Distribution of an Ultrastable Frequency Reference Using Optical Frequency Combs

by

Kevin W. Holman

B.S., Purdue University, 2000

M.S., University of Colorado, 2004

A thesis submitted to the

Faculty of the Graduate School of the

University of Colorado in partial fulfillment

of the requirements for the degree of

Doctor of Philosophy

Department of Physics

2005

This thesis entitled:
Distribution of an Ultrastable Frequency Reference Using Optical Frequency Combs
written by Kevin W. Holman
has been approved for the Department of Physics

Jun Ye

Steven T. Cundiff

Date ____________________

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

Holman, Kevin W. (Ph.D., Physics)

Distribution of an Ultrastable Frequency Reference Using Optical Frequency Combs

Thesis directed by Assoc. Prof. Adjoint Jun Ye

Frequency standards based on optical transitions in various atomic systems provide the potential for developing optical atomic clocks in the next few years that are several orders of magnitude more stable than the best existing clocks based on microwave-frequency references. The excellent stability that will be provided by optical clocks, and even the stability they currently offer, has applications in several fields, ranging from studies of fundamental physics to communication and remote synchronization. However, since optical clock systems are too complex to be portable, these applications rely on the ability to transfer over several kilometers a stable frequency reference that is linked to an optical clock. In this dissertation I present my research that enables the use of optical frequency combs from mode-locked lasers for transferring over optical fibers frequency references that are linked to optical standards.

This method of transfer requires that the optical frequency comb be stabilized to the optical standard. This is most easily accomplished by first stabilizing the comb from a mode-locked Ti:sapphire laser, which is described in detail. In particular, since intensity control of the laser is used for this stabilization, an analysis is presented on the intensity-related dynamics of the laser, enabling optimization of the control scheme. To minimize loss during transmission, it is necessary to transmit a 1550-nm comb instead of the 800-nm Ti:sapphire comb. Therefore, the stabilization of a 1550-nm mode-locked laser diode to the Ti:sapphire laser is discussed. This involves both the synchronization of the two lasers and the locking of their optical phases. The lowest reported timing jitter for a mode-locked laser diode is demonstrated.

Finally, measurements of the stability for transferring a microwave frequency over several kilometers of optical fiber using an optical frequency comb are presented. Without active stabilization, the comb provides an order of magnitude higher stability than is measured for existing methods of microwave-frequency transfer over fibers. With active cancellation of the transfer noise, the lowest timing jitter reported for the transfer of a frequency reference over several kilometers using optical fibers is achieved.

Dedication

To my father, Gary, who taught me the joy of understanding how things work.

Acknowledgements

Completing a Ph.D. thesis is not an individual achievement, and without the help and support from many others this would not have been possible. Here I acknowledge those who have made invaluable contributions without which my efforts over the last few years would have been far less fruitful, and I apologize in advance for any omissions I may make. First I would like to thank my advisor, Jun Ye, who provided me the opportunity to learn about ultrafast lasers and to be involved in making some of the most precise optical frequency measurements in the world. His enthusiasm for doing exciting science and performing world-class research has been an inspiration as I have defined myself professionally over the last several years. I also want to express my gratitude to Jan Hall and Steve Cundiff, who have always been available for discussions and to offer advice on any difficulties I have encountered.

I have been blessed with the opportunity to work with many talented people during my time in Jun's group. I want to thank Jin-Long Peng for his instruction from my very first day in the group until he left a few months later. Jason Jones joined the group as a postdoc shortly after I did, and his excellent support and guidance during the early stages of my graduate research and throughout my time at JILA have made an immeasurable impact on my development as a scientist. I am also grateful for the contributions that David Jones, Mark Notcutt, and Kevin Moll made to my research.

There are a number of students in Jun's group, and all contributed in some way or another during my time in the group. I want to acknowledge those who made the

most significant impact on my research, including Adela Marian, Seth Foreman, and Darren Hudson, without whom the research would have been far less productive and enjoyable. I also want to thank Lisheng Chen, who taught me a great deal about laser spectroscopy, as did the members of the group that work on the spectroscopy of laser-cooled strontium atoms. They taught me nearly everything I know about the exciting field of precision spectroscopy of ultracold atoms, and for this I want to thank Marty Boyd, Tom Loftus, Andrew Ludlow, and Tetsuya Ido.

I have enjoyed a very valuable collaboration with several researchers in the Time and Frequency Division of the NIST laboratories in Boulder. Leo Hollberg, Scott Diddams, and John Kitching have provided our lab in JILA with a signal referenced to the NIST hydrogen maser, which has been crucial for many of the measurements we have made. In addition, my research has directly benefited from many excellent discussions I have had with them, as well as those I have had with another NIST researcher, Nate Newbury.

As anyone at JILA knows, much of the excellent research that is done at JILA would not be possible if not for the world-class support staff, including the support from the electronics and machining staff. I am grateful for the contributions made by Terry Brown, James Fung-A-Fat, and Hans Green in devising ingenious solutions to some of the unique challenges I have encountered.

In addition to the technical contributions to my graduate research, I also want to acknowledge those contributions that often go unmentioned, but are no less significant. I am very grateful for the wonderful friends I have in JILA, including Eric Hudson, Marty Boyd, Jason Jones, Kevin Moll, Matt Stowe, Darren Hudson, Seth Foreman, and David Gaudiosi. They have not only made it a pleasure to work in JILA, but have really helped me enjoy my time in Boulder.

I am eternally grateful for the love and support my family has always given me. My wife, Hannah, has demonstrated unbounded patience and tolerance during the many

long nights I have been away from home while finishing my research and writing my thesis, as well as the months we have lived in separate parts of the country while she has finished her graduate studies in Indiana. I am also very thankful for the support and encouragement my parents, Gary and Brenda, and my younger brother, Brian, have provided throughout my life, without which it is likely I would never have chosen this path for my life.

Finally, I would like to acknowledge the financial support provided me by the Fannie and John Hertz Foundation. Without this assistance, I might not have had the freedom to pursue the field of research that most interested me.

Contents

Chapter

Appendix

Tables

Table

Figures

Figure

Chapter 1

Introduction to Optical Frequency Standards and Clocks

1.1 History of timekeeping: Interplay between science and clock technology

Measuring and recording the passage of time has always been a vital aspect of human societies. Ancient civilizations maintained records of the months and seasons to coordinate trade, community activities such as public meetings, and the planting of crops. In addition to enabling day-to-day activities to be synchronized among members of a community, the measurement of time has always been tightly intertwined with and interdependent on the progression of new scientific discoveries. Historically, improvements in time-keeping technologies and enhancements in the ability to more precisely subdivide the passage of a day into finer and finer increments have directly impacted the progress of new scientific discoveries. Conversely, recently new scientific discoveries have in turn enabled the development of better clocks, which can now be used to explore new problems of both scientific and technological interest.

Regardless of the sophistication of a timekeeping device, all clocks have one component in common — they must rely on something to provide the “tick.” At the heart of any clock is some periodic physical process, and by counting the oscillations of this process the duration of an interval of time can be measured with respect to the period of this oscillator. In ancient times, the most accessible oscillatory physical process was the rotation of the earth. As early as 20,000 years ago hunters during the ice age notched holes in sticks or bones to record the passage of days between the moon phases [73]. In the first or second century A.D., the Romans developed hemispherical sundials, which enabled them to use the position of the sun to subdivide the day into 24 segments, or hours [4]. Sundials were used by several other ancient cultures, including the Egyptians and Greeks, in conjunction with water clocks that measured the passage of time during the night by gauging the flow of water out of a basin. One of the most elaborate water clocks was built by the Chinese around 1000 A.D. [3]. A major technological advance was made in the 13th century with the development of a weight-driven mechanical clock. The period of the mechanical oscillations of the clock was adjusted to allow it to accurately demarcate the 24 hours of the day, and these clocks proved far more reliable than their predecessors. With the replacement of the weight-driven mechanism with one powered by a coiled spring, these clocks became portable and started appearing in households across Europe in the 15th century.

Though these clocks were sufficiently accurate to coordinate day-to-day activities, they were not yet suitable for scientific purposes. They would typically gain or lose 15 minutes a day. This situation changed, however, with the invention of the pendulum clock by Christian Huygens in 1656 [4]. Using the oscillations of a swinging pendulum as the periodic reference for measuring time, these clocks were accurate to about a minute over a week. Improvements of pendulum clocks soon allowed them to keep time to within a few seconds over a week. The precision provided by these new clocks was immediately used by astronomers to time the movement of stars and to create the most accurate maps of celestial bodies to date. The progress of clock technology had impacts in other fields as well. By the late 18th century, improvements in the spring-powered mechanical clock significantly improved sea navigation. With the ability to accurately keep time while at sea it became much easier to determine a ship's longitude, a feat which without accurate clocks had proven very difficult. In fact, the marine chronometer has not changed significantly to this day. Astronomy and navigation represent just the beginning of the impact new clock technology would have on scientific progress — with the ability to accurately measure fractions of a second, clock technology began to play a significant role in making scientific discoveries in a number of other fields as well.

Just as the ability to measure time more precisely was enabling new scientific discoveries, science was in turn providing keys to make even better time-measuring devices. In 1928, a researcher at Bell Laboratories, Warren A. Marrison, discovered that quartz crystals resonate at a very well-defined frequency when subjected to an oscillating voltage [4]. Soon clocks were developed that used as their oscillator the frequency reference provided by quartz crystals. These clocks could achieve an accuracy of one second over 30 years. It would not take long though before science would provide even more accurate ways to keep time. The development of quantum mechanics in the early part of the 20th century revealed that the electrons of an atom can occupy only discrete energy levels. Therefore, the transition between two electronic states of an atom corresponds to a well-defined energy difference, E, which in turn is equivalent to a specific frequency of electromagnetic radiation, ν: $\nu = E/h$, where h is Planck's constant. By using this transition frequency as a reference to guide the oscillator of a clock, extremely accurate measurements of time are possible. In 1955 an atomic clock was developed that was based on the hyperfine splitting of the electronic ground state of the cesium (Cs) atom, which corresponds to the energy difference involved in flipping the spin of one of its electrons and is equivalent to a transition frequency of $\sim$9.2 GHz [15]. Cesium atomic clocks enabled the precise measurement of unimaginably small increments of time and could demonstrate an accuracy of $\sim$1 ns over a day.

1.2 Redefining time

Even though the performance of clocks had improved dramatically, until 1956 the standard for measuring an interval of time was still the rotational period of the Earth. The basic unit of time, the second, was defined in terms of how many fit into a mean solar day. However, as the precision of clocks improved, it became obvious that this standard was changing over time — the rotational period of the Earth was in general getting longer, attributed in part to tidal friction. Therefore, in 1956 a new standard for measuring time was adopted. The new standard, the Ephemeris Second, was defined as a given fraction of the tropical year 1900 [15]. Although more stable than a standard based on the period of a day, this new definition was not very accessible, nor was it practical for making measurements of short intervals of time. The development of the Cs atomic clock provided a solution to this problem. The unrivaled precision for measuring time provided by this atomic frequency reference and the constancy of the atomic transition frequency made the Cs atom an excellent alternative to replace an astronomical standard for time. In 1967 the second was redefined as the time of 9,192,631,770 cycles of the ground-state hyperfine splitting of the unperturbed Cs atom [4], and this standard is still used today.

Improvements continued to be made on the Cs clock, and with the development of laser cooling of atoms in the late 1980s [56] came the Cs fountain clock. Early Cs clocks used a hot beam of Cs atoms, but fountain clocks laser cool the atoms to form a ball of Cs at microkelvin temperatures, thereby reducing Doppler shifts of the atomic resonance frequency. This ball of atoms is then launched upward to rise and fall ballistically under the force of gravity. Cooling of the atoms before launching them is also important to prevent thermal motions of the atoms from dispersing them during flight. The atoms are then interrogated by microwave radiation while travelling upward and again while falling. The relatively long time of flight of these atoms provides for an interrogation time of the atomic transition of $\sim$1 s, as compared to < 10 ms for the fast travelling Cs atoms in a beam clock. A longer interrogation time is important since the uncertainty in determining the center of the atomic resonance is reduced as the interrogation time is increased. Cesium fountain clocks currently provide the most accurate means to measure time. A common metric for measuring the performance of a clock is the fractional uncertainty of its frequency reference, $\delta\nu/\nu_0$, that guides the oscillations of the clock. For these systems, $\delta\nu/\nu_0$ can be as low as 6×10^{-16} [29], which corresponds to keeping time to within one second over more than 50 million years.

1.3 Motivations for precise timing

Though it may seem at first that keeping time to within a second over 50 million years has no practical value, the timing precision provided by atomic clocks plays a vital role in current scientific investigations and is crucial for many daily activities of the general public. For example, the Global Positioning System (GPS), which consists of 24 satellites transmitting synchronized coded signals, is reliant on accurate atomic clocks in each satellite to ensure the synchronization of the signals. A receiver on Earth that receives these signals from four or more satellites can determine its distance to each satellite from how much time lag there is between the synchronized signals, thereby pinpointing its exact location on Earth. A receiver that has an accurate timing mechanism can determine its position to within 8 m, and by averaging over a long period of time the measurement can be accurate to 1 mm [64]. GPS has become quite valuable in a number of scientific fields, including geology for measurements of Earth's crustal deformations and continental drifts, paleontology and archeology for recording the locations of fossils and artifacts, and civil engineering for monitoring the settling of manmade structures over time and land surveying. GPS is having an impact on the lives of the common person as well, as more people are relying on GPS navigators while hiking and even while driving their cars.

Precise timing that is provided by atomic clocks is also necessary for the satellite and high-speed optical communications that many rely on every day. As the need to transmit more and more information per second grows, the synchronization requirements for the components of communication networks become more stringent. Atomic clocks are necessary to support the large amount of cell phone traffic and to enable large computer networks, including the Internet, to function reliably. In addition to facilitating communication, atomic clocks help manage the electric power grid across the country by ensuring the oscillating current is maintained at exactly the right frequency and phase across different regions of the country.

Another field that has benefited from precise timing is that of metrology. The accuracy and precision of atomic clocks has made the second the most accurately realized unit of measurement of all physical quantities. For this reason, the definitions of other physical quantities are being changed to be expressed in terms of the second. For example, in 1983 the meter was redefined by specifying a defined value for the speed of light in vacuum. A meter became the distance light travels in a vacuum in a time interval of $1/299,792,458$ of a second. Other units whose definitions rely on the second include the ampere and the volt [15]. Currently there are even efforts to link the kilogram, the only remaining artifact standard, to the second [20]. Because of the reliance on the second for the definitions of several standards, the ability to accurately measure time is important in many scientific fields for the measurement of various physical quantities.

Atomic clocks have also proven very useful in radio astronomy. The size constraints for a single radio telescope limit its resolution, which is inversely proportional to the telescope aperture. However, by phase-coherently collecting data from radio telescopes on opposite sides of the globe, the effective aperture is the distance between the telescopes. For this scheme to work, which is referred to as "very long baseline interferometry" (VLBI), accurate atomic clocks must be used at each telescope to synchronize the data collection. Arrays of up to 12 radio telescopes have been used to achieve an angular resolution of 200 μarcs, which is 250 times better than the best optical telescope, including the Hubble Space Telescope [64]. This same principle is also used to track spacecraft as they travel through the solar system, using several radio telescopes.

Finally, careful measurements of time intervals have allowed the most rigorous tests to date of Einstein's special and general theories of relativity. The period of an atomic clock placed in a rocket travelling to an altitude of 10,000 km increased with speed and decreased with altitude by just the amount predicted by the special and general theories [79, 76]. In another experiment, the delays for radio waves passing near the Sun were found to agree very well with those predicted [64]. The most stringent test of general relativity has been carried out by carefully measuring the orbital periods of pulsars that are part of a binary pair [66, 12]. It was found that the orbital periods of these binaries are changing by just the amount expected from the loss of energy by the radiation of gravity waves, which were predicted by the general theory of relativity but never before experimentally observed. Furthermore, this experiment confirmed that gravity waves propagate with the velocity of light c, since the excellent agreement between the experiment and theory disappears when the velocity of propagation differs from c by more than 1%.

Although the excellent stability of the Cs atomic clock has enabled a number of scientific and technological improvements, even more could be accomplished with a better clock. For example, the averaging time needed to determine the relative positions of two nearby GPS receivers to within $\sim$1 cm could be significantly reduced using better clocks [64]. With this level of precision, one could envision using GPS to automatically guide vehicles on the road with no human intervention. A better clock would also enable utility and telecommunication companies to pinpoint faults in their networks. VLBI for radio astronomy could be improved to achieve higher resolution by using space-based telescopes that are separated much further than is possible on Earth. However, these would require better timing references than a Cs atomic clock. A clock exhibiting a femtosecond of stability over a measurement period of a few seconds would provide a ranging accuracy of $\sim$1 μm over millions of kilometers [15]. Finally, clocks with lower instability could be used to search for time variation of fundamental constants [48], such as the fine structure constant α [60]. General relativity requires that α be constant over time, but some competing theories of gravity such as some string theories predict that it should change over time. Careful measurements that put limits on the time variation of α are valuable for imposing tighter constraints on various theories. Though some astronomical data suggest that α may have changed by as much as 1 part in 10^5 over the last 10 billion years or so [81, 53], there have been some arguments raised over this conclusion. Also, evidence of a change in α over the last 10 billion years does not indicate whether or not there is currently any variation of α with time. Laboratory comparisons of the most precise clocks ever made that are based on different atomic transitions are now providing some of the tightest constraints on the current time variation of α [7, 55, 15].

Clearly many fields would benefit from the development of better clocks, so the question that arises is how to make a clock that is more stable than the Cs atomic clock. Since a clock's performance is ultimately limited by the fractional uncertainty of its frequency reference, this reference is what must be improved. We have seen how this has occurred throughout the history of clock technology: the earliest clocks were referenced to the frequency of the Earth's rotation, which was then replaced by oscillations of a mechanical systems, and finally by the hyperfine splitting of the Cs atom. What then can be used to improve upon this frequency reference?

1.4 Optical frequency standards

An obvious way to improve the fractional frequency uncertainty, $\delta\nu/\nu_0$, of an atomic reference is to use a transition with a much higher frequency, ν_0, for which the center of the transition can be determined with comparable uncertainty, $\delta\nu$. Therefore, there are currently several research efforts to develop clocks based on optical frequency transitions of 100s of THz, which have frequencies that are four orders of magnitude larger than the microwave transition on which the Cs atomic clock is based. These optical frequency transitions are then used as a reference to stabilize the laser that is serving as the clock's oscillator. There are several different atomic systems that are being investigated that offer different advantages for achieving the lowest fractional frequency uncertainty, including various neutral atoms and ions.

To reduce Doppler broadening of the transition, these investigations require the use of laser-cooled and trapped atoms or ions. Ultracold atoms and ions allow the separation of their internal degrees of freedom from their external center-of-mass motions. The uncertainty in determining the center frequency of the clock transition is then given by $\delta\nu = \frac{1}{2\pi\sqrt{NT_R\tau}}$, where N is the number of trapped atoms or ions, T_R is the interrogation time and is shorter than the lifetime of the transition, and τ $(\tau > T_R)$ is the total averaging time [84]. Since the fractional uncertainty of the transition frequency is reduced by using as many atoms as possible, neutral atoms allowing large values for N in the trap are attractive systems to study. Some neutral atoms that have been studied are calcium, strontium, ytterbium, magnesium, and hydrogen, some of which can achieve fractional inaccuracies of 10^{-14} or below [15, 72].

Although large values of N reduce the uncertainty of the transition frequency for neutral atoms, this can also lead to collisional shifts of the transition [39]. Another limitation of neutral atoms is that the trap must be turned off while probing the clock transition since the trapping lasers can shift this transition. This limits the interrogation time of the atoms and introduces Doppler-related systematic shifts of the transition. On the other hand, ions can be probed while trapped and so they offer very long interrogation times and virtually no Doppler-related effects. Trapped ions that have been studied include mercury, ytterbium, strontium, and indium ions, and currently these systems offer higher accuracy, but less stability, than neutral atoms [77]. The drawback of ions is that due to ion-ion interactions the number of ions in the trap is limited to only a few. However, this suggests that collisional shifts are not a problem for ions, though this must be demonstrated in each case.

It may be possible to develop a system that has the advantages of both neutral atoms and single ions. Neutral atoms can be trapped in an optical lattice whose laser frequency is carefully chosen to minimize the shift of the clock transition [50]. With this arrangement the atoms can be probed while trapped, enabling long interrogation times and the elimination of Doppler shifts while achieving a large number of atoms in the trap. Also, with a three-dimensional lattice it is possible to ensure that there is no more than one atom occupying each lattice site, eliminating collisional shifts. Recent experiments have begun to study a very narrow transition (∼1 mHz) of strontium atoms confined in an optical lattice [74].

Though it is not clear which system will ultimately provide the best performance for an optical atomic clock, it is expected that optical clocks will eventually be able to achieve uncertainties approaching a part in 10^{18} [63], nearly three orders of magnitude better than the best microwave atomic clocks. In addition to achieving a much lower uncertainty than microwave atomic clocks for long averaging times, optical clocks are expected to exhibit short-term instabilities of a few parts in 10^{17} for a 1-s averaging time [82], which is also nearly three orders of magnitude better than the best microwave clocks. The stability of a frequency reference for short averaging times less than one second is important for two reasons. First, in an experiment in which a certain level of uncertainty is required, a lower short-term instability allows this desired level of uncertainty to be reached with a shorter averaging time. Since in the quantum projection noise limit the instability of an optical clock averages down as $\tau^{-1/2}$ for an averaging time τ, achieving an order of magnitude improvement in the uncertainty requires averaging for a time two orders of magnitude longer. Therefore, it is very beneficial to begin with a short-term instability that is as small as possible. Second, some applications, such as the tight timing synchronization of various systems that relies on a high-stability frequency reference, cannot take advantage of averaging over an extended period of time. A frequency reference which is very stable over very short timescales (much less than one second) is crucial for the tight synchronization of these systems. It is important to consider what provides the short-term stability of an optical clock.

Typically information from the atomic transition is used to guide the frequency of the laser serving as the oscillator of an optical clock over timescales of ~10 s. Therefore, the short-term stability is determined by the laser itself. For a low short-term instability it is very important to develop extremely narrow-linewidth lasers. For example, for experiments with the mercury ion, a laser with a linewidth of 0.6 Hz has been developed, which provides a fractional frequency instability of 3×10^{-16} for a 1-s averaging time [89].

1.5 Gears for optical atomic clock

Although optical frequency transitions provide very stable references for the development of better clocks, they also present a significant technical challenge. The period of the oscillating electric field of the laser that is serving as the clock's oscillator is only a few femtoseconds, which means that existing methods for tracking the cycles of microwave clocks are not sufficient for optical clocks. New techniques are necessary that can count the oscillations of an optical frequency and produce a useful clock output. Essentially what is needed is the set of gears for the optical clock, figuratively speaking, which can phase-coherently divide the optical frequency down to a microwave frequency that can be processed and counted with existing microwave electronics. With this set of gears, the output of an optical clock would be a microwave signal that exhibits the same fractional frequency instability as the optical frequency reference. Phase coherence of these gears is important since phase errors in the clock output lead directly to timing error.

In the past, optical and microwave frequencies have been phase-coherently connected using elaborate harmonic phase-locked frequency chains [88]. In these chains, higher-frequency oscillators are phase coherently linked to harmonics of lower-frequency references, which are in turn stabilized to even lower-frequency oscillators in the same manner. Using various types of oscillators operating across a wide range of frequencies, these complex frequency chains provide a phase-coherent connection between a microwave and an optical frequency. However, these chains are typically designed to link the microwave region of the spectrum to a specific optical frequency. Therefore, to provide the flexibility to link to a variety of optical frequency references, other methods must be used to span gaps of many terahertz across the optical spectrum, such as frequency-interval bisection [75]. These complex systems were extremely expensive and difficult to build and run, and only a few were developed in national laboratories across the world. More importantly, phase coherence in these harmonic frequency chains is difficult to maintain, which limits the measurement precision they can provide.

In 1978, an alternative method of connecting microwave and optical frequencies was proposed, and an initial demonstration of this technique was performed. A dye laser emitting picosecond pulses was used to measure fine and hyperfine frequency intervals in sodium [17], with respect to a microwave reference. This can be understood by considering that the Fourier relationship between time and frequency dictates that evenly spaced laser pulses in time produce an array of discrete frequency components that are uniformly spaced by the pulse repetition frequency. These frequency components correspond to the longitudinal modes of the dispersion-compensated laser cavity. Careful measurement of the laser repetition frequency using microwave counting techniques allowed the frequency comb to be used as a ruler to measure the frequency intervals of sodium.

The discovery of Kerr-lens mode locking in titanium-doped sapphire (Ti:sapphire) lasers [71] and the development of Ti:sapphire laser systems with extremely short pulses in the early 1990s (∼10 fs) [5] greatly advanced the use of frequency combs for phase-coherently connecting microwave and optical frequencies. The short pulses from these systems provide a very broad and reliable optical frequency comb that can serve as the “gears” of an optical clock. Phase-locking the optical frequencies of the comb to the laser serving as the oscillator of the clock results in the phase-locking of the comb spacing (i.e., the laser repetition frequency) to the clock's oscillator as well. Therefore, the frequency comb provides the phase-coherent division from an optical frequency to a microwave frequency, and the clock output is the repetition frequency of the mode-locked laser producing the comb. In Chapter 2 I will discuss in a great deal more detail the very important process of distributing the stability of an optical clock across the visible spectrum and down to the microwave spectrum using an optical frequency comb.

1.6 Making optical frequency standards accessible: Frequency transfer

Investigations of potential optical frequency standards in various laser-cooled atoms and ions and the development of optical frequency combs promise to provide an optical clock which performs several orders of magnitude better than the best microwave atomic clocks today. However, these complex systems are not portable, which prevents the advantages offered by a more stable clock from being fully realized. Frequency standards based on atomic or molecular vapor cells, such as the one that will be discussed in Chapter 4, can serve as portable standards, but their performance is several orders of magnitude worse than the more complex optical standards based on cold atoms and ions. Therefore, what must be devised is a way to transfer to remote users a frequency reference that maintains the accuracy and stability of the best optical clocks throughout the transfer process. This is important for comparisons and performance verification of various clock systems and for the distribution of clock signals for other purposes.

Several of the applications discussed in Section 1.3 requiring a better clock will also rely on the ability to transfer over several kilometers an extremely stable frequency reference that is linked to this clock. For example, it was mentioned previously that comparing the frequency references of atomic clocks based on narrow transitions in different atomic species would enable the precise measurement of the time variation of fundamental constants such as α. If α were changing over time, the frequencies of these transitions would change with respect to each other. Therefore, it is necessary to make very precise comparisons of these frequencies over a period of several years. Since the development of an optical clock demands a great deal of time and resources, typically multiple clocks based on transitions in different atomic species are not built in the same laboratory. Therefore, the ability to transfer a frequency reference that reliably maintains the accuracy of these clocks between separate laboratories is crucial for these comparisons. The most accurate measurements to date have placed a constraint of a part in 10^{15} on the possible fractional variation of α in a year, and the future generations of optical frequency standards should provide a sensitivity at the level of 10^{-18} per year. The comparison of optical frequency standards also enables the evaluation of their performance by measuring their relative instability and systematic shifts, since there are no other frequency references stable enough against which these comparisons can be made.

Comparisons of optical frequency standards are accomplished by averaging the frequency measurements over an extended period of time, and so it is not important for the frequency transfer to exhibit high stability for extremely short timescales. However, as alluded to at the end of Section 1.4, there are some applications where it is necessary to achieve very tight timing synchronization among system components. For these applications, the distributed frequency reference to which all the components will be synchronized must have very high stability and low phase noise over short timescales, which is equivalent to saying that it must exhibit ultralow timing jitter (approximately tens of femtoseconds) over a broad bandwidth. One such application is high-speed communications. As higher and higher data rates are demanded, the tolerance for timing jitter between components of a communication network will become smaller and it will become necessary to synchronize them by distributing a frequency reference exhibiting low timing jitter. Another application of high-stability frequency transfer to communications is the implementation of a secure communication channel between two users that have access to a highly-stable reference frequency. The transmission of data by only slightly shifting the frequency of the transmitted signal would imply that an eavesdropper without access to the superior stability of an optical frequency reference would not be able to detect the slight modifications of the carrier frequency — they would be buried in the eavesdropper's measurement noise. To achieve a sufficiently high data rate, the transfer process would require excellent short-term stability so each bit could be detected without a significant averaging time.

Another application for distributing a low-jitter frequency reference is long-baseline interferometry for radio astronomy. This is similar to the VLBI discussed in Section 1.4 but over a much smaller scale, such as ∼20 km. Low-jitter transfer of a frequency reference could be used to distribute a signal from a master oscillator to each telescope in an array of ∼60 radio telescopes. This would enable all telescopes in the array to phase-coherently collect data, thereby simulating a single telescope with a very large aperture [70].

Finally, high-stability transfer could be used to transmit a low-jitter frequency reference throughout a linear accelerator facility for synchronization of its various components. Recently there has been considerable interest in producing ultrashort x-ray pulses to study ultrafast phenomena in several fields including chemistry, physics, biology, and materials science [2]. These studies will involve pump-probe experiments using visible lasers to pump the samples and the x-ray pulses to probe them. Transfer of a low-jitter frequency reference throughout an accelerator facility will be necessary to synchronize the visible pump pulses with the short x-ray pulses at the sample. It will also be crucial for the generation of the ultrashort x-ray pulses, since the components of the accelerator must be synchronized with the short bunches of accelerated electrons that produce the x-ray pulses.

A traditional method for transferring frequency and time standards over long distances has been common-view GPS, which is the method used to compare the frequency / time standards of national laboratories around the world [51]. In this scheme, the transmitter and the receiver both compare their times simultaneously with that of a common GPS satellite that is in view for both. With knowledge about their relative distances to the satellite, their relative time difference can be determined, as can their relative frequency difference with subsequent measurements. Common-mode fluctuations in the path lengths to the satellite and the actual time of the satellite cancel out and do not play a role in the relative time / frequency measurement. By averaging for about a day it is possible to reach accuracies of a part in 10^{14} [25]. However, this technique is limited by fluctuations in the paths that are not common-mode. It does not provide the short-term stability necessary for synchronization applications, nor is it practical in situations such as the distribution of a frequency reference throughout a linear accelerator facility. Furthermore, it would require $\sim$30 days of averaging to reach stabilities high enough to accurately compare the best current optical frequency standards at a part in 10^{15} [86].

An extremely promising alternative for stable distribution of a frequency reference is the transmission of a reference over optical fibers. One attractive feature of optical fibers is that an environmentally isolated fiber can be considerably more stable than free-space paths, especially over short time scales. Also, the advantages that optical fibers offer for communications (for example, low loss and scalability) are also beneficial for a frequency distribution system. A great deal of the infrastructure also exists already for disseminating frequency references over telecommunication optical fibers. In this dissertation I will present my research that enables the use of optical frequency combs from mode-locked lasers for transferring over optical fibers frequency references that are linked to optical frequency standards. As we will see, transmitting an optical frequency comb over optical fiber provides several advantages over existing methods for transferring frequency references over fibers.

After introducing some fundamental concepts and some metrics for quantifying the stability and timing jitter of a transfer system in Chapter 2, I will discuss in Chapters 3 and 4 the key issues related to the distribution of the stability of an optical frequency standard to the entire visible spectrum using an optical frequency comb from a passively mode-locked Ti:sapphire laser. This is an important first step before using the comb to transmit a frequency reference over an optical fiber. However, it is not optimal to directly transfer the comb from a Ti:sapphire laser over the fiber. To minimize loss during fiber transmission it is necessary to transmit an optical frequency comb centered at a wavelength of 1550 nm instead of 800 nm, which is the operating wavelength of Ti:sapphire lasers. Therefore, in Chapter 5 I will describe the process of transferring the stability of the frequency comb of a Ti:sapphire laser to the comb of a mode-locked laser diode operating at 1550 nm. Finally, in Chapter 6 we will see the superior performance for transfer of a frequency reference by transmitting the optical frequency comb, as compared to existing methods of frequency transfer over fibers. I will present results for transmission with and without active cancellation of the noise from the transfer process, and we will see that the lowest timing jitter reported to date for the transfer of a frequency reference over several kilometers using an optical fiber network has been achieved.

Chapter 2

Fundamentals of Stabilizing an Optical Frequency Comb

We understand from the previous chapter that a phase-coherent link between an optical frequency and microwave frequencies is crucial for the implementation of an optical atomic clock. Also, for the transmission of a frequency reference derived from an optical clock, it is important to be able to transfer the stability of an optical frequency standard across the visible spectrum as well as to microwave frequencies. An optical frequency comb from a mode-locked laser can provide the phase-coherent connection among these spectral regions. By stabilizing the comb to the optical frequency standard, it provides an array of stable frequencies across the visible spectrum and in the microwave region. In this chapter I will discuss in more detail the properties of a frequency comb and what parameters must be controlled to stabilize it to an optical frequency standard.

2.1 Free parameters of an optical frequency comb

Since optical frequency combs are produced by mode-locked lasers, it is important to understand in detail the output of a mode-locked laser for a complete understanding of the frequency comb. A mode-locked laser is a laser that relies on some physical process to force all the longitudinal modes of the laser cavity to have a fixed phase relationship. All of the phase-locked modes then add up coherently to produce a train of short pulses from the laser, which has a repetition frequency determined by the round-trip time of the cavity, or equivalently, the mode spacing of the laser cavity. For Ti:sapphire lasers the mode-locking mechanism is the Kerr nonlinearity of the Ti:sapphire crystal, which also serves as the gain medium of the laser. The Kerr nonlinearity is a contribution to the refractive index of the crystal that depends on the intensity of the light field in the crystal. Therefore, the index of refraction is $n_0 + n_2 I$, where n_0 is the contribution from the linear response of the material to the light field and n_2 is a coefficient that indicates the strength of the nonlinearity, or how strongly the light intensity affects the refractive index. For a Gaussian beam propagating in the cavity, the higher-intensity center of the beam will experience a larger refractive index while passing through the crystal than the surrounding area of the beam. This is identical to a beam passing through a convex lens, which focusses the beam. With the appropriate geometry of the laser cavity, this self-focussing in the crystal maximizes overlap of the laser beam with the region in the crystal being inverted by the pumping laser. Since the laser can reduce its loss by operating with short, intense pulses that self-focus and maximally overlap with the pump beam, it self-mode locks and produces a train of ultrashort pulses that are on the order of 10 fs in duration. The shortness of the pulses is due to the fact that the Kerr nonlinearity operates on an extremely fast time scale of femtoseconds, allowing mode locking across a large optical bandwidth.

Figure 2.1(a) shows the output of a mode-locked laser in the time domain, while Fig. 2.1(b) shows the output in the frequency domain. From Fourier theory, it is clear that an ultrashort pulse will be represented by a broad range of frequencies, the extent of which is proportional to the inverse of the pulse width, τ_{FWHM}. However, for an infinite periodic train of pulses, the frequency spectrum will consist of discrete, infinitely narrow frequency components, creating an optical frequency comb instead of a continuous band of frequencies. These components, corresponding to the modes of the laser cavity, will be spaced by the repetition frequency of the laser, f_{rep}, as discussed above.

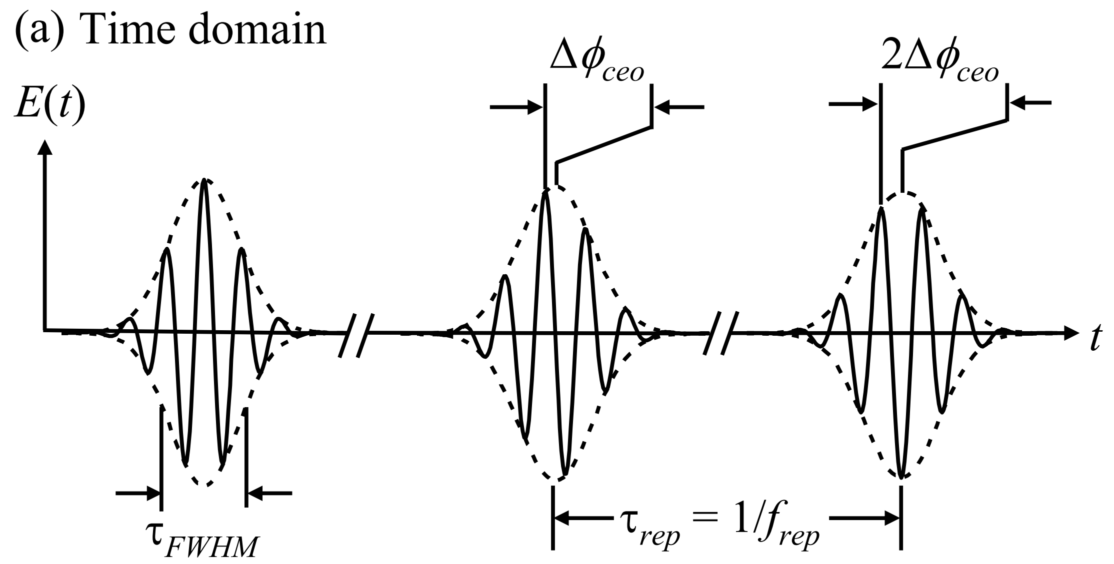

(a) Time domain

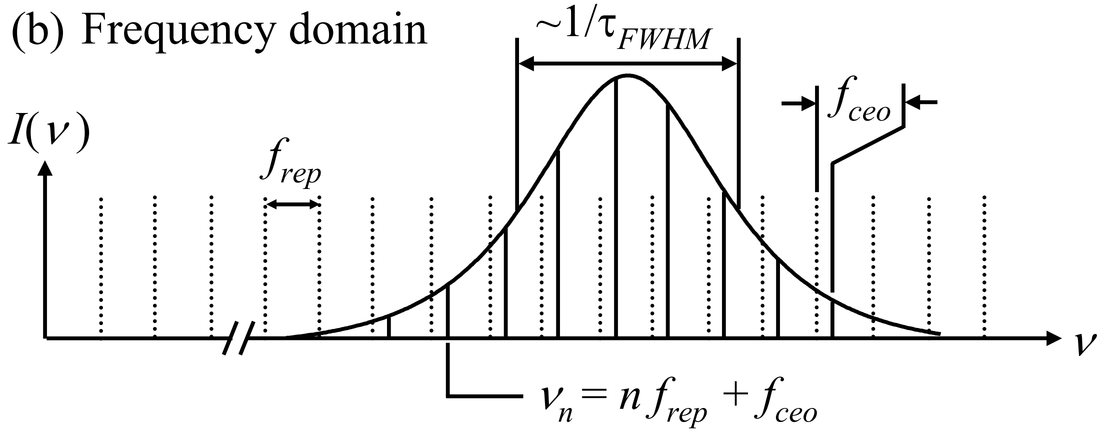

(b) Frequency domain

Figure 2.1: (a) The output of a mode-locked laser in the time domain. The laser produces a periodic train of ultrashort pulses having a duration of τ_{FWHM} and a repetition frequency of f_{rep}. The difference between the group and phase velocities in the laser cavity causes the peak of the oscillating carrier electric field to shift with respect to the peak of the envelope from one pulse to the next by an amount $\Delta\phi_{ceo}$, the carrier-envelope offset phase shift. (b) The corresponding frequency-domain output of a mode-locked laser. The periodic train of pulses produces a comb of discrete frequency components with an extent that is inversely proportional to the duration of each pulse and a spacing that is given by the repetition frequency of the pulse train. The carrier-envelope offset phase shift uniformly shifts the components of the optical frequency comb from integer multiples of the repetition frequency. This uniform shift, or carrier-envelope offset frequency, f_{ceo}, is proportional to $\Delta\phi_{ceo}$. Every component of the frequency comb is uniquely determined by the two degrees of freedom of the comb, f_{rep} and f_{ceo}.

Furthermore, because of dispersion in the laser cavity, the components of the frequency comb are not simply given by integer multiples of the repetition frequency. To see how dispersion affects the frequency comb, let us first consider its effect on the laser output in the time domain. In the laser cavity the pulse envelope propagates at the group velocity, v_g, whereas the oscillating carrier electric field travels at the phase velocity, v_p. The material dispersion in the cavity causes these two velocities to be different, and in the usual case of normal material dispersion $v_p > v_g$. Therefore, from one pulse to the next the peak of the electric field shifts with respect to the peak of the envelope, with the peak of the electric field arriving slightly earlier in time than that of the envelope as shown in Fig. 2.1(a). The additional phase that the carrier accumulates from the peak of the electric field to the peak of the envelope is referred to as the carrier-envelope offset phase shift, denoted as $\Delta\phi_{ceo}$, and can be expressed in terms of the group and phase velocities in the cavity. The time between the peaks of the electric field in consecutive pulses, τ_{phase}, is given by the time for the carrier to complete a cavity round trip of geometric length l_c.

$$\tau_{phase} = \frac{l_c}{v_p} \tag{2.1}$$

Likewise, the time between peaks of the envelope, τ_{rep}, is

$$\tau_{rep} = \frac{l_c}{v_g} \tag{2.2}$$

Therefore, the additional phase accumulated by the carrier can be expressed as

$$\begin{aligned} \Delta\phi_{ceo} &= \omega_c\left(\tau_{rep} - \tau_{phase}\right) \\ &= \omega_c l_c \left(\frac{1}{v_g} - \frac{1}{v_p}\right) \end{aligned} \tag{2.3}$$

where ω_c is the angular frequency of the carrier.

