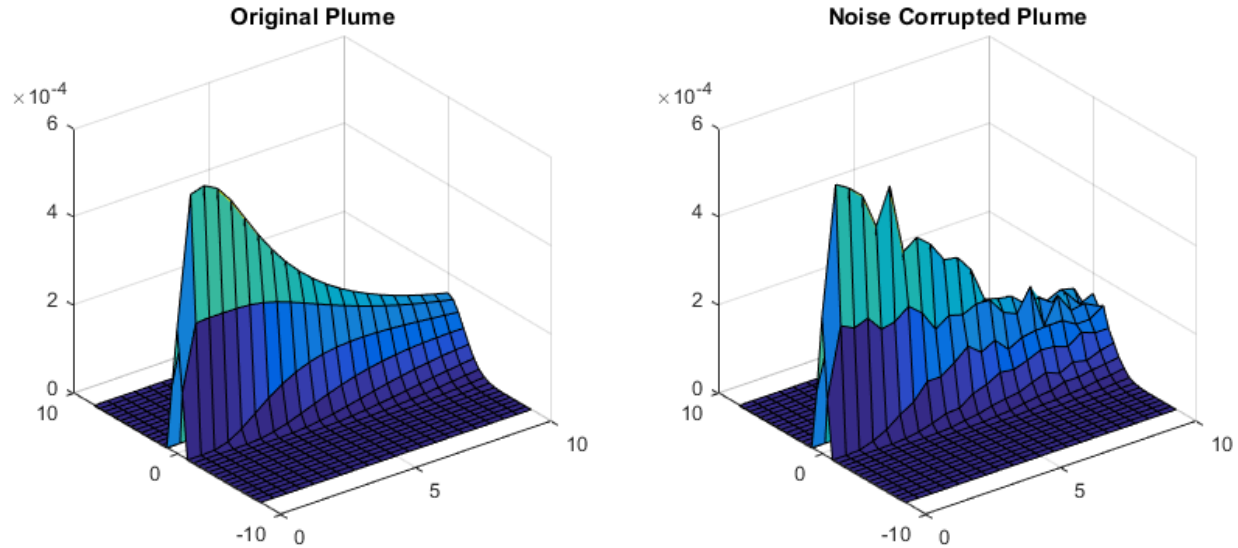


Figure 20: Example simulated plume interacting with a 1 meter diameter laser beam. The image on the left is the perfect result while the image on the right assumes Poisson shot noise.

Comparing this with the real world plumes, two plots showing the peak values vs. the bin centers are shown in Figure 21. The y-axis values represent the path concentration above the background concentration. It is seen that the plots share similarities with the expectations of 1) a sharp rise to the peak concentration and 2) a noisy, but discernible non-linear curve back to the floor. An example of how a simulated result compares is provided in Figure 22, which uses the same flight parameters as those used to capture the methane plume.

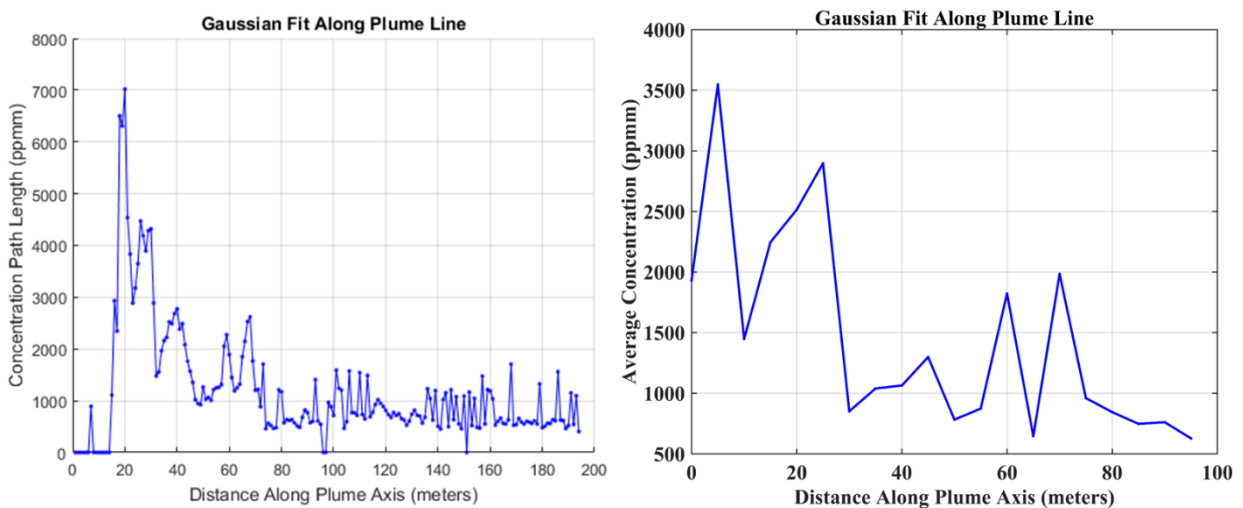


Figure 21 Along-axis plots with points representing Gaussian fit peaks. The first plots' source is the plume in Figure 16, the second is from a ~800 SCFH plume.

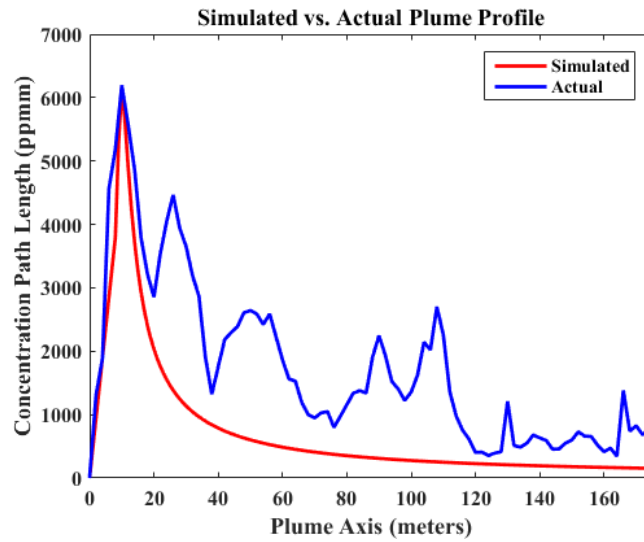


Figure 22: Actual plume profiles vs. simulated plume with matched peaks.

The overall conclusion is that Gaussian plume modeling works best near the leak source when considering snapshot views of a plume, before local effects cause “lumps” to form in the plume. It may be that further downstream, as things continue to disperse, that the plume again becomes smoothly distributed, but this cannot be confirmed with this approach. It may also be true that observing the plume over time would produce a more smoothed Gaussian-like result. The final observation is that this approach becomes noisier/less characteristic as plume release rates decrease as the number of applicable measurement points are reduced. It is believed, though, that the model is sufficient for testing instrument models against.

Chapter 4: Analysis Trade Studies

With a reliable model of plume characteristics established, it is possible to investigate engineering concepts that interact with the plumes. For instance, determining the maximum allowable beam diameter, or how densely measurements must be made. Section 4.1 introduces the signal to noise methodology used in this dissertation and investigates a minimum value. Section 4.2 considers beam diameter and shows that it strongly limits the detectability of plumes, regardless of an instrument's SNR value. Key plots are provided, allowing for estimating leak rate performance as a function of the noise floor and the beam diameter. Section 4.3 continues the discussion of the interplay of SNR and plume detectability by considering how the visible area of a plume changes with SNR. This leads to Section 4.4 and a discussion of the required scan density for a given plume area. Section 4.5 then investigates how scan patterns can be tailored to achieve the necessary scan density. The discussion changes gears slightly for a side-topic in section 4.6 related to how beam divergence interacting with a surface returns smaller effective beam diameters, even with a Lambertian surface assumption. Section 4.7 ends the chapter with a simulation of laser speckle and how it can interact with variations in surface reflectivity to introduce meaningful noise in the form of false positives and negatives.

4.1 Signal to Noise Ratio Calculation

Hardware planning started with error analysis of the final result as a function of traditional system parameters such as signal to noise ratio and dark current noise. The effort closely followed the excellent work of Tamer Refaat et. al. in their 2013 paper evaluating performance of ground, aircraft, and UAV methane DIAL systems operating in the region of $1.6 \mu\text{m}$ [Refaat, 2013]. There, the total error percent in the final answer is expressed as:

$$\varepsilon = \frac{\varepsilon_R}{\sqrt{S}} + \sqrt{\varepsilon_A^2 + \varepsilon_T^2} \quad \text{Eqn. 7}$$

, where ϵ_R is the error due to SNR impacting parameters such as optical efficiency, s is the number of samples being analyzed at once, ϵ_A is the error due to atmospheric interactions, and ϵ_T is the error related to laser variations. In their paper, they showed that the first term of the equation dominated the second term, especially if the systems were flown over vegetation rather than the ocean. The contributions that did come from the second term were primarily due to temperature-based variations in optical depth for the online and offline wavelength. For them, their UAV was expected to fly 8 km above the ground where the temperature variations between the ground and flight altitude were meaningful. Similarly, pressure and relative humidity are functions of altitude. However, for the low flying UAV discussed in this dissertation, these environmental parameters are considered more or less constant, further minimizing the impact of that error term.

Next, the laser error term looks at the impact of wavelength jitter and how it affects optical depth. For a ± 10 MHz jitter and 20 MHz line half-width on the online wavelength, the rolled-up error contribution was less than an order of magnitude of the first term's value. (The offline wavelength is even less sensitive to jitter and was two to three orders of magnitude less than the online value's contribution.) Based on the above, the first term is considered the most important for this application and will be used as the guiding parameter.

Expanding on ϵ_R , Refaat assumes that the speckle noise is sufficiently small enough to ignore, resulting in equation 8. (Note that the speckle assumption is visited in its own upcoming section.)

$$\epsilon_R \cong \frac{\sqrt{SNR_{on}^{-2} + SNR_{off}^{-2}}}{dOD_{mn}} \quad \text{Eqn. 8}$$

Here, the SNR must be calculated for both the online and offline. The denominator dOD_{mn} is the differential optical depth of the two wavelengths, taking care to double the range to account for the

distance to the ground and then back to the instrument. The primary parameter driving the differential optical depth is the differential cross-section for methane at the online and offline wavelengths.

Starting the analysis with the differential optical depth, the cross-section of methane in the 1,645 nm region was computed using the HITRAN database. A temperature of 296 K and 1 atmosphere of pressure of 101 kPa were used to establish the feature shapes. The cross-sections at the offline and online wavelengths of 1645.37 nm and 1645.55 nm were determined from the output using MATLAB's P-CHIP interpolation method. A dry air number density of 2.48×10^{25} was used following the Ideal Gas Law. A one-way path length of 300 feet was used as a baseline for a typical flight altitude of the UAV. Flying higher would increase the amount of methane the signal travels through, which would increase the differential optical depth, which would lower the error. Next, due to these low altitudes, variables that can be functions of range (such as pressure and temperature) were treated as constants.

To convert path-averaged ppm values to SCFH leak rates, the integrating Gaussian plume model was used. A beam diameter of 2 meters, a flight altitude of 300 feet, and standard temperature and pressure were used. A wind speed of 1 m/s was assumed. The peak path-averaged ppm value from a series of plumes was collected from the runs and used to make the conversion, again utilizing MATLAB's P-CHIP interpolation method to reduce the number of computationally expensive runs.

The trade study run on this setup involved varying the SNR value. For each run, the same value was used for both the offline and online wavelengths as a first order approximation. The result of this effort is seen below in Figure 23.

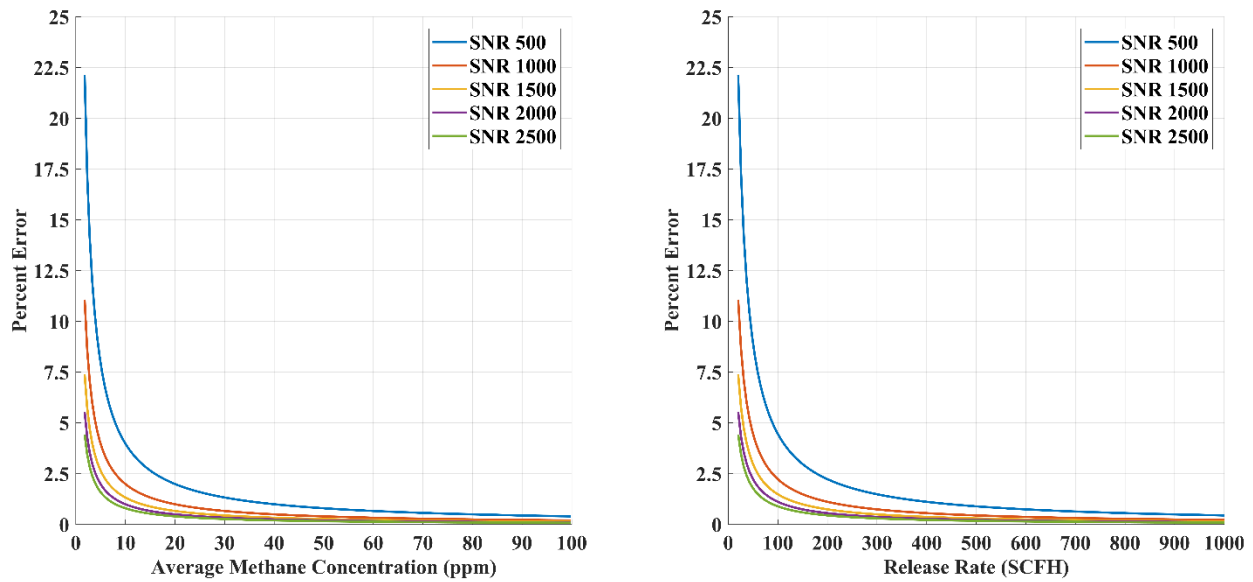


Figure 23: Plots demonstrating the relationship between the system SNR, the quantity of methane in the environment, and the ability to correctly resolve the methane. See the text for how the average methane concentration was converted to SCFH.

Perhaps the most noticeable result is that, for these parameters, it can be difficult to correctly quantify low leak rates as errors are greater than 20% for a SNR of 500. On the other hand, a 20% error on a 20 SCFH plume means reporting a leak rate uncertainty of 4 or 6 SCFH while, as discussed previously, industry representatives are more interested in order of magnitude estimates than highly accurate measurements. As another example, a SCFH of 100 means an error of about 4.3% which would mean reporting 100 ± 4 SCFH, a negligible difference.

In fact, the relationship between SCFH and the absolute error is, to a first order, constant for this configuration. No matter the leak rate, the absolute error for a SNR of 500 is about 4.5 SCFH. A SNR of 2,500 means an absolute error of about 1 SCFH regardless of leak rate. This mostly constant relationship exists because increasing the leak rate increases the optical depth at the same rate. Further, all of the equations driving the differential optical depth are either linear or have been linearized by assuming no range induced changes. It should be noted that many parameters make up the end constant such as beam footprint, flight altitude, temperature, and pressure. The latter two, in particular, drive a number of things

such as the absorption cross-section of methane and the number density of the atmosphere. So there is no one term, simple formula to estimate the constant error term.

The conclusion from this study is that, for this error term, a general target to aim for is a SNR above 500. Getting to 1,000 would halve the error and is desirable. However, it's not critical as a ~5 SCFH error is unimportant, even at the smallest leak rates.

Of course, there are other errors and effects to consider. Speckle and reflectivity will be discussed in an upcoming section. The next discussion, though, discusses beam diameters and how instrument sensitivity is strongly driven by footprint diameter.

4.2 Choosing a Beam Diameter

The beam diameter drives the sensitivity of the instrument because it relates to how much of the air around a leak source is included in the measurement, which can decrease the average methane within the volume. Figure 24 plots a plume (light blue) with two beam diameters sampling it (light red). It demonstrates that large beams are more likely to sample the plume, but the path averaged concentration drop as more of the beam is outside of the plume.

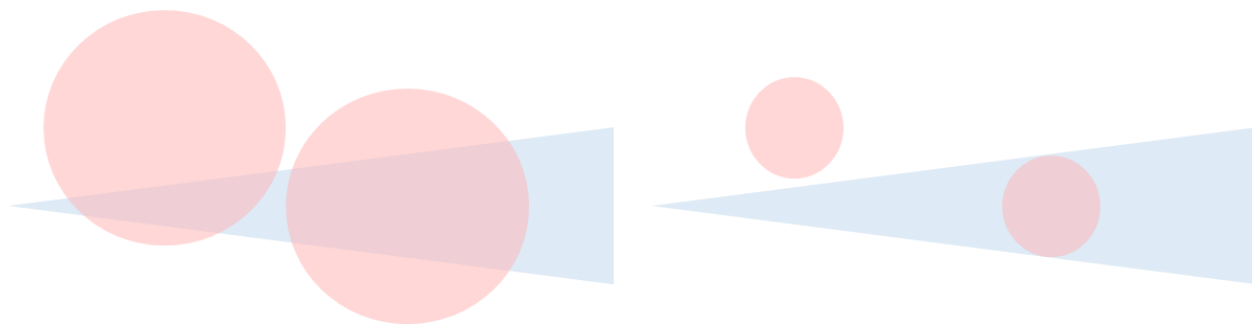


Figure 24: A larger diameter beam on a plume measures more of the surrounding air than a smaller beam.

The beam diameter of the laser system is driven by both mechanical and operation considerations. Mechanical drivers include the fiber diameter, a speckle reducing diffuser for multi-mode fiber systems, and collimating optics. Operation factors include flight altitude and the beam pointing angles that

contribute to the laser's flight path. Flight path drives the final expansion of the beam through the beam divergence angle. For example, if a 10 cm beam with a divergence angle leaving the instrument is 0.05 degrees and the distance to the target is 500 meters, the footprint will be $0.1 + 2 \cdot 500 \cdot \tan(0.05) = 0.97$ meters, while at 1,000 meters the beam diameter is 1.85 meters.

The plume simulation described in chapter 3 can help determine the beam diameter appropriate for the instrument's operation goals. The result of varying beam diameter while holding everything else constant is shown in Figure 25 and Figure 26. The constants are a 1,000-meter altitude, a gridded spatial resolution of 0.5 meters, and a leak rate of 300 SCFH. Figure 25 scales the z-axis to the measured values to better show how the shape of the measured plume changes. Figure 26 sets the scale of all plots to that of the 0.5-meter diameter beam to illustrate how much of a difference the beam diameter can make in magnitude.

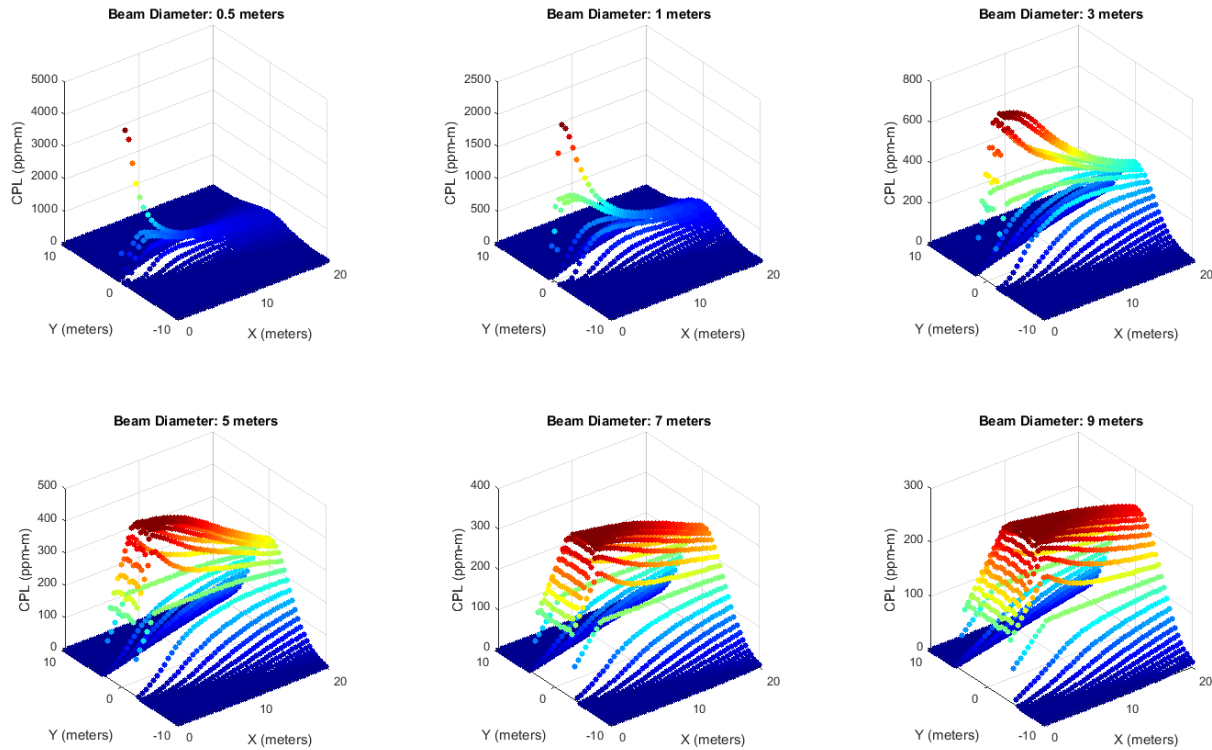


Figure 25: 300 SCFH plume as measured by different diameter beams. Scales relative to measurement.

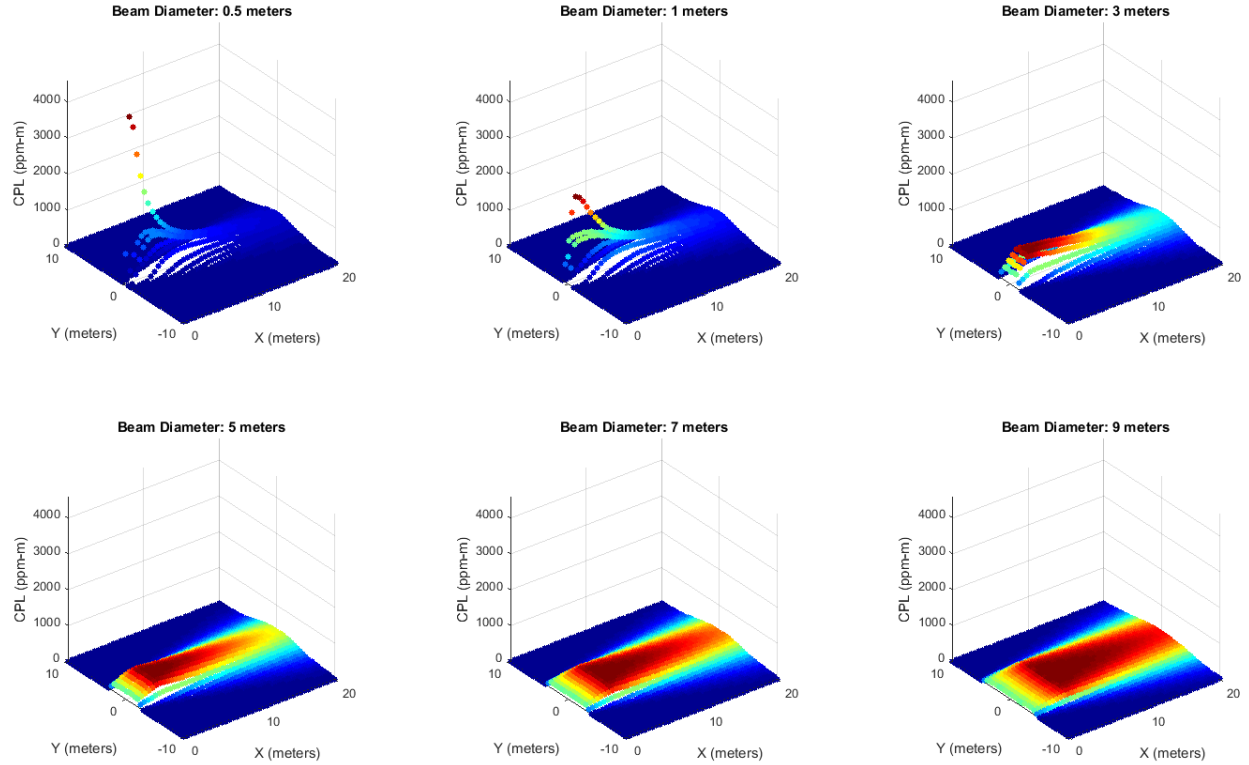


Figure 26: 300 SCFH plume as measured by different diameter beams. Scales relative to the 0.5-meter beam.

Although the reduction in measured CPL is evident as beam diameter increases, it may not be a problem. In fact, if the final instrument can differentiate small differences in CPL, then larger beam diameters may be preferable as it makes apparent plume sizes larger and thus easier to distinguish. On the other hand, if the instrument has a noise floor due to dark current, bit resolution of the analog to digital converter, or other myriad reasons then there may be a minimum CPL value above its noise floor that needs to be measured.

An example of how this can affect the final product is shown in Figure 27. There, the flow rate was set to 50 SCFH with a noise floor of 50 ppmm. This has the effect of making any simulated measurements less than the noise floor go to zero. It is seen that the smaller beam size returns higher value measurements in most measurements while the largest beam diameter of 9-meters returns no data. It is also seen that,

although the values are not much higher than the noise floor, the 3, 5, and 7-meter cases do have a wider profile. This may be useful or it may be too close to the threshold to rely on in day-to-day operations.

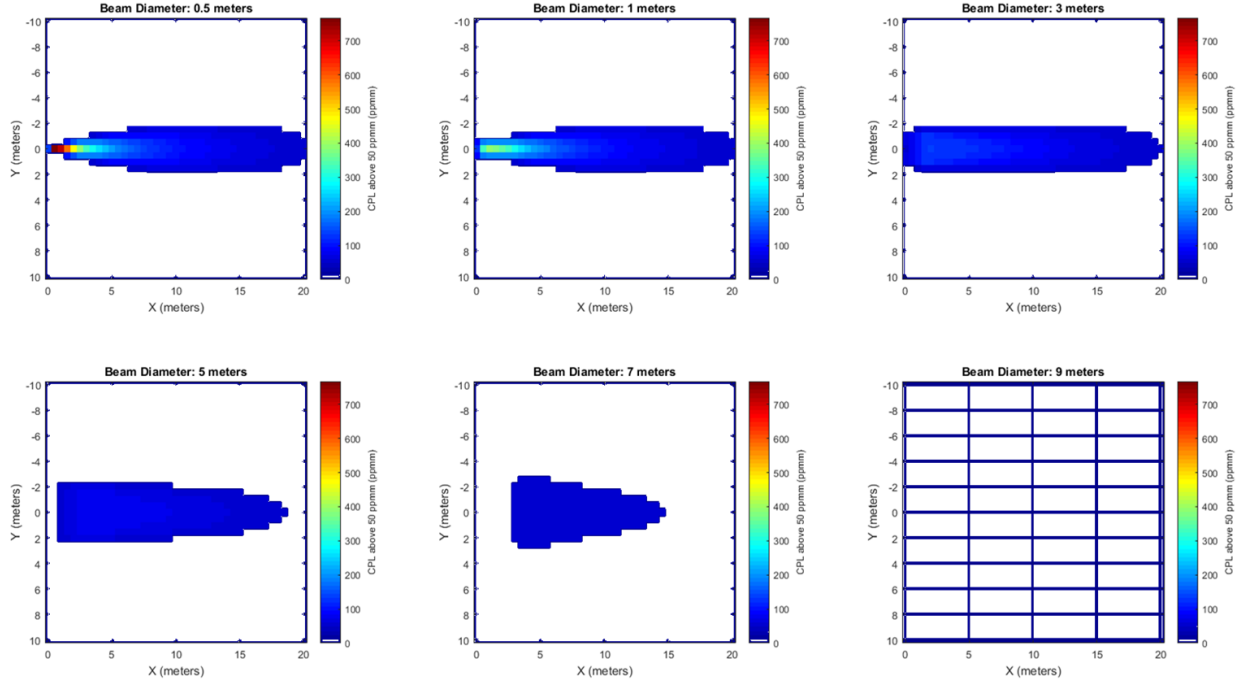


Figure 27: Simulation of beam diameters on a 50 SCFH flow simulation with 50 ppm noise floor. The color scale is fixed to that of the 0.5-meter case.

Extending Figure 27 to a wide range of plume and beam diameters produces Figure 28. This plot assumes the beam falls in the sub-optimal flank of each plume near its inflection point as it is unlikely the instrument will directly measure the peak. Because of the path-averaged nature of the plot, it is altitude independent. A log plot on the y-axis is used to better illustrate the relationships over a wide range of plume leak sizes. One key point is that the best performance is strongly non-linear for beam diameter values near zero and then roughly linearizes on the log-scale after about 5 meters.

An example use of the plot is to choose a beam diameter of 1 meter, and then a target leak rate of 50 SCFH. This gives a y-axis value of 0.3 ppm. This means that the maximum measured concentration above background for this plume would be 0.3 ppm. If the SNR and other measurement errors in the instrument create a best-case sensitivity of 0.2 ppm, then the plume would be detectable. If the errors set a

floor of 0.4 ppm, then the 50 SCFH plume would not be visible from 1,000 meters with a 1-meter diameter beam.

Another example is to choose an error level and determine the maximum beam diameter. For instance, if the background levels are 1.8 ppm and the instrument has a 10% error, then the noise floor is 0.18 ppm. If the goal is to detect a 50 SCFH plume, then the maximum allowable beam diameter is about 2 m.

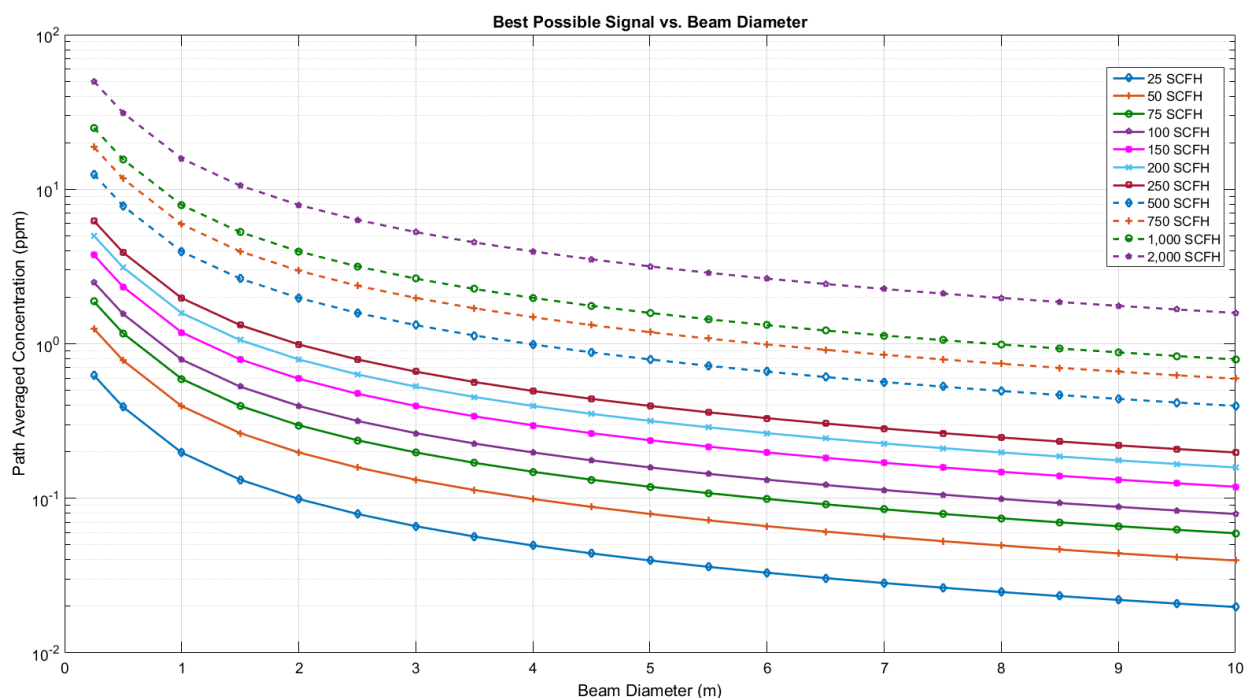


Figure 28: A variety of curves representing the plume sizes visible to an instrument with beams of different diameters. For this plot, all other variables are held fixed.

The data that enabled Figure 28 can be further mined to predict the Receiver Operator Curve (ROC) for a given configuration. An example is shown in Figure 29. There, a beam diameter of 0.5 meters and 3 meters are chosen, an altitude of 1,000 meters set, and a noise floor of 200 ppmm is used. Then the best possible performance of the instrument is estimated for a range of wind speeds. This is made possible because the coefficient in the front of the Gaussian plume equation has a $1/\text{Wind Speed}$ relationship with the concentration. The noise floor in ppmm can be used in Figure 28 by first converting to path-averaged ppm by dividing by the altitude and then drawing a line across the curves at that level. For instance, 200

ppmm / 1,000 m = 0.2 ppm. Looking at this value on the y-axis for a 25 SCFH plume, only beam diameters smaller than a meter will be able to detect the plume as all other points on the curve are below the noise floor.

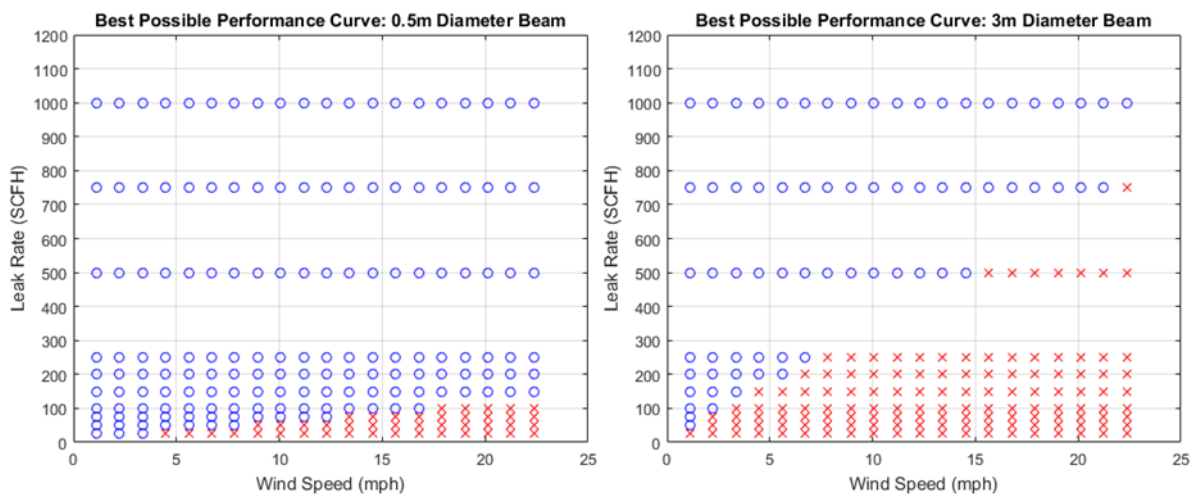


Figure 29: Example receiver operator curves for a set beam diameter and a range of wind speeds. Blue circles indicate detectable plumes. Red X's indicate non-detectable plumes.

In the two examples in Figure 29, both instruments are able to detect a 50 SCFH plume when the wind speed is negligible. However, even light winds of 5 mph are sufficient to make plumes smaller than 200 SCFH nearly undetectable to the 3-meter beam instrument.

The figures illustrate that the beam diameter and thus beam divergence angle, is a critical parameter. Designing without consideration of this factor may result in an instrument unable to find leaks that customers care about, even though the SNR equation may have instilled confidence in good performance. It may seem, then, that the best case solution is to design for the smallest possible beam diameter. However, another consideration is how the beam is scanned along the ground.

4.3 Impact of Error on Plume Area

Similar to the analysis of beam diameter and plume visibility is the impact of signal error on plume visibility. An understanding of the minimum measurable methane levels above the background can be

used to determine how much of a given plume is visible. In other words, as a plume disperses and the path-averaged ppm level drops, there will be a threshold set by the error where the plume cannot be distinguished from background. Calculating the visible area of the plume leads to an understanding of what kind of measurement density is required, which leads to laser pulse rates and scanning parameters.

This analysis again uses the Gaussian plume simulation approach. Plumes of various SCFH values are first calculated. Then, the boundaries of the plume are determined from the minimum detectable values. Finally, the area created by the boundary points is calculated. An example result is provided in Figure 30 with error thresholds of 10%, 20%, and 30% of a 1.8 ppm background, which are 0.18 ppm, 0.36 ppm, and 0.54 ppm respectively. For this plot, the instrument was at an altitude of 150 meters, the wind speed was 1 m/s, the beam diameter was 80 mm, and standard temperature and pressure values were used. A lateral resolution of 2.5 cm was used in the iterative search process and a 0.5 m resolution was used on the axial axis, leading to some small amount of rounding error.

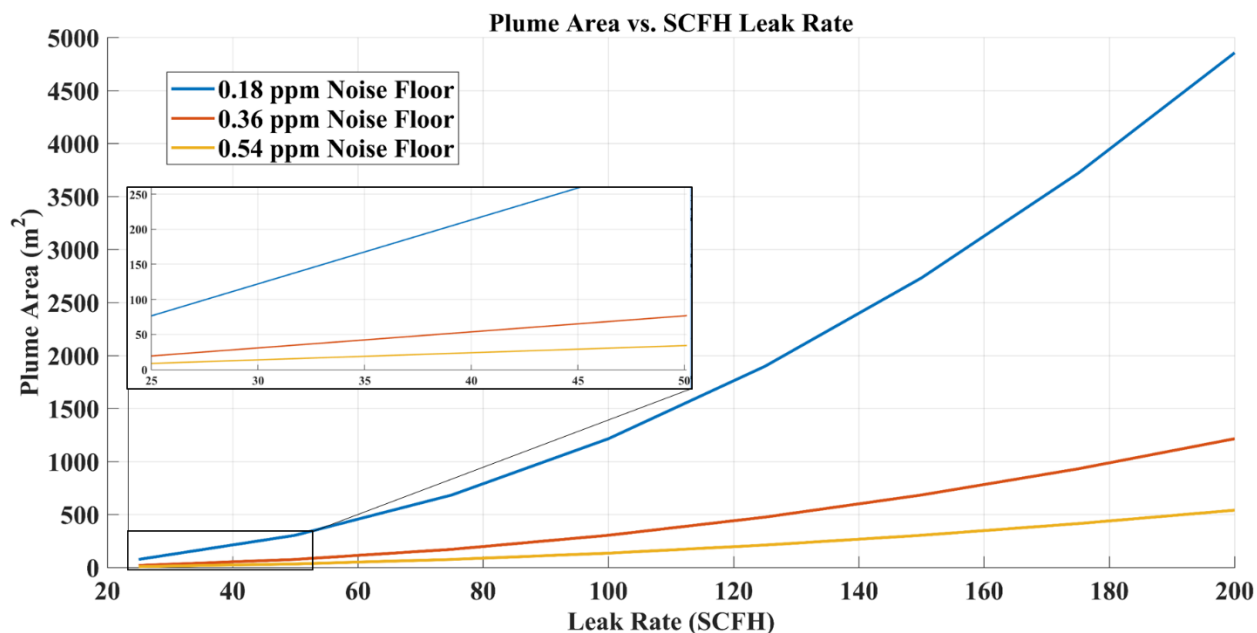


Figure 30: Detectable plume area vs. SCFH leak rate for three minimum observable methane concentration above a background concentration of 1.8 ppm. See text for assumptions.

The results show that scanning a 200 SCFH plume with sufficient numbers of points may not be that difficult with over a thousand square meters of methane for a well-designed system. On the other hand, a 50 SCFH plume can have less than 100 m² of plume area. These areas are not square, either. Plots of the plume area axial and lateral (not vertical) dimensions are shown in Figure 31. Note that the plumes are symmetric about the x-axis, but only one half of the lateral dispersion is shown to increase the scale. The plots show that leaks are much longer than wide. Some of the smaller leaks are just a meter or two wide.

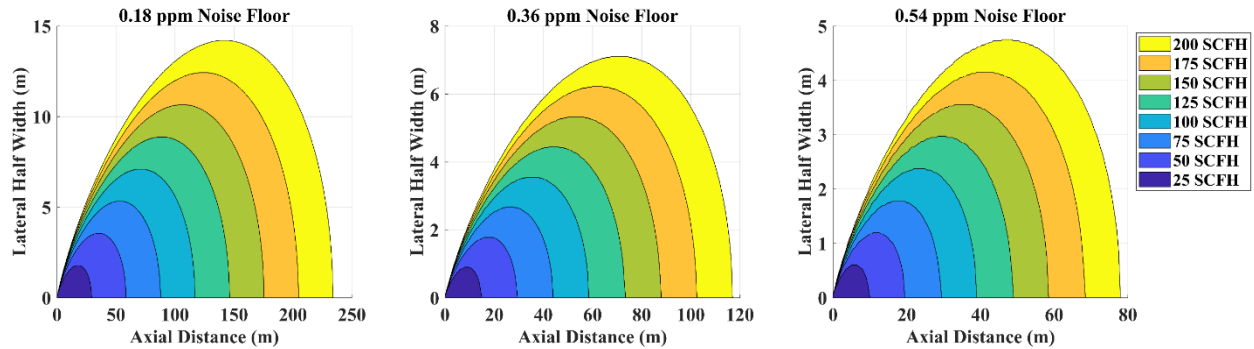


Figure 31: Plume dimensions as a function of noise floor and leak rate. The y-axis is the half-width of the plume.

The useful area is further reduced when the quality of the measurement points is considered. For instance, 10 points on the outskirts of a plume will not be as descriptive of the plume shape as 10 points clustered around the peak density. The topic of measurement quality is discussed in the next section.

4.4 Quality of Measurement Sets

The process of quantifying a leak rate and source is helped with both higher quantity and quality of data. Plumes naturally have their greatest concentration near the leak source and these are the most useful data to collect. Stated from the opposite point of view, the exponential spreading from the peak means the outer edges of a plume have minimal relative curvature and, therefore, minimal descriptive properties. Another cause for concern is that a large leak may have a 2 ppm concentration on its edge while a small leak may have a 2 ppm concentration at its peak. In this one-point example, if the large leak had been

sampled at its peak instead of its edges then there would be discriminating data to differentiate the two leak sizes.

Another Gaussian simulation effort was made to begin investigating the importance of scanning high quality points. To that end, a “Quality” ratio was calculated. The ratio is the measured volume of gas divided by the actual volume of gas:

$$Q = \frac{V_M}{V_A} \quad \text{Eqn. 9}$$

, where V_M is “Volume Measured” and V_A is “Volume Actual”. The actual volume is calculated by simulating a plume with high numbers of data points. The measured volume was calculated by taking varying numbers of points from across the plume, using MATLAB’ *delaunay* triangulation function to link the points, and then using the *alphaShape* function to calculate the volume of each wedge.

The first approach to developing an understanding of the required point density was taking a random number of points spread across random locations within an area. This approach is physically analogous to using a quadcopter’s hovering and tip/tilt ability to point the instrument in various directions, but perhaps without fine-control. For this first example, the beam diameter was set to 1m, the flight altitude to 1 km, the error to 10% of background (0.18 ppm), the wind speed to 1 m/s, and a standard temperature and pressure were used. The plume is from a 100 SCFH leak. The results of this approach are shown below in Figure 32 with 1,000 sample runs.

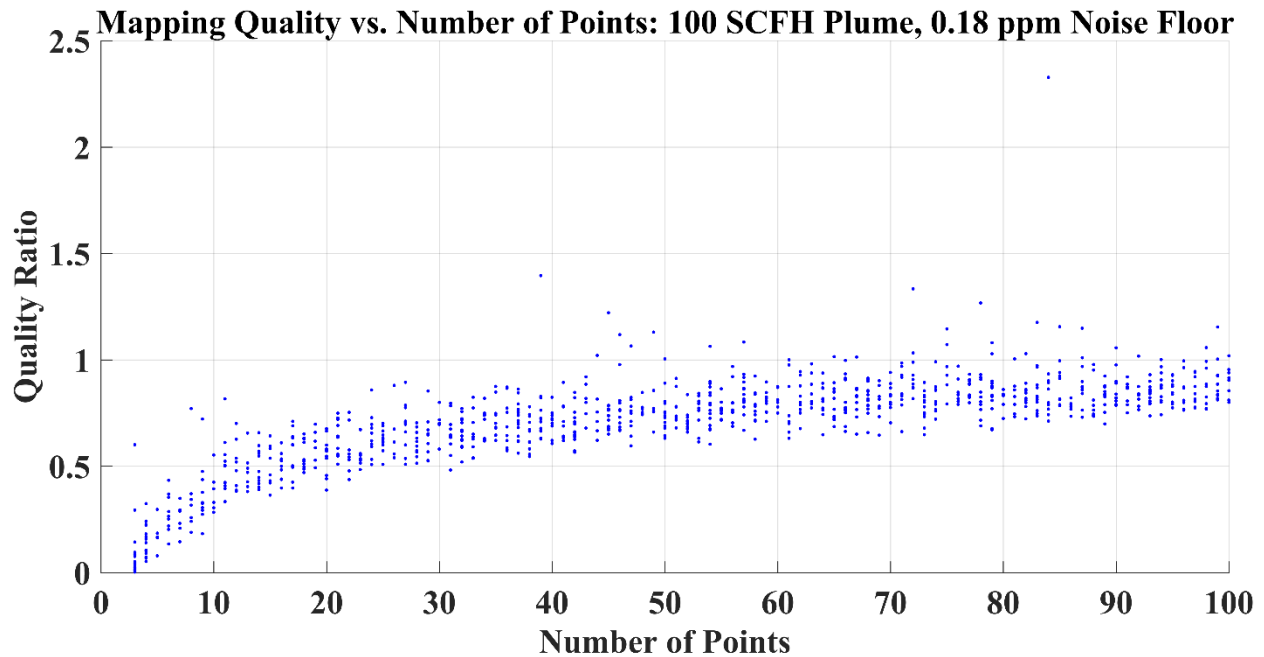


Figure 32: Plot of Quality ratios versus the number of random points measured within a 100 SCFH plume with an instrument noise floor of 0.18 ppm.

It is seen for this example that even 100 random points within a plume area will not reliably map the volume of a plume. It may, though, be close enough as many different algorithms could be applied to extrapolate the true volume given enough points.

The second approach is to consider scan resolution. In this scenario, a scan mechanism measures pulses at fixed lengths in both in-track and cross-track directions. Perfect scan resolution is not likely to be achievable in the real world, but this type of trade study can bound expectations from this direction.

The study took two approaches. The first looked at a 100 SCFH plume with a 0.18 ppm noise floor with an instrument at 500 feet and a wind speed of 1 m/s. The data in Figure 31 gave a visible plume width of about 14.2 meters and an axial length of 117 meters. An array of scan resolutions was setup, from 0.25 meters in the width and length directions to 5 meters. The points were aligned “worst case” in the plume’s lateral direction. This means that no measurements were taken on the most descriptive “spine” of the plume. Instead, the lateral points were set as far away from the spine as the resolution allowed.

The second approach introduced a “Scan Resolution Ratio”. This ratio is the scan resolution divided by the plume’s visible width. This makes the results plume leak rate agnostic, although the noise floor, flight altitude, and other parameters are still important. Within those bounds, though, the number of points required for good quality can be determined from the plot and then applied to Figure 31 to get a required scan resolution in meters.

The results of the two analysis runs are shown in Figure 33. The first thing to be noticed is the jagged nature of the data. This is due to “modes” of resolution in the lateral direction. The plume width is fixed so, as the resolution increases, the number of points that fit within the span decrease. Within a resolution range of a set number of points, the quality drops because the points nearest the center of the plume are moving away with increasing resolution. When the number of points that fit within the lateral direction drops, the quality slightly increases because the distance from the center spine of the plume drops.

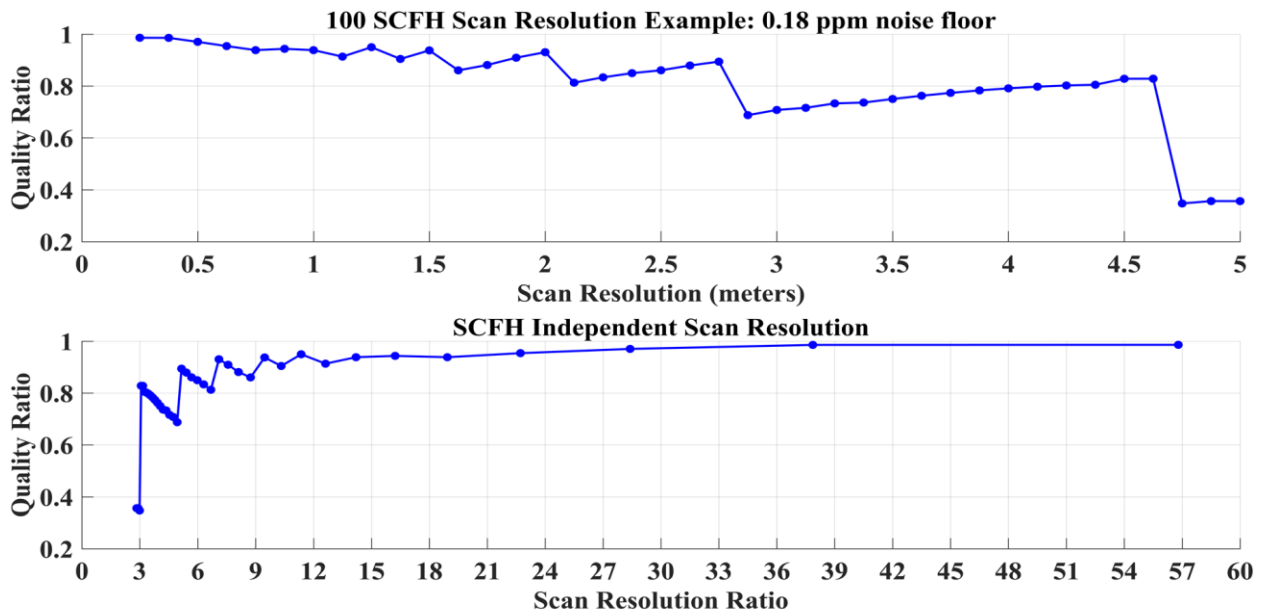


Figure 33: Plots showing how fixed scan resolutions map to measurement quality. The jagged nature of the curves is related to “modes” of scan resolution on fixed lengths.

The conclusion is that the likelihood of best measuring a plume starts with at least three measurement points across the width of a plume. Better is measuring at least 5 points. The benefit of this tier is that the

worst-case quality ratio is 0.8, whereas the quality of the second tier of 3 to almost 5 points can drop by about 15% depending on how unlucky you are when scanning. Even still, although not investigated here, advanced fitting algorithms may produce sufficient outcomes with quality around 50%.

This section looked at the number or resolution of points needed to get a good understanding of a plume shape. The next section considers achieving the resolution required when scanning from a fixed-wing aircraft.

