DISSERTATION

A Frequency-agile Na Lidar for the Measurement of Temperature and Velocity in the Mesopause Region

Submitted by
Michael Alton White
Department of Physics

In partial fulfillment of the requirements
For the Degree of Doctor of Philosophy
Colorado State University
Fort Collins, Colorado
Fall 1999

UMI Number: 9950922

UMI®

Bell & Howell Information and Learning Company
300 North Zeeb Road
P.O. Box 1346
Ann Arbor, MI 48106-1346

COLORADO STATE UNIVERSITY

October 27, 1999

WE HEREBY RECOMMEND THAT THE DISSERTATION PREPARED UNDER OUR SUPERVISION BY MICHAEL ALTON WHITE ENTITLED "A FREQUENCY-AGILE NA LIDAR FOR THE MEASUREMENT OF TEMPERATURE AND VELOCITY IN THE MESOPAUSE REGION" BE ACCEPTED AS FULFILLING IN PART REQUIREMENTS FOR THE DEGREE OF DOCTOR OF PHILOSOPHY.

Committee on Graduate Work

Stephen R. Lowder
____

Wayne Schubert
____

Chiao-yao She
____
Co-Adviser

David A. Krueger
____
Adviser

James R. Sites
____
Department Head

Abstract of Dissertation

A frequency-agile Na Lidar for the Measurement of Temperature and Velocity in the Mesopause Region

A clear understanding of the energetics and dynamics of the Earth's mesopause region, located between 80 and 110km, has been pursued for years by the atmospheric modeling and global change communities. These theoretical efforts, however, require long-timespan data sets including both temperature and velocity in the mesopause region.

Modification of Colorado State University's existing two-frequency Na-fluorescence temperature lidar to measure accurate line-of-sight velocity has been accomplished, incorporating several improvements to the original system.

To measure velocity from Doppler-shifting, as well as temperature from Doppler-broadening, at least three operating frequencies of the lidar transmitter are required. However, the time necessary to cycle between output frequencies in the two-frequency temperature lidar is long enough to allow error from atmospheric change between integrating periods to be non-negligible, complicating the addition of another frequency by the same tuning method. A new method, therefore, by which the lidar output frequency may be changed on a shot-to-shot basis has been developed using a unique tandem acousto-optic modulator system, allowing simultaneous integration of the three operating frequencies.

Vertical velocities derived from routine operation of the 3-frequency lidar suggest the presence of a systematic frequency shift in the pulsed output of the lidar. Therefore, a frequency monitoring system utilizing a fortuitously located iodine absorption line has been developed whereby the frequency lineshape-centers of the lidar output puses are measured and any shift accounted for in the temperature/velocity calculation.

With several nights of frequency-monitored lidar temperature and vertical wind data, greatly improved accuracy is observed, allowing for the use of this lidar for detailed energetics and wind field measurements.

Michael White
Department of Physics
Colorado State University
Fort Collins, Colorado 80523
Fall 1999

ACKNOWLEDGMENTS

Whomsoever it is that references this thesis long after its completion date in 1999, it is likely out of a sense of melancholy nostalgia for its subject matter, or in exasperation over a similar lidar system. In either case the situation is desperate. There is hope, however, since no terrestrial lidar work is ever done in a vacuum, and nature abhors a vacuum almost as much as she abhors being bombarded by photons and so forced to yield up her most precious secrets, which is precisely what lidar researchers intend. This present work, being exceptionally non-unique, was also not created in a vacuum, and is beholden to many people for its existence. To my parents, for encouraging me to pursue the sciences. To my advisers Dr.'s Krueger and She for never loosing enthusiasm in the research at the CSU lidar lab. To Brian Jones and Wally Acres for providing expertise in critical fields about which I had never wanted to know. And to lidar scientists past and present, Micah McCauley, Max Caldwell, John Hair, Vince Vasoli, Sean Madine, and Dr. Hu for fighting like god-kings and warrior-poets into the darkness of long lidar nights, and so through the watery abyss of intellectual endurance. Take heart, then, you who enter these pages bathed in the reek of human striving; I suspect you also are not alone.

Table of Contents

A List of Figures and their Captions	page
Fig.2.1 The sodium D_2 fluorescence spectrum for two cases: a. at a constant temperature with Doppler velocities of 0 and ±60m/s marked, b. at rest with three temperatures.	9
Fig.3.1 The sensitivity of the fluorescence spectrum to a 20m/s change in velocity and a 20K change in temperature as a function of frequency relative to the D_{2a} Doppler-free feature. The sum of the absolute value of each is also plotted to guide the eye in determining frequencies of dual sensitivity to temperature and velocity.	16
Fig.3.2 The 2-dimensional field of calibration curves for mapping points of known ratios (R_T, R_V) into temperature and radial velocity. The curves arise from theoretical spectra of sodium fluorescence.	18
Fig.4.1 A Block diagram of the 3-frequency lidar transmitter and receiver. Electronic control lines are shown dashed, and optical paths are solid.	26
Fig.4.2 The Doppler-free, sodium D_{2a} Lamb dip used for active frequency locking of the CW part of the lidar transmitter. The deepest dip is selected for locking in routine data collection, while the smaller troughs may be selected for diagnostic purposes.	29
Fig.4.3 Schematic of the active-locking process. The sinusoidal signal at top is applied to the fast tuning mechanism of the ring dye laser, while the baseline offset responds to the integrated low-noise lock-in output, which in turn drives the laser towards the peak.	31
Fig.4.4 The beam path through an acousto-optical frequency shifter. For constructive Bragg scattering, all of the marked angles are the same. The shown configuration results in an up-shift equal to the acoustic frequency, while selection of a negative angle on the input would result in a down-shift.	34
Fig.4.5 A schematic of the tandem acousto-optical frequency shifter used to alternately generate the up and down-shifted operating frequencies of the lidar. A selective optical chopper near the back mirror determines the current output frequency and switches microwave power to the crystals.	35
Fig.4.6 The beampaths through an acousto-optic modulator for double-pass shifting. The first pass, at top, shifts the beam in frequency and spatially out of parallel; the second pass, at bottom adds to the frequency shifting and returns the beam to the original beam line. If left unblocked, unconverted light is also returned along the original beam line.	36
Fig.4.7 The averaged pulsed dye amplifier output lineshape, as measured by a 750Mhz, finesse 70 confocal interferometer. Some asymmetry is apparent in the lineshape, suggesting the presence of frequency modulation.	39
Fig.5.1 A good, 3-frequency, minute-averaged photocount profile. The lower peak arises from Rayleigh and Mie scattering in the lower atmosphere, the higher peak represents fluorescence signal from the sodium layer. The shape of the lower peak is due to electronic gain-switching in the receiving electronics. The separation of the three frequencies in the lower atmosphere is due to power differences arising from imperfect frequency conversion.	47
Fig.5.2 Top: Typical hour-averaged profiles of temperature, Na density, and vertical velocity taken by the 3-frequency lidar. 8051 represents the 51st day of 1998. Bottom: Similar profiles acquired the following night (8052) using the 2-frequency lidar.	50
Fig.5.3 Five groups of nightly-averaged temperature, Na density, and vertical velocity measurements in which the 3-frequency lidar (bold profiles) and 2-frequency lidar (fine profiles) were operated in close succession over one and one-half years. Though night-to-night variation in the mesopause parameters can be large, no obvious trend is apparent to suggest that the two measurement techniques are systematically biased from each other. In the labeling convention, 97066 corresponds to the 66th day of 1997, etc. Over 80 nights similar to those shown have been acquired by the 3-frequency lidar since 1997, and hundreds using the 2-frequency system since 1993.	52

Chapter 1: Introduction

"The long road of transmigration is a road of pain for the traveler." Written sometime in the 3rd century B.C., Buddhist philosophy proves relevant even now in the field of laser remote sensing. In the 20th Century, the science of atomic spectroscopy has been refined to a fine art of delicacy and precision, and as a less Buddhist, but more recent sage has said, "Lidar is just spectroscopy from a distance." But to exact the transmigration of pure spectroscopy into the body of a working fluorescence lidar has indeed at times been "pain for the traveler", and there have been many on that road. I think, though, that the philosophical Buddhist would approve, despite the imperfections and incompleteness of the transformation, which seems capable of absorbing infinite amounts of human energy. After all, it was he who continued, "The long road of transmigration is a road of pain for the traveler: let him rest by the road and be free."

This first chapter of the thesis outlines the motivation for the research discussed herein, as well as a brief outline of the organization of its presentation. First a brief definition of lidar and typical lidar techniques will be presented, followed by a more detailed history of the sodium lidar. A review of interesting atmospheric phenomena measurable by Na lidar, and a historical sketch of Colorado State University's Na lidar research will precede a brief discussion of the intent and structure of the remainder of the work.

1a. What is a Lidar?

Lidar is an acronym for light detection and ranging. Though many varieties exist, in general a laser beam is directed through the atmosphere, and selected atmospheric information along the path of the laser beam may be retrieved from the returned signal. This information may range from thermodynamic parameters such as temperature, pressure, and density, to dynamic properties such as velocity and turbulence, to species detection such as pollutant concentration, water concentration, and cloud detection. There are a variety of methods for probing the atmosphere including balloon radiosonde, rocket radiosonde, satellites, radar and lidar; each with its applications, strengths, and weaknesses. The principal advantages of lidar include high temporal resolution, often on the order of minutes, high range resolution, often on the order of tens of meters, steerability of the detection area, and the feasibility of continuous measurements over extended periods. Disadvantages include lidar's sensitivity to outside light sources, especially the daytime sun, and its requirement for a generally cloud-free atmosphere.

Various lidar techniques have been developed since the advent of the laser, including differential absorption lidar (DIAL), Raman lidar, Rayleigh lidar, Mie or aerosol lidar, and fluorescence lidar. The principal difference between these methods lies in the scattering mechanism which provides the return signal, and therefore the experimental information, from the atmosphere. DIAL lidar employs Rayleigh-Mie scattering by which signals from two transmitted frequencies may be compared to determine the frequency-dependent optical depth of the atmosphere, and thereby selected species concentrations, and in the case of majority species like nitrogen, atmospheric state parameters as well. Raman lidar uses the up or down-shifted Raman scattering from some constituent of the atmosphere to determine its abundance or other state parameters. Rayleigh lidar uses quasi-elastic scattering from air molecules to deduce temperature, density, and other parameters from the spectrally broadened return signal. Mie lidar depends on Mie backscattering from larger constituents of the atmosphere, such as dust, and is used to determine temporal and spatial distributions of aerosol and clouds. Fluorescence lidar utilizes resonant scattering with atmospheric constituents to determine parameters of that constituent, and possibly generalizing to other species where equilibrium conditions apply.

In most cases the lidar transmitter is a pulsed laser and the receiver includes a telescope aligned with the transmitted beam. Range resolution is usually obtained by measuring the 'time of flight' or time interval between the exit of the laser pulse into the atmosphere and the arrival of the range-resolved data. In many cases the laser is directed into the atmosphere at some distance from the receiving telescope, while in other cases the same telescope is used for expanding the transmitted laser in addition to the collection of the return signal.

1b. The Sodium Lidar

Large fluorescence cross-sections in the visible spectrum of many common elements make fluorescence lidar appealing for precision measurements in the presence of a suitable atmospheric 'tracer'. It is therefore fortunate that the upper mesosphere exhibits layers of atomic metals, likely due to the ablation of meteors in the earth's atmosphere. Of particular interest here is the sodium layer which generally occupies a region between 80 and 110km above the earth's surface, referred to throughout this document as the 'mesopause region'. The peak density of neutral Na atoms of 10^3 to 10^4 cm^{-3} typically occurs around 92km, providing an ideal tracer for fluorescence measurements in this region of the atmosphere.

The basic technique of the Na lidar is to tune a laser in the lidar transmitter to the resonance frequency of the well-known absorption spectrum of atomic Na. Upon interaction with the laser beam, the Na atoms emit fluorescence light, or photons. A portion of this light is collected by the lidar receiving telescope and usually stored by time-of-flight range binning. With proper data analysis, the recorded photon signal may determine the Na density, related to fluorescence intensity, and a variety of other atmospheric parameters which affect the response of the Na atoms to the laser excitation, such as temperature and collective motion. Lidar measurements of the density of atmospheric sodium have been reported since 1972 (Gibson and Sanford, 1972). However, since the behavior of trace amounts of Na in the atmosphere is of less general interest than the characteristics of the majority species in the upper mesosphere, Na lidars are typically operated to derive the parameters of the surrounding environment, with which the Na atoms are assumed to be in equilibrium. For example, the trace Na may assumed to be at the same temperature as the surrounding nitrogen and oxygen, and carried along with the bulk motions of these more plentiful species.

Lidar measurement of Na temperature was first successfully made by Gibson et. al. in 1979 and advanced by Fricke and von Zahn (1985). The temperature measurement is derived from the temperature dependence of the Na absorption spectrum from Doppler broadening. Using their excimer pumped dye laser system, von Zahn's group performed observations at Andoya, Norway (69°N) between 1985 and 1990 and published several significant papers on the structure of polar mesopause region (Lübken and von Zahn, 1991). By probing the sodium layer with a laser at precisely determined frequencies near a Na resonance transition (She et. al., 1990), the Doppler-broadened width of the absorption spectrum, and therefore the temperature of the atmospheric Na atoms, may be determined. This technique offers high resolution observations of the temperature structure of the mesopause region and is being widely used for atmospheric studies.

Techniques for velocity measurements in the mesopause region have also been developed by Bills et. al. (1991) and by She and Yu (1994), in which additional probing of the absorption spectrum of atmospheric sodium resolves both the Doppler broadening and the Doppler shift arising from collective motion. Reliable profiles of both the thermal and dynamic properties of the mesopause region offer a fuller understanding of the energy sources which drive the middle and upper atmosphere.

The apparata of Na fluorescence lidars vary and continue to evolve as laser and detector technology develops. Progress in the construction of all-solid-state lidar transmitters and daytime operating Na lidars (Chen et. al., 1996) is promising for long-term, routine monitoring of the mesopause region.

1c. The Mesopause Region

The thermal environment in the mesopause region exhibits the lowest temperatures in the earth's atmosphere, falling as low as 170 K for the nightly mean in the midlatitudes (She et. al., 1993). The region's temperatures are paradoxically warmest in the winter and coolest in the summer (Lübken and von Zahn, 1991), suggesting that it is not solar flux, but rather dynamical or chemical processes which drive the region's heating (She et. al., 1995). The dynamical energy sources in the mesopause region include atmospheric tides, which are waves excited by solar heating and propagate through the atmosphere with periods equal to integer fractions of 24 hours, and atmospheric gravity waves, which are typically excited near the earth's surface by convection, orographic perturbations, and thunder storms, and propagate upwards with increasing amplitude. The signatures of such waves have been observed in the mesopause region in sodium density profiles (Batista et. al., 1985), temperature profiles (She et. al., 1993) and velocity profiles (She and Yu, 1994).

The velocity, or wind, in the mesopause region has been measured by radar (Vincent and Reid, 1983) and fluorescence lidar (She and Yu, 1994). Velocity in the vertical direction is small, at least on timescales resolvable by ground-based instrumentation, typically on the order of a few centimeters per second, although vertical velocities as high as 8m/s have been reported (Tao and Gardner, 1995). In the horizontal plane, however, large velocities, as high as 100m/s have been reported (Vincent and Fritts, 1987). Also, since the role of wave propagation through the mesopause region is regarded as a principal energy source to these altitudes, flux values, such as momentum flux and heat flux, which combine velocity measurements or velocity and temperature measurements in order to determine the source and quantity of incoming energy, have been determined.

There are a variety of chemical species in the altitudes around the mesopause, including various bands of atomic metals: sodium, potassium, calcium, and iron. In the case of sodium, the total abundance of free Na atoms is about 3×10^9 cm^{-2}, which, together with estimates of the average meteoric influx to the earth's atmosphere suggest an average lifetime of a sodium atom in the atmosphere of about 3 days (Clemsha et. al., 1992).

1d. Lidar Research at Colorado State University

Lidar research at Colorado State University in Fort Collins, Colorado (41°N, 105°W) began with Rayleigh lidar in the mid 1980's. The development of high spectral resolution Rayleigh lidars by Alvarez (1991), Caldwell (1995) and Hair (1999) for profiling temperature and other atmospheric parameters in the lower atmosphere has generated numerous publications and technical advances. In 1989 Colorado State's first Na resonance lidar was constructed in collaboration with scientists from the University of Illinois (She et al., 1990; Bills et. al., 1991). Development and routine data collection with the two-frequency temperature and Na density monitoring lidar was continued by Yu (1994) who temporarily introduced a third frequency using an acousto-optic modulator for measurement of wind in the mesopause region. Construction and integration of a Faraday anomalous dispersion filter for use with the two-frequency Na lidar to allow daytime operation was achieved by H. Chen (1997), along with the first 24-hour monitoring of temperature and density above Fort Collins (Chen et. al., 1996). Since 1989 hundreds of nights of routine lidar observations have been collected and analyzed, along with many campaigns of 24 or more hours conducted and interpreted especially for investigation of atmospheric tides by S. Chen (1999). In 1996 a tandem acousto-optic modulator system was incorporated into the lidar for 3-frequency wind/temperature/density observations, and routine observation of over 80 nights using this system has continued through 1999. Improvements of the 3-frequency transmitter design and analysis of an unexplained, systematic offset in the obtained vertical velocity profiles continued from 1996. In 1999 a frequency monitor based on an iodine optical absorption cell was successfully added to the lidar transmitter with which the frequency of output laser pulses could be accurately monitored for diagnostics and correction of the observed velocity (White, 1998).

1e. The Structure of This Dissertation

At the instigation of the original research reported here, the CSU Na lidar program, under guidance of Dr.'s She and Krueger, was already about four years old. Therefore, some findings are included in this dissertation as necessary background to aid in the understanding of those advancements which are new. The development of the 3-frequency lidar, its associated frequency shifting system, and electronic acquisition system as a stable instrument capable of routine atmospheric observations, the diagnosis of the resulting measured anomalous vertical wind, and the development and integration of a frequency monitor for concurrent operation with the lidar for correction of frequency offsets in the transmitter represent the principal original advancements to be presented here. This work has included a significant amount of original sodium and iodine spectroscopy and novel applications of the obtained results. In an effort to create a comprehensible treatment of the 3-frequency lidar as a whole, a significant portion of the presented information has been drawn from those previous Na lidar developments pertinent to the current hardware innovations.

This thesis outlines the progress and techniques of the 3-frequency Na temperature/wind lidar for the interested researcher in the remote sensing field. Wherever possible apparatus descriptions, part numbers, and advice has been included to aid in the duplication of presented work. The presentation is as follows:

Chapter 2 outlines the physical processes involved in laser-induced fluorescence of sodium at 589nm. The spectroscopy of atomic sodium has been studied for decades, and the most of the experimental parameters of interest for the Na lidar measurement had been reported in the atomic physics literature long before their application to remote sensing. Therefore the rigorous derivation of spectroscopic principles and constants is not among the intentions of the present work. The principal goal of this chapter, rather, is to develop a general vocabulary and formalism with which to discuss the observable parameters utilized for the lidar measurement.

Chapter 3 introduces the general experimental theory of the lidar measurement and design considerations which have been addressed in the development of the 3-frequency Na lidar. The information here is intended to apply generally to all lidar measurements and fluorescence lidar in specific, in order to aquaint the reader with the general problems of lidar design in anticipation of the solutions adopted for the 3-frequency lidar, presented in the following chapter.

Chapter 4 introduces the hardware and technical operation of the 3-frequency Na wind/temperature lidar. The design questions broached in Chapter 3 are addressed in terms of the specific lidar design and the operation of its various instrumental sub-systems. To highlight the necessity of development of the lidar frequency monitor, it is considered as an addition to the basic lidar design discussed in this chapter, and is discussed fully in chapter 7.

Chapter 5 discusses the analyzed data which is converted from the lidar's collected photon counts into measurements of atmospheric parameters. The accuracy and resolution of the obtained parameters are discussed along with an introduction to the 'anomalous measured vertical velocity', a systematic offset in the obtained vertical wind which is inconsistent with even simple atmospheric theory. Interpreted as a measurement bias, the presence of the anomalous measured vertical velocity has necessitated the integration of a pulsed frequency monitor into the lidar system.

Chapter 6 presents an analysis of the anomalous measured vertical velocity in an effort to identify the source of the velocity measurement's systematic offset. A brief discussion of laser chirping and shifting is presented to further justify the necessity of a lidar frequency monitor, together with hypothetical design options for correction of the lidar data.

Chapter 7 introduces the design and hardware of the frequency monitor developed for the 3-frequency wind/temperature lidar. A discussion of the monitor's calibration and accuracy is presented with details of its integration into the routine lidar acquisition system.

Chapter 8 presents final results of the temperature/wind lidar with corrections from the concurrently operating frequency monitor. The accuracy of the resulting data is discussed along with speculation on future improvements of the monitored lidar system.

Chapter 9 provides a brief conclusion with speculation on further studies enabled by the current developments.

Chapter 2: Laser Induced Fluorescence of Sodium at 589nm

The object of this chapter is to understand the quantitative consequences of exciting atmospheric sodium with a laser similar to the one used in the 3-frequency lidar transmitter. To do this, the fluorescence spectrum of sodium must first be considered, along with its temperature and velocity dependence. Since the lidar measurement depends upon fluorescence photons scattered back from atomic sodium atoms, fluorescence properties, such as the backscattering cross-section and fluorescence emission frequencies will be considered. Since many of the quantum-mechanical details of sodium excitation have been presented before, both in the general literature (e.g. Corney, 1977) and in dissertations of previous developers of the CSU Na lidar (Yu, H. Chen, S. Chen), this chapter will focus only on quantities and concepts necessary to understand and reproduce the 3-frequency lidar measurement.

2a. The Sodium Fluorescence Spectrum

Resonance Transitions

The spectroscopic region of primary interest to the operation of the 3-frequency lidar is the vicinity of the sodium D_2 electric dipole transition near 589nm. From the D_2 transition, hyperfine splitting separates the D_{2a} and D_{2b} groups, each associated with one of the two lower energy levels in the D_2 transition, and each allowing excitation to three of the four upper energy levels. Six distinct transitions, therefore, contribute to the fluorescence spectrum of the D_2 transition each with an experimentally determined resonance frequency and relative line strength. These parameters are presented in Table 1.

Table 2.1: Sodium D_2 Absorption Lines

n	name	ν_n(GHz)	S_n
1	$D_{2a,1}$	-0.7150	1/6
2	$D_{2a,2}$	-0.6806	5/6
3	$D_{2a,3}$	-0.6216	14/6
4	$D_{2b,0}$	1.0408	2/6
5	$D_{2b,1}$	1.0566	5/6
6	$D_{2b,2}$	1.0911	5/6

This is a table of the hyperfine transitions in the D_2 line of sodium assigning a number, name, and reporting the frequency relative to the weighted center and the relative line strength of each(Weise et. al 1969).

Transition Line-Broadening

Though the fluorescence spectrum of the D_2 transition arises from six discrete resonance lines, the spectrum is made continuous by various line-broadening mechanisms. Lifetime broadening arises from the finite lifetime of the excited state of a transition, and results in a Lorenzian line shape with full-width-half-maximum of $1/2\pi\tau$, where τ signifies the lifetime of the excited state. In the case of Na, τ = 16.4 ns (Corney, 1977), and so introduces a lifetime broadened width of only about 10MHz. Collisional broadening can also add to the overall linewidth of resonance transitions. However, due to the low pressure in the mesopause region, typically less than .01mbar, collisional broadening contributes negligibly to the overall linewidth of atmospheric Na atoms (Yu, 1994). Far more significant than either lifetime or collisional broadening to the overall linewidth of the sodium transitions is Doppler broadening.

Doppler broadening arises from the thermal velocity of the sodium atoms which serves to Doppler shift photons emitted from a laser at some off-resonance frequency to the resonance frequency in the atoms moving frame of reference. Since any thermally energized gas contains a broad distribution of particle speeds and directions, it is reasonable that for most frequencies near the atomic resonance frequency, there are some atoms with the proper speed and direction to shift the emitted frequency to resonance in their reference frame. Therefore, in the context of a gas of moving atoms, a very narrow resonance band of frequencies will not strictly include all of the transmitted frequencies capable of exciting a given atomic transition, but a band broadened based on the velocity distribution of the gas. The resonance band is therefore considered to be Doppler broadened.

For a volume of gas exhibiting thermal motion but not collective motion (that is, the center of mass of the gas is stationary), and probed by a laser assumed to have a very narrow bandwidth centered at frequency ν, the effect of Doppler broadening is a Gaussian lineshape centered at the transition frequency with a full-width-half-maximum given by Equation 2.1.

$$\Delta \nu_{Dop} = \nu_0 \left[\frac{8 k_B T \ln 2}{M c^2} \right]$$

where

k_B = Boltzmann's Constant

T = temperature Eq.2.1

M = mass of one Na atom

c = speed of light

ν_0 = line center frequency

For sodium at 200K the Doppler width is about 1GHz, making Doppler broadening the dominant source of the transition linewidths of atmospheric sodium. The resulting equation for a Doppler-broadened transition line is given by Equation 2.2 (Theon et. al. 1972).

$$g_n(\nu) = S_n \left[\frac{D}{\pi T}\right]^{\frac{1}{2}} \exp\left[\frac{-D[\nu - \nu_n]^2}{T}\right]$$

where

$$D = \frac{M\lambda_0}{2k_B} = 497.6\,Kns^{-2} \qquad \text{Eq.2.2}$$

S_n = linestrength (from Table 1)

T = temperature

ν_n = line center frequency

It should be noted that it is precisely the effect of Doppler broadening on the fluorescence spectrum of atmospheric sodium which is measured by the 3-frequency lidar to determine the temperature in the mesopause region.

Doppler Shifting

In addition to thermal motion, in which the center of mass of a volume of gas remains fixed, it is possible for atmospheric sodium atoms to exhibit collective motion in the presence of wind in the mesopause region. The effect of such motion is not line broadening, since the motion does not belong to any instantaneous distribution of random velocities, but rather an overall Doppler shift of the Gaussian transition lineshapes. The amount of frequency shifting from a given collective velocity is expressed in Equation 2.3.

$$\Delta\nu = -\frac{V}{\lambda_0} \qquad \text{Eq.2.3}$$

The radial velocity V expressed in Equation 2.3 is the collective velocity along the exciting laser beam. The sign convention in Equation 2.3 has been chosen such that a velocity away from the laser, or vertically up in the case of an azimuthally pointing laser, is considered a positive velocity and yields a negative frequency shift. The effect of Doppler shifting on the thermally broadened, Gaussian resonance lines given by Equation 2.2 is given in Equation 2.4.

$$g_n(\nu) = S_n\left[\frac{D}{\pi T}\right]^{\frac{1}{2}} \exp\left(\frac{-D\left[\nu - \nu_n - \frac{V}{\lambda_0}\right]^2}{T}\right) \qquad \text{Eq.2.4}$$

It should be noted that the velocity dependence of the sodium fluorescence lines in Equation 2.4 is precisely what is measured by the 3-frequency lidar to determine the Doppler shift, and therefore the wind velocity, in the mesopause region.

The Fluorescence Spectrum

To determine the temperature and velocity dependent sodium D_2 spectrum it is only necessary to sum the contributions of the six transitions presented in table 1 in their Gaussian form, given by Equation 2.4. The resulting equation for the sodium fluorescence in the vicinity of the lidar excitation is given by Equation 2.5.

$$g_n(\nu) = \left[\frac{D}{\pi T}\right]^{\frac{1}{2}} \sum_{n=1}^{6} S_n \exp\left(\frac{-D\left[\nu - \nu_n - \frac{V}{\lambda_0}\right]^2}{T}\right) = \sum_{n=1}^{6} g_n(\nu) \qquad \text{Eq.2.5}$$

Plots of the sodium fluorescence spectrum at different temperatures and radial velocities are shown in Figure 2.1a and 1b respectively (see next page).

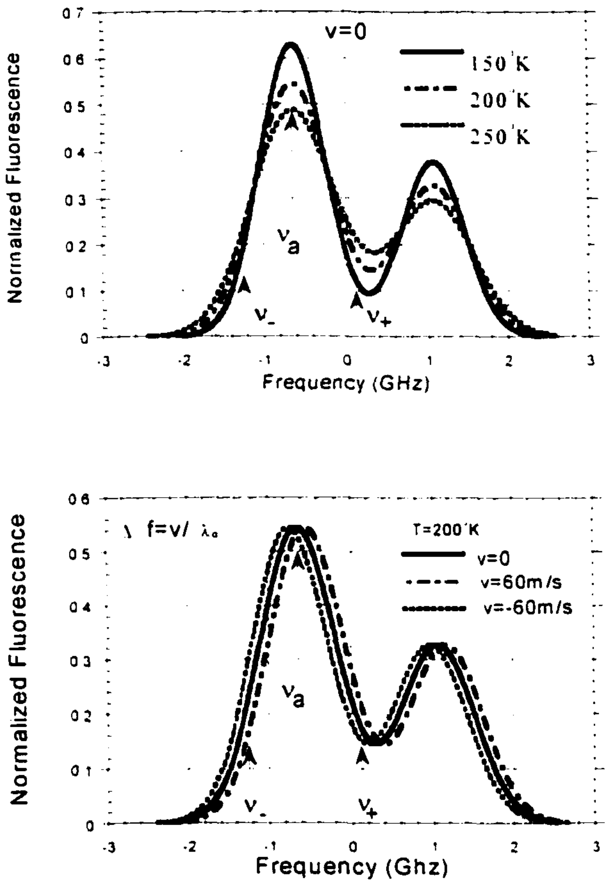

Fig.2.1 The sodium D_2 fluorescence spectrum for two cases: a. at a constant temperature with Doppler velocities of 0 and ±60m/s marked. b. at rest with three temperatures.

It is clear from Figure 2.1 that certain frequencies in the fluorescence spectrum are more sensitive to a change in the sodium temperature than others, and certain frequencies are more sensitive to a collective velocity than others. In the case of no collective velocity (Figure 2.1a), it is possible to determine the temperature of a given spectrum by comparing the relative fluorescence at two points on the spectrum, at the top of the left peak and near the well between the two peaks, for example. Similarly, in the case of constant temperature (Figure 2.1b), it is possible to determine the collective velocity of a given spectrum by comparing the relative fluorescence at two points on the spectrum, on the right and left wings of the left peak, for example. By sampling three points from the fluorescence spectrum, it is possible to determine both the temperature and the radial velocity corresponding to a given spectrum; this is the basis of the 3-frequency lidar measurement, details of which will be discussed in forthcoming chapters.

2b. Backscattering Signal Sources

To understand the experimental information imbedded in the photons returning from the sodium layer, it is necessary to first consider the various processes at work as the exciting laser beam propagates through the atmosphere to the sodium layer, then through the sodium layer to a given altitude, and interacting resonantly with the atmospheric sodium at that altitude. Then the details of the returning fluorescence photons propagating down through the sodium layer and lower altitude must be addressed.

Rayleigh Scattering

Before the light from the lidar transmitter reaches the bottom of the sodium layer, perhaps 80km above the earth's surface, it must pass through 80km of atmosphere. Though the lower atmosphere contains particulates and aerosols which scatter light in ways that are difficult to predict, the quasi-elastic or Rayleigh scattering from "clean air", containing few constituents with sizes comparable to the laser wavelength, is far more manageable. To determine the intensity of laser light which passed through the lower atmosphere to arrive at the sodium layer, it is useful to consider the Rayleigh scattering which the beam encounters.

Above about 30km the air may be considered "clean" to a very good approximation. This assumption is used in two ways, first to regard the atmosphere between about 30km and the sodium layer as effectively perfectly transmitting, and second to allow for the treatment of scattered photons from above 30km as arising from pure Rayleigh scattering, which is well enough behaved to yield atmospheric information, such as optical depth. In this regime the number of photons scattered back from a given region of altitude, collected by the lidar's receiving telescope, and counted by the receiver electronics may be written:

$$N_R = \frac{C}{z^2}\left[\rho_R(z,t)\sigma_R\right]$$

$$C = \left[\frac{\eta A_R}{4\pi}\right]\left[T_a^2\,\frac{E_L}{h\nu_L}\right]\Delta z$$

where

z = altitude

ρ_R = air density

σ_R = Rayleigh cross - section

η = combined efficiencies

A_R = telescope area

T_a = transmission of the atmosphere below 30km

E_L = transmitted laser energy

ν_L = laser frequency

Δz = range resolved bin size

t = time

Eq.2.6

The three terms grouped into C represent, in order, the combined electronic, optical, and collection efficiencies, the number of photons reaching the altitude range, and the size of the altitude range (or 'bin'). To determine the number of counted photons (or 'counts')which will be received from Rayleigh scattering, then, it is necessary to know the atmospheric density and Rayleigh cross-section at the altitude range from which the signal originates. Fortunately, for small laser frequency shifts, on the order of gigahertz, the Rayleigh cross section is approximately constant. Also, though the transmission of the "dirty" lower atmosphere is very changeable and difficult to determine, the atmospheric density above 30km, albeit somewhat changeable, may be successfully approximated using tabulated 'standard atmospheres' (Yu, 1994).

Sodium Fluorescence

Once the photons transmitted by the lidar arrive at the sodium layer, the resonance interaction described in section 2a above can proceed. Assuming that the altitude range of interest is small enough that only a small fraction of the photons arriving at its bottom are absorbed before moving to the next range bin, the photon counts from the first range bin of the sodium layer may be written:

$$N_s(z,\nu_L,t,T,V) = \frac{C}{z^2}\left[\rho_s(z,t)\sigma_{sh}(\nu_L,T,V)\right] \qquad \text{Eq.2.7}$$

Where C in Equation 2.7 is the same as C in Equation 2.6, serving to account for the number of photons involved in the interaction which go on to be counted in the laboratory. The sodium density, as was discussed in Chapter I, arises from meteor ablation and is both changeable and difficult to approximate theoretically. The total sodium back-scattering cross-section, $\sigma_{sb}(\nu_L, T, V)$, may be calculated approximately from the fluorescence lineshape of Equation 5 using:

$$\sigma_{sb}(\nu,T,V)=\frac{\sigma_{sa}(\nu,T,V)}{4\pi} \qquad \text{Eq.2.8}$$

Where $\sigma_{sa}(\nu_L,T,V)$, the total absorption cross-section, is given by:

$$\sigma_{sa}=A_{21}\left[\frac{g_2}{g_1}\right]\left[\frac{\lambda_0^2}{8\pi}\right]g_L(\nu,T,V)$$
$$g_L=COR\left[L(\nu),g(\nu,T,V)\right] \qquad \text{Eq.2.9}$$

For the purposes of the lidar measurement, g_1/g_2, A_{21}, and λ_0 are constants representing, in order, the degeneracy ratio and Einstein coefficient for the sodium D_2 transition (2 and 6.29GHz), and the mean wavelength (589nm) of the exciting laser (She et. al. 1992). The $L(\nu)$ in the definition of $g_L(\nu, T, V)$ is the lineshape of the exciting laser, normalized to unity area. Under the previous assumption of an idealized, purely monochromatic laser, the correlation function simply yields $g_L(\nu, T, V) = g(\nu,T,V)$ from Equation 2.5 and Figure 2.1. In the case of a pulsed laser, however, some Fourier-limited bandwidth is inevitable, and must be accounted in the laser fluorescence lineshape which therefore receives an added subscript L. Once the correlation has been performed, the ν in its functional dependence may be defined to refer to the 'centroid' of the laser lineshape, given by:

$$\text{centroid of } f(x)\equiv x_{cent}=\frac{\sum_{all\,x} xf(x)}{\sum_{all\,x} f(x)} \qquad \text{Eq.2.10}$$

The centroid serves as an unambiguous point on any spectrum by which to reference its overall location. Therefore, one may identify the laser's output spectrum in absolute frequency by knowing its lineshape and the location of its centroid. Since the effect of a frequency-tunable lineshape interacting with another spectrum is common in spectroscopy using tunable pulsed lasers, the identification of the frequency location of a spectrum using its centroid reoccurs in many forms throughout the lidar analysis.

From Equation 2.9 it is apparent that the temperature and velocity sensitivity of the backscattering cross-section, and therefore of the lidar experiment, arises entirely from that calculated for the fluorescence lineshape $g(\nu,T,V)$ in Equation 2.5 and Figure 2.1.

An approximation used to simplify Equation 2.8 assumes isotropic, or spherically distributed, re-radiation of the fluorescence photons from the sodium atoms. This simplification makes the meaning of the total backscattering cross-section, $\sigma_{sb}(\nu_L, T, V)$, intuitively clear, applying simply those photons from the uniform distribution which happen to radiate back along the direction of the laser beam. In fact, the scattering is not necessarily isotropic, especially if different radiated radiated photon polarizations are considered separately. A thorough quantum-mechanical treatment of the cross-section reveals more complicated quantum interference terms which Equation 2.8 ignores (H. Chen, 1997). This effect is often referred to as the Hanle Effect or the Zero-field Hanle Effect, and may account for lidar temperature errors of less than 1K and velocity errors of less than 2m/s (Frike and von Zahn, 1985).

2c. Beam Attenuation

As the laser beam ascends through the sodium layer, some of the photons it carries are absorbed by the atmospheric sodium and re-emitted in other directions. This, while providing the signal on which the lidar measurement is based, also serves to attenuate the beam as it ascends. Similarly, backscattered fluorescence signal from altitude bins high in the sodium layer must also traverse kilometers of sodium gas in which it can be absorbed and re-emitted, thereby attenuating the beam which carries information back to the lidar receiver. Details of each, the ascending attenuation and descending attenuation, will be discussed in turn below.

Ascending Transmission, $T_\uparrow$

Since the attenuation of the beam as it ascends through the sodium layer is caused by laser light being absorbed by the sodium and scattered in all directions, it is the total absorption cross-section which governs the details of the interaction. The fractional transmission of the beam as it ascends from an altitude z_L (designating the bottom of the observed sodium layer) to a given range bin at z may be written:

$$T_\uparrow = \exp(-\int_{z_L}^{z} \rho_s(z',t)\sigma_{su}(\nu_L,T,V)dz') \qquad \text{Eq.2.11}$$

At the lowest range bin of the observed sodium layer, the integral in Equation 2.11 yields zero, making $T_\uparrow$ unity. This suggests that before the beam enters the sodium layer the fractional upward transmission factor is unity, or full transmission, as expected. Any beam attenuation before reaching the sodium layer must be accounted elsewhere, such as from the total Rayleigh scattering. It is also interesting to note that the sodium density enters into Equation 2.11. Thus, to determine the beam strength reaching a given range z, the sodium density of all range bins below z must first be determined.

Descending Transmission, $T_\downarrow$

After the transmitted photons are absorbed and backscattered from a given range bin to be detected by the lidar receiver, as per Equation 2.7, the journey of the emitted photons is subject to attenuation by the sodium layer below. Additional complexity is leant to the analysis of this downward attenuation by the fact that the emitted photons arise from ten fluorescence pathways available in the sodium D_2 transition, and that the frequency and strengths of all ten lines depend on the frequency of the exciting laser.