The effect of $\Delta\phi_{ceo}$ on the frequency comb can be understood by considering that during one pulse period, each frequency component of the laser, ν_n, accumulates a phase that is the sum of an integer multiple of 2π and $\Delta\phi_{ceo}$.

$$2\pi\nu_n\tau_{rep} = 2\pi n + \Delta\phi_{ceo} \tag{2.4}$$

which can be rearranged to yield

$$\nu_n = nf_{rep} + f_{rep}\frac{\Delta\phi_{ceo}}{2\pi} \tag{2.5}$$

It is clear from this equation that $\Delta\phi_{ceo}$ uniformly shifts all the frequency components from the laser by an amount f_{ceo} as indicated in Fig. 2.1(b), where

$$f_{ceo} \equiv f_{rep}\frac{\Delta\phi_{ceo}}{2\pi} \tag{2.6}$$

This uniform shift is referred to as the carrier-envelope offset frequency. The repetition frequency of the laser and the carrier-envelope offset frequency are the only free parameters of the frequency comb. Once these two degrees of freedom are stabilized, every component of the frequency comb is determined by

$$\nu_n = nf_{rep} + f_{ceo} \tag{2.7}$$

Two alternative methods for deriving the structure of the frequency comb can be found in [11].

2.2 Stabilization of carrier-envelope offset frequency

From the previous section we can see that stabilizing every component of the optical frequency comb to an optical frequency standard can be accomplished by stabilizing two microwave frequencies, f_{ceo} and f_{rep}, to the optical reference. Actually, the first of these, f_{ceo}, can be stabilized independently using a self-referencing technique [44], without making use of the optical frequency reference. However, this technique is most easily implemented when the comb spans an entire octave of frequencies. For the frequency comb from a Ti:sapphire laser centered at a wavelength of 800 nm, this corresponds to it spanning from $\sim$ 500 nm to 1 μm, or $\sim$ 600 to 300 THz. Typically Ti:sapphire mode-locked lasers do not produce an output spanning an entire octave, and so a nonlinear process is necessary that can generate additional frequencies to broaden the comb to an octave while preserving the structure of the frequency comb. (In recent years a few Ti:sapphire systems have been developed with octave-spanning outputs [18, 6], enabling the stabilization of f_{ceo} without the need for a spectrum-broadening nonlinear process [23].) To broaden the spectrum from a typical Ti:sapphire laser, the pulses from the laser are coupled into a specially designed air-silica microstructure optical fiber [65]. The fiber consists of a solid silica core with an ~2-μm diameter surrounded by an array of ~1.7-μm-diameter air holes in a hexagonal close-packed configuration. The large index contrast between the silica core and the air cladding guides the pulses within the core of the fiber. Also, the waveguide contribution to the group-delay dispersion (GDD) significantly alters the GDD of the fiber from that of bulk silica and results in zero-GDD near the center wavelength of the Ti:sapphire laser (typically at ~760 nm) and negative (anomalous) GDD at its center wavelength. This allows the propagation of the pulses through the microstructure fiber for distances of tens of centimeters without significant temporal stretching of the pulses. Therefore, the ultrashort pulses that are confined within the small core of the fiber exhibit very high peak intensities and can experience strong nonlinear interactions with the silica core over the full length of the fiber. Through four-wave mixing and Raman scattering, the input spectrum is broadened to span an octave while preserving the comb structure.

With a frequency comb that spans an entire octave, f_{ceo} can be measured using the self-referencing scheme shown in Fig. 2.2. A comb component in the infrared region of the spectrum, at frequency $\nu_n = nf_{rep} + f_{ceo}$, is selected and frequency-doubled using second-harmonic generation in a crystal. The second-harmonic light has a frequency of $2\nu_n = 2nf_{rep} + 2f_{ceo}$. This second-harmonic light is then combined with a component in the high-frequency region of the spectrum at frequency $\nu_{2n} = 2nf_{rep} + f_{ceo}$ and the two co-propagating beams are focussed onto a photodetector. The signal from the

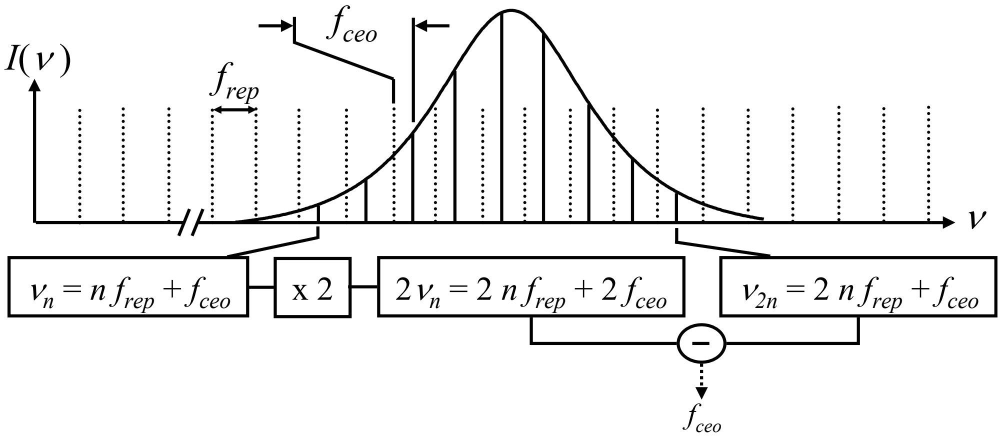

Figure 2.2: The carrier-envelope offset frequency is found by taking the difference of the frequencies of a comb mode at ν_{2n} and the second harmonic of a mode at ν_n.

photodetector oscillates at the difference (heterodyne beat) frequency between the two optical frequencies, which is exactly the microwave frequency f_{ceo}.

Figure 2.3 shows a simplified experimental setup for the measurement and subsequent stabilization of f_{ceo}. After passing the output from the mode-locked Ti:sapphire laser through the microstructure fiber to broaden the frequency comb to one octave, a dichroic mirror is used to separate the ν_n and ν_{2n} comb components. An ~3-mm-thick beta-barium-borate (BBO) crystal is used to generate the second harmonic of the ν_n component, which rotates the polarization by 90° since the phase matching of the crystal is Type I. This allows the two frequencies at ν_{2n} and $2\nu_n$ to be re-combined using a polarizing beamsplitter. An interference filter is used to pass only frequencies of the comb near ν_{2n}, and a polarizer is necessary before the photodetector to pass the common projections of the two orthogonal polarizations, since a heterodyne beat between two frequencies requires a common polarization state.

Once f_{ceo} is measured, it can be stabilized by phase-locking it to a microwave reference as shown in Fig. 2.3. By mixing the f_{ceo} and reference signals phase coherently in a double-balanced mixer (this is mathematically equivalent to multiplying the signals), two oscillatory signals are produced — one with a phase term that is the sum of the

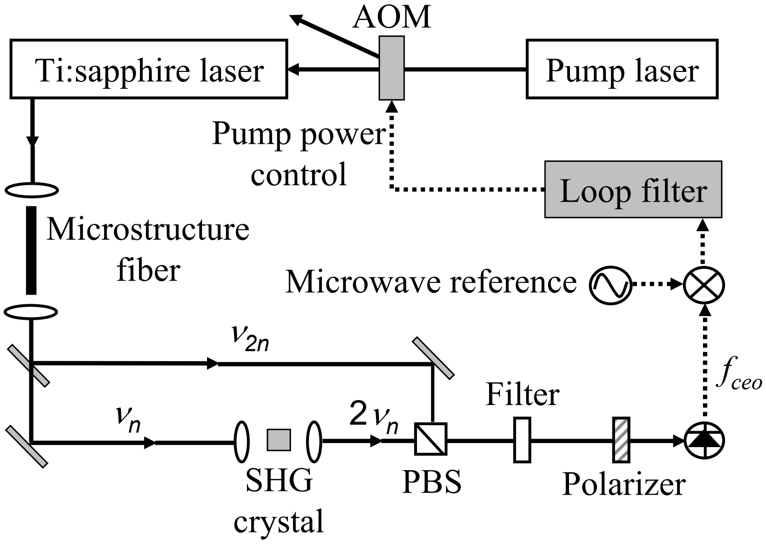

Figure 2.3: The carrier-envelope offset frequency is measured by combining the comb component at ν_{2n} with the second harmonic of the component at ν_n on a photodetector, which produces the heterodyne beat between the frequencies. It is phase-locked to a microwave reference by actuating the pump power using an acousto-optic modulator (AOM) in the pump beam. SHG, second-harmonic generation; PBS, polarizing beamsplitter.

phase terms of the two input signals, and one with a phase term given by the difference of the phases of the two input signals. By locking the difference signal to 0 V, the f_{ceo} signal is forced to track the phase of the microwave reference. The difference signal is filtered and amplified to produce a correction signal with the appropriate phase and gain to feed back to the laser and phase-lock f_{ceo} to the reference. The feedback is applied to the intensity of the Ti:sapphire laser, which is actuated by placing an acousto-optic modulator (AOM) in the pump beam and changing the power of the radio-frequency (RF) signal driving the AOM. This changes the amount of power diffracted by the AOM, thereby allowing adjustment of the power in the zeroth order beam used to pump the Ti:sapphire laser. We will see in Chapter 3 how the Ti:sapphire intensity affects the value of f_{ceo}.

It is important to note that although the offset frequency is stabilized to a microwave reference, its stability is sufficient to allow the optical frequency comb to be stabilized to even the best optical frequency standards. The best optical clocks are expected to reach short-term instabilities near 10^{-17} for a 1-s averaging time and uncertainties near 10^{-18}, which corresponds to 4 mHz and 0.4 mHz, respectively, for a 400 THz optical standard. To stabilize the optical frequency comb to an optical standard, the fluctuation of its offset should be limited to this amount. However, the measured value of f_{ceo} is typically ~100 MHz, and so 4 mHz and 0.4 mHz correspond to a 1-s instability and an uncertainty of 4×10^{-11} and 4×10^{-12}, respectively. This level of performance is easily attained by using as the microwave reference for f_{ceo} a synthesizer referenced to a commercial Cs beam clock.

2.3 Stabilization of repetition frequency: Locking comb to optical or microwave standards

Once the offset frequency of the frequency comb is stabilized, the only remaining degree of freedom is the spacing of the comb components, or the laser repetition frequency, f_{rep}. Using an optical frequency standard to stabilize f_{rep} transfers the stability of the optical standard to every component of the comb. Fig. 2.4 illustrates how an optical frequency reference can be used to stabilize f_{rep}. First, the optical frequency reference is combined with light from the comb onto a photodetector to produce the heterodyne beat, f_{beat}, between the optical reference, ν_{ref}, and the nearest comb component, ν_m. Using the broadened comb produced by the microstructure fiber allows it to be stabilized to any optical frequency reference across the visible spectrum. Fluctuations in f_{beat} are equivalent to fluctuations in the frequency of this comb component. However, because f_{ceo} has been independently stabilized, fluctuations in ν_m and f_{beat} can only be caused by changes in f_{rep}. Therefore, phase-locking f_{beat} to a microwave reference results in the stabilization of f_{rep} to the optical reference. Once an error signal is derived by mixing the f_{beat} signal with the microwave reference, it is applied to a piezoelectric transducer (PZT) that translates a mirror in the laser cavity and adjusts the cavity length, allowing stabilization of f_{rep}. Just as phase-locking f_{ceo} to a microwave reference provides sufficient stability for the stabilization of the frequency comb to an

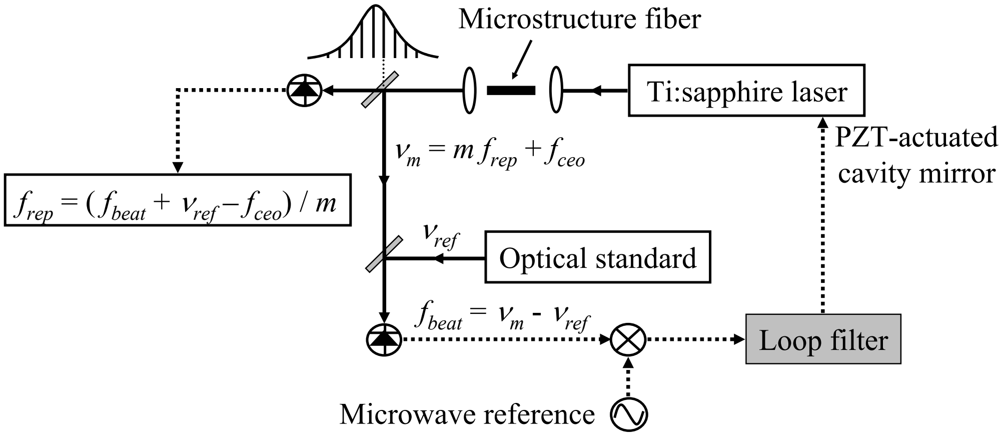

Figure 2.4: Once f_{ceo} is independently stabilized, the repetition frequency is stabilized to an optical standard by phase-locking the heterodyne beat between the optical standard and the nearest comb component to a microwave reference. It is controlled with a PZT-actuated mirror in the laser cavity. This provides f_{rep} and every optical frequency component of the comb with the same fractional instability as the optical standard.

optical standard, the instability of the microwave reference used to phase-lock f_{beat} is also not an issue. The requirement to limit the fluctuation of ν_m to 4 mHz for a 1-s averaging time and 0.4-mHz for long averaging times, calculated in the previous section, also corresponds to a 1-s instability and an uncertainty of only 4×10^{-11} and 4×10^{-12}, respectively, for the ~100 MHz f_{beat} signal.

With f_{rep} stabilized to an optical frequency standard in this way, the optical frequency comb provides a phase-coherent connection between the optical frequency standard and microwave frequencies, since f_{rep} exhibits the same fractional instability as the optical reference. To see this, consider that

$$\begin{aligned} f_{beat} &= \nu_m - \nu_{ref} \\ &= m f_{rep} + f_{ceo} - \nu_{ref} \end{aligned} \tag{2.8}$$

as shown in Fig. 2.4. Rearranging this, we have

$$f_{rep} = \frac{1}{m}\left(f_{beat} + \nu_{ref} - f_{ceo}\right) \tag{2.9}$$

Since it has already been shown that f_{ceo} and f_{beat} can be stabilized sufficiently well to ignore their contributions to the instability of the frequency comb, we have $\delta f_{rep} = \delta\nu_{ref}/m$. Therefore, $\delta f_{rep}/f_{rep} \approx \delta\nu_{ref}/\nu_{ref}$ since $\nu_{ref} \gg f_{beat}$ and f_{ceo} in Eqn. (2.9). Also, note that every optical frequency component of the comb also exhibits the same fractional instability as the optical reference, since $\nu_n \approx n f_{rep}$ and so $\delta\nu_n/\nu_n \approx \delta f_{rep}/f_{rep} \approx \delta\nu_{ref}/\nu_{ref}$. Therefore the frequency comb also provides a link to distribute the stability of an optical frequency standard across the visible spectrum. Experiments have verified that the frequency comb can be stabilized to an optical reference with a residual fractional frequency uncertainty for each comb component relative to the optical reference of close to a part in 10^{19} [53].

For situations where the comb is to be used for generating a microwave frequency that is phase coherent with an optical frequency standard but the individual components of the comb do not need to be stabilized, it is possible to stabilize f_{rep} to the optical reference without stabilizing f_{ceo}. One method for accomplishing this requires the comb to extend from the optical frequency reference to the second harmonic of this reference. It is possible to electronically combine the two heterodyne beats between the frequency comb and the fundamental and second harmonic of the reference in such a way so as to obtain a control signal that depends only on f_{rep} and the frequency of the optical reference, and not f_{ceo} [85]. A second method for stabilizing f_{rep} without controlling f_{ceo} involves difference-frequency generation from the high-frequency and low-frequency regions of the comb, which results in the cancellation of f_{ceo} and the creation of an infrared (IR) comb with a null f_{ceo}. By stabilizing the heterodyne beat between the IR comb and an IR optical frequency standard, f_{rep} is phase-locked to the optical standard [21].

Finally, the frequency comb can be used to transfer the stability of a microwave reference directly to optical frequencies. Though this will not provide the stability of locking the comb to an optical frequency standard, this can be useful for making optical frequency measurements that are linked to the Cs primary standard. In this scenario, instead of phase-locking f_{beat}, f_{rep} is phase-locked directly to the microwave frequency standard and the frequency of f_{beat} is measured with a frequency counter. The frequency ν_{ref}, which in this case represents the frequency of the unknown optical frequency, is determined by rearranging Eqn. (2.8) to obtain

$$\nu_{ref} = m f_{rep} + f_{ceo} - f_{beat} \tag{2.10}$$

In Chapter 4 I will discuss some experimental results for using the comb to link optical and microwave frequencies to measure optical frequency transitions in two different atomic species.

2.4 Transferring stability of one frequency comb to another

Although optical frequency combs from mode-locked Ti:sapphire lasers enable the distribution of the stability of an optical frequency standard across the visible spectrum, for the remote transfer of an optical frequency reference over optical fibers it is necessary to use a stabilized comb centered at a wavelength of 1550 nm for the reasons discussed in Section 1.6. However, depending on the source of the 1550-nm frequency comb, it is not always possible to directly stabilize the 1550-nm comb to the optical frequency standard. For example, a mode-locked laser diode produces a comb with a bandwidth of only ∼1 nm, which prevents the use of the self-referencing technique to stabilize f_{ceo}. In these cases it is necessary to stabilize the 1550-nm comb to the Ti:sapphire comb that is referenced to the optical standard.

The stabilization of one optical frequency comb (the slave comb) to another (the master comb) requires two conditions be met, which correspond to the two degrees of freedom of the comb just as for the stabilization of a single comb [69]. The stabilization of the slave comb to the master comb is accomplished by using the stability of the master comb to stabilize the repetition frequency and carrier-envelope offset frequency of the slave comb. Figure 2.5 shows the scheme for stabilizing a 1550-nm comb to a Ti:sapphire comb, where f_{rep}^{mst} and f_{ceo}^{mst} denote the degrees of freedom of the master comb, and f_{rep}^{slv} and f_{ceo}^{slv} denote the parameters for the slave comb. The repetition frequency of the slave comb is stabilized by phase-locking the nth harmonic of f_{rep}^{slv} with the mth harmonic of f_{rep}^{mst}. Locking the harmonics of the repetition frequencies allows for some freedom in their relationship, requiring them to simply have a rational ratio, while providing f_{rep}^{slv} with the same fractional instability as f_{rep}^{mst}. The harmonics are obtained directly from the photodetectors that detect the pulse trains for each comb. The generation of several harmonics of the repetition frequency by the photodetector can be explained using either a time-domain or a frequency-domain description. In the time domain, the

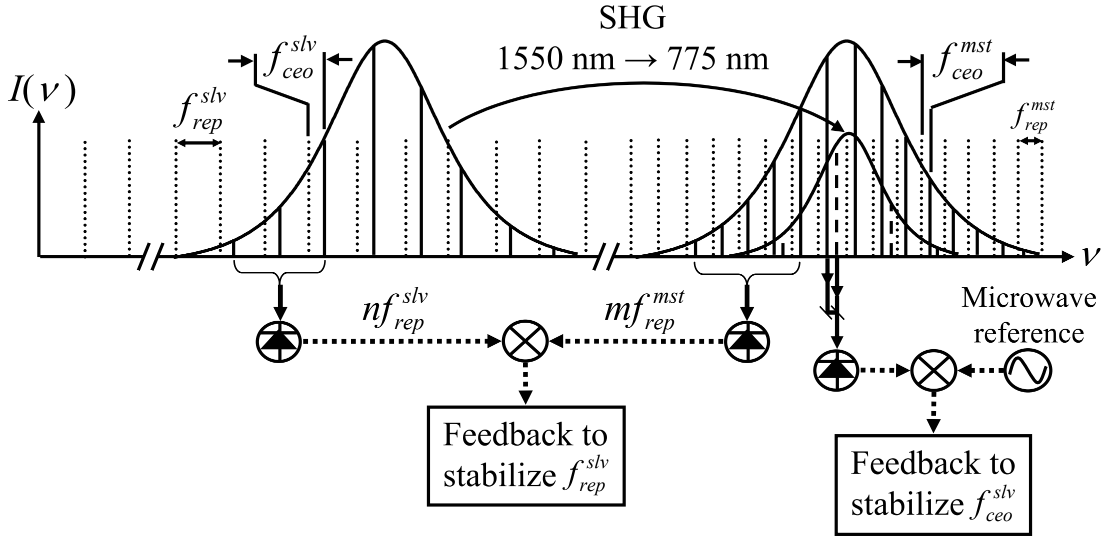

Figure 2.5: One optical frequency comb (the slave) can be stabilized to another (the master) by phase-locking harmonics of their repetition frequencies and by stabilizing the heterodyne beat between components from each comb to a microwave reference. For detection of the heterodyne beat between a 1550-nm comb and an 800-nm Ti:sapphire comb, the 1550-nm comb is frequency doubled using second-harmonic generation (SHG) to achieve spectral overlap.

harmonics of the repetition frequency represent all the Fourier components necessary to describe the train of ultrashort pulses, which are much too short to be resolved by the photodetector. Having only the fundamental repetition frequency with no higher-order harmonics would correspond to a sinusoidal amplitude modulation of the input to the photodetector, instead of a periodic train of very short pulses. In the frequency domain, the repetition frequency corresponds to the heterodyne beat between adjacent modes of the optical frequency comb, and the nth harmonic corresponds to the beat between every nth comb component. The number of harmonics is limited by the bandwidth of the detector, which is typically a few gigahertz. A microwave bandpass filter is then used to select the harmonic of interest for each comb. Phase-locking the nth harmonic of f_{rep}^{slv} to the mth harmonic of f_{rep}^{mst} corresponds in the time domain to synchronizing every mth pulse of the slave comb's pulse train to every nth pulse of the master comb's pulse train.

With the spacing between components stabilized for the slave comb, its offset frequency can be fixed by stabilizing the optical frequency of one of its components. This is accomplished by phase-locking the heterodyne beat between components from the slave and master comb to a microwave reference. However, for stabilizing a 1550-nm comb to a Ti:sapphire comb, there is no spectral overlap between the combs. Therefore, to detect a heterodyne beat between components from each comb it is necessary to frequency-double the 1550-nm comb using second-harmonic generation, as shown in Fig. 2.5. Every component of the slave comb, as well as its repetition frequency, then has the same fractional frequency instability as the optical frequency reference to which the master comb is stabilized. In Chapter 5 I will discuss the experimental details and results for stabilizing a frequency comb from a mode-locked laser diode to that of a Ti:sapphire laser.

2.5 Characterization of frequency stability, phase noise, and timing jitter

Now that we have seen how to stabilize frequency combs operating at various wavelengths to an optical frequency standard, it is important to consider how to quantitatively characterize the stability of a frequency source. This is crucial for determining how faithfully the comb distributes the stability of the frequency standard across the optical spectrum and down to microwave frequencies, as well as for analyzing the instabilities introduced when a frequency reference is transmitted long distances over optical fibers. There are typically two classes of characterization of a frequency source. The first involves calculating how the measured fractional frequency fluctuations of the source vary as a function of the time over which the frequency is averaged and is expressed using the Allan deviation. The second method for characterizing frequency stability is to measure the phase noise, which represents how much the phase of a signal is jittering. For analyzing the remote transfer of a microwave frequency reference, it is common to express this as timing jitter. Whereas the phase noise of the transmitted reference depends on the frequency of the reference, the timing jitter depends only on the magnitude of the environmental perturbations that are inducing phase noise during the transfer.

The Allan deviation, $\sigma_y(\tau)$, can be computed from a series of consecutive frequency measurements, each obtained by averaging over a period of time τ. This averaging time corresponds to the gate time of the frequency counter used to make the frequency measurements. To see how $\sigma_y(\tau)$ can be computed from these frequency measurements, it is necessary to make some preliminary definitions. The signal from the frequency source can be expressed as

$$V(t) = V_0 \cos[2\pi\nu_0 t + \phi(t)] \tag{2.11}$$

where V_0 is the amplitude of the signal, ν_0 is its nominal center frequency, and $\phi(t)$ represents time-varying deviations from the nominal phase $2\pi\nu_0 t$. The instantaneous frequency is then given by

$$\nu(t) = \nu_0 + \frac{1}{2\pi}\frac{\mathrm{d}}{\mathrm{d}t}\phi(t) \tag{2.12}$$

and the instantaneous fractional frequency deviation from the nominal center frequency is given by

$$y(t) = \frac{1}{2\pi\nu_0}\frac{\mathrm{d}}{\mathrm{d}t}\phi(t) \tag{2.13}$$

The Allan deviation for an averaging time τ is then defined as

$$\sigma_y(\tau) \equiv \left\langle \frac{1}{2}[\bar{y}(t+\tau) - \bar{y}(t)]^2 \right\rangle^{1/2} \tag{2.14}$$

where $\langle\rangle$ indicates an infinite time average and $\bar{y}$ represents the time average of $y(t)$ over a period τ [80]. $\sigma_y(\tau)$ can be estimated from a finite set of N consecutive average values of the fractional frequency deviation, $\bar{y}_i$.

$$\sigma_y(\tau) = \left[\frac{1}{2(N-1)}\sum_{i=1}^{N-1}(\bar{y}_i - \bar{y}_{i+1})^2\right]^{1/2} \tag{2.15}$$

Therefore,

$$\begin{aligned}
\sigma_y(\tau) &= \left\{\frac{1}{2(N-1)}\sum_{i=1}^{N-1}[(1+\bar{y}_i) - (1+\bar{y}_{i+1})]^2\right\}^{1/2} \\
&= \left\{\frac{1}{2(N-1)}\sum_{i=1}^{N-1}\left\{\frac{1}{\nu_0}\left[(\nu_0+\nu_0\bar{y}_i) - (\nu_0+\nu_0\bar{y}_{i+1})\right]\right\}^2\right\}^{1/2} \\
&= \left[\frac{1}{2(N-1)\nu_0^2}\sum_{i=1}^{N-1}(\bar{\nu}_i - \bar{\nu}_{i+1})^2\right]^{1/2}
\end{aligned} \tag{2.16}$$

where $\bar{\nu}_i$ denote consecutive measurements of the average signal frequency, averaged over a period τ. The Allan deviation for averaging times that are integer multiples of τ, $\sigma_y(m\tau)$, can then be calculated by forming a new set of N/m average frequency values from the original set of N values. The original set is sub-divided into adjacent, non-overlapping subsets. Each value of the new set of frequency values is computed by averaging the m values in each subset of the original data.

The Allan deviation is extremely useful for characterizing a frequency source because the type of phase noise present is revealed by the way in which $\sigma_y(\tau)$ depends on τ. For example, if $\sigma_y(\tau) \propto \tau^{-1}$ then white phase noise is the dominant noise process, whereas if $\sigma_y(\tau) \propto \tau^{-1/2}$ then white frequency noise is dominant [80]. However, for the Allan deviation to reliably indicate the type of noise present, it is crucial that there be no dead time between the consecutive average frequency measurements used to calculate $\sigma_y(\tau)$. The presence of dead time may bias the computed Allan deviation, depending on the type of phase noise present. For example, since dead time will result in a loss of coherence between data points, white phase noise may be detected as white frequency noise, depending on the length of the dead time.

In practice, the individual frequency measurements used to compute the Allan deviation are obtained with a frequency counter that is referenced to an auxiliary frequency source. Therefore, the measured instability is actually the relative instability between the frequency source being measured and the frequency reference, and is ultimately limited by the stability of the frequency reference and the resolution of the counter. However, if a frequency reference is available with a stability exceeding the resolution of the counter, it is still possible to use the counter to measure the relative instability between the frequency source and this more stable reference. The stable reference is used to synthesize a frequency exhibiting the same stability as the reference that is close to the frequency being measured. The synthesized frequency and the frequency being measured are then mixed, and their difference frequency is measured with the frequency counter. Since the counter is measuring a signal with a smaller frequency, its resolution is sufficient to detect small fluctuations of this frequency, which represent the same magnitude of fluctuations and a much lower fractional instability of the original frequency being measured. To illustrate with some specific numbers, suppose we have a frequency at 1 GHz, and we want to measure its fractional instability relative to another reference at a level of 10^{-15}. This corresponds to a deviation of 1 μHz. Measuring 1 μHz out of 1 GHz far exceeds the resolution of typical counters, but instead the reference is used to synthesize a frequency at 1 GHz + 10 kHz, which is mixed with the 1 GHz signal to produce a 10 kHz signal that is measured with the frequency counter. A 1 μHz deviation of the 1 GHz signal corresponds to a 1 μHz deviation of the 10 kHz signal, but only a fractional change of the 10 kHz signal of 10^{-10}, which can be measured easily with a good quality counter referenced to a commercial Cs beam clock.

The second method for characterizing the stability of a frequency source — directly measuring its phase fluctuations — is especially useful for determining the stability of a signal that is to be used for the synchronization of various system components, such as in the applications mentioned in Section 1.6. The phase noise is typically used to characterize only microwave-frequency signals, and is determined by measuring the phase fluctuations of a signal with respect to a frequency reference. If ν_0 is the frequency of the reference and the nominal center frequency of the signal to be measured, $\phi(t)$ represents deviations from $2\pi\nu_0 t$ in the phase of the signal, and ϕ_0 is an arbitrary fixed phase offset of the reference, then the reference frequency signal is $\propto \cos(2\pi\nu_0 t + \phi_0)$ and the signal to be measured is $\propto \cos(2\pi\nu_0 t + \phi(t))$. The deviations of the phase of the signal, $\phi(t)$, are measured with respect to the phase of the reference signal by mixing the two signals. The term in the mixing product which has a phase term given by the difference of the phase terms of the two mixed signals is isolated using a low-pass filter, which produces a signal that can be expressed as

$$A(t) = A_0 \cos[\phi(t) - \phi_0] \tag{2.17}$$

where A_0 is the amplitude of the mixer output and depends on the amplitudes of the two mixed signals. A phase-shifter in the signal path of the reference frequency before the mixer is used to set $\phi_0 = \pi/2$, which reduces $A(t)$ to

$$A(t) = A_0 \sin[\phi(t)] \tag{2.18}$$

Typically the deviations of $\phi(t)$ from zero are sufficiently small to allow the relationship between a deviation from zero of the phase error of the measured signal, $\Delta\phi$, and the deviation from zero of the voltage of the mixer output, ΔA, to be determined by

$$\begin{aligned}\frac{\Delta\phi}{\Delta A} &= \left(\left.\frac{\mathrm{d}A}{\mathrm{d}\phi}\right|_{\phi=0}\right)^{-1} \\ &= A_0^{-1} \qquad (2.19)\end{aligned}$$

A_0 is determined by using the phase-shifter in the path of the reference signal to adjust ϕ_0 from 0 to π and measuring the maximum voltage swing of the mixer output, $A(t)$. A fast Fourier transform (FFT) is then used to compute the Fourier transform of the mixer output, $\tilde{A}(f)$, which gives the root-mean-squared (rms) fluctuation of A about zero at each Fourier frequency f in a measurement bandwidth of 1 Hz and has units of Volts/$\sqrt{\text{Hz}}$. The rms fluctuation in the phase at each Fourier frequency in a 1-Hz measurement bandwidth, $\delta\tilde{\phi}(f)$, is computed using Eqn. (2.19).

$$\delta\tilde{\phi}(f) = \frac{\tilde{A}(f)}{A_0} \qquad \left[\frac{\text{rad}}{\sqrt{\text{Hz}}}\right] \qquad (2.20)$$

It is common to express the phase noise of a frequency source as the power spectral density (PSD) of phase fluctuations, $\mathrm{S}_\phi(f)$, which represents the mean squared phase fluctuation at Fourier frequency f from the carrier in a measurement bandwidth of 1 Hz [80]. It is defined as

$$\mathrm{S}_\phi(f) \equiv [\delta\tilde{\phi}(f)]^2 = \left[\frac{\tilde{A}(f)}{A_0}\right]^2 \qquad \left[\frac{\text{rad}^2}{\text{Hz}}\right] \qquad (2.21)$$

Note that $\mathrm{S}_\phi(f)$ includes contributions from both the upper and lower side bands of the carrier, since the Fourier transform of the mixer output, $\tilde{A}(f)$, folds the negative portion of the Fourier spectrum of the phase noise into the positive range of frequencies. Another quantity that is often used to specify the phase noise of a frequency source is the single side band phase noise, $\mathrm{L}(f)$, defined as $\frac{1}{2}\mathrm{S}_\phi(f)$ [80]. It is usually expressed in units of dBc/Hz, which means dB below the carrier in a 1-Hz measurement bandwidth, and can be computed from $\mathrm{S}_\phi(f)$ as

$$\mathrm{L}(f) = 10\log\left[\frac{1}{2}\mathrm{S}_\phi(f)\right] \qquad \left[\frac{\text{dBc}}{\text{Hz}}\right] \qquad (2.22)$$

It is also possible to relate these measurements of the phase noise to the Allan deviation, $\sigma_y(\tau)$, through the expression

$$\sigma_y(\tau) = \left(\frac{\sqrt{2}}{\pi\nu_0\tau}\right)\left[\int_0^{\infty} \mathrm{H}_\phi(f)\mathrm{S}_\phi(f)\sin^4(\pi f\tau)\,\mathrm{d}f\right]^{1/2} \tag{2.23}$$

where $\mathrm{H}_\phi(f)$ denotes the transfer function of the system used to measure the Allan deviation [80]. In practice, $\mathrm{H}_\phi(f)$ often has a low-pass behavior and rolls off sharply at some maximum Fourier frequency f_h. In this case, $\mathrm{H}_\phi(f)$ can be approximated by 1 for $f < f_h$ and by 0 for $f > f_h$, allowing the previous equation to be rewritten as

$$\sigma_y(\tau) = \left(\frac{\sqrt{2}}{\pi\nu_0\tau}\right)\left[\int_0^{f_h} \mathrm{S}_\phi(f)\sin^4(\pi f\tau)\,\mathrm{d}f\right]^{1/2} \tag{2.24}$$

When characterizing the remote transfer of a microwave frequency reference, it is more appropriate to consider the timing jitter of the transmitted signal instead of its phase noise. The typical noise processes that affect the transfer stability are due to environmental perturbations that alter the optical path length of the transmitted signal. The timing jitter is dependent only on the magnitude of these environmental perturbations, whereas the phase noise will be proportional to the frequency of the transmitted signal for a given perturbation. Also, for remote synchronization applications it is more appropriate to discuss the timing jitter for the transfer. The timing jitter spectral density, $\delta\tilde{T}(f)$, which represents the rms timing jitter at each Fourier frequency in a 1-Hz measurement bandwidth, is proportional to $\delta\tilde{\phi}(f)$. Since $\phi = 2\pi\nu_0 t$, then

$$\delta\tilde{T}(f) = \frac{\delta\tilde{\phi}(f)}{2\pi\nu_0} = \frac{\tilde{A}(f)}{2\pi\nu_0 A_0} \qquad \left[\frac{\text{seconds}}{\sqrt{\text{Hz}}}\right] \tag{2.25}$$

where Eqn. (2.20) has been used to express $\delta\tilde{T}(f)$ in terms of the FFT of the mixer output. The total rms timing jitter, T_{rms}, over a bandwidth from f_l to f_h, is then determined by

$$T_{rms} = \sqrt{\int_{f_l}^{f_h}\left[\delta\tilde{T}(f)\right]^2\,\mathrm{d}f} \qquad [\text{seconds}] \tag{2.26}$$

The limits of integration are determined by the specific application for which the frequency reference is being transferred. A higher sampling rate of the transmitted reference corresponds to a higher bandwidth over which the jitter spectral density must be integrated, and thus a higher value for f_h. Therefore, to achieve a small rms timing jitter over extremely short timescales, it is important that the high-frequency timing jitter be as small as possible.

In the specific situation where the pulse trains from two mode-locked lasers are to be synchronized, an alternate method for detecting the timing jitter between the two trains is available. This method involves measuring the cross-correlation between pulses from each train by overlapping the pulse trains in a nonlinear crystal and detecting the sum-frequency signal. The pulse trains can either be co-propagated through the crystal, or they can be overlapped in a crossed geometry. The latter configuration results in a sum-frequency signal that exits the crystal in a different direction than the original pulse trains. The intensity of the sum-frequency signal at a given time is proportional to the product of the intensities of the two pulse trains at that time. Therefore, in addition to oscillating at the repetition frequency of the lasers, the amplitude of the sum-frequency signal depends on the temporal overlap of the two pulse trains. By setting the temporal offset of the two trains such that the pulses overlap roughly where they attain half their maximum intensity, as shown by the pulse envelopes drawn as solid curves in Fig. 2.6, the amplitude of the sum-frequency signal is maximally sensitive to timing jitter between the pulse trains. The sum-frequency signal is illustrated with the dashed curve in Fig. 2.6. Introducing a known amount of delay between the two trains allows calibration of the sum-frequency signal, and variations in its amplitude provide a very sensitive method for measuring timing jitter between the pulse trains. The steep slope of pulses that have a duration on the order of 10–100 fs makes this technique for detecting timing jitter much more sensitive than the method previously discussed that depends on mixing two microwave frequencies, which is typically limited by noise in the mixer, the photodetectors that detect the repetition frequencies of the pulse trains and produce the microwave signals, and the amplifiers that are necessary to amplify

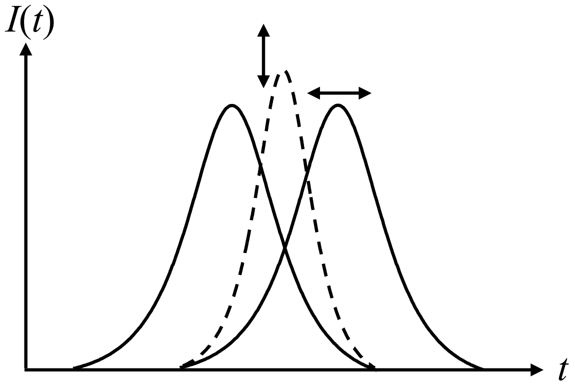

Figure 2.6: The timing jitter between the pulse trains from two mode-locked lasers can be detected by measuring their cross-correlation. The amplitude of the sum-frequency signal (dashed curve) produced by a nonlinear crystal is maximally sensitive to the relative timing jitter between the pulses (solid curves) when the temporal offset of the pulses is set such that they overlap roughly where they attain half their maximum intensity.

these microwave signals prior to mixing. This optical cross-correlation measurement technique can also be implemented with a two-photon detector instead of sum-frequency generation, since the two-photon-detector output is also proportional to the product of the incident pulse intensities at a given time. The cross-correlation technique is useful when the stability of one optical frequency comb is being transferred to another, as described in Section 2.4 [68]. It also has a potential application to the remote transfer of a microwave frequency reference over an optical fiber using the pulse train of a mode-locked laser. It could be used to detect the timing jitter introduced during transfer, and to synchronize a mode-locked laser at the remote end to the transmitted pulse train, provided sufficient peak power is present after fiber transmission.

Chapter 3

Intensity-related Dynamics of a Ti:sapphire Laser

The stabilization of the carrier-envelope offset frequency, f_{ceo}, of a mode-locked Ti:sapphire laser is a crucial step for using its frequency comb to distribute the stability of an optical frequency standard across the optical spectrum and to microwave frequencies for remotely transferring a frequency reference. Since f_{ceo} is controlled and stabilized using the intensity of the laser, as described in Section 2.2, it is important to understand how the intensity of the Ti:sapphire laser couples to its two degrees of freedom, f_{rep} and f_{ceo}. Understanding this coupling helps to identify a regime of operation where the free-running instability of f_{ceo} is minimized and enables optimization of the scheme for stabilizing f_{ceo}.

Previously, the effect of the laser intensity on f_{ceo} has been explained in terms of the dispersion of the nonlinear index of refraction of the Ti:sapphire crystal (also referred to as self-steepening) and beam pointing variations within the laser cavity caused by nonlinear refraction at the Ti:sapphire crystal surface [30]. However, we have observed that the dependence of f_{rep} and f_{ceo} on intensity is not monotonic. Instead, f_{rep} and f_{ceo} reach extrema as the intensity is increased and for higher intensities their dependence on the laser intensity experiences a sign reversal. This turning point cannot be explained by these previous coupling models, nor is the explanation in terms of nonlinear beam steering applicable for the particular Ti:sapphire laser we are using. In this chapter I will discuss the results of detailed experimental investigations of the intensity-related dynamics of both f_{rep} and f_{ceo}, in which we have found that the dominant mechanism for coupling the laser intensity to these parameters is an intensity-related spectral shift, as previously suggested by Xu et al [83]. The turning point in the dependence of f_{rep} and f_{ceo} on the laser intensity corresponds to an intensity-related turning point for the center of the laser spectrum. Our experimental findings are consistent with theoretical calculations predicting the intensity-related effects on the laser parameters [35, 36].