4.5 Scan Pattern Considerations

For a fixed-wing aircraft flying along a straight path, the frequency of the measurement combined with the beam diameter drives measurement overlap. Consider Figure 34 which assumes the instrument points straight down to the ground (nadir). The first case with a minimum beam diameter may miss the plume entirely. The larger beam diameter may measure lower values of CPL, but it will at least return a measurement of the plume.

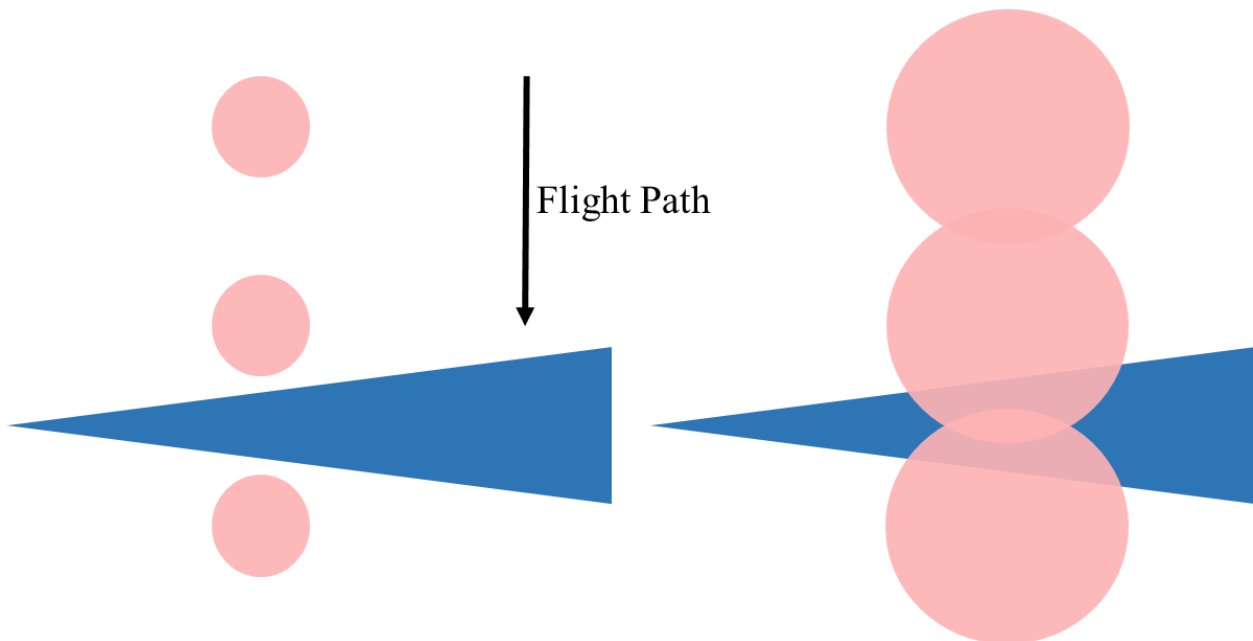


Figure 34: Overview of beam diameters on the ground (red) from an aircraft flying transverse to the plume (blue), demonstrating that beams that are too small may miss plumes.

Consider a case where the instrument points straight down from a UAV flying at 55 knots at the FAA maximum 500 feet. A low power laser is chosen and it can fire at 100 Hz. This gives a spatial distance between points of $55 \text{ knots} * (0.515 \text{ m/s} / 1 \text{ knot}) * (1/100 \text{ Hz}) = 0.28 \text{ meters}$. Meanwhile, back calculating the necessary beam divergence angle gives: $\text{Half angle} = \text{atan2}(\text{Beam Diameter}, \text{Range}) = 0.19 \text{ degrees}$. This is an achievable divergence angle with moderate effort and so the limiting factor of performance becomes either the aircraft speed, or the laser pulse rate.

Now consider a more complex case where the laser is fed through a rotating mirror system to cover more ground. An example is shown in Figure 35. The plotted pattern assumes a vehicle flying at 500 feet at 55 knots with a scan mirror pointing off-nadir at 6 degrees, rotating at 5 rotations per second and with a laser firing at 100 Hz. A pattern such as this might be used to more effectively scan an area rather than a line because it seems unlikely that: 1) a line will find the most optimal part of a plume, 2) that a UAV will stay level enough in roll to make the line straight, and 3) the goal is likely to map more of a plume than a couple of points.

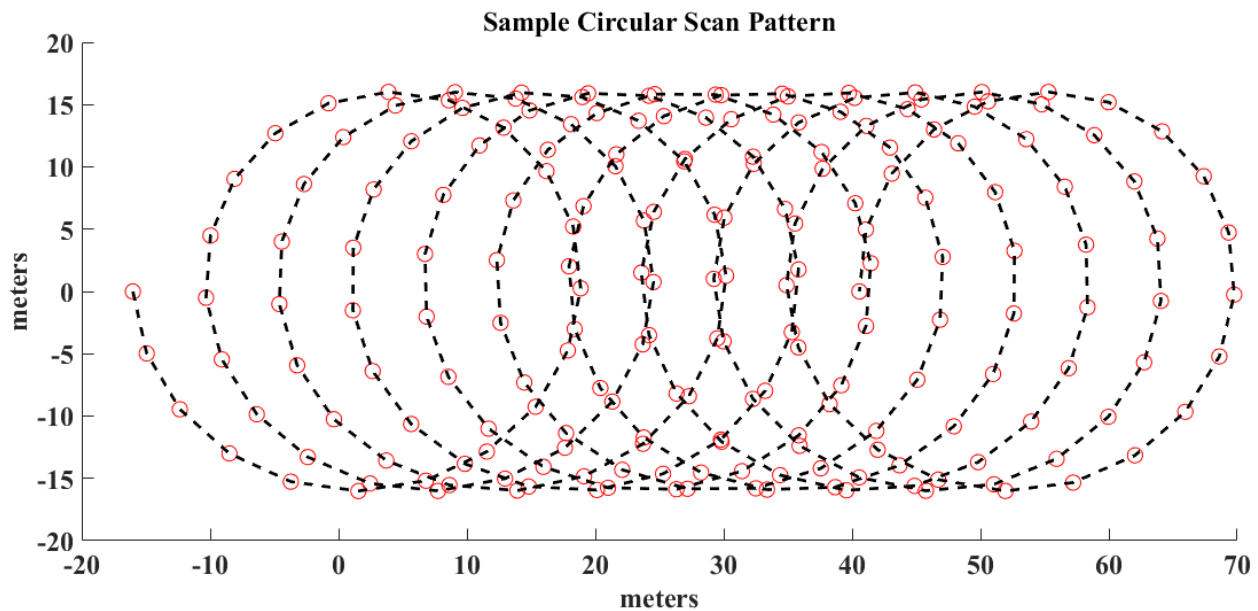


Figure 35: Sample circular scan pattern with red dots indicating pulse locations on the ground, following the scan pattern displayed as the dashed line. Scan rate sets the along-track point density while pulse rate sets the cross-track point density.

In the example above, the pulses fall about every 5.7 meters in the center of the scan pattern. Points at the edges have higher density. As a side note, a natural question would be “Can we optimize the system so that the circles fall neatly to optimize the distances between points at the center?”. The answer is yes, this can be done, but headwinds, pitch/roll, altitude uncertainty, and so on make lab optimizations irrelevant. It’s better to design with an assumption of unoptimized placement. Therefore, while the minimum distance could be about half of the 5.7 meters, the 5.7 meters is assumed. Comparing this result with the original no-scan resolution of 0.28 meters shows that a simple scan pattern increased the maximum point-to-point distance by more than order of magnitude.

With the goal of reliably scanning a 50 SCFH system in mind, section 4.3 showed that an area of 50 m² should be able to be reliably mapped. Figure 31 shows that a plume being measured with an instrument with a noise floor of 0.18 ppm has an axial length of about 30 meters and a width of about 4 meters in a 1 m/s wind. To add margin for winds of up to 10 mph (~ 3 m/s) the length and width become about 7 meters and 1.5 meters respectively. Using Figure 33 and appreciating how much of a challenge it will be to get very low scan resolution at this scale, three points are desired. This requires a scan resolution of $1.5/3 = 0.5$ meters. For a flight altitude of 500 feet and a scan angle of 6 degrees, the minimum pulse rate and scan rate of the laser can be determined.

Running a trade study on these parameters showed that a fast scan rate of 15 Hz and a laser pulse rate of 2,000 Hz would give a worst-case cross-track spacing of about 0.75 meters and an along track spacing of 1 meter. Therefore, even with these relatively high scan and pulse rates, seeing a 50 SCFH plume with 10 mph winds and this level of instrument noise will be a challenge. Although the plume may be detected, it is less likely that the leak rate could be determined without higher level algorithms.

One solution to better mapping would be to decrease the noise floor, which would increase the apparent width of the plume. Scanning faster seems unlikely based on scan hardware discussed in upcoming section 6.5. The pulse rate may be able to be increased based on Ball’s laser rate of 10 KHz, but that system is sized for an aircraft and shrinking it for UAV use may not be feasible at this time.

It should be noted that there are other factors which influence performance beyond the values discussed here. For instance, as seen in the Gaussian plumes section, the simulated plumes are smoothly distributed while real plumes are lumpy with pockets of higher density. The extra density in these pockets may allow an instrument to detect a plume that otherwise would be below the noise floor. Another consideration is that the noise floor is a standard deviation, which means there are instances when the actual noise floor is lower than that value. In those times, lower SCFH plumes may be visible.

The trade studies now turn to other considerations. The first is how the divergence of the laser beam can drive effective measurement diameters on the ground.

4.6 Influence of Beam Angle on Effective Beam Diameter

A study was done to investigate how beam width relative to the receive aperture diameter influences the effective diameter. In other words, can a beam footprint be so large that the angle into the receiver reduces observable energy at the edges? The hypothesis was that, even with Lambertian scattering, laser energy at the extreme points of a beam diameter have a lower return signal than points in the center. An exaggerated example of the concept is shown below in Figure 36. The heart of the issue is that the arc length between vectors to the edges of the aperture are shorter for points at the outer diameter than from the points in the center.

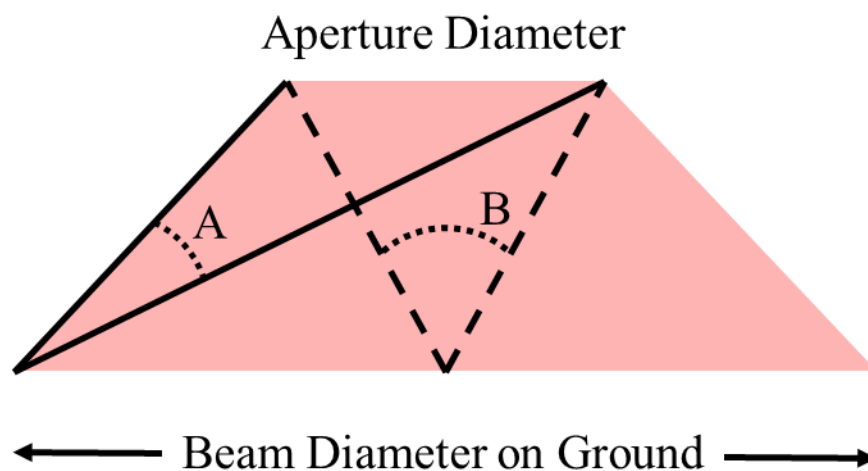


Figure 36: Exaggerated representation of subtended arcs created by rays reflecting from the ground.

An extreme test case was run assuming the instrument was 150 meters above the ground, the beam half-angle was 45 degrees, the aperture diameter was 10 cm, and the reflectance was Lambertian. The returned signal as a function of the maximum value returned is plotted in the left-side of Figure 37. This 300-meter diameter is nowhere near desirable, as described in the previous sections, but illustrates the concept well: The outer edges of the beam reflect half the energy back to the receiver that the center of the beam reflects. When this result is multiplied through a Gaussian profile, the right-side plot of Figure 37 is created. It is seen that the effective beam diameter is reduced relative to the beam that does not take into account reflection angles. For this example, the effective beam diameter is about 85% of the nominal beam diameter. Or, in terms of angles, the 45-degree half-angle is effectively a 40.5-degree half-angle.

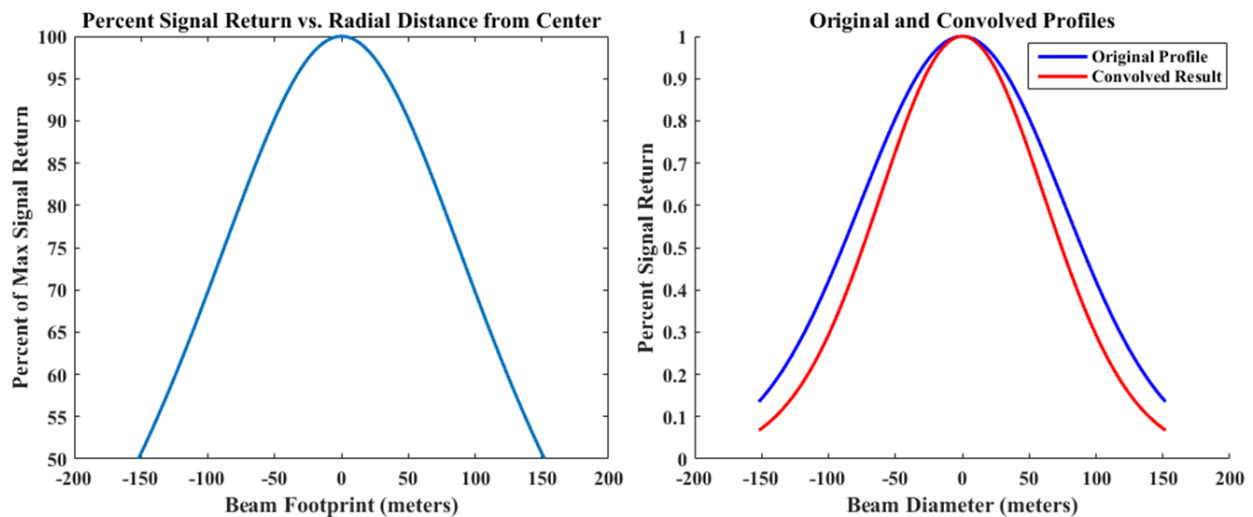


Figure 37: The plot on the right shows the percent of energy reflected to the instrument relative to the peak value. The plot on the right convolves this result with a Gaussian profile.

The loss of angle, even in this extreme beam divergence case, is relatively small because there wasn't much energy at the outer fringes of the Gaussian profile. This point is extended with the analysis of a wide range of initial beam half-angles to plot the sensitivity in Figure 38. The same assumptions of a 150-meter instrument height and 10 cm aperture diameter are used here. It is seen that a 1% reduction is not observed until the beam half angle is about 10 degrees. A collimated beam will not (should not) come

anywhere near that threshold. This conclusion shows that, while the effect is real, it's negligible.

Therefore, the consideration of beam diameter on energy loss is not carried through the rest of the analysis.

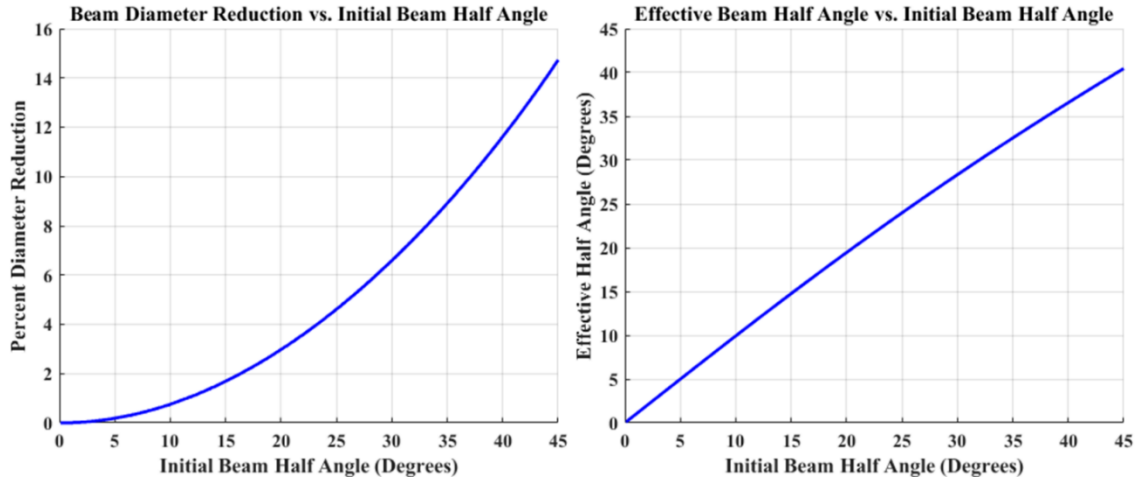


Figure 38: Figure on the left investigates how much the effective diameter is reduced as a function of percentage while the right image looks at the result in terms of resulting angles.

It is worth noting, though, that uncollimated beams could be impacted by this. For example, a numerical aperture value of 0.22, common in many fiber optics, translates to a half angle of about 12.5°. A numerical aperture of 0.37, also seen in commercial fibers, has a half angle of about 21.7°. These ranges of numerical aperture mean about 1% to 3.5% reductions in beam diameter, which is still small, but could be approaching meaningful for some applications.

4.7 Speckle Considerations

One of the contributors to measurement error with airborne DIAL systems is laser speckle noise [Goodman, 2007]. Speckle has two definitions: one where variations in return intensity are created by the reflecting surface and one where variations are caused by multimode fibers. The root cause is largely the same, though, as phase differences in the laser beam create interference. In multimode fiber, which is often used in DIAL instruments due to power limitations in single-mode fiber, the variations in path length and indices of refraction are the primary contributors. The challenge is that the speckle influences

measurements, creating false positives and false negatives. A brief speckle simulation demonstrates the effect.

Two laser beam profiles were simulated by stacking Zernike functions, introducing phase noise, and smoothing the result. A typical result is shown in Figure 39. The first assumption used was that the total energy in both figures is precisely the same. This does not have to be true in the real world because the variations are caused by phase differences with no law forcing two separate beams to have the same values. However, for illustrating the potential impact of speckle, it's informative to look at two equal energy beams to specifically investigate the impact of the distribution of energy. The second assumption is that the most important mode was the Gaussian mode. This had a weighting factor of one when stacking the Zernikes, before normalization. The other modes had a value of 0.5 applied and the sum of all the modes was normalized. In real world speckle, there could be significant variation in the distribution of modes. Again, though, the intent is to get a baseline for relative analysis.

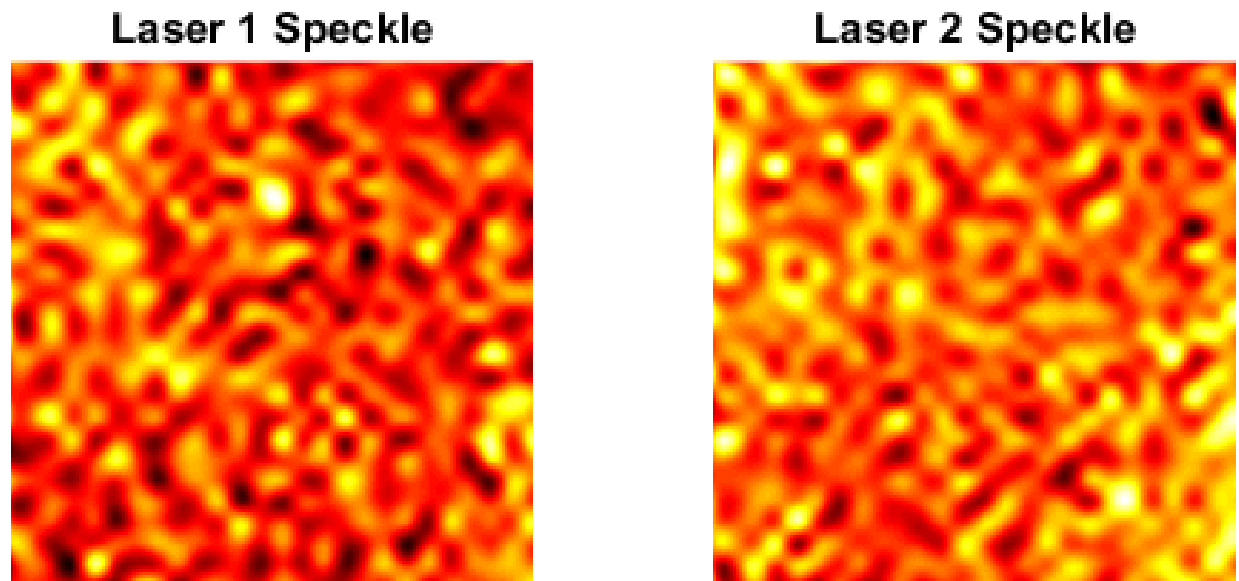


Figure 39: Simulated speckle patterns with equal total energies. The images were created by introducing phase noise to stacked Zernike functions to simulate multimode fiber.

A challenge with airborne DIAL instruments and speckle is that the reflective target is always changing. [Amediek, 2009; Lawrence, 2010]. The reflectivity values that interact with different parts of the beam profile thus change. This is a problem when a high intensity part of one laser hits, for instance a shiny surface while the beam profile of the second beam is less intense in that region. More energy comes back from the first signal. If the first beam was the offline wavelength and the second is the online wavelength with 20% less returned signal, then it would be mathematically equivalent to saying the two beams were equal and that 20% of the online wavelength was absorbed by methane. This is how speckle creates false positives. False negatives are created by the online being higher energy in the speckle region while going through a methane plume. The extra reflected energy offsets the absorption and the zero delta difference failed to show that methane existed on the path.

This result is demonstrated graphically in Figure 40. There, the beam profiles of Figure 39 are divided into sequentially smaller regions. The first image shows that if the ground is uniformly reflective and the sensor captures the entire beam, then speckle has no influence. However, as the beam is divided up, the magnitude of the potential impact changes. Even by the second iteration (2x2 grid), the lower right quadrant has an energy ratio difference of about 5%. By the 6x6 iteration there is a 30% difference in energy within the yellow regions. In some simulations, the highest ratios reached values of above 50%.

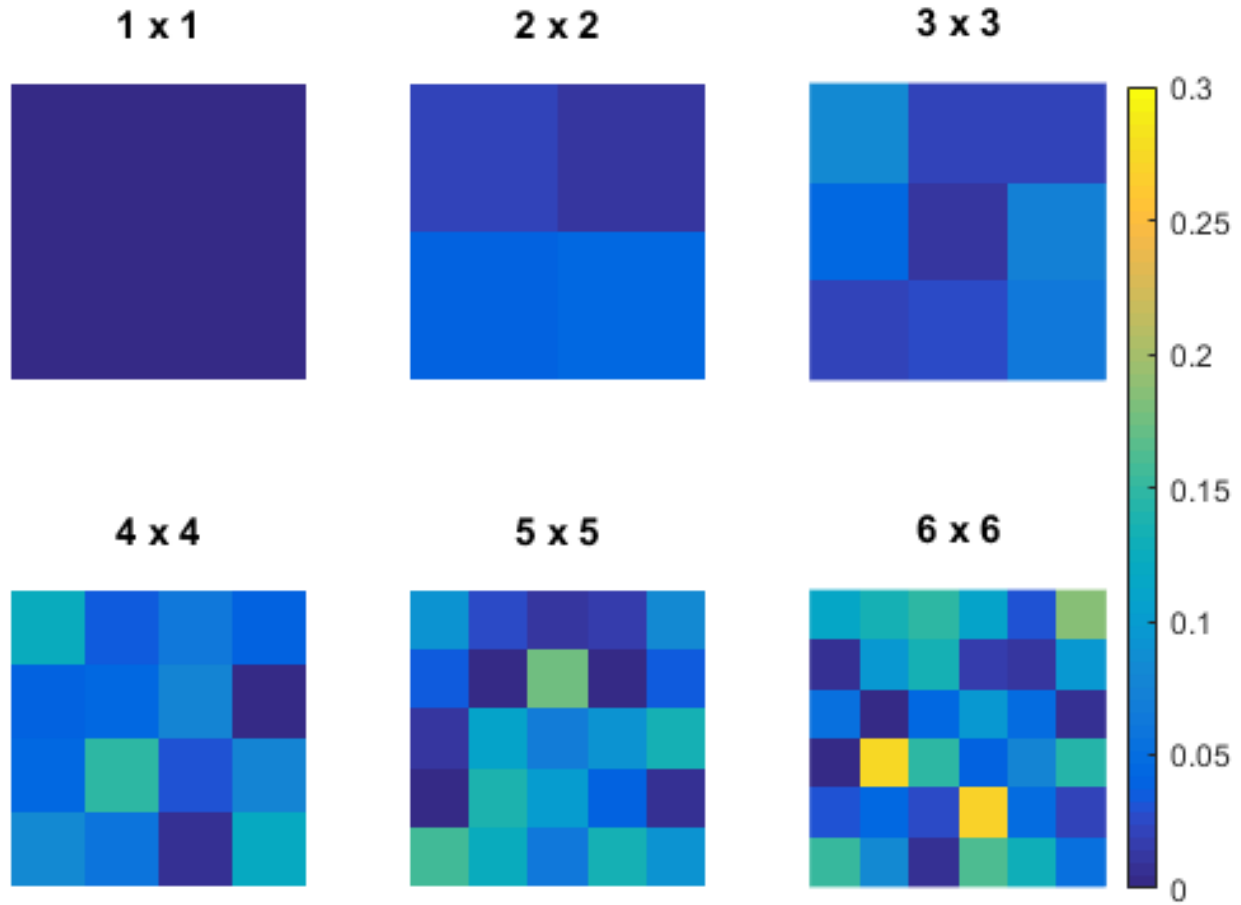


Figure 40: Ratios of energy from the speckle patterns. Values are absolute differences from unity. The first plot has a value of zero, demonstrating equal total energy.

Extending the analysis, a scenario is investigated where most of a scene is concrete, but one section has glossy paint. This example is relevant to the oil and gas industry as they use concrete well pads and paint their infrastructure to reduce corrosion. The differences in reflectivity for these two materials are shown in Figure 41 [Baldrige, 2009]. It is seen that, at a wavelength of 1,645 nm, concrete reflects about 33% of the laser energy while glossy paint is roughly double that at about 65%.

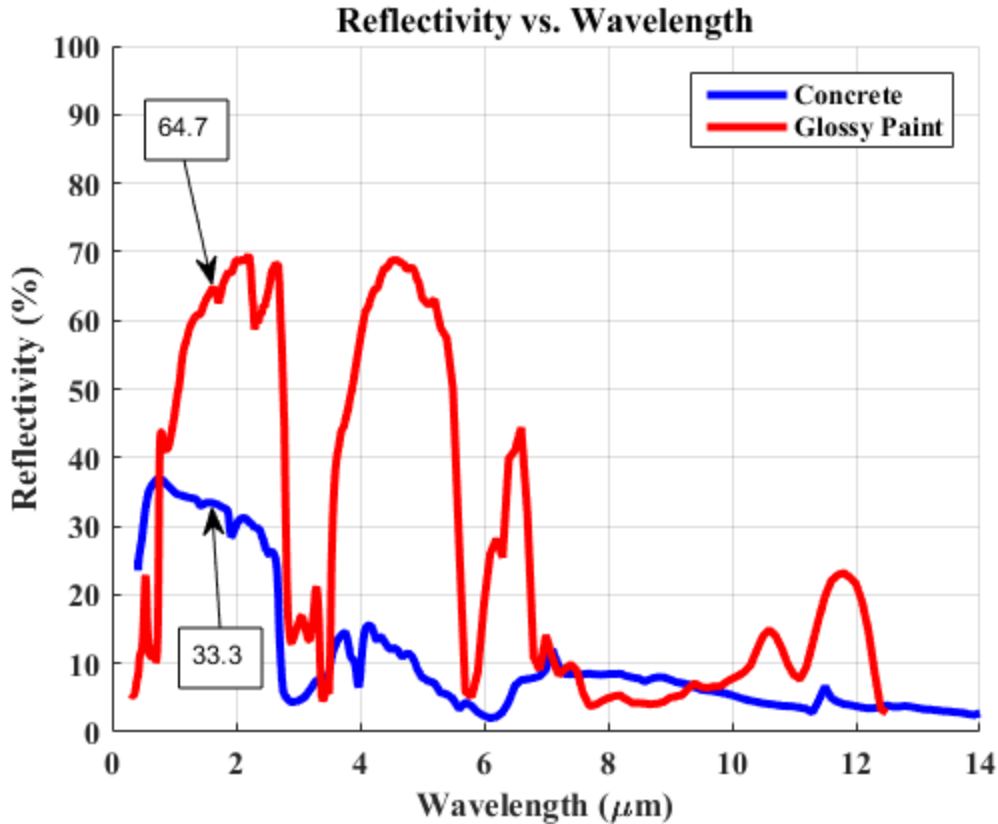


Figure 41: Reflectivity vs. wavelength plot for two materials commonly seen around oil and gas infrastructure. Data from JPL ASTER spectral library [Baldrige, 2009].

Figure 42 puts it all together by calculating the ratio of each sub-region, running a scenario where that region hits glossy paint while all other regions hit concrete, and then takes the natural log of the integrated results. (Recall from the DIAL equation 4 that the ratios are within a natural log function.) The colors of each region in Figure 42 represent the final result from this operation for that region. It is seen from this example that the peak error contribution due to speckle and reflectivity variations is on the order of 1%. Note that the smaller regions in the 6x6 image have higher ratios of energy, but their overall contribution is small relative to the total.

In other sample runs, values of up to 2% were seen, although those were rare. When looking at more extreme reflectivity variations, such as assuming 90% reflectivity, the errors were in the 5% to 6% region. Anecdotally, data provided by Ball Aerospace illustrates more extreme instances when flying over

parking lots. Asphalt has a reflectance of about 15% at their wavelength while cars are likely close to 90% or higher due to flat, shiny hoods and car tops [Baldrige, 2009]. The result is significant clusters of false positives. Flying over lakes and other water features also produces noise where wave slopes bounce some parts of the signal back, but not others.

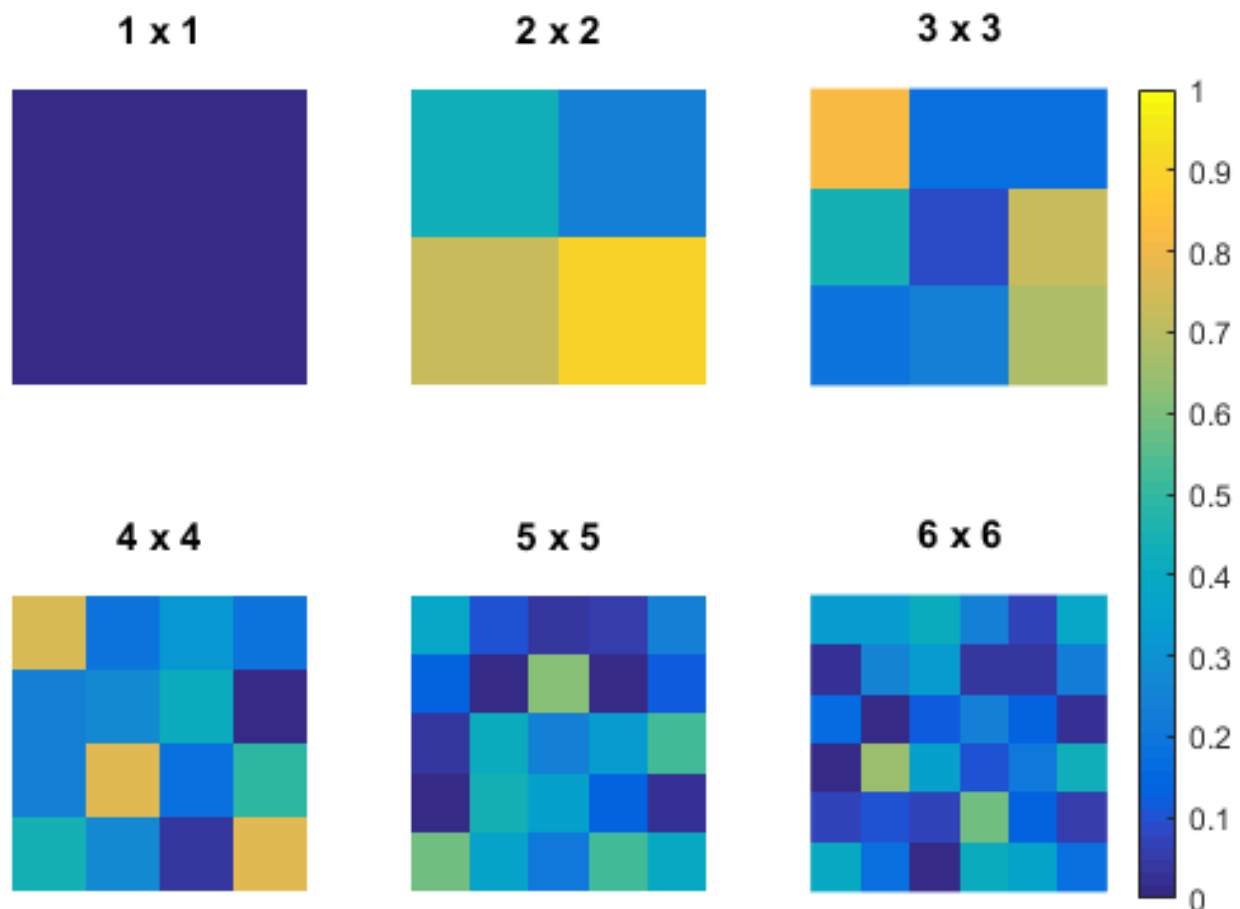


Figure 42: Percent error contribution due to reflectivity variations of concrete and glossy paint for the laser speckle examples.

The takeaway from this quantitative analysis is that, for the general scenario investigated over a well pad, the contribution is a small 1%. For scenarios with greater differences in reflectivity, the error contribution can approach 10%. This is an interesting result as actual flight campaigns over southwestern Europe showed that, for beam diameters on the order of 10 meters, reflectivity-based variations on the signal could be “several tens of percent” [Amediek, 2009]. They also showed that a beam diameter of 100

meters would drop their error to less than 1% due to effectively spatially averaging. Going with a smaller beam may also drop the error for areas with relatively homogenous reflectivity on short spatial scales.

For this dissertation, a value of 2.5% error due to speckle is carried. The value may be slightly conservative for most cases, but the evidence in Ball's real-world data suggests that when it does occur, the effect is quite noticeable and undesirable.

Chapter 5: Wavelength-Based Reflectivity Variations

Originally, investigating how reflectivity at different wavelengths impacts the final answer was going to be another quick analysis. This effort bloomed, though, into a full paper. The peer-reviewed work was accepted for publishing by SPIE's Journal of Remote Sensing in August, 2017 [Tandy, 2017]. This chapter presents and extends the effort.

5.1 Introduction

As seen in Equation 4, the third term covers the reflectivity difference between the online and offline wavelengths. As with all of the terms after the first, the goal of design is typically to make the argument of the natural log “1” so that the term goes to zero. For the third term, this means the reflectivity of the online and offline should be as equal as possible. Historically, this was addressed by making the wavelength as close as possible to each other [Schotland, 1964, Menzies, 1977]. Choosing nearby wavelengths also minimizes differential scatter, such as Rayleigh scatter, in the atmosphere.

This atmosphere effect is the one that has received the most attention, both with numerical analysis and field testing, especially in the UV region where wavelengths more strongly interact with molecules [Schotland, 1974; Browell, 1985]. In the first example, Schotland found that, even in the UV region, the spectral variation in backscatter was negligible for wavelength separations of a few tenths of a nanometer. In the second example, Browell et al found that the worst-case errors were less than 10 ppbv between the wavelengths of 286 and 298 nm, or roughly 10% of their original signal, while in some cases the error was less than 1%. This is not to say that the individual wavelength effects are not important. It is only demonstrating that the difference in these effects for narrow regions of wavelengths can be small.

Unfortunately, issues of reflectivity have not seemed to enjoy the same attention with no papers or textbooks being found covering the issue despite extensive searching both in literature and through researchers. No historical reason could be found explaining the discrepancy. However, as a theory for the shortfall, the issue only becomes relevant when the target is changing, which typically limits the effect to

airborne instruments. A lab or ground-fixed DIAL instrument can more readily investigate backscatter than surface variations. Numerical analysis might also have been historically difficult with few databases and the simulation of wide varieties of molecule absorption spectrums being a challenging solution path. Finally, it may be that the impact is not deemed particularly worth knowing about as there are only so many practical lasing wavelengths that match with absorption features. Designing the system with wavelengths as close as possible and accepting errors may be a more practical approach for developers with limited budgets and time.

The one relatively modern discussion found was started in an Ehret et al. 2008 paper [Ehret, 2008]. There, they investigated a range of theoretical impacts on instrument performance for a space-based, integrated path DIAL CO₂ mission. In their analyses, they assumed that the difference in reflectivity divided by the reflectance of the online wavelength ($\Delta\rho/\rho$) was less than 10^{-4} . From there, the discussion took a tangential direction. Amediek et al. followed up on Ehret's theories by flying airborne campaigns investigating ground reflectance for lasers with wavelengths near 1.57 μm [Amediek, 2009]. They observed that reflectance can vary significantly at a variety of spatial scales over both terrain and water. They calculated between 13% and 54% reflectivity differences for a 50-km flight path taken over snow free and partly snow-covered ground assuming a 10-m pulse-to-pulse ground separation distance. Averaging 350 measurements reduced the error to less than 0.11%. Spatial reflectivity was also the subject of a 2010 effort by Lawrence et al. where quantitative analysis using surface reflectance maps showed that the error contributions from mountains and coastlines may exceed 1% [Lawrence, 2010].

The paper found to be closest to addressing the wavelength-based reflectance of the surface is a 2016 study by on the Jet Propulsion Laboratory's (JPL) 2.05- μm laser spectrometer and its interaction with snow-covered surfaces [Spiers, 2016]. In their quantitative study of data collected in operations, they determined that the materials they fly over can have nearly order of magnitude variations in reflectance. Their resolution is even high enough to be able to estimate the relative ages of the snow as older snow tends to be coarser with higher reflectance values. Unfortunately, their discussion of how wavelength

changes might affect the outcome is limited to a brief comparison of trends with 1.57- μm systems and a statement that the general trends in bidirectional reflectance distributions should be similar.

Although not specifically investigating the relationship between online and offline wavelengths and reflectivity, these works do highlight the impact that varying reflectivity can have on the final result and that it may not be negligible. Other authors have also made advances by demonstrating that single-wavelength systems see improved classification results when considering intensity variations due to surface reflectivity [Burton, 2011; Yan, 2014]. For example, Burton et al. used measured reflectivity data to demonstrate that quartz-rich sandstones are more reflective than clay-rich mudstones at their wavelength of 1500 nm. If everything else about the instrument is consistent, then sweeping the beam over terrain and observing intensity variations in the signal allow them to narrow down the potential materials they are reflecting off. Papers such as these hint at the potential real-world influence of reflectivity on multiple wavelength systems but do not touch on the differences in reflectivity over a wavelength range.

In spite of the apparent lack of published analysis on wavelength-based variations in ground reflectivity, an ability to quantify their impact may be useful as there are cases where being able to choose an offline wavelength “far” away from the online one may be beneficial. For instance, there would be advantages to sending both the online and offline pulses at the same time: it cuts down on the differential terrain response between pulses for moving instruments and interfering atmospheric structures along the path would be temporally and spatially identical. However, sending both pulses at the same time can be difficult to implement if the online and offline pulses are close enough that wavelength separating optics in the receiver become problematic. Additionally, systems that measure more than one constituent with multiple wavelengths may benefit from having one offline signal common to multiple online signals. For instance, there are regions where both H₂O and CO₂ absorption peaks are nearby [PNNL, 2015]. A system able to measure both molecules concurrently would be collecting data on the two most important greenhouse gases with a single instrument. Finally, although not exhaustively, there may be a particularly

attractive online wavelength for myriad reasons, but the nearest offline wavelength could be nanometers away. It may be that, even though the separation is far, the reflectivity is favorable, so the instrument design process can move forward. Converse to the positive examples, it may be that the reflectivity variation gradient is particularly bad in the wavelength region of interest, so even common assumptions about sub-nm separation are not as robust as expected. Several of these specific examples are addressed in the applications section of this paper.

5.2 Data

The reflectivity data used in the analysis came from the JPL Advanced Spaceborne Thermal Emission and Reflection Radiometer (ASTER) spectral library [Baldrige, 2009]. The library has over 2,400 unique materials and is built from three different data sets with contributions from JPL, John Hopkins University, and the United States Geological Survey Spectral Library. The library is comprehensive and

“includes spectra of rocks, minerals, lunar soils, terrestrial soils, manmade materials, meteorites, vegetation, snow, and ice covering the visible through thermal infrared wavelength region (0.4–15.4 μm)”. [8]

The data are organized by type (i.e. Rocks, Vegetation), then by class (i.e. Igneous, Metamorphic, Sedimentary), and finally by material name (i.e. Granite, Basalt). Other material libraries exist, such as the ISRIC World Soil Information Database, but are excluded at this time to focus on the ASTER materials.

The data were brought into Matlab, taking care to catch miscellaneous issues such as having multiple reflectance values for the same wavelength, identifying data points with non-numeric values, and finding data sets that began with nonsensical wavelengths of 0 microns. Another issue was that most files were “reflectance vs. wavelength” while some were “transmission vs. wavelength”. The transmission files, of which there were 158, were removed, leaving 2,287 reflectance datasets.

Within this remaining set, there are a number of recurring materials. Duplications most often occurred within the “minerals” type. For example, “calcite” had samples with origins from both England and Italy. The other primary duplication was in having one type named “rocks” and one named “rock.” Johns Hopkins University produced the former and JPL the latter. Each had igneous, metamorphic, and sedimentary classes with duplications within the duplications. For example, each igneous class had multiple basalt entries. These types of duplications were merged into consistent type, class, and material names to create the final unique set. For instance, the types rock and rocks became just rocks and calcite was sorted into calcite 1, calcite 2, and so on to keep the subtle variations without confusing the data mining code. Table 6 provides a breakdown of the types, classes, and the number of relevant datasets.

Table 6: Type and class listing of the reflectance datasets. For compactness, the meteorite and minerals classes are summarized.

Types	Classes	# of Datasets
Manmade Materials	Road Asphalts and Tar	5
	Concretes	5
	General Construction	28
	Roofing Materials	18
Stony Meteorites	<i>Six Classes</i>	60
Rocks	Igneous	226
	Metamorphic	114
	Sedimentary	129
	Intermediate	4
Soils	Terrestrial	52
	Lunar	17
Vegetation	Grasses	2
	Trees	2
Water	Distilled Water	1
	Sea Water	2
	Tap Water	1
Frost, Snow, and Ice	Frost	1
	Ice	1
	Snow	3
Minerals	<i>19 Classes</i>	1,616

Clearly, not all types and classes are equally relevant for all DIAL applications. Stony meteorites and lunar soil samples are important, but not for airborne applications seeking CO₂ or methane. Therefore, in the analysis section we look at breaking the data down further into additional groups.

Important for all datasets is that data in the ASTER library cover two spectral regions: 0.4 - 2.5 μm and 2.0 - 15.4 μm . Figure 43 shows a few sample materials that span large wavelength ranges.

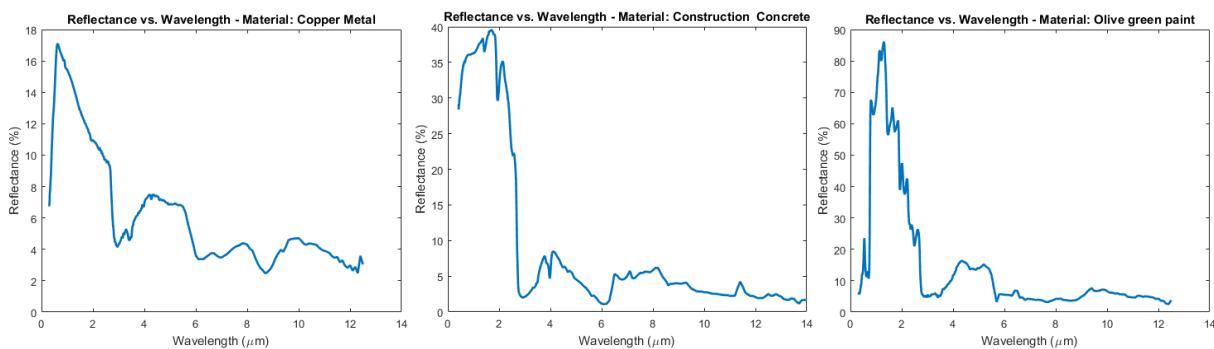


Figure 43: Wide varieties of materials demonstrate precipitous reflection increases until around 0.7 μm and decreases around 2.7 μm .

Detailed information on their data collection methods can be found in the Baldrige ASTER paper, but, briefly, they use gold as their reflectivity standard between database sources and water/pyrophyllite for the visible and infrared regions, respectively, within the JPL data set. In the 0.4- to 2.5- μm region, a single pass monochromator with a diffraction grating is used to collect data on a 1-nm scale from 400 to 800 nm and a 4-nm scale from 800 to 2500 nm. In the 2.5- to 15- μm region, a spectrometer with an integrating sphere took 1000 scans with a wave number delta of 4 cm^{-1} over a span of 15 min with background measurements being subtracted in post processing.

For our analysis, results are presented for two wavelength ranges: the first is 400 nm to 2,500 nm and the second is 2,500 nm to 4,000 nm. The first is designed to fall within the range of the data's natural boundary. The second picks up where the first left off and ends above where most DIAL systems operate. Although there are likely to be systems that operate at wavelengths longer than 4,000 nm, the trends are similar so that extending our analysis is a case of diminishing returns.

Within the datasets, the spectral library has varying wavelength resolutions. Sometimes sub-nm resolution is available while, in others, 10's of nm is available. The wavelengths of the data are not consistent either as some data is available at, for example, 3.9800 μm while others are at 3.9817 μm . To provide a consistent analysis, we used interpolation to align the analysis wavelengths. A quick trade study showed spline interpolation was not sufficiently consistent while Matlab's PCHIP interpolation scheme provided excellent results with less than $0.01 \pm 0.02\%$ error relative to a linear interpolation. Figure 44 graphically demonstrates an example interpolation.

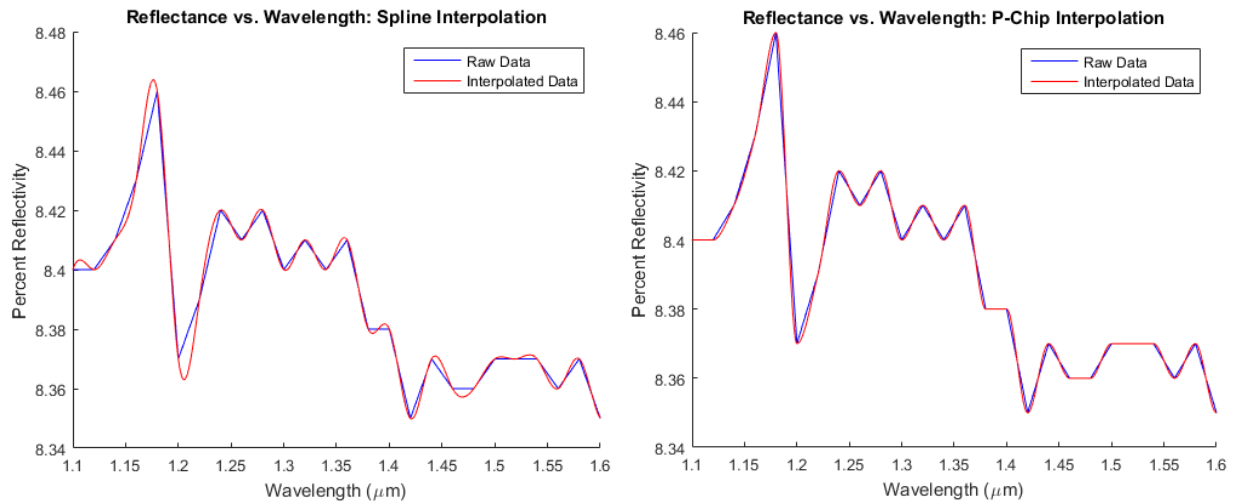


Figure 44: Spline interpolation incorrectly added reflectivity at peaks and shelves. MATLAB's "PCHIP" implementation more closely followed the underlying data and was used in the analysis

A final point worth discussing on the database is that the number of materials within each range varies. The number of materials is plotted in Figure 45, showing changes in step-wise fashion with increasing wavelength. Above 2,500 nm, the number of materials drops to 1,365 and maintains that level throughout the range. This is a natural result of combining databases from various organizations with different motivating interests.

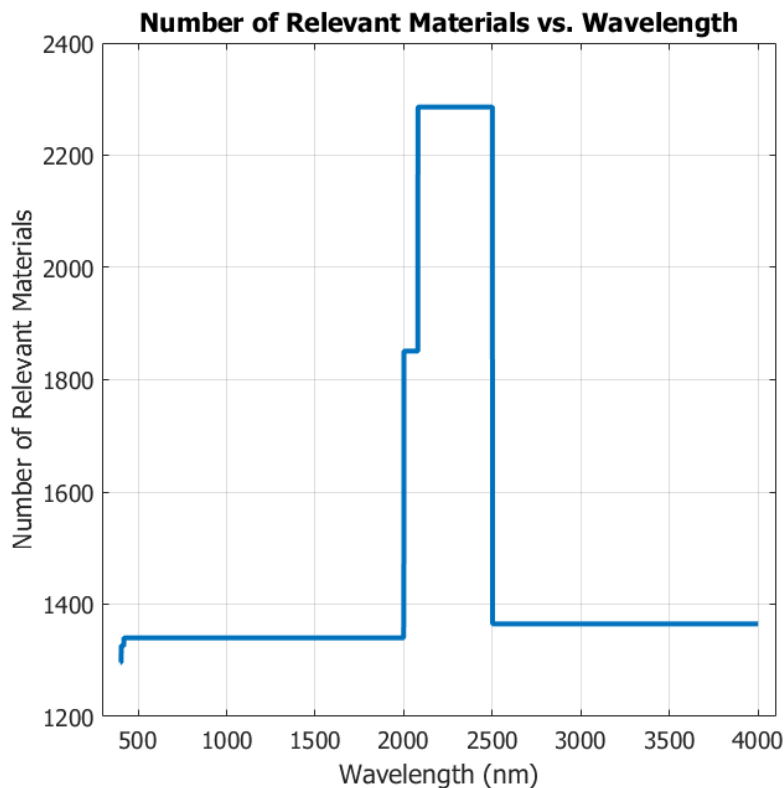


Figure 45: The materials in the database that are relevant to the analysis vary with wavelength.

The spectral data for absorption were taken from the Pacific Northwest National Labs' (PNNL) database. Their datasets are typically offered in several different temperature ranges. For this analysis, the 25°C files were used and no pressure broadening or other effects were attempted. The goal is to provide general trends, and it is expected that instrument developers likely will use their own absorption line values relevant to their specific application.

