INVITED REVIEW ARTICLE

Remote transfer of ultrastable frequency references via fiber networks

Seth M. Foreman,[a] Kevin W. Holman,[b] Darren D. Hudson, David J. Jones,[c] and Jun Ye[d]

JILA, National Institute of Standards and Technology and University of Colorado, Boulder, Colorado 80309, and Department of Physics, University of Colorado, Boulder, Colorado 80309

(Received 20 June 2006; accepted 26 November 2006; published online 28 February 2007)

Three distinct techniques exist for distributing an ultrastable frequency reference over optical fibers. For the distribution of a microwave frequency reference, an amplitude-modulated continuous wave (cw) laser can be used. Over kilometer-scale lengths this approach provides an instability at 1 s of $\sim 3\times10^{-14}$ without stabilization of the fiber-induced noise and $\sim 1\times10^{-14}$ 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 3×10^{-15} with active noise cancellation [Recent results reduce the instability at 1 s to 6×10^{-18}.] Finally, microwave and optical frequency references can be simultaneously transmitted using an optical frequency comb, and we expect the optical transfer to be similar in performance to the cw optical frequency transfer. The instability at 1 s for transfer of a microwave frequency reference with the comb is $\sim 3\times10^{-14}$ without active noise cancellation and $<7\times10^{-15}$ with active stabilization. The comb can also distribute a microwave frequency reference with root-mean-square timing jitter below 16 fs integrated over the Nyquist bandwidth of the pulse train ($\sim$50 MHz) when high-bandwidth active noise cancellation is employed, which is important for remote synchronization applications. © *2007 American Institute of Physics.*
[DOI: 10.1063/1.2437069]

TABLE OF CONTENTS

I. INTRODUCTION

Current research on optical atomic clocks[1,2] 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.[3–15] 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

[a]Electronic mail: foreman@jila.colorado.edu
[b]Present address: Lincoln Laboratory, MIT, Lexington, MA 02420.
[c]Present address: University of British Columbia, Vancouver, British Columbia V6T 1Z1, Canada
[d]Electronic mail: ye@jila.colorado.edu

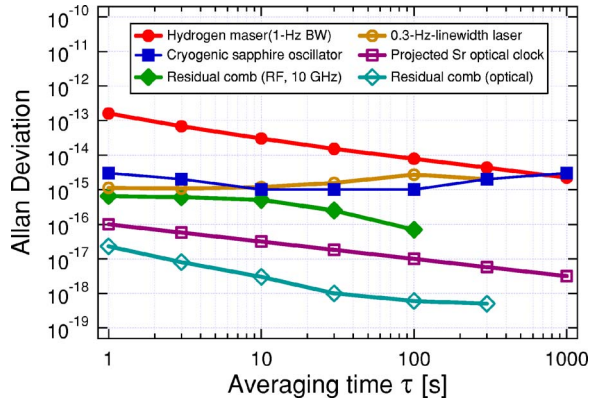

FIG. 1. (Color online) Fractional frequency instabilities of various high-stability microwave and optical frequency references. References are given in the text.

10^{16}.[16,17] However, frequency references based on optical transitions in laser-cooled and trapped atoms and ions are expected to eventually provide uncertainties approaching one part in 10^{18},[3] and short-term instabilities of a few parts in 10^{17} for a 1 s averaging time,[10] nearly three orders of magnitude better than the best microwave atomic clocks. The current performance of various high-stability oscillators is summarized in Fig. 1.

It shows the Allan deviation (a measure of the fractional frequency deviations of an oscillator as a function of averaging time, to be explained in detail in Sec. II A) of some of the most stable existing frequency references. In the microwave domain, a cryogenic sapphire oscillator[18] (closed squares) is shown, as well as a hydrogen maser[19] (closed circles) that is steered by a cesium-fountain clock. In the optical domain, a narrow 0.3 Hz linewidth laser[20,21] (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[22] (open squares). In addition, it has been shown that the excess instability associated with locking a mode-locked femtosecond laser-based optical frequency comb to an optical reference can achieve the level shown as open diamonds.[23] Finally, the rf signal provided by detecting the 10 GHz harmonic of a femtosecond laser's repetition rate has an excess instability[24] (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. 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.

The first 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,[25] such as the fine structure constant α.[26–31] 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

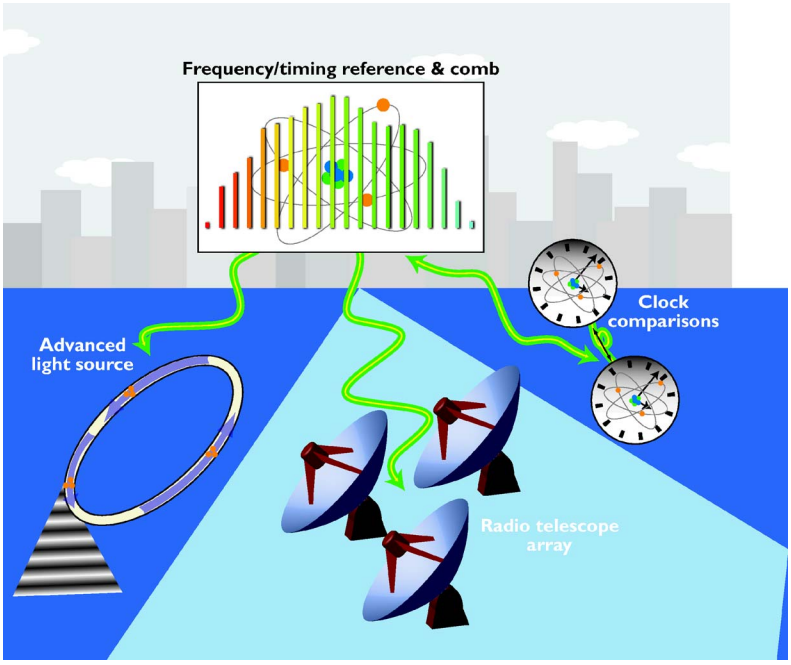

FIG. 2. (cover image) Ultrastable frequency/timing references are distributed through a kilometer-scale optical fiber network to facilities for several precision applications, including: atomic clock comparison and synchronization, also enabling mapping of the Earth's geoid; a long-baseline array of phase-coherent radio telescopes; and an accelerator-based advanced light source for generating x-ray pulses. Image courtesy of Jeff Fal, JILA.

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.

Comparisons of optical frequency standards 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 time scales 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 time scales. This is equivalent to saying that the transfer must exhibit ultralow timing jitter (approaching femtoseconds) over an extended Fourier frequency range.[32] One such application is long-baseline interferometry for radio astronomy.[33] 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.[33] 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.[34–37] 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.[35,37] 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 the common-view global positioning satellite system (GPS), which is the method used to compare the frequency and time standards of national laboratories around the world.[38] In this scheme, the transmitter and the receiver both compare their times simultaneously with that of a common GPS satellite that is in view of 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}.[39] 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 1 day.[40] 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 beneficial for a frequency distribution system. A great deal of the technology, components, and infrastructure also already exists for disseminating frequency references over telecommunication optical fibers, especially in urban environments as depicted in Fig. 2.

In this Review we discuss three distinct methods for transferring frequency references long distances over optical fibers for applications including those shown in Fig. 2. These methods represent the current state-of-the-art for high-stability frequency transfer. In Sec. III we 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 Sec. IV we will discuss the transfer of a microwave frequency reference using an amplitude-modulated continuous wave (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. In Sec. V 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. However, it has the drawback of not directly providing a timing signal for remote synchronization applications, since the modulation signal is the carrier itself, which is at hundreds of terahertz. An optical frequency comb must be employed to recover timing information. Finally, in Sec. VI we will discuss frequency transfer by directly transmitting an optical frequency comb, which is produced by the pulse train from a mode-locked laser.[41–43] The evenly spaced pulses in time correspond to an array of discrete frequency components that are uniformly spaced by the pulse repetition frequency. The frequency comb can be stabilized to either a microwave or optical frequency reference, and after trans-

mission it simultaneously provides optical frequency references (by detecting the individual comb components) spanning the spectrum and many microwave frequency references (via harmonics of the comb spacing or pulse repetition frequency). More details of the comb and its stabilization will also be provided in Sec. VI. However, before we discuss various transfer techniques, it is important to understand how to measure and characterize frequency stability and phase noise spectral density. This is the topic of the next section.

II. CHARACTERIZATION AND MEASUREMENT OF FREQUENCY STABILITY AND PHASE NOISE

There are typically two classes of characterization of the instability of a frequency source. The first involves determining how the measured fractional frequency fluctuations of the source vary as a function of the time over which the frequency is averaged, and is commonly expressed using the Allan deviation. The second method for characterizing frequency stability is to measure the phase noise spectrum, which represents the underlying frequency modulation process over a range of relevant Fourier frequencies surrounding the signal carrier. For analyzing the remote transfer of a microwave frequency reference, it is common to express this 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 a signal regardless of its frequency value. 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. Thorough discussions of the Allan deviation and phase spectral density can be found in many excellent references.[44–48] Here, we give a brief introduction to these methods for characterizing the stability and purity of precision frequency sources, as a useful start for a less-experienced reader.

A. Allan deviation

The Allan deviation,[49] $\sigma_y(\tau)$, 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 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. The signal from the frequency source can be expressed as

$$V(t) = V_0 \cos[2\pi\nu_0 t + \phi(t)], \tag{1}$$

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 fractional frequency deviation from the nominal center frequency is given by

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

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

$$\sigma_y(\tau) \equiv \left\langle \frac{1}{2}[\bar{y}(t+\tau) - \bar{y}(t)]^2 \right\rangle^{1/2}, \tag{3}$$

where $\langle\ \rangle$ indicates an infinite time average and $\bar{y}$ represents the time average of $y(t)$ over a period τ.[47] $\sigma_y(\tau)$ 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,

$$\sigma_y(\tau) \approx \left[\frac{1}{2(N-1)\nu_0^2} \sum_{i=1}^{N-1} (\bar{\nu}_i - \bar{\nu}_{i+1})^2 \right]^{1/2}. \tag{4}$$

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 subdivided into adjacent, nonoverlapping subsets. Each value of the new set of frequency values is computed by averaging the m values in each subset of the original data. When using this technique, estimation of the uncertainty associated with a given τ is not rigorous.[48] For the purposes of this Review, whenever an Allan deviation is shown without error bars, the total amount of frequency-counting data taken was ≈4 times as much as the longest τ shown.

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.[47] 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 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.

B. Phase noise and timing jitter

The second method for characterizing the stability of a frequency source—directly measuring its phase fluctuations—is especially useful for determining the short-term stability of a signal that is to be used for synchronization of various system components. Phase noise is determined by measuring the phase fluctuations of a signal with respect to a lower phase noise frequency reference.

The phase noise of a frequency source creates noise sidebands on the carrier frequency that are spaced from the carrier by the frequency at which the noise occurs. This weakens 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-squared phase fluctuation at Fourier frequency f from the carrier in a measurement bandwidth of 1 Hz.[47] It is defined as

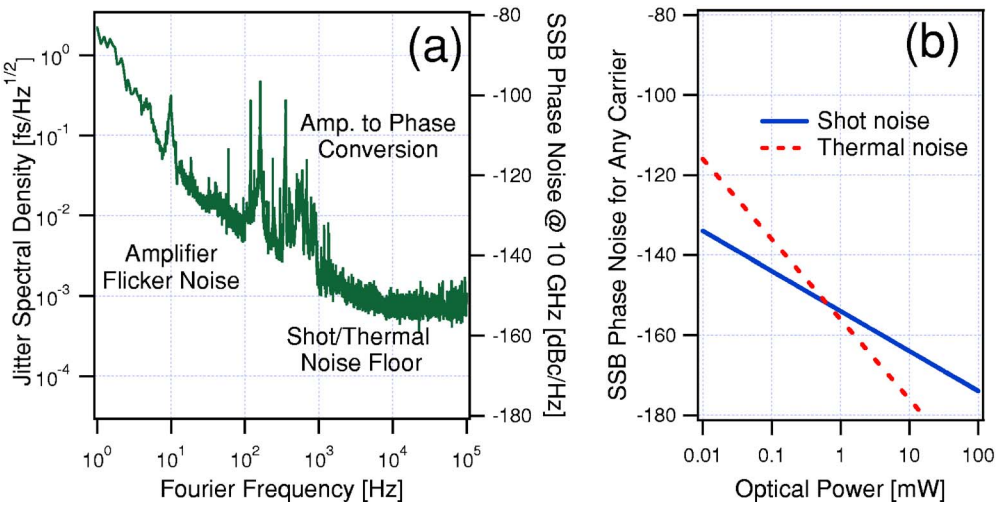

FIG. 3. (Color online) (a) Phase noise spectrum for ~2 mW of optical power from a femtosecond laser split equally onto identical photodetectors, and subsequently mixed to measure the excess phase noise in detection of the 10 GHz harmonics of the laser's repetition rate. For $f<1$ kHz, amplifier flicker 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.

$$S_\phi(f) \equiv [\delta\tilde{\phi}(f)]^2 \ (\mathrm{rad}^2/\mathrm{Hz}). \tag{5}$$

Note that $S_\phi(f)$ includes contributions from both the upper and lower sidebands of the carrier, since the Fourier transform 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-sideband phase noise, $L(f)$, defined as $\frac{1}{2}S_\phi(f)$.[47] It is usually written in units of dBc/Hz (decibels below the carrier in a 1 Hz measurement bandwidth) and can be logarithmically expressed as

$$L(f) = 10 \log\left[\frac{1}{2}S_\phi(f)\right] \ (\mathrm{dBc/Hz}). \tag{6}$$

In some applications, such as remote synchronization, it is more appropriate to express the phase noise of a transmitted signal in terms of its timing jitter. The timing jitter spectral density, $\delta\tilde{T}(f)$, which represents the root-mean-square (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

$$\delta\tilde{T}(f) = \frac{\delta\tilde{\phi}(f)}{2\pi\nu_0} \ (\mathrm{s}/\sqrt{\mathrm{Hz}}). \tag{7}$$

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

$$T_{\mathrm{rms}} = \sqrt{\int_{f_l}^{f_h} [\delta\tilde{T}(f)]^2 \, df} \ (\mathrm{s}). \tag{8}$$

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 time scales, 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 jitter spectral density within the locking bandwidth matters. Outside the locking bandwidth, the free-running noise of the oscillator is 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 Sec. II C) of the detection process. Therefore, the locking bandwidth should be limited to the crossover frequency where 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.