3.1 Theory of intensity-dependent mode-locking dynamics

Calculating how the intensity of the laser affects its parameters, f_{rep} and f_{ceo}, can be accomplished by first expressing f_{rep} and f_{ceo} in terms of the average group (v_g) and phase (v_p) velocities in the laser cavity. From Eqns. (2.3) and (2.6), we have

$$\begin{aligned} f_{ceo} &= f_{rep}\frac{\omega_c l_c}{2\pi}\left(\frac{1}{v_g}-\frac{1}{v_p}\right) \\ &= \frac{\omega_c}{2\pi}\left(1-\frac{v_g}{v_p}\right) \qquad (3.1) \end{aligned}$$

since f_{rep} is given by

$$f_{rep} = \frac{v_g}{l_c} \qquad (3.2)$$

where ω_c is the spectrally weighted center frequency and l_c is the round-trip cavity length. For this analysis the laser cavity is approximated as a homogeneous medium, with an average index of refraction $\bar{n}$ defined such that the product $\bar{n} l_c$ is equivalent to the true optical path length through the air and Ti:sapphire crystal during a cavity round trip. $\bar{n}$ contains the nonlinear contribution to the index of refraction as well, and can be written as $\bar{n} = \bar{n}_0 + \bar{n}_2 \hat{I}$ where $\bar{n}_0$ and $\bar{n}_2$ are also average values for the laser cavity and $\hat{I}$ represents the pulse peak intensity of the laser within the Ti:sapphire crystal. Both v_p and v_g can be expressed in terms of $\bar{n}$. v_p is expressed simply as

$$v_p = \frac{c}{\bar{n}} = \frac{c}{\bar{n}_0 + \bar{n}_2 \hat{I}} \qquad (3.3)$$

However, because of the nonlinearity, the expression for v_g is not simply $(\mathrm{d}k/\mathrm{d}\omega)^{-1}$, but instead is given by

$$v_g = \frac{c}{\bar{n}_0 + \omega_c \left(\frac{\mathrm{d}\bar{n}_0}{\mathrm{d}\omega}\right)_{\omega_c} + 3\left[\bar{n}_2 + \omega_c \left(\frac{\mathrm{d}\bar{n}_2}{\mathrm{d}\omega}\right)_{\omega_c}\right]\hat{I}} \tag{3.4}$$

Appendix A provides a derivation of this expression. Note that naively using $(\mathrm{d}k/\mathrm{d}\omega)^{-1}$ would result in a very similar expression, with the coefficient of the nonlinear contribution changed from a 3 to a 1.

The dependence of the frequencies f_{rep} and f_{ceo} on the pulse peak intensity in the crystal, $\hat{I}$, is found by differentiating Eqns. (3.1) and (3.2) with respect to $\hat{I}$, yielding

$$\begin{aligned}
\frac{\mathrm{d}f_{rep}}{\mathrm{d}\hat{I}} &= \frac{1}{l_c}\frac{\mathrm{d}v_g}{\mathrm{d}\hat{I}} && (3.5)\\
\frac{\mathrm{d}f_{ceo}}{\mathrm{d}\hat{I}} &= \frac{1}{2\pi}\frac{\partial\omega_c}{\partial\hat{I}}\left(1 - \frac{v_g}{v_p}\right) + \frac{\omega_c}{2\pi}\frac{v_g}{v_p}\left(\frac{1}{v_p}\frac{\mathrm{d}v_p}{\mathrm{d}\hat{I}} - \frac{1}{v_g}\frac{\mathrm{d}v_g}{\mathrm{d}\hat{I}}\right) && (3.6)
\end{aligned}$$

The term $\partial\omega_c/\partial\hat{I}$ is the intensity-related laser spectral shift. From Eqn. (3.4) it is clear that the dependence of v_g on $\hat{I}$ enters explicitly, as well as through its dependence on ω_c. Therefore

$$\frac{\mathrm{d}}{\mathrm{d}\hat{I}}\left(\frac{1}{v_g}\right) = \frac{\partial\omega_c}{\partial\hat{I}}\frac{\partial}{\partial\omega_c}\left(\frac{1}{v_g}\right) + \frac{3}{c}\left[\bar{n}_2 + \omega_c\left(\frac{\mathrm{d}\bar{n}_2}{\mathrm{d}\omega}\right)_{\omega_c}\right] \tag{3.7}$$

From Eqn. (3.3) we see that v_p also depends on $\hat{I}$ explicitly and through ω_c ($\bar{n}$ depends on ω_c through dispersion), and so

$$\frac{\mathrm{d}}{\mathrm{d}\hat{I}}\left(\frac{1}{v_p}\right) = \frac{1}{c}\left(\frac{\partial\omega_c}{\partial\hat{I}}\frac{\partial\bar{n}}{\partial\omega_c} + \bar{n}_2\right) \tag{3.8}$$

Rewriting Eqn. (3.5) and plugging in (3.7) gives the dependence of f_{rep} on $\hat{I}$.

$$\begin{aligned}
\frac{\mathrm{d}f_{rep}}{\mathrm{d}\hat{I}} &= -\frac{v_g^2}{l_c}\frac{\mathrm{d}}{\mathrm{d}\hat{I}}\left(\frac{1}{v_g}\right)\\
&= -\frac{1}{l_c}\frac{v_g^2}{c}\left[c\frac{\partial\omega_c}{\partial\hat{I}}\frac{\partial}{\partial\omega_c}\left(\frac{1}{v_g}\right) + 3\bar{n}_2 + 3\omega_c\left(\frac{\mathrm{d}\bar{n}_2}{\mathrm{d}\omega}\right)_{\omega_c}\right] && (3.9)
\end{aligned}$$

The same can be done for f_{ceo} using Eqns. (3.6), (3.7), and (3.8).

$$
\begin{aligned}
\frac{\mathrm{d}f_{ceo}}{\mathrm{d}\hat{I}} &= \frac{1}{2\pi}\frac{\partial\omega_c}{\partial\hat{I}}\left(1-\frac{v_g}{v_p}\right)+\frac{\omega_c}{2\pi}\frac{v_g}{v_p}\left[-v_p\frac{\mathrm{d}}{\mathrm{d}\hat{I}}\left(\frac{1}{v_p}\right)+v_g\frac{\mathrm{d}}{\mathrm{d}\hat{I}}\left(\frac{1}{v_g}\right)\right] \\
&= \frac{1}{2\pi}\frac{\partial\omega_c}{\partial\hat{I}}\left[1-\frac{v_g}{v_p}-\frac{\omega_c v_g}{c}\frac{\partial\bar{n}}{\partial\omega_c}+\frac{\omega_c v_g^2}{v_p}\frac{\partial}{\partial\omega_c}\left(\frac{1}{v_g}\right)\right] \\
&\quad + \frac{\omega_c}{2\pi c}\frac{v_g}{v_p}\left[3\bar{n}_2 v_g-\bar{n}_2 v_p+3v_g\omega_c\left(\frac{\mathrm{d}\bar{n}_2}{\mathrm{d}\omega}\right)_{\omega_c}\right] \qquad (3.10)
\end{aligned}
$$

All of the terms in Eqns. (3.9) and (3.10) are constants taken from the literature, except $\partial\omega_c/\partial\hat{I}$ and $\partial(v_g^{-1})/\partial\omega_c$. The first term in (3.9) and the first set of brackets in (3.10) reveal the dependence on the intensity-related spectral shift. Both equations are dominated by the term proportional to $(\partial\omega_c/\partial\hat{I})(\partial(v_g^{-1})/\partial\omega_c)$, explaining the near coincidence in the sign change of $\mathrm{d}f_{rep}/\mathrm{d}\hat{I}$ and $\mathrm{d}f_{ceo}/\mathrm{d}\hat{I}$ with that of $\partial\omega_c/\partial\hat{I}$. This also reveals that the spectral shift couples to the laser parameters through the group-delay dispersion (GDD) of the laser cavity, which is proportional to $\partial(v_g^{-1})/\partial\omega_c$. From the dependence of f_{rep} on the pulse peak intensity, given in Eqn. (3.9), an experimental measurement of $\mathrm{d}f_{rep}/\mathrm{d}\hat{I}$ and $\partial\omega_c/\partial\hat{I}$ will uniquely determine the dependence of the average group velocity on the laser spectrum, $\partial(v_g^{-1})/\partial\omega_c$. This value can then be used in conjunction with $\partial\omega_c/\partial\hat{I}$ to calculate the dependence of f_{ceo} on the peak intensity, using Eqn. (3.10), which can be compared with an experimental measurement of $\mathrm{d}f_{ceo}/\mathrm{d}\hat{I}$.

An evaluation of these expressions requires several parameters of the laser. The Ti:sapphire laser we are using has a repetition frequency of 750 MHz, corresponding to a round-trip cavity length of $l_c = 0.4$ m, and has a ring configuration. The generation of short pulses from a mode-locked Ti:sapphire laser requires that the GDD of the crystal and other intracavity elements be compensated, such that during operation the net intracavity GDD is zero. This laser utilizes specially designed mirrors for the dispersion compensation, so the only intracavity transmission element is the crystal. Therefore, the average linear contribution to the refractive index is calculated from $\bar{n}_0 l_c = (l_c - l_{Ti}) + n_0^{Ti} l_{Ti}$. Using 2.4 mm for the length of the crystal, l_{Ti}, and 1.76 for the index of refraction of the crystal, n_0^{Ti}, yields $\bar{n}_0 = 1.00456$. Its dispersion is found using $\mathrm{d}\bar{n}_0/\mathrm{d}\omega = (l_{Ti}/l_c)(\mathrm{d}n_0^{Ti}/\mathrm{d}\omega)$, which gives $\mathrm{d}\bar{n}_0/\mathrm{d}\omega = 6 \times 10^{-20}$ s. The contribution from the mirrors to $\mathrm{d}\bar{n}_0/\mathrm{d}\omega$ is approximated by attributing the pulse delay to a reasonable dielectric thickness and is found to be negligible. The average nonlinear contribution to the refractive index is computed by taking into account Gaussian beam propagation through the crystal. With a $\mathrm{TEM}_{0,0}$ laser mode, $\bar{n}_2\hat{I}l_c = n_2^{Ti}\int_{-l_{Ti}/2}^{l_{Ti}/2} I(z)\,\mathrm{d}z = n_2^{Ti}\hat{I}\int_{-l_{Ti}/2}^{l_{Ti}/2}\left[1+(z/z_0)^2\right]^{-1}\mathrm{d}z$. z_0 indicates the Rayleigh range of the mode [78] and can be computed from the beam waist inside the crystal, which is 10 μm. From this expression, it is found that $\bar{n}_2 \sim 7.3 \times 10^{-23}\ \mathrm{m}^2\,\mathrm{W}^{-1}$. The ratio $\bar{n}_2/n_2$ is used to scale $\mathrm{d}n_2/\mathrm{d}\omega$ to find $\mathrm{d}\bar{n}_2/\mathrm{d}\omega$. Using the value of $\mathrm{d}n_2/\mathrm{d}\omega$ calculated in [30] yields $\mathrm{d}\bar{n}_2/\mathrm{d}\omega \sim 1.2 \times 10^{-38}\ \mathrm{s}\,\mathrm{m}^2\,\mathrm{W}^{-1}$. Finally, conversion from the experimentally measured value of the average laser output power, P, to the pulse peak intensity within the crystal, $\hat{I}$, is accomplished with the expression $\hat{I} = (2P)/(T\Delta\tau f_{rep}\pi w_0^2)$, where T is the output coupling and $\Delta\tau$ is the temporal width of the pulse. For a sech^2 pulse, $\Delta\tau$ is given by $\Delta\nu\Delta\tau = 0.315$, where $\Delta\nu$ is the pulse bandwidth in frequency [16]. The output coupling and pulse bandwidth are respectively 3 % and ~22 nm, which give $\hat{I} = (10^{16}\ \mathrm{m}^{-2})\,P$.

3.2 Experimental investigation of intensity-related dynamics

The intensity-related dynamics of the Ti:sapphire laser are experimentally measured using two different approaches. For the measurement of the low-bandwidth response, a direct-frequency-counting technique is used to track f_{rep} and f_{ceo}. For determining the response of f_{ceo} at Fourier frequencies higher than a few Hertz, a frequency-to-voltage (f–V) converter is used to measure its fluctuations. These two schemes are summarized in Fig. 3.1. In both scenarios, the intensity of the laser is varied by adjusting the pump power using an AOM in the pump beam. Using a modulation signal summed with an adjustable DC bias to control the power of the RF signal driving the AOM allows the intensity of the laser to be modulated about an adjustable mean value.

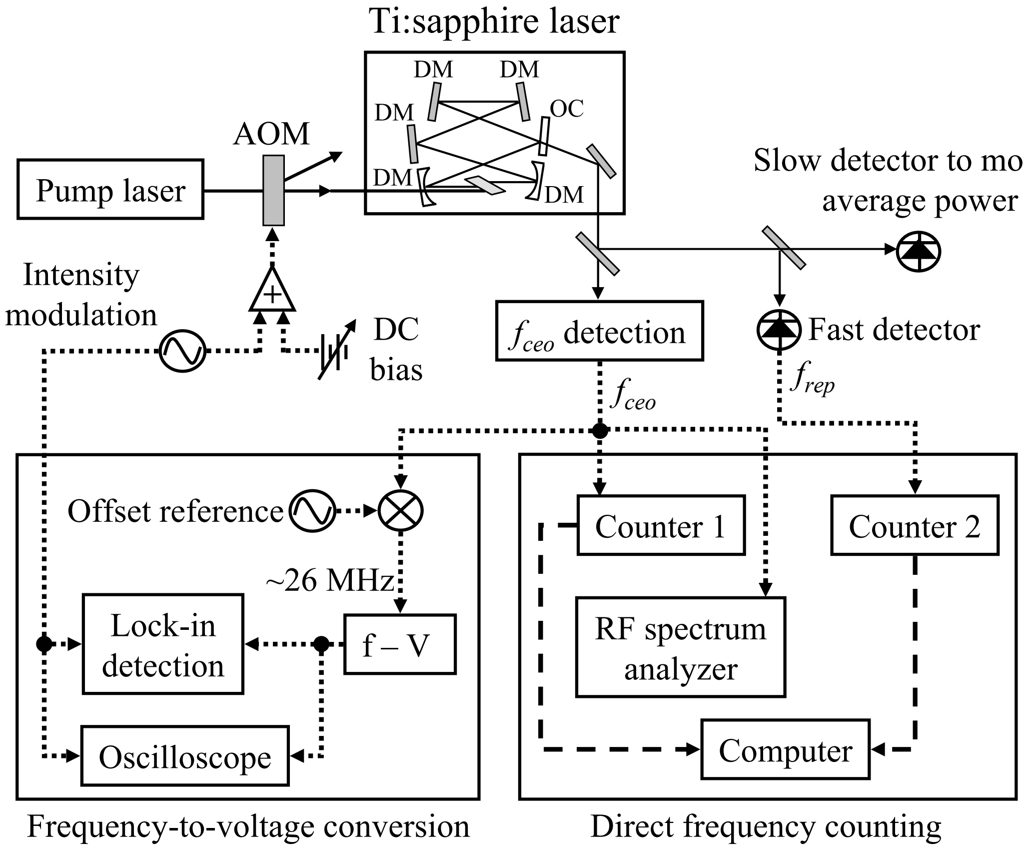

Figure 3.1: The intensity-related dynamics of the Ti:sapphire laser are measured by modulating the laser intensity and determining the response of f_{rep} and f_{ceo}. Direct frequency counting is used for low-frequency modulation and a frequency-to-voltage (f–V) converter is used to measure the fast dynamics of f_{ceo}. DM, dispersive mirrors to compensate dispersion of laser cavity; OC, output coupler of laser.

The carrier-envelope offset frequency is determined using the self-referencing technique described in Section 2.2, f_{rep} is obtained from a sufficiently fast photodetector, and the average laser output power is monitored with a low-bandwidth detector. For direct frequency counting, f_{rep} and f_{ceo} are recorded with two frequency counters operating at appropriately chosen gate times. An RF spectrum analyzer is utilized for analysis of the spectral linewidth and shape of f_{ceo}, as well as convenient readout of its mean value. The low-bandwidth response near DC is measured by sinusoidally modulating the laser intensity at ∼0.1 Hz so that the frequency counters with a gate time of 0.1 s can track the changing frequencies with sufficient resolution and speed. Measurement of the fast dynamics of f_{ceo} using the f–V converter is accomplished by mixing the f_{ceo} signal with a stable microwave frequency reference to produce a difference-frequency that falls within the input frequency range of the f–V converter. The response of f_{ceo} is determined from the amplitude of the oscillations of the voltage output from the f–V converter. This amplitude is measured using lock-in detection for sinusoidal modulation of the laser intensity at frequencies $<$ 100 kHz, and using an oscilloscope for frequencies $>$ 100 kHz up to 400 kHz. The frequency response (transfer function) of the f–V converter is carefully measured so that the true response of the laser can be extracted after removing the response of the f–V converter.

Figure 3.2 illustrates the intensity dependence of the center frequency of the pulse spectrum, ω_c, and f_{ceo}. These results are obtained by adjusting the DC bias in Fig. 3.1 that controls the pump power and recording the values of ω_c, f_{ceo}, and the Ti:sapphire laser average output power at each position. The values for ω_c are determined from the weighted averages of the pulse spectrum measured with an optical spectrum analyzer. Under certain conditions, variation of the pump power may result in the emergence of cw lasing components from the Ti:sapphire laser in addition to the pulse spectrum. These cw frequency components can have a significant influence on the pulse dynamics. However, the ring configuration of the laser cavity makes it straightforward to detect

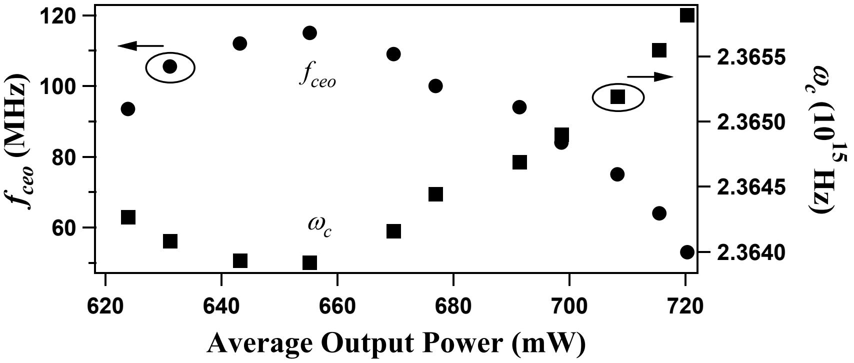

Figure 3.2: The dependence of ω_c (right axis) and f_{ceo} (left axis) on the Ti:sapphire laser average output power, which is directly related to its peak intensity, exhibits a sign reversal.

the presence of any cw components, since the pulses travel in only one direction around the cavity while the cw components will travel in both directions. Monitoring the laser output in the direction corresponding to propagation around the cavity opposite to that travelled by the pulses allows the detection of cw lasing components. The laser is adjusted to ensure that no cw components are present within the range of power variation. Also, an autocorrelation measurement confirms that the pulse width remains constant within this same range. Therefore, the measured changes in the average output power of the Ti:sapphire laser are directly related to changes in the peak intensity of the laser pulses within the laser crystal, $\hat{I}$, and the changes in the laser parameters can be attributed entirely to mode-locking dynamics related to the pulse peak intensity. Fig. 3.2 clearly shows that neither ω_c, shown as squares with respect to the vertical axis at right, nor f_{ceo}, displayed as circles with respect to the left vertical axis, is a monotonic function of the laser intensity. There is a change in the sign of $\mathrm{d}f_{ceo}/\mathrm{d}\hat{I}$, which is accompanied by a sign change of $\partial\omega_c/\partial\hat{I}$. A second-order polynomial fit to the ω_c data is used to find experimental values of the intensity-related spectral shift, $\partial\omega_c/\partial\hat{I}$.

Values for $\mathrm{d}f_{rep}/\mathrm{d}\hat{I}$ and $\mathrm{d}f_{ceo}/\mathrm{d}\hat{I}$ near DC are determined using the direct-frequency-counting method of Fig. 3.1. Shown in Fig. 3.3 are some representative data of the slowly modulated carrier-envelope offset frequency (pluses) with fits to a sine function (solid lines) for various average values of the laser output power. Note that the period of the oscillations indicated by the number of 0.1-s gate intervals per period is much shorter than that corresponding to the frequency of modulation of ∼0.1 Hz. This is due to the significant dead time of the frequency counter between adjacent data acquisition intervals. The data in Fig. 3.3 are for laser powers below that corresponding to the sign change of $\mathrm{d}f_{ceo}/\mathrm{d}\hat{I}$, where $\mathrm{d}f_{ceo}/\mathrm{d}\hat{I} > 0$, and are taken using nearly the same amplitude of power modulation. Hence, as the average laser power is increased, the mean value of f_{ceo} increases, whereas its modulation response amplitude decreases as $\mathrm{d}f_{ceo}/\mathrm{d}\hat{I}$ approaches zero. The original raw data for f_{ceo} contain a linear drift over time that is subtracted, since only the modulation amplitude is of interest. Similar data are obtained for f_{rep}. From the amplitudes of the sinusoidal fits and direct measurements of the modulation depth of the average laser output power, experimental values

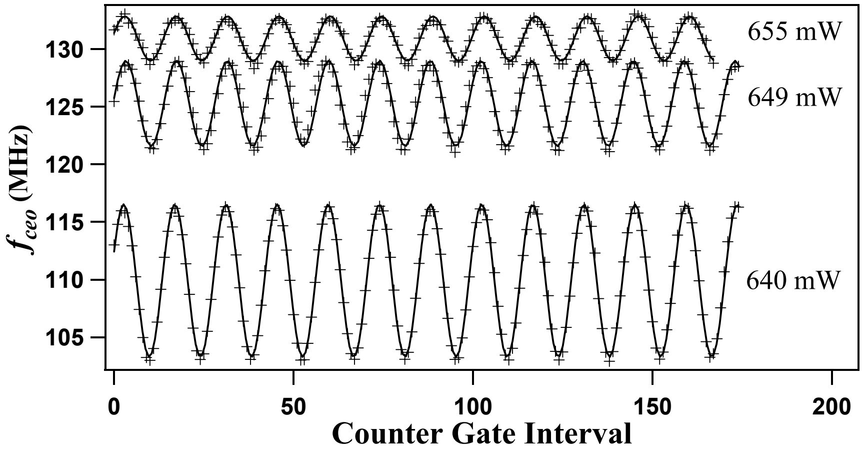

Figure 3.3: Direct frequency counting of the response of f_{ceo} to slow modulation of the laser intensity for laser powers below that corresponding to the sign change of $\mathrm{d}f_{ceo}/\mathrm{d}\hat{I}$ (pluses), and sine function fits (solid lines). The amplitude of power modulation is nearly constant. The mean value of f_{ceo} increases with power ($\mathrm{d}f_{ceo}/\mathrm{d}\hat{I} > 0$), whereas the response amplitude decreases as $\mathrm{d}f_{ceo}/\mathrm{d}\hat{I}$ approaches zero.

of $\mathrm{d}f_{rep}/\mathrm{d}\hat{I}$ and $\mathrm{d}f_{ceo}/\mathrm{d}\hat{I}$ are determined. These values are shown in Fig. 3.4, with filled circles for $\mathrm{d}f_{rep}/\mathrm{d}\hat{I}$ and open circles for $\mathrm{d}f_{ceo}/\mathrm{d}\hat{I}$. From the experimental values of $\mathrm{d}f_{rep}/\mathrm{d}\hat{I}$ and $\partial\omega_c/\partial\hat{I}$, the values of $\partial(v_g^{-1})/\partial\omega_c$ can be computed using Eqn. (3.9). The values of $\partial(v_g^{-1})/\partial\omega_c$ are found to be negative, which is consistent with the fact that the net linear contribution to the GDD in the laser cavity must be negative to compensate for the effective positive nonlinear contribution during operation from the self-phase modulation in the crystal. The calculated values of $\partial(v_g^{-1})/\partial\omega_c$ are used in combination with the experimental values of $\partial\omega_c/\partial\hat{I}$ to compute $\mathrm{d}f_{ceo}/\mathrm{d}\hat{I}$ via Eqn. (3.10), with the results shown as triangles in Fig. 3.4. These calculated values of $\mathrm{d}f_{ceo}/\mathrm{d}\hat{I}$ are in good agreement with the directly measured data, showing that the equations for $\mathrm{d}f_{rep}/\mathrm{d}\hat{I}$ and $\mathrm{d}f_{ceo}/\mathrm{d}\hat{I}$, (3.9) and (3.10), are consistent with the experimental data.

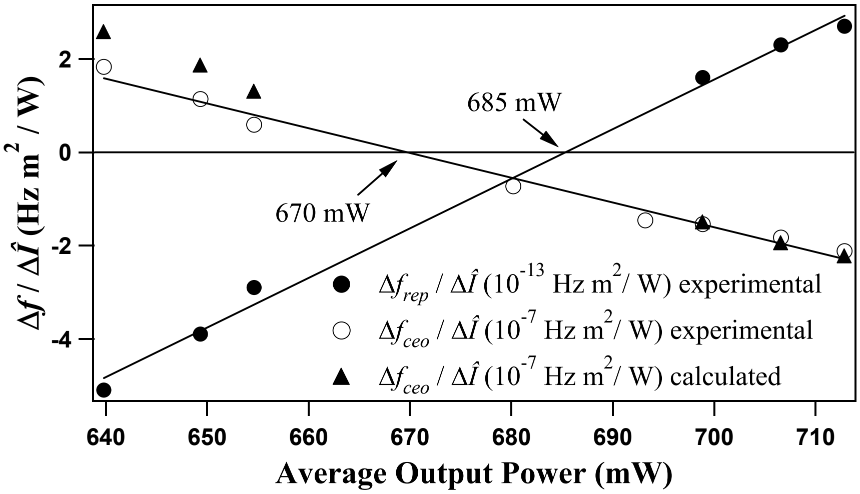

Figure 3.4: The experimental values of $\Delta f_{rep}/\Delta\hat{I}$ and $\Delta f_{ceo}/\Delta\hat{I}$ measured near DC, versus the Ti:sapphire laser average output power, confirm that there is a sign reversal of these coefficients. The values of $\Delta f_{ceo}/\Delta\hat{I}$ calculated using a parameter derived from the $\Delta f_{rep}/\Delta\hat{I}$ data are in close agreement with the experimental data.

The dynamic response of f_{ceo} to modulation of the Ti:sapphire laser intensity at frequencies higher than a few Hertz is determined using the f–V converter as shown in Fig. 3.1. The response of f_{ceo} and the depth of intensity modulation are measured for modulation frequencies ranging from 10 Hz to 400 kHz at three different values of average laser power. These data of the transfer function of f_{ceo}, $\Delta f_{ceo}/\Delta\hat{I}$ versus modulation frequency, for the three values of average laser power are shown in Fig. 3.5. (The data were taken at a different time than those in the previous figures, and f_{ceo} reaches a maximum at a higher average laser power of ~755 mW instead of ~660 mW.) The two sets of data taken at average powers of 711 mW and 744 mW (squares and circles, respectively) correspond to powers below the zero-crossing point of $\mathrm{d}f_{ceo}/\mathrm{d}\hat{I}$ ($\mathrm{d}f_{ceo}/\mathrm{d}\hat{I} > 0$). The third data set taken at an average power of 763 mW (triangles) corresponds to a region of operation where $\mathrm{d}f_{ceo}/\mathrm{d}\hat{I} < 0$. The steep rise in response toward DC for Fourier frequencies below 1 kHz is attributed to thermal effects in the Ti:sapphire crystal. Above 1 kHz, however, the response is roughly flat, at least over the measured frequency range extending to 400 kHz. Notice that although the fast response of f_{ceo} does change sign, following the sign of $\mathrm{d}f_{ceo}/\mathrm{d}\hat{I}$, the sign of the thermal response is unaffected by the sign change of $\mathrm{d}f_{ceo}/\mathrm{d}\hat{I}$. This can be seen by the fact that the response of f_{ceo} approaches zero for modulation frequencies approaching DC for the

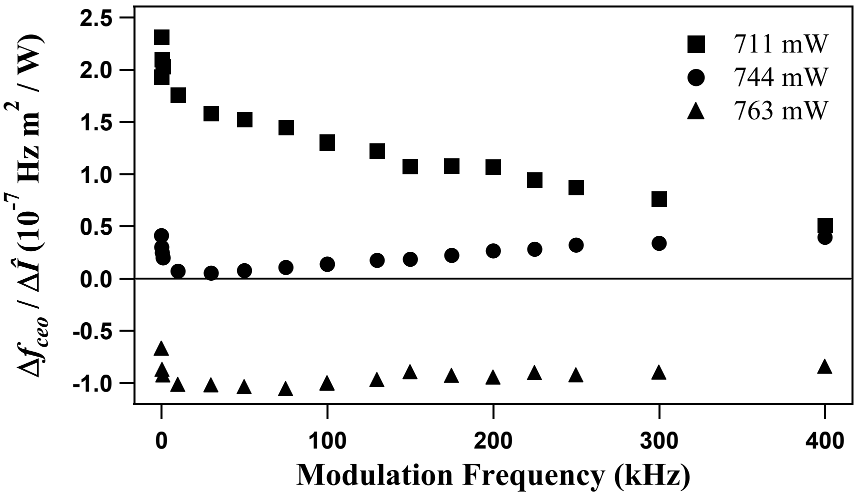

Figure 3.5: The dynamic response (transfer function) of f_{ceo} is shown for two laser powers where $\mathrm{d}f_{ceo}/\mathrm{d}\hat{I} > 0$ and for one value of the laser power where $\mathrm{d}f_{ceo}/\mathrm{d}\hat{I} < 0$. Notice that although the fast response changes sign accordingly, the thermal response near DC always provides a positive contribution.

763-mW data (where $\mathrm{d}f_{ceo}/\mathrm{d}\hat{I} < 0$). The increased positive thermal response partially cancels the negative response related to the fast mode-locking pulse dynamics.

3.3 Investigation of lineshape of carrier-envelope offset frequency

In addition to directly measuring $\mathrm{d}f_{ceo}/\mathrm{d}\hat{I}$, the dynamics of f_{ceo} can be studied by examining its spectral linewidth and shape. The value of f_{ceo} and its spectral linewidth are measured with an RF spectrum analyzer for several values of the average Ti:sapphire laser power, both above and below the turning point of f_{ceo}. These data, which represent a DC measurement of the intensity dependence of the mean value of f_{ceo}, are shown in the upper portion of Fig. 3.6. In the bottom panel of this figure are the experimental values of $\mathrm{d}f_{ceo}/\mathrm{d}\hat{I}$ versus average laser power obtained for intensity modulation at 10 kHz. From the data in Fig. 3.5, it is evident that when $\mathrm{d}f_{ceo}/\mathrm{d}\hat{I} = 0$ for modulation at 10 kHz, a DC measurement of $\mathrm{d}f_{ceo}/\mathrm{d}\hat{I}$ will produce a positive result because of the positive thermal response at DC. Therefore, a DC measurement will find that $\mathrm{d}f_{ceo}/\mathrm{d}\hat{I} = 0$ at a higher power, where the fast response is negative and cancels the positive thermal response. This is supported by Fig. 3.6, where the extremum of the f_{ceo} data taken at DC occurs at a higher power than that corresponding to the zero-crossing of $\mathrm{d}f_{ceo}/\mathrm{d}\hat{I}$ obtained with 10-kHz modulation. In addition, the lineshapes in the upper portion of Fig. 3.6 reveal that the phase noise of f_{ceo} increases dramatically when $\mathrm{d}f_{ceo}/\mathrm{d}\hat{I}$ deviates from zero, and is at a minimum when $\mathrm{d}f_{ceo}/\mathrm{d}\hat{I} = 0$ when measured at 10-kHz modulation (not at the extremum for the DC measurements). This indicates that the dominant source of phase variations of f_{ceo} is fast fluctuations of the pulse intensity, since the linewidth of f_{ceo} is at a minimum at the position where it is decoupled from such fast fluctuations. Attributing the phase variations of f_{ceo} to intensity fluctuations can be verified by considering the rms power fluctuations of the pump laser, provided by the laser manufacturer. The given value of 0.02 % for rms power fluctuations of the pump laser is converted into frequency noise using the measured

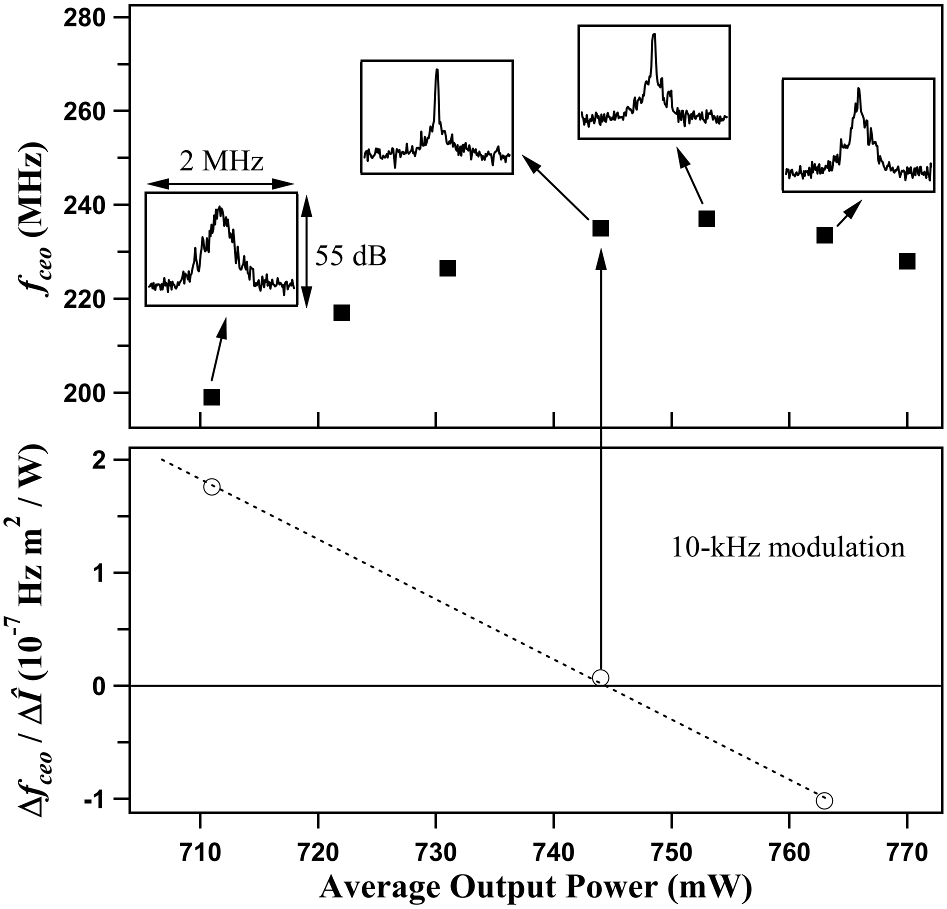

Figure 3.6: Top: The linewidth of f_{ceo} for average laser powers above and below that corresponding to $\mathrm{d}f_{ceo}/\mathrm{d}\hat{I} = 0$, with corresponding mean values for f_{ceo} obtained at DC. The lineshape traces are obtained with a 10-kHz resolution bandwidth. Bottom: The response of f_{ceo} to intensity measured at 10-kHz modulation.

coefficients for the dependence of f_{ceo} on the Ti:sapphire laser intensity and the dependence of the laser intensity on the pump power. For the first value of f_{ceo} in Fig. 3.6, at 197 MHz, this conversion produces an rms frequency noise of ~200 kHz. In comparison, the measured width of the corresponding lineshape in Fig. 3.6 is ~500 kHz. For the next-to-last f_{ceo} data point shown in Fig. 3.6, the pump power fluctuations predict an rms frequency noise of ~100 kHz, half that predicted for the first value. Indeed, the measured linewidth for this point is ~300 kHz, roughly half that measured for the first point. The linewidths predicted solely by power fluctuations of the pump laser are in rough agreement with those measured.

The universality of our results for different Ti:sapphire laser systems can be verified by studying a system that uses a different scheme for compensating the GDD in the laser cavity, which exhibits a linewidth for f_{ceo} that is much smaller than is typically achieved for the system I have been discussing. Whereas the system discussed until now uses dispersive mirrors for dispersion compensation, an alternate method for compensating dispersion relies on a pair of prisms inside the laser cavity [22, 71]. A typical lineshape of f_{ceo} for the Ti:sapphire laser system with prisms is shown on the right in Fig. 3.7. This is as much as an order of magnitude narrower than the typical lineshape for the prismless system, which exhibits a comparable linewidth only near the turning point of f_{ceo}. For comparison, the narrowest lineshape obtained for the prismless system is also shown in Fig. 3.7. It is often the case for the prismless laser that only near the position of the narrowest lineshape close to the turning point can f_{ceo} be phase-locked to a microwave reference and achieve a linewidth equivalent to that of the reference. Shown in the inset of Fig. 3.7 is such a phase-locked narrow linewidth approaching 10 mHz, which is not resolved by the 10-Hz resolution bandwidth of the figure. On the other hand, when $\mathrm{d}f_{ceo}/\mathrm{d}\hat{I}$ is tuned away from zero, it becomes increasingly difficult to phase-lock f_{ceo} to a microwave reference and achieve such a narrow lineshape.

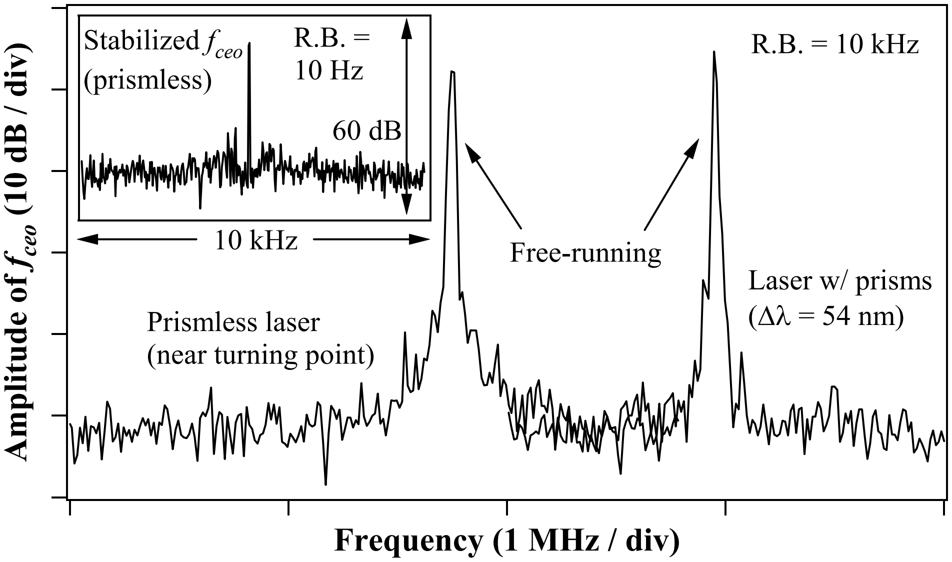

Figure 3.7: The typical lineshape of f_{ceo} for the Ti:sapphire laser incorporating prisms, recorded for a spectral bandwidth of 54 nm, is compared to the narrowest lineshape obtained for the prismless system. The inset shows the resolution-limited linewidth of f_{ceo} phase-locked to a microwave reference for the prismless system. This narrow lineshape is often only achieved when operating near the turning point. R.B., resolution bandwidth.

To understand the physical mechanisms responsible for the smaller linewidth of the system with prisms, its dynamics are studied as the insertion of one of the prisms into the cavity mode is varied. The insertion changes both the width of the pulse spectrum and the magnitude of the linear contribution to the net GDD of the laser cavity. Figure 3.8 shows the linewidths of f_{ceo} corresponding to various pulse spectral widths [full-width at half-maximum (FWHM)]. These data show that the f_{ceo} linewidth decreases as the width of the pulse spectrum increases, until the spectral width reaches about 47 nm. For spectra broader than ∼47 nm, the linewidth is nearly constant. This behavior is consistent with our model attributing the noise of f_{ceo} to intensity fluctuations of the laser, which becomes apparent when the response of f_{ceo} to intensity modulation is measured for various pulse spectral widths. This data is shown in Fig. 3.9, which reveals a threshold behavior where $\mathrm{d}f_{ceo}/\mathrm{d}\hat{I}$ decreases sharply with increasing spectral width up to a bandwidth of ∼47 nm, above which $\mathrm{d}f_{ceo}/\mathrm{d}\hat{I}$ remains constant

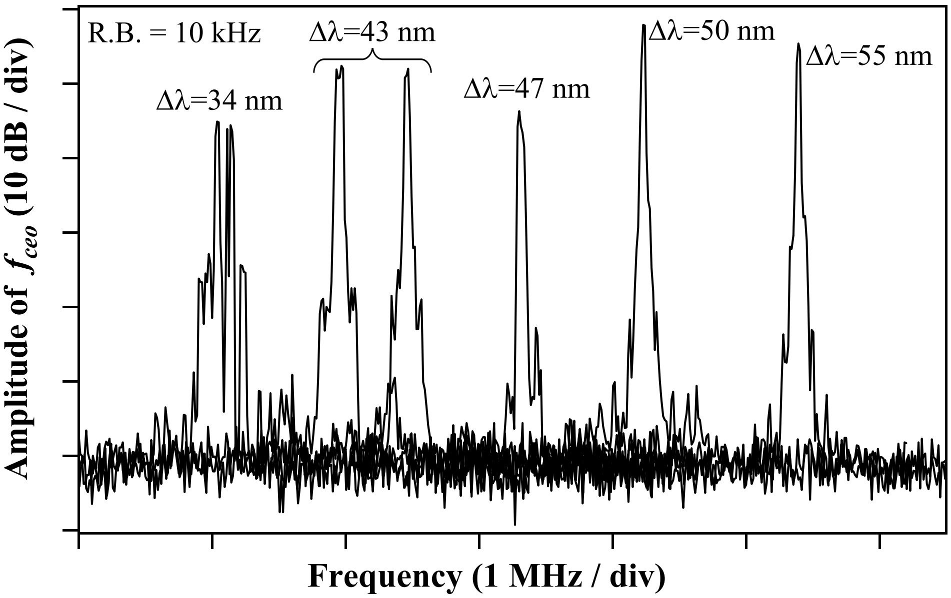

Figure 3.8: The linewidths of f_{ceo} for the Ti:sapphire laser system with prisms corresponding to various pulse spectral widths show that the linewidth decreases with increasing spectral width up to ∼47 nm, above which the linewidth is constant. R.B., resolution bandwidth.