5.3 Methodology

The first step in judging the acceptability of reflectance variations is to set a goal. For this analysis, the assumption is that an instrument developer would be satisfied with a 5% or less error contribution on the final gas concentration measurement due to reflectance variations of the online and offline wavelengths. This number was chosen based on a 1% error possibly being considered negligible or in the noise and 10% feeling like too much of an impact. It was felt that 5% was near the threshold of being important

enough to warrant extra attention. Instrument developers interested in tighter error bounds can expect shorter allowable wavelength deltas from those shown in the results section.

Next, as can be seen from Equation 4, the reflectance contribution is a subtraction from the power ratio; thus, the 5% threshold is a moving target. Therefore, the power ratio must be fixed for consistent analysis. Because molecular cross section, atmospheric environment, and other factors contribute to the received power ratio, the authors recognize that choosing a value is fraught with difficulties. Indeed, infinite combinations of transmitted power and percent absorption can result in the same ratio. As a starting point, however, a received offline power of 10% greater than the received online power is assumed. This value was chosen based on assuming nearly equal offline and online transmitted powers and that the selection of an online wavelength is made such that roughly 9% of the signal is absorbed in typical operations to provide a clean signal. Developers anticipating less online absorption can expect tighter restrictions on delta wavelength because the log ratio that the 5% error is measured from is reduced. Likewise, anticipating greater absorption allows for relaxed wavelength separation limits.

For the nominal case of equal transmitting power, a 9% online wavelength absorption, and a 5% error tolerance, we can express the acceptability threshold as

$$\ln(1.1/1.0) * .05 = 0.0048 \quad \text{Eqn. 10}$$

This 0.0048 threshold is applied by first taking the natural log ratio of an offline-to-online reflectance pair and then comparing the result against the threshold. If the value is equal to or less than 0.0048, then the error imposed by wavelength-dependent reflectance is less than 5% for the stated conditions. Values higher than this threshold mean that reflectance variations are likely to introduce appreciable noise in the final result. Note that factors other than wavelength-based reflectivity may also impact the final result since effects, such as Rayleigh scattering, are ignored in our analysis.

Three cases are presented in the analysis results section: the first demonstrates the concept of equal reflectivity not necessarily needing to come from nearby wavelengths, the second looks at all of the

possible online and offline wavelength combinations, and the third investigates a narrower region around a spectrum of online wavelengths. The second case is useful for the big picture view of the design space while the third is intended to be useful to designers looking for early-stage feedback on wavelength selection. In both of the latter cases, the plots are created by selecting an online wavelength and then calculating the natural log of the offline-to-online ratio for many different offline wavelengths.

Within the last two cases, the first uses an online wavelength that is incremented by 1 nm and an offline wavelength that is incremented by 0.5 nm to create just under 26 million offline-to-online ratios. In the second, the online wavelengths are incremented by 0.1 nm and the offline wavelengths are incremented by 0.05 nm within a range of -6.0 to $+6.0$ nm from the online wavelength to create a little over 8.7 million analysis cases. Both sets were processed using MATLAB's Parallel Processing Toolbox on an Intel i7-930 CPU with 6 GB of RAM, taking about 12 and 4 h, respectively.

For the latter case, the next step was to determine how far the online and offline wavelengths could be separated. Thresholds of the percentage of materials were set and the separation distance between wavelengths were determined that met these thresholds. The percentages chosen were 90%, 95%, and 99%, which mean, for instance, that 90% of the materials in the database would introduce 5% or less error for the calculated separation distance. These latter plots are intended to provide more specific guidance to instrument developers seeking information about trends at their wavelengths of interest.

The final step taken is to present the results by type. The first group is all materials and simply represents all of the materials in the database. The second is a grouping of the rocks, minerals, and soils types.

Within this set, the soils type had the lunar regolith class removed and the stony meteorites were removed from the rocks type as these were not deemed practical materials for most DIAL instruments. The third group consists of the water-based types of snow, ice, frost, and water. The vegetation type would be of keen interest, but it only has four samples and so is not further analyzed. Likewise, the manmade materials group suffers from only having 56 samples, which is not comprehensive enough to justify inclusion.

After the plots of the various results, a brief study is made of the impact of choosing different parameters for the error threshold. Instead of the offline-to-online signal power ratio of 1:1 and an allowable error contribution of 5%, a range of the two values are compiled for a given wavelength. The wavelength chosen is 1645.55 as it is used in two recent methane DIAL systems and is also the peak absorption value for methane in the near-infrared.^{6,8}

5.4 Analysis Results

Analysis of individual materials shows that there can be wide swaths of wavelengths where the reflectance of a material would support hundreds or even thousands of nanometers of wavelength separation. The plot in Figure 46 shows the reflectance of marble with the solid line highlighting ranges from which online and offline wavelengths could be chosen while meeting the 5% error contribution criteria. It is seen that, for single materials, satisfying the assumption of equal reflectance can be relatively straightforward.

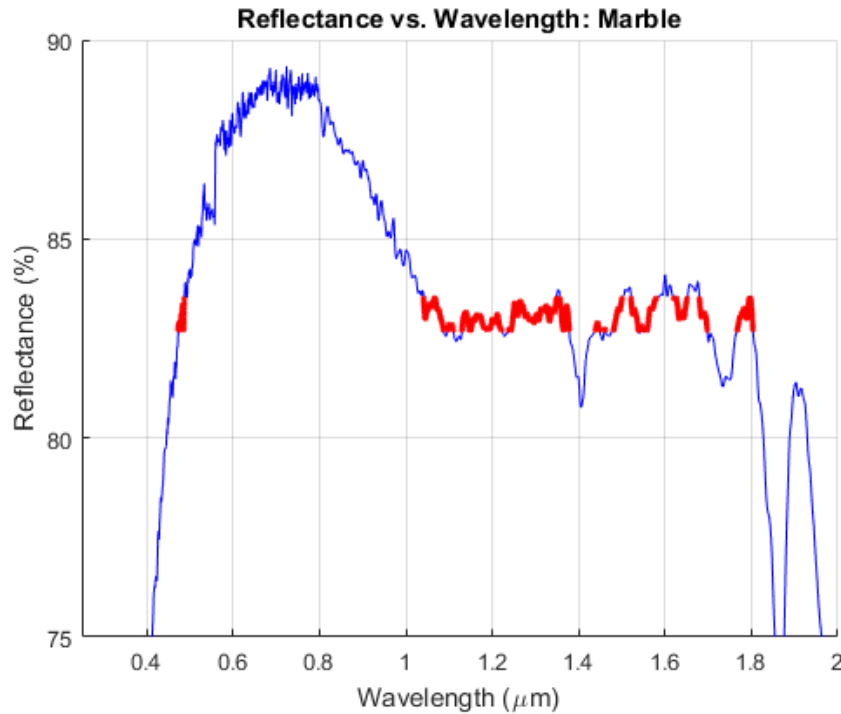


Figure 46: The distribution of reflectance (%) with wavelength. Many materials, such as marble, have multiple wavelengths over a wide range with similar reflectance. The red highlights show wavelengths that would produce less than 5% error in the overall DIAL result for a given reflectance.

A comparison of all wavelength combinations for all of the materials in the database is presented in Figure 47. A value of 100% (yellow) means that 100% of the materials at the specified offline and online combination will provide 5% or less error due to assuming equal reflectance. For these plots, the online wavelengths were analyzed with 1 nm increments while the offline was analyzed with 0.5 nm increments.

It can be seen that although the diagonal yellow value is narrow, some online wavelengths can tolerate a wider range of offline values. The regions where abrupt changes in the color gradient are observed, such as near 1.9 μm , are primarily due to the numbers of materials with values beginning/ending at those points. It is worth noting that the percentage is never 0%. There is always at least one material in the database where a given online and offline wavelength combination could work with it as a reflective target.

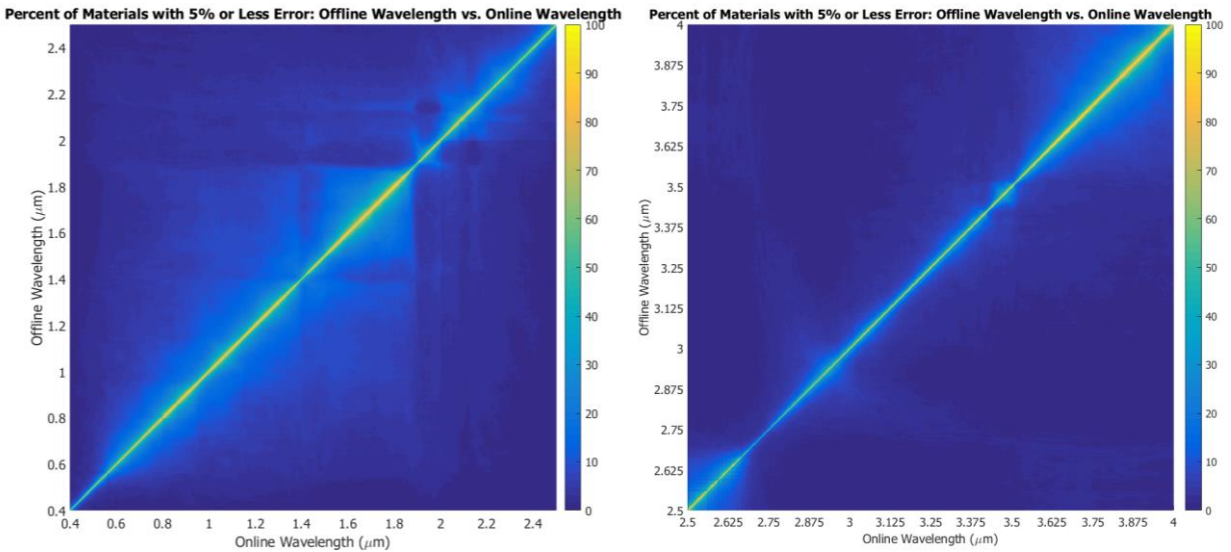


Figure 47: Complete comparison of all possible offline and online wavelength combinations reflectance that are in the database

A contour plot version of the left-side plot of Figure 47 is shown in Figure 48 to provide more visible demarcation. The contours are divided into 10% thresholds over the range of 10% to 100%. The region within a contour means that the offline and online combinations therein will provide 5% or less error due

to assuming equal reflectance. The outermost contour is the 10% line. Two insets demonstrate some of the fine grained structure within the dense region of the plot.

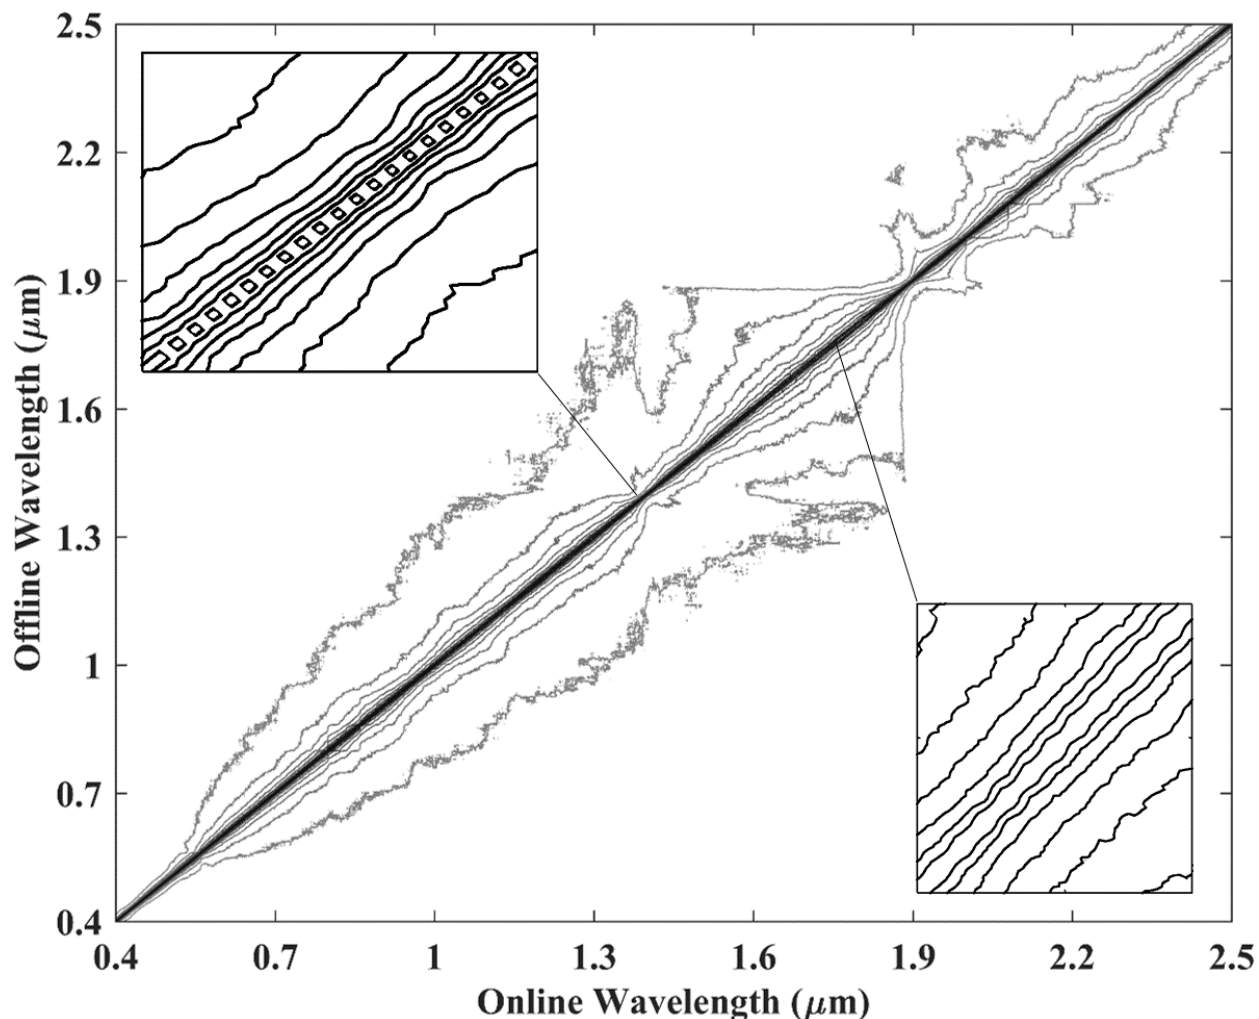


Figure 48: Contour plot of the percentage of materials within the JPL library that introduce less than 5% error in a DIAL signal. Contours are in 10% increments, starting with 10+% at the outermost contour and ending at 100% in the inner most regions. Insets demonstrate some of the fine-grained structure in the dense, central region.

While Figure 47 shows a texture of material reflectivities and the complex nature of photon and molecule interactions, it is somewhat impractical in practice. In an effort to create plots useful for instrument designers, a tighter range of offline wavelengths, specifically 0 to 6 nm around the online wavelengths, are presented next. As described in the methodology section, the values of 90%, 95%, and 99% mean that, for instance, 90% of the materials in the database meet or improve on the 5% error

threshold while 10% of the materials would induce errors greater than 5% of the overall signal. Because the allowable ranges are not symmetric about a given online wavelength, the minimum distance is used from either direction. This simplification sacrifices a mean absolute differential of 0.12 ± 0.21 nm while allowing for more detail in these already large plots. The simplification means that the plot is conservative by this amount, so, if a particular application is on the cusp of acceptability, it may be that analysis expands the envelope.

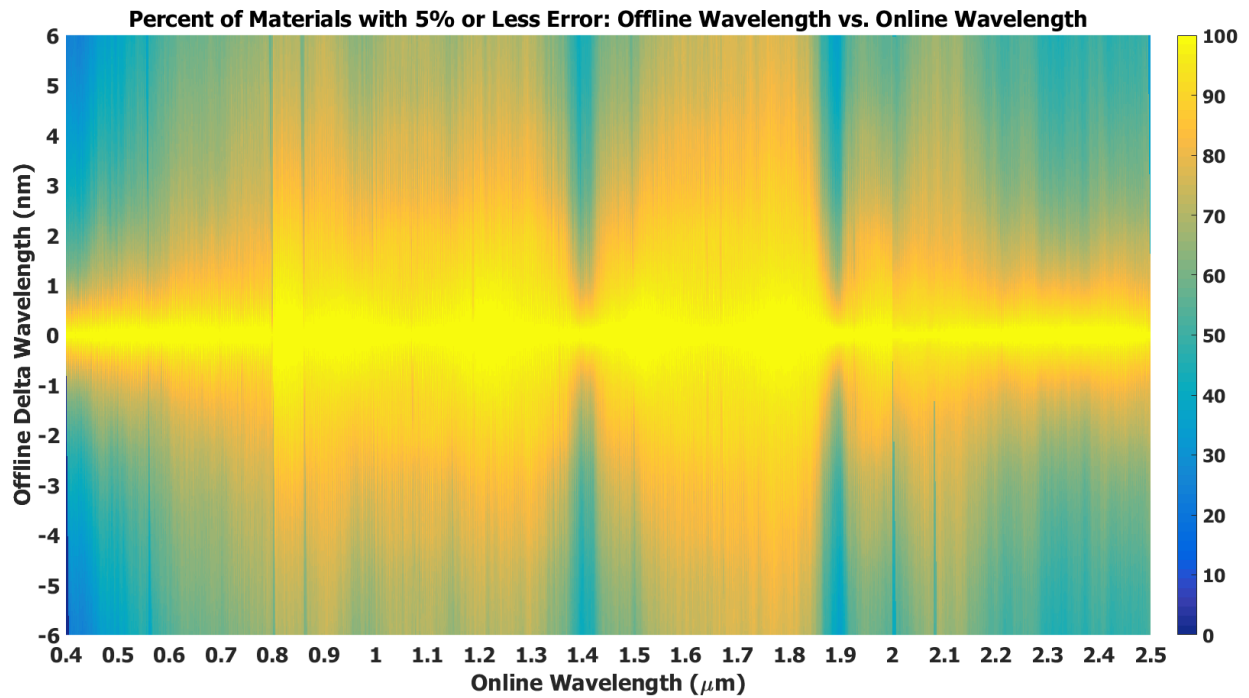


Figure 49: All the reflectance materials in the database analyzed from 400 nm to 2,500 nm with an offline range of ± 6 nm.

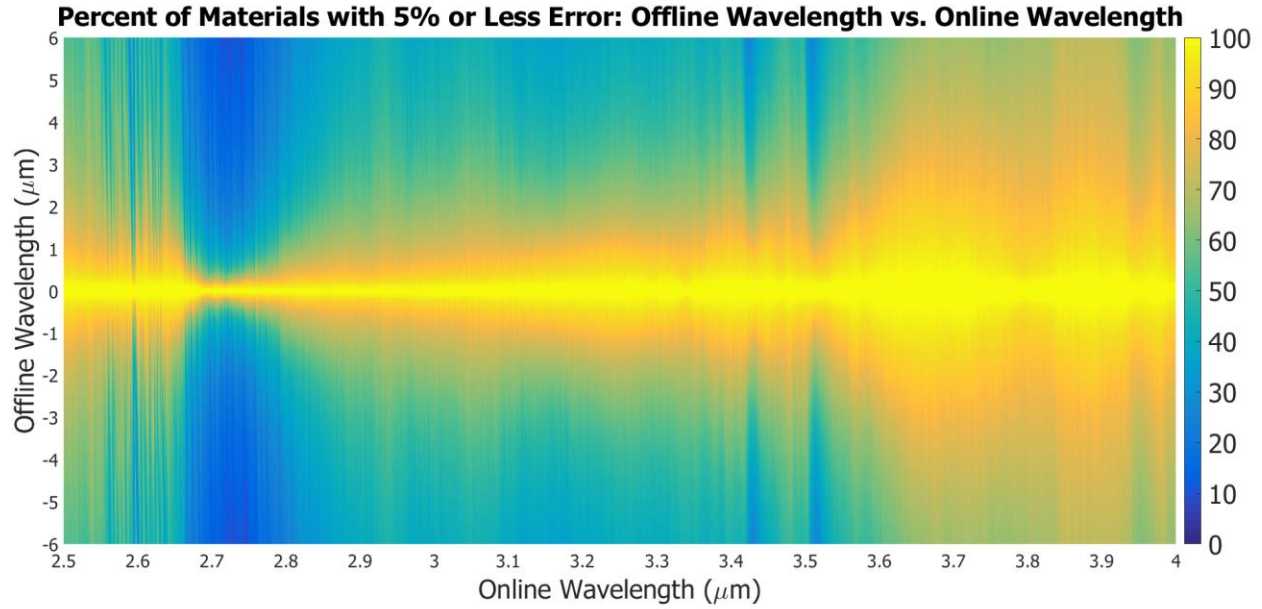


Figure 50: All the reflectance materials in the database analyzed from 2.5 μm to 4.0 μm with an offline range of $\pm 6 \text{ nm}$.

Setting specific thresholds for the percentage of materials that meet the criteria results in Figure 51 and Figure 52. This view may be more straightforward to interpret the design space for a given wavelength.

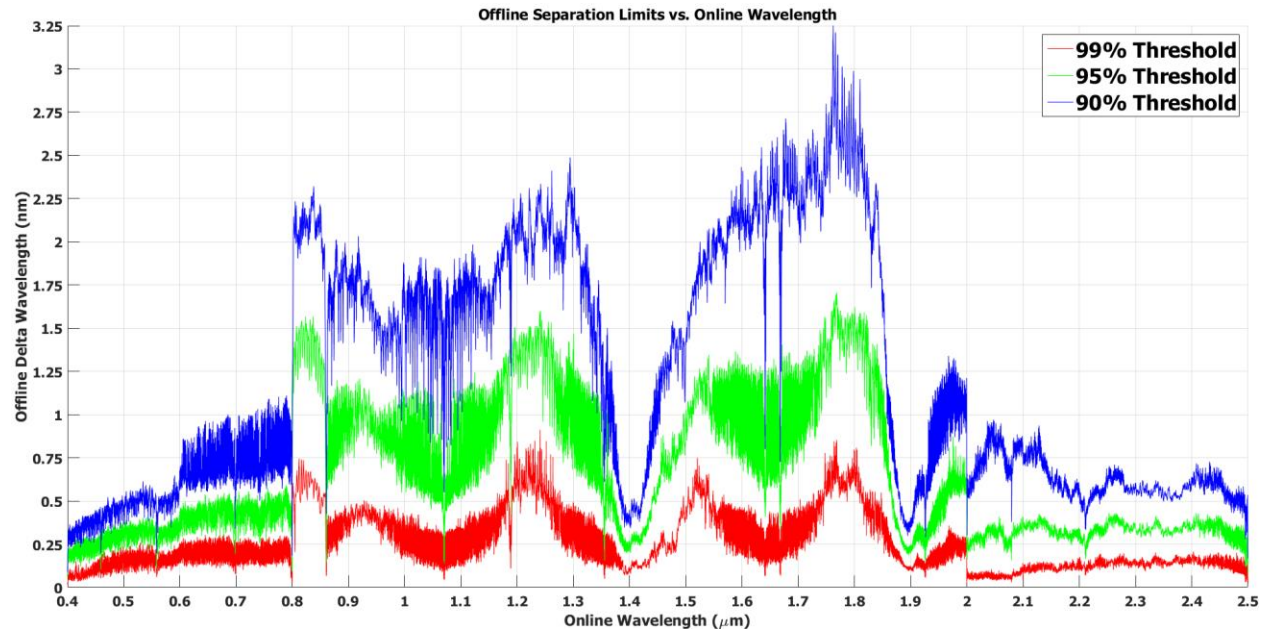


Figure 51: Absolute value of the offline wavelength separation for various thresholds for all materials from 0.4 μm to 2.5 μm with 2,287 sample sets.

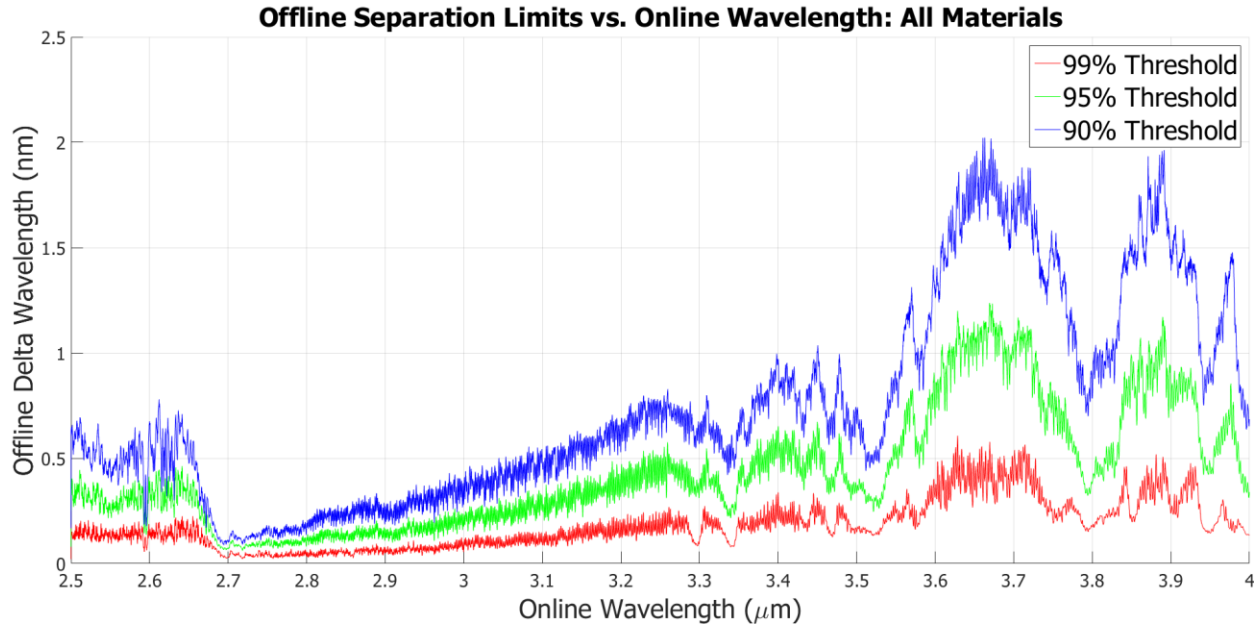


Figure 52: Absolute value of the offline wavelength separation for various thresholds for all materials from 2.5 μm to 4.0 μm with 2,287 sample sets.

It is seen that the results are non-uniform. In particular, the results above 2,000 nm in Figure 51 plummet. The source of this behavior is found in Figure 45, where the number of database materials above 2 μm jumps considerably. Many of the materials in this range have very low reflectivities. For instance, Sherwin-William “Flat black latex enamel spray paint” begins at 2 μm with 5.04% reflectivity and rarely exceeds 7% reflectivity through 14 μm . Low reflectivity materials are challenging in this percentage-based success criteria as a reflectance change from 5% to 4% is a larger percentage change than 90% to 89%.

Next, the same type of plot is presented for different types of materials in Figures 53 through 60. The types investigated are manmade, rock, minerals, and soils. The manmade class has had the “Reflectance Targets”, such as mirrors, removed from consideration as it’s not relevant to most applications. The soils type had the lunar regolith class removed for the same reason. The water and “frost, snow, and ice” types were combined. The vegetation type only has 4 sample sets and the stony meteorites are considered a fringe case so both are ignored here.

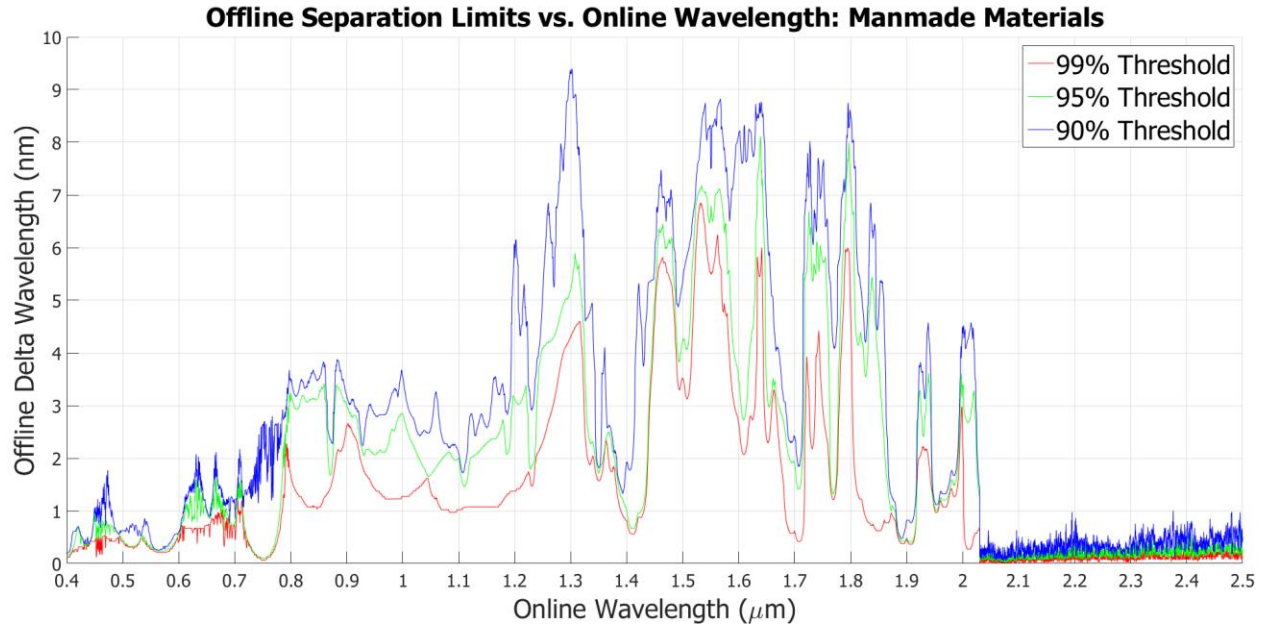


Figure 53: Absolute value of the offline wavelength separation for various thresholds for the manmade materials, with 56 sample sets.

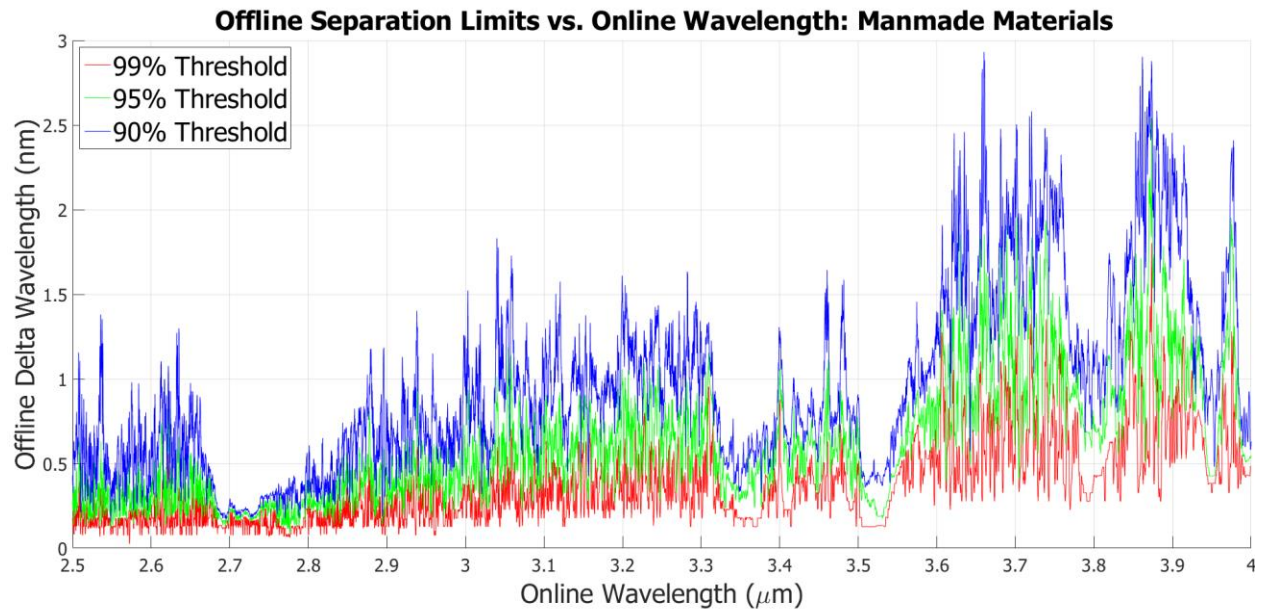


Figure 54: Absolute value of the offline wavelength separation for various thresholds for the manmade materials, with 56 sample sets.

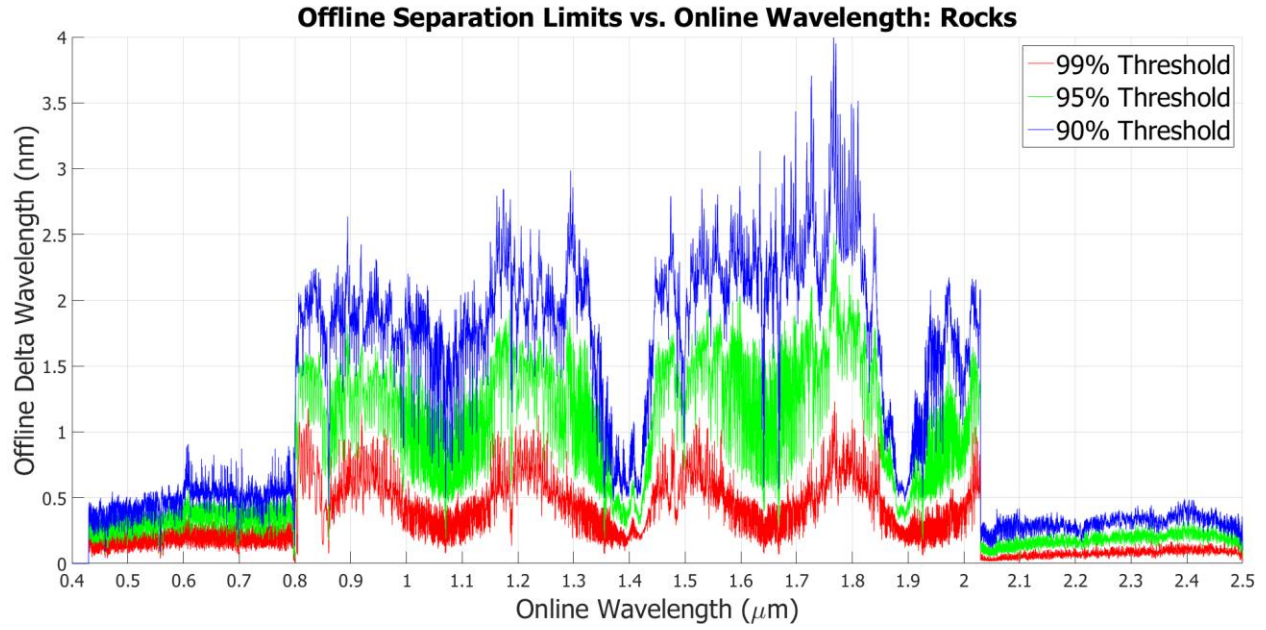


Figure 55: Absolute value of the offline wavelength separation for various thresholds for the rock type, with 473 sample sets.

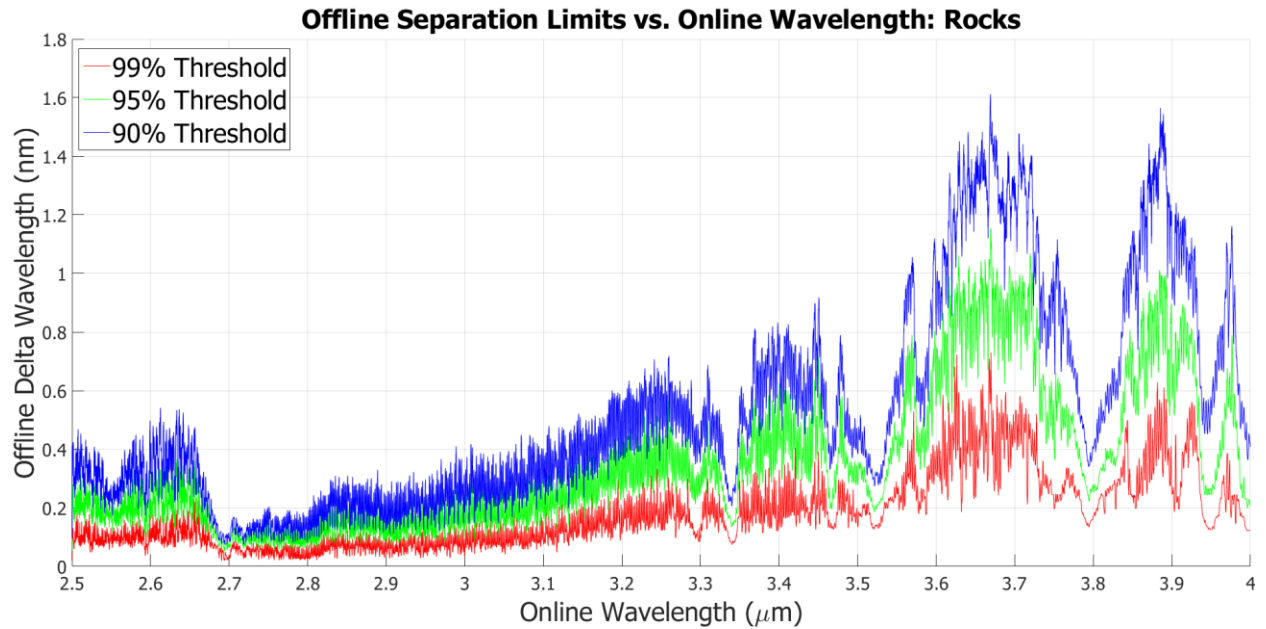


Figure 56: Absolute value of the offline wavelength separation for various thresholds for the rock type, with 473 sample sets.

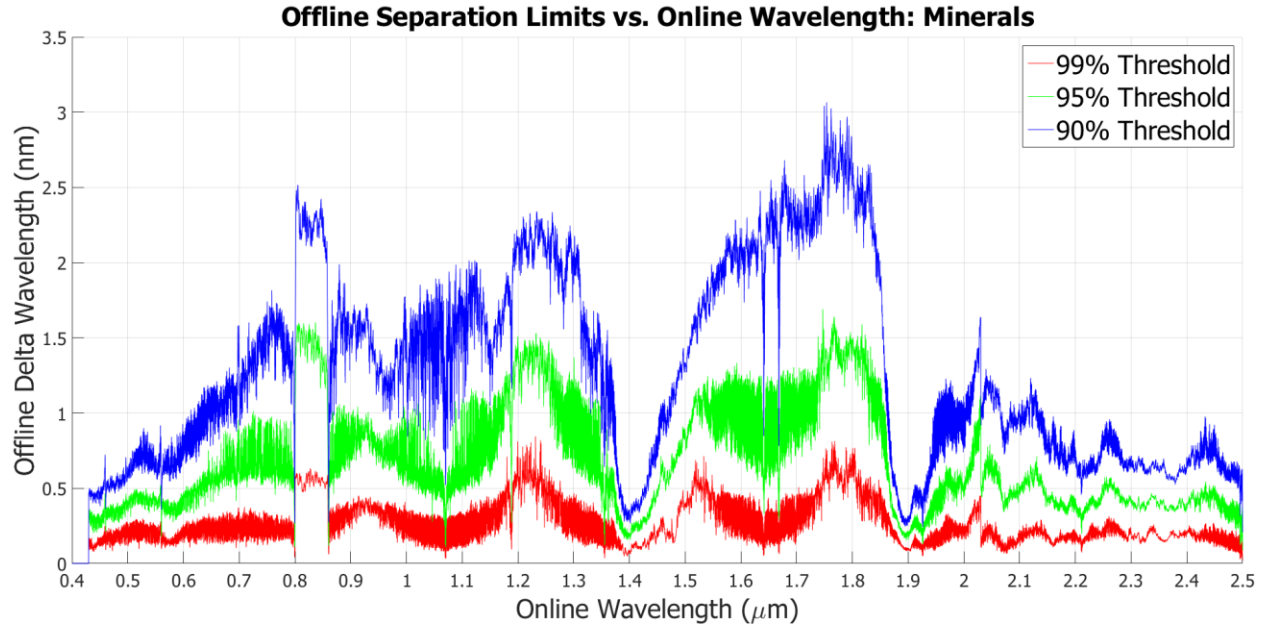


Figure 57: Absolute value of the offline wavelength separation for various thresholds for the minerals type, with 1,616 sample sets.

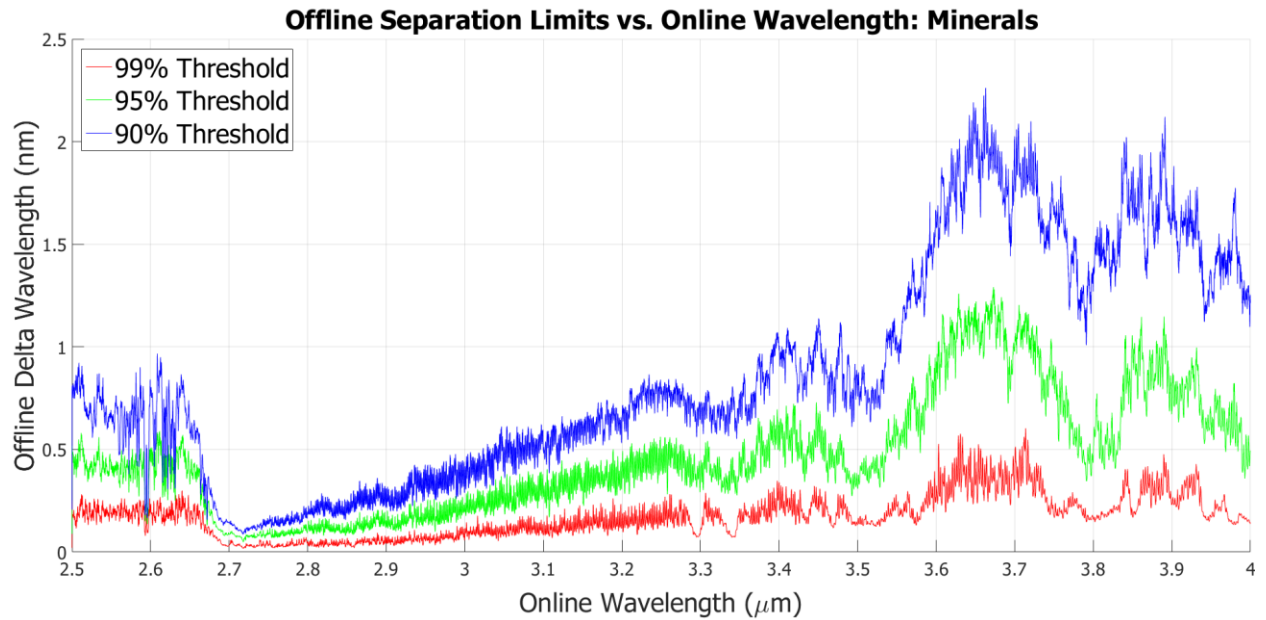


Figure 58: Absolute value of the offline wavelength separation for various thresholds for the minerals type, with 1,616 sample sets.

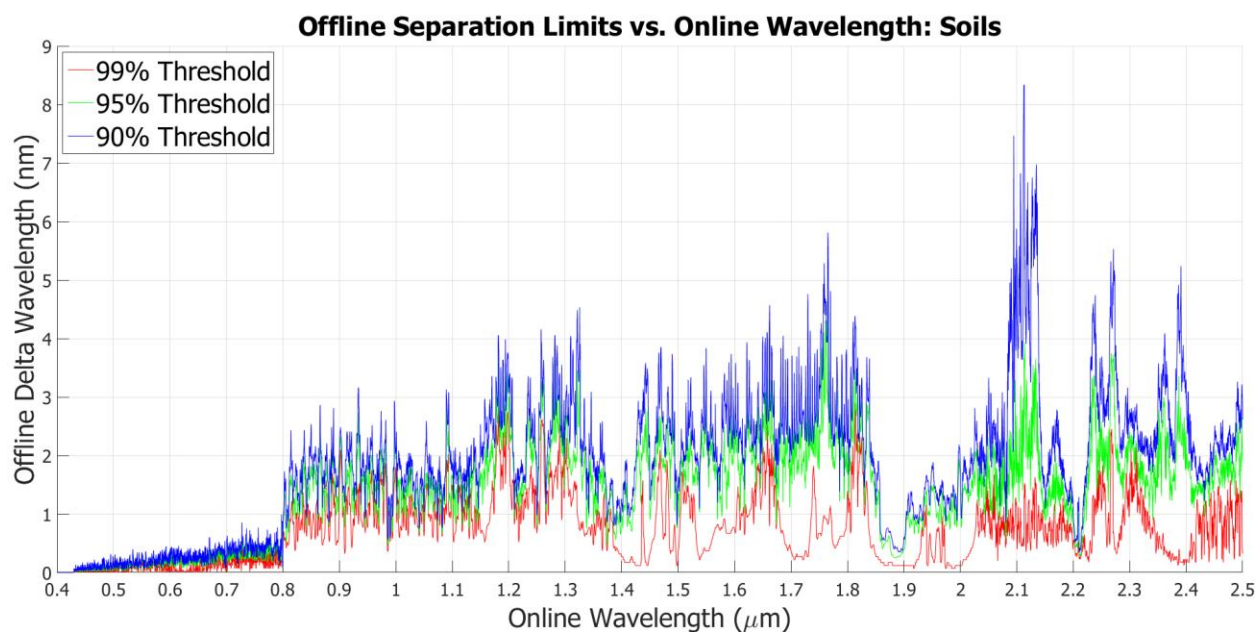


Figure 59: Absolute value of the offline wavelength separation for various thresholds for the soils type, with 52 sample sets.

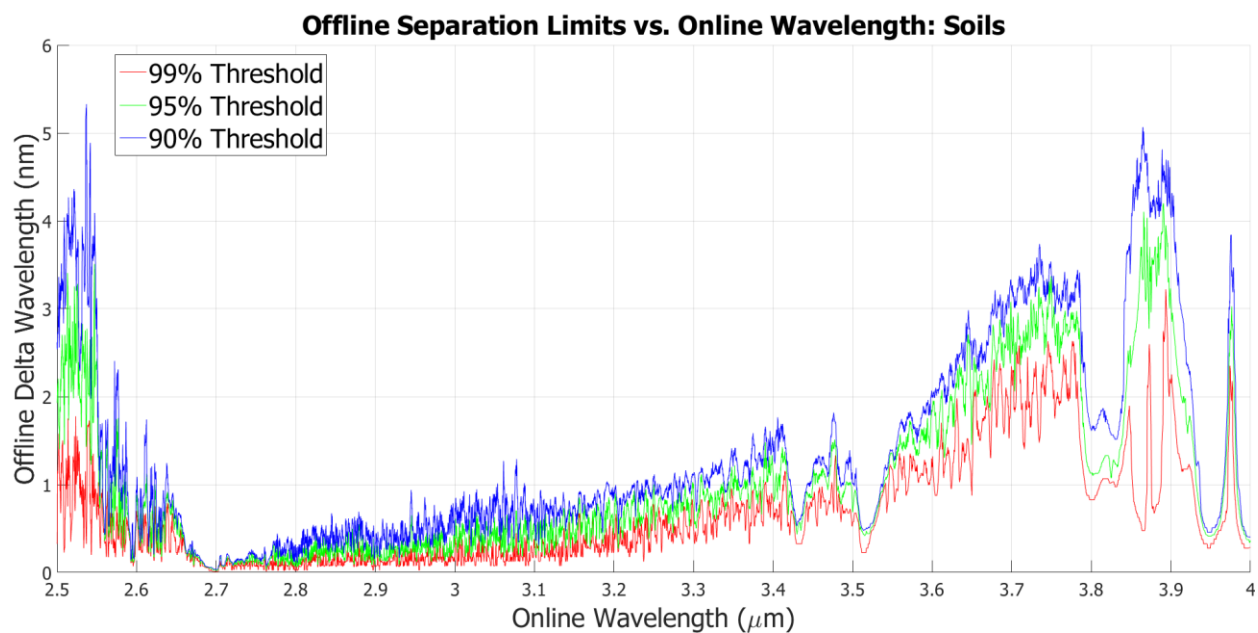


Figure 60: Absolute value of the offline wavelength separation for various thresholds for the soils type, with 52 sample sets.

For the water-centric categories, although there are various forms of water present in the database, it is still a single molecule; the 99% threshold is the only threshold provided as the other thresholds are

equivalent. It is seen that, if an instrument is designed to primarily interact with water, snow, or ice, there can be broad differences in the wavelengths.

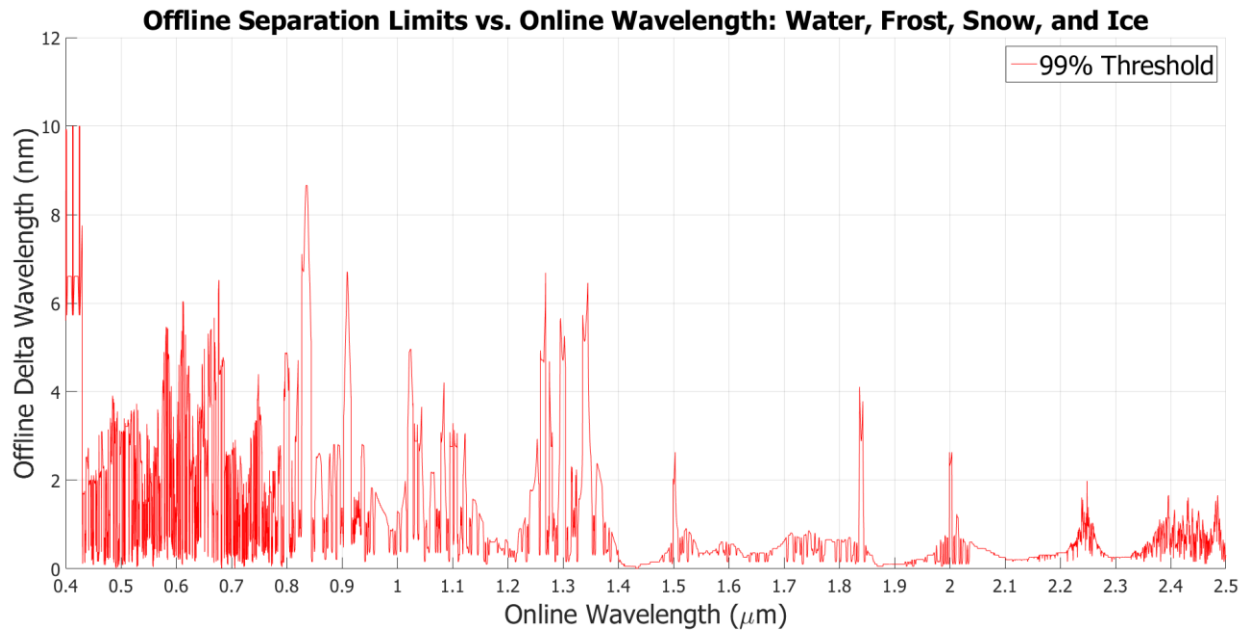


Figure 61: Absolute value of the offline wavelength separation for various thresholds for the water types, with 9 sample sets. The analysis only went to 10 nm of separation, cutting off some of the result around 400 nm.