C. Fundamental and technical limitations

The techniques for stable frequency transfer described in Secs. IV–VI rely on the ability to detect and subsequently cancel, with a high degree of precision, the phase noise accumulated over the transmission paths. 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 photodetection, 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 processes are evident in the data of Fig. 3(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 com-

pared 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.

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. 50; here, we 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 resistance is given by

$$\mathrm{L}_{\phi}^{\mathrm{thermal}}(f) = \frac{kTR}{2V_0^2} \propto \frac{1}{P_{\mathrm{rf}}} \propto \frac{1}{P_{\mathrm{opt}}^2}, \tag{9}$$

where k is Boltzmann's constant, T is the system's temperature, R is the characteristic system impedance, V_0 is the rms voltage level of the carrier signal, P_{rf} is the rf carrier power, and P_{opt} is the average 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[51]

$$\mathrm{L}_{\phi}^{\mathrm{thermal}}(f) = (-177 - P_{\mathrm{rf}})\ \ (\mathrm{dBc/Hz}), \tag{10}$$

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 photodetector, is given by

$$\mathrm{L}_{\phi}^{\mathrm{shot}}(f) = \frac{ei_{\mathrm{avg}}R}{P_{\mathrm{rf}}} \propto \frac{P_{\mathrm{opt}}}{P_{\mathrm{rf}}} \propto \frac{1}{P_{\mathrm{opt}}}, \tag{11}$$

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,[52] but Eq. (11) still yields results that are empirically valid for time averages of the photocurrent and shot noise power.[53] Note that since P_{rf} scales as the square of P_{opt}, $\mathrm{L}_{\phi}^{\mathrm{shot}}$ decreases as the inverse of P_{opt}, whereas $\mathrm{L}_{\phi}^{\mathrm{thermal}}$ decreases as the inverse *square* of P_{opt}. This results in the shot noise floor overtaking the thermal noise floor beyond a certain optical power, as shown in Fig. 3(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 Sec. II D) 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 will have $P_{\mathrm{rf}}=-21$ 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. 3(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/$\sqrt{\mathrm{Hz}}$, 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[46,54] as Allan deviation

$$\sigma_y(\tau) = \frac{\sqrt{3\mathrm{S}_{\phi}\Delta f}}{2\pi\nu_0\tau}, \tag{12}$$

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 $7\times10^{-17}\ \tau^{-1}$ level. It is important to remember that this fractional frequency instability is relative to the 10 GHz rf signal: in the case where the signal is derived from an optical heterodyne beat between, for example, two 100 THz optical signals, then the optical frequencies have a relative fractional instability of 10^4 better than that of the 10 GHz rf beat. Finally, note that in the case of fiber transfer, mechanical perturbations to the transfer 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 carrier frequency.

In addition to the fundamental limitations discussed so far, two important technical limitations remain: excess amplifier noise[55] and amplitude-to-phase conversion.[56] 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 milliwatts) optical powers.[57,58]

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.[55] Typical flicker noise is seen in the data shown in Fig. 3(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[59] or feedback[60] 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[61] (carrier suppression) have been used to operate amplifiers in the weak signal regime by interferometrically canceling 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.[62] Applying this technique to a system employing typical photodetection of an optical frequency comb, microwave frequen-

cies were synthesized from optical frequencies with −110 dBc/Hz of phase noise at a 1 Hz offset from the 10 GHz carrier.[53] Excellent discussion of design considerations for these microwave interferometric techniques can be found in Ref. 63.

The final technical limitation we discuss here concerns applying higher optical power to the photodetector. Increased incident power is desirable since shot noise, thermal noise, and amplifier noise limits can all be reduced when higher microwave signal power becomes available. However, typical photodetectors exhibit a power-dependent phase shift between the incident 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. 3(a), the sharp features of the optical beam's power spectrum have been converted to sharp features in the phase spectrum of the converted microwave signal. 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 photogenerated charge carriers by screening of the applied bias field.[64] 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.[65] For femtosecond pulse applications typical power-to-phase conversion coefficients are on the order of several picoseconds of time delay per milliwatt 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.[66] By carefully eliminating environmental perturbations that cause amplitude fluctuations at the photodetector, optical-to-microwave conversion exhibiting an excess instability of $<7\times10^{-16}$ at 1 s and phase noise below −98 dBc/Hz at 1 Hz offset from the 10 GHz carrier has been achieved.[24] It should be noted that a disadvantage of photodetecting a pulse train for phase stabilization 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, we 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.[67]

D. Optical techniques for phase detection

In a 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 in the optical domain.[68–71] 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 (SFG) signal. The pulse trains can either be copropagated through the crystal, or they can be overlapped in a crossed geometry. The latter configuration results in a SFG signal that exits the crystal in a different direction than the original pulse trains. 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. 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 homodyne rf mixing schemes limited by the noise processes discussed previously.

An interferometric (collinear) cross-correlation measurement can provide an even higher degree of sensitivity. In this case the intensity of the cross correlation oscillates as a function of delay between the two pulses at a frequency equal to the pulses' optical carrier frequency.[72] Thus, adopting the notation in Eq. (7), ν_0 is on the order of ~200 THz, leading to a substantially increased sensitivity for the jitter measurement. A further advantage of this measurement method over the noncollinear cross correlation is that the two pulses need not have perfect chirp compensation as the overall slope of the correlation envelope is unimportant. As described in Ref. 73 and in Sec. VI E, this method of jitter measurement was used to measure the remote locking of a mode-locked fiber laser operating at 1550 nm. Using an all-optical locking technique (shown in Sec. VI E as Fig. 20), an interferometric collinear cross correlation was performed with two-photon absorption in a silicon avalanche photodiode (Si-APD),[74] and was able to measure an out-of-loop jitter below 0.040 fs (1 Hz to 10 MHz) with a noise floor of 0.008 fs over the same frequency range.

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[75] in order to avoid amplitude-to-phase conversion; considerations for the design of such high-power photodetectors are complex.[76] The need for highly linear, high-current photodetectors is clear, since current performance levels for transferring frequency references over fiber networks are rapidly approaching the fundamental shot noise limit associated with milliwatt-level optical powers. When future optical clocks deliver the promised instability of $<1\times10^{-16}$ at 1 s, we will urgently need such advances to enable accurate transfer. Finally, whenever possible, higher order harmonics of the microwave signal, or even optical-based phase 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.

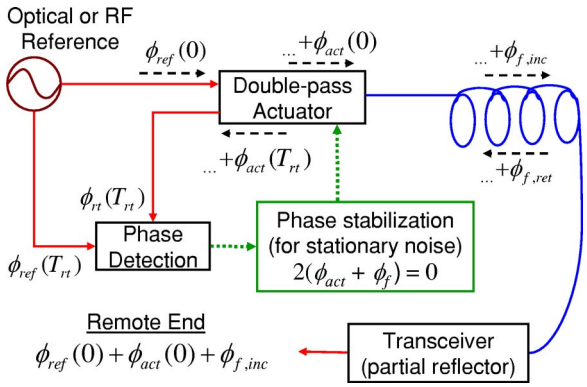

FIG. 4. (Color online) Principle of operation for a typical actively compensated fiber network for distributing frequency and timing references. A partial retroreflector or transceiver at the remote end returns the signal through a round trip of both the fiber network and the actuator. Accurate comparison 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 time scales 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.

III. 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 canceled 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.

A. Round-trip phase detection

Figure 4 shows the basic experimental scheme for fiber network phase stabilization. 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.[77] 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

$$\phi_{rt}(T_{rt}) - \phi_{ref}(T_{rt}) = \phi_{ref}(0) - \phi_{ref}(T_{rt}) + \phi_{act}(0) + \phi_{act}(T_{rt}) + \phi_{f,inc} + \phi_{f,ret}. \quad (13)$$

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 Eq. (13). 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 Eq. (13) are equal, and can be written as $2\phi_f$. Under this assumption, only noise processes with frequencies less than $1/2\pi T_{rt}$ 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 Eq. (13) 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 time scale longer than T_{rt}, and reference's coherence time longer than T_{rt}), Eq. (13) can be simplified as

$$\phi_{rt} - \phi_{ref} = 2(\phi_{act} + \phi_f), \quad (14)$$

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 Eq. (14) to a constant, the remote user will have a phase-coherent copy of the reference, up to the round-trip-limited bandwidth of $1/2\pi T_{rt}$. 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.

B. 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.[78,79]

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 $\phi_{f,\mathrm{inc}}=\phi_{f,\mathrm{ret}}$ is violated even for noise processes that are stationary during T_{rt}. As will be discussed in Sec. IV C, scrambling the polarization at time scales faster than T_{rt} offers a method for quickly averaging out PMD effects for microwave transfer.[77,80] 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 (Ref. 81) 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 Sec. III A, the bandwidth of the phase-locking servo is limited to $1/2\pi T_{\mathrm{rt}}$. Any attempt to cancel noise at higher bandwidths will simply write excess noise onto the link, since phase fluctuations on faster time scales 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, servos with low dynamic range (such as the PZT fiber stretchers described in Secs. IV and VI) 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 Eq. (12), use of an optical rather than microwave 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 cleanup oscillator at the remote end with a rapid phase noise roll-off beyond the servo bandwidth: by locking the cleanup 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 cleanup oscillator's output.

C. 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 canceled 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. 5(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 of the fiber.

Applications involving comparison of optical frequency standards can benefit from directly transferring a stable optical reference, as shown in Fig. 5(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. 5(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

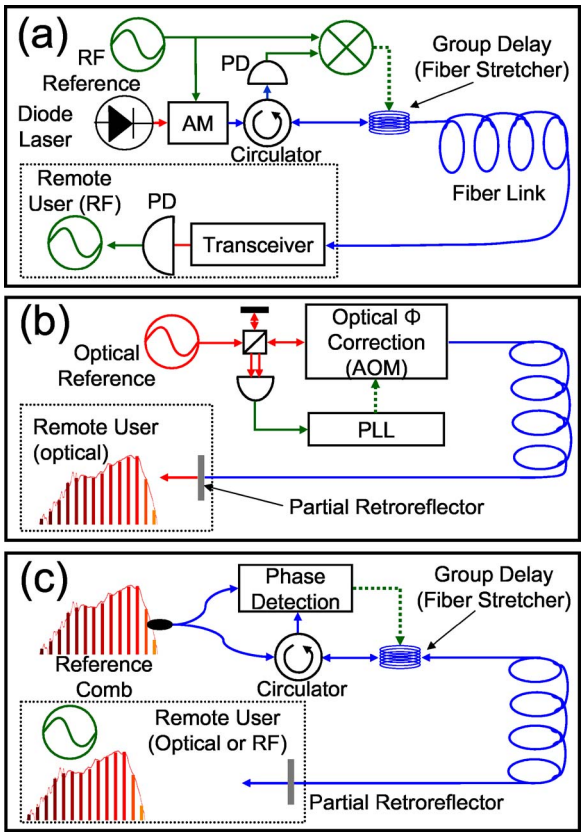

FIG. 5. (Color online) Various experimental schemes for transfer of frequency references. (a) Actively-stabilized transfer of a microwave reference by modulating a diode laser. (b) Direct transfer of the optical carrier by stabilizing the optical phase of the fiber link. (c) Transfer of simultaneous 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.

ultralow-jitter signal for timing to the remote user. In Sec. VI F we 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 Sec. III A 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 linewidth $\approx$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 reproduce a phase-coherent copy at the remote end. Therefore, especially for cw optical links as discussed in Sec. V, 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 colocating 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.[77,82] 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 (Sec. IV), optical carrier transfer (Sec. V), and frequency comb transfer (Sec. VI).

IV. MICROWAVE FREQUENCY TRANSFER WITH MODULATED cw LASER

A. 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. 5(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. 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 laboratories at JILA to the National Institute of Standards and Technology (NIST) Boulder,[83] 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.[84] 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 μm, and the modulation is performed external to the laser using a Mach-Zehnder intensity modulator.[85] In order to measure the fiber-induced frequency noise of the transmitted microwave signal[86] (typi-

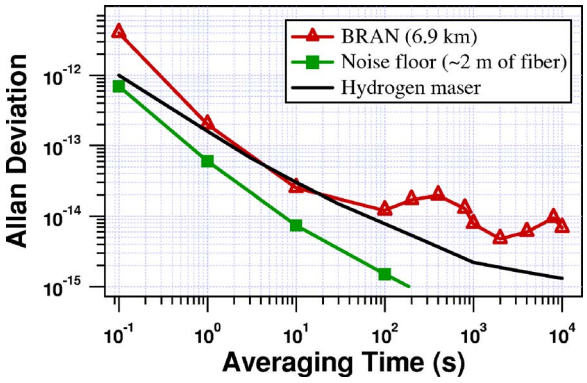

FIG. 6. (Color online) Measured instability of amplitude-modulated light transmitted through a short (~2 m) fiber (squares) and through a round trip of the BRAN fiber (open triangles). The rf modulation signal carried by the fiber-transmitted light is heterodyne detected against the original signal source, 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).