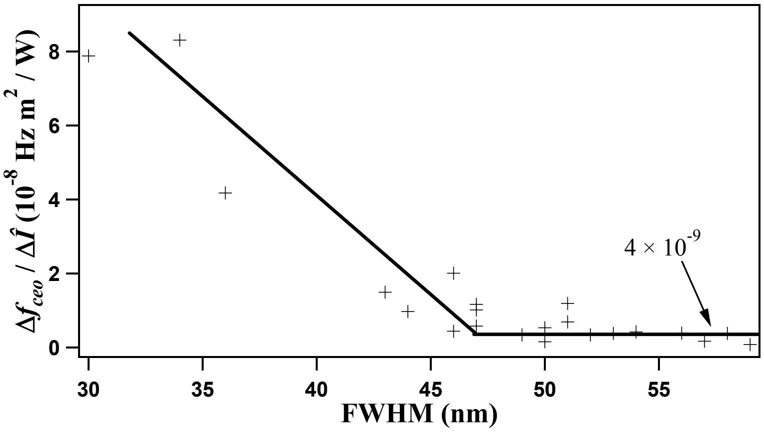

Figure 3.9: The intensity dependence of f_{ceo} for the laser system with prisms drops sharply with increasing pulse spectral width until ∼47 nm, above which it is constant.

at $\sim 4 \times 10^{-9}$ Hz m^2 W^{-1}. The broader spectral bandwidth and corresponding decrease in the magnitude of the linear contribution to the net cavity GDD both reduce the intensity dependence of f_{ceo} by reducing the magnitude of $\partial\omega_c/\partial\hat{I}$ and $\partial(v_g^{-1})/\partial\omega_c$, respectively. Therefore, for both the prismless system and the system with prisms, a decrease in the magnitude of the intensity-dependent spectral shift and / or the linear contribution to the net cavity GDD reduces the coupling of f_{ceo} to intensity fluctuations, thereby minimizing the free-running linewidth of f_{ceo}. The fact that for all recorded pulse bandwidths shown in Fig. 3.9, the magnitude of $\mathrm{d}f_{ceo}/\mathrm{d}\hat{I}$ for the system with prisms is approximately ten times smaller than it is for the prismless system, as shown in Fig. 3.4, explains why in general the f_{ceo} linewidth for the prism system is smaller than that of the prismless laser.

3.4 Conclusions

The experimental results and supporting theoretical calculations demonstrate that the peak pulse intensity of the Ti:sapphire laser couples to the laser parameters, f_{rep} and f_{ceo}, through an intensity-related shift of the pulse spectrum. Changes in the pulse spectrum directly affect f_{rep} and f_{ceo} through the linear contribution to the GDD of the laser cavity, and the turning point of the laser parameters with increasing pulse intensity is due to a corresponding turning point for the center of the spectrum. In a previous effort to explain the intensity-dependence of the parameters of the frequency comb produced by a mode-locked laser, the dependence of the group and phase velocities of solitons on the pulse energy was modeled [27]. However, although it was noted in that work that a spectral shift could cause a change in the group velocity, the model that was developed did not account for this possibility and only included effects from the Kerr nonlinearity. Our results clearly demonstrate that this model is not an accurate description of the intensity-related effects occurring in an active laser cavity and that the dominant mechanism for intensity coupling is an intensity-related spectral shift. There are two possible explanations for this intensity-induced spectral shift. The first is stimulated Raman scattering in the laser crystal, which causes a redshift of the pulse spectrum with increasing pulse energy [47, 28]. The second is a mismatch between the gain band of the laser and the loss band, determined by the mirrors, the output coupler, and reabsorption in the laser crystal. An increase in the laser intensity will reduce the gain due to multipass gain saturation, which will shift the net-gain maximum and the peak of the laser spectrum [47].

Also, our experimental investigations have identified fast fluctuations of the laser intensity as the dominant source of phase noise of f_{ceo}. For the laser system using a pair of intracavity prisms for dispersion compensation, it is possible to tune the GDD of the cavity to reduce the coupling of the intensity noise to f_{ceo} and achieve a narrow f_{ceo} lineshape. However, for the system using dispersive mirrors instead of prisms, the linear contribution to the net cavity GDD is fixed. The calculation of $\partial(v_g^{-1})/\partial\omega_c$ from our experimental data indicates this is ~ -400 fs^2. Therefore, for the prismless system the intensity can only be decoupled from f_{ceo} by reducing the intensity-dependent spectral shift. This explains why the free-running lineshape of f_{ceo} for this system occurs near the position where $\partial\omega_c/\partial\hat{I} = 0$. A reduction in the spectral shift may also contribute to the minimization of the intensity dependence of f_{ceo} for pulse bandwidths exceeding $\sim$47 nm in the system using prisms, since a broad bandwidth limits the amount of spectral shift that can occur.

Our results have important implications for the stabilization of f_{ceo} using the laser intensity. At first it would seem that it is desirable to operate with the narrowest free-running f_{ceo} linewidth. However, this corresponds to operation where $\mathrm{d}f_{ceo}/\mathrm{d}\hat{I} = 0$, which would make it impossible to stabilize f_{ceo} with the laser intensity. It is also important to not operate the laser too close to this position, since the intensity would have to be changed dramatically to affect f_{ceo}, resulting in a large amount of intensity noise on the laser. The stabilization of f_{ceo} is optimized when the laser is operated sufficiently close to the position of $\mathrm{d}f_{ceo}/\mathrm{d}\hat{I} = 0$ to minimize the dependence of f_{ceo} on the fast intensity fluctuations outside the bandwidth of the stabilization loop, while being far enough away to use the lower-bandwidth response of f_{ceo} for stabilization. In some cases it is also possible to select the sign of the intensity dependence of f_{ceo} such that the intensity noise on the laser is suppressed when f_{ceo} is stabilized, as demonstrated in [59]. However, operating with the incorrect sign for $\mathrm{d}f_{ceo}/\mathrm{d}\hat{I}$ could increase the intensity noise during stabilization of f_{ceo}.

Chapter 4

Experimental Results for Linking Optical and Microwave Frequencies with the Comb

Once the carrier-envelope offset frequency of a mode-locked laser is stabilized, all that is needed for the frequency comb produced by the laser to be used as a phase-coherent connection between microwave and optical frequencies is to stabilize the repetition frequency, f_{rep}, of the laser. The repetition frequency can either be stabilized directly to a microwave frequency standard, allowing the comb to be used for making optical frequency measurements referenced to the microwave standard, or an optical frequency standard can be used to stabilize f_{rep}, providing a microwave frequency that is phase-coherently linked to the optical standard. Both of these scenarios have been explained in detail in Section 2.3. In this chapter I will present experimental results for both methods of linking optical and microwave frequencies using the comb produced by a mode-locked Ti:sapphire laser. Optical transitions in two different atomic species are connected to microwave frequency references. The first is a transition in molecular iodine (I_2), contained in a vapor cell, and the second is one in ultracold, laser-cooled strontium-88 (^{88}Sr) atoms. From the experimental results, we will also see how the performance of a vapor cell compares to that of a laser-cooled sample of atoms.

4.1 Cell-based optical standard

An optical frequency standard based on atoms or molecules enclosed in a vapor cell is attractive for the compact size, reliability, and relative simplicity of the system. Molecular iodine provides several narrow transitions that can potentially serve as secondary, portable optical frequency standards. A system based on a frequency-doubled Nd:YAG laser stabilized to a transition near 532 nm of I_2 contained in a vapor cell has proven to be one of the most practical optical frequency standards, demonstrating an excellent stability for such a compact and simple system, with a fractional instability $< 5 \times 10^{-14}$ at 1 s [87]. To achieve even higher frequency stability, it is useful to explore I_2 transitions at wavelengths below 532 nm, where the natural linewidths of the transitions decrease at a rate faster than that for the transition strengths as molecular iodine approaches its dissociation limit [10]. There is a transition near 515 nm which may provide one of the better frequency standards based on I_2, but the transition frequency has previously only been measured using wavelength interferometry. Stabilizing the repetition frequency of a mode-locked laser to a microwave source referenced to the Cs standard allows the frequency comb to be used to measure the absolute frequency of this transition with respect to the primary standard defining the SI second [46].

For probing the I_2 transition, a cw Ti:sapphire laser operating near 1030 nm is first pre-stabilized to a mode of a medium-finesse, evacuated, vibration-isolated, thermally stabilized optical cavity. A more compact Yb:YAG laser has also been used for this experiment. The narrow-linewidth laser is frequency-doubled and stabilized to the transition of the I_2 contained in an 8-cm-long vapor cell using frequency-modulation (FM) saturated-absorption spectroscopy to minimize the contribution of Doppler effects to the measured linewidth [14]. The cell pressure is controlled by stabilizing its cold-finger temperature, which is set at -5 °C. The absolute frequency of the 515-nm transition is determined using the Cs-stabilized frequency comb to measure the fundamental frequency of the I_2-stabilized cw laser. Figure 4.1(a) shows an example record over $\sim$3300 s of the deviations from the mean of the heterodyne beat between the cw laser and the nearest comb component. These measurements are obtained by counting the heterodyne beat using a 1-s gate time for the frequency counter. The Allan deviation is used to express how the fractional fluctuations of the measured optical frequency depend on averaging time, as discussed in Section 2.5. It is calculated for the data in Fig. 4.1(a) using Eqn. (2.16) and normalizing by the optical frequency. Shown in Fig. 4.1(b), the Allan deviation for the optical frequency measurement is identical to that of the Cs standard, which is provided by a commercial Cs beam clock. This indicates that the measurement is limited by the stability of the Cs standard to which the frequency comb is referenced.

From the measurement of the heterodyne beat frequency between the cw laser and the frequency comb, the absolute frequency of the cw laser (and therefore the 515-nm I_2 transition to which it is stabilized) is determined using Eqn. (2.10). The mode number, m, of the comb component against which the cw laser is beating is determined by independently measuring the cw laser frequency with a resolution better than half the mode spacing of the comb. This corresponds to a resolution better than 375 MHz, which is achieved using a medium-resolution wavelength meter. The frequency of the 515-nm transition is measured over a period of two weeks for three different I_2 cells, all of which were prepared by the Bureau International des Poids et Mesures (BIPM). Two were made recently before the measurements were performed, while the third was made more than eight years prior to the measurements. The results are shown in Fig. 4.2 with respect to the value adopted for this transition by the Comité International des Poids et Mesures (CIPM) in 1997 from the previous measurements using wavelength interferometry. They reveal that there is no statistically significant difference among the three I_2 cells. The weighted mean of the results indicates the frequency of the transition measured with the Cs-referenced comb is 71.8 kHz higher than the CIPM recommended

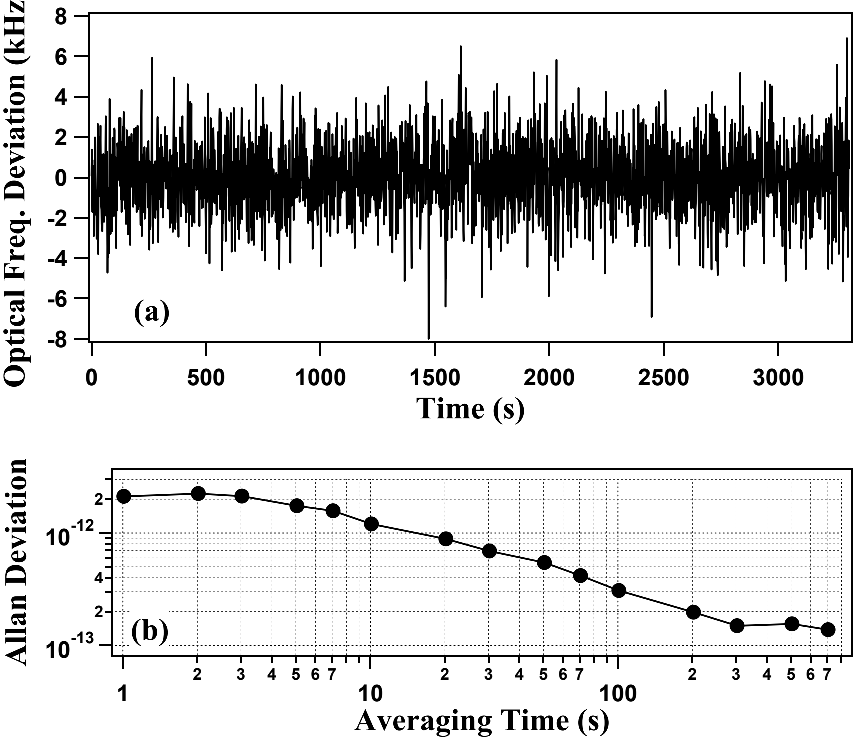

Figure 4.1: (a) The deviations from the mean of the heterodyne beat between the cw laser stabilized to the 515-nm I_2 transition and a mode of the Cs-referenced frequency comb. A gate time of 1 s is used for counting the beat frequency. (b) The Allan deviation of the I_2-stabilized cw laser frequency, calculated from the data in (a). It is limited by the stability of the Cs standard referencing the frequency comb for the averaging times shown here.

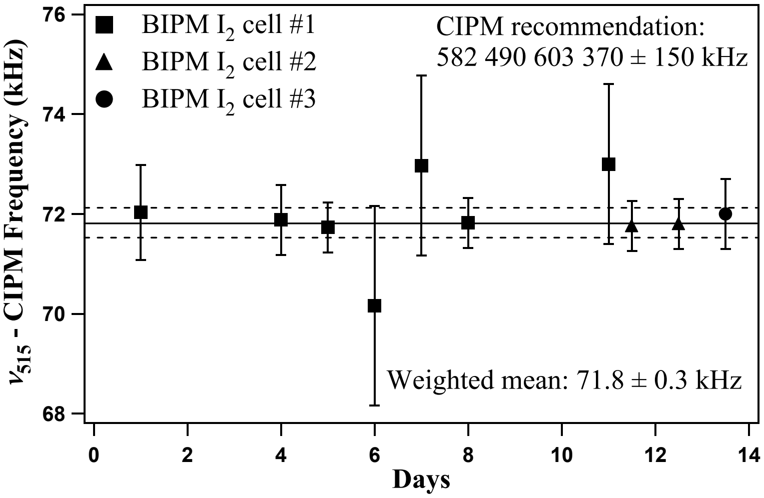

Figure 4.2: The absolute frequency of the 515-nm I_2 transition is measured with the Cs-referenced frequency comb over a period of two weeks for three different I_2 cells. A mean value is found that is 71.8 kHz above that recommended by the CIPM, which is within the uncertainty of the CIPM recommendation.

value of 582 490 603.37 MHz [62], with a 0.3 kHz statistical uncertainty. This offset is well within the 1-sigma uncertainty (150 kHz) of the recommended value. Allowing for various systematic effects on the transition frequency, the total uncertainty in the comb-based measurement is estimated to be 1.5 kHz, which represents an improvement by 2 orders of magnitude from the uncertainty in the CIPM recommendation.

Although Fig. 4.2 demonstrates that the frequency comb can be used to make consistent frequency measurements over a period of a few weeks, its reliability for providing accurate frequency measurements can be further verified by using it to measure the frequency of a Nd:YAG laser stabilized to the narrow I_2 transition near 532 nm and comparing the results with those obtained previously using a different comb system. The measurements of the I_2-stabilized fundamental laser frequency are shown in Fig. 4.3, with respect to the CIPM recommended frequency. Five different periods of measurement of this I_2 transition are shown, spanning more than two years. The last period contains the results using the current frequency comb produced by the mode-

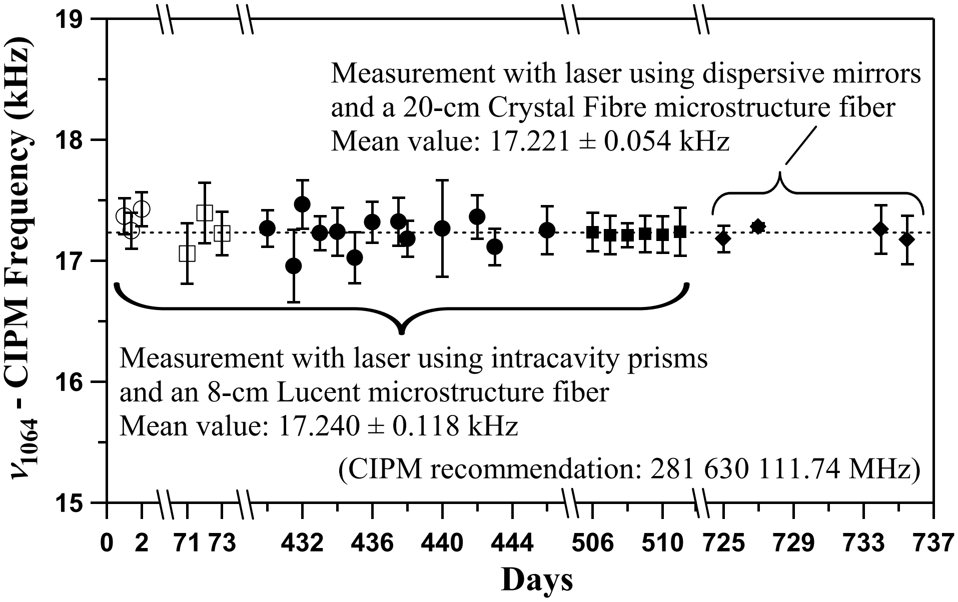

Figure 4.3: The results for measuring the 532-nm I_2 transition frequency with the current comb system agree very well with those obtained over a period up to two years prior using a significantly different comb system.

locked laser utilizing intracavity dispersive mirrors for dispersion compensation. The previous four sets of measurements were performed with a different comb system consisting of a mode-locked laser that uses intracavity prisms for dispersion compensation and a different microstructure fiber for broadening the spectrum to stabilize f_{ceo}. The agreement of these measurements made with two different comb systems over two years is excellent, with a difference between the mean value of the previous measurements and that for the results obtained with the current comb system that is less than 20 Hz. This is well within the standard deviations of the previous measurements and the current ones, which are 118 Hz and 54 Hz, respectively. As for the 515-nm measurement, the short-term frequency-measurement noise of the 532-nm system is limited by the instability of the Cs reference. The reproducibility of I_2 systems, however, is limited by the baseline fluctuations in the I_2 spectrometer.

Since the stability of the measurements of the I_2 transitions is being limited by the stability of the Cs standard, as evidenced by the Allan deviation of Fig. 4.1(b), it is useful to use a more stable microwave frequency reference for the comb. One such reference that is available is a hydrogen maser, which exhibits an instability of $\sim 2 \times 10^{-13}\,\tau^{-1/2}$ for an averaging time τ — roughly an order of magnitude better than the Cs standard. Stabilizing the repetition frequency of the comb to a microwave source referenced to the maser allows the stability of the I_2-stabilized laser to be explored at a level unattainable by the Cs-referenced comb. For an absolute frequency measurement, which must be referenced to the Cs primary standard that defines the second, the maser is calibrated relative to the NIST primary Cs frequency standard. Figure 4.4(a) shows the deviations from the mean of the heterodyne beat between the maser-referenced comb and a cw laser stabilized to the 532-nm I_2 transition, using a 10-s gate time for the frequency counter. The comb remains stabilized to the microwave reference during the entire measurement period shown here of $\sim$5 h. This level of reliability is critical for using the comb in the implementation of an optical clock. The Allan deviation for

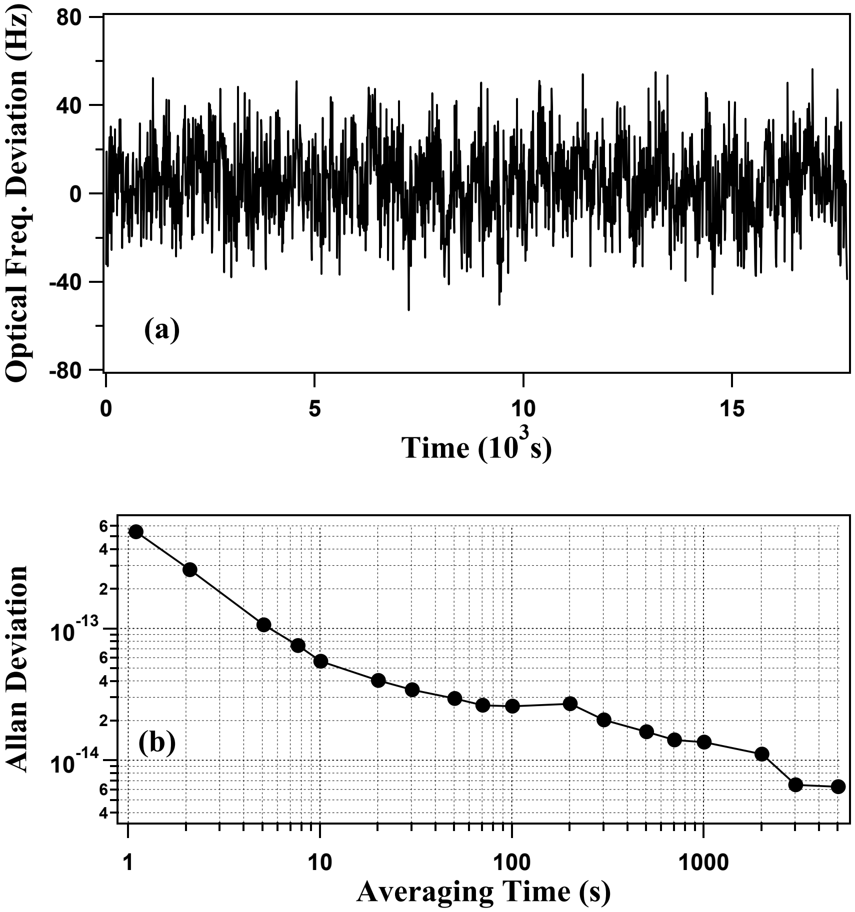

Figure 4.4: (a) The deviations from the mean of the heterodyne beat between the cw laser stabilized to the 532-nm I_2 transition and the maser-referenced comb, using a 10-s gate time for frequency counting. (b) The Allan deviation of the I_2-stabilized cw laser frequency, measured with the maser-referenced comb. It is limited by the stability of the microwave reference.

this optical frequency measurement [shown in Fig. 4.4(b)] at averaging times of 10 s and below is obtained from data taken with corresponding gate times to eliminate the effects of dead time between gate intervals, but at longer averaging times it is calculated by combining the 10-s data as described in Section 2.5. The maser-referenced microwave frequency to which the comb repetition frequency is stabilized must be transmitted from its source within the NIST laboratories in Boulder to our laboratory in JILA on the university campus [86]. Issues involved with this transfer will be discussed in detail in Chapter 6. The Allan deviation of the measured optical frequency corresponds to that of the transmitted maser-referenced microwave signal, revealing that the stability of this frequency measurement is still limited by the microwave reference.

The second method of using the comb to connect optical and microwave frequencies — establishing a phase-coherent link between f_{rep} and an optical frequency by stabilizing the heterodyne beat between the comb and the optical frequency — can also be demonstrated with the I_2 system. It is in this manner that the comb would be incorporated into an optical clock. After stabilizing the beat between the comb and the cw laser locked to the 532-nm I_2 transition, the fluctuations of the comb repetition frequency are measured by mixing f_{rep} with the microwave signal transmitted from NIST that is referenced to the maser. This generates a difference frequency that can be counted with sufficient resolution to determine the instability of f_{rep} with respect to the microwave reference. Of course, from the previous measurements, it is expected that the instability of the transmitted microwave reference will limit this measurement. Figure 4.5(a) shows a record over ~3 h of the fractional fluctuations of f_{rep} obtained in this way, using a gate time of 20 s. The Allan deviation of f_{rep} at averaging times of 20 s and below is obtained using data taken with corresponding gate times, and at longer averaging times it is determined by combining the 10-s data. It is shown in Fig. 4.5(b) and confirms that the measurement is limited by the stability of the microwave reference. To evaluate the true stability of this I_2-based optical clock, it is necessary to compare

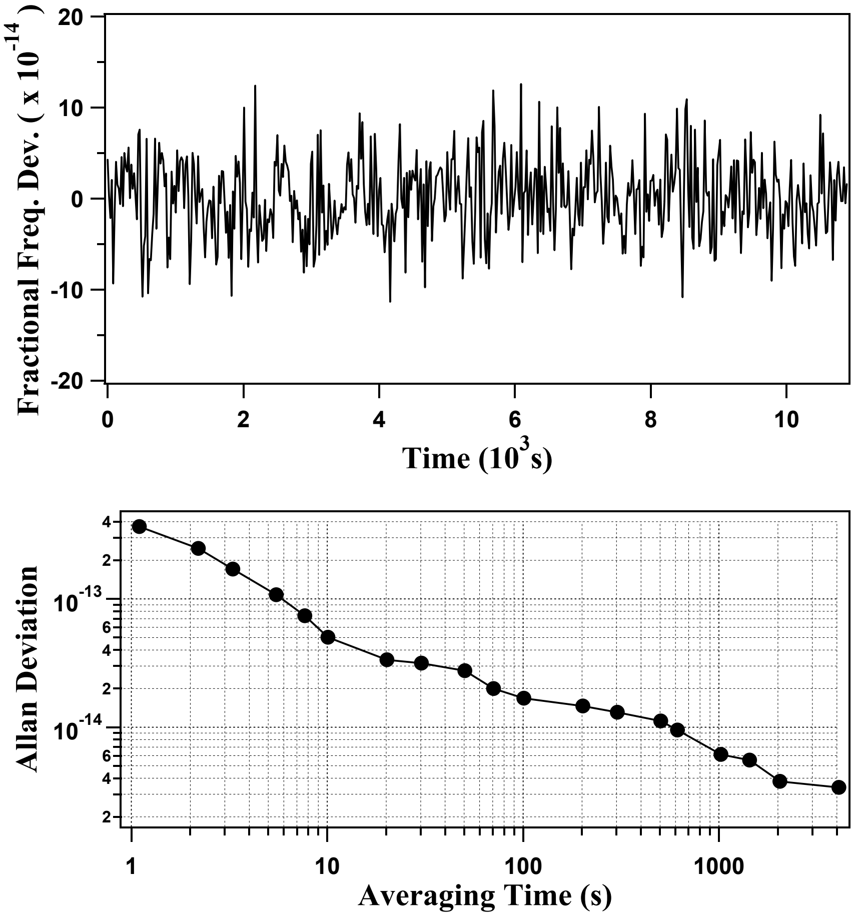

Figure 4.5: (a) The fractional frequency fluctuations of the repetition frequency of the comb when stabilizing the comb to the 532-nm I_2 transition. This represents the manner in which the comb would be used in an optical clock. The fluctuations are measured with respect to a maser-referenced microwave source, using a 20-s gate time for frequency counting. (b) The Allan deviation of f_{rep} for the I_2-stabilized comb. It is limited by the transmitted maser-referenced microwave signal against which the fluctuations of f_{rep} are measured.

its performance against another system which has a stability equal to or better than that of the I_2 clock.

The stability of the cw laser locked to the 532-nm I_2 transition is determined by comparing it to the highly cavity-stabilized laser at 563 nm that serves as the local oscillator for the single-mercury-ion-based optical standard at NIST [89]. To span the 250-THz frequency gap between the two lasers, a frequency comb at NIST is stabilized to the mercury-ion-standard local oscillator to transfer its stability across the visible spectrum. This enables the measurement of any visible frequency at an instability of 4×10^{-15} for a 1-s averaging time [86]. The stability of the I_2-stabilized laser is measured by transferring it over an optical fiber to NIST and counting the heterodyne beat between this laser and the stabilized frequency comb. Figure 4.6 shows the Allan deviation of the I_2-stabilized laser computed from data taken with a 1-s counter gate time, taken from [86]. It is well above the stability limit of the mercury-ion-standard local oscillator. Also, by comparing the outgoing frequency of the I_2-stabilized laser with that transmitted on a round trip to NIST and back, it is verified that the measured

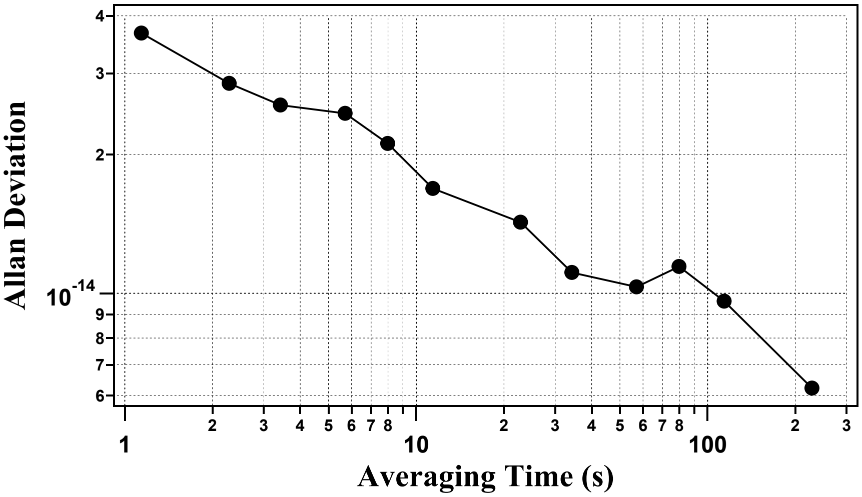

Figure 4.6: The Allan deviation of the I_2-stabilized cw laser is determined by comparing it with the local oscillator for the single-mercury-ion-based optical standard at NIST. This data is obtained from [86]

instability is greater than that introduced by the fiber transmission. Therefore, this represents the true stability of the I_2 system. The instability for a 1-s averaging time is $< 4 \times 10^{-14}$ and reaches 1×10^{-14} at 100 s.

4.2 Cold-atom-based optical standard

Although cell-based standards are compact and reliable, the motion of the atoms or molecules limits the performance of these systems. Laser-cooling the atoms [56] significantly reduces their thermal motion, reducing the Doppler-related linewidth of the transitions by a factor of $\sim 10^{-4}$–10^{-5}. The most stable optical frequency standards being developed are based on transitions in laser-cooled atoms and ions, as discussed in Section 1.4. Strontium is one example of an atom that can be laser-cooled [49, 52, 57] and offers a few extremely narrow transitions that could potentially be used in an optical clock [50, 37, 67]. A simplified energy level diagram of Sr is shown in Fig. 4.7.

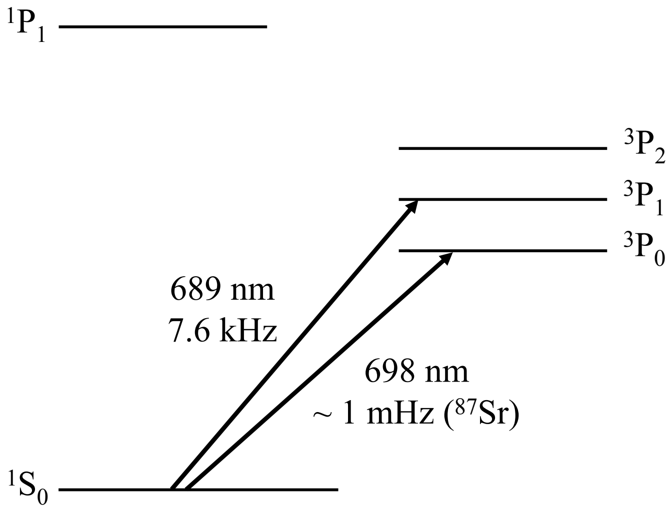

Figure 4.7: The weakly allowed $^1S_0$ – $^3P_0$ transition of ^{87}Sr could potentially be used in an optical clock. The broader $^1S_0$ – $^3P_1$ transition of ^{88}Sr is useful for demonstrating how a frequency comb can be used to link optical and microwave frequencies.

The $^1S_0$ – $^3P_0$ transition is weakly allowed in the fermionic Sr isotope, ^{87}Sr, and this narrow transition is currently being studied for use in an optical clock [74]. The $^1S_0$ –

$^3P_1$ transition of ^{88}Sr is not one of the ultra-narrow transitions, but it is suitable for demonstrating the use of the frequency comb to link optical and microwave frequency standards. Stabilizing the repetition frequency of the comb to a microwave source referenced to a hydrogen maser that is itself referenced to a Cs fountain clock allows the comb to be used for measuring the absolute frequency of this ^{88}Sr transition.

The $^1S_0 - {}^3P_1$ transition of ^{88}Sr is centered at 689 nm, and has a natural linewidth of 7.6 kHz. The cw laser used to probe this transition is pre-stabilized to a high-finesse, stable optical cavity, and is then locked to the $^1S_0 - {}^3P_1$ transition of ^{88}Sr atoms contained in a vapor cell. This produces a probe laser linewidth that is <100 Hz. An AOM is used to shift the frequency of the stabilized probe laser for scanning through the transition of laser-cooled ^{88}Sr atoms. At each position of the scan, the atoms are cooled to $\sim$1.3 μK in a magneto-optical trap, and then the trapping laser beams and the quadrupole magnetic field are turned off before probing the atoms and measuring the fluorescence. Measuring the absolute frequency of the probe laser with the maser-referenced comb during the entire scanning period and fitting a Voigt profile to the data of the measured lineshape provides the absolute frequency of the center of the transition [39]. Figure 4.8 shows the measurements of the $^1S_0 - {}^3P_1$ transition frequency obtained over a period of a little more than a month. These data have a statistical uncertainty of 20 Hz, which corresponds to a fractional uncertainty of $\sim 5 \times 10^{-14}$. This is an order of magnitude lower than the fractional uncertainty for the measurement of the 515-nm I_2 transition, which is $300/(582 \times 10^{12}) \sim 5 \times 10^{-13}$. Taking into account systematic uncertainties as well, this measurement is more than 200 times more accurate than previous measurements of this ^{88}Sr transition using a thermal atomic beam [19].

Since the trapping laser beams shift the transition frequency, they must be turned off before probing the atoms, which results in Doppler-related systematic shifts of the transition and limits the interrogation time. In future work, the Sr atoms will be confined in an optical lattice. As discussed in Section 1.4, carefully choosing the frequency

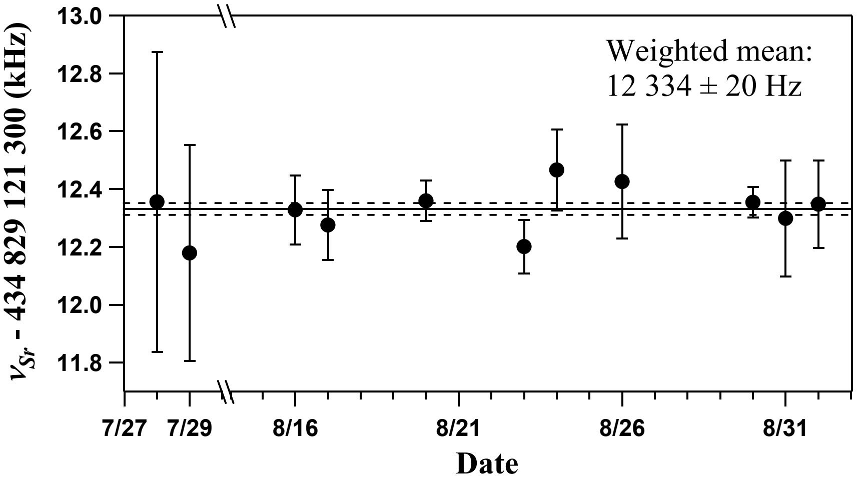

Figure 4.8: The absolute frequency of the $^1S_0$ – $^3P_1$ transition in ultracold ^{88}Sr atoms is measured with the maser-referenced frequency comb over a period of ~1 month. The maser is itself referenced to a Cs fountain clock. The fractional statistical uncertainty of these data is an order of magnitude lower than that for the measurement of the 515-nm I_2 transition.

of the trapping laser for the lattice minimizes the shift of the transition, allowing the atoms to be probed while trapped. This will allow the ultra-narrow $^1S_0$ – $^3P_0$ transition of ^{87}Sr to be studied.

Although the I_2-based optical clock offers portability and a level of simplicity not possible with a clock based on ultracold atoms, laser-cooled atoms obviously provide much higher stability and accuracy. Even the $^1S_0$ – $^3P_1$ transition of ^{88}Sr exhibits an order of magnitude better performance than the I_2 system, and this is not the narrowest clock transition available in Sr. Since systems of laser-cooled atoms or ions are too complex to be easily transported, making use of the superior performance provided only by these systems for the various applications mentioned in Section 1.6 requires the development of reliable transfer of a frequency reference linked to the optical standard. This is also necessary for studying the stability and accuracy of these systems, which can be accomplished only by comparing against similar systems that are often located in separate laboratories.

Chapter 5

Transferring Ti:sapphire Frequency Comb Stability to a Mode-locked Laser Diode Comb

With the optical frequency comb stabilized to an optical frequency standard, remote transfer of a frequency reference linked to the optical standard can be accomplished by transmitting the comb through an optical fiber. The details of this method of transfer and the advantages it provides over other techniques for frequency distribution over optical fibers is the topic of the next chapter. However, it is not appropriate to directly transmit the stabilized frequency comb from the Ti:sapphire laser because of the loss in optical fibers at 800 nm. Instead, a mode-locked laser that produces a comb at 1550 nm is used for this transmission to minimize the losses during transfer. A frequency comb at this wavelength is also useful for establishing a frequency reference grid for dense wavelength-division multiplexed telecommunication systems. Of course, the 1550-nm comb must also be stabilized to the optical frequency standard prior to transmission over the fiber. It has the same two degrees of freedom as the Ti:sapphire frequency comb: its repetition frequency ($f_{rep,1550}$) and carrier-envelope offset frequency ($f_{ceo,1550}$). Depending on the source of the 1550-nm comb, its stabilization process could very closely resemble the method used to stabilize the Ti:sapphire comb. If the 1550-nm comb covers a full-octave of bandwidth, the self-referencing technique described in Section 2.2 can be used to stabilize $f_{ceo,1550}$. Frequency-doubling the comb to achieve spectral overlap with a visible optical frequency standard allows the heterodyne beat between the standard and the comb to be used for stabilizing $f_{rep,1550}$.

An attractive laser for producing a 1550-nm comb is a mode-locked laser diode (MLLD) because of its compact size, efficiency, and reliability. However, the MLLD produces a comb with a bandwidth of only ~1 nm. This precludes the use of the self-referencing technique for stabilizing $f_{ceo,1550}$ and makes it very unlikely that frequency-doubling provides spectral overlap with the optical frequency standard. Therefore, an alternate technique must be used to stabilize the MLLD frequency comb to the optical standard. Instead of directly stabilizing the MLLD to the standard, the Ti:sapphire comb is used to transfer the stability of the standard to the MLLD. The Ti:sapphire laser is stabilized directly to the frequency standard, and the free parameters of the MLLD, $f_{rep,1550}$ and $f_{ceo,1550}$, are locked to those of the Ti:sapphire comb, $f_{rep,800}$ and $f_{ceo,800}$, as described in Section 2.4.

5.1 Synchronizing mode-locked laser diode to Ti:sapphire laser

The first step toward stabilizing the MLLD to the Ti:sapphire laser — phase-locking the two repetition frequencies, $f_{rep,1550}$ and $f_{rep,800}$ — is equivalent to synchronizing the pulse trains from the two lasers. A schematic of the MLLD is shown in Fig. 5.1. The semiconductor chip consists of a 500-μm gain section and a 50-μm saturable-absorber section that provides the mode-locking mechanism. An antireflection (AR) coating is applied to the facet adjacent to the gain section and a collimating graded-index lens is used to couple into an external cavity of adjustable length. The cavity length is set by a PZT-actuated high-reflecting (99.8 %) mirror. The facet adjacent to the saturable-absorber section serves as the output coupler. The parameters of the MLLD comb can be controlled by adjusting the diode injection current, I_d, and the DC bias voltage, V_s, applied to the saturable-absorber section. When a microwave frequency signal matching the extended-cavity mode spacing is applied to the saturable-absorber

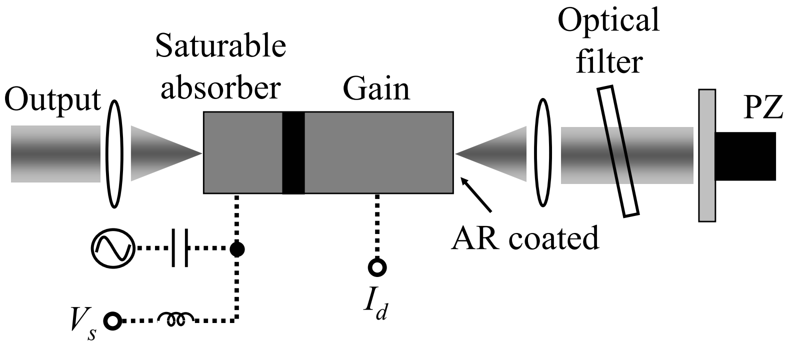

Figure 5.1: The MLLD consists of a 500-μm gain section and a 50-μm saturable-absorber section providing the mode-locking mechanism. The repetition frequency is set by the PZT-actuated external-cavity mirror. Applying a modulation signal to the saturable absorber results in hybrid mode locking. An optical filter is used to reduce the residual jitter between this applied signal and the laser. The MLLD parameters can be controlled using the diode injection current, I_d, and the DC bias voltage, V_s, applied to the saturable absorber. AR, antireflection.

section in addition to the DC bias, fundamental hybrid mode locking is established. This enables injection locking of the MLLD repetition frequency to the applied microwave signal. In the absence of a modulation signal the MLLD passively mode locks at the repetition frequency determined by the optical length of the extended cavity. Previous work has shown that the dominant source of residual timing jitter between the MLLD and a microwave reference to which it is injection locked is the amplified spontaneous emission quantum noise being projected onto the pulse train [42]. This is identical to the Gordon-Haus effect for soliton communication systems [24]. This effect is reduced by using an optical bandpass filter in the external cavity of the MLLD to restrict its bandwidth, and a residual timing jitter has been achieved as small as 40 fs when integrated from 10 Hz to 10 MHz and 60 fs when the integration is extended to 100 MHz [42]. We use an optical filter that has a bandwidth of ∼1 nm and produces pulses with a duration of ∼8 ps.