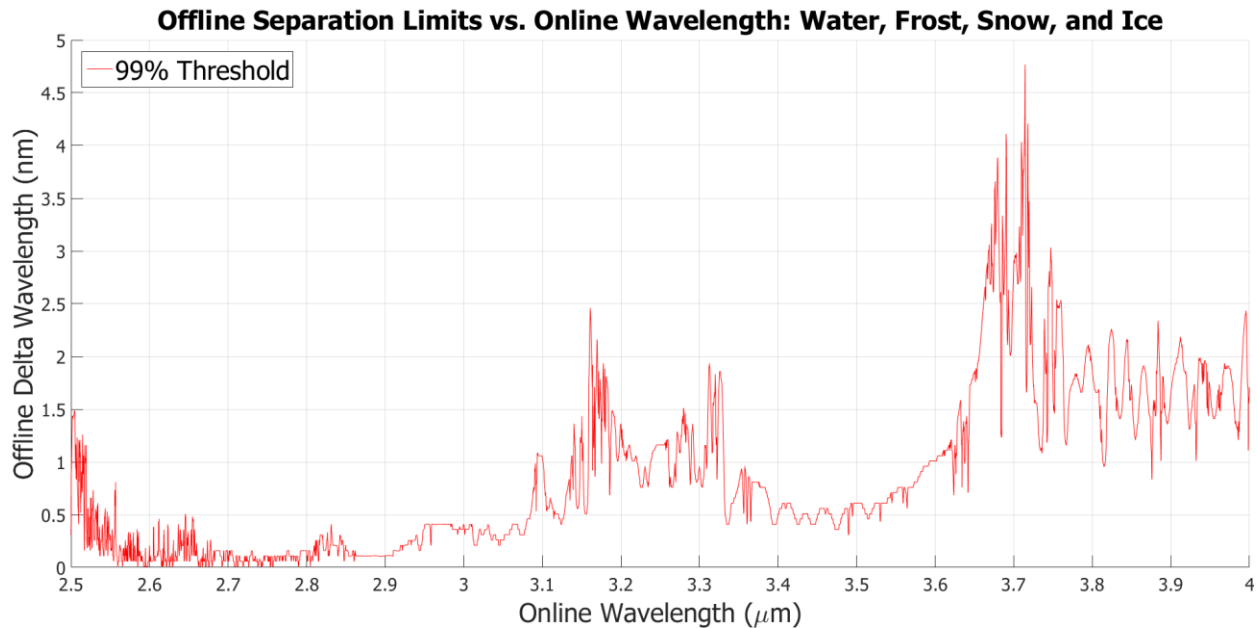


Figure 62: Absolute value of the offline wavelength separation for various thresholds for the water types, with 9 sample sets.

Throughout the paper, an offline-to-online ratio of 1.1 has been used with an allowable error contribution due to wavelength-based reflectivity variations of 5%. To demonstrate the sensitivity of the analyses to these values, Figure 63 plots the ratios and percentages for a wavelength of 1645.55nm. The plot on the left represents the result when 95% of the materials must meet the threshold while the plot on the right does the same for 99% of the materials. It is seen, for instance, that asking for 99% of the materials to contribute 1% or less error is virtually impossible with only about 0.06 nm of allowed offline separation within the threshold. On the other hand, larger wavelength separations are allowed if a system has a higher than the 1.1 offline to online signal power ratio used for the plots in this paper. Overall, the conclusion is that the error threshold relationship is not linear and a specific application will benefit from customized analysis. Also, instruments seeking no error from reflectivity variations either need very tight offline to online separation or a design with an eye toward significant online absorption. It is likely that a system will encounter at least a couple of percent error even with tight design requirements.

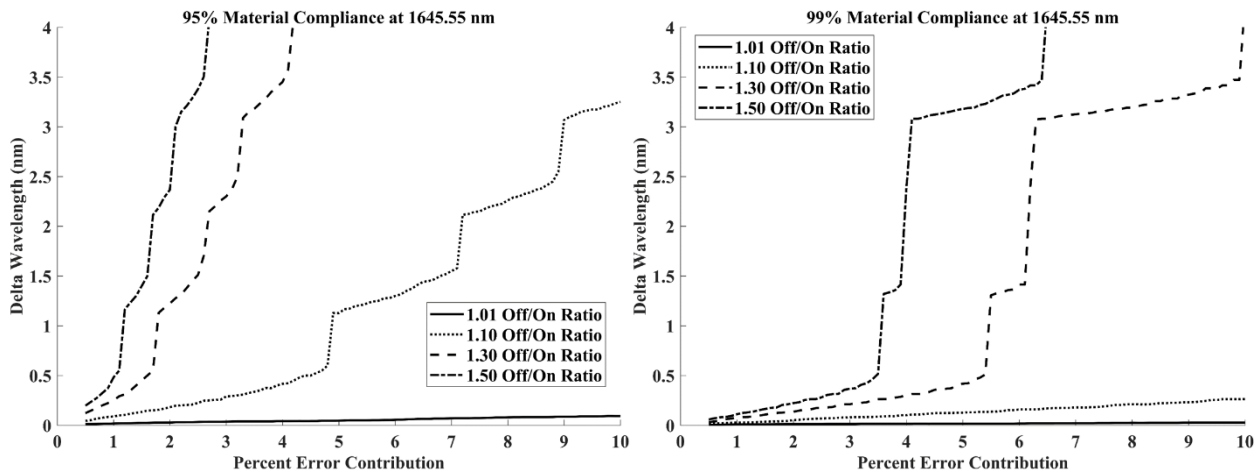


Figure 63: Sensitivity plots of the allowable percent error contributed by reflectivity variations with each line representing the ratio of offline to online received power. The nominal case used throughout the paper is 5% error with a ratio of 1.10.

5.5 Example Applications

Two example preliminary designs are provided to demonstrate how the plots above may be used to influence an initial wavelength choice. The first case studied is for an airborne instrument seeking to

detect methane in the atmosphere. The instrument designer begins by noticing that there are three main clusters of methane absorption lines below 4 μm . The peak absorptivities from these three regions are pulled and compared with the reflectivity data. The result is shown in Figure 64. The plots show that the first and third cases have strong potential with order of magnitude drops in absorptivity within the 99% regions and further reductions in absorptivity within the 95% regions. However, the 2.370 μm region is noisy, and it is not until the 90% region that a decent offline wavelength becomes available, according to reflectivity, even though it is less than 0.5 nm away.

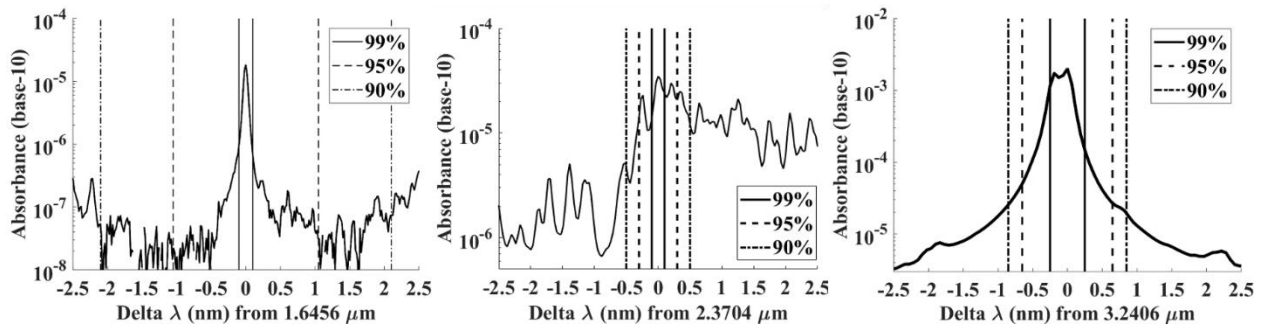


Figure 64: The center of the three plots represents the peak values from the three main clusters of absorptive online wavelengths for methane below 4 μm . The vertical bars span regions where 90%, 95%, and 99% of all the materials in the database contribute 5% error or less.

The second, more complex, application is a DIAL instrument seeking to measure both CO₂ and H₂O concentrations using a shared offline wavelength. In this example scenario, perhaps a low altitude unmanned aerial system is flying over arid deserts or polar zones (where water concentration is low) seeking to measure two of the most important greenhouse gases. The instrument designer starts by scanning the PNNL spectral absorption library and finds that there are two promising regions where the peaks of absorbance values of the two gases are near each other and there is potentially a reasonable nearby offline to warrant further consideration: 1.956 and 2.708 μm . In the first case, there are many lines within a 2-nm region with a nearby order of magnitude drop in absorptivity, indicating a potentially good offline selection. In the second case, there is only one peak for each gas, but they are within 0.9 nm of each other and there is more than an order of magnitude drop in absorptivity in their valley.

One of the next steps would be to then consider the potential impact that ground reflectivity would have on the results. The 1.956- μm wavelength region is shown in Figure 65. There, the CO₂ and H₂O absorption lines are shown with the vertical bars representing the percent of materials that would introduce minimal error due to reflectivity errors. It is seen in this example that CO₂ has two potential peaks and H₂O has one peak within the range where the vast majority of materials reflect with minimal error. If the designer feels that the 95% region is acceptable, then two more peaks of H₂O are available.

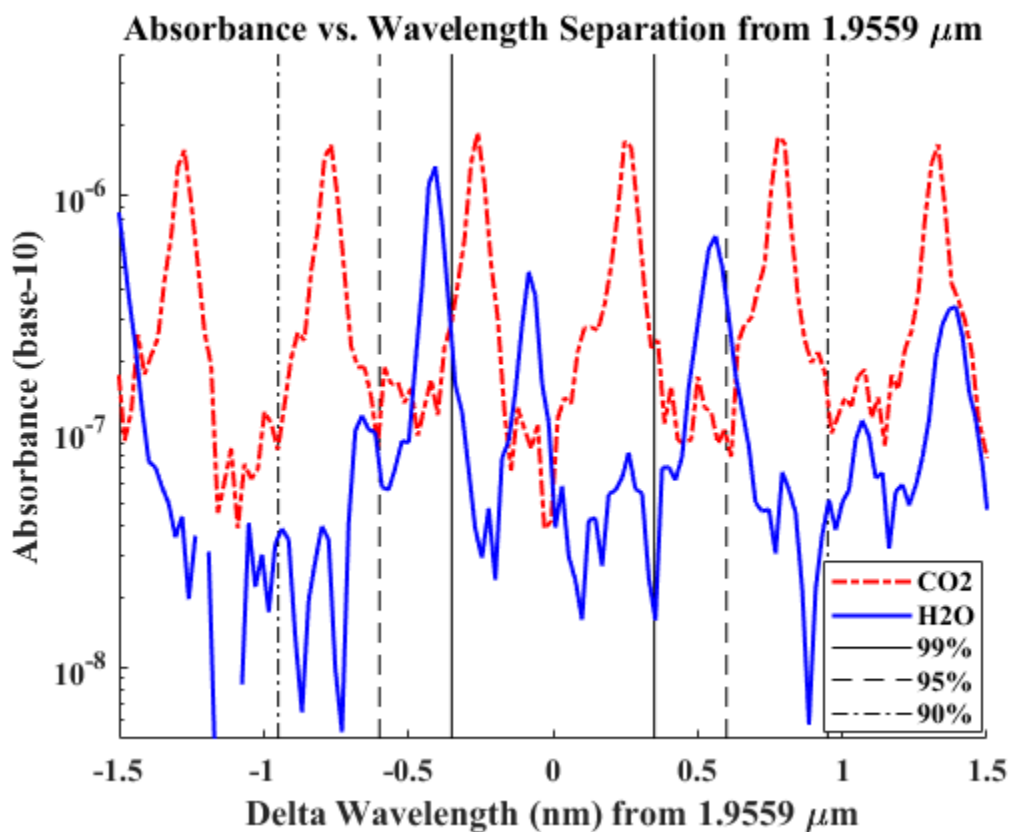


Figure 65: Plot of the CO₂ and H₂O absorption lines for the first wavelength region considered for the UAV instrument.

In the second case, the 2.708- μm region does not fare so well. It was seen in Figure 52 that the 2.7- μm region is one of the most difficult in terms of reflectivity and that impact is demonstrated in Figure 66: these two strongly absorbing features may not be suitable for airborne instruments as a significant percentage of materials impose meaningful error into the signals at this wavelength separation.

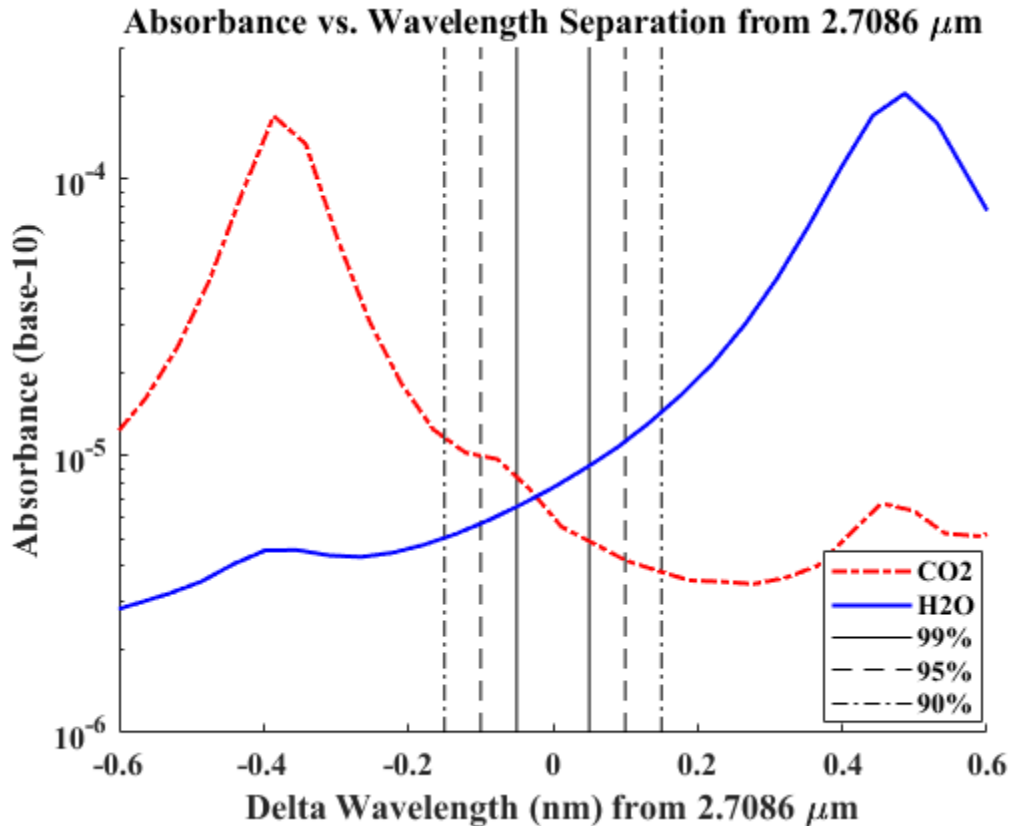


Figure 66: Plot of the CO₂ and H₂O absorption lines for the second wavelength region considered for the UAV instrument.

For the above cases, the central theme is that reflectivity can play a strong role in wavelength selection. The empirical analysis available with spectral libraries such as JPL's ASTER library means that instrument designers can make informed decisions instead of relying on general rules. The precomputed plots in Figures 53-60 can also be used as a first step to more detailed analysis.

5.6 Summary

Empirical analyses of thousands of materials have shown that, depending on their application, DIAL instrument developers utilizing unknown terrain to reflect signals may have flexibility in choosing their online and offline wavelengths. In many scenarios, several nanometers of separation may be achievable. In other cases, otherwise desirable DIAL wavelengths are shown to face challenges due to the potential for introduced errors from reflectivity mismatches. Although instrument designers are likely to want to

compute specific results for their applications, the plots in Figures 53-60 should prove useful as a first step to understanding the potential reflectivity benefits or challenges at a given wavelength. One area that would benefit greatly from future work would be extending the spectral library to include more vegetation samples. This is recognized as a challenge as vegetation reflectivity is likely to change with the time of year and moisture content. However, this type of data would be highly valuable to instrument developers that depend on vegetated ground reflectivity for their return signals. Along the same lines, greater numbers of manmade materials would be useful. In spite of these desired improvements, the authors recognize the great utility of the ASTER spectral library and are grateful for the efforts of its creators.

5.7 Application to Methane at 1.645 μm

The analysis code and database from the paper effort were used to investigate the 1645.37 and 1645.55 nm wavelengths for the dissertation effort. As with the study, it was assumed that the offline to online received power was 1.1. The percent error due to reflectivity was then calculated for all of the materials. The results were sorted in descending error and with the signs being kept to illustrate the direction of the error. (Negative percent error means that the offline reflectance is higher than the online reflectance, positive error means that the online reflectance is higher than the offline reflectance.) The result is displayed in Figure 67.

It is seen that there are about 100 materials that contribute more than 1% error to the overall result. The materials within the failing group are primarily crystals such as mica and pyrite, or composites of crystals such as granites and basalts. The molecules that make up the failing materials are primarily based on crystalline iron (Fe), Silicon (Si), and Carbon (C). None of the manmade materials, construction materials, water, or the vast majority of rocks and minerals introduce more than a fraction of a percent of error. The mean introduced error is -0.001% with a standard deviation of 1.06%.

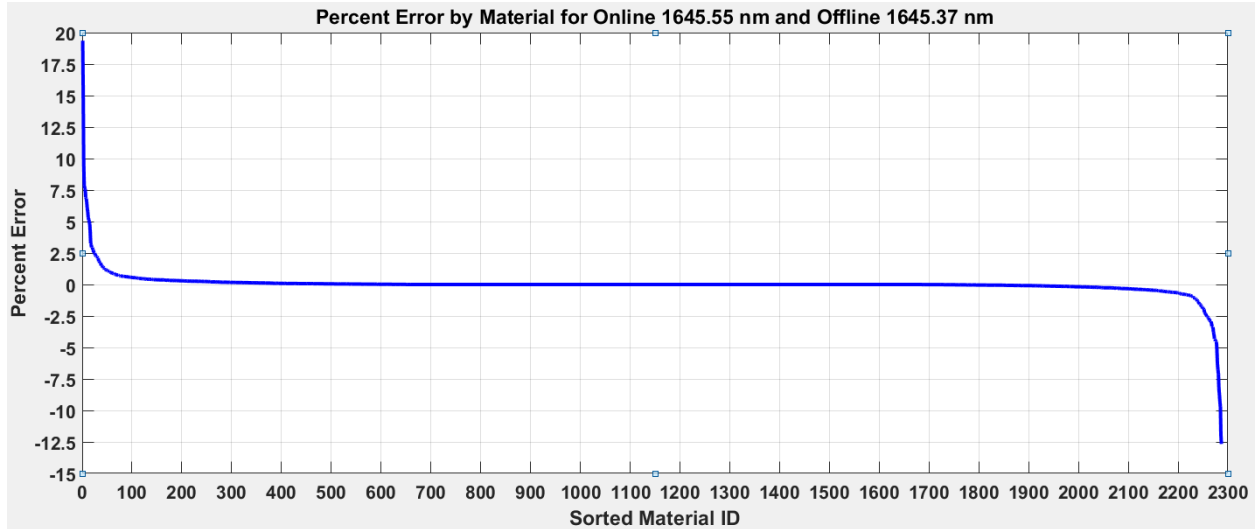


Figure 67: Plot of sorted percent error for all of the materials in the ASTER library for the online and offline wavelengths of 1645.55 and 1645.37 nm respectively, assuming that the offline to online received power has a ratio of 1.1.

The material with the highest error contribution is Tourmaline, a silicate crystal. A plot of its reflectivity is provided in Figure 68. The reason it can create 20% error is that it has very low reflectivity of less than 3%. When working with ratios, small variations of small numbers lead to large relative changes. The small reflectivity also means not much signal makes it back from the material to the instrument. The combination of these two things means that flying over materials like Tourmaline will create noise and error that could mean marginal performance.

Fortunately, for the materials in the ASTER library, there aren't too many of these types of materials in the targeted wavelength region and many of them, like Tourmaline, are relatively rare. For the purposes of this dissertation, a potential error of 1% is carried forward.

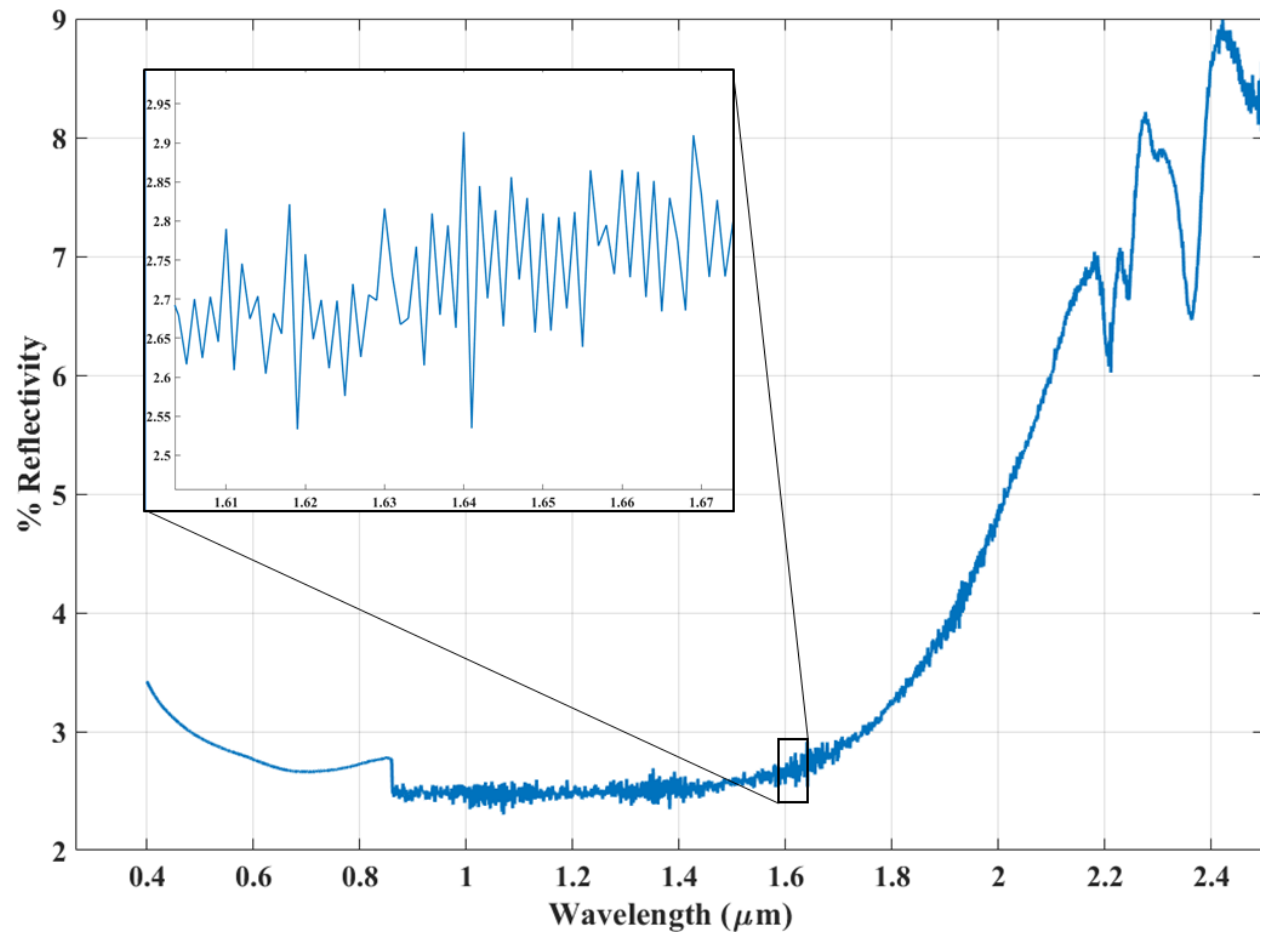


Figure 68: Reflectivity vs Wavelength for Tourmaline. The inset image is for roughly the 1.61 to 1.67 μm range.

Chapter 6: Hardware Planning and Considerations

This chapter moves the analysis process from plumes and external simulations to the hardware aspects of the instrument design. It starts in section 6.1 by introducing the two configurations considered in this dissertation. The first is a “traditional” design found in modern instruments such as Ball’s methane monitor. The second considers using filtering optics on the receive side to enable noisier, possibly cheaper lasers. Section 6.2 briefly introduces Volume Bragg Gratings as the filtering optic. Section 6.3 considers the pros and cons of choosing single mode or multimode fiber. Retail product datasheets are mined to establish practical rules of thumb as part of this process. Section 6.4 looks at beam pick-off issues. Both designs need a reliable way to measure a small amount of the light being sent out so the focus is on making the measurements of the two pulses as error free as possible. Section 6.5 is a small study that looks at the importance of minimizing optic count to maximize overall efficiency. A simple plot is provided to anchor discussions on optical efficiency later in the dissertation. Section 6.6 examines beam steering solutions for the fixed-wing UAV platform. A medium-sized, fast, and affordable beam steering mechanism for UAVs was not found commercially so a solution for a custom assembly was investigated. Section 6.7 looks at detectors, specifically investigating bandwidth considerations. It was found that rules of thumb offered by manufacturers are based on square wave pulses, while laser pulses are rounded. Fourier analysis suggests that there may be room for relaxation of the requirements and thus a more affordable option may be useable. Section 6.8 investigates pulse integration as a methodology for measuring the return energy. Using normalized, unitless methods, the timing of the integration windows and the resulting errors are plotted. Section 6.9 wraps up the technical discussion with a brief overview of the processing electronics. Although significant time and consideration was invested in this part of the analysis early on, budget limitations did not allow for pursuing this to demonstration levels. Finally, section 6.10 provides lessons learned from working with Ball’s instrument. Although details cannot be shared, many of the topics in this chapter were inspired from real-world problem solving with their production unit.

6.1 Two Instrument Concepts

Two instrument concepts are explored for feasibility in a UAV instrument. The first is referred to as a “Traditional Concept” and is shown in Figure 69. A recent example of this type of instrument design is Ball Aerospace’s methane detector [Bartholomew, 2017]. Note that the exact numbers and design may vary, but the general concept is the same: send the offline and online pulses at different times, minimize the optic count, and use multiple detectors. The benefits include relatively simple design, good optical efficiency, and the ability to fine tune the design of each component for its specific task. The primary drawback is that the design relies on three detectors, which means cost, instrument volume, weight, differences in noise levels, needing to buy and wire three detectors, have three ADC sampling channels, and the different pulses have different speckle profiles and different spatial placements on the ground.

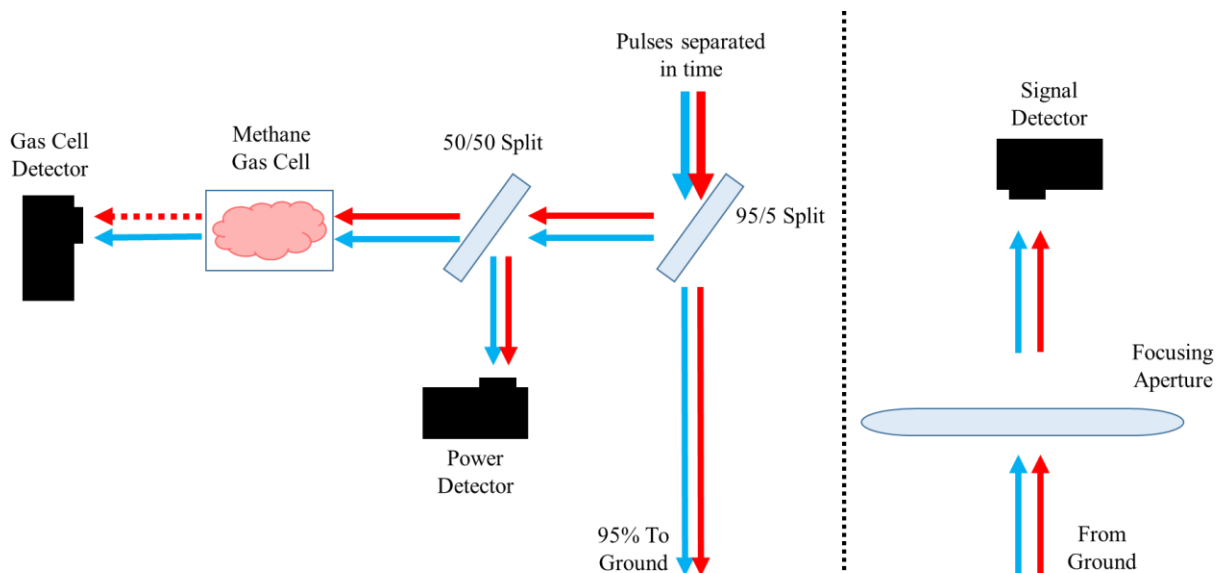


Figure 69: Simple schematic of a typical system that separate the online and offline pulses in time.

A second concept is provided in Figure 70 and is referred to as the “VBG Filtering Concept”. The general principle is to combine the wavelengths into a single pulse, preferably from a single laser source, and send all pulse paths to the same detector via path length extensions and spectral separation. The most visible benefit is the detector count reduction. As all signals are processed on the same detector, their

noise levels are more likely to be equivalent. The pulses themselves, if generated via a single laser source, also have the same speckle profile and spatial footprint on the ground, which eliminates two sources of error. Drawbacks include having additional optical interfaces (loss of efficiency), additional collimation and alignment challenges, extra volume needed for spectral separation, and extra cost for the spectral separation components.

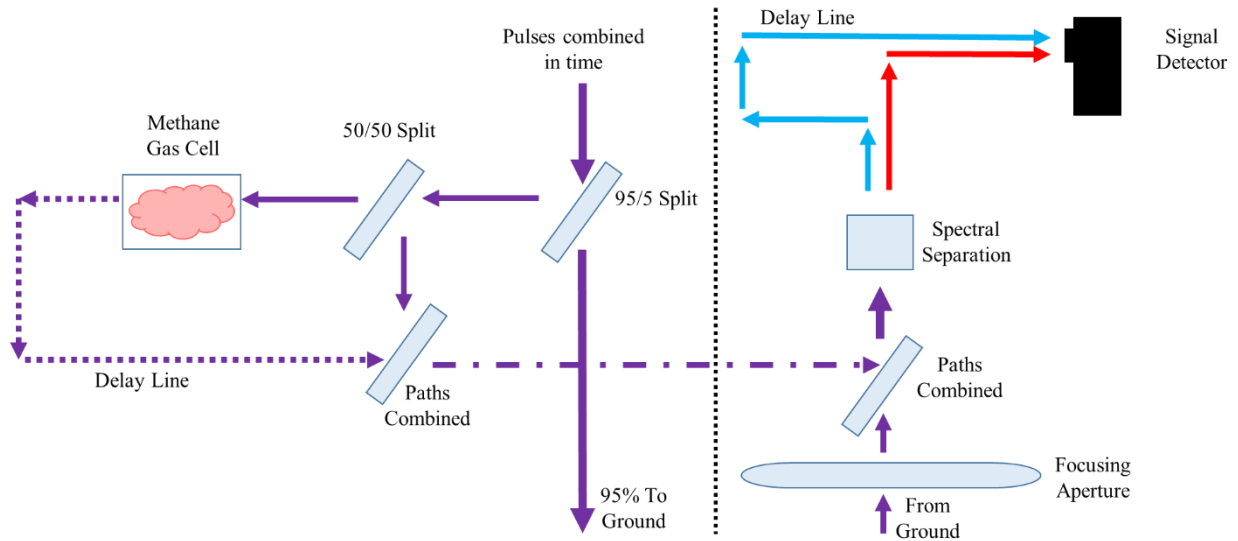


Figure 70: Simplified schematic of a system designed to process a single pulse with multiple wavelengths

The goal of analysis in the following sections is to quantify the subjective judgements of each design so that future decisions can be made from a more solid foundation. The next sections introduce key components.

6.2 Spectral Separation: Volume Bragg Grating

For this project, the spectral separation is being accomplished with a Volume Bragg Grating (VBG). A VBG is a solid piece of transparent material with the spectral separation gradient created internally. Its introduction to optics and lasers is relatively recent, with the first publication in 1999 [Efimov, 1999]. The technology has several advantages over traditional spectral filters, including the ability to operate with high power lasers (up to 10's of kilowatts/cm²), the ability to selectively reflect or transmit energy with a

bandwidth of 20 pm, survive temperatures up to 400 °C, and an ability to fine-tune the desired wavelengths through temperature control [OptiGrate, 2017]. In practice, it should be noted that the bandwidth selectivity is a function of wavelength, with narrower selectivity at lower wavelengths, and there is an upper wavelength limit of 2,700 nm.

The specific version of the technology investigated in this effort is the transmission grating, which deflects the wavelength of interest slightly as the other wavelengths are passed through [Ciapurin, 2006]. A pick-off mirror or other optical element is used to guide the deflected wavelength from the grating to its next destination and the remaining light can either be similarly guided or dumped. Figure 71 provides an illustration of the concept.

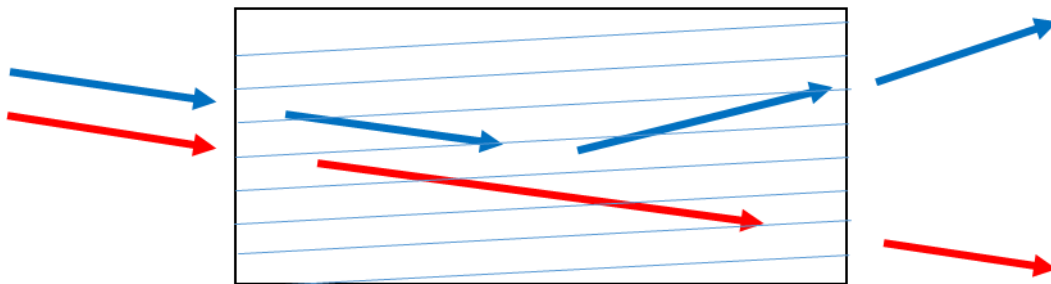


Figure 71: Simple diagram of a transmissive VBG redirecting a desired wavelength while other wavelengths pass through. See [Ciapurin, 2006] for more information.

The company transmitting VBG's are purchased from is OptiGrate. The company's technology team also wrote the founding papers (and most subsequent papers) on VBG's [Efimov, 1999, Ciapurin, 2006, Vorobiev, 2008]. They are an offshoot of the University of Central Florida and, although not the only VBG provider, they have the most experience as well as a discriminating patent portfolio within the United States.

Unfortunately, a condition of buying from them was that key technical details of our particular VBG's could not be shared. A few things can be shared, though. First, the spectral selectivity, besides being a function of wavelength, is also a function of length. Longer VBG's mean better selectivity. The practical result is that as the wavelengths increase, the VBG's length must also increase to compensate. The VBG's

used in this project are relatively long due to the 1,645 nm wavelength. Second, the gratings are anti-reflection coated with losses less than 1%. Third, the rotation orientation matters as the energy must interact with the gratings at the correct angle to be effective. In fact, one way to tune the diverted wavelength of interest is by slightly rotating the crystal. The angular selectivity is on the order of a milliradian.

6.3 Single Mode vs. Multimode Fiber

Single mode fibers are preferable over multimode fibers because 1) their single mode means speckle due to multiple modes is absent, 2) collimation is easier, again due to a single mode, and 3) they tend to be more affordable. In fact, just about the only reason to go with multimode fiber is that their larger diameters mean they can handle higher laser energies.

For instance, according to the fiber provider, ThorLabs, the “Practical Safe Value” for a continuous wave laser on unterminated silicon fiber is 250 kW/cm² and is 1 GW/cm² for a 10 ns pulsed laser [ThorLabs, 2017]. The diameter used in the single mode fiber analysis is the Mode Field Diameter (MFD). The MFD is defined as the diameter at which the beam power density falls to 1/e² of the peak and takes into account that some of the light traveling along the fiber also travels in the cladding. In a multimode fiber, the diameter used in the analysis is simply the diameter of the core.

Important to both fibers is that the MFD increases linearly as a function of wavelength, according to the Petermann diameter equation [Petermann, 1983]. This means that the effective area and thus damage threshold increase with the square of the wavelength, which is good news for the longer wavelength regions. A plot of the MFD increasing with wavelength for fiber optics sold by ThorLabs is shown in Figure 72. Plots demonstrating peak energy for the given damage thresholds are provided below in Figures 73, 74, and 75 for some of the fibers sold by ThorLabs as of June 2017.

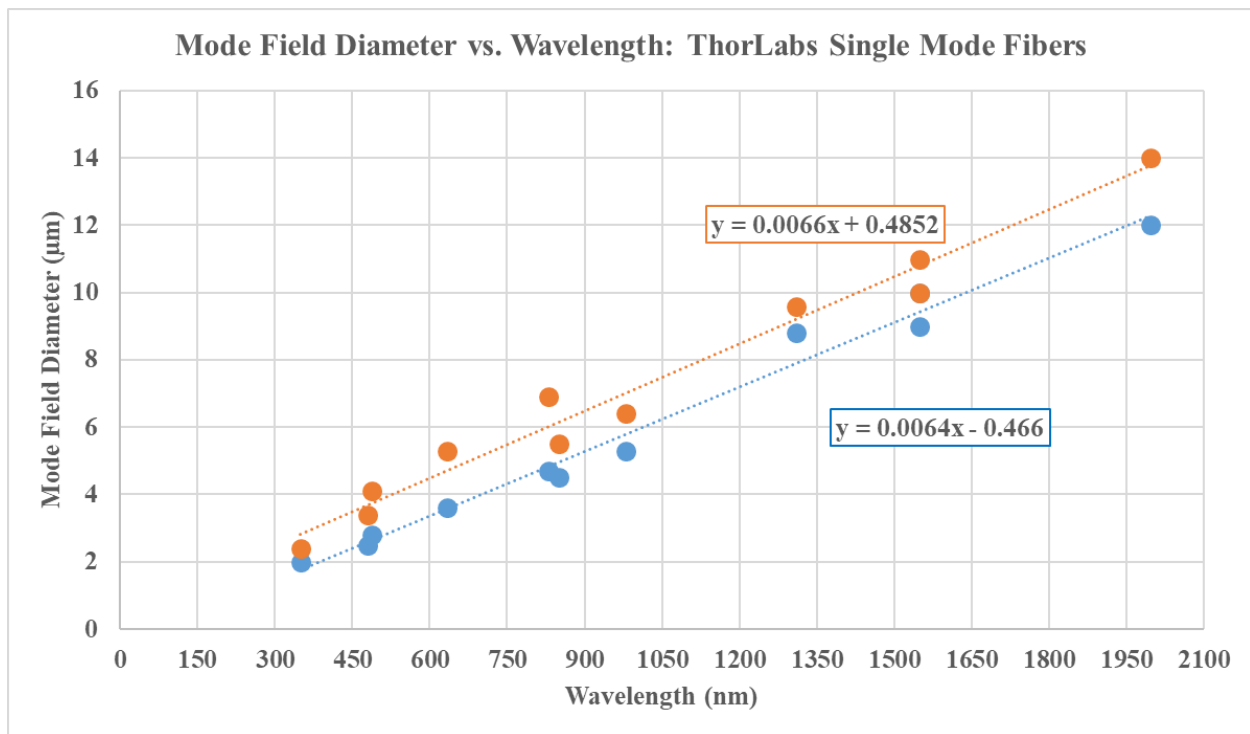


Figure 72: Plot of effective mode field diameters of single mode fibers sold by ThorLabs as of June, 2017. Curves represent max and min values in product properties. [ThorLabs, 2017]

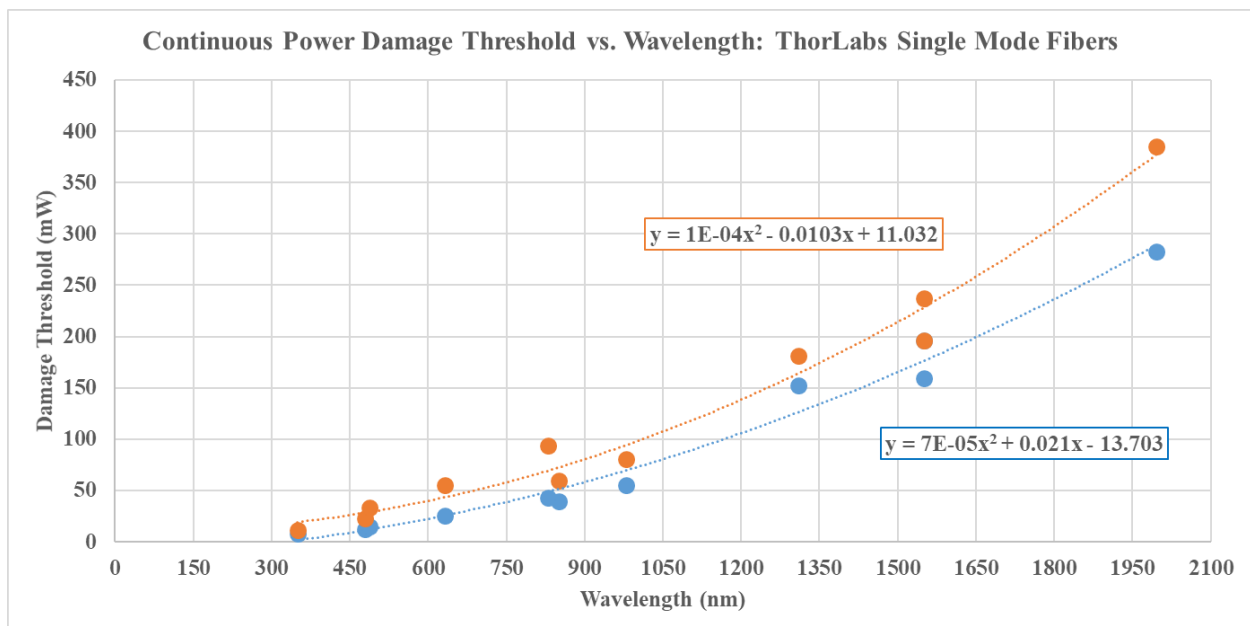


Figure 73: Plot of practical damage thresholds for ThorLabs' single mode fibers carrying continuous laser energy. Curves represent max and min values in product properties. [ThorLabs, 2017]

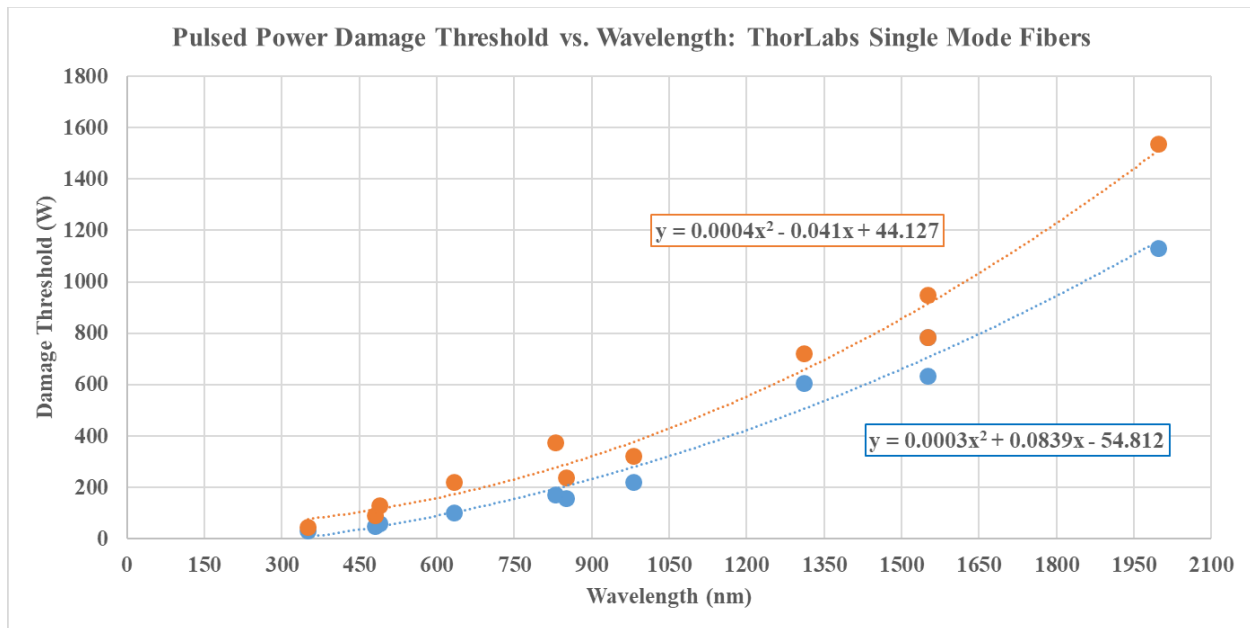


Figure 74: Plot of practical damage thresholds for ThorLabs' single mode fibers carrying 10 ns pulses. Curves represent max and min values in product properties. [ThorLabs, 2017]

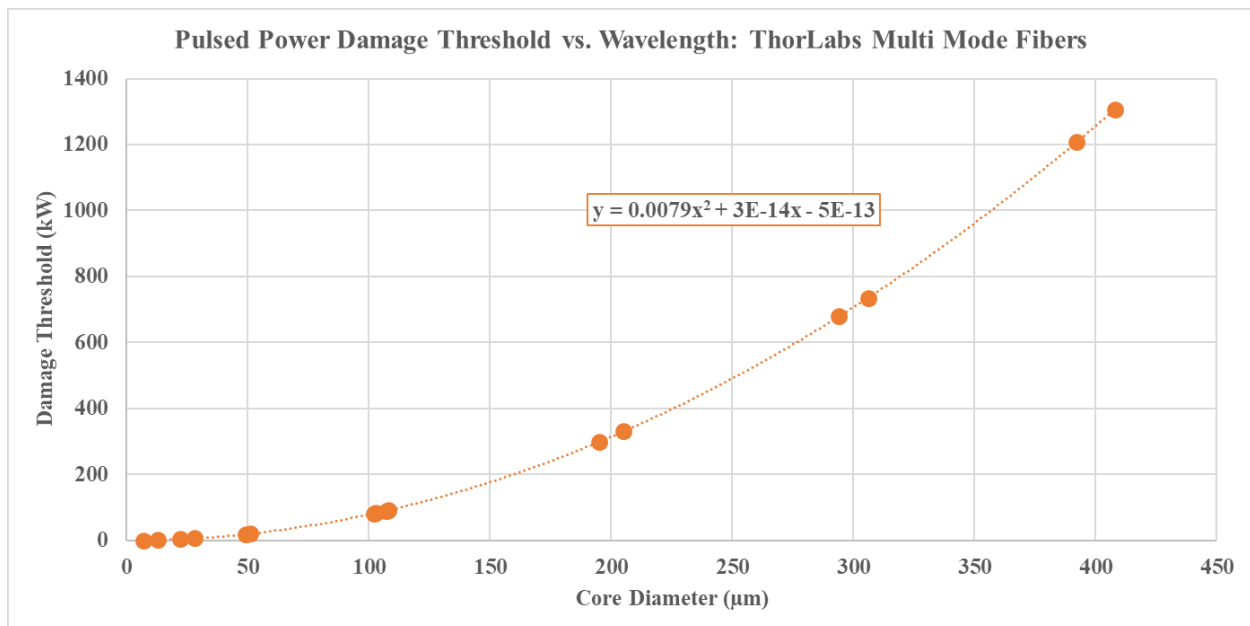


Figure 75: Plot of practical damage limit for ThorLabs' multi-mode fiber carrying 10 ns pulses. [ThorLabs 2017]

A rule of thumb noted by reading through vendor catalogs, such as RP-Photonics, ThorLabs and Newport, is that the damage threshold of a silica fiber carrying pulsed laser energy is typically related to

the square root of the pulse duration for pulses above 10 ps in length [Du, 1994, Stuart, 1995, RP-Photonics, 2017]. This means, for example, that a pulse that is 100 ns has a damage threshold about 10 times higher than the damage threshold of a 1 ns pulse. As vendors are quick to point out, though, these are just rules of thumb. One publication showed better than this rule of thumb with a 3x damage threshold improvement going from 8 ps to 14 ps for a silica fiber [Smith, 2008].

Unfortunately, ThorLabs is the only vendor found that published their damage thresholds. Most vendors point to articles with wavelength and pulse duration specific studies and the general rules of thumb for extending to your wavelength/duration of interest. All vendors encourage you to do your own testing. Therefore, for this study, the ThorLabs data is considered the general guideline and any improvement to damage thresholds due to duration differences are based on the square root of the duration, as long as the duration does not go below 10 ps.

One more point of interest on the topic of damage is that connectorized fibers have lower thresholds due to the epoxy holding the fiber in place. Most commercial, connectorized fibers use an optical epoxy which has good index of refraction properties, but poor thermal properties. Specific limits are hard to come by. ThorLabs gives a power limit of about 600 mW at 1,550 nm, but does not say if it's for a continuous wave or pulsed system. They do offer fibers with an air gap between the fiber and ferrule to reduce the amount of epoxy that can be burned up. However, the guidance again from manufacturers on this is to buy samples that you think are close to their general numbers and then test for acceptability.

6.4 Beam Pickoff

An important part of both design concepts is picking off a little bit of light for analysis. The first methodology covered is the use of beam splitters. This seems straightforward since the name of the optic is exactly what is needed. However, there is a difficulty in that the performance almost always has some polarization dependence. This can be problematic for DIAL systems because the offline and online wavelengths may have polarization states from upstream processes. A simple diagram in Figure 76

demonstrates how polarization can be used to combine beams from two sources so that they travel the same optic path with a beam splitter used in reverse.

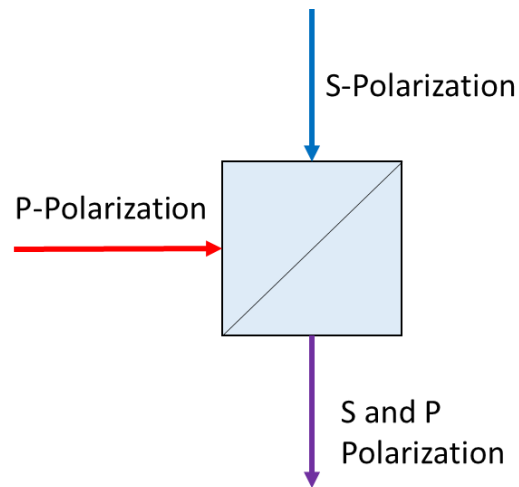


Figure 76: Simple schematic demonstrating how DIAL, which has two laser sources, can be combined to create a single laser output path.

A laser vendor may provide the laser as a unit with a single fiber output carrying both signals. In this case, the instrument developer will have to work with the different polarization states and this may make beam splitters problematic. An example polarization dependence from a standard plate beam splitter is shown in Figure 77 [ThorLabs, 2017].

