

**Femtosecond Frequency Combs for Optical Clocks
and Timing Transfer**

by

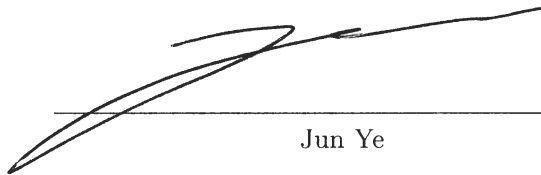
Seth M. Foreman

B.S., Harvey Mudd College, 2001

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

2007

This thesis entitled:
Femtosecond Frequency Combs for Optical Clocks and Timing Transfer
written by Seth M. Foreman
has been approved for the Department of Physics



Jun Ye



John L. Hall

Date July 10, 2007

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.

Foreman, Seth M. (Ph.D., Physics)

Femtosecond Frequency Combs for Optical Clocks and Timing Transfer

Thesis directed by Assoc. Prof. Adjoint Jun Ye

The rapid development of femtosecond optical frequency combs over the last decade has brought together ultrastable phase control of both cw and mode-locked lasers and ultrafast time-domain applications. Frequency-domain laser stabilization techniques applied to the ultrashort-pulse trains emitted by a mode-locked laser result in a level of optical phase control previously achievable only for radio frequencies and microwaves. I present our work extending such control to mode-locked lasers for both timing and frequency stabilization applications of optical frequency combs.

I first present a microwave technique for synchronizing two independent mode-locked lasers at a level of timing precision less than the duration of an optical cycle, below 1 fs of residual rms timing jitter. Using these synchronized pulses, simultaneous sum- and difference-frequency generation of 400-nm and tunable mid-infrared fs pulses is demonstrated, opening the door for broadband coherent control of atomic and molecular systems.

For frequency metrology, I report on an offset-free clockwork for an optical clock based on the 3.39- μm transition in methane. The clockwork's simplicity leads to a robust and reliable table-sized optical frequency reference with instability approaching a few parts in 10^{14} . Then I describe a directly-octave-spanning, self-referenced Ti:sapphire laser employed as the robustly-running phase-coherent clockwork for an ^{87}Sr optical lattice clock. The optical comb distributes the 2-s coherence time of the 698-nm ultrastable clock laser to its modes spanning the visible and near-IR spectrum, and is therefore simultaneously used as a hub for measuring absolute frequencies or frequency ratios between the Sr clock and other remotely-located microwave and optical atomic

standards.

Finally, I report on the transfer of ultrastable frequency references, both microwave and optical, through 10-km-scale optical fiber links. Actively stabilizing the optical phase delay of such a fiber link, we are able to transfer a cw optical frequency standard with a transfer instability of 6×10^{-18} at 1 s, more than two orders of magnitude lower than reported for any fiber link of similar length. Phase coherence between ends of the fiber link is preserved at the mHz linewidth level, and the transfer phase noise corresponds to less than 80 attoseconds of rms timing jitter integrated from 10 mHz to 30 MHz.

Dedication

to Dad: we could have a long “Important Talk” about this

and

to Mom: the most loving and giving person I know

Acknowledgements

I was fortunate enough to arrive on the scene in Jun's lab just as stabilized femtosecond frequency combs were being developed. He gave me a choice: would I like to work on an exciting new experiment with lasers that emit ultrashort bursts of light, or would I like to work on a system of cold Strontium atoms that could potentially be turned into a clock like the cold atomic fountain? As I have travelled from working with the frequency comb to learning about precision atomic spectroscopy and optical clocks, I've learned from and been encouraged by a great many people at JILA and in Jun's group. I'm extremely grateful to all these people for their significant contributions to my achievements and to furthering my development as both a skilled experimentalist and a simple, curious person.

First, thank you Jun, for maintaining the kind of environment where we can all stay interested in each others' work and keep discussing new, interesting, sometimes crazy, and once in a while great ideas with each other. I also very much appreciate being allowed to go at a project on my own terms, with just enough guidance and oversight to help me through the rough spots. The open discussions and keen interest in science you have fostered in your students keeps the excitement level high, and of course your technical expertise makes sure that our lofty plans can come to fruition – for that, I am certainly grateful!

In writing about the open and interesting discussions, nearly all the students and post-docs I've had the good fortune to work with come to mind. In the early days, Rob-

bie Shelton patiently taught me the secrets of the black art of Ti:sapphire lasers. Pat Kohl and I worked together on homework and various other (sometimes mischievous, always great fun) projects. Mark Notcutt is responsible for teaching me how to make technical measurements of noise. David Jones gave me the example of building experiments efficiently and expertly. Eric Hudson and especially Kevin Holman were always willing to think through the details of a problem with me, without losing focus on the bigger, more exciting picture. Then when Jason Jones arrived, I learned that Ti:sapphire lasers weren't **always** a black art. I owe a great debt to Tara Fortier, who showed me how to tweak a laser into perfection using patience and the right phraseology. Chikako Ishibashi is to thank for keeping up my interests and understanding of Japanese. Darren Hudson has the right perspective on how to perform good science and enjoy life, both at the same time. He also has excellent taste in movie characters. More recently, I've also been delighted to have thoughtful discussions with Matt Stowe and Mike Thorpe, who make sure I'm thinking about things correctly. Their ability to see to the fundamentals of an issue is a valuable skill and one I hope to emulate throughout my career. Dylan Yost is also keen to understand and discuss physics from a thoughtful perspective. Thomas Schibli's expertise in nearly anything I want to know about has certainly helped me to think about technical issues and make good technical choices. The entire Strontium "team," including Sebastian Blatt, Marty Boyd, Gretchen Campbell, Tetsumiya Ido, Andrew Ludlow, Thomas Zanon, and Tanya Zelevinsky, have my thanks for teaching me basic atomic physics and bringing my frequency comb and I into their fun. In particular I'm thankful to Andrew for keeping me honest and showing me how to be thorough when thinking through an issue. Gretchen Campbell is responsible for showing me how to get a post-doc position and finish up a thesis, and I'm grateful for her advice and motivation. And most recently of all, Marcio Miranda ("B" is for "Brazilian") and Mike Martin (Junior, Steve, Steve Martin, Mudd-Mike, etc.) have been a pleasure to teach and work with, especially when they ask me great questions I should have known

the answers to. Of course there have been many others I didn't interact as closely with; all help to keep up the interactive atmosphere we enjoy.

The general atmosphere of the frequency metrology community has also nearly always been quite friendly and open, and I hope that atmosphere continues to thrive in the face of some recent challenges. I'm thankful to Scott Diddams, Leo Hollberg, Jim Bergquist, Tom Parker, and Jason Stalnaker, all at NIST, for their friendly expertise and cooperation in all the fiber transfer experiments and remote clock comparisons we've performed. Another collaboration from the metrology community will always be ingrained in my memory: Oliver Mücke from MIT, and Mikhail Gubin and Evgeny Petrukhin from the Lebedev Institute in Moscow were a great pleasure to work with on the offset-free comb and methane clock experiments.

Fortunately, the laboratory environment isn't always just about science, and many people in the lab understand how a rowdy game of ping-pong can keep up good spirits. In order of decreasing ping-pong talent (no letters, please), in my estimation, they are: Xueren Huang (ping-pong Huang), Jun, Tetsuya Ido, Mike Thorpe, Jason Jones, Chris Ticknor, Marty and Darren or Darren and Marty, Brian Sawyer (Jazz), Kevin Holman, Sebastian, and the theorists. If you are reading this, my Chinese Pen-Style Grip can defeat your Western Handshake Style. (Thank you, Huang.)

Another huge support I've had while at JILA is the technical excellence of an electronics and machine-shop staff that I'm sure I will sorely miss the very next time I need to build anything. Terry Brown deserves special thanks for all his amazing work over the years, especially including the low-noise amplifiers and Z-matched photodiodes that still have unmeasurably-low phase noise. Also thanks go to James Fung-A-Fat, Mike Whitmore, Dave Tegart, Carl Sauer, and Paul Beckingham. In the instrument shop, I was always able to ask Todd Asnicar and Hans Green for anything and get the most helpful and immediate solutions. Blaine Horner, Kim Hagen, Leslie Czaia, and Dave Alchenberger also made important contributions to my experiments. The supply

office has also made my work more effortless through the diligence of Brian Lynch, Jeff Sauter, Dave Errickson, and in the early days, Ed Holliness and Maryly Dole.

I also would like to thank my thesis committee for their contributions to my work. Jan Hall has given me excellent technical advice during many of my projects and perspective on life as a physicist. Steve Cundiff has collaborated with many of my projects, Leo Hollberg is a model of a careful and open scientist, Konrad Lehnert spent hours-long discussions with me about amplifier phase noise and microwave circuitry, and Kelvin Wagner taught me nonlinear optics.

I've reserved the final paragraph for my family. My Dad (John Foreman) brought me up to be curious about things and think them through, usually with brainstorming ideas about how things **probably** work during "Important Talks" on the deck – we sometimes even got it right! From him I learned the Pythagorean Theorem, simple machines using "Work equals Force times Distance," how to build a deck, fix a toilet, and destroy the neighbors' gas line – all the things a man should know. My Mom (Marcia Foreman) made absolutely sure I always knew I was loved, no matter what, with her carefully-constructed heart-shaped bologna sandwiches on Valentine's Day, never-ending Jack-In-The-Box paraphernalia, and helping me take what had to seem to her to be the most boring of data whenever she visited me. These few short phrases can't capture all the ways my parents helped shape me into who I am today. The only way I can say thank you for the amazing environment they created for their children is to try to emulate it for my own one day. Of course, my brother and sister (Joel and Rachel) have also shaped me – Joel made sure I always enjoyed life while beating him at wiffle-ball or throwing football patterns until his 3-year age deficit stopped being significant. And Rachel was the person I could always count on to understand how good it can be to go enjoy and learn from a good book on our own without having to always be around people. Finally, most recently and more directly during my time at JILA, my wife Laura has allowed me to pursue this degree for six long years. It's a bit

ridiculous to ask a spouse to put up with low income and long hours for so long, but despite that stress and difficulty she has nearly always been ready to jump to my side in the most difficult times at work. She also made her own significant contributions to our work in the lab: drilling and tapping hundreds of holes in a breadboard, sewing zippers on plastic curtains for keeping optical tables clean, making delicious treats and hearty 11-PM meals for the Strontium team, and playing gutsy softball for the JilaMonsters, to name a few. Thank you, Laura. The Foreman Clan (now with added Kimura!) has been an amazingly closeknit and supportive family that I'm proud to be a part of.

Contents

Chapter

1	Introduction to Clocks: Optical Frequency Standards and Timing Transfer	1
1.1	A brief history of frequency	2
1.1.1	Sundials: the first frequency multipliers	2
1.1.2	Flow, frequency, and synchronization	3
1.1.3	Natural frequencies and isolation	5
1.1.4	Navigation and synchronization	6
1.1.5	Frequency measurement and scientific discovery	7
1.2	Modern motivations for precision timing	9
1.2.1	Atomic clocks and the current primary frequency standard	9
1.2.2	Applications of modern microwave atomic frequency standards	10
1.2.3	Applications of improved frequency and timing control	12
1.2.4	The need for frequency and timing transfer techniques	15
1.3	Next generation clocks: optical frequency standards	17
1.4	The case for the comb: precision clockwork	20
2	Characterization and measurement of frequency stability and phase noise	24
2.1	Characterization methods	24
2.1.1	Phase noise	26
2.1.2	Timing jitter	27

2.1.3	Allan deviation	28
2.2	Limitations to optoelectronic phase measurement	30
2.2.1	Thermal and shot noise	31
2.2.2	Electronic flicker noise and amplitude-to-phase conversion	34
2.3	Techniques for optical timing measurement	36
2.3.1	Optical-phase-insensitive techniques	37
2.3.2	Optical-phase-sensitive techniques	39
3	Fundamentals of Stabilizing an Optical Frequency Comb	41
3.1	Mode-locked Ti:sapphire laser	42
3.2	Frequency and time domain pictures	45
3.3	Servo control of frequency comb	50
3.3.1	Detection and stabilization of the offset frequency	50
3.3.2	Stabilization of the repetition frequency	52
4	Optical to Microwave Conversion	54
4.1	Measuring phase noise electronically	55
4.1.1	Homodyne phase detection with a mixer	55
4.1.2	Independent oscillators vs. single component measurement	57
4.1.3	Measuring the phase noise of repetition frequency harmonics	60
4.2	Noise floor of the microwave system	62
4.2.1	Low-Noise Amplifiers	64
4.2.2	Mixer Noise	64
4.2.3	Practical considerations	76
4.3	Noise floor of the optical system	77
4.4	Comparison of results	80

5	Synchronization of Femtosecond Lasers	84
5.1	Synchronization experiment	86
5.1.1	The laser system	86
5.1.2	Differential timing resolution for the servo	87
5.1.3	Synchronization performance	88
5.2	Application: infrared fs pulses	96
5.2.1	Experimental method for simultaneous DFG and SFG	98
5.2.2	Simultaneous difference- and sum-frequency generation	100
5.2.3	Broad mid-IR tunability with 500-fs pulses	101
5.2.4	Rapid wavelength switching using chirp and timing offset	104
6	Precision Infrared Spectroscopy: the Offset-Free Comb	107
6.1	Principle of offset-free spectroscopy	108
6.2	Experimental implementation of offset-free clockwork	110
6.2.1	Ti:sapphire laser with two separated spectral peaks	111
6.2.2	SFG-based offset-free IR frequency measurement	111
6.2.3	Evaluation of the OPO frequency measurement	115
6.3	Optical molecular clock based on CH ₄	117
6.3.1	Experimental setup for the clock	119
6.3.2	Evaluation of the optical molecular clock	121
7	The Octave-Spanning Comb: Clockwork and Distribution For an Optical Atomic Frequency Standard	128
7.1	Optical lattice clock using ⁸⁷ Sr atoms	129
7.1.1	The atomic clock's pendulum: cooling and trapping Strontium	130
7.1.2	Long interaction times: confinement in the optical lattice	132
7.1.3	The local oscillator: ultrastable cavity-stabilized probe laser	135
7.2	Octave-spanning Ti:Sapphire laser: optical design	137

7.3	Servo control	142
7.3.1	Control of the offset frequency	142
7.3.2	Control of another comb mode	145
7.3.3	In-loop performance of the stabilized comb	146
7.4	Out-of-loop experiments	148
7.4.1	Characterizing the clock local oscillator	149
7.4.2	Absolute frequency measurement of the ^{87}Sr optical lattice clock	152
7.4.3	Optical clock intercomparison	161
7.4.4	The comb's intrinsic coherence, stability, and accuracy	165
8	Frequency and Timing Transfer in Optical Fiber Networks	176
8.1	Motivation for ultrastable frequency reference transfer via fiber networks	177
8.2	Principles of cancellation of fiber-induced phase noise	183
8.2.1	Round-trip phase detection	185
8.2.2	Noise processes, bandwidth and dynamic range	187
8.2.3	Experimental schemes for frequency transfer	189
8.3	Microwave frequency transfer with modulated cw laser	193
8.3.1	Passive transfer technique	193
8.3.2	Passive microwave transfer for the ^{87}Sr clock	194
8.3.3	Actively-stabilized fiber transfer for 2006 ^{87}Sr measurement	199
8.4	Optical carrier transfer	203
8.4.1	Experimental techniques for cw optical carrier transfer	205
8.4.2	Coherence of optical carrier transfer	211
8.4.3	Stability of the transmitted optical carrier	215
8.4.4	Phase noise of the 7-km link	219
8.5	Remote synchronization of ultrafast lasers	222
8.6	Summary	227

Bibliography	230
---------------------	-----

Appendix

A A humorous note on units	246
-----------------------------------	-----

Figures

Figure

2.1	Phase noise spectrum showing typical fundamental and technical limits	32
3.1	Basic Ti:sapphire laser	43
3.2	Time- and frequency-domain pictures of the comb	47
3.3	Self-referencing scheme (f -to- $2f$)	51
4.1	Microwave phase measurement: saturated mixer	58
4.2	Optical-to-microwave conversion: phase noise measurement	61
4.3	Setup for measuring the microwave system's noise floor	63
4.4	Setup for measuring the mixer's noise floor	65
4.5	Amplitude-to-phase conversion and mixer flicker floor	67
4.6	Mixer output noise and dc voltage level vs. phase	69
4.7	Amplitude modulation leading to phase modulation	70
4.8	Amplitude-to-phase noise cancellation for various powers	74
4.9	Experimental setup for reducing excess noise in pulse train photodetection	78
4.10	Lowest-phase-noise optical-to-microwave conversion results	81
5.1	Experimental scheme for synchronization of independent fs lasers	89
5.2	Cross correlation for measuring timing jitter	91
5.3	Performance of the synchronization	92
5.4	Synchronization limited by mixer	93

5.5	Fast switching of the pulse trains' relative time delay	95
5.6	Experimental setup for sum- and difference-frequency generation	99
5.7	Simultaneous sum- and difference-frequency cross correlations	102
5.8	Average power of infrared pulse train from 7.5 to 12.5 μm	103
5.9	500-fs infrared pulse duration	104
5.10	Rapid switching of the SFG and DFG wavelengths	106
6.1	Scheme for offset-free precision IR spectroscopy	109
6.2	Two-peaked output spectrum of the Ti:sapphire laser	112
6.3	Setup for the offset-free SFG/DFG clockwork	114
6.4	Beat frequency for stabilizing the SFG-based clockwork	116
6.5	Instability of the OPO optical frequency	118
6.6	Spectral coverage of the offset-free IR comb	120
6.7	Heterodyne beats for locking the frequency comb to the HeNe/CH ₄ clock	122
6.8	Counting record and instability of the CH ₄ clock	124
6.9	Relative phase spectral density of the CH ₄ clock and I ₂ clock	126
7.1	Simplified Sr energy-level diagram	131
7.2	Lattice confinement and principle of the magic wavelength	134
7.3	Ultrastable, high-finesse, passive optical cavity for the Sr clock laser	136
7.4	Octave-spanning laser used as an optical clockwork	139
7.5	Octave spectrum, output coupler reflectivity and dispersion	141
7.6	In-loop measurement of the octave-spanning laser's linewidth	147
7.7	Octave-spanning comb allows cavity intercomparisons	150
7.8	⁸⁷ Sr clockwork with frequency drift compensation	153
7.9	Compensation of the clock laser's frequency drift	156
7.10	Highest quality factor for coherent spectroscopy	158
7.11	Sr optical clock counted against maser	160

7.12	Experimental setup for Sr-Ca remote clock intercomparison	162
7.13	Sr optical clock counted against Ca optical clock	164
7.14	Instability of Sr and Ca clock comparison	166
7.15	Experiment for measuring the comb's intrinsic coherence	168
7.16	Intrinsic instability of the frequency comb	172
7.17	Intrinsic linewidth of the frequency comb	173
8.1	Instabilities of various stable references to be transferred	178
8.2	Some applications of fiber network frequency- and timing-transfer	182
8.3	Principle of fiber network transfer phase stabilization	184
8.4	Microwave, cw optical, and frequency comb transfer	190
8.5	Passive instability of microwave transfer through the BRAN fiber	195
8.6	Phase fluctuations of a passively-transferred microwave signal	196
8.7	First accurate measurement of the ^{87}Sr optical lattice clock transition .	198
8.8	Phase of the passive microwave transfer over several days	202
8.9	Stabilized microwave transfer for the ^{87}Sr clock measurement	204
8.10	Three configurations for cw optical carrier transfer	208
8.11	Linewidths for cw optical transfer and 1-Hz remote cavity comparison .	214
8.12	Instability of the 7- and 32-km cw optical carrier fiber transfer	217
8.13	Phase noise and integrated timing jitter for optical carrier transfer . . .	220
8.14	Experimental scheme for remote synchronization of ultrafast lasers . . .	224
8.15	Timing jitter of remotely synchronized fs fiber lasers	226

Chapter 1

Introduction to Clocks: Optical Frequency Standards and Timing Transfer

How many clocks we have! Pause for a moment to consider the quiet space you occupy as you read this: chances are, you can hear some form of “tick” from one of the many clocks surrounding you. Whether it’s the actual 1-Hz ticks of an analog clock’s second hand, the characteristic 60-Hz harmonics of fluorescent lighting, or the 440-Hz used to tune the music you may be listening to, frequency “standards” are all around you. This is not to mention the wireless frequencies responsible for the efficient conveniences of modern life: remote controls, wireless internet, and cell-phones to name a few. These latter examples rely on the ability to tightly synchronize their signals for their utility. As we progress through the Chapters of this Thesis, we will see the close relationship between highly accurate frequency standards and precision timing (synchronization) of signals.

In this Chapter, I will first provide a brief history of the art of counting frequency, always with an eye toward measuring the phase of a periodic signal. While discussing this history, I will enjoy describing historical developments with some of the modern scientific “lingo” of the frequency metrology community, especially since it will help the less-experienced reader (early graduate students) to draw connections between our intuitive, time-obsessed world and the less-familiar concepts of phase coherence, synchronization, jitter, instability, etc. Then I will describe some motivations for our

obsession with ever-increasing frequency accuracy (17 digits!) and synchronization. Optical frequency standards afford the exciting potential to realize these levels, and I will describe them next. Finally, the need for the femtosecond frequency comb will become clear, and I will summarize the progress of the comb, which has grown from a new tool to a mature method of optical frequency measurement in the years I have been at JILA.

1.1 A brief history of frequency

The heart of every clock is its characteristic frequency – the thing that makes it “tick” so repetitively. But only in modern times are we aware of these ticks on a second-by-second basis. Of course in ancient times, the most obvious frequency was that of the sunrise. This most fundamental of frequencies now corresponds to only $1/84,600^{\text{th}}$ of a Hz, or $1/777,696,647,742,000^{\text{th}}$ of a Cesium atomic clock’s oscillation period, which shows what a long way we have come in subdividing our day (or frequency multiplying it!). In what follows I would like to acknowledge Ref. [1] for many illustrative anecdotes about the measurement of frequency through the ages.

1.1.1 Sundials: the first frequency multipliers

As civilizations formed and societies came together, farming was increasingly substituted for hunting, and the much lower frequency of Earth’s annual trip around the sun, responsible for the seasons, became more important. At the same time, the need to coordinate efforts between people meant the need for an accurate division of the day. The ratio between the frequency of the moon’s phases and that of Earth’s orbit (roughly 12) was a natural choice for the number of hours to divide daylight amongst. This was accomplished by means of the sundial, which, by placing unevenly spaced marks along the ground, allowed the shadow of the “gnomon” to tick off twelve equally spaced hours per day. In modern language, this was initially a method of phase-incoherent frequency multiplication, since it was an educated guess as to how to space the ticks so as to get

equal subdivisions in time. As is the usual case with such frequency multiplication, it was of course inaccurate. Refinements and an understanding of the concept of latitude eventually led to sundials which, on a clear day, could determine the hours to within 5 minutes, corresponding to an inaccuracy of 3.5×10^{-3} per day. Interestingly, a sundial made for one latitude gives incorrect time at another: in 264 B.C. a sundial from Sicily was brought back to Rome as wartime loot, and kept incorrect time for a century due to the change of latitude! This would be similar to, in modern times, failing to correct frequencies for altitude differences due to general relativity. Note that in this definition of time, summer hours are longer than winter hours, which represents a very different paradigm from its modern counterpart with exactly-equal hours, no matter the time of year or location on Earth.

1.1.2 Flow, frequency, and synchronization

The need to keep time at night was eventually stimulated by the practices of various religions. In Judaism and Islam, prayers were said 3 and 5 times per day, respectively. For these religions, prayer was a personal act and could be performed at roughly the prescribed time of day, without need for accuracy. However, in Christianity the group activity of worship required the synchronization of individuals' efforts; for example, the standard of the Benedictine monks required worship at eight "canonical hours," one of which (matins) was at night. This is where the concept of synchronization enters our story: monks would come together for worship at these hours, and the ringing of bells was their synchronization reference. How did the bell-ringer know which hour of the night to perform his duty? The answer lies in the first forms of man-made clocks, which operated on the principle of "flow" rather than "frequency." Water clocks, originally developed in China for commercial rather than religious reasons, relied on the regulated flow of water from one container to another, as did sandglasses. Alternatively, burning a candle or incense could be calibrated to measure time. (Different incense for

different hours meant a clock that could be smelled!) Of these, the water clock satisfied two requirements of the medieval monk: it lasted long enough into the night that he could sleep, and it could be made into an alarm clock to wake him up by tipping over a container when the water reached a certain height. The effect was to ring little bells to wake the monk so he could ring the large bells! In this way, the strict and orderly tasks of the monks could be well-synchronized.

The principle of flow for measuring time worked well enough, but inevitably gave way to the principle of counting periods of oscillation – in other words, the more modern concept of measuring frequency. Now we speak of “counting frequency” even though “counting periods” is perhaps a more accurate phrase. As waterwheels and windmills were developed to perform laborious, repetitive tasks, the concepts of the machine and in particular differential gearing were refined. The circular motion of such machines, where a complete revolution brings the system exactly back to its starting point for another identical trip, makes the shift in thought from “flow” to “frequency” seem immediate. As an important aside, although differential gearing is not fundamental to the idea of a clock, it has been necessary in the practical operation of every clock since the earliest mechanical versions. In Chapter 3 I will show how our modern frequency comb represents most of an optical clock’s “gearing” or clockwork.

The first man-made clocks to directly employ the idea of frequency were of the verge-and-folio type, developed some time in the 13th century. A heavy mass was allowed to fall down the height of a clock tower in small increments (the ticks of the clock) due to clever gearing. The period between successive increments was regulated by a horizontal bar (the folio), suspended by a silk rope, that rotated back and forth through a large angle. This period was dependent on the moment of inertia of the folio, but also dependent on the complicated degree to which the falling mass transmitted torque to the folio. Therefore the frequency was sensitive to temperature, lubricant (which was not well-developed at the time), and mechanical wear over time. These inaccuracies all

combined so that even a good verge-and-folio clock was off by 15 minutes per day, three times worse than a sundial on a clear day. Of even more significance, the new mechanical clocks didn't know the difference between a summer and a winter hour. This is still reflected when we say "o'clock," or "of the clock," since people had to specify which standard (sun or mechanical clock) they used. In fact, so many different standards existed, and were loudly proclaimed each night by the ringing of bells, that in 1370 the King of France, angry from the noise, decreed all bells had to ring at the same time as his own. Since his bells were "of the clock," the new synchronized frequency standard enjoyed 24 equal hours, no matter the time of year.

1.1.3 Natural frequencies and isolation

Although the early mechanical clocks succeeded in automating timekeeping, even through the night hours, the accuracy provided by the Earth's rotation was still superior – mechanical clocks were reset each day to noon, when the shadow of a sundial was shortest. Finally Galileo made his famous observation that a pendulum's period depended only on its length. The resulting new pendulum clocks, first built by Huygens in 1657, reflect a basic principle in the advancement of frequency standards that has endured ever since: "natural" frequencies of isolated phenomena make good standards; increase the isolation and the potential for accuracy increases as well. The old foliot design was a mechanical resonator that wasn't isolated from the rest of the clock at all. The new pendulums depended mostly on their length (and therefore temperature), and to a lesser degree on the escapements used to count their oscillations and counter friction. As a result of the isolation of the pendulum and the simple natural frequency built into it by nature, inaccuracies quickly approached 10 seconds per day, or 1.2×10^{-4} in one day.

The newfound ability to precisely define frequency led to one of the first examples of frequency measurement allowing a new discovery of nature: it was quickly found

that sundial time could be as much as 16 minutes ahead or behind pendulum-clock time, depending on the time of year and time of day. This results from the combined effect of the tilt of Earth's axis and the ellipticity of its orbit around the sun. In fact, conversion tables between clock and sun time existed already by 1660, just 3 years after the first pendulum clock was built! The pattern of improvements in frequency metrology resulting in rapid discoveries in many areas of science continues even now.

1.1.4 Navigation and synchronization

As exploration and discovery by sea became important, so did the need for accurate navigation. The famous story of the longitude problem is ultimately a story of building a clock that keeps accurate time on board a ship at sea. All the marvellous advances in clock design that led to John Harrison's successful marine chronometer in 1761 are much too detailed to be described here; eventually his clock was capable of losing only 5 seconds during a transatlantic voyage of many weeks. This precision tool represented a large portion of a ship's cost, but finally provided for accurate navigation at sea.

Although the precision frequencies of these clocks were now accurate enough to measure the predictable nonlinear drift of sundial time, no real need for tight synchronization of distant clocks existed. In fact, each city or town in the 18th century still tended to set its own local phase by the local definition of Earth's rotation time, measured by astronomical observations. As railroads were developed in the 19th century, however, the rapid proliferation of long-distance fast travel had important consequences for time-keeping: a much greater need for synchronization. Eventually complicated tables came into existence, comparing "railway" time to the local time of every city along each route. This was not synchronization – it was calibration! In New York it would be high noon, while in Columbus it would be 11:25, in Buffalo 11:40, and in Portland, Maine it might be 12:15. Eventually, the telegraph allowed time signals to be sent from

the Harvard College Observatory across the country. When the golden spike was driven for completion of the transcontinental railroad, a telegraph signal was supposed to have been sent through the country, representing the first synchronized moment from coast to coast. However, it was chaos for every town to tell its people at which local time to show up to listen for the signal!

The end result of the confusion was for the railroads to introduce time zones, with the zero of longitude at Greenwich, England. At noon, Eastern Time, on November 18, 1883, trains throughout the United States stood still for various amounts of time in order to synchronize with the new definition of time in each time zone. Most towns agreed to synchronize their times with the standard telegraphed from Harvard, and for the first time the nation was truly synchronized. Eventually other countries agreed to make Greenwich the standard Prime Meridian, thus synchronizing world time. Interestingly, France originally abstained unless Britain agreed to adopt the metric system! (Eventually in 1911, France finally defined its time as Paris Mean Time delayed by 9 minutes and 21 seconds.)

1.1.5 Frequency measurement and scientific discovery

With the world's time now synchronized and standardized, frequency measurement was poised for the rapid developments of the 20th century. Our main theme, then, is the give-and-take between technological development and advances in the measurement of frequency. The discovery of the piezoelectric effect led to the use of quartz crystals in radio broadcasting in the 1920's. The effect was eventually exploited for making a quartz clock in 1928, accurate to within a millisecond per day (10^{-8} in one day). This dramatic improvement over pendulum clocks illustrates two important principles in frequency metrology. First, a higher carrier frequency allows time to be divided into smaller units, and therefore measured with greater precision. This is fundamentally because a counting instrument will always have at least an inaccuracy of one cycle, so as

many periods that can be counted in a given amount of time should be used. Second, a clock oscillator should be well-isolated and insensitive to external perturbations. Quartz met both of these conditions much better than pendulums could. With the new accuracy, more discoveries were made: notably, variations in Earth's rotational frequency could now be measured. Seasonal melting of the ice caps change our planet's rotational frequency in a somewhat predictable way. More irregular fluctuations also occur, due to the fluctuating internal structure of the earth or storms above its surface.

Even the vibrational frequencies of quartz crystals changed as they aged. Clearly there was a need for a frequency standard – some natural frequency that could be reproduced in various measurements, always giving the same result. With the invention of the atomic clock in the 1950's, such a frequency standard was finally realized. Continued improvements to the measurement of the hyperfine splitting of the ground state of Cesium now result in clocks accurate to a fraction of a nanosecond per day (4×10^{-16} per day). Since the adoption of the Cs frequency as a standard in 1967, these incredible devices have defined the second, and are now responsible for improvements in technologies as varied as: keeping the power grid stable, GPS navigation, wireless communication, length metrology, radio astronomy, and tests of relativity.

Several insights come out of this tour of frequency measurements through history: the most obvious is that higher and higher carrier frequencies have always resulted in better accuracy. This principle is a large part of the motivation for optical atomic clocks, as I will soon describe. Another significant principle for the construction of a quality frequency standard is that of isolation: a truly free oscillator relies only on the laws of physics for its natural frequency. Exact duplication is also necessary: whereas pendulums and quartz crystals differ in size and shape from clock to clock, nature's most precise assembly line gives us atoms, with every atom of a given isotope built exactly the same. This eliminates the need for an "artifact" standard such as the cylinder of

platinum-iridium still used to define mass. A particular insight important to this Thesis is that of the clock gearing: as better clocks move to higher frequencies to improve their timing resolution, a greater requirement was placed on the ability to coherently divide these frequencies down to measurable, useful bits of time. A last insight is made clear from the close relationship between frequency metrology and navigation: the ability to synchronize frequency standards is necessary. Consider the ultimate in navigation – the Voyager mission would never have been possible without the precise measurement of its position. This measurement was enabled by accurately synchronized radio telescopes stationed around the world, sending signals back and forth through the 8-hour time delay due to Voyager's great distance.

1.2 Modern motivations for precision timing

Since the discovery of the atomic cesium clock, rapid advances in many applications of such precise time-keeping have been made. Here I describe some of the most important advances to come out of such precision frequency metrology, as well as how further advances in clock technology will continue to further our understanding of nature and deepen our bag of tricks for how we can manipulate nature for our own purposes.

1.2.1 Atomic clocks and the current primary frequency standard

After quartz oscillators and the first atomic clocks were built and compared against astronomical observations as well as against each other, it became clear that the rotation of the earth, or its revolution around the sun, was not as accurate as the frequency standards which exploited the Cesium atom. Therefore, in 1967 the definition of the SI second was changed to be the amount of time in which 9,192,631,770 cycles of the frequency corresponding to the ground-state hyperfine splitting of an unperturbed Cs atom. Since then, refinements have been continually made on the method used to probe this transition. With the rapid development of laser cooling techniques [2] in

the late 1980's, the first successful Cs fountain clock was born. The Cs atoms are now laser-cooled in order to reduce the Doppler shifts that used to pester the hot Cs beams of older atomic clocks. As is the usual case for spectroscopy, the longer the coherent interrogation time between light and matter, the more precisely the frequency can be known. Therefore the cold ball of atoms is launched upwards, and separated-field Ramsey spectroscopy is performed by microwave interrogation of the atoms once on their way up, and again on their way down. The laser cooling before launching the atoms upward also helps keep them from dispersing during their flight. For the best modern fountain clock, the fractional uncertainty in the clock transition is now reduced to 6×10^{-16} [3], and clocks under development are planning to push this down by a factor of three.

1.2.2 Applications of modern microwave atomic frequency standards

The levels of accuracy and precision offered by such clocks have stimulated many technologies and much fundamental scientific investigation. One of the first direct measurements of the predictions of special and general relativity was to measure the effect of high speeds and high altitudes on the timekeeping of atomic clocks [4, 5]. Launching the clocks in rockets up to altitudes of 10,000 km enabled measurements of the clock's period that agreed with the predictions of special relativity (time dilation due to high speed) and general relativity (time getting faster at higher elevation). In other experiments, radio waves were delayed by the relativistic prediction as they passed by the sun [6]. An extremely precise experiment measured the orbital periods of pulsars in a binary pair, and found them to change ever so slightly due to loss of energy into radiating gravity waves [7, 8].

Of great technological importance is the Global Positioning System (GPS) which utilizes 24 satellites, each transmitting precisely synchronized and coded signals. Each satellite contains its own atomic clock, so that a receiver anywhere on earth can compare

the arrival times of the synchronized codes in order to determine the relative delay times, and therefore determine its location. By averaging long enough, the receiver can determine its position to within 1 mm of accuracy, allowing geological measurements of continental drifts and plate tectonics, or allowing navigation programs in cars or ships at sea to effectively operate. It is certainly fascinating that the vibrations of tiny atoms in space can be used to measure the drift of tectonic plates or to navigate a car through the streets of Pittsburgh!

In metrology, frequency is the most accurately measured unit, and as a result, many other units are being measured by way of frequency. An excellent example is the 1983 redefinition of the meter by way of an accurate measure of the speed of light in vacuum [9, 10]. As a result, the meter is now defined as the distance travelled by light in $1/299,792,458^{\text{th}}$ of a second. Amperes and volts are also currently defined in terms of time [11], and the kilogram may be on its way as well [12].

Another entire field that benefits from highly accurate frequency standards is radio astronomy. By placing radio telescopes around the globe and phase-coherently collecting data from them, a synthesized aperture the size of their separation can be used. This very long baseline interferometry (VLBI) demands the data from each telescope be referenced precisely to an atomic clock in order to ensure the individual telescopes' data can be phase-coherently reconstructed. This scheme now works so well that widely-separated radio telescope arrays can achieve angular resolution 250 times better than the best optical telescopes [6]! And of course, as mentioned in the last Section, such arrays are also used for deep-space navigation such as in the Voyager mission.

As a final example of applications of highly stable frequency references, consider the relatively mundane issue of keeping the power grid across the country in phase with itself. The ability to synchronize the 60-Hz waveforms across the country is achieved by use of flexible delays built into the power grid. Atomic clocks are used to ensure these delays are controlled precisely such that whenever power is switched from one part of

the grid to another, its phase is correct. The Northeastern Blackout of 1965 was so widespread in part because of failure of these delays [1]! This kind of synchronization is also becoming increasingly important as high-speed communications networks try to pack more and more data into shorter amounts of time in order to achieve the bandwidths they require.

1.2.3 Applications of improved frequency and timing control

Having explored various technologies that the microwave atomic clocks have enabled, a more clear perspective on the desire for still-better frequency and timing control can be had. Whereas some of the technologies, such as stabilizing the phase across the power grid, might never require improved standards, many of the technologies stand to benefit from tighter synchronization and the improved accuracy of frequency standards.

The first, most obvious example (to a frequency metrologist) of an application directly benefitting from improved frequency standards is that of measuring certain fundamental constants. For instance, the Rydberg constant, Lamb shift, fine-structure constant, Planck's constant relative to the electron mass, and the proton-to-electron mass ratio could all be measured more precisely with improvements to the accuracy of frequency measurement. In addition to the constants themselves, there is the matter of how "constant" they actually are. Certain new theories going beyond the standard model suggest that the constants might actually change over time [13]. If the accuracies of clocks based on various different atoms are improved, a tighter constraint on these possible variations could be achieved. Equivalently, the amount of time it would take to measure at a given level of precision for the time-dependence of the physical constants would decrease. In particular, variations of the fine-structure constant have been suggested by astronomical observations at the 10^{-5} level over the course of the last 10 billion years [14]. With modern spectroscopic techniques in the laboratory, frequency comparisons are constraining the current (as opposed to historical or cosmological) variations

of the fine-structure constant at the highest level of any experiments [15, 16]. Other fundamental tests of physics benefitting from improved frequency metrology include tests of Lorentz invariance [17]. and special and general relativity.

Another example that comes to mind is that of GPS receivers. If fixed receivers were set up on the sides of roads, equipped with superior clocks, then a nearby travelling receiver could communicate with the fixed receiver in order to determine the relative time delay between the signals being received by each from the satellite constellation. With tight enough synchronization, it is conceivable that their relative positions could be determined in real-time to within a cm, perhaps enabling cars to navigate themselves on “auto-pilot” [6].

Somewhat related to the ideas of GPS is the mapping of the geoid. Currently the equipotential surfaces of Earth are mapped out by precision telemetry of satellites in orbit. The geoid is expected to fluctuate at any one location with respect to the Earth’s surface at a level that corresponds to a time dilation (or contraction) of a few parts in 10^{17} . At first glance this might seem to render clocks with even higher accuracies pointless, at least near the surface of Earth. However, one can envision establishing a network of highly accurate (better than 10^{-17}) clocks across the Earth. Comparisons among the clocks would be able to accurately measure the time-changing potential at their various locations. In this way the “choppy ocean” of the equipotential surface could be measured accurately. Of course, this is also an argument for precise and accurate frequency transfer over long distances, since comparisons of clocks separated by tens or hundreds of kilometers would be necessary.

The radio interferometry (VLBI) discussed in the previous Section would also benefit from tighter synchronization technologies [11]. Radio telescopes based on Earth can only be separated by at most the distance of Earth’s diameter. If instead they were located in space, they could be separated by millions of kilometers, greatly enhancing their resolution. Of course, this is only viable if the relative positions of the telescopes

were kept extremely constant, and if phase-coherent data collection were enabled between them. In principle, a frequency reference with fs stability over the course of several seconds would enable accurate ranging between the satellites at the level of a micron. A clever scheme employing the fs frequency comb has already been proposed for how this might work in practice [18].

The establishment of improved frequency standards is only one part of the story, however. There is a whole host of time-domain experiments for which tight timing synchronization of system components is perhaps more important than the long-term stability of the actual frequency reference being employed.

For instance, the precise synchronization of ps and fs lasers has led to the ability to conduct nonlinear surface probing at the single-molecule level with coherent anti-Stokes Raman scattering (CARS) spectroscopy [19]. As the synchronization improves, the SNR available for such systems also increases, leading to the ability to perform chemically-selective imaging of real-time biological dynamics.

The tight synchronization of fs lasers [20, 21] also leads to techniques for nearly-arbitrary waveform generation: phase-coherently combining the output of multiple pulsed sources [22, 23] allows each source to be spectrally shaped, with phase information encoded across its spectrum. The resulting phase-coherent total waveform can have frequency components from across the spectrum if sum and difference frequency generation are employed [24]. These technologies will be discussed in much greater detail in Chapter 5. Coherent control experiments stand to benefit from pulses of light shaped in amplitude and phase to selectively [25] drive chemical reactions [26], molecular vibrations [27], or the nonlinear up-conversion of light into the extreme ultraviolet region [28], to name a few examples.

Rather than considering the synchronization aspects of pulsed experiments, one can also consider the phase coherence of the underlying oscillating optical carrier field. Few-cycle pulses with stabilized carrier-envelope phase have led to new discoveries in

nonlinear light-matter interactions. For instance, controlled attoseconds-long bursts of soft x rays have been generated by amplified phase-stabilized pulses [29]. Another experiment demonstrating sensitivity to the electric field's phase with respect to the intensity envelope involves the control of injected photocurrent in a semiconductor [30].

The combination of fs pulse trains and external optical cavities is another exciting new application of ultrashort pulses that is made possible by precise frequency control methods. For instance, while amplifier systems operating at low repetition rates have been used for high-field experiments, the buildup of a fs pulse train inside an optical cavity now allows exciting high-field experiments to be performed at high repetition rates [31]. By constructing an external cavity geometry such that the pulses undergo a tight intracavity focus, high-harmonic generation of ultraviolet pulse trains has now been observed [31]. Or, instead of performing such high-field physics experiments, absorption spectroscopy can be performed on gas samples inside the cavity. This can be the familiar cavity-ringdown spectroscopy, except that the broad spectrum of the pulse train allows massively paralleled broadband spectroscopy to be performed [32].

All of these applications will benefit from the establishment of more accurate frequency standards, as well as from a higher degree of phase control across the spectra of broadband lasers. However, physicists perhaps expect the most exciting benefits of improved frequency metrology to be the unforeseen discoveries waiting just up ahead!

1.2.4 The need for frequency and timing transfer techniques

As frequency standards of greater accuracy and lower instability continue to be developed, the demand for distribution of their signals is increasing. An important example for this Thesis is illustrated in Chapter 7, where even the development of next-generation standards as discussed in Section 1.3 requires highly stable and accurate transmission of clock signals between various laboratories, since intercomparisons among remotely-located clocks is often necessary.

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 and time standards of national laboratories around the world [33]. 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 or frequency measurement. By averaging for about a day it is possible to reach accuracies of one part in 10^{14} [34]. GPS carrier phase and Two-Way Satellite Time and Frequency Transfer (TWSTFT) techniques can push the frequency transfer instability to the low parts in 10^{15} in one day [35]. However, these techniques are limited by fluctuations in the paths that are not common-mode. They do not provide the short-term stability necessary for synchronization applications, nor are they practical in situations such as the distribution of a frequency reference throughout a linear accelerator facility.

An extremely promising alternative for stable distribution of a frequency reference is transmission over optical fibers. The frequency reference (either optical or microwave) is encoded onto an optical carrier for transmission over a fiber network. Remote users are then able to recover the frequency reference by decoding the received optical signal. 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. In addition, 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 technology, components, and infrastructure also already exist for disseminating frequency references over telecommunication optical fibers, especially in urban environments.

In the microwave domain, signals in the form of amplitude modulation of an optical carrier have been transmitted over various fiber networks [36, 37]. However, a direct transfer of the optical carrier itself [38, 37, 39] should achieve even better stability, relying on the same advantage of high spectral resolution as optical clocks do. Fully coherent transfer of an ultrastable optical carrier signal over a large distance, along with the frequency-comb-based capability of optical coherence transfer across the entire visible spectrum, opens doors to a wide variety of exciting applications, including phase-coherent large arrays of radio telescopes [40], precisely synchronized advanced light sources based on large-scale accelerators [41, 42, 43], and precision optical interferometry over a long distance or encircling a large area.

1.3 Next generation clocks: optical frequency standards

As usual when it comes to building a more accurate clock, the next generation of frequency standards will be enabled by dividing time into even smaller intervals — in other words, higher carrier frequencies will be necessary to improve our measurements. It might seem that the sky is the limit in terms of how high our carrier frequency could be, but for the next generation of clocks we will be content with more than 4 orders of magnitude between the current ~ 10 -GHz standard and the several-hundred-THz regime of optical frequencies. However, in order to take full advantage of this large ratio, the phase of the optical carrier must be measured to just as great a fractional precision as it already is for the microwaves. It is not sufficient simply to count zero-crossings of the optical carrier; precision phase measurement must be achieved.

In the context of an atomic clock, it is simple to see how some relevant parameters influence the clock stability. Consider the simple case of Rabi excitation for a system of N atoms. If each atom ends the Rabi interrogation time T_R in a superposition of ground and excited clock states, then a fluorescent probe laser can be used to cycle many photons off of any atom whose wavefunction collapses into the excited state. If,

on average, N atoms are collapsed into the excited state, a signal-to-noise ratio (SNR) of \sqrt{N} could be achieved. The Rabi interrogation time could in principle take as long as the transition's natural lifetime τ_0 , but in general T_R will be shorter. Therefore the SNR of the measurement, in a bandwidth of 1-Hz, will be $\sqrt{N/2T_R}$. The natural linewidth of the transition divided by the SNR gives the accuracy to which the center of the transition can be found in one second: $\delta\nu \sim \Delta\nu/SNR$. Then the measurement can be repeated again and again for a total amount of interrogation time τ , so the instability should average down like $\tau^{1/2}$, giving:

$$\sigma_y(\tau) = \frac{\delta\nu}{\nu_0} = \frac{1}{\omega_0\sqrt{NT_R\tau}}. \quad (1.1)$$

Equation 1.1 explicitly shows a few important rules about making an accurate atomic clock. The most obvious rule is that moving to a higher carrier frequency, for a given linewidth, will lower the clock instability. Therefore optical transitions are of great importance, since the carrier frequency is some 4 orders of magnitude greater than the current microwave standard near 10 GHz.

Second, it is important to exploit the longest possible coherence time between the probing field and the atomic (or ionic) oscillator. This means several things: the transition lifetime should be as long as possible, the atomic sample should be in an environment such that decoherence is minimized, yet it should be confined somehow so as to allow long probe times, and the probing field should be phase-coherent for as long as can be achieved. Transitions with long lifetimes exist in several atoms and ions that are currently being studied, including Hg^+ , Yb^+ , Yb , Sr^+ , Sr , In^+ , Ca , Mg , and H , to name a few. For these atoms and ions, environments that minimize decoherence while maximizing the possible interrogation time must be carefully constructed. Laser-cooling and trapping must be performed in order to minimize motional effects such as Doppler broadening and so that T_R can be long. For ions, the external and internal degrees of freedom are well-separated, meaning that ion traps can be constructed to hold the ion

in an external trapping field for extremely long periods of time with minimal effect on the clock transition. For neutral atoms, this wasn't always the case; probing the atoms in free fall was an option, but the interrogation time would be limited and motional effects would still take place. Now optical lattice traps can be constructed that still affect the internal structure of the atoms, but in such a way as to perturb the ground and excited states by the same amount, independent of the trapping depth. These optical lattice clocks must operate at the “magic” wavelength so as not to perturb the transition frequency. More details of the lattice clocks will be given in Chapter 7, since in our laboratory we are developing an ^{87}Sr optical lattice clock.

Once the atoms or ions are held in a way that leaves the clock transition unperturbed while allowing long interrogation times, a probe laser must be constructed that remains coherent for as long as possible, in order to take advantage of the long-lifetime traps and transitions. Such stable lasers have indeed been developed [44, 45, 46], and enjoy coherence times longer than 1 s, enabling the coherent interaction time between light and matter to approach 1 s [47]. It is extremely important to stress that the development of such phase-stable local oscillators is of primary importance to the successful construction of any clock.

Another important consideration highlighted by Eqn. 1.1 is the size N of the ensemble. For ion systems, only a few ions can be trapped at a time due to the Coulomb interactions between them. However, for neutral atoms many thousands of atoms can be trapped simultaneously, potentially reducing the instability of such systems. Therefore the density shifts in neutral clocks must be evaluated, although for fermionic isotopes in three-dimensional lattices, the collisional shifts can be extremely small since it could be possible to ensure only one atom occupies each trap site. A last consideration for any atomic clock is that although in theory, the atoms or ions can be held stationary for long probe times in such a way as to leave them unperturbed, in practice there are always stray fields in the laboratory. These can be static or dynamic electric or magnetic

fields, as well as blackbody radiation. Therefore the clock transition should have low sensitivity to external stray fields.

1.4 The case for the comb: precision clockwork

Until now, I have only described how to move to a higher carrier frequency such that the stability and accuracy of higher timing resolution might be exploited. However, I have completely ignored one of the important features of any good clock — namely, the readout system. Unfortunately, the techniques for direct electronic counting of zero crossings that can be used in the microwave regime fail to work in the optical domain. Specifically, copper isn't transparent in the visible! No electronic system can directly keep up with optical oscillation frequencies. Therefore some scheme must be invented that phase-coherently and accurately divides down optical frequencies into the microwave regime.

Division from optical frequencies down to microwaves requires a set of phase-coherent “gears” for the optical clock. Due to the huge gearing ratio of several 10^4 that must be spanned, construction of such gears potentially poses a serious experimental challenge. As with any challenge, a handful of elaborate schemes have actually been constructed in the past in order to measure optical frequencies. These harmonic frequency chains used many successive steps of frequency multiplication and offset-locking to successively link higher-frequency oscillators to lower-frequency ones. Of course, since the many steps in the frequency chain included microwave, terahertz, far infrared, near infrared, and finally visible frequencies, a wide variety of technologies had to be mastered. In addition, a typical frequency chain was extremely frequency-specific, usually targeting a single optical frequency to be measured. This is partly because most lasers in the IR are gas lasers, tunable only within the gain linewidth of some specific Doppler-broadened transition. It is also because operation of such a complex chain is very difficult to maintain even for one frequency. Imagine tuning every component

simultaneously in a frequency chain involving upwards of 15 separate frequency offsets and multiplications. Due to this level of complexity, frequency chains for measuring only a few specific standards such as the 3.39- μm methane standard [48, 49] or the 657-nm Ca clock transition [50, 51] were constructed in several national laboratories around the world.

Other efforts were being developed as alternatives to harmonic frequency chains, and these methods can be classified as “frequency interval division” methods. The distinguishing feature of these techniques is that they rely on difference-frequency synthesis. One method that was often employed is frequency interval bisection. Essentially, if the optical frequency ν and its second harmonic 2ν are available, the frequency gap between them can be divided in two. Then one of the two smaller gaps is divided in two again. Eventually the interval is divided into small enough frequencies that a microwave readout is possible. The interval bisection is performed by making sure that the bisecting frequency, ν_3 , splits the original two frequencies ν_1 and ν_2 by enforcing the relationship $2\nu_3 = \nu_1 + \nu_2$. It is clear that second harmonic generation for the left side of this equation, and sum frequency generation for the right side, can be used along with heterodyne detection and a phase lock. Another method was the optical comb method. An intracavity electro-optic modulator (EOM) could be used to produce a wide spectrum of modulation sidebands on a laser, as long as the round-trip group delay in the cavity corresponded to the period of the EOM. This method has produced 30-THz-wide optical frequency combs [52].

If an optical comb could be broadened to cover the full octave from ν to 2ν , it could be used directly to measure the optical frequency. The modern optical frequency comb now provides this capability. Ti:sapphire lasers were discovered that can be passively mode-locked and subsequently they were demonstrated to form a phase-coherent optical comb as their output. Eventually the microstructure fiber was developed and shown to broaden the spectra from a Ti:sapphire comb to more than an octave, mak-

ing it possible in principle to measure an optical frequency directly. Once experiments showed that pulse-to-pulse phase coherence was preserved during the spectral broadening of the microstructure fiber, all the requirements for a precision optical clockwork were met, and soon the first absolute optical frequencies were being measured with this new technique. The most important aspect of the comb is that a single optical frequency comb can enable a vast spectrum of optical frequencies to be measured in a small, table-top experiment, unlike the harmonic frequency chains that used to measure only a very narrow range of frequencies and occupied several laboratories worth of space. The history of the femtosecond frequency comb's development and discovery is covered in excellent detail by several others, notably including J. Hall and T. Hänsch in References [34, 53, 54]. I will also describe some practical details of the optical comb in Chapter 3.

For now, it is important to communicate the remarkable effect the optical frequency comb has had on the fields of optical frequency metrology and ultrafast science. In the last 8 years, precise phase control over the comb has led to: simple and extremely accurate clockworks for optical clocks, precise synchronization of fs lasers, phase-coherent synthesis of fs lasers, optical frequency synthesizers, XUV comb generation, attosecond-duration x-ray generation, significant improvements to various pump-probe experiments such as CARS microscopy, remote transfer of ultrastable optical frequency references through long distance fiber links, remote synchronization of pulsed oscillators through long fiber links, broadband cavity ringdown spectroscopy, and phase control of multi-pathway quantum interference processes. These are just a handful of examples that I am personally familiar with to some degree, due to the work of many people in our laboratory.

It is the subject of this Thesis to concentrate on the direct application of frequency combs to optical clocks and frequency- and timing-transfer. The phrase “frequency- and timing-transfer” is purposely ambiguous — it could refer to long-distance transfer of pre-

cise timing or accurate frequency references, or it could refer to the transfer of these same references across large spectral gaps. In order to present the less-experienced reader with the tools necessary to begin speaking about characterizing the quality of such transfer, I begin in Chapter 2 by discussing phase noise and timing jitter, instability, and various fundamental and technical limitations to the measurement of optical (relative) phase. Then in Chapter 3 I will delve into much greater detail about the frequency comb, its generation, and its control. After these introductory Chapters, the rest of the Thesis is divided into 5 Chapters on various applications of the frequency comb, beginning with the important issue of detecting the frequency comb's repetition rate harmonics with as much accuracy and purity as possible, in Chapter 4. This ability leads to the precise synchronization of fs lasers, or in other words, the transfer of timing information accurately from one fs oscillator to another, as discussed in Chapter 5. We also demonstrated that with such synchronization capabilities, we could generate tunable-wavelength mid-infrared pulses with fast wavelength switching. Chapters 6 through 8 switch gears to some degree, from timing applications to clock frequency applications. In Chapter 6 I show how the frequency comb can be used in an extremely simple and elegant way to measure absolute frequencies in the infrared. This immediately led to the construction of an optical molecular clock based on the CH_4 transition at $3.39 \mu\text{m}$, representing a tabletop set of optical gears capable of replacing the harmonic frequency chains previously necessary for measuring this reference. In Chapter 7 I provide a detailed description of the octave-spanning femtosecond comb I constructed in order to make absolute frequency measurements of our ^{87}Sr optical lattice clock, including the usefulness of the comb for characterizing the clock laser and comparing it to other microwave and optical clocks in Boulder. Finally, in Chapter 8 I discuss the stable and low-phase-noise distribution of frequency and timing references through long-distance (~ 10 -km scale) fiber networks.

Chapter 2

Characterization and measurement of frequency stability and phase noise

Having described the intuitive principles behind making and distributing a frequency standard in Chapter 1, I will now turn to the more specific details of how to characterize the performance of such systems. In general, since the goal is to produce a waveform that matches a pure sine wave as closely as possible, it might seem that characterization of the three degrees of freedom (amplitude, frequency, and phase) would be straightforward. However, an arsenal of tools has been developed over time in order to measure and distinguish between the various types of noise sources that can arise, allowing the experimentalist greater insight into each system's performance. In this Chapter I will first formally introduce the main characterization tools, and then turn to some practical considerations for measuring the stability and purity of optical and microwave frequency sources. I happily confess that the clear line of reasoning developed in the early sections is due to the excellent discussion found in the paper by Hall and Zhu [55], which every new student of this field should read.

2.1 Characterization methods

The ultimate goal of a frequency standard is to repeatably and reliably produce an oscillation that can be represented mathematically as an ideal sine wave. Obviously this is never completely possible, but very close approximations to pure sine waves can

be generated. Assuming we are dealing with an electromagnetic oscillator, I will write the general form of such a nearly pure sinusoidal field as

$$\vec{E}(t) = \hat{e}E(t)e^{2\pi\nu_0 t + \phi(t)} + c.c.. \quad (2.1)$$

Here, “nearly pure” means I have assumed the polarization vector \hat{e} is constant, and both the amplitude $E(t)$ and phase $\phi(t)$ are slowly varying in time with respect to the carrier frequency ν_0 . In general, for frequency metrology we should be more concerned with the time-varying deviation $\phi(t)$ from the nominal phase $2\pi\nu_0 t$ than with the amplitude fluctuations, although later I will show the importance of processes leading to conversion from amplitude noise to phase noise. For now I will restrict the discussion to characterization of the phase (and therefore frequency) stability of the source.

There are typically two classes of characterization of the performance of a frequency source. The first is to measure the phase noise spectrum, which represents the underlying phase modulation processes over a range of relevant Fourier frequencies surrounding the signal carrier. The phase noise spectrum is an excellent way to think about the coherence of a source; integration over a range of Fourier frequencies reveals the rms phase deviation within that interval. It is also common to make the unit conversion from phase [rad] to time [s] by normalizing to the carrier frequency: the resulting quantity is known as timing jitter spectral density. Whereas the phase noise is usually specified with respect to a particular carrier reference frequency, the timing jitter specifies the total magnitude of perturbations to the carrier’s timing regardless of the carrier frequency. The phase noise and timing jitter spectral densities of the signal intrinsically contain the same information, but have different mathematical representations that are each more or less convenient depending on the physical situation involved. The second method for characterization of a frequency source involves determining how the measured fractional frequency fluctuations of the source vary as a function of the time over which the frequency is averaged. This is commonly expressed using the Allan deviation,

and can be interpreted as the immediate repeatability of the frequency measurement as a function of averaging time. Thorough discussions of the Allan deviation and phase spectral density can be found in many excellent references [55, 56, 57, 58, 59].

2.1.1 Phase noise

The first method for characterizing the performance of a frequency source — directly measuring its phase fluctuations — is especially useful for determining the coherence of a signal. Since only the relative phase between two sources is important, phase noise is determined by measuring the phase fluctuations of a signal with respect to a lower-phase noise reference. The coherence time T_{coh} can then be defined as the average amount of time before the signal acquires 1 rad of phase excursion relative to the reference.

This simple time-domain picture gives an excellent intuitive feel for the significance of phase noise. As usual, taking the Fourier transform of the time-domain picture makes clear more specific information about the relevant physical processes responsible for the phase deviations. These deviations create noise sidebands in the source's spectrum that are spaced from the carrier by the modulation frequency at which the noise occurs, at the same time weakening the power in the carrier by spreading it into the sidebands. It is common to express the phase noise of a frequency source as the power spectral density (PSD) of phase fluctuations, $S_\phi(f)$, which represents the mean square of phase fluctuations $\delta\tilde{\phi}(f)$ at Fourier frequency f from the carrier in a measurement bandwidth of 1 Hz [58]. Explicitly, it is defined as

$$S_\phi(f) \equiv [\delta\tilde{\phi}(f)]^2 \quad \left[\frac{\text{rad}^2}{\text{Hz}} \right]. \quad (2.2)$$

Note that $S_\phi(f)$ includes contributions from both the upper and lower sidebands of the carrier by construction. Another quantity that is often used to specify the phase noise of a frequency source is the single sideband phase noise, $L_\phi(f)$, defined as $\frac{1}{2}S_\phi(f)$ [58].