To improve upon these earlier results and further suppress the relative timing jitter between the MLLD and the Ti:sapphire laser, a wide-bandwidth high-gain electronic feedback loop is used to effectively replace the spontaneous emission quantum noise within the servo bandwidth with the lower measurement noise of the servo loop [45]. A multi-megahertz-bandwidth feedback loop acting on the diode injection current is developed to minimize the relative phase noise between the repetition frequencies of the two lasers. The setup is shown in Fig. 5.2 along with the components used to measure the residual jitter between the lasers. The repetition frequency of the MLLD is ~5.98 GHz (fundamentally mode locked), which corresponds to the eighth harmonic of the ~750-MHz repetition frequency of the Ti:sapphire laser stabilized to the optical standard. Variations of the diode injection current, I_d, and the saturable-absorber DC bias voltage significantly affect the MLLD repetition frequency, $f_{rep,1550}$. When the passively mode-locked $f_{rep,1550}$ is tuned to within the injection-locking range (~10 MHz) of the eighth harmonic of the Ti:sapphire repetition frequency, $f_{rep,800}$, applying the eighth harmonic signal to the saturable absorber results in the injection locking of $f_{rep,1550}$ to $f_{rep,800}$. The remaining phase noise between these frequencies is measured by detecting the pulse train from the MLLD to produce an $f_{rep,1550}$ signal and mixing this signal with the eighth harmonic of $f_{rep,800}$. Using a phase shifter in the $f_{rep,1550}$ path before the mixer to match the phases of the two mixed signals produces an output from the mixer that is proportional to the phase fluctuations between these signals, as described in Section 2.5. After appropriate filtering and amplification of this error signal, it is applied to the MLLD to correct the phase of $f_{rep,1550}$. A high-speed servo loop (~5-MHz unity gain frequency) acts on I_d, and a slow servo controls the external-cavity length using the PZT-actuated mirror. Actuating the external-cavity mirror and I_d through the laser diode current driver provides a 200-kHz feedback bandwidth, limited by the response of the current driver. Establishing a direct and physically short current-feedback path between the mixer generating the error signal and the laser diode current-injection port extends the bandwidth to 5 MHz. In this fast current path only a passive phase-compensating filter network and a resistive attenuator are used. The

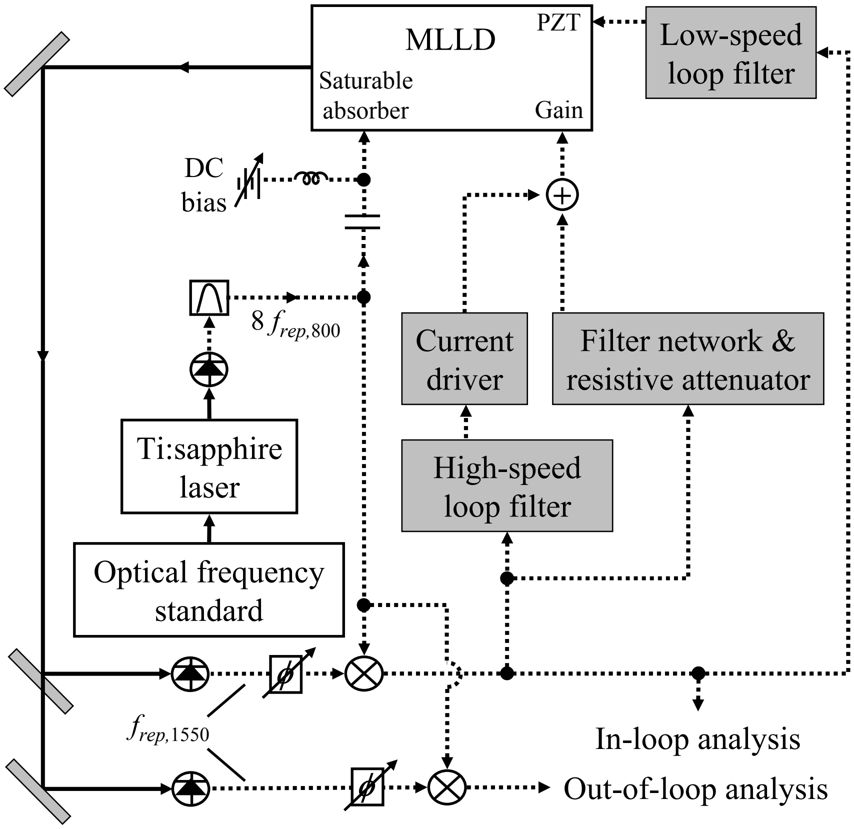

Figure 5.2: For stabilization of $f_{rep,1550}$, the eighth harmonic of $f_{rep,800}$ from the stabilized Ti:sapphire laser is first used to injection lock the MLLD. Mixing the two repetition frequency signals provides the residual phase noise of $f_{rep,1550}$. After appropriate filtering and amplification of this error signal, it is applied to the diode injection current and the PZT-actuated external-cavity mirror to reduce the timing jitter of the MLLD. A second detector, mixer, and phase shifter enable out-of-loop analysis of the MLLD stability.

mixer output signal is also used to perform in-loop measurements of the residual phase noise for characterizing the servo performance. A second photodetector, mixer, and matching phase shifter enable out-of-loop analysis of the remaining phase noise, providing the ultimate test of the reliability of transferring the Ti:sapphire stability to the repetition frequency of the MLLD.

The effect of the stabilization of the MLLD repetition frequency on its spectral lineshape is shown in Fig. 5.3. It shows the RF spectrum of the repetition frequency detected under various operating conditions. The lineshape shown with open circles corresponds to free-running, passively mode-locked operation. A large amount of the phase noise near the carrier is suppressed when the MLLD is injection locked to the repetition frequency of the Ti:sapphire laser, shown by the filled diamonds. When the 200-kHz feedback loop is engaged, the phase noise near the carrier is further suppressed, but visible servo noise bumps near 200 kHz from the carrier appear (solid trace). The 5-MHz loop completely eliminates the 200-kHz bumps and replaces them with servo bumps near 5 MHz from the carrier, shown with open triangles. Although the width of the carrier signal is not resolved by the resolution bandwidth of the RF spectrum analyzer, it is evident that the amplitude of the carrier is enhanced when this wide-bandwidth servo is activated.

The in-loop and out-of-loop spectral densities of the residual rms timing jitter between the MLLD and the Ti:sapphire laser are determined by analyzing the in-loop and out-of-loop mixer outputs with an FFT for Fourier frequencies up to 100 kHz and an RF spectrum analyzer for higher frequencies, and converting to the jitter spectral densities using Eqn. (2.25). An alternate way to express these quantities is the single side band phase noise, $\mathrm{L}(f)$, of the MLLD repetition frequency, which is computed using Eqns. (2.21) and (2.22). Figure 5.4(a) shows the residual rms jitter spectral density (left axis) and the single side band phase noise (right axis) when the MLLD is injection locked to the Ti:sapphire laser without any electronic feedback (solid trace), and the

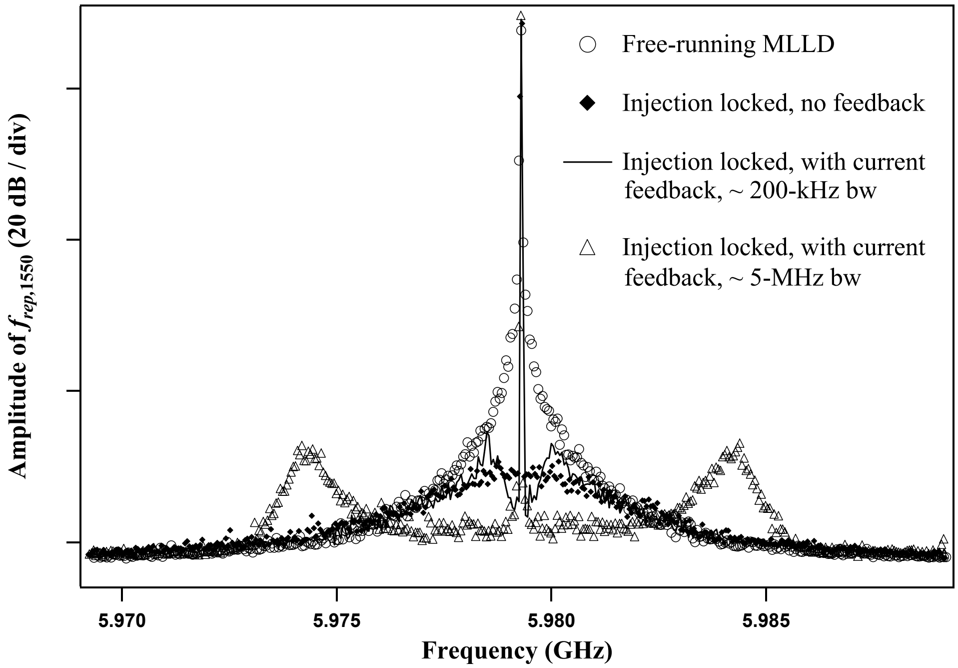

Figure 5.3: The RF spectrum of $f_{rep,1550}$ shows the progressive narrowing of its linewidth for free-running, passively mode-locked operation, then injection locking with no feedback, and then 200-kHz-bandwidth current feedback followed by 5-MHz-bandwidth feedback. bw, bandwidth.

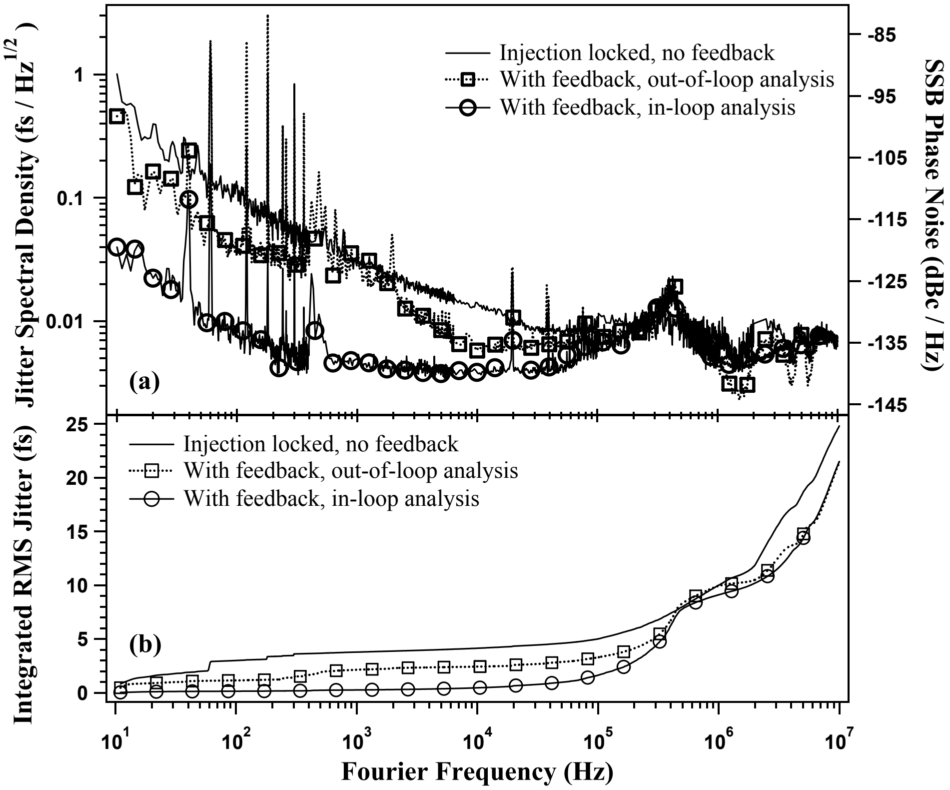

Figure 5.4: (a) The residual rms timing jitter spectral density (left axis) between the MLLD and the Ti:sapphire laser to which it is stabilized and the equivalent single side band (SSB) phase noise of $f_{rep,1550}$ (right axis). The jitter is measured using both in-loop and out-of-loop analysis. Feedback to the diode injection current reduces the jitter within the servo bandwidth. (b) The integrated rms jitter over a bandwidth from 10 Hz to f_h versus f_h, computed after removing the spikes in the spectral densities that are due to line noise. The majority of the jitter over the full bandwidth from 10 Hz to 10 MHz is contributed from 1 MHz to 10 MHz, resulting in 21 fs of total jitter for both in-loop and out-of-loop analysis.

in-loop (circles) and out-of-loop (squares) values with the wide-bandwidth servo activated. The residual jitter measured in-loop is suppressed by approximately one order of magnitude by the feedback over a Fourier frequency range extending from 10 Hz to ∼100 kHz. However, a true characterization of the stability of the MLLD repetition frequency is provided by the out-of-loop measurement. Over this same frequency range the out-of-loop jitter is approximately a factor of 2 lower than the injection-locked value. This higher jitter compared to the in-loop analysis is caused by the differential noise between the in-loop and out-of-loop optical path lengths and fiber-coupled photodetectors. This differential noise cannot be distinguished from instabilities of the MLLD and is attributed to mechanical vibrations. At ∼400 kHz there is a servo bump in both the in-loop and out-of-loop jitter spectral densities, caused by the feedback to the diode injection current through the current driver and corresponding to the set of servo bumps in Fig. 5.3 closest to the carrier. At frequencies above this bump the 5-MHz loop suppresses the in-loop and out-of-loop jitter by approximately the same amount from the solely injection-locked jitter.

Figure 5.4(b) shows the total rms timing jitters over a bandwidth from 10 Hz to f_h versus Fourier frequency f_h for the spectral densities in Fig. 5.4(a). These are determined using Eqn. (2.26) after removing the spikes from the spectral densities that are due to line noise. Over a bandwidth from 10 Hz to 100 kHz (below any effects from the servo bump) the in-loop rms timing jitter is reduced from an injection-locked value of 5 fs to 1.5 fs. The out-of-loop jitter is 3 fs. However, the majority of the total jitter is contributed from the noise in the frequency range from 1 MHz to 10 MHz, where the in-loop and out-of-loop spectral densities are equivalent. The total rms jitter from 10 Hz to 10 MHz is 21 fs, for both the in-loop and out-of-loop data. Above 10 MHz the injection-locked jitter rapidly falls off [42], and we find from data taken over a frequency range of 1 Hz to 100 MHz (not shown) with the wide-bandwidth servo activated that the total rms jitter is increased only slightly to ∼22 fs. Care must be taken to not apply too much gain to the 5-MHz loop leading to a servo bump at that offset frequency, as shown in Fig. 5.3. Although this produces greater noise reduction in the Fourier frequency range up to 1 MHz, the total rms timing jitter increases. For the data in Fig. 5.4, the feedback parameters are optimized to achieve the lowest possible total jitter.

For applications to frequency metrology, where the MLLD is used to transmit a frequency reference against which another frequency is measured, it is useful to determine the instability of the MLLD repetition frequency with respect to the Ti:sapphire laser for various averaging times. This is measured using a single-side-band generator [40, 41] to produce a signal with a frequency that is shifted by 10 kHz from the MLLD repetition frequency obtained with the out-of-loop photodetector. The design of the single-side-band generator is very similar to the one described in Appendix B, which is used in some of the measurements discussed in Chapter 6. The shifted repetition frequency and the eighth harmonic of the Ti:sapphire repetition frequency are then mixed to produce a 10-kHz signal that can be filtered and counted with sufficient resolution to detect the fluctuations of $f_{rep,1550}$. The Allan deviation of $f_{rep,1550}$ for averaging times less than 10 s is calculated by combining data taken with a 1-s gate time, and the Allan deviation for longer averaging times is obtained from data taken with a 10-s gate time. Figure 5.5 shows the Allan deviation of $f_{rep,1550}$ with (closed circles) and without (open circles) feedback to the laser, as well as the noise floor of the measurement system (open squares). These data indicate that the feedback reduces the instability of $f_{rep,1550}$ to the measurement noise floor for averaging times less than $\sim$20 s. The Allan deviation of the stabilized $f_{rep,1550}$ decreases with averaging time between τ^{-1} and $\tau^{-1/2}$ (for averaging time τ), indicating a mixture of white frequency and phase noise.

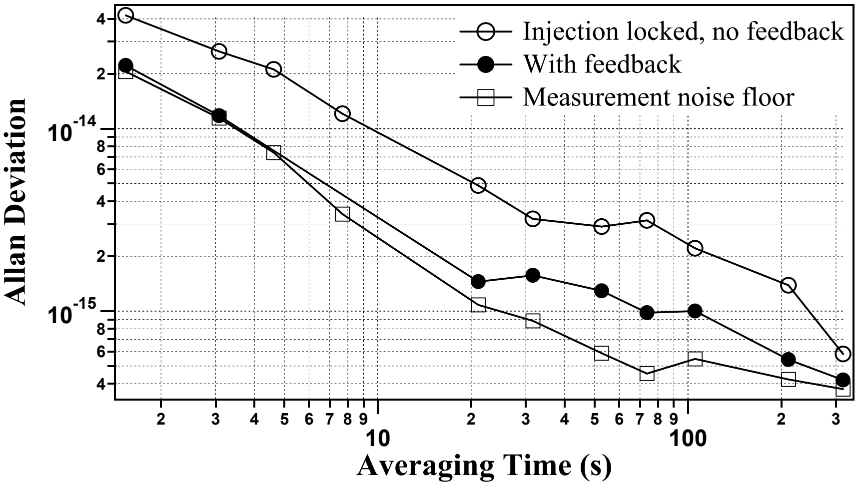

Figure 5.5: The Allan deviation of $f_{rep,1550}$ indicates that the feedback reduces the instability of $f_{rep,1550}$ to the measurement noise floor for averaging times < 20 s.

5.2 Detecting mode-locked laser diode carrier-envelope offset frequency

Once the repetition frequency of the MLLD has been stabilized to that of the Ti:sapphire laser, the second and final step toward transferring the stability of the Ti:sapphire comb to the MLLD comb is to stabilize the carrier-envelope offset frequency, $f_{ceo,1550}$, of the MLLD. Since its narrow bandwidth prohibits the use of the self-referencing technique to detect the fluctuations of $f_{ceo,1550}$, they must be determined by using the stabilized Ti:sapphire comb as a reference as described in Section 2.4. $f_{ceo,1550}$ is measured by frequency-doubling the MLLD comb to achieve spectral overlap with the Ti:sapphire comb and detecting the heterodyne beat between components from each comb. To see how $f_{ceo,1550}$ is given in terms of the heterodyne beat frequency, f_{beat}, and the repetition frequency, $f_{rep,800}$, and carrier-envelope offset frequency, $f_{ceo,800}$, of the Ti:sapphire laser, consider that each comb component of the MLLD is given by

$$\nu_{1550} = m f_{rep,1550} + f_{ceo,1550} \tag{5.1}$$

and each component of the Ti:sapphire comb is given by

$$\nu_{800} = n f_{rep,800} + f_{ceo,800} \tag{5.2}$$

Therefore, the heterodyne beat between a component of the frequency-doubled MLLD comb and the Ti:sapphire comb can be expressed as

$$\begin{aligned} f_{beat} &= 2\nu_{1550} - \nu_{800} \\ &= (16m - n) f_{rep,800} + 2 f_{ceo,1550} - f_{ceo,800} \end{aligned} \tag{5.3}$$

using the fact that the repetition frequency of the MLLD is stabilized to the eighth harmonic of the repetition frequency of the Ti:sapphire laser. From Eqn. 5.3 it is clear there will be a pair of beats at $nf_{rep,800} + [(2f_{ceo,1550} - f_{ceo,800})$ modulo $f_{rep,800}]$ and $nf_{rep,800} + [(f_{ceo,800} - 2f_{ceo,1550})$ modulo $f_{rep,800}]$ between every pair of consecutive $f_{rep,800}$ harmonics, $nf_{rep,800}$ and $(n+1)f_{rep,800}$. Without loss of generality, we can identify the lowest-frequency heterodyne beat with $2f_{ceo,1550} - f_{ceo,800}$ since the magnitudes of the fluctuations of all the beats are equal to the magnitude of the fluctuations of $2f_{ceo,1550} - f_{ceo,800}$. The residual rms fluctuation, σ_{rms}, of one of the heterodyne beats stabilized to a frequency reference is equivalent to the residual rms fluctuation of $2f_{ceo,1550} - f_{ceo,800}$, and assuming $f_{ceo,800}$ has been sufficiently well stabilized, the residual rms fluctuation of $f_{ceo,1550}$ is given by $\sigma_{rms}/2$.

The experimental setup for detecting the heterodyne beat between the stabilized Ti:sapphire comb and the frequency-doubled MLLD comb is shown in Fig. 5.6. The optical output of the MLLD is first amplified by two stages of erbium-doped fiber amplifiers (EDFAs) and then focused into a 2-cm-long piece of periodically poled lithium niobate (PPLN) for second-harmonic generation (SHG). The temperature of the PPLN is tuned with a manually-controlled thermo-electric cooler to maximize the efficiency of SHG, producing 35 μW centered at 775 nm. The frequency-doubled MLLD comb is then collinearly combined with the Ti:sapphire comb using a polarizing beamsplitter.

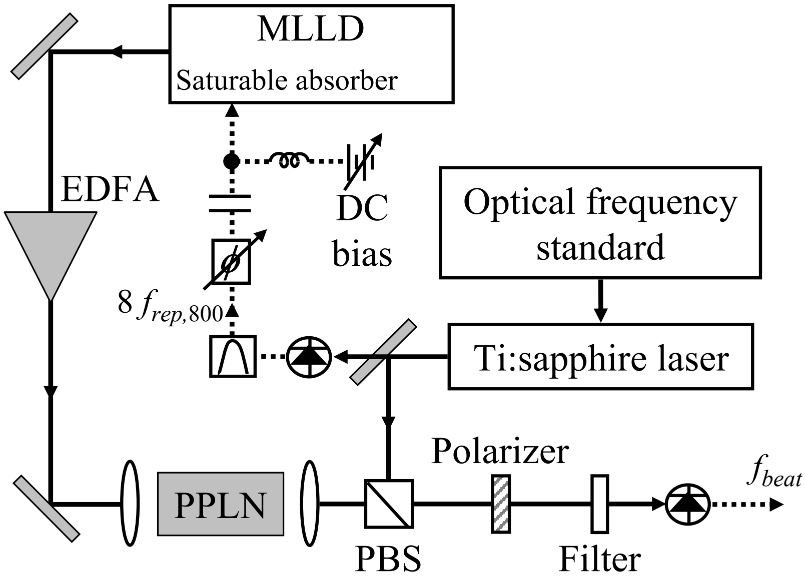

Figure 5.6: The fluctuations of $f_{ceo,1550}$ are determined by detecting the heterodyne beat, f_{beat}, between the frequency-doubled MLLD comb and the stabilized Ti:sapphire comb. Injection locking the MLLD to the Ti:sapphire laser with an adjustable phase shifter in the $f_{rep,800}$ path ensures the simultaneous arrival of pulses from both lasers at the photodetector. EDFA, erbium-doped fiber amplifier; PPLN, periodically poled lithium niobate; PBS, polarizing beamsplitter.

A polarizer is used to pass the common projections of the two orthogonal polarizations and an interference filter is used to block the portion of the Ti:sapphire spectrum not overlapping with the doubled MLLD spectrum. Detection of the heterodyne beat between the two combs is achieved once the pulses from each laser are synchronized and are incident simultaneously on the photodetector. Injection locking the MLLD repetition frequency to the eighth harmonic of the Ti:sapphire repetition frequency with an adjustable phase shifter in the path of the microwave signal driving the MLLD saturable absorber ensures that every eighth MLLD pulse arrives at the photodetector simultaneously with a pulse from the Ti:sapphire laser.

5.3 Orthogonal control of both mode-locked laser diode parameters

Using the Ti:sapphire comb as a reference for measuring $f_{ceo,1550}$ makes it possible to monitor the offset frequency of the MLLD while applying feedback to its injection current for stabilization of $f_{rep,1550}$ as described in Section 5.1. This reveals that the linewidth of the heterodyne beat signal between the MLLD and Ti:sapphire combs becomes substantially broadened when the synchronization servo is activated, indicating a strong coupling of the injection current to $f_{ceo,1550}$. Likewise, attempts to stabilize the heterodyne beat frequency using either the bias voltage, V_s, applied to the saturable-absorber section of the MLLD or the diode injection current, I_d, lead to a significant increase in the relative timing jitter between the two lasers. Since both of the adjustments available for the MLLD providing high-bandwidth feedback affect both $f_{rep,1550}$ and $f_{ceo,1550}$, the simultaneous stabilization of both of these parameters requires the design of an orthogonal control scheme that independently affects $f_{rep,1550}$ and $f_{ceo,1550}$ [34].

To design an orthogonal control scheme for the MLLD, it is necessary to determine the dependence of the MLLD free parameters on all the relevant control variables, I_d, V_s, and the adjustments, l_c, to the external cavity length provided by the PZT-actuated cavity mirror. Figure 5.7 shows the measured fractional changes of the MLLD parameters, $f_{rep,1550}$ and $f_{ceo,1550}$, versus the normalized changes in the three control variables. $f_{ceo,1550}$, which represents a uniform shift in the optical frequencies of all the MLLD comb components, is normalized by the optical frequency corresponding to 1550 nm. The control variables are normalized by their maximum operating values. The changes in $f_{ceo,1550}$ are measured with the repetition frequency of the MLLD injection locked to that of the Ti:sapphire laser, but for the measurements of $f_{rep,1550}$ the injection locking is disabled to observe its uncontaminated response to the control variables. When V_s is increased, the bandgap of the saturable absorber shifts to lower energy, resulting in a larger value of refractive index n at the operating wavelength. A smaller bandgap also leads to an increase in the band-edge curvature, and so $\mathrm{d}n/\mathrm{d}\omega$ increases. Therefore, the pulse group velocity decreases, leading to the reduction of $f_{rep,1550}$. This is confirmed by the experimental data, shown as circles in Fig. 5.7(a). The different fractional changes of n and $\mathrm{d}n/\mathrm{d}\omega$ also cause a change in $f_{ceo,1550}$ as V_s increases, as shown by the circles in Fig. 5.7(b). The exact form of the dependence of n and $\mathrm{d}n/\mathrm{d}\omega$ on V_s could be determined from the changes of $f_{rep,1550}$ and $f_{ceo,1550}$, just as the intensity dependence of the Ti:sapphire laser was determined in Chapter 3. When I_d is increased, the enhanced free-carrier density leads to a smaller value of n and therefore a larger value of $f_{rep,1550}$, as shown by the squares in Fig. 5.7(a). Again, measurements of the changes of $f_{rep,1550}$ and $f_{ceo,1550}$ enable the determination of the dependence of n and $\mathrm{d}n/\mathrm{d}\omega$ on I_d. The influence of I_d on $\mathrm{d}n/\mathrm{d}\omega$ is more complicated and is related to the wavelength and the current dependence of the linewidth-enhancement factor (α) of the specific diode structure. Experimentally it is found that increasing I_d causes a reduction in $f_{ceo,1550}$, as shown by the squares in Fig. 5.7(b).

The data in Figure 5.7 have some very important implications for the design of an orthogonal control scheme. First, they show that the dependence of both $f_{rep,1550}$ and $f_{ceo,1550}$ on all three control variables is linear over the full operating ranges of these

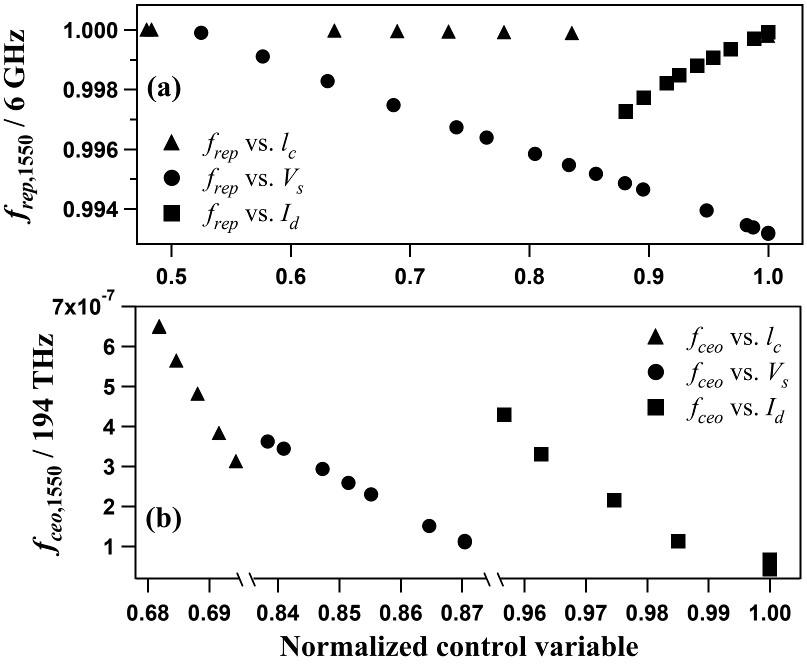

Figure 5.7: (a) The dependence of $f_{rep,1550}$ on the MLLD control variables: external-cavity length adjustments (l_c), saturable-absorber bias voltage (V_s), and diode injection current (I_d). This is determined with the injection locking disabled. (b) The dependence of $f_{ceo,1550}$ on the control variables with $f_{rep,1550}$ injection locked. $f_{ceo,1550}$ is normalized by the 1550-nm optical frequency, and the control variables are normalized by their maximum operating values. The dependence of both $f_{rep,1550}$ and $f_{ceo,1550}$ on all control variables is linear across their tuning ranges. l_c has a negligible effect on $f_{rep,1550}$, but its effect on $f_{ceo,1550}$ is comparable to that of V_s and I_d.

variables. Second, there is a relative opposite sign between the dependence of $f_{rep,1550}$ on V_s and its dependence on I_d, whereas the effects of these two control variables on $f_{ceo,1550}$ exhibit the same sign. This indicates it is possible to form a pair of orthogonal linear combinations of V_s and I_d, one of which only affects $f_{rep,1550}$, and the other which only affects $f_{ceo,1550}$. Third, the effect of l_c on $f_{ceo,1550}$ is of the same order of magnitude as the effects of the other control variables, whereas its effect on $f_{rep,1550}$ is dominated by those of V_s and I_d. At first this may appear alarming since the fractional change in the cavity length is equivalent to the fractional change of $f_{rep,1550}$. However, note that l_c indicates the adjustments in the cavity length provided by the PZT-actuated cavity mirror. Therefore, in Fig. 5.7 it is normalized by the maximum range of motion of the PZT, and so the fractional changes of the cavity are much smaller than the fractional changes of l_c shown in the figure. These conclusions suggest the following strategy for optimal control of the MLLD. The error signals associated with fluctuations of $f_{rep,1550}$ and $f_{ceo,1550}$ should be combined linearly with the appropriate signs and gains to produce two control signals, one for V_s and the other for I_d. This process represents a diagonalization of the 2×2 matrix describing the response of $f_{rep,1550}$ and $f_{ceo,1550}$ to V_s and I_d, leading to two orthogonal control loops for $f_{rep,1550}$ and $f_{ceo,1550}$. In addition, the PZT-actuated external-cavity mirror can be used to provide a higher loop gain in the low-frequency range for stabilizing $f_{ceo,1550}$. The negligible dependence of $f_{rep,1550}$ on l_c indicates the PZT control need not be incorporated into the orthogonalization of the control loop.

The complete servo control for simultaneously stabilizing $f_{rep,1550}$ and $f_{ceo,1550}$ to a stabilized Ti:sapphire laser is shown in Fig. 5.8. The repetition frequency of the MLLD is injection locked to the eighth harmonic of the Ti:sapphire laser and then mixed with this harmonic to detect the residual phase noise between the two repetition frequencies, just as shown previously in Fig. 5.2. It is observed that as the injection power to the MLLD saturable absorber is increased, the stability of $f_{rep,1550}$ is enhanced, but

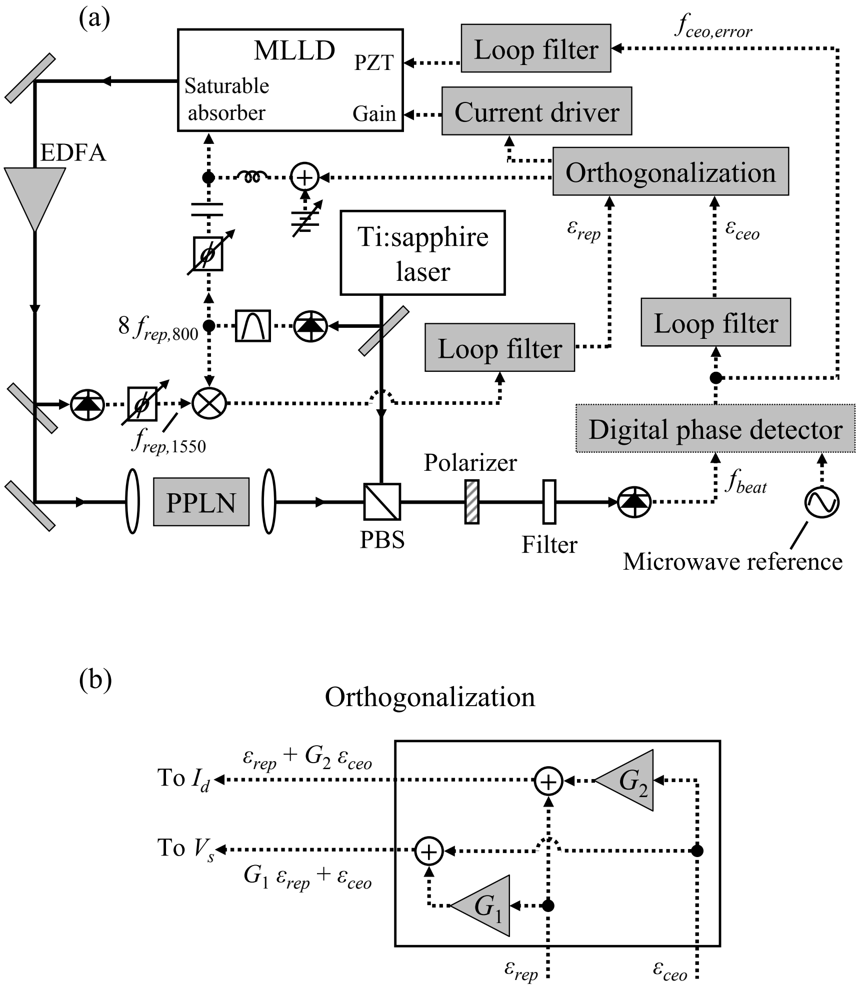

Figure 5.8: (a) The setup for simultaneously stabilizing $f_{rep,1550}$ and $f_{ceo,1550}$ using the orthogonal control scheme. The MLLD is injection locked and the fluctuations of $f_{rep,1550}$ and $f_{ceo,1550}$ are detected as before. Then the two error signals, ε_{rep} and ε_{ceo}, are combined in an orthogonalization stage to generate two linear combinations to be applied to the diode injection current, I_d, and the saturable-absorber bias voltage, V_s. A portion of the $f_{ceo,1550}$ error signal is also applied to the PZT-actuated external cavity mirror. (b) The details of the orthogonalization stage. Two linear combinations of the error signals are formed by choosing the appropriate values of G_1 and G_2, which can be positive or negative, to provide orthogonal control loops for $f_{rep,1550}$ and $f_{ceo,1550}$.

the free-running stability of $f_{ceo,1550}$ decreases dramatically. For the same reason that orthogonal feedback control is implemented within the loop bandwidth, the microwave injection signal should be simultaneously applied to the saturable absorber and the laser diode injection current with the proper phases and amplitudes such that the effect on $f_{ceo,1550}$ from the modulation applied to the saturable absorber can be minimized. This observation is important for future designs of MLLDs where high-frequency (6 GHz for the current MLLD) modulation capability should be accommodated for both the saturable absorber and the injection current. The heterodyne beat between the MLLD comb and the Ti:sapphire comb is detected after frequency-doubling the MLLD optical output using PPLN, as in Fig. 5.6. The phase fluctuations of the heterodyne beat, or equivalently $2f_{ceo,1550}$, are detected with a digital phase detector using a stable microwave frequency reference. A portion of the $f_{ceo,1550}$ error signal is filtered, amplified, and applied to the PZT-actuated external-cavity mirror of the MLLD for low-frequency corrections of $f_{ceo,1550}$, as explained earlier. Simultaneous stabilization of $f_{rep,1550}$ and $f_{ceo,1550}$ using V_s and I_d is accomplished by first filtering and amplifying the corresponding phase-noise signals to form error signals ε_{rep} and ε_{ceo}, respectively. The orthogonalization stage of Fig. 5.8(a) combines these error signals in two appropriate linear combinations to allow orthogonal control of $f_{rep,1550}$ and $f_{ceo,1550}$. One linear combination is applied to the diode injection current, I_d, and the other to the bias voltage, V_s, of the saturable absorber. Figure 5.8(b) shows the details of this orthogonalization stage, which consists simply of two adjustable gains, G_1 and G_2, that can be positive or negative, and two summing elements. The signal that is applied to I_d is given by $\varepsilon_{rep} + G_2\varepsilon_{ceo}$, while $G_1\varepsilon_{rep} + \varepsilon_{ceo}$ is applied to V_s. G_1 is adjusted such that regardless of the value of ε_{rep}, applying ε_{rep} to I_d while applying $G_1\varepsilon_{rep}$ to V_s has no effect on $f_{ceo,1550}$. The amplification that generates ε_{rep} from the measured phase fluctuations of $f_{rep,1550}$ can then be set to determine the gain of the $f_{rep,1550}$ servo loop without affecting the stabilization of $f_{ceo,1550}$. Likewise, G_2 is adjusted such that $f_{rep,1550}$ is unaffected by ε_{ceo}.

With the implementation of the orthogonal control loop, simultaneous stabilization of both $f_{rep,1550}$ and $f_{ceo,1550}$ can be achieved with essentially no reduction in the stability for either when compared with the results for optimized individual control. Figure 5.9 demonstrates the effect of the orthogonal control loop on $f_{rep,1550}$. The in-loop spectral density of the residual rms timing jitter between the MLLD and the Ti:sapphire laser is shown for three different operating conditions, all of which include injection locking the MLLD to the Ti:sapphire laser. The equivalent single side band phase noise, L(f), of the MLLD repetition frequency is also shown along the right axis. An in-loop measurement is performed because this is sufficient to determine the performance of the orthogonal loop, and any additional noise that would be measured with an out-of-loop analysis is irrelevant to the current discussion. The solid trace indicates the residual jitter when the MLLD is injection locked to the Ti:sapphire laser with no

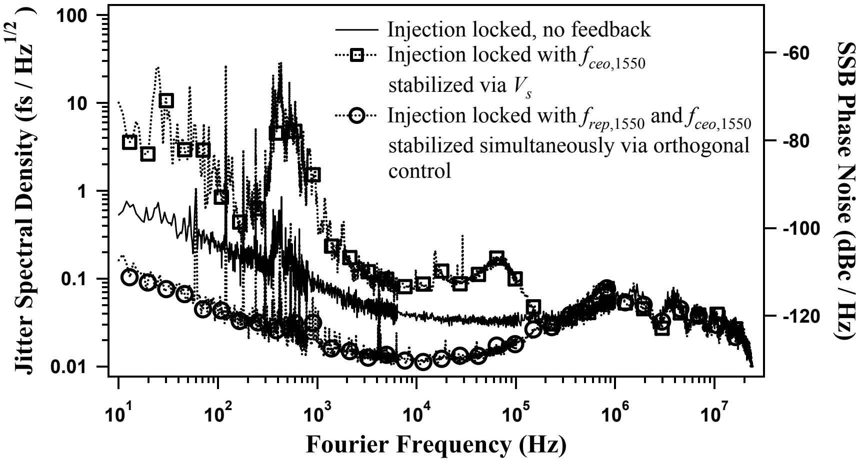

Figure 5.9: The in-loop residual rms timing jitter between the MLLD and the Ti:sapphire laser (left axis) reveals that the stability of $f_{rep,1550}$ becomes worse when it is attempted to stabilize $f_{ceo,1550}$ using the saturable-absorber bias voltage, V_s. Simultaneously stabilizing $f_{rep,1550}$ and $f_{ceo,1550}$ using the orthogonal control loop reduces the timing jitter from the injection-locked value by approximately the same amount as when only $f_{rep,1550}$ is stabilized with the diode injection current. The equivalent single side band (SSB) phase noise of $f_{rep,1550}$ is shown on the right axis.

servo control. Any attempt to stabilize $f_{ceo,1550}$ using only V_s results in an increase in the timing jitter, as shown by the trace denoted with squares. This illustrates the need for orthogonal control of $f_{rep,1550}$ and $f_{ceo,1550}$. When the orthogonal control loop is activated and optimized to simultaneously stabilize $f_{rep,1550}$ and $f_{ceo,1550}$, the timing jitter is reduced from the injection-locked-only value within the bandwidth of the servo loop, to that indicated by the trace denoted with circles. This level of performance for the synchronization loop is basically the same as when only $f_{rep,1550}$ is stabilized using I_d while leaving $f_{ceo,1550}$ unstabilized.