A moving sodium atom may only absorb a photon Doppler-shifted to the atoms resonance frequency. This serves to broaden the bandwidth of photons which may be absorbed. However, since the Doppler width is far broader than the laser line shape, the range of photon frequencies available to the moving atom in the presence of the lidar beam is limited by the laser line shape, and so the emitted lines can have linewidths only as large as that of the exciting laser. The frequency centroid of each line obeys the condition:

$$\nu_F = \left[\nu_{ge} + \nu_{eg'}\right] - \nu_L \qquad \text{Eq.2.12}$$

in which ν_f signifies the fluorescence centroid frequency, ν_{ge} signifies the photon frequency associated with the excitation, $\nu_{eg'}$ signifies the photon frequency associated with the fluorescence decay to one of the available ground states, and ν_L signifies the laser centroid frequency. Though the laser centroid frequency might change, the fluorescence line frequencies for the sodium D_2 transition may be tabulated in terms of the constant, numerical values of $\nu_{ge}+\nu_{eg'}$, and the unknown ν_L:

Table 2.2: D_2 Fluorescence Frequencies

j	ν_{fj}(MHz)
1	-1243.16 - ν_L
2	-1361.1 - ν_L
3	-1430 - ν_L
4	341.59 - ν_L
5	341.59 - ν_L
6	410.52 - ν_L
7	410.52 - ν_L
8	2081.6 - ν_L
9	2113.18 - ν_L
10	2182.14 - ν_L

Note: frequencies are reported relative to the weighted D_2 line center. ν_L signifies the frequency of the exciting laser.

For a known laser frequency, then, the frequencies of all ten return lines from the sodium D_2 transition may be calculated.

Since the relative strength of each backscattered line depends on the backscattering efficiency at its frequency. This may be determined from the back-scattering cross-section of Equation 2.8, $\sigma_{sb}(\nu, T, V)$ using Equation 2.13.

$$P_j(\nu_j,T,V) = \frac{\sigma_{sb}(\nu_j,T,V)}{\sum_{i=1}^{10}\sigma_{sb}(\nu_i,T,V)} \qquad \text{Eq.2.13}$$

$P_j(\nu_j, T, V)$ represents the fraction of the total scattered photons which exit from the j^{th} of the ten fluorescence pathways. If all the fractions are summed from j = 1 to 10 the result of unity shows that all of the backscattering photons have been accounted for.

Given the frequencies and relative strengths of the returning beams, it is possible to write the descending fractional transmission:

$$T_{\downarrow} = \sum_{j=1}^{10} P_j(\nu_j,T,V)\exp\left(-\int_{z_L}^{z}\rho_s(z',t)\sigma_{sa}(\nu_j,T,V)dz'\right) \qquad \text{Eq.2.14}$$

$T_{\downarrow}$ represents the fraction of the total returning beam which is transmitted through the sodium layer from a range bin at altitude z to the bottom of the measured sodium layer at z_L. As in $T_{\uparrow}$, if the range bin in question is at the bottom of the layer the exponential and the sum each yield unity, signifying full transmission. Also as in $T_{\uparrow}$, the integral contains both the sodium density for all range bins below the emission, and the total absorption cross-section by which photons are scattered out of the returning beam, in all directions. It may also be noted that, though not explicit in the notation of Equation 2.14, the determination of ν_j and therefore the values of the linestrengths and cross-sections, implicitly depends on the laser centroid frequency, ν_L, from Equation 2.12 and table 2.2.

With mathematical expressions for the physical processes involved in the altitude-resolved laser excitation of the sodium layer and retrieval of the resulting fluorescence scattering, it is possible to accurately address the theory of the 3-frequency lidar temperature and velocity measurement. This along with consideration of the problems faced in the design of the new lidar follow in Chapter 3.

Chapter 3: The Theory of Detection and Problems of Lidar Design

The object of this chapter is to discuss the method by which the 3-frequency lidar can obtain sodium density, temperature, and velocity information from the atmosphere. To do this, the expressions generated in the previous chapter for the physical phenomena of laser induced fluorescence and atmospheric attenuation of the laser beam will be inverted to yield the desired atmospheric parameters. The choice of laser frequency at which the lidar operates will be discussed based on the sensitivity to temperature and collective velocity of the resulting measurement, and the calibration of the lidar to accurately yield atmospheric parameters from the collected backscattering of the laser at these frequencies. Hardware design considerations and experimental uncertainties implied by the inversion method will then be discussed.

3a. The Lidar Operating Frequencies

From Equations 2.7, 2.8, and 2.9 of the previous chapter it was found that the temperature and velocity sensitivity of the fluorescence backscattering from the sodium layer follows from the absorption lineshape $g(\nu)$ found in Equation 2.5 and shown in Figure 2.1. The question which now arises is one of how to best capitalize on the sensitivities in $g(\nu)$ with the 3-frequency lidar. In order to better visualize these sensitivities, the change in $g(\nu)$ due to a 20K change in temperature from 200K, and due to a 20m/s change in Doppler velocity from 0m/s are plotted in Figure 3.1, together with the sum of the absolute values of each (see following page).

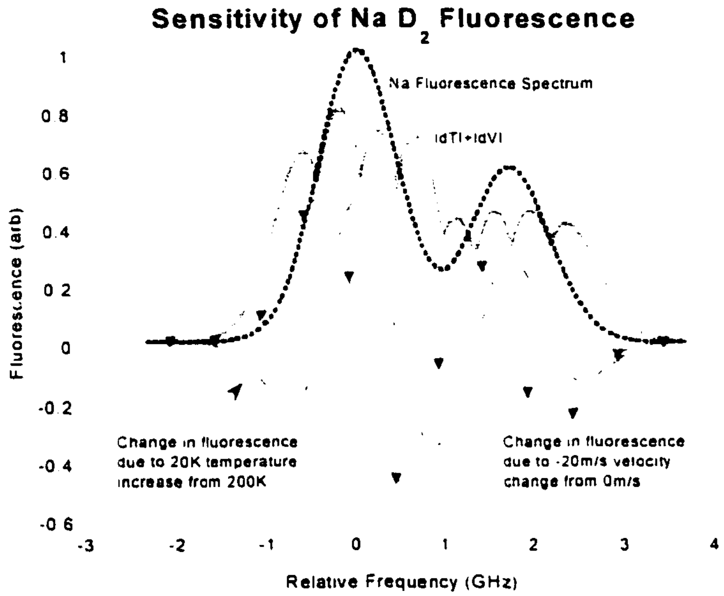

Fig.3.1 The sensitivity of the fluorescence spectrum to a 20m/s change in velocity and a 20K change in temperature as a function of frequency relative to the D_{2a} Doppler-free feature. The sum of the absolute value of each is also plotted to guide the eye in determining frequencies of dual sensitivity to temperature and velocity.

The frequency scale of Figure 3.1 is referenced against the sodium D_{2a} Doppler-free feature, which is used as the central locking point of the 3-frequency lidar. It is apparent from Figure 3.1 that the frequencies which yield the highest sensitivity to temperature do not necessarily have good sensitivity to velocity. This is especially apparent at the D_{2a} peak (0MHz) where the temperature sensitivity is high, but the velocity sensitivity disappears. Since two ratios will be generated from different combinations of the same three fluorescence response levels, it is necessary that both the temperature and velocity information be derivable from just these three values. The ratios, however should not each be sensitive to both temperature and wind to minimize the interdependence of the measurements. Luckily, the response to a change in temperature is symmetric about the D_{2a} peak while the response to velocity changes sign near the D_{2a} peak, and therefore, while the fluorescence responses at the three frequencies must be sensitive to both temperature and velocity, ratios can be formed which are predominantly sensitive to only one or the other.

For reasons inherent to the 3-frequency lidar, its frequency locking system, and frequency-shifter design, the choice of operating frequencies is limited to include the Doppler-free D_{2a} peak, and any equal shift both up and down from the D_{2a} peak: ν_{D2a} and $\nu_{D2a} \pm \Delta\nu_{shift}$. One consequence of this restriction which is apparent from Figure 3.1 is that almost all of the sensitivity to wind must be gained from the shifted frequencies. To aid in visualizing the simultaneous optimization of both sensitivities, the sum of the absolute values of both sensitivities is plotted in Figure 3.1. Several peaks in this sum are roughly equidistant in frequency from the center and therefore represent likely choices for the shifted frequencies. One set of peaks suggests a shifted frequency around 200MHz and the other suggests around 650MHz. Some differentiation of these choices comes from the fact that the outside peaks have slightly higher contributions from the velocity sensitivity than the inside peaks, which is an advantage since ν_{D2a} already carries high temperature sensitivity. Other considerations which may be weighed in selection of the lidar's operating frequencies include difficulties in the frequency shifter design, the signal-to-noise ratio arising from the frequency-dependent signal return strength from the atmosphere, and the compatibility of some frequencies with possible narrow-band filtering in the lidar receiver.

From more careful calculations conducted at CSU, an optimum shift frequency of 630MHz was determined, finalizing the lidar's three operating frequencies relatively at $\nu_p \equiv 0$MHz for the D_{2a} peak frequency (corresponding to a laser wavelength around 589.16nm), $\nu_+ = 630$MHz, and $\nu_- = -630$MHz. The pulsed laser transmitted by the lidar is intended to be centered at these three frequencies alternately, and signal returned from the atmosphere for each frequency is integrated separately since the fluorescence information which allows the determination of temperature and velocity is distinct for each.

3b. The Experimental Ratios and Calibration Curves

From Figure 3.1 it is apparent that the response at the two shifted frequencies to a positive velocity are in opposite directions, one yielding more signal and the other less, while the responses to a temperature increase are in the same direction, both yielding higher signal. From this information two ratios are suggested, R_T sensitive to temperature, and R_V sensitive to collective velocity:

$$R_T = \frac{g_L(\nu_-,T,V)+g_L(\nu_+,T,V)}{2g_L(\nu_p,T,V)} = \frac{\sigma_{sh}(\nu_-,T,V)+\sigma_{sh}(\nu_+,T,V)}{2\sigma_{sh}(\nu_p,T,V)}$$

$$R_V = \frac{g_L(\nu_-,T,V)-g_L(\nu_+,T,V)}{g_L(\nu_p,T,V)} = \frac{\sigma_{sh}(\nu_-,T,V)-\sigma_{sh}(\nu_+,T,V)}{\sigma_{sh}(\nu_p,T,V)} \quad \text{Eq.3.1}$$

Recall from Equations 2.8 and 2.9 that the subscript 'L' identifies that a correlation with the measured laser lineshape function has been enacted on $g(\nu)$ from Figure 3.1, and that the second set of equalities neglect the quantum interference contribution to the backscattering ratio. Since the pulsed laser lineshape is a measured quantity, a fully analytical, closed-form evaluation of the response of these ratios to changing temperatures and velocities is not possible. However, numerical calibration to high accuracy can be achieved, yielding a 2-dimensional calibration curve similar to that of Figure 3.2.

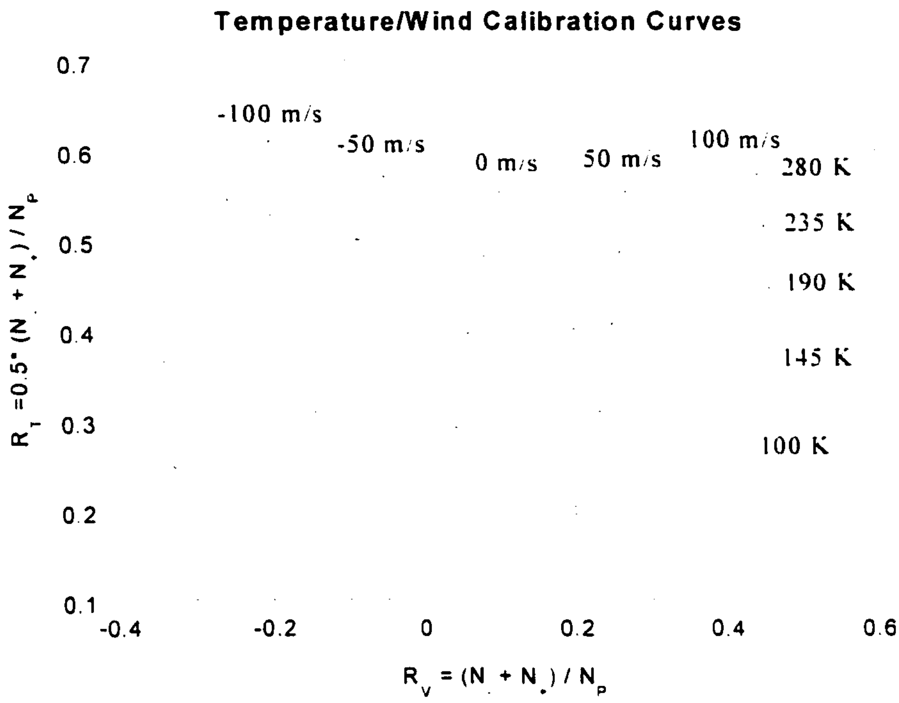

Fig.3.2 The 2-dimensional field of calibration curves for mapping points of known ratios (R_T, R_V) into temperature and radial velocity. The curves arise from theoretical spectra of sodium fluorescence.

Using a calibration such as shown in Figure 3.2, but with a higher resolution grid, any measured pair of R_T and R_V may be converted into an equivalent temperature and radial velocity pair. For ratio values which fall between grid lines, a polynomial expansion based on the temperature and velocity lines which form the grid cell in which the ratios lie may be computed, yielding temperature and velocity of higher precision than the chosen grid spacing. The process by which the ratios R_T and R_V pertinent to a certain range bin in the sodium layer are obtained is the subject of the following sections.

3c. The Lidar Equation

The experimental ratios of the previous section use the sodium absorption spectrum to extract the temperature and velocity from the backscatter cross-sections at three chosen operating frequencies. To extract the backscatter cross-sections pertinent to a given region of the atmosphere we use the so-called 'lidar equation', which, in the notation developed in Chapter 2 may be written:

$$N_S(z,\nu_L,t,T,V) = \frac{C}{z^2}\left[\rho_s(z,t)\sigma_{sb}(\nu_L,T,V)\right]T_\uparrow(\rho_s(z,t),\sigma_{sa}(\nu_L,T,V),z)T_\downarrow(\rho_s(z,t),\sigma_{sa}(\nu_L,T,V),z)$$

where

$$T_\uparrow = \exp\left(-\int_{z_L}^{z}\rho_s(z',t)\sigma_{sa}(\nu_L,T,V)dz'\right)$$

$$T_\downarrow = \sum_{j=1}^{10} P_j(\nu_j,T,V)\exp\left(-\int_{z_t}^{z}\rho_s(z',t)\sigma_{sa}(\nu_j,T,V)dz'\right)$$

Eq.3.2

Equation 3.2 combines the various phenomena investigated in Chapter 2. The first term, C/z^2, from Equation 2.6 accounts the total number of photons available to be counted by the lidar receiver and its efficiencies, this assumes that any background sources of photons have already been subtracted away before constructing Equation 3.2. The second term, $\rho_s\sigma_{sb}$, from Equation 2.7 determines the number of fluorescence photons created in the given range bin, the third term, $T_\uparrow$, from Equation 2.11, calculates the partial transmission of the exciting laser on its way through the sodium layer to the altitude of interest, and the fourth term, $T_\downarrow$, From Equation 2.14, using P_j from Equation 2.13, calculates the partial transmission of the backscattered fluorescence photons on their way through the sodium layer back to the lidar receiver. Equation 3.2 forms the basis for understanding the information imbedded in the laser excited, range-resolved, returning photons from the atmosphere.

To form the ratios of Equation 3.1, and thereby derive the temperatures and Doppler velocities via the calibration curves of Figure 3.2, Equation 3.2 may be inverted to determine the backscattering cross-section from the observed photocounts and other parameters for each operating frequency of the lidar.

$$\sigma_{sb}(\nu_L,T,V) = \frac{N_S(z,\nu_L,t,T,V)}{\left[\frac{C_i}{z^2}\right]\rho_s(z,t)T_\uparrow(\rho_s(z,t),\sigma_{sa}(\nu_L,T,V),z)T_\downarrow(\rho_s(z,t),\sigma_{sa}(\nu_L,T,V),z)} \equiv \frac{N_{Si}}{\left[\frac{C_i}{z^2}\right]\rho_s T_{\uparrow i} T_{\downarrow i}}$$

where $i = -, p, +$

Eq.3.3

The three forms of Equation 3.3 are all equivalent to Equation 3.2, the lidar equation, where the defined equalities omit the explicit functionality in favor of notational brevity. The three backscatter cross-sections in Equation 3.3 may be inserted directly into the experimental ratios of Equation 3.1:

$$R_T = \frac{\left[\dfrac{N_{S-}}{C_- T_{\uparrow -} T_{\downarrow -}}\right] + \left[\dfrac{N_{S+}}{C_+ T_{\uparrow +} T_{\downarrow +}}\right]}{2\left[\dfrac{N_{Sp}}{C_p T_{\uparrow p} T_{\downarrow p}}\right]}$$

$$R_V = \frac{\left[\dfrac{N_{S-}}{C_- T_{\uparrow -} T_{\downarrow -}}\right] - \left[\dfrac{N_{S+}}{C_+ T_{\uparrow +} T_{\downarrow +}}\right]}{\left[\dfrac{N_{Sp}}{C_p T_{\uparrow p} T_{\downarrow p}}\right]} \qquad \text{Eq.3.4}$$

The first barrier to evaluating the right-hand side of Equation 3.4, and thereby having experimental values for the two ratios R_T and R_V, is that the constant C_i terms are different for each frequency and not yet determined. The variation of C_i with the operating frequency of the laser is pronounced in the 3-frequency lidar transmitter due to the imperfect conversion of the laser power in the frequency shifting system which produces the + and – 630MHz beams. To experimentally determine the three C_i's, the Rayleigh signal from the aerosol-free lower atmosphere may be employed as a laser power and atmospheric transmission monitor, in the approximation that the atmosphere above 45km to the mesopause region introduces very little additional attenuation. From Equation 2.6 an experimentally determined value of C_i may be determined:

$$C_i = \frac{z_R^2 N_{Ri}}{\rho_R(z,t)\sigma_R} \quad \text{where } i = -, p, + \qquad \text{Eq.3.5}$$

These expressions for C_i may be inserted into Equation 3.4 to yield the final expression for the experimental ratios which may be determined iteratively from return signals originating in both the sodium layer and the 30 to 45km clean-air Rayleigh region.

$$R_T = \frac{\left[\dfrac{N_{S-}}{N_{R-} T_{\uparrow -} T_{\downarrow -}}\right] + \left[\dfrac{N_{S+}}{N_{R+} T_{\uparrow +} T_{\downarrow +}}\right]}{2\left[\dfrac{N_{Sp}}{N_{Rp} T_{\uparrow p} T_{\downarrow p}}\right]}$$

$$R_V = \frac{\left[\dfrac{N_{S-}}{N_{R-} T_{\uparrow -} T_{\downarrow -}}\right] - \left[\dfrac{N_{S+}}{N_{R+} T_{\uparrow +} T_{\downarrow +}}\right]}{\left[\dfrac{N_{Sp}}{N_{Rp} T_{\uparrow p} T_{\downarrow p}}\right]} \qquad \text{Eq.3.6}$$

3d. Iterative Determination of Temperature, Velocity, and Na Density

The First Range Bin

Equation 3.6 may be solved one range bin at a time starting from the lowest range bin of the sodium layer, the lowest bin which satisfies some minimum returned counts criteria. The altitude at the bottom of the sodium layer is designated z_L in the defining equations of $T\uparrow$ and $T\downarrow$. Since neither the transmitted laser beam nor the return signal from the first bin is subject to attenuation due to atmospheric sodium below the bin, $T\uparrow$ and $T\downarrow$ each equal unity. This, then, renders Equation 3.6 entirely in terms of measured photo counts, N_S and N_R, from the first bin of the sodium layer and the clean air Rayleigh scattering region, respectively, at each of the three operating frequencies.

$$R_{Tfirstbin} = \frac{\left[\frac{N_{S-}}{N_{R-}}\right] + \left[\frac{N_{S+}}{N_{R+}}\right]}{2\left[\frac{N_{Sp}}{N_{Rp}}\right]}$$

$$R_{Vfirstbin} = \frac{\left[\frac{N_{S-}}{N_{R-}}\right] - \left[\frac{N_{S+}}{N_{R+}}\right]}{\left[\frac{N_{Sp}}{N_{Rp}}\right]} \qquad \text{Eq.3.7}$$

From the ratios R_T and R_V, then, the temperature and the radial velocity may be calculated for the first bin of the sodium layer using a calibration curve such as shown in Figure 3.2, establishing the first data point on a range-resolved profile of temperature and velocity in the mesopause region.

To calculate the temperature and velocity of the next higher bin, however, the effects of imperfect transmission through the lower bin cannot be neglected. This is indicated formally by the fact that $T\uparrow$ and $T\downarrow$ are no longer unity because the integrals in their definitions (see Equation 3.2) include the Na density of the lowest bin. Therefore, it is necessary to determine the density of sodium atoms in the first bin in order to determine the temperature and radial velocity in the second bin, then to determine the density in the second bin to be used with the density in the first bin to determine the temperature and velocity in the third bin, and so on, until the calculation of the temperature and velocity of the final bin in the sodium layer, which utilizes the calculated Na density in all of the lower bins.

The Na Density Calculation

To solve for the Na density for any range bin in which the temperature and velocity has already been determined, Equation 3.2 may be inverted to isolate the sodium density in the second term, and C_i from Equation 3.5 substituted to yield Equation 3.8:

$$\rho_s(z,t) = \rho_R(z,t)\frac{N_{si}\sigma_R z^2}{N_{Ri} z_R^2 \sigma_{sh}(\nu_i, T, V)} \quad \text{where } i = -, p, + \qquad \text{Eq.3.8}$$

Equation 3.8 provides a means of calculating the density for a range bin given the measured photocounts at that range bin and in the Rayleigh return, the laser frequency, the temperature and radial velocity of the bin being calculated, the densities, temperatures, and velocities for all range bins below the bin being calculated (to determine T↑ and T↓), and the air density in the clean-air Rayleigh region. In the iteration process to determine temperature and radial velocity, all of the parameters required to determine the sodium density have been calculated with the exception of the air density in the clean-air Rayleigh region. This value must be approximated using a standard value according to the Rayleigh altitudes used. Uncertainty in this approximation is reflected directly in the sodium density calculation, but effects the calculation of wind and temperature less since the experimental ratio cancels the primary dependence on Na density, and the current calculation of density enters only in the transmittance calculations which vary over a smaller range, and tend also towards cancellation in the ratios.

Successive Range Bins

To clarify the density measurement and its place in the iterative procedure for determining the altitude-resolved temperature, radial velocity, and Na density, consider Equation 3.8 as applied to the lowest range bin in the sodium layer. At the first range bin T↑ and T↓ are each unity, so the temperature and velocity may be calculated directly from the returning photon counts using Equation 3.7. Once the temperature and velocity are established, the backscatter cross-section for a given laser frequency may be calculated (Equation 2.8), and so Equation 3.8 above may be evaluated because T↑ and T↓ are each unity. Note that the determination of Na density from Equation 3.8 requires only one operating frequency from the lidar transmitter.

Once the lowest range bin in the sodium layer has been evaluated for temperature, velocity, and Na density, then T↑ and T↓ may be evaluated for the next higher bin, since they represent the transmission encountered by photons in getting to and away from a given altitude. Once T↑ and T↓ have been evaluated, the temperature and velocity of the next higher bin are calculated using Equation 3.6, then the density using Equation 3.8, and so T↑ and T↓ for the next bin higher still. This is the iterative procedure which produces the lidar data profiles from the range-resolved photon counts received by the lidar acquisition system.

3e. Uncertainties and Lidar Hardware Considerations

The analysis of the measurement theory presented in Section 3d above suggests various sources of uncertainty which may be confronted in the conversion of lidar return signals into atmospheric data. Some such uncertainties may be minimized or eliminated with the help of far-sighted hardware design, while others are intrinsic to the measurement technique. A brief summary of some sources of uncertainty, together with the approximate magnitude of the contribution of each are presented below.

Ratio-timescale Fluctuations

In at least two instances in the analysis of Section 3d., atmospheric parameters were removed from expressions due to cancellation from the numerator to the denominator of a ratio: Once for sodium density in the step leading to Equation 3.4, and once for air density in the definition of C_1 leading to Equation 3.6. In both cases the cancellation occurs between terms representing the different operating frequencies of the lidar transmitter.

While mathematically valid as written, these cancellations assume that the variables in question have not changed on the timescale at which the lidar transmitter changes frequency, since at least one full cycle of the operating frequencies is required to create one experimental ratio. In some lidar transmitter designs it might be beneficial from an instrumentation point of view to remain at each operating frequency for the maximum time commensurate with the desired atmospheric temporal resolution, remaining at each of three frequencies for twenty minutes for a one-hour average of the atmosphere, for instance. In such a case, however, the cancellation of terms between operating frequencies is not necessarily an accurate approximation. Especially in the case of sodium density, which has been observed to exhibit large magnitude, short duration fluctuations (as in the presence of 'sporadic sodium layers') which also affect the variation with altitude of data from the mesopause region, long integration times may translate into large errors in temperature and velocity measurements.

The hardware solution for this problem is simply to increase the rate at which the lidar transmitter can switch between its operating frequencies. For this reason, as well as others, the CSU 3-frequency lidar has adopted an acousto-optical frequency shifting system which allows shot-by-shot cycling between the operating frequencies, effectively integrating all three data channels simultaneously and minimizing ratio-timescale fluctuations.

Laser Frequency Stability

The analysis of the detection theory made use of the three operating frequencies of the lidar, but in practice, locking a lidar transmitter precisely to a single frequency for long data collection periods can be problematic, and with the introduction of additional operating frequencies can prove more so. Many pulsed laser transmitters are prone to systematic frequency shifts and rapid frequency modulations or 'chirps', either of which might jeopardize assumptions made in calculating the calibration curves which transform the experimental ratios into atmospheric parameters.

The effect of an overall shift in all three operating frequencies has profound consequences on the measurement of radial velocity, but less impact on the measurement of temperature and Na density. The reason for this is that since the Doppler shift of the sodium absorption spectrum is used to determine the collective velocity, any laser frequency shift arising from sources other than a collective velocity are interpreted as additional velocity. The resulting error in the velocity of the wind, therefore, is directly proportional to the frequency offset, at a ratio of .6m/s per 1MHz of laser offset. Temperature and density measurements, however, benefit from the presence of the velocity measurement, which serves as a frequency monitoring measurement to minimize the effect of an overall shift on these measurements. Further discussion of this relationship follows in Chapter 7.

Another instance of laser frequency instability which would detrimentally affect the reliability of acquired lidar data arises in the presence of a laser shift or chirp which is not equal among the three operating frequencies. The extreme dependence of frequency chirp in some laser amplifiers on the exact alignment of the amplified beam suggests that for slight alignment discrepancies between the operating frequency beams, non-uniform frequency discrepancies might result. The response in the obtained temperature and velocity measurements to such a shift would depend on the operating frequencies which were shifted and in which direction, but inspection of the sodium absorption spectrum and calibration curves suggest that a 1MHz shift in one operating frequency could shift the resulting temperature by about 0.1K and the radial velocity by about 0.3 m/s.

Photon Noise

Since photon collection is a counting process, the statistical noise is governed by Poisson statistics, according to which the uncertainty in any sum N is the square-root of N. For this reason more signal, and therefore larger sums, is always preferable from the point of view of fractional uncertainty. In lidar signal collection photon noise is often the dominating error source, especially when ratios and differences of sums compounds the total fractional uncertainty through error propagation. The solution to unacceptably high photon noise is to acquire more photon counts, typically either by increasing the number of photons per laser pulse (pulse energy), transmitting more pulses per unit time (pulse repetition rate), or by increasing the total area of the lidar collecting telescope. For this reason, a quantity often quoted to represent the collection capability of a lidar is the 'power-aperture-product' in which the time-averaged power serves to express the number of photons transmitted, and the receiver telescope aperture diameter serves to express the lidar's collecting capacity.

One disadvantage of increasing a lidar transmitter's pulse energy is the possibility of saturating the absorption capacity of the atmosphere's sodium atoms. The result of such saturation is a lower signal level from the mesopause region than expected, and a disproportionate reduction between the peak frequency, at which absorption is high and therefore more easily saturated, and the shifted frequencies, which experience less saturation-induced reduction. The effect on the experimental ratios, then, is to reduce the denominator more than the numerator, resulting in erroneously increased temperatures and velocity magnitudes. Quantum mechanical calculations (H. Chen, 1997) suggest that the saturation intensity for the Na D_2 line is 38mW/cm^2. If twice the saturation intensity is transmitted, the temperature measurement is increased by about 1K and the magnitude of the measured wind is increased by about 0.1 m/s. Because of saturation effects, increasing the pulse repetition rate is often the preferred way to greatly increase the averaged power of a lidar transmitter.

Because the photon noise depends on the signal strength returning from the mesopause region, the absolute uncertainty of any range-resolved profile changes with altitude, with large uncertainties at the bottom and top of the sodium layer where low sodium density provides only weak signal strength, and the smallest uncertainties at the sodium density peak or centroid. Error bars are computed for each range bin depending on the strength of the signal originating there.

Transmitter and Receiver Bandwidth

The lineshape of the transmitted laser enters into the detection theory in the calculation of the backscattering cross-section, where a correlation between the laser lineshape and the sodium absorption spectrum is performed (Equations 2.8 and 2.9). Since the correlation function creates a smoothing effect on the absorption spectrum, it is clear that a laser lineshape of excessive width (such as on the scale of the absorption spectrum itself) could 'smooth away' all of the features in the resulting cross-section which lend the measurement its sensitivity. For this reason it is essential to have a narrow-band transmitter to ascertain the temperature and wind from sodium fluorescence. Similarly, a change in the lineshape of the laser changes the calibration curve which transforms the measured ratios to atmospheric parameters. Therefore, stability of the transmitter lineshape, often obtained by injection seeding the transmitter's pulsed laser, is also important. However, since the full-widths-at-half-maximum of Gaussian functions which have been correlated add in quadrature, the effect of a small change in a narrow band laser's lineshape on the outcome of a correlation with a much broader absorption spectrum is likely to be negligible.

For lidar operation in daylight hours, narrow-band filtering may be used in the lidar receiver, raising additional sources of uncertainty. The overlap of the filter's transmission function with the operating frequencies could reduce the transmission at one or more frequency to unacceptably low levels, additional complexity is required to correctly deduce the filter's transmission of the ten fluorescence return lines from the sodium D_2 transition, single-polarization transmitting filters both eliminate one-half of the returning, unpolarized fluorescence signal and increase the importance of quantum interference terms in the backscattering cross-section, and many established sources of uncertainty, such as laser frequency drift, are amplified. The scientific relevance of the daytime operation of sodium lidars, however, continues to encourage earnest effort in narrow-band filtering in Na fluorescence receiver designs (H. Chen, 1997, S. Chen, 1999, Chen et. al. 1996).

Chapter 4: The 3-Frequency Temperature Wind Lidar Design

4a. Basic Theory of Operation

The optical system of the lidar transmitter is most easily separated into a CW (continuous wave) and a pulsed subsystem. The outputs of these subsystems are combined in the Pulsed Dye Amplifier (PDA) which is the final stage in the transmitter before the beam exits for frequency monitoring and the atmosphere. Following the transmitter, the detection sub-system collects light from the atmosphere for photon counting and time-of-flight ranging.

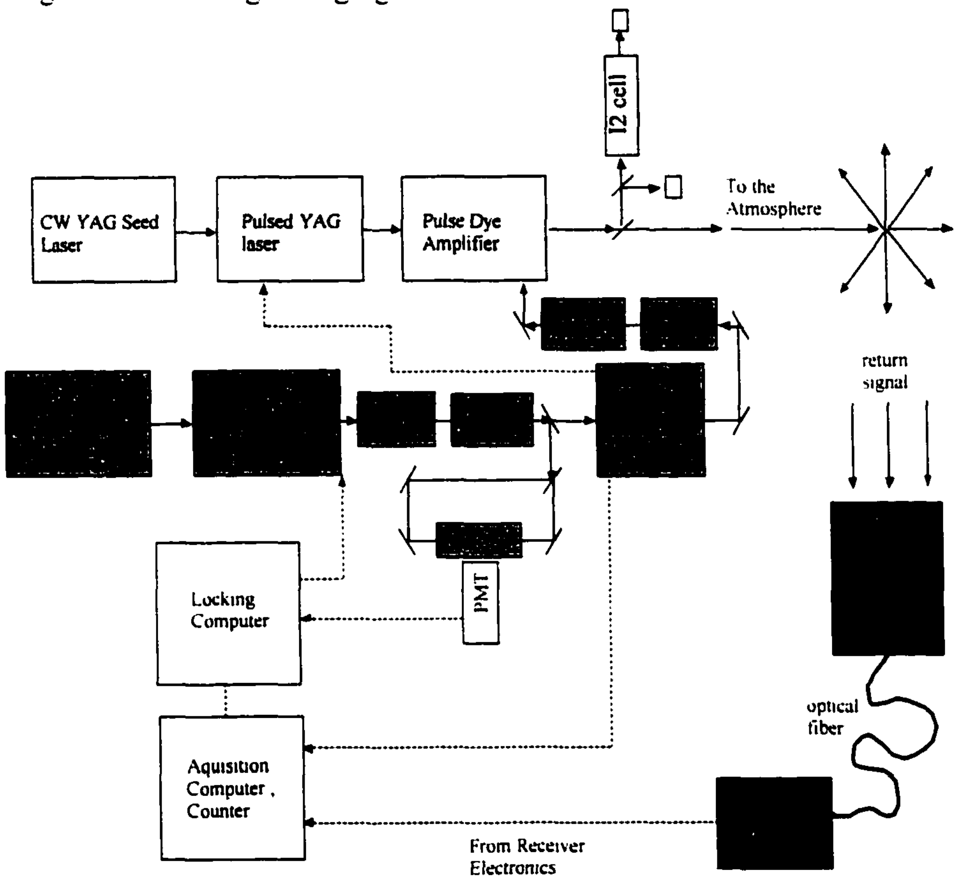

Fig.4.1 A Block diagram of the 3-frequency lidar transmitter and receiver. Electronic control lines are shown dashed, and optical paths are solid.

The CW Subsystem of the Lidar Transmitter

The CW system is comprised of a frequency-doubled CW YAG pump laser (Spectra-Physics Millenia model T40) which operates at a wavelength of 532nm with available power of 5 watts, but typically operated between 2 and 3 watts. The CW YAG output is the pump energy source for the Ring Dye Laser (Spectra-Physics model 380D) which is the heart of the CW subsystem in the lidar transmitter. The ring dye laser, together with its reference and slave etalon locking system (Spectra-Physics model 388) provides up to 500mw of widely tunable (>30GHz), narrow band (<1MHz bandwidth) output.

During data acquisition by the wind/temperature lidar, the ring-dye laser is actively locked to the center of the Lamb-dip located at the D_{2a} transition in the Doppler-free fluorescence spectrum of Na (589nm). The details of this active locking will be considered below.

Following the ring-dye laser, an optical isolator (OFR model IO-5-VIR), and collimating optics, about 4% of the laser beam energy is split out to a wavemeter for approximate frequency measurement (Burleigh WA-10), and the active-locking Na cell for diagnostic purposes. Following the principal beam, the CW transmitter subsystem incorporates a custom tandem acousto-optic modulator frequency-shifting unit. This AO unit utilizes, in order, a polarizing beam splitter, a periodic series of lenses and AO crystals, a quarter wave plate, and a back mirror to achieve double-pass frequency shifting by either AO crystal into the same exit beampath as is followed by unshifted light. A mechanical chopper wheel placed before the back mirror selects which beam order (shifted or unshifted) is allowed to traverse the unit, and provides the primary synchronization, switching microwave power to the crystals and triggering the pulsed transmitter subsystem through associated, custom logic circuitry. The purpose of the tandem AO unit is to provide the two additional frequencies at ±630MHz necessary for the wind/temperature measurement while allowing the ring-dye laser to remain actively locked to the single Na D_{2a} Lamb dip. More details on the tandem AO unit follow in Section 4c.

From the tandem AO system the CW beam, now cycling through the three operating frequencies, traverses another optical isolater to eliminate unwanted back reflections from the pulsed subsystem, and more beam conditioning optics for injection into the pulsed dye amplifier (PDA), and direct interaction with the pulsed part of the lidar transmitter.

The Pulsed Subsystem of the Lidar Transmitter

The heart of the pulsed part of the wind/temperature lidar transmitter is the pulsed, doubled Nd:YAG, a Spectra-Physics DCR-3. The DCR-3 fires at 20Hz with a time-averaged output power around 5W; its output pulses have been observed to be smooth and reproducible, possibly resulting from its injection seeded operation. The DCR-3 is seeded by a CW Nd:YAG which is coupled through an optical fiber, and Isolator, and collimating optics into the Q-switched cavity of the pulsed laser. The firing of the DCR-3 is controlled externally with the provided 'lamp' input, the control signal arriving from the custom control logic associated with the chopper wheel in the AO unit of the CW subsystem. The DCR-3 pulsed Nd:YAG laser provides the pump energy source for the pulsed dye amplifier.

The pulsed dye amplifier (PDA) combines the three-frequency CW beam from the CW subsystem with the pulsed Nd:YAG output from the pulsed subsystem in a single pass amplifier with three dye cell stages. The CW light constitutes the injected seed beam, the frequency of which the PDA output is intended to follow, while the pulsed laser output provides the pump energy lending the PDA the short pulse duration (about 4.5ns) and high pulse energy (30mJ/pulse) necessary for time-of-flight lidar applications. The PDA output is nearly Fourier-transform limited with a bandwidth of about 130MHz, and appears smooth and reproducible. The laser beam which leaves the lab for the atmosphere is about 1cm in diameter, with a full-angle divergence of about 1mrad. Details of the pulsed transmitter and its injection seeding are considered in Section 4d below.

Though the center frequency of the PDA output follows the center frequency of the injected seed light by well-known principles of laser gain competition in the amplifying path of the PDA, possible differences in the centroid frequency of the CW seed and the centroid frequency of the resulting output pulse has necessitated the use of detailed frequency monitoring of the PDA output pulse. This frequency monitor system is, for the purposes of this thesis, considered an addition to the 3-frequency lidar and will be considered separately in Chapter Seven.

The Detection sub-system

Once the transmitter has produced a beam of suitable quality directed towards the atmosphere, the return signal is detected and recorded by the detection sub-system.

The first stage of the detection is a 14 inch diameter Schmidt-Cassegrain telescope (Celestron-14) of commercial, amateur-astronomical quality. In typical operation, the telescope is leveled to point to zenith, while signal from the telescope is maximized by precise steering of the transmitted beam. The signal from the telescope is collected by an optical fiber optimally coupled to the output aperture of the telescope.

The optical signal traverses the optical fiber to the photon counting electronics, consisting first of a red-sensitive PMT, AC-coupled fast amplifier, and Discriminator. The return photons are thereby converted into well-behaved voltage pulses which continue to the counter for counting, time-of-flight binning, and summing over the selected integration time. The counter is triggered by the Q-switch output of the transmitter's pulsed laser to initiate the clock for the time-of-flight binning.

Since the frequency of the transmitted pulse alternates through each of the three operating frequencies at 20Hz during the integration time of lidar acquisition (typically 1 minute), the counting electronics must sort the data associated with each pulse according to operating frequency, so that summing of the signal in a given duration is done separately for each frequency. Similarly, additional care must be taken to ensure that if the firing sequence or counter triggering sequence is in some way perturbed, the frequency sorting is not likewise perturbed resulting in the summing of data of mixed frequencies for the remainder of the integration time and so invalidating it.

In addition to the lidar transmitter and detector, a sky-monitoring radar surveys the beam path for air traffic, and includes electronics to shut down the pulsed transmitter in the event of nearby aircraft, according to FAA guidelines. Details of these systems are provided in Section 4e below.