It is usually written in units of dBc/Hz (dB below the carrier in a 1-Hz measurement bandwidth) and can be logarithmically expressed as

$$L_\phi(f) = 10 \log \left[\frac{1}{2} S_\phi(f) \right] \quad \left[\frac{\text{dBc}}{\text{Hz}} \right]. \quad (2.3)$$

If the phase noise spectrum is known, integration of $S_\phi(f)$ from high Fourier frequency f_h to lower frequency f_l gives the rms phase deviation (squared) in a time $t_l = 1/f_l$. Then if this rms phase noise is 1 rad, we say that $t_l = T_{coh}$. Note that in performing this integration, an assumption is made that noise processes at different Fourier frequencies are independent of each other. Also, usually $\delta\tilde{\phi}(f)$ decreases with increasing f fast enough so that the exact choice of f_h is relatively unimportant.

2.1.2 Timing jitter

In some applications, such as those involving synchronization, a pulsed signal gives a sharp edge useful for synchronizing various system components. In these cases, it can be more useful to express the phase noise of a transmitted signal in terms of its timing jitter, which can be thought of as the deviation of the arrival time of the pulse from its nominal arrival time. 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

$$\delta\tilde{T}(f) = \frac{\delta\tilde{\phi}(f)}{2\pi\nu_0} \quad \left[\frac{\text{s}}{\sqrt{\text{Hz}}} \right] \quad (2.4)$$

The total rms timing jitter, T_{rms} , which must be specified over a bandwidth from f_l to f_h , is then determined by

$$T_{rms} = \sqrt{\int_{f_l}^{f_h} [\delta\tilde{T}(f)]^2 df} \quad [\text{s}] \quad (2.5)$$

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 short timescales, it is important that the high-frequency timing jitter be as small as possible. For an oscillator that is to be phase-locked to the reference signal, only the portion of the reference's jitter spectral density within the locking bandwidth matters. Outside the locking bandwidth, the free-running noise of the oscillator must be responsible for any further jitter. Indeed, most oscillators have intrinsic phase noise spectral densities that roll off at higher Fourier frequencies, whereas electronic detection of the reference signal's phase is typically subject to a white phase noise floor (see Section 2.2) of the detection process. Therefore the locking bandwidth should be limited to the crossover frequency above which the oscillator's intrinsic noise becomes less than the white phase noise floor of the measurement. The total rms timing jitter of the transmitted reference should be specified for the particular locking bandwidth of interest.

2.1.3 Allan deviation

Whereas the phase noise spectrum is especially useful for determining the characteristic frequencies of noise processes, the Allan deviation [60], $\sigma_y(\tau)$, is a convenient method for determining the characteristic timescales of various physical processes. Unfortunately, long-term noise is typically present in many experiments which only reveals itself upon making measurements well-separated in time. Short time-scale measurements may lead to over-optimism about an oscillator's performance, since various systematic effects can alter the frequency over longer times. Instead, Allan's idea is to focus on consecutive frequency measurements — by making a series of adjacent frequency measurements, each obtained by averaging over a period of time τ and separated in time by τ , the difference in frequency between consecutive measurements can be computed. Then plotting the rms frequency difference between adjacent measurements as a function of τ reveals how the oscillator's frequency is fluctuating over the various timescales.

As a result, long-term modulation processes at the different timescales can be cleanly separated and studied.

The averaging time corresponds to the gate time of a 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 introduce some definitions. Taking the real part from Eqn. 2.1 (ignoring amplitude noise and polarization), the signal from the frequency source can be expressed as

$$E(t) = E_0 \cos[2\pi\nu_0 t + \phi(t)]. \quad (2.6)$$

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

$$y(t) = \frac{1}{2\pi\nu_0} \frac{d}{dt} \phi(t) \quad (2.7)$$

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

$$\sigma_y(\tau) \equiv \left\langle \frac{1}{2} [\bar{y}(t + \tau) - \bar{y}(t)]^2 \right\rangle^{1/2} \quad (2.8)$$

where $\langle \rangle$ indicates an infinite time average and \bar{y} represents the time average of $y(t)$ over a period τ [58]. $\sigma_y(\tau)$ can be estimated from a finite set of N consecutive average values of the center frequency, $\bar{\nu}_i$, each averaged over a period τ .

$$\sigma_y(\tau) \approx \left[\frac{1}{2(N-1)\nu_0^2} \sum_{i=1}^{N-1} (\bar{\nu}_i - \bar{\nu}_{i+1})^2 \right]^{1/2} \quad (2.9)$$

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 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 [58]. 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 determine $\sigma_y(\tau)$. 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. At long timescales, the averaging of white noise will eventually meet a limit for most real oscillators, due to systematic shifts in the frequency. Also, a long term drift will eventually show up as a positive linear dependence of the Allan deviation on averaging time. Finally, sinusoidal frequency modulations of sufficient amplitude will result in a bump in the Allan deviation at half the modulation period. All of these well-known behaviors of the Allan deviation allow immediate insight into the physical mechanisms responsible for the long-term behavior of an oscillator.

2.2 Limitations to optoelectronic phase measurement

Phase stabilization of lasers relies on the ability to detect and subsequently cancel, with a high degree of precision, the phase noise of the source relative to the reference. In addition to the intrinsic phase fluctuations we want to measure, in the phase detection process several noise sources of optical or electronic origin can introduce additional phase noise, which limits the degree to which frequency transfer paths can be stabilized. These sources include thermal electronic (Johnson) noise, photon shot noise, amplifier flicker ($1/f$) noise, and amplitude-to-phase noise conversion during the photo-detection, amplification, and phase mixing processes. In general, for low signal strengths thermal noise and amplifier noise will dominate, and for higher signal strengths shot noise and amplitude-to-phase noise conversion become the main limitation. These four noise pro-

cesses are evident in the data of Fig. 2.1(a), which shows a state-of-the-art residual phase noise spectrum for the ~ 10 -GHz harmonic of a femtosecond laser's repetition rate when the optical beam is split onto two photodetectors; the microwave signals are subsequently phase-compared in order to measure the phase spectral density as shown. In this fashion, the excess noise associated with microwave detection, amplification, and mixing of the repetition rate is measured, not the intrinsic phase noise of the original optical pulse train that is common to both photodetectors. Such an experiment is described in much greater detail as the subject of Chapter 4.

2.2.1 Thermal and shot noise

Fundamentally, detection of phase errors is limited by thermal noise and shot noise. Excellent design considerations for laser phase noise measurements are discussed in Ref. [61]; here I introduce the basic concepts and give an example of the limitations imposed by a typical laser phase noise detection system. In a 1-Hz bandwidth, the single sideband phase noise from thermal fluctuations of current through a system's impedance is given by

$$L_{\phi}^{thermal}(f) = \frac{kTR}{2V_0^2} \propto \frac{1}{P_{rf}} \propto \frac{1}{P_{opt}^2} \quad (2.10)$$

where k is Boltzmann's constant, T is the system's temperature, R is the characteristic system impedance, V_0 is the measured rms voltage level of the rf carrier signal (the effects of quantum efficiency are already included in V_0), P_{rf} is the rf carrier power, and P_{opt} is the average optical power incident on the photodetector. For $R = 50 \Omega$ and at room-temperature, this phase noise power spectral density limit arising from thermal noise can be written as [62]

$$L_{\phi}^{thermal}(f) = (-177 - P_{rf}) \left[\frac{\text{dBc}}{\text{Hz}} \right] \quad (2.11)$$

where P_{rf} is given in units of dBm. This is to be compared with the phase noise contribution from shot noise which, for a time-independent optical power incident on a

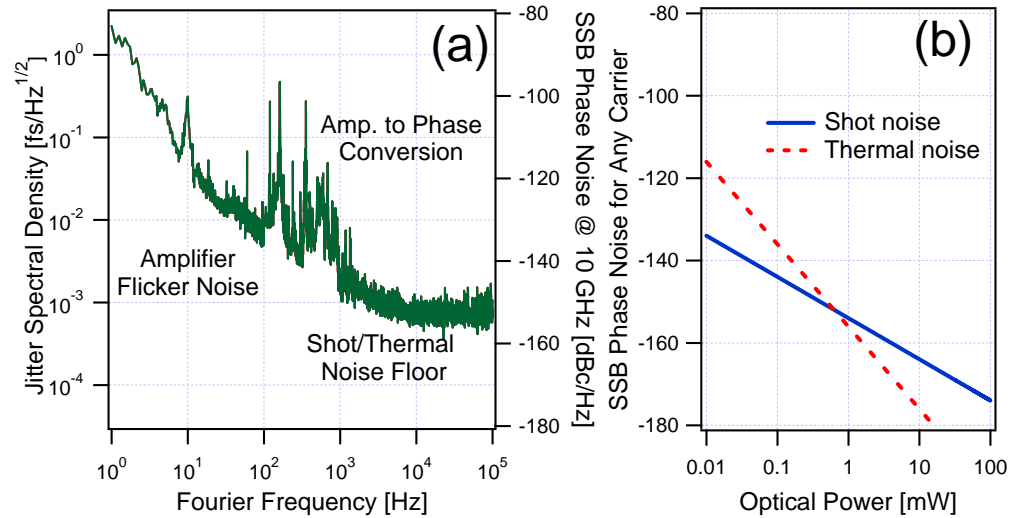


Figure 2.1: (Color online) (a) Phase noise spectrum for ~ 2 mW of optical power from a femtosecond laser split equally onto identical photodetectors, and subsequently mixed to measure the excess phase noise in detection of the 10-GHz harmonics of the laser's repetition rate. For $f < 1$ kHz, amplifier flicker noise behaves as $1/f$. The sharp features between 100 Hz and 1 kHz are due to amplitude-to-phase noise conversion. Above 10 kHz the shot/thermal noise floor is reached. (b) The shot noise and thermal noise limits are plotted as functions of the incident optical power. Note that shot noise and thermal noise introduce a white phase noise spectral density with magnitude independent of the carrier frequency; thus their relative contributions to frequency instability and timing jitter decrease with an increasing carrier frequency.

photodetector, is given by

$$L_{\phi}^{shot}(f) = \frac{ei_{avg}R}{P_{rf}} \propto \frac{P_{opt}}{P_{rf}} \propto \frac{1}{P_{opt}} \quad (2.12)$$

in a 1-Hz bandwidth, where e is the electronic charge and i_{avg} is the average DC photocurrent generated by P_{opt} . When pulses are incident on the photodetector, the assumption of time-independent power is invalid [63], but Eqn. (2.12) still yields results that are empirically valid for time averages of the photocurrent and shot noise power [64]. Note that since P_{rf} scales as the square of P_{opt} , L_{ϕ}^{shot} decreases as the inverse of P_{opt} , whereas $L_{\phi}^{thermal}$ decreases as the inverse **square** of P_{opt} . This results in the shot noise floor overtaking the thermal noise floor above a certain optical power, as shown in Fig. 2.1(b) for a typical GaAs photodetector with a response of 0.4 mA/mW near a wavelength of 800 nm. Note that the fundamental phase noise limitations imposed by shot noise and thermal noise do not depend on the carrier frequency of the microwave signal, and therefore the resulting instability and timing jitter limitations can be improved by use of a higher microwave carrier frequency. For the same reason, optical-based phase detection (as discussed in Section 2.3) provides much higher sensitivity.

As an example of the fundamental limitations imposed by shot noise and thermal noise, consider 1 mW of total average optical power (from two beams for the purpose of forming a heterodyne beat) incident on a typical photodetector with a sensitivity of 0.4 mA/mW, and a 50-ohm system impedance. The rf power obtained from the optical beat frequency could have $P_{rf} = -21$ dBm, which leads to single-sideband phase noise contributions of -156 dBc/Hz for thermal noise and -154 dBc/Hz for shot noise. These phase noise spectra both contribute to the shot/thermal phase noise floor shown in Fig. 2.1(a), for Fourier frequencies above 10 kHz. As discussed above, it is desirable to use high rf carrier frequencies whenever a system is limited by shot or thermal noise. For a 10-GHz carrier, -155 dBc/Hz represents a single sideband jitter spectral density of 2×10^{-4} fs/ $\sqrt{\text{Hz}}$, which leads to an integrated timing jitter of 0.04 fs in a 10-kHz

bandwidth (the total contribution is twice that from a single sideband). For a carrier subject only to white phase noise, written as the constant power spectral density S_ϕ , the measured instability can be expressed [57, 65] as Allan deviation

$$\sigma_y(\tau) = \frac{\sqrt{3S_\phi\Delta f}}{2\pi\nu_0\tau} \quad (2.13)$$

where Δf is the bandwidth of the measurement and ν_0 is the carrier frequency. For a bandwidth of 10 kHz, our example system would be limited by thermal noise and shot noise at approximately the $7 \times 10^{-17}\tau^{-1}$ level. It is important to remember that this fractional frequency instability is relative to the 10-GHz rf signal: in the case that the signal is derived from an optical heterodyne beat between, for example, two 100-THz optical signals, then the optical frequencies enjoy a white-phase-noise-limited relative fractional instability of 10^4 lower than that of the 10-GHz rf beat. Finally, note that in the case of laser stabilization, mechanical perturbations to the optical path will influence the fractional instability and jitter spectral density at a level independent of the carrier frequency, whereas thermal noise and shot noise limits, as well as any other limits from poor signal-to-noise ratio (SNR), can always be improved by use of a higher optical carrier frequency and a higher microwave heterodyne frequency.

2.2.2 Electronic flicker noise and amplitude-to-phase conversion

In addition to the fundamental limitations discussed so far, two important technical limitations remain: excess noise in the amplifiers or mixer [66], and amplitude-to-phase conversion [67]. Ultimately these problems stem from the photodetection process: typical photodiodes yield weak rf powers and are susceptible to saturation and pulse distortion when driven with higher ($>$ few mW) optical powers [68, 69].

The weak rf power available means that electronic amplification must be applied in order to drive subsequent rf circuitry for phase detection without accumulating a significant amount of additional noise. Unfortunately, rf amplifiers are susceptible to

$1/f$ flicker noise due to up-conversion of DC bias noise via nonlinear processes within the amplifiers [66]. Typical flicker noise is seen in the data shown in Fig. 2.1(a), as the smooth $1/f$ -shaped portion of the data for Fourier frequencies below 1 kHz. Attempts have been made to reduce this up-conversion by feedforward [70] or feedback [71] with ~ 20 dB reductions in the flicker noise for 1-GHz amplifiers. Rather than using feedback to suppress the flicker noise up-conversion, microwave interferometric techniques [72] (carrier suppression) have been used to operate amplifiers in the weak signal regime by interferometrically cancelling the carrier. Without a strong carrier driving the amplifier into nonlinearity, flicker noise from up-conversion has been reduced by as much as 24 dB at 10-GHz carrier frequencies [73]. Applying this technique to a system employing typical photodetection of an optical frequency comb, microwave frequencies were synthesized from optical frequencies with -110 dBc/Hz of phase noise at a 1-Hz offset from the 10-GHz carrier [64]. Excellent discussion of design considerations for these microwave interferometric techniques can be found in Ref. [74].

The final technical limitation I discuss here concerns applying higher optical power to the photodetector. Increased incident power is desirable since shot noise and thermal noise can be reduced, and the need for amplifiers can be eliminated, when higher microwave signal power becomes available. However, typical photodetectors exhibit a power-dependent phase shift between the incident optical light and the generated microwave power, leading to significant amplitude-to-phase conversion for higher incident optical powers. As seen between 100 Hz and 1 kHz in Fig. 2.1, the sharp features of the optical beam's power spectrum have been converted to sharp features in the phase spectrum of the homodyned microwave signals. This nonlinearity may be in part from saturation of the photodetector, which causes a high space-charge buildup in the depleted region to affect the velocity of photo-generated charge carriers by screening of the applied bias field [75]. Also, the temporal response of the photodetector can be highly spatially dependent, which allows beam-pointing fluctuations to affect the

microwave output phase at levels that limit system performance [76]. For femtosecond pulse applications, typical power-to-phase conversion coefficients are on the order of several picoseconds of time delay per mW change of average incident power. Stabilization of the amplitude fluctuations of the incident optical beam has reduced the excess phase noise by more than 20 dB for Fourier frequencies below 10 Hz [77]. By carefully eliminating environmental perturbations that cause amplitude fluctuations at the photodetector, optical-to-microwave conversion exhibiting an excess instability of $< 7 \times 10^{-16}$ at 1-s and phase noise below -98 dBc/Hz at 1-Hz offset from the 10-GHz carrier has been achieved [78]. It should be noted that a disadvantage of photodetecting a pulse train for phase stabilization (as opposed to heterodyning two cw optical carriers) is that, due to the higher peak powers present, such an approach is more susceptible to amplitude-to-phase conversion in the photodetection process than cw techniques would be. Finally, I have worked with a photodetector configured in a receiving circuit resonant with the target microwave carrier frequency, and it offers enhanced signal-to-noise by improving impedance matching and reducing noise bandwidth [79]. Details of this technique are discussed in Section 4.3.

2.3 Techniques for optical timing measurement

In Chapter 3 I will show how the time-domain picture of a mode-locked laser's output is of a cw optical carrier frequency modulated by a pulsed intensity envelope. As a result, an experimentalist can choose to access either information about the timing of the intensity envelope's steep rising edge (optical-phase-insensitive information), or information regarding the optical carrier's underlying phase (optical-phase-sensitive information).

2.3.1 Optical-phase-insensitive techniques

For a pulse train, optical-phase-insensitive techniques can be thought of as methods for detecting the timing separation between pulses, or in other words, the repetition rate of the pulses (see Chapter 3 for details about pulse trains). A simple microwave technique (homodyne mixing scheme) for accessing the repetition frequency and its associated microwave phase is to detect the pulse train on a photodetector, which will respond at the repetition rate and any harmonics within its detection bandwidth. A specific harmonic of the repetition rate can be electronically filtered and its phase can then be measured against some other reference microwave frequency by homodyne mixing of the two signals. Therefore the timing of the intensity envelope (repetition frequency), can be measured or controlled in the microwave domain. The reference frequency can be a stable microwave signal from some microwave oscillator such as quartz, a Cesium clock, Hydrogen maser, or sapphire-loaded cavity oscillator, or it could be another pulse train's repetition frequency harmonic. In fact, I have already been discussing this last case in Section 2.2 with the data shown in Fig. 2.1, where the pulse train was split onto two such photodetectors for the purposes of measuring the technical noise associated with photodetection and subsequent electronic filtering, amplification, and mixing. More details of this type of measurement are given in Chapter 4, where I show that the fundamental limitations described in Section 2.2 become extremely important. As a result of these noise limitations, higher harmonics (several GHz or higher) of the repetition frequency should be used.

Another method (fringe-unresolved cross correlation) for accessing the pulse train's intensity envelope timing is useful when more than one pulse train is involved. Then the two pulse trains can be focused into a nonlinear crystal for the purposes of sum-frequency generation (SFG) in a non-collinear geometry where the two beams cross through each other as they pass through the focus. Light at the sum frequency will be

generated as a third beam exiting the crystal between the two second-harmonic beams from each original pulse train. The intensity of the SFG 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 pulsing at the repetition frequency of the lasers, the amplitude of the SFG 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, the amplitude of the SFG signal is maximally sensitive to timing jitter between the pulse trains. Introducing a known amount of delay between the two trains while monitoring the SFG intensity allows its calibration, and variations in its amplitude provide a very sensitive method for measuring timing jitter between the pulse trains. This technique is illustrated in Section 5.1.3 where we used it for the purposes of characterizing the timing jitter between two independent mode-locked lasers when their repetition frequencies were phase-locked using the microwave homodyne mixing scheme. The steep slope of pulses that have a duration on the order of 10–100 fs makes the autocorrelation technique for detecting timing jitter much more sensitive than homodyne rf mixing schemes limited by the noise processes discussed in Section 2.2, although the dynamic range of the timing jitter that can be measured is limited to the order of the pulse duration. In fact, the steep slope arises from the contributions of $> 10^5$ modes of the frequency comb constructively interfering with each other to create the pulse. Therefore the cross-correlation technique can be viewed as similar to the homodyne mixing technique, but taking place at the 100,000th harmonic of the repetition frequency. However, intensity fluctuations of either input laser can lead to significant amplitude-to-phase conversion of such an autocorrelator; this problem has been successfully mitigated by implementation of a balanced cross-correlator in Ref. [80].

2.3.2 Optical-phase-sensitive techniques

More-powerful techniques are capable of measuring and controlling the optical phase of either cw lasers or the underlying carrier phase of a pulse train. The most direct way (heterodyne technique) to measure the relative optical phase between two cw lasers is to form a heterodyne beat between them by interfering them on a beamsplitter. The combined beam's electric field will have an intensity envelope that oscillates at the difference frequency between the two input optical carriers. If the heterodyne difference frequency lies within the bandwidth of a photodetector, then the photocurrent will respond at that rf or microwave difference frequency. In addition, the phase of the beat frequency depends directly on the phase difference between the input lasers. Therefore, phase-stabilization of the beat frequency to a quiet rf or microwave reference (by a homodyne phase-detection between the two) results in the same degree of relative phase stabilization of the two optical carriers. In addition, the heterodyne beat can be chosen at some high rf or microwave frequency to avoid typical noise present at lower frequencies. This heterodyne optical technique followed up with homodyne electronic phase detection is also a very powerful tool for phase-stabilizing a mode-locked laser to a cw optical reference if the frequency comb's offset frequency is independently stabilized (see Chapter 3 for details about such clockwork).

Finally, when two pulse trains are involved, a fringe-resolved cross-correlation can be formed. The two pulse trains are now focused into the nonlinear crystal in a collinear geometry, and the resulting SFG signal oscillates as a function of delay between the two pulses at a frequency equal to the pulses' optical carrier frequency [81]. Thus adopting the notation in Eqn. (2.4), ν_0 is on the order of ~ 200 THz leading to a substantially increased sensitivity for the jitter measurement [82]. A further advantage of this measurement method over the non-collinear cross-correlation is that the two pulses need not have perfect chirp compensation as the overall slope of the correlation

envelope is unimportant. The only disadvantage of the method is its dynamic range, since more than a few fs of integrated jitter will wash out the fringes.

In summary of the available techniques for phase detection, greater optical power is necessary for lowering thermal and shot noise, as well as for eliminating the need for electronic amplification and its associated flicker noise. However, photodetectors that are able to handle power levels > 10 mW must also perform with high linearity [83] in order to avoid amplitude-to-phase conversion; considerations for the design of such high-power photodetectors are complex [84]. The need for highly linear, high-current photodetectors is clear, since current performance levels for optical-to-microwave conversion (see Chapter 4) are rapidly approaching the fundamental shot noise limit associated with milliwatt-level optical powers. When future optical clocks deliver the promised instability of $< 1 \times 10^{-16}$ at 1 s, we will urgently need such advances to enable accurate low-phase-noise conversion into the microwave domain. Finally, whenever possible, higher order harmonics of a frequency comb's microwave repetition frequency, or even optical-phase-sensitive detection should be employed, since the long lever arm provided by using extremely high (optical) carrier frequencies drastically reduces the relative contributions of the technical and fundamental electronic noises to the total jitter spectral density of a signal.

Armed with the ideas of phase noise measurement and instability as characterized by the phase noise spectral density $S_\phi(f)$ and the Allan deviation $\sigma_y(\tau)$, we have the proper tools to discuss the performance and limitations of the optical frequency comb, as realized for several applications in Chapters 4 through 8. First, in Chapter 3 I will discuss the basic principles of the optical frequency comb generated by a mode-locked laser.

Chapter 3

Fundamentals of Stabilizing an Optical Frequency Comb

An ideal optical frequency comb would consist of a series of cw frequency modes, spaced by an exact constant frequency interval, with each mode taking the form of a spectral delta function. In addition, a fixed phase would exist between any two modes of the comb in order to provide temporal coherence across the spectrum. The most useful comb's spectrum would extend from microwave frequencies all the way into the UV regime and beyond, and the frequency spacing between comb modes would be required to lie within a reasonable photodetector's bandwidth so that heterodyne techniques would be possible against other optical frequency sources. Appropriate control of the phase and amplitude of each mode, as well as the comb spacing would in principle allow nearly arbitrary electromagnetic waveform generation. Of course, what I have described would be infinite coherence time for each mode, and is not possible in the presence of typical acoustic noise resulting in decoherence: a mere 300-nm physical displacement of some system component results in a π phase shift for an optical frequency. Nevertheless, the mode-locked laser's utility as an optical frequency comb was discovered at the recent turn of the century, immediately followed by rapid advances in its control and synthesis, leading us closer to the ultimate in control of the electromagnetic spectrum.

For the specific application of optical clocks and timing transfer, the mode-locked laser allows a phase-coherent connection to be made between the visible and microwave portions of the spectrum. In other words, it provides in a single step the long-sought-

after frequency division between optical and microwave frequencies, making possible the relatively direct measurement of optical frequencies. An important aspect of the mode-locked laser is that it allows these optical frequency measurements to be made not only at any single narrow frequency region, but across the spectrum. This great utility is responsible for much of the recent rapid advance in optical frequency metrology.

In this Chapter I will first describe the concept of mode-locking, specifically as it occurs in the modern workhorse of frequency combs, the mode-locked Ti:sapphire laser. Both the frequency- and time-domain pictures of the comb give insight as to how it can be controlled, and indeed the active stabilization of the comb relative to some reference is necessary in order to maximize its coherence. Therefore I will describe the various controls available to an experimentalist who wishes to influence the comb's two degrees of freedom. I would also like to point out that Refs. [85, 86, 87, 88] are excellent introductions to the stabilization of frequency combs.

3.1 Mode-locked Ti:sapphire laser

Mode-locking is, in general, the effect of any physical process which causes the many longitudinal modes of a laser cavity to enjoy a fixed phase relationship. As a result, the modes all constructively interfere to a significant degree at only one location in the cavity. This location is the pulse which propagates around the cavity with a characteristic group velocity. The resulting output of a mode-locked laser is a train of pulses, equally spaced in time by the round-trip time of propagation through the cavity. It is easiest to see in the time domain, that if some process in the cavity results in favorable gain for higher intensities, and if the process responds at sufficiently short timescales, then the dynamics of the laser cavity will lead to short, intense pulse propagation in the cavity. Many types of mode-locked lasers exist, with various mechanisms responsible for this mode-locking. These include low phase- and amplitude-noise mode-locked diode lasers at various wavelengths [89, 90], mode-locked Er^{+3} fiber lasers

at 1550 nm [91, 92, 93, 94, 95], and Cr:forsterite lasers near 1300 nm. For the measurements I describe in this Thesis, various forms of the Ti:sapphire mode-locked laser will be employed. Therefore I will restrict my discussion to the Kerr-lens-mode-locked Ti:sapphire laser and its control.

Figure 3.1 shows a typical such Ti:sapphire laser. Pumped by a solid-state laser at 532-nm, the titanium-doped sapphire crystal has an extremely broad (700- to 1000-nm) gain bandwidth, which in principle supports lasing in any of the cavity modes within this bandwidth. For this reason, cw Ti:sapphire lasers have been constructed allowing broad tunability and high power across the near-infrared portion of the spectrum. However, it is also possible to generate very short pulses at the output of a Ti:sapphire laser through careful design of the cavity's dispersion and the process of passive Kerr-lens mode-locking (KLM), which works as follows.




Figure 3.1: A Kerr-lens-mode-locked Ti:sapphire laser. Several watts of 532-nm pump light excite the Ti:sapphire gain medium. By misaligning the curved mirrors slightly from their positions for optimal cw lasing, the soft aperture of the pumping region combined with the Kerr-lensing induce self-mode-locking. Another requirement is for the net cavity dispersion to be close to zero (slightly negative). The prism pair accomplishes this, and at the same time allows tuning of the output spectral bandwidth, shape, and center. Typically the laser can produce transform-limited pulses of < 15 -nm duration, but before exiting the cavity they are broadened by dispersion of the crystal, a curved mirror, and the material of the output coupler.

The nonlinear Kerr-lens results in an intensity-dependent beam waist and waist location for a Gaussian pulse propagating through the cavity. This comes about because the transverse mode is Gaussian, and so the nonlinear index as seen by the pulse also exhibits a Gaussian profile in the transverse direction. By alignment of the cavity such

that a higher intensity pulse more favorably overlaps with the pumping laser than a lower intensity pulse, the more intense pulse will enjoy a favorable gain. A kind of “soft” aperturing of the beam is typically achieved by misaligning the curved mirrors from optimal cw laser alignment and into better alignment for high-intensity pulses. The process of pulsed operation is then initiated by introducing a brief intensity spike into the cw-laser (by hitting it!), such that the intensity spike enjoys more gain than the cw lasing mode. This initial process runs away until the pulse becomes so short that dispersion in the cavity can no longer keep the pulse together, or the gain bandwidth of the laser has been exploited, or more generally, the interplay of both effects [96]. Careful design of the intracavity dispersion has resulted in lasers with pulse durations below 5 fs. Note that the extremely fast timescale of the Kerr nonlinearity (femtoseconds) ultimately allows the cavity to discriminate between short, intense pulses and longer, less intense pulses.

For the laser shown in Fig. 3.1, a key concept is the intracavity dispersion control offered by the prism pair. The second- and third-order group-velocity dispersion can be controlled by choice of prism material, the separation between the prisms, and the degree of insertion of each prism into the propagating beam. This control over dispersion not only acts to keep the pulse together during a round trip through the cavity; it also can be used to influence the output spectrum of the laser. This fact will become important in the experiment described in Chapter 5.2, where the tunability of the output spectra of two lasers is used to control the spectrum of the infrared pulses they generate in a nonlinear crystal. It should also be mentioned that use of negative-dispersion mirrors is possible for designing lasers without the need for a prism pair; this allows the cavity to be made shorter and higher repetition frequencies to be achieved.

Another minor, yet important detail of the laser is its curved mirrors, which allow tight focusing into the crystal to maximize the nonlinear Kerr-lens effect. Since the crystal must be oriented at Brewster’s angle to minimize loss in the cavity, the

curved mirrors allow for compensation of the resulting astigmatism by reflecting from them at the proper angle of incidence.

The robust performance of current KLM Ti:sapphire lasers has ultimately led to their widespread use as the workhorse of optical frequency measurements. Indeed, these systems now can have spectra spanning over more than an octave of bandwidth, which we will see is an important property for phase stabilization of the pulses. The average powers available from such lasers are on the order of 1 W, with repetition frequencies ranging from 10 MHz to several GHz. They can be actively stabilized for periods of several days if need be, and are simple and quick to build as well as to operate.

3.2 Frequency and time domain pictures


Shortly after the invention of the laser, the optics community seems to have branched along two distinct lines of research. In one camp, researchers used their ingenuity to devise clever methods for phase and frequency control of cw lasers, leading to a variety of precision tests of fundamental physics with narrow-line laser spectroscopy. In the other camp, the ability to create increasingly short pulses of light led to pump-probe experiments resolving the nature of matter at extremely short timescales, or at extremely high field strengths. With the development of the frequency comb in the late 1990's, a shift of paradigm in both camps brought them together, resulting in a beautiful connection between the time- and frequency-domains and the resulting rapid progress in control of light fields. Before this paradigm shift, it was practically inconceivable that such ultrafast pulses in time would be capable of serving to make high-resolution spectroscopy. But now, working with ultrafast lasers in the frequency domain has made it possible to apply the tools of laser stabilization to such sources. Therefore it is important to make explicit the connection between time- and frequency-domains for the output of a mode-locked laser. A very insightful discussion of the mode-locked laser's output can be found in Ref. [97], with two alternative methods for derivation of

the frequency comb structure. Here, I give a slightly more simplistic treatment.

Start by considering the time-domain output of a single ultrashort pulse, shown as the first pulse in Fig. 3.2(a). The pulse consists of an envelope, which if squared would give the cycle-averaged intensity, and an underlying carrier oscillating at some average frequency ν_c . Clearly the Fourier transform of such a pulse with a Gaussian envelope results in a Gaussian spectrum, centered at the carrier frequency, shown as the envelope in the frequency domain picture of Fig. 3.2(b). Then the temporal and spectral widths must scale as inverses of each other in order to satisfy the Fourier uncertainty condition: shorter pulses can only be created through the interference of a broader spectrum of frequencies. If we are motivated to achieve a 15-fs pulse, for example, then our spectrum will have to be roughly 30 THz wide if centered at 300 THz. Now consider propagating the various frequency components making up our single pulse through a dispersive medium. In the case of normal material dispersion (the usual case for materials at Ti:sapphire wavelengths), the higher frequency portion of the spectrum (blue) will be delayed relative to the lower frequency (red) portion. As a result the envelope of the pulse will be broadened, as well as its peak power reduced. In fact, the broader the spectrum, the more exaggerated will be this temporal broadening due to “chirping” the pulse. In some cases where dispersion can not be set to zero, starting with a longer initial pulse will therefore lead to a shorter chirped pulse after the dispersion than if one had started with a shorter initial pulse.

Now, add the complexity of additional subsequent pulses from our laser, forming a train of pulses equally spaced in time. This situation is shown in Fig. 3.2(a) for the time-domain, where each pulse has a temporal full-width-half-maximum (FWHM) of τ_{FWHM} , and the time between successive pulses is τ_{rep} , which is the inverse of the repetition frequency f_{rep} for the pulse train. In the ideal case that the pulse train lasts for an infinite time, and the electric field repeats itself exactly every τ_{rep} , the frequency domain picture would consist of a spectral envelope identical to the single

(a) Time domain



(b) Frequency domain




Figure 3.2: (a) Time-domain picture of the mode-locked laser's output. Each pulse has duration τ_{FWHM} . The pulses are separated in time by τ_{rep} . From pulse to pulse, the underlying carrier shifts by the phase $\Delta\phi_{ceo}$ which is an amount of time $\Delta\phi_{ceo}/2\pi\nu_c$. (b) In the frequency domain, the spectral envelope corresponds to the Fourier transform of each pulse's temporal envelope. Modulating the spectral envelope are the narrow modes of the optical frequency comb, separated by f_{rep} and offset from zero frequency by f_0 .

pulse's spectrum, but modulated by a comb of infinitely narrow frequencies, each given as an integer multiple of f_{rep} . The fundamental limit to the spectral width of any one comb line is therefore governed by the duration of the entire pulse train, while the technical limit is set by fluctuations of τ_{rep} (and of f_0 as discussed in the next paragraph). It is less simple to discuss the limits to the pulse duration and spectral bandwidth, though contributing effects include the gain bandwidth of the laser and the dispersion which the Kerr-lens mode-locking must overcome.

The final degree of complexity for a real mode-locked laser involves the net dispersion of the cavity. Although the dispersion of the air, mirror coatings, prism material, and the linear material dispersion of the crystal should all be adjusted for a net dispersion close to zero, the pulses actually propagate in a medium with slightly negative net dispersion. Here, the time-domain helps picture the effect of dispersion. In general, the average phase velocity v_p is greater than the average group velocity v_g for a dispersive cavity. Since the envelope of the pulse propagates at v_g while the carrier propagates at v_p , the carrier is constantly slipping with respect to the frame of reference of the envelope. After an amount of time τ_{rep} , the intracavity pulse envelope will have propagated from the laser's output coupler, through one round trip of the cavity, and back to the output coupler, a total physical distance of L_c at average velocity v_g . If the carrier keeps up with the envelope (zero mismatch between phase and group velocity), it picks up a phase corresponding to some integer multiple m of π . However, due to dispersion it can pick up additional phase, resulting in the expression:

$$2\pi\nu_c\tau_{rep} = 2\pi m + \Delta\phi_{ceo}, \quad (3.1)$$

where an intuitive feel for $\Delta\phi_{ceo}$ is given by

$$\Delta\phi_{ceo} = 2\pi\nu_c L_c \left(\frac{1}{v_g} - \frac{1}{v_p} \right). \quad (3.2)$$

Since the carrier frequency ν_c picks up an integer multiple of 2π plus $\Delta\phi_{ceo}$ in phase, the Kerr lens forces every frequency component of the laser ν_n must pick up its own

integer multiple n of 2π plus $\Delta\phi_{ceo}$. Ultimately, the degree to which the Kerr lens forces every mode to enjoy this fixed phase relationship will be the degree to which our comb's accuracy is established. The resulting phase relationship can be written as

$$2\pi\nu_n\tau_{rep} = 2\pi n + \Delta\phi_{ceo} \quad (3.3)$$

or, rearranging,

$$\nu_n = nf_{rep} + f_{rep} \frac{\Delta\phi_{ceo}}{2\pi}. \quad (3.4)$$

So the dispersion which allowed the carrier-envelope phase slip (phase and group velocity mismatch) ultimately results in the frequency components of the comb being shifted by an amount f_0 , called the offset frequency, where we make the definition

$$f_0 \equiv f_{rep} \frac{\Delta\phi_{ceo}}{2\pi}. \quad (3.5)$$

Then the entire frequency comb, including the effect of dispersion, can be written in terms of just two degrees of freedom, both of which are microwave frequencies, as:

$$\nu_n = nf_{rep} + f_0. \quad (3.6)$$

Now the time- and frequency-domain correspondence shown in Fig. 3.2 is complete. In the time domain, the pulses are separated by the round-trip time of the pulse envelope, corresponding to the frequency separation of adjacent comb lines in the frequency domain. The carrier-envelope phase slip $\Delta\phi_{ceo}$ from pulse to pulse in the time domain corresponds to an offset of the entire comb from zero by the amount f_0 . The temporal width of each pulse's envelope in the time domain scales as the inverse of the spectral envelope's bandwidth. The frequency of any component of the comb (ν_n) is given by just the two microwave frequencies f_{rep} and f_0 . It will be the topic of the next Section to show how the comb's two degrees of freedom can be stabilized. For now, I emphasize that this amazingly simple relationship between every single optical frequency of the femtosecond laser and just two microwave frequencies is entirely due to the fixed phase

relationship established between modes of the comb from the passive Kerr lens self-mode-locking operation. This property of Kerr-lens–mode-locking is indeed responsible for the success of the Ti:sapphire femtosecond comb as a phase-coherent set of gears relating optical frequencies to each other as well as to the microwave domain.

3.3 Servo control of frequency comb

It is clear from Eqn. 3.4 that appropriate control of just the two degrees of freedom f_{rep} and f_0 will be sufficient to control the entire comb. In principle, fixing any two frequencies of the comb will achieve this goal; for instance, if f_{rep} is stabilized while another optical mode of the comb is locked to a reference laser, or if f_0 and an optical mode are simultaneously stabilized, or if two optical modes are fixed relative to absolute references, the full comb will be stable. Of course, choosing to fix two widely separated frequencies gives the system a longer lever arm on noise processes associated with detection and locking of the frequencies. These different methods each have their advantages and various applications are best-suited for different techniques.

3.3.1 Detection and stabilization of the offset frequency

For the most general applications of the comb, direct stabilization of f_0 is desired. It can be measured using the self-referencing technique [98] shown schematically in Fig. 3.3, provided the comb spans an entire octave. The heterodyne beat between the frequency-doubled infrared portion of the comb and the high-frequency portion provides the microwave frequency f_0 . This scheme for self-referencing is often referred to as the f -to- $2f$ scheme. The octave-wide spectrum needed for this scheme can be achieved either by use of highly nonlinear microstructure fiber [99] or by femtosecond lasers which directly emit octave-spanning spectra [100, 101] (such an octave-spanning comb will be discussed in Chapter 7). It is also possible to stabilize f_0 with a $2f$ -to- $3f$ scheme [102, 103], which only requires $2/3$ of an octave, or more-complicated schemes




Figure 3.3: 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 . As usually implemented, many thousand modes in each of the two (fundamental and frequency-doubled) beams beat against each other on a photodetector to cause the photodetector's photocurrent to oscillate at frequencies $|m f_{rep} \pm f_0|$ (for integer $m \geq 0$) within the detector's bandwidth.

involving various frequency ratios [104]. Once f_0 is detected, it can be stabilized to a frequency reference by controlling the difference of the group and phase velocities in the laser cavity using the pump laser's intensity [105, 106, 107] or by tilting the high-reflector mirror in the spatially dispersed arm of a cavity with prisms, using differentially-driven PZT's to tilt the mirror. Also, difference-frequency schemes allow an infrared comb to be formed that has no offset frequency; only the comb spacing need be stabilized [108]. Such an offset-free clockwork will be described in Chapter 6.

3.3.2 Stabilization of the repetition frequency

The stabilization of the comb to an optical frequency reference is completed by stabilizing the heterodyne beat between one of the comb lines and the optical reference, while f_0 is independently stabilized. Usually a PZT-actuated mirror in the mode-locked laser cavity, which adjusts the laser repetition frequency, is used for stabilizing the optical frequency of the comb line. Each beat frequency to be stabilized is mixed in a homodyne fashion with a stable and pure rf source in order to form the dc voltage error signal used for feedback. With every comb line now effectively phase locked to the optical frequency reference, the comb provides a phase coherent link from the optical reference to any other optical frequency within the comb bandwidth. This represents a very powerful “universal gearing” allowing any optical frequencies to be accurately compared with just one relatively simple setup. Even the microwave frequency domain can be accessed via the spacing of the comb lines (or equivalently, harmonics of f_{rep}). Use of the microwave information is one scheme for making the clockwork (gear reduction) of an optical atomic clock, allowing access to a countable clock readout and therefore time-keeping.

Alternatively, instead of locking one mode of the comb to an optical frequency reference, f_{rep} can be stabilized directly to a microwave reference. The photodetected comb repetition frequency is mixed with a microwave reference to produce an error signal

that can be applied to a PZT-actuated mirror in the laser cavity. With f_0 stabilized using the f -to- $2f$ scheme, the stability of the microwave reference is transferred to every comb line, and thus to the entire optical spectrum within the comb bandwidth.

Whether locking the comb in the microwave or optical domain, some finite residual locking error will always exist (perhaps set by thermal noise or shot noise in the best case). When stabilizing the comb using a heterodyne beat against some optical reference, this locking error contributes much less to the fractional instability than it would contribute in the case of stabilization of f_{rep} , as discussed in Section 2.3. Therefore it is best to consider optically locking the comb for deriving microwave information from stable optical references in order to take advantage of the large frequency division (and therefore division of the noise) from the optical reference to the microwave frequency of interest. In addition, optical references based on lasers stabilized to highly isolated passive cavities now represent the most stable and pure oscillators from any region of the spectrum – therefore directly locking the comb to such a source will provide the best available performance to the comb.

The simplicity of locking a femtosecond comb relative to some reference is made possible by the need to actively control only two degrees of freedom, allowing the mode-locking process of a femtosecond laser to enforce the rigid spacing of comb lines across the laser’s spectrum. The accuracy of Eqn. 3.4 has now been tested at the 1×10^{-19} level in comparisons of four independent combs, each locked to a common optical reference [109]. Indeed, stabilization of femtosecond frequency combs has become a routine and robust method for transferring the phase coherence and stability offered by a reference from anywhere in the electromagnetic spectrum (rf and microwave [105, 110, 104, 98, 111], infrared [24, 108, 112, 113], visible [114, 115], or even UV and XUV [31, 116]) across vast frequency gaps.

Chapter 4

Optical to Microwave Conversion

In Chapter 3 I showed how the frequency comb output from a femtosecond laser can be used to link optical and microwave frequencies together in a phase-coherent way. Essentially, some harmonic of f_{rep} can be detected as the microwave output of such an optical-to-microwave converter. Due to the inverse dependence of Eqn. 2.4 on the carrier frequency, the timing jitter from thermal and shot noise for such microwave detection is significantly greater than for optical heterodyne detection. Therefore great care must be taken to use high optical powers to detect high harmonics of f_{rep} , as discussed in Section 2.2. In this Chapter I will describe the efforts I have made toward extremely low-jitter detection of a femtosecond laser's stable repetition rate harmonics, out to large Fourier frequencies. As opposed to the case for measuring extremely low instabilities at long timescales, here I will concentrate on developing extremely pure (low phase noise at high Fourier frequencies) optical-to-microwave conversion. The applications for such pure microwaves include synchronization technology (as will be discussed in Chapter 5) for a wide variety of pump-probe applications at high sampling rates, and remote fiber transfer of femtosecond frequency combs as will be discussed in Chapter 8. In addition to creating extremely pure microwave signals at high (> 10 kHz) Fourier frequencies, the issue of dividing the optical frequency of a next-generation optical atomic clock down to the microwave regime is of prime importance for time-keeping. Even the world's best optical frequency reference is not useful for time-keeping without being able to count

its cycles (or better yet, keep track of its phase). Therefore the most stable conversion of a femtosecond laser's optical pulse timing into the timing of electronic pulses from a photodetector is necessary in order to faithfully transfer the optical frequency reference's stability into the countable microwave regime.

4.1 Measuring phase noise electronically

In Chapter 2 I defined the phase noise spectral density (Eqn. 2.2) and gave some examples of fundamental and technical limitations, such as shot noise and thermal noise, to the achievable phase noise purity. I concluded with the observation that optical techniques for phase detection were desirable due to the high timing resolution offered by optical frequencies. However, sometimes stable and pure microwaves must still be derived from optical references, such as in the clockwork described in Section 3.3 and Chapter 7. This is particularly true since state-of-the-art optical frequency references are now achieving the highest stability and purity of any electromagnetic oscillators; therefore it should be able to derive similarly-pure microwave references using the frequency comb. There is, however, one difficulty of using the comb: when measuring the phase of microwave information derived from pulsed lasers (harmonics of f_{rep}), great care must be taken to reduce the fundamental limits while avoiding technical noise sources (see Section 2.2).

4.1.1 Homodyne phase detection with a mixer

When two microwave signals are to be compared in order to measure their relative phase fluctuations, the traditional technique is to use a saturated mixer. Ignoring amplitude and phase noise on the input signals driving the mixer's local oscillator (LO) and radio frequency (RF) ports, they can be written as:

$$V_{LO}(t) = V_{LO} \sin(2\pi\nu_{LO}t + \phi_{LO}) \quad (4.1)$$

$$V_{RF}(t) = V_{RF} \sin(2\pi\nu_{RF}t + \phi_{RF}). \quad (4.2)$$

The mixer is a nonlinear device which (ideally) multiplies the two signals together, and outputs this product at the intermediate frequency (IF) port:

$$V_{IF}(t) = V_{max} \begin{bmatrix} \cos[2\pi(\nu_{LO} - \nu_{RF})t + (\phi_{LO} - \phi_{RF})] \\ -\cos[2\pi(\nu_{LO} + \nu_{RF})t + (\phi_{LO} + \phi_{RF})] \end{bmatrix}. \quad (4.3)$$

When used as a phase detector, $\nu_{LO} = \nu_{RF}$, and by using a low pass filter to dump the sum-frequency term in 4.3, we have the phase detector dc output voltage

$$V_{IF}(\Delta\phi) = V_{max} \cos(\phi_{LO} - \phi_{RF}) = V_{max} \cos(\Delta\phi). \quad (4.4)$$

Note that because the input power $(V_{LO}^2 + V_{RF}^2)/R$ (R is the system's characteristic impedance, usually 50 Ohms) is split into a sum frequency band and a difference frequency band, the output power in the difference frequency band (phase detection band) is at most half of the input power, with additional losses in the mixer itself. For typical powers at the RF and LO ports of 0 and +7 dBm, respectively, V_{max} is typically around 250 mV. The phase shifter in one arm of the measurement system is used to set the phases in quadrature at the mixer such that $\Delta\phi = 90^\circ$, which is where $V_{IF}(\Delta\phi)$ both crosses through zero and is the most sensitive to fluctuations in the relative phase.

For small deviations of the phase, the sensitivity of $V_{IF}(\Delta\phi)$ near this point is usually approximated as linear. The phase shifter can be tuned to determine the coefficient V_{max} in Eqn. 4.4. Then, by evaluating the phase-derivative of Eqn. 4.4 at the phase-quadrature (zero output voltage) point, the mixer's sensitivity to phase fluctuations is simply

$$K_\phi = \left[\frac{d}{d(\Delta\phi)} (V_{IF} \cos(\Delta\phi)) \right]_{\Delta\phi=90^\circ} = V_{max} [\text{V/rad}]. \quad (4.5)$$

Therefore, by analyzing the mixer's output with a fast-fourier-transform, which provides the double-sideband voltage spectral density $S_V^{IF}(f)$, we have knowledge of the phase

noise power spectral density of the DUT:

$$S_{\phi}^{DUT}(f) = S_V^{IF}(f) / K_{\phi}^2. \quad (4.6)$$

One practical consideration is the extremely low voltage noise levels to be expected from the mixer's output for state-of-the-art phase noise measurements. For the case of the typical level-7 mixer I have been discussing, a 10-GHz signal with, for example, -150 dBc/Hz of single-sideband phase noise in a 50-Ohm system will lead to a measured voltage power spectral density of $S_V^{IF}(f) = -159$ dBV/Hz. Such fluctuations are so weak that a low-noise dc amplifier should normally be used immediately at the output of the mixer. This ensures that any electronic noise pickup on a subsequent cable leading to the FFT machine or to locking electronics will be negligible compared to the measured phase noise. In fact, I believe low-frequency (sub-Hz) drifts of the out-of-loop synchronization measurements to be described in Chapter 5 could be attributed to such pickup, since at the time of those measurements we did not use such a low-noise amplifier close to the mixer output. Even just a little mechanical fluctuation of the BNC cable leading away from the mixer could lead to a significant timing fluctuation as a result.

4.1.2 Independent oscillators vs. single component measurement

It is important to understand the distinction between measuring the excess phase noise added to some pure signal by a microwave component, and measuring the relative phase noise of two independent oscillators.

The traditional saturated-mixer design for the measurement of the excess phase noise of a microwave component is shown in the left of Fig. 4.1. A strong, low-noise oscillator is split and driven through two differential paths, one containing the device-under-test (DUT), before driving the LO and RF ports of a double-balanced mixer (DBM). Because the power of the common-mode oscillator can be made very large, the excess noise of the DUT can be measured down to extremely low levels without




Figure 4.1: Traditional saturated-mixer systems for phase measurement. (Left) Common-mode phase noise measurement system for determining the excess phase noise of the DUT with a strong oscillator. (Right) Differential phase comparator for forming real-time phase comparisons of the two independent oscillators. For the case of optical-to-microwave conversion of harmonics of f_{rep} , the two signals in the phase comparator are relatively weak and must therefore be amplified, which represents an additional source of noise.

contamination of the measurement by thermal noise.

In contrast with the system used to measure the excess phase introduced by the DUT, a traditional phase comparator is shown on the right side of Fig. 4.1, for the purpose of measuring the relative phase noise between two independent oscillators. This is similar to the typical system we use for the phase comparison of microwave harmonics of f_{rep} after photodetection of pulse trains from femtosecond lasers. Two independent signals are formed, but are typically not strong enough to drive and saturate a mixer. As a result, amplification in each arm becomes necessary. The amplifier flicker noise can limit the performance of the traditional phase comparator, but microwave interferometric techniques can be used to prevent the flicker noise in some cases. See Section 2.2.2 for references discussing such methods.

This distinction between a phase noise measurement system and a phase comparator is very important. The phase noise measurement system assumes a strong, low-noise oscillator is available, so the DUT's phase noise properties can be measured with extremely low phase noise contributions from thermal noise. Another way of seeing this is that P_{rf} in Eqns. 2.10 and 2.12 can be made extremely large so as to reduce these limitations. The phase comparator is used for a different purpose, namely that of measuring the relative phase of the two oscillators. In the case of optical detection of harmonics of f_{rep} , often the oscillators themselves have weak signal strength and plenty of relative amplitude and phase noise, as opposed to the strong, relatively quiet and common-mode source used to drive the phase noise measurement system. Therefore the problem of making phase noise measurements of an optical-to-microwave conversion system, as discussed next in Section 4.1.3, is somewhat more difficult than the usual case of measuring noise of individual microwave components.

4.1.3 Measuring the phase noise of repetition frequency harmonics

I now turn to the specific case of detecting harmonics of f_{rep} from a femtosecond laser. As discussed in Section 2.3.1, this technique is insensitive to the optical phase and so does not take advantage of the large scaling factor between optical and microwave frequencies. Only the timing of the pulses, an intrinsically rf or microwave quantity, is detected, so considerations regarding excess noise sources are important.

In order to fully compare the microwave signals from two independent lasers, each must be individually photodetected and subsequently phase-compared. Figure 4.2 shows the simplest possible such system. Light from two independent femtosecond lasers is incident upon two high-speed photodetectors. The detectors respond at all harmonics of f_{rep} within their bandwidth (before their response becomes too small). The particular harmonic of interest is filtered with a narrow microwave bandpass filter, amplified, and used to drive a DBM as in the phase comparator scheme. A phase shifter before either of the mixer's inputs is used to tune the relative phase of the two oscillators to the appropriate quadrature point ($\Delta\phi = 90^\circ$) for making a phase measurement. Amplitude noise can also be measured by tuning the phase shifter so as to maximize the DC output of the mixer. For the extremely low level of voltage fluctuations we may hope to detect, it is important to use a low-noise preamplifier located close to the mixer, before analyzing the signal on an FFT machine or using it to feedback to one oscillator for synchronization. The preamplifier helps to reduce the effects of pickup from small dc voltage fluctuations on, for example, a BNC cable leading to the FFT machine or the feedback electronics.

Although the system shown in Fig. 4.2 is the simplest system for measuring the microwave phase noise of photodetected harmonics of f_{rep} , it still lends itself to many possible sources of additional, unwanted noise. These excess noise sources were described in Section 2.2, but here I will briefly list them again. For instance, the photodetectors

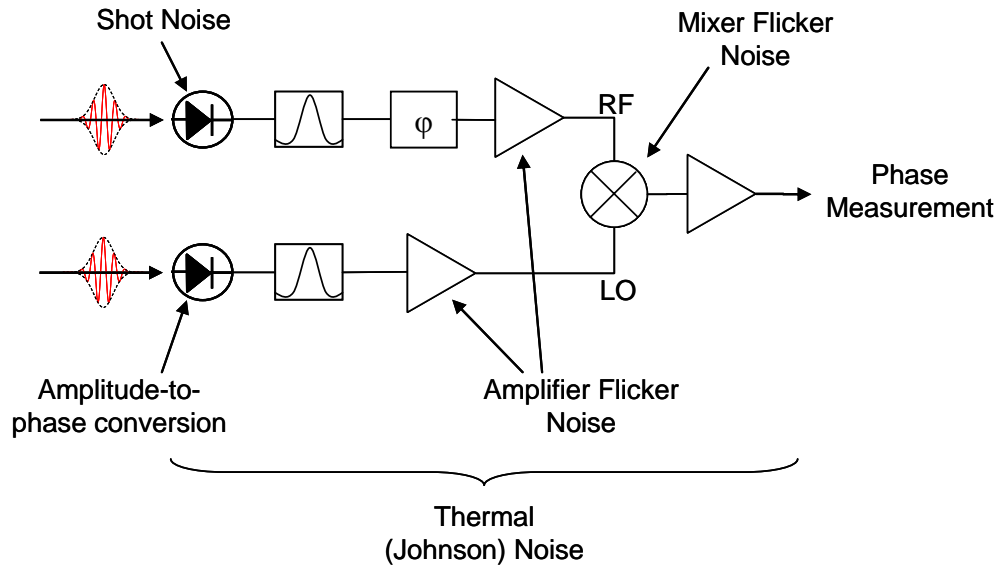


Figure 4.2: Measurement of the relative phase noise of two independent microwave signals available from photodetection of the pulse trains of fs lasers. The two pulse trains might come from separate lasers for the purposes of synchronization, or from the same laser for the purposes of measuring the optoelectronic measurement system's noise floor. Tunable bandpass filters are used to select the desired harmonic of f_{rep} , and amplification is necessary in order to appropriately drive the mixer. A preamplifier is used immediately after the mixer in order to avoid problems with pickup on any transmission cable leading to the FFT machine or servo electronics used to synchronize the two oscillators. Even though the optical pulse trains' repetition frequencies may have an exceedingly noiseless phase relationship, excess noise enters into the measurement system in several ways. The photodetectors allow for amplitude-to-phase conversion, or pointing-to-phase conversion, as well as the usual shot noise. The amplifiers and mixer all have an associated $1/f$ flicker noise. Thermal (Johnson) noise for the typical 50-ohm system is also present.

can introduce amplitude-to-phase conversion if any optical amplitude noise exists on the incident light. Pointing instability of the beams can create a similar effect. The usual shot noise from photodetection is present, in addition to thermal (Johnson) noise throughout the 50-ohm electronic system. Finally, the amplifiers and mixer can have excess noise, usually in the form of $1/f$ flicker noise. For all these reasons, extremely careful design of the optoelectronic system is critical. The subjects of the following two Sections will therefore be: first, how to build an appropriate electronic system, and second, how to construct an optical detection system minimizing unwanted excess noise. I will then describe how we combined the two systems for the purpose of measuring the noise floor of a complete optical-to-microwave conversion system, and show that we have now reached the combined shot and thermal noise floor for 1.4 mW of incident average optical power per photodetector, limited at Fourier frequencies below 10 kHz by the mixer's intrinsic flicker noise.

4.2 Noise floor of the microwave system

If the complete optical-to-microwave conversion system is to have the best-possible performance, a useful starting point is to make sure the microwave components are as noiseless as possible. In order to measure the noise floor of the microwave system, I modified the basic setup shown in Fig. 4.2 to include only microwave components by replacing the two photodiodes with a common-mode synthesizer split between the two arms of the experiment, as shown in Fig. 4.3. The 7.6-GHz microwave carrier frequency of the synthesizer was chosen because it is both the 10^{th} harmonic of a 760-MHz laser and the 80^{th} harmonic of a 95-MHz laser, both located on the same optical table in our laboratory. For 1 mW of incident optical power on the photodetector, we measured ~ -9 dBm of output microwave power at 7.6 GHz for the higher rep-rate laser. Therefore we simulated that same microwave power in each arm of the microwave phase noise measurement system.




Figure 4.3: Experimental setup for measuring the microwave system’s noise floor. A common synthesizer operating at 7.6 GHz is split into the two arms of the experiment. Attenuators are used just after the splitter in order to simulate the same power levels we expect from the photodiodes that will be eventually used in the optical-to-microwave conversion system. Low-phase-noise amplifiers (designed and built by Terry Brown) are then used to appropriately drive the mixer. JILA circuit design part numbers are given for the interested reader. We choose to drive the LO port with +8 dBm and the RF port with -2 dBm of microwave power. A phase shifter allows the mixer to be used at the phase quadrature point, but also allows exploration of interesting amplitude-to-phase conversion effects by slightly detuning the relative phase. A low-noise dc amplifier with a gain of 100 is used immediately at the mixer output so that electronic pickup between the mixer output and the FFT machine is made insignificant.

I will now briefly describe some low-phase-noise amplifiers we employ for the optical-to-microwave phase noise measurement, before a more lengthy description of a series of experiments showing that the mixer’s amplitude-to-phase conversion can be a limiting factor. Finally, even if the amplitude-to-phase conversion is solved, the mixer’s intrinsic flicker noise becomes the practical limit without resorting to microwave interferometric techniques.

4.2.1 Low-Noise Amplifiers

Even though we should achieve as much microwave power from the photodiodes as possible, the +7 dBm it would take to drive a level-7 double-balanced mixer would require far too much incident optical power for the photodetectors. Therefore amplification of the microwave signals is necessary. However, typical amplifiers exhibit a $1/f$ (for Fourier frequency f) flicker noise near the microwave carrier. This noise comes about due to nonlinearities in the amplifier which lead to upconversion of noise in the dc bias [66], as discussed in Section 2.2.2. A useful starting point for avoiding this problem is therefore to choose a highly linear amplifier chip. Terry Brown, in the JILA electronics shop, hunted down such a linear amplifier chip made by Sirenza, part number SBW-5089. With an appropriate micro-strip board layout, Terry was able to make one-, two-, and even three-stage amplifiers, each stage having roughly 10- to 12-dB of gain at 7.6 GHz. Although his initial design also included feedback to the bias voltage for the purpose of stabilizing its flicker noise, I have not yet been able to measure a contribution of excess phase noise from the amplifiers relative to the phase noise from the mixer. This assertion will be proven next.

