Chapter 2

Fundamentals of Stabilizing an Optical Frequency Comb

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

2.1 Free parameters of an optical frequency comb

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

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

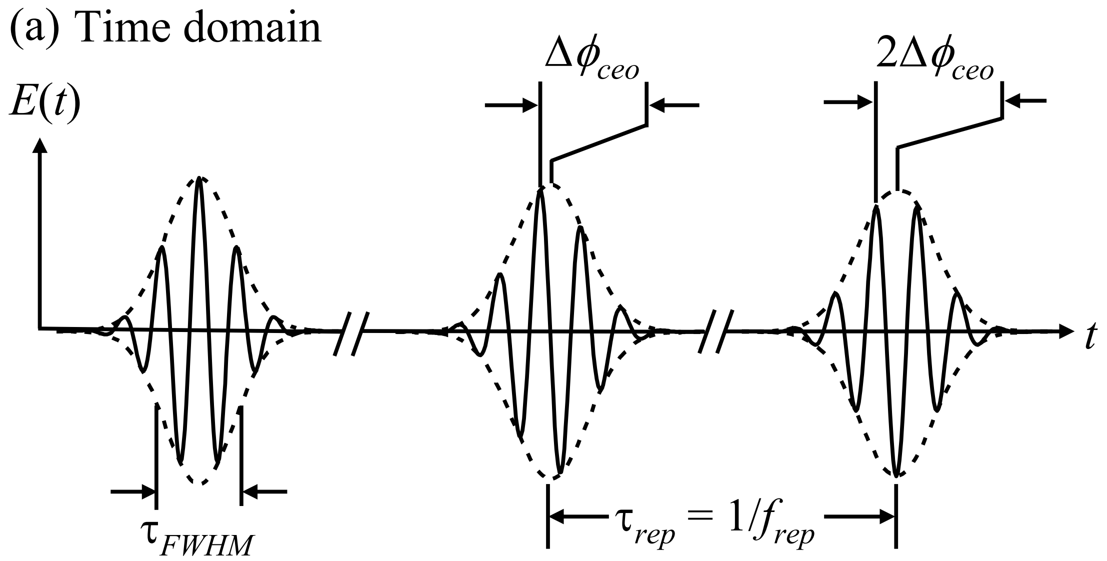
(a) Time domain

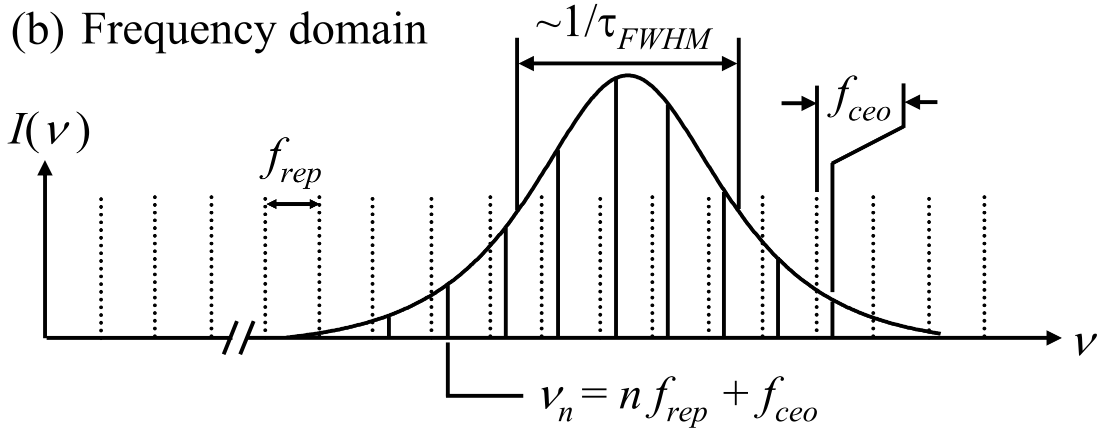
(b) Frequency domain

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

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

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

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

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

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

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

where ω_c is the angular frequency of the carrier.