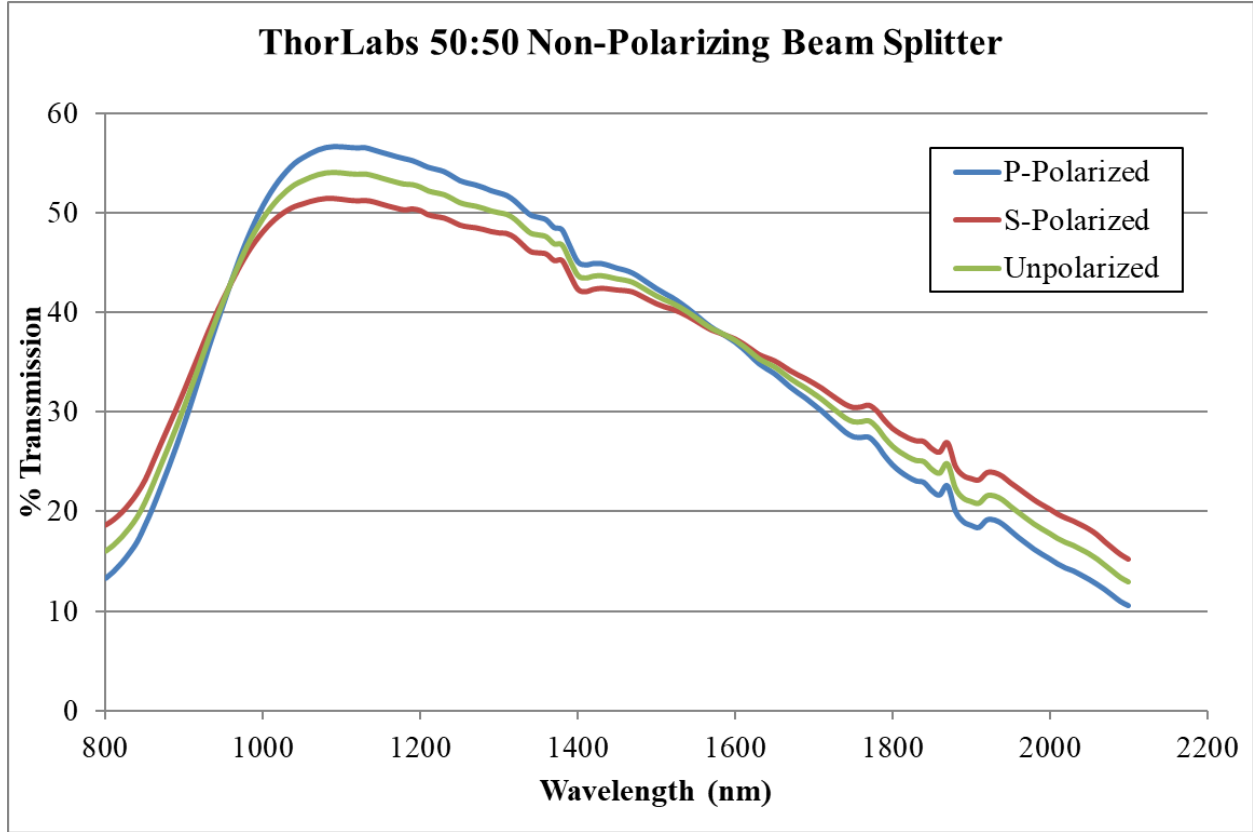


Figure 77: Example plot of how even “non-polarizing” beam splitters still have some polarization dependence. [ThorLabs, 2017]

The plot shows that two beams with opposite polarization will have different transmitted/reflected energy levels on the other side of the beam splitter. This is a problem if the power measurement is assuming that the transmitted/reflected energy for both beams is the same percentage. (Note that there is no assumption that the power levels are the same, only that the *percentage* of energy loss is the same.) From the second term of equation 4, the ratio of the optical efficiencies affects the final number density calculation and leads to errors in gas measurement:

$$\frac{1}{2\Delta\sigma_C(R)} \frac{d}{dr} \ln \left[\frac{\eta_d(\lambda_{off}) \cdot \eta_r(\lambda_{off,R})}{\eta_d(\lambda_{on}) \cdot \eta_r(\lambda_{on,R})} \right] \quad \text{Eqn. 11}$$

If the detector efficiency η_d is the same for both wavelengths, then the optical efficiency η_r is the driving parameter of the term.

A short trade study investigating polarization induced errors results in the plot of Figure 78. The approach was to calculate the ratio within the natural log term of equation 11 for a variety of online and offline efficiencies while assuming the detector efficiencies were equal. The natural log of the ratios creates the error coefficients on the y-axis that would feed into the rest of the term. Note for instance that if the online optical efficiency is 50% then an offline efficiency of 50% would create a ratio of 1 and the natural log error coefficient is zero. The error coefficients on the y-axis that would feed into the rest of the term. Note for instance that if the online optical efficiency is 50% then an offline efficiency of 50% would create a ratio of 1 and the natural log error coefficient is zero.

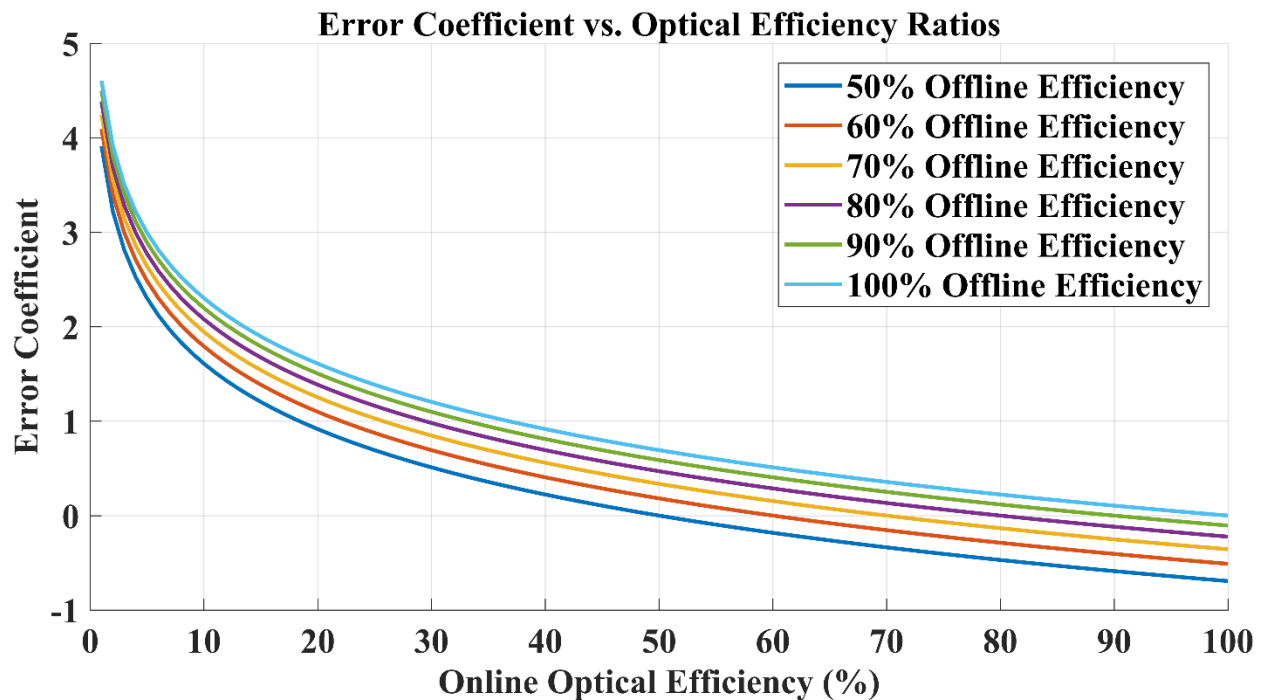


Figure 78: Plot of the error coefficient that would multiply through the rest of equation 11 as a function of optical efficiencies.

The theoretical analysis in Figure 78 shows that the error coefficient is likely to be small. It is then divided by the cross-section difference and, as with the reflectivity analysis, this is then subtracted from the first term of equation 4. Depending on the magnitude of the received signal, the error may be negligible or on the order of a few percent. As signal levels rise and fall, the optical efficiency error will be constant. This means the error's importance could fluctuate as the instrument flies over ground with

various reflectivity. It seems better, then, to try to minimize the polarization dependence to remove this type of error variability.

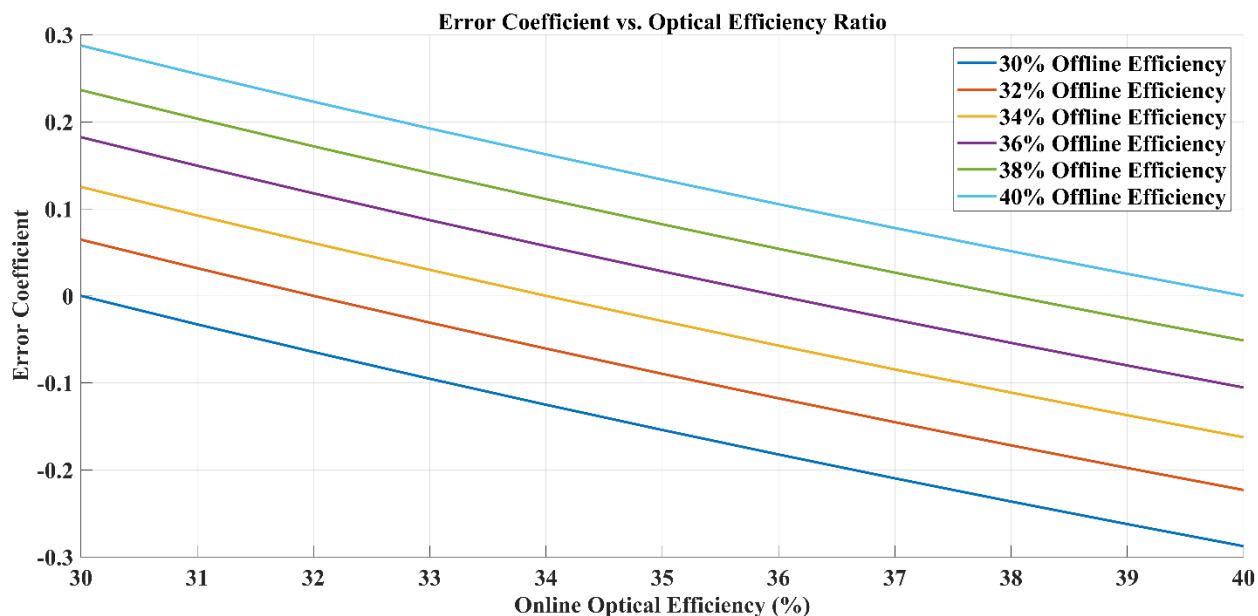


Figure 79: Plot of the error coefficient for ThorLabs' 50:50 beam splitting optic at a wavelength of $1.645 \mu\text{m}$. The range of values maps to their as-measured light splitting variations due to polarity.

As an example, a simulation run on ThorLabs' 50:50 beam splitting optic of Figure 77 at the target wavelength region of $1.645 \mu\text{m}$ created Figure 79. The plot shows that the error coefficient produced by polarization efficiency mismatches can be on the order of 0.1 to 0.2 for an optic with a 3% difference in polarization specific efficiency. If the ratio of the offline to online signal is on the order of 1 or 2 (a likely scenario) in the first term, then the error coefficient may represent a 10% modification.

Worth noting is that the efficiency mismatch is not an unavoidable error contribution! It's only a problem if it's an unknown, unaccounted for parameter. Careful analysis and testing may produce the correction values needed to negate the error source. The takeaway from this brief trade study, though, is that the differential optical efficiency is a non-trivial part of the approach and should be carefully considered in the instrument's design.

Another approach is to find optics that interact with different polarizations as evenly as possible. One such option is gold coated reflectors. The ThorLabs protected gold coating data is provided for the 1.645 μm wavelength region in Figure 80. The difference between S and P polarization reflectivity is a few tenths of a percent. Additionally, as a ratio of two high numbers, the natural log error coefficient is a low 0.0035. Using a reflector, though, to combine or split beams means mechanical separation with something like a knife edge prism. Splitting a beam in this way means precise alignment on the edge of the prism is required and scatter is an unavoidable effect. On the other hand, the polarization-induced losses are equivalent for both polarities and so does not impact the final result other than a bulk lowering of signal.

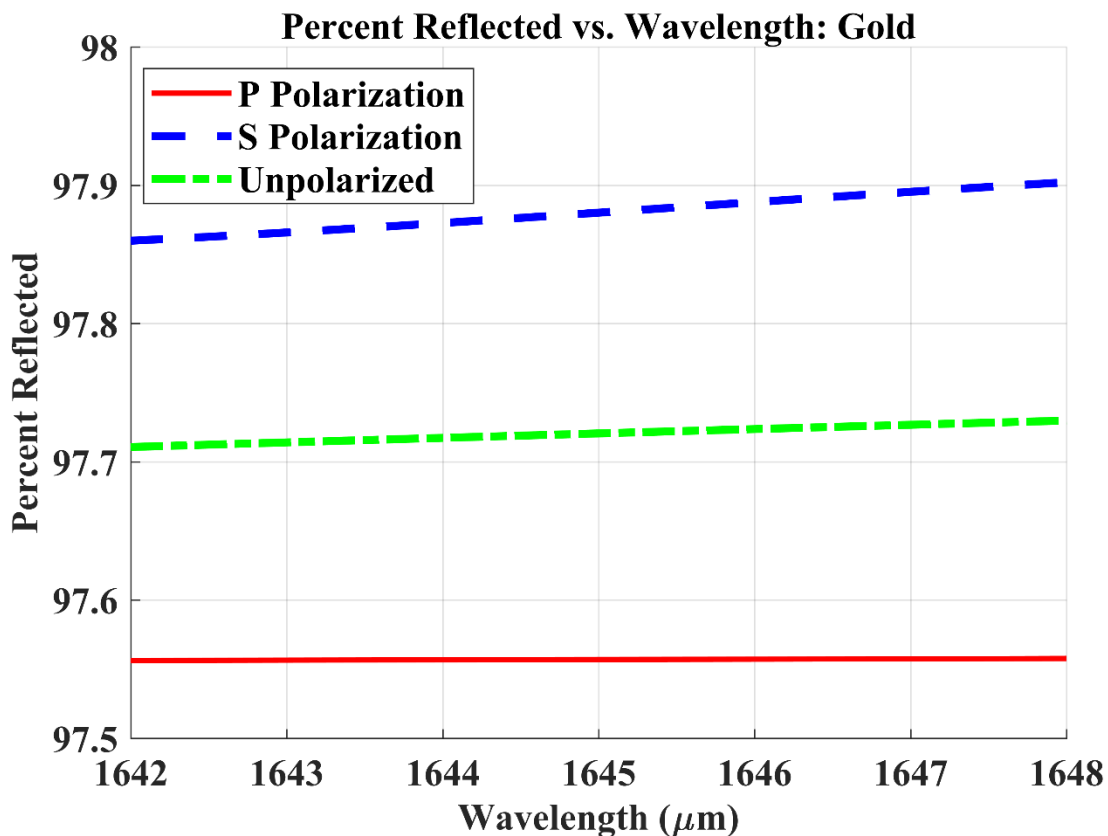


Figure 80: Plot of ThorLabs protected gold coating for the region around 1645 nm.

6.5 Minimizing Optic Count

Reducing optic count is an obvious goal to minimize both losses and alignment concerns. To help quantify it numerically, Figure 81 was created. It assumes that every optic in a system has the same efficiency. Although a simple plot, it helps ground the importance of carefully selecting the best optic and finding ways to minimize optical interfaces.

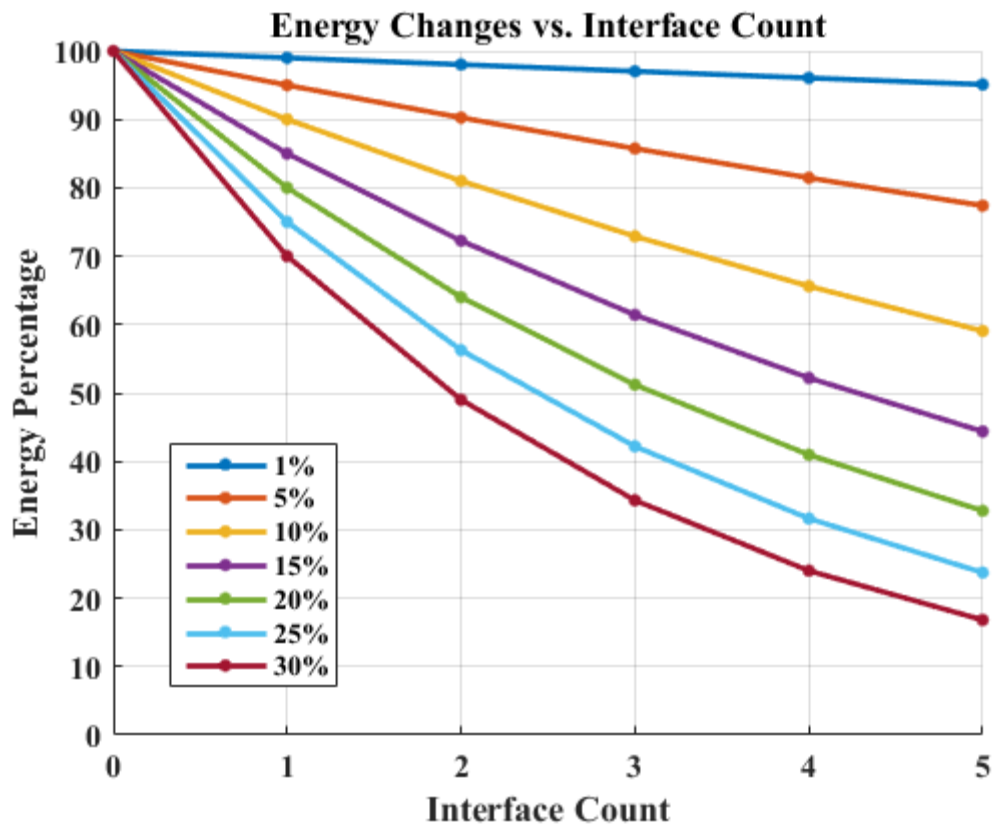


Figure 81: Plot of the reduction in energy/photon counts per optical interface given a fixed loss value per interface.

6.6 Beam Steering

Many beam deflecting technologies were considered including Risley prisms, voice coil controlled mirrors, Micro Electrical Mechanical (MEM) systems, and gimbals. The challenge with all of the above is that finding off-the-shelf systems with apertures greater than a couple of inches proved difficult. The largest diameter Risley prism pair found was 2" [Optra, 2017]. Additionally, Risley prisms rely on

refractive optics which introduce bulk losses as well polarization specific losses that the design is specifically working to avoid.

MEM systems are low power and weight, but reflected beam diameters on the order of a couple of centimeters at best as they are intended for steering the outgoing beam, not for collecting larger beams [OpusMicro, 2017; Hamamatsu, 2017]. On the other end of the spectrum, gimbal systems can be powerful, but personal professional experience has shown their pointing accuracy to be less than often stated. Further, their volume, weight, and complexity make them unattractive for a UAV system.

Voice coil controlled mirrors promise high accuracy and larger optics, but are expensive with a 3” assembly going for \$15,000 [Optics in Motion, 2017]. In spite of the cost, this technology offers possibly the best mix of large aperture and fine control in a relatively compact volume. Although extensive searching of commercial voice coils could not find a 4” or larger optic to get the needed collection area, it seems that this is an area that may see new commercial products. Further, at least one modern airborne DIAL system relies on the technology [Bartholomew, 2017].

The technology settled on for this application is something akin to a galvanometer assembly. The concept is a rotating axis with a customized mirror on one end. The axis is tracked so that the system knows its rotation in degrees. The mirror is oriented roughly 45° to the beam axis, but has a few degrees of relative tilt. The tilt causes incident beams to be steered in the pattern of a circle on the ground.

An additional aspect of this mirror design is that it has an outer diameter that is slightly tilted up relative to the center. This feature is used to deflect the outgoing beam without taking up any of the primary mirror surface in order to maximize the light gathering capability. The slight tilt accounts for wanting to steer the outgoing beam towards the center of the receiver’s optical footprint on the ground. This concept is demonstrated in Figure 82. The actual tilt is calculated in the final design as there are a number of input parameters that drive the design including expected flight altitudes, optic diameters, beam divergence, and aircraft angles that extend the path to the ground and back.

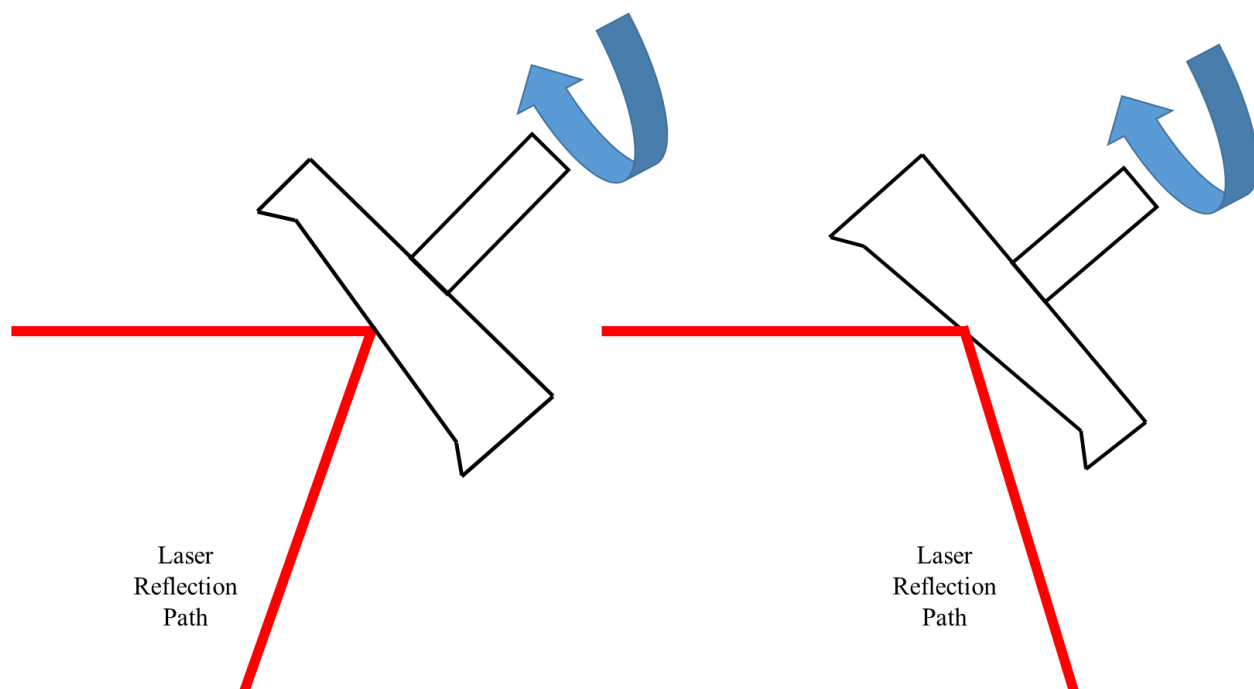


Figure 82: Simple schematic of how a rotating mirror with a wedge profile can trace out a path on the ground. Angles are exaggerated to illustrate the concept.

The primary tilt angle of the beam steering angle was set to 6° to give a beam diameter that varied between 63 feet and 105 feet between the altitudes of 300 and 500 feet. This range of coverage is based on in-person conversations with industry representatives who expressed interest in roughly 30 to 50 feet of coverage on either side of a pipeline in order to map leaks that may work their way through the soil at an angle [Wood, 2015]. The diameter of the central mirror section was set to 162 mm, which accounts for the 51° minimum mirror angle relative to the incoming beam angle and its collecting diameter of 100 mm. A few millimeters of margin were added in the mirror's diameter plane to account for ± 1 mm of translation of the optic in its plane ($3 \text{ mm}/\sqrt{2} \approx \pm 1 \text{ mm}$). The width of the outer surface used to deflect the outgoing laser beam was set to 19 mm. This is based on the beam having a diameter of 15.6 mm at the collimation optic, a beam divergence of 0.0048° , a path length of up to 250 mm, and the same compensation as above to account for the tilt.

Determining the relative tilt of the outer diameter to steer the outgoing beam into the receiving optics' ground footprint required a trade study. The above parameters were put into a MATLAB script that

calculated the ground footprints of both the outgoing and incoming energy for a range of altitudes. Next, the angle required to steer the smaller outgoing beam's footprint within the incoming beam's footprint was determined for a given altitude, assuming a 10% margin to keep the beam well within the bounding area. A plot from this script is provided in Figure 83. There, the minimum and maximum relative steering angles are provided.

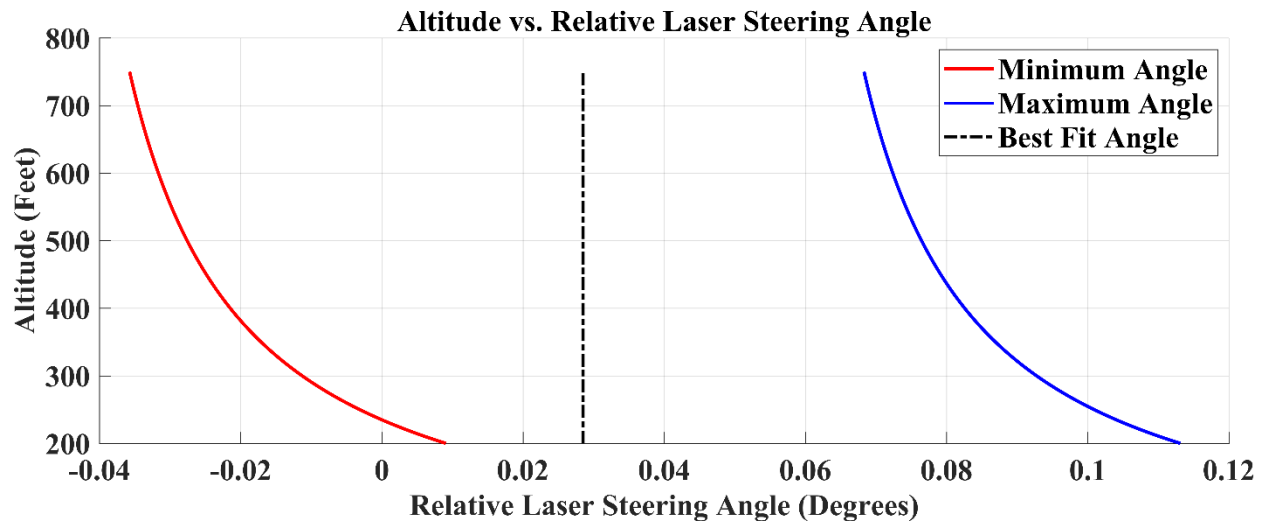


Figure 83: Plot of the minimum and maximum relative angles the outside diameter of the disk can be while keeping the laser footprint within the receiver footprint.

It is seen that the angle can go negative in order to account for the growing footprint of the receive optic system. The next observation is that the receive optic's footprint has a larger beam divergence angle, hence the curved nature of the minimum and maximum lines in the figure. Also, all of the magnitudes of angle are small, with just a couple of hundredths of a degree being necessary.

The two bounding lines were used to pick a final, fixed angle of 0.028° , which is the best fit for the smaller beam around the center point of the larger footprint for the altitudes between 300 and 750 feet. This angle keeps the beams collocated until within 150 feet of a surface, at which point the smaller beam is outside the ability of the receive system to see it.

It should be noted that the altitude can also be read as the flight path. A UAV turning at a 45° bank angle at 500 feet will produce a beam path length of $500/\cos(45^\circ) = 707$ feet. The footprints will no longer be circular, but their distortion into an ellipsoid does not alter their overall colocation. The 750 feet of altitude in Figure 83 was chosen based on an expectation of airplane orientation increasing path length to about that magnitude.

6.7 Detectors

In many ways, the detector is the heart of the system and is one of the largest drivers in overall design. Understanding the various parameters that influence the ability to capture the signal with minimum noise is the subject of lengthy, high quality papers such as Richards' 1994 paper with over 800 citations [Richards, 1994]. A few key parameters are discussed here with one detailed look into a common rule of thumb suggested by vendors.

The first parameter considered for choosing the detector is responsivity. This parameter strongly drives the performance of the system because it influences both the noise in the system as well as the amount of incoming energy converted to a signal. For detectors that output current, the units are Amps/Watts, which means that for every 1 Watt of incoming photon power, so many Amps of current are generated.

Because the incoming photon power is wavelength dependent, the responsivity is also wavelength dependent. The simple formula is:

$$R(\lambda) = \frac{\text{Generated Current}}{\text{Incoming Power}} \quad \text{Eqn. 12}$$

Sometimes, vendors report the quantum efficiency of the detector instead of the responsivity. Quantum efficiency is the percentage of the incoming photons that create charge carriers in the photodiode. This in turn is multiplied by the energy of the photon and the elementary charge. As a function of wavelength, the equation looks like:

$$R(\lambda) = \eta \frac{q\lambda}{hc} \quad \text{Eqn. 13}$$

, where η is the quantum efficiency, q is the elementary charge, λ is the wavelength, h is Plank's constant, and c is the speed of light.

The second key parameter in choosing the detector is understanding its Noise Equivalent Power (NEP).

The definition most commonly used by vendors is that NEP represents the incoming signal power required to equal the power created by the detector's noise when the detector has a 1 Hz bandwidth [Richards, 1994]. The normalization of 1 Hz is done to more effectively compare different detectors.

To determine the actual minimum power, two things must be done. First, the NEP is wavelength dependent because detector responsivity is wavelength dependent. The NEP provided by vendors is typically the lowest value across a detector's range of wavelengths and is defined at the maximum responsivity, so it must be converted using a simple ratio of responsivities [ThorLabs, 2017]:

$$NEP(\lambda) = NEP_{min} * \frac{R_{max}}{R(\lambda)} \quad \text{Eqn. 14}$$

Where R_{max} is the responsivity at the published maximum value and $R(\lambda)$ is the value of responsivity for the wavelength of interest and is usually provided by a vendor datasheet. Second, the square root of the detector's actual bandwidth must be multiplied through. The next section discusses bandwidth in detail, but for now the takeaway is that minimizing bandwidth is beneficial to reducing noise. The final formula for the minimum detectable signal power looks like [ThorLabs, 2017]:

$$P_{min} = NEP(\lambda) * \sqrt{Bandwidth} \quad \text{Eqn. 15}$$

This leads to the third key parameter, which is the detector's bandwidth. There are several bandwidths that might be found in vendor discussions including detector bandwidth, measurement bandwidth, and gain bandwidth. Briefly, detector bandwidth is the parameter of interest and relates to how fast the detector responds to a signal, measurement bandwidth is the bandwidth of the system after post-detection filtering has been applied, and gain bandwidth is the bandwidth of the amplifier system. Each of these are

important, but the detector bandwidth is the one that starts the process and is the one discussed further here.

The detector bandwidth is related to the rise time of the detector. The rise time is itself defined as the time it takes for the detector to go from measuring 10% of a signal to 90% of a signal [ATIS, 2017]. The “rule of thumb” for relating bandwidth to the rise time is:

$$Rise\ Time \cong \frac{0.35}{Bandwidth} \quad Eqn. 16$$

This relationship can be derived from the response to a unit step function in a low-pass RC filter [Orwiler, 1969]. A photodiode is precisely this type of system. However, finding a thorough discussion of this rule of thumb behind the linear analysis turned up little information. Out of curiosity, a deeper dive into the value of 0.35 was taken.

An interesting way to see the relationship between bandwidth and the rise time can be found by analyzing a square wave, which is a series of step functions. The first plot in Figure 84 is of a 20 ns square wave pulse with a 50% duty cycle and amplitude of unity. The second plot is the Fourier transform of the square wave into frequency space. The peaks of the second plot can be used to recreate the original square wave with sine waves of increasing frequency. The result of adding the first 6 peaks together is shown in sequence in Figure 85 with only the rising part of the signal plotted for clarity.

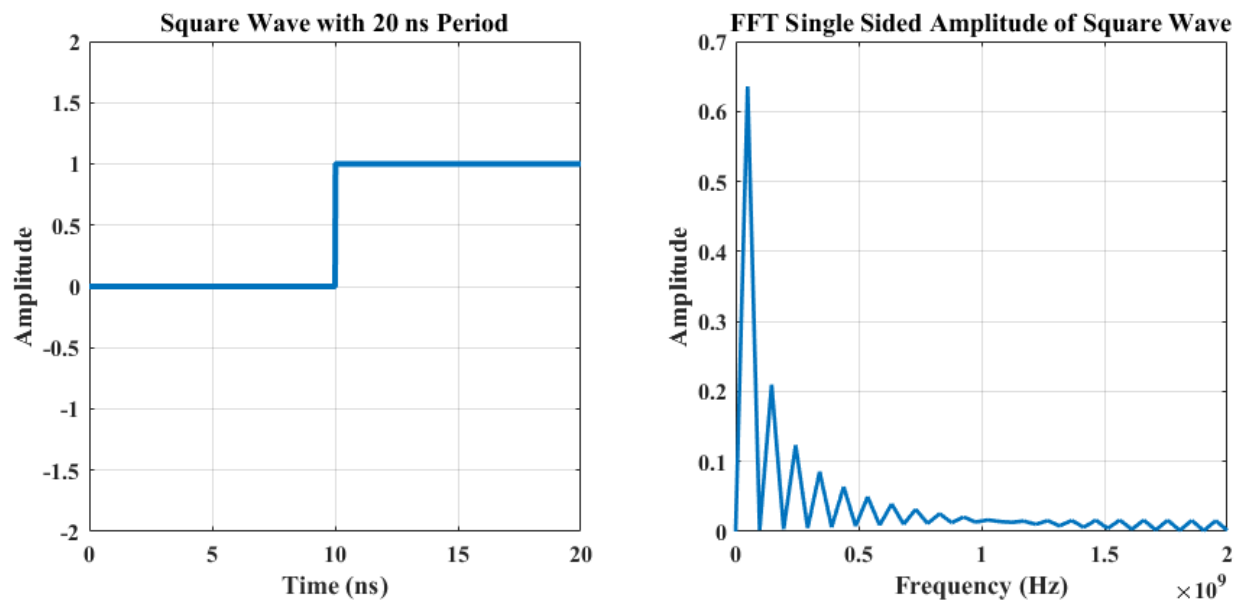


Figure 84: A square wave with 50% duty cycle and 20 ns period in both the time and frequency domains.

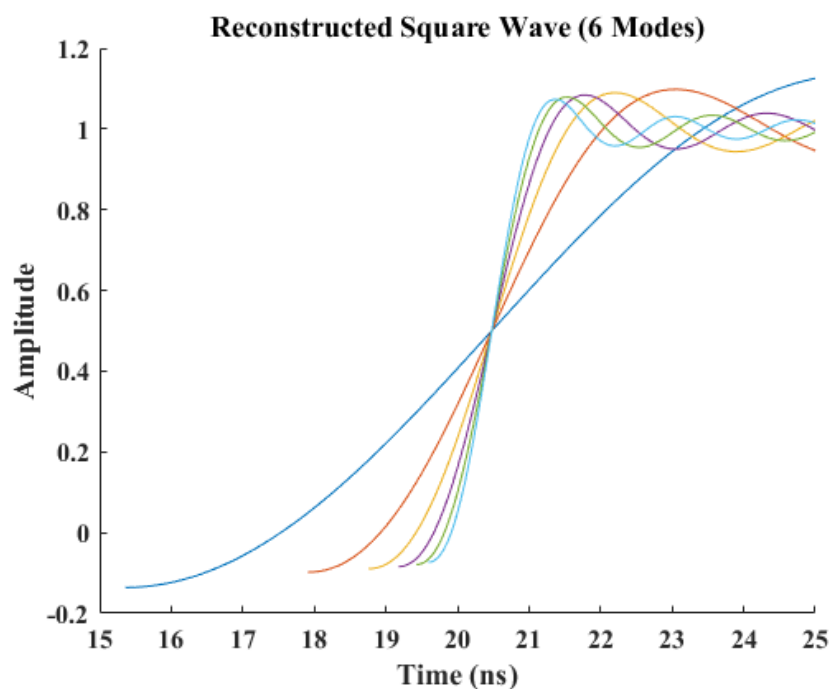


Figure 85: Each curve represents the signal created by combining the modes up to its value. For instance, the long periodic form has one mode, the next curve has the first two modes, and so on.

It is seen that increasing the number of modes increases the ability to resolve the original square wave. For a DC circuit where the lower frequency is zero, the upper frequency essentially defines the bandwidth of the system. Thus, if the bandwidth is set to the first mode's frequency, it is seen that the square wave is not well resolved. Increasing the highest frequency (bandwidth) to include the second mode provides for a significant percentage of the square wave's integrated area. The bandwidth can be increased to infinity, but it is seen, even with 6 modes, that there are diminishing returns.

In fact, this relationship holds for square waves, no matter what their period is. The plot of rise time as a function of mode number is shown in Figure 86 on the left. The plot in the middle is the delta decrease in rise time for each sequential mode. Finally, the plot on the right plots the rise time vs. the bandwidth. It is seen there that a value of 0.35 encompasses the first two modes, which is roughly where the linear range of the plot ends.

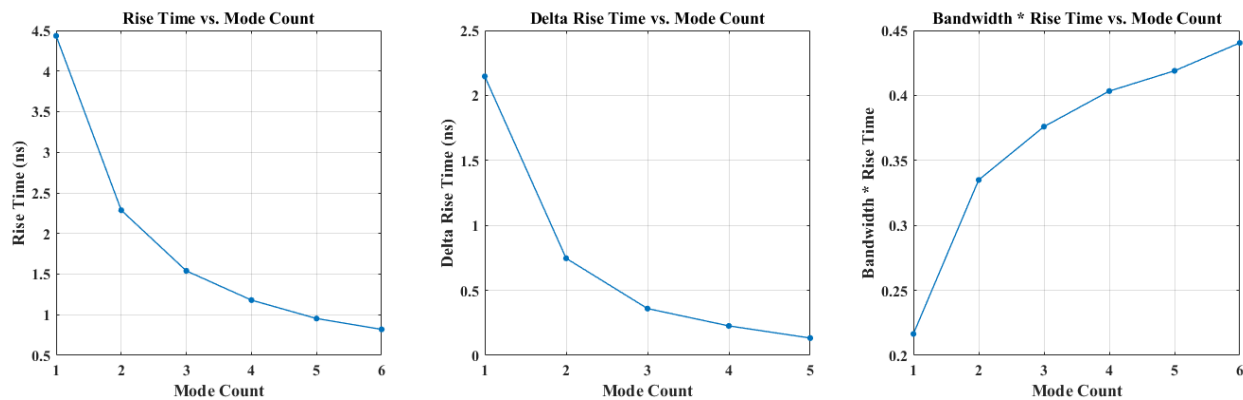


Figure 86: Plots of rise time demonstrating that adding additional modes (bandwidth) provides limiting returns in recreating the square wave pulse.

This type of analysis can be extended to consider a non-square wave pulse. For example, the physics of laser pulse creation makes a nicely square pulse unlikely. And, since the frequency based analysis is based on a Fourier transform, the limit of this methodology is a pulse shape similar to a half sine wave, such as that shown below in Figure 87. The Fourier transform to the frequency domain for a sine wave is

a single mode (the minor roll-off after the first mode is an artifact of the non-perfect discrete simulation of the sine wave).

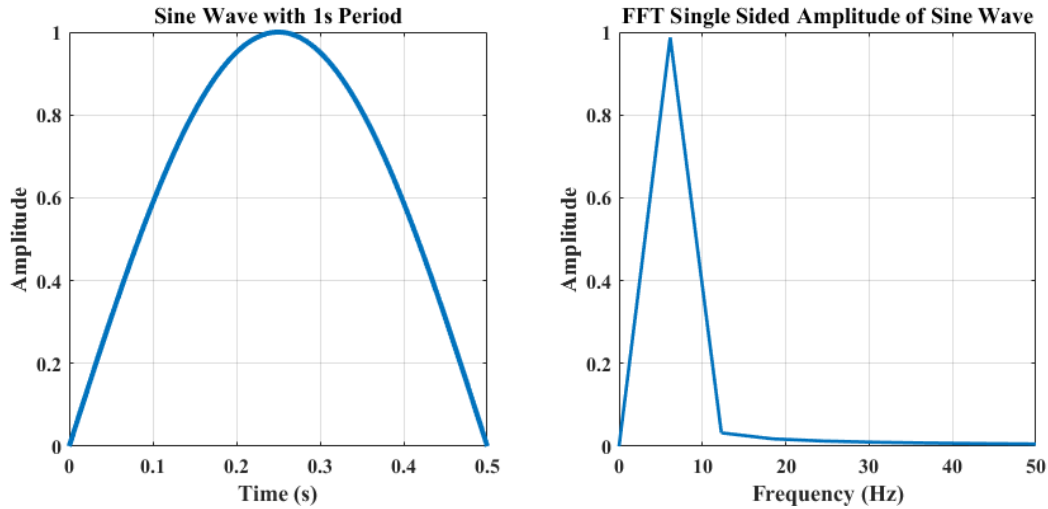


Figure 87: The plot on the left is a time series representation of a half sine wave. The plot on the right is the Fourier transform of the somewhat discretized sine wave.

When the rise time is calculated for the first three modes (two are just minor noise), the result is Figure 88. It is seen that the rise time levels off after the first mode as expected. In the plot on the right, it is seen that the bandwidth and rise time product make big jumps. Because rise time decreases with increasing bandwidth, the large jump is saying that the frequency has to increase significantly.

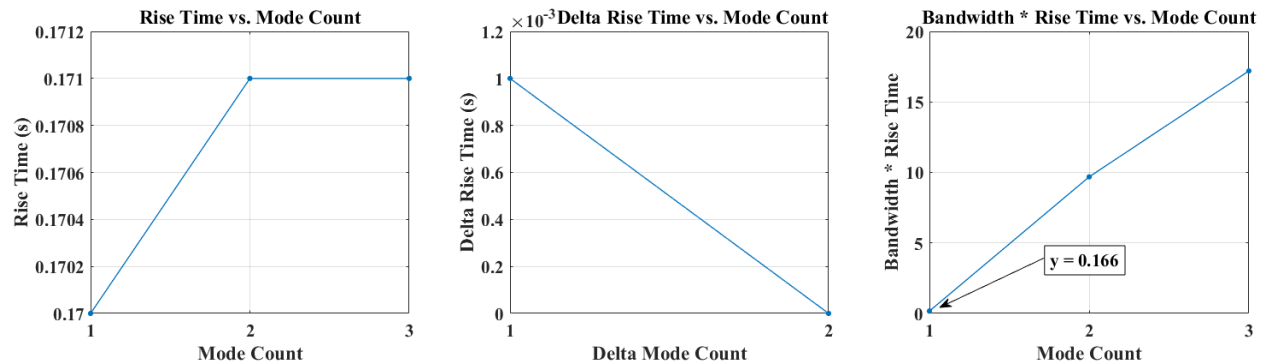


Figure 88: Plots of rise time demonstrating that increasing the mode count for a sine wave Fourier series reconstruction has limited utility.

The main takeaway from the plot, though, is that the majority of the pulse shape is captured with a bandwidth and rise time product of just ~ 0.17 . This is roughly half of the value for the square wave pulse. In practical terms, this means that the detector bandwidth can be a little narrower for applications measuring pulses that are closer to a sine wave shape than a square wave. This is great news since reduced bandwidth means reduced noise. It should be noted that pulses are not likely to look exactly like a sine wave and so specific analysis needs to be done for specific pulse shapes. However, if a detector is on the threshold of being acceptable for the square wave assumption then confidence in its applications for rounded pulses is reasonable.

6.8 Signal Measurement Via Integration

Two approaches were considered for converting the analog signal to digital: one where a pulse is traced with sufficiently fast sampling to map the shape at discrete points and one where the pulse energy is integrated by a capacitor which is then read before being reset. These two choices are designed to integrate the total energy in a pulse rather than taking the peak pulse value. Before continuing, a brief detour is made to describe the challenge of the simplest approach of taking the peak value.

The DIAL equation works with any consistent measurement of the various channels. One of the easiest measurements to take is the peak of each pulse. For an idealized shape, such as seen in the left-side plot of Figure 89, finding the peak is as simple as looking for the maximum point. Perhaps the algorithm takes all the data in the pulse and then finds the max, or stops taking data after the values start going down again. Curves can also be fit to the data to fine tune the peak estimate. The challenge, though, is that real world returns from a noisy surface are not likely to be smooth.

For instance, consider a pulse where part of the beam strikes a utility shed roof, part strikes the ground, and part strikes a parked truck. These various elevations corrupt the pulse to look something like the plot on the right-hand side of Figure 89. Determining what the peak return should be is difficult, if not impossible. It's more likely that a pulse like this might be tossed as being too noisy. This may be a

reasonable solution when flying along a well maintained pipeline right-of-way as the expectation is minimal variations in elevation. However, overgrown trees, storage tanks, well sites, and other common industry activities could mean that a system using just the peaks becomes functionally untenable.

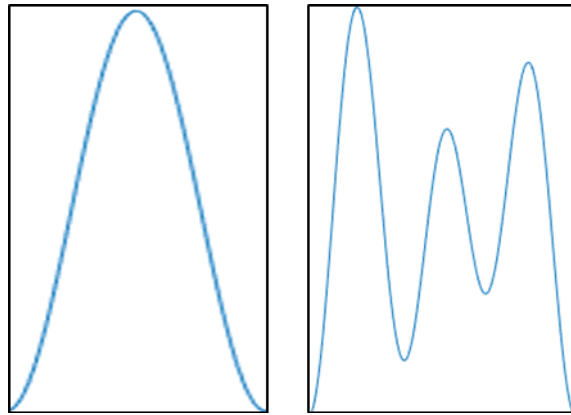


Figure 89: Simulated pulse shapes with an idealized return pulse on the left and a multiple reflecting pulse on the right.

Instead of trying to make a robust peak finding algorithm, it seems better to design with integration in mind. Capturing all of the energy means that it doesn't matter how the pulse is spread out or shaped. It is true, though, that new challenges are introduced. For instance, if the noise on or in the detector is not white noise then it will impact the integrated measurement in non-distinguishable ways. Also, the capacitor needs to be fully discharged before the next pulse so that previous pulses do not inflate subsequent pulses. A third thing to consider is how to determine the range to target.

The concept design is taken from a 1996 paper on the topic of a low-cost “boxcar” system [Collier, 1996]. (The term boxcar refers to the way rectangular integrating periods seem to form a train of box cars.) The approach is also sometimes referred to as a “gated integrator” [TN 1005, 2004]. A summary is shown in Figure 90.

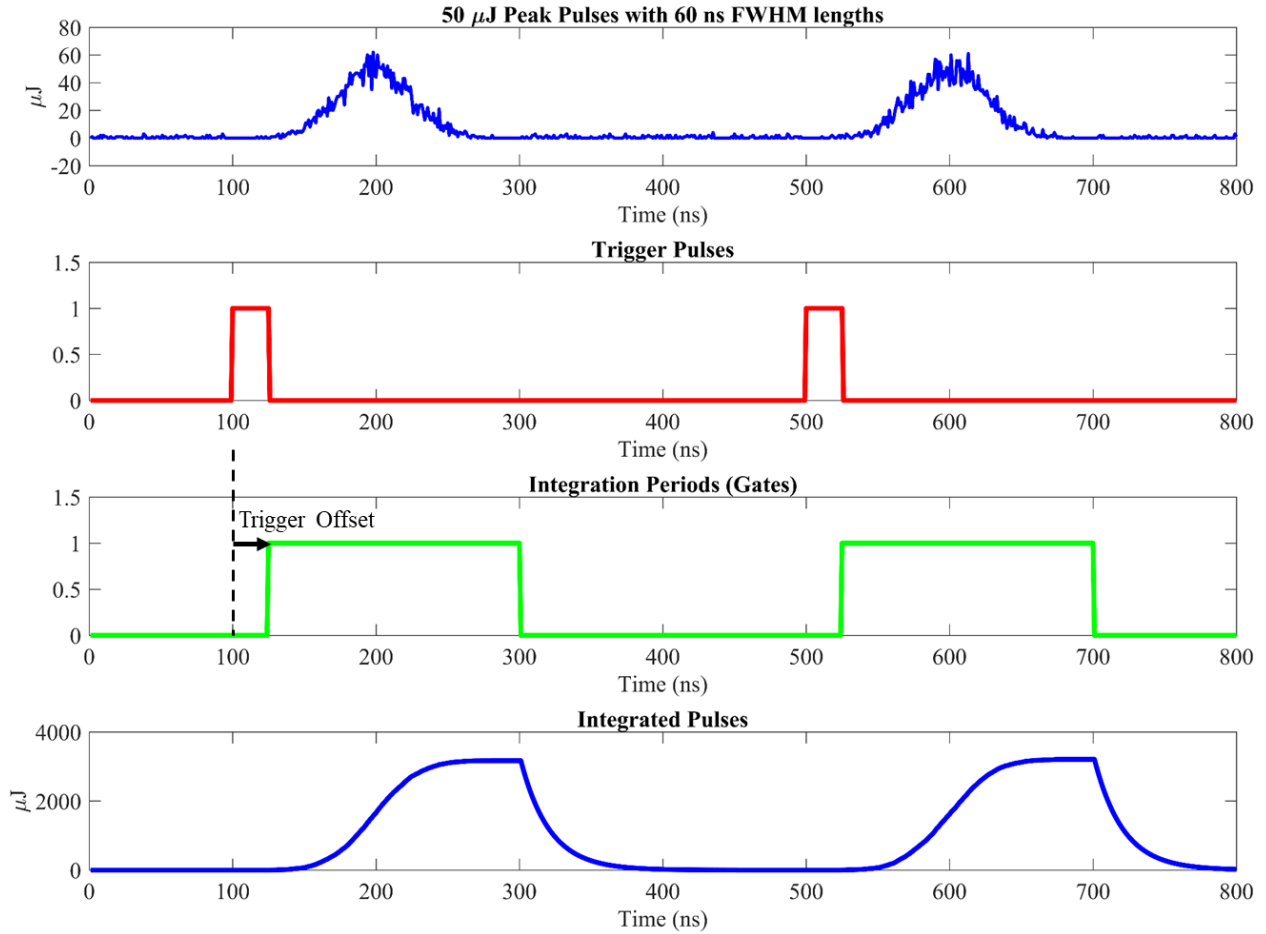


Figure 90: Overview of the boxcar integration technique where pulse energy is cumulatively summed within the integration period. The top plot is two Gaussian shaped plumes with simulated Poisson shot noise. The second plot shows external trigger pulses being sent to begin the process of integration. The third plot shows that the integration gates begin summing energy a fixed time after the trigger is sent. The bottom plot shows the results of the summing operation with RC roll-off. The ADC reads the energy from the top of the integrated pulse before the gate ends. Triggering the start of the integration and using an appropriate integration period are two key parts of the approach.