4.2.2 Mixer Noise

In order to determine whether the noise floor of the system shown in Fig. 4.3 is limited by the amplifiers or the mixer, I constructed the system shown in Fig. 4.4. With-

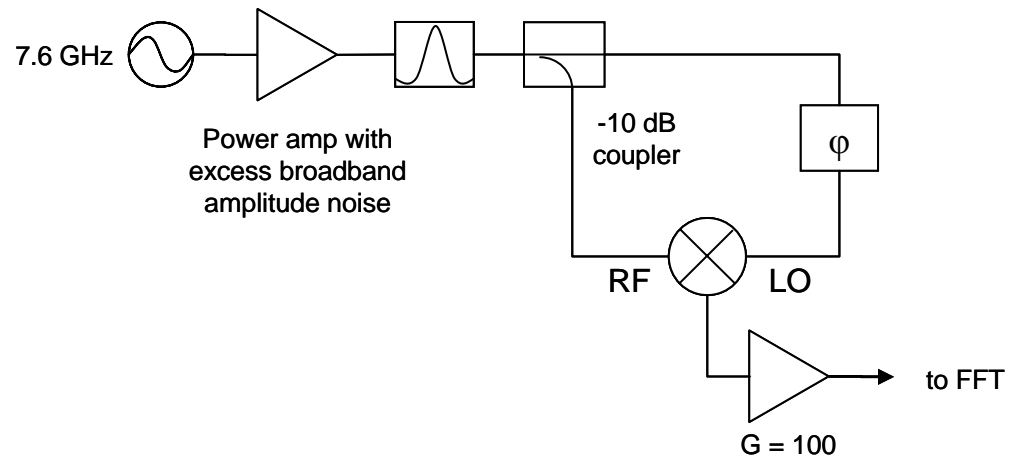


Figure 4.4: Experimental setup for measuring the mixer's noise floor. With no amplifiers, the measured noise must come from either the directional coupler, microwave cables, the phase shifter, or the mixer. Of these, the coupler and the mixer are the most suspect. Since the synthesizer could not drive the system with enough power, a power amp was used before the coupler. This led to very interesting results, since the power amp had significant broadband amplitude noise. Through amplitude-to-phase conversion in either the coupler (differential amplitude-to-phase noise between its two outputs) or the mixer, an excess of broadband phase noise was measured at the mixer output when the phase shifter tuned the system to phase quadrature. However, tuning slightly away from quadrature led to cancellation between amplitude noise and the phase noise it generated.

out amplifiers, the noise floor of the mixer could be directly measured. Unfortunately, the common-mode synthesizer could not output enough power to directly drive both mixer inputs, so I used a power amplifier before splitting the signal, and then ensured that the power driving the two mixer inputs was identical to the system with amplifiers.

This eventually led to a very interesting set of measurements regarding the mixer's amplitude-to-phase conversion, since the power amp had considerable amplitude noise. At first I tuned the system's phase shifter in one arm so that $V_{IF}(\Delta\phi)$ was set to zero DCV, and measured the residual voltage fluctuations with an FFT machine in the usual way. Surprisingly, the voltage fluctuations were greater than for the more-complicated system with amplifiers. This is shown in Fig. 4.5 in the upper plot as the upper curve. However, I noticed that by tuning the phase shifter slightly away from this point, I could decrease the voltage noise by a significant amount, as shown in the second highest curve in the upper plot of Fig. 4.5. This "sweet spot" did not occur at the level of a few mV that might be expected for a mixer with a small dc offset – rather, it occurred when the dc output voltage was 70 mV ($V_{max} \sim 260mV$). As a result, I postulated that either the mixer or the splitter was undergoing amplitude-to-phase conversion. If the amplitude noise and the resulting converted phase noise were therefore coherent with respect to each other, it would in principle be possible to find a particular value of $\Delta\phi$ for which the mixer output would have equal and opposite voltage fluctuations due to the amplitude and phase noise. This should lead to cancellation of the coherent noise processes, leaving only voltage fluctuations from the residual incoherent amplitude and/or phase fluctuations. Note that the level of cancellation is about 25 dB from 1 kHz to 10 kHz, but at lower frequencies the cancellation is less. The residual flicker floor is due to the intrinsic mixer flicker noise as shown in the lower plot of Fig. 4.5, which compares the mixer-only measurement at the sweet spot to the full microwave system measurement including microwave amplifiers. The two curves differ only by a few dB across the Fourier spectrum. This difference can be explained by the different values of

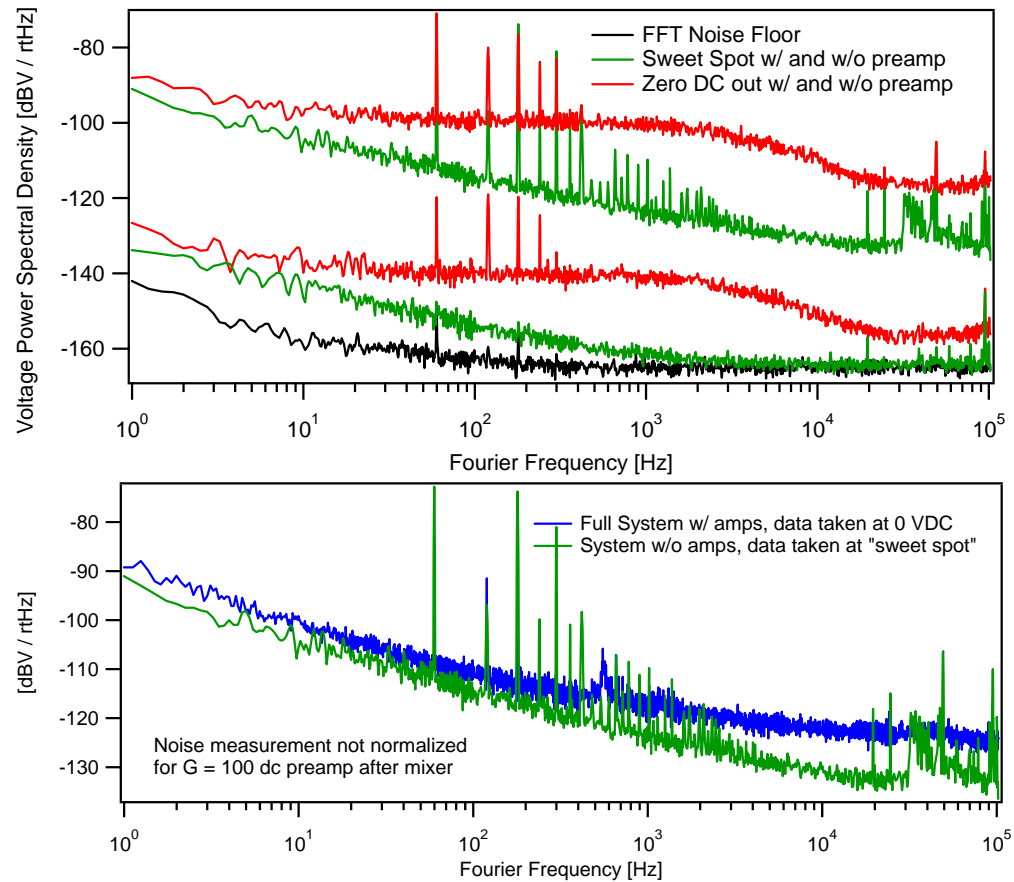


Figure 4.5: Amplitude-to-phase conversion and the mixer flicker noise floor. The upper plot shows several measurements of noise at the output of the system shown in Fig. 4.4. The upper two curves are data taken at the output of the dc preamplifier (Gain of 100) used at the mixer’s IF output. The highest curve in red is data taken at the phase quadrature point (0 VDC) while the next lower curve in green is the noise when the phase shifter is tuned to the “sweet spot” where amplitude noise cancels the phase noise it generates. The same data was also taken without the preamplifiers, again shown in red and green as the third and fourth lowest curves. Note that above 10 kHz, the preamplifier allows data to be taken below the noise floor of the FFT machine, shown in black. The lower plot shows voltage fluctuations out of the full system from Fig. 4.3 including microwave amplifiers (in blue) compared to the sweet spot data from the upper plot. Due to slightly different input powers to the mixer, the data differ by a few dB, but this provides strong evidence that the full microwave system is limited by the mixer’s flicker noise out to beyond 10 kHz where thermal noise begins to dominate.

V_{max} for the mixer in the two different setups.

I tested the idea of amplitude noise cancelling the phase noise it generates (through amplitude-to-phase conversion) in several ways. First, I made a measurement of the mixer’s output voltage noise at 10-kHz Fourier frequency as a function of the phase shifter’s value. The results are plotted in Fig. 4.6, along with the average dc value of V_{IF} as a function of the relative phase. The noise at 10 kHz achieves a minimum when $\Delta\phi$ is slightly below the phase quadrature points, corresponding to dc output voltages of $\sim \pm 70$ mV. I choose to denote the phase of these “sweet spots” as ϕ_{ss} and the dc voltage at the sweet spots as V_{ss} . The fact that the sweet spots occur at points along the dc voltage curve with both opposite sign and opposite slope is consistent with the amplitude-to-phase conversion idea. Also, note that the local maxima of noise occur slightly shifted from the amplitude quadrature point (where $V_{IF} = V_{max}$) by the same magnitude of phase shift between the noise minimum and the phase quadrature point. This is because the amplitude and phase noise can add constructively for these values of the phase.

Another way to show the amplitude-to-phase conversion is really occurring is to purposely amplitude-modulate the synthesizer output at a particular frequency; here I choose a modulation frequency of 10 kHz. Figure 4.7 shows the voltage spectral density of V_{IF} when the average value of V_{IF} is both zero, and when the phase is tuned to ϕ_{ss} , demonstrating a 30-dB cancellation between the coherent amplitude and phase fluctuations. Note that in addition, the depth of modulation was increased to give 40 dB of extra phase noise (as measured at the phase quadrature point) so that the resulting noise at the sweet spot could even be measured. Therefore the sweet spot cancellation between amplitude noise and phase noise coherently stemming from amplitude-to-phase conversion can be as high as 70 dB. Also, the broadband noise did not decrease by 70 dB at all Fourier frequencies, implying that some residual source of noise becomes the new limitation once the phase noise from amplitude-to-phase conversion is cancelled by the




Figure 4.6: Plot demonstrating the cancellation of amplitude noise and the phase noise it generates through amplitude-to-phase conversion. The blue sinusoidal curve plotted on the left axis shows the mixer’s dc output voltage as a function of the relative phase between its two input signals. The red data plotted on the right axis is the level of voltage noise at 10 kHz Fourier frequency. At the “sweet spots” near 0.4π and 1.4π radians of relative phase, coherent cancellation of the amplitude and phase noise occurs, indicating that a form of amplitude-to-phase conversion is occurring in the system, most likely in the mixer itself.

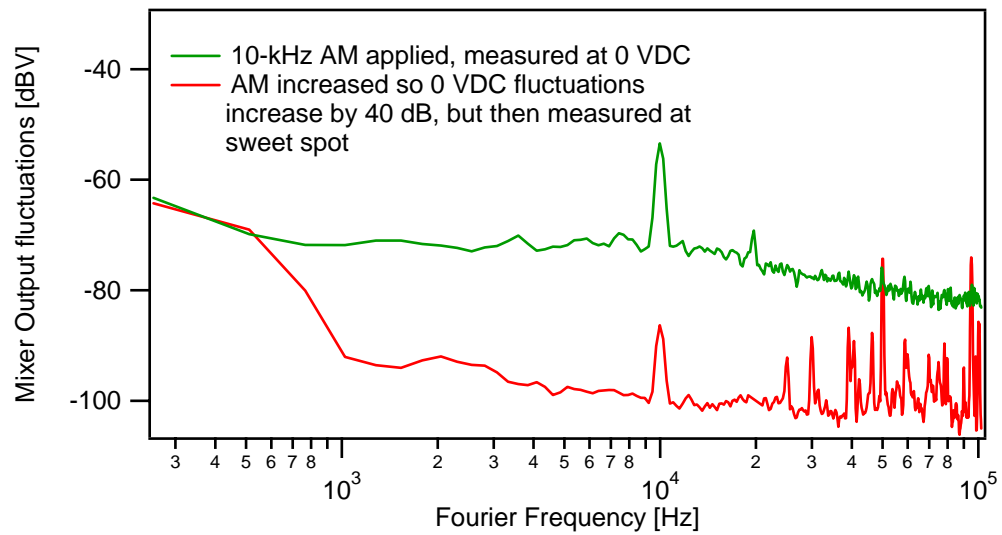


Figure 4.7: Amplitude modulation leading to phase modulation through amplitude-to-phase conversion. The two curves shown are measured using the setup from Fig. 4.4, but with amplitude modulation at 10 kHz applied to the common-mode synthesizer’s output. The upper curve was measured at the phase quadrature point where the mixer’s dc output voltage is zero. The lower curve was measured at the sweet spot where amplitude and phase fluctuations cancel. For the lower curve, the amplitude modulation was purposely increased by an amount that raised the phase-quadrature signal by 40 dB, so the cancellation at the sweet spot was as good as 70 dB.

amplitude noise it came from, at the sweet spot.

Mathematically, the cancellation between coherent amplitude and phase noise can be described as follows: let the output of the mixer be corrupted by both amplitude and phase fluctuations. We can write it as

$$V_{IF} = (V_{max} + \delta V_{max}) \cos(\Delta\phi + \delta\phi) \quad (4.7)$$

where δV_{max} and $\delta\phi$ represent the amplitude and phase fluctuations from the average amplitude V_{max} and average relative phase $\Delta\phi$. Then for small fluctuations of the phase and amplitude, we can expand V_{IF} around the point $\Delta\phi_0$ as follows:

$$\begin{aligned} V_{IF} &\approx (V_{max} + \delta V_{max}) \cos(\Delta\phi_0) - (V_{max} + \delta V_{max}) \sin(\Delta\phi_0) \delta\phi \\ &= V_{max} \cos(\Delta\phi_0) \\ &\quad + \delta V_{max} \cos(\Delta\phi_0) - \delta\phi \sin(\Delta\phi_0) V_{max} \\ &\quad - \delta\phi \delta V_{max} \sin(\Delta\phi_0) \end{aligned} \quad (4.8)$$

The first term is the constant dc term and is zero at exactly the phase quadrature point, ignoring any small dc offset of the mixer's response. The second term can be used to measure pure amplitude fluctuations if $\Delta\phi_0$ is tuned to π , the amplitude quadrature point, since the third and fourth terms will go to zero.

The third term is what we normally think of when we want to measure phase noise, since setting $\Delta\phi_0$ to the phase quadrature point of $\pi/2$ gives maximum sensitivity to $\delta\phi$ with zero contribution from the second term. Then only the third and fourth terms contribute, and they can be combined as $\delta\phi (V_{max} + \delta V_{max})$. As long as $V_{max} \gg \delta V_{max}$ this is an accurate measure of phase noise. Also, the phase quadrature point must be found accurately enough so that the second term is negligible. This condition is satisfied when $\delta\phi \gg (\delta V_{max}/V_{max}) \cos(\Delta\phi_0)$.

The sweet spot cancellation comes about when the second and third terms add to zero. This can happen if the amplitude and phase noises are correlated. Suppose

the phase fluctuations $\delta\phi$ can be written as the sum of $\delta\phi_a$, processes coherent with the amplitude noise, and $\delta\phi_p$, processes incoherent with the amplitude noise representing the true phase fluctuations we want to measure. Then at the sweet spot, $\Delta\phi_0 = \phi_{ss}$ and Eqn. 4.8 can be rewritten as

$$\begin{aligned}
V_{IF} &= V_{max} \cos(\phi_{ss}) \\
&+ \delta V_{max} \cos(\phi_{ss}) - \delta\phi_a [\sin(\phi_{ss}) (V_{max} + \delta V_{max})] \\
&- \delta\phi_p [\sin(\phi_{ss}) (V_{max} + \delta V_{max})]
\end{aligned} \tag{4.9}$$

Now if $\delta V_{max} \ll V_{max}$, the terms on the second line will cancel coherently at the sweet spot since $\delta\phi_a$ and δV_{max} are related by the amplitude-to-phase conversion coefficient $C_{ap} \equiv \delta\phi_a/\delta V_{max}$. Aside from the dc offset (first term), only the phase noise unrelated to amplitude noise will be measured as the fourth term. Then for small amplitude noise ($\delta V_{max} \ll V_{max}$) the remaining phase noise measurement term is just

$$-\delta\phi_p (V_{max} \sin(\phi_{ss})). \tag{4.10}$$

Note that after measuring this voltage spectral density, the phase noise information $\delta\phi_p$ is calculated by dividing the measured linear voltage noise spectrum by $V_{max} \sin(\phi_{ss})$, not just by the usual V_{max} .

It is simple to show from Eqn. 4.9 and from the definition of C_{ap} that in the limit of small amplitude noise, knowing ϕ_{ss} and V_{max} immediately gives the amplitude-to-phase conversion coefficient as

$$C_{ap} = \frac{1}{V_{max} \tan(\phi_{ss})} \tag{4.11}$$

For example, with the data from Fig. 4.6, ϕ_{ss} was 0.416π rad and V_{max} was 261.9 mV, so C_{ap} is 1.03 rad/V. Another way to arrive at this value is to note that the pure amplitude noise measured at the amplitude quadrature was -145 dBV/ $\sqrt{\text{Hz}}$ while the phase noise from amplitude-to-phase conversion, measured at the phase quadrature point, was -156 dBV/ $\sqrt{\text{Hz}}$. Then using $V_{max} = 261.9$ mV to convert from phase to

amplitude noise, C_{ap} is calculated to be 1.08 rad/V, in excellent agreement with the first calculation.

Another important observation is that as I applied more power to the mixer by simply turning up the synthesizer's power, ϕ_{ss} moved further from the quadrature point. Fig. 4.8 demonstrates this point in two ways. The upper left plot shows the mixer's dc voltage output as a function of the relative phase $\Delta\phi$, for various microwave power levels at the mixer's RF port. Here, 0 dB of attenuation corresponds to +5 dBm of power at the RF port, and for all curves the LO was driven by +8 dBm of microwave power. Note the saturation occurring as the RF power is increased. The upper right plot shows the voltage fluctuations at the mixer's IF output, in a 256-Hz bandwidth centered at 10 kHz Fourier frequency, for the same four powers driving the RF port. The middle plot shows this same data on a linear vertical axis. A calculation of just the terms in Eqn. 4.8 reveals that the functional form of the data in the middle plot should be the absolute value of a sinusoid. At the lowest level of power at the mixer's RF port, this becomes true, but for higher powers higher-order terms begin to matter and the functional form becomes much more complicated. Finally, the bottom graph directly plots the voltage fluctuations against the dc voltage output. The sweet spot cancellation clearly moves further from the phase quadrature point as the driving power is increased, consistent with amplitude-to-phase conversion that depends on the magnitude of the amplitude fluctuation, not on the fractional amplitude noise. The mixer's small dc offset of ~ 5 mV can also be seen, so even when $V_{IF} = 0$, the mixer is not at the true phase quadrature point.

Finally, the measurements do not pinpoint whether the source of the amplitude-to-phase conversion is in the mixer or in some other component such as the directional coupler used to split the signal to the two mixer inputs. It is unlikely the phase shifter or microwave cables are responsible. By trying a few different mixers, it became clear that since the value of ϕ_{ss} changed for each mixer, the amplitude-to-phase conversion

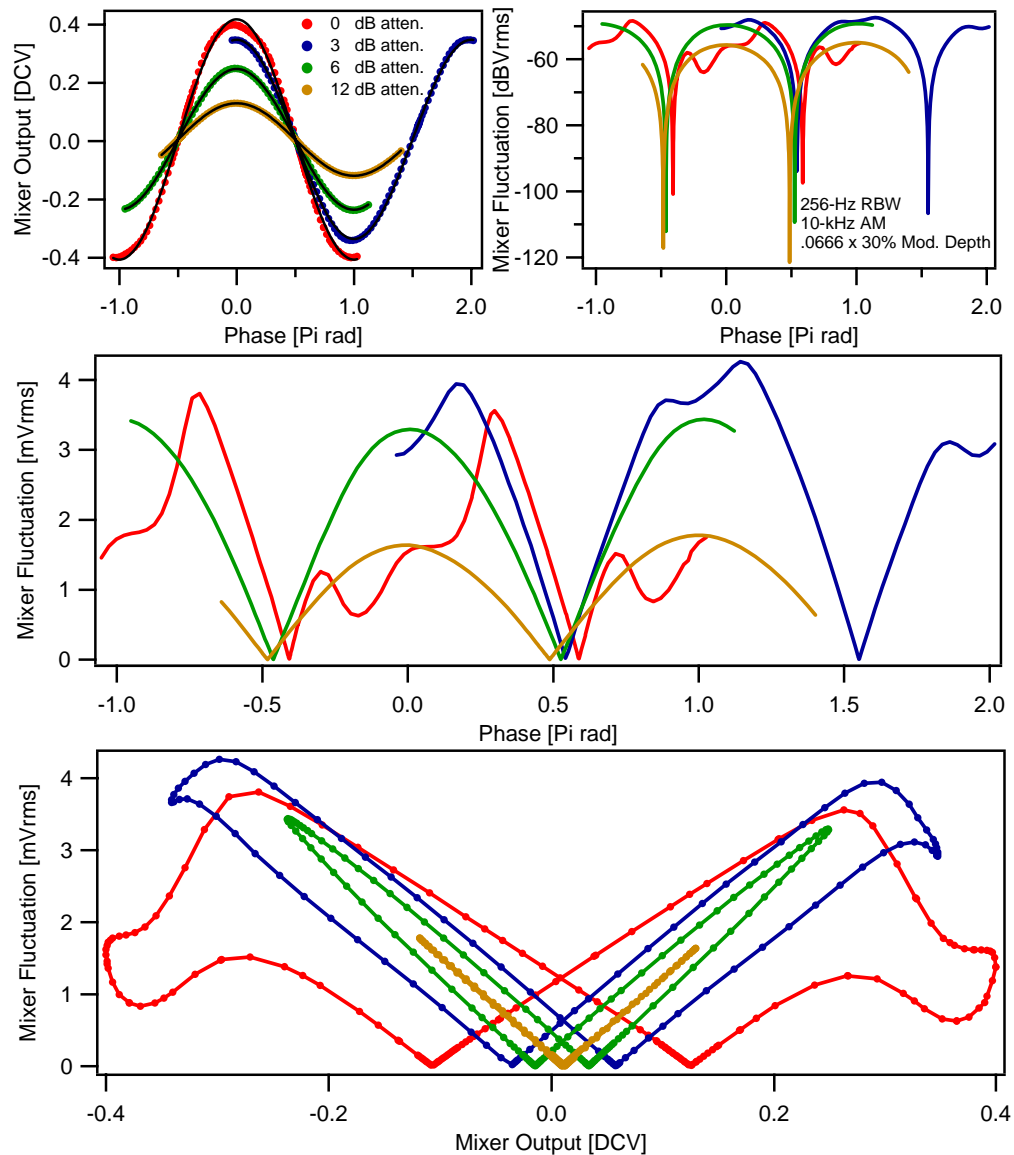


Figure 4.8: Amplitude-to-phase noise cancellation for various powers at the mixer's RF port, using the setup of Fig. 4.4. The LO port was always driven with +8 dBm of power at 7.6 GHz, but the RF port was driven into saturation as high as +5 dBm for the red curves. As labelled, 0 dB of attenuation corresponds to this +5 dBm of power into the RF port. The top left graph shows the mixer's dc output voltage as a function of the relative phase $\Delta\phi$ for the four power levels between -7 and +5 dBm. The upper right curve shows the voltage fluctuations of $V_{IF}(\Delta\phi)$ at 10 kHz Fourier frequency in a 256-Hz bandwidth vs $\Delta\phi$, on a log vertical axis. The middle plot shows the same data on a linear vertical axis. The bottom plot directly plots the voltage noise against the dc voltage output. As the mixer is driven higher into saturation, the sweet spot cancellation occurs further and further from the phase quadrature point.

was mixer-dependent. It would be useful to verify that ϕ_{ss} does not depend on the coupler used to split the common-mode synthesizer input.

With these problems of amplitude-to-phase conversion now well understood, recall that the green data shown in the bottom plot of Fig. 4.5 is a measurement of only the phase noise uncorrelated with the amplitude noise, for the system in Fig. 4.4 using no amplifiers. When compared with the phase noise of the full system from Fig. 4.3 including microwave amplifiers (shown in blue), it seems believable that the limiting flicker noise floor is associated with the mixer, not the amplifiers, since both systems exhibit nearly the same magnitude of flicker noise.

This conclusion, that the mixer flicker floor (rather than the amplifiers) limits the system's performance below 10 kHz, means that any further improvements to the system can only be made by interferometric techniques to get around the intrinsic mixer noise. At higher Fourier frequencies, we of course experience the usual thermal noise limit from the amount of power we use to drive the mixer. Clearly some sufficient amount of amplitude noise on the microwave signals from the photodetectors in Fig. 4.2 will degrade the phase noise measurement of the system through these electronic amplitude-to-phase conversion processes. Therefore it can become important to stabilize the amplitude noise of the input microwave signals, as I will discuss in Section 4.3. Note that the amplitude-to-phase conversion I discussed in this Section is most likely only the differential conversion between the two mixer input ports, as opposed to the absolute conversion of either port. If amplitude noise is uncorrelated at the two ports, the effect of amplitude-to-phase conversion may be even worse than suggested by my measurement of the coefficient C_{ap} , and so stabilization of the microwave amplitude noise may become even more important.

4.2.3 Practical considerations

It is important to briefly mention a few practical considerations for making the electronic phase noise measurements. First, for repeatability of the input levels to the mixer, it was fairly important to use a microwave power meter (essentially a high-frequency rectifier) to measure the power levels. Also, a torque wrench for sma connectors was important to ensure repeatability of the various connections. Second, reducing ground loops became important for the microwave system. All the microwave components were electrically isolated from the breadboard so that ground currents had only one path to use. Unfortunately, the dc preamplifier used at the mixer's IF port contaminated the signal with 60-Hz pickup noise, which is the disadvantage of using such a preamp. This contamination can easily be seen by comparing the upper and lower green curves of the upper plot in Fig. 4.5. However, as has already been explained, the preamp is necessary to avoid the FFT machine's noise floor and will be absolutely important in any system using the signal for active feedback, to reduce the effect of pickup (especially at low frequencies) on any BNC cable leading away from the mixer toward feedback electronics. Third, the microwave components should always be fixed to some common mount such as a small breadboard to give them mechanical stability. The level of timing fluctuations for these measurements correspond to extremely small path length fluctuations in the electrical system. It may even become important to thermally stabilize the microwave electronics. Finally, looking toward the future, it may become necessary to use ac encoding of the phase noise measurement by microwave interferometric techniques, combined with lock-in detection at the mixer output, in order to shift the measured noise to higher Fourier frequencies and therefore lower on the mixer's intrinsic flicker noise floor. This is an idea put forward by Eugene Ivanov, who always has wonderful ideas when it comes to phase noise measurement in the microwave regime.

4.3 Noise floor of the optical system

With the microwave system performing at the limit of the mixer’s intrinsic noise floor, I now turn to the optical system we use. Figure 4.9 shows the experimental setup we use to both reduce the contributions of technical noise and measure the excess phase noise associated with the optical-to-microwave conversion processes of photodetection and amplification, as well as the phase noise measurement using a mixer. In the experiment, we attempt to convert the optical timing of a pulse train from a 760-MHz Ti:sapphire laser into a microwave frequency at the 10th harmonic, or at 7.6 GHz, with as little excess phase noise as possible. The Ti:sapphire laser is left free-running for these measurements, since in the case of negligible microwave dispersion in the electronics, its frequency fluctuations are common-mode and therefore should not lead to measured phase noise.

One potential problem associated with photodetection of the pulse train is the pointing fluctuations of the beam incident on the photodetector. Therefore we use a short (0.5-m) piece of single-mode fiber just in front of the optical beamsplitter in order to both clean up the mode and reduce the distance the beam travels between the physically stable output of the fiber and the photodetectors. The fs pulse train is split on a polarizing beamsplitter so that two photodetectors can independently detect the 7.6-GHz harmonic of the repetition frequency. As seen in Section 4.2.2, amplitude noise of the microwave signals can lead to measured phase noise through amplitude-to-phase conversion. In a similar way, amplitude-to-phase conversion at the photodiode can occur as discussed in Section 2.2.2. For both reasons, we attempt to stabilize the optical intensity of each beam. In Ref. [76], where a similar system is used to measure the excess noise of optical-to-microwave conversion, the intensity at one photodiode was stabilized by feedback to an AOM before the beamsplitter. However, because of the dependence of the beamsplitter’s splitting ratio on pointing, stabilizing the intensity

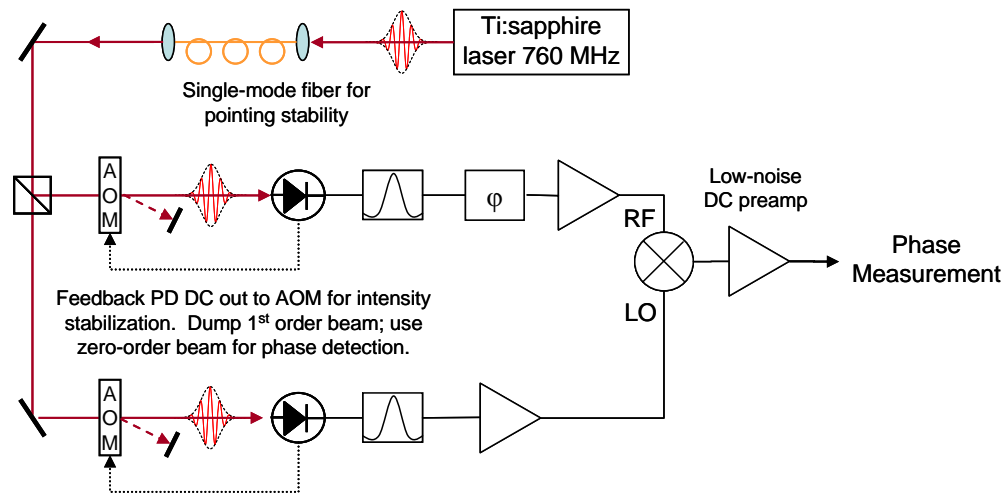


Figure 4.9: Experimental setup for reducing and measuring the excess noise in detection of microwave harmonics of a 760-MHz Ti:sapphire laser. The pulse train is first coupled into a short (0.5-m) single-mode fiber in order to minimize fluctuations of pointing on the photodetectors. After splitting a few mW toward each detector, independent AOM's in each path are used to stabilize the intensity of the beams. The AOM's dump a small fraction of the power into their first-order diffracted beams, leaving an intensity-stabilized zero-order beam for photodetection. Feedback to the AOM's is by comparing each photodetector's DC voltage output to a low-noise voltage reference. The 7.6-GHz 10th harmonic of the repetition frequency from each photodetector is then input to the phase noise comparison system from Fig. 4.3 in order to measure the excess phase noise associated with the independent photodetection, amplification, and mixing scheme.

on one photodetector had the effect of doubling the intensity fluctuations on the other. Therefore we choose to independently stabilize the intensities of the two optical beams after the photodiode. Although it is still an open question as to whether the rf power in the particular microwave harmonic should be stabilized, or whether it is sufficient to use the dc photocurrent, we chose to use the latter. Therefore the dc photocurrent from each detector is compared to a stable voltage regulator, and the resulting error signal is used to feed back to each AOM's rf driving power in order to dump more or less optical power into the diffracted orders. The zero-order beam is then photodetected.

Since the amount of microwave power available from each photodetector will determine the amount of phase noise from thermal and shot noise, as much microwave power as can be generated is desirable. Typical commercially available photodetectors have a 50-Ohm resistor wire-bonded to the diode, in order to provide wideband impedance matching to 50-Ohm electronics. However, this represents a voltage divider and so only one-fourth of the available power from photocurrent gets to the subsequent electronics. We therefore chose to use a photodetector without such a 50-Ohm resistor, in favor of using a microstrip layout in a resonant LC tank configuration to impedance match the photocurrent to subsequent amplifiers at our desired frequency of 7.6 GHz. Again, Terry Brown's design of the LC impedance-matching network worked as advertised – the available microwave power for the same optical incident power increased by 8 to 9 dB relative to the 50-Ohm resistor system. With just a few mW of average power available from a 760-MHz Ti:sapphire laser, we have achieved microwave power of -5 dBm at 7.6 GHz directly from the photodiode and impedance matching network. The photodetector is made by Electro-Optics Technology, model ET-4000, and they are more than happy to sell it without the wire-bonded 50-Ohm resistor. These impedance-matching methods are a simple way to push the fundamental phase noise floor limitation down by many dB for a given optical power incident on the photodiodes.

4.4 Comparison of results

Figure 4.10 shows a comparison of our various measurements to some of the best measurements made by others of the phase noise of microwave signals from femtosecond lasers. The light blue data labelled “Low optical power” and the light green data labelled “High optical power” were measured using the setup shown in Fig. 4.2, without any intensity stabilization of the optical beams. Average powers of 0.2 and 2 mW were used on each photodetector for the respective measurements. With the higher average power, the shot noise floor was reduced, but extremely structured noise from 100 Hz to 1 kHz was present due to amplitude-to-phase conversion. This was confirmed by measuring the power fluctuations at the output of the Ti:sapphire laser with a low-speed photodetector, which revealed intensity fluctuations with the same spectrum as the phase noise measurement.

By applying the intensity stabilization described in Section 4.3 and shown in Fig. 4.9, at a similar average power of 1.4 mW, the black data labelled “With amplitude stable” was achieved. With the intensity fluctuations greatly reduced, less noise between 100 Hz and 1 kHz was present, even for such a high optical power. We also measured the noise floor of the electronics without photodetectors as shown in Fig. 4.3 and discussed in Section 4.2, being careful to simulate the microwave powers available from the photodetectors. The result is shown in red, labelled “Full RF noise floor.” The comparison between the red “Full RF noise floor” data and the black “With amplitude stable” data demonstrates that the full optoelectronic system is essentially limited by the electronics, not the photodetectors, at these power levels. As discussed in Section 4.2.2, the mixer’s flicker noise is responsible for the flicker floor limit below 10 kHz, while the high frequency limit is due to shot noise from the optical power we can apply to the photodetector. When taking the “With amplitude stable” data, -12 dBm of microwave power was available from both photodetectors at 7.6 GHz. The noise floor

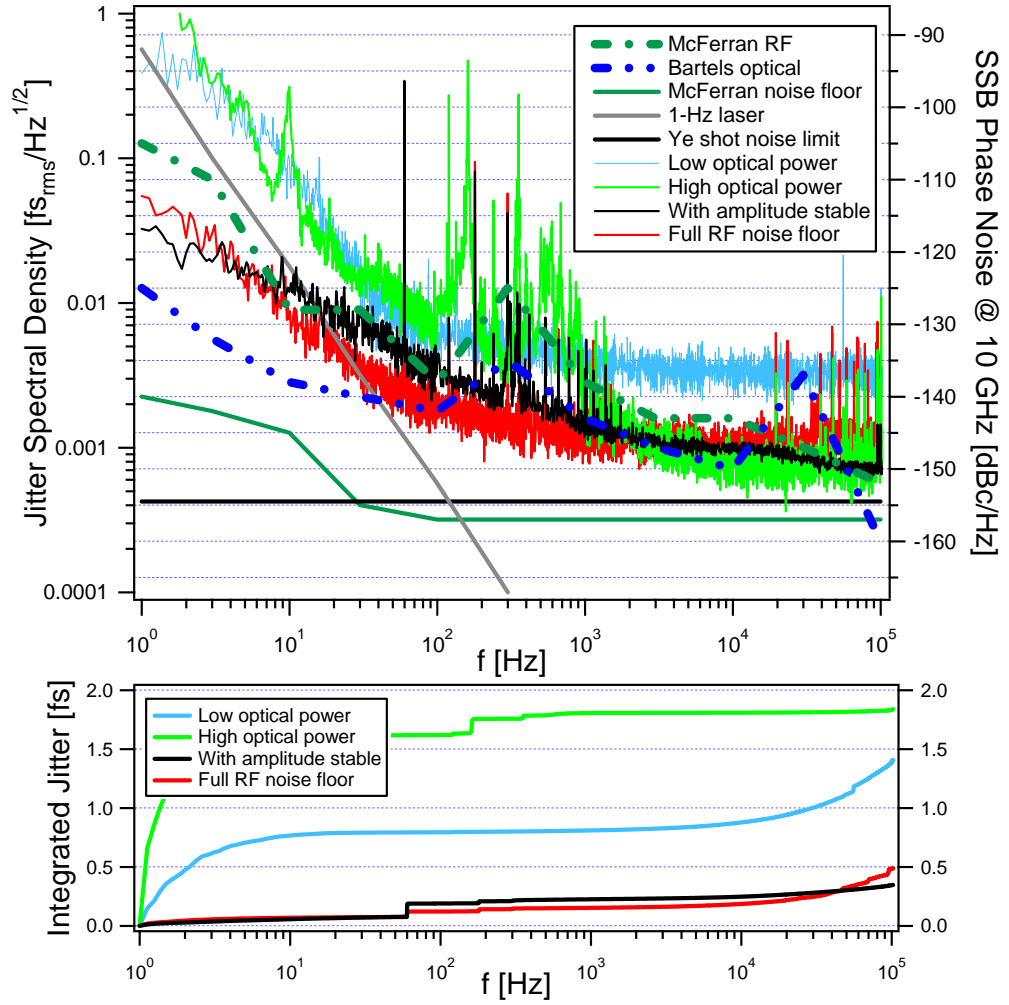


Figure 4.10: Several results of various experiments to detect the phase noise of microwave signals generated by detecting harmonics of f_{rep} from femtosecond combs. The traces labelled “Low optical power” and “High optical power” are from the detection scheme shown in Fig. 4.2, with 0.2 and 2 mW of average power per detector, respectively. The black data labelled “With amplitude stable” is from a similar system, where amplitude stabilization of the optical beams has been employed prior to photodetection. This leads to a reduced level of amplitude-to-phase conversion in the 100-Hz to 1-kHz regime. The “Full RF noise floor” is measured with an identical system, but replacing the two photodetectors with a common synthesizer, simulating the powers that would be available from the photodetectors. This demonstrates that the photodiodes, with amplitude stabilization, are at least as quiet as the microwave system. For comparison, “McFerran RF” and “Bartels optical” show the microwave and optical phase noise of 10-GHz harmonics of f_{rep} for some of the best systems yet measured. For references, see the text. Also for comparison, the optical phase noise from a sub-Hz ultrastable laser is shown, scaled to a 10-GHz carrier.

we reach near 100 kHz is only 4 dB higher than one can calculate for this amount of photocurrent in the system. Note from Fig. 2.1(b) that with the 1.4 mA we used, the phase noise contribution from thermal noise is only slightly less than the phase noise contribution from shot noise. Therefore it is reasonable to believe that the noise floor we achieve is limited by a combination of both thermal and shot noise. Higher incident optical powers will be necessary to further push these limitations down, and with the intensity stabilization we should be able to apply significantly more power (approaching 10 mW) to the photodetectors while keeping amplitude-to-phase noise to a minimum.

For comparison, I also show two dashed lines representing the performance levels achieved by other systems. The green dashed line labelled “McFerran RF” is taken from Ref. [64]. In their system, microwave interferometric techniques were employed in order to reduce the effective noise from flicker of the microwave amplifiers. As a result, the noise floor of their measurement (shown as a solid green line labelled “McFerran noise floor”) was exceedingly low, but other noise sources still led to the performance shown by the dashed green line. It is important to note that we achieve similar performance without the need for the complexity of microwave interferometry, due to the excellent design of our low-phase-noise amplifiers. The blue dashed line was measured in Ref. [78] using optical methods as opposed to microwave methods, in order to measure the excess phase noise of a femtosecond laser’s optical pulse train when it locks to a stable cw laser. By locking two femtosecond lasers to the same cw laser, and then optically (as opposed to with the optical-to-microwave phase detection methods described in this Chapter) detecting the heterodyne beat (at the difference between their offset frequencies, Δf_0) between the two frequency combs in a narrow spectral region, the phase noise of the optical pulse train was measured. Our best optical-to-microwave conversion’s phase noise results are approaching this intrinsic optical performance of the Ti:sapphire lasers at low frequencies (below 1 kHz) and exceeding the performance in the 1-to-100-kHz range.

Integrating our intensity-stabilized system's timing jitter spectral density from 1 Hz to 100 kHz gives an rms jitter of 0.4 fs, significantly less than the duration of an optical cycle. Also, this contribution presumably represents the incoherent (quadrature) sum of independent noise from each optical-to-microwave conversion arm of the experiment. Therefore, a reasonable claim is that the performance of either photodetector-plus-amp system has $\sqrt{2}$ less noise than shown in Fig. 4.10. Such low levels of timing jitter pave the way for synchronization experiments that preserve the coherence between two fs lasers as described in Chapter 5.

Finally, I have included a grey solid line labelled "1-Hz laser" to indicate the phase noise performance of the ultrastable clock lasers that will be described in Section 7.1.3. At the lowest Fourier frequencies (below 10 Hz) our optical-to-microwave conversion can take advantage of the ultrastable performance of such a laser, but clearly at higher Fourier frequencies there is much work to be done! The measurements described in this Chapter have focused mostly on pushing down the thermal and shot noise floors by turning up the photodetected optical power. It will be equally important to measure the stability of optical-to-microwave conversion at long timescales when considering whether it performs at a level suitable to preserve the stability of next-generation optical clocks. Such a microwave readout is indeed a necessary component of the clock for time-keeping; since the optical cycles can not be directly counted, optical-to-microwave conversion that preserves the coherence and stability of the optical frequency reference will allow the generation of a countable microwave clock output!

Chapter 5

Synchronization of Femtosecond Lasers

At the time I began working at JILA with Jun, the merger of ultrafast laser techniques and precision frequency metrology had already resulted in recent dramatic progress in a number of fields, including optical frequency measurement [104, 110], optical clocks [117, 118], and carrier-envelope phase control [98, 119]. An exciting area of research is “coherent control”, which employs pulses of light shaped in amplitude and phase to selectively [25] drive chemical reactions [26], molecular vibrations [27], or the nonlinear up-conversion of light into the extreme ultraviolet region [28] to name a few examples. Ideally, coherent control would need phase-coherent, temporally shaped-pulses over a wide spectral range (therefore of femtosecond-scale duration), so that all possible intermediate states in a multistep quantum pathway could be accessed in order to selectively drive a particular outcome. Some nonlinear techniques are available for generation of coherent light over a broad spectrum, such as white-light continuum generation [120, 99] and parametric amplification [121]. However, it would be more flexible and more efficient to bring two or more independently tailorable pulsed laser systems together, and precisely synchronize as well as phase-stabilize them so as to form a coherent field. Synchronization, therefore, is the first step toward this goal. Many other technologies also benefit from precise synchronization, including rapidly tunable near-UV and mid-IR generation by nonlinear mixing [24, 122] as discussed in 5.2, synchronization of laser light and x rays or electron beams from synchrotrons [41],

pump-probe experiments such as coherent anti-stokes raman spectroscopy (CARS) [19], and synthesis of light pulses shorter than can be achieved with a single laser [123].

Jun's group had already demonstrated that combined time and frequency active stabilization can allow one to tightly synchronize two separate, passively mode-locked femtosecond lasers [20]. The remaining rms timing jitter during an observation time of tens of seconds was shown to be less than 30 fs within a 50-kHz bandwidth and less than 5 fs within a 160-Hz bandwidth. This capability has allowed phase locking between the carrier waves of the two synchronized Ti:sapphire lasers [22]. However, since the characteristic timing jitter was larger than the optical cycle period, phase locking could occur for only a few tens of microseconds at a time. Coherent optical pulse synthesis from these lasers was subsequently demonstrated [22].

The ultimate goal of this research would be to demonstrate an arbitrary light waveform generator capable of synchronizing and phase locking arbitrary, separate mode-locked lasers with distinct optical properties. It is thus desirable to enhance the precision level of synchronization such that the remaining timing jitter between different lasers would be less than the oscillation period of the optical carrier wave, namely, 2.7 fs for Ti:sapphire lasers centered around 800 nm. One possible approach uses cross-phase modulation to passively synchronize two mode-locked lasers that share the same intracavity gain medium [124]. However, the requirement of sharing an intracavity element limits the flexibility and general applicability of this technique. Other techniques are to construct a balanced optical cross-correlator [80] or lock on the fringes of an interferometric cross-correlation [82] as discussed in Section 2.3. Both of these optical methods have greatly relaxed requirements on the electronic noise as compared to the microwave method I will describe in Section 5.1.2, but they also suffer from a limited dynamic range of the relative timing delay between the two pulse trains. A final method is to stabilize both femtosecond lasers to a common cw optical reference, and at the same time stabilize their relative offset frequencies [125]. This method has the added com-

plexity of an additional laser and the relative offset frequency lock as compared to the method of Section 5.1.2, but automatically results in relative optical phase stabilization of the pulse trains in addition to just synchronization of the repetition frequencies.

In this Chapter I will describe a system using two independent femtosecond Ti:sapphire lasers to demonstrate synchronization of their pulse trains [21]. We achieved fs-level rms timing jitter between the two lasers, with the ability to rapidly and repeatedly switch (60- μ s settling time, 1-fs accuracy) between arbitrary time delays between the pulse trains, up to 10-ps. In a further experiment, we used the synchronized lasers to simultaneously generate rapidly tunable mid-IR (7.5 to 12.5 μ m) and near-UV (400-nm) fs pulses with more than 10 μ W of average power in the infrared [24]. For the first time, synchronization between two independently adjustable femtosecond lasers was achieved, satisfying the requirement that the residual timing jitter be below an optical cycle's duration, opening the door for phase-coherently stitching lasers together across the visible and infrared spectrum.

5.1 Synchronization experiment

5.1.1 The laser system

We employ two low-threshold Kerr-lens mode-locked Ti:sapphire lasers that each produce a mode-locked average power of more than 300 mW with 2.5-W pump power at 532 nm [126]. The beam height inside the laser cavities is only 6 cm above the baseplate, which should enhance stability. We find that to achieve the lowest possible timing jitter it is critical to enclose both lasers and their pump beams to shield them from blowing dust and convection currents. The two lasers are both located on a temperature-controlled, 5-cm-thick, solid aluminum baseplate, which is decoupled from the vibration noise of the table by a set of supporting feet made from rubber. In addition, asymmetrically positioned lead plates are tightly bonded to the bottom side

of the baseplate to damp the vibration modes. The vibration noise measured on top of the baseplate is generally reduced by 10 dB or more within the frequency range of 100 Hz to 6 kHz compared with the noise on the optical table.

The two Ti:sapphire lasers are pumped with a Coherent Verdi V-5 solid state pump laser, with a 5.5-W output power split equally between them. Each fs laser produces ~ 200 mW of average mode-locked output power, at a repetition frequency of 100 MHz. The design of each laser is essentially the same as in Fig. 3.1, employing prism pairs for intracavity dispersion compensation. By adjusting the prism pairs, the output spectra can be tuned 20 to 40 nm to either side of 800 nm. Although the bandwidth of either laser corresponds to a transform limit of < 20 fs pulse durations, the pulses are chirped after exiting each output coupler to roughly 100 fs, depending on the prism pair alignment.

For these experiments, the master fs laser is left free-running, while the slave fs laser is actively synchronized to the master. The servo action on the slave laser is carried out by a combination of transducers, including a fast piezo-actuated small mirror (4 mm in diameter and 2.5 mm thick), a regular mirror mounted on a slow piezo with a large dynamic range (2.7 mm), and an acousto-optic modulator placed in the slave laser's pump beam to help with the fast noise. The unity-gain frequency of the servo loop is ~ 200 kHz, and the loop employs three integrator stages in the low-frequency region.

5.1.2 Differential timing resolution for the servo

To synchronize the two lasers, we use two phase-locked loops (PLLs) working at different timing resolutions. Figure 5.1(a) shows a diagram of the feedback system. As seen directly in Eqn. 2.4, for a given level of phase noise due to the electronic phase detection system, the timing jitter spectral density noise floor scales as the inverse of the microwave carrier frequency. Therefore, in order to relax the requirements on the electronics, detection of a high harmonic of f_{rep} must be employed; for this experiment

we used the 140th harmonic at 14 GHz. However, locking the two 14-GHz harmonics of the lasers' repetition rates together leads to a random timing offset between the two pulse trains. We overcome this by use of a coarse PLL to compare and lock the fundamental repetition frequencies (100 MHz) of the two lasers. A phase shifter operating on one of the two 100-MHz signals can be used to control the coarse timing offset between the two pulse trains with a full dynamic range of 5 ns. The second, high-resolution, PLL compares the phase of the 140th harmonic of the two repetition frequencies, i.e., at 14 GHz. This second loop provides enhanced phase stability of the repetition frequency when it supplements and then replaces the first PLL. A transition of control from the first to the second PLL can cause a jump in the timing offset by at most 35.71 ps (one half of one 14-GHz cycle), whereas the adjustable range of the 14-GHz phase shifter is 167 ps.

Figure 5.1(b) shows this transfer of control between the two PLL's explicitly. In the first panel, only the 100-MHz coarse PLL is activated. The slave laser can be adjusted temporally with respect to the master by tuning the 100-MHz phase shifter, as shown in the second and third panels. As in the third panel, once the pulses are within 150 ps of each other, the gain of the coarse PLL is slowly decreased while the gain of the fine 14-GHz loop is increased. After this smooth transition from coarse to fine PLL, a 14-GHz phase shifter is used to finish aligning the pulse trains in time, to the desired time delay. The whole system can be thought of as a “differential micrometer” in time, for the two pulse trains — the coarse servo allows for the full dynamic range in timing offset between the pulse trains, and the fine servo provides greater timing resolution and therefore greater phase stability for the synchronization.

5.1.3 Synchronization performance

To characterize the timing jitter, we focus the two pulse trains so that they cross in a thin β -Barium-Borate (BBO) crystal cut for sum frequency generation (type I SFG).

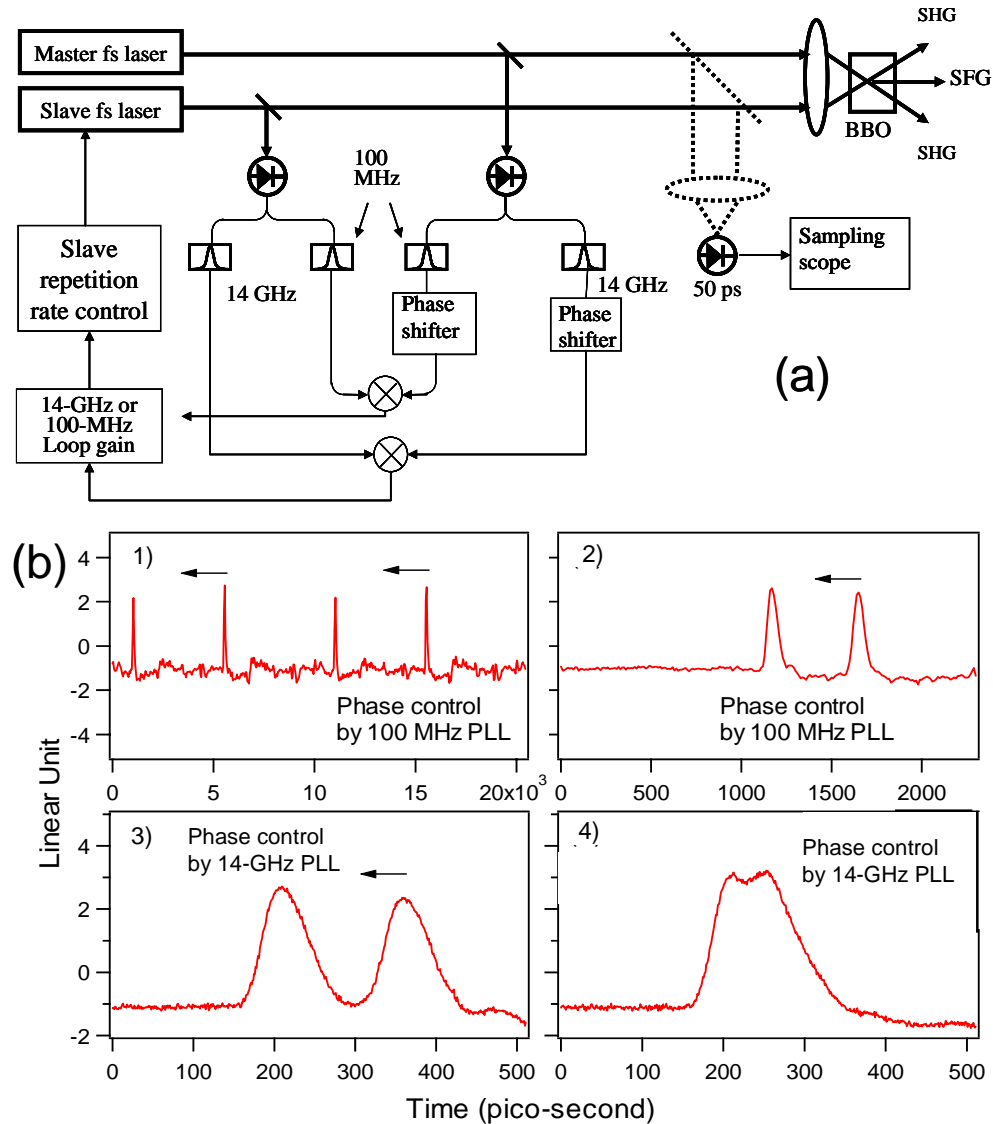


Figure 5.1: (a) Experimental setup for synchronizing the pulse trains from two independent fs lasers. Two servos operating with differential timing resolutions allow tight synchronization and an arbitrary time delay between the pulse trains. The servos feed back to a slow piezo-actuated mirror with large dynamic range, a fast piezo-actuated mirror, and an AOM in the slave laser's pump beam, resulting in a servo bandwidth of ~ 200 kHz. Characterization of the synchronization performance is made with an auto-correlation by sum frequency generation in a nonlinear BBO crystal. (b) A sampling scope is used to monitor the coarse time delay between the pulse trains. At first, the 100-MHz PLL is locked, allowing the two pulse trains to be tuned a full dynamic range of 5 ns. In the second panel shown, the pulse trains are coarsely overlapped using a 100-MHz phase shifter in the coarse PLL. The third step is to shift control from the 100-MHz coarse PLL to the 14-GHz fine PLL. Finally, the 14-GHz phase shifter is used to fine-tune the timing offset until SFG is observed, meaning the pulses are overlapped temporally in the BBO crystal.

The crossed-beam geometry produces a non-fringe-resolved SFG cross-correlation signal in between SHG signals from the two input beams, as shown in Fig. 5.2(b). Note that the SFG signal is proportional to the square of the product of the two incident fields; therefore it will only be observable when the pulses are overlapped in space and time in the BBO crystal. Figure 5.2(a) shows the two SHG signals that are always present, when the two pulse trains are not temporally overlapped. Figure 5.2(b) shows the case when the pulse trains do overlap in time, revealing the central SFG beam. By spatially filtering to detect only the SFG beam, the Gaussian cross-correlation peak (obtained when the two lasers are free running) is ~ 160 fs FWHM. (No extracavity dispersion compensation is used, so the would-be 20-fs laser pulses are broadened.) As shown in Fig. 5.2(c), the SFG signal strength depends on the temporal overlap of the two incident pulses, and can therefore be used to precisely measure the timing jitter between the pulse trains. The time axis of the cross correlation is calibrated by measurement of the difference between the two laser repetition rates when the cross correlation is recorded. The calibrated slope of the cross-correlation signal near the half-maximum can be used to determine the relative timing jitter between the two lasers from the corresponding intensity fluctuations. We record the intensity fluctuations of the SFG signal over a period of several seconds, using two different low-pass bandwidths, 160 Hz and 2 MHz, which suppress the pulsed nature (100 MHz) of the SFG signal amplitude and permit the study of intensity noise on a cw basis.

The top trace in Fig. 5.3(a) shows that the SFG signal, recorded with a 2-MHz bandwidth, shows only digitalization noise when the two laser pulses overlap maximally (at the top of the cross-correlation peak). In this case, the SFG signal is least sensitive to the timing jitter and its intensity noise is limited by each individual laser's amplitude fluctuation. The two lower traces are recorded with 2-MHz and 160-Hz bandwidths when the timing offset between the two pulse trains is adjusted to yield the half-intensity level of the SFG signal. Timing jitter is calculated from the intensity noise by use of




Figure 5.2: (a) When the two pulse trains cross through the BBO crystal, each pulse train undergoes second harmonic generation. (b) If the two pulses are also temporally overlapped in the crystal, they mix with each other, forming sum-frequency light between the SHG beams. (c) The central SFG signal can therefore be used to measure the timing jitter, by adjusting the time delay between the two pulses such that the SFG signal is at its half-maximum. Then any timing fluctuations between the pulse trains will be measured as variations in the SFG intensity.

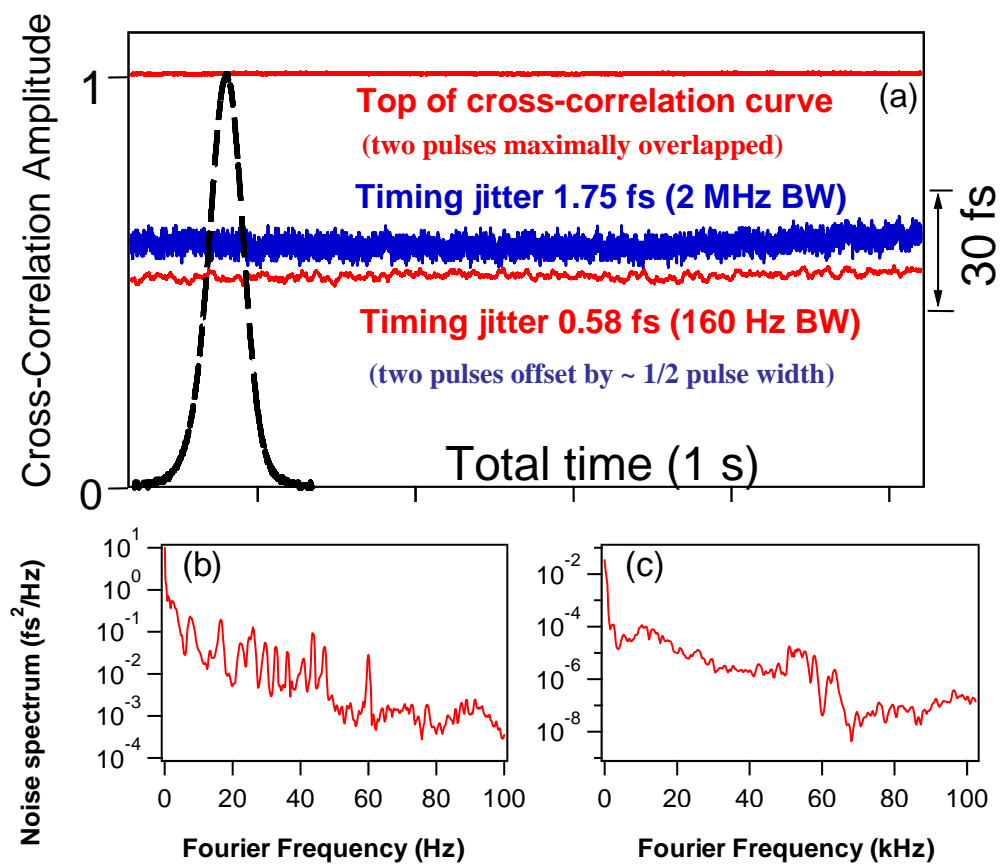


Figure 5.3: (a) One-second-long trace showing the cross-correlation signal when the pulse trains are synchronized. The top trace shows the signal when the pulses are maximally overlapped in time. The middle trace and bottom trace show rms timing jitters of 1.75 fs and 0.58 fs in bandwidths of 2 MHz and 160 Hz, respectively. For clarification, the cross-correlation curve is shown at left. (b) SFG intensity-noise power spectral density from dc to 100 Hz. (c) SFG intensity-noise power spectral density from dc to 100 kHz.

the slope of the correlation peak, with the scale of the jitter indicated on the vertical axis of Fig. 5.3(a). The rms timing noises are thus determined to be 1.75 fs at a 2-MHz bandwidth and 0.58 fs at a 160-Hz bandwidth. For detection bandwidths above 2 MHz, the observed jitter does not increase. We have recorded such stable performance over several seconds. The synchronization lock can be maintained for durations of several hours. However, the intensity stability of the SFG signal is found to correlate strongly with the temperature variations in the microwave cables.