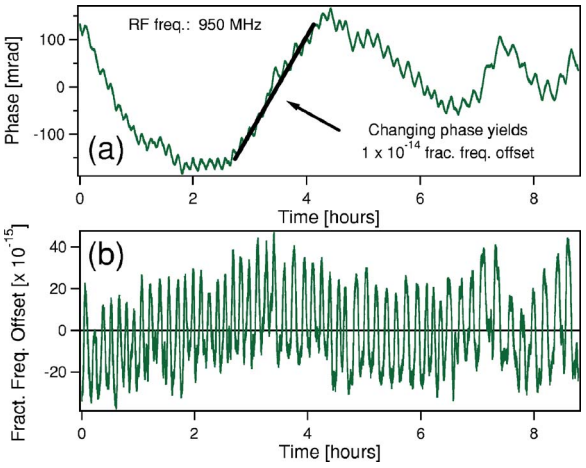

FIG. 7. (Color online) (a) Phase of the 7-km round-trip BRAN fiber for a transferred rf frequency of 950 MHz, showing long-term fractional frequency offsets of $\sim 10^{-14}$. (b) Time derivative of the phase data, with a rolling box average of 100 s applied, showing the fractional frequency offset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

cally operating at about 1 GHz), we first frequency shift a local copy of the reference signal at NIST by 10 kHz with a single-sideband generation microwave interferometer.[87] Then, we mix the round-trip signal with this shifted local copy, and count the frequency fluctuations of the resulting ~10 kHz beat signal, which is suitable for high-resolution frequency counting. The 10 kHz source used to generate the single sideband must have absolute frequency fluctuations lower than those we measure on the 10 kHz beat. When this condition is satisfied, measured fluctuations of the 10 kHz beat frequency are 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. 6. Also shown are results obtained when the BRAN fiber is 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 $1-2\times10^{-14}$ 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 7(a) shows this phase record over the course of a 9 h period with a 1 s gate time for the measurement. The steep long-term slope between 2.5 and 4 h corresponds to a fractional frequency offset of $\sim1\times10^{-14}$ 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. 7(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.

B. Application to precision spectroscopy

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. V, is the case of absolute frequency measurements of a Sr-based optical 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[16,19] from NIST to JILA for the initial work on absolute optical frequency measurements of clock transitions in neutral Sr atoms.

Large ensembles of ultracold alkaline earth atoms have provided impressive short-term clock stability.[9] So far, 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.[9,11,88] A promising approach is to explore the ultranarrow optical transitions of atoms held in an optical lattice.[12,89,90] The atoms are tightly localized so that Doppler and photon-recoil related effects on the transition frequency are eliminated.[91] 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.[92–94] 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 $^1S_0-^3P_0$ transition in the fermionic Sr isotope, ^{87}Sr.[12–14,95] Similar work on Yb is also in progress.[15,96]

Absolute frequency measurements were recently reported for the ^{87}Sr $^1S_0-^3P_0$ clock transition.[12,13] Reference

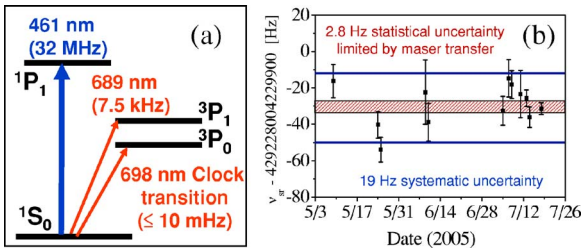

FIG. 8. (Color online) (a) Simplified energy level diagram for neutral Sr atoms, showing the two cooling transitions at 461 and 689 nm, and the clock transition at 698 nm. (b) Frequency-counting record for ^{87}Sr, showing the achieved 2.8 Hz statistical uncertainty (maser-transfer-limited) in the red-striped box and the 19 Hz systematic uncertainty as the outer blue lines.

12 used a GPS-based frequency reference with instability $\sim 10^{-10}$ at 1 s averaging time, while the JILA results[13] used a fiber-transferred reference. A troublesome 4σ disagreement between these two results was eventually resolved by a third independent investigation using an on-cite Cs-fountain reference.[97] A revised report by the Tokyo group[98] has since 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 Sec. IV A. 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.[13,99] An octave-spanning Ti:sapphire laser-based frequency comb[100] is phase-locked to the optical local oscillator for the ^{87}Sr clock,[21] as well as self-referenced using the standard f-to-2f scheme[101] as discussed in Sec. VI A. Then, the ~ 1 GHz harmonic of the Ti:sapphire laser's repetition rate is frequency counted against the rf reference received over the fiber network from NIST, in order to determine the absolute frequency of the ^{87}Sr $^1S_0 - ^3P_0$ optical clock transition. The Cs-calibrated, maser-based rf standard we transfer from NIST has an instability, as shown in Fig. 6, of $\sim 2.4 \times 10^{-13}$ at 1 s.[19] Passive transfer of this frequency from NIST to JILA is achieved by amplitude modulation as described in Sec. IV A, and reaches an instability $< 1 \times 10^{-14}$ for averaging times $> 10^3$ s as shown in Fig. 6. With the ability to average down to such a low fractional uncertainty in this relatively short time, 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 summarizes the measurement of the transition frequency within a period of 3 months. The statistical uncertainty of these measurements is 2.8 Hz, and systematic uncertainties are 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. 7, 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. For the current needs of ultrastable and accurate optical clocks, we have adopted active cancellation of the microwave phase drift that achieves instabilities below 10^{-16} for times $> 10^4$ s as described in the following section. Coupled with other improvements in the Sr system,[102] this approach has resulted in the improved absolute frequency measurement uncertainty of 1.1 Hz (2.5×10^{-15}).[99] An even higher-resolution approach using direct phase stabilization of the optical carrier will be discussed in Sec. V, and it will be required to connect the JILA Sr standard to the single ion- and neutral atom-based optical standards located at NIST.[4,15]

C. Active cancellation of the fiber-link-induced group-delay noise

It is clear from measurements of the excess phase noise introduced by fiber links (see Figs. 6 and 7) that, while the maser reference is not degraded for short time scales (< 100 s), the passive transfer is not suitable for attempting to transfer the Cs-referenced maser signal with long time scale instabilities below 5×10^{-15}. For this reason, several projects have employed active stabilization of the fiber-induced phase error for systems using rf modulation of an optical carrier.

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.[103,104] In that work, a stabilized fiber-optic distribution system was developed for microwave frequency dissemination. An optical carrier was amplitude modulated with the rf standard and sent to the remote site through the fiber link. Environmental perturbations on the fiber, such as vibration or temperature variation, introduced group delay variations and degraded the frequency stability of the rf standard. To actively eliminate the induced fiber noise, a signal was sent back to the transmitting site with a proper encoding to distinguish the round-trip signal from the input reference. The round-trip signal contained the fiber-induced noise and was used to derive an error signal to control the group delay of the transmitted path using a temperature-controlled fiber reel as the actuator. 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. The temperature-controlled fiber reel had a small bandwidth and was only able to reduce the instability over time scales longer than 100 s.

A similar technique is used for the transfer of a 100 MHz reference between Laboratoire de Physique des Lasers (LPL) and the Systèmes de Rèfèrence Temps Espace (SYRTE) for the comparison of an optical frequency at LPL and a microwave primary frequency standard at SYRTE.[77,82] Two parallel optical links, each 43 km, are used for bidirectional transmission of a 100 MHz reference signal by modulating the junction current of a distributed feedback semiconductor laser emitting at 1.55 μm. However, in contrast to NASA's Deep Space Network, actuators are present in both paths that allow higher bandwidth control. The two fiber links are in-

dependently stabilized by slightly different schemes.[77] One path electronically adjusts the phase of the modulation source to compensate for path length fluctuations, and the second path uses a piezoelectric-actuated fiber stretcher for high bandwidth control and a temperature-controlled spool of fiber for compensation of slow perturbations. The control bandwidth is about 300 Hz, limited by the round-trip delay (0.3 ms). The piezo-actuated link uses the stable signal it receives via the modulator-actuated link as its source, so that an out-of-loop comparison of the total performance of the bidirectional link can be achieved. This results 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 comparison, the instability in the absence of active noise cancellation is 3×10^{-14} at 1 s, and never goes below 1×10^{-15} for any time scale up to 10^5 s. A notable feature of the fiber links in Paris is the use of a slightly wavelength-shifted (by $\approx$4 nm) laser for the return signal. This allows optical filters to isolate the incident and return signals from false signals due to reflections in the fiber. It also assumes that the group delay dispersion fluctuations are negligible across the 4 nm span, an important result to keep in mind when considering the effects of group delay dispersion fluctuations for comb transfer in Sec. VI.

One important issue uncovered by the results in Ref. 77 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 $\phi_{\mathrm{rt}}=2\phi_{f,\mathrm{inc}}$ 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 has reduced the effects of PMD in a new system where the two 43 km links are 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 colocated. The polarization scrambling, in addition to using 1 GHz modulation frequencies to reduce the effects of SNR, has allowed achievement of transferred instabilities of 5×10^{-15} at 1 s and 2×10^{-18} after 1 day.[80]

In another experiment, time transfer over a 200 km optical fiber between Phoenix and Tucson, Arizona has been demonstrated with subpicosecond rms fluctuations over a 3.5 day period.[105] Similar to the previously described work, the optical carrier of a 1.55 μm laser diode was amplitude modulated with the 155.52 MHz rf signal. Based on the measured phase differences between a portion of the transmitted signal and a return signal, the overall time delay was actively controlled to compensate for fluctuations in optical path length of the fiber. Though the actuator time constant was only 10 s, the rms jitter was reduced to 630 fs over a 100 Hz bandwidth.

These efforts have been primarily focused on the transfer of a frequency reference for measurements that are averaged over an extended period of time. However, for some remote synchronization applications, the short-term (high-frequency) timing jitter of the transmitted microwave frequency reference is of fundamental importance. In links with low servo bandwidth, cleanup oscillators are often employed at the remote site in order to achieve the desired short-term frequency instability. However, when short-term coherence is desired between ends of the fiber link, higher-bandwidth servos must be used while obeying the strict bandwidth limit imposed by T_{rt}. In addition to applying active stabilization of the fiber link's group delay to the rf modulation technique described so far, at JILA we have adopted other techniques such as direct cw optical carrier transfer and microwave transfer with a frequency comb, using active stabilization of the fiber optical phase or the fiber group delay, respectively. In Sec. V we will show that, due to the improved timing resolution available when using a much higher reference frequency (optical instead of microwave), direct optical transfer achieves an instability below 1×10^{-17} at 1 s of averaging. In Sec. VI we show that using the high microwave carrier frequencies offered by the harmonics of a femtosecond laser's repetition rate allows very low timing jitter integrated over a large bandwidth, and that use of the optical frequency comb itself is expected to provide the same stability as direct cw optical transfer. In addition, the latest generation of optical clocks provides much more stable optical frequency references than existing microwave references, which can reduce the amount of averaging time needed to achieve a given level of precision by several orders of magnitude. Therefore, techniques which can directly transfer optical frequencies at high levels of precision (both low instability and low timing jitter in a large bandwidth) are being developed, as discussed in Secs. V and VI.

V. 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 Sec. VI A 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 canceled.

A. Fiber-optical phase noise cancellation

A noise-cancellation scheme similar to that shown in Fig. 5(b) and discussed in Sec. III C is used to cancel the fiber-induced phase noise introduced during the transmission of a cw optical frequency reference from JILA to

NIST.[86,106,107] 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 canceling 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. An important limitation of the optical heterodyne technique is that the use of non-polarization-maintaining fiber results in significant amplitude fluctuations of the optical heterodyne beat from which an error signal is derived. A tracking filter can be used to clean up this amplitude noise before deriving the error signal.

The performance of the active noise cancellation scheme is determined by measuring its effect on the round-trip phase noise.[86] We start by analyzing the line shape of the heterodyne beat signal between the original laser beam (before the AOM used for phase correction) and the returned light after double-passing through this AOM. For these experiments there are three separate time scales relative to T_{rt} that are of importance, as discussed in Sec. III A. The first is the time scale 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 time scale is the inverse of the servo bandwidth, which must be longer than T_{rt}. Another important time scale 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}. In this case, the line shape of the heterodyne beat takes the form of a spectral "delta" function modulated by phase delay fluctuations of the transfer fiber. However, when the delay time and the source's coherence time are comparable, the beat line shape takes the form of a spectral delta function sitting on top of a modified Lorentzian pedestal, since the fiber delay acts as a high-pass filter for the optical reference's intrinsic phase fluctuations. 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.