At least two parameters are impactful: the width of the integration period and how well centered the pulse is. To understand how they can influence the outcome, two trade studies were run. The first looked at how the integrated energy changes as a function of the integration period (gate width) when the gate is centered at the center of the pulse. The goal was to isolate all other effects, therefore contributions, such as dark noise, are set aside for now. However, Poisson shot noise was added to determine if it would

meaningfully change the curves. To scale the result agnostically of pulse size or width, the x-axes are ratios of the gate width and the FWHM time. An example outcome of this study is shown in Figure 91.

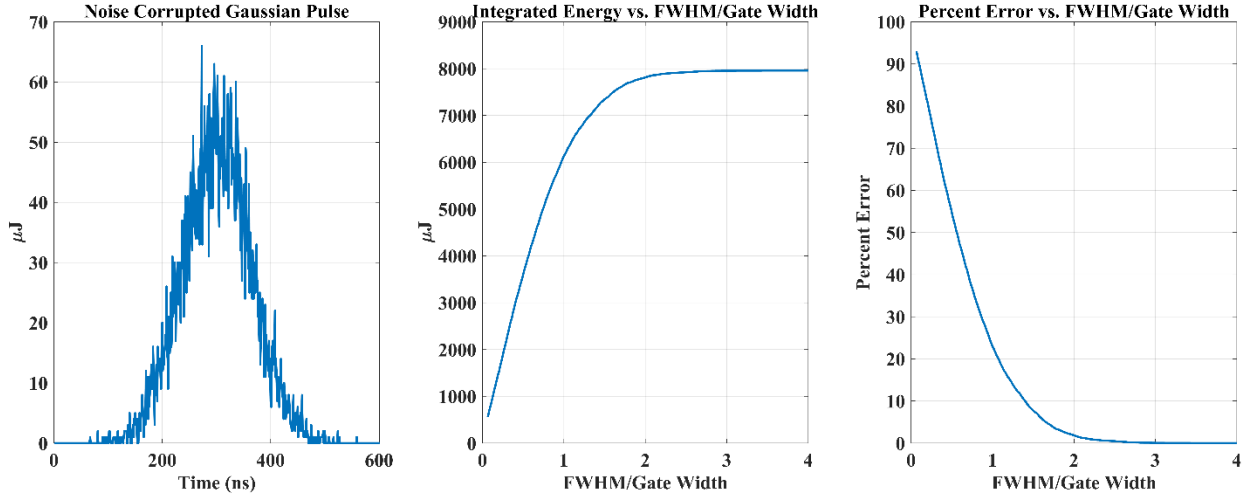


Figure 91: Example integration result for a series of centered integration gate widths. The plot on the left is a Gaussian pulse with a 150 ns FWHM duration. The middle plot shows the integrated energy as a function of the number of FWHM time periods. The plot on the right shows the percent error relative to the true energy within the pulse.

It is seen that, regardless of pulse size and duration, 99% of the energy in a pulse is captured with a gate width of 2.18 times the FWHM pulse duration. Getting the last 1% of energy requires a gate width of about 4 times the FWHM period. It was also determined that simulated Poisson shot noise had less than 1% impact on the outcome, although the zero-mean nature of the noise coupled with the number of measurements makes this result somewhat expected.

The second trade study investigated the sensitivity to how well the gate is centered on the pulse. For instance, if the integration period starts and ends too early then it would miss energy. Because this result also depends on the gate width, that was added as an extra parameter. As with the previous study, the results are normalized relative to the FWHM time. Figure 92 provides the outcome of the study.

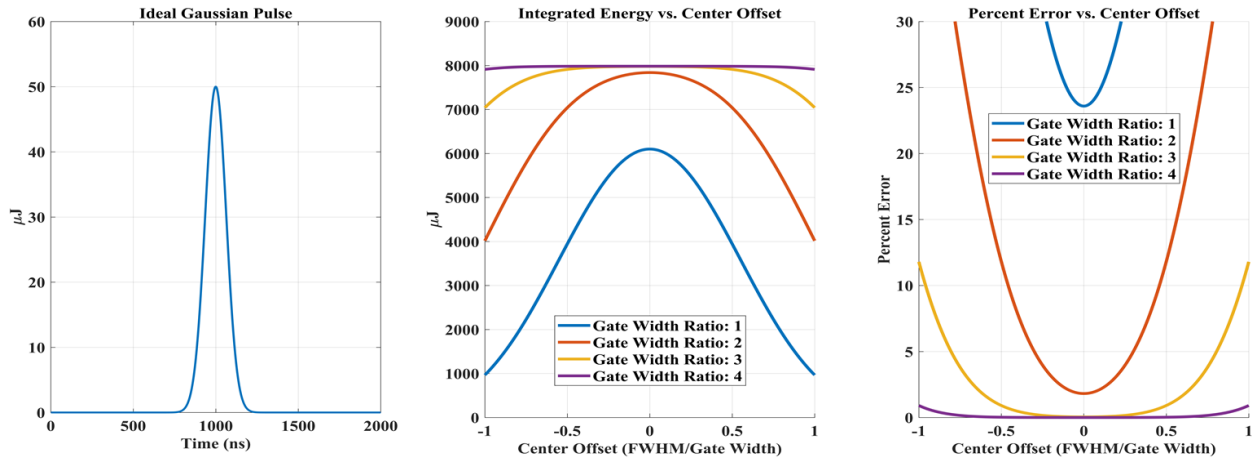


Figure 92: Plots results for a study investigating the sensitivity to gate location relative to the center of the pulse. The plot on the left is an ideal Gaussian pulse with FWHM time of 150 ns. The middle graph plots the integrated energy for four gate widths that slide left and right of the center. The x-axis value of zero means the pulse is centered in the gate width period. The plot on the right provides the percent error for the same parameters as the middle plot.

It is seen, in this idealized scenario, that the least error comes from a large integration window that is well centered on the pulse. However, even with this noise free study, it is seen that the errors can quickly escalate. In the previous study, a normalized gate width of 2.18 captured 99% of the signal. This study shows that a gate width of 3x FWHM is somewhat robust with a 50% offset returning about 1% error. If the gate width is instead 2x FWHM and 50% offset, then the error is 10%.

A key issue with getting the integration period well centered is range uncertainty. It is possible to know when pulses are sent, but terrain and instrument orientation lead to return time uncertainties. For the example pulse of 150 ns FWHM, a 50% missed centering means 75 ns, which is about 75 feet of unexpected time shift. If the system is at 500 feet, this is about 30 degrees of unexpected roll or some combination of rolling and a terrain cliff. That seems like a rare case. Likewise, if the pulse center is missed by closer to 60 ns, then this would be about 19 degrees of roll, which is still more than what might be considered a typical scenario.

More practical is a roll of 5 degrees or terrain drops/rises of 5 feet or so. For the more stringent case of a 40 ns pulse, this would mean a 2 ns to 5 ns shift. In normalized terms, this would be a ratio of 0.1 to 0.25

for Figure 92's x-axis. For low error, choosing a gate width ratio of 3x the FWHM seems the most robust. If both pulses have the same error, then it's a constant that divides out in the ratio of the two. However, the above assumes no noise integration and that the returned pulses shapes are the same.

The final study looks at adding noise to the floor of the signal. Three noise floor values are used: 1/10, 1/100, and 1/1000 of the peak pulse value, representing SNR values of 10, 100, and 1000 respectively. Poisson shot noise is added to the signals. Three example pulses are shown in Figure 93 to demonstrate what is being integrated. Because wider pulses now mean more integrated noise and thus an expectation that the results will change with pulse width, the calculations are first done for a 50 ns FWHM pulse and then for a 150 ns FWHM pulse. The three computations for each set follow the approach in the third plot of Figure 92 in that percentage errors are provided as a function of normalized gate width.

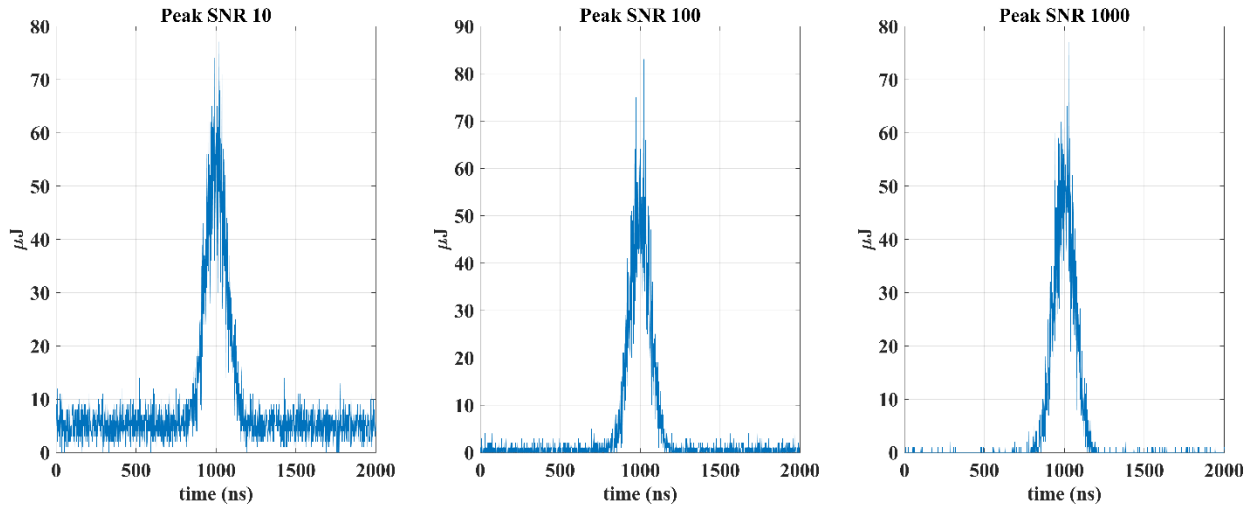


Figure 93: Example 150 ns, 50 μJ peak pulses with peak SNR values of 10, 100, and 1000. Poisson shot noise has been applied to the signal.

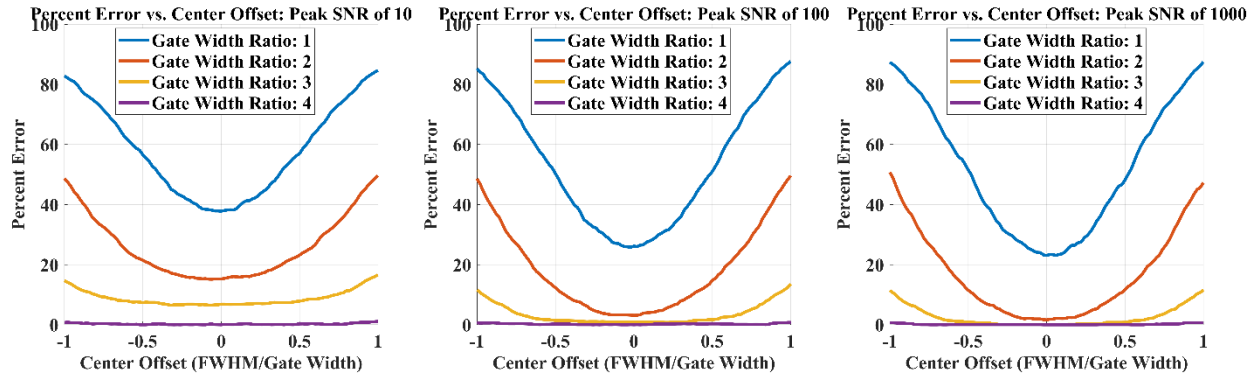


Figure 94: Percent errors as a function of the center offset and gate width for the 50 ns pulses and peak SNR ratios of 10, 100, and 1000.

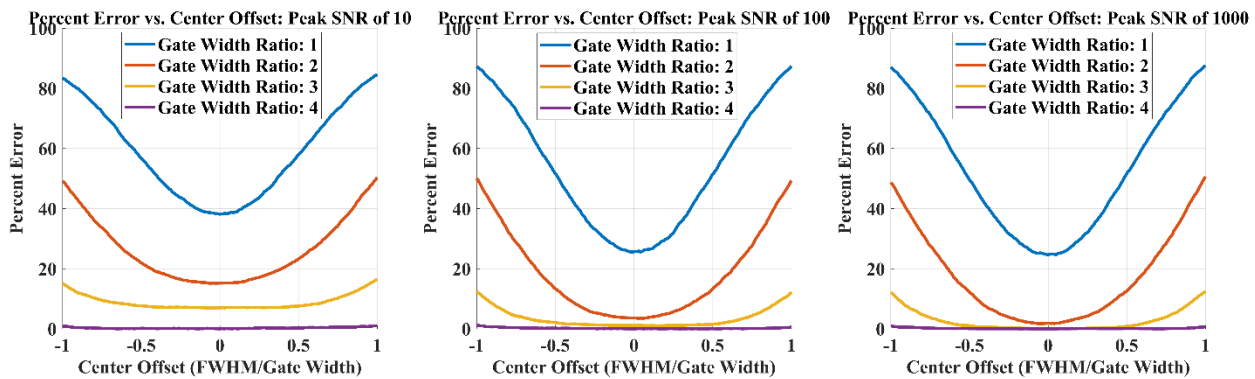


Figure 95: Percent errors as a function of the center offset and gate width for the 150 ns pulses and peak SNR ratios of 10, 100, and 1000.

A number of observations are made from the plots in Figure 94 and Figure 95. The first is that, perhaps contrary to expectations, the pulse width does not show notable differences in outcomes, even for the low SNR case. Second, even with a low SNR of 10, as long as the entire pulse is captured, the noise does not introduce statistically important error. (Note that a gate width of 4 and a center offset of 1 means that a gate width of 3 is mostly covering the pulse.) Third, having a SNR of 100 drives the error related to this type of measurement just about as low as it can go with the middle and right plots of either Figure being about the same. The SNR of 1000 does provide a little more margin, but there are stronger reasons for wanting higher SNR than integration windows. For this dissertation, a gate width of 3x to 4x the FWHM is being used with a center offset not greater than about 0.5x the FWHM.

One phenomena not discussed here is that true pulses are captured by a detector with a capacitor and resistor, which will mean at least a pulse-distorting roll-off and possibly ringing of the signal. Another consideration is that real-world pulses are stretched by interactions with things such as optics, the environment, and detector bandwidth. Expecting a 50 ns pulse and actually receiving a 70 ns pulse would drop the gate width ratio by about 30%.

6.9 Analog to Digital Conversion, FPGA, and Storage

This is an important topic, but detailed design of the analog to digital conversion (ADC), FPGA, and data storage are beyond the scope of this effort. The shortcoming is twofold: robust analysis requires electrical engineering and the budget does not support buying the laser let alone hardware to test a laser. Ball Aerospace did allow access to their National Instruments processing chain including the ADC, FPGA, data storage, and collection tools [Karcher, 2017]. Testing their laser system with various detector chains and algorithms helped provide a fundamental understanding of some of the concepts. Some of these are sketched out here.

First, the ADC has a bit resolution that determines the granularity of the signal measurement. The equation is:

$$\Delta V = \frac{V_{max}}{2^{bit-resolution}} \quad \text{Eqn. 17}$$

, where the ΔV is the voltage step size, V_{max} is the maximum voltage that can be input into the ADC, and the bit resolution is the number of bits on the ADC. The maximum voltage is specified by the manufacturer as well as the bit resolution. The voltage step size matters because the ratios in the DIAL equation will in turn have their own granularity that may introduce error. Most modern ADC chips are both fast and high resolution, so this is not a primary concern. For instance, Texas Instruments sells ADC chips meant for the JESD204B interface with 12-bit resolution at a sampling rate of 3.2 billion samples per second with a 0.8 V peak-to-peak input range for about \$20 [Texas Instruments, 2017]. A 12-bit resolution on 0.8 V means a voltage step of 1.95E-4 V, which is likely to be well below the noise floor.

The JESD204B interface is likely to be the choice for any new development in converters and related signal handling chains [Analog Devices, 2017]. It was designed (after several revisions) to handle the next generation of high resolution, high speed conversion needs. Among other things, it reduces the number of routes on the board, leading to weight, volume, and efficiency improvements. Both Texas Instruments and Analog Devices (two of the largest component developers) have created end-to-end product solutions supporting the interface including ADC, clocks, and FPGAs. Designing the needed layouts for the instrument is (definitely) well beyond the scope of this effort. However, off-the-shelf solutions are available from both providers in breadboard form to test with.

6.10 Lessons Learned from Ball Hardware

Although details cannot be shared, there are some experiences worth sharing from lab trials of different hardware available for testing. In no particular order:

1. Instruments are a team effort. Even a “simple” instrument requires the expertise of multiple people in disciplines from purchasing/accounting to optics to electrical engineering to operations. Working with the Ball team demonstrated how valuable the right personnel can be.
2. Detectors control nearly everything. A lot of the analysis in this dissertation assumes well behaved pulses, both in amplitude and time. The problems come when the bandwidth of the detector is not well matched to the pulses. If the detector stretches pulses in time, then the as-measured peak can fall non-linearly with variations in pulse length. Ball’s suite of detector options allowed for a number of tests where this behavior was observed and figuring out what was going on with the signal was a real challenge.
3. Within reason, longer pulses seem to be better than shorter pulses. Shorter pulses reduce smearing on the ground and atmospheric interaction differences. On the other hand, they make detector bandwidth matching more challenging and percent-based magnitude/time variations are more dramatically impactful. Testing of two lasers with different cavity lengths seemed to point to the longer pulses being more stable.

4. Integrating pulses seems better than taking the peak. The DIAL equations work with any ratio of measurements. Finding the peak value in a pulse seems an easy way to get a measurement. However, pulse stretching from detectors, interaction with the ground, and non-Gaussian noise over short time periods lead to distorted pulses and measurement error. Integrating pulses offers new challenges in timing, but removes many of the problems associated with assuming ideal behaviors.
5. Pointing control is hard. Ball uses a voice-coil mirror, which works great, but it's a complicated assembly of hardware, electronics, and software that an entire team at Ball was dedicated to. This is not something that a one-person developer can hope to do on their own while developing the rest of the instrument. It would be great if a company were to offer an off-the-shelf solution for steering large beams, but that seems unlikely in the near-term at cost points reasonable for a UAV.
6. Little non-linearity issues stack up. There often isn't one big problem, but a series of little problems. Perhaps the detector has a slight non-linear amplitude response. Or the RC circuit's roll-off just slightly impacts a subsequent pulse. Or the aforementioned detector bandwidth response ever so slightly changes the energy profile. Finding and tracking these things can burn a lot of lab time. Solving them takes even more time.

Chapter 7: Baseline Design and Calculation

This chapter uses the information and studies to this point to derive a baseline configuration to work towards. The effort here will drive the hardware design. The discussion starts with a brief section setting performance goals. The second section outlines the approach to meeting that goal. The best-effort implementation of the plan is laid out in Chapter 8.

7.1 Goals and Requirements

In verbal discussions with PHMSA and various industry partners, a goal of consistently finding 50 SCFH leaks is considered valuable [PHMSA, 2015; PG&E, 2015]. It is also desirable to be effective in winds of at least 5 mph, with a 10 mph ceiling being ideal. The FAA flight altitude limit is in effect so the instrument altitude cannot exceed 500 feet. Finally, a double-path, averaged sensitivity of 0.1 ppm is desired, which is less than 10% of a background measurement of 1.8 ppm. For a flight altitude of 500 feet (~150 meters), this is a CPL of 60 ppm.

In a sentence:

“The goal is to reliably measure a 50 SCFH plume in winds up to 10 mph with an instrument flying at an altitude up to 500 feet above the target.”

7.2 Design Process

The first step is to determine a maximum beam diameter using the expected minimum detection threshold of 0.1 ppm. From the plot of Figure 28, it is seen that the beam diameter must be no larger than 3 meters for a 0.1 ppm noise floor. However, that assumes a 1 m/s wind. A 10 mph wind is about 4.5 m/s. This gives an equivalent leak rate of $50/4.5 = 11$ SCFH. Extrapolating down by this factor gives a maximum beam diameter of about 0.9 meters. Smaller would be better.

With a beam diameter no larger than 0.9 meters chosen, the design process can now determine the approximate SNR required of the system to meet the noise floor goal. To do this, the SNR terms of

equation 8 are considered equal for the online and offline paths. While this will not be true, it will return a ballpark number. Standard temperature and pressure values are assumed in order to determine the differential cross section from HITRAN.

The errors from the speckle and reflectivity also need to be accounted for. They take 3.5% off the overall budget, leaving 6.5% for SNR related performance. Then the SNR becomes:

$$SNR = \frac{\sqrt{2}}{\epsilon_R \cdot dOD} \quad \text{Eqn. 19}$$

The differential optical depth is .019 and the error is 0.065. The SNR target is found to be about 1,130.

From here, assumptions need to be made since many parameters affect this value.

The first assumption is that single mode fiber will be used in the instrument to launch the laser. Choosing single mode fiber means minimizing the effect of speckle, makes the beam less divergent before collimation, and keeps the weight and cost lower. This also sets a damage threshold, which sets the maximum laser power in a pulse. Next, a pulse width of 70 ns is also assumed as a reasonable value. This set of information plus the wavelength can be used to determine the maximum allowable pulse energy using the plot in Figure 74. It is found that the peak pulse power can be about 2.1 kW and the peak energy is about 150 mJ.

The next set of calculations is to determine the atmospheric impacts. For solar photons, the irradiance is set to 0.005 W/(m²·nm·sr) which is the vegetation coverage value used by Kiemle et al in their space-based methane instrument analysis at the same target wavelengths [Kiemle, 2011]. A solar filter 10 nm wide is assumed, limiting the number of photons to 0.05 W/(m²·sr). Thermally, the Planck equation estimates a few photons at best for this low wavelength. Next, backscatter is likewise assumed to be zero for these short distances. For transmittance, HITRAN was used to determine that the online and offline values are 0.92 and 0.99 respectively. Since an average value is being determined as a rough estimate of parameters, a value of 0.95 is used here.

The two parameters needed for the ground are the reflectivity and the height difference within the beam's footprint. For the first, a value of 50% is used. This covers a wide variety of materials in the ASTER library and is slightly conservative in most cases. The height difference takes into account pulse stretching as one part of the surface reflects earlier than the other part. For this analysis, a 10-meter difference is used, which roughly simulates a structure on a well pad or a tree along a pipeline.

For the detector, an InGaAs APD is assumed. ThorLabs' APD 410C detector is used as a likely candidate. The maximum responsivity at 1,500 nm is 18 A/W and the minimum NEP is 0.12 pW/Hz^{0.5}. Converting these values to the 1,645 nm region, the values become 17 A/W and 0.13 pW/Hz^{0.5}. The pulse width is nominally slated to be about 70 ns. Using the bandwidth rule of thumb gives a rise time of $0.35/70 = 5$ MHz. This is half the 10 MHz bandwidth of the detector so should work.

The missing piece is the optical efficiency of the system. Given all of the other parameters, this can be solved for. The series of equations and values were programmed in MATLAB. Plugging things in, the total optical efficiency of the system must be about 75%.

For fixed-wing applications, a scan mirror should be used to provide better confidence that an area will be sufficiently mapped. For a quad-copter design, the expectation is that it can simply tip and tilt to point the beam. Both should provide a scan density of about 1 per meter. This means about a 15 Hz scan rate with 2 kHz pulse rates.

In summary, to meet the 50 SCFH plume design goal, the system should have approximately:

1. 70 ns pulses of 150 mW of energy at 2 kHz
2. A beam diameter of no greater than 0.9 meters
3. A SNR of about 1,130.
4. An optical efficiency better than 75%.

Chapter 8: Final Hardware Design

This chapter presents three instrument concepts, analyzes their optical efficiency, predicts their SNR, and then predicts their overall gas detection performance. Of the three systems, two are for a quad-copter platform and one is for a fixed-wing platform with scanning. The fixed-wing and one of the quadcopter platforms have a “traditional” approach and one of the quadcopters has the “VBG Filtering” concept.

8.1 Design Concept 1: Quad-Copter “Traditional” Instrument

The traditional instrument follows the general outline of Figure 69, which is recreated below in Figure 96 with path labels used in this analysis section. This design is about as simple optically as can be designed. The trade-off is that three detectors and the associated power/cabling/ADC/computation assemblies make the system more expensive and software becomes more complex.

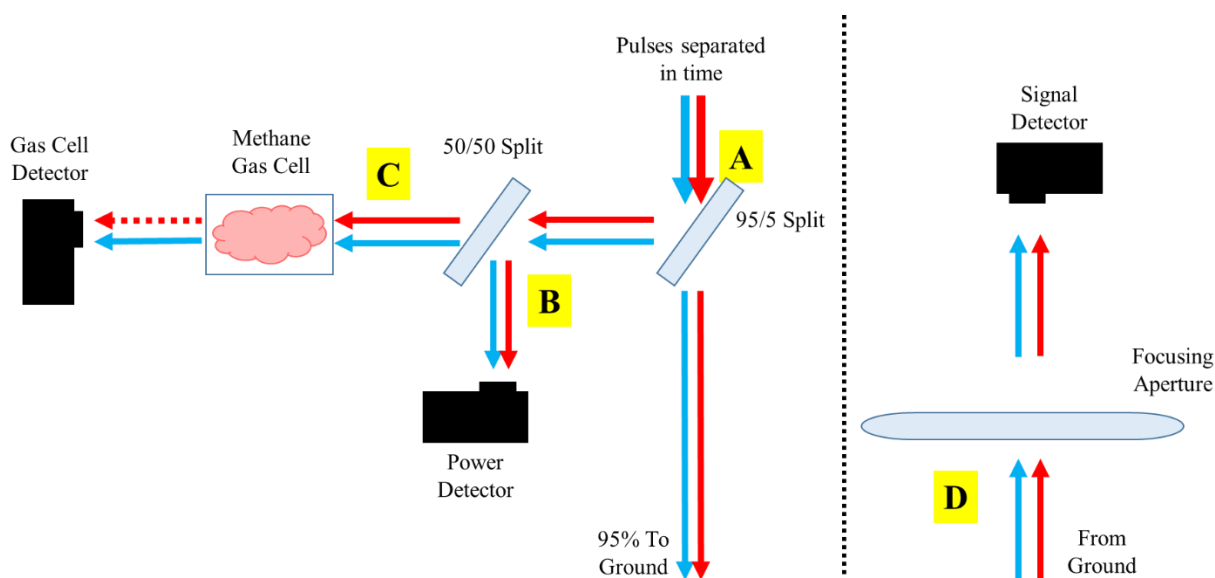


Figure 96: General overview of the traditional design concept. The letters in yellow boxes indicate optical paths discussed in this section.

Path “A” is the first analysis case and follow the light from input to collimation and through the beam splitter, and to the ground. The online and offline pulses are separated in time. An approximate delta time is 600 ns to give 300 ns on either side of a pulse to deal with ring down on the detectors and give

integration window margin. The pulses first interact with an aspheric lens collimator, modeled as ThorLabs AL2550, a 1” diameter optic with a 50-mm focal length. The efficiency from the fiber, through the lens was determined to be 98.34% with OpticStudio modeling Fresnel reflection losses and lens material absorption.

OpticStudio was also used to calculate a beam diameter on the ground of about 15 mm at 150 meters (~500 feet). A trade study showed the aspheric lens can only be about 0.2 mm out of position before the beam diameter on the ground exceeds the 100-mm diameter of the receive telescopes’ collimated view diameter.

From the aspheric lens, the light travels to a beam splitter, which in this case, is a gold-coated, knife-edged prism. As shown in Figure 80, the minimum efficiency is 97.5%. From there the light is bounced off a 45° gold mirror oriented 90° relative to the first mirror to evenly distribute polarization losses. The light reflects from the mirror to the ground. The total path efficiency is $.9834 * .975 * .975 = .9348$, or about 93.5%.

Path “B” follows the power measurement optical path. The fiber launch, collimation, first prism, and second prism have the same optical efficiencies as path “A”. The differences are the addition of a solar filter to block stray sunlight, a neutral density filter to prevent detector saturation, and a final aspheric lens to focus the light onto the detector. Working backwards, the aspheric lens is a ThorLabs AL2520 aspheric lens with a 1” diameter and a 20 mm focal length. In perfect alignment, it will collimate the light from the AL2550 aspheric lens to a spot size of about 8 μm in diameter, which is much smaller than the advertised area of the ThorLabs DET10N detector’s active area of 1 mm active diameter. A trade study in OpticStudio showed the focusing lens could be nearly a mm out of position in the axial direction and by 5 mm in the orthogonal directions before the beam fell outside of the detector area. Figure 97 provides an example result.

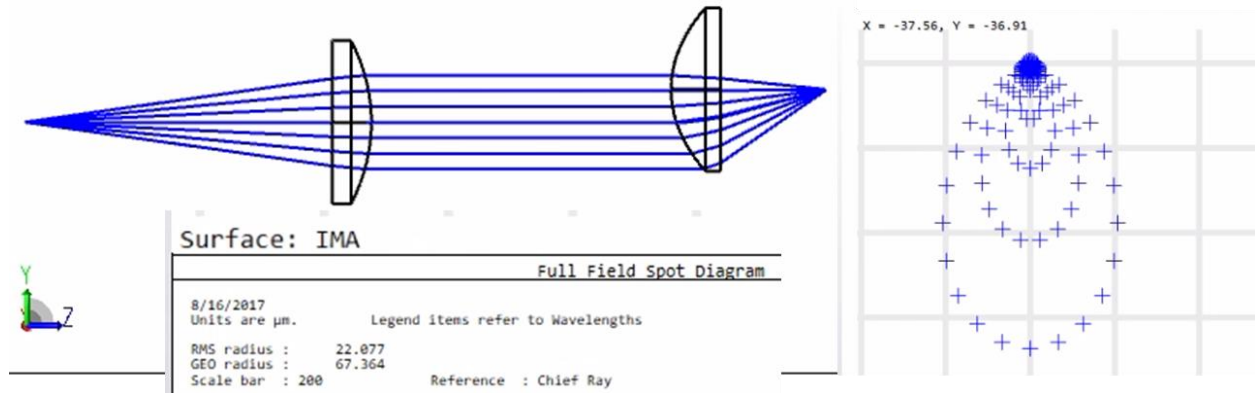


Figure 97: Trade study example of displacing optics in OpticStudio and calculating the resulting focus spot diameter. In this case, the lens is translated by 5 mm and the focus spot radius is 0.67 mm.

Next, the neutral density filter is meant to reduce optical efficiency so is left out of the general efficiency calculations. The solar filter, if well designed, should have less than 0.5% transmission losses. The total path efficiency is thus $.9834 \cdot .975 \cdot .995 \cdot .975 \cdot .9834 = 91.5\%$ before intentional losses from the neutral density filter.

Path “C” is identical to path “B” except for a gas cell filled with methane. The gas cell has four surfaces, of which only the two outside surfaces can be anti-reflection coated. Instead of coating regular glass, a quartz reference cell is assumed which has about 94% transmission in the 1650 nm wavelength region [ThorLabs, 2017]. Borosilicate is a more affordable option by several hundred dollars, although the efficiency is lower at 90%. The efficiency, assuming the higher quality quartz option is $.915 \cdot .94 = 86\%$.

Path “D” follows the light from the ground into the detector. A number of different optic designs were considered, including plano-convex lenses and aspheric lenses of various sizes and focal lengths. Two example focal spot studies are shown in Figure 98. Ultimately, a 100 mm diameter aspheric lens with a 100 mm focal length was chosen. This diameter was chosen based on the SNR guidance from section 8.2. The aspheric lens was chosen based on its superior ability to focus light on the 0.2 mm diameter detector relative to the plano-convex options. The 100 mm focal length was chosen because it offered sufficient focusing performance in the smallest volume possible.

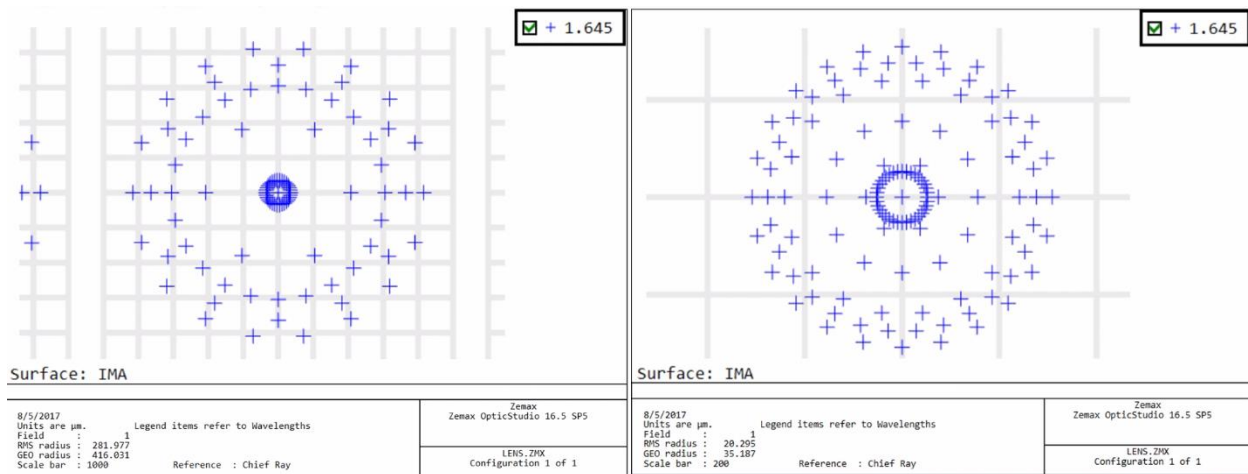


Figure 98: Zemax Optics Studio plots of the focal spot for the 125 mm Edmund Optics 26-506 convex-plano optic with a 375 mm focal length (left) and the 100 mm Asphericon ALL100-100-S-U aspheric optic with a 100 mm focal length (right). The best-case spot radii are 416 and 35 μm respectively.

The primary drawback of this lens is the cost with prices ranging from \$850 to \$1,100 depending on the vendor, while plano-convex options were roughly an order of magnitude less expensive. However, OpticStudio analysis of the light collected showed that the low performance with a plano-convex would drive a need for a separate amplifier, which could cost over \$1,000 from a company like Femto. The decision was made to go with the more expensive optic over additional electronics.

The final optic in this sub-assembly is a solar filter. The filter is designed to be high-pass with a cut-off as close to 1.645 μm as possible. The natural and steep reduction in quantum efficiency of an InGaAs detector limits background photon collection above the cut-off. A well-designed filter should have less than 0.5% efficiency loss.

Calculating the efficiency of this path in OpticStudio, a value of 91.63% was found if ThorLabs “C” coating was assumed. The focal spot radius when the system is in perfect alignment is an optimistic 0.3 μm . A trade study on positioning found the detector could be out of alignment in the axial direction by about 0.3 mm before the spot size grew to overfill the detector’s active area.

For detector efficiency, there are two types being used: ThorLabs DET10N and APD410C. The former is a biased InGaAs detector with a 70 MHz bandwidth, a 0.8 mm² active detector area, a nominal NEP of 0.02 pW/Hz^{1/2}, and a dark current of 1.5 nA. No amplifier is used as the signals are high enough that neutral density filters are expected to be used to prevent saturation. The DET20C and DET10C were considered, but the first has 55 nA of dark current and a larger detector area while the second had higher NEP and junction capacitance.

The APD410C is an avalanche photodiode with a variable-gain, built-in amplifier to scale up to 20x the input. The active area diameter is 0.2 mm, the bandwidth is 10 MHz, and the NEP is 0.12 pW/Hz^{1/2}.

There is an APD430C which has a higher bandwidth of 400 MHz that increases the NEP to 0.45 pW/Hz^{1/2}. The higher bandwidth allows for less pulse distortion at the expense of higher noise. However, if the background is measured and the noise is Gaussian in nature then the primary drawback is the reduction in dynamic range for the ADC path. Both detectors are InGaAs and so have quantum efficiency of about 70% in the 1.645 μ m region. Accounting for wavelength variations, the NEP for the two detectors is 0.13 pW/Hz^{1/2}.

The summary of the various efficiencies is provided in Table 7. The average efficiency for a well-assembled system is expected to be between 55% and 60% efficient. For the gas cell and power path, this efficiency is not strongly relevant since attenuation will intentionally reduce photon count. Discussions at the next level of design will cover this in more detail. For the transmit path, the optical efficiency is more critical. A 60% efficiency is lower than hoped for. However, each individual component is about as high as can be achieved at this wavelength.

Table 7: List of individual and cumulative losses in the optical system. In cases where polarization differentially affects losses, the minimum value of the two is displayed.

Optical Assembly	Efficiency (%)	Cumulative Efficiency (%)		
		Transmit Path	Gas Cell Path	Power Path
Fiber Input to Launch	93.5	93.5	-	-
Fiber Input to Gas Cell Fiber	86.0	-	86.0	-
Fiber Input to Power Fiber	91.5	-	-	91.5
Receive Optic to Detector	91.6	85.6	78.8	83.8
Detector Efficiency	70.0	60.0	55.1	58.6
Total Optical Efficiency		60.0	55.1	58.6

Representative views of the hardware are provided in Figures 99-101. The assembly is about 11”x8”x7”.

The optic paths of the gas cell and power are joined by 1” optical tubes both to block stray light and to provide stiffness. The receive telescope is designed into a tube for stiffness reasons as well. All of the subassemblies are mounted with three-point contacts and spherical washers to minimize the potential of over constraining assemblies that may relax unpredictably in a vibration environment. Not modeled are the laser and the electronics. The exception is the Applanix APX-15 IMU designed for UAVs. The volume in the mechanical model is based on their dimensions. The IMU allows the quad-copter vehicle to act as the pointing and control mechanism with tip/tilt maneuvering with the engines.

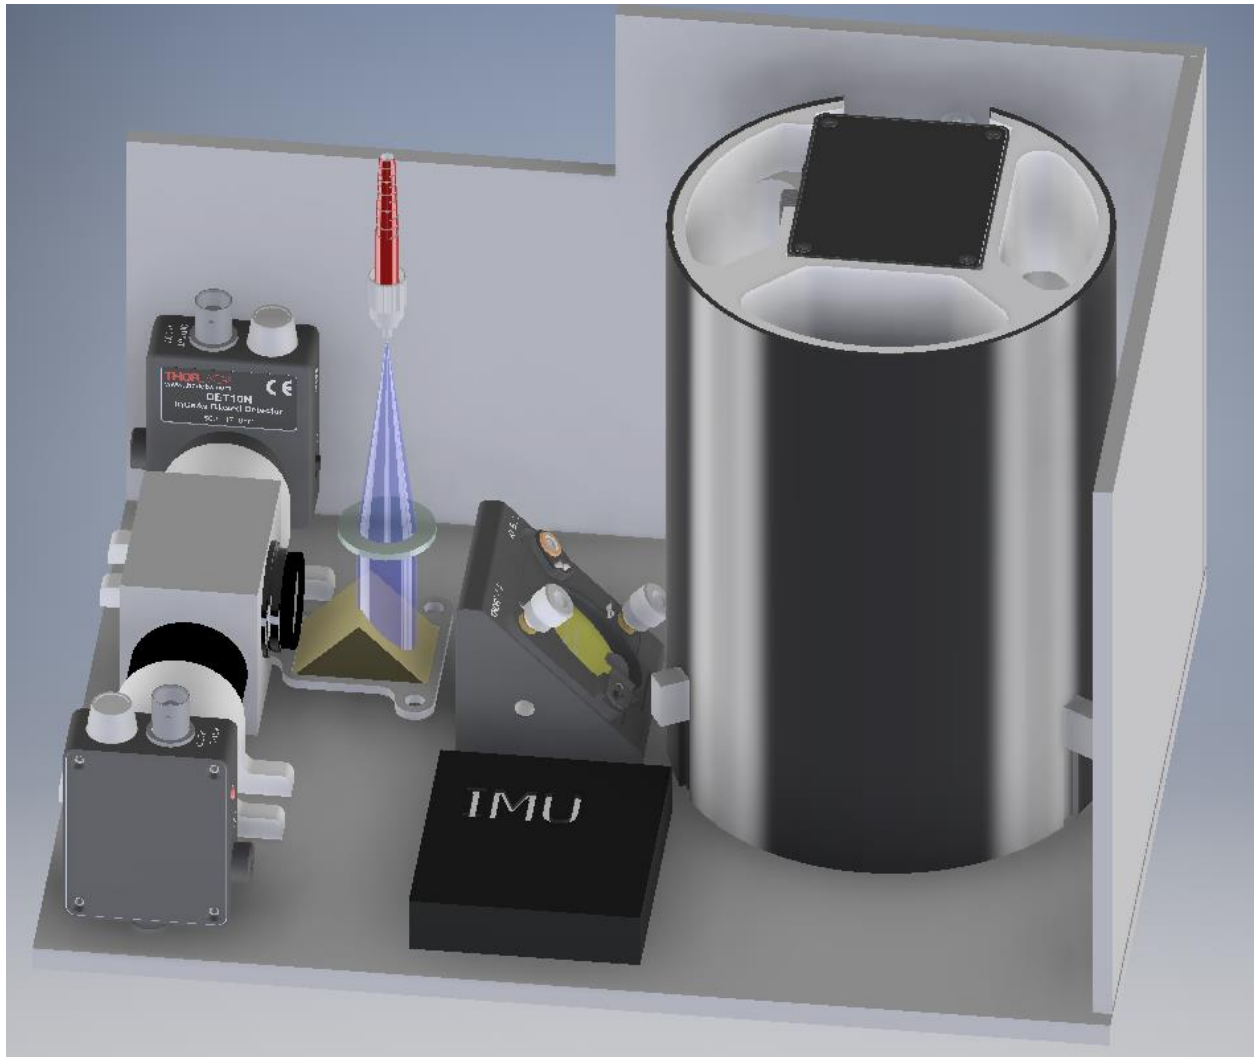


Figure 99: View of the optical layout for the quad-copter design. The base plate is 11"x8". The assembly is 7" tall. The receiver telescope optic is ~4" (100 mm) in diameter.

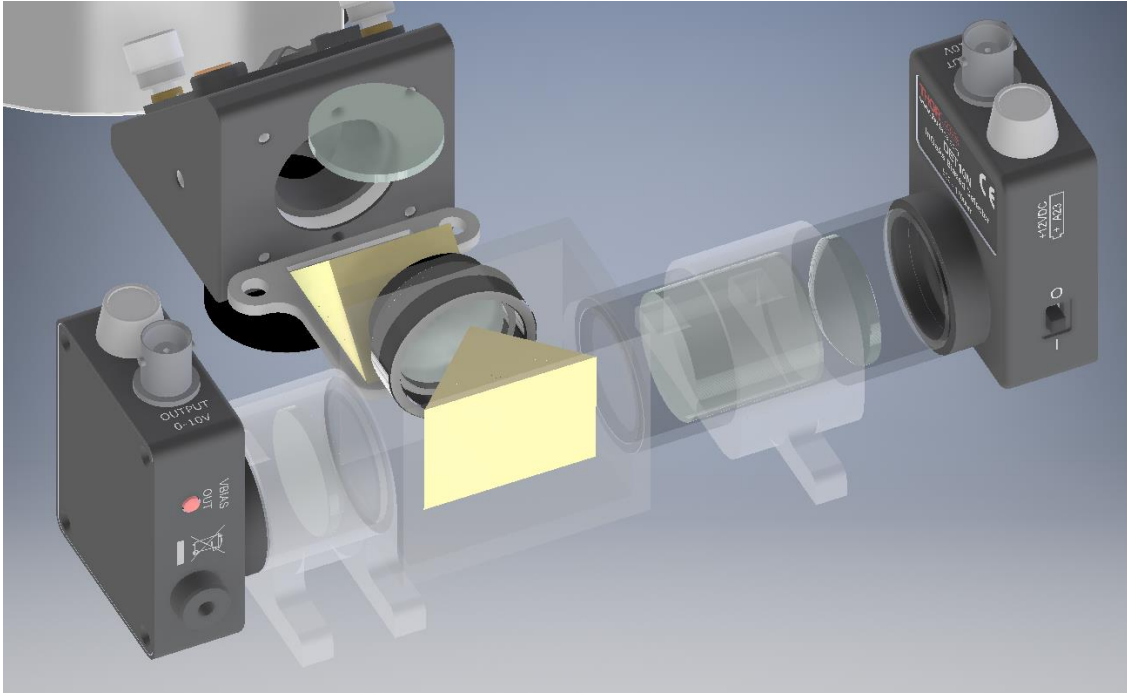


Figure 100: View of the beam splitting assemblies for the transmit, gas cell, and power paths. The supporting structures are set to semi-transparent.

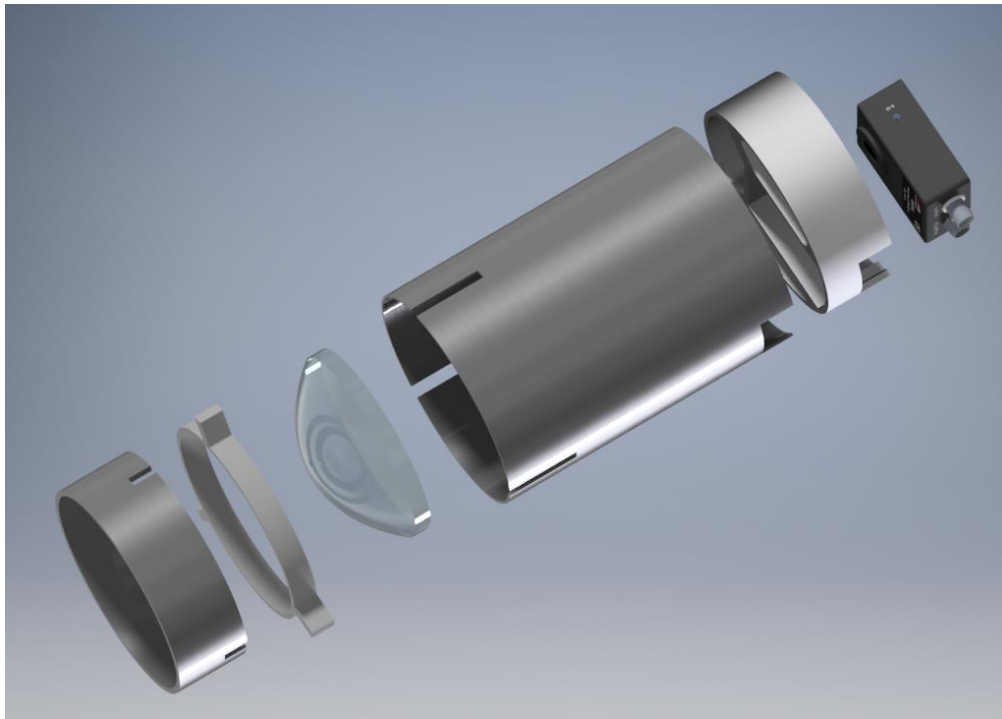


Figure 101: The receive telescope assembly. The lens mounts to a ring with three tabs. The tabs fit into the lower left tube and are epoxied in place after alignment. The longer tube slides over this sub-assembly. A bracket holds the detector in place and slides into the tube where it is also epoxied in place. Fasteners may be a better attachment option to facilitate modifications.

8.2 Design Concept 2: Fixed-Wing “Traditional” Instrument with Scanning

The second design concept modifies the first by extending it to a fixed-wing platform. In lieu of the vehicle acting as the pointing control system, a scan mirror is installed.

For the purposes of optical efficiency, though, this system is identical to the first. The scan mirror is gold-plated and replaces the fixed gold-plated mirror from the quad-copter concept. The optical path is also a little longer for the transmitted beam as it exits the instrument, but a few inches compared to several hundred feet is a negligible impact.

An outline of how the components are assembled is provided in Figure 102. For the purposes of showing the optical path, a tube is not shown on the receive detector and optic. The tube, similar to the previous design, offers stiffness and stray light protection.

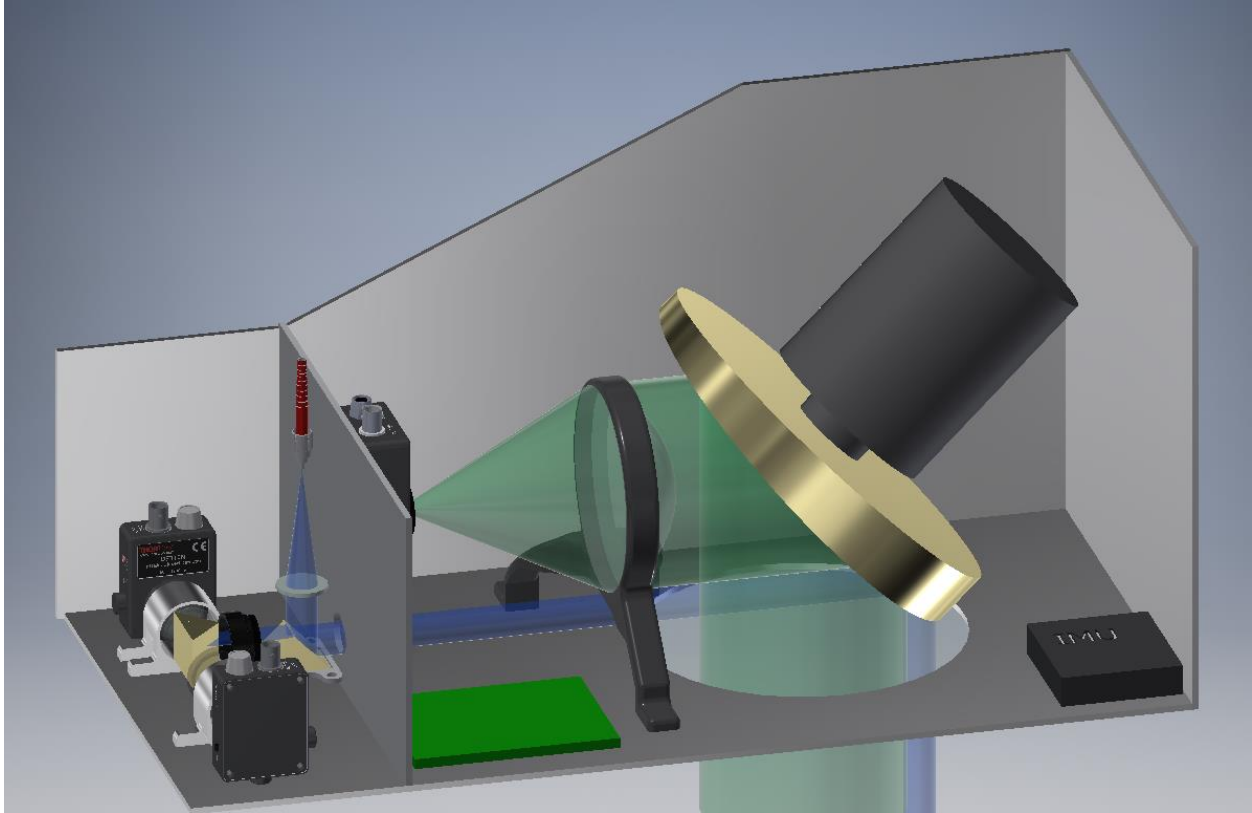


Figure 102: An outline of how the fixed-wing instrument could be assembled. The dimensions are 9.5" wide, 19.5" long, and 9.5" tall at the tallest end (5.5" at the shorter end). The transmitted beam (blue) travels under the received beam (green). The transmitted beam strikes the outer edge of the scan mirror which has subtle angles machined into it to reflect the transmitted beam into the receive telescope's footprint. The mirror rotates with a larger angle to reflect the light in a circular pattern on the ground.