To understand noise contributions at various time scales, we study the SFG signal with a fast Fourier-transform frequency spectrum analyzer. Figures 5.3(b) and 5.3(c) display the power spectral density (resolution bandwidth normalized) of the timing noise up to Fourier frequencies of 100 Hz and 100 kHz, respectively. The spikes near dc are artifacts from the analyzer itself. When the spectral noise from Fig. 5.3(b) is integrated from 1 Hz to 100 Hz (avoiding the near-dc artifacts), we calculate a timing jitter of ~ 0.6 fs. When the data set in Fig. 5.3(c) is integrated from 1 Hz to 100 kHz, the resultant rms timing jitter is 1.4 fs. These results are consistent with the direct time-domain data shown in Fig. 5.3(a).

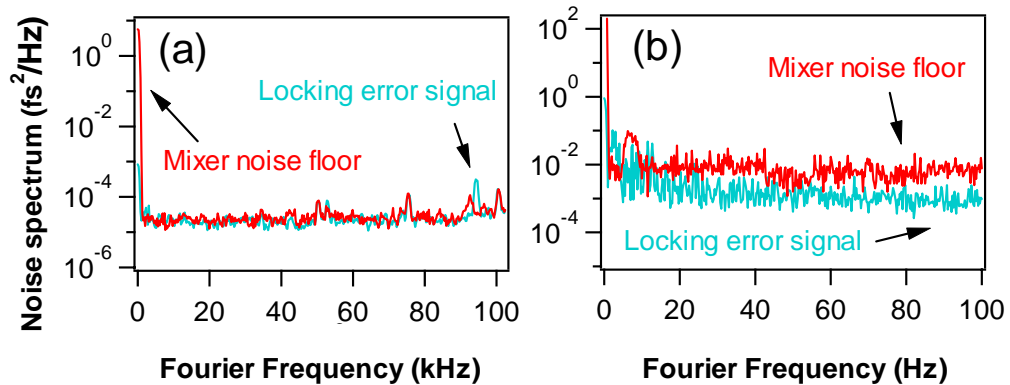


Figure 5.4: (a) Power spectral density of the servo error signal and the intrinsic noise of the mixer from dc to 100 kHz and (b) from dc to 100 Hz. These results show that the performance was limited by the intrinsic noise floor of the microwave electronics. It was later determined the amplifiers were responsible.

Since the SFG signal represents a measure from outside the servo loop, the intensity analysis of this signal reflects the true evaluation of the relative timing jitter between the two synchronized lasers. A careful study of the servo error signal inside the feedback loop reveals that a major part of the limitation to the current attainable level of performance is actually due to the intrinsic noise of the 14-GHz phase detector, a double-balanced mixer. Figure 5.4 displays the power spectral density of the error signal under tight lock, along with that of the mixer noise. The microwave electronic system's phase-noise power spectral density was measured to be $2 \times 10^{-3} \text{ rad}^2/\text{Hz}$ at 14 GHz. This phase-noise floor corresponds to a timing noise of $2.5 \times 10^{-5} \text{ fs}^2/\text{Hz}$. Figure 5.4(a) shows that the synchronization loop error signal and the microwave system's noise floor basically overlap across the entire 100-kHz Fourier frequency range. At low frequencies, the servo loop gain is actually slightly too high, such that the noise level of the servo error becomes less than the microwave electronics' noise (Fig. 5.4(b)). Integration of this (resolution-bandwidth-normalized) intrinsic noise level produces the rms timing jitter noise floor limit for the synchronization loop. For example, from Fig. 5.4(b) we determine the rms timing jitter limit for the 1-Hz to 160-Hz frequency range to be $(2.5 \times 10^{-3} \text{ fs}^2/\text{Hz} \times 160 \text{ Hz})^{1/2} \approx 0.63 \text{ fs}$.

For practical ultrafast applications, it is important to show that this low level of timing noise can be obtained at any predetermined timing offset between two laser pulse trains. It is also desirable to realize a reliable setting of the timing offset at high speed and with excellent repeatability. Figure 5.5 demonstrates such capabilities of our synchronization system. Figure 5.5(a) shows fast switching of the pulse delay time by $\sim 78 \text{ fs}$, indicated by a jump in the SFG signal. The switching signal applied to the synchronization PLL is also shown. Associated with servo bandwidth limitations, there is a 6.3-ms delay between the command signal applied to the servo loop and the actual jump in the pulse timing offset. The servo attack time is $\sim 10 \text{ ms}$ and the pulse timing is settled in $\sim 60 \text{ ms}$. The precision level of the switched pulse timing is $\sim 1 \text{ fs}$. For

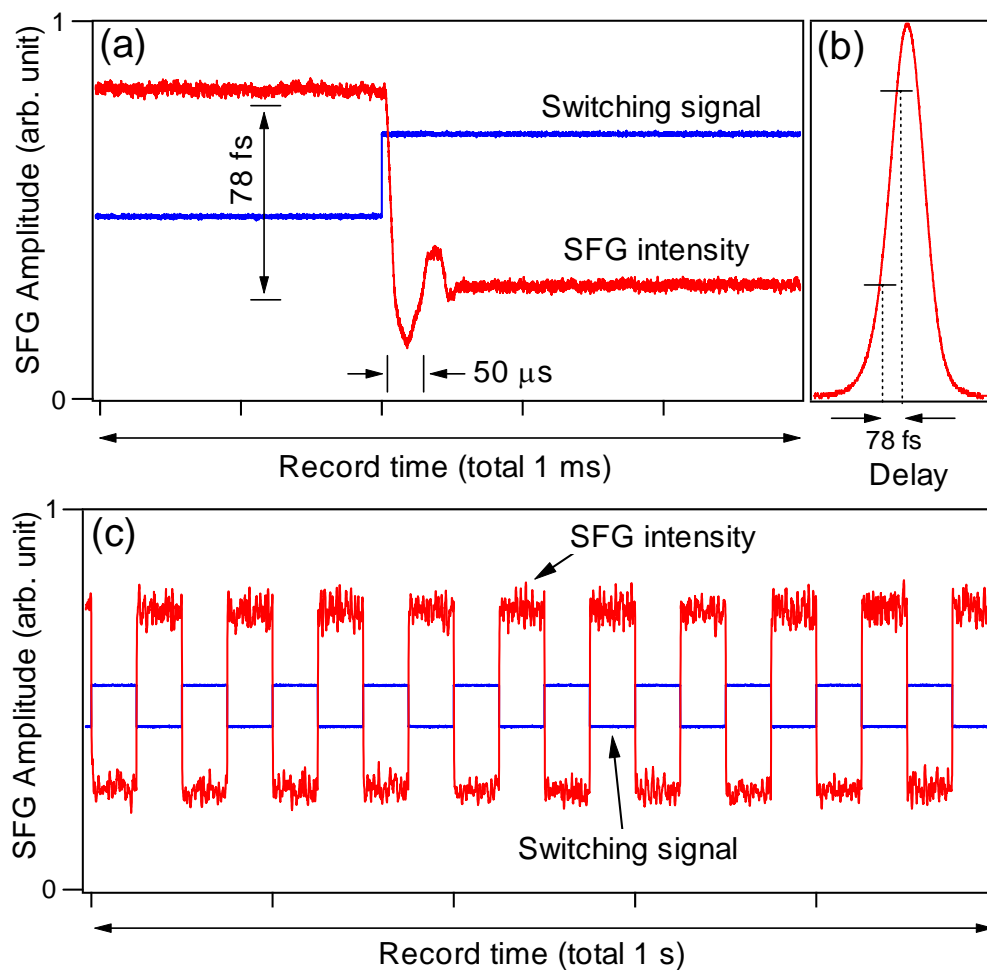


Figure 5.5: Controlled switching of time delays between the two pulse trains. (a) Switching dynamics showing a 6.3-ms delay, a 10-ms servo attack time, and ~60-ms settling time in the real signal after the control signal in the feedback loop is applied. (b) Cross-correlation signal between the two pulse trains. (c) 10-Hz switching sequences recorded in a 1-s period.

clarity, we also show the magnitude of the cross-correlation signal in Fig. 5.5(b). Figure 5.5(c) displays a 1-s record of continuous switching at a frequency of 10 Hz between two predetermined pulse delay times. Again, repeatability in timing offset ~ 1 fs is demonstrated. Compared with a mechanical scanning system, this electronic tuning method thus offers vastly superior performance in terms of repeatability, reliability, and speed, with no noticeable hysteresis.

The all-electronic control of the timing offset between separate lasers is flexible and generally applicable. Combined with the ultralow level of the timing noise, the techniques that we present here should have a wide range of applications. Further reduction of the timing noise can come from detection of even higher harmonics of the repetition frequency. Improvements to the optical-to-microwave detection of harmonics of the repetition rate, as discussed in Chapter 4, have the potential to decrease the rms timing jitter to below 0.4 fs integrated in a 100-kHz bandwidth. However, we will quickly run out of the practical range of microwave technology. An attractive alternative would be to use optical means. One can use either a single highly stable cw laser or a common stable optical resonator to control a high-order harmonic of the repetition frequency, well into the terahertz or tens or hundreds of terahertz frequency range, as discussed in the introductory Section of this Chapter.

5.2 Application: infrared fs pulses

After having demonstrated such tight synchronization between independent femtosecond lasers, a natural application for the technology was to generate infrared fs pulses by difference frequency generation (DFG) from the two pulse trains in a nonlinear crystal. Mid-infrared (MIR) ultrashort light pulses are of primary importance for the study of ultrafast dynamics in chemical reactions, molecular vibrations, and the application of femtosecond IR spectroscopy to problems in solid-state physics. In particular, experiments in coherent control require a flexible, tunable source of MIR

femtosecond pulses. New techniques for controlled pulse shaping in the MIR [127, 128] in combination with such adaptable sources promise to make possible abundant progress in the science of coherent control.

High-repetition-rate sources are desirable for high-speed acquisition of data and signal averaging. Two methods have been employed for the generation of MIR pulses at high repetition rates, usually by frequency mixing in nonlinear crystals such as AgGaS₂, AgGaSe₂, or GaSe. In one type of experiment, 800-nm pulses directly from a Ti:sapphire laser are used. An 82-MHz two-color Ti:sapphire laser mixed in AgGaS₂ yielded 5-mW pulses tunable from 7.5 to 12.5 μm [129]. Mixing two ends of the spectrum from an 88-MHz broadband Ti:s laser produced 9- to 18- μm femtosecond pulses with an average power of 1 mW [130]. A 1-MHz cavity-dumped Ti:s laser employed the same technique to produce 5 mW of average power that was tunable from 7 to 20 μm [131]. These methods are limited by the fixed amount of power available in the spectral region accepted by the crystal for the mixing process and are tunable only via nonlinear crystals' phase-matching angles. A second type of experiment involves the use of an optical parametric oscillator system, in which the signal and idler waves, centered near 1.6 μm , are used for the difference-frequency generation (DFG) process. An 82-MHz (76-MHz) optical parametric oscillator system generated 500 mW (2 mW) of average power that was tunable from 2.5 to 5.5 μm (5.2 to 18 μm) [132, 133]. These systems have the advantage of producing a larger MIR power at the cost of increased complexity and greater initial pump power.

In the work I performed [24], we took advantage of the synchronized, independently tunable lasers discussed in the previous Section. The individual tunability of the two Ti:sapphire lasers would in principle make possible DFG generation with spectral coverage from 3 μm down to the terahertz region. The pulse trains, operating independently at separate wavelengths, were used for DFG in a GaSe crystal. Tuning the crystal's phase matching angle as well as the wavelength separation of the two input

pulses generated MIR pulses from 7.5 to 12.5 μm , with the measured upper (lower) tuning range limited by the MIR monochromator (the nonlinear crystal). A direct measurement of the MIR pulse duration shows 670-fs pulses at a wavelength of 11.5 μm . The average power of the pulses exceeds 15 μW , which at the time represented the highest power in the 7.5- μm spectral range obtained from direct DFG of Ti:sapphire pulses. More importantly, precision setting of the time delay between the two original chirped-pulse trains allowed rapid switching of the MIR wavelength as well as programmable amplitude modulation, both extremely useful features for experiments in coherent control.

5.2.1 Experimental method for simultaneous DFG and SFG

A schematic of the setup is depicted in Fig. 5.6. The same two fs lasers from Section 5.1.1 were operated at center wavelengths 20-40 nm on either side of 800 nm, with (uncompressed) pulse durations of ~ 140 fs. The shorter-wavelength laser (carrier frequency ω_{SW}) served as the free-running master, whereas the longer-wavelength laser (carrier frequency ω_{LW}) was operated as the slave. The synchronization technique was identical to the methods described in Section 5.1.2.

The two Ti:sapphire beams were combined collinearly with orthogonal polarization into two arms. In one arm of the experiment, the pulse trains were focused into a thin BBO crystal cut for type I SFG, rotated 45° so that polarization components of both pulse trains propagate as ordinary beams in the crystal. The resultant SFG signal was used to characterize the timing jitter of the pulse trains when the timing offset corresponded to half a pulse width, as described in Section 5.1.3. In the other arm of the experiment, the pulses were given a slight negative chirp by use of two negatively chirped mirrors and focused ($f = 12$ cm) onto a GaSe crystal for DFG.

The 1-mm-thick uncoated GaSe crystal is cleaved normal to its z axis, since it is not yet possible to polish under arbitrary angles. This results in external incident angles

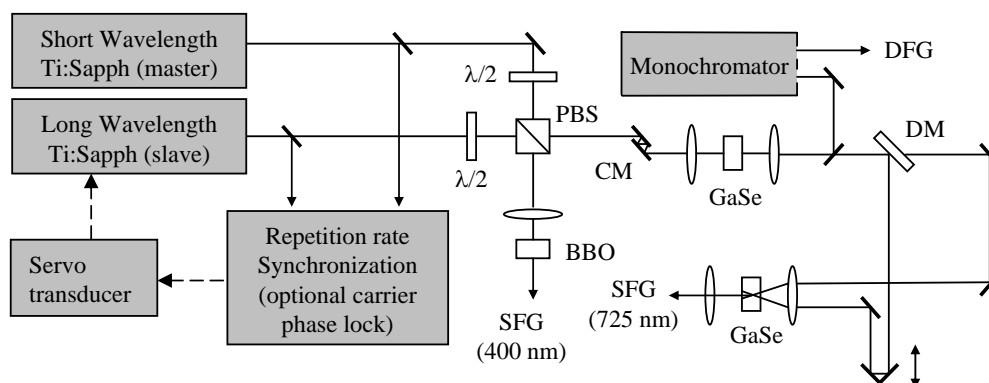


Figure 5.6: Experiential system used to generate and characterize femtosecond mid-IR pulses: Two fs Ti:sapphire lasers with separate center wavelengths are tightly synchronized. The pulse trains are combined into two paths, one for sum-frequency generation in BBO (β -barium borate) and one for difference-frequency generation in 1-mm thick GaSe. An IR monochromator is used to characterize the spectra of the resulting mid-IR pulses, which are tunable from 7.5 to 12.5 μm . For measurement of the IR pulse-width, the short-wavelength Ti:sapphire is combined with the IR pulses into a cross-correlator based on sum-frequency generation. CM, chirped mirror pair; DM, dichroic beam splitter; PBS, polarizing beam splitter

of $51\text{-}74^\circ$ for type I phase matching ($e - o \rightarrow o$ for $\omega_{SW} - \omega_{LW} \rightarrow \omega_{IR}$) of the 800-nm light. Even though such steep external angles entail large ($> 40\%$) reflection losses at the surfaces of the crystal, GaSe exhibits a higher nonlinear susceptibility and a wider transparency range than crystals such as AgGaS_2 or AgGaSe_2 . Recent research with indium-doped GaSe also indicates a developing capability of producing arbitrary-cut GaSe crystals as well as antireflection coating [134].

After the GaSe crystal, the MIR light was separated from the two collinear near-IR beams with a long-pass filter. The MIR beam's intensity was then measured by use of a liquid-nitrogen-cooled HgCdTe detector. An IR monochromator was sometimes used before the HgCdTe detector for spectral measurements of the MIR pulses. For temporal measurements of the MIR pulses, the shorter-wavelength Ti:sapphire and the MIR DFG pulses were mixed in a second GaSe crystal by use of type I phase matching ($o + o \rightarrow e$ for $\omega_{IR} + \omega_{SW} \rightarrow \omega_{XC}$) for another SFG-based cross-correlation measurement. A photomultiplier tube combined with lock-in detection and suitable interference filters was used to measure the resultant SFG power as a function of the delay between the Ti:sapphire and MIR pulses, and hence the duration of the MIR pulses.

5.2.2 Simultaneous difference- and sum-frequency generation

Figure 5.7 shows the cross-correlation data between the two Ti:sapphire pulse trains, measured simultaneously in the SFG and DFG arms of the experiment. The cross correlation is measured to be slightly shorter when the DFG signal is used because of the compression offered by the chirped-mirror pair; the differing optics in either arm of the experiment also cause slight differences in the temporal profiles of the pulses. The insets show timing jitter measurements made while the locking system is activated, demonstrating similar timing jitter of 1.3 and 1.5 fs in a 160-Hz low-pass bandwidth when measured with SFG and DFG, respectively. Flexible DFG repetition rates are achievable by setting of the Ti:sapphire lasers at different commensurable repetition

frequencies [21]. Phase locking the carrier frequencies of the two Ti:s lasers would permit generation of optical spectra in three different regions with the desired phase coherence [22].

5.2.3 Broad mid-IR tunability with 500-fs pulses

The broad tunability in wavelength of our DFG scheme is depicted in Fig. 5.8(a). The Ti:sapphire lasers are tuned to various spectral separations, and the maximum MIR power obtainable at each separation is recorded as a function of the MIR wavelength. Figure 5.8(b) shows how this process is executed: At a given wavelength separation of the Ti:sapphire beams, the spectra of the MIR pulses are recorded for various external angles (ranging from 59° to 71°) of the GaSe. Then the external angle that generates the greatest power is used to record a point in Fig. 5.8(a). In Fig. 5.8(a), the squares represent the raw data, corrected only for the long-pass filter's response and the IR detector's response. The circles are normalized to the average power of each Ti:sapphire beam just before the GaSe crystal, and the triangles are normalized to the average powers just after the GaSe crystal.

The data normalized to the input powers of the Ti:sapphire beams (circles) represent the true capability of our system. When the system was optimally aligned, more than 15 mW of MIR power could be generated above $9 \mu\text{m}$, although less power is available at shorter wavelengths because of higher reflection losses at the GaSe surfaces at the steeper external angles necessary for phase matching. This effect is shown in Fig. 5.8(c), where the triangles (right-hand axis) represent the product of transmissions of both Ti:sapphire beams through the GaSe crystal. The development of AR coating for GaSe [134] would allow for at least three times enhancement of the conversion efficiency of this system. The squares (left-hand axis) of Fig. 5.8(c) show the external phase-matching angle as a function of the desired MIR wavelength.

To illustrate the duration of the DFG pulse, in Fig. 5.9 we show the SFG-based

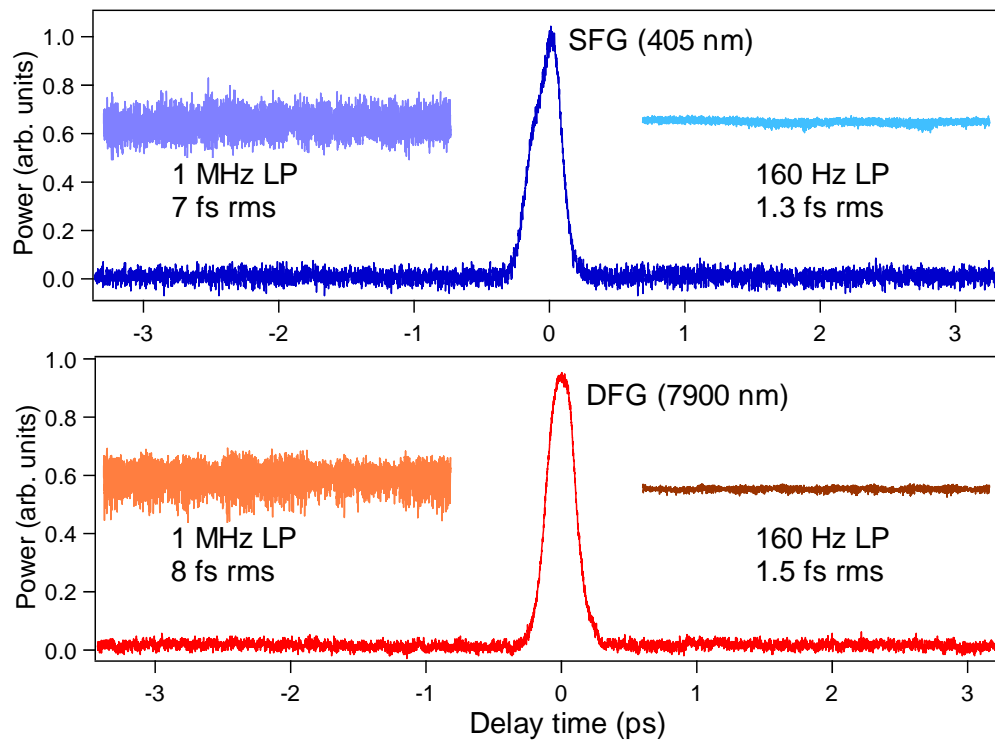


Figure 5.7: Simultaneous cross-correlation measurements of SFG (top) and DFG (bottom). The insets show records of timing jitter (of a few seconds duration) determined from intensity fluctuations of the respective signals, observed through 1-MHz (left) and 160-Hz (right) low-pass filters when the two pulse trains are offset by half a pulse width while synchronized.

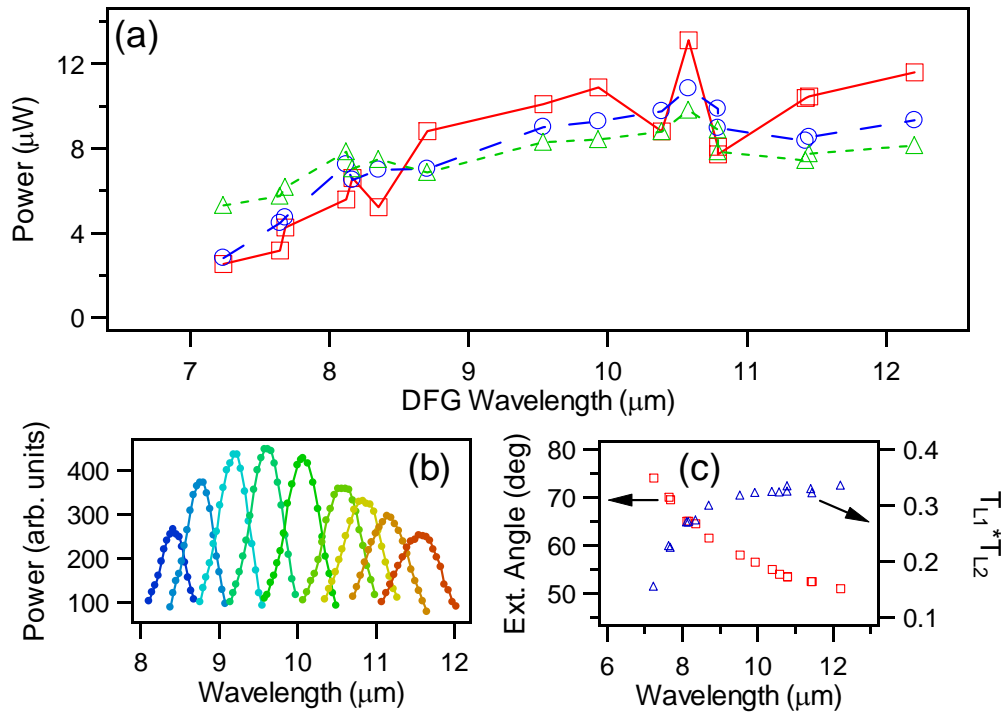


Figure 5.8: (a) MIR average power, as a function of wavelength: squares, raw data corrected only for detector response; circles, normalized to the incident power of each Ti:s laser; triangles, normalized to exiting power of each Ti:s laser. (b) Spectra for various external phase-matching angles, all taken for the same separation of the Ti:sapphire center wavelengths. Here, the global maximum corresponds to one data point in (a). (c) Squares, external phase-matching angles versus MIR wavelength; triangles, product of the two laser transmissions through GaSe as a function of MIR wavelength, showing the need for AR coatings and arbitrary cuts of GaSe.

cross-correlation signal (typically near 720 nm) between the DFG (11.5 μm) and the shorter-wavelength Ti:sapphire (770-nm) pulses, with a measured FWHM of 670 fs. Taking into account the Ti:sapphire pulse width before the second GaSe crystal, the group-velocity dispersion at 11.5 μm , and the group-velocity mismatch between the Ti:sapphire and DFG pulses inside the GaSe, we note that the true DFG pulse width is estimated to be less than 500 fs.

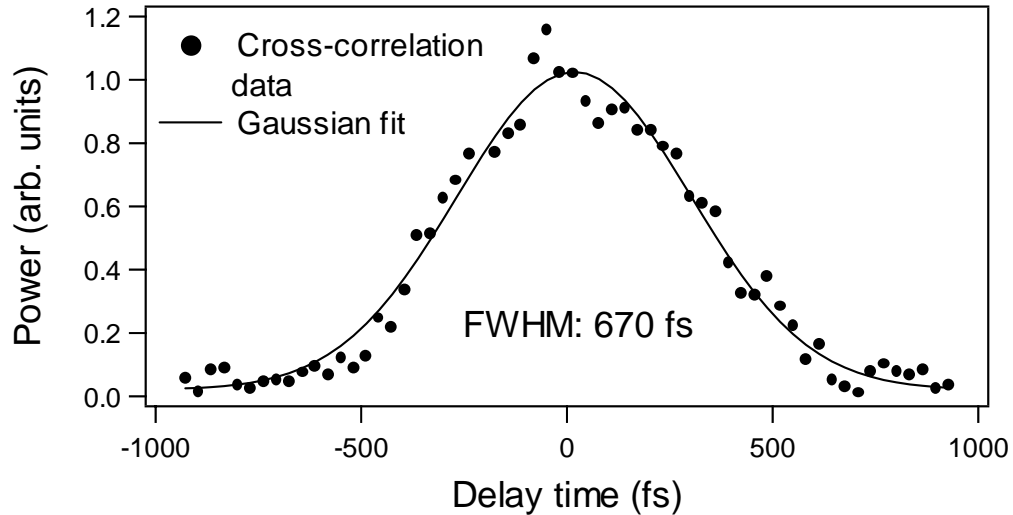


Figure 5.9: Cross-correlation measurement between the 11.5- μm MIR pulse and the short-wavelength Ti:sapphire laser. Taking into account the Ti:sapphire pulse duration, group-velocity dispersion, and group-velocity mismatch between the pulses inside the GaSe, the true MIR pulse duration has an upper limit of 500 fs. MIR pulse versus the shorter-wavelength Ti:s laser, showing a FWHM of 670 fs.

5.2.4 Rapid wavelength switching using chirp and timing offset

Using rapid switching of the delay time between the two original pulse trains, the dual-laser approach makes possible significant and unique flexibilities in the control of the DFG and SFG signals. When the Ti:sapphire pulses are well-compensated for with minimal chirps, changing the delay offset between these pulses will lead to controlled variations in the nonlinear signal amplitudes, while leaving the associated wavelength intact. This situation is illustrated in Fig. 5.10(a), showing the possibility

of fast and precise switching of the nonlinear signal amplitudes. However, when the two lasers are chirped, we can use this to our advantage by precisely setting and (or) switching the DFG and SFG wavelengths, again with control of the timing offset. When the two lasers have similar chirp sizes and signs, a maximum amount of wavelength switching can be produced in the DFG signal. For SFG, chirps of opposite signs are preferred. Figures 5.10(b) and 5.10(c) show the wavelength-switching results in SFG and DFG, respectively. In this case, both lasers are positively chirped. We note that the wavelength-switching capability is limited chiefly by the phase-matching restrictions in the GaSe crystal.

These synchronization experiments have only brushed the surface of the wide variety of applications that can benefit from tight synchronization of fs-pulsed lasers. For instance, the tight synchronization of pump- and probe-pulses for coherent anti-Stokes Raman scattering microscopy has allowed greatly enhanced signal-to-noise leading to real-time, chemically sensitive microscopy [19]. Also, stabilization of the group delay of a 7-km fiber link allowed for the remote synchronization of two fs fiber lasers (1550-nm wavelength) at a level of 19 fs, an important step for remote optical frequency comb distribution [135]. See Chapter 8 for more details on this. Ultimately, phase stabilization of the relative offset frequencies between the two combs in addition to this level of synchronization will allow phase coherence across the electromagnetic spectrum, from 350-nm all the way into the THz regime. One can imagine coherently addressing the electronic, vibrational, and rotational states of a molecule simultaneously with a completely phase-coherent pulse sequence using this kind of system. I'm excited to see this type of system used for wide-ranging coherent control schemes beyond these initial demonstration experiments!

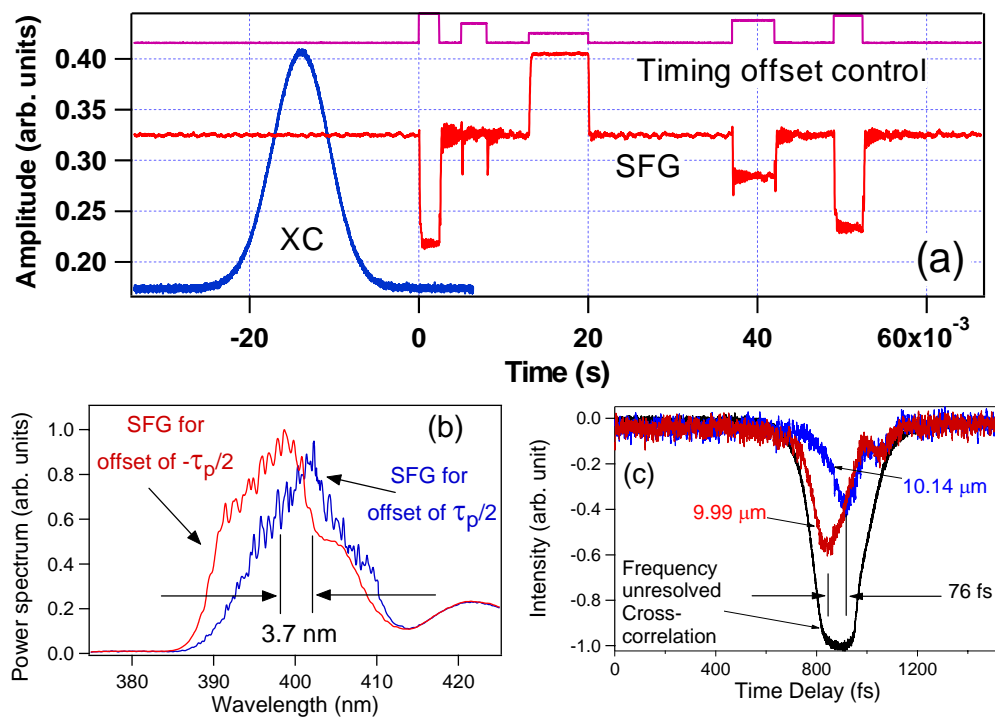


Figure 5.10: (a) Arbitrary and programmable amplitude modulation of the SFG (in this case) or DFG light. (b) Spectra of the SFG light for two different timing offsets by the differently chirped Ti:sapphire pulse trains, illustrating a fast and programmable wavelength jump. (c) Frequency-resolved cross-correlations by use of the DFG signal for two different monochromator settings, while triggering on the frequency-unresolved signal. Different spectral information is contained at different timing offsets during the cross-correlation sweep, demonstrating that timing offset control will influence the IR spectrum.

Chapter 6

Precision Infrared Spectroscopy: the Offset-Free Comb

One of the most striking and immediate applications of femtosecond combs has been the ability to easily and accurately make absolute optical frequency measurements with respect to the current microwave atomic frequency standard. In this Chapter, I will be discussing the specific case of using the comb to derive a microwave output from an optical clock. As discussed in Chapter 3, the optical frequencies ν_m of a femtosecond comb are related to two radio frequencies: f_{rep} , the repetition frequency of the pulse train, and f_0 , the carrier-envelope offset frequency. The relationship $\nu_m = mf_{rep} + f_0$, where integer m is $\sim 10^6$, allows optical-microwave phase coherence to be achieved in a single step when controlling the comb's two degrees of freedom. To establish an optical clock with a countable microwave readout, the best method would be to derive f_{rep} phase coherently from an optical frequency standard ν_s . A heterodyne beat signal f_b between ν_s and ν_m is used to establish a connection between this phase coherent connection, provided that f_0 is either (1) independently stabilized or (2) eliminated as a variable. In this Chapter I will focus on the second option, and describe a simple method for the measurement of infrared frequencies without the need for any measurement or control of the offset frequency. The technique we employed was then used to construct an optical molecular clock based on a HeNe laser stabilized to CH_4 , demonstrating the application of the scheme to metrology in the infrared region.

First it is useful to describe the various schemes that have been employed for

using the comb to measure absolute optical frequencies. The most common of these involves the use of self-referencing to directly detect and stabilize f_0 [98]. This technique is most easily implemented with an octave-spanning spectral bandwidth, which is usually achieved through spectral broadening in a microstructure or tapered fiber. More recently, lasers that directly emit octave-spanning spectra [101, 136] have made possible the elimination of such highly nonlinear fibers from the clockwork while maintaining a stable f_0 lock. A second scheme involves orthogonalizing error signals derived from two optical heterodyne beats of the comb against a reference laser ν_s and its second harmonic $2\nu_s$ [118]. If the two optical beats are expressed as $f_{b1} = \nu_s - (mf_{rep} + f_0)$ and $f_{b2} = 2\nu_s - (2mf_{rep} + f_0)$, then the difference between the two beats is $f_{b2} - f_{b1} = \nu_s - mf_{rep}$, permitting control of f_{rep} from $n\nu_s$ without stabilizing f_0 . To avoid the use of highly nonlinear fiber or the need for special octave-spanning lasers, Amy-Klein et al. [137] locked two modes of the same femtosecond comb to two cw diode lasers. The difference frequency between the two cw lasers was locked to a stable infrared reference. This technique is insensitive to f_0 but involves numerous locks and lasers.

6.1 Principle of offset-free spectroscopy

I will now discuss our demonstration of the operation of an optical molecular clock based on either sum- or difference-frequency generation (SFG or DFG) that does not rely on microstructure fiber, additional cw lasers, or any stabilization of f_0 [113, 112]. A schematic diagram of the concept is shown in Fig. 6.1. If a normal Ti:sapphire frequency comb has strong power in two widely-separated regions of its spectrum, say near ν_{blue} and ν_{red} , then the comb modes in these two regions can be written as $\nu_{m1} = m_1 f_{rep} + f_0$ near ν_{blue} and $\nu_{m2} = m_2 f_{rep} + f_0$ near ν_{red} . Narrow spectral regions near ν_{blue} and ν_{red} are then focused into a nonlinear crystal for DFG, which results in the formation of an

infrared comb centered at ν_{IR} . The IR frequency comb has components

$$\nu_n = \nu_{m_1} - \nu_{m_2} = (m_1 - m_2) f_{rep} = n f_{rep}. \quad (6.1)$$

Notice that the resulting IR comb has no offset frequency, which ultimately will make it possible to measure absolute infrared frequencies without any measurement or control of f_0 .

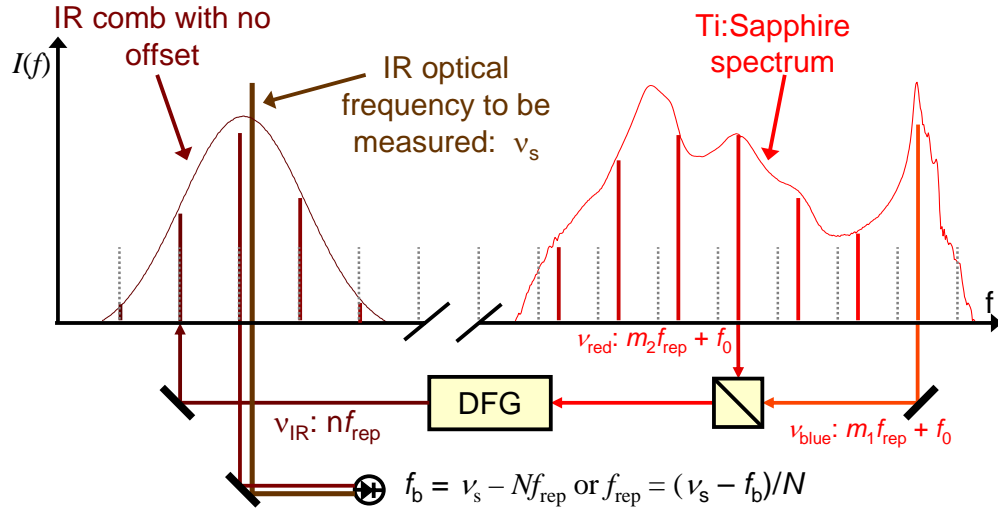


Figure 6.1: Scheme for offset-free precision IR spectroscopy. Difference-frequency-generation between red and blue portions of a single Ti:sapphire comb results in an offset-free IR comb. Stabilization of the heterodyne beat between a cw infrared laser (ν_s) and the IR comb is performed by feedback to the Ti:sapphire cavity length. Then absolute frequency measurement of the microwave readout frequency f_{rep} from the Ti:sapphire laser exactly determines the absolute infrared frequency ν_s .

By careful choice of the spectral regions ν_{blue} and ν_{red} , the IR comb centered at ν_{IR} can be made to overlap with the reference cw optical frequency ν_s . A heterodyne beat between some N^{th} component of the IR comb and ν_s has the frequency

$$f_b = \nu_s - N f_{rep} \quad (6.2)$$

which can be stabilized relative to some microwave reference by feedback to a piezo-actuated mirror to control f_{rep} by adjusting the Ti:sapphire laser's cavity length. Once f_b is stabilized in this manner, f_{rep} can be detected as the microwave readout of the

absolute frequency measurement. Since

$$f_{rep} = (\nu_s - f_b) / N, \quad (6.3)$$

the optical frequency must be

$$\nu_s = f_b + N f_{rep}. \quad (6.4)$$

Measurement of the absolute microwave frequencies f_b and f_{rep} , combined with knowledge of the integer N (which can be found from knowing ν_s to within $f_{rep}/2$, using a wavemeter if need be) therefore exactly measures the absolute frequency ν_s .

A very similar method employing SFG instead of DFG can also be used, where the cw IR beam at ν_s is summed with the comb near ν_{red} , producing a narrow frequency comb centered at ν_{SFG} near the original comb at ν_{blue} . Subsequently, a heterodyne beat between all the components of the two combs ν_{SFG} and ν_{blue} is detected and stabilized, again by feedback to the Ti:sapphire laser's cavity length. This is really the same frequency math, but a different order of operations for the optical beams. In both schemes, the offset frequency f_0 of the Ti:sapphire laser is left unmeasured and uncontrolled. Therefore this scheme can not be used to relate the infrared ν_s to frequencies in the visible portion of the spectrum. However, the simplicity of the scheme should enjoy great utility in making absolute frequency measurements throughout the infrared, as well as making atomic clocks using natural resonances located in the infrared.

6.2 Experimental implementation of offset-free clockwork

We began these experiment with an eye toward constructing an optical molecular clock based on a methane-stabilized HeNe laser at $3.39 \mu\text{m}$ [138, 139], constructed by M. Gubin from the Lebedev Physical Institute in Moscow. Due to his limited availability, we first built a system [112] using an optical parametric oscillator (OPO) operating at $3.39 \mu\text{m}$ as a proof-of-principle concept. In this Section I describe the results of our first experimental demonstration. Then in Section 6.3 I will describe the clock as operated

with the true optical molecular reference [108]. A testament to the simplicity of the offset-free scheme is that the first experiment with the OPO was built in just under 2 months.

6.2.1 Ti:sapphire laser with two separated spectral peaks

The X-folded Ti:sapphire laser operates at a repetition frequency f_{rep} of 78 MHz and emits an average output power of typically 150 mW. It employs a 2.3-mm-thick Ti:sapphire gain crystal and seven bounces on double-chirped mirrors for dispersion compensation. The two concave double-chirped mirrors, each with a focal length of 50 mm, produce a beam waist inside the gain crystal, which is oriented at Brewster's angle of $\sim 60^\circ$. Approximately 6.7 W of 532-nm pump light emitted by a frequency-doubled Nd:YVO₄ laser is focused into the gain crystal by an achromatic doublet lens (76.2-mm effective focal length). The output spectrum, shown in Fig. 6.2 is spectrally shaped by a custom-designed narrow-band output coupler consisting of five pairs of SiO₂—TiO₂ layers. The resultant transmission is $\sim 0.5\%$ in the center of the output coupler at ~ 800 nm; it increases strongly to 5% at the designated wavelengths of 685 and 866 nm. This enhances the peak at 670 nm and gives rise to a shoulder in the spectrum at 880 nm (Fig. 6.2). The frequency difference between the 670-nm peak and a portion of the spectrum near 830 nm corresponds to a wavelength of $3.39 \mu\text{m}$, an important frequency for metrology due to the nearby methane standard. Approximately 12.3 mW of average power is contained within the 7-nm FWHM bandwidth of the 670-nm peak, and a 10-nm-wide spectral region centered at 831 nm has 7.8 mW of average power.

6.2.2 SFG-based offset-free IR frequency measurement

The experimental setup for both the OPO SFG-based frequency measurement [112] described in this Section, and the CH₄-stabilized HeNe optical molecular clock [108] described in Section 6.3, is portrayed in Fig. 6.3. The cw OPO used for the first proof-

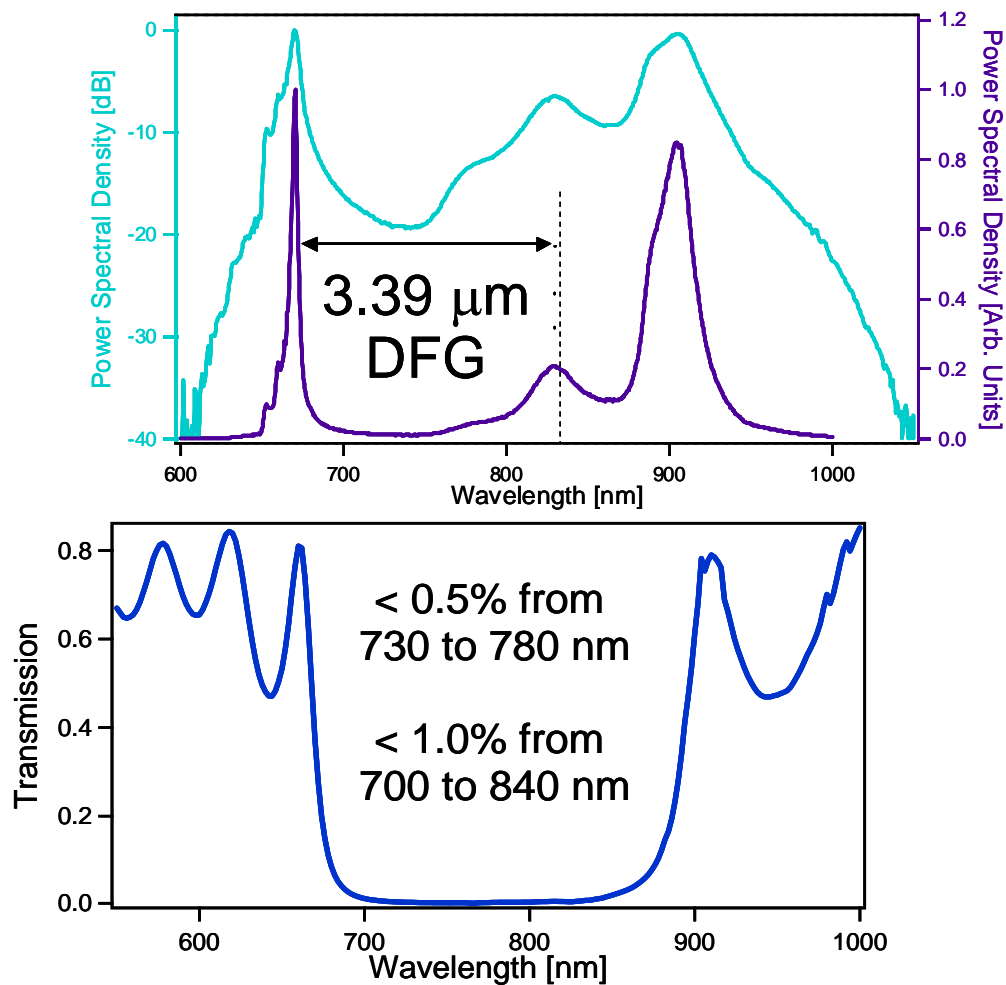


Figure 6.2: (Top) Output spectrum of the Ti:sapphire laser, with two peaks near 670 and 831 nm used to generate 3.39- μm light. A third parasitic peak at 902 nm is not used for the experiment. (Bottom) Transmission of the output coupler used to force the laser to emit the two-peaked spectrum.

of-principle experiment is a commercially available system (OS 4000, Linos Photonics) pumped by a Nd:YAG ring laser at $\lambda_p = 1064$ nm. This singly resonant OPO with a resonated pump is based on a periodically poled lithium niobate (PPLN) crystal, and it features a narrow linewidth (< 150 kHz) and an idler output power of more than 20 mW [140]. The idler wavelength can be tuned from 2.3 to 3.8 μm . The OPO's cavity length is locked to the pump laser by the Pound–Drever–Hall stabilization technique. One can conveniently tune the signal wavelength, which is monitored with a wavemeter, to $\lambda_s = 1550$ nm by varying the PPLN crystal's temperature and using the intracavity etalon of the OPO. The resultant idler wavelength is $\lambda_I = 3390$ nm (frequency ν_I).

After the Ti:sapphire output and the OPO idler were combined on a beam splitter, 28-mW of Ti:sapphire output near 834 nm and 16 mW of the cw OPO idler were focused into a PPLN crystal by a 40-mm focal-length calcium fluoride lens. This 5-mm-long PPLN crystal has a quasi-phase-matching period of 16.2 mm and was heated to 130°C. The resultant phase-matching bandwidth is ~ 0.75 THz. The output light was collimated with a 38-mm calcium fluoride lens. The resultant SFG comb and the (adequately attenuated) Ti:sapphire comb at 670 nm were temporally and spatially overlapped by a prism-based delay line (two SF10 prisms with an apex-to-apex separation of 33 cm; see Fig. 6.3), prefiltered by a narrow-band interference filter at 670 nm, and detected by a photomultiplier (Hamamatsu R7400U-20). In the experiment we observed SFG powers on the order of 870 pW. The signal-to-noise ratio of the heterodyne beat $f_{b,SFG}$ between the two combs was 23 dB in a 100-kHz resolution bandwidth, as shown in Fig. 6.4. Note that $f_{b,SFG}$ results from the coherent superposition of many corresponding comb lines in each of the two combs:

$$f_{b,SFG} = (\nu_I + m_2 f_{rep} + f_0) - (m_1 f_{rep} + f_0) = \nu_I - (m_1 - m_2) f_{rep}, \quad (6.5)$$

since the mode indices m_1 and m_2 can take on any values corresponding to comb modes within the relevant spectral bandwidths contributing to the beat. In fact, an rf comb

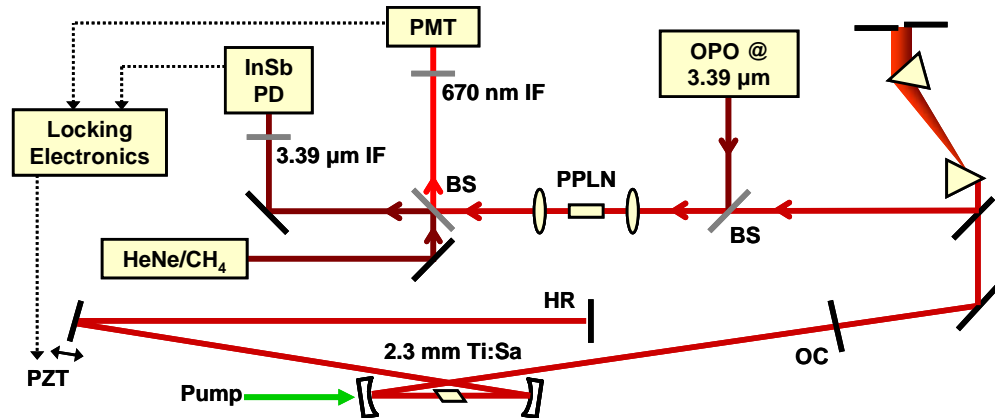


Figure 6.3: Experimental setup for the offset-free SFG- or DFG-based clockwork. The initial proof of principle experiment used an OPO and the SFG scheme $(3.39\mu\text{m})^{-1} + (834\text{nm})^{-1} = (670\text{nm})^{-1}$. Then the beat near 670 nm was detected on the PMT. For the optical molecular clock, a CH_4 -stabilized HeNe was used for the optical reference in the DFG-based scheme $(670\text{nm})^{-1} - (834\text{nm})^{-1} = (3.39\mu\text{m})^{-1}$. In this scheme, the beat was detected on the liquid-nitrogen-cooled InSb photodiode. In either case, the beat was used to feedback to the Ti:sapphire cavity length in order to phase-coherently connect the IR reference of interest to f_{rep} , which was subsequently detected as the microwave readout for the measurements. OC, output coupler; HR, high-reflector; PZT, piezoelectric transducer; BS, beam splitter; IF, interference filter; PMT, photomultiplier tube; PD, photodiode. All intracavity mirrors for the Ti:sapphire laser are double-chirped mirrors.

of beat frequencies $f_{b,SFG}$ and conjugate beat frequencies $f_{rep} - f_{b,SFG}$ is detected as is usual with heterodyne beats involving optical frequency combs. Importantly, $f_{b,SFG}$ does not depend on carrier-envelope frequency f_0 . Because of its marginal signal-to-noise ratio, we regenerated the beat signal with an rf tracking oscillator, which consists of a voltage-controlled oscillator (VCO) and analog phase-locked loop to keep the VCO phase-locked to the input beat signal. By then phase locking the tracking filter's output to an rf reference (any synthesizer with absolute fluctuations below the optical frequency fluctuations being measured), we established a direct phase-coherent link between f_{rep} and the optical OPO idler frequency ν_I . The feedback was applied to a small PZT on which the fold mirror of the laser cavity was mounted. Ideally, measurement of f_{rep} is therefore an accurate and phase-coherent measurement of ν_I , via Eqn. 6.5. In an attempt to stabilize fluctuations of ν_I in order to better-characterize the system, we also phase locked the OPO pump laser to an iodine-stabilized Nd:YAG laser at 1064 nm [118], which was delivered to our setup via a polarization-maintaining optical fiber.

6.2.3 Evaluation of the OPO frequency measurement

To evaluate the stability of our optical clockwork we made a comparison of the tenth harmonic of f_{rep} and an rf reference signal. The reference was derived from a low-phase-noise rf synthesizer referenced to the National Institute of Standards and Technology (NIST) ST-22 hydrogen maser (instability of $\sim 2.4 \times 10^{-13}$ at a 1-s counter gate time) and transferred from the NIST to JILA via the Boulder Research and Administrative Network single-mode fiber [37, 141], as will be discussed in 8.3.1. Measured frequency-counting data are depicted in Fig. 6.5. Under unlocked conditions, we typically observe standard deviations of 240 Hz (relative to the 10th harmonic frequency of 760 MHz) at a 1-s gate time. In time interval A in Fig. 6.5, $f_{b,SFG}$ is phase locked to an rf reference, while the OPO pump laser is phase locked to the cw laser used in the




Figure 6.4: Heterodyne beat between the SFG comb and the Ti:sapphire comb at 670 nm. The signal-to-noise ratio of the beat signal is 23 dB measured in a 100-kHz resolution bandwidth.

iodine optical frequency standard. The data reveal a standard deviation of 0.837 Hz at a 1-s gate time. At 108 s, we deliberately unlock the OPO pump laser; however, the heterodyne beat signal remains phase locked. In subsequent time interval B, the standard deviation is correspondingly increased to 5.09 Hz. These data present convincing evidence that the instability of f_{rep} under the locked condition arises directly from the frequency fluctuations of the OPO idler. Although the OPO cavity length, measured at 1064 nm, is stabilized to that of the pump laser, which is in turn stabilized to the iodine transition, the intracavity etalon used to tune the OPO idler wavelength causes cavity length instabilities at 3.39 μm , which limit the stability of the idler frequency and subsequently that of f_{rep} . With the in-loop error signal indicating a tight phase lock between the optical beat signal and the rf reference, the full potential of this carrier-envelope-frequency independent clockwork could only be revealed if the OPO would be further stabilized by methods such as use of a 2:1 frequency divider [142].

Without an independent way to accurately measure the optical frequency ν_I , our measurement of the clockwork instability was therefore limited by the intrinsic instability of ν_I , even when the OPO pump was frequency-stabilized. The instability was $\sim 1 \times 10^{-9}$ at 1-s averaging time. Therefore it will be the subject of the next Section to show that use of a much more stable IR frequency reference allows the offset-free clockwork to be characterized to four orders of magnitude better precision. The new reference will be a HeNe laser stabilized on the methane $F_2^{(2)}$ at 3.39 μm , which has demonstrated a frequency instability of 10^{-14} to 10^{-15} for 1–1000-s measurement times and a frequency reproducibility approaching 10^{-14} [139].

6.3 Optical molecular clock based on CH₄

After having demonstrated a working offset-free clockwork, we turned toward the goal of employing it for an optical clock based on a CH₄-stabilized HeNe. In principle this would provide a compact, robust, and relatively inexpensive optical clock featuring




Figure 6.5: Comparison of the tenth harmonic of the Ti:sapphire repetition frequency with the hydrogen maser-based rf reference at 780 MHz f_{maser} (at a 1-s gate time). In time interval A, $f_{b,SFG}$ is stabilized and the OPO pump laser is phase locked to the iodine frequency standard, resulting in a standard deviation of 0.837 Hz at a 1-s gate time. At 108 s, the optical beat signal remains stabilized, but we deliberately unlock the OPO pump laser. In subsequent time interval B, the standard deviation is correspondingly increased to 5.09 Hz.

both high accuracy and high stability. It also allows for better characterization of the offset-free clockwork than the OPO does.

6.3.1 Experimental setup for the clock

For this experiment, the OPO shown in Fig. 6.3 was now blocked, as well as the SFG beams incident on the PMT. Instead of the sum-frequency method described in Section 6.2.2, we chose to use the original difference-frequency idea discussed in Section 6.1. Therefore, for DFG between the 670- and 831-nm portions of the Ti:sapphire laser's spectrum, its output was focused into the temperature-controlled PPLN crystal. For efficient DFG inside the PPLN the Ti:sapphire beam traversed a prism-based delay line placed before the PPLN for an optimal temporal overlap of the 670- and 834-nm spectral components. The resultant offset-free IR frequency comb passed through a 50-nm FWHM bandwidth interference filter centered at $3.39 \mu\text{m}$ that transmitted ~ 10 mW of average power, corresponding to nearly 600 pW per mode of the IR comb, with modes given by Eqn. 6.1. With an IR monochromator and lock-in detection we observe that the IR comb has a bandwidth of 270 nm (7 THz) centered near $3.4 \mu\text{m}$ (88.5 THz) as shown in Fig. 6.6. Again, the IR comb modes have no dependence on the offset frequency.

The optical frequency standard was a compact, portable, double-mode HeNe laser stabilized to the $F_2^{(2)}$ transition in methane (HWHM ~ 200 kHz, hyperfine structure is unresolved) at $3.39 \mu\text{m}$. Saturated absorption and saturated dispersion resonances are both used to lock the reference laser frequency, providing optical radiation with a linewidth of < 100 Hz. This type of reference HeNe laser is being routinely incorporated in transportable HeNe/CH₄ systems with telescopic beam expanders for providing resolution of methane's hyperfine structure and frequency repeatability of 2×10^{-13} over several years [138]. Measurements with a frequency chain indicate an instability at 1 s of 4×10^{-13} , limited by a hydrogen maser. A direct optical comparison of two versions




Figure 6.6: Spectral coverage of the offset-free IR comb generated by difference-frequency mixing of the Ti:sapphire laser comb in a PPLN crystal. The 270-nm FWHM is centered near $3.4 \mu\text{m}$ in order to overlap with the CH_4 -stabilized HeNe optical molecular frequency reference.

of the system (one was a non-transportable, resolved-hyperfine-structure version) has a demonstrated Allan deviation of 2×10^{-14} at 1 s. A second HeNe laser (heterodyne laser) was phase locked to the reference laser with a fixed 600-kHz frequency offset and provides the optical field used to stabilize the IR comb. An amplifier HeNe tube was also used to increase the beam's power from 300 μ W to 1 mW.

After the cw HeNe beam and IR comb were combined on a 60/40 beam splitter, they passed through a 50-nm FWHM interference filter centered at 3.39 μ m. A 1.8-cm focal-length CaF₂ lens focused the beams onto a nitrogen-cooled InSb photodiode with a 250- μ m-diameter active area. A heterodyne beat between the cw HeNe beam at ν_{HeNe} and the N^{th} line of the IR comb was detected and amplified. As shown in Fig. 6.7, the beat frequency $f_{b,DFG} = Nf_{rep} - \nu_{HeNe}$ had a signal- to-noise ratio of 25 dB in a 100-kHz bandwidth. It was subsequently phase locked by means of a rf tracking oscillator to a 70-MHz synthesized signal (with a negligible noise contribution) derived from a stable cesium reference. Control of $f_{b,DFG}$ was achieved by feedback of the phase-locked loop error signal to a piston-mode piezoelectric transducer on which the Ti:sapphire laser's fold mirror was mounted. Now rf frequency f_{rep} was directly expressed in terms of optical frequency ν_{HeNe} as $f_{rep} = (\nu_{HeNe} + f_{b,DFG})/N$. The optical clock microwave readout can be any harmonic of f_{rep} , and was detected independently at the beam splitter's second port, where the PMT used to be located for the OPO measurement in 6.2.2.