The effect of the orthogonal control loop on the stability of $f_{ceo,1550}$ can be seen by examining the in-loop spectral lineshape of the heterodyne beat between the MLLD and Ti:sapphire combs. Figure 5.10 shows the RF spectrum of this beat for three different experimental conditions. The upper trace represents the beat signal when the MLLD is injection locked to the Ti:sapphire laser with no servo control. When the synchronization is further improved using feedback to I_d as discussed in Section 5.1, the

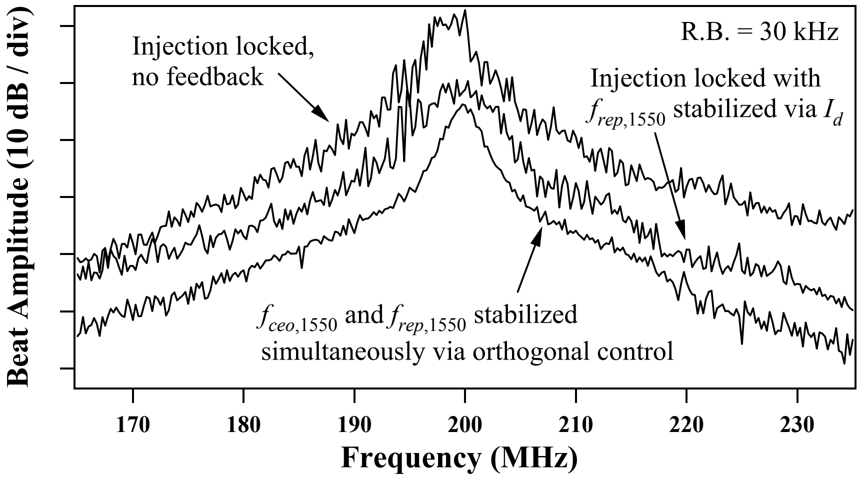

Figure 5.10: The RF spectrum of the in-loop heterodyne beat between the MLLD and the Ti:sapphire laser shows that stabilizing $f_{rep,1550}$ using the diode injection current, I_d, decreases the stability of $f_{ceo,1550}$ as compared with injection-locked-only operation. Simultaneously stabilizing $f_{ceo,1550}$ and $f_{rep,1550}$ using the orthogonal control loop provides the narrowest linewidth for $f_{ceo,1550}$. The traces are vertically displaced for viewing clarity. R.B., resolution bandwidth.

linewidth of the heterodyne beat, and therefore $f_{ceo,1550}$, is broadened. This is shown by the middle trace in Fig. 5.10 and again illustrates the need for orthogonal control. When the orthogonal control loop is activated, not only is the stability of $f_{rep,1550}$ improved as shown in Fig. 5.9, but the stability of $f_{ceo,1550}$ is also improved as shown by the narrowed lineshape of the bottom trace in Fig. 5.10.

Although the linewidth of $f_{ceo,1550}$ is considerably narrowed by the orthogonal control loop, it is clear from Fig. 5.10 that it is still quite broad for applications that require a narrow linewidth or a long coherence time of the MLLD comb components. However, for frequency metrology applications in which the MLLD comb components are used to provide frequency references, the linewidth of $f_{ceo,1550}$ is not as important as how well its center value can be determined by measuring its frequency for a given averaging time. The impact of the orthogonal control loop for this type of application is shown in Fig. 5.11. Frequency-counting records of the deviations from the mean of the in-loop heterodyne beat between the MLLD and Ti:sapphire combs with and without servo control are shown in Fig. 5.11(a). With a 1-s gate time, the rms fluctuation (σ_{rms}) of the heterodyne beat is 1.5 MHz without any stabilization. The orthogonal control loop reduces σ_{rms} to 3.2 mHz. Monitoring the $f_{ceo,1550}$ error signal produced by the digital phase detector ensures that no cycles have slipped over this measurement period. The Allan deviation of the stabilized beat frequency record is shown in Fig. 5.11(b), normalized by the 775-nm optical carrier frequency. The fractional instability is $\sim 1 \times 10^{-17}$ for a 1-s averaging time and decreases with averaging time between τ^{-1} and $\tau^{-1/2}$ (for averaging time τ), indicating a mixture of white frequency and phase noise in the stabilization loop.

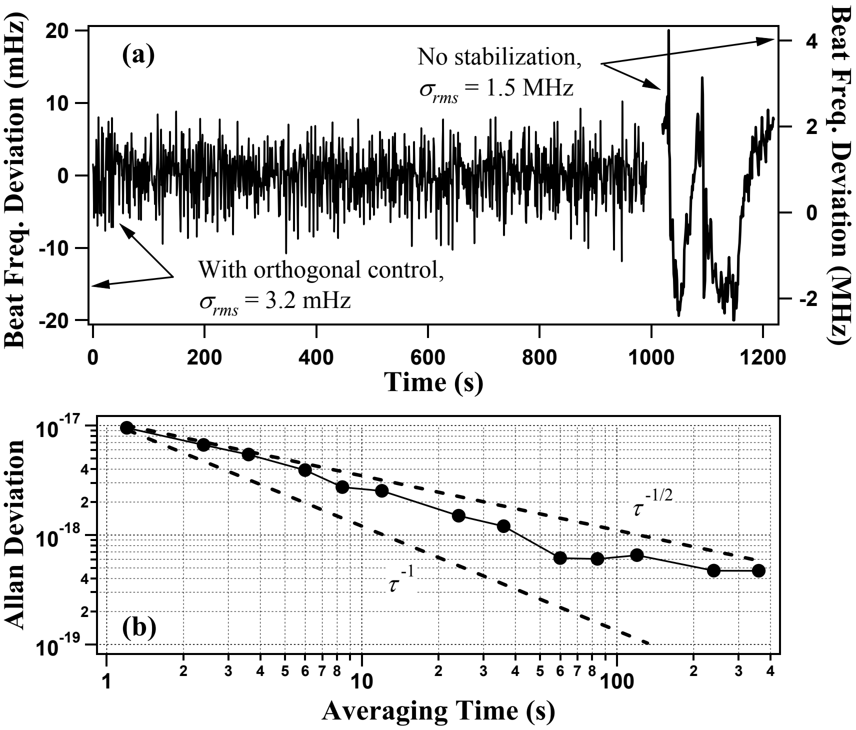

Figure 5.11: (a) The deviations from the mean of the in-loop heterodyne beat between the MLLD and the Ti:sapphire laser, using a 1-s gate time for frequency counting. The orthogonal control loop reduces the rms fluctuation, σ_{rms}, by almost nine orders of magnitude. (b) The Allan deviation of the heterodyne beat using the orthogonal control loop, normalized by the 775-nm optical frequency. It indicates a mixture of white frequency and phase noise.

5.4 Conclusions

Transferring frequency references over optical fibers using an optical frequency comb requires a stabilized comb at 1550 nm. A MLLD is an attractive source of such a comb, but since it cannot be stabilized directly to an optical frequency standard it must be stabilized to a Ti:sapphire frequency comb that is itself stabilized to the optical standard. Using the diode injection current to reduce the residual phase noise between the repetition frequencies of the MLLD and Ti:sapphire laser below that present with injection locking produces a residual relative timing jitter between the two lasers of ∼22 fs, integrated from 1 Hz to 100 MHz. This represents what is to my knowledge the lowest reported residual jitter for a MLLD. However, to transfer the stability of the optical standard to every component of the MLLD comb, its carrier-envelope offset frequency must also be stabilized to that of the Ti:sapphire laser. Since the two high-bandwidth control variables of the MLLD, the diode injection current and the saturable-absorber bias voltage, both affect the repetition frequency and carrier-envelope offset frequency of the MLLD, an orthogonal control scheme is necessary to simultaneously stabilize both free parameters of the MLLD comb.

Careful measurements of the dependence of the MLLD parameters on its control variables indicate that this dependence is linear across the operating ranges of these variables and that it is possible to design an orthogonal control scheme that independently affects the repetition frequency and the carrier-envelope offset frequency. The implementation of this orthogonal control loop results in the simultaneous stabilization of both comb parameters with a performance that is essentially as good as that obtained for optimized independent control. One way in which the orthogonal control could be improved is by incorporating frequency dependence into the gain coefficients [G_1 and G_2 in Fig. 5.8(b)] that determine the linear combinations of the error signals applied to the injection current and saturable-absorber bias voltage. The dependence of the MLLD parameters on its control variables has been measured at DC, and presently the gain coefficients of the linear combinations are constant for all Fourier frequencies. A more sophisticated servo loop would use information on the frequency-dependent response of the MLLD comb to its control variables to determine the frequency dependence of the gain coefficients that would optimize the orthogonal control loop across the entire servo bandwidth.

Finally, although the orthogonal control loop produces a linewidth of the MLLD carrier-envelope offset frequency that is too broad for applications requiring a narrow linewidth or a long coherence time of the comb components, it is important to distinguish the phase noise of the offset frequency and that of the comb components from that of the repetition frequency. Although the entire comb may be uniformly shifting on fast time scales because of a broad carrier-envelope offset-frequency linewidth, the phase noise of the comb spacing, equivalent to the repetition frequency, can still be quite small. The low timing jitter of the stabilized MLLD makes it appropriate for transmitting an ultrastable microwave frequency reference even for applications that require a narrow linewidth.

Chapter 6

Stable Distribution of a Frequency Comb over Optical Fiber

The use of optical fibers for distributing an ultrastable frequency reference offers several of the advantages that fibers provide for communications, such as low loss and scalability. This also allows the distribution of frequency references to take advantage of the existing telecommunication infrastructure. In addition, environmentally isolated fibers can be considerably more stable than free-space paths, especially over short time scales. Currently, there are essentially two distinct methods of transferring a frequency reference over an optical fiber. An optical frequency reference can be distributed by transmitting a stabilized cw laser over the optical fiber, or a microwave frequency can be distributed by amplitude modulating a cw laser at the frequency of the reference and then transmitting the modulated light over the fiber. The microwave frequency reference is recovered at the remote end by detection of the modulation frequency.

For optical-frequency transfer of a reference linked to an optical standard, the cw laser is either directly stabilized to the optical standard, or it is stabilized to a frequency comb that is itself stabilized to the optical standard. Optical-frequency transfer exhibits excellent stability, with a fractional frequency instability introduced by the transfer process of 2×10^{-14} for a 1-s averaging time without active noise cancellation. Actively cancelling the noise introduced during the transfer reduces the fractional instability to only 3×10^{-15} at 1 s [86]. However, despite the low transfer-induced instability, optical-frequency transfer has a few significant limitations. The transmission of an optical frequency offers very little flexibility, so if the remote user wishes to use the frequency reference to measure an optical frequency that is even ∼100 GHz away from the reference (corresponding to a 0.05 % deviation for a 1550-nm reference), an optical frequency comb would be needed by the remote user to span this gap and make this measurement. If a microwave frequency reference is required by the remote user, an optical frequency comb would again be needed. The frequency comb would be stabilized to the optical frequency reference to provide a microwave frequency phase-coherent with the optical reference.

Microwave-frequency transfer of a reference linked to an optical standard can be accomplished by using the microwave frequency generated by a comb stabilized to an optical standard for modulating the transmitted cw laser. Microwave-frequency transfer provides considerably more flexibility than optical transfer since any other microwave frequency can be measured with the microwave reference using readily available microwave electronics. It is this method of transfer that is used to transmit the maser-referenced signal from within the NIST laboratories to our JILA laboratory for some of the measurements discussed in Chapter 4 [86]. Although microwave transfer provides greater flexibility, measurements have shown that the transfer process introduces an order of magnitude more instability as compared to that introduced for optical-frequency transfer [86]. This does not present a significant problem for the transfer of the maser-referenced signal since the instability of the maser is a few parts in 10^{13} at 1 s, but this is a problem for transferring an ultrastable reference linked to an optical standard. Another drawback to the use of microwave transfer is that for the remote user to make an optical frequency measurement, a frequency comb would again be necessary to phase-coherently connect the transmitted microwave reference to the optical frequency.

Transmitting an optical frequency comb over an optical fiber provides a transfer technique with the greatest flexibility that can supplant these two current methods for frequency transfer. It provides the simultaneous transfer of both a microwave and an optical frequency reference. A microwave reference can be obtained by a remote user with a photodetector that provides harmonics of the repetition frequency of the transmitted comb, or equivalently, the spacing of the comb components. Alternatively, it is likely that a component of the transmitted comb or its second or third harmonic is sufficiently close to any optical frequency the remote user wishes to measure. Therefore, after generating the necessary harmonic of the comb, an appropriate optical frequency reference is available to the user. When the transmitted comb is stabilized to an optical frequency standard, both the microwave and optical frequency references it provides exhibit the stability of this optical standard, and the remote user does not have a need for an additional comb as in the two current transfer techniques. In this chapter I will discuss the stability of transferring a microwave reference using an optical frequency comb. We will see that without actively cancelling the noise introduced during transmission (passive transfer), the comb provides a lower instability for microwave transfer than is measured for transfer with a modulated cw laser, reaching the level of stability provided by optical-frequency transfer with a cw laser [33]. We will also see how the stability can be further enhanced using active cancellation of the transfer-induced noise [32].

6.1 Passive transfer of a frequency comb over optical fiber

To study the instability introduced on the repetition frequency of an optical comb during its transfer over an optical fiber, the comb provided by a passively mode-locked fiber laser is used. The fiber laser is a polarization additive pulse mode-locked laser with a ring configuration [58]. A schematic of the laser is shown in Fig. 6.1. The mode-locking mechanism is nonlinear polarization rotation in the fiber coupled with the polarizing beamsplitter, which provides an intensity-dependent loss allowing the laser to minimize intracavity losses by operating with short, intense pulses. Mode-locked operation occurs when the three quarter-wave plates and the half-wave plate in the free-space section are adjusted to the proper orientations. In addition to providing a polarization-dependent

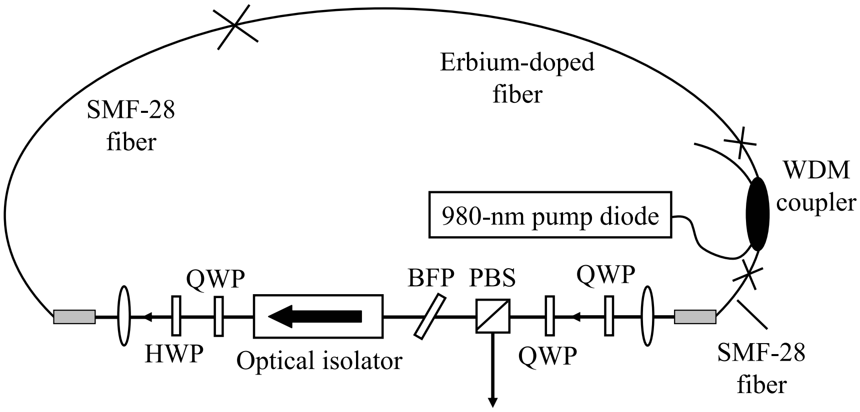

Figure 6.1: The passively mode-locked fiber laser, used to study the instability for transfer of an optical frequency comb, operates by nonlinear polarization rotation. Mode locking occurs with the appropriate orientations of the wave plates in the free-space section. The birefringent filter plate allows tuning of the laser spectrum and the erbium-doped fiber provides the gain. The alternating sections of positive- (erbium-doped fiber) and negative- (SMF-28 fiber) GDD fiber produce a near-zero net GDD, enabling stretched-pulse mode locking. HWP, half-wave plate; QWP, quarter-wave plate; PBS, polarizing beamsplitter; BFP, birefringent filter plate; WDM, wavelength-division-multiplexed.

loss in the cavity, the polarizing beamsplitter also serves as the output coupler. The less-intense wings of the pulses that do not have the appropriate polarization to pass through the beampslitter are rejected from the cavity and constitute the laser output. The optical isolator in the free-space section ensures unidirectional operation in a clockwise orientation, and the birefringent filter plate provides wavelength-dependent polarization rotation, allowing tuning of the laser spectrum. Gain is provided by the section of erbium-doped fiber, which is pumped by a laser diode operating at 980 nm. The pump is coupled into the erbium-doped fiber using a wavelength-division-multiplexed (WDM) coupler. The erbium-doped fiber has positive (normal) group-delay dispersion (GDD), which must be compensated for mode-locked operation. This compensation is provided by the fiber pigtails of the WDM coupler and a section of standard SMF-28 fiber, both which exhibit negative (anomalous) GDD. The alternating sections of large positive and negative GDD cause the pulses to be alternately stretched and compressed as they circulate in the cavity, and so this type of mode-locked fiber laser is often referred to as a stretched-pulse laser. It provides a frequency comb centered at 1550 nm with a bandwidth of ∼10–20 nm and pulses with durations of a few hundred femtoseconds. The fundamental repetition frequency of the laser is ∼97 MHz.

The transfer instability for the repetition frequency of the comb is determined by transmitting the fiber laser comb over two dark, standard telecom fibers installed in the Boulder Research and Administration Network (BRAN) [1]. These two fibers provide a 6.9-km round trip from our JILA laboratory to the NIST laboratories and back. The fibers are installed in underground conduits and steam tunnels in an urban area, and each direction of transmission passes through five breakout panels. Figure 6.2 shows the experimental setup used to determine the instability introduced during transfer. The output of the fiber laser is split into two portions. One portion is transmitted through the BRAN fiber while the second portion provides a reference against which to measure noise introduced by the transmission. The reference and transmitted pulse

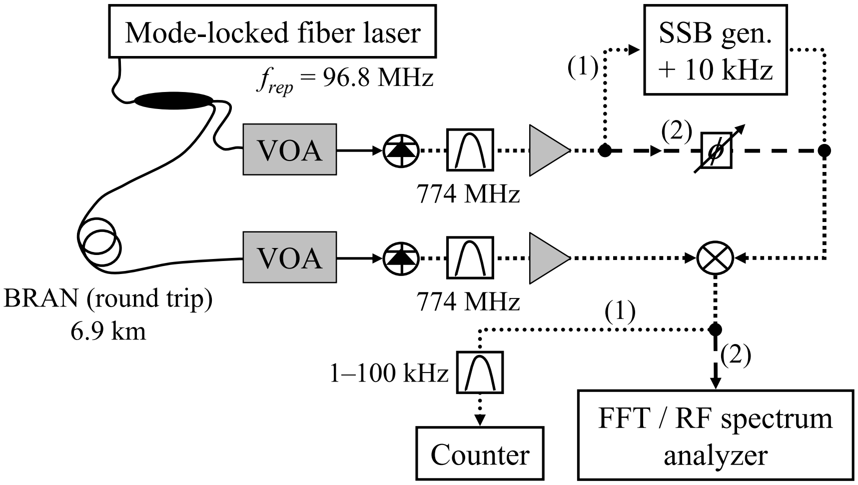

Figure 6.2: The frequency instability for transfer of the comb repetition frequency is measured by counting the 10-kHz mixing product between the shifted reference signal from the mode-locked fiber laser and the transmitted signal [path (1)], using the eighth harmonic of the repetition frequency. Spectral analysis of the timing jitter is performed using the mixing product from the unshifted reference and transmitted signals [path (2)]. SSB gen., single-side-band generator; VOA, variable optical attenuator.

trains are detected with two photodetectors, and the resultant microwave signals are filtered to select the eighth harmonic (774 MHz) of the comb repetition frequency. As we will see, the optical power incident on each photodetector plays a crucial role in determining the detected instability. Fiber-pigtailed variable optical attenuators are used to control the power incident on the photodetectors. To measure the transfer instability for various averaging times, which is relevant for frequency metrology applications, the reference signal is shifted by 10 kHz using a single-side-band generator [40, 41]. The frequency fluctuations introduced by the transfer are measured by filtering and counting the 10-kHz mixing product of the shifted reference signal and the transmitted signal, as described in Section 5.1 for determining the instability of the MLLD repetition frequency. The details of the single-side-band generator are given in Appendix B. This measurement is illustrated by path (1) in Fig. 6.2. For measuring the timing jitter introduced during transfer, relevant for remote synchronization applications, the unshifted reference signal and transmitted signal are mixed. A phase shifter in the path of the reference signal is adjusted so the mixer output provides the timing jitter (or phase noise) introduced during transfer, as described in Section 2.5. This is shown by path (2) in Fig. 6.2. An FFT is used for spectral analysis of the timing jitter for Fourier frequencies up to 100 kHz, and an RF spectrum analyzer is used for higher frequencies. Since the instability introduced by the transfer is measured by comparing the microwave signal from the portion of the fiber laser output transmitted over the BRAN fiber with a reference signal also obtained from the fiber laser output, fluctuations of the fiber laser are common to both signals and do not contribute to the detected noise. Therefore, for these measurements the fiber laser is left free-running.

The Allan deviation of the transmitted repetition frequency is computed from the results of counting the 10-kHz frequency difference between the transmitted signal and the shifted reference. This provides the fractional frequency instability introduced by the transfer process. The Allan deviation for all averaging times is computed by combining data taken with a 1-s gate time. The effect of dead time between the counter gate intervals on the computed Allan deviation is avoided by acquiring the data from the counter in such a way so as to minimize this dead time. This results in a dead time of only ∼5 ms for counting the 10-kHz signal. Data taken with a 3-s gate time and a 10-s gate time are also used to verify that the Allan deviation computed from the 1-s data is unaffected by dead time. Figure 6.3 shows the Allan deviation for the transfer of the eighth harmonic of the mode-locked fiber laser repetition frequency as filled circles. Also shown are the Allan deviations for transferring a microwave frequency using external modulation of a cw laser (open circles) and for transferring an optical frequency using a cw laser (open squares) over the BRAN fiber, taken from [86]. Transfer of a microwave frequency using the comb has an instability of $< 3 \times 10^{-14}$ for a 1-s averaging time, an order of magnitude lower than is measured for microwave transfer with a modulated cw laser. The stability provided by the comb matches the excellent stability exhibited by the transfer of an optical frequency using a stabilized cw laser for averaging times ≤ 30 s. For small averaging times the instability for transfer with the comb decreases as $\tau^{-1/2}$ (for averaging time τ), indicating white frequency noise. The increase in the instability at higher averaging times is caused by long-term length variations of the BRAN fiber due to environmental perturbations.

It is found that to obtain the results shown in Fig. 6.3 for microwave-frequency transfer with the comb, it is crucial to apply the correct level of optical power to the photodetectors producing the transmitted and reference signals. The lowest instability is achieved when the average optical powers are reduced as low as possible while still achieving a sufficient signal-to-noise ratio (SNR) for the recovered microwave signals. This is attributed to there being amplitude noise on the laser pulse train that is converted to phase noise by the photodetectors [41]. Reducing the average optical power reduces the magnitude of the power fluctuations, thereby reducing the magnitude of phase fluctuations produced by the photodetectors. This explains why microwave transfer

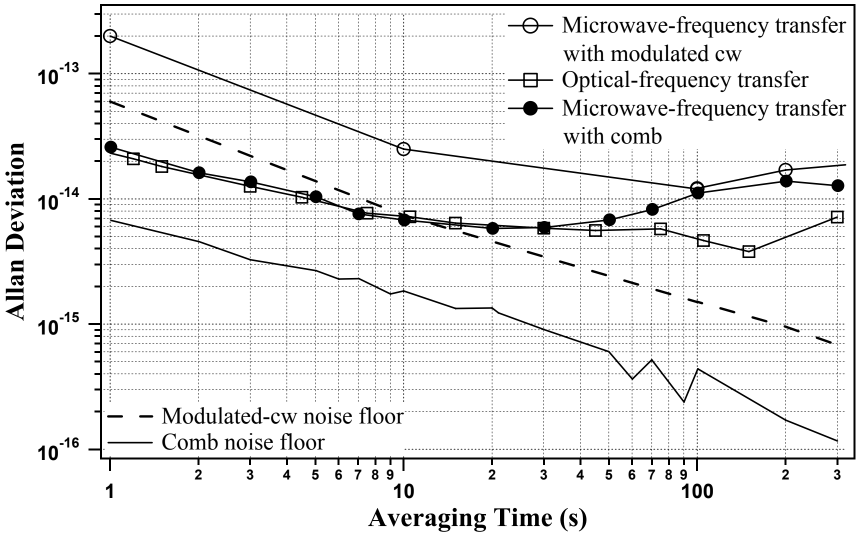

Figure 6.3: The Allan deviation of the microwave frequency transmitted over the BRAN fiber round trip using the comb demonstrates an order of magnitude improvement as compared to transfer using a modulated cw laser and reaches the stability provided by optical-frequency transfer with a cw laser. The noise floors for comb transfer and modulated-cw transfer, obtained using a short piece of fiber, indicate the increased stability for comb transfer is related to noise introduced by the photodetectors. The data for modulated-cw and optical-frequency transfer over the BRAN fiber are obtained from [86].

using the comb has better stability than is measured for a modulated cw laser, since for a given average optical power, the pulse train from the mode-locked fiber laser produces a microwave signal with a higher SNR than is provided by the sinusoidally modulated cw laser (which has a modulation index of only 10 %). To achieve the same SNR, transfer with the modulated cw laser requires the use of a higher average optical power than for transfer with the comb. Attributing the significant difference in the stabilities of these two microwave-frequency transfer methods to noise introduced by the photodetectors and not to noise processes in the transmission path is supported by the measurement noise floors for both transfer methods. The noise floors are obtained by measuring the instabilities after replacing the BRAN fiber with a short piece of fiber. The noise floors for transfer with the comb and with the modulated cw laser are shown in Fig. 6.3 as a solid and dashed curve, respectively. The latter noise floor is obtained from [86]. The noise floor for transfer with the comb is approximately an order of magnitude lower than that for transfer with the modulated cw laser, indicating that the difference in the stabilities of these two transfer methods is indeed due predominately to noise introduced by the photodetectors.

Because of the need to achieve a sufficient SNR for the microwave signal recovered from the transmitted pulse train using the smallest possible average optical power, it is necessary to keep the temporal width of the pulses as small as possible. Shorter pulses provide a larger SNR for harmonics of the repetition frequency than longer pulses when using the same average optical power. However, the 6.9-km BRAN fiber round trip possesses considerable dispersion at 1550 nm and significantly stretches the fiber laser pulses to ~1 ns during transfer. Because of this dispersive stretching, the temporal length of the transmitted pulses can be minimized by reducing the spectral bandwidth of the pulses. Though this results in longer pulses initially, the dispersion of the BRAN fiber does not have as large an impact, resulting in shorter pulses after transfer compared to when broader-bandwidth pulses are used. Therefore, it is expected that the operating parameters of the fiber laser play an important role in minimizing the instability for transfer of harmonics of its repetition frequency.

The effect of the laser parameters on the transfer stability is investigated by varying the pulse bandwidth and average optical power, which is accomplished by adjusting the pump power and polarization components of the fiber laser, thus modifying the mode-locking condition. The results are shown in Fig. 6.4 and summarized in Table 6.1.

Table 6.1: The dependence of the microwave-frequency transfer stability using a comb on the average optical power incident on the receiving photodetector and the SNR of the recovered microwave signal. The SNR is measured with a 1-kHz resolution bandwidth. The symbols correspond to those in Fig. 6.4(a).

Power	**SNR**	**Instability at 1 s**
220 μW	70	6×10^{-14} ($\circ$)
660 μW	80	6×10^{-14} ($\Box$)
160 μW	80	3×10^{-14} ($*$)
30 μW	85	6×10^{-15} (-)

With a pulse bandwidth of 15 nm and 220 μW of average optical power incident on the photodetector receiving the transmitted pulses, the SNR of the recovered microwave signal is only 70 dB, measured with a 1-kHz resolution bandwidth. Because of the relatively low SNR, the measured transfer instability is $\sim 6 \times 10^{-14}$ for a 1-s averaging time, shown by the open circles in Fig. 6.4(a). The corresponding pulse spectrum is shown by the dotted trace in Fig. 6.4(b). For a similar bandwidth of 12 nm and a higher average optical power of 660 μW, it is possible to obtain a SNR of 80 dB. However, because of the large average optical power incident on the photodetector, the instability for a 1-s averaging time is still $\sim 6 \times 10^{-14}$, shown by the open squares in Fig. 6.4(a). This pulse spectrum is shown by the solid trace in Fig. 6.4(b). Reducing the pulse bandwidth to 5.5 nm enables an 80-dB SNR to be achieved with only 160 μW of average optical power. This small optical power provides a transfer instability of $\sim 3 \times 10^{-14}$ at 1 s, shown by the asterisks in Fig. 6.4(a). The pulse spectrum is shown by the dashed trace in Fig. 6.4(b). Because of this dependence on the pulse spectrum, the best results for

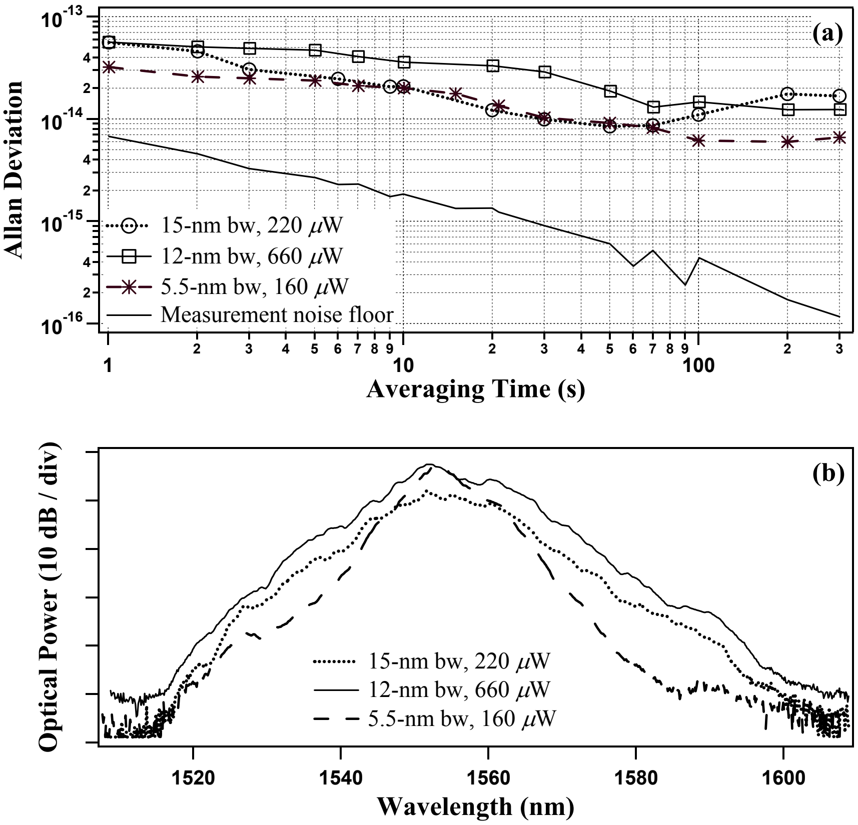

Figure 6.4: (a) The Allan deviation for microwave-frequency transfer over the BRAN fiber using various bandwidths for the transmitted fiber laser pulses and several different average optical powers incident on the receiving photodetector. The lowest instability is achieved with the narrowest bandwidth, which provides a sufficient SNR for the received microwave signal with the smallest optical power. The measurement noise floor, obtained with a short piece of fiber, is the same for all operating conditions. The Allan deviations are computed from data taken with a 1-s gate time. Data taken with a 3-s gate time and a 10-s gate time are used to verify that the Allan deviations computed from the 1-s data are unaffected by dead time. (b) The pulse spectra for the three mode-locking conditions used in (a). bw, bandwidth.

transfer of a microwave frequency with the comb shown in Fig. 6.3 are obtained using the narrowest pulse bandwidth. Without the dispersive BRAN fiber, the measured instability is no longer dependent on the pulse bandwidth, since for all cases the pulses are shorter than the timing resolution of the photodetector. This is indicated by the fact that the measurement noise floor, shown in Fig. 6.3 and again by the solid curve in Fig. 6.4(a), is the same regardless of the operating parameters. For measurement of the noise floor, a SNR of 85 dB is achieved with an average optical power of only 30 μW.

In applications where it is necessary to transfer a frequency reference over optical fiber for long distances, it will be necessary to use erbium-doped fiber amplifiers (EDFAs) to compensate for optical losses in the transmission fiber. Therefore, it is important to determine how an EDFA affects the stability of transferring a microwave reference using a comb. To accomplish this, an EDFA is inserted in the path of the transmitted fiber laser pulses, immediately before the BRAN fiber in Fig. 6.2 and after splitting the laser output. The EDFA provides a factor of ∼4 of optical gain. The effect of the EDFA on the transfer stability of the fiber laser repetition frequency for averaging times greater than 1 s is shown by the Allan deviations in Fig. 6.5. They are computed from

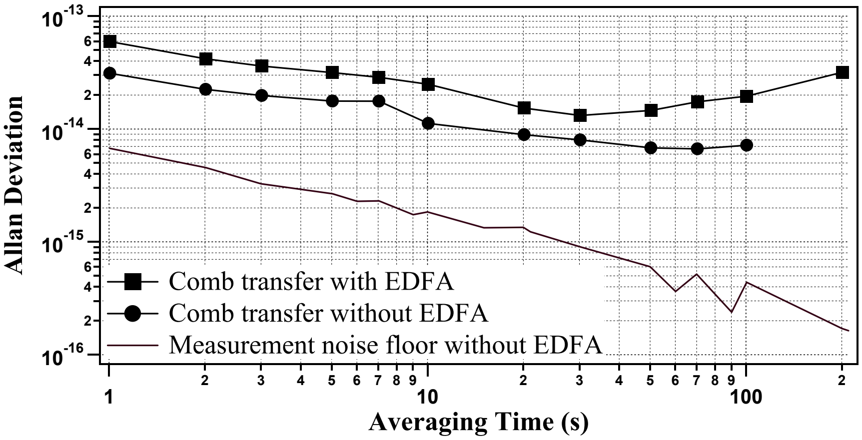

Figure 6.5: The Allan deviation for microwave-frequency transfer over the BRAN fiber using the comb indicates that the introduction of an EDFA providing an optical gain of ∼4 doubles the instability for averaging times ≥ 1 s.

data taken with a 1-s gate time, and data taken with 3-s and 10-s gate times are used to verify that the Allan deviations computed from the 1-s data are unaffected by dead time. For a 1-s averaging time, the fractional frequency instability is increased from $\sim 3 \times 10^{-14}$ (filled circles) to $\sim 6 \times 10^{-14}$ (filled squares) with the introduction of the EDFA. The increased instability is caused by noise introduced by the EDFA itself and is not related to the effects discussed previously of the incident optical power on the photodetector and the SNR of the microwave signal. This is evident since the average optical power incident on the photodetector is less and the recovered SNR is greater with the EDFA in place. With the use of the EDFA the dominant noise process for transfer is still white frequency noise for averaging times less than 30 s.

For remote synchronization applications, it is also important to determine the timing jitter introduced during transfer of the microwave signal, both with and without the use of the EDFA. Spectral analysis of the phase noise obtained by mixing the transmitted and reference microwave signals in Fig. 6.2 provides the spectral density of the transfer-induced timing jitter, using Eqn. (2.25). Figure 6.6(a) shows the spectral density of the rms timing jitter without the EDFA (open circles), with the EDFA (open squares), and for the measurement noise floor without the EDFA (solid trace), obtained by replacing the BRAN fiber with a short piece of fiber. The equivalent single side band phase noise, L(f), of the transmitted eighth harmonic of the fiber laser repetition frequency is shown on the right axis. As expected from the Allan deviations in Fig. 6.5, the EDFA increases the rms timing jitter at low Fourier frequencies. However, for frequencies > 100 Hz the EDFA suppresses the measured jitter. This is not reflected by the Allan deviations for averaging times greater than 1 s, since the behavior at high frequencies is averaged out by the Allan deviation for these longer averaging times. The high-frequency suppression of the jitter by the EDFA is attributed to the long time-constant of its gain dynamics, allowing it to serve as a low-pass filter for amplitude fluctuations of the fiber laser pulse train. This supports the previous conclusion that

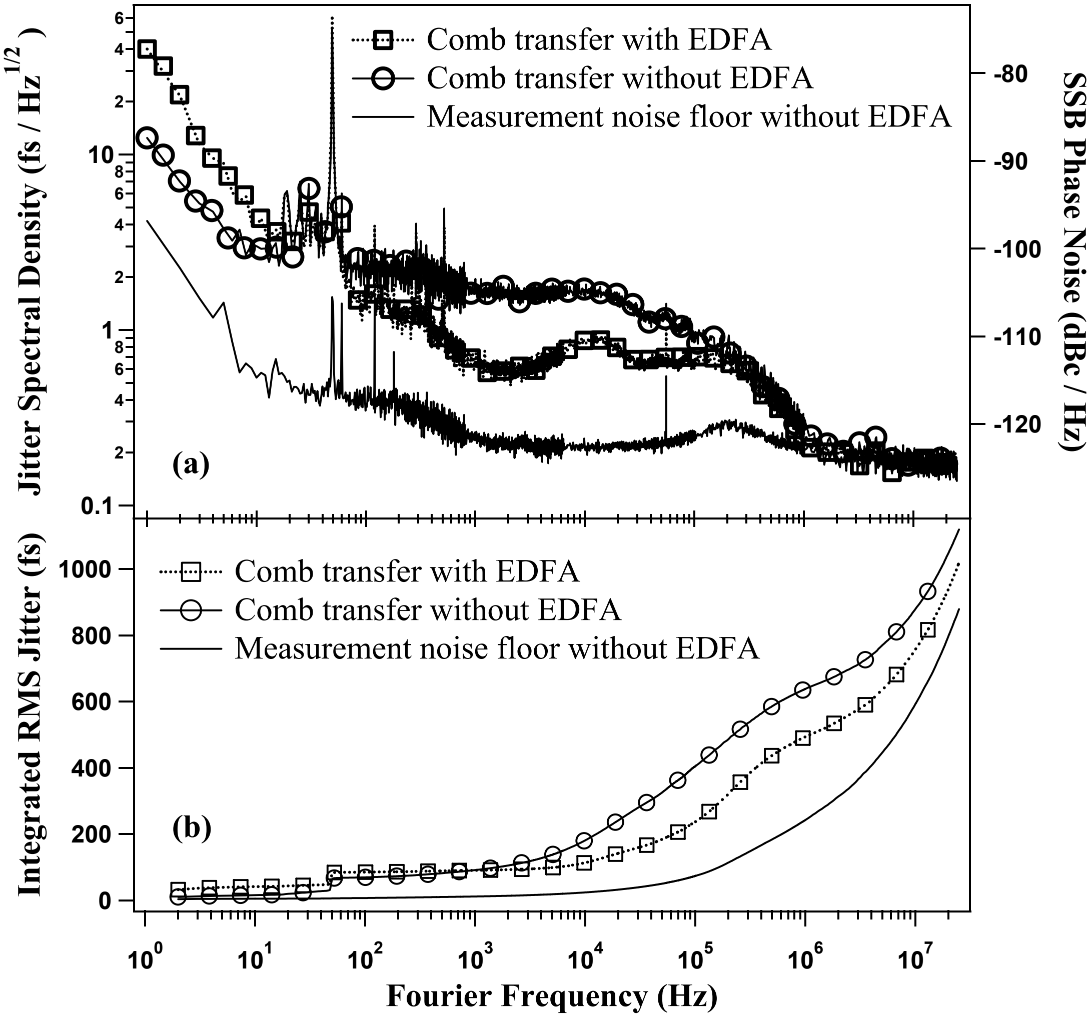

Figure 6.6: (a) The rms timing jitter spectral density (left axis) for transfer over the BRAN and the equivalent single side band (SSB) phase noise of the transmitted eighth harmonic of the comb repetition frequency (right axis), with and without an EDFA. At low Fourier frequencies, the EDFA increases the jitter, but above 100 Hz the EDFA suppresses the jitter. (b) The integrated rms jitter over a bandwidth from 1 Hz to f_h versus f_h. The high-frequency suppression of the jitter by the EDFA results in a lower total jitter integrated over the full measurement bandwidth. The measurement noise floor is determined using a short piece of fiber.

the dominant source of noise is the conversion of amplitude noise on the pulse train to phase noise by the photodetectors. For Fourier frequencies above ∼200 kHz the timing jitter rolls off with a slope of ~ -6 dB / decade, with or without the use of the EDFA. Above 1 MHz the measured jitter is limited by the noise floor of the measurement system. The total rms timing jitters over a bandwidth from 1 Hz to f_h versus Fourier frequency f_h are shown in Fig. 6.6(b), which are obtained from the spectral densities using Eqn. (2.26). The high-frequency suppression of the jitter by the EDFA reduces the total timing jitter over the bandwidth from 1 Hz to 100 kHz from ∼410 fs to ∼240 fs, and when the bandwidth is extended to 1 MHz the total jitter is reduced from ∼640 fs to ∼490 fs.