The experimental condition in Ref. 86 corresponds to the situation where both the fiber delay and the laser coherence time are $\sim 33\ \mu$s. Because of the phase noise accumulated during the 6.9 km round trip through the BRAN fiber, the central delta function component of the beat line shape is broadened to about 2 kHz. This effect is shown as the dotted curve in Fig. 9, where a fast Fourier transform (FFT) signal analyzer has been employed to reveal the linewidth of the heterodyne beat signal. The 50 kHz line center frequency of the signal is obtained by mixing the optical heterodyne beat

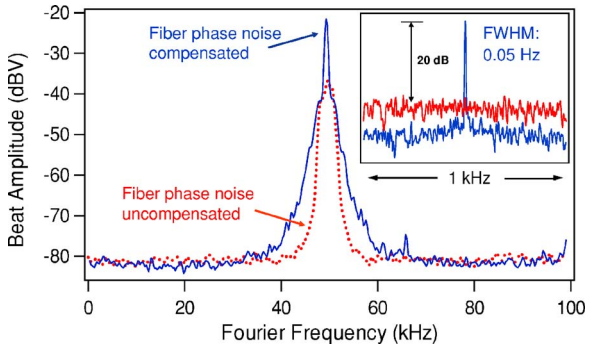

FIG. 9. (Color online) Fourier spectral analysis of the heterodyne beat linewidth between the original laser beam (before the AOM) and the returned light after a round trip through the fiber link. Without phase compensation, the beat linewidth is broadened by the fiber to ~2 kHz. The inset shows the beat signal with a much narrower 1 kHz span and a 0.048 Hz resolution bandwidth when the fiber noise cancellation is activated. This is to be contrasted with the white noise floor showing in the same spectral window when phase noise is not canceled. Notice the noise level is lowered by ~7 dB due to the emergence of the recovered carrier signal when cancellation is activated. (Adapted from Ref. 86.)

with an appropriately tuned auxiliary rf signal. This step of signal processing is necessary for an enhanced resolution in spectral analysis. When the fiber noise cancellation system is activated, we recover a narrow spectral peak on top of the pedestal, as shown in Fig. 9. Within the cancellation servo bandwidth of ~3 kHz, the round-trip and reference laser fields are mainly coherent, so the noise cancellation acts primarily on the fiber-induced noise. At higher Fourier frequencies, however, the round-trip and reference fields are incoherent at the detector. Therefore, the cancellation servo loop writes this phase incoherence back onto the outgoing signal. The result is the increased noise level at these higher Fourier frequencies, leading to formation of a pedestal. One can appreciate the effect of fiber noise cancellation from a high-resolution scan of the beat spectrum. As the inset in Fig. 9 indicates, when the beat signal was measured with a full span of only 1 kHz and a resolution bandwidth of 0.048 Hz, only a white noise floor was observed when the round-trip fiber noise was uncanceled. When the active noise cancellation was turned on, a sharp coherent linewidth of 0.05 Hz was recovered, essentially limited by the analyzer resolution bandwidth. Noise was deliberately added to the original laser to further decrease its coherence time, which led to further broadening of the pedestal without any effect on the coherence of the peak in the middle. The peak height could be reduced, of course, due to the spread of the central carrier power to the noise sidebands in the pedestal.

Frequency counting of the heterodyne beat signal was also recorded, with the results shown in Fig. 10. The top panel at left shows a counting standard deviation of 5.4 Hz at 1 s counter gate time, when the fiber noise was uncompensated. The bottom panel shows the counting record when the fiber noise was canceled, now with a standard deviation of 0.9 Hz. Allan deviations determined from these two time records are shown to the right in Fig. 10. After phase compensation, the round-trip transfer process exhibits an instability of 3×10^{-15} at 1 s integration time, and averages down as $\tau^{-1/2}$. The performance is limited by the actual SNR avail-

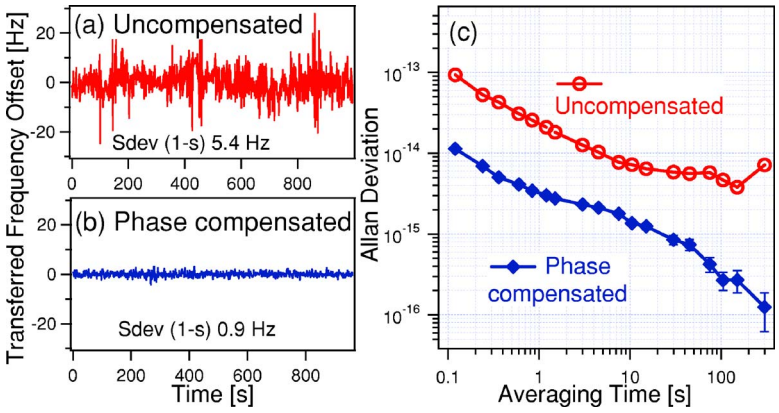

FIG. 10. (Color online) Counted heterodyne beat signal between the original laser and the returned light after a round trip. The top left panel corresponds to the case when the fiber noise is uncompensated, and the bottom left panel shows the compensated case, both with a 1 s gate time. Allan deviations determined from the time records are shown at right. (Adapted from Ref. 86.)

able in the heterodyne beat, which in turn is limited by the amount of power that can be transmitted and detected without inducing significant amplitude-to-phase conversion, as discussed in Sec. II C. In practice, the amount of cw optical power that can be sent should be less than about 90 mW in order to avoid the threshold for Brillouin scattering. The performance in Ref. 86 was also limited by use of a 1064 nm optical beam, which suffered 95% loss through each single pass of the fiber link, with multiple interconnectors. Choosing a more optimal wavelength such as 1320 or 1550 nm would dramatically improve the SNR available.

Despite the fact that the SNR was limited by the amount of optical power transmitted, the performance for actively stabilized transfer of a cw optical beam represents nearly an order of magnitude improvement over the best results for stabilized transfer of a microwave frequency reference through amplitude modulation, over the same fiber link. 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 Eq. (12), use of a higher transfer frequency allows the phase noise to be significantly higher, 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 $<1\times10^{-15}$ at 1 s integration time, if only the fundamental phase noise limit due to SNR is considered. In Ref. 86 the performance for passive microwave transfer by amplitude modulation of a cw beam was limited by the achievable SNR of the frequency-counting measurement. This is also responsible for the passive instability of microwave transfer appearing an order of magnitude worse than the passive instability of optical transfer. In the case of the optical transfer, the SNR was primarily limited by back-reflections from various connectors along the transmission fiber, which may also be the reason the instability averages down as $\tau^{-1/2}$ like white frequency noise, instead of τ^{-1} for white phase noise. Since the results of Ref. 86, 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, instabilities better than 10^{-17} at 1 s integration time (confirmed outside of the servo loop) are being achieved.[112] Another method for reducing the effects of these back-reflections involves slightly frequency shifting the optical carrier at the remote end before returning it. In this way, the round-trip signal of interest can be easily distinguished from undesirable in-line back-reflections. Despite methods for improving the SNR of direct optical transfer, use of non-polarization-maintaining fiber still can in principle degrade performance by leading to significant amplitude-to-phase conversion of the heterodyne beat from which an error signal is derived, as well as from PMD.

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 kilometers), sudden large phase delays can occur. If the servo does not 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. 108 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.

B. Remote comparison of two optical frequencies

The transfer of an optical frequency reference over the BRAN has enabled the comparison of an optical frequency standard in our JILA laboratory based on an iodine-stabilized Nd:YAG laser at 1064 nm,[2,109–111] with the 563 nm local oscillator of the single mercury-ion-based optical standard operated by the Bergquist group at NIST.[86] The 250 THz gap between the Hg^+ and the I_2 standards is spanned by a frequency comb (see Sec. VI A). The Hg^+ standard local oscillator is used to stabilize the frequency components of the comb in the Diddams' laboratory, and the comb component at 1064 nm is then used to measure the instability of the I_2 standard transmitted over the BRAN with respect to the Hg^+ reference. The Allan deviation of the I_2 standard measured with respect to Hg^+ is shown as circles in Fig. 11. Also shown is the Allan deviation determined from heterodyne beats between two similar I_2 standards at JILA (open

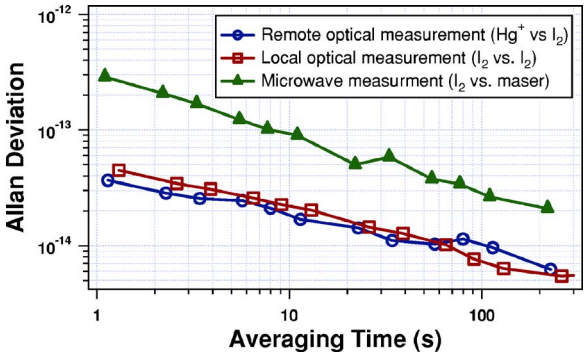

FIG. 11. (Color online) Optical measurement between the transmitted I_2 and the Hg^+ standards is shown as filled circles. The local measurement between twoI_2 standards is shown as open squares, and the measurement between the transmitted I_2 standard and the NIST maser reference is shown as triangles. (Adapted from Ref. 86.)

squares). It is satisfactory to observe that the two optical diagnostics of the I_2 standard are in such close agreement, with no observable effects of transmission-induced instability. Note that for these measurements the active noise cancellation of the BRAN was not implemented, since the passive stability it exhibits at time scales $\leq$100 s is sufficient for reliably transmitting the stability of the I_2-based standard. For comparison, the Allan deviation of the transmitted I_2 standard measured with respect to the maser reference at NIST is also shown (triangles). This clearly demonstrates the advantage of optical frequency standards for comparison, as opposed to microwave standards.

For cell-based standards exhibiting instabilities above 10^{-14} at 1 s, passive transfer of the optical carrier is sufficient. However, for primary cold-atom- or single-ion-based standards with instabilities approaching 10^{-16} at 1 s, stabilization of the transfer fiber's optical phase is necessary and has indeed reached 6×10^{-18} instability at 1 s.[112] Ultimately the most useful type of transfer will involve direct, actively stabilized transmission of an optical frequency comb, as will be discussed in the following section. Stable transfer of the complete frequency comb would achieve simultaneous transfer of optical and microwave frequency references, providing the end user with a set of carrier frequencies spanning the electromagnetic spectrum.

VI. TRANSFER WITH FREQUENCY COMB

Although direct transfer of an optical frequency reference over fiber provides better short-term stability than microwave frequency transfer using amplitude modulation, there are a few limitations of direct optical frequency transfer. The transmission of an optical frequency offers very little flexibility, so if the remote user wishes to use the frequency reference to measure an optical frequency that is even ~100 GHz away from the reference (corresponding to a 0.05% deviation for a 1550 nm reference), an optical frequency comb would be needed by the remote user to span this gap and make this measurement. If a microwave frequency reference is required by the remote user, an optical frequency comb would again be needed. Likewise, for microwave-frequency transfer using amplitude modulation a frequency comb is also needed for making an optical frequency measurement. Hence, a natural choice is to directly transmit the optical frequency comb, which can be stabilized to either an optical or microwave frequency reference prior to transfer. This provides the end user with a set of both microwave and optical frequency references spanning the spectrum. A microwave reference can be obtained by a remote user with a photodetector that provides harmonics of the repetition frequency of the transmitted comb, or equivalently, the spacing of the comb components. Alternatively, it is likely that a component of the transmitted optical comb or its second or third harmonic is sufficiently close to any optical frequency the remote user wishes to measure. Therefore, after generating the necessary harmonic of the comb, an appropriate optical frequency reference is available to the user. When the transmitted comb is stabilized to an optical frequency standard, both the microwave and optical frequency references it provides exhibit the stability of this optical standard.[23] Of course, the use of an optical frequency comb does bring other technical complexities such as amplitude-to-phase conversion, dispersion fluctuations in the transfer path, and frequency comb stabilization with two degrees of freedom.

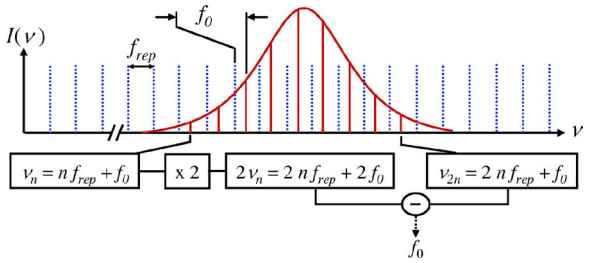

FIG. 12. (Color online) The carrier-envelope offset frequency is found by taking the difference of the frequencies of a comb mode at ν_{2n} and the second harmonic of a mode at ν_n.

A. Stabilization of frequency comb

To understand the use of a frequency comb for frequency distribution, it is important to understand some basic principles of the comb and its stabilization.[22,43,113] The frequency of every component of the comb produced by a mode-locked laser is given by two free parameters: the laser repetition frequency (f_{rep}), and the carrier-envelope offset frequency (f_0), which is related to the difference between the group and phase velocities in the laser cavity. The frequency of each component, ν_n, is then given by

$$\nu_n = nf_{\mathrm{rep}} + f_0. \tag{15}$$

Once the two microwave frequencies f_{rep} and f_0 are stabilized, every component is stabilized. The offset frequency f_0 can be measured using the self-referencing technique[101] shown schematically in Fig. 12, 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. 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.[41,114,115] 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[116] or by femtosecond lasers which directly emit octave-spanning spectra.[100,117] It is also possible to stabilize the f_0 with a $2f$-to-$3f$ scheme[118,119] which only requires 2/3 of an octave, or more complicated schemes involving various frequency ratios.[120] Also, difference-frequency schemes allow an infrared comb to be formed that has no offset frequency; only the comb spacing need be stabilized.[121]

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 is used for this, which adjusts the laser repetition frequency. 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, and even to the microwave frequency domain via the spacing of the comb lines (or, equivalently, f_{rep}).

Alternatively, instead of locking 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 in the optical domain, this locking error contributes much less to the fractional instability than it would contribute in the case of microwave stabilization of the comb. Therefore, it is best to consider optically locking the comb to derive microwave information from stable optical references in order to take advantage of large frequency ratios between the optical frequency comb modes and the microwave frequency of interest.