The remainder of this chapter will focus on design details of the lidar system. The intention is to provide the necessary information for researchers intending to construct the equivalent of the system described here, with a minimum of unnecessarily repeated investigation.

4b. The Doppler-free Locking System

Purpose

Since the Na lidar depends upon laser induced fluorescence of atmospheric Na atoms, the strength of the optical signal returning from the sodium layer depends on precise tuning of the laser frequency. The choice of optimum operating frequencies for temperature and velocity sensitivity outlined in Chapter 2 presupposed the possibility of maintaining those frequencies with good accuracy throughout lidar data acquisition. It is to this end that absolute frequency locking of the CW transmitter subsystem is intended. The use of a Na vapor cell as an absolute frequency reference in the vicinity of the operating frequencies is required, and when heated to around 70°C the vapor pressure of Na is suitable for easy detection. However, the absorption spectrum of Na, with a FWHM of about 1.4 GHz at 70°C, appears far too broad to be used as a frequency marker to lock with accuracies of a few MHz. The Doppler-free spectrum of Na offers a narrow (about 40MHz) "Lamb dip" spectral feature at the frequency location of the D_{2a} transition in sodium which is more suitable for accurate locking of the CW subsystem in absolute frequency to the determined operating frequencies within ~ 2MHz.

The Lamb Dip

The Lamb dip at the Na D_{2a} transition becomes easily observable, superimposed upon the Doppler-broadened fluorescence spectrum of Na, when counterpropegating beams are aligned inside the Na vapor cell with a detector collecting light emitted orthogonally to the beams. The Lamb dip is comprised of three distinguishable wells arising from independent hyperfine transitions, together yielding a FWHM around 40MHz. During routine active locking, the deepest, central well is selected as the locking point; however, diagnostic experiments may be carried out by selecting any one of the three wells as the locking point.

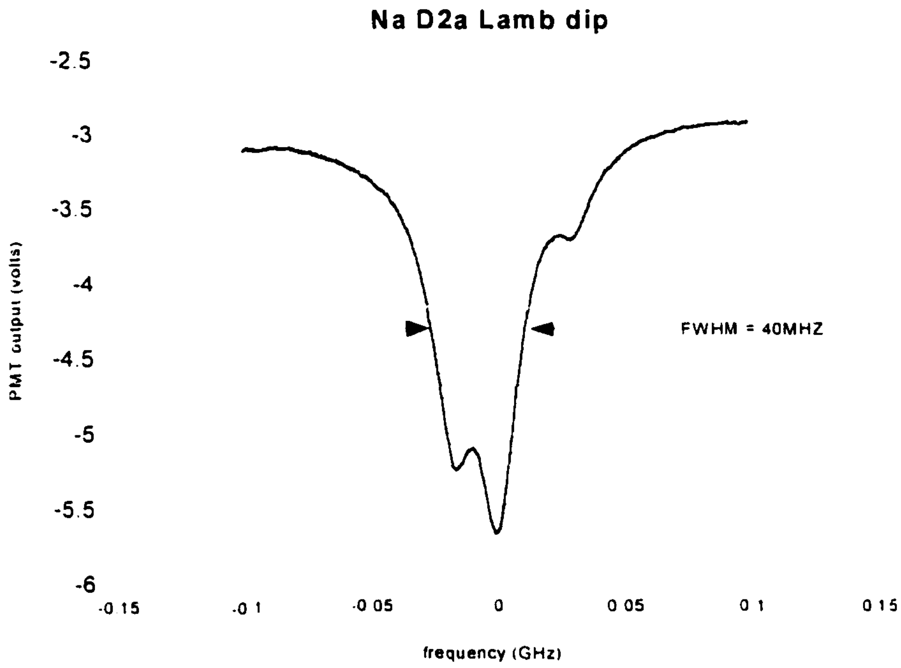

Fig.4.2 The Doppler-free, sodium D_{2a} Lamb dip used for active frequency locking of the CW part of the lidar transmitter. The deepest dip is selected for locking in routine data collection, while the smaller troughs may be selected for diagnostic purposes.

The principle by which the Lamb dip is formed is qualitatively simple. In the case that the laser is detuned from the Na D_{2a} resonance, each of the two counterpropegating beams interacts with a different velocity group of atoms in the cell. If, for example, the laser is tuned slightly below resonance, each beam is resonant with atoms moving toward the incoming beam due to Doppler shifting, and since the atoms are not shared by both beams the effect of saturation is negligible. If, however, the beam is tuned precisely to resonance, both beams interact with the same group of stationary atoms, which therefore see a substantially increased intensity and tend to be pumped out of the excitation circuit. A reduction in measured fluorescence arising form this saturation is therefore observed as the Lamb dip at resonance, and since only the stationary velocity group is saturated, the Lamb dip is unaffected by the thermally driven velocity distribution, and so, is free of the Doppler broadening responsible for the much broadened width (~ 1.4 GHz) of the Na fluorescence spectrum.

The Active Locking Apparatus

In the active locking portion of the CW sub-system of the lidar transmitter, 1 to 2mW of CW light is split from the principle seed beam before the introduction of additional operating frequencies by the A-O modulators. This light is then split evenly between the two counterpropegating beams necessary to observe the Lamb dip. The counterpropegating beams are co-aligned through a 8cm Na vapor cell, temperature stabilized at approximately 70°C. Set orthogonally to the cell and beams is a high-power focusing lens and the cathode of a PMT (Thorn EMI) used for detecting the Doppler free spectrum. The PMT is used in current mode with an external, high-voltage supply. The output of the PMT, suitably impedance matched, is sent to an analog-to-digital converter for observation of the Doppler-free spectrum on a personal computer.

The intention of the active locking circuitry is to provide constant feedback to the CW light source, in this case a tunable dye laser, to maintain a frequency corresponding precisely to the bottom of the lowest well of the Lamb dip. First, it is necessary to tune the laser to the vicinity of the D_{2a} transition using manual, course tuning and a wavemeter (Burleigh WA-10) for frequency bearings. Once the laser is manually tuned to the vicinity of the Lamb dip (about 589.158 nm) the externally driven mechanical scanning of the dye laser may be used to tune to the Lamb dip itself. Mechanical tuning, however, is limited by speed and hysteresis, and unsuitable for the rapid feedback necessary to maintain an absolute frequency stability of a few MHz. Instead, modifications have been made which allow for the introduction of a fast-tuning voltage into the etalon locking system of the dye laser. Tuning from this input has a limited tuning range of several hundred MHz, but the laser frequency was observed to follow a sinusoidal input voltage with a frequency well in excess of 1 Khz, far above the rate at which mechanical tuning would be reliable.

The technique for actively locking the dye laser frequency to the Lamb dip consists of modulating the laser frequency, by use of the fast tuning input, with a modulation amplitude less than 1MHz at a modulation frequency of about 300Hz. This is easily achieved by delivering a 300Hz sinusoidal AC voltage to the fast tuning input. This AC voltage and the output of the PMT monitoring the Doppler free Na spectrum are then routed to a simple lock-in amplifier (Thorlabs LIA-100) as reference and input, respectively. The output of the lock-in amplifier, derived from the Na Doppler-free spectrum, exhibits a low-noise zero-crossing at the wells of the Lamb dip. The noise-reduction characteristics of the lock-in amplifier allows for the use of a small-amplitude modulation or "dither" signal.

The presence of a low-noise zero-crossing for use as a locking point simplifies the generation of a feedback signal to the dye laser, since a signal proportional to a non-zero lock-in output, if applied to the fast tuning input, has the effect of tuning the laser towards the zero-crossing at the well of the Lamb dip. Therefore, given that the laser has been mechanically tuned to the vicinity of the Lamb dip, a DC voltage applied to the fast tuning (also called the dither) input, in addition to the AC oscillation necessary for operation of the lock-in amplifier, can tune the dye laser to the well if the DC signal is a simple integration of the lock-in output signal.

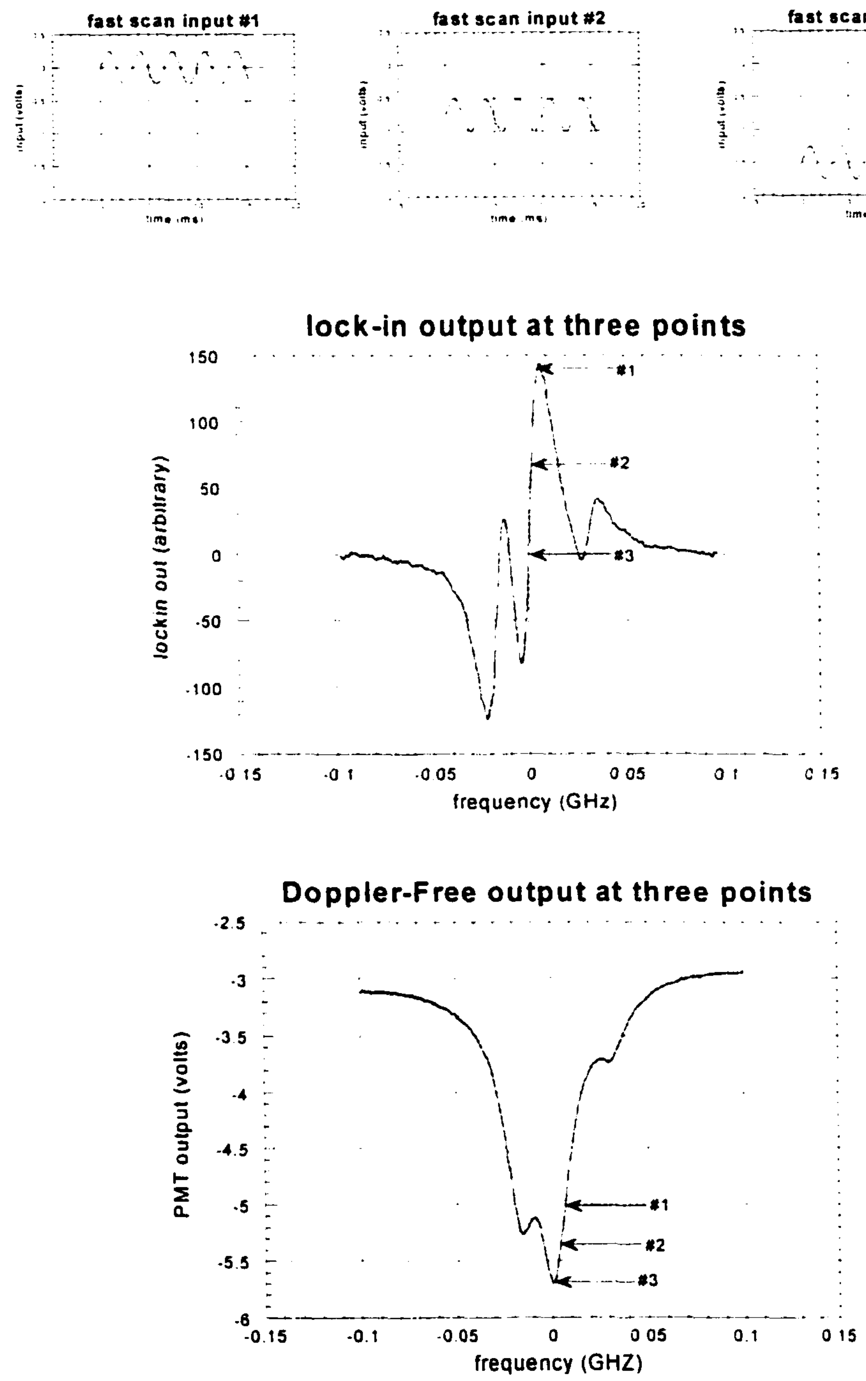

Fig.4.3 Schematic of the active-locking process. The sinusoidal signal at top is applied to the fast tuning mechanism of the ring dye laser, while the baseline offset responds to the integrated low-noise lock-in output, which in turn drives the laser towards the peak.

In normal operation of the active locking, then, a zero-baseline AC signal initiates the locking sequence, which persists superimposed upon a DC signal from an integrator which tunes the laser to the well. One would like the amplitude of the AC sinusoid as small as possible since it determines the effective bandwidth of the CW subsystem, and the speed of the integrator response as fast as possible to move quickly to the D_{2a} dip and adjust to perturbations. In practice, a balance exists between these options since the smaller the dither amplitude, the greater averaging time required by the lock-in to achieve high signal-to-noise in determining the location of the well from the Na Doppler-free signal.

Locking Accuracy

An absolute frequency accuracy of several megahertz is difficult to assess quantitatively since the typical temperature drift of the interferometers available in the lab is prohibitively large. Similarly, the Doppler-free signal which exhibits the spectral features to which the laser is locked includes some intrinsic variation which is difficult to distinguish from variation due to true frequency deviation from the Doppler-free dip. From the perspective of typical operation, the output of the Doppler-free monitoring PMT is compared throughout the one minute lidar signal integration time to the most recent measurement of the PMT voltage corresponding to the well of the D_{2a} Lamb dip. In a typical minute after locking the laser, a maximum variation of about .1 volt about the Lamb dip well voltage may be observed, compared to the total Lamb dip depth of about 3 volts. In a worst-case analysis, in which all of the observed variation is considered due to frequency movement of the laser, and none due to detection noise, this corresponds to about a 5MHz bandwidth around the D_{2a} central well. In terms of averaged lock accuracy, for each integration minute of the lidar, a variance in voltage from the D2a well voltage given by:

$$\sigma^2 = \frac{\sum_{n=1}^{N} [V_n - V_L]^2}{N} \qquad \text{Eq.4.1}$$

where

V = PMT voltage

V_L = last measured well voltage

N = number of shots in the integration time (typically 1200)

is computed and recorded. A typical value of the variance or "lock accuracy" for a minute after locking the laser is $10^{-3}V^2$ suggesting good adherence to the central D_{2a} Doppler-free well throughout the lidar integration time.

4c. The Tandem Acousto-Optical Modulator System

Purpose

Once the CW dye laser is reliably locked in frequency to the central operating frequency corresponding to the Na Doppler-free D_{2a} Lamb dip, the CW sub-system must generate the remaining two operating frequencies necessary for the lidar temperature/wind-velocity measurement, as outlined in chapter 2. To do this, a unique, dual acousto-optic modulator device has been developed by which the CW seed laser at 589nm is modulated to alternate through a cycle of unshifted, up-shifted (+630MHz), and down-shifted (-630 MHz) frequencies. Within each AO crystal the CW laser interacts with a traveling compression wave, the resulting light/soundwave interaction converting both the frequency and direction of the laser beam. The optics of the tandem AO system selects the desired beam, blocks unconverted light which would contaminate the spectral purity of the seed beam, and directs the resulting, frequency-shifted beam along the same beampath which couples the CW subsystem with the pulsed sub-system in the pulsed dye amplifier. Since the Tandem AO system selects the current operating frequency of the CW subsystem, it must also trigger the pulsed subsystem for synchronized interaction in the pulsed dye amplifier, and regulate the detection electronics for correct sorting and integrating of time-of-flight data from pulses seeded by the same operating frequency. Therefore, the tandem AO system serves as the timing center for the entire 3-frequency lidar, as well as the source of its frequency-agility.

The Operation of Acousto-Optic Modulators

An acousto-optic crystal is typically a piece of transparent material chosen for its propensity to change refractive index with applied compression or expansion and fitted on one side with a small acoustic transducer. The acoustic transducer, when receiving a sinusoidal driving voltage, applies a strain to the crystal which creates a traveling acoustic wave which traverses the crystal. Since a change in strain changes the index of refraction of the crystal, a traveling wave of sinusoidally varying index is initiated in the AO crystal.

Given a laser beam traversing a sinusoidally varying index of refraction, Bragg scattering may be observed when the angle of the beam's travel to the variation's phase front is set to an angle θ such that:

$$\sin\theta_B = \frac{\lambda}{2\Lambda} \qquad \text{Eq.4.2}$$

where λ signifies the laser light wavelength and Λ signifies the wavelength of the index variation. At this angle the incremental reflections from the like-index planes separated by the variation's wavelength have a phase shift of 2π so that they interfere constructively. The result is an angular dependence of the outgoing laser beam strongly peaked at θ as long as the index variation persists.

Given a sinusoidal index variation traveling through the crystal at its sound velocity, another phenomenon in addition to stationary Bragg scattering is observed. A frequency shift, which can be thought of as arising from a Doppler shift, whereby the laser frequency becomes slightly higher viewed from the frame of the moving acoustic wave than it did from the frame of the laser, is introduced into the Bragg scattered beam with a magnitude equal to the frequency of the acoustic wave which scattered it. Therefore, when the acoustic transducer is switched on, the beam which exhibits an intensity maximum at the Bragg angle theta is not only redirected angularly, but also exits the AO crystal with a frequency greater or less than that with which it entered by exactly the frequency of the sinusoidal voltage applied to the transducer. The selection of 'upshifting' and 'downshifting' is given by a positive or negative theta, respectively, and may be visualized in terms of a Doppler shift as either head-on phase fronts of the light and sound for an upshift, or pursuant phase fronts for a downshift.

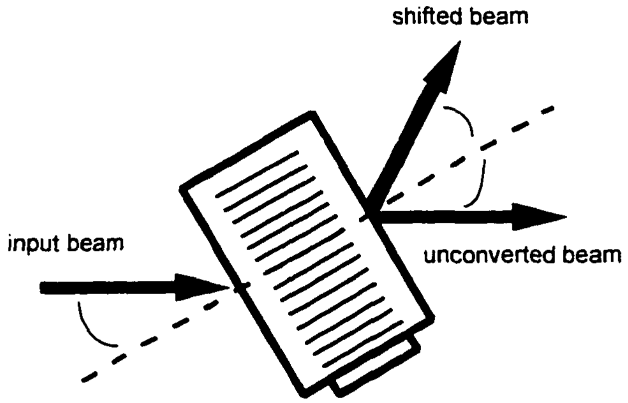

Fig.4.4 The beam path through an acousto-optical frequency shifter. For constructive Bragg scattering, all of the marked angles are the same. The shown configuration results in an up-shift equal to the acoustic frequency, while selection of a negative angle on the input would result in a down-shift.

In general, the intensity of the Bragg scattered light will be greater with a higher amplitude acoustic wave, and therefore a greater index gradient, inside the crystal. The amount of frequency shift, however, remains fixed by the transducer frequency regardless of acoustic amplitude or input beam polarization. The efficiency with which input light is converted to Bragg scattered light is increased proportionally to acoustic intensity, but some unconverted light, which follows a path as though the crystal had no acoustic wave present, is inevitable. Therefore, for a laser beam incident at the Bragg scattering angle upon an AO crystal with an active acoustic wave present, two principle exit beams are observed, the 'unconverted' beam following the beampath of the incident laser and at the same frequency, and the 'converted' beam separated angularly from the other by 2θ and with a frequency difference equal to the frequency of the acoustic wave.

Optical Layout of the Tandem AO System

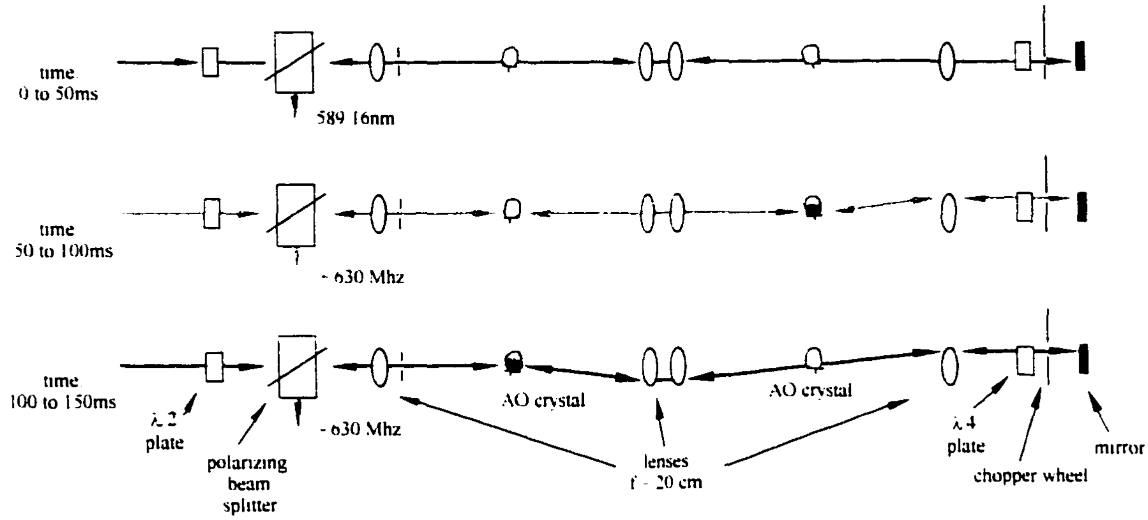

Fig.4.5 A schematic of the tandem acousto-optical frequency shifter used to alternately generate the up and down-shifted operating frequencies of the lidar. Selective optical chopping near the back mirror determines the current output frequency and switches microwave power to the crystals.

The pulsed dye amplifier, in which the CW seed beam interacts with a pulsed pump beam, requires the seed light to be focused to slightly overfill a 50 micron aperture located roughly one meter from the tandem AO modulator system. This requirement imposes a strict angular intolerance upon the modulators, which, as was discussed above, operate in both an angular and frequency shifting capacity. The optical system must therefore gather all of the beams for the lidar's three operating frequencies into the same exit beampath, to within a few microradians. This is accomplished with a periodic system of lenses and a double-pass configuration of the AO crystals.

Double-pass shifting capitalizes on the geometry of Bragg scattering which implies that if a beam scattered 2θ from its entrance direction is retroreflected to retrace its path, it can be shifted again, back into its original path, but in the opposite direction. Furthermore, since the Bragg angle theta is of the same sign on both traversals of the crystal, the frequency shift associated with the two passes add, rather than cancel. The result is a beam which travels along the same exit path as a simply retroreflected beam, but with a frequency shift equal to twice the acoustic frequency in the crystal.

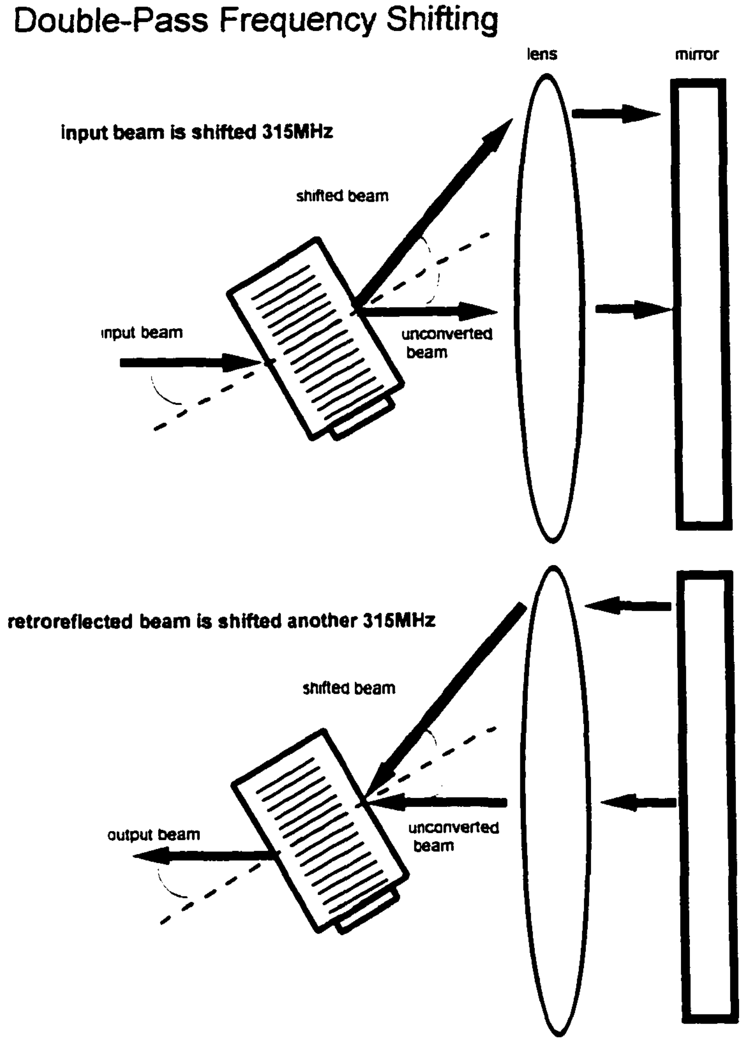

Fig.4.6 The beampaths through an acousto-optic modulator for double-pass shifting. The first pass, at top, shifts the beam in frequency and spatially out of parallel; the second pass, at bottom adds to the frequency shifting and returns the beam to the original beam line. If left unblocked, unconverted line is also returned along the original beam line.

The redirection of the shifted light back into the AO crystal along the same path requires that the active region of the crystal be set at the focal point of a collecting lens which collimates any light coming from the crystal for retroreflection, then redirects the reflected light back into the crystal. Since AO crystals typically require small beam diameters, on the order of tens of microns, the focal-point arrangement of the lenses and crystals serves also to set the waist size of the beam in the interaction region of the crystal. Additionally, good conversion efficiency within the crystal typically requires the beam to be reasonably well collimated through the interaction region, which favors the use of slower (longer focal length) lenses, and may limit the possible miniaturization of such a system.

The 3-frequency lidar's tandem AO modulator system utilizes a periodic system of 20cm focal length lenses to direct the beam through each of the two AO modulator crystals. The crystals operate with an acoustic frequency of 315MHz, half of the frequency necessary to shift the incoming beam to the lidar's operating frequencies of nuD_{2a}

±630MHz. The beam waist inside the crystals, set by the dye laser divergence, an external reducing telescope, and the 20 cm lenses is 30um, with the diameter maintained at within root2 of the waist for about 5mm centered on the crystal. With these specifications, a conversion efficiency of about 0.7 is typical, signifying that 70 percent of the unshifted light would be retroreflected with a frequency equal to one of the ± 630MHz operating frequencies.

The procedure for retroreflecting a sequence of unshifted, upshifted, and downshifted light arises from the AO switching circuitry, which determines which crystal, if either, is receiving a sinusoidal driving voltage at any given time. With both crystals 'off', or receiving no voltage, the input light from the actively locked CW dye laser traverses the lenses and crystals without Bragg scattering; returning along the same path. With only the right hand crystal turned 'on' (receiving a 315 MHz driving voltage) Bragg scattering with a positive frequency shift occurs, resulting in the second, upshifted operating frequency returning, and passing through the left-hand crystal with no further Bragg scattering. Similarly for only the left-hand crystal on, except a negative frequency shift occurs, resulting in the downshifted operating frequency. At no point during normal operation are both crystals turned on at the same time.

As noted before, the conversion efficiency of each crystal can never reach 100%, and so, unless the remaining, unconverted light is somehow removed from the system, each of the shifted output beams would include significant portions of unshifted light as well, jeopardizing the spectral purity of the entire transmitter. Fortunately, there is one plane, just before the back mirror at which the two frequency-shifted beams are spatially separated from any unshifted or unconverted beams. At this point, therefore, a rotating chopper wheel (Stanford Research SRS540) with a custom-machined blade is introduced into the system, the apertures of which allow either shifted or unshifted beams only to continue through the optical path. This chopper wheel, then sets the synchronization for the crystal switching logic, and indeed for the entire lidar transmitter, since the pulsed laser must be triggered in accordance with the seed beam's frequency cycling.

In order to direct the modulated output beam of the AO system towards the pulsed dye amplifier, rather than back into the CW source laser, a polarizing beam splitter is used for turning the output beam orthogonal to the input beam. A half-wave plate at the input of the AO system sets the incoming beam's polarization for maximum transmission through the splitter, and a quarter-wave plate before the final mirror which is necessarily traversed twice, results in an output beam orthogonal to the input beam, and maximum turning at the polarizing beam splitter. In this way a seed beam of each of the three operating frequencies alternately is made to follow the same exit beampath to the pulsed dye amplifier and direct interaction with the pulsed subsystem.

Lidar Synchronization from the AO System

As was noted above, the chopper wheel in the tandem AO modulator system becomes the synchronizing clock source for the entire lidar transmitter. From the chopper wheel's output, a custom electronic logic circuit generates the necessary trigger signals to orchestrate the operation of each crystal's transducer, the pulsed laser trigger, and data sorting in the detection sub-system. The chopper wheel itself generates output pulses based on built-in photogates which monitor the actual rotation rate of the wheel, which is set to be commensurate with the intended 20Hz repetition rate of the pulsed laser. Therefore, every 50ms the following sequence recurs: An aperture rotates into place to allow one of the operating frequencies to traverse the AO system, the corresponding AO transducer is switched on or off depending on the intended operating frequency, and the pulsed laser is triggered to initiate the interaction between the CW seed and pulsed pump beams in the pulsed dye amplifier. After the cycle of all three operating frequencies is complete, a reset pulse is supplied by the chopper wheel logic. This reset is required by the collection electronics to initiate the data collection at the start of a new operating frequency cycle, and to maintain proper sorting of the time-of flight data from each frequency, especially in the event of electrical perturbations to the collection electronics.

4d. The Pulsed Transmitter

Overview

The light with which the 3-frequency lidar probes the atmosphere is made up of short (approximately 5ns duration) pulses with megawatt peak powers. Although the transmitter is considered narrow-band, the short pulse length imparts a theoretical lower limit to the spectral bandwidth of over 100MHz. Given the sensitivity of both the temperature measurement and especially the velocity measurement to laser frequency, as well as the accuracy with which the CW seed beam is locked to the specified operating frequencies, it becomes necessary to fully characterize the spectrum of the pulsed output beam in order to correctly deduce the meaning of the optical signal which it produces. Two seeding processes take place in the pulsed transmitter. The first is the injection seeding of the pulsed YAG laser by a narrow-band CW YAG laser, to insure smooth, repeatable pump pulses of well defined frequency (about 75MHz bandwidth) for optical pumping. The second is the seeding of the pulsed dye amplifier by the CW lidar sub-system. The effect of each of these processes on the signal collected by the lidar must be determined in order to correctly analyze the resulting data.

The Seeded, Pulsed Nd:YAG Laser

The heart of the pulsed transmitter is a Spectra Physics DCR-3D pulsed Nd:YAG laser, which is doubled to output visible green light at 532nm. The repetition rate is 20Hz with an averaged power of around 6 watts and a pulse length of about 7ns. The principal purpose of the pulsed YAG laser is to pump the seeded dye cells in the pulsed dye amplifier. Since it is important that the pump beam be as spectrally compact as possible, in order to avoid unwanted intensity modulation (mode-beating) in the excitation of the dye medium, it is advantageous to injection-seed the pulsed YAG laser with a narrow-band CW YAG laser, so that the pulsed YAG will lase in a single-mode. For this purpose, a Lightwave model 142 YAG laser was fiber coupled and injected into the lasing cavity of the pulsed laser, using feedback circuitry in the DCR-3 to actively adjust the cavity length to match the seed frequency. Monitoring of the effectiveness of the injection seeding to improve the bandwidth of the pulsed output takes two forms. First, the active electronics which adjust the pulsed cavity length do so by monitoring the time between the Q-switch initiating the build up of resonating excitation of the lasing medium, and the output pulse. In the case of effective seeding, gain competition among cavity modes within the lasing medium should be greatly in favor of the mode which closely matches the seeding frequency, leading to faster gain buildup, and less delay before the output pulse. A voltage proportional to this lag time is available from the feedback circuitry, and may be monitored during routine lidar operation. Second, a small amount of unused green light from the pulsed laser has been coupled into an optical fiber, the output of which leads to a thin etalon. In the case of effective seeding, bright, stationary rings may be observed from the etalon output, indicating a single-mode operation in contrast to hazy or jittering rings in the case of imperfect seeding. The effect of poor seeding on the operation of the pulsed YAG laser is the lasing of several axial modes in the resonance cavity simultaneously, introducing additional bandwidth in the green output, which translates into a broadband pedestal (nearly 1 GHz) in the pulsed dye amplifier output.

The Pulsed Dye Amplifier

The last stage of the lidar transmitter before the beam is sent into the atmosphere is the pulsed dye amplifier (PDA), in which the CW sub-system's tunability meets the pulsed subsystem's high peak power. The PDA's 20Hz pulse repetition rate is derived from the YAG pump beam, along with its short pulse length, around 5ns FWHM. The output frequency is broadly tunable with the frequency of the CW seed beam, and has a typical pulse power of 30mJ or averaged power of around 600mW. Under normal conditions of good seeding of the YAG laser, the output pulse of the pulsed dye amplifier is seen to be smooth, with a nearly Fourier-limited bandwidth around 130MHz FWHM.

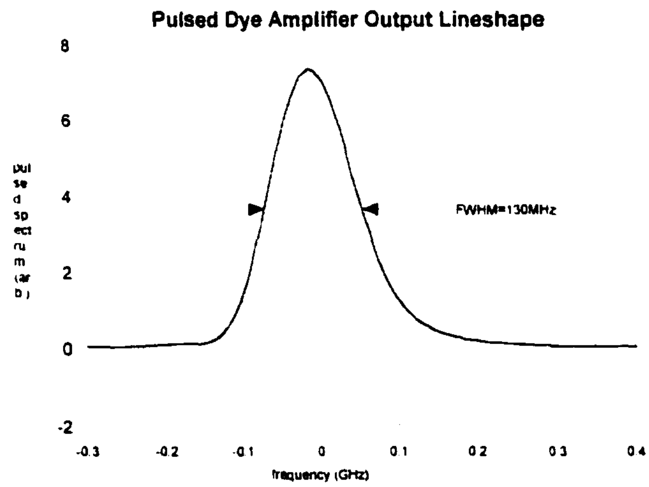

Fig.4.7 The averaged pulsed dye amplifier output lineshape, as measured by a 750Mhz, finesse 70 confocal interferometer. Some asymmetry is apparent in the lineshape, suggesting the presence of frequency modulation.

The pulsed lineshape of the pulsed dye amplifier has been measured using a temperature-stabilized, 750MHz free spectral range, finesse 70, confocal interferometer (Burleigh CFT100). The result is a slightly asymmetrical spectrum, with a centroid about 13MHz higher than the peak frequency. Similarly, the intensity spectrum of the pulse as a function of time can be measured using a fast diode and oscilloscope. From these spectra the root-mean-square width of each may be calculated and the product compared to the minimum allowable for a Fourier limited pulse:

$$\sigma_T^2 = \frac{\int [T - \bar{T}]^2 |f(T)|^2 dT}{\int |f(T)|^2 dT}$$

$$\sigma_\nu^2 = \frac{\int [\nu - \bar{\nu}]^2 |f(\nu)|^2 d\nu}{\int |f(\nu)|^2 d\nu} \qquad \text{Eq.4.3}$$

$$\sigma_T \sigma_\nu = \frac{1}{4\pi} \approx 0.08$$

The same product calculated for the PDA output pulse yields about .1 rather than 0.08, confirming that the output is not quite transform-limited. This suggests that some frequency modulation may be present in the pulsed output, a likelihood which will be treated again in following sections of this thesis. For most purposes of data analysis, the pulsed frequency lineshape spectrum measurement is sufficient as a transmitter characterization, though the location of the spectrum's centroid in absolute frequency is an essential addition for accurate velocity measurements, as will be discussed below.

CW Seeding of the Pulsed Dye Amplifier

Since the purpose of the CW lidar sub-system, including shifting to the three operating frequencies of the lidar, is to set the frequency center of the output pulses of the pulsed dye amplifier, the effectiveness of this seeding process is of critical importance. The typical power with which the CW beam reaches the pulsed dye amplifier ranges from 70 to 120mW, with lower powers belonging to the up and downshifted operating frequencies which suffer from the imperfect conversion in the acousto-optic crystals. The dependence of the pulsed output power on CW input power, fortunately, is less sensitive due to gain saturation beyond 50mW. Above 70 mW cw input, an increase or decrease of CW power of 25%, for example, is unlikely to change the pulsed output power by more than 10%, leading to a smaller variation in output powers among the operating frequencies than the acousto-optic conversion efficiencies might suggest.

In an imagined, idealized pulsed dye amplifier the output pulse would be a transform-limited Gaussian output centered precisely on the input frequency of the CW seed. Three principal factors threaten such utopian fidelity in the actual lidar transmitter: Amplified spontaneous emission, imperfect YAG laser seeding, and unknown nonlinear lineshape effects. The first of these, amplified spontaneous emission (ASE), is present in the pulsed dye amplifier even in the complete absence of the CW seed beam. It is the result of spontaneous emission, which is the natural result of YAG pump excitation of laser dye, which by chance moves along the same beampath as would the CW seed. The result is that the spontaneous emission, which has a broad spectral bandwidth, is amplified by successive YAG pumped dye cells in just the same way the seed beam would be when it is present. The effect, therefore, on the seeded output of the pulsed dye amplifier is to add a broad band platform under the usual pulsed lineshape, increasing the overall baseline. With careful alignment of the pulsed dye amplifier to eliminate reflections along the seed path, ASE is typically reduced to a few percent of the output power. Imperfect YAG laser seeding, which has been discussed above, is more problematic. Unlike ASE, poor seeding of the YAG pump source produces a strong modulation of the PDA ouput with a bandwidth of about 1GHz, leading to incorrect temperature and wind measurements. Fortunately, as mentioned above, YAG seeding may be monitored through Q-switch build-up time and thin etalon fringes, and typical, well-seeded performance of the PDA can be assured. The final, catch-all category, unknown nonlinear lineshape effects, revisits the statements above that the lidar output pulse is, in fact, neither gaussian nor truly transform-limited. Therefore, the somewhat arbitrary assumption that the pulsed output has its centroid, rather than its peak, say, precisely aligned with the CW seed frequency is not necessarily realized in practice. A number of non-linear effects which might endanger the fidelity of the lidar output to the seed input, such as frequency chirping, have been documented in connection with pulsed dye systems, and without a pulsed lineshape measurement accurate in absolute frequency such effects cannot be entirely discounted. This issue will be of principal interest in Chapter 6 of this thesis.

4e. The Detector System

The Fiber Coupled Telescope

The receiving telescope for the 3-frequency lidar is a commercial 14 inch diameter Schmidt-Cassegrain (Celestron 14) with a focal length of 4 meters. The telescope is located in a separate room, about 5 meters from the lidar transmitter, beneath a retractable roof which provides it a clear field of view to the sky over a wide range of angles, if desired. In the case of routine, vertical observation, the telescope is leveled using a bubble balance to point to zenith; a procedure accurate to probably better than 1 degree in any direction. Close to the focal plane of the telescope a broadband interference filter with a bandwidth of 10nm centered at 589nm, and a peak transmission of about 70 percent is mounted. This filter is used to reject off-resonance light from city lights, stars, and the moon which would add unwanted noise into the laser-induced signal. After the interference filter is a 5cm focal length lens for focusing the return light into an optical fiber. The fiber is fitted with a threaded connector which screws into a 3-dimensionally translating mount. The mount can be adjusted to maximize signal, which is usually accomplished with the fiber face approximately 4cm from the focusing lens. Since the lidar information does not depend on imaging the telescope return, but rather on the total amount of signal and the time elapsed from the laser's firing to the signal's arrival at the telescope, a single fiber can be used for light collection. Since the current 3-frequency lidar has been designed for nighttime operation, the optical fiber was chosen purely for coupling efficiency, rather than field of view reduction, as might be required for a daytime system in order to reduce bright-sky background. Therefore, a 1mm fiber was chosen with a numerical aperture, or sine of the maximum collection half-angle, of 0.37. This choice of fiber provides significant tolerance of the coupling efficiency to focal length or lateral alignment drift, situations which must be anticipated given the extremes of year-round temperature at which the open-air telescope system must operate.

