Chapter 1

Introduction to Optical Frequency Standards and Clocks

1.1 History of timekeeping: Interplay between science and clock technology

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

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

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

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

1.2 Redefining time

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

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

1.3 Motivations for precise timing

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

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

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

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

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

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

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

1.4 Optical frequency standards

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

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

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

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

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

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

1.5 Gears for optical atomic clock

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

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

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

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

1.6 Making optical frequency standards accessible: Frequency transfer

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

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

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

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

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

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

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

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