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

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

which can be rearranged to yield

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

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

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

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

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

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

2.2 Stabilization of carrier-envelope offset frequency

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

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

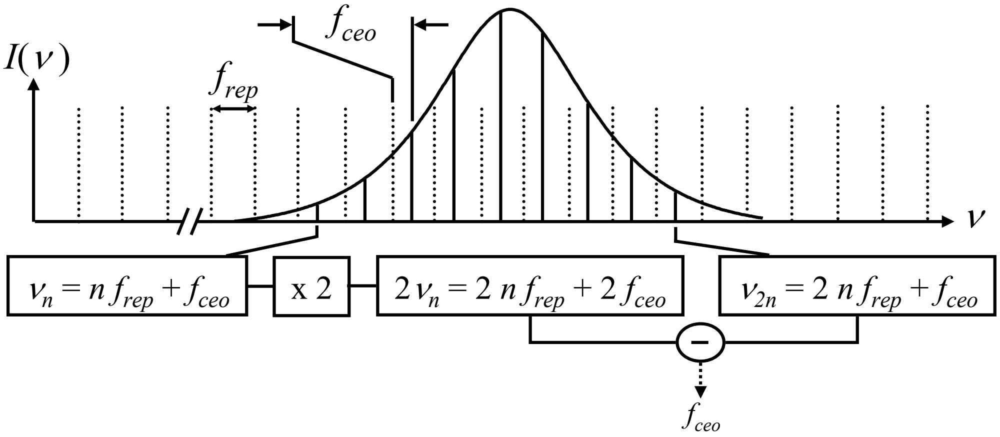

Figure 2.2: The carrier-envelope offset frequency is found by taking the difference of the frequencies of a comb mode at ν_{2n} and the second harmonic of a mode at ν_n.

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

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

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

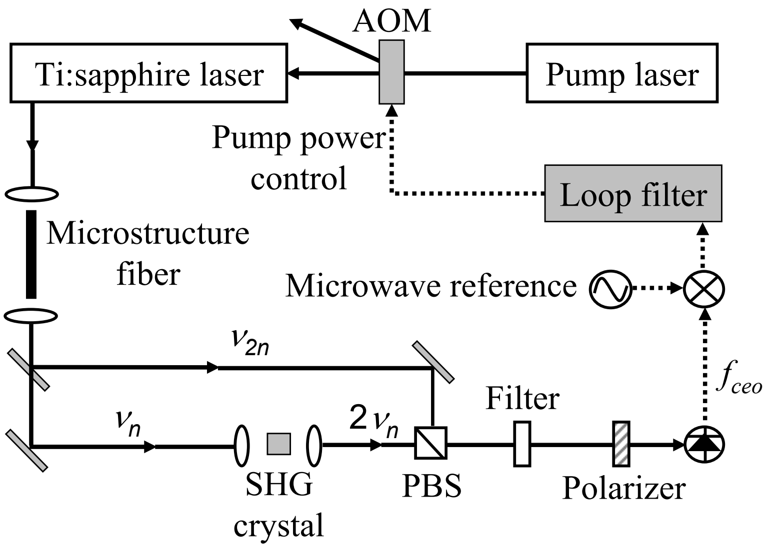

Figure 2.3: The carrier-envelope offset frequency is measured by combining the comb component at ν_{2n} with the second harmonic of the component at ν_n on a photodetector, which produces the heterodyne beat between the frequencies. It is phase-locked to a microwave reference by actuating the pump power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