The Photon Counting Electronics

Since the optical signal which is collected by the telescope each time the laser fires is very small, photon counting must be used to distinguish between signal strengths at different altitudes and operating frequencies. The 1mm optical fiber which collects the returning photons from the telescope returns them to the main lidar facility in which the photon counting electronics are maintained in an environment with less temperature variability. The fiber is terminated with a focusing lens and screwed into another 3-dimensional positioner in order to ensure good optical coupling between the fiber output and the photomultiplier tube (PMT) cathode detector. The PMT (Hamamatsu R934-02) is cooled inside a thermally controlled housing (Products For Research "Photocool") to reduce thermal background noise, and has a quantum efficiency of 15% at 589nm. The PMT is operated in photon counting mode, which means that a detected photon will result in a narrow current pulse which may be amplified, conditioned, and counted. The lidar's PMT output pulses are amplified using a preamplifier (Phillips Scientific 6954B-100) which linearly amplifies signals of up to 1.8GHz frequency with a gain of 100. The amplified pulses are then conditioned with a discriminator (Phillips Scientific 6904) which reproduces those pulses with heights in excess of some discrimination level (usually about –60mV), in order to reject low-level noise which might otherwise be counted as photons. The discriminator outputs a train of identical, low-noise pulses, converting input pulses arriving at up to 300MHz with negligible delay.

The Counter and Synchronizing Electronics

The electrical pulses corresponding to photons incident upon the PMT cathode are finally sent to the counter, where they are sorted into bins according to their inferred time of flight, and summed with other laser pulses of the same operating frequency throughout a preset integration time. The 3-frequency lidar uses a PC card style counter board (Optech fdc700m) which receives the conditioned PMT output pulses as well as a trigger signal from the pulsed YAG Q-switch, from which all times of flight are referenced. The counter uses two fast first-in-first-out memory buffers to temporarily store the data corresponding to the most recent triggers, adding the contents of one to the proper operating frequency array while the other buffer is receiving data associated with the current shot. In this way 1200 shots, corresponding to 400 shots at each operating frequency transmitted at 20Hz for one minute, are collected into separate arrays for each operating frequency signifying the total photon counts at each time-of-flight altitude bin for that minute. At the end of the one minute integration time, collection is momentarily halted while the arrays, along with archiving heading information such as date and time, are saved to the PC's hard disk. The saving sequence requires only a few seconds before data acquisition begins again for the next minute.

The use of only one counterboard channel has the advantage of cost effectiveness and interfacing simplicity. However, the use of three counterboards, one for each operating frequency, incorporates additional flexibility in that the counter, and therefore the array, for each operating frequency, can be more explicitly triggered by the electronics which determine the operating frequency of the transmitter. In a one channel measurement the dependability of the frequency cycling is the only assurance that the operating frequency of each laser pulse is being correctly interpreted and integrated. If, for instance, a pulse of electrical noise, perhaps from an air conditioner turning on, mistriggers the counter board, given no further diagnostics, the next true laser pulse which triggers the counter will be interpreted incorrectly and for the remainder of the minute after the perturbation the integration will mix data from different operating frequencies, invalidating the entire integration period. It is necessary, therefore, to have additional synchronizing electronics which monitor the cycling of the laser frequency and insure that the proper order is being maintained before the counter is triggered.

It is for this purpose that the AO switching electronics include a reset output, which signals the beginning of a new sequence of the three operating frequencies. In the detector system additional electronics between the YAG laser Q-switch output and the counter trigger input have been developed in order to insure proper cycling using this reset signal. The first stage in these electronics checks that exactly three counter triggering pulses have occurred between successive reset pulses. If more than three pulses arrive before the reset pulse, all additional pulses after the first three are blocked, and no further counter triggering occurs until a new reset pulse arrives. If fewer than three pulses arrive before the next reset pulse, the first pulses after the reset are used to close the short cycle, and any further pulses are blocked until counter triggering commences again after the following good reset pulse. The effect of this circuit is that regardless of whether a perturbation to the system adds or looses counter triggering pulses, only a single cycle of the three operating frequencies will be affected by a single perturbation, leading to erroneous data in only one fourhundredth of the entire integrating period rather than the entire subsequent integrating period. The second stage in the cycling diagnostics electronics is an independent counter which displays the total number of triggers sent to the counter. At the end of the integrating period of 1200 laser pulses, the electronics' counter can be compared to the incrementing display on the photon counter board interface to insure against spurious triggers from within the counter board PC. With these safeguards in place, the one channel counter synchronization gains robustness comparable to the three counter design, while retaining its advantages in cost and simplicity.

Radar Monitoring for Air Safety

Since the pulsed laser output of the 3-frequency lidar is not considered eye-safe until its divergence has expanded the beam manifold, the possibility of danger to air traffic in the area requires monitoring for possible aircraft in the beam path. In order to allow single-operator lidar operation, a nautical radar (Furuno1831) designed for alerting against shipborne collisions has been customized to survey the beam area in the sky. The radar unit was intended to rotate in order to sweep out a circular scanning area about a ship, and return distance and size information to a CRT monitor sweep display. A built in system of menus allow the operator to set an alarm region which would activate an internal relay meant to sound a ship's horn if an obstacle was detected within selected inner and outer range boundaries.

Since the divergence of the radar beam is far greater than that of the laser, it was decided to disable the rotation of the transmitter, and to fix it pointing vertically. Mounted to the roof of the lidar facility, the radar monitors only a rectangular portion of the sky containing the lidar beam, reducing the circular sweep display on the radar monitor unit to a altitude-only measurement. All other functions of the radar unit were left intact, so that the guard alarm function, intended to monitor a circular region of landscape, could be set instead to guard a region of altitude. Since the small angular field of view of the radar would leave little time to manually investigate a radar detection if a low, fast aircraft entered the guarded region, a small external circuit was built using the radar unit's internal horn relay which switches the YAG Q-switch off whenever the guard zone is breached. During typical lidar operation the guard zone is set from near zero to about 20km, above which eye-safety is not an issue. The system was tested with radio controlled airplanes for fast shutoff time, and initially monitored with actual air traffic for range and reliability. With active radar monitoring, single-operator lidar acquisition may be conducted in accordance with FAA guidelines for air safety.

Chapter 5: Analysis and the 3-Frequency Lidar Data

The true test of the hardware design of the 3-frequency lidar, of course, is in the atmospheric data acquired by it. Therefore, in this chapter the data, along with a description of the data analysis procedure applied to the collected photon counts, the uncertainty of the measurements, and a comparison with the 2-frequency temperature lidar will be presented. Since over 80 nights of data have been acquired, a detailed presentation of the entirety of the measurements is not feasible. Instead, a subset of the full data base is selected to illustrate the various features common to all of the data, and to suggest the operational parameters of the lidar.

5a. The Data Analysis Procedure

Since the details of laser-induced fluorescence of atmospheric sodium and the inversion process by which it yields the temperature, radial velocity, and sodium density in the mesopause region have been discussed in Chapters 2 and 3, only a brief outline of the process specific to the analysis of the 3-frequency lidar data is needed here. The programming which converts a night of data into range-resolved profiles of atmospheric variables has been discussed in previous work (She and Yu, 1994; H. Chen 1997) but continues to evolve on a nearly daily basis due to the indefatigable David Krueger and other lidar programmers. The principal analysis code is written in IDL, the interactive data language.

Reading the Saved Photocount Data

The range-gating counter which sums the electrical pulses generated by returning photons in the photo-multiplier tube generates hundreds of one-minute integration files in a typical night of lidar operation. In each file are three columns representing the photo-count sums for each of the lidar's operating frequencies for the minute referenced to each range bin. For the 3-frequency lidar data, the range bins are not of equal size. This allows smaller bins to be used for altitudes in which high range-resolution is desired, such as the mesopause region, while using fewer, large bins to monitor less variable regions of the atmosphere. Since each file represents one minute of hardware-summed lidar data collection, this is the smallest timescale resolvable in the acquired data. Additional averaging of multiple one-minute profiles may be desired in the data analysis, but the raw data is archived in its one-minute form to maintain the maximum range of observable phenomena.

Since the size and arrangement of the range binning applied by the hardware counter is user-defined and somewhat arbitrary, this information must also be stored with the lidar data along with other user-defined experimental parameters such as integration time and the type of lidar used to take the data. Similarly, some additional experimental information is also required by the data analysis software and must be provided with the raw data. This information includes the pulsed laser lineshape. which is measured periodically and used in the analysis of data taken around the same time. and the characterization of the photomultiplier tube gain versus altitude. Since much of this information corresponds to certain observation time periods over which many nights of data were taken. is important that it is saved in standardized formats so the analysis software is able to correctly interpret the received photon counts based on the current experimental parameters.

Once the raw data and all of the necessary information required to perform the data analysis is read into the computer. the data can be scrutinized for errors. summed into larger time intervals. and referenced to the correct altitudes. The inversion procedure for determining temperature. Na density. and radial velocity progresses after the photocount profiles have been summed to the desired time resolution intervals. to provide the best possible counting statistics.

Bad File Selection

Before a photocount profile for a given time interval is converted into atmospheric parameters. a variety of tests are performed on the individual one-minute files in that interval. The purpose of these tests is to verify that the lidar transmitter. receiver. and atmospheric conditions during a given minute conform to some minimum operating standard necessary for their inclusion into the data analysis. In a typical night of data which includes 300 or more one-minute files. there are typically between 2 and 20 files which fail to be included possibly due to laser malfunction resulting in either frequency instability or loss of power. temporary cloudiness which reduces the received signal. or momentarily high background levels perhaps due to electronic interference.

So that these rare instances do not adversely effect the statistical quality of the summed photocounts. the analysis software first recognizes and removes them. Tests which perform this function include trend tests. which make sure the overall photocounts do not vary from one minute to the next by more than some physically realistic value (usually around 20%). ratio tests which calculate approximate layer-averaged temperatures and winds to insure the results are within the range of physical possibility. and background tests which insure that the signal level is measurably higher than the measured sky background. The intention of these tests is not to force an expected result from the atmospheric profiles, but rather to insure that accidentally produced statistical outliers do not systematically bias the time-averages. The test criteria, therefore. are chosen to apply generally to the lidar data, despite seasonal changes in the response of the mesopause region and the presence of highly variable atmospheric phenomena.

Time Interval Averaging of the Raw Photocount Profiles

Once the bad files are selected and discarded, the remaining files in each time interval are averaged. Typically, one-hour time intervals are selected as a compromise between high statistical confidence and the atmospheric phenomena which can still be resolved. Other time intervals may be selected, and can be useful in studies of atmospheric effects which appear at shorter timescales, such as "sporadic sodium layers". In addition to whatever time averaging is selected in the analysis software, an average of the entire night is also computed, which may be used with other nightly averages to distinguish long-scale phenomena, such as seasonal variations. A typical raw data profile is shown in Figure 1.

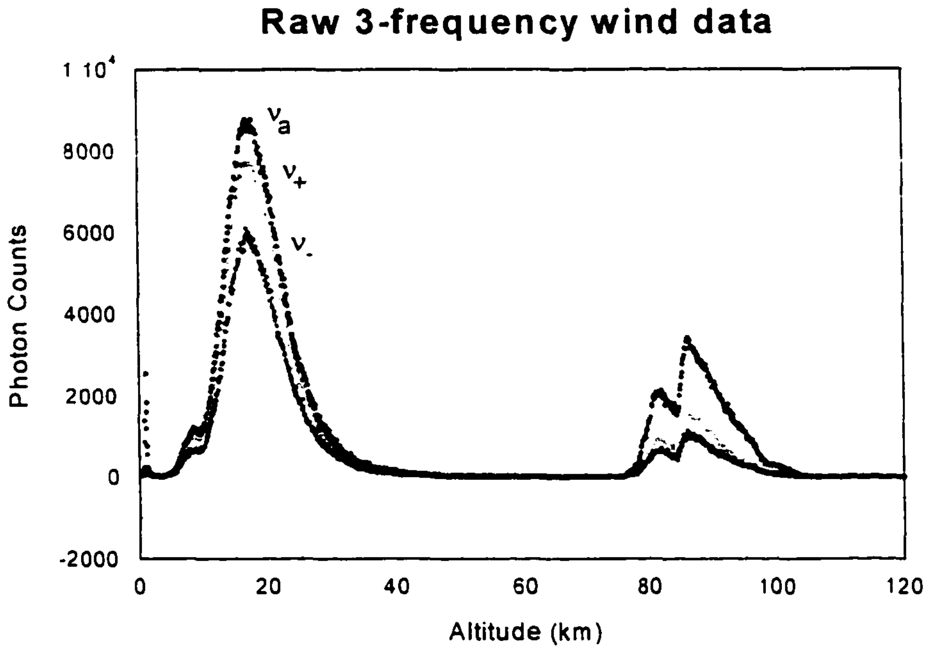

Fig.5.1 A good, 3-frequency, minute-averaged photocount profile. The lower peak arises from Rayleigh and Mie scattering in the lower atmosphere, the higher peak represents fluorescence signal from the sodium layer. The shape of the lower peak is due to electronic gain-switching in the receiving electronics. The separation of the three frequencies in the lower atmosphere is due to power differences arising from imperfect frequency conversion.

Corrections to the Photocount Profile

Once time-averages with sufficiently high counting statistics are formed, two corrections to the raw data must be applied before the data inversion can be executed. The first corrects the range-resolved counts for the time-dependent gain of the photomultiplier detector. Because the returning photon signal from the atmosphere is very intense in the lower atmosphere from 0 to 20km, errors can arise in the detector's sensitivity to the much weaker, higher atmosphere if the photomultiplier's full gain is applied to converting the photons received from low altitudes into an electrical pulse. Therefore, the photomultiplier tube is electronically 'blanked' for a time starting with the laser pulse and persisting the equivalent of about 20km in round-trip time-of-flight. However, since the photomultiplier is a high-voltage device being subjected to rapid switching, some time is required for full settling of the receiver gain, which translates to a range-dependent gain over the photocount profile. To correct for this hardware-induced variation of the received data, then, the range-dependent gain is characterized in a laboratory experiment, and the time-averaged photocount data is normalized to this measured gain profile to yield the true, relative photocounts.

Once the correct relative photocount profiles are obtained, to ready the profiles for conversion to atmospheric parameters, the signal must be corrected for non-laser-induced background light sources. Since background sources such as the moon and dispersed city light are not pulsed synchronously with the lidar transmitter, their effect is the same on all altitudes of the lidar profile. To subtract the proper level of signal from all altitudes corresponding to background photon counts, very high altitude range bins are saved in the lidar data files, in which laser-induced return signal is negligible in comparison to background signal. From these high-altitude bins, the correct background contribution to the variously sized bins in the lidar profiles can be determined and subtracted from the range-resolved photocount profiles.

Rayleigh Normalization

As was described in the discussion of the inversion method in Chapter 3, the Rayleigh signal from the 'clean air' scattering region between about 30 and 45km is used in two capacities. First, the total counts from the region are used as a monitor of the laser power transmitted through the changeable lower atmosphere at each of the lidar's operating frequencies. Used in this way, the specific Rayleigh returns at each altitude are not important but rather a statistically well-determined sum from the entire range best establishes the relative laser power for each frequency. Second, the Rayleigh signal from a single reference altitude is necessary as a standard for determining the sodium density at range bins in the mesopause region. Used in this way, a range-resolved Rayleigh profile is necessary to accurately determine the Rayleigh return from a single altitude at which the air density is referenced to a standard atmosphere.

To obtain an accurate value of the Rayleigh backscattered signal at a single reference altitude, the measured photocounts in the 'clean air' scattering region from 30 to 45km are fit to the formula $[B_i/z^2]\exp[-A_i z]$ corresponding to the theoretical Rayleigh scattering profile for clean air. From this fit the two coefficients A_i and B_i are determined and the signal from a specific altitude in the region may be accurately evaluated. A further application of the resulting Rayleigh backscattering fit is to determine the small Rayleigh contribution to photocounts from the mesopause region and subtract them, leaving only photocounts from fluorescence as required in the lidar equation inversion procedure.

Smoothing and Determination of Atmospheric Parameters

Once the corrected, relative photon counts have been normalized to the Rayleigh sums for the temperature and velocity calculations, and the Rayleigh reference altitude for the density calculation, these ratios are smoothed in altitude with a Hanning window with a full width of 99 bins, for a final vertical resolution of 3.1km (She and Yu, 1994). These smoothed ratios are the measured quantities which are used in the inverted lidar equation to determine the temperature, radial velocity, and sodium density in the mesopause region as outlined in Chapter 3.

5b. Typical Data from the 3-frequency Lidar

To illustrate the information obtained from the 3-frequency lidar, one night which is considered typical of the database of over 80 nights of data acquired by the lidar will be discussed first. In this case, the chosen night is the 51st of 365 in 1998, or the 20th of February, 1998. This particular night of data has been chosen partially due to the existence of a 2-frequency lidar acquisition night measured by the CSU lidar lab on the following night. Though the variability of phenomena in the mesopause region may be substantial from one night to the next, some comparison of the operation of the two lidars is possible.

The 2-frequency lidar has been the subject of several publications (She and Yu, 1994; Bills et. al., 1991) and the design details will not be reproduced here. However, the 2-frequency Na lidar has been the standard system for routine monitoring of the mesopause region at CSU and elsewhere, and the obtained results have been compared with those of other measurement systems favorably (She et. al., 1995). Therefore, comparison of the 2 and 3 frequency lidars, though imperfect due to hardware limitations precluding simultaneous operation of the two systems at CSU, does provide some gauge of the performance of the newer, 3-frequency lidar.

The hour-averaged profiles of temperature, Na density, and vertical velocity for day 51 of 1998 from the 3-frequency lidar are shown in Figure 5.2, together with temperature and Na density measurements for day 52 of 1998 from the 2-frequency lidar (see following page).

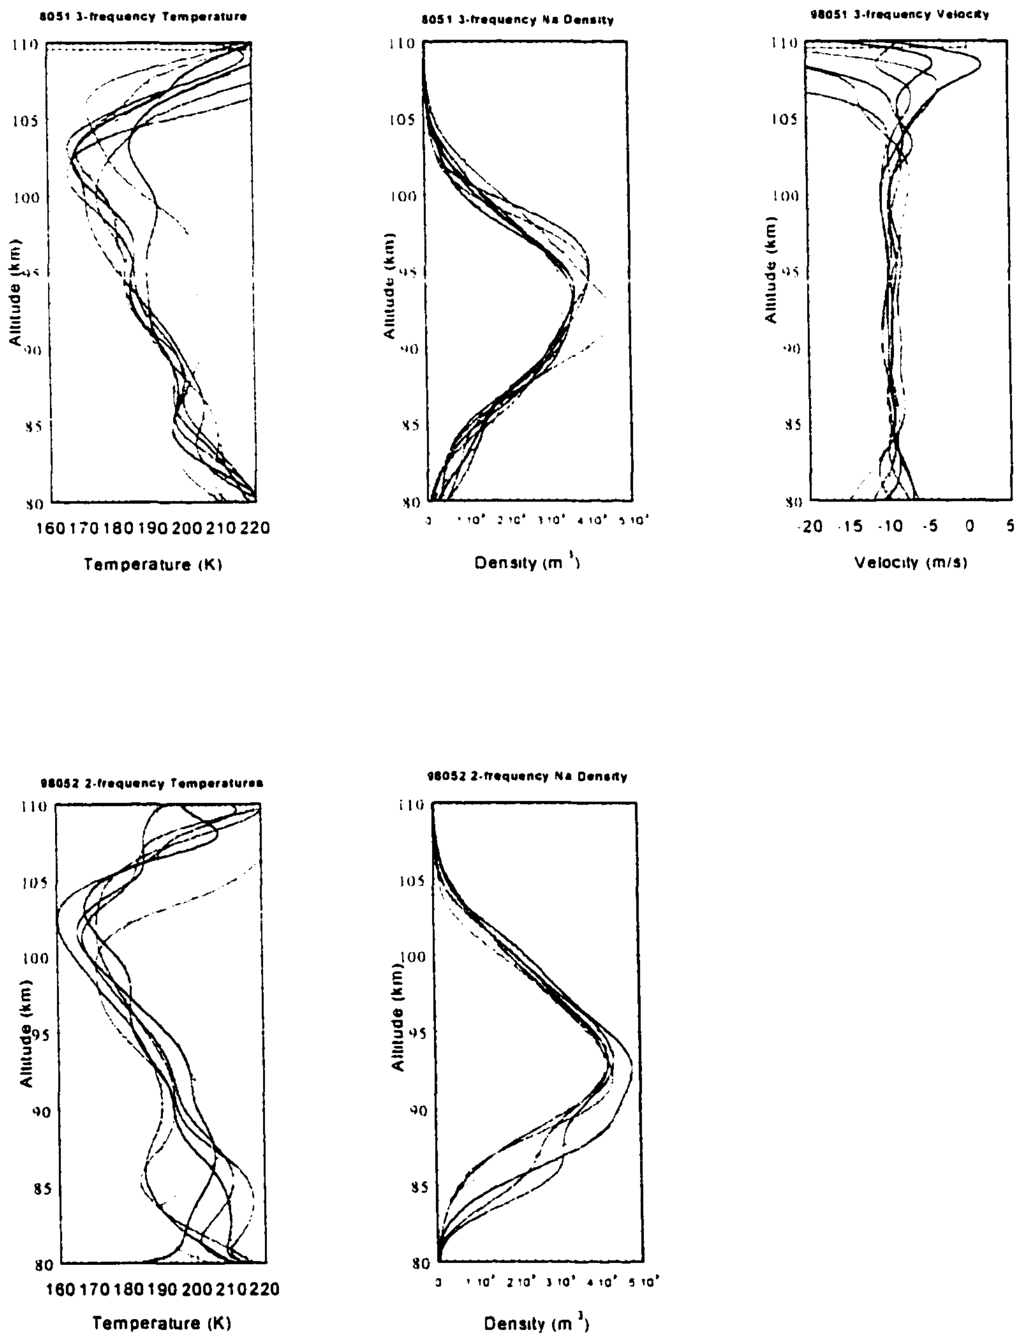

Fig.5.2 Top: Typical hour-averaged profiles of temperature, Na density, and vertical velocity taken by the 3-frequency lidar. 8051 represents the 51st day of 1998. Bottom: Similar profiles acquired the following night (8052) using the 2-frequency lidar.

From the hour-averaged measurements on these consecutive days it is apparent that the temperature structure of the mesopause region remained fairly stable. Both nights exhibit a high-altitude mesopause, the coldest point in the profile, around 101km and a small local minimum around 85km. Seasonal studies of the mesopause thermal structure suggests this to be typical behavior for the winter mesopause above Fort Collins. The all-night averages for each date show an agreement on the altitude of the mesopause within 1km and agreement on the overall temperature of the sodium layer within 2%, both values well below the typical night-to-night variation for measurements conducted with the same lidar system.

The density measurements show slightly poorer agreement in absolute magnitude, likely due to the measurement's required assumption that the air density in the lower atmosphere agrees with a tabulated 'standard atmosphere'. This approximation, however, only effects the overall scaling of the density measurement. So, in the measurements of Figure 5.2 there is an agreement of only 15% between the nightly-averaged total abundance of atomic Na between the two nights, while there is a better than 1km agreement in the altitude of the centroid and RMS thickness of the density profile. The disparity in absolute density and agreement in absolute temperature of the two measurements suggests the overall insensitivity of the temperature measurement on the absolute density, as was asserted in the previous discussion of the role of the density measurement in the temperature and velocity inversion.

To illustrate the night-to-night variation of measurements of the mesopause region and provide additional comparison between the 3-frequency and 2-frequency lidar systems, several additional instances in which the two lidar systems were employed on consecutive, or nearly consecutive dates have been selected for Figure 3 and the all-night averages of temperature and density corresponding to the different systems are plotted together, along with the vertical velocity profiles from the 3-frequency lidar (see following page).

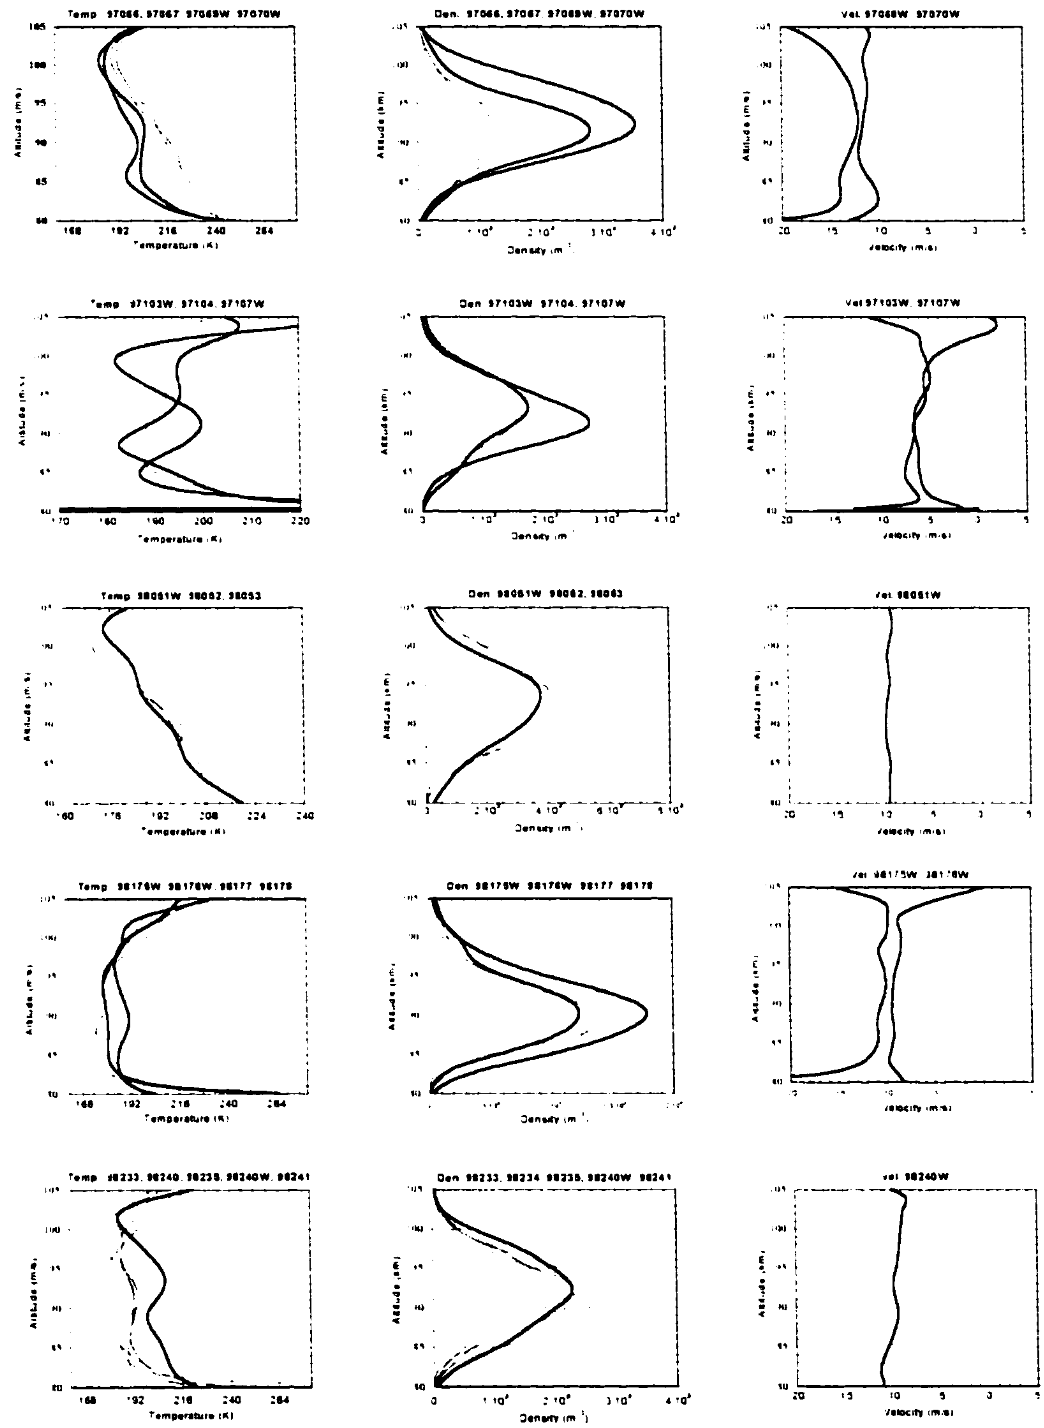

Fig.5.3 Five groups of nightly-averaged temperature, Na density, and vertical velocity measurements in which the 3-frequency lidar (bold profiles) and 2-frequency lidar (fine profiles) were operated in close succession over one and one-half years. Though night-to-night variation in the mesopause parameters can be large, no obvious trend is apparent to suggest that the two measurement techniques are systematically biased from each other. In the labeling convention, 97066 corresponds to the 66[th] day of 1997, etc. Over 80 nights similar to those shown have been acquired by the 3-frequency lidar since 1997, and hundreds using the 2-frequency system since 1993.

A careful look at the many plots in Figure 5.3 reveals that the nightly averages of atmospheric parameters measured with the 3-frequency lidar agree with the measurements of the 2-frequency lidar to within the night-to-night variability of same-system measurement nights. There does not appear to be any observable trend with which measurements of either temperature or Na density systematically differ between the two systems. The profiles seem to track each other with respect to slope and local minima, and in the case of the 1997 nights 103, 104, and 107 the similarity in shape and magnitude of wave activity in the temperature profiles is surprisingly good. Figure 5.3 suggests that simultaneous measurements using the 2-frequency and 3-frequency systems would be necessary to discern any systematic differences in the atmospheric profiles obtained by the two systems. Unfortunately, this is not possible at the CSU lidar facility due to shared hardware between the two systems.

Error Estimation

Since the uncertainty in the lidar profiles depends on the signal level from a given altitude range, the error bars for the measurements of temperature, Na density and radial velocity must be computed on a range-resolved basis. Figure 5.4 shows the calculated error bars for the 3-frequency lidar measurements compared with the 2-frequency measurements for a sample hour and the nightly average, again for 1998 nights 51 and 52 (see following page).

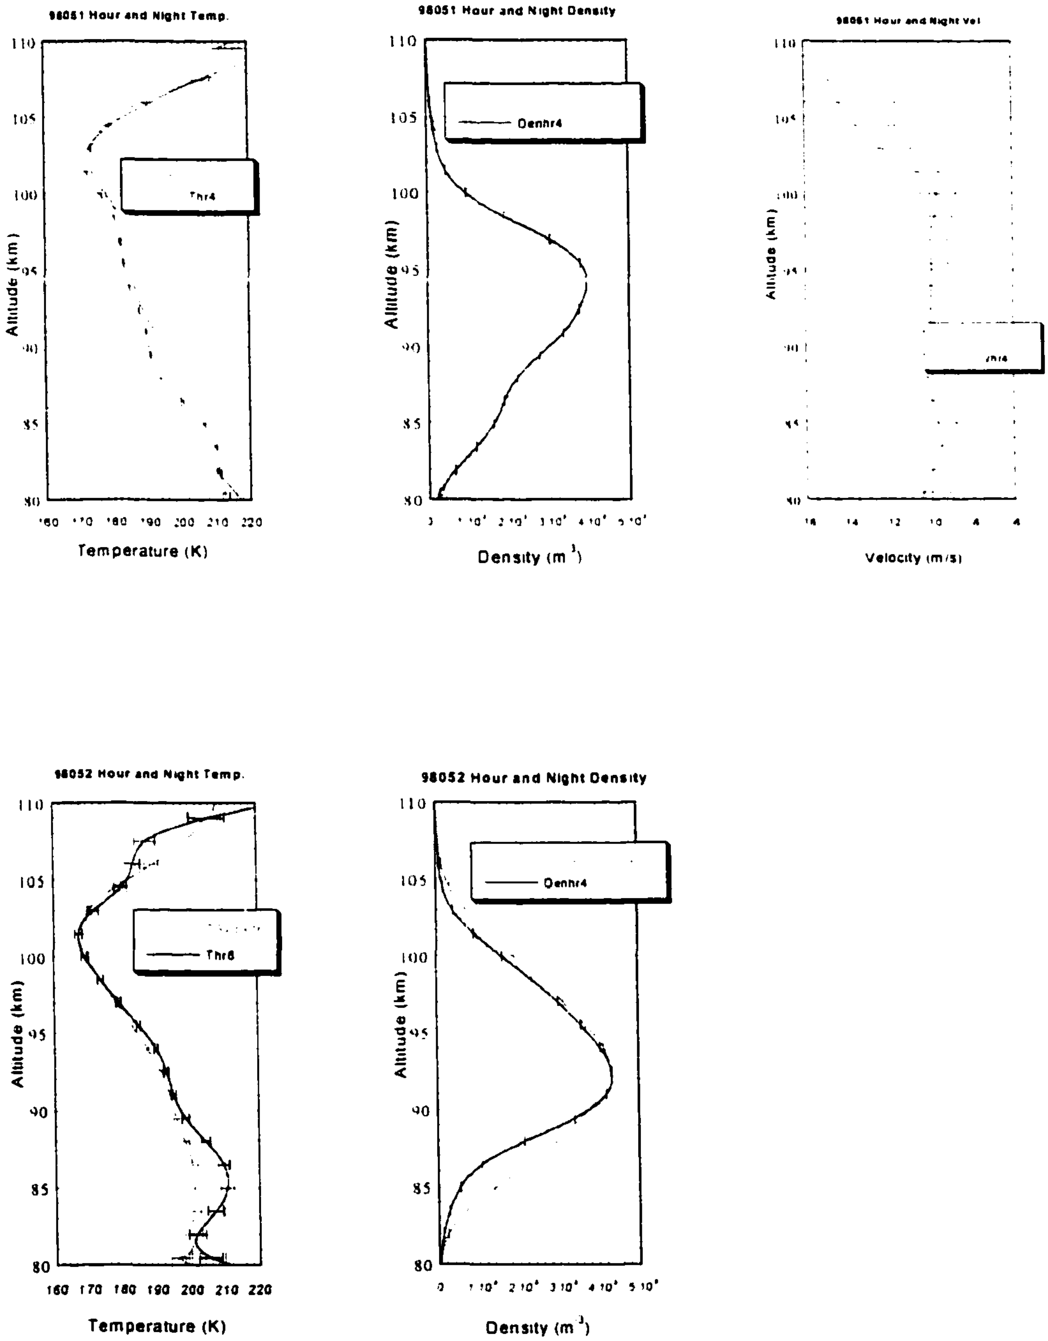

Fig.5.4 Typical profiles from the 3-frequency lidar (top) and 2-frequency lidar (bottom) showing the error bars associated with 1-hour averages (dark profile) and nightly averages (light profile).

The error bars for the 3-frequency lidar data arise from counting noise and laser tuning error, the first of which scales as the square-root of the total photon counts, while the second is not considered randomly distributed, and therefore does not benefit from additional signal or averaging. The error bars of the temperature profile have a minimum of about ±0.3K at the sodium density peak for the one-hour averaged profile and about ±0.17K for the nightly average, which is dominated by laser frequency uncertainty. These error bars increase with the distance from the density peak, having doubled at an altitude approximately 7km higher or 13km lower. The error bars on the density measurement have a minimum of about $\pm 4 \times 10^6$ at the density peak or 0.1%. The velocity measurement, unlike the temperature and density measurements, is dominated by the laser frequency error over a large range of the mesopause region. The minimum error bar size is set by the laser tuning uncertainty at ±0.75m/s which is nearly achieved close to the sodium density peak. The velocity error at the density peak is ±0.8m/s for the hour averages and ±0.76m/s for the nightly average, doubling at about 15km below and 11km above the peak.

Comparison with the error bars calculated for the 2-frequency lidar data in Figure 5.4 shows very comparable accuracies for the temperature and Na density measurements between the two systems. For the two nights shown, the 3-frequency temperature measurement has slightly smaller error bars, but different by only a few tenths of 1K. The density measurements have essentially the same accuracy. These results suggest very similar performance between the 2-frequency and 3-frequency lidars.

5c. The Anomalous Measured Vertical Velocity

One consistent feature of the data presented in Figures 5.2, 3, and 4 above is the negative, or downward, vertical velocity of between 6 and 12m/s, with little vertical structure. Both the literature of vertical velocity measurements in the mesopause region and simple atmospheric theory disagree with the size and constancy of the 3-frequency lidar velocity measurement, which in this sense is termed 'anomalous'.

Unlike the horizontal plane, the vertical structure of the atmosphere is stabilized by gravity, which generates the steady-state, background distribution of atmospheric density and pressure. For this reason, a movement of air from one altitude to another is accompanied by a restoring force which directs the air back to its original altitude. Ascending air, for example, is instantaneously more dense than its surroundings, and therefore must either sink back to its equilibrium altitude or expand to match the density of its environment, requiring work from the moving air, and therefore expending energy. It is reasonable, therefore, that processes which generate vertical wind loose energy rapidly as the air ascends, and vertical winds in the middle and upper atmosphere would be less than horizontal winds which do not encounter such a restoring force.

Similarly, the presence of the sodium layer in the mesopause region, the basis of the Na lidar measurement, also suggests the impossibility of large downward winds in its vicinity. The yearly average influx of sodium in the form of meteors burning in the high atmosphere, estimated from meteor measurements, combined with the sodium density measurements of radars and even the 3-frequency lidar itself, suggests an average lifetime of about 3 days for an atmospheric sodium atom (Clemsha et. al.1992). In the presence of a constant, downward 10m/s wind, however, a sodium atom would traverse the entire 40km sodium layer in about one hour. Therefore, it is clear that the lidar velocity measurement is not commensurate with the observed structure of the atmosphere.

Radar measurements (Vincent and Reid, 1983) and circulation models (Roble and Ridley, 1994) suggest the velocity in the mesopause region to be on the order of centimeters per second. This figure is in reasonable agreement with the observed variation with altitude of the vertical velocity measured by the 3-frequency lidar, but in stark disagreement with the overall magnitude, and lies far outside the error bars which represent the known sources of uncertainty intrinsic to the lidar measurement. Because of the solid physical basis for expecting small vertical winds in the mesopause region, the vertical wind measurement provides a test of the 3-frequency velocity measurement which must be passed before off-vertical radial winds containing components from possibly large horizontal winds can be accurately measured. A more careful analysis of the assumptions made in the calculation of the velocity from the lidar return photocounts must be made to account for the observed anomaly and determine techniques for correcting the velocity measurement. Such an analysis is the subject of the following chapter.

Chapter 6: Analysis of the Anomalous Measured Vertical Velocity

In this chapter, the interpretation of the velocity measurements presented in the previous chapter, along with a discussion of strategies for correcting implied inaccuracies are undertaken. Since the measured velocities are inconsistent with a reasonable understanding of the atmosphere, as was discussed above, the focus in determining the locus of the problem will be on the lidar transmitter and assumptions based on its operation which are made in the data analysis. Since it is the instrument which finally produces the laser pulses which interact with the atmosphere, the Pulsed Dye Amplifier's input and output parameters will be investigated in some detail in a discussion of the fidelity of such injection seeded systems. Since the lidar velocity measurement is based on Doppler shifting, any frequency shift in the lidar transmitter will have a proportional effect on the measured velocity. Once the likelihood of such a shift being present in the lidar transmitter is established, its effects on the lidar temperature, Na density, and velocity measurements will be discussed and preliminary experimental support for the existence of such a shift will be presented. Finally in this chapter, techniques for correcting the measured anomaly will be considered, with an introduction to the spectral-edge frequency monitor applied in the CSU lidar system; the experimental details of which are given in the following chapter.