For transmission of frequency references over optical fibers, a stabilized comb at 1550 nm is required for minimizing loss and dispersion during fiber transmission. Mode-locked Er^{+3} fiber lasers are a good source of 1550 nm frequency combs.[122–126] They can be frequency broadened to span an entire octave, allowing the use of the f-to-$2f$ stabilization technique for f_0.[101] As an alternative, the 1550 nm comb can be stabilized to the comb of a mode-locked Ti:sapphire laser, which is itself stabilized to the frequency reference. For example, a mode-locked laser diode has a bandwidth of only ~1 nm, preventing the use of the self-referencing technique for stabilizing f_0. It is stabilized to a Ti:sapphire comb by first phase-locking the two repetition frequencies together when they are independently detected.[127] This results in the tight synchronization of the two lasers, as is described in more detail in Sec. VI D for two mode-locked fiber lasers. Then, the 1550-nm comb is frequency doubled to achieve spectral overlap with the Ti:sapphire comb, enabling the detection of the heterodyne beat between the combs in a manner similar to Ref. 128. The stabilization of this heterodyne beat completes the stabilization of the 1550 nm comb to the Ti:sapphire comb.[129] Low phase- and amplitude-noise mode-locked diode lasers[130,131] can take advantage of this general scheme to stabilize their optical frequency comb outputs.

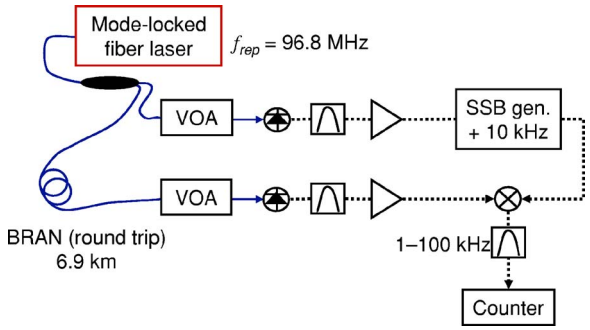

FIG. 13. (Color online) The frequency instability for transfer of the comb repetition frequency is measured by counting the 10 kHz mixing product between the shifted reference signal from the mode-locked fiber laser and the transmitted signal. SSB gen., single-sideband generator; VOA, variable optical attenuator. (Adapted from Ref. 141.)

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. 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,[41,101,120,132,133] infrared,[121,134–136] visible,[137,138] or even UV and XUV [139,140] across vast frequency gaps.

B. Passive microwave transfer: Effect of fiber dispersion

The instability introduced during transfer of a microwave frequency reference using a frequency comb was determined using the comb produced by a passively mode-locked fiber ring laser with a fundamental repetition frequency of ~97 MHz.[141] The setup used is shown in Fig. 13. The output of the fiber laser is split into two portions. One portion is transmitted through the BRAN fiber while the second portion provides a reference against which to measure noise introduced by the transmission. The reference and transmitted pulse trains are detected with two photodetectors, and the resultant microwave signals are filtered to select the eighth harmonic (774 MHz) of the repetition frequency. Because of amplitude-to-phase noise conversion in the photodetectors, as discussed in Sec. II C, the optical power incident on each photodetector plays a crucial role in determining the detected instability. Fiber-pigtailed variable optical attenuators are used to control the power incident on the photodetectors. Since Allan deviation measurements need a frequency-shifted local copy for beat frequency counting (while phase noise measurements can be made in a homodyne fashion), the reference signal is first shifted by 10 kHz

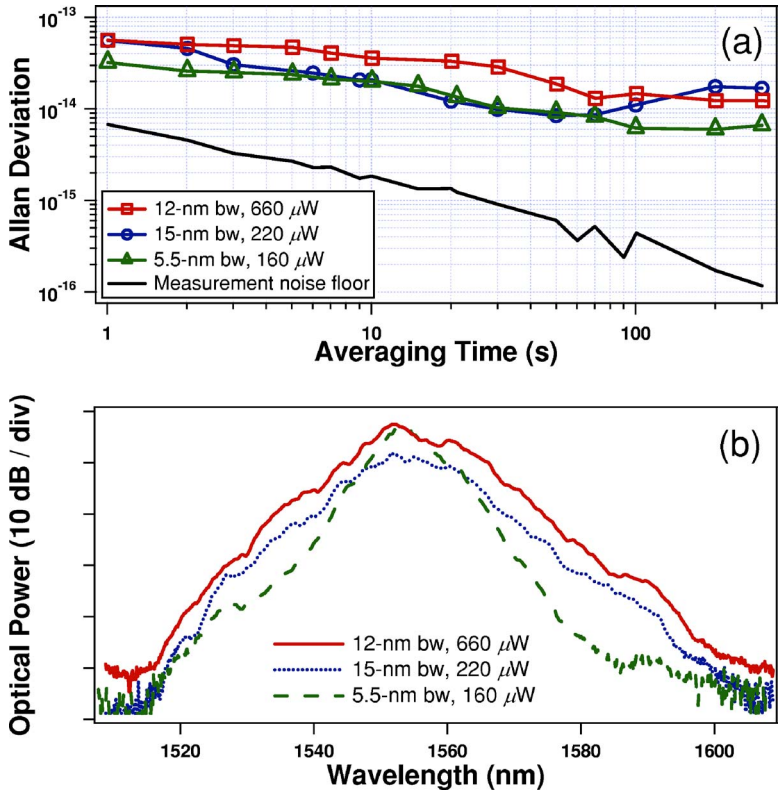

FIG. 14. (Color online) (a) The Allan deviation for microwave-frequency transfer over the BRAN fiber using various bandwidths for the transmitted fiber laser pulses and several different average optical powers incident on the receiving photodetector. The lowest instability is achieved with the narrowest bandwidth, which provides a sufficient SNR for the received microwave signal with the smallest optical power. The measurement noise floor, obtained with a short piece of fiber, is the same for all operating conditions. (b) The pulse spectra for the three mode-locking conditions used in (a). bw, bandwidth. (Adapted from Ref. 141.)

using a single-sideband generator. Then, the frequency fluctuations introduced by the transfer are measured by filtering and counting the 10 kHz mixing product of the shifted reference signal and the transmitted signal. Since both signals originate from the same optical source, fluctuations of the fiber laser's f_{rep} are common mode and do not contribute to the detected noise. Therefore, for these measurements the fiber laser is left free-running. As opposed to the case of the free-running cw optical source in Sec. V, the microwave coherence time of the free-running fiber laser's f_{rep} harmonic is significantly longer than T_{rt}. This prevents the servo from being confused between incoherence of the source and incoherence of the fiber link, as discussed in Sec. III C.

It is found that the lowest instability for transfer with the comb is achieved when the average optical powers on the photodetectors producing the transmitted and reference signals are reduced as low as possible while still achieving a sufficient SNR for the recovered microwave signals. This is attributed to there being amplitude noise on the laser pulse train that is converted to phase noise by the photodetectors,[65] as described in Sec. II C. A good solution is to stabilize the optical power incident on the detectors. Without such implementation, reducing the optical power decreases the magnitude of the power fluctuations, thereby lowering the magnitude of phase fluctuations produced by the photodetectors.

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

The effect of these laser parameters on the transfer stability is investigated by varying the pulse bandwidth and average optical power, which is accomplished by adjusting the pump power and polarization components of the fiber laser, thus modifying the mode-locking condition. The results are shown in Fig. 14. The lowest instability is obtained with the narrowest pulse bandwidth, when a high SNR is achievable with the minimum average optical power. Broader pulse bandwidths require more optical power to recover the same SNR, which increases the instability because of amplitude-to-phase conversion of the photodetectors. Reducing the optical power but maintaining a broad bandwidth

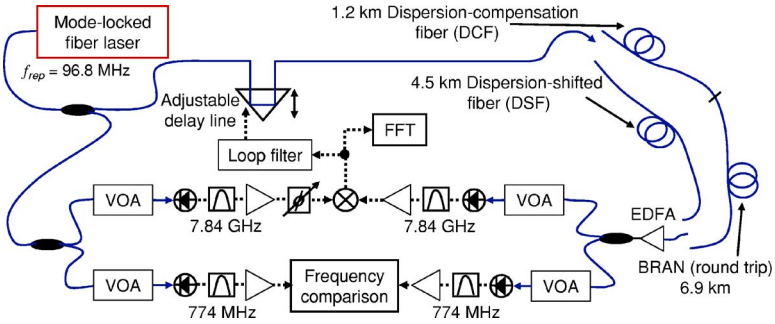

FIG. 15. (Color online) To cancel the phase noise introduced during transfer of the comb repetition frequency, the phase noise is first measured by mixing the 81st harmonic of the repetition frequency obtained from the transmitted pulses and a reference portion of the fiber laser output. The resulting error signal is then applied to an adjustable delay line after appropriate filtering and amplification. The transmission fiber is either the dispersion-compensated BRAN fiber, or the DSF. A FFT provides spectral analysis of the in-loop timing jitter. The out-of-loop frequency instability is determined from the eighth harmonic of the repetition frequency obtained from the lower pair of photodetectors as in Fig. 13. VOA, variable optical attenuator. (Adapted from Ref. 142.)

still produces poor results, since the SNR is insufficient. Clearly, highly linear, high-power photodetectors would be helpful to keep the amplitude-to-phase conversion very low, while allowing the use of higher incident powers for achieving a sufficient SNR, even with temporally stretched large bandwidth pulses. Active stabilization of the optical power would further help to prevent amplitude-to-phase conversion while allowing the increased SNR offered by higher optical powers. Without the dispersive BRAN fiber, the measured instability is no longer dependent on the pulse bandwidth, since for all cases the pulses are shorter than the timing resolution of the photodetector. This is indicated by the fact that the measurement noise floor, using a short, 2 m fiber and shown by the solid curve in Fig. 14(a), is the same regardless of the operating parameters.

At 1 s, the best stability provided by the comb for passive microwave frequency transfer[141] is comparable to that reported for the LPL-SYRTE link[82] and NASA's Deep Space Network[104] without stabilization. However, the comb provides nearly an order of magnitude improvement over the performance demonstrated for modulated cw transfer over the same BRAN fiber, and provides stability similar to that obtained for optical frequency transfer. This is likely because the modulated cw transfer over the BRAN made use of only a 10% modulation index. The average optical power present for a given SNR of the rf signal from modulated cw transfer would have been much higher than that needed to produce the same SNR from a pulse train.

C. Active noise cancellation using group delay

The stability of microwave frequency transfer with the optical frequency comb can be further improved by implementing active cancellation of the fiber noise.[142] An adjustable delay line is used to compensate for fluctuations in the effective transmission length and therefore group delay of the fiber. This active noise cancellation reduces the frequency instability for metrology applications and the broadbandwidth timing jitter for remote synchronization applications. Such control of the timing jitter has only recently been attempted,[142,143] although interesting techniques for measuring the phase noise of pulse trains have previously been explored.[144,145] For comparison, Ref. 105 reports the timing jitter in an active system (microwave transfer by modulating a cw optical carrier) that was limited by a feedback-loop time constant of ~10 s and a measurement resolution of 400 fs, for a limited bandwidth extending from dc to only 100 Hz. The active noise cancellation reported by Ref. 143 provided a residual timing jitter after transfer of ~20 fs over a bandwidth from 1 Hz to 50 MHz, the Nyquist frequency, for transmission over several kilometers. Depending on the application, the entire Nyquist bandwidth may not always be relevant since some transducers at the remote end will not have the speed or bandwidth to track the fast noise. Therefore, in order to achieve low phase noise at Fourier frequencies outside the tracking bandwidth of local oscillators at the remote end, oscillators with excellent noise properties at these higher Fourier frequencies are required. For example, the phase noise of a free-running mode-locked Ti:sapphire laser goes as f^{-4} above 10 kHz, allowing it to achieve very low phase noise at high frequencies.[146]

In order to achieve the best transfer with respect to phase noise or instability, dispersion compensation of the fiber should be implemented. It was previously explained that to reduce the instability from amplitude-to-phase noise conversion it is crucial to minimize the average optical power on the photodetector while maintaining a sufficient SNR of the microwave signal, which requires transmission of the shortest pulses possible. However, the dispersion of the BRAN fiber stretches the pulses to ~1 ns. The impact of this dispersion was mitigated for passive transfer in Ref. 141 by reducing the spectral bandwidth of the fiber laser's pulses, but the reduced spectral coverage potentially renders the optical comb less useful at the remote end. A more useful strategy for minimizing the pulse length after transfer is to implement dispersion control of the transmission fiber.[78,79] This enables the transmission of pulses with durations below the photodetector resolution of 40 ps, which gives a high SNR for the detected repetition frequency harmonics, thereby reducing the phase noise limits associated with the SNR as discussed in Sec. II C. Figure 15 shows the setup used in Ref. 142 for transfer of the repetition frequency harmonics using active stabilization and dispersion compensation. The fiber laser output is again split into two portions, one of which is used for microwave-frequency transfer, while the other provides a reference against which to measure the noise introduced during transfer. The transmitted portion

passes first through an adjustable delay line before entering the distribution fiber. The adjustable delay line, which will be described in detail later, provides the means to compensate for path-length fluctuations of the distribution fiber and actively cancel the transfer-induced noise. The distribution fiber consists either of the BRAN round trip preceded by an appropriate length of dispersion-compensation fiber (DCF), or simply a long dispersion-shifted fiber (DSF) without the BRAN. The DCF has a group delay dispersion (GDD) with a sign opposite to that of the BRAN, but with equal magnitude for a carefully measured length, yielding a transmission path with a net GDD of magnitude <4 ps^2. This produces transmitted pulses that are $\sim$60 ps in duration, limited by uncompensated higher-order dispersion. The DSF has had its 0 GDD point shifted to the wavelength of the fiber laser (1550 nm) and has a net GDD of magnitude <23 ps^2. It provides transmitted pulses shorter than 40 ps. Both transmission media transmit pulses that are short enough to yield a high SNR for the received microwave signal with an average optical power sufficiently low ($\sim 30\ \mu$W) to significantly reduce noise caused by amplitude-to-phase conversion in the photodetection process.