8.3 Design Concept 3: VBG Filtering

The third concept attempts to push the design envelope to gain system efficiencies. The theory is that if both the online and offline pulse could be sent at the same time then beam overlap is perfect. This could mean reductions in reflectivity variations. In a multi-mode fiber system, this would reduce variations due to speckle too. It may be possible to use a somewhat broad band laser source as well, perhaps reducing laser cost. With the right design, it would be possible to use just a single detector instead of three (one for each path), reducing cost, weight, and volume mechanically. Electrically, just a single ADC path would be required. The possible advantages are tempered by the difficult in temporally separating the pulses before the detector. However, the potential advantages were considered worthwhile enough to further

explore the idea. The outline concept in Figure 70 is reproduced in Figure 103 with interface letters to more easily trace the analysis in this section.

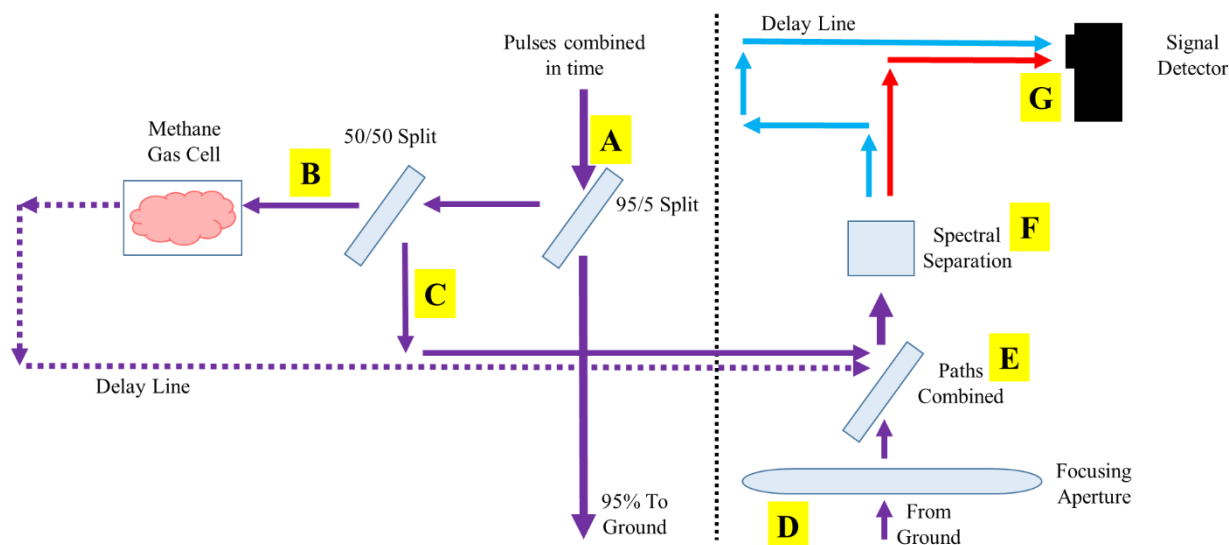


Figure 103: General overview of the VBG filtering design concept. The letters in yellow boxes indicate optical paths discussed in this section.

Analysis starts with the incoming beam path splitting at “A”. Three paths diverge from this point: the pulses sent to the ground, the pulses used to measure power, and the pulses used to measure the gas cell. Path “A” follows the energy from input until transmission into the atmosphere. These optics set the beam spot diameter on the ground so limit the flight altitude of the instrument. Numerous trade studies were run on the types of optics, their positioning/path lengths, and their impact on beam size. The final solution was found with OpticStudio optimization set to minimize divergence.

The result is that no GRIN lens is used on the optic fiber input. Instead, the 10 μm mode field diameter is allowed to expand at the fiber's NA of 0.14. An aspheric lens (ThorLabs AL108) is used to collimate the beam to a diameter of 2.0 mm. The beam is then split by a knife-edge prism (not modeled) to pickoff energy for the monitor assembly. The bulk of the light is reflected towards a plano-convex lens (LC1975-C) to expand the beam. A plano-convex lens (LA1509-C) with a 100 mm focal length is used to collimate the beam to a launching diameter of ~ 6.0 mm. Simulating full field ray traces with polarization effects

creates a beam diameter of about 14 mm at a range of 150 m (~500 feet). This diameter will easily allow a 100 mm telescope optic on the receiver with no overflow of the detector.

The beam divergence is low due to two influences: first, the beam is expanded with fairly low NA by 3+ orders of magnitude and, second, the lenses are perfectly positioned in the software. To more fairly estimate performance, a sensitivity study was undertaken where the optics were pushed out of position by 0.2 mm in any direction and/or rotated by 2 mrad. It was found that the beam spot size could expand to nearly a meter if the first aspheric lens from the fiber was allowed to be translated out of position. The beam expander and collimating lenses showed minimal impact of less than 5% for the same movements. The conclusion was that the aspheric lens should be an assembled unit, similar to ThorLabs F240FC-1550. This is a fiber-fixed, aspheric system designed for a 1,550 nm wavelength and an 8.18 mm focal length. This part is insufficient due to an 8 mm focal length expanding the beam too large for GRIN lenses in the optical path as well as being the wrong wavelength. ThorLabs and Newport offer custom options along these lines or the assembly could be built with the lens component from the optimization study.

The optical efficiency for the optimized system, not accounting for the knife-edge prism and mirror losses, is 95.54% assuming ThorLabs “C” coatings on all surfaces. Two gold-coated mirrors are used in the path, each with an efficiency of about 97.5% regardless of polarization. The final efficiency is thus $0.9854 \times 0.975 \times 0.975 = 93.6\%$. Figure 104 provides a view of the results from OpticStudio. A rendering of the beam splitting prisms and the three paths are in Figure 105.

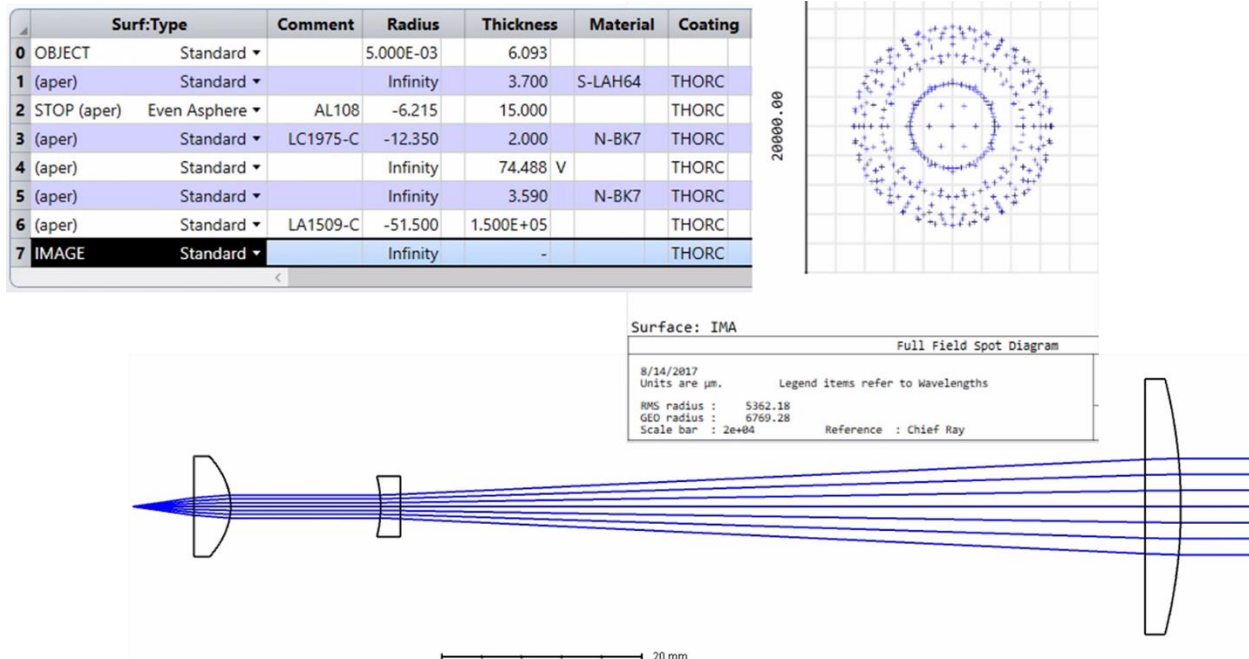


Figure 104: OpticStudio rendering of the launch optics. A single mode fiber with a NA of 0.14 is collimated with an aspheric lens before interacting with a gold knife-edge prism (not shown). A plano-concave lens then diverges the beam until a plano-convex optic with a long focal length collimates the expanded beam. A perfectly aligned system creates a beam size on the ground, accounting for polarization effects, of about 14 mm in diameter at a range of 150 m (~500 feet). Minor variations of two tenths of a millimeter from the ideal alignment can expand the beam to 88 mm in diameter at 50 meters (~150 feet)

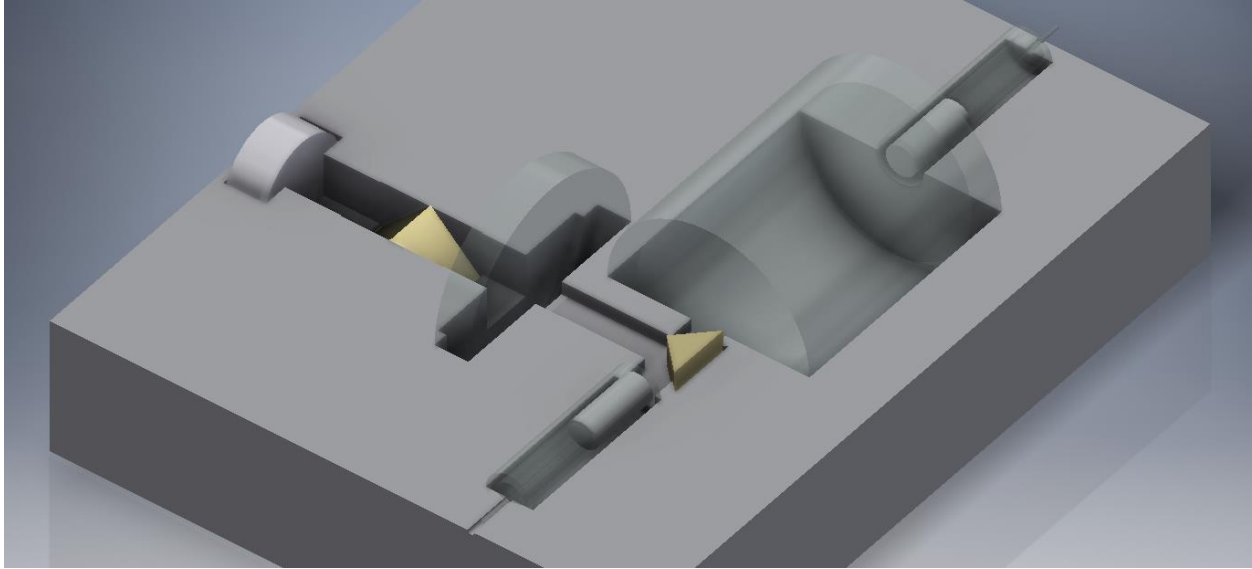


Figure 105: View of the light splitting and monitor box. A collimated beam is split by the gold prism on the left with the majority of the light being sent to a beam expanding optic on the left. A small percentage is sent through a neutral density filter before being split by the small prism on the right. One half of the path goes into a GRIN optic for focusing onto a single mode fiber. The other half goes through a gas cell before being captured by another GRIN/fiber combination.

This first sub-assembly, as with all sub-assemblies in this design, is built on its own optics bench. This approach should allow careful integration of the components and lock them in place irrespective of most of the other components. It also allows for more exotic and expensive bench material selection in case a material with a higher stiffness or lower CTE is desired. However, unlike the other sub-assemblies, this box does feature a free-space optical input, which will require alignment. The other assemblies benefit from all-fiber inputs and outputs.

Transitioning to interface “B”, it follows the laser input through beam splitting, through the gas cell, and then into a GRIN/fiber assembly. The primary consideration here is ensuring the beam does not expand beyond the focusing abilities of the GRIN lens. The beam coming from the aspheric fiber collimator at the input point has a diameter of about 2 mm, which exceeds the 1.8 mm limit as-is. However, the light going through this path is first cut from one side of the original diameter by the first prism, then cut in half by the second prism.

The nominal efficiency of this design is based on estimates. The aspheric lens has an efficiency of about 99.75% for the ThorLabs “C” coating. The gold prisms are assumed to have an efficiency of 97.5% each. The optical efficiency of the GRIN assembly is expected to be better than 95% based on the near-zero divergence angle from the aspheric lens collimation. As with the previous assemblies, a quartz gas cell is used. The unknown in the path is the neutral density filter with its purposeful reduction in optical efficiency. Everything else combines for an optical efficiency of about 85.6%.

Path “C” follows the transmitted power measurement energy. The path is the same as the gas cell sans the cell. The efficiency is 90.1%.

Path “D” is the receive telescope optic to fiber coupling path. Using OpticStudio, a trade study was conducted on the maximum coupling efficiency for a “perfect” 100 mm diameter aspheric lens with a 200 mm focus. The NA is 0.24 for the lens. Polarization effects were accounted for including Fresnel reflection losses and absorption of the N-BK7 glass. ThorLabs “C” coating for the near infrared was assumed for all surfaces. The minimum waist diameter for the optic is about 2.5 μm . The best performance a single mode fiber can achieve with this lens is 50.3% due to single mode fibers having maximum NAs of ~ 0.14 . Even if the NA of the fiber could be increased to 0.25 to match the optic, the maximum efficiency is about 65.5% due to losing the other modes. Investigating multimode fibers, the maximum theoretical efficiency is 88.45%. Losses are due to polarization effects and the “C” coating not being optimized for the 1645 μm wavelength.

The closest suitable fiber found for purchase is ThorLabs’ GIF625 graded-index multimode fiber. It has a diameter of 62.5 μm and a NA of 0.275 ± 0.015 giving a minimum of NA of ~ 0.26 , which exceeds the lens’ NA of 0.24. The cost is about \$2.20 per meter.

A trade study on alignment sensitivity showed that the interface is not tolerant of more than a couple hundredths of a mm in side-to-side motion and about a tenth of a millimeter in axial movement. The mechanical support was designed with this in mind. A carbon fiber tube is used to increase the moment of

inertia as high as possible and aluminum structure is used at both ends to help keep the tube round. The fiber itself is held in place with a three-point mount on spherical washers to help lock the assembly in place. It is also possible to add braces to the side of the tube. A full design would look at fundamental structural frequencies to ensure dynamic motions are within tolerances. A render of the sub-assembly is shown in Figure 106.

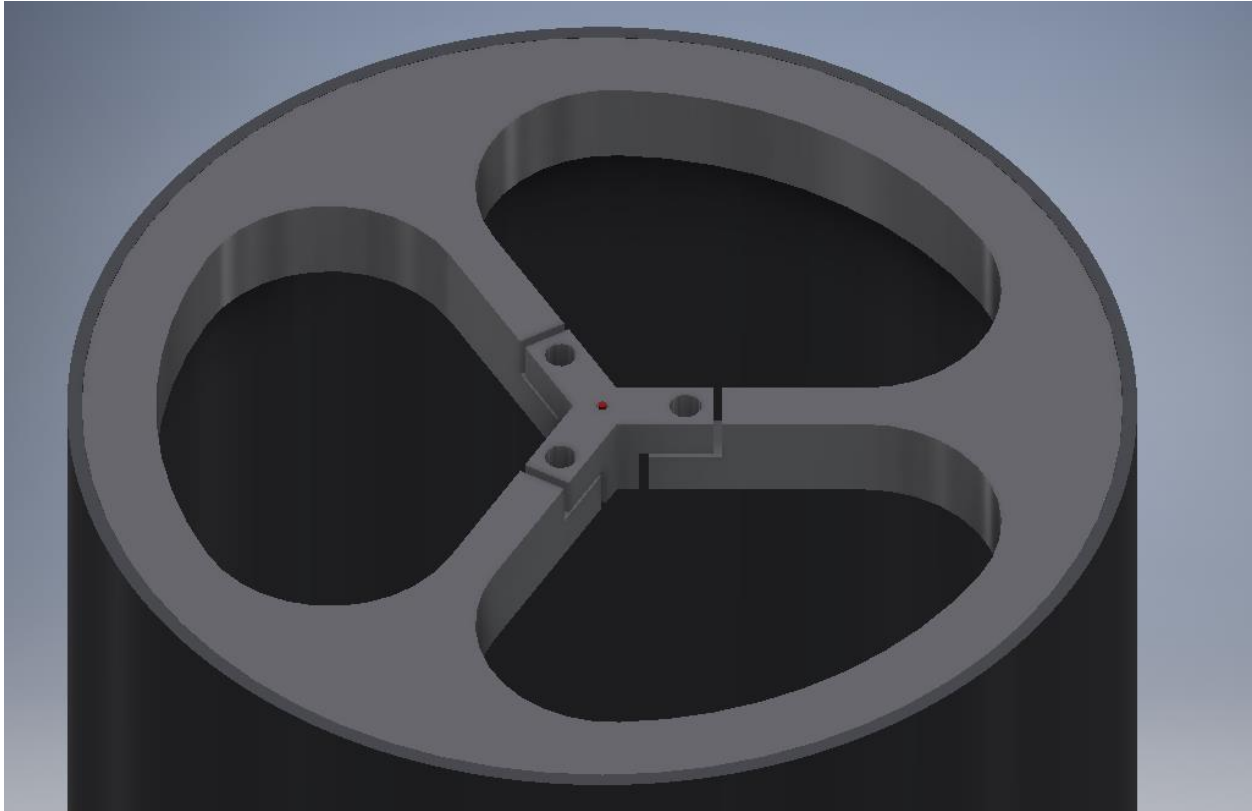


Figure 106: The fiber from the receive telescope is held in place by a three-point mount. Spherical washers are meant to be used between the mount and support. The real outer support would not have the three large holes in order to increase stiffness and block stray light.

Analysis is now at path “E”. Three fibers carry the three sets of pulses and the instrument is on the threshold of filtering with the VBGs. The next step is to combine the three sets of energy into a single fiber. This is important because the collimated beam diameter from a GRIN lens for the $62.5\text{ }\mu\text{m}$ multi-mode fiber at this wavelength is 1.1 mm and the GRIN lens itself is 1.8 mm in diameter. Three of these

beams together would mean the output from the VBG would be too large to capture. In this case, the combination of three fibers into one is done with an aspheric lens.

The ThorLabs “High Precision” AL1225H aspheric lens is a 12.5 mm diameter optic with a 25 mm focal length. The same multimode fiber is used on the focusing side. Each incoming GRIN assembly is offset from the center point of the aspheric lens by 2.0 mm. Taking into account the collimated beams are about 1.1 mm in diameter, the combined beam diameter is 5.2 mm, including the gap, which gives an effective NA on the 25 mm focal length of 0.10, which easily fits within the 0.26 NA of the fiber. Using OpticStudio, the coupling efficiency for well-positioned fiber and optics is 96%, assuming that all surfaces are coated with ThorLabs “C” coating. This value agrees well with advertised GRIN losses of less than 0.35 dB. A rendering of the analysis is provided in Figure 107. A mechanical rendering is provided in Figure 108. Of note is that the beam diameter only leaves about 10 μm of translation margin, further stacking the alignment challenges. Another option is to go with a larger diameter fiber.

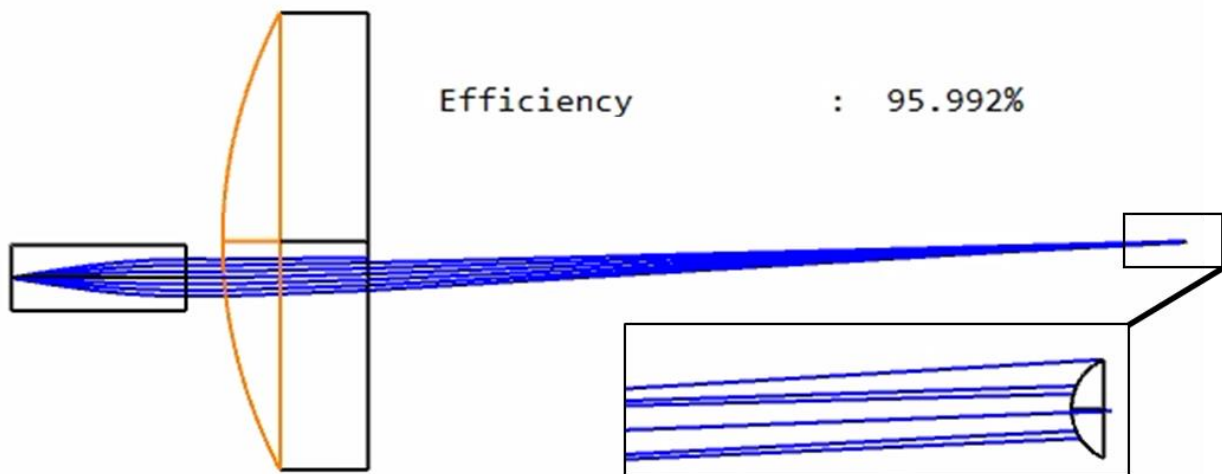


Figure 107: OpticStudio rendering of launching a 0.26 NA fiber into a GRIN collimator then into an offset aspheric lens and finally into another 0.26 NA fiber. The fiber diameter on both sides is 62.5 μm . ThorLabs “C” coatings are assumed on all surfaces and polarization effects have been accounted for.

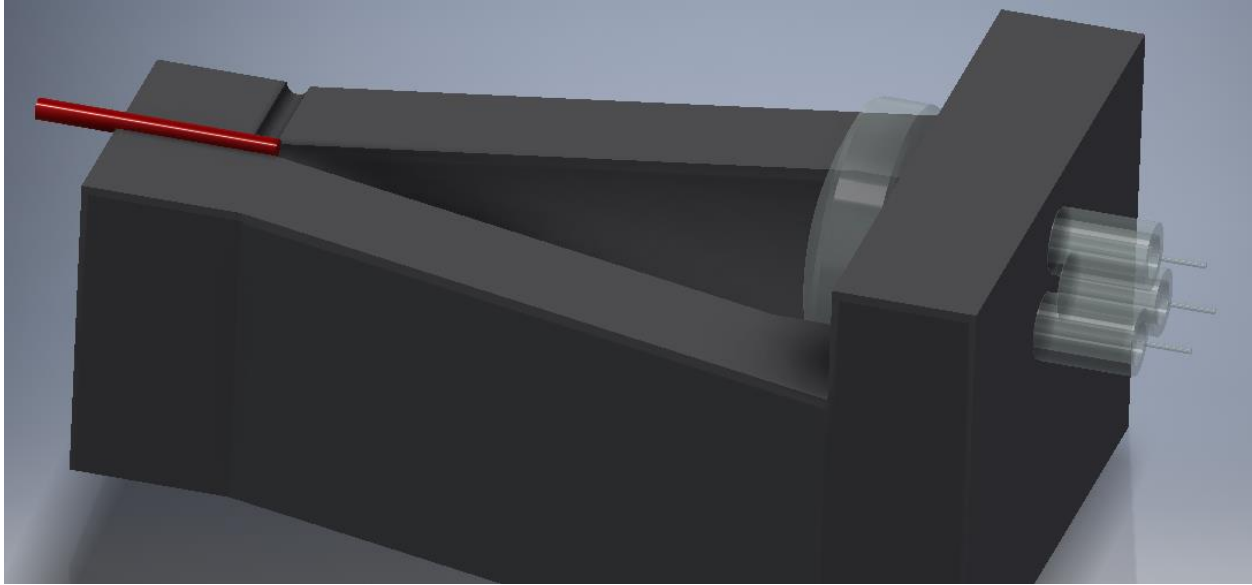


Figure 108: View of the path combining assembly. Three GRIN collimating lenses feed into an aspheric lens that focuses onto a multi-mode fiber. The notch near the fiber (and a hidden hole under the fiber) are meant for aligning the fiber before being staked in place with epoxy.

Analysis moves to path “F”. The instrument now has all paths combined and is ready for filtering by the VBGs. One challenge is that GRIN assemblies have coupling efficiencies tied to the free-space distance between them. A path length less than 25 mm gives better than 95% efficiency, but it quickly falls off past that distance. The VBG has a length of 45 mm. Further, the glass casing used in the GRIN assemblies to couple the lens to a pig-tailed fiber is wide enough that the narrow 10° beam separation needs a long path length before two assemblies can be close enough not to interfere with each other. A solution used to minimize the distance is to not use the glass-casing and instead mount both the lens and fiber ferrule in the optics bench directly, saving about 10 mm of path length.

The total path length from input to output is ~ 50 mm, which means expected efficiency of around 84% [ThorLabs, 2017]. The as-purchased VBGs from OptiGrate have reflectivity losses per surface of less than 0.5%. The total efficiency is, theoretically, about 83.5% although this does not account for diffraction losses. Testing will be required to determine the true value, which may be meaningfully lower. Also, two of these assemblies are used in series for the offline path so the total efficiency for the VBG

filtering is about 69.8%. (Note that the online wavelength only goes through one VBG assembly so is more efficient.)

Unique to this assembly is the inclusion of a TEC to keep the VBG tuned to the appropriate wavelengths. The top surface of a 2A TEC is flush against a cut-out on the bottom of the optics bench. There is just 1 mm of material between the TEC and the VBG to keep thermal capacitance low. The VBG is not mounted directly to the TEC because a) the TEC is not optically flat or controlled and b) the thermal paste on the TEC will not migrate onto the expensive glass. The bottom of the TEC is mounted to a copper surface for heat spreading. The expectation is that the TEC will only need to drive the VBG a couple of degrees in most conditions so exotic cooling/heating solutions are not required. Finally, the optics bench does not contact the copper heat spreader to minimize heat transfer between the two. Instead, Titanium or Ultem-type spacers are used. Figure 109 provides a rendering of this sub-assembly.

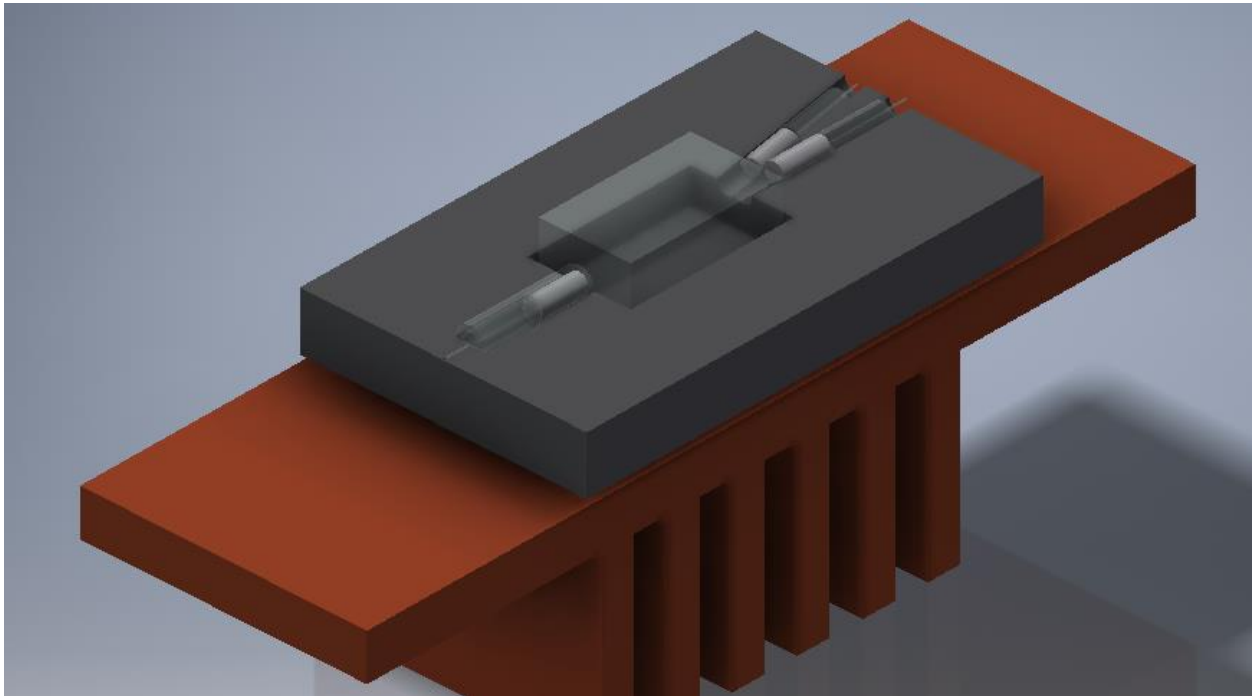


Figure 109: View of the VBG wavelength splitting assembly. Multiple wavelength light is introduced from the left by a collimating GRIN assembly. The VBG splits the wavelength of interest off at an angle of 10° into a GRIN/fiber assembly. The remaining light is captured by a GRIN/fiber assembly for further filtering downstream. The assembly is mounted on a TEC for temperature control of the VBG. The TEC itself is mounted on a copper heat spreader. Titanium or Ultem spacers are used between the copper and optics bench to reduce heat transfer.

The pulses are measured in path “G”. The goal here is to get the light from the online and offline pulses into the detector. Similar to the previous path combining assembly (Path “E”), GRIN assemblies collimate light that is focused via an aspheric lens onto the detector. The ThorLabs APD410C and 430C have a detector diameter of 0.2 mm, although they suggest staying away from the edges to minimize non-linearity. The AL2520-C 1” lens has a focal length of 20 mm and a NA of 0.54. The optical efficiency is 96% based on the OpticStudio analysis of path “E”. Either detector’s quantum efficiency is 70% based on the InGaAs detector discussion. Figure 110 shows the mechanical assembly, which uses a part directly threading into a 1” diameter and 1” long optic tube to hold the GRIN assemblies. The aspheric lens is also within the tube. The tube threads onto the detector.



Figure 110: Light enters the detector from two GRIN/fiber assemblies. One path is the “online” wavelength, one is the “offline” wavelength. An aspheric lens focuses the light from both optics onto the detector. The GRIN/fiber assemblies are mounted on a 1” diameter part that threads into an optic tube mounted directly to the detector for maximum stiffness.

The final discussion is on the fiber optics. Separation of the pulses requires both wavelength and temporal groupings. The temporal separation is done by using longer lengths of fiber optics on some paths. The goal is to separate the pulses in time by at least 6x a pulse’s expected width. This value is based on the

pulse integration discussion in section 6.7. If the pulses are 100 ns in length (to be conservative), then 600 ns of time is required between pulses. Then, there are three groups of these pulses that also must be separated in time.

Sending pulses to the ground and back introduces timing uncertainty. With the rough rule of thumb that 1 ns is about 1 foot of light traveling distance, the two-way path could be anywhere between 200 and 1,000 ns (100-500 feet of altitude). The upper bound of 1,000 ns is assumed to be conservative. However, if 600 ns of time is required between pulses, then the power and gas cell measurements cannot be squeezed in before the return pulse. This means that the monitoring pulses are measured after the return pulses come back. Rolling all of this into a final measurement time, the 6 total pulses, each 100 ns wide and with 600 ns of time between them and assuming 1,000 ns for the ground return pulse time, the total delay time is ~4,600 ns.

To a first order, these delays means more than a half-mile of optical fiber! This is a lot of glass, but the small diameter keeps the weight and volume lower than what might be expected from sticker shock. For instance, ThorLabs sells 1,000 meters of single mode fiber for ~\$475 and it weighs less than 5 lbf [ThorLabs, 2017]. Multi-mode fiber is more expensive at about \$5.50 per meter from either ThorLabs or Newport, although shorter lengths of multi-mode are required than the single mode fiber initially carrying the power and gas cell pulses.

The longest path is the gas cell's online pulse. This pulse travels 1,600 ns before the path combining assembly and then a further 700 ns to accommodate the gas cell's offline pulse. 2,300 feet of optical fiber is about 700 meters. ThorLabs advertises fiber losses of about 0.9 dB/km (~8%), which would mean this pulse would lose about 7.1% of its original energy. The power cell is next with about 500 meters and 4.0% loss. The transmitted pulses lose the least with ~180 meters of fiber optic and 1.4% loss. Although the online and offline lose separate amounts, the higher losses of the two are reported.

There are also options for reducing the lengths. For instance, choosing to fly no more than 200 feet above the ground could cut hundreds of feet of fiber (although the tip/tilt of the quadcopter means path lengths will be more than 200 feet). The time difference between pulses could also be reduced from 600 ns to perhaps 300 ns or less with tight control of the integration hardware. Finally, the gas cell and power paths could also be combined earlier instead of making both travel 1,000 feet on separate fibers before being combined with the received pulses. All of these solutions could improve the optical efficiency as well as saving money, weight, and volume. However, the overall concept is on the feasibility of the VBG filtering and the fiber's optical losses are not the limiting factor on overall performance.

The summary table of the various optical paths is Table 8. The cumulative efficiency of the three pulse paths (transmitted, gas cell, and power) all come in around the mid-30% levels. These numbers compare to values of around 55% for the traditional design, nearly a 40% decrease.

Table 8: List of individual and cumulative losses in the optical system. In cases where polarization differentially affects losses, the minimum value of the two is displayed.

Optical Assembly	Efficiency (%)	Cumulative Efficiency (%)		
		Transmit Path	Gas Cell Path	Power Path
Fiber Input to Launch	93.6	93.6	-	-
Fiber Input to Gas Cell Fiber	85.6	-	85.6	-
Fiber Input to Power Fiber	90.1	-	-	90.1
Receive Optic to Fiber	88.45	82.7	75.7	79.7
3-Fiber Input to Single Fiber	96.0	79.4	72.6	76.5
Fiber Input to VBG to Fibers	69.8	55.4	50.6	53.4
Fibers to Detector	96.0	53.1	48.5	51.2
Detector Efficiency	70.0	37.2	34.0	35.8
Fiber Optics	98.6/92.9/96.0	36.6	31.6	34.4
Total Optical Efficiency		36.6	31.6	34.4

A few renderings of the mechanical layout are in Figures 111-113. The assembly is small with room left over for computing hardware and space for coiling fiber optics.

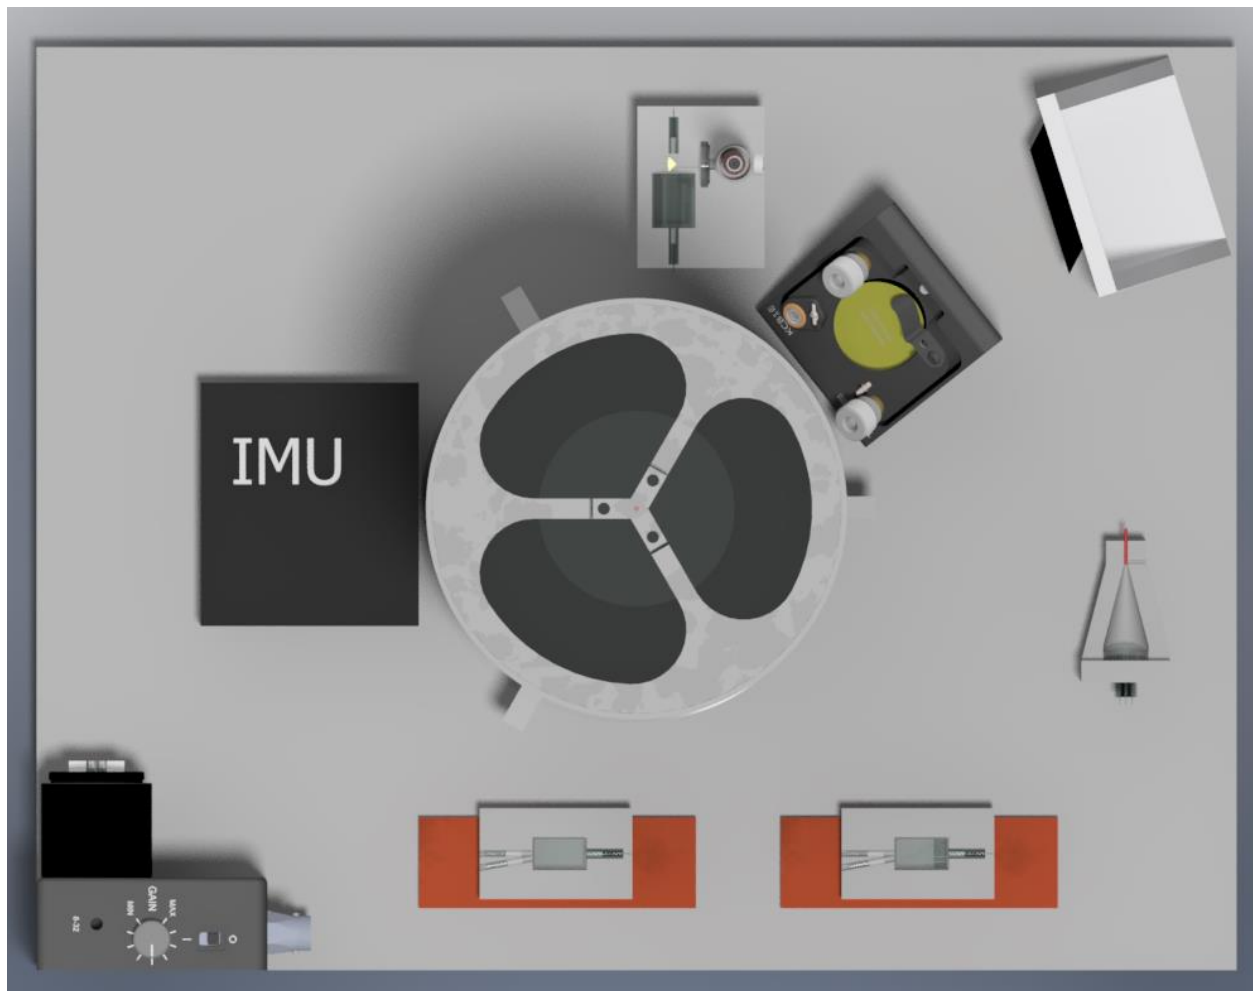


Figure 111: Overhead view of the concept instrument. The platform is 13" wide and 10" tall. The structure is 4" (100mm) in diameter and 9.8" tall.

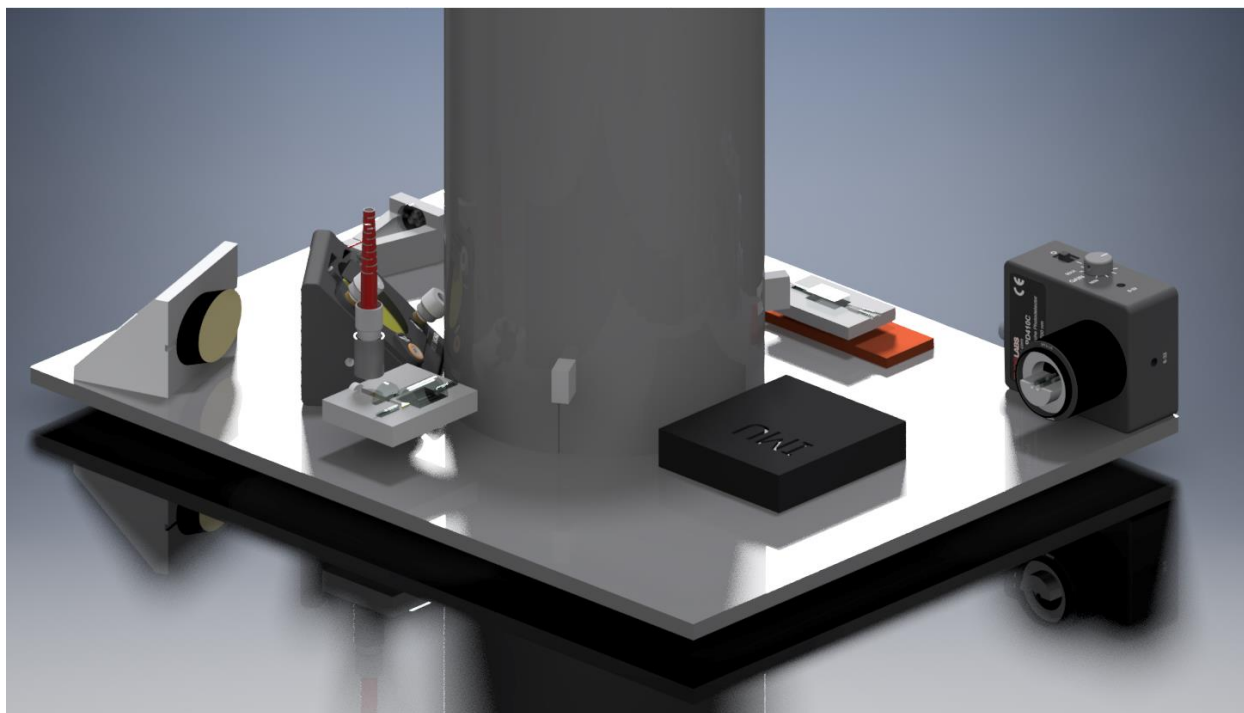


Figure 112: Side view of the concept instrument. The center optical structure is a tube for maximum stiffness and light isolation. It has a height of about 9.8" (248 mm)

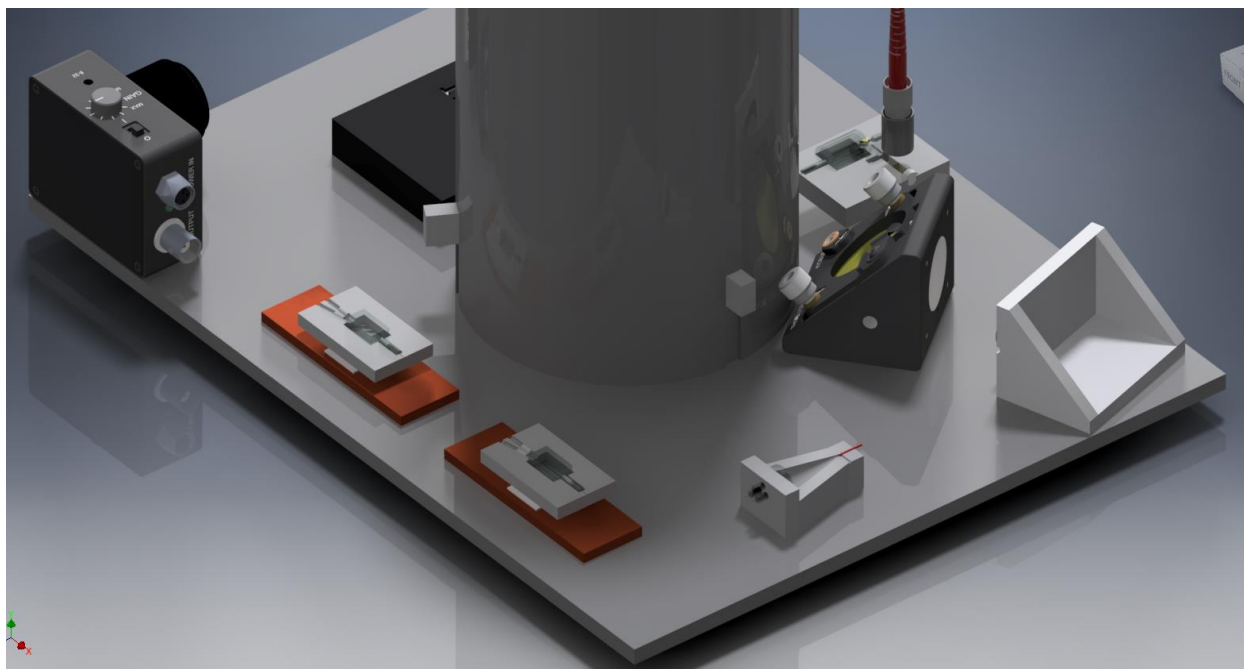


Figure 113: Opposite side view.

8.4 SNR Analysis

With the optical efficiencies characterized, the final analysis can be performed. The effort includes calculating the SNR of the system, considering plume behaviors, and providing final ROC plots as a function of wind speed. For these purposes, the first two design concepts are considered functionally the same and so the comparison is between the “traditional” system and the VBG filtering system.

An important assumption is that, with 2,000 Hz pulse rates, the plume will be well characterized. The quad-copters, especially, can loiter over an area gathering dense data sets. For the fixed-wing version, the analysis in section 4.5 shows that it is likely that a plume being measured will have sufficient numbers of points to detect leaks.

Table 9: Input parameters for the SNR equation. Values are for the “traditional” design.

Parameter	Value	Comment
Flight Altitude	150 m	~500 ft.
Optical Properties		
Fiber Input MFD	10 μm	Single Mode Fiber
Beam Area	.005 m^2	80 mm diameter
Receive Telescope Area	.008 m^2	100 mm diameter
Transmit Efficiency	93.5%	
Receive Efficiency	64.1%	
Solar Filter Width	10 nm	
Laser Properties		
Laser Pulse Energy	150 μJ	
Pulse Length	80 ns	
Detector Properties		
Efficiency	70%	
Responsivity	17 A/W	at 1.645 μm
Bandwidth	10 MHz	
Noise Equivalent Power	0.13 $\text{pW/Hz}^{1/2}$	
Dark Current		No Reverse Bias
Ground Properties		
Reflectivity	50%	
Height Difference	5 m	
Atmosphere Properties		
Background Methane	1.8 ppm	
Solar Irradiance	5 $\text{mW}/(\text{m}^2 \cdot \text{nm} \cdot \text{sr})$	Vegetation
Transmittance - Online	92%	Background methane, HITRAN w/ 1 km path
Transmittance - Offline	99%	Background methane, HITRAN w/ 1 km path

Table 10: Output from model for input from Table 9

Parameter	Value
Fiber Properties	
Allowable Pulse Power	2.25 KW
Expected Pulse Power	1.88 KW
Percent of Damage Threshold	83.6%
SNR Terms	
SNR Online	969
SNR Offline	1,006
Differential Optical Depth	0.019
Error Terms	
SNR/dOD Error	7.5%
Reflectance Error	1.0%
Speckle Error	2.5%
Total % Error	11%
Absolute Error (of 1.8 ppm)	0.20 ppm

Table 11: Input parameters for the SNR equation. Values are for the “VBG Filtering” design.

Parameter	Value	Comment
Flight Altitude	150 m	~500 ft.
Optical Properties		
Fiber Input MFD	10 μm	Single Mode Fiber
Beam Area	.005 m^2	80 mm diameter
Receive Telescope Area	.008 m^2	100 mm diameter
Transmit Efficiency	93.6%	
Receive Efficiency	39.1%	
Solar Filter Width	10 nm	
Laser Properties		
Laser Pulse Energy	150 μJ	
Pulse Length	80 ns	
Detector Properties		
Efficiency	70%	
Responsivity	17 A/W	at 1.645 μm
Bandwidth	10 MHz	
Noise Equivalent Power	0.13 $\text{pW/Hz}^{1/2}$	
Dark Current		No Reverse Bias
Ground Properties		
Reflectivity	50%	
Height Difference	5 m	
Atmosphere Properties		
Background Methane	1.8 ppm	
Solar Irradiance	5 $\text{mW}/(\text{m}^2 \cdot \text{nm} \cdot \text{sr})$	Vegetation
Transmittance - Online	92%	Background methane, HITRAN w/ 1 km path
Transmittance - Offline	99%	Background methane, HITRAN w/ 1 km path

Table 12: Output from model for input from Table 11

Parameter	Value
Fiber Properties	
Allowable Pulse Power	2.25 KW
Expected Pulse Power	1.88 KW
Percent of Damage Threshold	83.6%
SNR Terms	
SNR Online	749
SNR Offline	778
Differential Optical Depth	0.019
Error Terms	
SNR/dOD Error	9.7%
Reflectance Error	1.0%
Speckle Error	2.5%
Total % Error	13.2%
Absolute Error (of 1.8 ppm)	0.24 ppm

8.5 Leak Detection Performance

The instrument performance values can now be used to calculate leak detection performance. The quadcopter tip/tilt, somewhat random approach is considered first. The approach follows the analyses in Chapter 4. The outcome is a ROC plot that shows leak rate detection vs. wind speed with good, marginal, and no detection given the instrument properties and noise floor. For these purposes, “Good” detection means that there are at least 40 samples within the plume. These points are colored green. Marginal means there are between 10 and 40 samples and these points are yellow. No detection means there are less than 10 samples and these points are red. Figure 114 and 115 provide the plots.

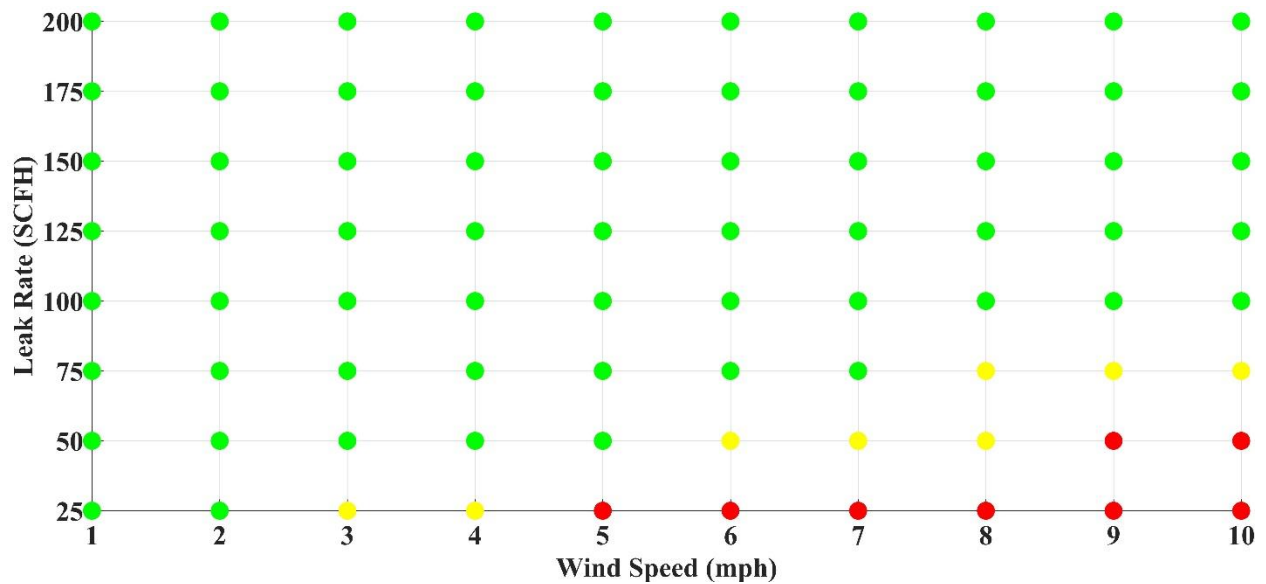


Figure 114: Traditional instrument designs' ROC plot for a quadcopter scanning with somewhat random tip/tilt. Green means "good" probability of detecting and mapping the plume, yellow means "marginal", and red means "unlikely".