6a. Fidelity of Seeded Systems

As was discussed in chapter 4, the output of the pulsed dye amplifier falls short of an idealized, Gaussian pulse. The transform product is only about 80% of the transform limited value, a result commonly interpreted to suggest the presence of phase modulation or frequency chirping. Frequency chirp is a name given to the variation of the instantaneous frequency of a laser throughout the lifetime of a single pulse, presumably due to the phenomenon's similarity to a bird's chirp which changes frequency (or tone) rapidly over a short burst.

Pulsed amplifiers have the unfortunate effect of reducing the coherence of the injected seed beam. One unavoidable result of this loss in coherence is the transform-limited bandwidth, which accounts for a symmetric output pulse centered at the seed frequency with a bandwidth of about 100MHz. A larger bandwidth and asymmetrical shape suggests the presence of time-dependent phase distortions observed in pulsed dye amplifiers even when pumped by a very symmetric pump pulse. The sources of this phase modulation are variously attributed to heating of the dye solvent, intensity dependence of the refractive index of the dye, and, most convincingly, time dependence of the gain in the amplifying dye cells (Eyler et. al., 1994a). Heterodyne measurements of the instantaneous frequency throughout the duration of a amplifier's output pulse have revealed average shifts on the order of 10MHz and excursions of up to hundreds of megahertz. Studies also suggest the amount and predictability of chirping in an output pulse depend on the input seed power and the alignment of the pump and seed beams in the amplifier dye cells (Eyler et. al., 1994b).

Since the 120MHz bandwidth of the lidar pulsed output is not far from transform limited, and since the averaged output lineshape seems robustly repeatable, the presence of a chirp of hundreds of megahertz is unlikely. However, due to the unpredictability of the effect, especially in transversely-pumped amplifiers such as the one used in the lidar transmitter, chirping in the lidar output pulses seems a likely cause of the anomalous measured vertical velocity in the atmosphere.

The effect of imperfect seeding of the YAG pump laser could also account for additional bandwidth, asymmetry, and phase modulation in the pulsed dye amplifier's output. Studies have reported additional sidebands in the spectrum of pulses amplified with an unseeded YAG pump, along with greater shot-to-shot variability of frequency chirping and average shift; both effects attributed to modulation in the pump beam spectrum due to mode beating (Eyler et. al. 1994b). Since the seeding accuracy of the pump laser is monitored by two independent methods during routine lidar operation, and a high degree of seeding is suggested, improvement of this element of the lidar amplifier to eliminate the anomalous measured velocity has not been actively pursued. Similarly, the CW seed beam is carefully locked in absolute frequency and monitored throughout lidar operation in order to eliminate it as the cause of any frequency offset.

The leading hypothesis, therefore, is that the anomalous measured velocity offset is due to chirping and/or shifting in the pulse spectrum of the pulsed dye amplifier and arises as a natural, if not predictable, consequence of the amplifier's configuration. In order to investigate this possibility, it is necessary to consider the effect such transmitter behavior would have on the analyzed atmospheric data and determine if it is consistent with the presented observations.

6b. Effects of a Frequency Shift on the Lidar Measurement

The analysis of the 3-frequency lidar's data depends on good knowledge of the output pulse frequency of the lidar transmitter. The effect of inaccuracy in this knowledge, however, is easily determined.

As discussed earlier, since the measurement of line-of-sight velocity depends on the Doppler shift of the laser frequency with respect to atmospheric sodium atoms, the lidar's velocity measurement is most responsive to a shift in transmitter frequency. Moving sodium atoms interact with the laser pulses as though the laser frequency were shifted according to the equation:

$$\Delta\nu = \left[\frac{V}{c}\right]\nu = \frac{V}{\lambda} \qquad \text{Eq.6.1}$$

where $\Delta\nu$ is the frequency shift, V is the collective line-of-sight velocity of the atoms, c is the speed of light, and ν is the operating frequency of the laser in the laboratory reference frame. The use of a negative sign in Equation 6.1 conforms to the convention that upward velocity is positive; a positive shift is obtained for atoms moving toward the incoming photons, and a negative shift for atoms moving away. For the operating frequency of the lidar transmitter, this equates to about 1.7MHz per 1 m/s of wind velocity in the mesopause region.

It is apparent, therefore, that since the velocity measurement is based on the exact frequency of the lidar pulse which interacts in the atmosphere, a shift of this frequency in the laboratory reference frame merely implies an additional velocity with the same ratio of about +- .6m/s per 1MHz. Given this sensitivity, a typical measurement taken with the 3-frequency lidar which exhibits a -12m/s vertical velocity, when essentially zero velocity is expected, could be explained by a shift in the laser frequency of +20MHz from the value assumed in the data analysis. Such a shift may be reasonably explained based on chirps and shifts measured elsewhere for similar laser amplifier arrangements.

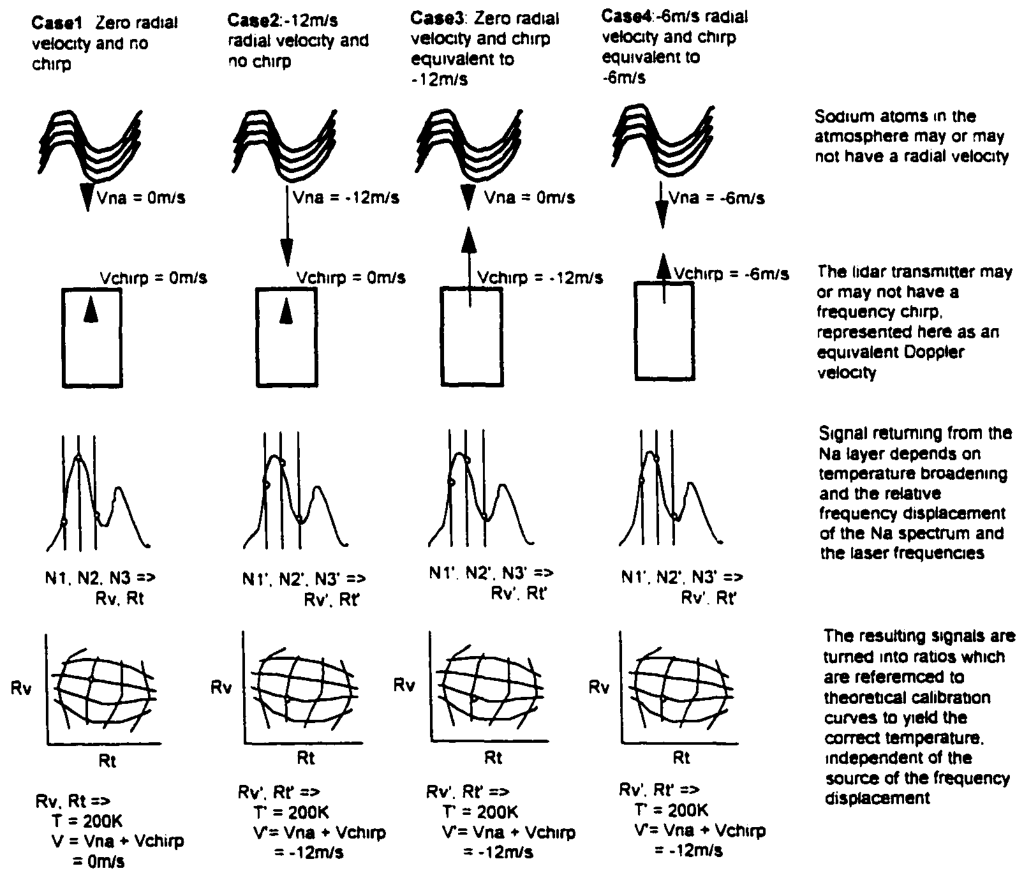

Fig.6.1 Diagram of the relationship between the temperature measurement, the velocity measurement, and possible laser frequency shifts. Note that in all four cases the temperature of the atmospheric sodium atoms are equivalent, imparting equivalent Doppler broadening to the fluorescence spectrum in each case, and therefore equivalent temperature measurements in each case.

Fortunately, the lidar measurements of temperature and Na density are far less sensitive to a frequency shift on the order of 20MHz than is the velocity measurement. In fact, the velocity measurement provides the crucial determination of the laser frequency interacting with the sodium layer which serves to correct the lidar temperature measurement for offsets in the transmitted laser frequency. Figure 6.1 suggests why this is the case.

Statistical noise aside, the returned fluorescence signals from a given altitude in the sodium layer depends only on atmospheric temperature, the transmitter frequencies, and the line-of-sight velocity. Four cases of these parameters are illustrated in Figure 6.1. If the intended laser frequencies centered at the Na D_{2a} transition are transmitted, as in cases 1 and 2, the acquired photon counts (N_1, N_2, and N_3) may be manipulated into experimental ratios (R_v and R_T), which, with the theoretical calibration curves, determine the atmospheric temperature and line-of-sight velocity (T and V). Note that a change in velocity in the Na layer in case 2 results in a relative displacement between the absorption spectrum and the transmitted laser frequencies, resulting in different photon counts (N_1', N_2', and N_3'), and different ratios (R_v' and R_T') which are interpreted in the calibration curves to yield the correct T and V as expected.

Now, if the transmitted laser frequencies are offset from their intended values by some shift $\Delta\nu_{chirp}$ (i.e. $\nu_{D2a} + \Delta\nu_{chirp}$, $\nu_{D2a} + \Delta\nu_{chirp} \pm 630$MHz) the photon counts and ratios change in the same way as in the presence of an equivalent Doppler velocity of $-\lambda\Delta\nu_{chirp}$. For this reason, the laser frequency shift is represented in Figure 6.1 as an equivalent velocity v_{chirp} where $v_{chirp} = -\lambda\Delta\nu_{chirp}$. Using the same calibration curves as in cases 1 and 2 in which no chirp was present, the photon returns and ratios obtained in the presence of transmitter chirping yield T and V + V_{chirp}. The fact that shifted output frequencies gives rise to a bias in the velocity measurement only, and not also in the temperature measurement, may be understood by realizing that the photocounts resulting from the laser stimulation of the sodium absorption spectrum are a function of the temperature and the laser frequency actually interacting with the moving atoms. For this reason, a physical sodium velocity V, and a laser chirp equivalent to V_{chirp}= V, produce the same counts and ratios, and are therefore for the purposes of the lidar experiment mathematically equivalent. Using calibration curves generated assuming the transmitted laser frequencies are precisely centered at the Na D_{2a} transition, a velocity equivalent to the total relative frequency displacement between the Na spectrum and the lidar's operating frequencies, V + V_{chirp}, and the correct temperature, T, based on this displacement are obtained. Given no additional information (such as could be provided by a frequency monitor), the correct atomic velocity in the sodium layer cannot be derived from the measured sum V + V_{chirp} in cases 3 and 4 of Figure 6.1, but in all cases the measured temperature is correct. In this sense, the atmospheric velocity and the laser frequency shift are measured together to determine the total frequency shift present at the moving sodium atoms, effectively correcting the temperature measurement.

This argument does not imply that the temperature measurement is without error. The statistical uncertainties intrinsic to the lidar data still creates uncertainty in both temperature and velocity ratio measurements, which together contribute to the overall error bars of the obtained temperature. What is suggested here is that the 3-frequency temperature/wind measurement frees its temperature measurement of systematic effects of laser frequency offset which would be present in measurements without the measurement of the laser frequency in the moving frame of the sodium atoms by the velocity ratio. Therefore, no systematic effects of a laser frequency shift present in the lidar transmitter appear in the present temperature measurement.

The Na density measurement, while also corrected by use of the frequency information included in the velocity measurement, is intrinsically insensitive to laser frequency shifts on the order of 20MHz. As outlined in Chapter 3, the Na density measurement is in principle a single-frequency measurement in which a density-sensitive ratio is formed of return signals from Rayleigh scattering at low altitude (around 25km) and resonance fluorescence in the higher sodium layer. Since the Rayleigh scattering cross-section is very insensitive to frequency shifts on the order of 20MHz and the resonance cross section at the D_{2a} peak is only weakly sensitive, the resulting change in measured density for a 20MHz change in laser frequency is on the order of only a few percent. Also, since the absolute Na density measurement is referenced to a standard atmospheric model to estimate the Rayleigh scattering strength, the measurement uncertainty greatly exceeds contributions from such a small laser frequency offset. Therefore, it is not surprising that the Na density measurement exhibits no trend which correlates with laser frequency changes on this scale.

The lidar data therefore reflects the presence of a frequency shift in the transmitter primarily in the velocity measurement, where the anomalous measured vertical velocity is seen, but the temperature and density data are considered to be consistent with the presence of such a shift, if not in direct support. To follow the postulate further, it is necessary to characterize the frequency of the lidar transmitter further and then seek experimental verification of a frequency shift.

6c. Characterization of the Implied Frequency Shift

In addition to the hour-averaged, altitude-resolved data presented in Chapter 5, it is possible to analyze the lidar data with greater temporal resolution by sacrificing vertical resolution. In order to further characterize the anomalous measured vertical velocity, one-minute averaged velocities, obtained by summing photons returning from the entire sodium layer (75 to 110km) together before calculating the experimental ratios, may be graphed. These retain high enough signal levels to provide additional insight into the temporal structure of the velocity measurement and, therefore, the frequency shift it implies.

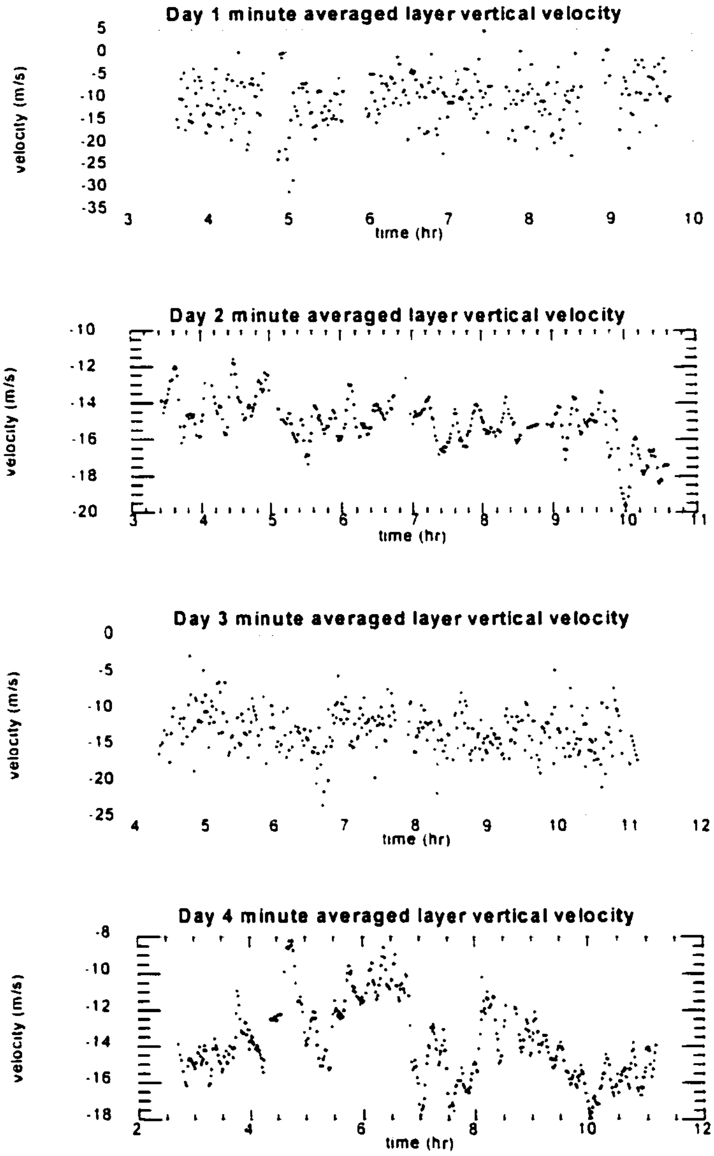

Fig.6.2 The minute-by-minute measurements of vertical velocity of the entire sodium layer over 4 lidar observation nights. Both the magnitude and variation exhibited are difficult to explain based on the atmosphere. A measurement bias is the likely cause.

Observable in several nights of minute-averaged layer velocity data is variability of the velocity value, which often averages from 10-12 m/s for an entire night, while on other nights it varies more widely. This seems consistent with characteristics of laser amplifier chirp which has been determined to be sensitive to seed laser power and amplifier alignment, two parameters which vary widely on some nights of lidar operation and little on others. Further support of this association is seen in jumps of the velocity's average value after a gap in data which designates a realignment of the optical system. Following such a pause the seed laser power and exact alignment would have very likely changed.

Adjustment of the seed laser locking is carried out throughout the night with no obvious effect on the measured velocity. No change to the amplifier pump laser is required in a typical night of lidar operation, providing no correlation with the observed variations.

Supposing a frequency chirp is responsible for both the observed vertical velocity and the asymmetry of the pulsed laser lineshape, it is informative to consider the form it might take. An entirely unchirpped light pulse, as would arise from a chopped CW beam, may be written:

$$E(t) = E_p(t)\cos(2\pi\nu_0 t)$$

$$\text{where } E_p(t) = \exp\left(-\left[\frac{t}{\tau}\right]^2\right) \qquad \text{Eq.6.2}$$

in which $E_p(t)$ contains the envelope with which the pulse was turned on and off again, and ν_0 is the carrier light frequency, in our case corresponding to a wavelength of 589nm. Now, in the case of a chirped pulse, a potentially complicated modulation, $\varphi(t)$, is introduced into the carrier frequency:

$$E(t) = E_p(t)\cos(2\pi\nu_0 t + \phi(t)) \qquad \text{Eq.6.3}$$

where

$$\phi(t) = at + bt^2 + ct^3 + \cdots \qquad \text{Eq.6.4}$$

and the instantaneous frequency chirp may be expressed as:

$$f_{ins} = \frac{1}{2\pi}\frac{d\phi}{dt} \qquad \text{Eq.6.5}$$

Since the Fourier transforms of cos(at) and cos(bt^2) are symmetric, the first term in the modulation $\varphi(t)$ which contributes to the asymmetry of the pulsed lineshape about its centroid is ct^3. It is possible therefore to construct a simple model of a chirp modulation function which approximately recreates the measured lineshape of the lidar pulsed output in bandwidth and general asymmetry, while introducing a center of mass offset of about 20MHz, similar to that inferred from the vertical velocity measurements:

$$\phi_{model}(t) = at + ct^3$$
$$a = .212ns^{-1} \qquad \text{Eq.6.6}$$
$$b = -0.014ns^{-3}$$

Which yields the simulated pulsed power spectrum in fig. 6.3.

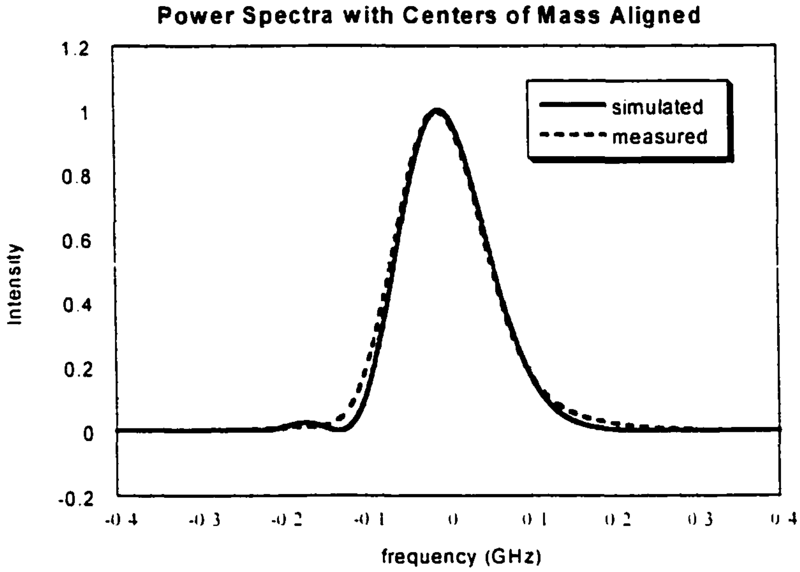

Fig.6.3 The measured power spectrum of the pulsed dye amplifier (dashed) and the power spectrum arising from a simple model in which constant and cubic terms of a frequency modulation function have been chosen to yield an offset and asymmetry similar to the measured spectrum.

By computing the derivative of the modulation function according to eq. 6.5. the instantaneous frequency evolution corresponding to this simple model is obtained in Figure 6.4.

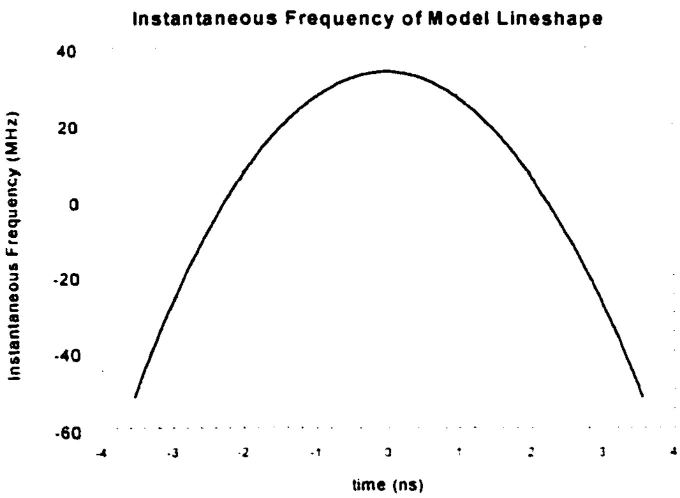

Fig.6.4 The instantaneous frequency function of a modulation consistent with the measured pulsed power spectrum. The frequency is seen to change over 80MHz from the 1/e rise to 1/e fall of the laser pulse in time.

Figure 6.4 suggests that instantaneous frequency excursions in excess of 40MHz could contribute to a power spectrum with similar bandwidth, asymmetry, and center of mass offset as is believed to be present in the output of the lidar's pulsed dye amplifier. Comparison of this simply modeled frequency progression with measured instantaneous frequency progressions from similar Rhodamine dyes suggests that the 40MHz magnitude to be reasonable or even modest (Eyler et. al., 1994a). The consequence of these investigations is to suggest that the anomalous velocity measured by the lidar is compatible with the observed pulsed lineshape and current understanding of pulsed dye amplifier chirping.

Given that the pulse output of the lidar is consistent with the presence of a frequency chirp, it becomes useful to attempt to experimentally verify the existence of such a chirp. The lineshape measurement alone, while consistent with a chirp which would shift the center of mass 20MHz and explain much of the anomalous measured velocity, is not a measurement in absolute frequency. That is, the center of mass of the pulsed lineshape is not referenced to spectroscopic features or the CW seed laser frequency. A simple experiment for comparing the absolute frequency of the pulse lineshape centroid with the CW frequency using a laboratory iodine vapor cell was developed. In this experiment a sharp absorption feature in the iodine spectrum was selected and absorption spectra recorded simultaneously with both the seeded lidar pulsed output and a small amount of the CW seed beam. Also recorded simultaneously were a 750MHz free specral range interferometer and the sodium Doppler-free spectrum for frequency scaling of the acquired absorption spectra. Since the pulsed output has a bandwidth of about 120MHz whereas the CW laser is only KHz, the spectra should not be identical unless the CW spectrum is first correlated with the measured pulse line shape. Once this is done the two iodine spectra should have the same shape and any discrepancy in absolute frequency location should correspond to the difference in frequency between the two lasers.

The results of this experiment indeed suggest a frequency discrepancy between the CW seed laser and the pulsed center of mass of the correct sign and within a factor of two of the 20MHz necessary to adequately explain the measured vertical velocity. Limitations in the accuracy of the experiment, mostly due to imperfect frequency scaling with the 750MHz interferometer, along with the long timespans necessary to carry it out, preclude this method as a viable monitor for the correction of the lidar velocity measurement. However, it suggests that given the proper technique, whereby an accurate measurement of the center of mass frequency is obtained on the same timescale as the lidar integration time, the vertical wind measurement could be corrected from its geophysically impossible value.

6d. Techniques for Correcting the Measured Anomaly

If the systematic error apparent in the vertical wind measurement is indeed attributable to a shift in the centroid of the pulsed line shape due to frequency chirping in the pulsed dye amplifier, either a reduction of the chirp, or a measurement of the location of the pulsed centroid relative to the CW seed laser, or an independent measurement of the pulsed centroid location in absolute frequency is required. In either of the latter two cases, the measurement would be the most applicable if executed simultaneously with the lidar operation. in order to be used to correct the lidar data as the chirp changes. if necessary.

As was previously discussed. the amount of chirp exhibited by a pulsed dye amplifier depends on a variety of parameters. many of which are intrinsic to the geometry and dye employed. Eyler et.al. have shown that a mix of dyes may be successfully formulated in order to reduce the chirp in a given amplifier without adversely affecting the gain of the amplifier. This technique. while still an interesting option for the lidar community. requires an accurate measurement technique for measuring the frequency chirp. and may not prove robust over the long collection times required by lidar. or for transversely pumped amplifiers. More research is required to fully assess the applicability of this to lidar. an option which the CSU lidar lab has not yet fully addressed.

A simple monitoring option involves simply using the vertical velocity measurement in the sodium layer as a measurement of the transmitter's frequency offset (Tao and Gardner. 1995). In this scheme it is assumed that the vertical velocity is zero for the purposes of the lidar. and the frequency offset of the lidar transmitter is identified as that which reduces the vertical velocity measurement to zero. For continuous correction this would require one beam to be maintained at zenith as the frequency monitor. For most lidar systems this is a discouraging prospect due to the beam power and telescope area required to accurately measure wind 90km above the earth. and for even powerful transmitters the geophysical variability intrinsic to the wind measurement reduces the accuracy of the frequency monitor. For near-continuous correction. a rapidly steerable telescope could be used which moves to zenith periodically to supply the frequency reference. The availability and cost of such receiving systems however is prohibitive for most research programs. A final disadvantage of the 'atmospheric monitor' approach is that such a system does not scale well to very powerful transmitters. Short integration timescales could weaken the required assumption of zero vertical velocity.

The principal method for detailed analysis of the instantaneous frequency progression in pulsed lasers is the heterodyne measurement. In this technique a small portion of the CW seed light is modulated in frequency. hopefully to a frequency beyond the Fourier bandwidth of the pulsed light. and superimposed on the collector of a fast detector with the pulsed light which is to be analyzed. The resulting signal exhibits a beat signal at a frequency equal to the amount the CW light was upconverted plus the difference in the carrier frequencies of the pulsed and CW beams. The analysis of the structure of the beat signal provides the most accurate way to probe the instantaneous frequency of the pulse, any overall offset, and in principle allows repetition the measurement with every pulse. The disadvantages of this experiment as a lidar frequency monitor is primarily the required signal processing speed. In order to move the important beat signal away from contamination by the wings of the pulsed bandwidth. an upconversion of several hundred megahertz may be required. necessitating the use of detection electronics capable of analyzing the (not-necessarily sinusoidal) beat signal of at least that frequency. This requirement necessitates very high sampling rates, careful filtering, and therefore high cost. As we have seen, the non-symmetric lineshape of the lidar output pulse implies a possibly complex frequency progression between the start and end of the pulse, suggesting that the beat signal might not easily yield the averaged offset without some effort, perhaps making such a system difficult to automate. Finally, this method provides only the difference in frequency between the CW seed frequency and the pulse output's carrier frequency, making it incapable of correcting for imperfect CW laser locking.

A final method for monitoring the frequency of the lidar pulsed output, and the one adopted for the CSU 3-frequency lidar, improves upon the iodine absorption spectroscopy experiment outlined earlier. By capitalizing on the fortuitous location of an iodine spectral feature in the vicinity of the sodium absorption spectrum, and therefore the three operating frequencies of the lidar, the absolute frequency of the output pulse can be determined from a ratio of the transmission through the iodine cell at each of the operating frequencies. A sensitivity at the operating frequencies can be obtained with proper heating and temperature control of the iodine cell, providing in principle a good measurement of the absolute frequency for each complete cycle of the three operating frequencies. One advantage of an absolute frequency measurement is the monitor's ability to correct for CW laser locking inaccuracy as well as chirp-induced frequency shifts. The disadvantages and intricacies of this method include the necessity of two detection channels, gated integrators, and repetition rate doubling electronics, which will be discussed in detail in the following chapter.

Chapter 7: A Frequency Monitor for the 3-Frequency Na Lidar

7a. Basic Theory of Operation

The basis of the frequency monitoring system developed for the 3-frequency lidar is the absorption spectrum of iodine in the vicinity of the lidar's operating frequencies. Quantum mechanical transitions in the fine structure of iodine produce strong absorption bands about 1GHz higher and 1.5 GHz lower in frequency than the sodium D_{2a} Doppler-free peak which marks the center frequency of the lidar output. The transmission of a tunable, narrow-band laser through a laboratory iodine cell, shown in Figure 7.1, is a smooth transmission peak centered about 400MHz less than the D_{2a} Lamb dip (chosen as the reference frequency), with a full width at half maximum of about 1.1 GHz when heated to 78°C. A truly fortunate alignment of this transmission peak at the proper temperature and both the D_{2a} peak and the -630MHz operating frequency allows both edges of the peak to be employed for frequency determination.

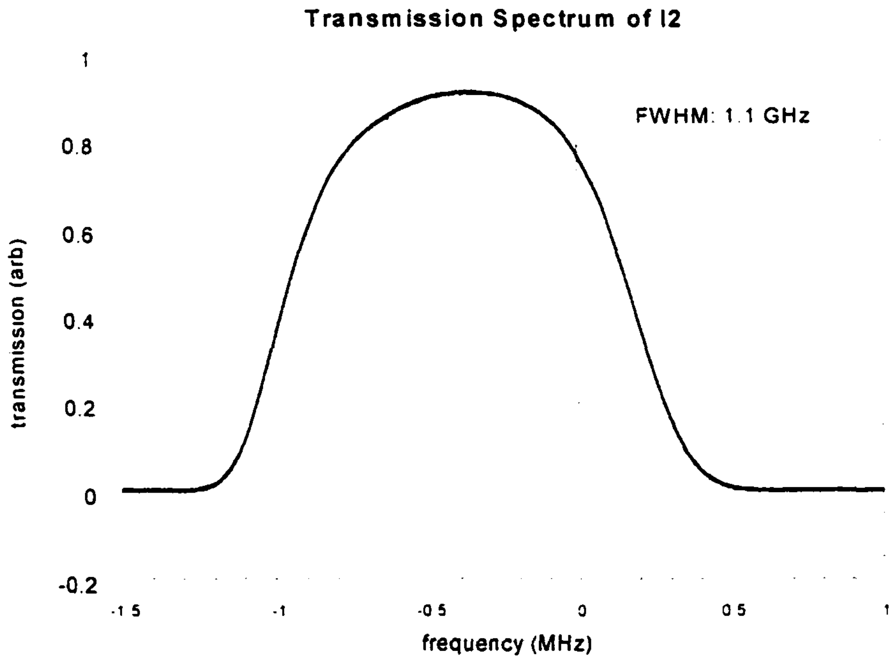

Fig.7.1 The location of an iodine transmission line with respect to the three operating frequencies of the lidar. The spectrum was acquired using a 4-inch iodine cell at 78 C and tunable CW laser light.

Since the frequency monitor depends on the two spectral edges of this transmission peak, for the purposes of this thesis it will be referred to as 'the iodine peak' despite the fact that two absorption lines conspire to form it.

Since the three operating frequencies of the lidar, ν_{D2a} and $\nu_{D2a}\pm630MHz$, are generated by the actively locked ring dye laser and the tandem AO modulator system, their separation is assumed to be fixed regardless of any frequency shift imparted by the pulsed dye amplifier. This justification for this assumption is based on the AO construction discussed in Chapter 4 and the fact that the scale of the amplification bandwidth of the pulsed dye amplifier is large in comparison to which the 630MHz AO-shift frequency. The assumption will be discussed further in Chapter 8 of this thesis.

If the separation between the operating frequencies is assumed to be constant, then the frequency monitor may use all three to detect an overall frequency shift present in the lidar transmitter. From Figure 7.1 it is clear that for a positive frequency shift, the transmission of the -630MHz pulses should increase, while the transmission at the D_{2a} peak frequency should decrease. The difference in the relative transmissions at each of these two operating frequencies, therefore, is sensitive to any overall shift in the transmitter. Though the difference in transmission corresponding to a 1MHz shift is small for each frequency, a good average of each is obtainable in an integration time comparable to the integration time used to accumulate the lidar signal. Moreover, since the location of the iodine peak is a quantum mechanical certainty (given a stable iodine temperature among other parameters), the frequency shift is measured in absolute frequency rather than referenced to the CW seed laser frequency.

The third operating frequency, +630MHz, while too far from the steep (and therefore frequency-sensitive) edges of the iodine peak to contribute significantly to the monitor's sensitivity, is employed based on its location near an absorption maximum of the iodine spectrum. In the event of broadband transmission of the lidar transmitter, arising from amplifier ASE or imperfect YAG laser seeding for example, the +630MHz operating frequency's relative transmission through the iodine cell changes most perceptibly of the three. This frequency, therefore, can provide monitoring and possible correction based the on spectral purity of the lidar transmitter.

The short pulse duration, high peak power, and uneven mode shape of the lidar output laser pulses contribute some experimental difficulty to the precise measurement of the relative transmissions at the three operating frequencies, and some spectroscopic precision is required to characterize the iodine transmission in absolute frequency accurately enough to correctly resolve the shift of the 140MHz pulsed lineshape's centroid by 1MHz. The experimental engagement of these problems is discussed in the remainder of this chapter.

7b. The Experimental Theory

To discuss the necessities of the monitor's experimental design, it is informative to first consider various factors which contribute to the measurement of the pulsed output of the 3-frequency lidar. Since small changes in transmission through an iodine vapor cell carries the crucial frequency information, the essence of the frequency monitor is a pulse intensity measurement. The iodine cell and two photodiode detectors are arranged in a light-tight box as shown in Figure 7.2.

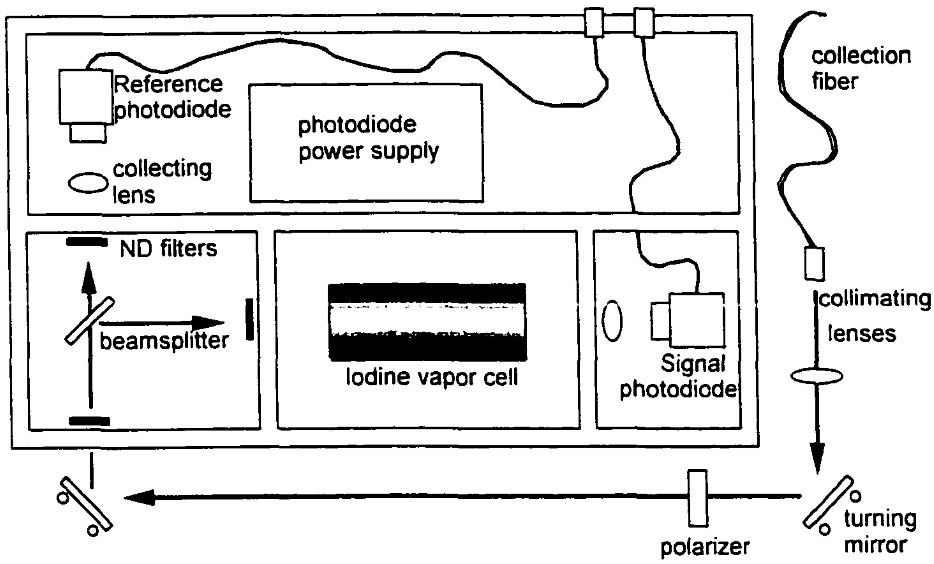

Fig.7.2 A schematic of the experimental apparatus for the frequency monitor measurement. The double walls represent a light-tight housing with sub-chambers. the beam path is in black and electronic signal lines in gray.

Signal Sources

Designating the diode which measures light transmitted through the iodine cell as the 'signal' diode, and the second diode required for laser power normalization the 'reference' diode, we may write the signal obtained by each diode respectively as signal, S_i, and reference, R_i,:

$$S_i = FP_iGCQ_i + O; \quad i = -, p, +$$

$$R_i = fP_ig + o; \quad i = -, p, + \qquad \text{Eq7.1}$$

where:

F,f = fractional split
P_i = laser power
G,g = overall gain
C = cell transmission
Q_i = iodine transmission
O,o = DC offset

The variables on the right hand side of each equation represent sources which contribute to the measurable signals, and, with the exception of Q_i, contaminate the intended frequency measurement. Therefore, each must be eliminated in the experimental design and data analysis. Each will be discussed in turn.

The variable i designates the frequency of the laser pulse in the cycle of operating frequencies (- = ν_{D2a}-630MHz, p = ν_{D2a}, + = ν_{D2a}+630MHz). Since the lidar transmitter cycles at its 20Hz firing rate through the three frequencies, each diode measures all three and the frequency monitor depends on the measurement of all three. Some contributions to the diode signals cycle with frequency, such as the laser power and iodine filter transmission, while others, such as the electronic gain of each diode, do not.

The variables F and f represent the fraction of laser power from the optical fiber which is split to each diode channel. The two beampaths to the detectors traverse a glass splitter and neutral density filters. The fractions are therefore certainly unequal; but the passive filters and beamsplitter are considered insensitive to the transmitter's comparatively small frequency shifts, and so the fractions are not referenced to the laser frequency. A polarizer is placed before the light-tight box to insure that polarization variations do not contaminate the stability of the beam splitter's split-ratio.

The variable P_i represents the laser pulse power coupled out of the collection fiber. Since the three operating frequencies are generated by acousto-optical crystals with imperfect conversion efficiency, the lidar output power varies considerably from pulse to pulse, especially with the cycling of the output frequency given by i.

The variables G and g designate the electronic gain of each diode. The detection bandwidth of the diodes is orders of magnitude greater than the 630MHz frequency shift generated by the acousto-optical modulators, but systematic effects can still arise from the detectors' interaction with a changing transverse mode shape in the laser correlated with its frequency shifting. This is one principal reason for using an optical fiber for collecting the laser light used in the frequency monitor. The fiber output exhibits a smoothed, Gaussian output mode for each of the operating frequencies. Therefore, the electronic gain is assumed to be independent of frequency. However, though the diodes and their amplifiers and digitizers are of the same model, differences within the models, and especially between individual gated integrators and A to D channels, as well as diode alignment, requires that the gains of the two diode channels to be considered unequal.

The variable C represents the frequency independent part of the iodine cell transmission, included separately due to the propensity for the invisible iodine vapor to solidify on the windows of the cell, often creating a time dependent transmission function, yet independent of the laser frequency.

The variables O and o designate DC offsets added to the diode signals, primarily in or after the gated integrators where the fast AC signals from the diodes are converted to DC values for digitizing and analysis.

Finally, the variable Q_i indicates the frequency-dependent, shift-sensitive transmission of the iodine cell as determined by the quantum mechanics of the I_2 molecule and the length and temperature of the I_2 cell. It is the fractional change in this variable with a change in transmitter frequency given by i which alone facilitates the monitor's measurement of frequency. All other variables in Equation 7.1 must be eliminated in the experimental design and data analysis in order for Q_i to accurately determine the frequency shift.