6.2 Active noise cancellation with dispersion control for transfer of comb

The stability for microwave-frequency transfer using an optical frequency comb can be improved by addressing the two significant noise sources contributing to the transfer instability. First, amplitude fluctuations of the transmitted pulse train are converted to phase noise by the receiving photodetector, as was discussed previously. In the previous discussion it was explained that the instability introduced by this phenomenon is reduced by minimizing the average optical power incident on the photodetector while maintaining a sufficient SNR of the microwave signal, which requires the transmission of the shortest pulses possible. However, the dispersion of the BRAN fiber stretches the pulses to ∼1 ns. The impact of this dispersion has been mitigated by reducing the spectral bandwidth of the fiber laser pulses, but a more effective strategy for minimizing the pulse length is to implement dispersion control of the transmission fiber. This enables the transmission of pulses with durations near the photodetector resolution of 40 ps, eliminating photodetector contributions to the measured instability.

The second significant noise source for microwave-frequency transfer is environmental perturbations of the transmission fiber, causing fluctuations of its optical path length. These produce variations in the group delay of the transmitted pulses, introducing timing jitter and frequency deviations of the transmitted microwave signal. Environmental perturbations of the fiber include temperature variations that induce changes of the fiber length over relatively long time scales, and acoustic noise that introduces high-frequency stresses and length variations. However, an adjustable delay line can be used to compensate for fluctuations in the effective transmission length and reduce the variations in the group delay of the transmitted pulses. This active noise cancellation reduces the frequency instability for metrology applications and the broad-bandwidth timing jitter for remote synchronization applications. Though active stabilization has been implemented by others to reduce the frequency instability of transfer over optical fibers for time scales >1 s [13, 61], to my knowledge the only previously reported control of the timing jitter was limited by a feedback-loop time constant of $\sim$10 s and a measurement resolution of 400 fs, for a bandwidth extending from DC to only 100 Hz [9]. The active noise cancellation reported here provides a residual timing jitter after transfer of $\sim$20 fs over a bandwidth from 1 Hz to 100 kHz for transmission over several kilometers. Noise occurring at frequencies above 100 kHz is irrelevant, as the typical tracking bandwidth of a local oscillator at the receiving end of the frequency-distribution network is < 100 kHz.

To demonstrate the performance of the noise-cancellation servo, the feedback is applied to reduce the jitter of the signal after a round-trip path through the transmission fiber. In an actual implementation for transferring the microwave signal to a remote user, the detected noise introduced after a round-trip transmission would be used to cancel the perturbations for one-way transfer. This would be accomplished by applying half the correction signal to the adjustable delay line that would be necessary to eliminate the round-trip jitter, as has been done previously to cancel the phase noise introduced during the transfer of a cw optical frequency over fiber [54]. For this to be successful, the noise processes must be stationary over the round-trip time of the laser pulses. This fundamentally limits the bandwidth, ν, of the noise that can be effectively cancelled to

$$\nu = \frac{\phi}{2\pi}\frac{v_g}{2d} \tag{6.1}$$

where ϕ is the nominal maximum phase deviation allowed for the noise process to be considered stationary, v_g is the group velocity of the pulses in the fiber, and d is the geometric one-way length of the fiber. For $\phi = 1$ radian and a one-way fiber length of 5 km, this yields a bandwidth of $\nu \approx 3$ kHz.

To address the first significant noise source, amplitude- to phase-noise conversion by the receiving photodetector, two different strategies are employed to control the dispersion of the transmission fiber. The first is to measure the group-delay dispersion (GDD) of the BRAN fiber round trip and to introduce an appropriate amount of dispersion-compensation fiber (DCF) having a GDD of equal magnitude and opposite sign prior to the BRAN, producing a transmission path with a net zero GDD. Transmitting the fiber laser pulses through the DCF initially prechirps the pulses, temporally stretching them. Their subsequent passage through the BRAN fiber introduces a chirp of the opposite sign, re-compressing them and resulting in short pulses after transmission. The dispersion of the installed BRAN fiber is determined by measuring the dependence of the round-trip travel time on the wavelength of the transmitted pulses. The fiber laser pulses are spectrally filtered, providing 1-nm bandwidth pulses. As the center of the spectral window is adjusted, the variations in the round-trip travel time are measured by a fast photodetector (~40 ps timing resolution) with respect to a time reference provided by a second photodetector and a second portion of the fiber laser output. The filled circles in Fig. 6.7 show the dependence of the BRAN round-trip travel time on the transmitted wavelength. A linear fit (solid curve) provides the GDD of the 6.9-km BRAN fiber round trip — 124 ps / nm. This can be expressed in more

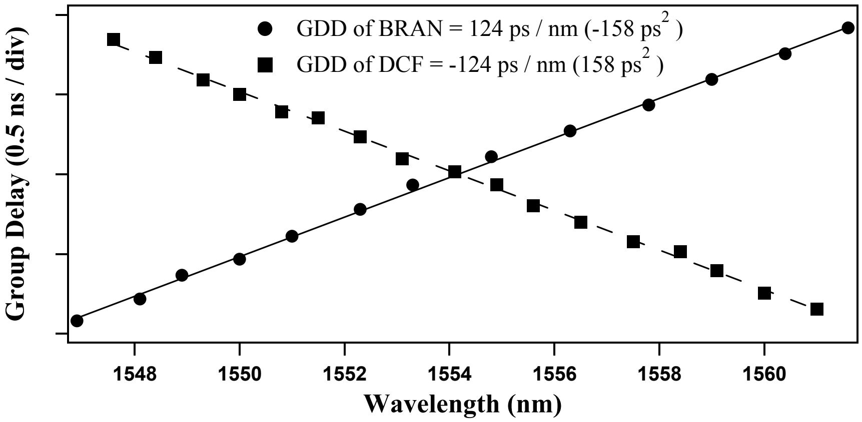

Figure 6.7: The dependence of the group delay on pulse wavelength for transmission of pulses through the BRAN fiber round trip and the DCF provides the GDD of these fibers. The linear fits to the data indicate the GDD of the BRAN fiber and DCF are equal in magnitude but opposite in sign, as required for dispersion control of the transmission fiber.

familiar units by considering that

$$\frac{\mathrm{d}T}{\mathrm{d}\omega} = \frac{\mathrm{d}T}{\mathrm{d}\lambda}\frac{\mathrm{d}\lambda}{\mathrm{d}\omega} \tag{6.2}$$

where T is the group delay, and ω and λ are the optical frequency and wavelength of the transmitted pulses, respectively. Since $\lambda = 2\pi c/\omega$, where c is the speed of light in vacuum, then

$$\frac{\mathrm{d}T}{\mathrm{d}\omega} = -\frac{\mathrm{d}T}{\mathrm{d}\lambda}\frac{2\pi c}{\omega^2} = -\frac{\mathrm{d}T}{\mathrm{d}\lambda}\frac{\lambda^2}{2\pi c} \tag{6.3}$$

With $\lambda = 1550$ nm, this produces a GDD for the BRAN fiber round trip of -158 ps^2. With this value, it is possible to measure the GDD of a known length of DCF and cut the appropriate length to compensate the GDD of the BRAN fiber. The squares in Fig. 6.7 show the measured dispersion of the 1.2-km-long DCF necessary for compensating the BRAN dispersion. The linear fit (dashed curve) indicates that the GDD of the DCF is -124 ps / nm, or 158 ps^2, as required. The use of the DCF prior to the BRAN yields a transmission path with a net GDD of magnitude $<$ 3 ps / nm, or 4 ps^2, and produces transmitted pulses that are ~60 ps. The pulse length is most likely limited by uncompensated higher-order dispersion. These pulses are short enough to yield a high

SNR for the received microwave signal with an average optical power sufficiently low (∼30 μW) to eliminate noise caused by the photodetection.

The second strategy for controlling the dispersion of the transmission fiber is to replace the dispersive BRAN fiber with dispersion-shifted fiber (DSF). DSF exhibits negligible dispersion at 1550 nm since its 0-GDD point has been shifted to this wavelength. The stability provided by microwave-frequency transfer over DSF is studied using a laboratory-based network comprised of DSF with a 4.5-km round trip. The net GDD of the DSF round trip has a magnitude < 18 ps / nm, or 23 ps^2. The DSF provides transmitted pulses shorter than 40 ps.

For actively cancelling the noise introduced by environmental perturbations during transfer, the transfer-induced noise is detected in the same manner described in Section 6.1 for studying the instability of passive transfer. The fiber laser output is again split into two portions, one of which is used for microwave-frequency transfer, while the other provides a reference against which to measure the noise introduced during transfer. This is shown in Fig. 6.8. The transmitted portion passes first through an adjustable delay line before entering the distribution fiber. The adjustable delay line, which will be described in detail later, provides the means to compensate for path-length fluctuations of the distribution fiber and actively cancel the transfer-induced noise. The distribution fiber consists either of the BRAN round trip preceded by the DCF, or the DSF round trip, as previously discussed. After the distribution fiber, an EDFA is used for compensation of losses in the fiber link. The reference and transmitted pulse trains are each detected on two photodetectors, with the incident average optical powers adjusted with variable optical attenuators. The microwave signals from the upper pair of detectors in Fig. 6.8 are filtered to select the 81st harmonic (7.84 GHz) of the fiber laser repetition frequency and are mixed with the appropriate relative phase to provide the timing jitter introduced during transfer. The measured jitter is applied to the adjustable delay line with appropriate gain and filtering to minimize the transfer-induced

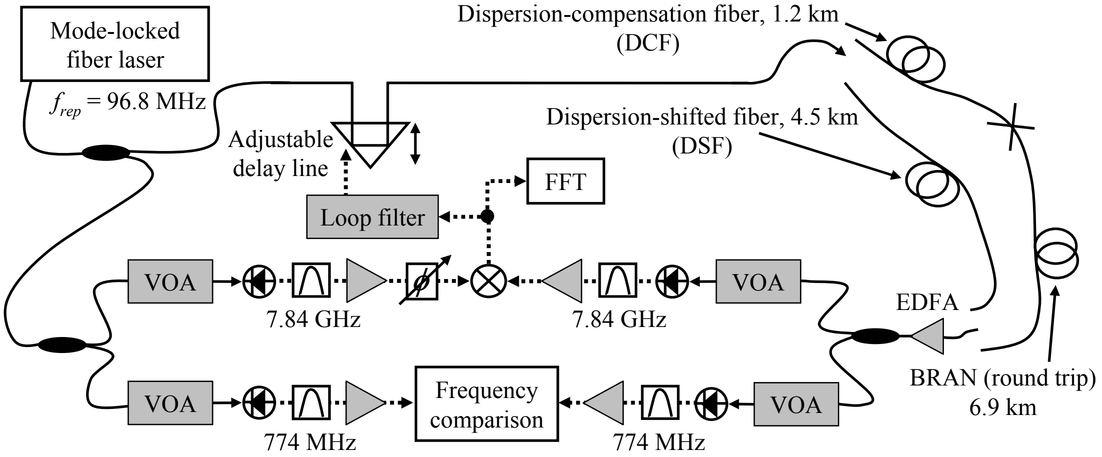

Figure 6.8: To cancel the timing jitter introduced during transfer of the comb repetition frequency, the jitter is determined by mixing the 81st harmonic of the repetition frequency obtained from the transmitted pulses and a reference portion of the fiber laser output. It is then applied to an adjustable delay line after appropriate filtering and amplification. The transmission fiber is either the dispersion-compensated BRAN fiber, or the DSF. An FFT provides spectral analysis of the in-loop timing jitter. The out-of-loop frequency instability is determined from the eighth harmonic of the repetition frequency obtained from the lower pair of photodetectors as in path (1) of Fig. 6.2. VOA, variable optical attenuator.

noise. An FFT provides spectral analysis of the residual in-loop jitter. The second pair of microwave signals is filtered to select the eighth harmonic (774 MHz) of the repetition frequency. The frequencies of these two signals are compared for averaging times $\geq$1 s by shifting the reference signal by 10 kHz with a single-side-band generator [40, 41] and counting the 10-kHz mixing product of the shifted reference and unshifted transmitted signals, as described for the analysis of the passive transfer stability and shown in path (1) of Fig. 6.2. This provides an out-of-loop measurement of the residual frequency instability present with active cancellation of the transfer noise.

The adjustable delay line, shown in Fig. 6.9, consists of two elements in series. A free-space path length terminated with a retroreflector actuated by an electrodynamic shaker is followed by a PZT-actuated fiber stretcher. The shaker operates by the same principle as a speaker and provides a large range of motion (a few millimeters). The fiber stretcher consists of optical fiber tightly wound 40 times around an oval mount with two PZTs that vary the smaller diameter of the oval and induce stretching of the fiber. It extends the servo bandwidth to ~100 Hz. The gain and bandwidth of the servo are limited by a resonance in the fiber stretcher at ~5 kHz. The performance of

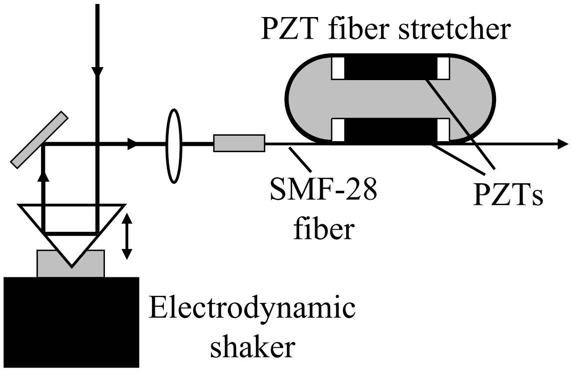

Figure 6.9: The adjustable delay line used for cancelling the transfer noise includes a free-space path length terminated with a retroreflector actuated by an electrodynamic shaker that provides a few millimeters of motion. This is followed by a PZT-actuated fiber stretcher, which extends the servo bandwidth to ~100 Hz.

the stabilization servo is not yet affected by the fundamental limitation imposed by the round-trip transit time through the transmission fiber.

Figure 6.10 shows the fractional frequency instability introduced during transfer through the dispersion-compensated BRAN fiber (open triangles) and the DSF (open circles) for averaging times ≥ 1 s, without active noise cancellation. These data are obtained from the frequency comparison of the 774-MHz signals from the pair of out-of-loop photodetectors in Fig. 6.8, using a 1-s gate time for the frequency counter. Data taken with 3-s and 10-s gate times are used to verify that the Allan deviation computed from the 1-s data is not affected by dead time. With the activation of the noise cancellation, the instability for transfer through the BRAN is reduced from $> 3 \times 10^{-14}$ to $< 9 \times 10^{-15}$ (filled triangles) for a 1-s averaging time and reaches 1×10^{-15} at 100 s, just slightly above the noise floor of the measurement system (open squares). The noise cancellation reduces the instability for transfer through the DSF from $> 6 \times 10^{-14}$ to $< 7 \times 10^{-15}$ (filled circles) for a 1-s averaging time, limited by the measurement noise floor. Transfer through the DSF reaches an instability of a few parts in 10^{16}

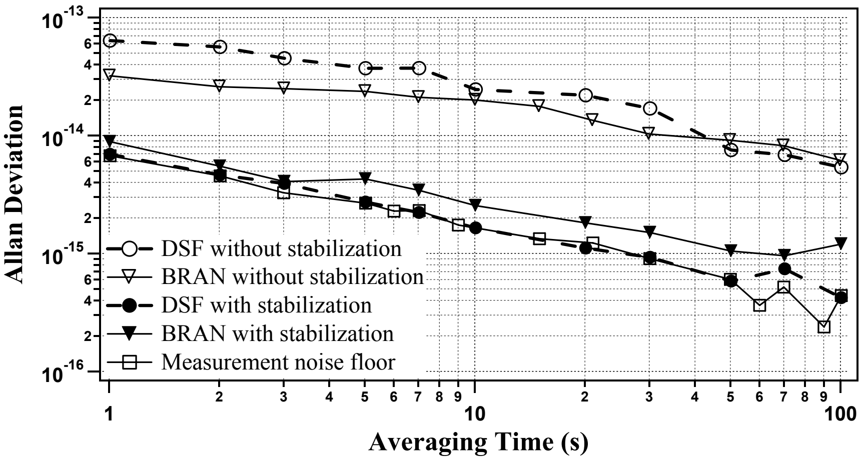

Figure 6.10: The Allan deviation for microwave-frequency transfer through the DSF and dispersion-compensated BRAN fiber indicates that the active noise cancellation reduces the transfer instability for transfer through the DSF to the measurement noise floor, and nearly to the noise floor for transfer through the dispersion-compensated BRAN fiber. The noise floor is determined using a short piece of fiber.

at 100 s with the noise cancellation. The measurement noise floor is determined by replacing the transmission fiber with a short length of fiber. These results, obtained with independent detectors outside the stabilization loop, confirm that the photodetectors do not contribute any significant instability above the level of the noise floor. The stability achieved with active noise cancellation is very similar to that recently reported for the stabilized transfer of microwave signals over optical fibers between LPL and SYRTE in Paris [13]. For that work microwave signals are transmitted over 43-km optical fibers by modulating the injection current of diode lasers at the microwave frequency to be transferred. Adjustment of the modulation signal phase, a PZT-actuated fiber stretcher, and a temperature-controlled spool of fiber are used to cancel the transmission noise.

For frequency metrology applications, it is sufficient to reduce the long-term ($\geq$1 s) transfer instability as described above. However, for time-domain experiments that are sensitive to noise occurring on faster time scales and for those that involve the remote synchronization of various system components, such as those mentioned in Section 1.6, it is necessary to consider the timing jitter introduced over a broad bandwidth. Figure 6.11(a) shows the spectral density of the rms jitter introduced during transfer of the fiber laser repetition frequency over the dispersion-compensated BRAN fiber with (open triangles) and without (open circles) active stabilization of the transmission path. The corresponding single side band phase noise, L(f), of the transmitted 81st harmonic of the fiber laser repetition frequency is shown on the right axis. The active noise cancellation reduces the residual jitter up to a Fourier frequency of ~100 Hz, and the jitter rolls off at ~8 kHz with a slope of ~ -10 dB / decade until reaching the measurement noise floor (solid trace) near 20 kHz. Figure 6.11(b) shows the corresponding total rms timing jitter over a bandwidth from 1 Hz to f_h versus Fourier frequency f_h. The total timing jitter over the bandwidth from 1 Hz to 100 kHz is only ~49 fs without stabilization, compared with a total jitter over the same bandwidth of ~240 fs [Fig. 6.6(b)] obtained under similar operating conditions, but without compensation of the BRAN

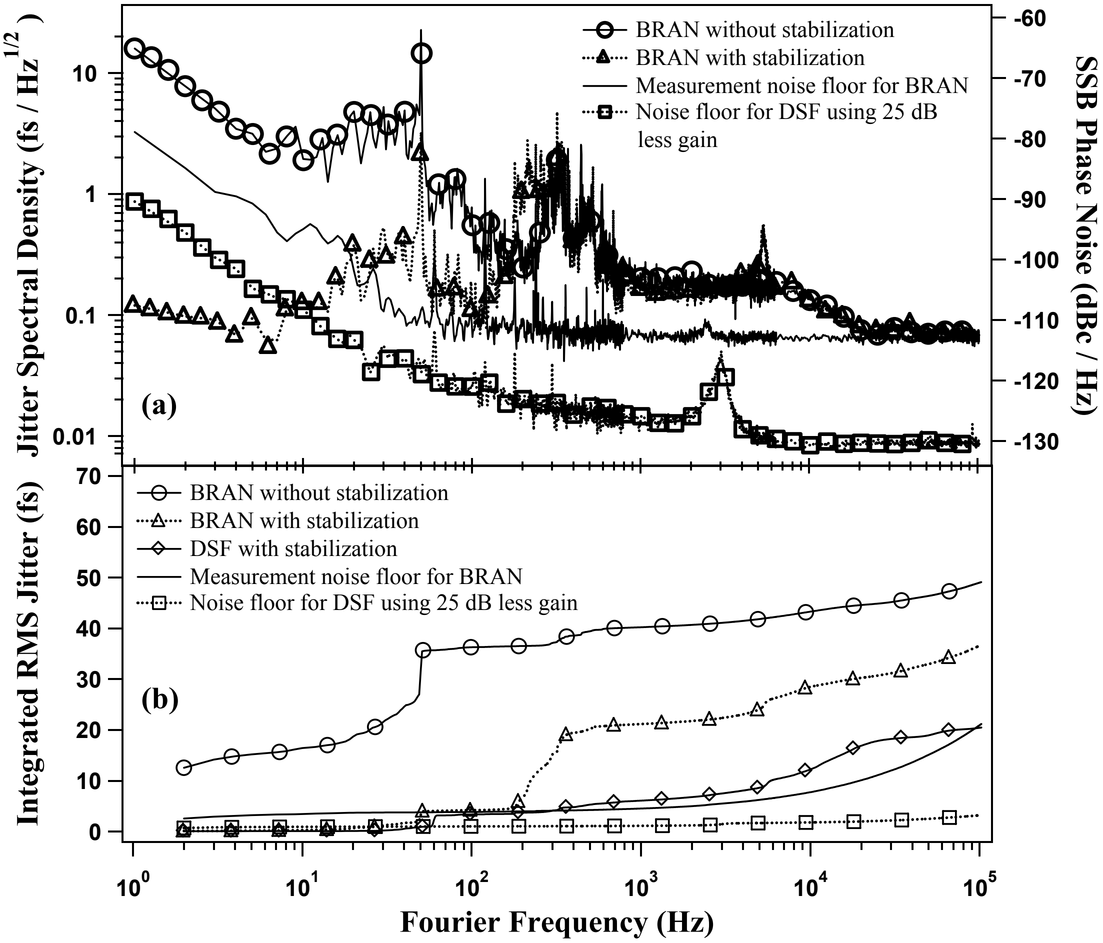

Figure 6.11: (a) The rms timing jitter spectral density (left axis) for transfer over the dispersion-compensated BRAN fiber and the equivalent single side band (SSB) phase noise of the transmitted 81st harmonic of the comb repetition frequency (right axis) with and without active noise cancellation. The noise cancellation reduces the jitter for Fourier frequencies below ∼100 Hz. (b) The integrated rms jitter over a bandwidth from 1 Hz to f_h versus f_h for transfer over the BRAN fiber and the actively stabilized DSF. The measurement noise floors are obtained by replacing the corresponding transmission fiber with a short piece of fiber and are limited by the thermal noise input to the microwave amplifier that amplifies the photodetector signal.

dispersion. This comparison provides a clear demonstration of the impact made on the transfer stability by implementing dispersion control. Active cancellation of the transfer noise further reduces the jitter over this bandwidth to ∼37 fs. The jitter for transfer over the DSF is reduced to ∼20 fs with the active noise cancellation, shown as open diamonds.

The measurement noise floor for the BRAN-transfer jitter is determined by replacing the BRAN fiber and the DCF with a short piece of fiber. The light incident on the photodetector receiving the transmitted pulses is attenuated to produce the same photocurrent and microwave signal power that are obtained for transmission over the dispersion-compensated BRAN. It is useful to calculate the contributions to the noise floor from the shot noise of the photocurrent and the thermal noise input to the microwave amplifier that amplifies the photodetector signal, to determine if either of these fundamental noise processes is limiting the performance of the measurement system. The shot noise of the photocurrent produces a white current noise spectral density (in $\mathrm{A}/\sqrt{\mathrm{Hz}}$) of $i_{sn} = \sqrt{2ei_{avg}}$, where e is the charge of the electron and i_{avg} indicates the average DC photocurrent [31]. This current noise produces RF noise on the signal from the photodetector that has a power spectral density (PSD) given by

$$\mathrm{S_{RF}}(f) = \frac{2i_{sn}^2 R}{P_{signal}} \tag{6.4}$$

where R is the load impedance for the photodetector and P_{signal} is the microwave power of the desired signal from the photodetector. The factor of 2 comes from the fact that $\mathrm{S_{RF}}(f)$ includes the contributions to the RF noise from the noise sidebands on both sides of the signal frequency [31]. Since the random RF noise is distributed equally between amplitude-modulation noise and phase-modulation noise [26], the PSD of phase fluctuations due to the shot noise is given by

$$\begin{aligned} \mathrm{S}_\phi(f) &= \frac{1}{2}\mathrm{S_{RF}}(f) \\ &= \frac{2ei_{avg}R}{P_{signal}} \end{aligned} \tag{6.5}$$

During transfer over the dispersion-compensated BRAN fiber the average photocurrent is $\sim$ 10 μA, and a microwave signal power of $\sim$ 1 nW for the transmitted 81st harmonic of the fiber laser repetition frequency is obtained with a 50 Ω load impedance. Using Eqn. (2.22) to obtain the single side band phase noise, $\mathrm{L}(f)$, from $\mathrm{S}_\phi(f)$, the contribution from the shot noise to the single side band phase noise of the measurement system for transfer over the BRAN fiber is found to be $\sim$ -131 dBc/Hz. This is well below the measured noise floor of $\sim$ 113 dBc/Hz shown in Fig. 6.11(a), indicating the measurement system is not limited by the photocurrent shot noise.

The other source of noise that may be fundamentally limiting the performance of the measurement system is the thermal noise input to the microwave amplifier that amplifies the signal from the receiving photodetector. This thermal noise produces a white voltage noise spectral density (in $\mathrm{V}/\sqrt{\mathrm{Hz}}$) on the input to the microwave amplifier that is given by $v_{tn} = \sqrt{4kTR}$, where k is Boltzmann's constant, T is the temperature in degrees Kelvin, and R is the input impedance of the amplifier [26]. The PSD of RF noise caused by this voltage noise on the photodetector signal that is input to the amplifier is given by an expression similar to the one for the shot noise contribution [Eqn. (6.4)] [26]. For the thermal noise contribution,

$$\mathrm{S}_{\mathrm{RF}}(f) = \frac{2v_{tn}^2/R}{P_{signal}} \tag{6.6}$$

where P_{signal} again denotes the power of the signal from the photodetector. The PSD of phase fluctuations caused by the thermal noise is then given by

$$\begin{aligned} \mathrm{S}_\phi(f) &= \frac{v_{tn}^2/R}{P_{signal}} \\ &= \frac{4kT}{P_{signal}} \end{aligned} \tag{6.7}$$

At room temperature ($T \approx 300$ K), the $\sim$1 nW of microwave power provided by the photodetector after transfer of the 81st harmonic of the fiber laser repetition frequency over the dispersion-compensated BRAN fiber results in a thermal contribution to the single side band phase noise of the measurement system of $\sim$ -111 dBc/Hz. This agrees very well with the measured noise floor shown by the solid trace in Fig. 6.11(a), and it is evident that the noise floor of the measurement system for transfer over the BRAN is limited by this fundamental thermal noise. From Eqn. (6.7) it is clear that without cooling the system, the only way to reduce the noise floor is to increase the signal power obtained from the photodetector. However, because of the optical loss through the BRAN and DCF, this is not possible at this time. The total jitter of the measurement system over the bandwidth from 1 Hz to 100 kHz is $\sim$21 fs, shown in Fig. 6.11(b).

Although the optical loss through the BRAN and DCF prevents the power of the received microwave signal from being increased, the relatively lower loss through the DSF provides a signal from the receiving photodetector that has 25 dB higher power for the 81st harmonic of the fiber laser repetition frequency. Therefore, for measuring the jitter introduced during transfer over the DSF, a measurement system involving 25 dB less amplification is used. Replacing the DSF with a short piece of fiber and attenuating the light incident on the receiving photodetector to achieve the same photocurrent and microwave signal power that are obtained with the DSF allows the noise floor of this measurement system to be determined. The noise floor is shown by the trace denoted with open squares in Fig. 6.11(a) and is significantly less than the noise floor for the dispersion-compensated BRAN fiber. Transfer over the DSF provides an average photocurrent in the receiving photodetector of $\sim$ 160 μA and a microwave signal power of $\sim$ 250 nW with a 50 Ω load impedance. Using Eqn. (6.5), the shot noise contribution to the single side band phase noise of the measurement system for transfer over the DSF is found to be $\sim$ -143 dBc/Hz. Since this is again below the measured noise floor of $\sim$ -130 dBc/Hz, the measurement system is still not limited by the shot noise. From Eqn. (6.7), the thermal noise contribution to the single side band phase noise of this measurement system is calculated to be $\sim$ -135 dBc/Hz, which is in very close agreement with the measured noise floor for Fourier frequencies above $\sim$10 kHz.

Although the measurement system is not limited by fundamental noise processes for Fourier frequencies below ~10 kHz, above 10 kHz it is still limited by the thermal noise. The total jitter of this measurement system over the bandwidth from 1 Hz to 100 kHz is ~4 fs, shown in Fig. 6.11(b). Since the total jitter (measured in-loop) over the full measurement bandwidth for transfer over the BRAN fiber and the DSF is not suppressed below the jitter of the corresponding measurement system by the active noise cancellation, the in-loop results are identical to what would be obtained with a pair of out-of-loop detectors.

The active noise cancellation can be improved by increasing the gain and bandwidth of the adjustable delay, which is demonstrated using a different type of PZT-actuated fiber stretcher. Approximately 40 m of optical fiber wound around a cylindrical PZT provide enough dynamic range to stabilize the path-length fluctuations of the dispersion-compensated BRAN fiber without using the free-space retroreflector actuated by the electrodynamic shaker. Also, the first resonance of this fiber stretcher occurs at ~20 kHz, as compared to ~5 kHz for the fiber stretcher used in obtaining the previous data. Figure 6.12(a) shows that for microwave-frequency transfer over the dispersion-compensated BRAN fiber, this fiber stretcher suppresses the jitter to the measurement noise floor over the entire bandwidth from 1 Hz to 100 kHz, except for the Fourier frequency range from ~1 kHz to ~10 kHz. This jitter spectral density is obtained from in-loop data, and the original data falling below the measurement noise floor is replaced with the higher spectral density of the noise floor. The total integrated jitter over the bandwidth from 1 Hz to 100 kHz is reduced to ~16 fs, nearly limited by the measurement noise floor.

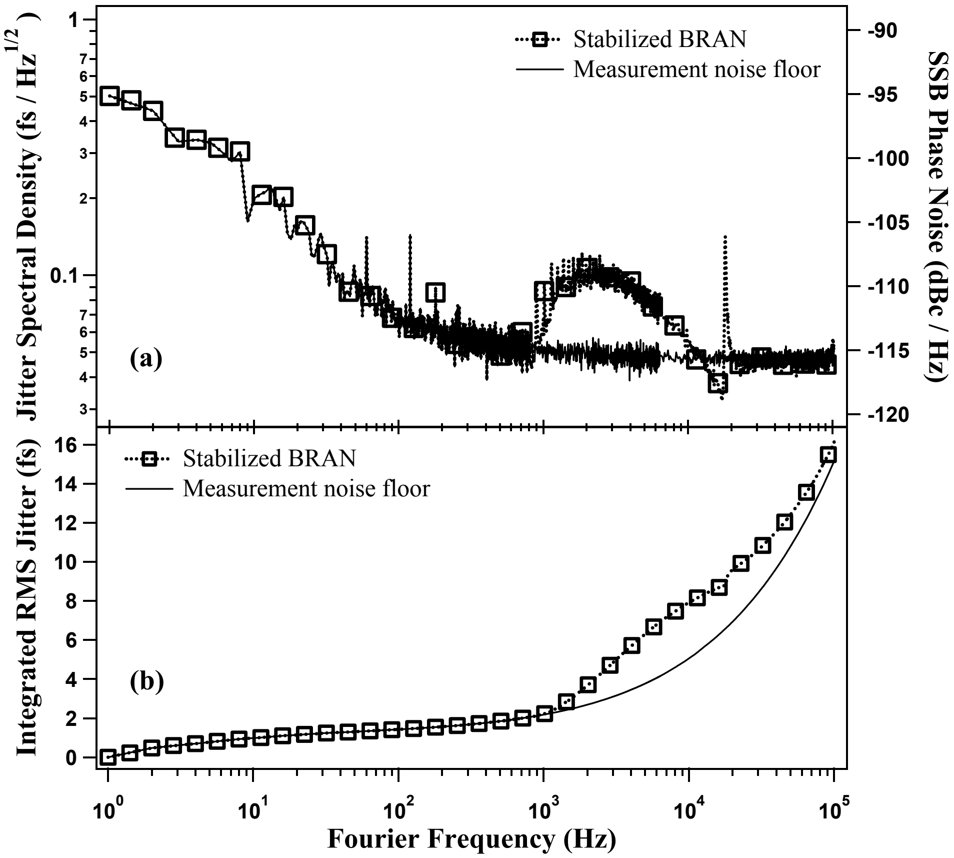

Figure 6.12: (a) The rms timing jitter spectral density (left axis) for transfer over the dispersion-compensated BRAN fiber and the equivalent single side band (SSB) phase noise of the transmitted 81st harmonic of the comb repetition frequency (right axis) with the transfer noise actively cancelled by a cylindrical PZT-actuated fiber stretcher. (b) The integrated rms jitter over a bandwidth from 1 Hz to f_h versus f_h for transfer over the BRAN fiber. The measurement noise floor is obtained by replacing the transmission fiber with a short piece of fiber.

6.3 Summary and future outlook for frequency transfer with combs

In Section 1.6, several applications of the excellent stability provided by optical atomic clocks were described that will rely on the ability to transfer over several kilometers a stable frequency reference that is linked to an optical standard. Studies of fundamental physics will benefit from this ability, including the search for time variation in the fine structure constant, which requires precise comparisons of different optical frequency standards. Because of the lack of portability of the most stable optical clock systems, the transfer of a reference linked to the optical standard is necessary. The distribution of a stable reference for tight timing synchronization among the components of a system is useful for radio astronomy and pump-probe experiments using ultrashort x-ray pulses produced by a linear accelerator. High-speed and secure communications will also benefit from the transfer of a stable frequency reference.

In this chapter, we have seen that an optical frequency comb provides the flexibility of simultaneously transmitting over several kilometers of fiber an optical and a microwave frequency reference linked to an optical standard, replacing the two separate methods currently used for accomplishing this. Furthermore, the instability introduced during transfer of a microwave reference using the comb is an order of magnitude lower than that measured for transfer with a modulated cw laser. Dispersion control of the transmission fiber and active cancellation of the transfer noise reduce the frequency instability for microwave transfer with the comb to the measurement noise floor for averaging times ≥ 1 s. The timing jitter over the bandwidth from 1 Hz to 100 kHz is reduced to approximately 20 fs for the dispersion-compensated BRAN fiber and the laboratory-based DSF network. To my knowledge, this represents the lowest amount of jitter reported for the transfer of a timing signal over several kilometers using optical-fiber networks.

Of course, for the microwave and optical frequencies transmitted by the comb to be linked to an optical standard, the comb itself must be stabilized to the optical standard. We have seen that the 1550-nm comb (at the appropriate wavelength for fiber transmission) from a MLLD can be stabilized to the comb of a mode-locked Ti:sapphire laser with a residual relative timing jitter between the two lasers of $\sim$22 fs over the bandwidth from 1 Hz to 100 MHz, which is the lowest reported jitter for a MLLD. In addition to controlling the timing of the MLLD pulses, the stabilization of its comb also requires control of its optical phase, which is equivalent to stabilizing its carrier-envelope offset frequency, f_{ceo}. The implementation of an orthogonal control scheme for simultaneously controlling both the MLLD repetition frequency and f_{ceo} enables the complete stabilization of its comb.

With the MLLD stabilized to the Ti:sapphire laser, the stabilization of the MLLD to the optical standard is completed by stabilizing the Ti:sapphire comb to the optical standard. The carrier-envelope offset frequency of the Ti:sapphire laser is controlled using the laser intensity, which couples to the comb parameters through an intensity-dependent shift of the laser spectrum and the linear contribution to the GDD of the laser cavity. Furthermore, it has been demonstrated that the dominant source of phase noise of f_{ceo} is fast fluctuations of the peak pulse intensity, $\hat{I}$, which has important implications for controlling f_{ceo} using the laser intensity. The stabilization of f_{ceo} is optimized when the laser is operated sufficiently close to the position of $\mathrm{d}f_{ceo}/\mathrm{d}\hat{I} = 0$ to minimize the dependence of f_{ceo} on the fast intensity fluctuations outside the bandwidth of the stabilization loop, while being far enough away to use the lower-bandwidth response of f_{ceo} for stabilization. The Ti:sapphire comb has been stabilized to a portable optical standard based on a transition in I_2 contained in a vapor cell and to a transition in ultracold, laser-cooled ^{88}Sr atoms. The ^{88}Sr transition exhibits an order of magnitude lower uncertainty than the I_2 transition, and future work with ultracold Sr atoms should reduce the uncertainty even further. However, the complexity of the Sr system precludes it from being portable, which illustrates why stable frequency transfer is necessary for using the superior performance provided by this and other similar ultracold atomic systems.

The active cancellation of the noise introduced during transfer of a microwave reference using an optical frequency comb could be further improved by using an actuator that is faster than the PZT fiber stretcher to compensate for fluctuations in the transmission path length. One possibility is to use an electro-optic modulator (EOM). Applying a voltage across an EOM modifies its refractive index, changing its optical path length. Also, placing the EOM within the cavity of the mode-locked laser that provides the comb to be transmitted would prevent the stabilization loop from ever exceeding its possible range of transmission-delay compensation. An actuator outside the laser cavity must change by a static amount to compensate for a change in the transmission path length, and for a monotonic drift in this path length the actuator will inevitably run out of range. However, with the actuator inside the laser cavity, the cavity would effectively integrate the variations of the actuator. For example, a drift in the transmission path length at a constant rate would be compensated by a fixed deviation of the actuator, resulting in a constant offset of the laser repetition frequency. For a sufficiently slow drift of the path length, the actuator would never run out of range. Of course, this compensation scheme implies that the transmitted comb would no longer be rigidly locked to the optical standard, since its repetition frequency must change to compensate for fluctuations in the path length. When a Ti:sapphire laser is first stabilized to the optical standard, and then the stability of the Ti:sapphire comb is transferred to a 1550-nm comb, detection of the noise introduced during transmission of the 1550-nm comb with respect to the Ti:sapphire comb would enable the derivation of a feedback signal to apply to the 1550-nm comb source for cancelling the transmission noise. Control of the repetition frequency of a mode-locked fiber laser with an intracavity EOM [38] would also enable broad-bandwidth stabilization of this laser to a stable transmitted frequency reference for the remote synchronization of two fiber lasers.

Once the feedback for stabilizing the transfer of a microwave reference is optimized, it will be necessary to improve the noise floor of the measurement system for detecting the timing jitter and frequency instability introduced during transfer. The noise floor is currently limited by the fundamental thermal noise on the photodetector signal input to the microwave amplifiers. Therefore, without cooling the system, the noise floor can only be reduced by increasing the signal power produced by the photodetector receiving the transmitted pulses. From Eqn. (6.7), it is seen that the contribution from the thermal noise decreases with increasing microwave signal power, P_{signal}, as P_{signal}^{-1}. On the other hand, the contribution from the photocurrent shot noise decreases as $P_{signal}^{-1/2}$, as shown by Eqn. (6.5) since the average photocurrent, i_{avg}, is proportional to $\sqrt{P_{signal}}$. Therefore, as the signal power is increased, the thermal noise contribution will decrease faster than the shot noise contribution, and for signal powers above some threshold the shot noise contribution will be dominant. It is desirable to operate in this regime, where the decrease in the phase noise for a given increase in signal power is not as large as when thermal noise is dominant.

However, the signal from the receiving photodetector is currently limited by optical loss in the transmission fiber, as well as the need to minimize the optical power incident on the photodetector because of the conversion of amplitude fluctuations of the pulse train to phase noise by the photodetector. The first limitation can be removed with more optical amplification in the transmission path, and it also suggests the use of DSF for transmission since it exhibits lower loss than the dispersion-compensated BRAN fiber. The amplitude fluctuations of the pulse train can be eliminated with active stabilization of the pulse amplitude. To further reduce the noise floor, the noise introduced on the microwave signal can be completely avoided by using an optical cross-correlation to detect the noise introduced during transfer, as described at the end of

Section 2.5. This will also require stabilization of the pulse amplitude. Since the sensitivity provided by a cross-correlation relies on the pulses being as short as possible, careful dispersion compensation will need to be implemented to cancel the remaining dispersion that the pulses experience upon passing through the DSF or the BRAN compensated by the DCF. This can be accomplished using a programmable pulse shaper, as described in [43]. Optimized dispersion control and the cross-correlation of a pulse train from one laser with a pulse train transmitted over fiber will also be useful for remote synchronization of two lasers.

The elimination of these technical limitations on the stability of frequency transfer leaves only the fundamental limitation imposed by the finite transit time of the laser pulses through the fiber. As explained in Section 6.2, the noise introduced during transfer is determined after the signal completes a round trip, and the successful cancellation of the noise for one-way transfer requires that it be stationary over the round-trip time of the pulses. This implies that for a one-way transmission distance of 5 km, the noise up to a bandwidth of $\sim$3 kHz can be successfully cancelled. The bandwidth over which noise exists that must be cancelled determines the maximum distance a frequency reference can be transmitted, with this distance scaling inversely with the bandwidth. However, it is possible to distribute a frequency reference over larger distances using multiple stages, each consisting of a mode-locked laser and transmission fiber with a length given by the bandwidth of the noise that must be cancelled. The transfer of a frequency reference through the fiber of each stage using the mode-locked laser comb would be stabilized by detecting the noise introduced after a round trip through the stage. The individual stages would then be chained together, with the laser for each stage stabilized to the frequency reference transmitted through the previous stage. Ultimately, as the distribution of frequency references is implemented over longer and longer distances with ever increasing stability, it will be important to consider other factors such as the effect gravity has on the transmitted reference. For example, a frequency shift of a part in 10^{18} accompanies a change in elevation of only 1 cm.