6.3.2 Evaluation of the optical molecular clock

The rf clock signal's stability was compared with both a hydrogen maser and an independent optical clock based on a molecular iodine transition [118]. For comparison with the maser, the tenth harmonic of f_{rep} was compared with a 780-MHz signal derived from the National Institute of Standards and Technology (NIST) ST-22 hydrogen maser and transferred to JILA through a single-mode fiber [37, 141], as will be discussed in Sec-

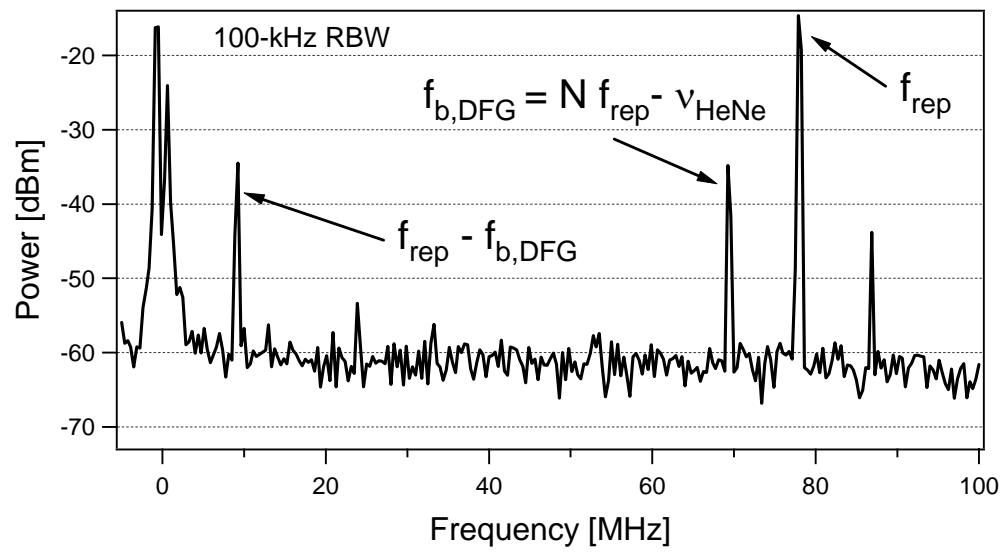


Figure 6.7: Heterodyne beats between the CH_4 -stabilized HeNe laser and neighboring IR comb components, with a signal-to-noise ratio > 25 dB in a 100-kHz RBW.

tion 8.3.1. The I₂ clock used a second Ti:sapphire laser with a 100-MHz repetition rate as clockwork in order to generate a microwave readout (harmonic of its repetition rate) in the usual way. Its frequency comb was fixed by locking f_0 by use of self-referencing and by phase locking one comb line near 1064 nm to a Nd:YAG laser that was stabilized to molecular iodine [118]. The seventh harmonic of this I₂-referenced Ti:sapphire laser's repetition frequency is compared with the ninth harmonic of the CH₄-referenced Ti:sapphire laser's repetition frequency at ~ 702 MHz. We deliberately introduced an offset frequency of 10 kHz between the two clock signals. A double-balanced mixer detected this 10-kHz difference frequency, which was subsequently filtered and counted with a 1-s gate time to determine the Allan deviation of the CH₄ clock compared with either the I₂ clock or the maser. Figure 6.8(a) shows representative frequency counting for both comparisons. Denoting the comparison frequency (780 MHz for comparison against the maser, 702 MHz for comparison against the I₂ clock) as f_c , frequency fluctuations Δf_c were normalized to f_c for fair evaluation of the two fractional instability measurements.

Parasitic back-reflections from surfaces inside the HeNe resonator could shift the centers of the saturated absorption and saturated dispersion resonances of the reference system [143]. As the position of the reflecting surface moved because of temperature changes, both a slow drift and a slow oscillation of ν_{HeNe} were observed. A quadratic drift of Δf_c tightly correlated with recorded temperature changes of the HeNe system was subtracted from the raw data. An oscillation with a period of >20 s was still present, however. Figure 6.8(b) shows the Allan deviations determined from various frequency-counting records. Measurements against the maser (filled diamonds) were limited by the maser's intrinsic short-term instability of $\sim 3 \times 10^{-13}$ at 1 s. Data run I (bow ties) against the I₂ clock used free-space photodiodes for detection of both clocks' repetition rates, whereas data run II (circles) used fiber-coupled detectors that exhibit greater phase noise and amplitude-phase conversion. The noisy plateau in the Allan deviation

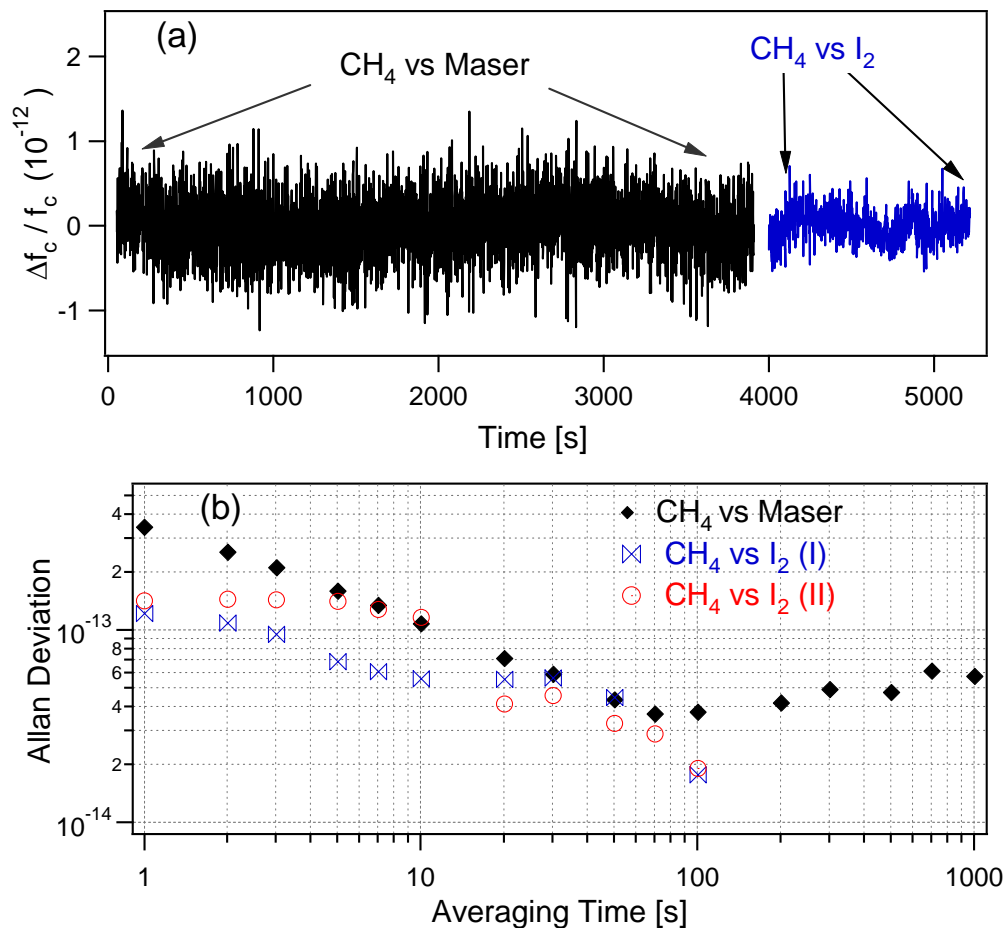


Figure 6.8: Comparison of the rf signal from the CH₄ clock with NIST's hydrogen maser and with the I₂ clock's rf signal. (a) Fractional fluctuations of the comparison frequency in each case. (b) Allan deviations for comparison with the maser (diamonds), the I₂ clock with fiber-coupled detection of each clock's f_{rep} (circles, and the I₂ clock with free-space-coupled detection of each clock's f_{rep} (bow ties).

for data run II for $t > 20$ s strongly suggests that amplitude-to-phase conversion in the detection process limited our instability for short time scales [76]. However, for time scales of $20 \text{ s} < t < 100 \text{ s}$, data run II exhibits stability comparable with data run I, even though they were taken under very different conditions of temperature fluctuations in the laboratory. Therefore the instability of our CH_4 clock as compared with the I_2 clock was $\sim 1.2 \times 10^{-13}$ at 1 s, averaging down as $\tau^{-1/2}$ for averaging times $\tau < 100$ s, limited by excess noise in the photodetection process as discussed in Section 2.2.2. Both before and after the HeNe laser was brought to JILA, it demonstrated an instability in the optical domain of 2×10^{-14} at 1-s averaging time. The I_2 optical frequency has demonstrated 5×10^{-14} instability at 1-s averaging time. Therefore, with the improved techniques described in Chapter 4, it should be possible to recover these levels of performance for both clocks in their respective microwave outputs by appropriate optical-to-microwave conversion of microwave harmonics of the combs' repetition frequencies.

To characterize the phase noise of our CH_4 clock signal, we examined the 10-kHz signal derived from comparison with the I_2 clock (using free-space detection of the clock signals) with a fast Fourier transform spectrum analyzer. The measurements represent an upper limit to the single-sideband (SSB) phase noise, because they were sensitive to amplitude noise as well (we used a heterodyne scheme rather than a homodyne phase-detection). We therefore take the power spectral density (relative to the carrier power, in units of [dBc/Hz]) to be an upper limit to the SSB phase noise as shown in Fig. 6.9. In addition, we scaled the phase noise up to a 1-GHz carrier frequency for ease of comparison against other oscillators. We include curves representing the typical SSB phase noise of other extremely stable microwave sources [144]. Less than 10 Hz from the carrier, our SSB phase noise is superior to some of the lowest-phase-noise microwave sources available, achieving -93 dBc/Hz at a 1-Hz offset.

Unfortunately, in our haste to complete the experiment due to the Moscow team's limited (only 2 weeks!) stay, we failed to realize that the amplifiers used in the detection




Figure 6.9: Upper limit to the SSB relative phase noises of the CH₄ and I₂ clock signals. Additional curves show the approximate SSB phase noises of other extremely low-phase-noise microwave sources for comparison. All data were scaled for a 1-GHz carrier.

of f_{rep} harmonics for both clocks were being undersupplied and therefore overdriven far into saturation. With the latest techniques for optical-to-microwave conversion of fs pulse trains described in Chapter 4, using detection of a much higher harmonic of f_{rep} , low-phase-noise amplifiers, and impedance-matched photodetectors in an LC-tank configuration, the phase noise and instability of the CH₄ offset-free clock could be much improved. Even so, without the need for stabilization of f_0 and coupled with the use of a portable highly stable HeNe/CH₄ optical standard, our setup represents a compact, reliable clock with high stability and exceedingly low phase noise, which can in principle be operated for long >24-hour periods.

Chapter 7

The Octave-Spanning Comb: Clockwork and Distribution For an Optical Atomic Frequency Standard

In addition to cell-based optical frequency references such as I_2 and CH_4 , in our laboratory we have, in operation, a cold-atom based optical clock that has the potential to become a primary frequency standard in the future. The clock is based on the doubly-forbidden $^1S_0 - ^3P_0$ transition in ^{87}Sr atoms, which has a lifetime longer than 100 seconds. More about the clock is described in Section 7.1. The development and characterization of the clock depends on the ability to compare its optical frequency to other frequency standards; these standards can be the microwave Cs standard, or they can be other optical clocks. In order to cover the large spectral gaps between the ~ 430 -THz clock transition and these other standards, a reliable and robust clockwork that can operate for many hours, or even days, at a time is desired.

In this Chapter I will describe just such a clockwork, realized by the high-bandwidth control of a directly-octave-spanning Ti:sapphire laser. The extremely broad spectrum available from the laser means that the self-referencing technique described in Section 3.3.1 can be achieved without the need for external broadening of the spectrum in a microstructure fiber. The use of microstructure fiber normally entails difficulty with maintaining the lock of f_0 due to stringent alignment requirements and frequent damage of the fiber. In contrast, our octave-spanning laser has demonstrated the ability to maintain a continuous phase lock to the Sr clock laser for longer than 12 hours with

no operator intervention. This is achieved even in a laboratory environment where the temperature fluctuates by several degrees Celsius, with no thermal control of the laser's breadboard.

Such a robust clockwork has enabled the development and characterization of the Sr optical lattice clock in several ways that will be discussed in Section 7.4: the octave-spanning laser can be phase-locked to the Sr local oscillator (clock) laser with a demonstrated absolute coherence time (for the comb, and therefore the clock laser as well) exceeding one second. This has allowed the initial characterization of the Sr clock laser relative to another ultrastable laser operating with a known performance at 1064 nm [145, 146]. With the clock laser performance suitably characterized, absolute frequency measurements of the Sr optical lattice clock were also performed, by measuring the repetition frequency of the octave-spanning comb against a Cs-calibrated microwave reference made available to us from NIST [146]. As a result, the clock transition's absolute frequency is now known to a fractional uncertainty of 2.5×10^{-15} , or 1.1 Hz at the clock frequency, a record for neutral-atom-based optical standards [147]. In addition, the octave-spanning comb is capable of transferring the stability and coherence of the 698-nm Sr clock laser to a 1064-nm cw laser for the purpose of long-distance fiber transfer for comparisons with other clocks operating at NIST [148]. The methods for the actual transfer will be described more thoroughly in Chapter 8.

7.1 Optical lattice clock using ^{87}Sr atoms

Large ensembles of ultracold alkaline earth atoms have provided impressive short-term clock stability [149]. Previously, interrogation of neutral atom based optical standards has been carried out primarily in free space, unavoidably including atomic motional effects that typically limit the overall system accuracy [150, 149, 151]. A promising approach is to explore the ultranarrow optical transitions of atoms held in an optical lattice [152, 153, 154]. The atoms are tightly localized so that Doppler and photon-recoil

related effects on the transition frequency are eliminated [155]. Meanwhile, the trapping potential is created at a carefully chosen laser wavelength (λ_{magic}) such that it has essentially no effect on the internal clock transition frequency [156, 157, 158]. Additionally, the increased atom-probe laser interaction time enabled by the lattice confinement will permit full utilization of the narrow natural linewidth. This optical lattice approach using neutral atoms may provide the best possible combination of clock stability and accuracy. Such a proposal has been under intensive investigation in the case of the doubly forbidden $^1\text{S}_0 - ^3\text{P}_0$ transition in the fermionic Sr isotope, ^{87}Sr [159, 152, 146, 160]. Similar work in Yb is also in progress [161, 162]. In this Section I will describe the various components of our ^{87}Sr optical lattice clock.

7.1.1 The atomic clock's pendulum: cooling and trapping Strontium

The development of optical atomic clocks requires the ability to reach ultracold temperatures, and the stability of such clocks benefits from large atom numbers (N), as described in Section 1.1. By achieving sufficiently low temperatures, atoms can be tightly confined in the Lamb-Dicke regime in a lattice potential, virtually eliminating adverse Doppler effects on the clock's performance. In addition, the stability of the clock can be improved by \sqrt{N} as larger numbers of atoms are employed. In the case of neutral Sr, we use two cooling transitions (the 32-MHz-wide $^1\text{S}_0 - ^1\text{P}_1$ transition at 461 nm followed by the 7.4-kHz-wide $^1\text{S}_0 - ^3\text{P}_1$ transition at 689 nm) in order to cool the atoms to μK temperatures before loading them into the optical lattice. The relevant energy levels are shown in Fig. 7.1. Complete details of the laser cooling apparatus are given in References [163] and [164].

The $^1\text{S}_0 - ^3\text{P}_0$ transition in ^{87}Sr is an excellent clock transition for several reasons. The most obvious is its extremely narrow natural linewidth, corresponding to a lifetime exceeding 100 seconds. Such a narrow line offers an extremely high quality factor, allowing the transition to be probed for long times in order to achieve high spectral




Figure 7.1: Important energy levels for ^{87}Sr . The 461- and 689-nm energy levels are used for Doppler-cooling before the Sr atoms are loaded into an optical lattice, where the 698-nm clock transition frequency is probed by a cavity-stabilized sub-Hertz-linewidth diode laser.

resolution. Another feature of the Sr system is that the cooling and trapping transitions are accessible by diode lasers, and the 7.4-kHz-linewidth $^1S_0 - ^3P_1$ cooling transition allows regular Doppler cooling to achieve 1 μ K temperatures directly. In addition, the clock transition frequency is insensitive to several potential systematic shifts from stray fields; this is because both the ground and excited states of the clock transition have no net angular momentum (quantum number $J = 0$). As a result, coupling to external magnetic fields is minimal. In fact, ignoring the hyperfine structure, no linear Zeeman shift exists for the transition frequency. Even including the hyperfine magnetic sublevel structure, the Zeeman shift has been measured to be on the order of 100 Hz / Gauss, per magnetic projection quantum number m_F . This is a dramatic level of insensitivity compared to typical Zeeman shifts for other transitions on the order of multiple MHz / Gauss. In fact, with resolution of the magnetic hyperfine sublevels, the linear Zeeman shift can be experimentally cancelled by locking the clock to two symmetric sublevels and taking the average frequency. The frequency's dependence on collisions (density of the trapped atoms) is also small, although it must be experimentally determined for a real clock. To date, the collisional shift is consistent with zero at the 3×10^{-16} level for typical operating trap densities.

For all these reasons, the clock transition in fermionic Sr makes an excellent candidate for a frequency standard. Combination of the extremely narrow natural linewidth with insensitivity to stray fields and collisions allows for a clock “pendulum” with a natural oscillation frequency that is currently being accurately reproduced by three independent laboratories around the world, with agreement on the order of 2×10^{-15} [147, 165, 166].

7.1.2 Long interaction times: confinement in the optical lattice

Another important feature of the Sr clock transition is that it allows a trap to be constructed that is capable of holding the atoms still for long times without perturbing

the clock transition. The narrow sub-Hz transition can therefore be probed for times on the order of seconds, allowing high spectroscopic resolution to be achieved.

Free-space spectroscopy of Sr [151] yields troublesome Doppler effects which limit the achievable accuracy. An attractive solution is the tight confinement offered by an optical lattice, which allows long interaction times while essentially eliminating any Doppler or recoil effects. We use a one-dimensional lattice as shown on the left of Fig. 7.2 in a vertical orientation, so that the atoms can be supported from gravity and confined tightly in the same direction as the probe laser’s pointing. Figure 7.2 also shows the ac stark shift of the lower and upper states of the clock transition, as a function of the lattice wavelength, for a typical lattice intensity. When the optical lattice is operated at the “magic” wavelength, λ_{magic} , the relatively large AC stark shifts induced by the lattice field become equal for the 1S_0 and 3P_0 energy levels. In this manner, the clock transition frequency is unperturbed by the presence of the lattice field. For typical operating intensities of our trap, we measured the clock shift to be on the order of 1 Hz per 1-GHz change of lattice frequency. Using the frequency comb to stabilize the lattice frequency should therefore keep the repeatability of the clock’s frequency shift due to the lattice below the level of 10^{-18} . In fact, the intensity-related AC Stark shift of the system has already been characterized to $< 6 \times 10^{-16}$ fractional uncertainty [147]. Our one-dimensional lattice is constructed from a ~ 300 mW standing wave of optical power at $\lambda_{magic} \sim 813.4$ nm, and is oriented nearly parallel to the direction of gravity. After transfer from the second laser-cooling stage to the lattice, $> 10^4$ atoms at ~ 1.5 μ K temperatures are confined with a Lamb-Dicke parameter of ~ 0.3 , allowing Doppler- and recoil-free spectroscopy of the clock transition along the axis of the lattice beam. We have experimentally demonstrated the coherent interaction time between the probe laser and the atoms to be nearly 1 second with this scheme [47], while keeping the systematic shift due to the AC stark shift of the lattice beam below 6×10^{-16} [147].




Figure 7.2: AC stark shifts for the ground and excited states of the ^{87}Sr clock transition. The relative stark shift between the two states is tunable by adjusting the lattice wavelength. At one particular wavelength, λ_{magic} , the relative stark shift can be zeroed such that the lattice confinement does not induce a frequency shift for the transition, relative to an unperturbed atom. The typical lattice operating intensity corresponds to a trap depth of ~ 100 lattice recoils.

7.1.3 The local oscillator: ultrastable cavity-stabilized probe laser

For the purpose of operating a clock with the best possible stability, it is necessary for the probe laser to have both a narrow intrinsic linewidth as well as small frequency drift and low instability for at least several tens of seconds. We probe the extremely narrow natural linewidth of the clock transition in fermionic Sr with a cavity-stabilized diode laser operating at 698 nm, constructed by Andrew Ludlow [46]. By locking the diode laser to the highly-isolated ultrastable passive cavity, laser light with a narrow 0.5-Hz linewidth is generated, suitable for probing the clock transition.

The diode laser is operated in an external cavity (Littman configuration) tuned to operate at 698 nm. Using feedback to the diode current and a PZT on the laser cavity, the laser is prestabilized to an optical cavity with a finesse of about 10^4 . Then the light passes through an AOM and is incident on the ultrastable cavity. The Pound-Drever-Hall stabilization technique is used to stabilize the laser relative to the ultrastable cavity by feedback to the AOM and a PZT on one of the prestabilization cavity mirrors, with a servo bandwidth of 100 kHz.

The ultrastable cavity is mounted inside a temperature-stabilized vacuum chamber, which is in turn fixed to a passive vibration-isolation platform. The entire vibration isolation platform is enclosed in a foam-lined (for acoustic damping), temperature-controlled outer box. The high-finesse (finesse $\sim 250,000$), 7-cm-long cavity is mounted in a vertical orientation, supported near its geometrical midpoint in order to reduce sensitivity to vibrations in the vertical direction [45, 46]. Such vertical acceleration insensitivity is achieved because the upper half of the cavity stretches by the amount the lower half compresses, or vice-versa, for vertical accelerations, allowing the sensitivity to be tuned and therefore minimized. The cavity is constructed from ultralow expansion glass (ULE) in order to minimize the cavity length's coefficient of thermal expansion. Figure 7.3 shows the cavity in its final resting position, with the vacuum can removed




Figure 7.3: The ultrastable, high-finesse, passive optical cavity for the Sr clock laser is made of ultralow-expansion glass. It is supported at its midplane in a vertical orientation in order to minimize the cavity length's sensitivity to accelerations in the vertical direction. The cavity will be enclosed in a temperature-controlled vacuum chamber mounted on a passive vibration-isolation stage in its final configuration. The orange tape covering the vacuum ports leading to the central bore of the cavity is first removed.

so the picture could be taken.

A second, independent cavity-stabilized laser system was also constructed for the purpose of comparison against the first system in order to characterize their performances [46]. The heterodyne beat between the two systems was first frequency counted, and after removing the linear drift (<1 Hz/s) from the data and dividing by $\sqrt{2}$ in order to characterize the instability of either system, the Allan deviation was measured to be 1×10^{-15} for timescales from 0.2 to 300 seconds, with the exception of a small (2×10^{-15}) bump at 100-s averaging times. The phase coherence between the two lasers was preserved at the level of 1 rad for longer than 2 s, at the optical frequency of ~ 430 THz. The limit to the cavity's performance was consistent with a theoretical model of the thermal-mechanical noise [167, 145] contributions from the mirror substrates and their dielectric coatings.

Put together, the lattice-trapped ^{87}Sr atoms and probe laser form a very stable and accurate frequency standard. With the “pendulum” of this clock now described, I will turn to its gears. The optical clockwork formed by the octave-spanning comb provides great flexibility to the kinds of measurements that can be performed with the lattice clock, allowing the absolute frequency to be measured, or allowing comparisons against many other optical clocks operating at various wavelengths, located in Boulder.

7.2 Octave-spanning Ti:Sapphire laser: optical design

As discussed at the beginning of this Chapter, it is desirable for the femtosecond comb employed as the optical clockwork to span an octave of bandwidth directly at its output, without the need for external broadening in microstructure fiber. In this context, “octave-spanning” should be taken to mean that there will be enough power at the octave points ν and 2ν in order to use the self-referencing technique to measure f_0 with a sufficient SNR for its electronic stabilization. Unfortunately, even Ti:sapphire

doesn't have an octave of gain bandwidth. Therefore some source of nonlinear spectral broadening inside the cavity must be exploited. In our octave-spanning laser, self-phase modulation inside the Ti:sapphire crystal itself is believed to be responsible for the spectral broadening the laser achieves. This has the added advantage that the action of the passive KLM mode-locking mechanism forces the pulses to remain transform-limited in the crystal, whereas microstructure fiber has higher-order dispersion that can limit the system's performance.

Figure 7.4 shows the optical design of the octave-spanning laser, similar to the design from Ref. [101]. A Verdi V-6 solid state laser outputs 6.84 W of 532-nm wavelength pump light, which is coupled into the laser cavity through a curved mirror (CM1) and focused into the Ti:sapphire crystal by a 10-cm focal length lens. The Ti:sapphire crystal is 2.6-mm thick, and higher-doped than a more-standard 5-mm thick crystal in order to achieve reasonable absorption of the pump laser. A second curved mirror (CM2) collimates the light from the crystal, and reflects it to the near prism (NP) of a CaF₂ prism pair for dispersion compensation. A fold mirror then reflects the light to the far prism (FP) which collects the spatially dispersed spectrum into a non-diverging, spatially-chirped beam. This beam is retroreflected from the cavity's end-mirror high-reflector (HR) in order to make a linear cavity geometry. After returning through the crystal, CM1 reflects the beam toward the retro-reflecting output coupler (OC). Model numbers of all the intracavity optics are given in Fig. 7.4. In order to generate an octave of spectral bandwidth, the laser is initially aligned so as to maximize the amount of cw power possible, typically as high as 800 mW. Then CM2 is translated in toward the crystal while maintaining the pointing alignment of the end-mirrors until the cavity is on the verge of stability for cw operation. Finally, both CM1 and CM2 are translated away from the incoming pump beam until mode-locking is achieved, while the insertion of prism FP is adjusted to maximize the spectral bandwidth. In its final configuration that gives the strongest f_0 beat from the self-referencing interferometer, the laser

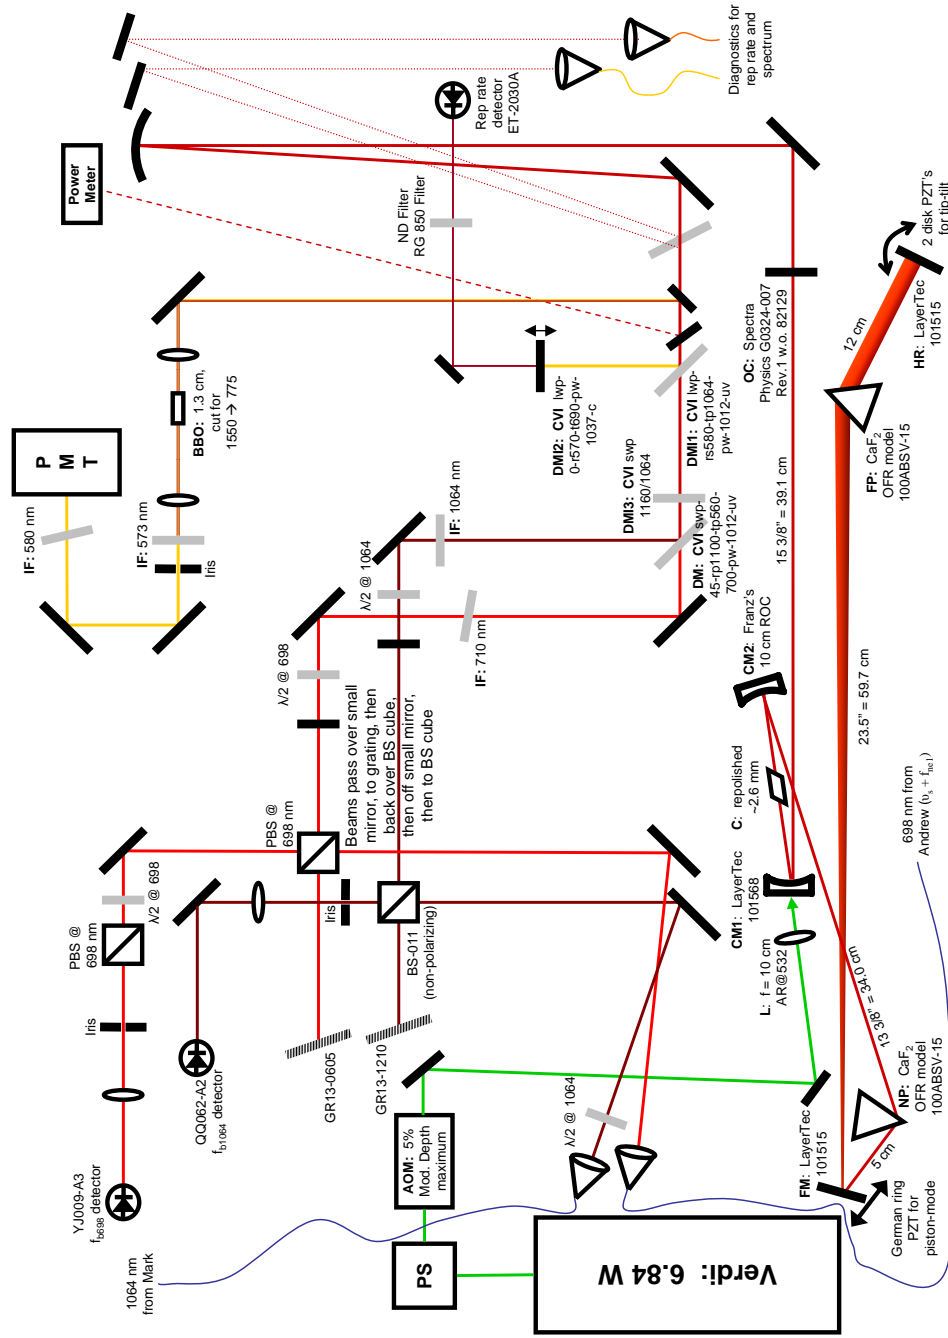


Figure 7.4: The octave-spanning laser and the optical system used to form the ⁸⁷Sr optical lattice clockwork. Abbreviations and details are given in the text. Dashed lines indicate beams from pick-off mirrors only used for diagnostics. The pick-off mirrors are removed during operation of the clockwork.

outputs ~ 80 mW of cw power and 350 mW of mode-locked power. Initiation of mode-locking is performed by tapping on the NP mount so as to decrease its insertion in the beam briefly, or by moving the translation stage for the OC briefly so as to increase the cavity length. Sometimes mode-locking can be initiated by pushing FP briefly into the beam.

The output from the laser is rapidly diverging, so after a steering mirror just outside the OC, a long-focal-length curved mirror is used to collimate the light. After another steering mirror, the spectrum and pulse timing can be measured by using the 4% reflections from the front and back surfaces of a thick glass plate. A flipper mirror after the plate also allows the power to be measured. By coupling one of the 4% pickoffs into a multimode fiber, the spectrum can be measured on a fiber-coupled optical spectrum analyzer (OSA). Figure 7.5(a) shows the output spectrum from the laser. Figures 7.5(b) and (c) show the transmission and dispersion of the OC, both of which are important for octave-spanning operation. Note that the OC is only highly reflective from roughly 760 to 980 nm, whereas it transmits the majority of any incident power of wavelength shorter than 720 nm or longer than 1050 nm. Within the OC's high reflectivity bandwidth, light is therefore resonant with the cavity, whereas outside that bandwidth light does not have to conform to the resonant spatial modes of the cavity. As a result, the observed profile of the beam is Gaussian for the central part of the spectrum, but is distorted outside this bandwidth. More details are given in Ref. [101]. Since the majority of the intracavity power in the wings of the spectrum leaks out of the OC on every round trip of the cavity, it must be generated in a single pass of the Ti:sapphire crystal. Self-phase modulation (SPM) is responsible for this spectral broadening. Finally, we investigated the dispersive properties of the OC by measuring its group velocity dispersion (GVD) using a white-light interferometer. As shown in Fig. 7.5(c), the dispersion is kept between -10 and $+20$ fs² from 800 to 930 nm. This means that the central portion of the spectrum returning from the OC to the Ti:sapphire crystal is minimally chirped, allowing the formation of

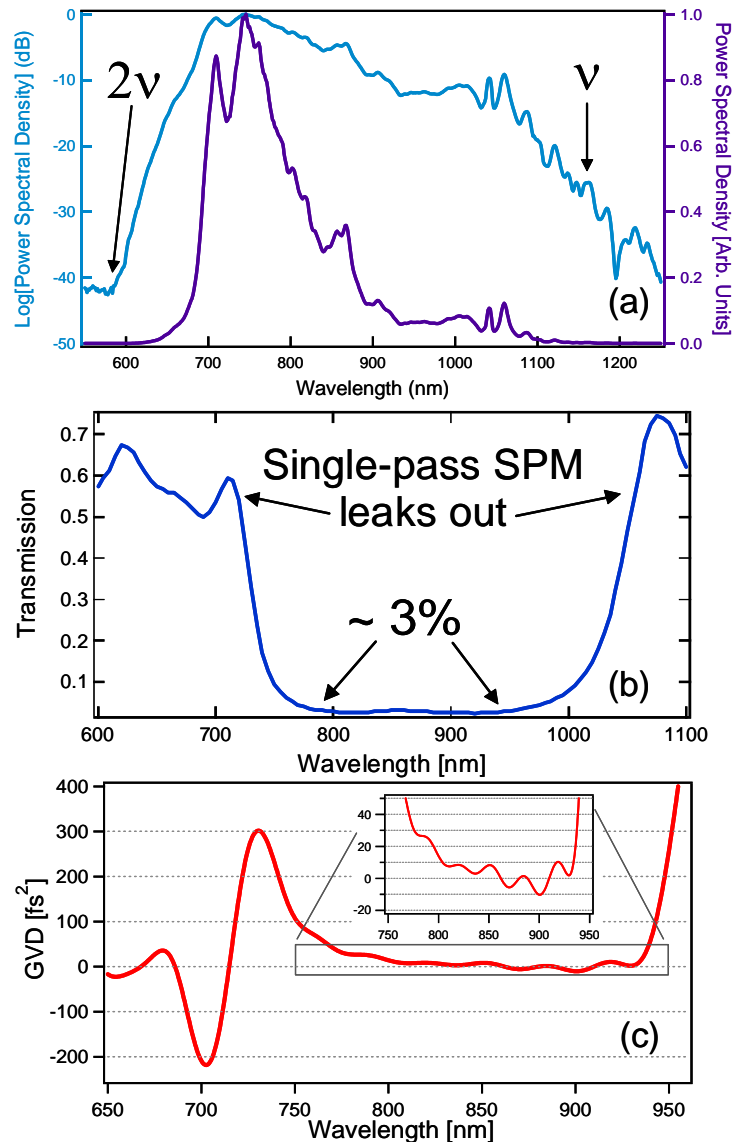


Figure 7.5: (a) Output spectrum of the octave-spanning laser, plotted on a logarithmic scale (left axis) and a linear scale (right axis) for the power spectral density. The spectral points used for self-referencing are at 575- and 1150 nm. (b) Transmission of the output coupler. The steep edges at 740 and 1040 nm allow the majority of power in the spectral wings, generated by self-phase modulation in the crystal, to leak out of the cavity in one round trip. (c) Group velocity dispersion of the output coupler, measured using white-light interferometry. The low dispersion in the output coupler's high-reflectivity bandwidth allows the intracavity pulse to remain short on its reverse path through the crystal, keeping the nonlinearity of the Kerr lens high in each direction.

a strong Kerr-lens effect as it passes through the crystal in this direction. Since the arm of the cavity containing the prisms can have its dispersion independently fine-tuned, the pulses inside the crystal can be very short as they travel in either direction through the crystal. This keeps the nonlinearity extremely high, making for a very robust mode-locking mechanism as well as the large amount of SPM responsible for the spectral broadening.

7.3 Servo control

The broadband spectrum provided by the octave-spanning laser is able to heterodyne with several important lasers for the characterization and development of the optical lattice clock. In this Section I will first describe the self-referencing optics and electronics which use the 575- and 1150-nm parts of the spectrum. Then I will describe the servo loop and optics needed to stabilize the comb to the 698-nm sub-Hz-linewidth clock laser. Finally, I will describe the in-loop measurements showing that these various servos have enough gain and bandwidth to, in principle, stabilize the entire comb relative to the clock laser within one mHz of relative linewidth at 698 nm.

7.3.1 Control of the offset frequency

In order to control the octave-spanning comb's offset frequency, f_0 , the self-referencing scheme described in Section 3.3.1 is used. Referring to Fig. 7.4, after the output of the fs laser is collimated with a curved mirror and reflected with a steering mirror, it passes just over the top of a mirror and is incident on a dichroic Michelson interferometer. The dichroic beamsplitter of the interferometer (DMI1) separates light below 600 nm from the rest of the spectrum. The majority of the spectrum (at wavelengths longer than 600 nm) is transmitted by DMI1, while the wavelength range near 575 nm is reflected to an end mirror of the interferometer (DMI2). DMI2 retroreflects the 575-nm portion of the spectrum with a slight downward tilt so that it can reflect from

the short mirror located just in front of the interferometer. A small amount of longer-wavelength power that reflected from DMI1 transmits through DMI2 and is detected on a fast photodiode in order to measure f_{rep} and its harmonics. In the meantime, a third dichroic mirror (DMI3) in the long-wavelength port of the interferometer reflects the 1150-nm portion of the spectrum, again with a slight downward tilt, while transmitting the central portion of the comb's spectrum between 600 and 1100 nm. The two retro-reflected beams from DMI2 and DMI3 are combined again on DMI1 before being picked off by the short mirror and sent to a focusing lens. The lens focuses both the 575-nm and 1150-nm beams into a BBO crystal, which is angle-tuned to phase-match the 1150-nm second harmonic generation process. After collimating the original 575-nm comb (with offset frequency f_0) and the new SHG comb (with offset frequency $2f_0$), also centered at 575-nm, a pair of interference filters are used to select only light at 575-nm for detection on a photomultiplier tube. Although the spectral bandwidth of the detected 575-nm light is quite narrow (~ 3 nm), it still represents pulsed light. Therefore DMI2 can be translated in order to account for the system's group delay between the original 575-nm pulses and the SHG pulses generated from the original 1150-nm pulses. This small group-delay correction is the reason for building the dichroic Michelson interferometer in the first place.

The PMT detects f_{rep} from the two pulse trains, but more importantly it detects the f_0 and $f_{rep} - f_0$ heterodyne beats between adjacent comb lines of the fundamental and SHG combs. The SNR of these heterodyne beats is typically 25 dB in 100-kHz, and has been observed as high as 35 dB in the same bandwidth whenever the optical alignment of both the laser and the self-referencing system have recently been perfected. After electronically filtering the f_0 beat at 70 MHz and amplifying it to a suitable level, we phase-lock an rf tracking filter to it. The rf tracking filter has a sufficiently high bandwidth to stabilize its voltage-controlled oscillator (VCO) to the f_0 beat in an extremely robust fashion. Then the VCO output at 70 MHz, which tracks f_0 robustly, is digitized

and prescaled to 1.75 MHz by a digital frequency divider. This frequency-divided square wave is effectively mixed against the output of a stable external synthesizer that has also been digitized, using a digital phase detector, in order to measure the relative phase between the f_0 beat and the reference synthesizer. The voltage signal corresponding to the relative phase is then integrated by a loop filter in order to have proportional-integral control of the error signal. The shaped error signal is split into two copies, and one is inverted. Both copies are summed with independently-controllable dc offsets and amplified to as much as 150 V in order to differentially drive two PZT's that control the tilt of the fs laser's HR. Whenever one PZT's voltage swings high, the other drops low in order to tilt the HR, thus changing the group delay dispersion of the laser, and therefore shifting f_0 . When the servo loop is closed, the feedback therefore phase-locks f_0 to the 70-MHz external synthesizer reference. The servo loop has a bandwidth of 30 kHz, limited by mechanical resonances of the differential-dual-PZT system.

Note that to obtain 1 mHz of stability or accuracy of f_0 , corresponding to 1 mHz stability or accuracy of every optical frequency mode of the comb, the 70-MHz reference synthesizer must have 1 mHz of stability or accuracy. This is only a fractional stability or accuracy requirement of $1 \text{ mHz} / 70 \text{ MHz} = 1.4 \times 10^{-11}$, easily within reach of a good microwave frequency reference's specifications. We stabilize all our reference synthesizers to a commercial Cs-beam microwave frequency reference with a stability of better than 10^{-12} at 1 s, and an accuracy that is maintained by steering it to a GPS receiver and therefore coordinated universal time (UCT). The steering is performed as a step function frequency shift, often enough (on the order of weeks between corrections) to keep the integrated time within certain bounds. The Cs oven-based reference is also continuously monitored against GPS so the data is available for post-processing if it is needed.

7.3.2 Control of another comb mode

The comb's remaining degree of freedom, f_{rep} , is controlled by phase-locking a neighboring comb mode to the 698-nm ultrastable clock laser. The 600-to-1000-nm spectral components transmitted by DMI3 are split on another dichroic mirror (DM) that transmits 698-nm light while reflecting 1064-nm light. The 698-nm light passes through a narrow interference filter and a half-wave plate, before it is sent to a grating. The grating is used to spatially disperse the beam so that less than a nm of spectral bandwidth will be heterodyned with the 698-nm clock laser, in order to prevent saturation of the photodetector's ac electronics due to the signals at harmonics of f_{rep} . After the grating, the comb's light is combined with the cw 698 clock laser on a polarizing beamsplitter cube (PBS). The clock laser's light is delivered to the optical table by a ~ 15 -m optical fiber with a fiber transfer phase noise cancellation scheme like the one to be shown in Fig. 8.4. After combining it with the comb's light, the orthogonal polarizations are collapsed and the collinear beams are detected on a silicon photodetector. The output of the photodetector responds to the heterodyne beats between the clock laser and all nearby modes of the frequency comb, as well as f_{rep} and its harmonics that roll off after the 180-MHz bandwidth of the photodiode. The two lowest-frequency beats are defined as f_{b698} and $f_{rep} - f_{b698}$, respectively.

After filtering f_{b698} near 25 MHz, it is phase-stabilized to another rf frequency reference at 25 MHz using the same digital phase detection technique as described for the f_0 lock, with a prescaling factor of 20. Several actuators are available for controlling f_{rep} , including an AOM in the pump laser for its intensity control, a fast PZT for driving the laser's intracavity fold mirror (FM), and a slow, long range PZT for driving the translation stage that the OC is mounted on. The error signal from the digital phase detector is first split; one part is integrated once and used to drive the AOM with a bandwidth of ~ 150 kHz. The other part is integrated once before driving the high-speed

PZT, and twice before driving the long range PZT in order to keep the average voltage of the high-speed PZT within its dynamic range. The resulting phase-locked-loop has a unity gain bandwidth of 150 kHz, limited by the delay time of the acoustic wave through the AOM to the location of the pump beam. Other mechanical resonances occur for the fast PZT at 50 kHz and for the long-range PZT below 100 Hz.

A special feature of the octave-spanning Ti:sapphire laser that sets it apart from more typical systems is that control of the pump’s intensity affects Nf_{rep} about a factor of 20 more than it affects f_0 for the comb mode near 698 nm, for modulation frequencies of several-hundred Hz. It has been shown that operating Ti:sapphire lasers with broader spectra leads to less intensity-dependence of f_0 [106], and the octave-spanning laser takes this to the extreme. Since f_0 is therefore more intrinsically stable for the broad octave-spanning laser, the relatively low 30-kHz bandwidth used to stabilize f_0 is more than sufficient, despite the pump intensity being driven with the high servo bandwidth of 150 kHz in order to control f_{rep} .

7.3.3 In-loop performance of the stabilized comb

In order to verify that our servo loops have enough gain to transfer the full phase-coherence of the 698-nm clock laser’s sub-Hz linewidth to each mode of the frequency comb, we measured the in-loop linewidths of both f_0 and f_{698} . They are plotted on linear and logarithmic scales in Fig. 7.6. For these measurements, both the f_0 and f_{698} servos were engaged in order to measure the full experimental situation of the comb locking to the 698-nm clock laser while simultaneously being self-referenced. This measurement reveals in-loop linewidths of 1 mHz, limited by the resolution bandwidth of the measurement. Therefore, the servo gains and bandwidths are sufficient to stabilize the comb extremely tightly to the clock laser.




Figure 7.6: In-loop linewidths of f_0 (bottom graphs) and f_{b698} (top graphs) on linear (left graphs) and logarithmic (right graphs) scales for power spectral density. All measurements were taken while both degrees of freedom were phase-locked. This demonstrates the feedback loops have enough gain and bandwidth to transfer the coherence of the 698-nm clock laser to the rest of the comb's modes across the spectrum.

7.4 Out-of-loop experiments

After constructing and testing the octave-spanning laser as described in the previous Section, we put it to use for several applications involving the ^{87}Sr optical lattice clock. First, the comb proved to be an invaluable tool for the development of the first 698-nm clock laser, since the only other stable optical reference we could compare it to was an independently-characterized cavity-stabilized laser operating at 1064-nm. The large spectral gap between the two sub-Hz lasers had to be covered with the octave-spanning comb.

Then the comb was used for both the initial 5×10^{-14} [146] and subsequent 2.5×10^{-15} [147] accurate measurements of the clock transition's absolute frequency by comparing the tenth harmonic of its repetition frequency to a ~ 1 -GHz signal referenced to a Hydrogen maser calibrated to the Cs fountain primary standard at NIST.

Next, we used the comb to stabilize a 1064-nm cw laser to the 698-nm clock laser, so that the 1064-nm transfer laser could be transmitted over a 3.5-km optical fiber to NIST for comparisons against other optical clocks. Details of the fiber transfer itself will be discussed in Section 8.4, but here I will briefly present some preliminary data highlighting the utility of making frequency intercomparisons directly in the optical domain between optical clocks. Essentially, the significantly-lower instability offered by the optical clocks (as compared to the maser-referenced signal) allows extremely rapid averaging when making measurements of systematic clock shifts. Another result from these experiments is a 1-Hz-linewidth intracavity comparison between a NIST cavity-stabilized laser and the Sr lattice clock laser, demonstrating the ability to preserve phase coherence over large-spectral- and long-distance- gaps at the level that current technology can produce optical phase coherence!

Finally, I will present preliminary results from a new measurement of the octave-spanning comb's ability to phase-lock to a cw reference laser. The comb is locked to a

cw laser operating at 1064 nm, and heterodyned against the cw laser's second harmonic at 532 nm (while independently locking f_0). Then the beat frequency from the 532-nm beams allows an out-of-loop measurement of the comb's ability to both lock to the 1064-nm reference and phase-coherently transfer the reference's exact timing across the spectrum to the 532-nm mode of the comb. Preliminary results indicate that a 1-mHz relative linewidth is preserved across the comb's spectrum, with inaccuracy of the comb lines constrained to below a level approaching 10^{-19} . Such results imply that the octave-spanning comb is fully capable of accurately and phase-coherently distributing the quality of our best optical frequency standards across the visible and near-infrared spectrum.

All four of these results, taken together, prove that the octave-spanning frequency comb can accurately and phase-coherently distribute an optical frequency reference to the microwave, infrared, and visible parts of the spectrum, even over distances as long as 32 km. Such a versatile and accurate tool represents the fully-characterized hub of an active and robustly-working optical frequency metrology laboratory.

7.4.1 Characterizing the clock local oscillator

The first direct application for the octave-spanning frequency comb in our laboratory was to characterize the clock laser being developed for the ^{87}Sr optical lattice clock. At the time, only one cavity was completed for the 698-nm wavelength, and to characterize the laser's performance we needed to bridge the spectral gap between it and another cavity-stabilized laser operating at 1064 nm [45] The 1064-nm laser was already well-characterized against two other systems operating close enough to its frequency that direct heterodyne beats could be used for intercomparison.

In order to make the measurement, light from both the 698-nm laser and the 1064-nm laser was brought to the octave-spanning comb via fiber links with active cancellation of the phase noise introduced by each ~ 15 -m-long fiber (using the technique

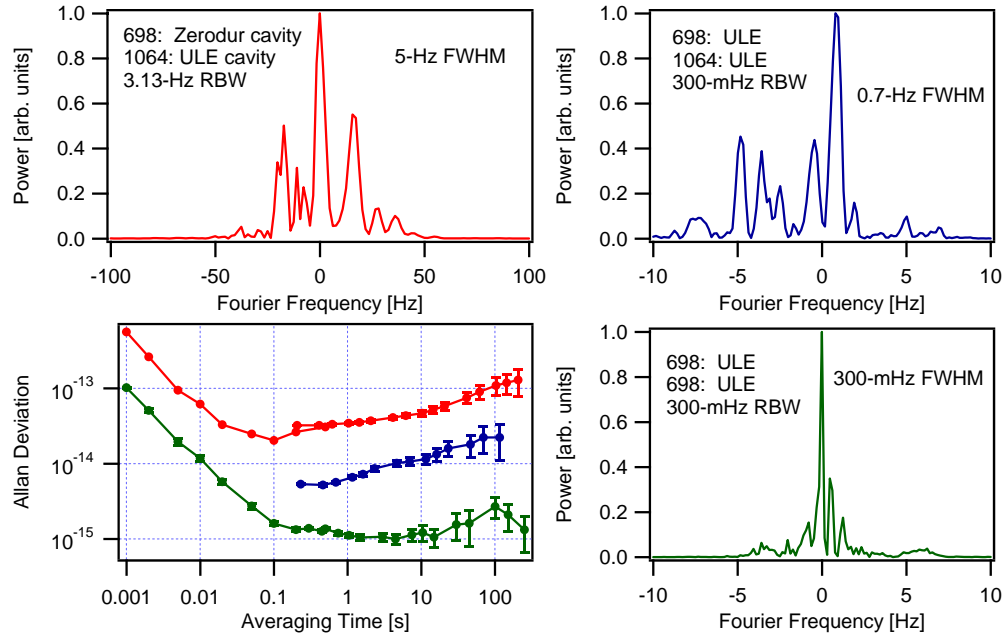


Figure 7.7: Characterizing the 698-nm clock laser against a known 1064-nm narrow-linewidth reference laser. Top-left (red): 5-Hz linewidth of the first 698-nm laser stabilized to a Zerodur cavity, as compared to the ULE 1064-nm cavity via the octave-spanning comb. The 698-nm laser’s frequency noise limited this measurement. Top-right (blue): 0.7-Hz linewidth of the second 698-nm laser stabilized to the ULE cavity shown in Fig. 7.3, when effectively heterodyned against the ULE 1064-nm cavity via the octave-spanning comb. The 1064-nm laser’s frequency noise limited this measurement. Note the factor of 10 reduction in the frequency axis scale between the top-left plot and the other linewidth plots. The narrow features represent the “linewidth” quoted here, since sometimes the energy is seen to compress into these features. Shown here, typical laser frequency noise spreads energy into other noisy features. Bottom-right (green): Two independent 698-nm lasers beat together to record this 300-mHz linewidth, as in Ref. [46]. No octave-spanning comb is involved, but the measurement indicates the octave-spanning laser, when locked to the 300-mHz line, likely can transfer the full phase coherence to every mode of the comb. Bottom-left: Allan deviations for the same three cases, color-coded respectively. The octave-spanning comb enjoys the 1×10^{-15} stability provided by the 698-nm ULE cavity, distributed to each mode of the optical frequency comb.

to be described in Section 8.4). The frequency comb was self-referenced and locked to the 698-nm laser as described in Section 7.3. Then another mode of the comb was heterodyned against the 1064-nm laser as the out-of-loop effective beat between the two ultra-stable lasers. A typical measurement of the linewidth of the effective beat is shown in the top-left plot of Fig. 7.7. The central feature is only 5 Hz wide, although significant energy is spread into a 50-Hz span. We also frequency-counted the effective beat note in order to measure the relative instability of the two cavities. After subtracting a linear drift term from the data, the Allan Deviation was calculated as shown in the red curve in the bottom-left plot of Fig. 7.7. The $\sim 4 \times 10^{-14}$ fractional frequency instability and the linewidth were both limited by the 698-nm cavity, which used a different geometry from the cavity shown in Fig. 7.3, and the material Zerodur for the spacer and mirror substrates. Nevertheless, our first absolute frequency measurement of the ^{87}Sr clock transition [146] was performed using the Zerodur cavity.

After the first absolute frequency measurement was finished, Andrew Ludlow constructed the cavity shown in Fig. 7.3 and described in Section 7.1.3. Again, the octave-spanning comb was used to compare the new 698-nm laser to the same 1064-nm laser as before. A typical linewidth measurement is shown in the top-right plot of Fig. 7.7 in blue. The 0.7-Hz central feature again shows the measurement's potential, but more often than not significant energy is spread into a ~ 5 -Hz span. The fractional instability of the measurement at 1064-nm is shown as the middle curve (blue) in the bottom-left plot of Fig. 7.7. The significant improvement over the Zerodur cavity is due to the smaller thermal noise contribution from ULE, as well as some engineering issues (for details, please see Andrew Ludlow's thesis). This time, the measurement was likely limited by the 1064-nm laser. In order to prove it, Andrew built a second 698-nm laser stabilized to a similar ULE cavity and directly compared the two 698-nm lasers [46] without need for the octave-spanning comb. A 300-mHz linewidth of the heterodyne beat between the two independent 698-nm clock lasers is shown in the bottom-right

plot of Fig. 7.7. Also, the instability of the lasers (after dividing the frequency-counting data by $\sqrt{2}$) is shown as the green curve in the bottom-left plot. The 1×10^{-15} flicker floor for averaging times longer than 0.2 s agrees well with theoretical predictions of the thermal noise floor for such a cavity [167, 46].

The first two cavity comparisons made with the octave-spanning laser were an encouraging start for demonstrating the laser could phase-coherently transfer the ultrastable performance of the clock laser across its spectrum. In addition, the octave-spanning comb was a valuable tool for characterization of the clock lasers, so that rapid development could be made while knowing what the performance levels were. Presumably, the comb was also able to lock tightly enough to enjoy the 1×10^{-15} performance of the ULE cavity, although only in-loop data existed to support this conclusion until recently. See the new measurements described in Section 7.4.4 for further out-of-loop proof.

7.4.2 Absolute frequency measurement of the ^{87}Sr optical lattice clock

With the clock's local oscillator suitably characterized, we began to make absolute frequency measurements of the clock transition. The clock laser always has some degree of linear and nonlinear frequency drift that must be controlled to a suitable level. In the case of our first absolute frequency measurement [146], we were scanning lines that were several-hundred-Hz wide, and fitting the measured lineshapes to find the center frequencies. For our next measurement [147], the clock laser had been improved and we scanned 10-Hz-wide lines. Therefore, the drift rates had to be kept significantly below 500 Hz and 10 Hz over the course of the 60 and 30 seconds it took to scan the line for the two measurements, respectively.

In order to keep this drift to a minimum, we stabilized the long-term drift to a Cs-calibrated H-maser located at NIST, using the drift compensation scheme shown in Fig. 7.8. (The scheme shown was used for the measurements in Ref. [147]. The earlier

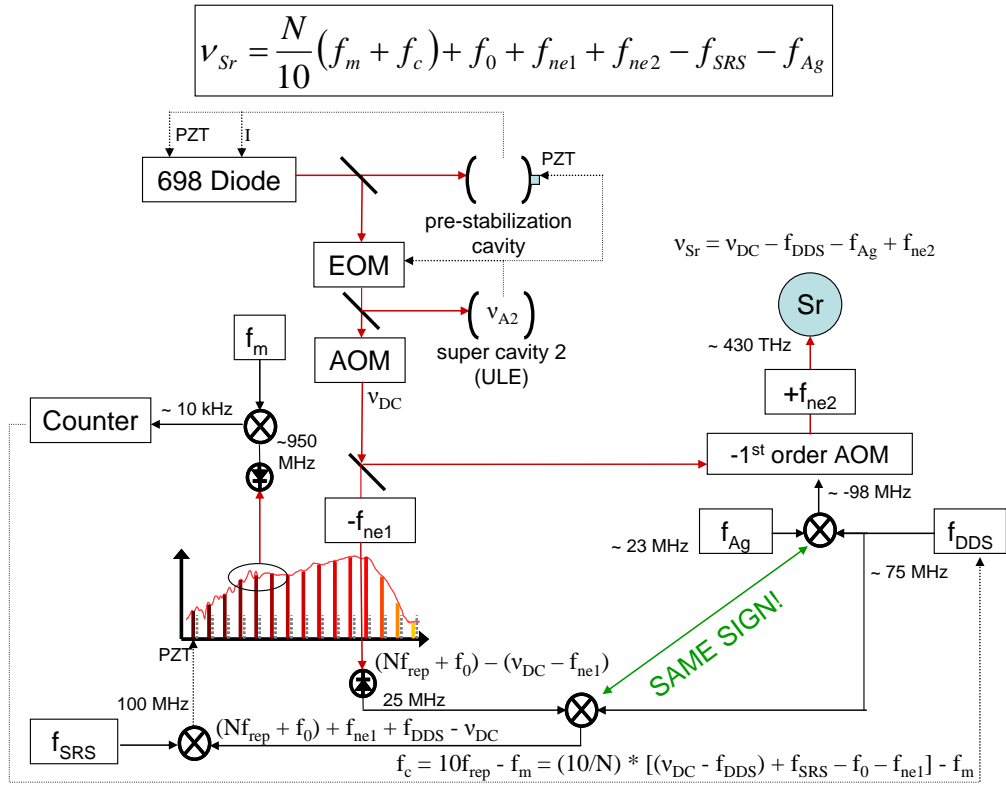


Figure 7.8: Schematic showing the ^{87}Sr clockwork, including the various frequency offsets used to steer the laser onto the clock transition. The frequency f_m is referenced to the maser, at ~ 950.5 MHz. Mixing the tenth harmonic of f_{rep} against f_m while stabilizing the comb to the clock laser yields a 50-kHz signal for counting, f_c . Measuring the drift of f_c determines the drift of the cavity frequency ν_{A2} with respect to the H maser. Then f_{DDS} is used to compensate that drift so that the laser's frequency ν_{A2} is stable with respect to the Sr transition frequency, ν_{Sr} . Measurement of f_c also allows the ratio ν_{Sr}/f_m to be determined accurately, and since f_m is calibrated by the Cs fountain clock, the absolute frequency of the Sr transition frequency is determined.

measurement in Ref. [146] used a similar scheme with different gain signs and offset frequencies.) The microwave reference signal from NIST was transmitted over optical fiber by amplitude modulation of a fiber-coupled laser diode, as will be described in Section 8.3.3. A synthesizer at NIST provided the modulation frequency of $f_m \sim 950$ MHz, close to the tenth harmonic of f_{rep} for the octave-spanning comb when locked to the clock laser.

Measurement of the clock transition frequency was performed as shown in Fig. 7.8, and worked as follows. First, the clock laser was stabilized to its pre-stabilization cavity and the ultrastable cavity, and its light was brought through a fiber noise canceller's AOM down to a breadboard from which the light could be split to the atoms and the comb. I denote this frequency at the "Distribution Center" as ν_{DC} . The light that was sent to the comb passed through the -1 order of another AOM for fiber noise cancellation ("noise eater") of a 15-m fiber, which frequency-shifted the light by $-f_{ne1}$. Light at the output of the fiber was heterodyned against the next-highest-frequency mode of the comb (beat frequency $Nf_{rep} + f_0 - (\nu_{DC} - f_{ne1})$). Then a programmable DDS synthesizer (frequency f_{DDS}) was summed into the beat frequency, for the purpose of cancelling the long-term drift of the clock laser relative to f_m from the maser. The sum of f_{DDS} and the beat frequency was stabilized to a synthesizer (frequency f_{SRS} by feedback to the octave-spanning laser as described in Section 7.3.

With the comb locked to the clock laser, including the possibility for summing in a frequency drift through control of f_{DDS} , the clock laser frequency was counted against the maser: first, a portion of the comb's light was detected in order to derive a signal at $10f_{rep}$. Then f_m was controlled to be ~ 10 kHz less than $10f_{rep}$ so that the two signals could be mixed. The mixer output was low-pass filtered to include only the difference frequency term at the frequency we then counted:

$$f_c = 10f_{rep} - f_m = \frac{10}{N} [(\nu_{DC} - f_{DDS}) + f_{SRS} - f_0 - f_{ne1}] - f_m. \quad (7.1)$$

Note that feedback to the value of f_{DDS} could then be used to stabilize the long-term drift of $\nu_{DC} - f_{DDS}$ relative to the maser. Such feedback was applied by counting f_c for 30 s at a time and then fitting the data to a line to determine the linear drift rate. The linear drift was calculated by averaging (typically four) of these 30-s measurements, and then the drift rate of f_{DDS} was updated in order to keep the drift of f_c well-zeroed.