Measurement Design and Analysis

The first variables which require special hardware design to overcome are the DC offsets O and o. The gated integrators used in the 3-frequency lidar frequency monitor (SRS SR245) are prone to large and rapid DC baseline fluctuations apparently at random. For this reason it is often necessary when using gated integrators to initiate a measurement between data pulses in order to evaluate and subtract the baseline, and SRS gated integrators include averaging modes in which this is carried out automatically. However, since the lidar transmitter changes operating frequency with every shot, the pulses must remain distinct in order to average together only pulses arising from the same transmitter frequency. For this reason a circuit was built which doubles the trigger frequency from the pulsed transmitter which would ordinarily initiate the gated integrators. With the gated amplifiers triggered at twice the repetition rate of the laser, a baseline measurement is made between each laser shot, at a time when no laser light is present in the monitor's light tight box. Using this frequency doubling or trigger 'echo' box, a good average offset is measured for the monitor's integration time. For these between-pulse measurements, Equation 7.1. becomes:

$$S_{echo} = O \qquad \text{Eq7.2}$$
$$R_{echo} = o$$

Therefore, by subtracting the averaged echo signals from S_i and R_i signals which arise from actual laser pulses, the new diode signals become:

$$s_i \equiv S_i - S_{echo} = FP_iGCQ_i : \quad i = -, p, +$$
$$\text{eq.7.3}$$
$$r_i \equiv R_i - R_{echo} = fP_ig : \quad i = -, p, +$$

The next, obvious choice is to use the reference channel in the manner it was intended and eliminate the frequency dependent power term:

$$\Sigma_i \equiv s_i / r_i = (F/f)\,(G/g)\,C\,Q_i \quad i = -, p, + \qquad \text{eq.7.4}$$

By dividing the signals, after digitization and integration, from the two diodes, the frequency dependence has been eliminated from all but the iodine vapor transmission function, where it is intended to remain. However, the value of the measurement quantity s_i / r_i still depends on the exact transmission of each beam path, the electronics particulars of each channel, and perhaps most difficult to measure, the exact transmission of the iodine cell, which, as noted above, tends to change randomly in time. The final step, then, is to ratio the result of Equation 7.4 based on the frequency of the transmitter.

$$\Sigma_{-} / \Sigma_{p} = Q_{-} / Q_{p} \qquad \text{eq.7.5}$$

The quantity in Equation 7.5 is the frequency sensitive function upon which the monitor design relies: it depends only on the iodine transmission lineshape and the absolute frequencies which the subscripts represent. One might wonder about the cancellation of C, the time-dependent iodine transmission function, carried out between Equations 7.4 and 5. In the case of a lidar system which integrates at each frequency individually for some time before moving to the next operating frequency, a time dependent quantity could introduce error in the frequency measurement due to imperfect cancellation in this step. Fortunately, since the CSU 3-frequency lidar shifts frequency on a shot by shot basis, the three frequencies used in the monitor are summed and averaged simultaneously over the integration time, restoring the accuracy of Equation 7.5.

One disadvantage of this approach apparent from the above discussion is the propagation of operations, particularly ratios, performed among measured values, each with some level of noise and uncertainty. With each operation the uncertainty of the result increases, necessitating very low noise levels in the initial measured quantities. Low noise measurements are obtainable with proper experimental design, some details of which are provided in the following discussion.

7c. The I_2 Cell and Oven Design

The principal parameter which affects the shape of the iodine transmission spectrum is the temperature at which the cell is maintained. The cell itself is a 4 inch long and 1 inch diameter glass cylinder with a tapered finger or nipple protruding about 1 inch radially near the halfway point. Since iodine is an opaque solid near room temperature, it is important to have a reservoir which is maintained at a lower temperature than the rest of the cell so that the iodine which is not vaporized, and therefore not participating in the filtering application of the cell, does not solidify on the windows and reduce the peak transmission. Once the cell is built, filled, and sealed, the users only control of the iodine vapor pressure is the temperature of the solid iodine enclosed in the tapered finger. It is important that the cell is filled with enough iodine before the cell is sealed to ensure the presence of solid iodine in the tapered finger during operation.

To maintain the cell at a selected temperature, an aluminum oven was constructed into which the glass cell is inserted. Good heat conduction between the aluminum and glass is obtained with thermally conducting paste. In order to maintain a uniform temperature over the cell, a substantial amount of aluminum was used to disperse the heat from the heater elements, and the oven was designed to maintain as much cylindrical symmetry as possible. About one inch of aluminum separates the heating elements at the surface of the aluminum oven and glass of the cell wall. The temperature control system is based on a Fuji digital PID (Fuji PXV3-NAYI-4V) temperature controller with an RTD detector (Minco S308PD) as its feedback. The programming of the temperature controller calculates the proportional, integral, and derivative coefficients based on user-defined parameters, such as temperature set point and allowed hysteresis, and based on an automated test of the temperature response to an applied heating current of the system that is being controlled. To heat the oven, the controller modulates the duty cycle of slow DC voltage pulses applied to heating element strips fixed around the cylindrical oven. It was determined that for most accurate control of the oven, and therefore the cell temperature, the feedback RTD should be imbedded in the aluminum about midway between the heater elements and the cell. In this way a balance between responsiveness of the feedback detector to applied current and feedback lag due to the thermal mass of the oven/cell system is achieved. A second thermal detector is placed in contact with the cell wall using thermal paste for monitoring of the resulting cell temperature.

Using the Fuji temperature controller, an instantaneous temperature uncertainty within ±.05°C of the selected temperature set point was achieved, with very good adherence to the set point over long time scales. An average of the cell temperature over one minute, corresponding to the typical lidar integration time, typically equals the setpoint to better than ±.01°C, even in the winter when the thermal environment of the frequency monitor during lidar operation may vary by as much as 30°C. To relate this temperature controller performance to the accuracy of the lidar frequency monitor it is necessary to consider in some detail the temperature dependence of the iodine transmission peak. Figure 7.3 shows iodine transmission spectra measured by scanning the CW ring dye laser, each at a different iodine cell temperature.

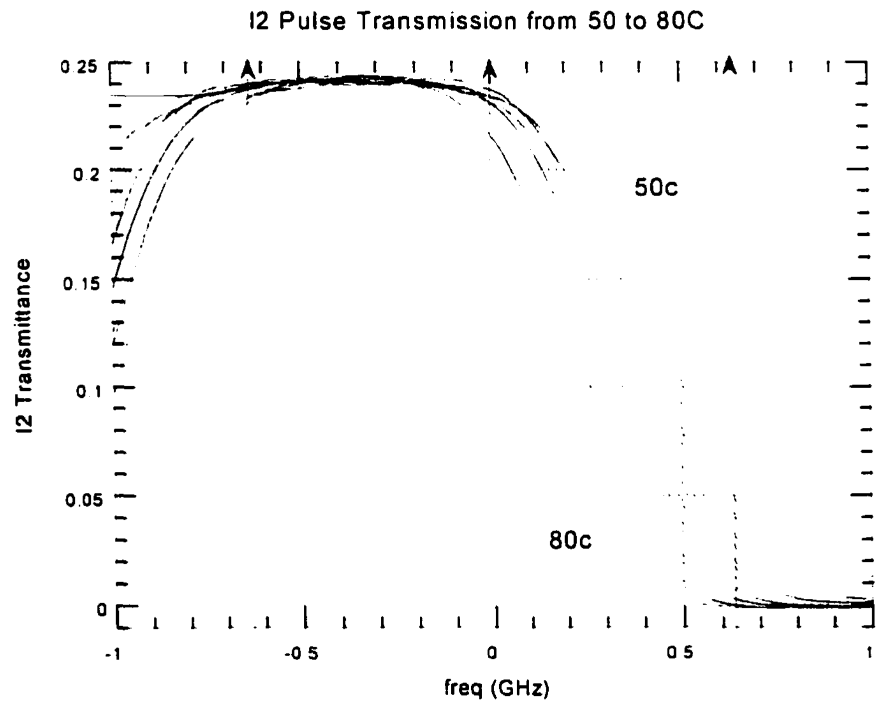

Fig.7.3 The temperature dependence of the iodine transmission spectrum with changing cell-temperatures referenced to the operating frequencies of the lidar. The spectra have not been re-normalized to preserve the effect of imperfect extinction at lower temperatures. The absolute transmittance scale is arbitrarily chosen.

The transmissions at the three operating frequencies may be deduced from Figure 7.3 and plotted versus temperature in Figure 7.4.

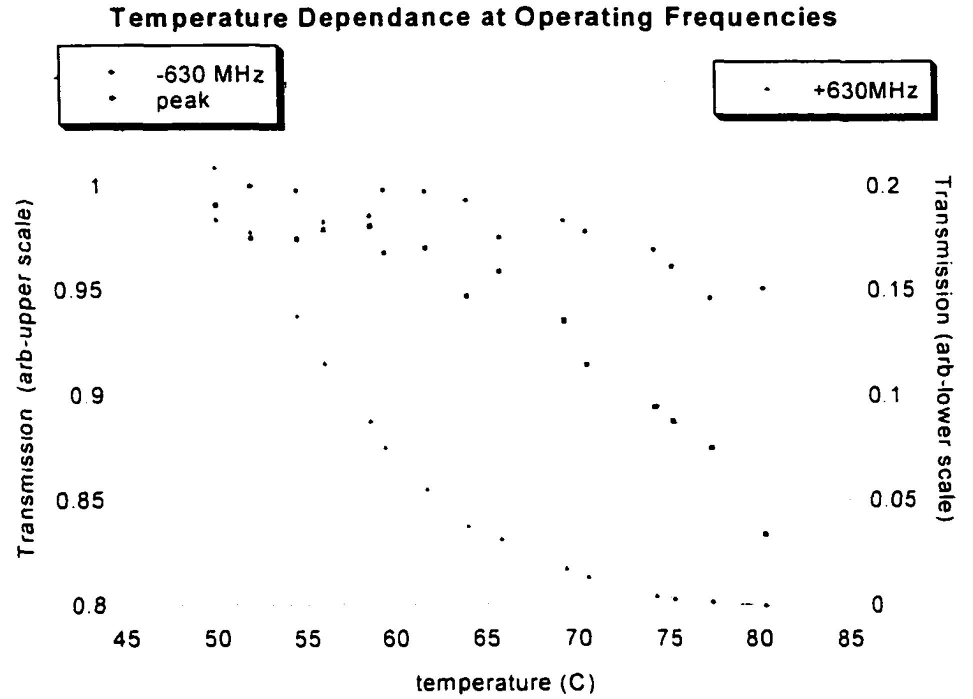

Fig.7.4 Detailed look at the transmittance change at each of the operating frequencies with changing temperature. Despite point-to-point uncertainties of 5% or more, trends are apparent and cubic fits have been added to aid the eye. Due to it's much smaller magnitude the -630MHz response is plotted on an expanded, but still relative, scale.

From Figure 7.4, the error in transmission due to temperature drift around a given setpoint can be deduced. Since the temperature used in the lidar frequency monitor is 78°C, it is informative to focus our attention on that set point. The slopes of the curves in Figure 7.4 around 78°C for each of the lidar's operating frequencies suggest a sensitivity of relative transmission to temperature of about 1% per °C for the D_{2a} peak frequency, .3% per °C for the +630MHz edge and .01% per °C for the -630MHz edge. It is interesting to note that these sensitivities are not equal, of different signs, and are approximated from curves which are not genuinely linear. Therefore ratios of relative transmission values generated from all or some of these operating frequencies may differ in their overall sensitivity to temperature. Even for the peak frequency, the iodine transmission which varies the most dramatically, the temperature uncertainty of .01°C exhibited by the lidar frequency monitor's iodine cell in one minute suggests that temperature-induced error is not an important contribution to the overall error budget. It is, however, interesting to note from Figure 7.4 that the temperature sensitivity for the peak and -630MHz frequencies increases as the setpoint temperature increases. This, combined with the likelihood that the temperature controller will operate less effectively at higher temperatures suggests that at very high cell temperatures, the induced error might become a problem.

7d. The Detection and Integration Electronics

The frequency monitoring measurement depends not only on carefully stabilized optical filtering, but also on accurate conversion of the optical signal into electronic signals to be processed, analyzed, and stored. The chain of electronics used for acquisition of the pulsed information is outlined in Figure 7.5.

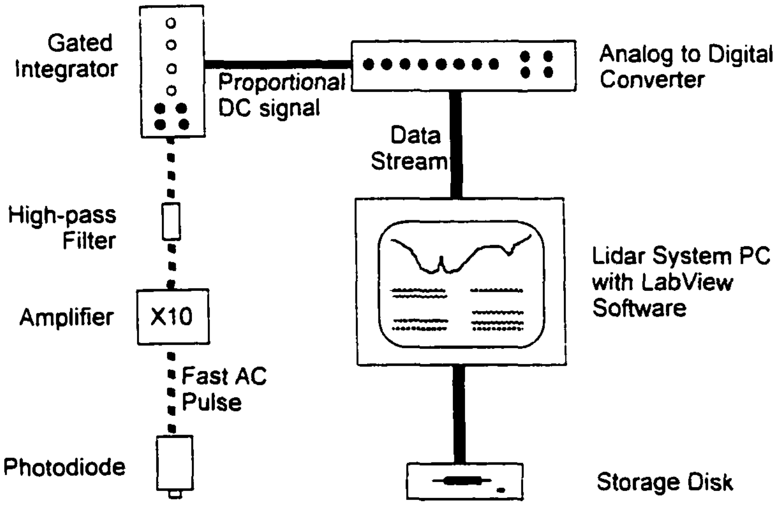

Fig.7.5 A block diagram of the detection electronics used for the frequency monitor experiment. The data flow begins at the photodiode and ends as a time-averaged signal saved to disk during routine lidar measurements. Two channels represented by this diagram are used for both the signal and reference measurements.

The diodes (EG&G FND-100) are reverse biased with -90 volts for a faster response time. Since it is important for the measurement that the two detector channels remain absolutely uncoupled from one another, two separate supplies are used for the reverse bias voltages of the two detectors. The best, fast, most noiseless response was found using batteries for these voltage supplies, and good results were obtained by interlocking ten standard 9-volt batteries. Good linearity over a wide range of incident light powers was observed, a crucial element for the lidar frequency monitor due to the transmitter power fluctuations which often occur in the course of taking lidar data over periods in excess of 10 hours. The light level incident on the monitoring photodiodes was typically adjusted to initiate voltage pulses of 10 to 20mV from the diodes into a 50Ω impedance; this corresponded to much less than 1mW of continuous-averaged power. Even with small levels of incident light, some residual ringing after the nanosecond-scale pulse was observed. The majority of this spurious signal can be rejected by the gated integrators further along the detection chain, but some ringing is inextricably mixed with the pulse signal and may contribute to the overall noise of the signal. This eventuality will be discussed further below.

From the photodiodes the voltage pulse response travels less than 1 meter before being amplified for longer distance transmission. The amplifiers (Phillips 6950) have a gain of ten, yielding a typical output pulse between 100 and 200mV. Longer coaxial cables then carry the pulses to the remainder of the detection chain.

Upon arrival to the gated integrators, the electronic pulses are AC coupled through a passive RC high-pass filter with about a one kilohertz cutoff frequency to eliminate any electronic line noise which might have been superimposed. The gated integrators require two inputs: the pulses to be integrated and a trigger signal to initiate the gating window delay. As discussed above, the trigger signal originates from the pulsed YAG laser, the Q-switch output of which is coincident with the laser output pulse to within a few nanoseconds. The Q-switch output frequency is then doubled by the 'echo' circuitry described before. So, for each shot of the pulsed laser, the gated integrator is triggered twice, once to integrate the voltage pulses arising from the interaction of the laser pulse and the frequency monitor, and once to integrate the monitor output lines with no laser light present. Referenced from the arrival of the trigger, both the delay before integration begins and the duration of integration can be precisely set and monitored from the gated integrator unit. The utility of such an integrator is twofold. First, since the voltage pulses are only nanoseconds long, determining a peak height proportional to the intensity of the light incident on the diode is a very time-critical operation, dependent on one point in time on an already short signal. Since the area of the voltage pulse should also be proportional to the light intensity, an integration of the area under the pulse uses all of the information supplied by the diode, thereby increasing the signal-to-noise ratio of the detection. Second, since the window of integration is fully selectable with widths down to a few nanoseconds, all of the time not containing information of interest, but very likely containing some level of unwanted noise, is rejected, again to the advantage of detection signal-to-noise. The output of the gated integrators is an amplified DC voltage proportional to the integrated area, which is held until the next trigger again initiates integration. Therefore, the digitizer which follows the gated integrator in the detection electronics chain has almost the entire 25ms between echo triggers to acquire a low noise value proportional to the light intensity collected by the photodiode.

Following the gated integrators for both detection channels, the DC outputs are sent to an eight channel, 16-bit analog to digital (A-to-D) converter (National Instruments Daq AT-MIO-XE-10). Since the output of the integrators is DC, speed is not a primary concern of the digitizer, as it would be if the nanosecond voltage pulse were to be resolved by the digitizer. Instead, the more modestly priced 16-bit digitizer provides good voltage resolution of the integrated results, for analysis and storage by a personal computer. The A-to-D unit also manages several digital inputs and analog and digital outputs which are managed by the computer software to orchestrate other parts of the lidar system such as tuning the ring dye laser, initiating active locking feedback, and synchronizing the software with the echoed triggers.

All the information which courses through the A-to-D converter is organized by custom LabView software responsible for scanning, locking and monitoring the laser systems of the lidar transmitter. The frequency monitor information arrives on two analog input channels from the A-to-D and is separated according to channel, (signal or reference) and operating frequency (peak, +630MHz, or -630MHz). Throughout the period of the lidar integration time the following twelve quantities from the frequency monitor are separately averaged: On the first channel, three signals corresponding to light pulses of each operating frequency arriving at the reference photodiode, and three signals corresponding to background levels after each of those pulses, and on the second channel three signals corresponding to light pulses of each operating frequency arriving through the iodine cell to the signal photodiode and three signals corresponding to background levels after each of those pulses. These twelve quantities are all used in the experimental ratio which determines the lidar output frequency, and all twelve are saved to disk at the end of the integration time for each lidar photocount profile. The software also calculates the frequency for that integration time and displays it for monitoring during routine lidar operation.

Though the detection process detailed above is conceived to minimize noise in the frequency monitor measurement, some variation and uncertainty is unavoidable. Examples of the measurements which together determine the lidar transmitter frequency, along with an analysis of the accuracy with which they are made follow in the next section.

7e. Detection Accuracy

Due to the statistics of the atmospheric measurement, a compromise is necessary between resolution in time and resolution in altitude in the analysis of the lidar data, usually one hour temporal averages and vertical smoothing by a Hanning window of 7.4km full width. Typically, the 3-frequency temperature/velocity lidar has been operated with a 20Hz pulse repetition rate and a one minute hardware integration time (the time between saving acquired sums to the computer); this represents the smallest possible timescale which can be resolved by the lidar. Of the 1200 laser shots sent into the atmosphere in one minute, 400 are at each of the operating frequencies. Therefore, each of the twelve measured quantities used in the lidar frequency determination represent a minimum of 400 averaged shots, and often more depending on the temporal averaging desired in the data analysis. For the purposes of error analysis, investigating the 400 shot averages will provide results which are easily extended to an arbitrary number of minutes. Typical shot-by-shot measurements of the twelve quantities are shown in Figure 7.6.

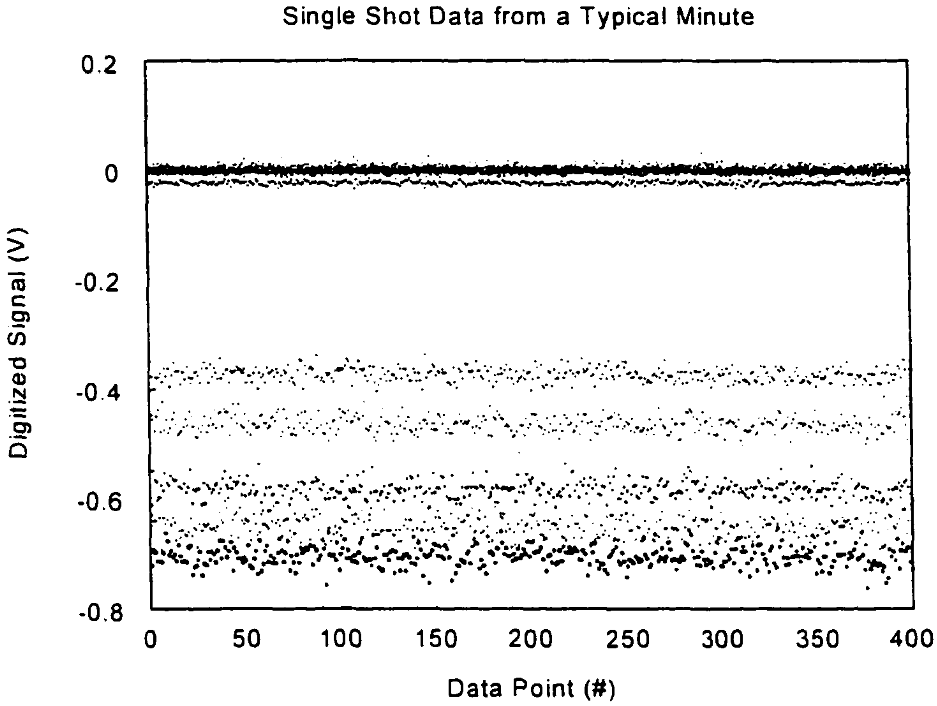

Fig.7.6 Typical shot-by-shot measurements of the 12 parameters needed for the frequency monitor measurement. Five values are the non-zero rows, one value is close to zero, and six are overlaid at zero. These correspond to the signal from the three operating frequencies, the reference from the three frequencies, and six background measurements taken between shots.

In Figure 7.6, the six values from the monitor which correspond to laser pulses are seen as the six non-zero bands: the +630Mhz pulse which travels through the iodine filter falling close to zero, as would be expected from the transmission curves above. The six values which correspond to "echo" readings without laser light present are superimposed around zero volts. It is apparent that the standard deviation of the readings increase with the mean value as would be expected for normally distributed noise, and as is typical of most electronic measurements, the standard deviation remains non-zero, even for a mean very close to zero. In order to characterize the standard deviation over a wide range of mean values corresponding to widely varied laser power, about one hundred measurements similar to, and including, the twelve in Figure 7.6 are plotted verses their standard deviation in Figure 7.7.

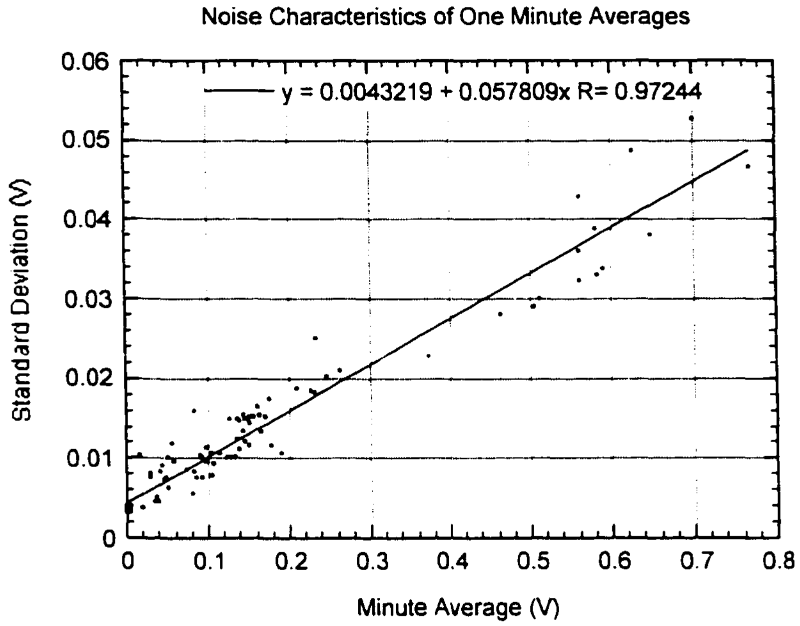

Fig.7.7 Based on shot-by-shot measurements, the standard deviation of the frequency monitor measurements correspond to a roughly linear function of signal strength.

From Figure 7.7 a clear linear trend is discerned, along with a well determined y-intercept. It should be noted that the standard deviations represented in Figure 7.7 include both electronic noise present in the detection chain and laser amplitude and frequency variation over the minute integration time, frequency variation appearing only in quantities filtered by the iodine cell. Laser amplitude (or power) variation is correlated between the measurements of the signal and reference channels, reducing the contribution of amplitude instability to the uncertainty of ratios of the two channels. Given the standard deviation, the fractional uncertainty of the mean may be calculated using:

$$\frac{\Delta M}{M} = \frac{\sigma}{\sqrt{N}}\frac{1}{M}$$

where

ΔM = uncertainty of th mean Eq.7.6

σ = standard deviation

N = number of samples in the average

M = mean

The fractional error of the mean can be determined for any mean from the equation of the linear trend in Figure 7.7. This is plotted verses the mean values in Figure 7.8.

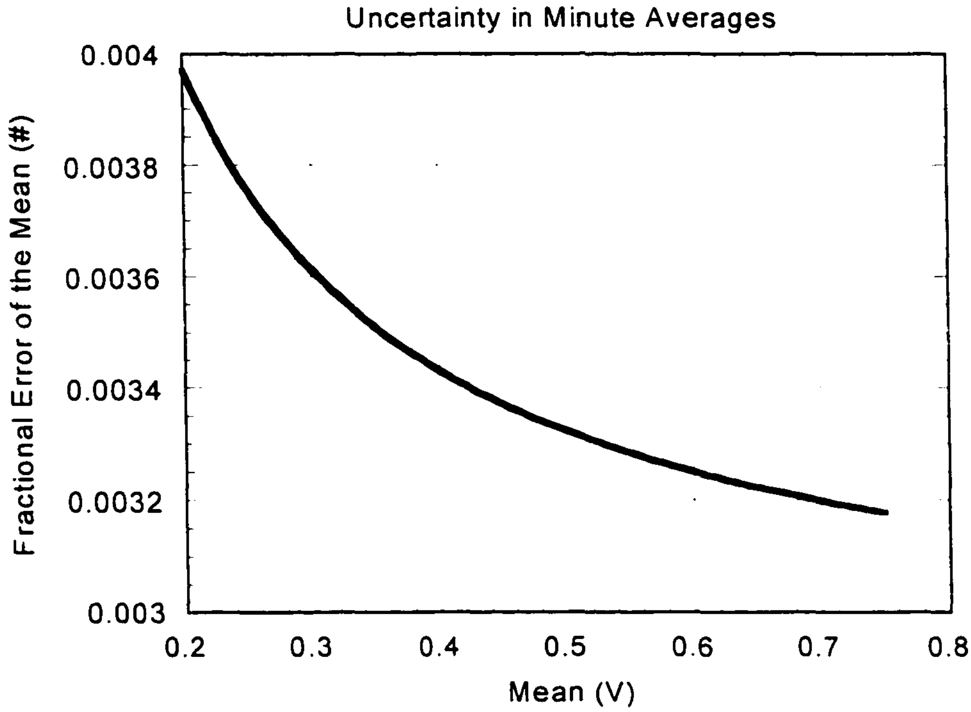

Fig.7.8 The linear dependence of the standard deviation of the frequency monitor measurements on signal level translates to an inverse relationship between the fractional error of a mean of such measurements and the signal level.

Due to the large power variations of the pulsed laser due to both the incomplete conversion of the CW seed light by the acousto-optic modulator system and laser power sag over long lidar operation times, the mean value used for a realistic error analysis should not be that of an idealized laboratory experiment. The typical minute showed in Figure 7.6 exhibits mean measurements of the principal pulses between -0.8 and -0.4 volts. A worst-case error analysis will assume a 50% drop in transmitter power (beyond this the lidar operator would certainly take action to rectify the situation). The result, -0.2 volts, corresponds to a .4% uncertainty (or ±.8mV) in the mean for a one minute average.

Since the experimental theory requires subtraction of the echo measurements to correct for offsets and division of the signal channel by the reference channel for power normalization, it is worthwhile to estimate the uncertainty from these operations. Subtraction warrants adding the absolute errors in quadrature, and division adding the fractional errors in quadrature. For offset subtraction ±.8 mV will be used for the principal pulse and ±.2 mV for the small offset, derived from the y-intercept in Figure 7.7. Adding in quadrature yields ±.825 mV or about .41%. We see, therefore, that despite its valuable role in the frequency measurement, the echo measurement offset correction costs little in added error. For signal division two .41% fractional errors are added in quadrature yielding .58% error in the resulting dimensionless ratio. This again is a worst-case value since the two measurements in the ratio are assumed to be entirely uncorrelated. Correlated amplitude variations from the signal and reference channels would contribute less, or, for perfect correlation, not at all.

As discussed above, the one minute hardware integration time represents the smallest averaging period allowed during typical lidar operation. The averages calculated for one minute may themselves be averaged in the data analysis, dividing the one minute uncertainties by the square root of the number of minutes averaged. This process reduces the 0.58% one minute uncertainty to 0.14% for 15 minutes of averaging, and 0.075% for one hour. The impact of these uncertainties on the accuracy of the absolute frequency measurement depends on the sensitivity of the iodine filter to a changing frequency. To determine this sensitivity, and thereby calibrate the frequency measurement, detailed spectroscopy of the iodine transmission peak is required: a discussion of which follows.

7f. I_2 Spectroscopy for Monitor Calibration

The frequency monitor depends on a measurement of the iodine transmission in the region of the lidar's operating frequency accurate to at least the intended accuracy of the frequency measurement. As was discussed in Chapter 6, Doppler shifting of the lidar transmitter occurs at a ratio of about 1.7MHz per 1m/s of radial velocity. Therefore, to correct a frequency shift in the lidar transmitter such that a velocity accuracy of 1m/s is theoretically obtainable, measurement of the iodine transmission curve in absolute frequency to an accuracy of 1.7MHz is necessary. Furthermore, since the frequency of the pulsed laser is itself the source of uncertainty and variation in the lidar output frequency, it can not be used to obtain a transmission spectrum in absolute frequency. Therefore, it is necessary to use the tunable CW laser, the absolute frequency of which can be more accurately determined, in order to characterize the iodine filter. Once the CW spectrum of the iodine transmission is obtained to desired accuracy, it must be adjusted to account for the much wider bandwidth of the pulsed laser, resulting in an experimentally based, simulated transmission spectrum accurately applicable to the pulsed laser. This iodine spectrum, referred to as the 'pulsed transmission function' contains the necessary information to calibrate the frequency monitor for correction of the lidar output.

CW Spectroscopy

The optical set-up for the CW iodine transmission function measurement is the same as for the pulsed frequency measurement shown in Figure 7.2 above, except that the fiber-collected pulsed laser light is replaced with the freely propagating CW beam from the ring dye laser. Because the CW laser has no short pulses to acquire, the detection electronics are simplified greatly from the pulsed detection chain in Figure 7.5. Instead of the fast, AC coupled, reverse-biased diodes used for the signal and reference channels in the pulsed measurement, slower, DC coupled diodes are used, the output signals of which are amplified and sent directly to the analog to digital converter. Since the iodine transmission spectrum changes with temperature as discussed above, the CW transmission spectra are conducted at the intended operating temperature for the frequency monitor. and the temperature is monitored for the duration of the CW measurement. As a compromise between measurement sensitivity, which increases with temperature, and both temperature control stability and the iodine transmissions sensitivity to temperature variation, 78°C was chosen as the operating temperature of the frequency monitoring cell.

Because the ring dye laser scans in frequency by mechanically changing its effective cavity length in response to an externally applied voltage, the frequency response to a given change in voltage is not always constant. Consequently, additional frequency scaling is required if all parts of the 2GHz iodine transmission spectrum are to be accurately located to within about 1MHz. For this purpose two additional data channels are added to the CW transmission function experiment, the temperature stabilized, 750MHz FSR interferometer (Burleigh CFT-100-VIS1) for relative scaling of the laser scan rate, and the sodium Doppler-free spectrum for secondary scan rate scaling and absolute frequency location. Figure 7.9 shows an example of a typical data scan.

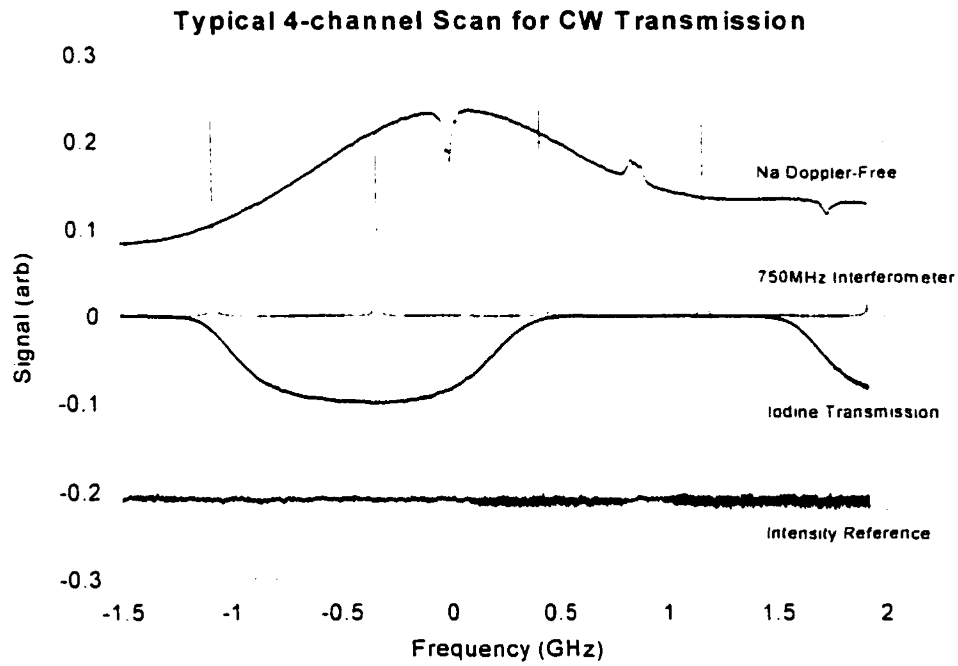

Fig.7.9 A typical 4-channel scan for determining a frequency-scaled iodine transmission function. The interferometer fringes and sodium Doppler-free spectrum are used for frequency scaling. The transmission function appears inverted due to the response of the photodiode used.

All four channels, the signal diode, the reference diode, the interferometer, and the Doppler-free spectrum, are collected simultaneously through the analog to digital converter to be displayed and saved to disk by custom LabView software. Additional software is then used to read the data and interpret the interferometer fringes. Each fringe is located and fitted to a Gaussian profile, from which the centroid is determined. Though the interferometer is temperature stabilized, some drift with temperature, sometimes as large as 7MHz per minute, is still observed. In order to minimize the effect of this drift as well as random inconsistency of the laser scan rate, an ensemble of scans are taken, each requiring only a few minutes to complete. Each is independently frequency scaled from the interferometer as described, then the Na Doppler-free features recorded for each scan are referenced against the their theoretical locations.

The Na Doppler free spectrum exhibits three principal spectral features over about 3GHz: the D_{2a}, crossover, and D_{2b}, relatively located at 0, 839.2MHz and 1719.4MHz respectively. The first is used as an absolute frequency reference, while the additional two serve to check and, if necessary, correct slight errors in the interferometer scaling. A total of ten or more four-channel scans are taken, each is frequency scaled using the custom interferometer scaling software, and the zero-frequency point of each is set to the largest hyperfine peak of the D_{2a} transition. If the two remaining Doppler-free features do not fall within a few megahertz of their theoretical value, that scan is discarded. If the crossover peak is within a few megahertz, the entire scan is multiplied by a scaling factor to correct it to exactly its proper value. Once this is done the D_{2b} peak is checked for agreement with its theoretical value, and the scan is discarded if reasonable agreement is not observed. Once this process has been repeated for all scans, the undiscarded, well-scaled iodine spectra are divided by the reference channel spectra to eliminate the effects of laser power variation which may have occurred during the course of the scan. The resulting power-normalized spectra are summed and the result normalized to unity area.

Following this process it is estimated that every point on an individual iodine transmission spectra is known in frequency to about 5MHz, and assuming that seven spectra are added together and their errors are uncorrelated, that a point on that averaged spectrum has an uncertainty of 5MHz divided by the square root of seven, or about 1.9MHz. The error in the voltage measurements is only hundredths of one percent due to the power normalizing process, the 16-bit resolution of the analog to digital converter, and seven scan average. The uncertainty in the value at any point on the spectrum is therefore dominated by its frequency uncertainty. The resulting CW iodine transmission spectrum is shown in Figure 7.10.

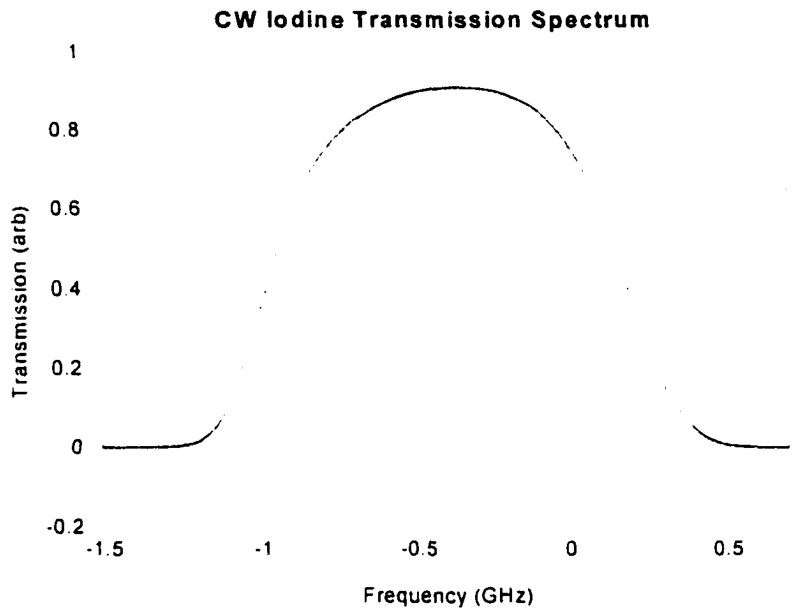

Fig.7.10 The measured CW iodine transmission spectrum resulting from multiple 4-channel scans and frequency scaling procedures. A correlation with the pulsed laser lineshape must still be performed to simulate the transmission function used in the pulsed frequency measurement.

Pulsed Spectroscopy

Due to the pulse bandwidth of the lidar output, the CW iodine transmission spectrum does not correspond exactly with the transmission function for the pulsed beam. To account for frequencies in the wings of the pulsed lineshape being simultaneously transmitted through different parts of the iodine transmission spectrum, it is necessary to determine a pulsed transmission spectrum, the frequency scale of which is measured relative to the centroid frequency of the transmitted lineshape. The shape of the pulsed transmission spectrum is the result of a correlation integration of the CW transmission function and the pulsed line shape.

The pulsed line shape is measured using the 750MHz interferometer and detection electronics identical to those used by the frequency monitor (Figure 7.5). The pulsed lidar output beam is split for signal and reference channels, the signal beam is directed through the interferometer and both are detected by reverse-biased, fast diodes. As in the CW lineshape measurement, the temperature-stabilized interferometer fringe provides the frequency discrimination while the laser is scanned across it. In this measurement, however, an absolute frequency reference is not required since the frequencies of the pulsed lineshape are referenced only to its centroid; but the Na Doppler-free spectrum is still simultaneously measured to provide a check of the frequency scaling. The process for frequency scaling, summing, and normalizing is the same as that described above for the CW iodine measurement. The final result of the pulsed lineshape measurement is shown in Figure 7.11.