After the distribution fiber, an erbium-doped fiber amplifier (EDFA) is used for compensation of losses in the fiber link. It has been confirmed that the EDFA does not add additional instability during transfer.[141] The reference and transmitted pulse trains are each detected on two photodetectors, with the incident average optical powers adjusted with variable optical attenuators. The microwave signals from the upper pair of detectors in Fig. 15 are filtered to select the 81st harmonic (7.84 GHz) of the fiber laser's f_{rep} and are mixed with the appropriate relative phase to provide an error signal for the phase noise introduced during transfer. Using a high harmonic helps reduce possible phase noise limits associated with shot noise and thermal noise, as predicted by Eq. (7). The measured jitter is applied to the adjustable delay line with appropriate gain and filtering to minimize the transfer-induced noise. A FFT provides spectral analysis of the residual in-loop jitter. The second pair of microwave signals is filtered to select the eighth harmonic (774 MHz) of f_{rep} for out-of-loop frequency instability measurements. The frequencies of the local and round-trip signals are compared for averaging times ≥ 1 s by shifting the reference signal 10 kHz with a single-sideband generator and subsequently counting the 10 kHz mixing product of the shifted reference and unshifted transmitted signals, as shown in Fig. 13. This provides another out-of-loop measurement of the residual frequency instability present with active cancellation of the transfer noise.

For the microwave frequency transfer considered here, it is necessary to stabilize the group delay of the fiber, which has path length fluctuations on the order of millimeters over the course of several hours. Therefore, an adjustable delay line is employed to stabilize the microwave phase of the 81st harmonic of f_{rep}. It consists of two elements: a free-space path length terminated with a voice coil-activated retroreflector and a PZT-actuated fiber stretcher. The voice coil operates by the same principle as an audio speaker and provides a large range of motion (a few millimeters), but with low

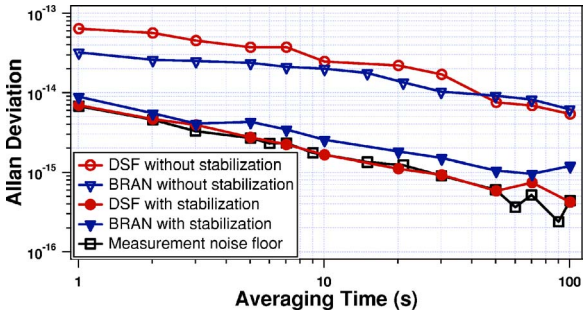

FIG. 16. (Color online) The Allan deviation for microwave-frequency transfer through the DSF and dispersion-compensated BRAN fiber indicates that the active noise cancellation reduces the transfer instability for transfer through the DSF to the measurement noise floor, and nearly to the noise floor for transfer through the dispersion-compensated BRAN fiber. The noise floor is determined using a short piece of fiber. (Adapted from Ref. 142.)

bandwidth. The fiber stretcher consists of optical fiber tightly wound 40 times around a cylindrical PZT. It extends the servo bandwidth to $\sim$1 kHz. The gain and bandwidth of the servo are limited by a resonance in the fiber stretcher at $\sim$20 kHz. To transfer the microwave signal to a remote user, the detected noise introduced after a round-trip transmission was used to cancel the perturbations for one-way transfer. This was done by simply retroreflecting a portion of the optical comb into a return trip along the fiber. As long as the light double-passes the adjustable delay line before detecting an error signal, as in Fig. 5(c), stabilizing the round-trip phase will also have the effect of stabilizing the one-way phase at the remote end.

Figure 16 shows the fractional frequency instability introduced during transfer through the dispersion-compensated BRAN fiber (open triangles) and the DSF (open circles) for averaging times ≥ 1 s, without active noise cancellation. These data are obtained from the frequency comparison of the 774 MHz signals from the pair of out-of-loop photodetectors in Fig. 15. With the activation of the noise cancellation, the instability for transfer through the BRAN is reduced from $>3\times10^{-14}$ to $<9\times10^{-15}$ (filled triangles) for a 1 s averaging time and reaches 1×10^{-15} at 100 s, just slightly above the noise floor of the measurement system (open squares). The noise cancellation reduces the instability for transfer through the DSF from $>6\times10^{-14}$ to $<7\times10^{-15}$ (filled circles) for a 1 s averaging time, limited by the measurement noise floor. Transfer through the DSF reaches an instability of a few parts in 10^{16} at 100 s with the noise cancellation. The measurement noise floor is determined by replacing the transmission fiber with a short length of fiber. Since these results were obtained with independent detectors outside the stabilization loop, they represent an upper limit to any contributions the in-loop photodetectors (operating at the 81st harmonic of the repetition rate) may have made toward instability.

While a wide servo bandwidth benefits the frequency instability measurement as shown in Fig. 16, which is useful for frequency metrology applications where the long-term (>1 s) transfer instability could be important, it is the time-domain experiments that benefit the most from a fast servo, as the adverse short time scale noises can be suppressed to

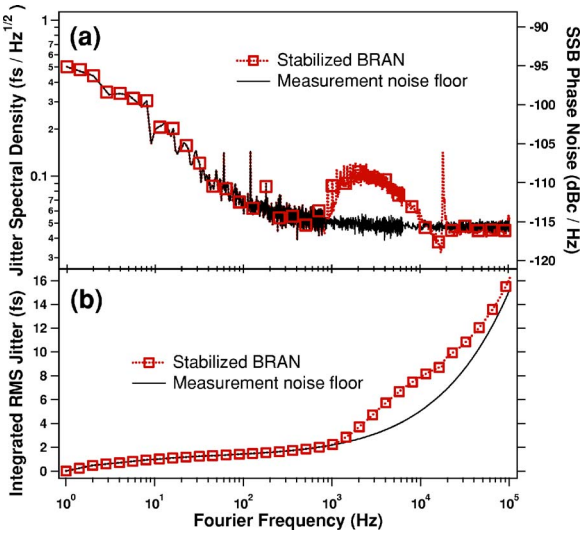

FIG. 17. (Color online) (a) The rms timing jitter spectral density (left axis) for transfer over the dispersion-compensated BRAN fiber and the equivalent single-sideband (SSB) phase noise of the transmitted 81st harmonic of the comb repetition frequency (right axis) with active noise cancellation. (b) The integrated rms jitter over a bandwidth from 1 Hz to f_h vs f_h for transfer over the BRAN fiber. The measurement noise floor is obtained by replacing the transmission fiber with a short piece of fiber. (Adapted from Ref. 142.)

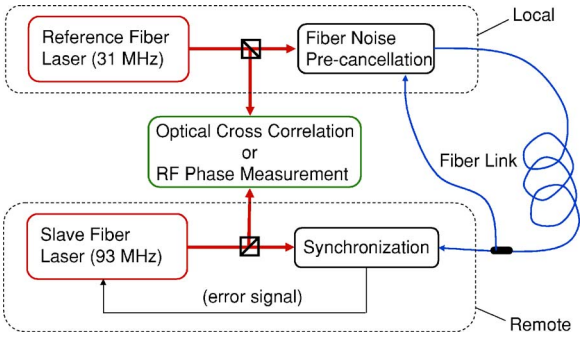

FIG. 18. (Color online) High-level diagram showing the setup for remote synchronization of two femtosecond fiber lasers. The repetition rate of the reference laser is transmitted through 4.5 km of dispersion-shifted fiber with active stabilization of the group delay. The remote laser is synchronized 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. 143.)

preserve the true phase coherence. Figure 17(a) shows the spectral density of the rms jitter (measured in-loop) introduced during transfer of the fiber laser's f_{rep} over the dispersion-compensated BRAN fiber with active stabilization of the transmission path (open squares). The corresponding single-sideband phase noise, $L(f)$, of the transmitted 81st harmonic of the fiber laser repetition frequency is shown on the right axis. The active noise cancellation reduces the residual jitter to the measurement noise floor over the entire bandwidth from 1 Hz to 100 kHz, except for the Fourier frequency range from ~1 to ~10 kHz. This jitter spectral density represents the upper limit of the combined in-loop noise and measurement noise floor. The measurement noise floor is determined by replacing the BRAN fiber and the DCF with a short piece of fiber and attenuating the light incident on the photodetector to produce the same photocurrent and microwave signal power that are obtained for transmission over the dispersion-compensated BRAN. Figure 17(b) shows the corresponding total rms timing jitter integrated over a bandwidth from 1 Hz to f_h, versus Fourier frequency f_h. The total integrated jitter over the bandwidth from 1 Hz to 100 kHz is reduced to ~16 fs, nearly limited by the measurement noise floor. This result was confirmed by an out-of-loop measurement, presented in the following section, by optical cross correlation.

By compensating the effects of fiber dispersion, attenuating optical powers to avoid amplitude-to-phase conversion, detecting high harmonics of the comb's repetition rate, and actuating on the group delay with a high servo bandwidth, frequency transfer using the comb has reached a level of $<7\times10^{-15}$ at 1 s for kilometer-scale fiber links, limited by the measurement noise floor. Also, as a result of the high servo bandwidth, timing applications such as the remote synchronization discussed next are now able to achieve very low residual jitter (<20 fs over the Nyquist bandwidth) even when systems are separated by great distances.

D. Remote synchronization of ultrafast lasers

Actively stabilized microwave frequency transfer using a frequency comb can be employed to synchronize two remotely located femtosecond lasers.[143] To achieve remote synchronization, the pulse train from the reference laser is transmitted through several kilometers of fiber using the stable transfer technique previously discussed. Then, the remote slave laser is synchronized to the stable incoming pulse train with a high-bandwidth actuator. The setup used for remote synchronization experiments is shown in Fig. 18. In our experiment, two mode-locked fiber lasers were constructed so that harmonics of their repetition frequencies coincide at ~93 MHz. The transmitting laser is free-running, while the slave laser has an electro-optic modulator (EOM) (Ref. 147) and a PZT inside its cavity. The slave laser is 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 μm) that allows for long time scale locking (>12 h), while the EOM provides a high bandwidth actuator to achieve tight synchronization.

To characterize the residual timing jitter of the synchronization we use a crossed-beam, background free, optical cross correlation of the two lasers' pulse trains, as discussed in Sec. II D. The two pulse trains are focused onto a $LiIO_3$ crystal (type I phase matching), which generates sum frequency light (SFG) when the two pulses overlap in time and space. To achieve temporal overlap, we use two phase-locked loops that operate at two different timing resolutions.[68,69] A course-timing loop operates at the fundamental frequency of 93 MHz, while a higher-resolution loop operates at 7.6 GHz (80th harmonic of f_{rep}). A phase shifter in the fundamental frequency loop allows course-timing adjustments such that temporal overlap between the two pulse trains can be found. Once a SFG signal is observed on a photomultiplier tube (PMT), we measure the total cross-correlation width to calibrate SFG intensity fluctuation to timing jitter conversion. We then transfer control from the fundamental frequency loop to the 7.6 GHz loop. A phase shifter in the high-harmonic loop allows fine-tuning of the

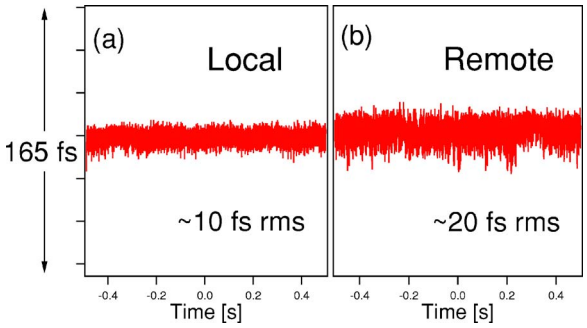

FIG. 19. (Color online) Optical cross-correlation measurement of the timing jitter for (a) local synchronization, and (b) slave synchronization through a fiber link. Both traces are shown with a 50 MHz low-pass filter that allows us to observe the timing jitter up to the Nyquist frequency. (Adapted from Ref. 143.)

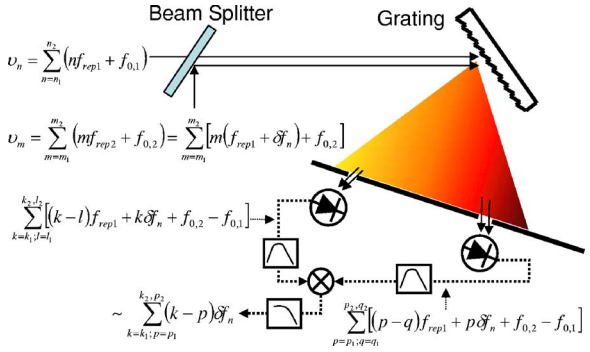

FIG. 20. (Color online) All-optical generation of an error signal for synchronization of two combs. When the combs are close in repetition rate, we can express one (the slave) as having a slight amount of frequency noise on its repetition rate (δf_n) as compared to the master comb. By spectral leveraging, this noise can be greatly amplified during the detection process to produce an improved error signal.

timing overlap of the pulses such that we can 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 monitor 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 is performed on the same optical table that holds the two lasers used in the experiment. This configuration allows a direct comparison of the two lasers, which reveals 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 allows a relatively narrow cross-correlation width of 165 fs because the pulses have 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. 19(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. 19(b), which was directly observed from the out-of-loop cross-correlation measurement. This result is in agreement with the in-loop measurement discussed in Sec. VI C.