Bibliography

[1] "BRAN - Boulder Research and Administration Network," http://www.branfiber.net .

[2] "Linac Coherent Light Source - LCLS," http://www-ssrl.slac.stanford.edu/lcls .

[3] "Timekeeping history," http://www.beaglesoft.com/maintimehistory.htm .

[4] W. J. H. Andrewes, "A chronicle of timekeeping," Sci. Am. **287**, 76–85 (2002).

[5] M. T. Asaki, C.-P. Huang, D. Garvey, J. Zhou, H. C. Kapteyn, and M. M. Murnane, "Generation of 11-fs pulses from a self-mode-locked Ti:sapphire laser," Opt. Lett. **18**, 977–979 (1993).

[6] A. Bartels and H. Kurz, "Generation of a broadband continuum by a Ti:sapphire femtosecond oscillator with a 1-GHz repetition rate," Opt. Lett. **27**, 1839–1841 (2002).

[7] S. Bize, S. A. Diddams, U. Tanaka, C. E. Tanner, W. H. Oskay, R. E. Drullinger, T. E. Parker, T. P. Heavner, S. R. Jefferts, L. Hollberg, W. M. Itano, and J. C. Bergquist, "Testing the stability of fundamental constants with the $^{199}\mathrm{Hg}^+$ single-ion optical clock," Phys. Rev. Lett. **90**, 150 802 (2003).

[8] R. W. Boyd, Nonlinear Optics (Academic Press, San Diego, 2003).

[9] T. P. Celano, S. R. Stein, G. A. Gifford, B. A. Mesander, and B. J. Ramsey, "Sub-picosecond active timing control over fiber optic cable," in Proceedings of the 2002 IEEE International Frequency Control Symposium and PDA Exhibition (Institute of Electrical and Electronics Engineers, Piscataway, New Jersey, 2002), pp. 510–516.

[10] W.-Y. Cheng, L. Chen, T. H. Yoon, J. L. Hall, and J. Ye, "Sub-Doppler molecular-iodine transitions near the dissociation limit (523-498 nm)," Opt. Lett. **27**, 571–573 (2002).

[11] S. T. Cundiff, "Phase stabilization of ultrashort optical pulses," J. Phys. D: Appl. Phys. **35**, R43–R59 (2002).

[12] T. Damour and J. H. Taylor, "Strong-field tests of relativistic gravity and binary pulsars," Phys. Rev. D **45**, 1840–1868 (1992).

[13] C. Daussy, O. Lopez, A. Amy-Klein, A. Goncharov, M. Guinet, C. Chardonnet, F. Narbonneau, M. Lours, D. Chambon, S. Bize, A. Clairon, G. Santarelli, M. E. Tobar, and A. N. Luiten, "Long-distance frequency dissemination with a resolution of 10^{-17}," Phys. Rev. Lett. **94**, 203 904 (2005).

[14] W. Demtröder, Laser Spectroscopy: Basic Concepts and Instrumentation (Springer-Verlag, Berlin, 1996), 2nd edn.

[15] S. A. Diddams, J. C. Bergquist, S. R. Jefferts, and C. W. Oates, "Standards of time and frequency at the outset of the 21st century," Science **306**, 1318–1324 (2004).

[16] J.-C. Diels and W. Rudolph, Ultrashort Laser Pulse Phenomena: Fundamentals, Techniques, and Applications on a Femtosecond Time Scale, Optics and Photonics (Academic Press, San Diego, 1996).

[17] J. N. Eckstein, A. I. Ferguson, and T. W. Hänsch, "High-resolution two-photon spectroscopy with picosecond light-pulses," Phys. Rev. Lett. **40**, 847–850 (1978).

[18] R. Ell, U. Morgner, F. X. Kärtner, J. G. Fujimoto, E. P. Ippen, V. Scheuer, G. Angelow, T. Tschudi, M. J. Lederer, A. Boiko, and B. Luther-Davies, "Generation of 5-fs pulses and octave-spanning spectra directly from a Ti:sapphire laser," Opt. Lett. **26**, 373–375 (2001).

[19] G. Ferrari, P. Cancio, R. Drullinger, G. Giusfredi, N. Poli, M. Prevedelli, C. Toninelli, and G. M. Tino, "Precision frequency measurement of visible intercombination lines of strontium," Phys. Rev. Lett. **91**, 243 002 (2003).

[20] J. Flowers, "The route to atomic and quantum standards," Science **306**, 1324–1330 (2004).

[21] S. M. Foreman, A. Marian, J. Ye, E. A. Petrukhin, M. A. Gubin, O. D. Mücke, F. N. C. Wong, E. P. Ippen, and F. X. Kärtner, "Demonstration of a HeNe/CH_4-based optical molecular clock," Opt. Lett. **30**, 570–572 (2005).

[22] R. L. Fork, O. E. Martinez, and J. P. Gordon, "Negative dispersion using pairs of prisms," Opt. Lett. **9**, 150–152 (1984).

[23] T. M. Fortier, D. J. Jones, and S. T. Cundiff, "Phase stabilization of an octave-spanning Ti:sapphire laser," Opt. Lett. **28**, 2198–2200 (2003).

[24] J. P. Gordon and H. A. Haus, "Random-walk of coherently amplified solitons in optical fiber transmission," Opt. Lett. **11**, 665–667 (1986).

[25] J. L. Hall and T. W. Hänsch, "History of optical comb development," in Femtosecond Optical Frequency Comb Technology: Principle, Operation, and Applications, J. Ye and S. T. Cundiff, eds. (Springer, New York, 2005), pp. 1–11.

[26] A. Hati, D. A. Howe, F. L. Walls, and D. Walker, "Noise figure vs. PM noise measurements: A study at microwave frequencies," in Proceedings of the 2003 IEEE International Frequency Control Symposium and PDA Exhibition Jointly with the 17th European Frequency and Time Forum (Institute of Electrical and Electronics Engineers, Piscataway, New Jersey, 2003), pp. 516–520.

[27] H. A. Haus and E. P. Ippen, "Group velocity of solitons," Opt. Lett. **26**, 1654–1656 (2001).

[28] H. A. Haus, I. Sorokina, and E. Sorokin, "Raman-induced redshift of ultrashort mode-locked laser pulses," J. Opt. Soc. Am. B-Opt. Phys. **15**, 223–231 (1998).

[29] T. P. Heavner, S. R. Jefferts, E. A. Donley, J. H. Shirley, and T. E. Parker, "Recent improvements in NIST-F1 and a resulting accuracy of $\delta f/f = 0.61 \times 10^{-15}$," IEEE Trans. Instrum. Meas. **54**, 842–845 (2005).

[30] F. W. Helbing, G. Steinmeyer, U. Keller, R. S. Windeler, J. Stenger, and H. R. Telle, "Carrier-envelope offset dynamics of mode-locked lasers," Opt. Lett. **27**, 194–196 (2002).

[31] L. Hollberg, C. W. Oates, E. A. Curtis, E. N. Ivanov, S. A. Diddams, T. Udem, H. G. Robinson, J. C. Bergquist, R. J. Rafac, W. M. Itano, R. E. Drullinger, and D. J. Wineland, "Optical frequency standards and measurements," IEEE J. Quantum Electron. **37**, 1502–1513 (2001).

[32] K. W. Holman, D. D. Hudson, J. Ye, and D. J. Jones, "Remote transfer of a high-stability and ultralow-jitter timing signal," Opt. Lett. **30**, 1225–1227 (2005).

[33] K. W. Holman, D. J. Jones, D. D. Hudson, and J. Ye, "Precise frequency transfer through a fiber network by use of 1.5-μm mode-locked sources," Opt. Lett. **29**, 1554–1556 (2004).

[34] K. W. Holman, D. J. Jones, J. Ye, and E. P. Ippen, "Orthogonal control of the frequency comb dynamics of a mode-locked laser diode," Opt. Lett. **28**, 2405–2407 (2003).

[35] K. W. Holman, R. J. Jones, A. Marian, S. T. Cundiff, and J. Ye, "Detailed studies and control of intensity-related dynamics of femtosecond frequency combs from mode-locked Ti:sapphire lasers," IEEE J. Sel. Top. Quantum Electron. **9**, 1018–1024 (2003).

[36] K. W. Holman, R. J. Jones, A. Marian, S. T. Cundiff, and J. Ye, "Intensity-related dynamics of femtosecond frequency combs," Opt. Lett. **28**, 851–853 (2003).

[37] T. Hong, C. Cramer, W. Nagourney, and E. N. Fortson, "Optical clocks based on ultranarrow three-photon resonances in alkaline earth atoms," Phys. Rev. Lett. **94**, 050 801 (2005).

[38] D. D. Hudson, K. W. Holman, R. J. Jones, S. T. Cundiff, J. Ye, and D. J. Jones, "Mode-locked fiber laser frequency-controlled with an intracavity electro-optic modulator," Opt. Lett. (in press).

[39] T. Ido, T. H. Loftus, M. M. Boyd, A. D. Ludlow, K. W. Holman, and J. Ye, "Precision spectroscopy and density-dependent frequency shifts in ultracold Sr," Phys. Rev. Lett. **94**, 153 001 (2005).

[40] E. N. Ivanov, design provided by S. A. Diddams.

[41] E. N. Ivanov, S. A. Diddams, and L. Hollberg, “Analysis of noise mechanisms limiting the frequency stability of microwave signals generated with a femtosecond laser,” IEEE J. Sel. Top. Quantum Electron. **9**, 1059–1065 (2003).

[42] L. A. Jiang, M. E. Grein, E. P. Ippen, C. McNeilage, J. Searls, and H. Yokoyama, “Quantum-limited noise performance of a mode-locked laser diode,” Opt. Lett. **27**, 49–51 (2002).

[43] Z. Jiang, S.-D. Yang, D. E. Leaird, and A. M. Weiner, “Fully dispersion-compensated ∼500 fs pulse transmission over 50 km single-mode fiber,” Opt. Lett. **30**, 1449–1451 (2005).

[44] D. J. Jones, S. A. Diddams, J. K. Ranka, A. Stentz, R. S. Windeler, J. L. Hall, and S. T. Cundiff, “Carrier-envelope phase control of femtosecond mode-locked lasers and direct optical frequency synthesis,” Science **288**, 635–639 (2000).

[45] D. J. Jones, K. W. Holman, M. Notcutt, J. Ye, J. Chandalia, L. A. Jiang, E. P. Ippen, and H. Yokoyama, “Ultralow-jitter, 1550-nm mode-locked semiconductor laser synchronized to a visible optical frequency standard,” Opt. Lett. **28**, 813–815 (2003).

[46] R. J. Jones, W.-Y. Cheng, K. W. Holman, L. Chen, J. L. Hall, and J. Ye, “Absolute-frequency measurement of the iodine-based length standard at 514.67 nm,” Appl. Phys. B: Lasers Opt. **74**, 597–601 (2002).

[47] V. L. Kalashnikov, E. Sorokin, and I. T. Sorokina, “Mechanisms of spectral shift in ultrashort-pulse laser oscillators,” J. Opt. Soc. Am. B-Opt. Phys. **18**, 1732–1741 (2001).

[48] S. G. Karshenboim, “Some possibilities for laboratory searches for variations of fundamental constants,” Can. J. Phys. **78**, 639–678 (2000).

[49] H. Katori, T. Ido, Y. Isoya, and M. Kuwata-Gonokami, “Magneto-optical trapping and cooling of strontium atoms down to the photon recoil temperature,” Phys. Rev. Lett. **82**, 1116–1119 (1999).

[50] H. Katori, M. Takamoto, V. G. Pal’chikov, and V. D. Ovsiannikov, “Ultrastable optical clock with neutral atoms in an engineered light shift trap,” Phys. Rev. Lett. **91**, 173 005 (2003).

[51] J. Levine, “Introduction to time and frequency metrology,” Rev. Sci. Instrum. **70**, 2567–2596 (1999).

[52] T. H. Loftus, T. Ido, A. D. Ludlow, M. M. Boyd, and J. Ye, “Narrow line cooling: Finite photon recoil dynamics,” Phys. Rev. Lett. **93**, 073 003 (2004).

[53] L.-S. Ma, Z. Bi, A. Bartels, L. Robertsson, M. Zucco, R. S. Windeler, G. Wilpers, C. Oates, L. Hollberg, and S. A. Diddams, “Optical frequency synthesis and comparison with uncertainty at the 10^{-19} level,” Science **303**, 1843–1845 (2004).

[54] L.-S. Ma, P. Jungner, J. Ye, and J. L. Hall, "Delivering the same optical frequency at two places: Accurate cancellation of phase noise introduced by an optical fiber or other time-varying path," Opt. Lett. **19**, 1777–1779 (1994).

[55] H. Marion, F. P. Dos Santos, M. Abgrall, S. Zhang, Y. Sortais, S. Bize, I. Maksimovic, D. Calonico, J. Grünert, C. Mandache, P. Lemonde, G. Santarelli, P. Laurent, A. Clairon, and C. Salomon, "Search for variations of fundamental constants using atomic fountain clocks," Phys. Rev. Lett. **90**, 150 801 (2003).

[56] H. J. Metcalf and P. van der Straten, Laser Cooling and Trapping (Springer, New York, 1999).

[57] T. Mukaiyama, H. Katori, T. Ido, Y. Li, and M. Kuwata-Gonokami, "Recoil-limited laser cooling of ^{87}Sr atoms near the Fermi temperature," Phys. Rev. Lett. **90**, 113 002 (2003).

[58] L. E. Nelson, D. J. Jones, K. Tamura, H. A. Haus, and E. P. Ippen, "Ultrashort-pulse fiber ring lasers," Appl. Phys. B **65**, 277–294 (1997).

[59] A. Poppe, R. Holzwarth, A. Apolonski, G. Tempea, C. Spielmann, T. W. Hänsch, and F. Krausz, "Few-cycle optical waveform synthesis," Appl. Phys. B: Lasers Opt. **72**, 373–376 (2001).

[60] J. D. Prestage, R. L. Tjoelker, and L. Maleki, "Atomic clocks and variations of the fine-structure constant," Phys. Rev. Lett. **74**, 3511–3514 (1995).

[61] L. E. Primas, R. T. Logan, Jr., and G. F. Lutes, "Applications of ultra-stable fiber optic distribution systems," in Proceedings of the 43rd Annual Symposium on Frequency Control, 1989 (Institute of Electrical and Electronics Engineers, Piscataway, New Jersey, 1989), pp. 202–211.

[62] T. J. Quinn, "Practical realization of the definition of the metre (1997)," Metrologia **36**, 211–244 (1999).

[63] R. J. Rafac, B. C. Young, J. A. Beall, W. M. Itano, D. J. Wineland, and J. C. Bergquist, "Sub-dekahertz ultraviolet spectroscopy of ^{199}Hg^{+}," Phys. Rev. Lett. **85**, 2462–2465 (2000).

[64] N. F. Ramsey, "Applications of atomic clocks," in Laser Physics at the Limits, H. Figger, D. Meschede, and C. Zimmermann, eds. (Springer, Berlin, 2002), pp. 3–8.

[65] J. K. Ranka, R. S. Windeler, and A. J. Stentz, "Visible continuum generation in air-silica microstructure optical fibers with anomalous dispersion at 800 nm," Opt. Lett. **25**, 25–27 (2000).

[66] L. A. Rawley, J. H. Taylor, M. M. Davis, and D. W. Allan, "Millisecond pulsar PSR 1937+21: A highly stable clock," Science **238**, 761–765 (1987).

[67] R. Santra, E. Arimondo, T. Ido, C. H. Greene, and J. Ye, "High-accuracy optical clock via three-level coherence in neutral bosonic ^{88}Sr," Phys. Rev. Lett. **94**, 173 002 (2005).

[68] R. K. Shelton, S. M. Foreman, L.-S. Ma, J. L. Hall, H. C. Kapteyn, M. M. Murnane, M. Notcutt, and J. Ye, “Subfemtosecond timing jitter between two independent, actively synchronized, mode-locked lasers,” Opt. Lett. **27**, 312–314 (2002).

[69] R. K. Shelton, L.-S. Ma, H. C. Kapteyn, M. M. Murnane, J. L. Hall, and J. Ye, “Phase-coherent optical pulse synthesis from separate femtosecond lasers,” Science **293**, 1286–1289 (2001).

[70] B. Shillue, S. AlBanna, and L. D’Addario, “Transmission of low phase noise, low phase drift millimeter-wavelength references by a stabilized fiber distribution system,” in 2004 IEEE International Topical Meeting on Microwave Photonics Technical Digest (Institute of Electrical and Electronics Engineers, Piscataway, New Jersey, 2004), pp. 201–204.

[71] D. E. Spence, P. N. Kean, and W. Sibbett, “60-fsec pulse generation from a self-mode-locked Ti-sapphire laser,” Opt. Lett. **16**, 42–44 (1991).

[72] U. Sterr, C. Degenhardt, H. Stoehr, C. Lisdat, H. Schnatz, J. Helmcke, F. Riehle, G. Wilpers, C. Oates, and L. Hollberg, “The optical calcium frequency standards of PTB and NIST,” C. R. Phys. **5**, 845–855 (2004).

[73] G. Stix, “Real time,” Sci. Am. **287**, 36–39 (2002).

[74] M. Takamoto, F.-L. Hong, R. Higashi, and H. Katori, “An optical lattice clock,” Nature **435**, 321–324 (2005).

[75] H. R. Telle, D. Meschede, and T. W. Hänsch, “Realization of a new concept for visible frequency division: Phase locking of harmonic and sum frequencies,” Opt. Lett. **15**, 532–534 (1990).

[76] J. P. Turneaure, C. M. Will, B. F. Farrell, E. M. Mattison, and R. F. C. Vessot, “Test of the principle of equivalence by a null gravitational red-shift experiment,” Phys. Rev. D **27**, 1705–1714 (1983).

[77] T. Udem, S. A. Diddams, K. R. Vogel, C. W. Oates, E. A. Curtis, W. D. Lee, W. M. Itano, R. E. Drullinger, J. C. Bergquist, and L. Hollberg, “Absolute frequency measurements of the Hg^+ and Ca optical clock transitions with a femtosecond laser,” Phys. Rev. Lett. **86**, 4996–4999 (2001).

[78] J. T. Verdeyen, Laser Electronics, Prentice Hall Series in Solid State Physical Electronics (Prentice Hall, Englewood Cliffs, New Jersey, 1995).

[79] R. F. C. Vessot, M. W. Levine, E. M. Mattison, E. L. Blomberg, T. E. Hoffman, G. U. Nystrom, B. F. Farrel, R. Decher, P. B. Eby, C. R. Baugher, J. W. Watts, D. L. Teuber, and F. D. Wills, “Test of relativistic gravitation with a space-borne hydrogen maser,” Phys. Rev. Lett. **45**, 2081–2084 (1980).

[80] F. L. Walls, “Phase noise issues in femtosecond lasers,” in Laser Frequency Stabilization, Standards, Measurement, and Applications, J. L. Hall and J. Ye, eds. (SPIE, Bellingham, Washington, 2001), Vol. 4269 of Proceedings of SPIE, pp. 170–177.

[81] J. K. Webb, M. T. Murphy, V. V. Flambaum, V. A. Dzuba, J. D. Barrow, C. W. Churchill, J. X. Prochaska, and A. M. Wolfe, “Further evidence for cosmological evolution of the fine structure constant,” Phys. Rev. Lett. **87**, 091 301 (2001).

[82] G. Wilpers, T. Binnewies, C. Degenhardt, U. Sterr, J. Helmcke, and F. Riehle, “Optical clock with ultracold neutral atoms,” Phys. Rev. Lett. **89**, 230 801 (2002).

[83] L. Xu, C. Spielmann, A. Poppe, T. Brabec, F. Krausz, and T. W. Hänsch, “Route to phase control of ultrashort light pulses,” Opt. Lett. **21**, 2008–2010 (1996).

[84] J. Ye and S. T. Cundiff, “Optical frequency combs and their applications,” in Femtosecond Optical Frequency Comb Technology: Principle, Operation, and Applications, J. Ye and S. T. Cundiff, eds. (Springer, New York, 2005), pp. 12–53.

[85] J. Ye, L.-S. Ma, and J. L. Hall, “Molecular iodine clock,” Phys. Rev. Lett. **87**, 270 801 (2001).

[86] J. Ye, J. L. Peng, R. J. Jones, K. W. Holman, J. L. Hall, D. J. Jones, S. A. Diddams, J. Kitching, S. Bize, J. C. Bergquist, L. W. Hollberg, L. Robertsson, and L.-S. Ma, “Delivery of high-stability optical and microwave frequency standards over an optical fiber network,” J. Opt. Soc. Am. B-Opt. Phys. **20**, 1459–1467 (2003).

[87] J. Ye, L. Robertsson, S. Picard, L.-S. Ma, and J. L. Hall, “Absolute frequency atlas of molecular I_2 lines at 532 nm,” IEEE Trans. Instrum. Meas. **48**, 544–549 (1999).

[88] J. Ye, H. Schnatz, and L. W. Hollberg, “Optical frequency combs: From frequency metrology to optical phase control,” IEEE. J. Sel. Top. Quantum Electron. **9**, 1041–1058 (2003).

[89] B. C. Young, F. C. Cruz, W. M. Itano, and J. C. Bergquist, “Visible lasers with subhertz linewidths,” Phys. Rev. Lett. **82**, 3799–3802 (1999).

Appendix A

Derivation of Group Velocity with Kerr Nonlinearity

To find the group velocity of an ultrashort pulse in the presence of the Kerr nonlinearity, we first express the electric field in terms of the slowly varying field amplitude $A(\boldsymbol{r}, t)$.

$$E(\boldsymbol{r}, t) = A(\boldsymbol{r}, t)\mathrm{e}^{\mathrm{i}(k_0 z - \omega_c t)} + \text{c.c.} \tag{A.1}$$

where ω_c is the center frequency of the pulse bandwidth and $k_0 = n_0\omega_c/c$ is the linear part of the wavevector at ω_c. $A(\boldsymbol{r}, t)$ satisfies the general form of the nonlinear Schrödinger equation given in Eqn. (13.2.25) in [8].

$$\left[\left(1 + \frac{\mathrm{i}}{\omega_c}\frac{\partial}{\partial\tau}\right)^{-1}\nabla_\perp^2 + 2\mathrm{i}k_0\frac{\partial}{\partial z'} + 2k_0 D\right]A(\boldsymbol{r}, t) = -\frac{4\pi\omega_c^2}{c^2}\left(1 + \frac{\mathrm{i}}{\omega_c}\frac{\partial}{\partial\tau}\right)p \tag{A.2}$$

In this equation, D is the differential operator

$$D = \sum_{n=2}^{\infty}\frac{1}{n}k_n\left(\mathrm{i}\frac{\partial}{\partial\tau}\right)^n \tag{A.3}$$

where k_n indicates the nth derivative of the linear part of the wavevector with respect to ω, evaluated at ω_c. D represents second- and higher-order dispersion. p is the slowly varying amplitude of the polarization of the material through which the pulse is propagating. Finally, Eqn. (A.2) is expressed in a retarded time frame specified by

$$z' = z \quad \text{and} \quad \tau = t - k_1 z \tag{A.4}$$

To calculate the group velocity we can ignore diffraction of the pulses, which is equivalent to dropping the $\nabla_\perp^2$ term in Eqn. (A.2). In a mode-locked laser we can also ignore higher-order dispersion and drop the D term. The Kerr nonlinearity in a material is produced by a third-order nonlinearity. Therefore, including the effects of dispersion up to first-order, we have for p

$$p = 3\left[\chi^{(3)}(\omega_c) + \frac{\mathrm{d}\chi^{(3)}}{\mathrm{d}\omega}\mathrm{i}\frac{\partial}{\partial\tau}\right]|A|^2 A \tag{A.5}$$

which is given in Eqn. (13.2.28) in [8]. Using this in Eqn. (A.2), we have

$$\begin{aligned} 2\mathrm{i}k_0\frac{\partial}{\partial z'}A(\boldsymbol{r},t) &= -\frac{12\pi\omega_c^2}{c^2}\left(1+\frac{\mathrm{i}}{\omega_c}\frac{\partial}{\partial\tau}\right)\left[\chi^{(3)}(\omega_c) + \frac{\mathrm{d}\chi^{(3)}}{\mathrm{d}\omega}\mathrm{i}\frac{\partial}{\partial\tau}\right]|A|^2 A \\ &= -\frac{12\pi\omega_c^2}{c^2}\left\{\chi^{(3)}(\omega_c) + \mathrm{i}\left[\frac{\mathrm{d}\chi^{(3)}}{\mathrm{d}\omega} + \frac{1}{\omega_c}\chi^{(3)}(\omega_c)\right]\frac{\partial}{\partial\tau}\right\}|A|^2 A \end{aligned} \tag{A.6}$$

where the term higher than first-order in $\partial/\partial\tau$ has been dropped. From this we can write

$$\frac{\partial}{\partial z'}A(\boldsymbol{r},t) = \frac{6\pi\omega_c^2}{k_0c^2}\left\{\mathrm{i}\chi^{(3)}(\omega_c) - \left[\frac{\mathrm{d}\chi^{(3)}}{\mathrm{d}\omega} + \frac{1}{\omega_c}\chi^{(3)}(\omega_c)\right]\frac{\partial}{\partial\tau}\right\}|A|^2 A \tag{A.7}$$

Now, to obtain the propagation equation for the intensity envelope of the pulse from Eqn. (A.7), consider first that

$$\frac{\partial}{\partial\tau}|A|^2 A = 2AA^*\frac{\partial A}{\partial\tau} + A^2\frac{\partial A^*}{\partial\tau} \tag{A.8}$$

and so

$$\begin{aligned} A^*\frac{\partial}{\partial\tau}|A|^2 A &= 2|A|^2 A^*\frac{\partial A}{\partial\tau} + |A|^2 A\frac{\partial A^*}{\partial\tau} \\ &= |A|^2 A^*\frac{\partial A}{\partial\tau} + |A|^2\frac{\partial}{\partial\tau}|A|^2 \end{aligned} \tag{A.9}$$

Multiplying Eqn. (A.7) by A^* and using Eqn. (A.9) produces

$$\begin{aligned} A^*\frac{\partial}{\partial z'}A(\boldsymbol{r},t) &= \frac{6\pi\omega_c^2}{k_0c^2}\mathrm{i}\chi^{(3)}(\omega_c)|A|^4 \\ &\quad - \frac{6\pi\omega_c^2}{k_0c^2}\left[\frac{\mathrm{d}\chi^{(3)}}{\mathrm{d}\omega} + \frac{1}{\omega_c}\chi^{(3)}(\omega_c)\right]\left[|A|^2 A^*\frac{\partial A}{\partial\tau} + |A|^2\frac{\partial}{\partial\tau}|A|^2\right] \end{aligned} \tag{A.10}$$

Adding each side of this equation to its complex conjugate provides the propagation equation for the intensity envelope.

$$\frac{\partial}{\partial z'}|A|^2 = -\frac{18\pi\omega_c^2}{k_0c^2}\left[\frac{\mathrm{d}\chi^{(3)}}{\mathrm{d}\omega} + \frac{1}{\omega_c}\chi^{(3)}(\omega_c)\right]|A|^2\frac{\partial}{\partial\tau}|A|^2 \tag{A.11}$$

To determine the group velocity, which corresponds to the velocity of propagation of the intensity envelope, we need to express Eqn. (A.11) in terms of coordinates in the laboratory frame instead of the retarded time frame. From Eqn. (A.4), we have

$$\frac{\partial}{\partial z'} = \frac{\partial}{\partial z} + k_1 \frac{\partial}{\partial t} \quad \text{and} \quad \frac{\partial}{\partial \tau} = \frac{\partial}{\partial t} \tag{A.12}$$

With this coordinate transformation, the propagation equation in the laboratory frame becomes

$$\frac{\partial}{\partial z} |A|^2 + \left\{ k_1 + \frac{18\pi\omega_c^2}{k_0 c^2} \left[\frac{\mathrm{d}\chi^{(3)}}{\mathrm{d}\omega} + \frac{1}{\omega_c}\chi^{(3)}(\omega_c) \right] |A|^2 \right\} \frac{\partial}{\partial t} |A|^2 = 0 \tag{A.13}$$

This differential equation has the form $\frac{\partial}{\partial z} |A|^2 + \gamma \frac{\partial}{\partial t} |A|^2 = 0$, which has a solution of the form $|A|^2 \sim \mathrm{e}^{\mathrm{i}(z - \gamma^{-1} t)}$ that propagates at the velocity γ^{-1}. Therefore, from Eqn. (A.13) the group velocity, v_g, can be identified as

$$v_g = \left\{ k_1 + \frac{18\pi\omega_c^2}{k_0 c^2} \left[\frac{\mathrm{d}\chi^{(3)}}{\mathrm{d}\omega} + \frac{1}{\omega_c}\chi^{(3)}(\omega_c) \right] |A|^2 \right\}^{-1} \tag{A.14}$$

From Eqn. (4.1.16) in [8], the intensity of the pulse, I, is given by

$$I = \frac{n_0 c}{2\pi} |A|^2 \tag{A.15}$$

Also, the nonlinear coefficient of the index of refraction, n_2, is given by Eqn. (4.1.19) in [8].

$$n_2 = \frac{12\pi^2}{n_0^2 c}\chi^{(3)} \tag{A.16}$$

Using

$$k_1 = \left[\frac{\mathrm{d}}{\mathrm{d}\omega} \left(\frac{\omega n_0}{c} \right) \right]_{\omega_c} = \frac{n_0 + \omega_c \left(\frac{\mathrm{d}n_0}{\mathrm{d}\omega} \right)_{\omega_c}}{c} \tag{A.17}$$

the expression for the group velocity can be rewritten as

$$v_g = \frac{c}{n_0 + \omega_c \left(\frac{\mathrm{d}n_0}{\mathrm{d}\omega} \right)_{\omega_c} + 3 \left[n_2 + \omega_c \left(\frac{\mathrm{d}n_2}{\mathrm{d}\omega} \right)_{\omega_c} \right] I} \tag{A.18}$$

where the term in the denominator proportional to $\omega_c (n_2/n_0)(\mathrm{d}n_0/\mathrm{d}\omega)_{\omega_c} I$ has been dropped since it is negligibly small for the parameters given in Section 3.1.

Appendix B

Single-side-band Generator

A single-side-band generator provides a way to effectively shift the frequency of a signal by a relatively small amount (for example, 10 kHz) [40]. This is useful for measuring the frequency instability of a given signal when it is to be compared against a reference signal of the same frequency [41]. Shifting the reference signal by 10 kHz and mixing the shifted reference and the signal to be measured provides a 10-kHz signal that can be counted with sufficient resolution to detect the frequency fluctuations of the measured signal. The single-side-band generator operates by mixing the original signal, or carrier, with a 10-kHz signal to produce an upper and lower side band on the carrier, and then one of the side bands is interferometrically cancelled.

Figure B.1 illustrates the schematic of a single-side-band generator, with the specific parameters of some of the components given for the single-side-band generator used in the measurements discussed in Chapter 6. The signal to be shifted is first set to a power level of 10 dBm and then split, producing two in-phase 7-dBm signals that each drive the LO port of a frequency mixer. This original signal, at frequency ν, can be considered proportional to $\cos(2\pi\nu t)$. Any phase offset is irrelevant, since it would be equivalent for the two signals produced after splitting the original signal, and is therefore neglected. The two signals driving the mixer LO ports are mixed with two signals at frequency ν' that have a relative phase difference of ϕ_0, where ν' is the frequency difference between the original signal and the side band to be generated. For

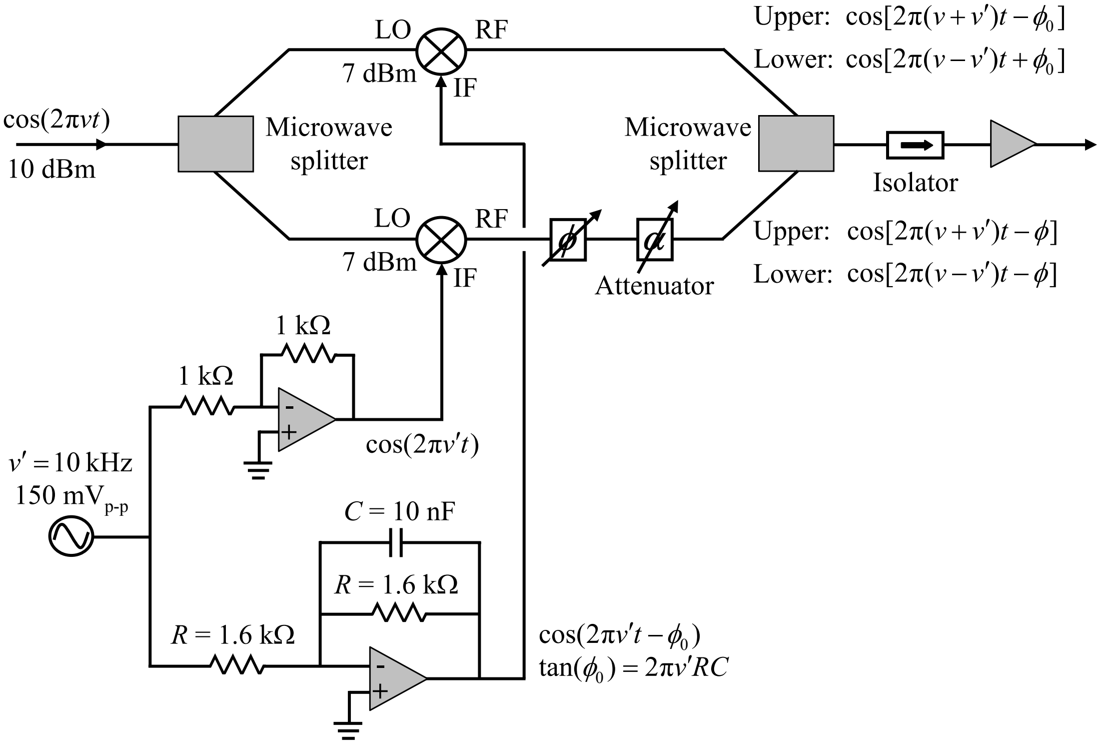

Figure B.1: The single-side-band generator operates by splitting the carrier signal and mixing the two portions with two 10-kHz signals that have a relative phase difference, producing 10-kHz upper and lower side bands on the carrier signals. Combining the mixer outputs and introducing the appropriate relative phase difference between them results in the cancellation of one of the side bands.

the scenario presented here, $\nu' = 10$ kHz. A stable synthesizer operating at 10 kHz and two op-amp circuits produce the two 10-kHz signals with a relative phase difference of ϕ_0. The upper op-amp circuit in Fig. B.1 has unity gain and produces a signal proportional to $\cos(2\pi\nu' t)$. The lower op-amp circuit exhibits a gain of < 1 and produces a signal proportional to $\cos(2\pi\nu' t - \phi_0)$, where ϕ_0 is given by

$$\tan(\phi_0) = 2\pi\nu' RC \tag{B.1}$$

R and C are the resistance and capacitance values of this circuit, as shown in Fig. B.1. Injecting these two 10-kHz signals into the IF ports of the two mixers produces an output from the RF port of each mixer consisting of an upper side band with a phase term that is the sum of the phase terms of the mixer inputs, and a lower side band with a phase term given by the difference of the phase terms of the mixer inputs. An adjustable phase shifter and attenuator after one of the mixers allows the upper side bands of the two mixer outputs to be set to the same amplitude and 180° out of phase, so that combining the two outputs results in the complete cancellation of the upper side band.

With the adjustable phase shifter placed after the lower mixer in Fig. B.1 and introducing a phase delay of ϕ, the upper and lower side bands of the lower mixer output are proportional to

$$\text{Upper}: \quad \cos[2\pi(\nu + \nu')t - \phi] \tag{B.2}$$

$$\text{Lower}: \quad \cos[2\pi(\nu - \nu')t - \phi] \tag{B.3}$$

The side bands of the upper mixer output, which is produced using the 10-kHz signal that is lagging in phase by ϕ_0, are proportional to

$$\text{Upper}: \quad \cos[2\pi(\nu + \nu')t - \phi_0] \tag{B.4}$$

$$\text{Lower}: \quad \cos[2\pi(\nu - \nu')t + \phi_0] \tag{B.5}$$

From these expressions, it may at first seem that to cancel the upper side band the optimal choices for ϕ_0 and ϕ are $\phi_0 = 90\,^\circ$ and $\phi = 270\,^\circ$, since the upper side bands will be $180\,^\circ$ out of phase and the lower side bands will be completely in phase. However, in the limit that ϕ_0 approaches $90\,^\circ$, the gain of the lower op-amp circuit approaches 0, eliminating the 10-kHz signal that is delayed by ϕ_0. Instead, a good choice for ϕ_0 is $45\,^\circ$, which corresponds to an attenuation by only a factor of $1/\sqrt{2}$. Then R and C must satisfy the condition $RC = (2\pi\nu')^{-1}$, which they do for the values given in Fig. B.1 when $\nu' = 10$ kHz. With $\phi_0 = 45\,^\circ$, the upper side band is cancelled with $\phi = 225\,^\circ$, which corresponds to a phase difference between the lower side bands of $90\,^\circ$. Note that selecting $\phi = 135\,^\circ$ would instead cancel the lower side band while preserving the upper side band. After the cancellation of one side band, the remaining side band is amplified to a level suitable for use in subsequent measurements. A microwave isolator is placed before the amplifier to prevent reflections from disturbing the destructive interference of the cancelled side band.

Figure B.2 shows an example spectrum of the output from the single-side-band generator, obtained with an RF spectrum analyzer. The original signal has a frequency of 760 MHz, and from Fig. B.2 it is evident that this carrier frequency is still present in the output of the single-side-band generator. This is because of leakage of the carrier through the mixers. The largest peak, 10 kHz below the carrier, is the uncancelled lower side band. The upper side band has been nearly completely cancelled and is only slightly above the noise floor of the spectrum analyzer. After mixing the output from the single-side-band generator with a signal at the carrier frequency for which the frequency fluctuations are to be measured, any low-frequency noise on the 10-kHz mixing product caused by the carrier mixing with the signal to be measured is easily eliminated with a 1-kHz high-pass filter. The peaks in the spectrum that are 20 and 30 kHz from the carrier are due to additional frequency mixings in the two mixers. The ratio of the lower-side-band power to the power of these peaks is maximized by varying

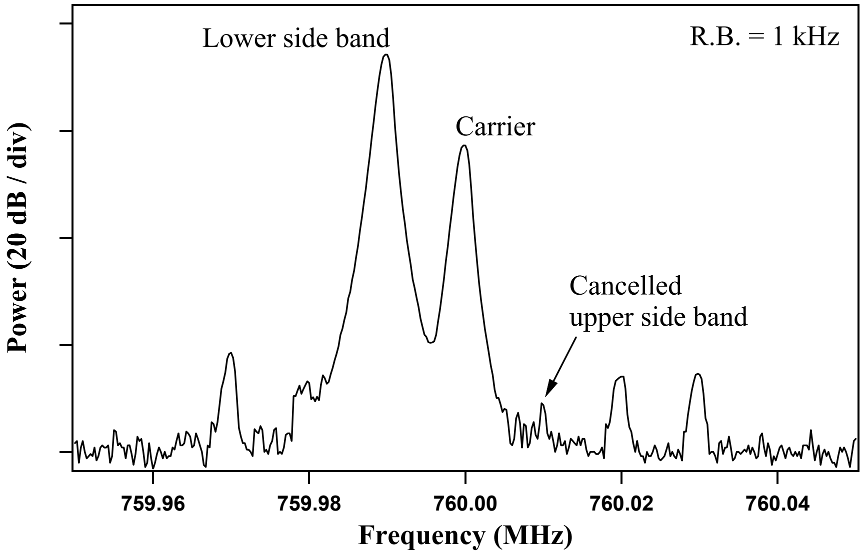

Figure B.2: The spectrum of the output from the single-side-band generator shows that the largest spectral peak is the desired lower side band. Because of leakage through the mixers, the carrier frequency is also present in the output. The upper side band has been nearly completely cancelled. R.B., resolution bandwidth.

the amplitude of the 10-kHz synthesizer used to produce the 10-kHz signals that are mixed with the original signal in the single-side-band generator. The optimal amplitude for the synthesizer is 150 mV peak-to-peak, as indicated in Fig. B.1.

When the single-side-band generator is used for determining the frequency instability of a signal as described earlier, it is important to consider what limits the stability of the single-side-band generator itself. Suppose the frequency of a 1-GHz signal is to be measured with respect to a 1-GHz reference with a sensitivity of a part in 10^{15} for an averaging time of 1 s. This corresponds to a 1-μHz deviation of the reference. Therefore, the 10-kHz frequency, which shifts the 1-GHz reference for comparison against the signal to be measured, must have fluctuations less than 1 μHz over a 1-s averaging time. Though this corresponds to an easily attainable fractional frequency deviation of the 10-kHz source of 1×10^{-10} for a 1-s averaging time, few commercially available high-stability 10-kHz sources provide an absolute frequency precision better than 1 μHz. Therefore, the fractional frequency instability of a single-side-band generator designed for a 1-GHz reference is typically $> 1 \times 10^{-15}$ over a 1-s averaging time, limited by the precision of the 10-kHz source.