Figure 7.9 shows the performance of the drift compensation when used to compensate the drift of the original Zerodur cavity-stabilized laser used for the measurements in Ref. [146]. At first, drift compensation was not applied, and the average drift was ~ 7 Hz/s. Then as the drift compensation was turned on and feedback to f_{DDS} was applied every 100 s, the drift rate rolled over to an average value close to zero. Since the net frequency displacement at any time is given by the integrated drift rate, and the lock has a finite error with respect to the drift rate, the frequency displacement could still execute a somewhat random walk. However, the “instantaneous” frequency drift at any time (say averaged for the neighboring 30 s, which was our scan time) was kept below 1 Hz/s. The bottom plot of Fig. 7.9 shows how the linear frequency drift, if left uncompensated, would have drifted very far from resonance after just a few hundred seconds. With the new ULE cavity-stabilized laser described in Section 7.1.3, the average uncompensated drift rate is below 1 Hz/s, and the nonlinear fluctuations that lead to the ~ 100 -s ripple in Fig. 7.9 no longer appear.

In order to make sure the atoms saw the same drift-free probe laser, f_{DDS} was also summed in to an AOM used to tune the probe laser frequency to the atomic resonance. From the Distribution Center again, a portion of the light was sent through the -1 order of this AOM being driven at the frequency $f_{Ag} + f_{DDS}$. The signal at frequency f_{Ag} came from an Agilent synthesizer that was used to scan through the atomic resonance. In addition, another AOM operating in the $+1$ order at frequency f_{ne2} was used to cancel the fiber noise of the 15-m fiber between the Distribution Center and the lattice-trapped




Figure 7.9: Compensation of the clock laser's frequency drift. The upper plot shows the optical frequency of the clock laser as a function of time in red. For the first 200 s, no drift compensation is applied. Then feedback is applied by fitting the slope every 100 s and applying the opposite drift to f_{DDS} . The lower plot illustrates (in blue) how quickly the 7 Hz/s drift of the probe laser would have walked off resonance with the atoms if left uncompensated. Note that in addition to compensating drift, the system also allows real-time measurement of the absolute frequency of the probe laser relative to the Hydrogen maser located at NIST.

atoms. Therefore the probe frequency at the atoms was

$$\nu_{Sr} = \nu_{DC} - f_{DDS} + f_{ne2} - f_{Ag}. \quad (7.2)$$

Since the terms $\nu_{DC} - f_{DDS}$ were kept drift-free relative to f_m by feedback from f_c , the light at the atoms was also free of drift during a scan of the atomic resonance's line shape. By recording the changing frequencies f_c and f_{Ag} , as well as knowing the stable frequencies f_m , f_0 , f_{ne1} , f_{ne2} , and f_{SRS} , the frequency of light seen by the atoms at any point in time was determined: solving Eqn. 7.1 for $\nu_{DC} - f_{DDS}$ and inserting the expression into Eqn. 7.2, we have

$$\nu_{Sr} = \frac{N}{10} (f_m + f_c) + f_0 + f_{ne1} + f_{ne2} - f_{SRS} - f_{Ag}. \quad (7.3)$$

Note this expression conveniently does not depend on f_{DDS} , which was only used to compensate the drift of the clock laser relative to the maser.

Then after scanning through the line and fitting to determine its center frequency, ~ 30 s of f_c data (from times symmetrically around the time at which the center frequency was measured) was used to determine ν_{Sr} for each line center. The accuracy of this determination was limited by the 30-s instability of the maser to $\sim 6 \times 10^{-14}$. Averaging of many scans of the line over the course of a 24-hour period (see Section 8.3.3 and Fig. 8.9) allowed us to determine the absolute frequency to an uncertainty of 1.1 Hz, or 2.5×10^{-15} [147], limited mainly by the calibration of the maser to the Cs-fountain frequency standard.

In addition to using the octave-spanning comb to make the absolute frequency measurements by accurately dividing the optical clock frequency down to microwave frequencies, the drift cancellation has also allowed the highest-quality-factor coherent spectroscopy ever to be performed [47]. With the low residual drift rate, typically ≤ 0.1 Hz/s, scans of individual resolved hyperfine magnetic sublevels (with a bias magnetic field applied) were able to measure Fourier-limited linewidths of 1.8 Hz. An

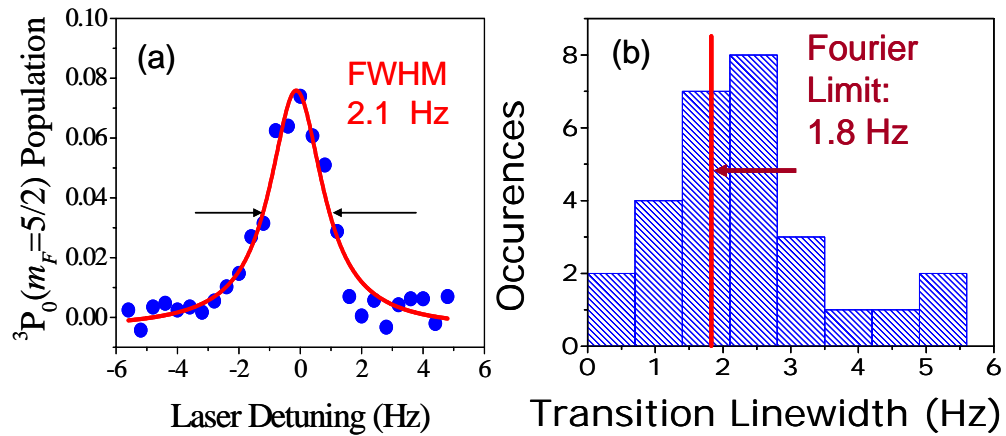


Figure 7.10: The highest quality factor yet demonstrated for coherent spectroscopy. (a) Lineshape of a single hyperfine magnetic sublevel for the clock transition in ^{87}Sr under a bias magnetic field, demonstrating a 2.1-Hz linewidth. Note each point corresponds to a single cooling and trapping cycle, so no signal averaging is used. (b) Histogram of measured linewidths, demonstrating that probe laser frequency noise can lead to narrower and broader linewidths than the 1.8-Hz Fourier limit set by the interrogation time. Long-timescale compensation of the probe laser's frequency drift was necessary in order to separate the linear frequency drift from the nonlinear fluctuations (from thermal noise of the ultrastable cavity) that limit the linewidth.

example of a 2.1-Hz linewidth is shown in Fig. 7.10(a). Figure 7.10(b) shows a histogram of many measurements demonstrating that the probe laser frequency noise can decrease or increase the measured linewidth, but that on average the Fourier limit is achieved. Such a narrow resonance at 430 THz represents a record-high quality factor for coherent spectroscopy of 2.4×10^{14} . Attempts at longer probe times were limited by the probe laser's frequency noise (due to the ultrastable cavity's thermal noise limit). Again, these linewidths were recovered when the octave-spanning comb was used to compensate the probe laser's linear drift relative to the maser; only the nonlinear frequency fluctuations from thermal noise remain to limit the linewidth.

Recently we also operated the system as a true clock rather than by taking spectroscopic data of the line center. By applying a bias magnetic field and alternately probing the two stretched hyperfine sublevels, the linear Zeeman frequency shift for the clock was cancelled when the probe laser was locked to the average frequency of the two stretched transitions. Locking the probe laser to the average frequency was achieved by controlling f_{Ag} , but in order to also lock the comb's output to the atoms, f_{SRS} was also controlled by the same amount with the appropriate sign. Then, as described above, the frequency comb derived a true countable microwave clock output at $10f_{rep}$, which was mixed against f_m to form the ~ 50 -kHz signal at f_c containing the relative frequency fluctuations between the maser and the Sr clock's microwave output.

The upper plot of Fig. 7.11 shows frequency-counting data when f_c was counted with a 1-s gate time and using a 100-kHz low-pass filter. The Sr optical lattice clock operated continuously for nearly 8 hours with only a brief interruption at about the 3-hour mark. The interruption allowed the clock laser to drift before an updated value of f_{Ag} was used upon re-acquisition of the lock to the atoms. Due to the particular programming of our system, f_{SRS} was left unchanged during the un-locked time and so did not pick up the accumulated frequency offset that f_{Ag} did. Therefore the comb's microwave output experienced a net frequency offset due to the event. The lower plot of




Figure 7.11: The Sr optical lattice clock’s performance, as directly compared against the Hydrogen maser signal. Top: Frequency-counting the signal at frequency f_c , from mixing the clock’s microwave output near 950-MHz with the Hydrogen maser signal. Fractional frequency fluctuations between the two ~ 950 -MHz signals are displayed on the vertical axis. The counting gate time was 1 s, and f_c was near 50 kHz, filtered with a low-pass filter at 100 kHz. The Sr clock laser was locked to the atoms continuously for nearly 7.5 hours with the exception of a brief unlock at the 3-hour mark, which allowed the probe laser to drift before it was re-locked to the atoms with a new offset (from f_{SRS}) as seen by the comb. Bottom: the Allan Deviation, calculated from the last 4.5 hours of the data shown in the top plot. The stability is limited at all timescales by the maser.

Fig. 7.11 shows the Allan Deviation calculated from the last 4.5 hours of the frequency-counting data; the instability of the comparison is limited by the maser at all timescales shown.

Operation of the optical atomic clock for such a continuous span of time is an encouraging indication that the system can eventually be made robust enough to operate for long spans of time without cycle slips or loss of information. The ability to continuously count the microwave output of such a fully-functioning clock for so long is of fundamental importance to the application of time-keeping for such an accurate frequency standard. In the future, measurements of the frequency comb’s microwave output stability should be made in order to determine whether the low instability of the clock in the optical domain (to be described in Section 7.4.3) is fully preserved in the microwave output. Measurements of the frequency comb’s ability to transfer an optical reference’s stability into the microwave domain have already been performed in Ref. [78], indicating that the excess instability introduced by the optical-to-microwave conversion of the frequency comb was 6.5×10^{-16} for a 1-s averaging time. The low-noise optical-to-microwave conversion techniques described in Chapter 4 should be able to further reduce this level of instability.

7.4.3 Optical clock intercomparison

With the clock laser locked to the atoms, direct comparisons between the ^{87}Sr optical clock and other optical clocks could be performed. Figure 7.12 shows a high-level diagram of how we compared the Sr optical lattice clock to the Ca optical clock located at NIST [150]. As described in Section 7.4.2, the 698-nm clock laser was locked to the Sr atoms using the stretched magnetic hyperfine sublevel transitions under a small magnetic bias field. Then a mode of the octave-spanning comb was locked to the clock laser, while f_0 was locked, as described in Section 7.3. With every mode of the optical frequency comb now referenced to the Sr atoms, a cw laser (“transfer laser”)

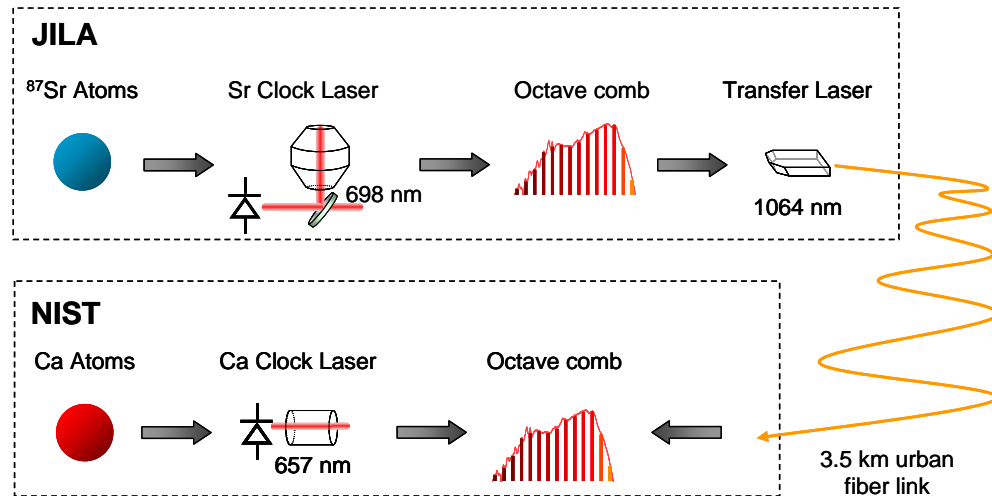


Figure 7.12: High-level diagram of the experimental setup for remote comparison of the ^{87}Sr optical lattice clock with the neutral-atom Ca clock. At JILA, the 698-nm clock laser is locked to the ^{87}Sr atoms, and the octave-spanning comb is self-referenced and locked to the 698-nm clock laser, as discussed in Section 7.4.2. A cw transfer laser operating at 1064 nm is stabilized to the comb and transmitted to NIST through 3.5 km of fiber installed in an urban environment. Details of the fiber transfer will be given in Section 8.4. At NIST, another ultrastable clock laser operating at 657 nm is locked to the Ca atoms, and another octave-spanning comb is again self-referenced and stabilized to the Ca clock laser. Then the cw transfer laser at 1064 nm coming from JILA is heterodyned against a nearby mode of the frequency comb at NIST. This heterodyne beat is the effective beat between the two optical frequency standards separated by 3.5 km of distance and by a spectral gap of 27 THz and therefore contains all the information about their relative frequency fluctuations.

operating at 1064 nm was locked to the comb. Then the transfer laser was transmitted to NIST via a 3.5-km urban fiber link, as will be described in detail in Section 8.4. The noise introduced by the fiber link was actively cancelled to an insignificant level for these measurements. Meanwhile, at NIST another ultrastable laser (Ca clock laser) was locked to the Ca clock transition [150] and another independent octave-spanning frequency comb [168] was self-referenced and stabilized to the Ca clock laser. Therefore every mode of the octave-spanning comb at NIST was accurately referenced to the Ca atoms, while the incoming cw transfer laser from JILA was accurately referenced to the Sr atoms. The cw transfer laser was heterodyned against a nearby mode of the NIST octave-spanning comb in order to generate a beat note containing the relative frequency fluctuations between the two optical frequency standards.

The effective beat between the Sr and Ca optical clocks was frequency-counted in order to measure the relative frequency fluctuations between the optical clocks. At the same time, the usual frequency f_c comparing the Sr optical clock's microwave frequency output to the maser was also counted. Figure 7.13 shows frequency counting measurements for a particular night where the Ca clock was being used as a stable reference against which to measure systematic frequency shifts of the Sr system. We used a 1-s gate time for all the frequency-counting measurements shown. The upper plot shows the fractional frequency fluctuations of the microwave clock output (relative to the maser) in red, on the order of a part in 10^{12} , consistent with the usual $\sim 3 \times 10^{-13}$ instability of the maser at 1-s averaging time. On the same scale, the fractional frequency fluctuations between the Sr and Ca optical clocks are shown in blue. The microwave data does not usually suffer from excess noise such as at times beyond 7 hours, but for this data it was having problems; its performance is usually similar to the data between 4 and 6 hours, or the data shown in Fig. 7.11. However, the real point of Fig. 7.13 is to illustrate how the optical frequency standards have nearly 2 orders of magnitude less fluctuations than the microwave comparison limited by the maser. The bottom plot of




Figure 7.13: Frequency counting the Sr clock against both the Hydrogen maser and the Ca optical clock at NIST. The upper plot shows fractional fluctuations of the Sr optical clock's microwave output (in red) at $10f_{rep}$ (near 950 MHz) relative to the Hydrogen maser, which limits the measurement. Also, on the same scale (in blue) the fractional frequency fluctuations between the Sr and Ca optical clocks (as measured at 282 THz) are shown. Plotting the two measurements on the same vertical axis scale dramatically shows the advantage of the optical clocks: the short-term instability is ~ 2 orders of magnitude smaller than for the maser. The bottom plot shows a magnified axis for the direct optical clock intercomparison. The frequency ripple in the optical comparison is because large systematic shifts were being checked for the Sr clock by comparing it to the Ca clock; every 100 s the systematic parameter was varied. The ability to compare two remotely-located optical clocks continuously for hours at a time is a powerful initial demonstration of the robust operation of these new frequency standards.

Fig. 7.13 shows the optical comparison on a scale magnified by this factor of 100, for comparison to the microwave data. The frequency ripple of the optical comparison is due to intentional changes of the Sr clock's operating parameters, for the purposes of measuring systematic frequency shifts against the stable Ca optical reference. Every 100 seconds a systematic effect in the Sr clock was varied.

By calculating the Allan Deviation from the optical frequency comparison measurements, the relative instability of the two optical clocks was also determined as a function of averaging time. The two clocks were operated continuously without intentionally changing any systematic effects, while the beat frequency between the NIST octave-spanning comb and the cw transfer laser was counted. The calculated Allan Deviation is displayed in Fig. 7.14, along with a typical measured Allan Deviation of the microwave Sr clock output relative to the maser. The optical clocks demonstrate a factor of nearly 50 less instability than the microwave comparison (limited by the maser) at all averaging times shown.

The ability to compare two remotely-located optical frequency standards continuously for hours at a time is a powerful initial demonstration of the robust operation of these optical frequency standards. It also illustrates the importance of the optical frequency comb as the information hub of a robustly-working optical atomic clock. Indeed, through two frequency combs and a 3.5-km phase-stabilized fiber link, the two clocks (Sr and Ca) separated by 27 THz are now being routinely compared at the 10^{-16} level of stability. This capability is already being exploited at the time of this writing to check systematic shifts of the ^{87}Sr optical lattice clock at the 10^{-16} level of accuracy.

7.4.4 The comb's intrinsic coherence, stability, and accuracy

Through all of the frequency comparisons described in the last two Sections, we assumed that the octave-spanning comb did not introduce any inaccuracy in the optical-to-microwave conversion or in distributing the 698-nm clock laser's frequency




Figure 7.14: Instabilities of the microwave and optical comparisons. The Sr vs Maser comparison, performed in the microwave domain at 950 MHz, is limited by the performance of the Hydrogen maser out to 800 s before the passive microwave fiber transfer introduces significant excess instability at longer averaging times. The Sr vs Ca comparison has nearly a factor of 50 less instability for all averaging times. The optical comparison is recently-achieved preliminary data.

to other optical modes of the frequency comb. This assumption was reasonable to make, since comparisons of many other frequency combs have already shown that the optical accuracy across the comb is as low as 10^{-19} [109, 169]. Another important question is regarding the coherence of the frequency comb. When locked to an optical frequency reference and self-referenced, the frequency comb should coherently distribute the phase coherence of the optical reference to each of its modes. Previously, this phase coherence has been measured in two ways. By comparing the frequency comb passively to both a cw laser and its second harmonic, the intrinsic coherence of the comb was demonstrated to lead to a relative linewidth between comb modes (separated by an octave of spectral bandwidth) of 9 mHz [169]. By locking two frequency combs to the same optical reference and directly heterodyning the combs in various narrow spectral bandwidths, the ability of the frequency comb to lock to an optical reference was additionally tested in Ref. [170]. The result was a relative linewidth between the two combs of as little as 23 mHz, but a phase-noise measurement indicated an rms phase reached 1 rad at a 0.2-Hz Fourier frequency, corresponding to a 5-s coherence time.

Recently we constructed an experiment similar to the one described in Ref. [169], but with sensitivity to the comb's ability to lock to an optical reference in addition to measuring the intrinsic coherence of the comb. The experimental setup is shown in Fig. 7.15. The octave-spanning comb is phase-locked to a cw laser at 1064 nm using the exact same servo electronics (but a different photodetector) as are used to lock the comb to the 698-nm Sr clock laser as described in Section 7.3. Also, the comb's offset frequency is locked in exactly the same way as it is for the Sr clock experiments. In order to test the ability of the comb both to lock to the 1064-nm reference, and to transfer the reference's phase coherently across its spectrum, the cw reference is also frequency-doubled and compared against an appropriate mode of the comb near 532 nm. The out-of-loop beat at 532 nm therefore contains any wavelength-dependent intrinsic fluctuations of the frequency comb structure as well as both the residual locking errors

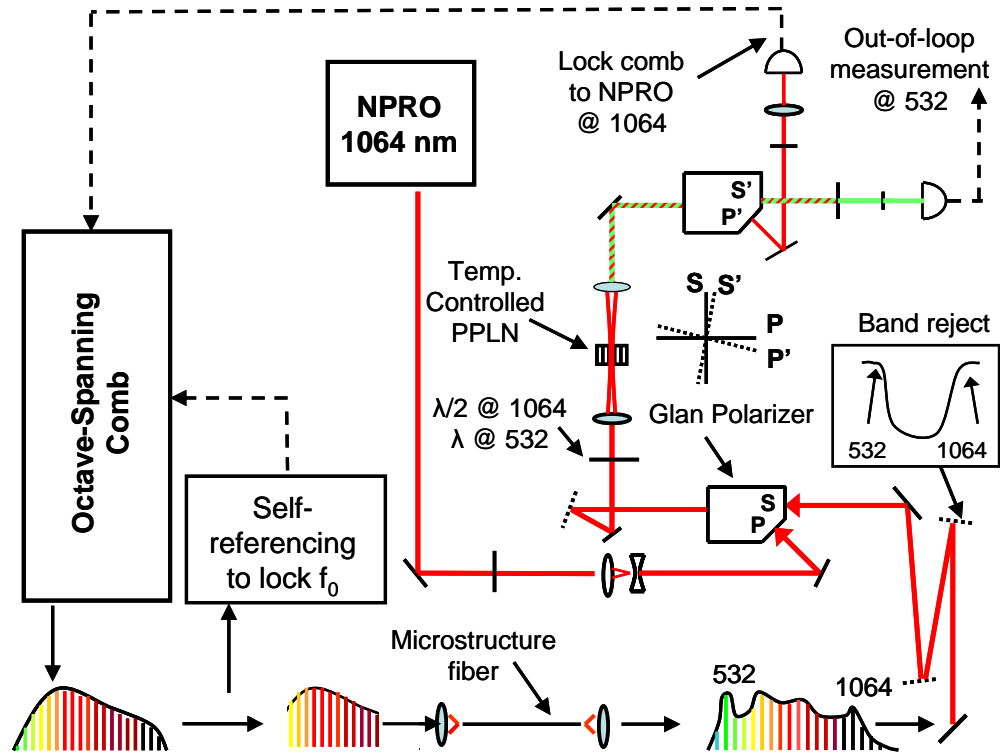


Figure 7.15: Experimental scheme for measuring the octave-spanning frequency comb's intrinsic coherence and accuracy. The frequency comb is first locked to a 1064-nm cw laser (NPRO) and simultaneously self-referenced to lock f_0 to a stable rf synthesizer. Then the second harmonic of the cw laser is generated in a PPLN crystal, and heterodyned against the appropriate mode of the comb near 532 nm in order to measure the comb's ability to phase-coherently lock to the 1064-nm optical reference and transfer the phase coherence across its spectrum. The accuracy of the out-of-loop heterodyne frequency from the 532-nm beat is also tested. Note that since the octave-spanning comb does not directly generate power at 532 nm, a microstructure fiber must be used to broaden its spectrum. Some band-reject filters are used to dump the power in the center of the spectrum before the frequency comb enters the PPLN, as a precaution against high peak intensities (and therefore undesired nonlinear consequences) in the PPLN. Also, details are given in the text of a few polarization tricks employing Glan polarizers in order to maximize the amount of power available for the two heterodyne signals while ensuring common-mode paths are maximized for the two wavelengths.

of the 1064-nm beat and the f_0 lock.

Unfortunately, the octave-spanning comb does not directly output power at 532 nm, so we had to broaden the spectrum in a short ~ 10 -cm piece of microstructure fiber. In order to ensure the optical paths in the measurement were kept as common-mode as possible, the output from the microstructure fiber was combined with the cw 1064-nm reference before the combined beam was sent through a PPLN (periodically-poled Lithium-Niobate) crystal for frequency-doubling of the 1064-nm cw light. With 250 mW of cw power from the NPRO (non-planar ring oscillator) at 1064 nm, nearly a mW of 532-nm cw light was generated. The NPRO has a fairly narrow linewidth of 1 kHz in 1 ms, which makes for a clean source to lock the comb to, ensuring any residual locking error we measured would not be due to attempting to lock to a noisy reference.

A clever polarization scheme was employed in order not to waste half the power in any of the beams. After combining the cw 1064-nm beam with the comb on a Glan polarizer, a dual-wavelength waveplate was used to rotate the cw and fs 1064-nm beams' polarizations relative to the 532-nm fs beam's polarization. This ensured that the SFG-generated 532-nm cw light would be polarized parallel to the comb's 532-nm light, even though the 1064-nm cw and fs beams were orthogonally polarized. However, since nearly 250 mW of cw 1064-nm light was present, we could afford to dump most of it in a second Glan polarizer used to separate the 532-nm and 1064-nm wavelengths. The second Glan polarizer was aligned at a slightly different angle relative to the first, such that most of the 1064-nm cw power exited the same port (S' in Fig. 7.15) as the 532-nm light, but was dumped using an interference filter. About 1 mW of 1064-nm cw power exited the P' port of the second Glan polarizer along with nearly all the 1064-nm power available from the fs comb. An interference filter centered at 1064-nm was used to dump the very small amount of 532-nm light from both the comb and the doubled-cw laser that also exited the P' port. This polarization scheme satisfied the dual requirement of keeping the paths of all the beams as common-mode as possible while using the available power

from the two comb modes as efficiently as possible.

In order to gain insight into what the 532-nm out-of-loop beat is really measuring, I will now write out the comb math behind the measurement. The modes of a frequency comb can be described in the most general possible form as:

$$\nu_n(t) = n f_{rep} + f_0 + \delta f_0(t) + \epsilon_n(t). \quad (7.4)$$

In this expression, n labels the particular mode of the frequency comb and $\delta f_0(t)$ represents any residual error of the offset frequency when it is locked to a stable rf reference. The term $\epsilon_n(t)$ accounts for the possibility that the comb may not exactly be described by only two degrees of freedom; perhaps the comb has optical-frequency-dependent noise or frequency shifts associated with its structure. Indeed, the measurement we are making is to ensure that $\epsilon_n(t)$ is as small as possible, such that the comb's intrinsic phase noise it describes can be shown to be negligible for frequency comb experiments requiring stringent phase-coherence across the comb's spectrum. If the self-referenced comb described by Eqn. 7.4 is then heterodyned against the cw optical reference near 1064 nm (and locked to it), the beat frequency $f_{b,1064}$ will be given by:

$$f_{b,1064} + \delta f_{b,1064}(t) = N_{1064} f_{rep} + f_0 + \delta f_0(t) + \epsilon_{1064}(t) - \nu_{1064}(t). \quad (7.5)$$

Here, $\nu_{1064}(t)$ is the time-dependent frequency of the 1064-nm cw optical reference, N_{1064} is the comb tooth to which the comb is locked, and $\delta f_{b,1064}(t)$ describes the residual locking error of the 1064-nm beat when locked to a stable rf reference. While locked to the 1064-nm reference, if some mode of the comb (say the mode $n = N_{532}$) is also heterodyned against the second harmonic of the 1064-nm reference at 532 nm, the beat frequency at 532 nm will be given by:

$$f_{b,532}(t) = N_{532} f_{rep} + f_0 + \delta f_0(t) + \epsilon_{532}(t) - 2\nu_{1064}(t). \quad (7.6)$$

Solving Eqn. 7.5 for f_{rep} and inserting it into Eqn. 7.6 results in the expression:

$$f_{b,532}(t) = R f_{b,1064} + (1 - R) f_0 + (R - 2) [\nu_{1064} + \delta \nu_{1064}(t)]$$

$$+ R\delta f_{b,1064}(t) + (1 - R)\delta f_0(t) + \epsilon_{532}(t) - R\epsilon_{1064}(t). \quad (7.7)$$

In this expression, $R \equiv N_{532}/N_{1064}$. If $R \neq 2$, then $f_{b,532}(t)$ will have a dependence on the frequency fluctuations $\delta\nu_{1064}(t)$ of the cw optical reference. For instance, if N_{532} corresponded to the next mode over from the mode satisfying $R = 2$, then for our 100-MHz repetition frequency comb, $(R - 2) \sim 10^{-6}$, which would result in a 1-mHz frequency fluctuation of $f_{b,532}$ for every 1-kHz fluctuation from $\delta\nu_{1064}(t)$.

If the condition $R = 2$ is properly enforced in the measurement scheme (the correct beat frequency is measured) then Eqn. 7.7 reduces to:

$$f_{b,532}(t) = 2f_{b,1064} - f_0 + 2\delta f_{b,1064}(t) - \delta f_0(t) + \epsilon_{532}(t) - 2\epsilon_{1064}(t). \quad (7.8)$$

Then if $f_{b,1064}$ and f_0 are locked to the rf frequencies 25 MHz and 70 MHz, respectively, the correct out-of-loop beat frequency to measure for $f_{b,532}$ will be -20 MHz. It will fluctuate by twice the locking error of the 1064-nm in-loop lock, plus the locking error of the f_0 lock, plus the two ϵ terms describing the intrinsic comb noise we hope to measure an upper limit to.

In fact, we used these values to lock the frequency comb to the 1064-nm reference, and found the 532-nm beat signal at 20 MHz as calculated. By mixing the 20-MHz out-of-loop beat frequency to a low frequency against another stable rf reference, we generated a 10-kHz signal that could be frequency-counted. An Allan Deviation calculated from the frequency counting data, with a 1-s gate time, is shown in Fig. 7.16. Since the instability includes terms related to the locking error and the intrinsic instability across the comb's spectrum, it represents an upper limit to both. Recently we have also counted the frequency for long enough to measure the accuracy of the frequency comb. Fluctuations of the 20-MHz out-of-loop signal ($f_{b,532}$) corresponding to the frequency comb's instability at 584 THz averaged down to a fractional optical frequency offset of $(1.2 \pm 1.8) \times 10^{-19}$.

A first attempt to measure the intrinsic coherence of the comb is to measure




Figure 7.16: Upper limit to the intrinsic instability of the frequency comb. The Allan deviation shown here is for the out-of-loop heterodyne beat at 532 nm, which carries information about the intrinsic instability of the frequency comb structure, in addition to information about the comb's lock to the 1064-nm optical reference, its f_0 lock, and any out-of-loop noise in the optical measurement paths. This is the first preliminary indication that our octave-spanning comb has instability at the few parts in 10^{-19} level for a few thousand seconds of averaging time.

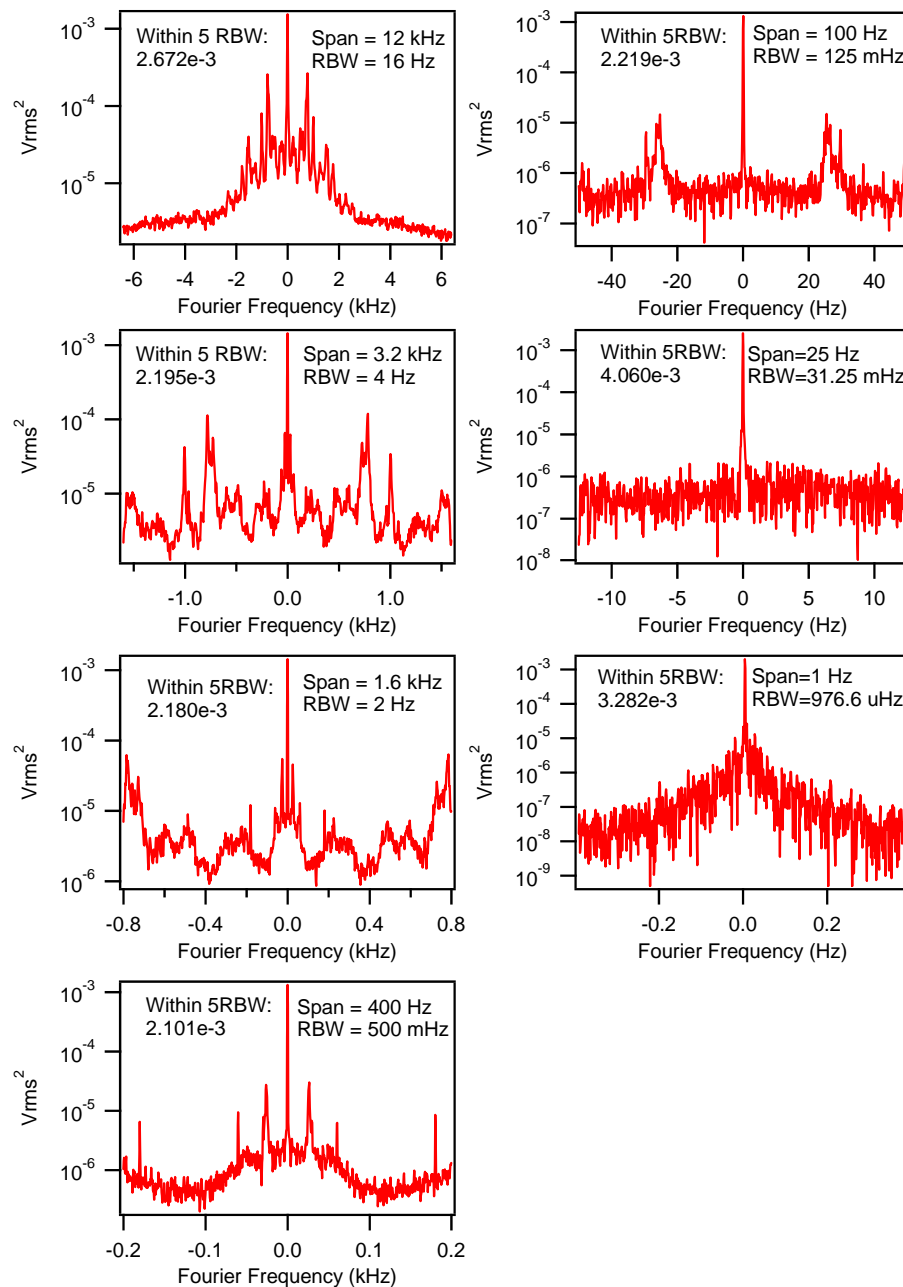


Figure 7.17: Upper limit to the intrinsic linewidth of the frequency comb. The linewidth is shown for various spans and resolution bandwidths. Note that between the 100-Hz span and the 25-Hz span, the frequency to which the 532-nm heterodyne beat was mixed was changed, and therefore the signal strength changed. We confirmed that the factor of 1.8 increase in the carrier's power was due to this change. These preliminary results demonstrate that the power in the central carrier remains unchanged as we measure using narrower and narrower bandwidths, suggesting that an upper limit to the comb's intrinsic linewidth (including locking error) could be as low as 1 mHz.

the linewidth of $f_{b,532}$ using an FFT machine after mixing it down to, for example, 10 kHz against a stable rf source. Figure 7.17 shows several measurements of the lineshape as the span and resolution bandwidth are successively reduced. The power in the central carrier is unchanged as the resolution bandwidth decreases from 16 Hz to less than 1 mHz, indicating that the comb remains coherent at the 1-mHz level across its spectrum, relative to the 1064-nm optical reference to which it is tightly locked. A more stringent test of the comb's coherence is to mix $f_{b,532}$ down to DC in a homodyne phase measurement scheme. By directly measuring the phase fluctuations, the phase noise spectral density can be integrated to discover the coherence time at which the rms phase noise reaches 1 rad. Since significant phase noise exists in the sidebands near 1 kHz, this measurement has so far proven difficult to make. However, recent preliminary measurements indicate that we are able to keep the integrated phase noise below 1 radian when integrating all the way from tens of kHz down to 1 mHz. When these results are confirmed, they will have proven that the coherence time of the comb relative to an optical reference frequency to which it is locked can be longer than 1000 s.

The experiments described in Section 7.4, taken together, demonstrate the frequency comb's importance for the field of optical frequency metrology. In Section 7.4.1 I described how our octave-spanning comb allowed the Sr clock's local oscillator to be characterized against another ultrastable reference, even though they were spectrally separated by 148 THz, at the 0.5-Hz-linewidth level. In Section 7.4.2 I showed how our comb routinely allows robust operation of our Sr optical clock, with a microwave readout, for hours at a time without interruption. This capability ultimately led to an absolute frequency measurement at the uncertainty level of 2.5×10^{-15} , limited by the Hydrogen maser's calibration to the Cs-fountain primary standard. Then in Section 7.4.3 I described how our octave-spanning comb represents the distribution hub of the ^{87}Sr optical lattice frequency standard, with the ability to compare directly in the

optical domain with other optical frequency standards located several km away and operating at optical frequencies removed from the Sr transition by hundreds of THz. The direct optical frequency comparisons could be made with a 1-s instability even lower than the 10^4 -s instability offered by comparison against the maser in the microwave domain. Finally, in Section 7.4.4 I described a preliminary measurement we made to measure the phase-coherence of the comb relative to an optical frequency reference to which it is locked. The comb offers a 1-mHz relative linewidth across an octave of optical bandwidth, as well as an accuracy better than 2×10^{-19} . Therefore the comb is fully capable of transferring the coherence, stability, and accuracy of today's best optical frequency references across its full optical spectrum.

The octave-spanning comb has truly fulfilled its purpose to provide a hub for the laboratory against which optical frequencies across the visible spectrum can accurately be measured, their phase noises characterized, and their linewidths revealed. The absolute performance of the comb is at the level of the ultrastable 698-nm clock laser, which is one of the most stable, pure, and when locked to the ^{87}Sr atoms, accurate oscillators available at any frequency!

Chapter 8

Frequency and Timing Transfer in Optical Fiber Networks

In the last two Chapters I showed two ways in which a frequency comb can be used to form a microwave readout for an optical clock. The offset-free clock of Chapter 6 had the advantage of simplicity, since f_0 did not need to be controlled. However, the octave-spanning comb in Chapter 7 proved to be the workhorse of a potential future primary frequency standard by allowing phase coherence to be established not only from the clock transition to the microwave readout, but also across the visible portions of the spectrum. This enables such frequency combs to become the “hubs” of metrology laboratories, allowing multiple optical clocks operating at various frequencies to be directly compared.

In this Chapter I will continue to show how frequency combs can be used for the purposes of both frequency- and timing-transfer in optical fiber networks. First I will give an extensive motivation of the need for precision transfer through fiber networks. Then I will discuss two technologies for such transfer (microwave transfer and cw optical carrier transfer), each motivated by the needs of our Sr optical lattice clock to compare performance with other clocks in Boulder. Finally, I give a brief discussion about frequency transfer of the fs comb itself for the purpose of remote synchronization of fs lasers.

Three distinct techniques exist for distributing an ultrastable frequency reference over optical fibers. For the distribution of a microwave frequency reference, an

amplitude-modulated cw laser can be used. Over km-scale lengths this approach provides an instability at 1 s of $\sim 3 \times 10^{-14}$ without stabilization of the fiber-induced noise and $\sim 1 \times 10^{-14}$ with active noise cancellation. An optical frequency reference can be transferred by directly transmitting a stabilized cw laser over fiber and then disseminated to other optical and microwave regions using an optical frequency comb. This provides an instability at 1 s of 2×10^{-14} without active noise cancellation and 6×10^{-18} with active noise cancellation. Finally, an optical frequency comb has been transmitted through a km-scale fiber network, allowing a remotely-located fs comb to be synchronized relative to the master comb to within a residual rms timing jitter below 16 fs integrated over the Nyquist bandwidth of the pulse train (~ 50 MHz) when high-bandwidth active noise cancellation is employed, which is important for remote synchronization applications.

8.1 Motivation for ultrastable frequency reference transfer via fiber networks

Current research on optical atomic clocks [117, 118] offers the potential to produce frequency references with orders of magnitude lower instability and inaccuracy than is provided by the best existing microwave frequency references [171, 172, 173, 174, 175, 176, 149, 177, 151, 152, 146, 160, 162]. For accuracy, the best microwave frequency reference is provided by a cesium fountain clock, which is able to measure the frequency of the hyperfine splitting of the cesium electronic ground state with an accuracy better than six parts in 10^{16} [3, 178]. However, frequency references based on optical transitions in laser cooled and trapped atoms and ions are expected to eventually provide uncertainties approaching a part in 10^{18} [171], and short-term instabilities of a few parts in 10^{17} for a 1-s averaging time [177], nearly three orders of magnitude better than the best microwave atomic clocks. The current performance of various high-stability oscillators is summarized in Fig. 8.1. It shows the Allan deviation of some of the most stable




Figure 8.1: Fractional frequency instabilities of various high-stability microwave and optical frequency references. References are given in the text.

existing frequency references. In the microwave domain, a cryogenic sapphire oscillator [144] (closed squares) is shown, as well as a hydrogen maser [179] (closed circles) that is steered by a cesium-fountain clock. In the optical domain, a narrow 0.3-Hz-linewidth laser [45, 46] (open circles) that serves as the local oscillator for an optical clock based on ultracold Sr atoms is compared to the projected performance for the Sr optical clock [87] (open squares). In addition, it has been shown that the frequency comb from a mode-locked femtosecond laser can be slaved to an optical reference at the level shown as open diamonds [109]. Finally, the rf signal provided by detecting the 10-GHz harmonic of a femtosecond laser's repetition rate has an excess instability [78] (relative to the stability of the comb's optical frequencies) shown as closed diamonds. From these results it can be seen that future needs for the stability of fiber-transferred frequency references will approach the 10^{-16} level at 1 s and 10^{-18} at longer time scales, in order to preserve the quality of these references after transfer.

Further development of optical frequency references will require unprecedented levels of stability for characterization. The initial research-oriented systems are complex and not necessarily portable, so the ability to remotely transfer frequency references without introducing any additional instability is of urgent concern for optical clock development. In fact, the current state-of-the-art microwave frequency standards already demand improved-stability transfer protocols for signal comparison and clock synchronization. In addition, as listed in the following paragraphs and as shown in Fig. 8.2, a number of exciting applications taking advantage of the superior stability of optical references will benefit a great deal from the capability of high-stability frequency transfer. They include the test of fundamental physical principles, development of next generation accelerator-based X-ray sources, long baseline coherent radio telescope arrays, and the accurate mapping of the Earth's geoid, to name a few important examples.

At one of the frontiers of modern physics, an exciting application requiring ultrastable transfer of frequency references is the comparison of frequency references

based on transitions in different atomic species. This would enable measurements of the time variation of fundamental constants [13], such as the fine structure constant α [180, 15, 16, 181, 14, 182]. If α were changing over time, the frequencies of these transitions based on different atomic systems would change with respect to each other. 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 reliably transfer a frequency reference is crucial for these comparisons. The most accurate measurements to date have placed a constraint of one 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. These comparisons are accomplished by averaging the frequency measurements over an extended period of time. Thus, the spectral distribution of frequency fluctuations may not be explicitly displayed, since phase information at short timescales has been averaged over. However, for applications that involve very tight timing synchronization among system components, the distributed frequency reference to which all the components will be synchronized must have very high stability and low phase noise for short timescales. This is equivalent to saying that the transfer must exhibit ultralow timing jitter (fs-level or better) over an extended Fourier frequency range [183, 21]. One such application is long-baseline interferometry for radio astronomy [184]. 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 over a distance of ~ 20 km. This would enable all telescopes in the array to phase-coherently collect data, thereby simulating a single telescope with a very large aperture [184]. Low-jitter transfer could

also be used to distribute a frequency reference throughout a linear accelerator facility over distances of 2 – 5 km for synchronization of its various components. Recently there has been considerable interest in producing ultrashort x-ray pulses generated in accelerator-based facilities to study ultrafast phenomena in several fields including chemistry, physics, biology, and materials science [185, 41, 42, 43]. These studies will involve pump-probe experiments using visible lasers to pump the samples and the x-ray pulses to probe them, or vice-versa [41, 43]. 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.

Here I will discuss three distinct methods for transferring frequency and timing references long distances over optical fibers for applications including those shown in Fig. 8.2. These methods represent the current state-of-the-art for high-stability frequency transfer. In Section 8.2 I will present the basic principles for the design of a fiber transfer network, including discussion of noise processes and relevant design considerations for active stabilization of the link. The subsequent three Sections will then delve into specific experimental techniques for the fiber transfer: in Section 8.3 I will discuss the transfer of a microwave frequency reference using an amplitude modulated cw laser. In practice this provides the most straightforward method of transferring a microwave frequency reference or timing signal, though it does not provide a straightforward means of transferring a more stable optical frequency reference. We have used this technique in combination with the octave-spanning comb to measure the absolute transition frequency of the ^{87}Sr optical lattice clock as discussed in Chapter 7. In Section 8.4 the direct transfer of an optical frequency reference with a cw laser will be presented. This is the most appropriate approach for remote comparisons of optical frequency standards. As such, we have employed it in combination with the octave-




Figure 8.2: By transmitting pulses from a femtosecond fiber laser through a km-scale fiber network, a frequency/timing reference is distributed from a master optical atomic clock, using an optical frequency comb, to facilities for several precision applications, including: other atomic clocks for inter-clock frequency comparisons, also enabling mapping of the Earth's geoid; a phase-coherent radio telescope array; and components of a synchrotron advanced light source for generating x-ray pulses.

spanning comb in order to compare the Sr clock against other optical clocks at NIST. However, for other applications such as remote synchronization, the cw transfer method has the drawback of not directly providing a timing signal, since the modulation signal is the carrier itself, which is at hundreds of terahertz. Therefore an optical frequency comb must be employed to recover timing information, once again demonstrating the need for optical frequency combs when it comes to timing transfer. Finally, in Section 8.5 I will briefly discuss a remote synchronization experiment where an optical frequency comb is directly transmitted through the fiber network [186, 187, 135]. More details of the transfer of the frequency comb for stability purposes (at timescales > 1 s) rather than synchronization purposes (Fourier frequencies 1 Hz to 100 MHz) are given in Kevin Holman's Ph. D. Thesis [188].

8.2 Principles of cancellation of fiber-induced phase noise

The goal of a fiber transfer network is to faithfully (phase-coherently) reproduce a local frequency or timing reference signal at the remote output of the fiber network. Unfortunately, various noise processes in the optical fiber degrade the quality of the transmitted signal by perturbing its phase. This excess phase noise must be detected and subsequently cancelled in order to properly distribute the frequency reference signal. Ultimately, the technology should employ a sufficiently robust servo with a large enough dynamic range to cancel the phase noise present in a typical urban fiber network installation, since in many cases the existing infrastructure provides the most convenient opportunity for frequency dissemination, even though it was not necessarily designed with high-stability transfer in mind. The need for large dynamic range is particularly true for fiber links approaching 100 km in length, since the integrated noise increases with length.

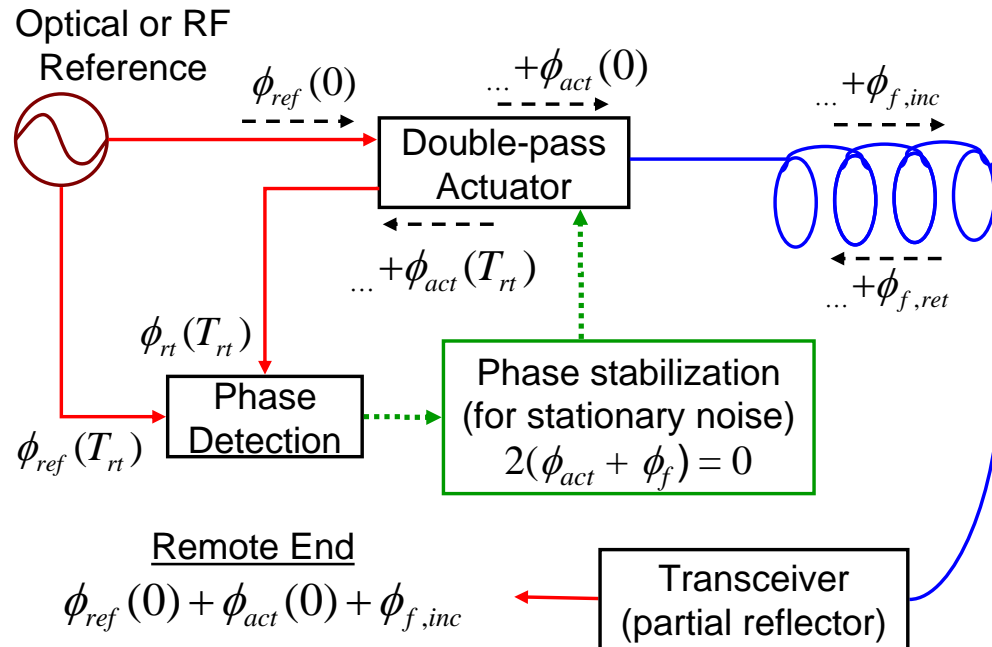


Figure 8.3: Principle of operation for a typical actively compensated fiber network for distributing frequency and timing references. A partial retroreflector or transceiver at the remote end returns the signal through a round trip of both the fiber network and the actuator. Accurate comparison between the phase of the round-trip signal and the original reference allows the double-pass actuator to compensate for the fiber link's phase perturbations for timescales longer than T_{rt} . The actuator can act on the group delay for modulated-cw transfer or frequency comb (pulse-timing) transfer, or it can act on the optical phase for direct optical carrier transfer.

8.2.1 Round-trip phase detection

Figure 8.3 shows the basic experimental scheme for fiber network phase-stabilization. The idea essentially comes from the work in Ref. [4], where a hydrogen maser on board a rocket was compared to a ground-based maser as a test of relativity. The doppler shift from the high-velocity spacecraft was cancelled using the round-trip phase detection techniques described in this Section. After the work of Vessot, Jim Bergquist employed a similar scheme in free space, but in the optical frequency domain, to eliminate the doppler phase noise between two ultrastable cavity-stabilized lasers [189]. Then Jan Hall applied a similar scheme to cancellation of the phase noise from a short-distance optical fiber link [38]. The scheme works as follows: a reference signal which could be a modulated cw optical carrier, the optical carrier itself (in which case the modulation frequency is also the carrier frequency), or a frequency comb, is transmitted to the remote user. However, the phase perturbation from the fiber link, ϕ_f , must be precompensated by some actuator. This can be done by reflecting a portion of the transmitted light from the remote end, detecting the phase ϕ_{rt} of the round-trip signal, and comparing it to the reference phase ϕ_{ref} . Sometimes, instead of a simple retroreflection a transceiver is used at the remote end in order to boost the power on the return trip. Also, the transceiver can slightly shift the carrier frequency or return a harmonic of the modulation frequency in order to distinguish the return signal from parasitic reflections of the incident signal at interconnects along the fiber link [36]. The precompensation is performed by an actuator acting once on both the outgoing and incoming signals, introducing a phase ϕ_{act} on each pass. Typical actuators are piezo-actuated fiber-stretchers or thermally controlled spools of fiber for acting on the group delay, or acousto-optic-modulators (AOM) for indirect control of the optical phase delay. Since an AOM acts on the optical frequency, it provides unlimited dynamic range for optical phase compensation.

We can detect and subsequently compare ϕ_{rt} , the signal's phase after making a

round trip through both the link and the actuator, to ϕ_{ref} . Importantly, this phase comparison occurs after a round-trip time T_{rt} twice that of the single-pass delay. If a particular reference signal is launched into the fiber link (including the actuator) with phase ϕ_{ref} at time $t = 0$, then at time T_{rt} the measured phase difference between the reference and the round-trip signal will be

$$\begin{aligned} \phi_{rt}(T_{rt}) - \phi_{ref}(T_{rt}) &= \phi_{ref}(0) - \phi_{ref}(T_{rt}) \\ &+ \phi_{act}(0) + \phi_{act}(T_{rt}) \\ &+ \phi_{f,inc} + \phi_{f,ret} \end{aligned} \quad (8.1)$$

Here, $\phi_{f,inc}$ and $\phi_{f,ret}$ are the phases introduced by the incident and return trips through the fiber link, respectively.

Important assumptions can be made to simplify Eqn. (8.1). First, if the noise processes responsible for fluctuations of $\phi_{f,inc}$ and $\phi_{f,ret}$ are assumed stationary between the incident and return transmissions through the link, and the propagation through the fiber samples the same polarization modes on the return trip as on the incident trip, then the last two terms of Eqn. 8.1 are equal, and can be written as $2\phi_f$. Under this assumption, only noise processes with frequencies less than $\frac{1}{2\pi T_{rt}}$ are considered. Second, the actuator bandwidth must also be limited to this value, so that the third and fourth terms on the right side of Eqn. (8.1) will be equal. Finally, the first two terms will cancel up to an overall constant if the coherence time of the frequency reference (microwave or optical) is longer than T_{rt} . Under these three assumptions (equal noise on incident and return trips, servo timescale longer than T_{rt} , and reference's coherence time longer than T_{rt}), Eqn. (8.1) can be simplified as

$$\phi_{rt} - \phi_{ref} = 2(\phi_{act} + \phi_f) \quad (8.2)$$

and the phase at the remote end of the link is simply $\phi_{ref} + \phi_{act} + \phi_f$. Therefore, if the local detection feeds back to the actuator to stabilize the left side of Eqn. (8.2)

to a constant, the remote user will have a phase-coherent copy of the reference, up to the round-trip-limited bandwidth of $\frac{1}{2\pi T_{rt}}$. We emphasize that T_{rt} , and therefore the fiber link's length, sets the limit on the bandwidth that can be stabilized as well as the minimum quality of the reference that can be transmitted (due to the coherence time assumption.) In order to transmit over very long distances with high noise-cancellation bandwidths, repeater stations will become necessary.

8.2.2 Noise processes, bandwidth and dynamic range

The main source of excess phase noise in the fiber comes from mechanical and thermal perturbations to the optical path length of the transmission fiber. Other sources of noise may come from internal reflections at fiber interconnects along the transmission fiber, stimulated Brillouin scattering (SBS), and polarization mode dispersion (PMD) fluctuations, all arising from mechanical perturbation of the fiber.

When designing a system for frequency transfer over fiber links, the designer has control over any of the following: group delay of the link, optical phase, dynamic range of the actuator, the optical power employed, group-velocity-dispersion (GVD) of the link, the launch polarization, and the servo bandwidth. Choosing between actuating on the group delay or the optical phase depends on whether the frequency reference being transmitted is microwave or optical. Knowledge of the noise properties of the particular fiber link of interest will help to determine what actuator dynamic range is necessary. Optical power should generally be as high as possible in order to maintain excellent SNR, but SBS sets limits on the practical optical power level that can be transmitted via typical single-mode fibers. GVD is an important consideration only for frequency comb transfer, and can be adjusted with dispersion-compensation fiber or fine-tuned with a pulse-shaper [190, 191].

Polarization purity is an important concern for fiber links, due to PMD fluctuations. Unfortunately, most preexisting fiber networks utilize non-polarization-maintaining

(PM) fiber due to the expense and optical loss associated with PM fiber. The birefringence of typical non-PM fiber can lead to problems due to PMD fluctuations for microwave systems, or due to slow polarization rotation leading to AM fluctuations in the heterodyne beat used for direct optical transfer. This is because the polarization modes are sampled differently by the return beam than by the incident beam, so the assumption that $\phi_{f,inc} = \phi_{f,ret}$ is violated even for noise processes that are stationary during T_{rt} . As will be discussed in Section 8.3.3, scrambling the polarization at timescales faster than T_{rt} offers a method for quickly averaging out PMD effects for microwave transfer [36, 192]. However, the phase coherence of an optical carrier would be washed out by first-order PMD if the polarization were scrambled. The effects of second-order (wavelength-dependent) PMD [193] could be especially complex when frequency comb transfer is considered, although averaging over millions of comb lines could potentially reduce the effect.

Finally, bandwidth limitations due to the interplay between SNR, dynamic range, and integrated noise of the fiber link must be considered. As discussed in Section 8.2.1, the bandwidth of the phase locking servo is limited to $\frac{1}{2\pi T_{rt}}$. Any attempt to cancel noise at higher bandwidths will simply write excess noise onto the link, since phase fluctuations on faster timescales than T_{rt} violate the assumption that the round-trip phase accumulation is twice that of the one-way phase. Another bandwidth limitation depends on the relative magnitudes of the servo dynamic range and the integrated noise of the fiber link. A servo with sufficient dynamic range and high enough SNR will be able to cancel noise within the bandwidth set by causality. However, a servo without sufficient dynamic range will require a lower cancellation bandwidth. Since the overall noise increases as a function of the fiber length, sometimes worse than linearly, servos with low dynamic range (such as the PZT fiber stretchers described in Sections 8.3 and 8.5) can suffer for longer fiber links. An additional concern for longer links is attenuation of the signal. The available signal will decrease with length while the noise

increases, and the amount of signal that can be input to the link is limited by SBS. From Eqn. (2.13), use of an optical rather than microwave carrier frequency ν_0 will greatly relax the requirements for S_ϕ , and therefore for the SNR as well. These bandwidth considerations can also depend on the specific target application of the transferred frequency signal. For instance, consider a clean-up oscillator at the remote end with a rapid phase noise roll-off beyond the servo bandwidth: by locking the clean-up oscillator to the transferred reference with the same bandwidth as the fiber stabilization servo, the excess noise of the fiber link at frequencies higher than the servo bandwidth can be avoided in the clean-up oscillator's output.

8.2.3 Experimental schemes for frequency transfer

These complex bandwidth limitations must be considered when choosing the type of frequency reference to transmit over the fiber link (cw microwave, cw optical, or optical frequency comb). Indeed, the manner in which the excess phase noise is cancelled depends on which kind of reference is being transmitted. In some cases, the excess phase noise ϕ_f of the fiber network is tolerable by the end-user's requirements. Such a passive transfer scheme is often used when transmitting microwave frequencies over the fiber network. The microwave reference is used to amplitude modulate a cw optical beam, which is then launched into the fiber for transfer. The remote user simply demodulates the sidebands on the optical carrier by photodetection of the optical beam in order to recover the microwave information. However, the excess phase noise introduced in the fiber typically limits the overall system performance to 10^{-15} at long time scales. For more demanding applications involving comparison of optical frequency standards or remote synchronization work, active cancellation of the fiber-induced phase noise is necessary, as shown in Fig. 8.4(a). Since the phase of the microwave information must be stabilized, a PZT fiber stretcher (small dynamic range) or temperature-controlled fiber spool (larger dynamic range but lower bandwidth) acts directly on the group delay

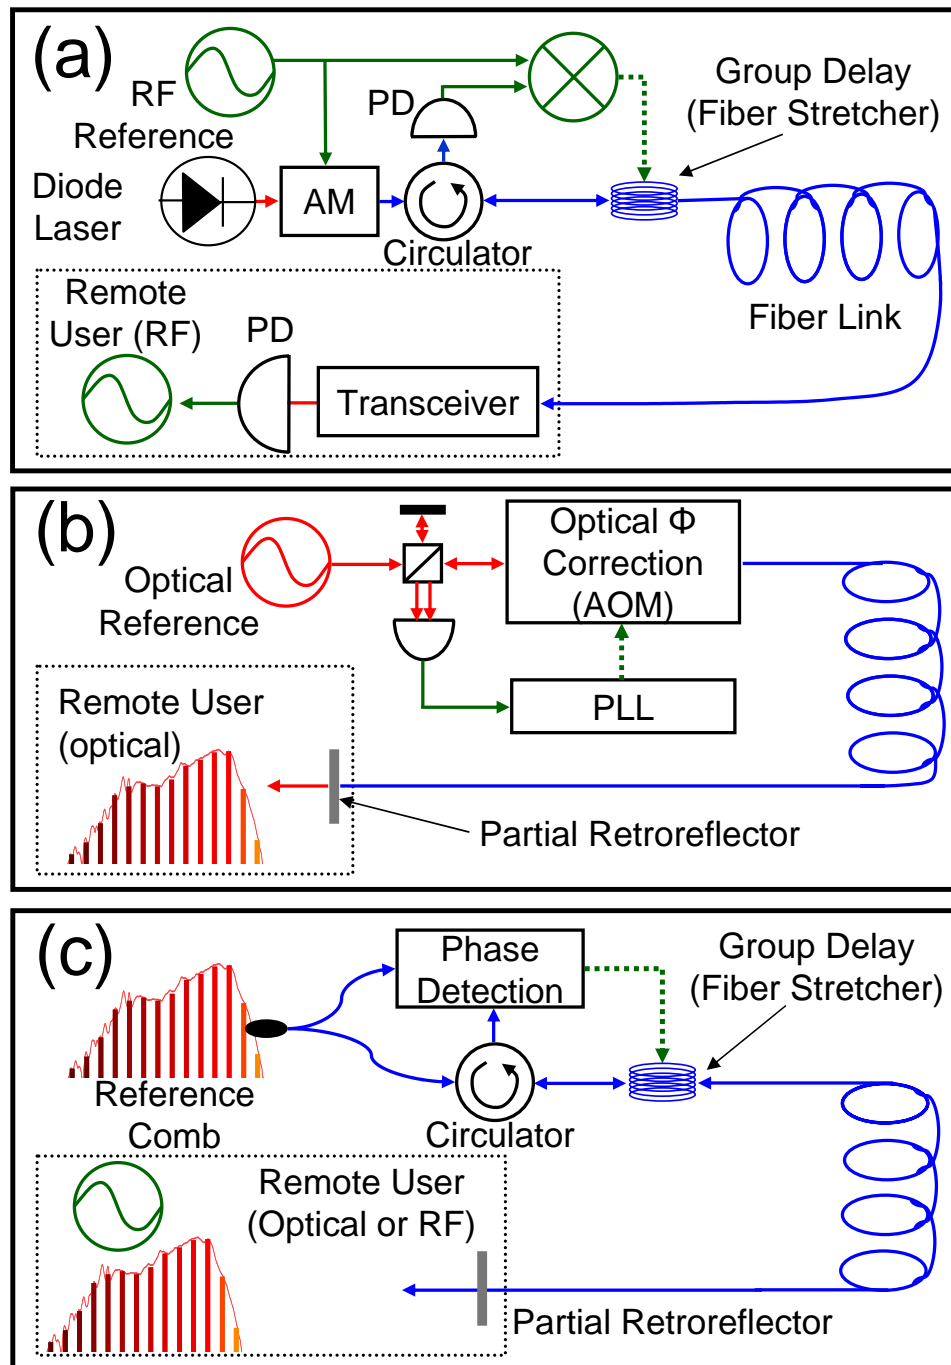


Figure 8.4: Various experimental schemes for transfer of frequency references. (a) Actively-stabilized transfer of a microwave reference by modulating a diode laser. (b) Direct transfer of the optical carrier by stabilizing the optical phase of the fiber link. (c) Transfer of simultaneous microwave and optical frequency references by transferring an optical frequency comb. Either optical or microwave information from the round trip can be used to stabilize the group delay and/or optical phase delay of the fiber link.

of the fiber.