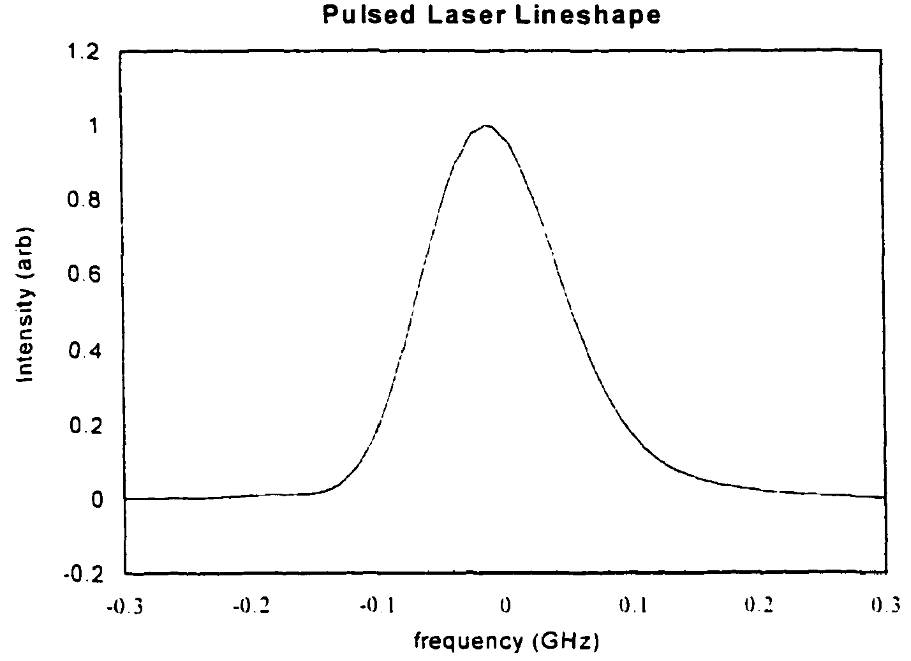

Fig.7.11 The pulsed lineshape which must be correlated with the measured CW iodine transmission function to calibrate the pulsed frequency measurement.

The accuracy of the frequency scaling for the pulsed lineshape measurement is considered to be the same as for the CW measurement, or good to about 2MHz. The voltage measurement, while intrinsically more noisy than the CW measurement due to the pulsed electronics, can be made to an accuracy of a few tenths of one percent with conservative use of smoothing. It should be noted that despite its asymmetry and variability in absolute frequency, the pulsed lineshape seems very robust in width and shape, and measurements separated by years routinely measure a very similar lineshape. This detail is important since the pulsed transmission function of the iodine filter, and therefore the frequency monitor measurement would change if the averaged pulse lineshape changed significantly.

The correlation of the CW iodine transmission function (Figure 7.10) and the pulsed lineshape (Figure 7.11) results in the pulsed transmission function shown in Figure 7.12.

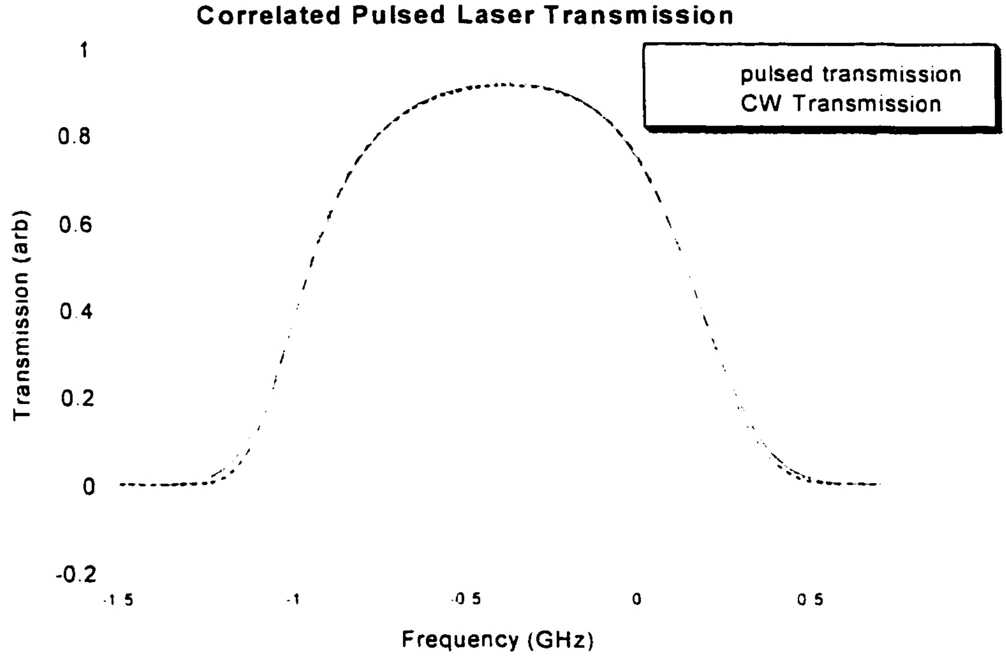

Fig.7.12 The result of the correlation of the measured CW transmission spectrum and the laser lineshape (solid) and the uncorrelated transmission spectrum (dashed). A small overall effect is observed.

Since the two correlated functions have widths which differ by almost a factor of ten, the overall width of the pulsed transmission function differs little from its CW counterpart (note: the FWHM's of correlated Gaussian functions add in quadrature), however, the details of the shape, especially in the low-transmission wings, does change noticeably.

Since the correlation function performs a weighted average of points across the associated functions, quickly varying, statistically uncorrelated noise is smoothed away. In this sense the correlation function performs a running average. Since the noise associated with the actual voltage measurements of the two functions is rapidly varying and statistically uncorrelated from point to point, its effect in the correlated output is greatly reduced. Since the randomness of the laser scanning over the frequency ranges of the functions is difficult to determine, its exact effect in the correlated output is unclear. Cases in which systematic change in the frequency scaling occur, such as the possibility of a monotonically increasing laser scan rate between interferometer fringes, may exhibit the same type of error in the correlated output, while the effect of a sinusoidally varying scan rate would likely be reduced. A reasonable degree of pessimism suggests that the frequency uncertainty in the correlated output is similar to that of each of the constituent functions.

The frequency scale of the pulsed transmission function in Figure 7.11 represents the centroid frequency location of the transmitted laser pulse referenced against the principal hyperfine peak of the Na D_{2a} Lamb dip. This spectrum may be shifted in frequency by + and - 630MHz to illustrate the transmission at the shifted operating frequencies referenced to the location of the centroid of the unshifted operating frequency (Figure 7.12).

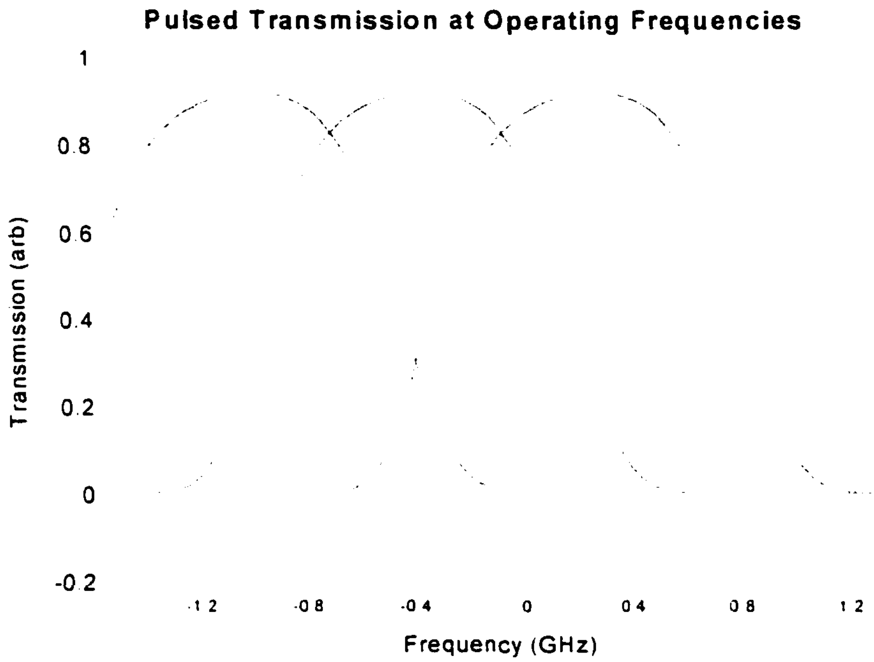

Fig.7.13 Shifting the pulsed transmission function by ±630MHz, the relative transmissions at each of the lidar's three operating frequencies are superimposed at 0MHz. The effect of an overall shift in the three operating frequencies on the relative transmissions can be observed on the graph at a frequency location equal to the overall shift frequency.

This again makes use of the assumption that all three operating frequencies are shifted by amplifier chirp in exactly the same way, so that the zero of the frequency scale in Figure 7.13 means that the centroid of the unshifted operating frequency is exactly aligned with the D_{2a} peak and the centers of mass of the two additional operating frequencies are located + and - 630MHz from the D_{2a} peak respectively. Figure 7.14 shows the relative transmission of pulses at each of the three operating frequencies over a smaller frequency scale, at which the relative change in transmission corresponding to a shift of the centers of mass by several megahertz is apparent.

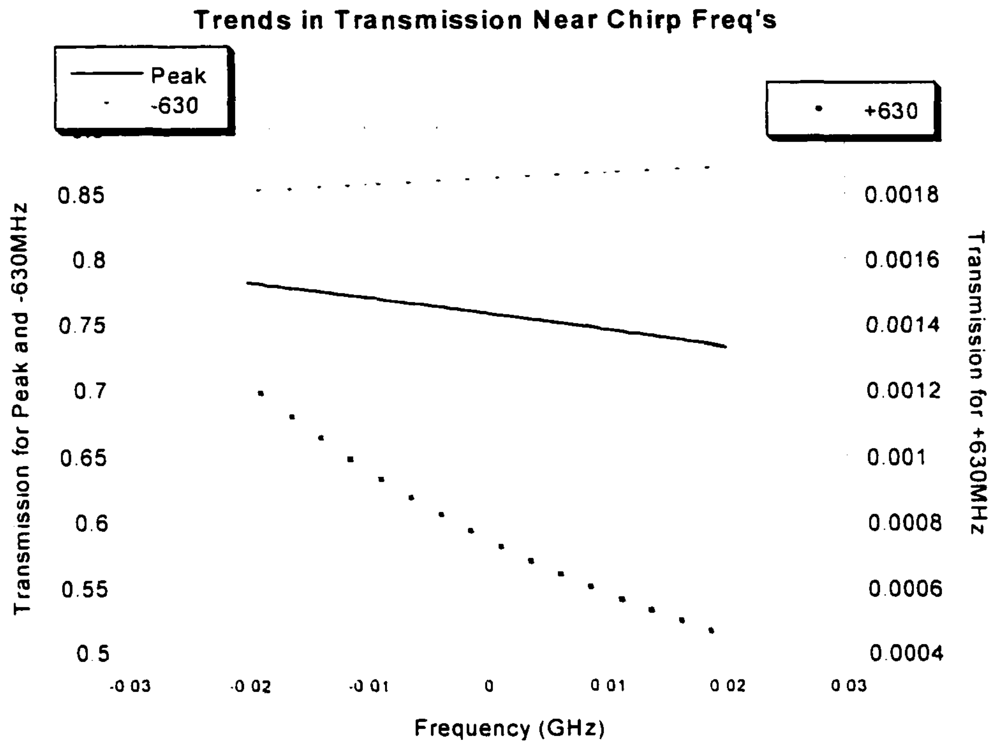

Fig.7.14 A detailed view of the change in relative transmissions at the lidar's three operating frequency with a small shift in overall frequency of the three. Due to its smallness, the -630MHz transmission is plotted on an expanded scale, still relative to the other two.

A difference of several percent is observed in the high-transmission frequencies for a shift of just 10MHz. This sensitivity is the basis of the frequency monitoring measurement. A ratio generated from the curves in Figure 7.13 may be formed which best capitalizes on the information they provide, and this ratio becomes the principal experimental result of the 7.12 values measured on a minute by minute basis by the monitor. The selection and characterization the best ratio to form from these 7.12 values follows directly from Figure 7.13.

7g. The Experimental Ratio and Calibration Curve

The information the lidar frequency monitor uses to determine a shift in the centroid frequency of the pulsed transmitter output is held entirely in the 12 values stored by the monitor. However, some analysis is necessary to deduce the actual centroid frequency from the numerical data. To this end, it is necessary to select a ratio which best capitalizes on the information stored by the monitor, and generate a calibration curve based on the selected ratio and the theoretical iodine response illustrated in Figure 7.13.

Choosing an Experimental Ratio

The 12 values stored by the monitor experiment include the echo measurements for DC offset corrections and measurements from both the signal and the reference channels.

Since the offsets must be subtracted and power normalized before a meaningful ratio can be formed, the 12 values reduce to just 3 values to be used in the ratio: the echo-subtracted, reference-normalized transmission at each of the three operating frequencies. The only requirement on the form of the experimental ratio is that it be, in fact, a ratio: since this removes the dependence of its result on the absolute gain of the detector chain. The principal criteria for selecting the best form of the ratio are propagation of experimental error, sensitivity to a overall shift in frequency, and an insensitivity to additional uncertainties intrinsic to the experiment.

The principal sensitivity to frequency is exhibited by the high-transmission frequencies at the D_{2a} and -630MHz frequencies, as is clear from Figure 7.14. The simplest ratio suggested by this observation is to divide the -630MHz frequency by the peak frequency, denoted in Equation 7.7.

$$R_1 = \frac{NEG}{PK}$$

where

$$NEG \equiv \frac{[\text{signal at -630MHz}] - \textit{offset}}{[\text{reference at -630MHz}] - \textit{offset}}$$

$$PK \equiv \frac{[\text{signal at } D_{2a}] - \textit{offset}}{[\text{reference at } D_{2a}] - \textit{offset}} \qquad \text{Eq.7.7}$$

$$POS \equiv \frac{[\text{signal at +630MHz}] - \textit{offset}}{[\text{reference at +630MHz}] - \textit{offset}}$$

No experimental effect or loss of generality results if unity, or any constant, is subtracted from the ratio and Equation 7.7 is rewritten as Equation 7.8.

$$\frac{NEG}{PK} - 1 = \frac{NEG - PK}{PK} \qquad \text{Eq7.8}$$

On the basis of our selection criteria, the ratios in Equations 7.7 and 7.8 are equivalent: Equation 7.8 being preferable due to its tendency towards zero rather than any other constant. The form of Equation 7.8 also elucidates its principal error sources and the error propagation through the experiment in general. Noticing that NEG and PK have values equal to within about 15%, their difference, which has an absolute uncertainty added in quadrature about 1.4 times as large as either subtractand, has a fractional error about 10 times larger than either. So, since qualitatively the fractional error of the entire ratio is some sum of the fractional errors of numerator and denominator, it is dominated by the numerator. This insight suggests that the form of the denominator is relatively unimportant as long as it is large compared with the numerator, and that the numerator be similar to that in Equation 7.8 to avoid introduction of further error into the ratio. Equation 7.9 lists some likely ratios for consideration.

$$
\begin{aligned}
Ratio1 &= \frac{NEG - PK}{PK} \\
Ratio2 &= \frac{NEG - PK - POS}{NEG + PK + POS} \\
Ratio3 &= \frac{NEG - PK - POS}{PK + POS} \\
Ratio4 &= \frac{NEG - PK}{NEG + PK} \\
Ratio5 &= \frac{NEG - PK}{NEG + PK - 2POS}
\end{aligned}
\qquad \text{Eq.7.9}
$$

Two qualitative comments about the uncertainties of these ratios are worth noting. First, since both the value and absolute uncertainty of POS is much smaller than those of PK and NEG, its effect on the uncertainty of larger quantities, such as in the denominators of ratios 2, 3, and 5 is reduced, while the error it introduces to smaller values, such as the numerators of ratios 2 and 3 may be non-negligible. Second, a larger denominator is advantageous from a fractional uncertainty standpoint since the errors of uncorrelated variables will add in quadrature, yielding a smaller fractional error for a sum than for either summand due to averaging of random noise. Therefore, when the uncertainty is more rigorously calculated, ratio 1 is at a disadvantage due to its small denominator, and ratios 2 and 3 are more uncertain due to the additional term in the small numerator. These effects, however, account for a difference in the error in the measured ratios of only about .2%, and therefore these disadvantages may be balanced by other advantages for these ratios.

Considering the sensitivity of the ratios to a shift in frequency of the three measurements, it is important to recognize the differences shown in Figure 7.13. The normalized transmission measurement by the +630MHz lidar operating frequency is about 700 times smaller than the other transmission curves. Therefore, the contribution of POS to a frequency sensitive sum or difference is negligible. Once the appearance of POS is removed from the ratios for the purposes of analyzing the sensitivity, each of the five ratios can be reduced to some linear combination of integers and the quantity NEG/PK or its reciprocal. Since this is the basic quantity which carries the frequency shift information, any transformation in the ratio must eventually be decoded by the calibration curve only to yield the same information. For this reason, under a more detailed analysis of the effective sensitivity to frequency, all five of the ratios in Equation 7.9 are found to be essentially identical.

Unlike frequency shift, the effect of a changing temperature on the value of POS can be dramatic, as was illustrated before in Figure 7.3. Since the temperature is controlled only to some acceptable level, and possibly subject to change under conditions of unusual environmental variation, a robustness of the measurement in the case of temperature change is a desired quality of the experimental ratio. From the temperature dependence curves in Figure 7.4, the approximate effect of changing temperature on each of the ratios under consideration can be determined, as shown in Figure 7.15.

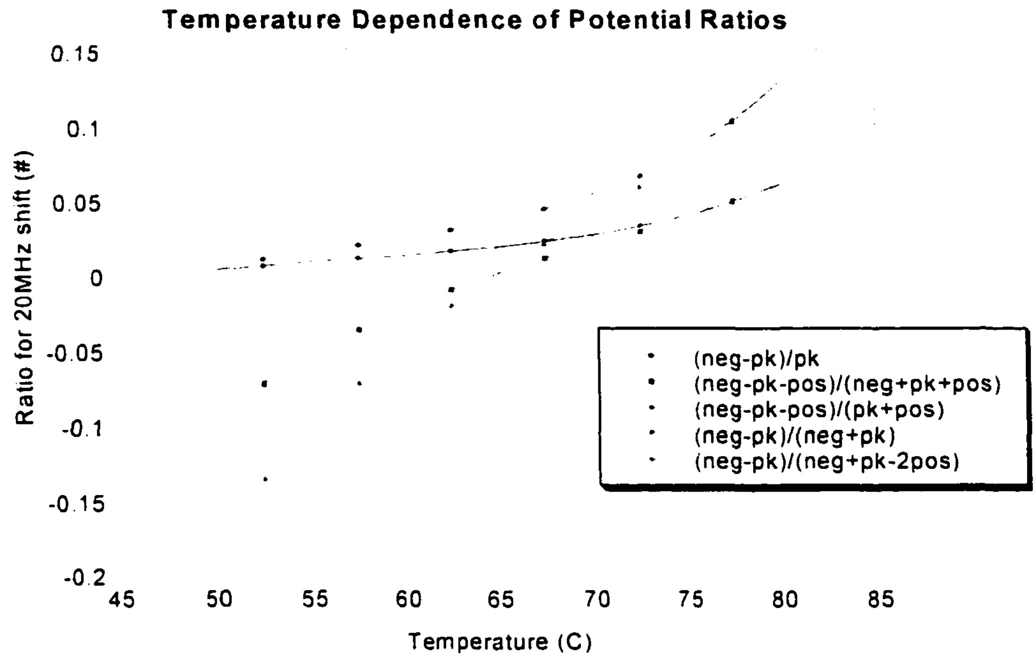

Fig.7.15 The sensitivities of the four potential ratios (listed in order) to a changing cell temperature. Two families of slopes emerge at the intended operating temperature of 78°C.

In the vicinity of the operating temperature of 78°C, the ratios fall into two distinct families, ratios 1 and 3 having a higher slope through all temperatures than the other three ratios. Of the remaining three, ratio 5 has the smallest slope, but is almost identical to ratio 4, as might be expected since they differ only by a factor of 2POS in a large denominator. Fitting the ratio 5 curve to determine the fractional change with a small change in temperature around 78°C, a value of about 6% per 1°C is calculated. For a typical one minute averaged cell temperature stability better than ±.005°C the temperature dependence represents a small contribution to the total error. However, atypical temperature controller performance could introduce fluctuations in the frequency monitor measurement which, when propagated through the lidar data analysis, would be misinterpreted as velocity fluctuations or waves in the atmosphere. Therefore, insensitivity to cell temperature is valuable to the overall reliability of the frequency monitor.

Another experimental parameter which could compromise the accuracy of the frequency monitor is a light-induced DC offset. Recall from the experimental theory that a varying offset, such as that which would be produced in the gated integrators, is corrected in the monitor experiment by subtracting the echo signal which is taken when no laser light is present in the monitor. For offsets which change no faster than the timescale of the pulse repetition rate, or 20Hz, this is an accurate correction. However, certain laser-induced phenomena, such as amplified spontaneous emission (ASE), imperfect YAG laser seeding, and detector diode ringing could create an added offset to the ratio quantities NEG. PK. and POS which would last only as long as the laser pulse itself. and therefore would not be corrected by the echo subtraction technique. It is informative. then. to consider the effect on each prospective ratio of a added constant to each of its constituent values. Ratios 1. 4. and 5 with a simple difference in the numerator cancel an offset in the numerator whereas ratios 2 and 3 have a remaining negative contribution. and ratios without a difference in the denominator. which includes all but ratio 5. would exhibit a positive offset in the denominator. The overall effect on all ratios except ratio 5 would be a low result with ratio 2 the lowest. followed by ratios 3. 4. 1. and finally the unaffected 5. This. of course. was the initial rationale for including the 2POS term in the denominator of ratio 5. Since the usually small value of POS contributes little to the frequency sensitivity. it is used instead to correct for reduced spectral purity and other offset inducing phenomena.

The Calibration Curve

Given the criteria for selection of an experimental ratio. it seems clear that ratio 5 is the best of those considered. From the pulsed transmission curves in Figure 7.13 it is easy to use ratio 5 to create a calibration curve which references the result of the measured ratio to the implied frequency chirp of the lidar output laser. This calibration curve is shown in Figure 7.16.

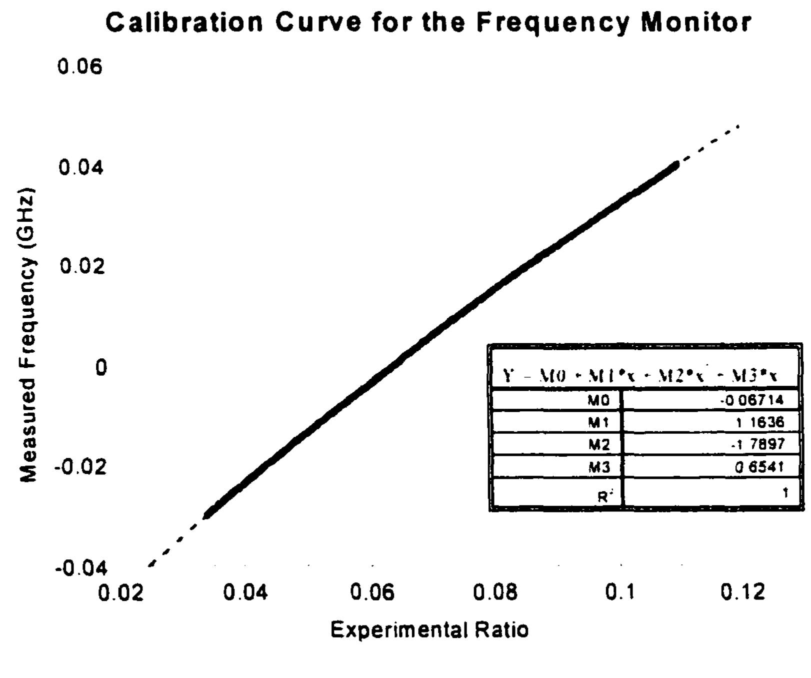

Fig.7.16 The calibration curve for mapping the measured experimental ratio of (NEG-PK)/(NEG+PK-2POS) into a the absolute frequency of the pulsed laser lineshape's centroid. The frequency is measured relative to the Doppler-free locking peak.

The frequency in Figure 7.16 is again referenced to the Doppler-free D_{2a} Lamb dip at which the CW seed laser is intended to be frequency locked. This is not to be confused with a calibration referenced to the CW seed beam. In the 3-frequency lidar monitor if the seed beam wanders from its lock-point at the D_{2a} transition, the frequency monitor will track the resulting frequency movement in the pulsed output of the lidar. The uncertainty in the calibration curve is, as discussed, dominated by the frequency scaling error in the scanning measurements which produced it. Since each point on the calibration curve is the result of all three operating frequencies, if we assume that frequency scaling errors were uncorrelated over the 630MHz separating the operating frequencies, the resulting frequency uncertainty of the calibration curve in Figure 7.16 is about 1MHz.

Since the calibration curve is slowly varying over the range of frequencies at which the monitor will operate, it may be fit to a third-order polynomial with essentially no loss of accuracy. This fit is shown in Figure 7.16, and the resulting equation is reproduced in Equation 7.10.

$$\nu(R) = -0.06714 + 1.1636R - 1.7897R^2 + 0.6541R^3$$

where

$$\nu = \text{offet frequency in GHz} \qquad \text{Eq.7.10}$$

$$R = \frac{NEG - PK}{NEG + PK - 2POS}$$

Equation 7.10, then, is the key to the frequency measurement and the net result of all of the spectroscopic effort outlined in this chapter. To correct the 3-frequency lidar data for any frequency shift, the frequency monitor acquires 12 measured values, each averaged over one minute, performs echo subtraction and power normalization to produce three principal transmission measurements, NEG, PK, and POS. These three values are divided according to Equation 7.9, ratio 5, the result is finally interpreted by Equation 7.10 to provide a measurement of the centroid of the output laser pulse in absolute frequency calculated relative to the D_{2a} Lamb Dip. This process, together with the error associated with each step, and the total uncertainty is outlined in Figure 7.17 and Table 7.1.

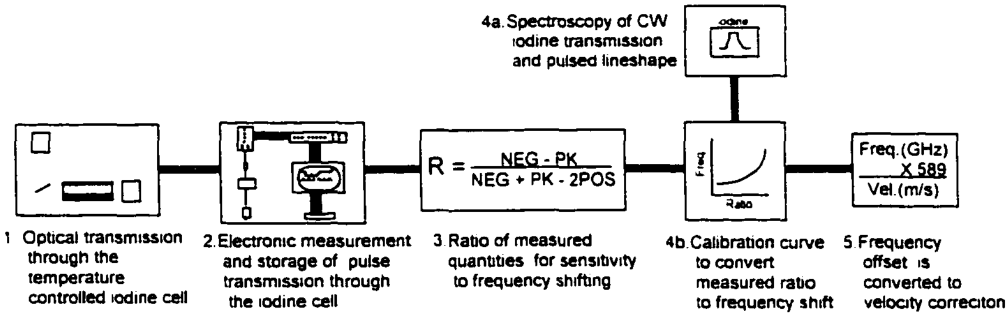

Fig.7.17 A block diagram outlining the processes by which the calibration of the frequency monitor is achieved. Each step carries with it some level of experimental uncertainty which may be combined to determine the overall expected accuracy of the frequency measurement.

Table 7.1. Uncertainty Propagation in the Frequency Monitor

Averaging	1. Optical	2. Electronic	3. Ratio	4. Calib.	5. Velocity	Notes
1 Minute	Cell Temp. ±.01°C	Noise and Normalizing ±.58%	Propagation ±7%	±5 MHz Temp. effect: <±.1MHz Scaling: ±2MHz (Systematic)	±3 m/s Scaling: ±1 m/s (Systematic)	Smallest integrating time allowed in normal lidar operation
15 Minutes	Cell Temp. <±.01°C	Noise and Normalizing ±.15%	Propagation ±1.8%	±1.3 MHz Temp. effect: <±.1 MHz Scaling: ±2MHz (Systematic)	±.8 m/s Scaling: ±1 m/s (Systematic)	Systematic offset correctable from lidar data
1 Hour	Cell Temp. <±.01°C	Noise and Normalizing ±.075%	Propagation ±.9%	±.6 MHz Temp. effect: <±.1 MHz Scaling : ±2 MHz (Systematic)	±.4 m/s Scaling: ±1 m/s (Systematic)	Averaging time for typical data analysis

Table7.1

Once the frequency monitor has acquired and stored the measurements necessary for the frequency measurement, any number of minutes may be averaged in the data analysis to provide a frequency centroid value for that period with least possible uncertainty. The frequency centroid measurement is easily converted to a velocity correction which can be applied to the lidar velocity measurement averaged over the same period. The results of such frequency-monitored lidar data, and further details concerning the integration of the frequency monitor into the routine lidar system, are presented in Chapter 8 of this thesis.

Chapter 8: Frequency Monitored Data

The operation of the 3-frequency lidar with the concurrently operating frequency monitor is essentially the same as without the frequency monitor. The inversion of the lidar equation to yield atmospheric parameters from acquired, range-resolved photocount profiles proceeds as though the frequency monitor was not present. As has been discussed, the velocity obtained without the frequency monitor is the relevant velocity for generation of the temperature and velocity dependent backscatter cross-sections needed in the data analysis. The purpose, then, of the frequency monitor is to determine what part of the measured velocity arises from motion of atmospheric sodium atoms, and what part arises from frequency shifts inherent to the lidar transmitter, the first quantity being of interest in understanding the structure of the middle atmosphere, while the second quantity is of interest principally only to lidar operators.

The correction of the measured velocity is determined from the frequency monitor data which is averaged for the same interval as the atmospheric velocity measurement. Once the velocity profile for a given time interval is determined, the velocity equivalent to the frequency shift measured by the monitor can be applied, added equally to all altitudes of the profile.

This chapter presents the frequency corrected lidar data, first in the context of a night of data acquisition with unusual laser operation conceived to test the performance of the lidar velocity measurement and its correction by the frequency monitor, second in context of routine lidar data acquisition, and finally in the context of a two-beam, simultaneous measurement of vertical and horizontal velocity in the mesopause region. A discussion of the accuracy of the corrected velocity and the implications of both the weaknesses and strengths of the final lidar system follows.

8a. The Hyperfine Locking Test

Since the lidar frequency monitor is designed to measure the location of the centroid of the pulsed laser output lineshape, and since the uncorrected velocity measured in the atmosphere should respond linearly to a change in the output centroid according to the ratio .6m/s per megahertz, a change of the locking frequency of the lidar transmitter by tens of megahertz should be noticeable in both the atmospheric data and the frequency monitor data intended to correct it. For purposes of testing the lidar measurement and its frequency correction, then, it is fortunate that there exist additional spectral features in the Doppler-free sodium fluorescence spectrum of well-defined absolute frequencies at which the laser may be locked to check for the expected response in the obtained data. Figure 8.1 shows the three peaks to which the laser may be locked.

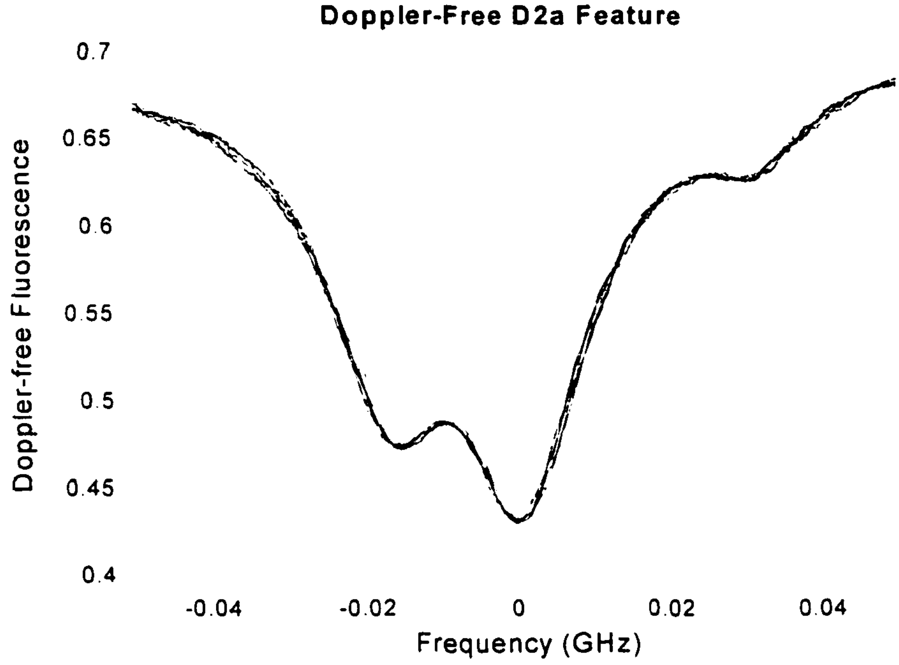

Fig.8.1 A detailed plot of the sodium D_{2a} Doppler-free peak from five independently scaled CW scans. Locking at each of the three peaks which are located relatively at about -16, 0, and 30MHz, is possible and the shift in frequency of the lidar transmitter should be obvious in both the atmospheric velocity measurement and the frequency monitor experiment.

The Doppler-free spectrum near the D_{2a} peak exhibits the overlapping hyperfine transitions and associated cross-over features, with the largest peak at 0MHz representing the standard peak used for locking the lidar for routine data collection. The smaller peaks are located –15.6MHz and +29.9 MHz from the central peak. Locking the laser alternately on each of the three peaks, velocity profiles from the sodium layer should exhibit their typical negative velocity offset of around 10m/s while locked on the central peak, a velocity 9.1m/s less while locked to the left peak and a velocity 17.6m/s greater while locked to the right peak. Similarly, if the frequency monitor is correctly measuring the centroid frequency of the lidar output, it should provide correction velocities of equal and opposite magnitude to consistently correct the atmospheric velocity measurements for the intentional lidar transmitter frequency shifts.

To achieve high statistical certainty, the lidar transmitter was locked to each frequency for one hour before changing for a total of seven hours corresponding to the central peak, the left peak, the central peak, the right peak, central, left, and right again. To resolve the changes, the data was analyzed at 12-minute time intervals. Figure 8.2 shows the uncorrected layer-averaged vertical velocity and vertical velocity at the Na density centroid representing low-noise atmospheric velocity measurements, the correction velocities from the lidar frequency monitor, and the corrected layer-averaged velocities and centroid velocities, which should represent the sum of the uncorrected and correction velocities.

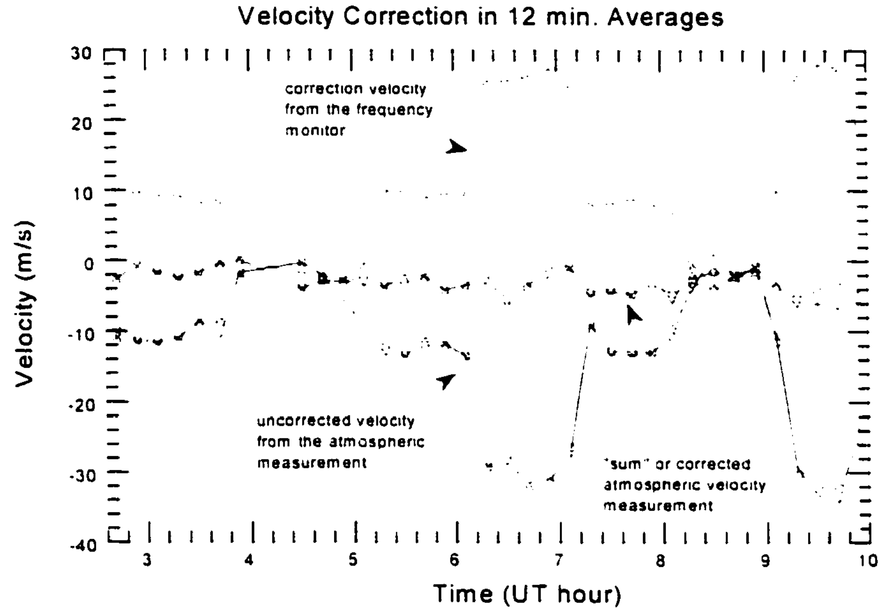

Fig.8.2 12-minute averages from the night of 9089 (March 30, 1999) during which the laser was locked on a different hyperfine peak for each successive hour. Both the layer (circles) and centroid (x's) velocity values are shown for the best statistical significance. The expected variation in atmospheric velocity and correction velocity from the frequency monitor are observed, but overall correction is not complete, as shown by the corrected layer and centroid velocities which exhibit a small residual velocity after the correction of induced shifts.

It is clear from Figure 8.2 that both the atmospheric measurement and the frequency monitoring measurement track the intentional changes in output frequency, and both show the correct magnitude of shifting to within the limits of statistical noise and uncertainty in the laser locking. The sum of the atmospheric measurement and the correction velocity yields the corrected atmospheric velocity measurement, shown for both the layer and density centroid. Both measurements exhibit a residual velocity after correction of about 1m/s at the start of the night and growing to about 2.5m/s at the end of the night. The observed residual velocity of about 2m/s may be compared to the typical uncertainty expected from the frequency monitor, about 0.8m/s random with a possible additional 1m/s systematic offset, and from the atmospheric measurements due to photon noise, about 1m/s.

The data shown in Figure 8.2 may also be analyzed in hourly profiles and displayed in their range-resolved format.

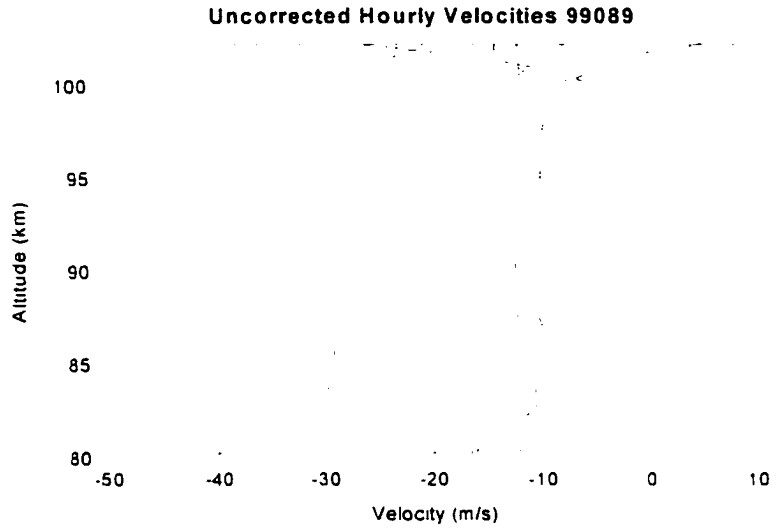

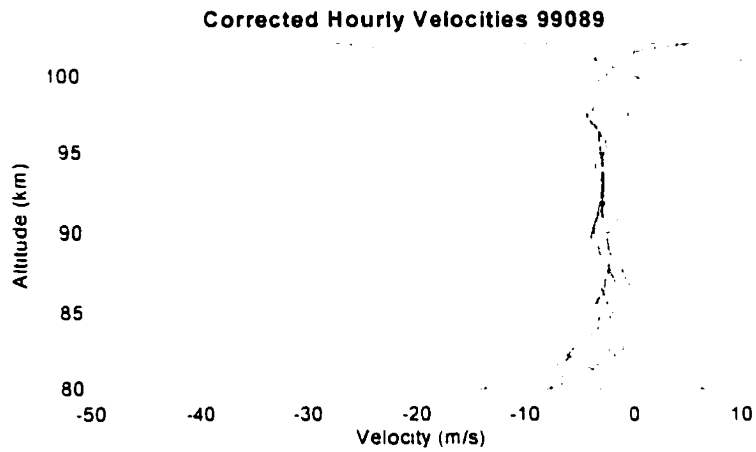

Fig.8.3 Altitude-resolved 1-hour averaged profiles from the hyperfine experiment night 9089. The three locking frequencies are clearly visible, and the relative correction eliminates the applied frequency shifting.