E. All-optical detection of synchronization error signal

The noise floor limits present in microwave detection schemes (discussed in Sec. II C) can be circumvented by turning to optical detection of the error signal for synchronization. A modified optical cross-correlation arrangement was employed to derive a dispersion-type error signal for local synchronization of a fs Ti:sapphire laser with a fs Cr:LiSaf laser.[70] Alternatively, an optical frequency standard was employed at NIST to synchronize two local fs Ti:sapphire lasers.[71] While in principle both of these techniques could be used to derive an error signal for remote synchronization, their implementation in this setting would be rather complex. In the former case, near-perfect chirp compensation of the transmission pulse would be required, while the latter case, similar to the work discussed in Sec. V, would require stabilized transmission of an optical frequency standard.

A third approach is shown in Fig. 20. The master and slave combs are combined colinearly and then spectrally resolved. An optical heterodyne beat is detected at each spectral extreme. The difference of these two optical heterodyne beats contains an error signal proportional to the difference in the comb spacing between the two lasers (δf_n in Fig. 20), multiplied by the spectral separation of the beats ($k-p$ in Fig. 20). A complete derivation of this result appears in Ref. 73. The principle of this approach has been demonstrated by detecting and canceling the noise from a 100 m fiber link with a mode-locked fiber laser operating at 1550 nm. The pulse train from the free-running laser is split into two copies, one of which is sent through the fiber link. After transit the transferred copy is overlapped with the local copy, and the system shown in Fig. 20 employs a spectral separation of 50 nm in order to form an error signal. With this spectral separation ($\sim$6 THz) and the repetition rate of the lasers (50 MHz), the optical heterodyne technique is effectively generating a synchronization signal at the 120 000th harmonic. Using an interferometric cross correlation as the diagnostic technique (described in Sec. II D), the residual jitter is measured to be $<$0.04 fs (1 Hz to 10 MHz).[73]

F. Future work and considerations

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

fication of the optical heterodyne setup shown in Fig. 20 must be made. Instead of taking a difference between the heterodyne beats, each beat would need to be separately locked, thereby producing two error signals. While not completely orthogonal, one of these error signals could be used to lock the group delay (via similar actuators discussed previously). The second error signal would then be fed back to an adjustable dispersion compensation device (such as a PZT-equipped chirped fiber Bragg grating). Thus, by locking the comb at two distinct spectral regions both f_{rep} and f_0 would be stabilized. Correction of higher-order dispersion variations would require sampling of the interference spectra at more than two points to generate additional error signals.

VII. 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 ($\geq$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[112] at 3×10^{-15} at 1 s, while microwave transfer using the comb performs at $<7\times10^{-15}$ 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.

Fiber transfer experiments are starting to aid the frequency standards community to transfer the highest quality optical atomic clock signals at this point. It remains to be seen which type of frequency transfer will result in the highest stability links. In particular, longer data runs for both direct optical transfer and frequency comb transfer will be necessary in order to determine how they average down with time, once other technical noise limitations are overcome. On the other hand, it is already a useful discovery that for short time scale synchronization applications ($\tau<10$ s); both technologies yield superior results to microwave transfer by modulation of a cw optical carrier. Other factors that may play into the decision of which method to use include whether a microwave reference is needed, or an optical reference, or both simultaneously, and whether the remote user has access to an optical frequency comb or would prefer to have the comb itself used for distribution. Finally, as discussed in Sec. III, signal attenuation, integrated noise of the link, and PMD must all be considered when designing a particular fiber link's noise cancellation system, since their relative scalings with length of the link will determine particular system design parameters.

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 canceled. Therefore, as the distribution distances get longer, the bandwidth within which noise can be canceled 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.

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 in a recent Reference Frame article by D. Kleppner.[148] 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[149] 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!

ACKNOWLEDGMENTS

The authors are indebted to J. C. Bergquist, Y. F. Chen, S. T. Cundiff, S. A. Diddams, L. W. Hollberg, S. Jefferts, and T. Parker for helpful collaborations and stimulating discussions. The authors also thank A. Ludlow, M. Boyd, and J. L. Peng for work on fiber transfer and M. Boyd, A. Ludlow, T. Zelevinsky, S. Blatt, and T. Ido for the Sr work. This work was supported by the U.S. Office of Naval Research, NASA, the National Institute of Standards and Technology, and the National Science Foundation. K. W. Holman was supported by the Hertz Foundation.

[1] S. A. Diddams *et al.*, Science **293**, 825 (2001).
[2] J. Ye, L.-S. Ma, and J. L. Hall, Phys. Rev. Lett. **87**, 270801 (2001).
[3] R. J. Rafac, B. C. Young, J. A. Beall, W. M. Itano, D. J. Wineland, and J. C. Bergquist, Phys. Rev. Lett. **85**, 2462 (2000).
[4] P. O. Schmidt, T. Rosenband, C. Langer, W. M. Itano, J. C. Bergquist, and D. J. Wineland, Science **309**, 749 (2005).
[5] H. S. Margolis, G. P. Barwood, G. Huang, H. A. Klein, S. N. Lea, K.

Szymaniec, and P. Gill, Science **306**, 1355 (2004).
[6] T. Schneider, E. Peik, and C. Tamm, Phys. Rev. Lett. **94**, 230801 (2005).
[7] P. Dube, A. A. Madej, J. E. Bernard, L. Marmet, J. S. Boulanger, and S. Cundy, Phys. Rev. Lett. **95**, 033001 (2005).
[8] F. Ruschewitz, J. L. Peng, H. Hinderthur, N. Schaffrath, K. Sengstock, and W. Ertmer, Phys. Rev. Lett. **80**, 3173 (1998).
[9] U. Sterr, C. Degenhardt, H. Stoehr, C. Lisdat, H. Schnatz, J. Helmcke, F. Riehle, G. Wilpers, C. Oates, and L. Hollberg, C. R. Phys. **5**, 845 (2004).
[10] G. Wilpers, T. Binnewies, C. Degenhardt, U. Sterr, J. Helmcke, and F. Riehle, Phys. Rev. Lett. **89**, 230801 (2002).
[11] T. Ido, T. H. Loftus, M. M. Boyd, A. D. Ludlow, K. W. Holman, and J. Ye, Phys. Rev. Lett. **94**, 153001 (2005).
[12] M. Takamoto, F.-L. Hong, R. Higashi, and H. Katori, Nature **435**, 321 (2005).
[13] A. D. Ludlow, M. M. Boyd, T. Zelevinsky, S. M. Foreman, S. Blatt, M. Notcutt, T. Ido, and J. Ye, Phys. Rev. Lett. **96**, 033003 (2006).
[14] A. Brusch, R. Le Targat, X. Baillard, M. Fouche, and P. Lemonde, Phys. Rev. Lett. **96**, 103003 (2006).
[15] Z. W. Barber, C. W. Hoyt, C. W. Oates, L. Hollberg, A. V. Taichenachev, and V. I. Yudin, Phys. Rev. Lett. **96**, 083002 (2006).
[16] T. P. Heavner, S. R. Jefferts, E. A. Donley, J. H. Shirley, and T. E. Parker, IEEE Trans. Instrum. Meas. **54**, 842 (2005).
[17] S. Bize *et al.*, J. Phys. B **38**, S449 (2005).
[18] Advertised data for model SLCO of Poseidon Scientific Instruments, Australia. Trade name is mentioned for scientific clarity. We do not endorse this product; others might be equally or better suited.
[19] Hydrogen maser: ST-5, NIST Time and Frequency Division (T. Parker), Boulder.
[20] M. Notcutt, L.-S. Ma, J. Ye, and J. L. Hall, Opt. Lett. **30**, 1815 (2005).
[21] A. D. Ludlow, X. Huang, T. Zanon, S. M. Foreman, M. M. Boyd, S. Blatt, and J. Ye, Opt. Lett. **32**, 641 (2007).
[22] J. Ye, H. Schnatz, and L. W. Hollberg, IEEE J. Sel. Top. Quantum Electron. **9**, 1041 (2003).
[23] L.-S. Ma, Z. Bi, A. Bartels, L. Robertsson, M. Zucco, R. S. Windeler, G. Wilpers, C. Oates, L. Hollberg, and S. A. Diddams, Science **303**, 1843 (2004).
[24] A. Bartels, S. A. Diddams, C. W. Oates, G. Wilpers, J. C. Bergquist, W. H. Oskay, and L. Hollberg, Opt. Lett. **30**, 667 (2005).
[25] S. G. Karshenboim, Can. J. Phys. **78**, 639 (2000).
[26] J. D. Prestage, R. L. Tjoelker, and L. Maleki, Phys. Rev. Lett. **74**, 3511 (1995).
[27] H. Marion *et al.*, Phys. Rev. Lett. **90**, 150801 (2003).
[28] S. Bize *et al.*, Phys. Rev. Lett. **90**, 150802 (2003).
[29] M. Fischer *et al.*, Phys. Rev. Lett. **92**, 230802 (2004).
[30] 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, Phys. Rev. Lett. **87**, 091301 (2001).
[31] E. R. Hudson, H. J. Lewandowski, B. C. Sawyer, and J. Ye, Phys. Rev. Lett. **96**, 143004 (2006).
[32] J. Ye *et al.*, Appl. Phys. B: Lasers Opt. **74**, S27 (2002).
[33] B. Shillue, S. AlBanna, and L. D'Addario, in *2004 IEEE International Topical Meeting on Microwave Photonics Technical Digest* (Institute of Electrical and Electronics Engineers, Piscataway, NJ, 2004), pp. 201–204.
[34] http://www-ssrl.slac.stanford.edu/lcls
[35] R. W. Schoenlein, W. P. Leemans, A. H. Chin, P. Volfbeyn, T. E. Glover, P. Balling, M. Zolotorev, K. J. Kim, S. Chattopadhyay, and C. V. Shank, Science **274**, 236 (1996).
[36] M. F. DeCamp *et al.*, Nature **413**, 825 (2001).
[37] A. L. Cavalieri *et al.*, Phys. Rev. Lett. **94**, 114801 (2005).
[38] J. Levine, Rev. Sci. Instrum. **70**, 2567 (1999).
[39] J. L. Hall and T. W. Hänsch, in *Femtosecond Optical Frequency Comb Technology: Principle, Operation, and Applications*, edited by J. Ye and S. T. Cundiff (Springer, New York, 2005), pp. 1–11.
[40] A. Bauch *et al.*, Metrologia **43**, 109 (2006).
[41] J. Reichert, R. Holzwarth, T. Udem, and T. W. Hänsch, Opt. Commun. **172**, 59 (1999).
[42] S. T. Cundiff, J. Ye, and J. L. Hall, Rev. Sci. Instrum. **72**, 3749 (2001).
[43] S. T. Cundiff and J. Ye, Rev. Mod. Phys. **75**, 325 (2003).
[44] J. L. Hall and M. Zhu, in *Laser Manipulation of Atoms and Ions, Proceedings of the International School of Physics*, "Enrico Fermi:" Course 118, edited by E. Arimondo, W. D. Phillips, and F. Strumia (North Holland, Amsterdam, 1992), pp. 671–702.
[45] M. Zhu and J. L. Hall, J. Opt. Soc. Am. B **10**, 802 (1993).
[46] J. Rutman, Proc. IEEE **66**, 1048 (1978).
[47] F. L. Walls, in *Laser Frequency Stabilization, Standards, Measurement, and Applications*, Vol. 4269 of Proceedings of SPIE, edited by J. L. Hall and J. Ye (SPIE, Bellingham, WA, 2001), pp. 170–177.
[48] *Characterization of Clocks and Oscillators: NIST Technical Note 1337*, edited by D. B. Sullivan, D. W. Allan, and F. L. Walls (U.S. GPO, Washington, DC, 1990).
[49] D. W. Allan, Proc. IEEE **54**, 221 (1966).
[50] R. P. Scott, C. Langrock, and B. H. Kolner, IEEE J. Sel. Top. Quantum Electron. **7**, 641 (2001).
[51] A. Hati, D. A. Howe, F. L. Walls, and D. Walker, in *Proceedings of the 2003 IEEE International Frequency Control Symposium and PDA Exhibition Jointly with the 17th European Frequency and Time Forum* (Institute of Electrical and Electronics Engineers, Piscataway, NJ, 2003), pp. 516–520.
[52] P. J. Winzer, J. Opt. Soc. Am. B **14**, 2424 (1997).
[53] J. J. McFerran, E. N. Ivanov, A. Bartels, G. Wilpers, C. W. Oates, S. A. Diddams, and L. Hollberg, Electron. Lett. **41**, 650 (2005).
[54] L. Hollberg *et al.*, IEEE J. Quantum Electron. **37**, 1502 (2001).
[55] F. L. Walls, E. S. Ferre-Pikal, and S. R. Jefferts, IEEE Trans. Ultrason. Ferroelectr. Freq. Control **44**, 326 (1997).
[56] 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, *IEEE Proceedings* (Microwave Photonics, Piscataway, NJ, 2004).
[57] P.-L. Liu, K. J. Williams, M. Y. Frankel, and R. D. Esman, IEEE Trans. Microw. Theory Tech. **47**, 1297 (1999).
[58] D. A. Tulchinsky and K. J. Williams, IEEE Photon. Technol. Lett. **17**, 654 (2005).
[59] J. K. A. Everard and C. Broomfield, IEEE Int. Frequency Control Symposium, pp. 156–160 (2001).
[60] E. S. Ferre-Pikal, IEEE Trans. Circuits Syst., I: Fundam. Theory Appl. **51**, 1417 (2004).
[61] E. N. Ivanov, M. E. Tobar, and R. A. Woode, IEEE Trans. Ultrason. Ferroelectr. Freq. Control **45**, 1526 (1998).
[62] E. Rubiola, E. Salik, N. Yu, and L. Maleki, Electron. Lett. **39**, 1389 (2003).
[63] E. Rubiola, V. Giordano, and J. Groslambert, Rev. Sci. Instrum. **70**, 220 (1999).
[64] K. J. Williams, R. D. Esman, and M. Dagenais, IEEE Photonics Technol. Lett. **6**, 639 (1994).
[65] E. N. Ivanov, S. A. Diddams, and L. Hollberg, IEEE J. Sel. Top. Quantum Electron. **9**, 1059 (2003).
[66] E. N. Ivanov, S. A. Diddams, and L. Hollberg, IEEE Trans. Ultrason. Ferroelectr. Freq. Control **52**, 1068 (2005).
[67] T. Brown (2006), Work performed at JILA (unpublished).
[68] L.-S. Ma, R. K. Shelton, H. C. Kapteyn, M. M. Murnane, and J. Ye, Phys. Rev. A. **64**, 021802 (2001).
[69] R. K. Shelton, S. M. Foreman, L.-S. Ma, J. L. Hall, H. C. Kapteyn, M. M. Murnane, M. Notcutt, and J. Ye, Opt. Lett. **27**, 312 (2002).
[70] T. R. Schibli, J. Kim, O. Kuzucu, J. T. Gopinath, S. N. Tandon, G. S. Petrich, A. Kolodziejski, J. G. Fujimoto, E. P. Ippen, and F. X. Kärtner, Opt. Lett. **28**, 947 (2003).
[71] A. Bartels, S. A. Diddams, T. M. Ramond, and L. Hollberg, Opt. Lett. **28**, 663 (2003).
[72] J.-C. Diels and W. Rudolph, *Ultrashort Laser Pulse Phenomena: Fundamentals, Techniques, and Applications on a Femtosecond Time Scale, Optics and Photonics* (Academic, San Diego, 1996).
[73] Y.-F. Chen, J. Jiang, and D. J. Jones, Opt. Express **14**, 12134 (2006).
[74] J. K. Ranka, A. Gaeta, A. Baltuska, M. S. Pshenichnikov, and D. A. Wiersma, Opt. Lett. **22**, 1344 (1997).
[75] A. Pauchard, M. Bitter, Z. Pan, S. Kristjansson, L. A. Hodge, K. J. Williams, D. A. Tulchinsky, S. G. Hummel, and Y. H. Lo, IEEE Photonics Technol. Lett. **16**, 2544 (2004).
[76] K. J. Williams and R. D. Esman, J. Lightwave Technol. **17**, 1443 (1999).
[77] F. Narbonneau, M. Lours, S. Bize, A. Clairon, G. Santarelli, O. Lopez, C. Daussy, A. Amy-Klein, and C. Chardonnet, Rev. Sci. Instrum. **77**, 064701 (2006).
[78] C.-C. Chang and A. M. Weiner, IEEE J. Quantum Electron. **33**, 1455 (1997).
[79] C.-C. Chang, H. P. Sardesai, and A. M. Weiner, Opt. Lett. **23**, 283 (1998).
[80] O. Lopez, C. Daussy, A. Amy-Klein, C. Chardonnet, F. Narbonneau, M. Lours, and G. Santarelli, in International Workshop on Advances In Precision Tests and Experimental Gravitation in Space, 28 September 2006, Galileo Galilei Institute, Firenze, Italy, pp. 14–16.
[81] P. Ciprut, N. Gisin, R. Passy, J. P. Von der Weid, F. Prieto, and C. W.