Applications involving comparison of optical frequency standards can benefit from directly transferring a stable optical reference, as shown in Fig. 8.4(b). In this case, a portion of the optical beam is used as a reference arm, while the rest of the beam is first frequency shifted by an AOM and then launched into the fiber network. A heterodyne beat signal between the round-trip optical beam and the local reference reveals the optical phase noise of the fiber link. The error signal is used to feed back to the AOM in order to precompensate for the phase of the link. Since the AOM acts directly on frequency, it has extremely large dynamic range with respect to phase. The transmitted optical signal can provide a microwave reference with use of an optical frequency comb.

When using a frequency comb to simultaneously transfer optical and microwave references, a similar scheme to the microwave transfer can be employed as shown in Fig. 8.4(c). In this case usually the group delay of the fiber is actively stabilized in order to stabilize the repetition rate of pulses exiting the fiber at the remote end, and can provide an ultralow-jitter signal for timing to the remote user. In Section 8.5 I will also discuss stabilization of the group delay and optical phase simultaneously by controlling both the transmission path length and the dispersion therein.

Another practical consideration involving T_{rt} , especially for test-bed experiments employing relatively noisy sources, is the coherence time of the signal being transmitted. Recall from Section 8.2.1 that the reference's coherence time must be longer than T_{rt} . For transfer of microwave frequencies, even a relatively poor microwave signal usually has a coherence time much longer than T_{rt} . However, if optical frequencies are being transferred, even a source with a relatively narrow linewidth ≈ 1 kHz will not remain phase coherent with a round-trip copy of itself after the copy has been delayed by a link longer than 100 km. The servo will treat this phase incoherence as being due to the fiber link itself, and add the phase incoherence of the source to the phase of the remotely transferred signal rather than reproducing a phase-coherent copy at the remote

end. Therefore, especially for cw optical links as discussed in Section 8.4, and for very long links in general, it is important to ensure that the source being transmitted has a coherence time longer than T_{rt} in order for the active cancellation to perform as desired.

The final consideration for frequency transfer involves the practical issue of achieving an out-of-loop measurement of the excess instability induced by the fiber-link on the frequency reference at the remote end of the system. Clearly, since the end user is remotely located, the remote copy of the reference cannot be directly compared to the original reference. Initial research experiments demonstrate the stability of the transfer over a long fiber link by simply co-locating the “remote” end with the local reference in order to perform a direct comparison, before the full system is implemented at remote locations. Another way, which preserves the ability to measure a system’s performance even while operating, is to build two transfer systems and operate them in opposite directions [36, 194]. The reference can be transmitted via a stabilized fiber link to the remote user, who then returns a copy via an independently-stabilized second fiber link. The second fiber link uses the signal it received from the first link as its source reference. This kind of handshake configuration allows the original reference to be compared to the round-trip signal to determine the overall system performance, which represents an upper limit to the instability of either one-way transfer.

Keeping in mind these general principles for the practical implementation of fiber network frequency and timing dissemination, we can now turn our attention to the specific experimental realizations of microwave transfer (Section 8.3), optical carrier transfer (Section 8.4), and pulse timing transfer (Section 8.5).

8.3 Microwave frequency transfer with modulated cw laser

8.3.1 Passive transfer technique

The most straightforward method of transferring a microwave frequency reference over optical fibers is to amplitude modulate a cw laser at the frequency of the reference and to transmit this modulated light over the fiber. Then the microwave frequency reference is recovered at the remote end by detection of the modulation frequency. A general schematic for this method of transfer is given in Fig. 8.4(a), which shows an optical source being amplitude-modulated at the reference frequency, then being transmitted over a long-distance fiber link, and finally being detected at the remote end with a photodetector for recovering the reference frequency by demodulation of the sidebands surrounding the optical carrier. At first we will discuss operation of such a link without feedback to stabilize the fiber-induced phase fluctuations.

For the case of a fiber link connecting our labs at JILA to the National Institute of Standards and Technology (NIST) Boulder [195], we use two dark fibers of the Boulder Research and Administration Network (BRAN) that are installed in underground conduits and steam tunnels in an urban area [196]. Transmission through the 3.45-km (one-way) link is accomplished by amplitude modulating the light from a high-power (30-mW) single-mode distributed-feedback semiconductor laser, lasing at $1.3\ \mu\text{m}$, and the modulation is performed external to the laser using a Mach-Zehnder intensity modulator [197]. In Ref. [37], in order to measure the fiber-induced frequency noise of the transmitted microwave signal [37] (typically operating at about 1 GHz), the local copy of the reference signal at NIST was first shifted by 10 kHz with a single-sideband generation microwave interferometer [198]. Then the round-trip signal was mixed with this shifted local copy, and the frequency fluctuations of the resulting ~ 10 -kHz beat signal, which is suitable for high-resolution frequency counting, were counted. The 10-kHz source used to generate the single sideband must have absolute frequency fluctuations

lower than those measured on the 10-kHz beat. When this condition was satisfied, measured fluctuations of the 10-kHz beat frequency were due to the instability of the transfer. The Allan deviation of these frequency measurements provides the fractional frequency instability of the transfer as a function of averaging time, and is shown in Fig. 8.5. Also shown are results obtained when the BRAN fiber was replaced with a 2-m-long reference fiber. The additional fractional frequency noise from the passive microwave transfer over the BRAN is approximately 2×10^{-13} at 1-s averaging time and $1 - 2 \times 10^{-14}$ between 100 and 1000 s.

Another method for characterizing the excess instability of the fiber link is to mix the round-trip signal with a local copy of the reference in a homodyne scheme, in order to directly measure the phase fluctuations on the 7-km round-trip signal. Figure 8.6(a) shows this phase record over the course of a 9-hour period with a 1-s gate time for the measurement. The steep long-term slope between 2.5 and 4 hours corresponds to a fractional frequency offset of $\sim 1 \times 10^{-14}$ for the 950-MHz frequency being transferred. By taking the time-derivative of the phase data and applying a rolling box average of 100-s width, the fractional frequency offset of the passively-transferred signal is determined as plotted in Fig. 8.6(b). Since both the outgoing and incoming fibers are contained in the same bundle, it is reasonable to assume that long-term noise processes in the fibers are correlated. Therefore the one-way transfer is expected to have half the phase fluctuations and associated frequency offsets. Even so, the average fractional frequency offset during certain times of the day can be as high as several parts in 10^{-14} , which is insufficient for the purposes of comparing the best frequency standards at remote locations.

8.3.2 Passive microwave transfer for the ^{87}Sr clock

An important illustration of the kind of application that can benefit from passive or active microwave transfer, or the higher-resolution optical transfer described in Sec-




Figure 8.5: Measured instability of amplitude-modulated (modulation frequency 1-GHz) light transmitted through a short (~ 2 m) fiber (squares) and through a round trip of the BRAN fiber (open triangles). The rf modulation signal carried by the fiber-transmitted light is heterodyne detected against the original signal source, which has been frequency shifted by 10 kHz via a single-sideband generation rf interferometer. Also shown for comparison is the stability of NIST's Cs-referenced hydrogen maser (solid line).

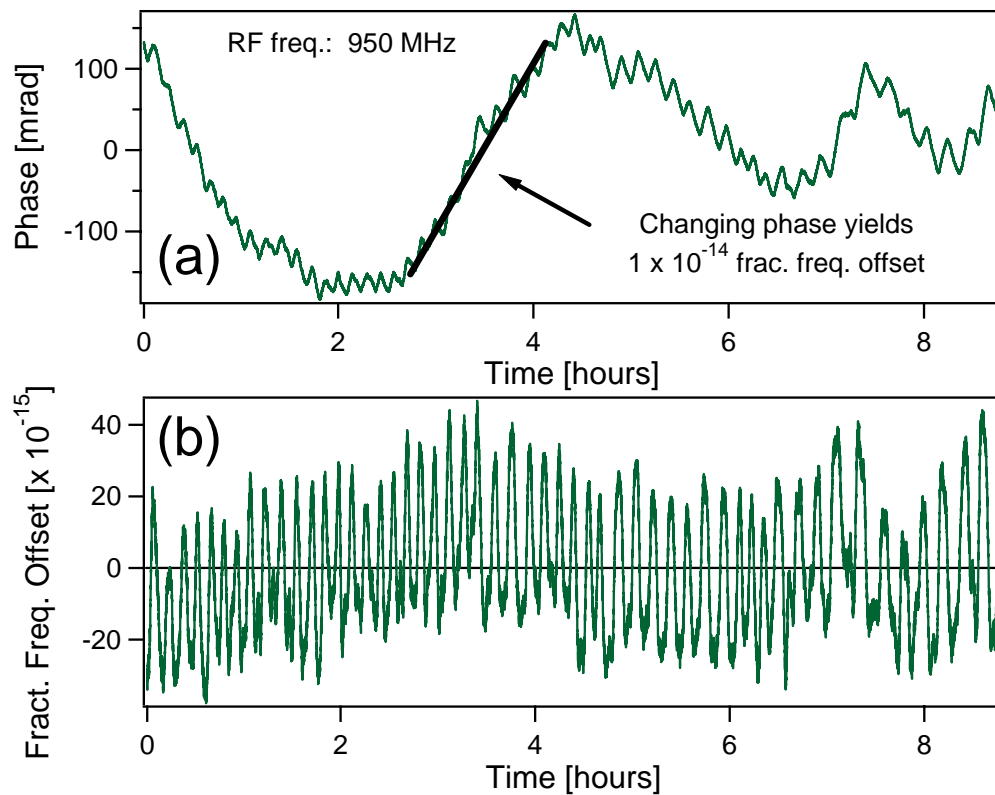


Figure 8.6: (a) Phase of the 7-km round-trip BRAN fiber for a transferred rf frequency of 950 MHz, showing long-term fractional frequency offsets of $\sim 10^{-14}$. (b) Time-derivative of the phase data, with a rolling box average of 100 s applied, showing the fractional frequency offset induced by the passive fiber transfer setup.

tion 8.4, is the case of absolute frequency measurements of a Sr-based optical clock. In Chapter 7 I focused the discussion on the octave-spanning comb used as the “gears” of the JILA ^{87}Sr optical lattice clock. In this Section I will discuss the microwave transfer we used to measure the absolute frequency relative to the NIST Cs fountain clock. In some cases the stability provided by the passive transfer of a modulated cw laser over optical fiber is sufficient, without needing to implement active noise cancellation. This is the case for the transfer of a Cs-referenced, hydrogen maser-based rf standard [3, 179] from NIST to JILA for the initial work on absolute optical frequency measurements of the Sr lattice clock.

Absolute frequency measurements were reported in 2005 for the $^{87}\text{Sr } ^1\text{S}_0 - ^3\text{P}_0$ clock transition [152, 146]. Ref. [152] used a GPS-based frequency reference with instability $\sim 10^{-10}$ at 1-s averaging time, while the JILA results [146] used a fiber-transferred reference. A troublesome 4σ disagreement between these two results was eventually resolved by a third independent investigation using an on-site Cs fountain reference [199]. A revised report by the Tokyo group [200] then appeared. This situation makes clear the benefits of fiber links for frequency distribution. The Sr clock experiments at JILA take advantage of the much more stable fiber link described in Section 8.3.1. This allows absolute frequency measurements relative to the NIST F1 Cs fountain clock via a hydrogen maser-based rf signal, which has yielded the smallest statistical uncertainty to date for this ultranarrow clock transition in a magic-wavelength optical lattice [146, 147].

For the first accurate measurement of the clock transition [146] in 2005, the octave-spanning Ti:sapphire laser-based frequency comb discussed in Chapter 7 was phase-locked to the optical local oscillator for the ^{87}Sr clock [46], as well as self-referenced using the standard f -to- $2f$ scheme [98] as discussed in Section 3.3.1. Then the ~ 1 -GHz harmonic of the Ti:sapphire laser’s repetition rate was frequency-counted against the rf reference received over the fiber network from NIST, in order to determine the absolute frequency of the $^{87}\text{Sr } ^1\text{S}_0 - ^3\text{P}_0$ optical clock transition. The Cs-calibrated,

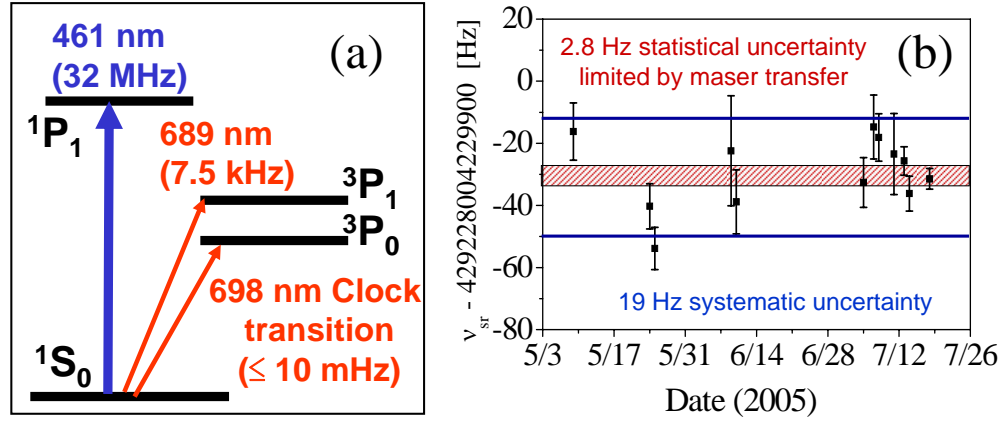


Figure 8.7: (a) Simplified energy level diagram for neutral Sr atoms, showing the two cooling transitions at 461 nm and 689 nm, and the clock transition at 698 nm. (b) Frequency-counting record for ^{87}Sr , showing the achieved 2.8-Hz statistical uncertainty (maser-transfer-limited) in the red-striped box and the 19-Hz systematic uncertainty as the outer blue lines.

maser-based rf standard we transfer from NIST has an instability as shown in Fig. 8.5 of $\sim 2.4 \times 10^{-13}$ at 1 s [179]. Passive transfer of this frequency from NIST to JILA is achieved by amplitude modulation as described in 8.3.1, and reaches an instability $< 1 \times 10^{-14}$ for averaging times $> 10^3$ s as shown in Fig. 8.5. With the ability to average down to such a low fractional uncertainty in this relatively short time (as compared to GPS frequency transfer), several systematic frequency shifts were investigated, including those originating from atomic density, wavelength and intensity of the optical lattice, residual magnetic field, and probing laser intensity. Figure 8.7 summarizes the measurement of the transition frequency within a period of three months. The statistical uncertainty of these measurements was 2.8 Hz, and systematic uncertainties were below the 20-Hz level, all made possible because of the reliable and robust passive transfer over the BRAN fiber network of the precision microwave reference, calibrated to the Cs-fountain clock.

Nevertheless, as seen in Fig. 8.6, the overall accuracy of the transferred reference at certain times of the day can be limited to $\sim 1 \times 10^{-14}$ because of the long-term phase

drift of the fiber transfer. Therefore, active cancellation of the microwave phase drift that achieves instabilities below 10^{-16} for times $> 10^4$ s, as described in the following Section, became necessary to improve the accuracy of our frequency measurement.

8.3.3 Actively-stabilized fiber transfer for 2006 ^{87}Sr measurement

It is clear from measurements of the excess phase noise introduced by fiber links (see Figs. 8.5 and 8.6) that while the maser reference is not degraded for short timescales (< 100 s), the passive transfer is not suitable for attempting to transfer the Cs-referenced maser signal with long timescale instabilities below 5×10^{-15} . In order to take full advantage of this reference's accuracy, it became necessary to actively stabilize the group delay of the fiber link.

We were not the first group to use active cancellation of the group delay fluctuations in order to transfer a microwave reference: NASA's Deep Space Network relies on this technique to distribute an ultrastable reference frequency from their Signal Processing Center at the Jet Propulsion Laboratory to multiple antenna sites for gravity wave searches, occultation science, and other radio science experiments [201, 202]. Instability $< 10^{-14}$ at 1 s was achieved for a 16 km-long fiber buried 1.5 m beneath the earth's surface, with the instability falling to below 10^{-17} for 10^4 s. A cryogenic sapphire oscillator was needed at the remote site for the short-term requirement, as an instability of about 5×10^{-15} at 1 s is required for the desired gravitational wave searches. Their servo had a small bandwidth and was only able to reduce the instability over timescales longer than 100 s. A similar technique was used for the transfer of a 100-MHz reference between Laboratoire de Physique des Lasers (LPL) and the Systèmes de Référence Temps Espace (SYRTE) for the comparison of an optical frequency at LPL and a microwave primary frequency standard at SYRTE [194, 36]. Two parallel optical links, each 43 km, were used for bidirectional transmission of a 100-MHz reference signal. However, in contrast to NASA's Deep Space Network, actuators were present in both

paths that allowed higher bandwidth control. This resulted in an instability over the combined 86-km round trip of one part in 10^{14} at 1 s, and several parts in 10^{17} at 10^4 s. In another experiment, time transfer over a 200-km optical fiber between Phoenix and Tucson, Arizona was demonstrated with sub-picosecond rms fluctuations over a 3.5-day period [203]. Though the actuator time-constant was only 10 s, the rms jitter was reduced to 630 fs over a 100-Hz bandwidth.

One important issue uncovered by the results in Ref. [36] is the effect of polarization mode dispersion. The fiber employed in many experiments is typically an urban-environment installation, and for reasons of cost and loss it is usually non-polarization-maintaining. As a result, the light launched into the fiber propagates differently through the fiber's polarization modes than the light that is returned for deducing the phase error of the link. Therefore, the principle that $\phi_{rt} = 2\phi_{f,inc}$ is violated because of PMD, even for frequencies below the bandwidth imposed by T_{rt} . Polarization scrambling of the light faster than the control bandwidth eventually reduced the effects of PMD in a new system, where the two 43-km links were connected to make one 86-km link for the purposes of an out-of-loop experiment with the local and remote ports of the link co-located. The polarization scrambling, in addition to using a higher 1-GHz modulation frequencies to enhance the effective SNR, decreased the transfer instabilities to a level of 5×10^{-15} at 1 s and 2×10^{-18} after one day [192].

For our 2006 measurement of the ^{87}Sr clock's frequency [147], we wanted to ensure the accuracy of the link would be better than 10^{-15} while keeping the instability below that of the H-maser reference being transmitted. Even though fast oscillations of the fiber link's frequency fluctuations could be averaged over, we were concerned that making a frequency measurement only during the warm or cold part of the day could lead to an overall frequency shift of the whole measurement. In order to characterize the typical magnitudes of such shifts, we retro-reflected a portion of the transmitted light from JILA back to NIST, and compared its phase with the local frequency ref-

erence being input to the fiber link. The phase comparison was made with a mixer in the usual homodyne phase detection way, with the mixer's output in Volts given by $0.22\text{Sin}(\phi_{ref} - \phi_{rt})$, up to a constant phase offset that could be adjusted with an electronic delay line. Figure 8.8 shows the voltage output of the mixer over the course of several days. Also plotted are solid blue and green lines with slopes representing (near zero mixer output voltage) a phase drift corresponding to a frequency shift of 10^{-14} or 10^{-15} , respectively. During some times of the day, a measurement of the clock frequency could be wrong by as much as 10^{-14} if only the passive transfer were used, even if several hours of data were averaged over.

Because of these large frequency shifts associated with the passive transfer of the microwave frequency reference, we decided to implement active cancellation of the fiber's group delay fluctuations. As shown in Fig. 8.4(a), the fiber length was stabilized using a fiber stretcher controlled by comparison of the local microwave phase at NIST with that of modulated light reflected back from JILA. The fiber stretcher's dynamic range was only sufficient to stabilize the drift for roughly an hour at a time, and then had to be briefly unlocked and reset. For our purposes this was sufficient to allow a 24-hour long measurement of the transition frequency, keeping the fiber link's phase stabilized during every spectroscopic scan of the transition. The 24-hour long record of the experiment is displayed in Fig. 8.9. The individual data points in red represent ~ 30 -s long scans of the transition, which was operated with a 10-Hz linewidth during spectroscopy. In green, referenced to the righthand axis, the output voltage of the in-loop mixer used to stabilize the fiber link's group delay fluctuations is shown. Note that on average, once per hour the fiber stretcher had to be reset due to its limited dynamic range. Nevertheless, the phase of the fiber link could be well-stabilized during these measurements, with a robust and easy-to-operate performance. Although the in-loop phase measurement shows that the frequency transfer was stabilized to a few parts in 10^{17} , we assigned a conservative uncertainty of 1×10^{-16} to the fiber transfer in our

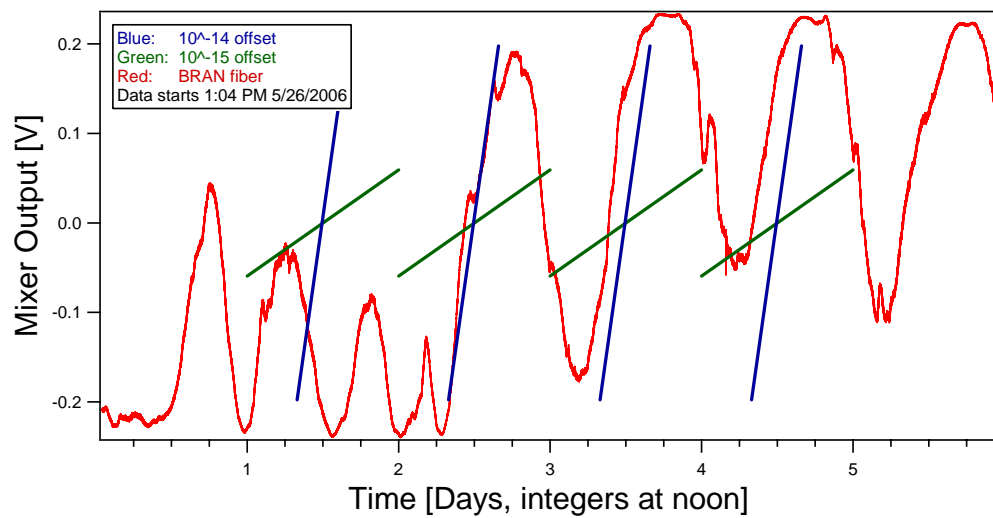


Figure 8.8: Mixer output voltage showing the microwave phase drift of the passive fiber link between NIST and JILA. At certain times of the day, the phase drift leads to average frequency shifts as high as 10^{-14} , which could lead to an incorrect measurement of the optical clock frequency.

final evaluation of the absolute frequency [147].

Coupled with other improvements in the Sr system [47], this approach resulted in the improved absolute frequency measurement uncertainty of 1.1 Hz (2.5×10^{-15}) [147]. An even higher-resolution approach using direct phase stabilization of the optical carrier will be discussed next, in Section 8.4, and it will be required to connect the JILA Sr standard to the single ion- and neutral atom-based optical standards located at NIST [172, 162].

8.4 Optical carrier transfer

The development of optical frequency references with short-term instabilities far lower than those of existing microwave frequency references will rely on the ability to transfer the frequency references with lower instability than that demonstrated by the microwave frequency transfer techniques previously discussed. Taking advantage of an optical frequency reference instead of microwaves gives much higher resolution for measuring phase noise fluctuations of the fiber link. The simplest approach to optical transfer is to directly transmit a cw optical frequency over optical fiber. At the remote end the cw reference can be directly compared with nearby optical frequencies, or it can be used to stabilize an optical frequency comb that can span the gap between the transmitted reference and other optical or even microwave frequencies. (See Chapter 3 for more details on using a frequency comb to phase-coherently link optical and microwave frequencies.) However, phase noise accumulated during transmission will dramatically decrease the reference's stability and must be actively cancelled.

Our particular motivation for optical carrier transfer is to compare our ^{87}Sr optical lattice clock against other optical clocks located 3.5 km away, at NIST. Although the BRAN fiber network employs typical telecommunications fiber with a single-mode cutoff wavelength around 1250-nm, the clocks at NIST and JILA have significantly lower wavelengths. One possible approach is to use a frequency comb to transfer the stability




Figure 8.9: 24 hours of continuous spectroscopy of the ^{87}Sr optical lattice clock transition frequency. Data points in red are plotted against the left axis, showing the absolute frequency measurement of the line-center of each 30-s scan of the optical transition. The green trace plotted against the righthand axis is the output voltage of the in-loop mixer used to stabilize the microwave phase of the fiber link. Notice the fiber stretcher had to be reset roughly once per hour due to its limited dynamic range. Keeping only the data when the phase was stabilized, the in-loop accuracy of the transfer system was a few parts in 10^{17} .

of one of the optical clocks' frequency to a cw laser operating in the telecommunications wavelength regime (1320 or 1550 nm). Then the cw laser could be transferred through the fiber network, with active cancellation of the fiber's optical phase fluctuations as shown in Fig. 8.4(b). However, the octave-spanning Ti:sapphire comb described in Chapter 7 does not provide spectral coverage of the telecommunications wavelengths, and so the scheme would require frequency doubling of the cw laser or difference frequency generation of the comb (see Chapter 6). A simpler approach that works for shorter (< 50 -km) fiber links is to use a frequency comb to transfer the stability of the clock laser to a cw laser operating at 1064 nm. Although it could in principle pose the problems of greater attenuation of power during the course of transfer and potential multi-mode transfer through the link, we decided to try this method as a first approach to direct optical clock comparisons over a fiber link. In particular, we were motivated by the demonstrated in-loop 3×10^{-15} (1-s averaging time) performance of just such a system [37] that had already been used to compare our I_2 standard to a stable optical reference at NIST at the 4×10^{-14} level (1-s averaging time). In what follows I will describe the incredible 6×10^{-18} at 1-s performance of our first simple system, and its extension to links much longer than just 3.5 km.

8.4.1 Experimental techniques for cw optical carrier transfer

We made measurements on several different schemes for cw optical carrier transfer, as summarized in Fig. 8.10. Three different lengths of fiber link were stabilized in the various measurements, corresponding to experimental configurations A (3.5-km one-way transfer to NIST), B (7-km transfer: JILA \rightarrow NIST \rightarrow JILA), and C (32-km transfer as in B but with an additional 25-km spool of fiber inserted in the link) as shown in Fig. 8.10.

In all three configurations, the 1064-light being sent through the fiber links is provided by a non-planar ring oscillator (NPRO 1) which is capable of providing 200

mW of output power, operates in a single-wavelength mode with a wide mode-hop-free tunability, and is optimistically specified to have a narrow (< 1 -kHz in 1 ms) linewidth. NPRO 1, operating at frequency ν_{NPRO1} , can be stabilized to the octave-spanning frequency comb using the usual heterodyne beat technique: a portion of its light is combined with a portion of the frequency comb's spectrum centered near 1064 nm onto photodiode (PD) 1. The resultant heterodyne beat between ν_{NPRO1} and a nearby comb mode is stabilized to an rf reference with frequency f_{b1} by feeding back to the PZT of NPRO 1, which has the effect of phase-locking ν_{NPRO1} to the octave spanning frequency comb. The frequency comb is in turn phase-locked to the clock laser of the ^{87}Sr lattice clock, operating at 698 nm, as discussed in Chapter 7. This provides the entire comb, and therefore ν_{NPRO1} , with the 10^{-15} instability and 0.5-Hz linewidth of the 698-nm clock laser. The majority of the power from NPRO 1 (~ 40 mW) is frequency-shifted into the -1 order of an AOM before being collapsed onto a 45° polarizer and launched into the fiber link. A small portion of the power (~ 2 mW) is first picked off by a polarizing beamsplitter (PBS), and then retroreflected through a double-pass quarter-wave plate in order to be photodetected on PD3. When light is returned from the remote end of the fiber link, it makes a second pass through the AOM in the reverse direction before being combined with the local copy on PD3 in order to make a heterodyne beat of frequency f_{b3} which is used to stabilize the fiber link's optical phase fluctuations. The stabilization is performed by feedback to the drive frequency of the AOM, f_{AOM} , in order to phase-lock f_{b3} to a rf reference. For active stabilization of the optical phase at the remote end of the fiber, it is convenient to use an AOM as the actuator. Unlike a fiber stretcher, which is typically used to stabilize the group delay of the link for microwave transfer, an AOM offers nearly limitless dynamic range for stabilizing the optical phase delay. This is because the AOM is able to act on the frequency directly, thereby cancelling any frequency offsets caused by phase drifts of the transmission link. On the other hand, a fiber stretcher acting directly on the path length can only stretch

the glass a finite and small amount, which can easily limit the locking range of a fiber transfer link. Once f_{b3} is stabilized, the light exiting the remote end of the fiber is also stabilized, and the remote user can enjoy the full stability of the 698-nm reference, wavelength-shifted by the frequency comb to 1064 nm.

Configuration A will be the system used to compare the ^{87}Sr optical clock at JILA to other optical references at NIST, and has already been used to compare stable sub-Hertz-linewidth lasers against each other, as will be discussed in Section 8.4.2. In this configuration, the flat remote tip of the BRAN fiber located at NIST is coated with a thin layer of gold so that it transmits 10% of the light sent from JILA, and retro-reflects the majority of the rest of the power back toward PD3 for the fiber phase-noise cancellation.

Configuration B is used in order to evaluate the performance of fiber transfer directly. In this case, the two BRAN fibers are connected to each other at NIST so that the resulting fiber link can be thought of as a single 7-km fiber. The “remote” end of the link is again coated with gold in order to retro-reflect most of the power toward PD3. However, in this configuration the output of the remote end can be directly compared against either ν_{NPRO1} or the comb (or both) on PD1 in order to evaluate the 7-km one-way transfer in an out-of-loop fashion.

For configuration C, a 25-km fiber spool is inserted between the two BRAN fibers at NIST in order to form a 32-km (each direction) link. In principle, the same gold tip experiment could be performed as in configuration B; however, the attenuation of 32 km of telecommunications fiber at 1064 nm results in too weak a signal making it back to PD3 after the retro-reflection. Each pass of the 32-km link attenuates the power by a factor of ~ 500 , and < 50 mW must be launched in order to avoid stimulated Brillouin scattering. At best only 200 nW could make it back to PD3, where it would have to compete against other small back-reflections from much earlier in the fiber link. In order to avoid this problem, a second laser (NPRO 2) operates at the remote end of the 32-km

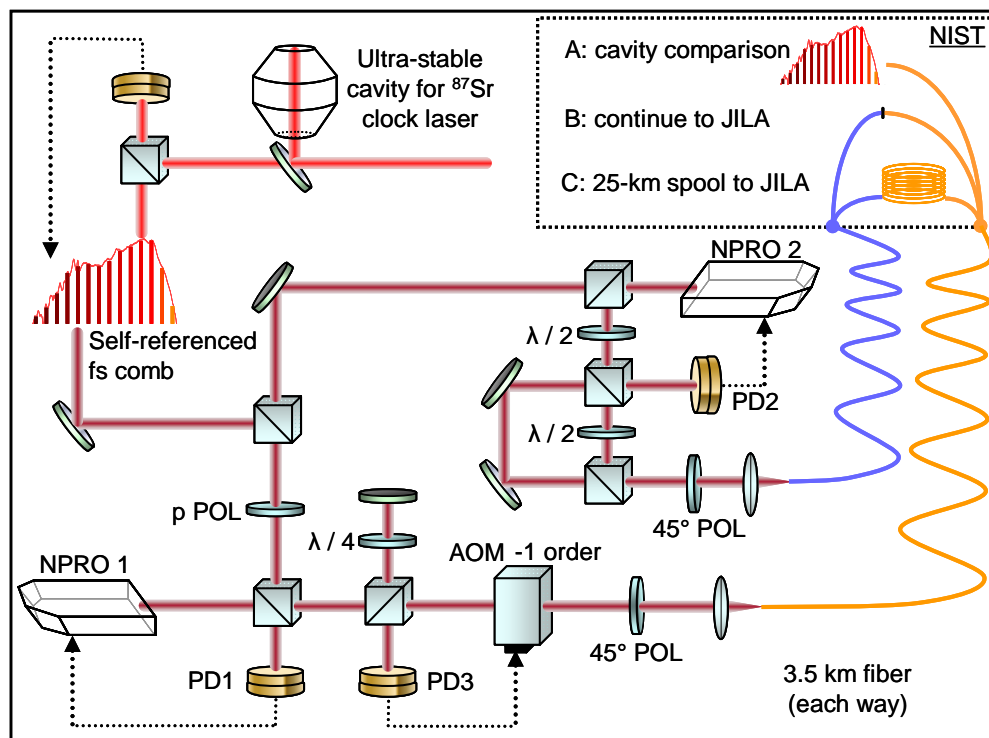


Figure 8.10: Experimental configurations for 3.5-, 7-, and 32-km fiber network transfer of a cw optical carrier: Laser NPRO 1, operating at 1064 nm, is stabilized to the 698-nm clock laser for an ^{87}Sr optical lattice clock. An octave-spanning, self-referenced frequency comb is used to span the 150-THz spectral gap. NPRO 1 is frequency-shifted by an AOM before traversing the fiber network. In configuration A, the gold-coated remote tip of the fiber, 3.5 km away, reflects 90% of the incident power on a return trip through the fiber for cancellation of the fiber's one-way optical phase noise. The remaining phase-stabilized light exiting the fiber at NIST is used to compare the JILA and NIST clock lasers via another frequency comb. In configuration B, the fiber is linked to a second fiber for a return trip to JILA, forming a 7-km one-way fiber link with local and remote ends co-located at JILA. The link's phase noise is cancelled in the same way as for configuration A, and an out-of-loop comparison is made between the light exiting the remote end of the 7-km fiber and either the local comb or NPRO 1. In configuration C, a 25-km fiber spool is inserted halfway into the 7-km link, at NIST. NPRO 2, operating at the remote tip of the fiber, forms a transceiver by locking to the signal exiting the 32-km link at the remote end. Most of the power from NPRO 2 is returned through the fiber link in order to beat with the local NPRO 1 on PD3, and stabilize the fiber link's optical phase. Independently heterodyning NPRO 2 (at the remote end) with the comb (at the local, but co-located, end) on PD1 allows out-of-loop characterization of the full 32-km fiber link. PD, photodiode; POL, linear polarizer

link in a transceiver configuration. Most of the power from NPRO 2 (~ 40 mW) (with frequency ν_{NPRO2}) is sent back into the fiber in order to heterodyne with ν_{NPRO1} on PD3 after passing through the AOM on the input side of the fiber link. A small part of the power (2 mW) is used to heterodyne with light exiting the fiber's remote end from NPRO 1. The resulting beat frequency, f_{b2} , is phase-locked to a rf reference by feedback to the PZT of NPRO 2, with f_{b2} carefully chosen so that the return light is offset in frequency from the input light inside the fiber. This scheme not only serves to boost the power for the return trip, but also lets light from the remote end be distinguished from other small reflections of the input light which make it back to PD3, since electronic filtering can distinguish the beat between NPRO 1 and NPRO 2 from the beat between NPRO 1 and its own back-reflected light. Finally, another 2 mW of power from NPRO 2 can be used to heterodyne against ν_{NPRO1} or the comb on PD1 in order to characterize the performance of the 32-km (one-way) link in an out-of-loop way.

We also operated configuration B with the transceiver for some experiments, rather than the retro-reflecting gold tip, in order to make more fair comparisons between the two schemes. In the transceiver scheme, f_{b1} on PD1 is first used to lock NPRO 1 to the comb. Then the remote laser NPRO 2 is locked to the output from the fiber link at the remote end by stabilizing f_{b2} from PD2. Finally, the fiber phase-noise is cancelled by stabilizing f_{b3} on PD3, such that ν_{NPRO2} is remotely phase-locked (with a frequency offset) to ν_{NPRO1} , ready to be used for any variety of experiments that can take advantage of long-distance ultrastable optical frequency transfer.

A few experimental details are important to mention here. The most important was that for the transceiver link, it was absolutely essential to include the 45° (relative to the polarizing axes of the nearby polarizing beamsplitters) polarizers at the input and output ends of the fiber link. Without them, the f_{b3} could not be locked at all, but simply by inserting the polarizers, the system worked as described in 8.4.2. We believe that polarization-mode-dispersion is responsible for this effect. Without the polarizer

at the remote end, the light exiting the fiber, used for f_{b2} to stabilize ν_{NPRO2} , would be orthogonally polarized with respect to the light returned through the fiber link. Since the polarization modes have a relative time-varying phase between them, the system would maximally violate the principle that $\phi_{f,inc}$ and $\phi_{f,ret}$ differ by only a constant as discussed in 8.2. Collapsing the polarization at the input and output of the fiber link helps to ensure that the same combination of polarization modes are sampled on the incident and return trips through the fiber link.

Another important detail is that the heterodyne beats are all polarization-sensitive, but the birefringence of the non-polarization-maintaining fiber link is time-dependent. Therefore we installed a set of fiber polarization paddles at the input end of the link in order to help the polarization states at the two 45° input and output linear polarizers to map onto each other as much as possible. For all of the experiments, it was sufficient to monitor the power in f_{b3} and use the polarization paddles to maximize it by hand once or twice per day. Finally, it became important to splice a small (5-m long) section of 1064-nm single-mode fiber at the input and output ends of the fiber link, in order to filter out any light that had been propagating in the few other transverse modes allowed by the telecommunications fiber at 1064 nm. One further experimental consideration for actively-stabilized direct transfer of optical frequencies is burst fluctuations of the fiber link's phase delay. Especially in the case of longer links (tens of km), sudden large phase delays can occur when incoherent noise processes distributed throughout the fiber randomly interfere constructively. If the servo doesn't have enough dynamic range or slew rate, cycle slips can occur when such a burst results in the sudden shift of the locking point to another stable point. This problem has been addressed in Ref. [40] by careful design of a digital phase-frequency discriminator which allows a phase error signal to have a wide discrimination range. As a result, cycle-slip-free operation could be achieved over a 25-km link, even with length excursions as large as 1.5 mm and rates as fast as 9.3 mm/s. In our case, we pre-scale f_{b3} by dividing its frequency by 50 before

phase-stabilizing it to the rf reference. This gives the servo a factor of 50 more range to phase slip and recover without an actual cycle slip.

With these improvements, each of the configurations and schemes has operated robustly for hours at a time. In fact, the main limit to the amount of time the system could be operated was usually either the clock laser or the comb's lock to it. Configuration B with retro-reflection from the gold tip could be operated almost indefinitely when NPRO 1 was not stabilized to the comb, as long as the polarization paddles were used often enough to compensate for the fiber link's slow variations of birefringence.

8.4.2 Coherence of optical carrier transfer

Since all of the servo loops in the experiment are phase-locks, it is most appropriate to begin characterization of the transfer experiments by measuring the phase noise of light at the remote end of the fiber relative to either NPRO 1 or the comb (at 1064 nm) at the input end, in configurations B and C. For these experiments there are three separate timescales relative to T_{rt} that are of importance, as discussed in Section 8.2.1. The first is the timescale of noise processes in the fiber link; cancellation of the one-way phase noise of the transfer path using ϕ_{rt} relies on the assumption that the noise processes are stationary during T_{rt} . The second important timescale is the inverse of the servo bandwidth, which must be longer than T_{rt} . The most important timescale for direct optical transfer is the source's coherence time, and the line-shape of the heterodyne beat depends on the source's coherence time relative to T_{rt} . Since a local portion of the source laser is being used as a reference to measure the fiber-induced phase noise, intrinsic phase fluctuations of the laser itself should be stationary during T_{rt} . This is equivalent to saying that the inverse of the optical reference's linewidth, i.e. the source's coherence time, should be longer than T_{rt} . Indeed, if T_{rt} is significantly longer than the reference's coherence time, the recovered beat is simply a measure of the reference's intrinsic line-shape.

A useful way to measure the effect of phase noise (although not phase noise directly) is therefore by measuring the lineshape of the out-of-loop beat on PD1 between the remote light (either light exiting the remote gold tipped fiber in configuration B or ν_{NPRO2} in configuration C) and the local input optical reference. In all of the following lineshape measurements, the heterodyne beat was mixed down from typical values between 50 and 100 MHz to 50 kHz against a low-phase-noise rf reference. The 50-kHz signal was then characterized on a FFT machine.

Figure 8.11(a) shows the first out-of-loop line-shape measurement between ν_{NPRO1} at the local end of the fiber and $\nu_{NPRO1} + f_{AOM}$ at the output of the 7-km fiber (configuration B with retro-reflection), as measured on PD1 when NPRO 1 is left free-running. The coherence time of the free-running NPRO 1 is $\sim (\pi \times 1 \text{ kHz})^{-1} = 318 \mu\text{s}$, whereas for the 7-km link (assuming phase velocity $= c/1.5$), $2\pi T_{rt} \sim 220 \mu\text{s}$. Since these two timescales are similar, we expect the line-shape to take the form of a spectral “delta” function sitting on top of a modified Lorentzian pedestal. This is indeed what we observe: the blue trace is taken with no active feedback to the AOM, showing how the fiber link introduces enough phase noise on the optical carrier to broaden it by 800 Hz (FWHM, relative to the input optical carrier). When active cancellation of the fiber link-induced phase noise is applied, the red trace is recovered, clearly showing the 1.5-kHz unity-gain bandwidth of the servo.

Figure 8.11(c) shows the narrow coherent central feature in a 400-Hz span with 0.5-Hz RBW. Notice the power spectral density for the passive transfer case (blue) appears as white noise on this scale, and is reduced by > 20 dB when the fiber is actively cancelled. Therefore the servo is successfully taking power which had been spread from the carrier to the fiber noise-induced sidebands and putting it back into the carrier.

In order to measure the full ability of the fiber transfer system, we then stabilized ν_{NPRO1} to the sub-Hz 698-nm clock laser in order to give it a coherence time longer than

1 second. For this measurement, we detected the beat between light exiting the remote tip of the fiber and a nearby frequency comb mode, as an out-of-loop measurement of the system’s ability to transfer the purity of a comb mode through 7 km of fiber. The resulting line-shape is shown in Fig. 8.11(b) with a 1-mHz RBW-limited linewidth. With such a long coherence time for ν_{NPRO1} , and due to the rapid roll-off of fiber-induced phase noise beyond 1-kHz, the line-shape looks like an exceedingly pure δ -function at ranges including Fourier frequencies as high as 50-kHz. The 7-km fiber network was therefore phase-stabilized well-enough that a 1000-s coherence time could be established between the frequency comb at one end of the fiber and the cw light at the other end.

Pushing the system one step further, we explored the limits imposed on fiber transfer through longer links by making the same line-shape measurement for configuration C, the 32-km link with a transceiver at the remote end. The line-shape displayed in Fig. 8.11(d) was measured by mixing ν_{NPRO2} at the remote end of the link against the same comb mode used to stabilize ν_{NPRO1} at the local end. In this case, a narrow coherent feature is still visible with a linewidth < 30 mHz (here only measured down to 0.5 Hz for comparison against Fig. 8.11(c)). However, the SNR for a 0.5-Hz RBW decreased from 50 dB, for the 7-km link, to only 15 dB for the 32-km link. Although the central carrier now only contains $< 10\%$ of the total power in a 10-kHz bandwidth, narrow-band filtering should allow the coherence of the reference to be recovered at the remote end by a sufficiently-low-phase-noise cleanup oscillator. Therefore, the bandwidth limitation imposed by the 32-km link still allows for coherent transfer of an ultrastable optical reference.

As a practical demonstration of the application for such coherent transfer, the 698-nm ultrastable laser, serving as the local oscillator for the JILA ^{87}Sr lattice clock, was compared against ultrastable 1126-nm light from the NIST Hg^+ optical clock system. Both clocks’ local oscillators have demonstrated sub-hertz linewidths in other independent measurements [46, 44]. The fiber link was set up in configuration A, and light

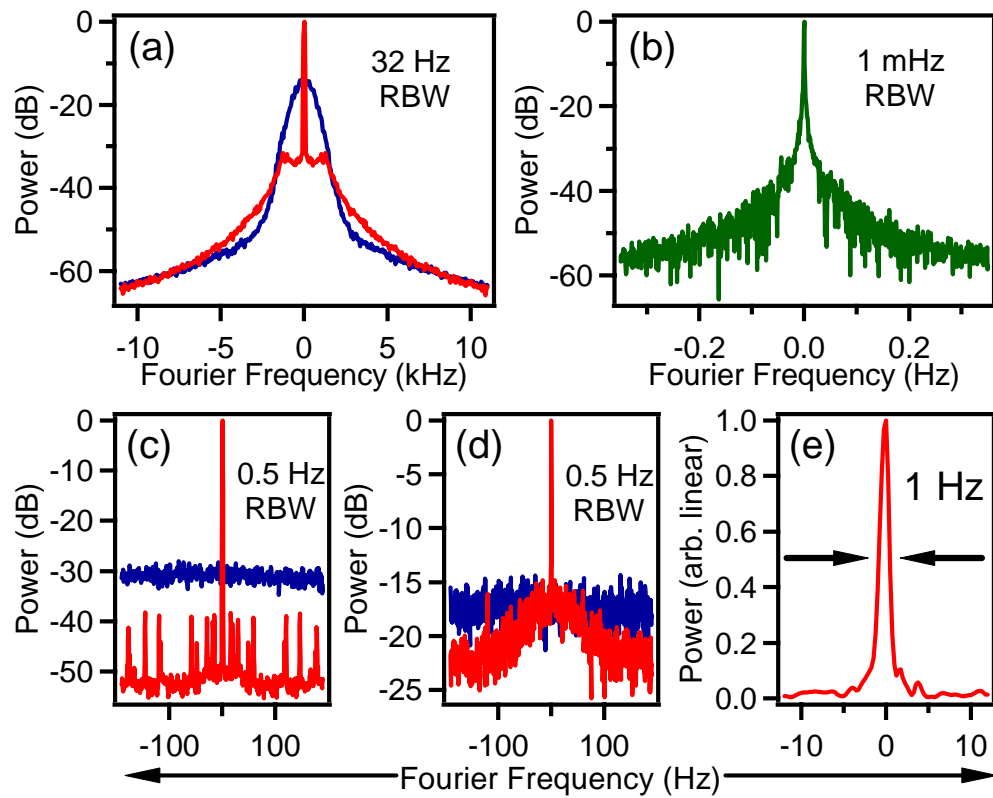


Figure 8.11: (a) Linewidth of heterodyne beat between NPRO 1 at the local end of the fiber link and the optical carrier after being transmitted 7 km. (b) NPRO 1 after 7-km transmission, heterodyned with a nearby comb mode to which it was stabilized at the input end. (c) 0.5-Hz-resolution display of the beat shown in (a). (d) NPRO 2 heterodyned against the frequency comb for the 32-km system employing NPRO 2 as a transceiver at the remote end. (e) 1-Hz linewidth for the effective beat between independent stable cavities separated by 3.5 km and 163 THz.

exiting the gold-tipped fiber at NIST was heterodyned against another octave-spanning comb that was phase-locked to the 1126-nm reference, as well as self-referenced. The heterodyne beat at 1064-nm is displayed in Fig. 8.11(e), and reveals a 1-Hz RBW-limited linewidth between the two stable lasers separated by 3.5 km. This measurement clearly shows the ability of the optical carrier transfer system to coherently compare some of the most stable and pure optical sources that can currently be made.

8.4.3 Stability of the transmitted optical carrier

In addition to measuring the coherence of the actively-stabilized fiber links, we directly counted the frequencies of the various out-of-loop heterodyne beats used to characterize the link. The upper two curves (open green circles and open brown diamonds) in Fig. 8.12 show the passive instabilities of the 7- and 32-km links, respectively. The magnitudes and shapes of these Allan deviations could change by as much as a factor of 3 at various timescales, depending on the time of day, weather, and other environmental factors. Nevertheless, it seems clear that for a passive fiber link, the instability is an order of magnitude worse than modern optical references such as the sub-Hz clock lasers can achieve, and therefore needs to be improved by active cancellation techniques.

The blue, closed circles in Fig. 8.12 are the Allan deviation of the heterodyne beat between ν_{NPRO1} and light exiting the gold-tipped remote end of the 7-km fiber link. The beat frequency was mixed down to 50-kHz against a stable rf reference, in order to generate a signal suitable for high-resolution frequency counting. Then the absolute frequency fluctuations were normalized to 282 THz in order to report fractional frequency instabilities of transferred optical carrier relative to ν_{NPRO1} , the optical input reference. The portion of the data for averaging times $\tau < 1000$ s was measured with a 1-s gate time for counting, while data for $\tau > 1000$ s was measured with a 10-s gate time. For this measurement, NPRO 1 was not stabilized to the frequency comb. Therefore the Allan Deviation represents an out-of-loop measurement of the fiber transfer instability

for transmission of a cw laser through a 7-km fiber link. It is also important that the system remained locked and robust over the course of the 70 continuous hours used to make the 10-s gate time measurement. Of the 25,036 points, only 78 were outliers and were removed before calculating the Allan deviation. The vast majority of the outliers occurred during two particular times when the fiber link's birefringence had drifted enough that spurious reflections in the link began to compete with the light that was retro-reflected from the remote gold-coated tip. Both times, when the problem was discovered it was quickly solved by using the in-fiber polarization paddles to change the birefringence so as to maximize the power of f_{b3} , the in-loop signal used to stabilize the optical phase of the link.

The red, closed diamonds in Fig. 8.12 are a direct out-of-loop measurement of the instability of the 32-km fiber link. This time, the signal being frequency-counted was a heterodyne beat between ν_{NPRO2} and the same comb mode at 1064 nm used to stabilize ν_{NPRO1} to the 698-nm ultrastable cw optical reference. The heterodyne beat was again mixed to 50-kHz for high-resolution counting. The data was measured only with 1-s gate times. Due to the additional complexity of maintaining several extra phase-locks associated with the 698-nm laser and the frequency comb, the data could only be continuously acquired for a few hours at a time. Either the comb would lose lock from the 698-nm laser due to its limited dynamic range, or the 698-nm laser would lose lock from the ultrastable cavity due to large thermal drifts in the lab causing the 698-nm diode laser to become unstable near a mode-hop. Despite this, we were able to count the frequency fluctuations during a period of 14 hours when the system only briefly (for less than 3 minutes each time) came unlocked a few times.

These results demonstrate that the 7-km transfer of an optical frequency reference can have an instability of 6×10^{-18} at 1-s averaging time, averaging down like $\tau^{-1/2}$ until $\tau \sim 4000$ seconds. Also, the transceiver system sent light stabilized to an optical frequency comb over a 32-km link with a demonstrated out-of-loop performance of




Figure 8.12: Instability of fiber transfer in various configurations. Open green circles and open brown diamonds are the passive instabilities of the 7- and 32-km fiber links, respectively. Closed red diamonds are the instability of the 32-km system, as measured by heterodyning the “remote” NPRO 2 against the comb-mode to which the “local” NPRO 1 is phase locked. Closed blue circles represent the instability of the 7-km link, measured by heterodyne beat between light exiting the remote fiber tip and a local copy of NPRO 1 before transmission. Black triangles represent the in-loop instability of the beat used to stabilize the fiber link, for all distances and configurations of the link.

$2 \times 10^{-17} \tau^{-1/2}$ for $\tau < 4000$ seconds. Ignoring for a moment the diurnal “bump” in the Allan deviation at $\sim 40,000$ s, we addressed the question of what limits this instability performance as follows. First, in order to make sure the electronics were not the limit, we counted the in-loop signal f_{b3} during the measurements. Its Allan deviation, which was the same for both the 7- and 32-km fiber links, is shown in Fig. 8.12 as the closed black triangles. The frequency fluctuations of f_{b3} represent fluctuations of light having passed through a round trip (14 or 64 km) of the fiber link, and were therefore divided by two before calculating the Allan deviation as plotted. Since the in-loop instability is a factor of 6 lower than the 7-km out-of-loop instability, the electronics must not be the main limit to the instability. We confirmed that the main limit was the out-of-loop optical paths used to make the measurement in several ways. First, opening the box enclosing the system increased the instability by an order of magnitude, even though the in-loop instability did not change. Second, a 5-m piece of fiber was substituted for the 7-km fiber in the retro-reflecting scheme, and the results were the same. Third, the 32-km transceiver system has a factor of 5 more out-of-loop optical components than the 7-km retro-reflecting system, which is consistent with its Allan deviation being essentially a constant multiple of 2.5 greater than the 7-km system. In order to confirm this last point, the 7-km fiber was operated in a transceiver configuration, and the out-of-loop instability increased to that of the 32-km system. We believe the diurnal frequency drift as evidenced at timescales $\tau > 4000$ s is therefore also due to the slow daily temperature cycle of our laboratory environment.

Finally, the accuracy of the fiber links was characterized by averaging the full 60 hours of data for the 7-km link together. The result was accurate to less than 1×10^{-19} , consistent with the trend of the Allan deviation at the longest timescales. For the 32-km link, the accuracy was 5×10^{-19} , also consistent with the measured instability.

8.4.4 Phase noise of the 7-km link

In order to make a true measurement of the residual phase noise of the 7-km actively-stabilized link, we mixed the out-of-loop beat (between ν_{NPRO1} and light exiting the gold-tipped remote end of the fiber) down to dc against a low-phase-noise rf reference, in a homodyne scheme. The voltage fluctuations of the mixer's output were measured by an FFT machine for Fourier frequencies below 100 kHz, and by an rf spectrum analyzer for Fourier frequencies above 100 kHz. These voltage fluctuations were normalized to the mixer output range and to the optical carrier frequency in order to characterize the optical phase fluctuations of the remotely-transferred frequency. The phase power spectral density is displayed in Fig. 8.13(a) as a red curve for the 7-km link, and as a blue curve when a short 2-m piece of fiber was substituted into the system. When the 2-m fiber was used, no changes to the servo gains or corners was made. However, the optical powers used were attenuated in order to simulate the same powers available from the 7-km system. Note that the noise between several Hz and 20 kHz is larger for the 7-km system. This excess noise must average out at timescales longer than 1 s, since for $\tau > 1$ s the instabilities of the 2-m and 7-km links performed identically. Finally, the sharp drop in phase noise spectral density beyond 3 MHz is due to the bandpass filter used to filter the original heterodyne beat signal near 77 MHz. Even if the noise floor had remained flat (limited by the white phase noise floor of the photodetection and subsequent phase-measurement electronics) at the 10^{-12} rad²/Hz level out to 50 MHz, the contribution to rms phase fluctuations would be insignificant.

We also converted the phase noise spectral density to units of timing jitter spectral density, and integrated it from 10 mHz to 30 MHz, as displayed in Fig. 8.13(b). The total rms timing jitter integrated in this bandwidth was 0.085 fs, or less than 0.024 of an optical cycle at 1064 nm. Although the cw field exiting the remote end of the fiber link does not directly provide a sharp rising edge for synchronization pur-




Figure 8.13: (a) Relative phase noise spectral density between the local cw reference and the signal exiting the 7-km fiber link. (b) The phase noise shown in (a) is integrated and displayed as rms timing jitter, demonstrating 7-km timing transfer with 0.85 fs of jitter integrated from 10 mHz to 25 MHz.

poses, a frequency comb could be employed to recover this level of timing jitter for its pulse train's repetition rate. This extremely low level of timing jitter for fiber transfer should find application for advanced light sources at synchrotrons and phase-coherent radio telescope arrays requiring orders of magnitude improvement over previous remote synchronization techniques.

The performance for actively-stabilized transfer of a cw optical beam represents several orders of magnitude improvement over the best results for stabilized transfer of a microwave frequency reference through amplitude modulation. This is made possible by the choice of a much higher (optical as opposed to microwave) modulation frequency. In fact, for the direct transfer of an optical beam, the optical carrier frequency is the modulation frequency to be recovered. As directly manifested in Eqn. (2.13), use of a higher transfer frequency allows the phase noise (as measured in radians) to be significantly higher for a given level of instability, putting less stringent requirements on the SNR needed to achieve a certain instability. For instance, for a counting bandwidth of 100 kHz and a carrier frequency of 1 GHz, 90 dB of SNR in a 100-kHz bandwidth is necessary to achieve an instability below 1×10^{-14} at 1 s integration time. If an optical frequency corresponding to 1064 nm is used for transfer instead, only 10 dB of SNR should be sufficient to push the instability an order of magnitude lower, to $< 1 \times 10^{-15}$ at 1 s integration time, if only the fundamental phase noise limit due to SNR is considered. (However, more than 12 dB of SNR is necessary to ensure cycle-slip-free counting.) In Ref. [37], the performance for passive microwave transfer by amplitude modulation of a cw beam was limited by the achievable SNR of the frequency-counting measurement. Such insufficient SNR is also responsible for the passive instability measurement of microwave transfer in Fig. 8.5 appearing a factor of ten worse for 1-s averaging than the passive instability of optical transfer.

A final very important point needs to be made regarding the real-world fiber links

employed, For Ref. [37], the SNR for direct optical carrier transfer was primarily limited by back-reflections from various connectors along the transmission fiber. Since those results we have fusion-spliced the fiber at all the previously-connectorized junctions in order to eliminate the effects of these back-reflections. Now with a much higher effective SNR available, the simple retro-reflection system performs with two orders of magnitude less instability. In this work we also demonstrated the transceiver method for distinguishing the signal of interest (returning from the remote end) from reflections in the fiber link. This increased the effective SNR by at least 30 dB in the case of the 7-km link. A similar scheme was used in Ref. [38], involving an AOM at the remote end of the fiber link. Here, the transceiver additionally serves as an amplifier at the remote end so that extra power can serve as the repeater signal, to be used for the next stage of a much longer multi-stage fiber frequency distribution link. The main point is that back-reflections in the fiber links are of extreme importance and should be minimized in any future installation of a fiber link to be used for ultrastable frequency transfer.