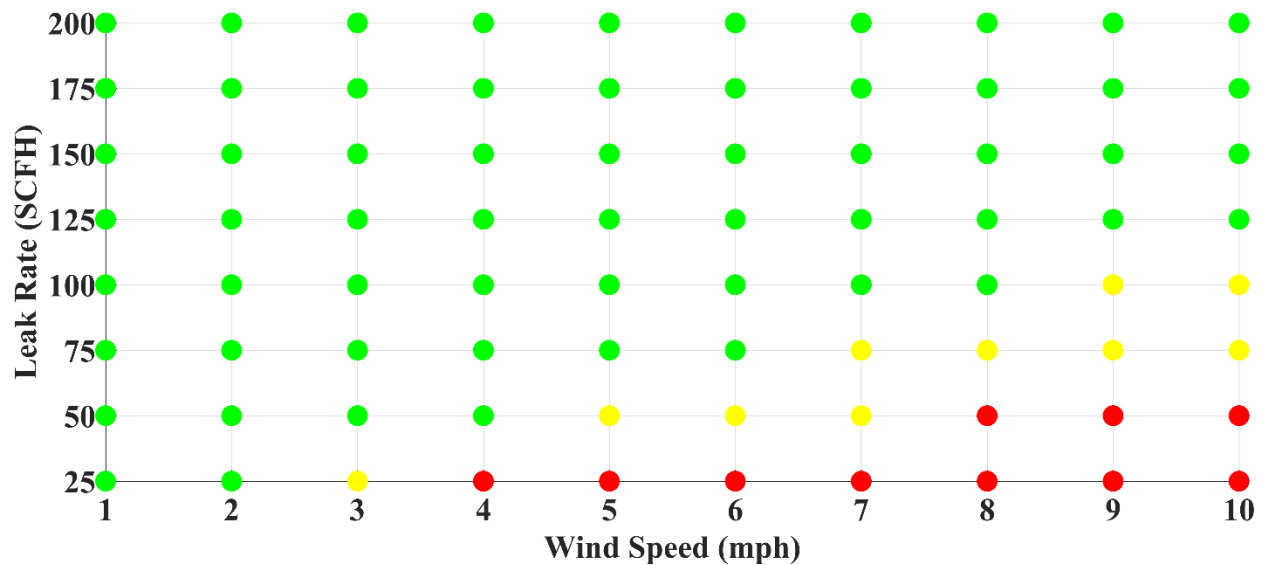


Figure 115: VBG Filtering ROC plot for a quadcopter scanning platform.

In the traditional design, the results show that a 50 SCFH plume can be seen with good certainty in up to 5 mph winds. If things work out well, good detection probability exists up to 8 mph. However, it is unlikely that this instrument concept can perform in winds above 9 mph. There is good news, though, in

that even 25 SCFH plumes can be detected in mild wind conditions. Even the bad news about not making it to 10 mph winds is tempered somewhat in that the quadcopter concept may not fly well in winds of those speeds.

The VBG filtering design performs worse, as expected, although not by much. It shows decent performance up to 5 mph for the 50 SCFH target plume. And the 25 SCFH plume performance is just about on par with the traditional design concept.

The next ROC set assumes a scanning mechanism. Based on section 4.5, a resolution of 1 meter is set.

The plot in Figure 33 is used to set a requirement of having three points in the plume's cross-axis. For this case, there is no "marginal" condition as the number of points threshold is either met or not. The resulting plots are shown in Figure 116 and 117.

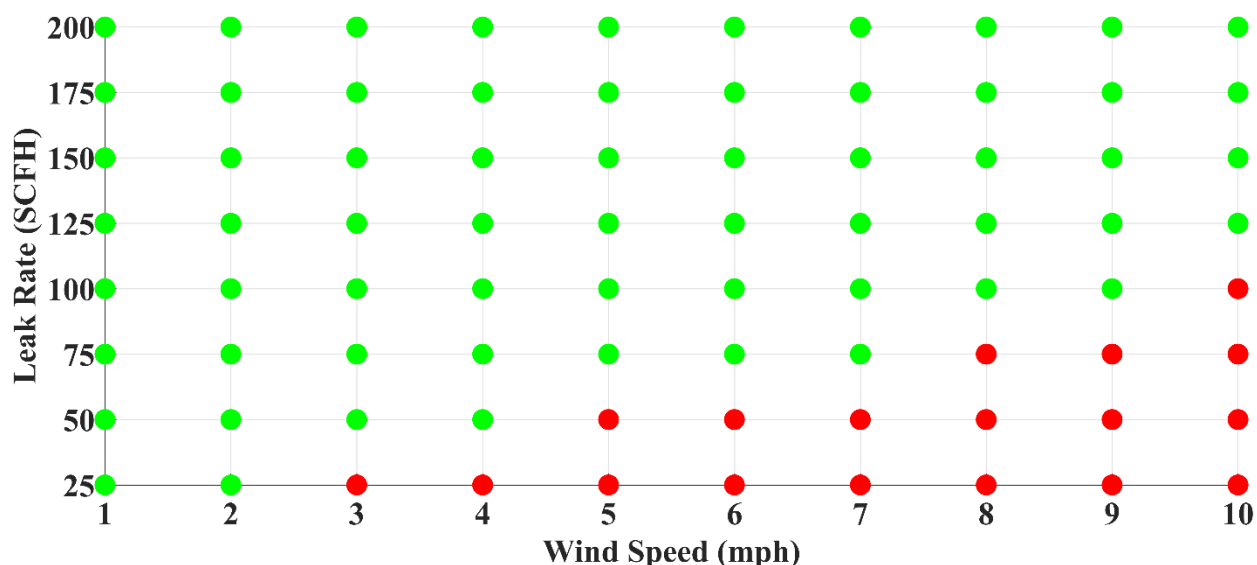


Figure 116: ROC plot of the traditional instrument design with 1 meter resolution scanning. Green dots represent plumes with sufficient numbers of point to reliably detect and map them. Red dots represent plumes with insufficient points.

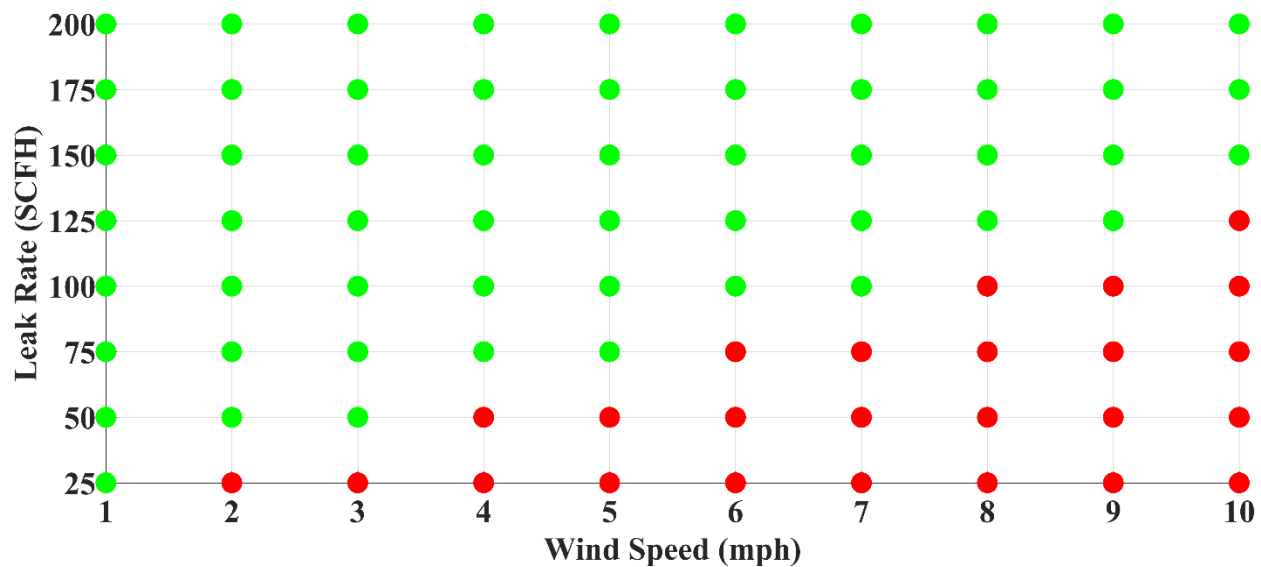


Figure 117: VBG Filtering ROC plot for a scanning system with 1 meter resolution.

8.6 Approximate Cost and Weight

The cost of the two instrument configurations is provided in Table 1 of Appendix 1. A summary is provided in Table 13. Prices for off-the-shelf components were current as of September 2017. Parts machined from Aluminum 6061 are based on a \$1.00 per cubic inch from compiling a rough average from online prices. Machinable Copper C145 is about \$5 per cubic inch. Volumes for all parts are based on pre-machined volumes. For each part, a minimum \$10 is assumed for material costs. The carbon parts are based on off-the-shelf tubes. An assumption of \$200 labor per machined part is made as an average. Some parts are just sheets with holes drilled in them and others are more touch intensive.

Where the actual desired part doesn't exist and requires custom vendor manufacturing, such as GRIN lenses for the 1,645 nm wavelength, the nearest match was used as a stand-in. An effort was made to contact vendors for quotes, but they were understandably hesitant to provide estimates for hardware that was not likely to be purchased. One optics vendor of GRIN lenses offered a price of about \$10 a lens, which is cheaper than the ThorLabs price of ~\$35, but they had a minimum order of 50 lenses when the instrument only requires nine of them. This type of inefficiency is likely to continue for similar custom hardware. On the other hand, the cost efficiency for multiple instruments would be favorable.

The weights for the assemblies are provided in Table 2 of Appendix 1 and summarized in Table 13. Most of the weights came from manufacturer specifications. ThorLabs, for instance, provides a shipping weight that is good to two decimal places. For small parts with weights less than their resolution, a value of 0.01 pounds was assumed. For machined parts, the volume of the final configuration was used. A density of 2.7 g/cc (0.0975 lbs./in³) was assumed for aluminum, a density of 8.94 g/cc (0.323 lbs./in³) was used for C145 copper.

Both the cost and weight analyses favor the “traditional” design. As the traditional design also performed better in finding plumes, this is an unfortunate three strikes against the filtering concept. Because the cost/weight of the laser itself is not known, there may yet be some savings for the VBG filtering concept. It seems unlikely, though, that the laser would be 10 lbs. lighter, even if it was \$7,000 cheaper.

Table 13: Summaries of the cost and weight of the two concepts. The Traditional concept is from the quadcopter iteration.

Cost		Weight in kg (lbs.)	
Traditional	VBG Filtering	Traditional	VBG Filtering
\$9,119	\$16,376	6.9 (15.20)	11.5 (25.45)

The cost and weight of the laser system is unknown and represents an unfortunate gap in knowledge. This is also true of the electronic processing system and scanning assembly. All are custom solutions and expected to be *significant* drains on a budget.

Table 14 provides a list of the unknown assemblies.

Table 14: List of major system components with unknown costs and weights

Primary Systems with Unknown Costs and Weights
DIAL Laser System 2 kHz, 150 μ m pulse energy, 80 ns pulse width
Scanning Mirror Assembly 15 Hz rotation rate 7.8" Gold-Plated with 6° Wedge
Processing Chain JESD204b ADC & FPGA Chain
Applanix IMU APX-15

Looking back at the original UAV design goals, the weight target was 20 kg (~44 pounds) for quadcopters and 16 kg (~35 pounds) for fixed-wing aircraft. (Quadcopters can lift more, but last a fraction of the time.) A comparison with the estimated weight shows that both designs are close fits within a medium-scale, fixed-wing or quadcopter vehicle as a laser, computer, processing hardware, and electrical supply are likely to weigh more than 20 pounds. The fixed-wing version must also carry a scan mechanism and the quadcopter must carry batteries. Mass savings could be had by changing out the aluminum panels with composites to save perhaps 2.5 kg.

Weight comparison with Ball's lasers shows that the quadcopter design might just be doable with some engineering rigor. Their laser is designed for aircraft operations with minimal effort spent on minimizing weight and volume. They also have higher power requirements. Perhaps closely working with a laser vendor could get the weight to around 8 kg, leaving 3 kg for electronics and batteries. That's still a stretch, although the weight limit is a trading variable for flight time in quadcopters.

Taking into account the instrument complexity, weight, volume, and cost, it seems likely that the first generation instrument should be built for a quadcopter platform. This would negate the need for the development effort of a scanning system. It also makes it easier to take short flights in the early testing phase. Personal experience also points to quadcopters tending to survive crashes a little better. The

minimized flight time of a quadcopter would likely eliminate pipeline operations to start, but the instrument could find utility over other assets such as wells, compressor stations, and storage farms. Then, as the design and technology matured, it could be expanded to fixed-wing platforms.

Chapter 9: Summary, Conclusions, and Next Steps

This chapter wraps up the dissertation with a summary of the major topics. Some of the top conclusions follow, and then some thoughts on the next steps.

9.1 Summary

The original intent of the dissertation was to determine if wavelength filtering on the receive side could be a better solution than filtering on the transmit side. In the three categories of performance, cost, and weight, the answer is “no”. Consider the author wistful. On the other hand, developing the tools and analysis necessary to make the comparisons proved educational and, in a couple of areas, broadly useful to the community.

Starting with the introduction content, analysis of trends in the methane industry highlighted the confluence of climate, industry, and political interests. Now seems to be a good time to be developing efficient, next-generation methane instrumentation. And, at the same time, UAV capabilities have somewhat stabilized for fixed-wing platforms, leading to increased certainty in design guidelines. Quadcopters are still an emerging technology with tighter restrictions, chief among them flight duration and power. Even still, the conclusion is that the technology required to get first generation DIAL instruments into the air on either UAV platform seems sufficient.

The first broadly applicable analysis was discussed in chapter 3 with Gaussian plume modeling and correlation. The discussion started with a broad overview of the equations and the ubiquity of the plume’s use in plume characterization approaches. This was followed with a series of steps to convert the model from providing simple point mass concentrations to an integrated version that can be used to calculate total column densities in units of ppm-m or ppm for DIAL instruments.

Using Ball Aerospace’s dense data sets, the model was correlated with an Eigenvector approach where the plume concentrations were used in a weighted approach. The ability to identify a plume’s primary axis allowed for further analyses. In particular, the five conclusions were:

1. The overall shape predictions were accurate, including the characteristic rise and fall of a plume's profile due to beam overlap with the plume and its environment.
2. The plume's cross-section closely followed a Gaussian distribution, perhaps even better than what might be expected.
3. The axial profile did not taper quite as quickly as the model would predict.
4. The correlation was higher at the leak source than further downstream where real-world atmospheric turbulence and variation corrupts the smoothly varying predictions.
5. The model is sufficient for general instrument predictions.

With a suitable plume model in place, attention next turned to various smaller trade studies. The series kicked off with the SNR approach, which is strongly based on previous efforts. Some modifications to account for low altitude flight were made. Next came a study of a laser beam's diameter and its influence on plume detectability. It was found that, in general, smaller beam diameters mean better certainty of detecting small plumes, although wider beams tend to make plumes apparently larger, which may aid detection.

The foundation of SNR, beam diameter impacts, and plume simulation led to analyses of required instrument performance to be able to detect plumes of various sizes. It was found that a SNR of around 10% of background (0.18 ppm) would provide a good design-to target, if the scan approach provided sufficiently dense data sets and the beam diameter could be made less than a meter. A study of scan patterns showed that a 2 kHz laser scanning a circle on the ground at 15 Hz from 500 feet with a 6-degree angle could provide a worst-case spatial density of about 1 point per meter, which could meet the instrument performance goals.

The chapter concluded with two small trade studies on the effective beam diameter of a Gaussian beam and the influence of speckle on error. In the former case, it was found that Lambertian surfaces do indeed reduce the apparent diameter of a beam via shrinking subtended angles from the ground to the receiver. However, the impact is negligible until the beam's divergence angle exceeds 25 degrees. The speckle

analysis also showed fairly small influences of a couple of percent or less in most cases. Although it was observed that there could be local variations of 10% or more within a beam, the overall integrated result was harder to influence. This result was corroborated with somewhat analogous papers based on field work.

Chapter 5 covered the influence of wavelength-based reflectivity variations on DIAL instrument error budgets. The content was expanded into a peer-reviewed paper published with SPIE's Journal of Remote Sensing after extensive searching found no similar quantitative discussions. It was found that, for the most part, errors due to reflectivity variations were 1% or less while some materials could introduce errors up to 10%. The types of materials causing large errors are primarily those with low, but wildly varying reflectivity in the wavelengths of interest.

Chapter 6 moved the discussion from simulations and analysis of the environment to foundation work in hardware concepts. The first section introduced the two primary instrument configurations: a "traditional" concept with multiple detectors and a "VBG filtering" concept that relied on narrow-wavelength filtering on the receive side and time-delaying fiber loops. The hope of the VBG filtering concept was that cost, weight, and performance efficiencies could be gained by using a single optical analysis path and detector.

The second section introduced Volume Bragg Gratings as wavelength filtering optics that are often used in laser assemblies. It was observation of their excellent performance in the literature that first seeded a central question of this dissertation: Could VBG's be used on the receive side of an instrument to reduce costly engineering rigor requirements on the transmission side?

Discussion next moved to comparisons of single-mode and multi-mode fiber with the motivation being an understanding of damage thresholds. Single-mode fiber is both smaller and cheaper than multi-mode fiber, two enabling benefits that were deemed necessary if the VBG filtering concept could be made to work. Plots from data provided by ThorLabs were used to determine the trade-offs in core diameter, pulse

time, and pulse energy in later analyses. The conclusion was that single-mode fiber could be made to work with careful design.

Beam splitting was the next topic of discussion with polarization dependent losses being the principal concern. Each instrument input beam, regardless of design, needs to have a little energy picked-off for monitoring. Beams input from the laser may have different polarizations and the influence of differential absorption leads to errors in the final result. The pursued solution was the use of gold knife-edge prisms based on gold's high and nearly polarization-independent reflectivity properties. The drawback is a requirement for stable optical alignment and some loss of energy due to scattering.

Analysis then switched gears with a section on scanning hardware. The unfortunate conclusion was that there are few off-the-shelf solutions available for accurate and fast scanning of beams larger than a couple of centimeters. A proposed solution was a rotating mirror with a wedged face to scan a circle on the ground. A slight increase in the wedge angle is designed for the outer diameter in order to co-locate the transmitted and received signals. However, although the concept is simple, expectation is that implementing a viable solution may be as complex as the rest of the transceiver design.

Detectors were the next focus. Although somewhat buried in this chapter, the more analysis and hands-on time with Ball's instrument was achieved, the more a feeling that the detector is at the heart of design was realized. As part of the various studies in this area, a common "rule of thumb" about how to choose the bandwidth was pursued. It was found that the bandwidth and rise time product of 0.35 is designed for a square wave pulse and that there is some leeway to be had with rounded laser pulses.

The last major section of chapter 6 looked at pulse integration methods. An integration scheme is chosen instead of discrete measurements taken at set time intervals to provide measurement robustness as the pulse interacts with ground structures and vegetation. The subsequent analysis of this approach showed minimal noise contributions from setting a conservative integration window's width and center time. Specifically, a normalized integration window of 4x the pulse width and centered within 0.5x of the

pulses' center time should provide good certainty of capture with low noise. The chapter concluded with lessons learned from the hundreds of hours of hands-on time with Ball's system.

Chapter 7 took all of the disparate analyses and lessons learned to set design guidelines. These were 70 ns pulses of 150 mW of energy at 2 kHz, a beam diameter of no greater than 0.9 meters, a SNR of about 1,130, and an optical efficiency better than 75%. If a system could be made to those specifications, then the system should be able to just about see a 50 SCFH leak in 10 mph winds.

Chapter 8 provided an end-to-end design and analysis of the two concepts by building on top of the foundation setup in the previous chapters. Optic analysis used Zemax OpticStudio software and included various trade studies on optical alignment stability. Autodesk Inventor was used to create CAD renderings useful for laying out a general configuration of the components. The outcome of the design portion of the chapter was a set of optical efficiency predictions for each instrument concept. It was found that the traditional design outperformed the VBG filtering design ~55% to ~35%, with small variations depending on the specific optical path.

The next section used these efficiencies as well as an on-the-ground beam diameter prediction of 80 mm to calculate predicted ROC plots. These plots estimate the ability to reliably detect leak rates of a given SCFH in various wind conditions. For the quadcopter implementations, the plots were broken into categories of "likely", "maybe", and "not likely". The fixed-wing platform with the scan mirror is more coarsely defined as "likely" and "not likely" due to the more certain nature of its scan density. The general conclusion was that a 50 SCFH plume could be likely seen in winds up to 7 mph, but became dicey at higher wind speeds. The ideal goal was to detect these plumes in winds up to 10 mph, so analysis helped limit expectations. The primary limitation was reduced SNR relative to the goal, even though the beam diameter was smaller than the design goal.

The chapter concluded with an analysis of the cost and weight of the designs. As with optical efficiency, the traditional concept returned better numbers than the VBG filtering concept. The VBG filtering

concept is approximately \$7,000 more expensive and weighs about 10 pounds more. For both systems, the large cost and weight unknowns for the laser, processing electronics, and navigation leaves unfortunate holes in the overall story. Attempts were made to work with vendors, but did not progress far with their understanding that there was no intent to buy the hardware.

9.2 Conclusions

The primary conclusion is straightforward: All three categories of price, weight, and performance are won by the traditional design concept. There is nothing about the filtering concept to recommend it in practice, given the current information. It may be that the laser could be noisier and thus cheaper as a result of the back-end filtering, but it seems unlikely that it would be so much cheaper that the performance and weight penalties could be overcome.

Some of the other conclusions that may be beneficial to DIAL efforts in general include:

1. Integrating the Gaussian plume model is a good way to make instrument performance predictions.
2. Wavelength-based reflectivity variations are often small, but can become non-negligible for low reflectivity materials.
3. Speckle-based errors are usually small due to beam integration. However, they can reach 10% in extreme cases.
4. A DIAL beam's diameter strongly drives the performance ceiling of an instrument, no matter how good the SNR is.
5. Scan density matters as much as the beam diameter. A spatial density of about 1 measurement per meter is a good target.
6. A leak rate target of 50 SCFH probably means sub-meter beam diameters and 2 kHz+ pulse rates, regardless of how far or near the instrument is to the ground.
7. Detector bandwidth rules of thumb can be relaxed somewhat for curved laser pulses.

8. A beam's divergence angle can lead to measurable losses in effective beam diameter if it exceeds 25 degrees.
9. Polarization dependent absorption of beams can lead to errors in gas measurement. Using reflective gold surfaces reduces these types of losses while introducing scattering losses and stiffer alignment requirements.
10. A transceiver design for UAV operations is likely to cost around \$25,000 and weigh around 20 pounds. This is without accounting for the laser and electronics. A scanning assembly adds additional cost and weight. And, of course, none of this assumes the cost of labor, which usually costs much more than the hardware.

9.3 Next Steps

The last conclusion above highlights an unfortunate relationship where a modern DIAL instrument for UAVs is likely to cost much more than the platform it flies on. Although only lightly corroborated with industry sources, it seems questionable to put something relatively expensive on a sometimes unpredictable and crash-prone vehicle. The first next step, then, would be to be ruthlessly efficient in reducing cost and weight. Every attempt must be made to remove and reduce components or find cheaper stand-ins. The approach of moving the expensive filtering to the back-end failed in this dissertation, but perhaps there are other ways.

Another next step is to correlate these overall results with hardware and publish the results. The hardware budget for this effort prevented some of this type of work that would have proven insightful. Many hours of hands-on time with Ball's methane instrument did show that early conceptions and models were wrong. Those types of experiences were invaluable for guiding thought processes here, but cannot be directly shared due to trade secret limitations. This is too bad since numerous paper searches showed a lack of useful hardware discussions in most all the published DIAL work.

Overall, it seems that a practical and powerful UAV-based instrument is on the threshold of being achievable. Through careful consideration of physical plume properties and the instrument parameters that drive their detection, an instrument designer should be able to just make it work. Therefore, a third next step would be to build a direct partnership with a large natural gas company. Easier access to industry personnel, assets, and funding would be the best way to develop a UAV instrument. The approach, once funded, would be to design a version for the quadcopter platform and target industry assets that could be mapped on the order of ten minutes.

Bibliography

Alliance for Telecommunications Industry Solutions (ATIS),
<http://www.atis.org/glossary/definition.aspx?id=2014>, Last Accessed Aug 21st, 2017

Amediek, A., et al. "Airborne lidar reflectance measurements at 1.57 μm in support of the A-SCOPE mission for atmospheric CO₂." *Atmospheric Measurement Techniques* 2.2 (2009): 755-772.

Analog Devices Technical Article MS-2374. <http://www.analog.com/media/en/technical-documentation/technical-articles/JESD204B-Survival-Guide.pdf>. Accessed August 3, 2017.

Australia, Data measured at Cape Grim, Australia: 1984 AD to 2017 AD. NOAA (National Oceanic and Atmospheric Administration). 2014. Monthly mean CH₄ concentrations for Cape Grim, Australia. Accessed April 30, 2017.

Baldrige, A. M., et al. "The ASTER spectral library version 2.0." *Remote Sensing of Environment* 113.4 (2009): 711-715.

Barger, R. L., and J. L. Hall. "Wavelength of the 3.39- μm laser-saturated absorption line of methane." *Applied Physics Letters* 22.4 (1973): 196-199.

Bartholomew, Jarett; Lyman; Phillip; Weimer, Carl; Tandy, William; "Wide area methane emissions mapping with airborne IPDA lidar", Proc. SPIE 10406, Lidar Remote Sensing for Environmental Monitoring 2017, 1040607 (2017/08/30); doi: 10.1117/12.2276713

Bartholomew, Jarett, Ball Aerospace & Technologies. Personal Interview. November 2015-September 2017.

Beechcraft Queen Air Overview. https://en.wikipedia.org/wiki/Beechcraft_Queen_Air. Last Accessed Jan 5th, 2016

Boycott, A. E., G. C. C. Damant, and J. S. Haldane. "The Prevention of Compressed-Air Illness." *The Journal of Hygiene* 8.3 (1908): 342-443.

Browell, Edward V., Syed Ismail, and Scott T. Shipley. "Ultraviolet DIAL measurements of O₃ profiles in regions of spatially inhomogeneous aerosols." *Applied Optics* 24.17 (1985): 2827-2836.

Burton, Darrin, et al. "Lidar intensity as a remote sensor of rock properties." *Journal of Sedimentary Research* 81.5 (2011): 339-347.

Christopher Lawrence, The power and the glory: Humphry Davy and Romanticism, reference in Andrew Cunningham and Nicholas Jardine, Romanticism and the Sciences Cambridge: University Press, 1990 page 224

Chu, Xinzhaoh, "Lidar Remote Sensing" class notes.

<http://cires1.colorado.edu/science/groups/chu/classes/Lidar2016/>, Last accessed Sep 4, 2017.

Ciapurin, Igor V., Leonid B. Glebov, and Vadim I. Smirnov. "Modeling of phase volume diffractive gratings, part 1: transmitting sinusoidal uniform gratings." *Optical Engineering* 45.1 (2006): 015802-015802.

Collier, J. L., et al. "A low-cost gated integrator boxcar averager." *Measurement science and technology* 7.9 (1996): 1204.

Colorado, Colorado Department of Public Health and Environment, Air Quality Control Commission, Regulation Number 7, "Control of Ozone via Ozone Precursors and Control of Hydrocarbons via Oil and Gas Emissions", 5 CCR 1001-9

Colorado Oil & Gas Conservation Commission. <http://cogcc.state.co.us/#/home>. Last accessed April 30, 2017.

Davy, H. (1816). "On the Fire-Damp of Coal Mines, and on Methods of Lighting the Mines So as to Prevent Its Explosion". *Philosophical Transactions of the Royal Society of London* 106: 1.
doi:10.1098/rstl.1816.0001

Denver Post. "Colorado Oil and Gas Wells by the Numbers". Available at http://www.denverpost.com/datacenter/ci_27519246/colorado-oil-and-gas-wells-by-numbers, Last accessed April 30, 2017.

Dobler, Jeremy T., et al. "Atmospheric CO₂ column measurements with an airborne intensity-modulated continuous wave 1.57 μ m fiber laser lidar." *Applied optics* 52.12 (2013): 2874-2892.

Du, D., et al. "Laser-induced breakdown by impact ionization in SiO₂ with pulse widths from 7 ns to 150 fs." *Applied physics letters* 64.23 (1994): 3071-3073.

Efimov, Oleg M., et al. "High-efficiency Bragg gratings in photothermorefractive glass." *Applied Optics* 38.4 (1999): 619-627.

Ehret, G., et al. "Space-borne remote sensing of CO₂, CH₄, and N₂O by integrated path differential absorption lidar: a sensitivity analysis." *Applied Physics B* 90.3-4 (2008): 593-608.

U.S. Environmental Protection Agency, "Inventory of U.S. Greenhouse Gas Emissions and Sinks: 1990-2013", April 15, 2015.

U.S. Environmental Protection Agency, "Climate Change Indicators", <http://www3.epa.gov/climatechange/science/indicators/ghg/ghg-concentrations.html>. Last accessed Sep. 4, 2017.

U.S. Environmental Protection Agency Press Release. Available at [https://www.epa.gov/controlling-air-pollution-oil-and-natural-gas-industry/new-source-performance-standards-and#Final rules](https://www.epa.gov/controlling-air-pollution-oil-and-natural-gas-industry/new-source-performance-standards-and#Final%20rules), Accessed October 15, 2015

Federal Aviation Agency, https://www.faa.gov/regulations_policies/rulemaking/media/021515_sUAS_Summary.pdf. Last accessed May 7, 2017.

Fritsche, Haro, et al. "Increased efficiency of Er: YAG lasers at 1645 nm using narrow bandwidth diode lasers and dual-wavelength resonant pumping." SPIE LASE. International Society for Optics and Photonics, 2014.

Geiger, Allen R. "Mid-infrared light hydrocarbon dial lidar." U.S. Patent No. 5,250,810. 5 Oct. 1993.
Goodman, Joseph W. Speckle phenomena in optics: theory and applications. Roberts and Company Publishers, 2007.

Hamamatsu Photonics, <http://www.hamamatsu.com>. Accessed August 6, 2017.

Hawaii, Data measured at Mauna Loa, Hawaii: 1987 AD to 2017 AD. NOAA (National Oceanic and Atmospheric Administration). 2017. Monthly mean CH₄ concentrations for Mauna Loa, Hawaii. Accessed April 30, 2017.

Hubbert, R. G., and G. J. Troup. "Single-mode 3.39 μm laser study of methane absorption lineshape." Physics Letters A 61.1 (1977): 33-34.

Intergovernmental Panel On Climate Change. "Climate change 2007: The physical science basis." *Agenda* 6.07 (2007): 333.

Intergovernmental Panel On Climate Change, "IPCC, 2014: Climate Change 2014: Mitigation of Climate Change. Contribution of Working Group III to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change." (2014).

Karcher, S., "Embedded DIAL System for Measuring Fugitive Natural Gas Emissions" NIWeek, <http://sine.ni.com/cs/app/doc/p/id/cs-17379#>, (2017).

Kiemle, Christoph, et al. "Sensitivity studies for a space-based methane lidar mission." *Atmospheric Measurement Techniques* 4.10 (2011): 2195-2211.

Kim, J. W., et al. "Fiber-laser-pumped Er: YAG lasers." *IEEE Journal of Selected Topics in Quantum Electronics* 15.2 (2009): 361-371.

Kohl, Claus-Dieter, Wagner, Thorsten. *Gas Sensing Fundamentals*. Heidelberg: Springer Berlin, 2014. Print.

Lawrence, J. P., R. J. Leigh, and P. S. Monks. "The impact of surface reflectance variability on total column differential absorption LiDAR measurements of atmospheric CO₂." *Atmospheric Measurement Techniques Discussions* 3.1 (2010): 147-184.

Le Quéré, 2013] Le Quéré, C., et. al.: "Global carbon budget 2013", *Earth Syst. Sci. Data Discuss.*, 6, 689-760, doi:10.5194/essdd-6-689-2013, 2013.

Lumeau, Julien, et al. "Ultra-narrow bandpass filters based on volume Bragg grating technologies." *Proc. SPIE*. Vol. 7675. 2010.

Maiman, Theodore H. "Stimulated optical radiation in ruby." (1960): 493-494.

Menzies, Robert, and M. Shumate. "Ozone monitoring with an airborne laser absorption spectrometer." *Quantum Electronics, IEEE Journal of* 13.9 (1977): 907-907.

Murdock, Darryl G., et al. "Applications of real-world gas detection: Airborne Natural Gas Emission Lidar (ANGEL) system." *Journal of Applied Remote Sensing* 2.1 (2008): 023518.

Murray, Edward R. "Remote measurement of gases using differential-absorption lidar." *Optical Engineering* 17.1 (1978): 170130-170130.

Olympus Medical, <http://medical.olympusamerica.com/products/lithotripsy/holmiumyag-30w-laser-dur-hl20> Last accessed August 17, 2017.

Optics In Motion LLC, <http://www.opticsinmotion.net>. Accessed August 6, 2017.

Optigrate, Corporation, <http://www.optigrate.com/technology.html>, Last Accessed Sep. 4, 2017

OPTRA, Inc., <http://www.optra.com>. Accessed August 6, 2017

Opus Microsystems, <http://www.opusmicro.com/>. Accessed August 6, 2017

Orwiler, Bob, "Vertical Amplifier Circuits", Circuit Concepts, Tektronix, 062-1145-00, 1969

Pacific Northwest National Library. "Vapor Phase Infrared Spectral Library". Used under license. Accessed, December 2015.

Petermann, K. "Constraints for fundamental-mode spot size for broadband dispersion-compensated single-mode fibres." Electronics Letters 19.18 (1983): 712-714.

PG&E Industry Representative. Personal Interview. November 2015.

Pipeline and Hazardous Materials Safety Administration, "Pipeline Safety Update, October 2012". Available at <http://opsweb.phmsa.dot.gov/>. Last accessed February 10, 2016.

Pruitt, Jeff, et al. "High-speed CW lidar retrieval using spectral lock-in algorithm." Lidar Remote Sensing for Environmental Monitoring IV. Vol. 5154. 2003.

Rao, S. Trivikrama, Jia-Yeong Ku, and K. Shankar Rao. "Sampling strategies for toxic air contaminants." Risk Analysis 11.3 (1991): 441-451

Rao, K. Shankar. "Source estimation methods for atmospheric dispersion." Atmospheric Environment 41.33 (2007): 6964-6973.

Refaat, Tamer F., et al. "Performance evaluation of a 1.6- μ m methane DIAL system from ground, aircraft and UAV platforms." Optics express 21.25 (2013): 30415-30432.

Richards, P. L. "Bolometers for infrared and millimeter waves." Journal of Applied Physics 76.1 (1994): 1-24.

Riken Keiki History. <http://www.rkiinstruments.com/pages/history.htm> Last accessed February 15, 2016.
Setzler, Scott D., et al. "Resonantly pumped eyesafe erbium lasers." IEEE Journal of Selected Topics in Quantum Electronics 11.3 (2005): 645-657.

Schotland, Richard M. "The determination of the vertical profile of atmospheric gases by means of a ground based optical radar." Proceedings, Third Symposium on Remote Sensing of the Environment (Environmental Research Institute of Michigan, Ann Arbor, 1964). 1964.

Schotland, Richard M. "Errors in the lidar measurement of atmospheric gases by differential absorption." (1974).

Scientific Materials. <http://www.scientificmaterials.com/products/ho-yag.php> Last accessed August 17, 2017.

Shuman, Timothy, et al. "Efficient 1.6 micron laser source for methane DIAL." SPIE Optical Engineering+ Applications. International Society for Optics and Photonics, 2013.

Smith, Arlee V., and Binh T. Do. "Bulk and surface laser damage of silica by picosecond and nanosecond pulses at 1064 nm." Applied optics 47.26 (2008): 4812-4832.

Sorokin, Evgeni, et al. "Ultrabroad continuous-wave tuning of ceramic Cr: ZnSe and Cr: ZnS lasers." Advanced Solid-State Photonics. Optical Society of America, 2010.

Spiers, Gary D., Robert T. Menzies, and Joseph C. Jacob. "Lidar reflectance from snow at 2.05 μm wavelength as measured by the JPL Airborne Laser Absorption Spectrometer." Applied optics 55.8 (2016): 1978-1986.

Steven Vincent Stearns, Joseph Lawrence Lippert, inventor; ITT Manufacturing Enterprises, Inc., assignee. Determination of gas flux using airborne dial lidar. US patent 8,010,300. August 30, 2011.

Steele, L.P., P.B. Krummel, and R.L. Langenfelds, Data measured at Shetland Islands, Scotland: 1993 AD to 2001 AD. Atmospheric methane record from Shetland Islands, Scotland (October 2002 version). In: Trends: A compendium of data on global change. Oak Ridge, TN: U.S. Department of Energy. Accessed September 13, 2005.

Steinhart, John S., and Stanley R. Hart. "Calibration curves for thermistors." Deep Sea Research and Oceanographic Abstracts. Vol. 15. No. 4. Elsevier, 1968.

Stuart, B. C., et al. "Laser-induced damage in dielectrics with nanosecond to subpicosecond pulses." Physical review letters 74.12 (1995): 2248.

Sun, Xiaoli, and James B. Abshire. "Comparison of IPDA lidar receiver sensitivity for coherent detection and for direct detection using sine-wave and pulsed modulation." *Optics express* 20.19 (2012): 21291-21304.

Tang, Pinghua, et al. "Robust wavelength-locked narrow-linewidth Er-doped yttrium aluminum garnet laser." *Applied Physics Express* 8.1 (2014): 012703.

Tandy, William D., Jarett Bartholomew, William J. Emery, Ashwin Yerasi, "Analysis of the impact of wavelength separation on reflectivity error for differential absorption lidar using the ASTER spectral library," *J. Appl. Remote Sens.* 11(3), 036008 (2017), doi: 10.1117/1.JRS.11.036008.

Technical Note TN 1005, "Signal Recovery", 2004,
<http://sunnytek.net/admin/xiazaifiles/2010114103920522.pdf>. Accessed August 1, 2017.
Texas Instruments catalog. <http://www.ti.com>. Accessed August 8, 2017.

Thoma, E. D., et al. "Assessment of methane and VOC emissions from select upstream oil and gas production operations using remote measurements, interim report on recent survey studies." *Proceedings of 105th Annual Conference of the Air & Waste Management Association, Control*. No. 2012-A. 2012.

ThorLabs, Incorporated, <http://www.thorlabs.com>, Last Accessed Sep. 4, 2017

Verena Mackowiak (GmbH), Jens Peupelmann (GmbH), Yi Ma (USA), and Anthony Gorges (USA)
"ThorLabs Whitepaper: NEP-Noise Equivalent Power,
https://www.thorlabs.com/images/TabImages/Noise_Equivalent_Power_White_Paper.pdf, Last Accessed Sep. 4, 2017

Volta, Alessandro (1777) *Lettere del Signor Don Alessandro Volta ... Sull' Aria Inflammabile Nativa delle Paludi* [Letters of Signor Don Alessandro Volta ... on the flammable native air of the marshes], Milan, Italy: Giuseppe Marelli.

von Fischer, Joseph C., et al. "Rapid, Vehicle-Based Identification of Location and Magnitude of Urban Natural Gas Pipeline Leaks." *Environmental Science & Technology* 51.7 (2017): 4091-4099.

Vorobiev, N., L. Glebov, and V. Smirnov. "Single-frequency-mode Q-switched Nd: YAG and Er: glass lasers controlled by volume Bragg gratings." *Optics express* 16.12 (2008): 9199-9204.

White, Kenneth O., and Stuart A. Schleusener. "Coincidence of Er: YAG laser emission with methane absorption at 1645.1 nm." *Applied Physics Letters* 21.9 (1972): 419-420.

White, Kenneth O., and Wendell R. Watkins. "Erbium laser as a remote sensor of methane." *Applied optics* 14.12 (1975): 2812-2813.

Wood, Dave. In person interviews. Fall 2015.

Yan, Wai Yeung, and Ahmed Shaker. "Radiometric correction and normalization of airborne LiDAR intensity data for improving land-cover classification." *IEEE Transactions on Geoscience and Remote Sensing* 52.12 (2014): 7658-7673.

Young, York E., et al. "Efficient 1645-nm Er: YAG laser." *Optics Letters* 29 (2004): 1075-1077.
"Separation of Organic Molecular in Carbon Nanotubes & Bundles : Molecular Dynamics Simulation",
Zugang Mao & Susan B. Sinnott, *J. Phys. Chem. B* 105, (2001) 6916 – 6924

Appendix A: Cost and Weight Tables

Table 15: List of approximate component costs.

Part	Cost	Traditional		VBG Filtering	
		#	Ext. Cost	#	Ext. Cost
Optics					
OptiGrate Volume Bragg Grating Custom	\$3,000.00	-	-	2	\$6,000
ThorLabs GRIN Sleeve 51-2800-1800	\$6.05	-	-	9	\$54.45
ThorLabs GRIN Lens 2315A	\$36.58	-	-	13	\$475.54
ThorLabs SMF Pig Tailed Ferrule SMPF0115	\$18.40	-	-	13	\$239.20
Asphericon 100 mm Aspheric Lens, f 100 mm ALL100-100-S-U	\$818.00	1	\$818.00	-	-
Asphericon 100 mm Aspheric Lens, f 200 mm ALL100-200-S-U	\$818.00	-	-	1	\$818.00
ThorLabs 10 mm Aspheric Lens: f 8 mm AL108-C	\$161.00	-	-	1	\$161.00
ThorLabs 25 mm Aspheric Lens: f 20 mm AL2520-C	\$234.00	2	\$468.00	1	\$234.00
ThorLabs 12.5 mm Aspheric Lens: f 25 mm AL1225-C	\$183.00	-	-	1	\$183.00
ThorLabs 1” Aspheric Lens: f 60 mm LA1134	\$21.40	1	\$21.40	-	-
ThorLabs 1” Plano-Convex Lens: f 100 mm LA1509-C	\$33.50	-	-	1	\$33.50

ThorLabs Plano-Concave Lens: f -24 mm LC1975-C	\$35.75	-	-	1	\$35.75
ThorLabs Gold Coated Knife-Edge Prism MRAK25-M01	\$123.00	2	\$246.00	2	\$246.00
ThorLabs Gold Coated Elliptical Mirror PFE10-M01	\$67.50	1	\$67.50	2	\$135.00
ThorLabs Optical Cut-Off Filter FELH1500	\$161.00	2	\$322.00	1	\$161.00
ThorLabs Neutral Density Filter NENIR40B-C	\$55.50	1	\$55.50	-	-
Detectors					
ThorLabs Avalanche Photodetector APD410C	\$2,301.00	1	\$2,301.00	1	\$2,301.00
ThorLabs InGaAs Biased Detector DET10N	\$503.00	2	\$1,006.00	-	-
Fiber Optics					
ThorLabs SMF Patch Cable, 10.0 μ m 1-meter length P1-1550A-FC-1	\$71.50	-	-	2	\$143.00
ThorLabs SMF Bare Cable, 10.0 μ m 1000-meter length SMF-28-1000	\$474.00	-	-	1	\$474.00
ThorLabs MMF Bare Cable, 200.0 μ m 500-meter length FP200ERT	\$555.00	-	-	1	\$555.00
Thermal Control					
ThorLabs 2.5A Thermal Electro Coolers TEC3-2.5	\$31.00	-	-	2	\$62.00

ThorLabs 10 kOhm Thermistor TH10K	\$4.19	-	-	2	\$8.38
ThorLabs TEC Controller MTD415T	\$26.75	-	-	2	\$53.50
Miscellaneous					
ThorLabs Quartz Gas Cell GC19075-CS	\$615.00	1	\$615.00	-	-
ThorLabs 1" Lens Tube: 1" Long SM1L10	\$14.25	2	\$28.50	1	\$14.25
ThorLabs 1" Lens Tube: 2.5" Long SM1L25	\$20.70	1	\$20.70	-	-
ThorLabs Elliptical Optic Support KCB1E	\$202.00	1	\$202.00	1	\$202.00
Custom Hardware					
Custom Beam Splitting Bench Al6061-T651, 1.2 in ³	\$210.00	-	-	1	\$210.00
Custom Elliptical Optic Mount Al6061-T651, 5.00 in ³	\$210.00	-	-	1	\$210.00
Custom Path Combining Bench Al6061-T651, 1.50 in ³	\$210.00	-	-	1	\$210.00
Custom VBG Optics Bench Al6061-T651, 0.5 in ³	\$210.00	-	-	2	\$420.00
Custom 100 mm lens mount Al6061-T651, 13 in ³	\$213	1	\$213.00	1	\$213.00
Custom Telescope Fiber Mount Al6061-T651, 13 in ³	\$213	-	-	1	\$213.00
Custom	\$210	2	\$420.00	-	-

Prism Holder Al6061-T651, 2 in ³					
Custom Detector Path Combining Al6061-T651, 1 in ³	\$210	-	-	1	\$210.00
Custom 1" Lens Tube Holder Al6061-T651, 6.25 in ³	\$210	2	\$420.00	-	-
Custom Detector Holder Al6061-T651, 24 in ³	\$224	1	\$224.00	-	-
Custom Various Panels Al6061-T651, 13 in ³	\$220	6	\$1,320.00	6	\$1,320.00
Custom TEC Heat Sink Copper, 3 in ³	\$215	-	-	2	\$430.00
Custom Telescope Lens Tube, 200 mm Long Carbon Fiber, 0.2" thick	\$300	-	-	1	\$300.00
Custom Telescope Lens Tube, 100 mm Long Carbon Fiber, 0.2" thick	\$300	1	\$300.00	-	-
Custom 100 mm lens support Carbon Fiber, 1.5"	\$50	1	\$50.00	1	\$50.00
Totals			\$9,118.60		\$16,375.57

Table 16: List of component weights

Part	Weight (lbs.)	Traditional		VBG Filtering	
		#	Ext. Weight	#	Ext. Weight
Optics					
OptiGrate Volume Bragg Grating Custom	0.01	-	-	2	0.02
ThorLabs GRIN Sleeve 51-2800-1800	0.01	-	-	9	0.09
ThorLabs GRIN Lens 2315A	0.01	-	-	13	0.13
ThorLabs SMF Pig Tailed Ferrule SMPF0115	0.01	-	-	13	0.13
Asphericon 100 mm Aspheric Lens, f 100 mm ALL100-100-S-U	1.28	1	1.28	-	-
Asphericon 100 mm Aspheric Lens, f 200 mm ALL100-200-S-U	1.28	-	-	1	1.28
ThorLabs 10 mm Aspheric Lens: f 8 mm AL108-C	0.04	-	-	1	0.04
ThorLabs 25 mm Aspheric Lens: f 20 mm AL2520-C	0.07	2	0.14	1	0.07
ThorLabs 12.5 mm Aspheric Lens: f 25 mm AL1225-C	0.05	-	-	1	0.05
ThorLabs 1” Aspheric Lens: f 60 mm LA1134	0.05	1	0.05	-	-
ThorLabs 1” Plano-Convex Lens: f 100 mm LA1509-C	0.05	-	-	1	0.05
ThorLabs Plano-Concave Lens: f -24 mm	0.03	-	-	1	0.03

LC1975-C					
ThorLabs Gold Coated Knife-Edge Prism MRAK25-M01	0.09	2	0.18	2	0.18
ThorLabs Gold Coated Elliptical Mirror PFE10-M01	0.06	1	0.06	2	0.12
ThorLabs Optical Cut-Off Filter FELH1500	0.04	2	0.08	1	0.04
ThorLabs Neutral Density Filter NENIR40B-C	0.04	1	0.04	-	-
Detectors					
ThorLabs Avalanche Photodetector APD410C	2.65	1	2.65	1	2.65
ThorLabs InGaAs Biased Detector DET10N	0.34	2	0.68	-	-
Fiber Optics					
ThorLabs SMF Patch Cable, 10.0 μm 1-meter length P1-1550A-FC-1	0.20	-	-	2	0.40
ThorLabs SMF Bare Cable, 10.0 μm 1000-meter length SMF-28-1000	4.51	-	-	1	4.51
ThorLabs MMF Bare Cable, 200.0 μm 500-meter length FP200ERT	5.00	-	-	1	5.00
Thermal Control					
ThorLabs 2.5A Thermal Electro Coolers TEC3-2.5	0.03	-	-	2	0.06
ThorLabs 10 kOhm Thermistor	0.01	-	-	2	0.02

TH10K					
ThorLabs TEC Controller MTD415T	0.03	-	-	2	0.06
Miscellaneous					
ThorLabs Quartz Gas Cell GC19075-CS	0.13	1	0.13	-	-
ThorLabs 1" Lens Tube: 1" Long SM1L10	0.04	2	0.08	1	0.04
ThorLabs 1" Lens Tube: 2.5" Long SM1L25	0.07	1	0.07	-	-
ThorLabs Elliptical Optic Support KCB1E	0.31	1	0.31	1	0.31
Custom Hardware					
Custom Beam Splitting Bench Al6061-T651, 1.2 in ³	0.07	-	-	1	0.07
Custom Elliptical Optic Mount Al6061-T651, 5.00 in ³	0.16	-	-	1	0.16
Custom Path Combining Bench Al6061-T651, 1.50 in ³	0.05	-	-	1	0.05
Custom VBG Optics Bench Al6061-T651, 0.5 in ³	0.03	-	-	2	0.05
Custom 100 mm lens mount Al6061-T651, 13 in ³	0.10	1	0.10	1	0.10
Custom Telescope Fiber Mount Al6061-T651, 13 in ³	0.24	-	-	1	0.24
Custom Prism Holder Al6061-T651, 2 in ³	0.02	2	0.03	-	-

Custom Detector Path Combining Al6061-T651, 1 in ³	0.02	-	-	1	0.02
Custom 1" Lens Tube Holder Al6061-T651, 6.25 in ³	0.06	2	0.11	-	-
Custom Detector Holder Al6061-T651, 24 in ³	0.51	1	0.51	-	-
Custom Various Panels Al6061-T651, 13 in ³	1.40	6	8.40	6	8.40
Custom TEC Heat Sink Copper, 3 in ³	0.29	-	-	2	0.58
Custom Telescope Lens Tube, 200 mm Long Carbon Fiber, 0.2" thick	0.40	-	-	1	0.40
Custom Telescope Lens Tube, 100 mm Long Carbon Fiber, 0.2" thick	0.20	1	0.20	-	-
Custom 100 mm lens support Carbon Fiber, 1.5"	0.10	1	0.10	1	0.10
Totals			15.20		25.45