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

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

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

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

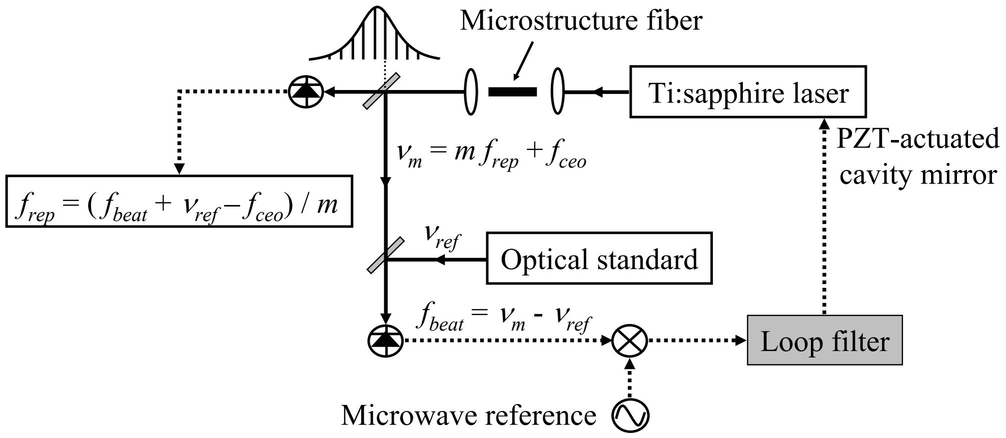

Figure 2.4: Once f_{ceo} is independently stabilized, the repetition frequency is stabilized to an optical standard by phase-locking the heterodyne beat between the optical standard and the nearest comb component to a microwave reference. It is controlled with a PZT-actuated mirror in the laser cavity.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

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

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

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

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

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

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

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

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

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

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

2.4 Transferring stability of one frequency comb to another

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

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

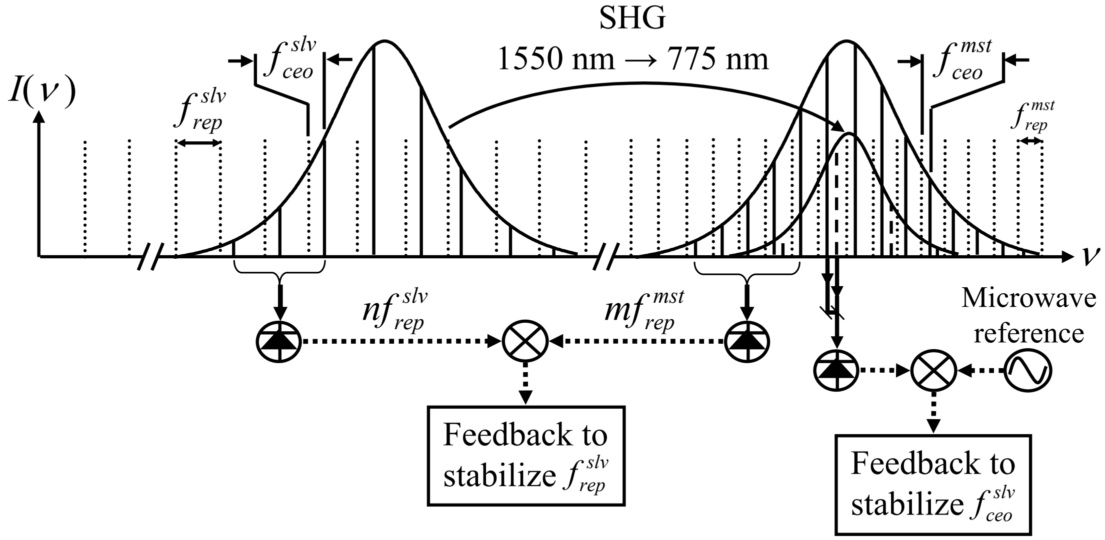

Figure 2.5: One optical frequency comb (the slave) can be stabilized to another (the master) by phase-locking harmonics of their repetition frequencies and by stabilizing the heterodyne beat between components from each comb to a microwave reference. For detection of the heterodyne beat between a 1550-nm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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Therefore,

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

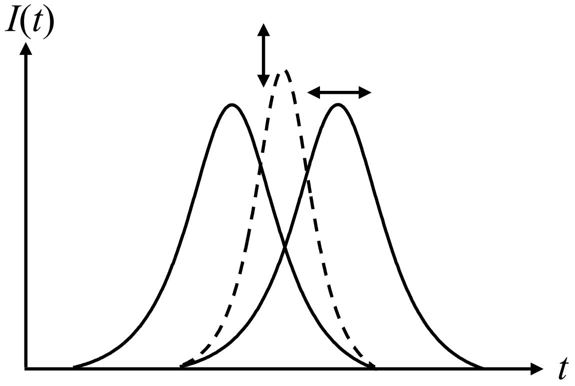

Figure 2.6: The timing jitter between the pulse trains from two mode-locked lasers can be detected by measuring their cross-correlation. The amplitude of the sum-frequency signal (dashed curve) produced by a nonlinear crystal is maximally sensitive to the relative timing jitter between the pulses (solid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

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