8.5 Remote synchronization of ultrafast lasers

In addition to the microwave and optical fiber transfer measurements I discussed in the previous two Sections, I had a small role to play in a remote synchronization experiment involving two femtosecond fiber lasers. Whereas the frequency transfer experiments were motivated by the needs of our Sr optical clock, the remote synchronization experiment was motivated by being the first step toward fiber transfer of the entire frequency comb, which would ultimately represent the most useful form of fiber transfer since the end user would be directly provided with a set of carrier frequencies spanning the electromagnetic spectrum. The first step toward this goal would be to stabilize fluctuations in the group delay of the transmission fiber, to produce a stabilized repetition frequency at the remote end. Even if the group delay is stabilized, dispersion would still be present in the transmission path, significantly broadening the

pulses. Therefore, the remote user might like to synchronize his own, transform-limited pulse train to the incoming reference. This exactly corresponds to the case of remote synchronization of the pulse trains of two fiber lasers.

We performed actively stabilized microwave frequency transfer using a frequency comb in order to synchronize two remotely located femtosecond lasers [135]. The frequency transfer was performed with a setup similar to that of Fig. 8.4(c). To achieve remote synchronization, the pulse train from the reference laser was transmitted through several kilometers of fiber using the stable transfer technique of Ref. [187]. Then the remote slave laser was synchronized to the stable incoming pulse train with a high-bandwidth actuator as in Ref. [204]. The setup used for remote synchronization experiments is shown in Fig. 8.14. In our experiment, two mode-locked fiber lasers were constructed so that harmonics of their repetition frequencies coincided at ~ 93 MHz. The transmitting laser was free running, while the slave laser had an electro-optic modulator (EOM) [204] and a PZT inside its cavity. The slave laser was synchronized to the reference laser by way of a microwave phase lock of the repetition frequencies. In this scheme, the PZT has a large dynamic range ($14 \mu\text{m}$) that allows for long timescale locking (> 12 hours), while the EOM provides a high bandwidth actuator to achieve tight synchronization.

To characterize the residual timing jitter of the synchronization we used a crossed beam, background free, optical cross-correlation of the two lasers' pulse trains, as discussed in Section 2.1.1. The two pulse trains were focused onto a LiIO_3 crystal (type I phase matching), which generated sum frequency light (SFG) when the two pulses overlapped in time and space. To achieve temporal overlap, we used two phase-locked loops that operate at two different timing resolutions [20, 21]. A coarse-timing loop operated at the fundamental frequency of 93 MHz, while a higher resolution loop operated at 7.6 GHz (80^{th} harmonic of f_{rep}). A phase shifter in the fundamental frequency loop allowed coarse-timing adjustments such that temporal overlap between the two pulse

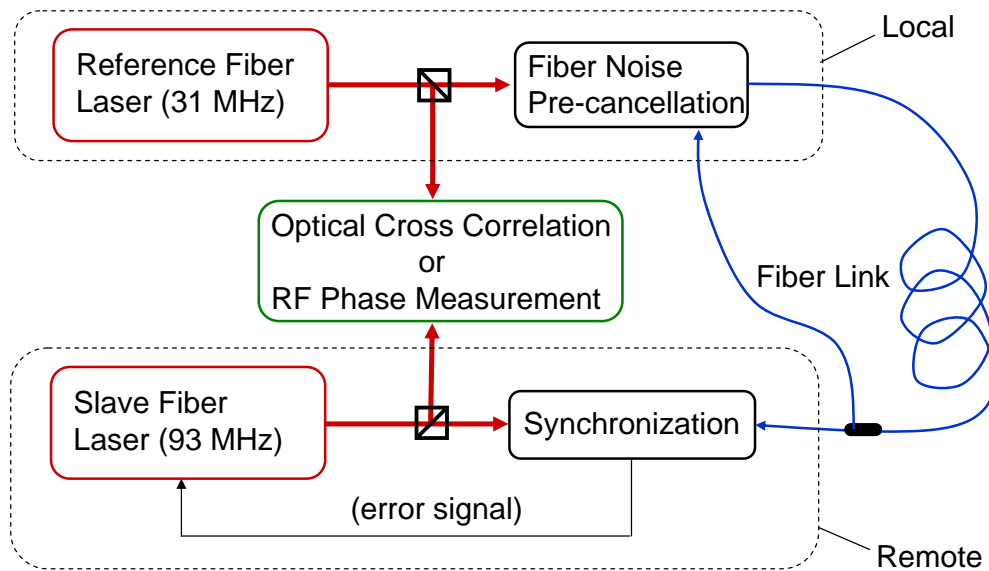


Figure 8.14: High-level diagram showing the setup for remote synchronization of two femtosecond fiber lasers. The repetition rate of the reference laser is transmitted through 4.5 km of dispersion-shifted fiber with active stabilization of the group delay. The remote laser is synchronized to the pulse trains exiting the fiber link using rf techniques, and an out-of-loop optical cross correlation is used to characterize the net performance of the system. (Adapted from Ref. [135])

trains could be found. Once an SFG signal was observed on a photo-multiplier tube (PMT), we measured the total cross correlation width to calibrate SFG intensity fluctuation to timing jitter conversion. We then transferred control from the fundamental frequency loop to the 7.6-GHz loop. A phase shifter in the high-harmonic loop allowed fine tuning of the timing overlap of the pulses such that we could position the SFG signal at the steepest point of the cross correlation slope to obtain the most sensitive measurement of pulse timing jitter, which is proportional to the amplitude fluctuations of the SFG signal. We monitored the amplitude fluctuations of the SFG signal through a 50-MHz low-pass filter to determine the timing jitter within an integration bandwidth up to the Nyquist frequency.

It is important to note that the cross-correlation was performed on the same optical table that holds the two lasers used in the experiment. This configuration allowed a direct comparison of the two lasers, which revealed the timing jitter due to both the transmission path and the slave laser's locking ability. Also, the proximity of the two lasers with the cross correlator allowed a relatively narrow cross-correlation width of 165 fs because the pulses had not spread significantly in time due to dispersion.

Initially, synchronization was performed with the two fiber lasers collocated (i.e. no fiber link between them). The synchronization was within a level of 10 fs over the entire Nyquist bandwidth [Fig. 8.15(a)], which agreed with our rf-domain in-loop measurement of the timing jitter. A 4.5-km dispersion-shifted fiber was then inserted between the two lasers, with active cancellation of the fiber noise implemented. The transfer process introduced 16 fs of jitter; this jitter adds in quadrature with the synchronization jitter of 10 fs. The total timing jitter is thus around 20 fs, as can be seen in Fig. 8.15(b), which was directly observed from the out-of-loop cross-correlation measurement. This result is in agreement with the in-loop measurement discussed previously.

To date, distribution of a fully stabilized frequency comb (meaning stable, simul-

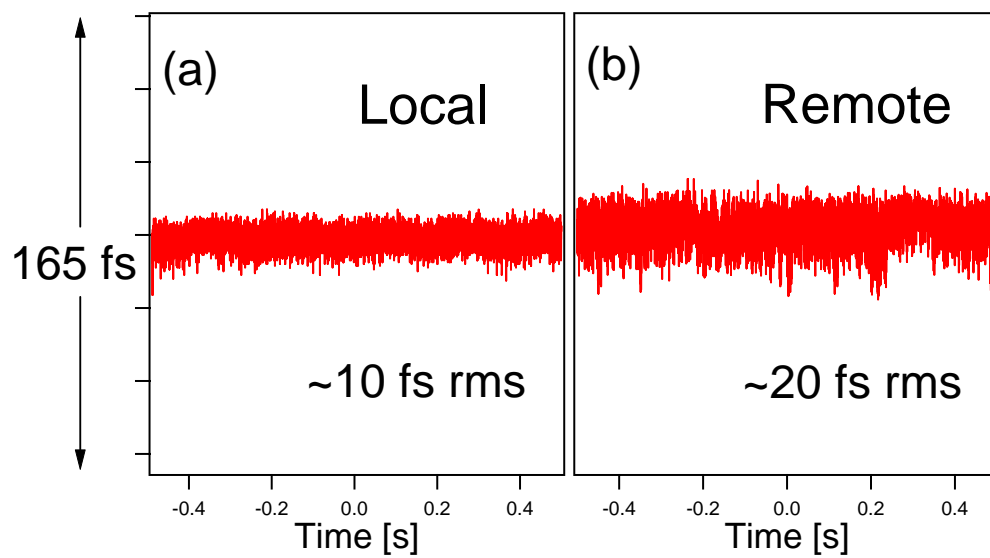


Figure 8.15: Optical cross-correlation measurement of the timing jitter for (a) local synchronization, and (b) slave synchronization through a fiber link. Both traces are shown with a 50-MHz low-pass filter that allows us to observe the timing jitter up to the Nyquist frequency. (Adapted from Ref. [135])

taneous delivery of both f_{rep} and f_0) has not been achieved. As the comb's capability to transmit both a microwave and optical frequency is one of the motivating factors for its use, it is worth discussing some considerations for accomplishing this goal. After fluctuations in the group delay of the transmission fiber (near a particular narrow spectral bandwidth) have been corrected (to produce a stabilized f_{rep} at the remote end), f_0 will experience shifts if there is a time-varying dispersion in the transmission path. Such fluctuations in f_0 are certainly possible with environmental perturbations to the optical fiber. In order to generate an error signal sensitive to these shifts, multiple error signals at multiple wavelengths could be formed, each similar in principle to the case of cw optical carrier transfer. Then the relationships between the various wavelength-dependent error signals could be used to compensate for various orders of dispersion fluctuations. For example, in the simplest case the group delay could be stabilized for a narrow spectral window, and the phase delay could be independently stabilized using information from another spectral region. As the number of orders of dispersion to be cancelled increases, the number of error signals and actuators also increases. The fiber transfer of the comb would be complete to as many orders of dispersion fluctuation as desired, at the cost of extreme complexity in the limit of stabilizing each comb line independently with line-by-line pulse shaping!

8.6 Summary

After studying in detail the three fiber-based techniques for distribution of a frequency reference, it is apparent that optically based approaches benefit from the higher phase resolution offered by optical frequencies, although the microwave-based approaches may be more straightforward and are already adequate for many applications. Ultimately the choice of which technique to use will depend on the application. For stability purposes (≥ 1 s) the modulated cw transfer, direct optical transfer, and the comb transfer all offer instabilities of a few parts in 10^{14} at 1-s without any active

stabilization. Once stabilization is utilized the direct optical transfer performs the best at 6×10^{-18} at 1 s, while microwave transfer using the comb performs at $< 7 \times 10^{-15}$ and the modulated cw beam technique achieves 5×10^{-14} . It is expected that stabilization of the optical phase during transfer of the comb would enable it to perform as well as the direct optical transfer for distributing an optical frequency reference. On the other hand, if the remote site has a frequency comb available to connect optical and microwave frequencies, it is expected that the direct optical transfer would work as well as frequency comb transfer for remote synchronization applications requiring low timing jitter at high sampling rates.

As frequency distribution networks of greater and greater lengths are developed, it is important to keep in mind some of the fundamental limitations to the stability that can be achieved. As we discussed earlier, thermal noise and shot noise present fundamental limits to the achievable measurement noise floor. To reduce the impact of these limits it is important to have at least enough light power incident on the system's photodetectors to reduce the thermal noise floor below the shot noise floor. However, this requires the development of highly linear, high-current photodetectors to avoid amplitude-to-phase conversion. Also, since stabilization techniques rely on the detection of round-trip phase noise to deduce and cancel one-way phase noise, it is important to keep in mind that only noise that is stationary during the round-trip time can be effectively cancelled. Therefore, as the distribution distances get longer, the bandwidth within which noise can be cancelled becomes smaller. One way to overcome this limitation would be to have "repeater" stations that would receive a stabilized frequency reference and phase lock a local oscillator to this reference, and then retransmit this reference to the next repeater in the chain. In this scenario the maximum transmission length that must be stabilized never exceeds that dictated by the bandwidth of noise present. To this extent, the transceiver configuration C discussed in Section 8.4 already provides a suitable repeater station, since NPRO 2 can also be used as the input to the next segment of a multiple-

repeater link.

An excellent question can be formed, as to how long a single fiber link can be made, given that an extremely pure “cleanup” oscillator could in principle be placed at the remote end. For instance, such an oscillator might have a 1-s coherence time, and therefore only need to be slaved to the transferred signal with a 1-Hz bandwidth in order to maintain phase coherence with the original reference. Then as long as the fiber noise cancellation has sufficient gain to cancel noise below 1-Hz, such a system should in principle work. Unfortunately, for our 7- and 32-km links, the stability performance was not yet limited by the fiber link and its associated round-trip bandwidth, and so we can not yet extrapolate the performance to longer links in order to determine at what length the uncanceled noise within a 1-Hz bandwidth becomes too large. Such an investigation would indeed be a worthy pursuit for future experiments!

As we look forward to the future of optical clock-based technology, we quickly come to the realization that characterization of optical clocks and distribution of their signals on the surface of the Earth will become seriously limited by our knowledge of the local gravitational potentials, as pointed out by a recent Reference Frame article by D. Kleppner [205]. An interesting solution might be to place a network of optical clocks in space, where frequency combs used as clockwork can function also for precise and absolute distance ranging [18] as well as remote transfer of optical clocks, such that a stable geoid could be defined from space. In that scenario, we are more than happy that the fiber-based approach might retire!

Bibliography

- [1] Jo Ellen Barnett. Time's Pendulum. Plenum Press, New York, 1998.
- [2] H. J. Metcalf and P. van der Straten. Laser Cooling and Trapping. Springer, New York, 1999.
- [3] 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(2):842–845, 2005.
- [4] 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(26):2081–2084, 1980.
- [5] 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(8):1705–1714, 1983.
- [6] Norman F. Ramsey. Applications of atomic clocks. In H. Figger, D. Meschede, and C. Zimmermann, editors, Laser Physics at the Limits, pages 3–8. Springer, Berlin, 2002.
- [7] L. A. Rawley, J. H. Taylor, M. M. Davis, and D. W. Allan. Millisecond pulsar psr 1937+21: A highly stable clock. Science, 238(4828):761–765, 1987.
- [8] T. Damour and J. H. Taylor. Strong-field tests of relativistic gravity and binary pulsars. Phys. Rev. D, 45(6):1840–1868, 1992.
- [9] K. M. Evenson, J. S. Wells, F. R. Petersen, B. L. Danielson, and G. W. Day. Accurate frequencies of molecular transitions used in laser stabilization: the 3.39- μm transition in ch_4 and the 9.33- and 10.18- μm transitions in co_2 . Appl. Phys. Lett., 22(4):192–195, 1973.
- [10] R. L. Barger and J. L. Hall. Wavelength of the 3.39- μm laser-saturated absorption line of methane. Appl. Phys. Lett., 22(4):196–199, 1973.
- [11] 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(5700):1318–1324, 2004.

- [12] Jeff Flowers. The route to atomic and quantum standards. Science, 306(5700):1324–1330, 2004.
- [13] S. G. Karshenboim. Some possibilities for laboratory searches for variations of fundamental constants. Can. J. Phys., 78(7):639–678, 2000.
- [14] 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(9):091301, 2001.
- [15] 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(15):150801, 2003.
- [16] 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 $^{199}\text{Hg}^+$ single-ion optical clock. Phys. Rev. Lett., 90(15):150802, 2003.
- [17] D. Hils and J. L. Hall. Improved kennedy-thorndike experiment to test special relativity. Phys. Rev. Lett., 64(15):1697–1700, 1990.
- [18] J. Ye. Absolute measurement of a long, arbitrary distance to less than an optical fringe. Opt. Lett., 29(10):1153–1155, 2004.
- [19] E. O. Potma, D. J. Jones, J.-X. Cheng, X. S. Xie, and J. Ye. High-sensitivity coherent anti-stokes raman scattering microscopy with two tightly synchronized picosecond lasers. Opt. Lett., 27(13):1168–1170, 2002.
- [20] L.-S. Ma, R. K. Shelton, H. C. Kapteyn, M. M. Murnane, and J. Ye. Sub-10-femtosecond active synchronization of two passively mode-locked ti:sapphire oscillators. Phys. Rev. A, 64(2):021802, 2001.
- [21] 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(5):312–314, 2002.
- [22] 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(5533):1286–1289, 2001.
- [23] A. Bartels, N. Newbury, I. Thomann, L. Hollberg, and S. Diddams. Broadband phase-coherent optical frequency synthesis with actively linked ti:sapphire and cr:forsterite femtosecond lasers. Opt. Lett., 29(4):403–405, 2004.
- [24] S. M. Foreman, D. J. Jones, and J. Ye. Flexible and rapidly configurable femtosecond pulse generation in the mid-ir. Opt. Lett., 28(5):370–372, 2003.
- [25] R. S. Judson and H. Rabitz. Teaching lasers to control molecules. Phys. Rev. Lett., 68(10):1500–1503, 1992.

- [26] A. Assion, T. Baumert, M. Bergt, T. Brixner, B. Kiefer, M. Seyfried, V. Strehle, and G. Gerber. Control of chemical reactions by feedback-optimized phase-shaped femtosecond laser pulses. *Science*, 282:919–922, 1999.
- [27] A. M. Weiner, D. E. Leaird, G. P. Wiederrecht, and K. A. Nelson. Femtosecond pulse sequences used for optical manipulation of molecular motion. *Science*, 247:1317–1319, 1990.
- [28] R. Bartels, S. Backus, E. Zeek, L. Misoguti, G. Vdovin, I. P. Christov, M. M. Murnane, and H. C. Kapteyn. Shaped-pulse optimization of coherent emission of high-harmonic soft x-rays. *Nature*, 406:164–166, 2000.
- [29] R. Kienberger, E. Goulielmakis, M. Uiberacker, A. Baltuska, V. Yakovlev, F. Bammer, A. Scrinzi, Th. Westerwalbesloh, U. Kleineberg, U. Heinzmann, M. Drescher, and F. Krausz. Atomic transient recorder. *Nature*, 427:817–821, 2004.
- [30] P. Roos, L. Xiaoqin, R. Smith, J. Pipis, T. Fortier, and S. Cundiff. Solid-state carrier-envelope phase stabilization via quantum interference control of injected photocurrents. *Opt. Lett.*, 30(7):735–737, 2005.
- [31] R. J. Jones, K. D. Moll, M. J. Thorpe, and J. Ye. Phase-coherent frequency combs in the vacuum ultraviolet via high-harmonic generation inside a femtosecond enhancement cavity. *Phys. Rev. Lett.*, 94(19):3201–3204, 2005.
- [32] M. J. Thorpe, K. D. Moll, R. J. Jones, B. Safdi, and J. Ye. Broadband cavity ringdown spectroscopy for sensitive and rapid molecular detection. *Science*, 311:1595–1599, 2006.
- [33] J. Levine. Introduction to time and frequency metrology. *Rev. Sci. Instrum.*, 70(6):2567–2596, 1999.
- [34] J. L. Hall and T. W. Hänsch. History of optical comb development. In J. Ye and S. T. Cundiff, editors, *Femtosecond Optical Frequency Comb Technology: Principle, Operation, and Applications*, pages 1–11. Springer, New York, 2005.
- [35] A. Bauch, J. Achkar, S. Bize, D. Calonico, R. Dach, R. Hlavac, L. Lorini, T. Parker, G. Petit, D. Piester, K. Szymaniec, and P. Urich. Comparison between frequency standards in europe and the usa at the 10^{-15} uncertainty level. *Metrologia*, 43:109–120, 2006.
- [36] F. Narbonneau, M. Lours, S. Bize, A. Clairon, G. Santarelli, O. Lopez, Ch. Daussy, A. Amy-Klein, and Ch. Chardonnet. High resolution frequency standard dissemination via optical fibre metropolitan network. *Rev. Sci. Instrum.*, 77:064701, 2006.
- [37] 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(7):1459–1467, 2003.

- [38] 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(21):1777–1779, 1994.
- [39] I. Coddington, W. C. Swann, L. Lorini, J. C. Bergquist, Y. Le Coq, C. W. Oates, Q. Quraishi, K. S. Feder, J. W. Nicholson, P. S. Westbrook, S. A. Diddams, and N. R. Newbury. Coherent optical link over hundreds of metres and hundreds of terahertz with subfemtosecond timing jitter. *Nature Photonics*, 1(5):283–287, 2007.
- [40] M. Musha, Y. Sato, K. Nakagawa, K. Ueda, A. Ueda, and M. Ishiguro. Robust and precise length stabilization of a 25-km long optical fiber using an optical interferometric method with a digital phase-frequency discriminator. *App. Phys. B*, 82(4):555–559, 2006.
- [41] R. W. Schoenlein, W. P. Leemans, A. H. Chin, P. Volfbeyn, T. E. Glover, P. Balling, M. Zolotarev, K. J. Kim, S. Chattopadhyay, and C. V. Shank. Femtosecond x-ray pulses at 0.4 angstrom generated by 90 degrees thomson scattering: A tool for probing the structural dynamics of materials. *Science*, 274(5285):236–238, 1996.
- [42] M. F. DeCamp, D. A. Reis, P. H. Bucksbaum, B. Adams, J. M. Caraher, R. Clarke, C. W. S. Conover, E. M. Dufresne, R. Merlin, V. Stoica, and J. K. Wahlstrand. Coherent control of pulsed x-ray beams. *Nature*, 413(6858):825–828, 2001.
- [43] A. L. Cavalieri, D. M. Fritz, S. H. Lee, P. H. Bucksbaum, D. A. Reis, J. Rudati, D. M. Mills, P. H. Fuoss, G. B. Stephenson, C. C. Kao, D. P. Siddons, D. P. Lowney, A. G. MacPhee, D. Weinstein, R. W. Falcone, R. Pahl, J. Als-Nielsen, C. Blome, S. Dusterer, R. Ischebeck, H. Schlarb, H. Schulte-Schrepping, T. Tschentscher, J. Schneider, O. Hignette, F. Sette, K. Sokolowski-Tinten, H. N. Chapman, R. W. Lee, T. N. Hansen, O. Synnnergren, J. Larsson, S. Techert, J. Sheppard, J. S. Wark, M. Bergh, C. Caleman, G. Hultdt, D. van der Spoel, N. Timneanu, J. Hajdu, R. A. Akre, E. Bong, P. Emma, P. Krejcik, J. Arthur, S. Brennan, K. J. Gaffney, A. M. Lindenberg, K. Luening, and J. B. Hastings. Clocking femtosecond x rays. *Phys. Rev. Lett.*, 94(11):114801, 2005.
- [44] B. C. Young, F. C. Cruz, W. M. Itano, and J. C. Bergquist. Visible lasers with subhertz linewidths. *Phys. Rev. Lett.*, 82(19):3799–3802, 1999.
- [45] M. Notcutt, L.-S. Ma, J. Ye, and J. L. Hall. Simple and compact 1-hz laser system via an improved mounting configuration of a reference cavity. *Opt. Lett.*, 30(14):1815–1817, 2005.
- [46] A. D. Ludlow, X. Huang, T. Zanon, S. M. Foreman, M. M. Boyd, S. Blatt, and J. Ye. Compact, thermal-noise-limited optical cavity for diode laser stabilization at 1×10^{-15} . *Opt. Lett.*, 32(6):641–643, 2007.
- [47] M. M. Boyd, T. Zelevinsky, A. D. Ludlow, S. M. Foreman, S. Blatt, T. Ido, and J. Ye. Optical atomic coherence at the one-second time scale. *Science*, in press.

- [48] V. P. Chebotayev, V. G. Goldort, V. M. Klementyev, M. V. Nikitin, B. A. Timchenko, and V. F. Zakharyash. Development of an optical time scale. Appl. Phys. B, 29(1):63–65, 1982.
- [49] C. O. Weiss, G. Kramer, B. Lipphardt, and E. Garcia. Frequency measurement of a ch_4 hyperfine line at 88 thz/“optical clock”. IEEE J. Quantum Electron., 24(10):1970–1972, 1988.
- [50] H. Schnatz, B. Lipphardt, J. Helmcke, F. Riehle, and G. Zinner. First phase-coherent frequency measurement of visible radiation. Phys. Rev. Lett., 76(1):18–21, 1996.
- [51] J. E. Bernard, A. A. Madej, L. Marmet, B. G. Whitford, K. J. Siemsen, and S. Cundy. Cs-based frequency measurement of a single, trapped ion transition in the visible region of the spectrum. Phys. Rev. Lett., 82(16):3228–3231, 1999.
- [52] K. Imai, M. Kourogi, and M. Ohtsu. 30-thz span optical frequency comb generation by self-phase modulation in an optical fiber. IEEE J. Quantum Electron., 34(1):54–60, 1998.
- [53] J. L. Hall. Nobel lecture: Defining and measuring optical frequencies. Rev. Mod. Phys., 78:1279–1295, 2006.
- [54] T. W. Hänsch. Nobel lecture: Passion for precision. Rev. Mod. Phys., 78:1297–1309, 2006.
- [55] J. L. Hall and M. Zhu. An introduction to phase-stable optical sources. In E. Arimondo, W. D. Phillips, and F. Strumia, editors, Laser Manipulation of Atoms and Ions, Proc. of the International School of Phys. “Enrico Fermi”: Course 118, “Enrico Fermi” Courses, pages 671–702. North Holland, Amsterdam, 1992.
- [56] M. Zhu and J. L. Hall. Stabilization of optical phase/frequency of a laser system: application to a commercial dye laser with an external stabilizer. J. Opt. Soc. Am. B, 10(5):802–816, 1993.
- [57] J. Rutman. Characterization of phase and frequency instabilities in precision frequency sources: Fifteen years of progress. Proc. IEEE, 66(9):1048–1074, 1978.
- [58] F. L. Walls. Phase noise issues in femtosecond lasers. In J. L. Hall and J. Ye, editors, Laser Frequency Stabilization, Standards, Measurement, and Applications, volume 4269 of Proceedings of SPIE, pages 170–177. SPIE, Bellingham, Washington, 2001.
- [59] D. B. Sullivan, D. W. Allan, and F. L. Walls, editors. Characterization of Clocks and Oscillators: NIST Technical Note 1337. U.S. GPO, Washington, DC, 1990.
- [60] D. W. Allan. Statistics of atomic frequency standards. Proc. IEEE, 54:221, 1966.
- [61] R. P. Scott, C. Langrock, and B. H. Kolner. High-dynamic-range laser amplitude and phase noise measurement techniques. IEEE J. Sel. Top. Quantum Electron., 7(4):641–655, 2001.

- [62] 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, Proceedings of the 2003 IEEE International Frequency Control Symposium and PDA Exhibition Jointly with the 17th European Frequency and Time Forum, pages 516–520. Institute of Electrical and Electronics Engineers, Piscataway, New Jersey, 2003.
- [63] P. J. Winzer. Shot-noise formula for time-varying photon rates: a general derivation. J. Opt. Soc. Am. B, 14(10):2424–2429, 1997.
- [64] J. J. McFerran, E. N. Ivanov, A. Bartels, G. Wilpers, C. W. Oates, S. A. Diddams, and L. Hollberg. Low-noise synthesis of microwave signals from an optical source. Electron. Lett., 41(11):650–651, 2005.
- [65] L. Hollberg, C. W. Oates, E. A. Curtis, E. N. Ivanov, S. A. Diddams, Th. 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(12):1502–1513, 2001.
- [66] F. L. Walls, E. S. Ferre-Pikal, and S. R. Jefferts. Origin of $1/f$ pm and am noise in bipolar junction transistor amplifiers. IEEE Trans. Ultrason. Ferroelectr. Freq. Control, 44(2):326–334, 1997.
- [67] L. Hollberg, S. A. Diddams, A. Bartels, J. J. McFerran, E. N. Ivanov, G. Wilpers, C. W. Oates, W. H. Oskay, and J. C. Bergquist. Generation of microwaves with ultra-low phase-noise from an optical clock. IEEE Proceed. Microwave Photonics, 2004.
- [68] P.-L. Liu, K. J. Williams, M. Y. Frankel, and R. D. Esman. Saturation characteristics of fast photodetectors. IEEE Trans. Microwave Theory Tech., 47(7):1297–1303, 1999.
- [69] D. A. Tulchinsky and K. J. Williams. Excess amplitude and excess phase noise of rf photodiodes operated in compression. IEEE Photon. Technol. Lett., 17(3):654–656, 2005.
- [70] J. K. A. Everard and C. Broomfield. Flicker noise removal in microwave oscillators using gaas based feedforward amplifiers. IEEE Int. Freq. Cont. Symp., pages 156–160, 2001.
- [71] E. S. Ferre-Pikal. Reduction of phase noise in linear hbt amplifiers using low-frequency active feedback. IEEE Trans. on Circuits and Systems – I, 51(8):1417–1421, 2004.
- [72] E. N. Ivanov, M. E. Tobar, and R. A. Woode. Microwave interferometry: Application to precision measurements and noise reduction techniques. IEEE Trans. Ultrason. Ferroelectr. Freq. Control, 45(6):1526–1536, 1998.
- [73] E. Rubiola, E. Salik, N. Yu, and L. Maleki. Phase noise measurements of low power signals. Elec. Lett., 39(19):1389–1390, 2003.

- [74] E. Rubiola, V. Giordano, and J. Gros Lambert. Very high frequency and microwave interferometric phase and amplitude noise measurements. Rev. Sci. Instr., 70(1):220–225, 1999.
- [75] K. J. Williams, R. D. Esman, and M. Dagenais. Effects of high space-charge fields on the response of microwave photodetectors. IEEE Photon. Tech. Lett., 6(5):639–641, 1994.
- [76] 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(4):1059–1065, 2003.
- [77] E. N. Ivanov, S. A. Diddams, and L. Hollberg. Study of the excess noise associated with demodulation of ultra-short infrared pulses. IEEE Trans. Ultrason. Ferroelectr. Freq. Control, 52(7):1068–1074, 2005.
- [78] A. Bartels, S. A. Diddams, C. W. Oates, G. Wilpers, J. C. Bergquist, W. H. Oskay, and L. Hollberg. Femtosecond-laser-based synthesis of ultrastable microwave signals from optical frequency references. Opt. Lett., 30(6):667–669, 2005.
- [79] T. Brown, 2006. unpublished work performed at JILA.
- [80] T. R. Schibli, J. Kim, O. Kuzucu, J. T. Gopinath, S. N. Tandon, G. S. Petrich, A. Kolodziejewski, J. G. Fujimoto, E. P. Ippen, and F. X. Kärtner. Attosecond active synchronization of passively mode-locked lasers by balanced cross correlation. Opt. Lett., 28:947–949, 2003.
- [81] Jean-Claude Diels and Wolfgang Rudolph. Ultrashort Laser Pulse Phenomena: Fundamentals, Techniques, and Applications on a Femtosecond Time Scale. Optics and Photonics. Academic Press, San Diego, 1996.
- [82] Y.-F. Chen, J. Jiang, and D. J. Jones. Remote distribution of a mode-locked pulse train with sub 40-as jitter. Opt. Exp., 14:12134–12144, 2006.
- [83] A. Pauchard, M. Bitter, Z. Pan, S. Kristjansson, L. A. Hodge, K. J. Williams, D. A. Tulchinsky, S. G. Hummel, and Y. H. Lo. Infrared-sensitive ingaas-on-si p-i-n photodetectors exhibiting high-power linearity. IEEE Photon. Tech. Lett., 16(11):2544–2546, 2004.
- [84] K. J. Williams and R. D. Esman. Design considerations for high-current photodetectors. J. Lightwave Technol., 17(8):1443–1454, 1999.
- [85] Th. Udem, R. Holzwarth, and T. W. Hänsch. Optical frequency metrology. Nature, 416(6877):233–237, 2002.
- [86] S. T. Cundiff and J. Ye. Colloquium: Femtosecond optical frequency combs. Rev. Mod. Phys., 75(1):325–342, 2003.
- [87] 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(4):1041–1058, 2003.

- [88] J. Ye and S. T. Cundiff. Optical frequency combs and their applications. In J. Ye and S. T. Cundiff, editors, Femtosecond Optical Frequency Comb Technology: Principle, Operation, and Applications, pages 12–53. Springer, New York, 2005.
- [89] C. M. DePriest, T. Yilmaz, P. J. Delfyett, Jr., S. Etemad, A. Braun, and J. Abeles. Ultralow noise and supermode suppression in an actively mode-locked external-cavity semiconductor diode ring laser. Opt. Lett., 27(9):719–721, 2002.
- [90] F. Quinlan, S. Gee, S. Ozharar, and P. J. Delfyett. Ultralow-jitter and -amplitude-noise semiconductor-based actively mode-locked laser. Opt. Lett., 31(19):2870–2872, 2006.
- [91] J. Rauschenberger, T. M. Fortier, D. J. Jones, J. Ye, and S. T. Cundiff. Control of the frequency comb from a modelocked erbium-doped fiber laser. Opt. Express, 10(24):1404–1410, 2002.
- [92] H. Hundertmark, D. Wandt, C. Fallnich, N. Haverkamp, and H. Telle. Phase-locked carrier-envelope-offset frequency at 1560 nm. Opt. Express, 12(5):770–775, 2004.
- [93] F. Tauser, A. Leitenstorfer, and W. Zinth. Amplified femtosecond pulses from an er:fiber system: Nonlinear pulse shortening and self-referencing detection of the carrier-envelope phase evolution. Opt. Express, 11(6):594–600, 2003.
- [94] B. R. Washburn, S. A. Diddams, N. R. Newbury, J. W. Nicholson, M. F. Yan, and C. G. Jrgensen. Phase-locked, erbium-fiber-laser-based frequency comb in the near infrared. Opt. Lett., 29(3):250–252, 2004.
- [95] F.-L. Hong, K. Minoshima, A. Onae, H. Inaba, H. Takada, A. Hirai, H. Matsumoto, T. Sugiura, and M. Yoshida. Broad-spectrum frequency comb generation and carrier-envelope offset frequency measurement by second-harmonic generation of a mode-locked fiber laser. Opt. Lett., 28(17):1516–1518, 2003.
- [96] E. P. Ippen. Principles of passive mode locking. Appl. Phys. B, 58(3):159–170, 1994.
- [97] S. T. Cundiff. Phase stabilization of ultrashort optical pulses. J. Phys. D: Appl. Phys., 35(8):R43–R59, 2002.
- [98] 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(5466):635–639, 2000.
- [99] 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(1):25–27, 2000.
- [100] R. Ell, U. Morgner, F. X. Krtner, 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(6):373–375, 2001.

- [101] T. M. Fortier, D. J. Jones, and S. T. Cundiff. Phase stabilization of an octave-spanning ti:sapphire laser. Opt. Lett., 28(22):2198–2200, 2003.
- [102] U. Morgner, R. Ell, G. Metzler, T. R. Schibli, F. X. Kärtner, J. G. Fujimoto, H. A. Haus, and E. P. Ippen. Nonlinear optics with phase-controlled pulses in the sub-two-cycle regime. Phys. Rev. Lett., 86(24):5462–5465, 2001.
- [103] T. M. Ramond, S. A. Diddams, L. Hollberg, and A. Bartels. Phase-coherent link from optical to microwave frequencies by means of the broadband continuum from a 1-ghz ti:sapphire femtosecond oscillator. Opt. Lett., 27(20):1842–1844, 2002.
- [104] Th. Udem, J. Reichert, R. Holzwarth, and T. W. Hänsch. Absolute optical frequency measurement of the cesium d_1 line with a mode-locked laser. Phys. Rev. Lett., 82(18):3568–3571, 1999.
- [105] J. Reichert, R. Holzwarth, Th. Udem, and T. W. Hänsch. Measuring the frequency of light with mode-locked lasers. Opt. Comm., 172:59–68, 1999.
- [106] K. W. Holman, R. J. Jones, A. Marian, S. T. Cundiff, and J. Ye. Intensity-related dynamics of femtosecond frequency combs. Opt. Lett., 28(10):851–853, 2003.
- [107] 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(4):1018–1024, 2003.
- [108] 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/ch₄-based optical molecular clock. Opt. Lett., 30(5):570–572, 2005.
- [109] 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(5665):1843–1845, 2004.
- [110] S. A. Diddams, D. J. Jones, J. Ye, S. T. Cundiff, J. L. Hall, J. K. Ranka, R. S. Windeler, R. Holzwarth, Th. Udem, and T. W. Hänsch. Direct link between microwave and optical frequencies with a 300 thz femtosecond laser comb. Phys. Rev. Lett., 84(22):5102–5105, 2000.
- [111] R. Holzwarth, Th. Udem, T. W. Hänsch, J. C. Knight, W. J. Wadsworth, and P. St. J. Russell. Optical frequency synthesizer for precision spectroscopy. Phys. Rev. Lett., 85(11):2264–2267, 2000.
- [112] O. D. Mücke, O. Kuzucu, F. N. C. Wong, E. P. Ippen, F. X. Kärtner, S. M. Foreman, D. J. Jones, L.-S. Ma, J. L. Hall, and J. Ye. Experimental implementation of optical clockwork without carrier-envelope phase control. Opt. Lett., 29(23):2806–2808, 2004.
- [113] M. Zimmerman, C. Gohle, R. Holzwarth, T. Udem, and T. W. Hänsch. Optical clockwork with an offset-free difference-frequency comb: accuracy of sum- and difference-frequency generation. Opt. Lett., 29:310–312, 2004.

- [114] J. D. Jost, J. L. Hall, and J. Ye. Continuously tunable, precise, single frequency optical signal generator. *Opt. Express*, 10(12):515–520, 2002.
- [115] T. R. Schibli, K. Minoshima, F.-L. Hong, H. Inaba, Y. Bitou, A. Onae, and H. Matsumoto. Phase-locked widely tunable optical single-frequency generator based on a femtosecond comb. *Opt. Lett.*, 30(17):2323–2325, 2005.
- [116] C. Gohle, T. Udem, M. Hermann, J. Rauschenberger, R. Holzwarth, H. A. Schuessler, F. Krausz, and T. W. Hänsch. A frequency comb in the extreme ultraviolet. *Nature*, 436(14):234–237, 2005.
- [117] S. A. Diddams, Th. Udem, J. C. Bergquist, E. A. Curtis, R. E. Drullinger, L. Hollberg, W. M. Itano, W. D. Lee, C. W. Oates, K. R. Vogel, and D. J. Wineland. An optical clock based on a single trapped $^{199}\text{Hg}^+$ ion. *Science*, 293(5531):825–828, 2001.
- [118] J. Ye, L.-S. Ma, and J. L. Hall. Molecular iodine clock. *Phys. Rev. Lett.*, 87(27):270801, 2001.
- [119] A. Apolonski, A. Poppe, G. Tempea, Ch. Spielmann, Th. Udem, R. Holzwarth, T. W. Hänsch, and F. Krausz. Controlling the phase evolution of few-cycle light pulses. *Phys. Rev. Lett.*, 85(4):740–743, 2000.
- [120] R. R. Alfano. *The Supercontinuum laser source: fundamentals with updated references*. Springer-Verlag, New York, 2006.
- [121] P. B. Corkum, N. H. Burnett, and M. Y. Ivanov. Subfemtosecond pulses). *Opt. Lett.*, 19(22):1870–1872, 1994.
- [122] R. A. Kaindl, M. Wurm, K. Reimann, P. Hamm, A. M. Weiner, and M. Woerner. Generation, shaping, and characterization of intense femtosecond pulses tunable from 3 to 20 μm . *J. Opt. Soc. Am. B*, 17(12):2086–2094, 2000.
- [123] J. Zhou, G. Taft, C.-P. Huang, M. M. Murnane, H. C. Kapteyn, and I. P. Christov. Pulse evolution in a broad-bandwidth ti:sapphire laser. *Opt. Lett.*, 19(15):1149–1151, 1994.
- [124] A. Leitenstorfer, C. Furst, and A. Laubereau. Widely tunable two-color mode-locked ti:sapphire laser with pulse jitter of less than 2 fs. *Opt. Lett.*, 20(8):916–918, 1995.
- [125] A. Bartels, S. A. Diddams, T. M. Ramond, and L. Hollberg. Mode-locked laser pulse trains with subfemtosecond timing jitter synchronized to an optical reference oscillator. *Opt. Lett.*, 28(8):663–665, 2003.
- [126] K. Read, F. Blonigen, N. Riccielli, M. M. Murnane, and H. C. Kapteyn. Low-threshold operation of an ultrashort-pulse mode-locked ti:sapphire laser. *Opt. Lett.*, 21(7):489–491, 1996.
- [127] F. Eickenmeyer, R. A. Kaindl, M. Woerner, T. Elsaesser, and A. M. Weiner. Controlled shaping of ultrafast electric field transients in the mid-infrared spectral range. *Opt. Lett.*, 25(19):1472, 2000.

- [128] N. Belabas, J. P. Likforman, L. Canioni, B. Bousquet, and M. Joffre. Coherent broadband pulse shaping in the mid infrared. *Opt. Lett.*, 26(10):743–745, 2001.
- [129] M. R. X. de Barros, R. S. Miranda, T. M. Jedju, and P. C. Becker. High-repetition-rate femtosecond mid-infrared pulse generation. *Opt. Lett.*, 20(5):480–482, 1995.
- [130] R. A. Kaindl, D. C. Smith, M. Joschko, M. P. Hasselbeck, M. Woerner, and T. Elsaesser. Femtosecond infrared pulses tunable from 9 to 18 μm at an 88-mhz repetition rate. *Opt. Lett.*, 23(11):861–863, 1998.
- [131] R. A. Kaindl, F. Eickemeyer, M. Woerner, and T. Elsaesser. Broadband phase-matched difference frequency mixing of femtosecond pulses in gas: Experiment and theory. *Appl. Phys. Lett.*, 75(8):1060–1062, 1999.
- [132] A. Lohner, P. Kruck, and W. W. Rühle. Generation of 200 femtosecond pulses tunable between 2.5 and 5.5 μm . *Appl. Phys. B*, 59(2):211–213, 1994.
- [133] S. Ehret and H. Schneider. Generation of subpicosecond infrared pulses tunable between 5.2 μm and 18 μm at a repetition rate of 76 mhz. *Appl. Phys. B*, 66(1):27–30, 1998.
- [134] D. R. Suhre, N. B. Singh, V. Balakrishna, N. C. Fernelius, and F. K. Hopkins. Improved crystal quality and harmonic generation in gas doped with indium. *Opt. Lett.*, 22(11):775–777, 1997.
- [135] D. D. Hudson, S. M. Foreman, S. T. Cundiff, and J. Ye. Synchronization of mode-locked femtosecond lasers through a fiber link. *Opt. Lett.*, 31(13):1951–1953, 2006.
- [136] L. Matos, D. Kleppner, O. Kuzucu, T. R. Schibli, J. Kim, E. P. Ippen, and F. X. Kaertner. Direct frequency comb generation from an octave-spanning prismless ti:sapphire laser. *Opt. Lett.*, 29:1683–1685, 2004.
- [137] A. Amy-Klein, A. Goncharov, C. Daussy, C. Grain, O. Lopez, G. Santarelli, and C. Chardonnet. Absolute frequency measurement in the 28-thz spectral region with a femtosecond laser comb and a long-distance optical link to a primary standard. *Appl. Phys. B: Lasers Opt.*, 78(1):25–30, 2004.
- [138] M. A. Gubin, D. A. Tyurikov, A. S. Shelkovnikov, E. V. Koval'chuk, G. Kramer, and B. Lipphardt. Transportable he-ne/ch₄ optical frequency standard and absolute measurements of its frequency. *IEEE J. Quantum Electron.*, 31(12):2177–2182, 1995.
- [139] N. G. Basov and M. A. Gubin. Quantum frequency standards. *IEEE J. Sel. Top. Quantum Electron.*, 6:857–868, 2000.
- [140] E. V. Kovalchuk, D. Dekorsy, A. I. Lvovsky, C. Braxmaier, J. Mlynek, A. Peters, and S. Schiller. High-resolution doppler-free molecular spectroscopy with a continuous-wave optical parametric oscillator. *Opt. Lett.*, 26(18):1430–1432, 2001.
- [141] S. M. Foreman, K. W. Holman, D. D. Hudson, D. J. Jones, and J. Ye. Remote transfer of ultrastable frequency references via fiber networks. *Rev. Sci. Instr.*, 78:021101–1 – 021101–25, 2007.

- [142] D. Lee and N. C. Wong. Tunable optical frequency division using a phase-locked optical parametric oscillator. Opt. Lett., 17(1):13–15, 1992.
- [143] D. D. Krylova, A. S. Shelkovnikov, E. A. Petrukhin, and M. A. Gubin. Effect of intracavity back reflections on the accuracy and stability of optical frequency standards. Quantum Electron., 34(6):554–558, 2004.
- [144] Poseidon Scientific Instruments sapphire loaded cavity oscillator and Frequency Electronics, Inc., FE-102A-100 and FE-205A.
- [145] M. Notcutt, L.-S. Ma, A. D. Ludlow, S. M. Foreman, J. Ye, and J. L. Hall. Contribution of thermal noise to frequency stability of rigid optical cavity via hertz-linewidth lasers. Phys. Rev. A, 73:031804R, 2006.
- [146] A. D. Ludlow, M. M. Boyd, T. Zelevinsky, S. M. Foreman, S. Blatt, M. Notcutt, T. Ido, and J. Ye. Systematic study of the ^{87}sr clock transition in an optical lattice. Phys. Rev. Lett., 96:033003, 2006.
- [147] M. M. Boyd, A. D. Ludlow, S. Blatt, S. M. Foreman, T. Ido, T. Zelevinsky, and J. Ye. ^{87}sr lattice clock with inaccuracy below 10^{-15} . Phys. Rev. Lett., 98(8):083002, 2006.
- [148] S. M. Foreman, A. D. Ludlow, M. H. G. de Miranda, J. E. Stalnaker, S. A. Diddams, and J. Ye. Coherent optical phase transfer over a 32-km fiber with 1-s instability at 10^{-17} . Phys. Rev. Lett., submitted.
- [149] U. Sterr, C. Degenhardt, H. Stoeckl, Ch. 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(8):845–855, 2004.
- [150] C. W. Oates, E. A. Curtis, and L. Hollberg. Improved short-term stability of optical frequency standards: approaching 1 hz in 1 s with the ca standard at 657 nm. Opt. Lett., 25(21):1603–1605, 2000.
- [151] 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(15):153001, 2005.
- [152] M. Takamoto, F.-L. Hong, R. Higashi, and H. Katori. An optical lattice clock. Nature, 435(7040):321–324, 2005.
- [153] S. G. Porsev and A. Derevianko. Hyperfine quenching of the metastable $^3p_{0,2}$ states in divalent atoms. Phys. Rev. A, 69:042506, 2004.
- [154] R. Santra, K. V. Christ, and C. H. Greene. Properties of metastable alkaline-earth-metal atoms calculated using an accurate effective core potential. Phys. Rev. A, 69:042510, 2004.
- [155] T. Ido and H. Katori. Recoil-free spectroscopy of neutral sr atoms in the lambda-dicke regime. Phys. Rev. Lett., 91:053001, 2003.

- [156] H. Katori, T. Ido, and M. Kuwata-Gonokami. Optimal design of dipole potentials for efficient loading of sr atoms. J. Phys. Soc. Japan, 68:2479, 1999.
- [157] J. Ye, D. W. Vernooy, and H. J. Kimble. Real time tracking and trapping of single atoms in cavity qed. Phys. Rev. Lett., 83(24):4987–4990, 1999.
- [158] H. J. Kimble, C. J. Hood, T. W. Lynn, H. Mabuchi, D. W. Vernooy, and J. Ye. In R. Blatt, J. Eschner, D. Leibfried, and F. Schmidt-Kaler, editors, Laser Spectroscopy XIV, Proceedings of the International Laser Spectroscopy Conference 1999, Proceedings of ICOLS'99, page 80. World Scientific, Singapore, 1999.
- [159] I. Courty, R. P. Quessada, P. Kovacich, A. Brusch, D. Kolker, J. Zondy, G. D. Rovera, and P. Lemonde. Clock transition for a future optical frequency standard with trapped atoms. Phys. Rev. A, 68:030501(R), 2003.
- [160] A. Brusch, R. Le Targat, X. Baillard, M. Fouche, and P. Lemonde. Hyperpolarizability effects in a sr optical lattice clock. Phys. Rev. Lett., 96(10):103003, 2006.
- [161] T. Hong, C. Cramer, E. Cook, W. Nagourney, and E. N. Fortson. Studies of the 1s_0 - 3p_0 transition in atomic ytterbium for optical clocks and qubit arrays. Opt. Lett., 30:2644–2646, 2005.
- [162] Z. W. Barber, C. W. Hoyt, C. W. Oates, L. Hollberg, A. V. Taichenachev, and V. I. Yudin. Direct excitation of the forbidden clock transition in neutral yb-174 atoms confined to an optical lattice. Phys. Rev. Lett., 96(8):083002, 2006.
- [163] 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(7):073003, 2004.
- [164] T. H. Loftus, T. Ido, M. M. Boyd, A. D. Ludlow, and J. Ye. Narrow line cooling and momentum-space crystals. Phys. Rev. A, 70:063413, 2004.
- [165] Lemonde et. al., ESA Worskshop, March (2007).
- [166] M. Takamoto, F.-L Hong, R. Higashi, Y. Fujii, M. Imae, and H. Katori. Improved frequency measurement of a one-dimensional optical lattice clock with a spin-polarized fermionic ^{87}sr isotope. J. Phys. Soc. Japan, 75:104302, 2006.
- [167] K. Numata, A. Kemery, and J. Camp. Thermal-noise limit in the frequency stabilization of lasers with rigid cavities. Phys. Rev. Lett., 93:250602, 2004.
- [168] T. M. Fortier, A. Bartels, and S. A. Diddams. Octave-spanning ti:sapphire laser with a repetition rate > 1 ghz for optical frequency measurements and comparisons. Opt. Lett., 31(7):1011–1013, 2006.
- [169] J. Stenger, H. Schnatz, C. Tamm, and H. R. Telle. Ultraprecise measurement of optical frequency ratios. Phys. Rev. Lett., 88(7):073601, 2002.
- [170] A. Bartels, C. W. Oates, L. Hollberg, and S. A. Diddams. Stabilization of femtosecond laser frequency combs with subhertz residual linewidths. Opt. Lett., 29(10):1081–1083, 2004.

- [171] R. J. Rafac, B. C. Young, J. A. Beall, W. M. Itano, D. J. Wineland, and J. C. Bergquist. Sub-dekahertz ultraviolet spectroscopy of $^{199}\text{Hg}^+$. Phys. Rev. Lett., 85(12):2462–2465, 2000.
- [172] P. O. Schmidt, T. Rosenband, C. Langer, W. M. Itano, J. C. Bergquist, and D. J. Wineland. Spectroscopy using quantum logic. Science, 309(5735):749–752, 2005.
- [173] H. S. Margolis, G. P. Barwood, G. Huang, H. A. Klein, S. N. Lea, K. Szymaniec, and P. Gill. Hertz-level measurement of the optical clock frequency in a single sr-88(+) ion. Science, 306(5700):1355–1358, 2004.
- [174] T. Schneider, E. Peik, and C. Tamm. Sub-hertz optical frequency comparisons between two trapped yb-171(+) ions. Phys. Rev. Lett., 94(23):230801, 2005.
- [175] P. Dube, A. A. Madej, J. E. Bernard, L. Marmet, J. S. Boulanger, and S. Cundy. Electric quadrupole shift cancellation in single-ion optical frequency standards. Phys. Rev. Lett., 95(3):033001, 2005.
- [176] F. Ruschewitz, J. L. Peng, H. Hinderthur, N. Schaffrath, K. Sengstock, and W. Ertmer. Sub-kilohertz optical spectroscopy with a time domain atom interferometer. Phys. Rev. Lett., 80(15):3173–3176, 1998.
- [177] G. Wilpers, T. Binnewies, C. Degenhardt, U. Sterr, J. Helmcke, and F. Riehle. Optical clock with ultracold neutral atoms. Phys. Rev. Lett., 89(23):230801, 2002.
- [178] S. Bize, P. Laurent, M. Abgrall, H. Marion, I. Maksimovic, L. Cacciapuoti, J. Grunert, C. Vian, F. P. dos Santos, P. Rosenbusch, P. Lemonde, G. Santarelli, P. Wolf, A. Clairon, A. Luiten, M. Tobar, and C. Salomon. Cold atom clocks and applications. J. Phys. B: At. Mol. Opt. Phys., 38(9):S449–S468, 2005.
- [179] Hydrogen maser: ST-5, NIST Time and Frequency Division (T. Parker), Boulder.
- [180] J. D. Prestage, R. L. Tjoelker, and L. Maleki. Atomic clocks and variations of the fine-structure constant. Phys. Rev. Lett., 74(18):3511–3514, 1995.
- [181] M. Fischer, N. Kolachevsky, M. Zimmermann, R. Holzwarth, T. Udem, T. W. Hänsch, M. Abgrall, J. Grunert, I. Maksimovic, S. Bize, H. Marion, F. P. Dos Santos, P. Lemonde, G. Santarelli, P. Laurent, A. Clairon, C. Salomon, M. Haas, U. D. Jentschura, and C.H. Keitel. New limits on the drift of fundamental constants from laboratory measurements. Phys. Rev. Lett., 92(23):230802, 2004.
- [182] E. R. Hudson, H. J. Lewandowski, B. C. Sawyer, and J. Ye. Cold molecule spectroscopy for constraining the evolution of the fine structure constant. Phys. Rev. Lett., 96(14):143004, 2006.
- [183] J. Ye, S. T. Cundiff, S. Foreman, T. M. Fortier, J. L. Hall, K. W. Holman, D. J. Jones, J. D. Jost, H. C. Kapteyn, Kahv Leeuwen, L.-S. Ma, M. M. Murnane, J. L. Peng, and R. K. Shelton. Phase-coherent synthesis of optical frequencies and waveforms. Appl. Phys. B: Lasers Opt., 74:S27–S34, 2002.

- [184] 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, 2004 IEEE International Topical Meeting on Microwave Photonics Technical Digest, pages 201–204. Institute of Electrical and Electronics Engineers, Piscataway, New Jersey, 2004.
- [185] <http://www-ssrl.slac.stanford.edu/lcls>.
- [186] 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(13):1554–1556, 2004.
- [187] 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(10):1225–1227, 2005.
- [188] K. W. Holman. Ph. D. Dissertation, University of Colorado, 2005.
- [189] J. C. Bergquist, W. M. Itano, and D. J. Wineland. In Hänsch T. W. and Inguscio M., editors, Frontiers in Laser Spectroscopy, Proceedings of the International School of Physics "Enrico Fermi": Course 120, "Enrico Fermi" Courses, pages 359–376. North Holland, Amsterdam, 1994.
- [190] C.-C. Chang and A. M. Weiner. Fiber transmission for sub-500-fs pulses using a dispersion-compensating fiber. IEEE J. Quantum Electron., 33(9):1455–1464, 1997.
- [191] C.-C. Chang, H. P. Sardesai, and A. M. Weiner. Dispersion-free fiber transmission for femtosecond pulses by use of a dispersion-compensating fiber and a programmable pulse shaper. Opt. Lett., 23(4):283–285, 1998.
- [192] O. Lopez, C. Daussy, A. Amy-Klein, C. Chardonnet, F. Narbonneau, M. Lours, and G. Santarelli. Fiber frequency dissemination with resolution below 10^{-17} . In International Workshop on Advances In Precision Tests and Experimental Gravitation in Space, Galileo Galilei Institute, Firenze, Italy, pages 14–16, Sep. 28, 2006.
- [193] P. Ciprut, N. Gisin, R. Passy, J. P. Von der Weid, F. Prieto, and C. W. Zimmer. Second-order polarization mode dispersion: impact on analog and digital transmissions. J. Lightwave Technol., 16(5):757–771, 1998.
- [194] 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(20):203904, 2005.
- [195] Scott Diddams' and Leo Hollberg's laboratories, Time and Frequency Division, NIST.
- [196] <http://www.branfiber.net>.
- [197] J. Kitching made the initial setup of this system.

- [198] E. N. Ivanov. design provided by S. A. Diddams.
- [199] R. Le Targat, X. Baillard, M. Fouché, A. Brusch, O. Tcherbakoff, G. D. Rovera, and P. Lemonde. Accurate optical lattice clock with ^{87}Sr atoms. Phys. Rev. Lett., 97:130801, 2006.
- [200] M. Takamoto, F.-L. Hong, R. Higashi, Y. Fujii, M. Imae, and H. Katori. Improved frequency measurement of a one-dimensional optical lattice clock with a spin-polarized fermionic ^{87}Sr isotope. J. Phys. Soc. Japan, 75(10):104302, 2006.
- [201] T. P. Krisher, L. Maleki, G. F. Lutes, L. E. Primas, R. T. Logan, J. D. Anderson, and C. M. Will. Test of the isotropy of the one-way speed of light using hydrogen-maser frequency standards. Phys. Rev. D, 42(2):731–734, 1990.
- [202] M. Calhoun, R. Sydnor, and W. Diener. A stabilized 100-megahertz and 1-gigahertz reference frequency distribution for cassini radio science. In The Interplanetary Network Progress Report, number 48-142, 15 February 2002.
- [203] 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, Proceedings of the 2002 IEEE International Frequency Control Symposium and PDA Exhibition, pages 510–516. Institute of Electrical and Electronics Engineers, Piscataway, New Jersey, 2002.
- [204] 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., 30:2948, 2005.
- [205] D. Kleppner. Time too good to be true. Phys. Today, 59(3):10–11, 2006.

Appendix A

A humorous note on units

Many interesting and varied units have been used in various pursuits throughout time. These include the crocodile, pony, barn, lumen, becquerel, sievert, katal, candela, dalton, tonne (not the ton), knot, SWU, fathom, palm, shaftment, CNRHK (Chuck-Norris-Round-House Kick), cubit, pennyweight, dram, hogshead, bushel, peck, tarefa, quincena, and foe, just to name a few fun examples. This Thesis wouldn't be a quality work of physics without at least one appendix inventing a new problem with units and then "solving" it.

Sometimes units are mis-used by physicists simply for the sake of making things easier to say. A prime example of this in the field of ultrafast science is the $[\text{fs}^2]$ unit for measuring the group-delay dispersion of an optical system. In reality, the unit should be quoted as $[\text{fs}/\text{PHz}]$, which would make it clear that one is speaking about a certain number of femtoseconds of group delay per PetaHertz of optical bandwidth, as measured near a particular optical frequency. Of course, since the "per PetaHertz" part of this would also be equivalent to fs, and since "fs-squared" is easier to say, we adopt the latter terminology.

Inspired by this kind of creativity, and finding myself always in the position of measuring frequency drift rates of optical clock oscillators, I have also come up with a new twist on a unit that is somewhat painful to say. Often, the ^{87}Sr clock laser described in Section 7.1.3 would drift by an amount that could be measured and compensated

for as described in Section 7.4.2. The drift rate for the first cavity made from Zerodur was on the order of several Hz of frequency shift per second. Of course, the natural SI unit would therefore be [Hz/s], but this is egregiously cumbersome to enunciate, much like this sentence. While watching my automated clockwork for hours on end one day, I realized that the “per second” part of the unit also corresponds to “Hz,” and that therefore the unit could be described as “Hz-squared.” However, it becomes even easier to say “square-Hz,” which of course led me to a rare “Eureka!” moment: I introduce to you, the reader, the new unit of “Squirtz” – appropriately denoted as [Sqz], which is untaken by any other unit I can find. Obviously it becomes great fun to speak of “milliSquirtz” when the drift rate is as well-cancelled as it can be for the new ULE cavities and their lower thermal noise limit.

So, the next time you consider using the official SI unit for frequency drift, please consider using the official SF unit instead – long live the Sqz!