Figure 8.3 shows the same experiment averaged to one-hour profiles. At top, the three hyperfine locking points are clearly distinguishable, while the corrected data at bottom show good grouping of the profiles, as would be expected from near-zero vertical velocities, but a clearly visible offset from the expected value of a few cm/s. The implication of these plots is that the sensitivities of both the atmospheric velocity measurement and the frequency monitor are sufficient to measure velocities on the order of a few meters per second, and that the responses of both measurements are properly calibrated to correctly deduce a range of velocities from 0 to 30m/s. However, even the corrected measurement exhibits a physically unlikely assessment of the vertical wind in the mesopause region. The most likely explanation for this result is that a possibly time-varying offset is present in one or both measurements, giving rise to imperfect measurement of the atmospheric velocity, which is assumed to be very close to 0m/s.

An additional check of the 3-frequency lidar analysis provided by the application of large, controllable frequency shifts in the lidar transmitter frequency is a test of the expected independence of the temperature and density measurements on the locking frequency. Figure 8.4 shows the 12-minute-averaged layer and centroid temperatures, and total sodium abundance on the night of the hyperfine locking experiment, with the frequency monitor's correction velocity measurement superimposed to mark the changes in output frequency.

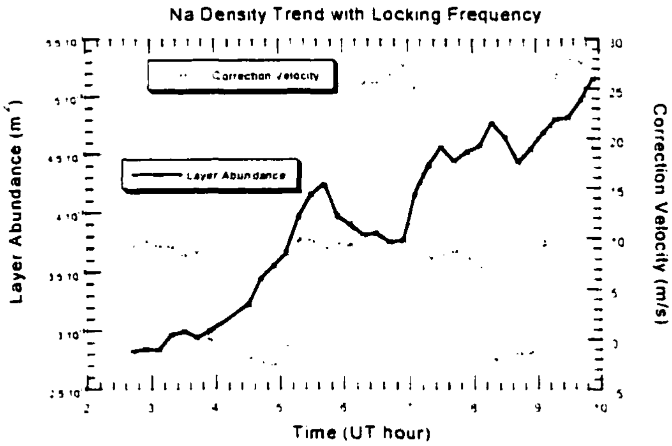

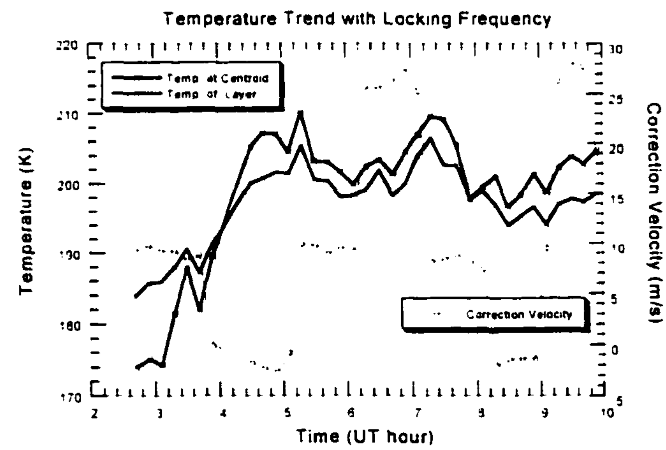

Fig.8.4 Trends of the layer and centroid temperature and total Na abundance are shown with the frequency shifting information from the frequency monitor for the night of 9089, again exhibiting user-imposed frequency offsets. No correlation is observed, affirming the expected independence of the temperature and wind measurements on transmitter frequency offsets.

No correlation between the Na density or temperature measurements is evident with the controlled change in laser output frequency. This result is as expected as the shift in laser frequency will be reflected in the measured line-of-sight velocity, but does not impact the velocity distribution which determines the measurement's temperature sensitivity, or the abundance of atmospheric atoms.

8b. Frequency Monitored Routine Data

Following from the first successful operation of the lidar frequency monitor on the night of March 18, 1999, or the 77th day of 1999, eight nights of routine data collection have been operated in the frequency-monitored configuration of the 3-frequency lidar. On these nights, the lidar transmitter was locked to the largest Doppler-free feature of the sodium D_{2a} transition for the entirety of the night, and therefore a smaller variation of uncorrected atmospheric velocities were observed than on the night of the hyperfine locking test. Figure 8.5 shows the hourly-averaged, vertical velocity profiles of a typical data night before and after frequency-monitor correction (see following page).

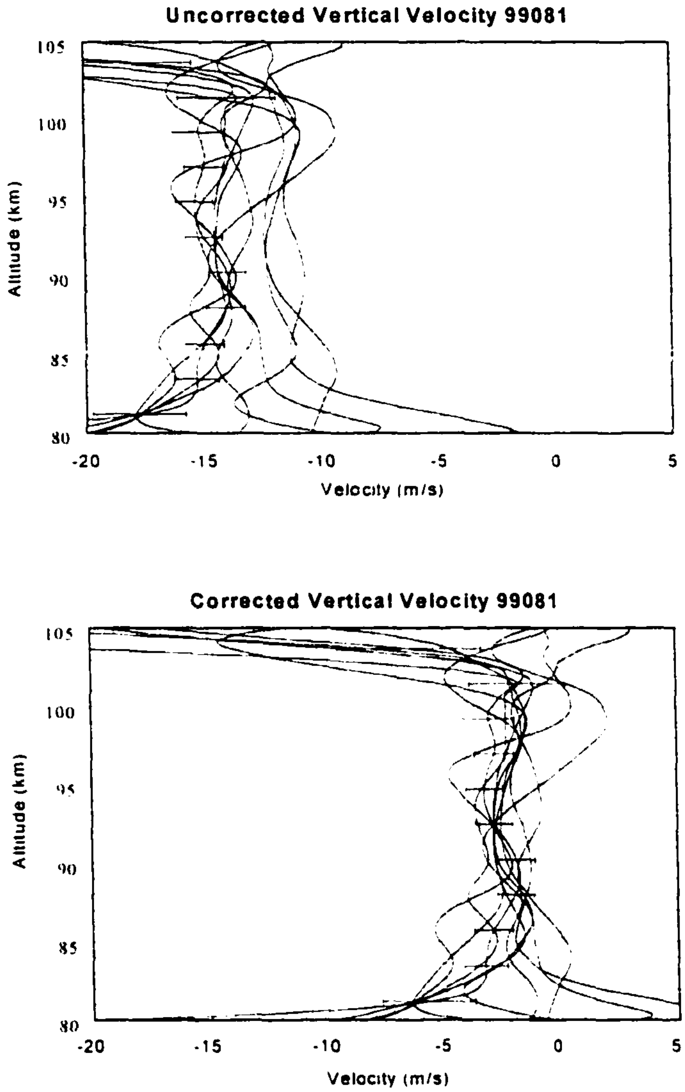

Fig.8.5 Uncorrected (top) and corrected (bottom) hour-averaged profiles from a typical monitored 3-frequency lidar night. Improvement of the variation in velocity over the night is achieved with the frequency monitor.

The night of data shown in Figure 8.5 was taken on the 81st day of 1999, or March 22, and exhibits several features common to the velocity-corrected measurements. First, the grouping of the hourly profiles is seen to improve greatly from the uncorrected (top) to corrected (bottom) data. This suggests that at least one contribution to the instability of the lidar transmitter frequency changes with time over the night and is tracked by the frequency monitor adequately. This may be due to laser locking error in the CW subsystem of the lidar transmitter or possibly a time dependence of the pulsed laser chirp linked to the age of the laser dye used. Second, a considerable amount of wave structure in the vertical profile of the velocity is observed, some on scales larger than the calculated error bars of the measurement. This suggests that either the statistics of the measurement have been calculated over-optimistically, or else there are occasional vertical winds in the mesopause region with meters-per-second amplitudes. The later proposition gains some credibility based on the apparent correlation between waves observed in the velocity profile and waves observed in the temperature profile at the same time. Third, after velocity correction some residual unexpected mean velocity remains in the velocity profiles. This corrected measurement varies in magnitude, but is more often negative than positive. Had the corrected vertical velocity measurement been a non-zero constant for all nights, the obvious explanation would have been an offset present in the frequency scaling of the frequency monitor's calibration curve, which may still be present to some degree. However, the variability of the measurement suggests a more dynamic source. Discounting the possibility that the measurement is an accurate representation of the vertical velocity in the mesopause region, the more likely explanation is that the residual velocity represents a systemic offset which could arise from either the atmospheric measurement or the frequency monitoring measurement. In addition, the analysis details leading to the calibration curves for the simultaneous atmospheric measurements of temperature and wind should be revisited for possible improvement.

To better characterize the unexplained measurement, referred to here as a "residual velocity offset" or simply "offset", the nightly averages of all eight, frequency-corrected data nights are shown in Figure 8.6, first before correction then after (see following page).

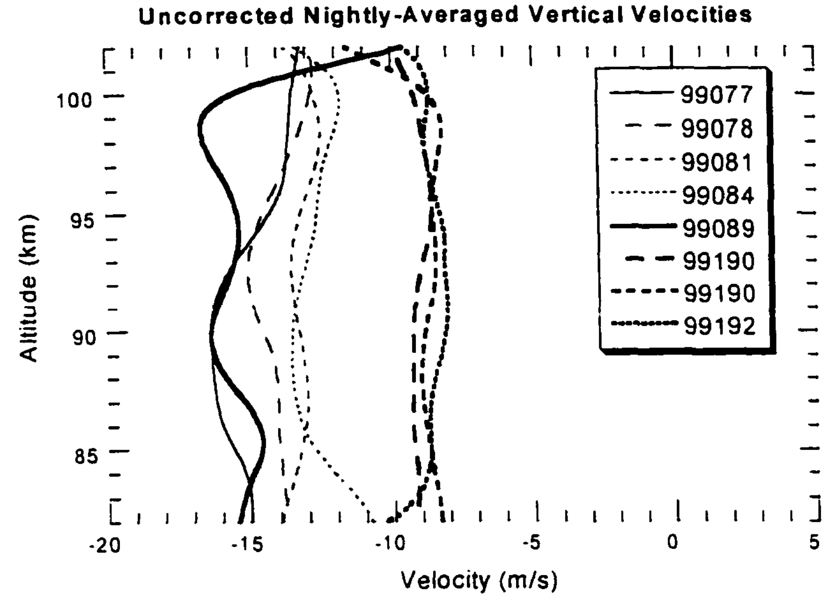

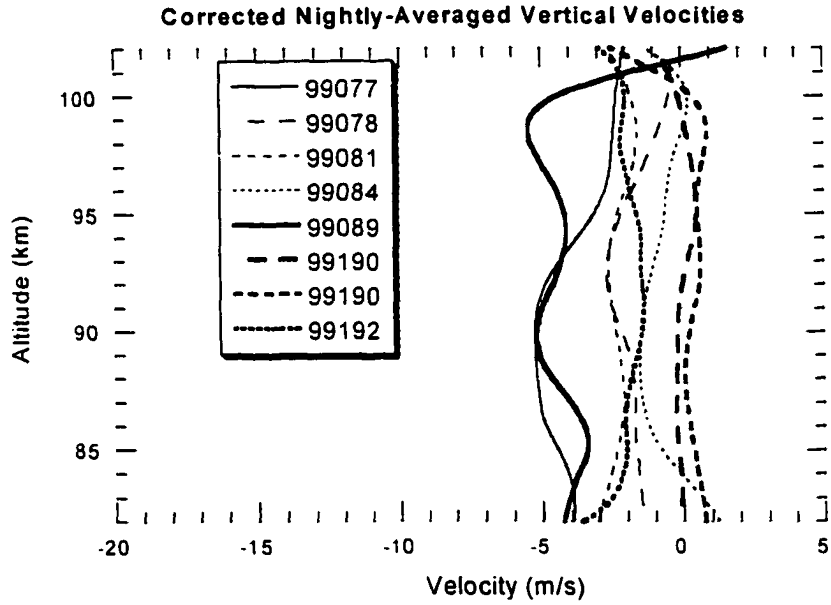

Fig.8.6 Nightly averages for frequency monitored lidar nights. The uncorrected profiles (top) show a larger variation than the corrected profiles (bottom), but a widely variable residual offset still remains. The duplicated 99190 label represents two independent profiles taken simultaneously that night.

The profiles shown in Figure 8.6 show some improvement in grouping after frequency-correction, but not to the same degree as hourly profiles for each night. This suggests that while the frequency of the lidar transmitter does vary over the data collection period, the source of the residual velocity offset changes less over one night than between nights.

This characteristic still allows the possibility of a wide range of sources, both in the lidar velocity measurement, such as laser alignment dependencies, and the frequency monitoring measurement, such as amplifier non-linearities. Regarding Figure 8.6 it is also interesting to note that most of the nights exhibiting the largest negative residual velocities were among the first to be measured, and that the most recent data exhibits residual velocities less than 1m/s. Recent instrument recalibration may have reduced the residual velocity in the frequency-monitored measurements. More nights of data are required to determine this. A number of additional changes not directly related to the 3-frequency lidar measurements, but which could affect the operation of the lidar, have also been made in the lidar laboratory. Additional data is required to assess the night-to-night behavior of the measured residual velocity in the presence of these changes.

8c. The Horizontal Velocity Measurement

Since the vertical wind velocity is nearly zero, its measurement offers an accurate gauge for the sensitivity and performance of the 3-frequency velocity measurement. Due to the relatively small power-aperture product of our lidar, the mean vertical wind is difficult to determine except with long integration and further improvement of the overall accuracy of the velocity measurement. Such measurements could be useful in characterizing ubiquitous gravity waves in the mesopause region. Lidar systems with increased power-aperture statistics could extend the same measurement principle to provide direct experimental determination of heat flux due to gravity waves in the mesopause region (Tao and Gardener, 1995).

Given the current statistical limitations of the CSU lidar system, the measurement of horizontal winds, which are much larger than the vertical winds and the yet unresolved bias of 2m/s, represents a more profitable application. With ongoing upgrades to two-beam and daytime operation, the simultaneous measurement of mesopause temperature, zonal, and meridional winds on a 24-hour basis is in sight, suggesting the possibility of climatology studies of the thermal and dynamic properties of the mesopause region. Such a system could be used to assess the dynamics of gravity and tidal waves with key parameters such as temperature lapse rate, Brunt-Vaisala frequency, and horizontal wind shear (or Richardson Number) determined simultaneously and experimentally.

One principal goal of the lidar frequency monitor was to allow off-vertical lidar measurements to yield accurate horizontal velocities without expending valuable laser power for a vertical measurement with limited atmospheric relevance, but required for correction of lidar transmitter frequency offsets. To determine the performance of the current frequency-monitored lidar in such an application, a 3-frequency lidar experiment with one off-axis beam collecting frequency-monitored temperature, Na density, and radial velocity data and one vertical beam for determining the accuracy of the obtained frequency correction was proposed. With the replacement of the 20Hz pulsed YAG laser used for all previous lidar measurements with a 50Hz model, the additional power necessary to split the beam into two directions, each capable of yielding experimental statistics equivalent to the older system, was achieved, and the experiment was performed on the night of July 10, 1999 (day 192 of 1999).

A successful experiment, then, was one in which a residual frequency offset negligibly effecting the accuracy of the deduced horizontal velocity was determined from the vertical beam, suggesting that it was not necessary for supplementing the frequency monitor with additional frequency correction information, and that in principle, full laser power could be used for off-vertical measurements in the future.

The angle of the off-axis beam was chosen as a compromise between sensitivity to horizontal wind, which increases with the beam angle from vertical, and beam attenuation, which increases with the same angle due to the additional range of lower atmosphere which must be traversed by the laser beam. An angle of 30° was chosen, but due to initial hardware limitations a maximum angle of 25.8°±3° was obtained. Since the lidar measures radial velocity, a projection of both vertical and horizontal velocities are included in the measurement. The vertical velocity, however, is assumed to be negligible, so the horizontal velocity is directly proportional to the measured radial velocity by the ratio shown in Equation 8.1.

$$V'_H = \frac{V'_R}{\sin(25.8^\circ)} = \frac{V'_R}{0.435}$$

where

V'_H = Horizontal Velocity (Eq.8.1)

V'_R = Radial Velocity measured by the off - vertical beam

For the compass direction of the off-axis beam, due west was desired because of large zonal (east-west) velocities expected in the summer, and to face away from the rising sun and thereby extend the possible lidar observation as much as possible. Again due to hardware limitations typical of first-time measurements, a beam direction 20° south of due west was adopted.

Profiles of the hour-averaged vertical velocity, horizontal velocity, and temperature are shown in Figure 8.7.

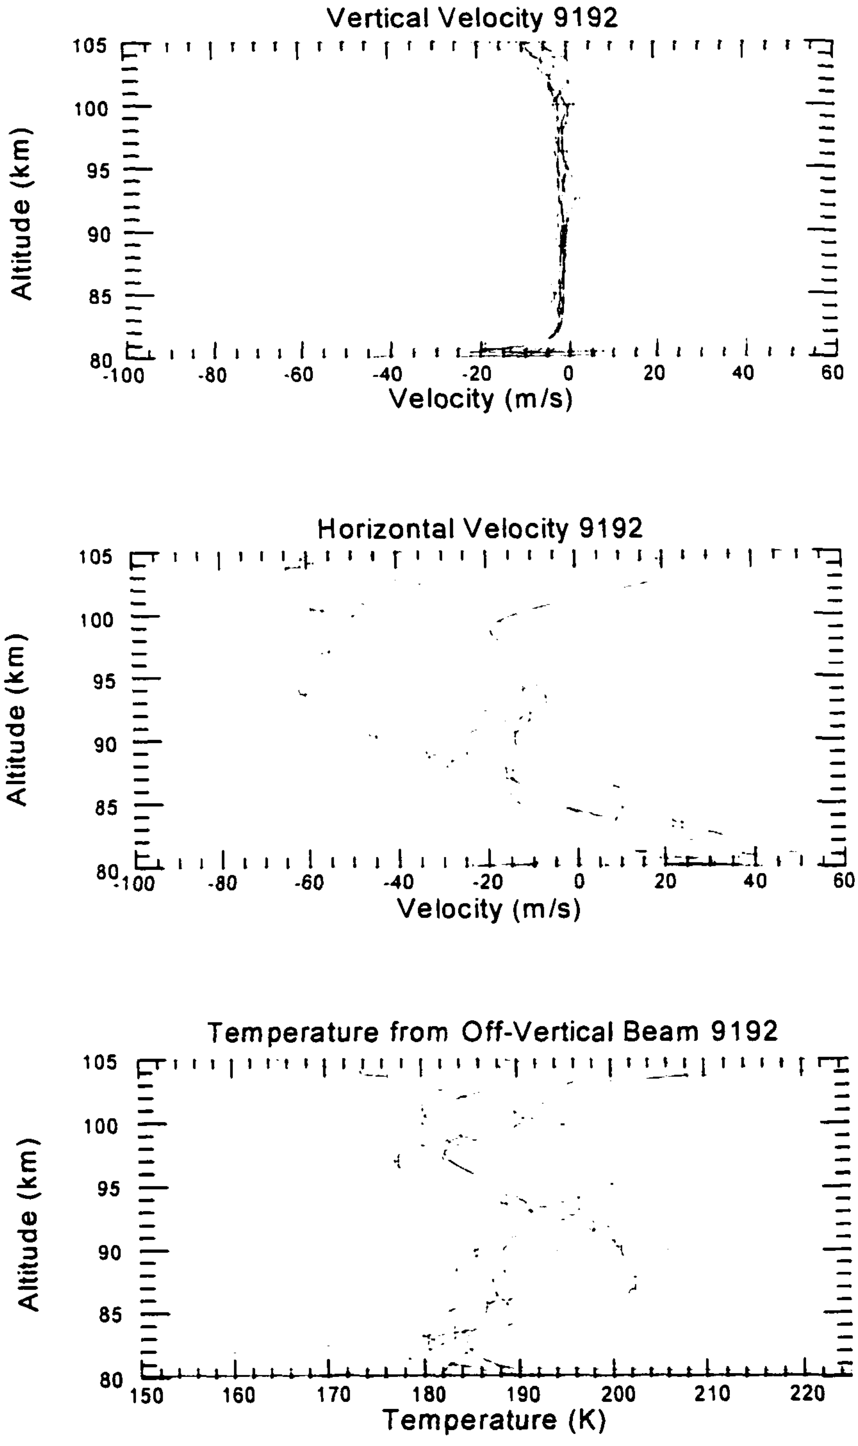

Fig.8.7 The vertical velocity (top), horizontal velocity (middle), and temperature (bottom) profiles from the two-beam horizontal velocity night 9192. Positive horizontal velocities are predominantly westward.

The temperature profile shown in Figure 8.7 was determined from the off-vertical beam, to facilitate comparison with the co-located horizontal velocity measurements. All altitudes in Figure 8.7 have been adjusted to represent vertical velocity above sea-level rather than the off-vertical range corresponding to the range bins of the off-vertical data. A residual vertical velocity of about 1m/s is observed in Figure 8.7, suggesting a small offset assumed present in the horizontal velocities as well. However, the offset is mostly within the calculated error bars of the velocity measurement and generally represents a small part of the horizontal velocities which range from about –60m/s (predominantly eastward) to +30m/s (predominantly westward).

The horizontal velocities in Figure 8.7 agree well in magnitude with zonal winds measured by MF radar (Vincent and Fritts, 1987) and in both direction and profile with combined simulations and meteor radar (Williams, 1997) which show an eastward wind above about 85km which changes to westward below in the summer near 40° latitude. Despite the small, negative offset implied by the vertical wind data, the accuracy of the horizontal wind measurements is within about ±2m/s at the sodium density peak, or about 2% of the full range of the measured horizontal wind. This result is encouraging for the desired elimination of the vertical beam. However, owing to the previously observed variability of the measured velocity offset, additional data is required before real confidence in a one-beam, off-vertical measurement is justified.

8d. Uncertainty and Implications of the Monitored Lidar Data

With respect to temperature and Na density measurements, the 3-frequency lidar may be considered a second-generation instrument, owing the primary development of these capabilities to the 2-frequency lidar which came before it. With respect to the radial velocity measurement, the 3-frequency lidar is closer to being a first-generation instrument, and the uncertainties in the measurements somewhat reflect this. The sources and sizes of these uncertainties, along with the implications of the strengths and weaknesses of the monitored 3-frequency lidar measurement are considered below.

Uncertainty

Table 8.1 outlines the calculated uncertainties in the monitored 3-frequency lidar measurements of temperature, Na density, and radial velocity.

The values shown in Table 1 represent cases of standard operation in which the lidar transmitter is locked to the sodium D_{2a} Doppler-free peak with an accuracy of about 1MHz. In the case of a variable transmitter chirp or imperfect locking of the laser frequency, the correction provided by the velocity measurement and frequency monitor prevents the errors attributed to laser locking from increasing proportionately. The level of uncertainty does become more complex, however, since the uncertainty of the frequency location of the laser for purposes of calculating Na density and temperature then depends on the uncertainty of the velocity measurement, which in turn depends on the photon signal level. A simple calculation of the sensitivities of the temperature and velocity measurements to laser frequency suggests that the error imparted to the temperature measurement, even in the presence of extreme laser frequency uncertainty, is about .25K per 1m/s of uncertainty in the velocity measurement, but without including frequency induced uncertainty in the velocity uncertainty. So, at the sodium density centroid, for example, in the case of a large and variable frequency shift in the lidar transmitter of 20MHz ± 2MHz, the shift induced in the velocity measurement is about 12m/s ± 1.2m/s, but the photon-noise induced uncertainty is only about 0.1m/s due to high signal levels. The temperature uncertainty induced by the frequency uncertainty in the lidar transmitter, then, is about 0.25 x 0.1m/s or 0.025K, which would increase with the altitude difference from the Na density centroid. The effect of such a shift on the atmospheric velocity measurement is independently minimized by the frequency monitor and would be accurate to about .4m/s plus contributions from photon noise and any residual offset.

Table 8.1: Error Budget of the 3-Frequency Lidar Measurements

	Na Density	Temperature	Velocity
Typical Value	4×10^9 m^{-3}	200K	0m/s (vertical)
Laser Locking Error	$< 10^5$ m^{-3}	< .05K	.4m/s Frequency Monitor Uncertainty
Photon Noise Error	$4 \times 10^6 (N_0/N)^{1/2}$ m^{-3}	$0.3(N/N_0)^{1/2}$K	$0.1(N_0/N)^{1/2}$m/s
Other Sources	up to 10% scaling error due to reliance on standard atmosphere tables		-4 to 1m/s night-to-night systematic offset (unexplained residual velocity)

Note: Calculated for 1 hour averages. N signifies total photocounts from a given altitude, N_0 signifies peak photocounts from an altitude near the Na density peak, assuming the same bin size, integration time, and smoothing for N and N_0.

In profiles of both temperature and vertical velocities, it is often apparent that offsets at the top and bottom of the sodium layer, where photon noise is high, varies more than the calculated error bars can account for. The calculated error bars generally rely on the presence of high enough signal to form the experimental ratios and calculate the uncertainties from simple Poisson statistics and error propagation considerations. In the case of low statistical significance of the signal, additional factors such as biasing of the ratio averages due to denominators close to zero, and biasing from the boundaries of the calculated calibration curves may become important issues. The analysis of the uncertainties presented here are applicable to most ranges in the presented data. However, there is certainly room for a great deal more statistical evaluation of the 3-frequency lidar data in the future.

The Residual Velocity Offset

The figures presented in Table 8.1 suggest that the dominant sources of uncertainty in the atmospheric velocity measurement are the currently unknown causes of the residual offset discussed above. Since the vertical velocity measurement provides the only stable, known reference in the mesopause region, namely that the vertical velocity should, on average, be very small, it is unknown whether the factors contributing to the measured velocity offset contribute to the temperature measurement as well. There are a number of possible causes for the measured residual offset, some of which arise in the frequency monitor experiment, some in the atmospheric velocity measurement, and some in both. The details of the analysis that leads to the temperature/velocity calibration curves should also be reviewed for a possible source of offsets in the measured temperature and velocity.

One possible explanation which has not been entirely eliminated is that the measured offset arises purely from a bias in the frequency monitor. A hopeful explanation would be that the ±1m/s systematic error in the calibration of the frequency monitor could account for the offset. However, this explanation is not commensurate with the night-to-night variation of the offset, which could not arise purely from a calibration error. It may be noted, though, that a systematic error from the calibration curves may be present, and given more nights of data could be eliminated based on the averaged vertical velocity measurements from several nights. Other time-varying offsets could be imagined arising from the frequency monitoring experiment, especially from the gated integrators which measure small voltages at high speeds. This conjecture led to the recalibration of the gated integrators by the manufacturer, since which only two nights of useful vertical data have been acquired, both of which exhibit much smaller residual offsets. The statistical dependability of these results is yet to be determined by continued measurements.

Other explanations focus on assumptions made for both the atmospheric and monitor measurements. One such possibility is that the laser lineshape is not as stable as has been assumed. This seems plausible based on the variability of the laser chirp, which is closely related to the shape of the laser lineshape. One problem with this explanation is that these suggested variations of the laser lineshape have never been measured, despite several interferometer-based lineshape measurements. Also, a fairly large change would be required, since a small change in laser lineshape would affect both the monitor and atmospheric measurements similarly, and therefore would tend to be eliminated by the monitor correction. Another possibility is that the frequency differences between the peak and shifted operating frequencies of the lidar are not exactly as assumed. This could be caused by a alignment dependence of the frequency chirp in the pulsed dye amplifier, which might be subject to somewhat random change on a night-by-night basis. Again, the change would be small, and might affect the monitor and atmospheric measurements similarly, but perhaps not identically.

A final possibility is that the frequency monitor experiment is accurate, but the residual velocity offset arises from a night-by-night variation in the atmospheric velocity measurement. One possibility for such an offset would be a frequency-dependent pointing instability in the lidar transmitter. Since the pulsed dye amplifier utilizes several apertures, one of which is only 50μm in diameter, it has been assumed in the data analysis that the pulses which arise from differently shifted seed beams from the acousto-optic modulators have precisely the same output path through the atmosphere. If this is not the case, there could be a frequency-dependent collection efficiency term omitted from the analysis which represents the collecting telescope's efficiency from each of the different beam paths, especially with regard to the difference in collection efficiency between the altitude of the sodium layer and the employed Rayleigh region. This again could be variable on a night-by-night basis since the exact alignment of the seed beams into the pulsed dye amplifier is adjusted for each data collection night, and one hint of such an effect may lie in slight differences in the measured density calculated separately from each of the lidar's output frequencies. Such a pointing offset would not affect the frequency monitor since it is fiber coupled before the beam leaves for the atmosphere, giving it a large angular tolerance.

Chapter 9: Implications and Conclusion

In summary, based on eight nights of atmospheric vertical wind measurements, the use of the frequency monitor has resulted in the reduction of biases in the vertical wind measurement from about 15m/s to about 2m/s. Two nights of data have been acquired after recalibration of the system's gated integrators, and their corrections exhibit smaller residual offsets. Therefore, the feasibility of using iodine absorption spectroscopy to monitor the centroid of the lidar transmitter output lineshape to correct wind measurements has been clearly demonstrated. The causes of the remaining measurement bias are not yet clear. Continued collection of frequency-monitored data with persistent investigation into the properties and possible causes of the offset will hopefully isolate the dominant sources and suggest the means of their elimination. Even with the observed bias, the lidar is clearly accurate enough to perform single-beam simultaneous temperature and horizontal wind measurements as intended.

The current performance of the monitored, 3-frequency lidar suggests a great deal of additional work, both to improve its performance, and use its advantages to advance the understanding of the middle and upper atmosphere. Further characterization and stabilization of the 3-frequency velocity measurement is necessary, and continued measurement of vertical velocity is probably the best way to pursue this. Off-vertical measurements, however, promise the avenue of greatest scientific interest, especially measurements using two orthogonally tilted beams to characterize the range-resolved 2-dimensional horizontal wind field in the mesopause region. A series of such measurements could be used to generate a climatology of winds and wave effects in the region and possibly answer a variety of theoretical questions about the energy budget of the atmosphere. The first step toward this goal is to finalize dependable single-beam off-vertical measurements of the radial wind and temperature.

Another goal which is being rapidly approached with the current development of the 3-frequency lidar is 24-hour range-resolved windfield measurements. With the ongoing progress in ultra-narrow-band filtering for the rejection of daylight background (Chen et. al., 1996), the CSU 3-frequency lidar has acquired preliminary data including temperature and horizontal wind through daylight and nighttime hours. Since the background noise during daylight hours is much greater than at night, optimization of the counting statistics for daytime measurements is required. Again, a stable, single beam measurement of the radial velocity is necessary.

Other lidar applications are also suggested by the successes of the monitored 3-frequency lidar. First, with the implied accuracy of the frequency monitor for correction of velocity measurements in the presence of transmitter frequency instability, and the resulting accuracy of velocity, temperature, and density measurements in the presence of such instability, the traditional requirement of megahertz-scale frequency locking for such measurements is somewhat relaxed. This may open the field of mesopause-region temperature and wind measurements to new classes of transmitting lasers which have been traditionally excluded due to their lack of precise tunability derived from a CW seed laser. Therefore, much simpler systems which are less costly and more amenable to automation may be used, making the possibility of a worldwide network of mesopause monitoring stations more feasible. Similarly, since the frequency shifting system used in the 3-frequency lidar imposes minimal tuning requirements from the CW laser system, solid-state replacements for the CW pumped dye laser are more practicable, and such a laser is currently under development at CSU (Vance et. al., 1998).

The development of the 3-frequency lidar and its associated frequency monitor has been the result of the creative energies of many CSU scientists, and reflect especially the inventiveness of Dr.'s She and Krueger. The intention of this thesis was to review the operational principles of the sodium lidar (Chapter 2), the lidar inversion method for determining the Na density temperature and radial velocity (Chapter 3), the hardware and acquired data of the lidar (Chapters 4 and 5), and the justification, apparatus, and results associated with the lidar's associated frequency monitor (Chapter 6, 7, and 8). This document, however, is not as much a summary of a project begun and completed, encapsulated for historical evidence, as a continuation of previous efforts, a report of current progress, and a pointer to continuing work which spans far more than the tenure of a single, albeit long-lasting, graduate student. For every researcher, though, there comes a time to heed the continued wisdom of the 3rd century B.C. Buddhist philosopher from the introduction of this dissertation and "Leave the past behind; leave the present behind; leave the future behind".

References

P. P. Batista, B. R. Clemsha, D. M. Simonich, V. W. J. H. Kirchhoff, "Tidal Oscillations in the Atmospheric Sodium Layer", J. Geophys. Res., 90, 3881-3888 (1985)

R. E. Bills, C. S. Gardner, C. Y. She, "Narrowband Lidar Technique for Sodium Temperature and Doppler Wind Observations of the Upper Atmosphere", Opt. Eng., 30, 13-21 (1991)

L. M. Caldwell, "A High Spectral Resolution Rayliegh Lidar Using an Iodine Vapor Filter at 589nm", Thesis, Colorado State University, (1995)

H. Chen, "A Sodium Fluorescence Lidar for Daytime Operation Using a Dispersive Faraday Filter", Thesis, Colorado State University (1997)

H. Chen, M. A. White, D. A. Krueger, and C. Y. She, "Daytime Mesopause Temperature Measurements using a Sodium-Vapor Dispersive Faraday Filter in a Lidar Receiver", Opt. Lett., 21, 1093-1095 (1996)

S. Chen, "24 Hour Lidar Campaigns and Tidal Analysis", Thesis, Colorado State University (1999)

B. R. Clemsha, D. M. Sinonich, H. Takashi, P.P. Batista, Y. Shahai, "The Annual Variation of the Height of the Atmospheric Sodium Layer at 23 S: Possible Evidence of Convective Transport", J. Geophys. Res., 97, 5981-5985 (1992)

A. Corney, *Atomic and Laser Spectroscopy*, Oxford, 39, 1977

E.E. Eyler, A. Yiannopoulou, S. Gangopadhyay, N. Melikechi, "Chirp-free Nanosecond Laser Amplifier for Precision Spectroscopy", Opt. Lett., 22, 49-51 (1997)

K. H. Frike, U. von Zahn, " Mesopause Temperature Derived From Probing the Hyperfine Structure of the D_2 Resonance Line of Sodium by Lidar", J. Atmos. Terr. Phys., 47, 499 (1985)

S. Gangopadhyay, N. Mlikechi, and E. E. Eyler, "Optical Phase Perturbations in Nanosecond Pulsed Amplification and Second-harmonic Generation", J. Opt. Soc. Am. B, 11, 231-233, (1994a)

A. L. Gibson, M. C. W. Sandford, "Daytime Lidar Measurements of the Mesospheric Sodium Layer", Nature, 33, 509-511 (1972)

J. W. Hair, "High Spectral Resolution Rayleigh Lidar at 532nm", Thesis, Colorado State University (1999)

F. J. Lubken, U. von Zahn, "Thermal Structure of the Mesopause Region at Polar Latitudes", J. Geophys. Res., 96, 20841-20857 (1991)

N. Melikechi, S. Gangophdhyay, E. E. Eyler, "Phase Dynamics in Nanosecond Pulsed Dye Laser Amplification", J. Opt. Soc. Am. B, 11, 2402-2411 (1994b)

H. Moosmuller, J. D. Vance, "Sum-frequency Generation of Continuous-wave Sodium D_2 Resonance Radiation", Optics Letters, 22, 1135-1137 (1997)

G. C. Papen, W. M. Pfeninger, D. M. Simonich, "Sensitivity Analysis of Na Narrowband Wind-Temperature Lidar System", App. Opt., 34, 480-497 (1995)

R. Roble, Ridley, "A Thermosphere-Ionosphere-Mesosphere-Electrodynamics General Circulation Model (TIME-GCM)", Geophys. Res. Lett., 21, 417-420 (1994)

C. Y. She, J. R. Yu, D. A. Krueger, R. Roble, P. Keckhut, A. Huchecorne, M-L. Chanin, "Vertical Structure of the Midlatitude Temperature From Stratosphere to Mesopause (30-105km)", Geophys. Res. Lett. 22, 377-380 (1995)

C.Y. She, J. R. Yu, "Simultaneous Three-Frequency Na Lidar Measurements of Radial Wind and Temperature in the Mesopause Region", Geophys. Res. Lett., 21, 1771-1774 (1994)

C. Y. She, J. R. Yu, H. Chen, "Observed Thermal Structure of a Midlatitude Mesopause", Geophys. Res. Lett., 20, 567-570 (1993)

C. Y. She, J. R. Yu, H. Latifi, R. E. Bills, "High-spectral-resolution Fluorescence Light Detection and Ranging for Mesospheric Sodium Temperature Measurements", Appl. Opt., 31, 2095-2106 (1992)

C. Y. She, J. R. Yu, J. W. Huang, C. Nagasawa, C. S. Gardner, "Na Temperature Lidar Measurements of Gravity Wave Perturbations of Wind, Density, and Temperature in the Mesopause Region", Geophys. Res. Lett., 18, 1329-1331 (1991)

C. Y. She, "Remote Measurements of Atmospheric Parameters: New Applications of Physics with Lasers", Contemp. Phys., 31, 247-260 (1990)

C. Y. She, H. Latifi, J. R. Yu, R. J. Alvarez II, R. E. Bills, C. S. Gardner, "Two-Frequency Lidar Technique for Mesospheric Na Temperature Measurements", Geophys. Res. Lett., 17, 929-932 (1990)

X. Tao and C. S. Gardner, "Heat Flux Observations in the Mesopause Region above Haleakala", Geophys. Res. Lett., 22, No. 20, 2829-2832 (1995)

J. S. Theon et. al., "The Mean Observed Structure and Circulation of the Stratosphere and Mesosphere", NASA Tech. Rep. TR-375 (1972)

J. D. Vance, "Sum-Frequency Generation of Continuous-wave Sodium Resonance Radiation for Mesospheric Temperature and Wind Measurements", Thesis, University of Nevada, Reno (1997)

J. D. Vance, C. Y. She, H. Moosmuller, "Continuous-wave, All-solid-state, Single-frequency 400mW Source at 589nm Based on Doubly-resonant Sum-frequency Mixing in a Monolithic Lithium Niobate Resonator", Applied Optics, 37, 4891-4896 (1998)

R. A. Vincent and I. M. Reid, "HF Doppler Measurements of Mesospheric Gravity Wave Momentum Fluxes", J. Atmos. Sci., 40, 1321-1333 (1983)

R. A. Vincent, D. C. Fritts, "A Climatology of Gravity Wave Motions in the Mesopause Region at Adelaide, Australia", J. Atmos. Sci., 44, 748-760 (1987)

M. A. White, D. A. Krueger, C. Y. She, "A Frequency-agile Lidar with Concurrent Frequency Monitor for Simultaneous Measurement of Temperature and Line-of-sight Wind in the Mesopause Region", 19th ILRC Abstracts, 363 (1998)

W. L. Wiese, W. M. Smith, B. M. Miles, *Atomic Transition Probabilities, Vol.II-Sodium*, NSRDS-NBS 22, 2 (1969)

B. Williams, "Observations of the Atmospheric Tide, Mean Wind, and Sodium Nightglow near the Mesopause with the Magneto-Optic Doppler Analyzer", Cooperative Thesis No. 161, University of Colorado and The National Center for Atmospheric Research, p.13 (1997)

J. R. Yu, "A Sodium Wind/Temperature Lidar and Observed Mesopause Thermal Structure over Fort Collins, CO", Thesis, Colorado State University (1994)

INDEX

G

H

I

L

M

N

O

P

R

S

T

U

W