Zimmer, J. Lightwave Technol. **16**, 757 (1998).
[82] C. Daussy *et al.*, Phys. Rev. Lett. **94**, 203904 (2005).
[83] S. Diddams' and L. Hollberg's Laboratories, Time and Frequency Division, NIST.
[84] http://www.branfiber.net
[85] J. Kitching made the initial setup of this system.
[86] J. Ye *et al.*, J. Opt. Soc. Am. B **20**, 1459 (2003).
[87] E. N. Ivanov, design provided by S. A. Diddams.
[88] C. W. Oates, E. A. Curtis, and L. Hollberg, Opt. Lett. **25**, 1603 (2000).
[89] S. G. Porsev and A. Derevianko, Phys. Rev. A **69**, 042506 (2004).
[90] R. Santra, K. V. Christ, and C. H. Greene, Phys. Rev. A **69**, 042510 (2004).
[91] T. Ido and H. Katori, Phys. Rev. Lett. **91**, 053001 (2003).
[92] H. Katori, T. Ido, and M. Kuwata-Gonokami, J. Phys. Soc. Jpn. **68**, 2479 (1999).
[93] J. Ye, D. W. Vernooy, and H. J. Kimble, Phys. Rev. Lett. **83**, 4987 (1999).
[94] H. J. Kimble, C. J. Hood, T. W. Lynn, H. Mabuchi, D. W. Vernooy, and J. Ye, in *Laser Spectroscopy XIV, Proceedings of the International Laser Spectroscopy Conference 1999*, Proceedings of ICOLS'99, edited by R. Blatt, J. Eschner, D. Leibfried, and F. Schmidt-Kaler (World Scientific, Singapore, 1999), p. 80.
[95] I. Courtillot, R. P. Quessada, P. Kovacich, A. Brusch, D. Kolker, J. Zondy, G. D. Rovera, and P. Lemonde, Phys. Rev. A **68**, 030501(R) (2003).
[96] T. Hong, C. Cramer, E. Cook, W. Nagourney, and E. N. Fortson, Opt. Lett. **30**, 2644 (2005).
[97] R. Le Targat, X. Baillard, M. Fouché, A. Brusch, O. Tcherbakoff, G. D. Rovera, and P. Lemonde, Phys. Rev. Lett. **97**, 130801 (2006).
[98] M. Takamoto, F.-L. Hong, R. Higashi, Y. Fujii, M. Imae, and H. Katori, J. Phys. Soc. Jpn. **75**, 104302 (2006).
[99] M. M. Boyd, A. D. Ludlow, S. Blatt, S. M. Foreman, T. Ido, T. Zelevinsky, and J. Ye, Phys. Rev. Lett. **98**, 083002 (2007).
[100] T. M. Fortier, D. J. Jones, and S. T. Cundiff, Opt. Lett. **28**, 2198 (2003).
[101] D. J. Jones, S. A. Diddams, J. K. Ranka, A. Stentz, R. S. Windeler, J. L. Hall, and S. T. Cundiff, Science **288**, 635 (2000).
[102] M. M. Boyd, T. Zelevinsky, A. D. Ludlow, S. M. Foreman, S. Blatt, T. Ido, and J. Ye, Science **314**, 1430 (2006).
[103] T. P. Krisher, L. Maleki, G. F. Lutes, L. E. Primas, R. T. Logan, J. D. Anderson, and C. M. Will, Phys. Rev. D **42**, 731 (1990).
[104] M. Calhoun, R. Sydnor, and W. Diener, The Interplanetary Network Progress Report, 15 February 2002, pp. 48–142.
[105] T. P. Celano, S. R. Stein, G. A. Gifford, B. A. Mesander, and B. J. Ramsey, in *Proceedings of the 2002 IEEE International Frequency Control Symposium and PDA Exhibition* (Institute of Electrical and Electronics Engineers, Piscataway, NJ, 2002), pp. 510–516.
[106] L. S. Ma, P. Jungner, J. Ye, and J. L. Hall, Opt. Lett. **19**, 1777 (1994).
[107] J. C. Bergquist, W. M. Itano, and D. J. Wineland, in *Frontiers in Laser Spectroscopy, Proceedings of the International School of Physics*, "Enrico Fermi:" Course 120, edited by T. W. Hänsch and M. Inguscio (North Holland, Amsterdam, 1994), pp. 359–376.
[108] M. Musha, Y. Sato, K. Nakagawa, K. Ueda, A. Ueda, and M. Ishiguro, Appl. Phys. B: Lasers Opt. **82**, 555 (2006).
[109] J. Ye, L. Robertsson, S. Picard, L.-S. Ma, and J. L. Hall, IEEE Trans. Instrum. Meas. **48**, 544 (1999).
[110] J. L. Hall, L. S. Ma, M. Taubman, B. Tiemann, F. L. Hong, O. Pfister, and J. Ye, IEEE Trans. Instrum. Meas. **48**, 583 (1999).
[111] R. J. Jones, W.-Y. Cheng, K. W. Holman, L. Chen, J. L. Hall, and J. Ye, Appl. Phys. B: Lasers Opt. **74**, 597 (2002).
[112] *Note added in proof.* Direct cw optical transfer has now been demonstrated at instability of 6×10^{-18} at 1 s, with $\approx$0.1 fs of timing jitter integrated from 15 to 10 MHz. (Foreman *et al.*, to be published)
[113] T. Udem, R. Holzwarth, and T. W. Hänsch, Nature **416**, 233 (2002).
[114] K. W. Holman, R. J. Jones, A. Marian, S. T. Cundiff, and J. Ye, Opt. Lett. **28**, 851 (2003).
[115] K. W. Holman, R. J. Jones, A. Marian, S. T. Cundiff, and J. Ye, IEEE J. Sel. Top. Quantum Electron. **9**, 1018 (2003).
[116] J. K. Ranka, R. S. Windeler, and A. J. Stentz, Opt. Lett. **25**, 25 (2000).
[117] R. Ell *et al.*, Opt. Lett. **26**, 373 (2001).
[118] U. Morgner, R. Ell, G. Metzler, T. R. Schibli, F. X. Kärtner, J. G. Fujimoto, H. A. Haus, and E. P. Ippen, Phys. Rev. Lett. **86**, 5462 (2001).
[119] T. M. Ramond, S. A. Diddams, L. Hollberg, and A. Bartels, Opt. Lett. **27**, 1842 (2002).
[120] T. Udem, J. Reichert, R. Holzwarth, and T. W. Hänsch, Phys. Rev. Lett. **82**, 3568 (1999).
[121] 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, Opt. Lett. **30**, 570 (2005).
[122] J. Rauschenberger, T. M. Fortier, D. J. Jones, J. Ye, and S. T. Cundiff, Opt. Express **10**, 1404 (2002).
[123] H. Hundertmark, D. Wandt, C. Fallnich, N. Haverkamp, and H. Telle, Opt. Express **12**, 770 (2004).
[124] F. Tauser, A. Leitenstorfer, and W. Zinth, Opt. Express **11**, 594 (2003).
[125] B. R. Washburn, S. A. Diddams, N. R. Newbury, J. W. Nicholson, M. F. Yan, and C. G. JØrgensen, Opt. Lett. **29**, 250 (2004).
[126] F.-L. Hong, K. Minoshima, A. Onae, H. Inaba, H. Takada, A. Hirai, H. Matsumoto, T. Sugiura, and M. Yoshida, Opt. Lett. **28**, 1516 (2003).
[127] D. J. Jones, K. W. Holman, M. Notcutt, J. Ye, J. Chandalia, L. A. Jiang, E. P. Ippen, and H. Yokoyama, Opt. Lett. **28**, 813 (2003).
[128] R. K. Shelton, L.-S. Ma, H. C. Kapteyn, M. M. Murnane, J. L. Hall, and J. Ye, Science **293**, 1286 (2001).
[129] K. W. Holman, D. J. Jones, J. Ye, and E. P. Ippen, Opt. Lett. **28**, 2405 (2003).
[130] C. M. DePriest, T. Yilmaz, P. J. Delfyett, Jr., S. Etemad, A. Braun, and J. Abeles, Opt. Lett. **27**, 719 (2002).
[131] F. Quinlan, S. Gee, S. Ozharar, and P. J. Delfyett, Opt. Lett. **31**, 2870 (2006).
[132] S. A. Diddams, D. J. Jones, J. Ye, S. T. Cundiff, J. L. Hall, J. K. Ranka, R. S. Windeler, R. Holzwarth, T. Udem, and T. W. Hänsch, Phys. Rev. Lett. **84**, 5102 (2000).
[133] R. Holzwarth, T. Udem, T. W. Hänsch, J. C. Knight, W. J. Wadsworth, and P. S. J. Russell, Phys. Rev. Lett. **85**, 2264 (2000).
[134] S. M. Foreman, D. J. Jones, and J. Ye, Opt. Lett. **28**, 370 (2003).
[135] 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, Opt. Lett. **29**, 2806 (2004).
[136] M. Zimmermann, C. Gohle, R. Holzwarth, T. Udem, and T. W. Hänsch, Opt. Lett. **29**, 310 (2004).
[137] J. D. Jost, J. L. Hall, and J. Ye, Opt. Express **10**, 515 (2002).
[138] T. R. Schibli, K. Minoshima, F.-L. Hong, H. Inaba, Y. Bitou, A. Onae, and H. Matsumoto, Opt. Lett. **30**, 2323 (2005).
[139] R. J. Jones, K. D. Moll, M. J. Thorpe, and J. Ye, Phys. Rev. Lett. **94**, 193201 (2005).
[140] C. Gohle, T. Udem, M. Hermann, J. Rauschenberger, R. Holzwarth, H. A. Schuessler, F. Krausz, and T. W. Hänsch, Nature **436**, 234 (2005).
[141] K. W. Holman, D. J. Jones, D. D. Hudson, and J. Ye, Opt. Lett. **29**, 1554 (2004).
[142] K. W. Holman, D. D. Hudson, J. Ye, and D. J. Jones, Opt. Lett. **30**, 1225 (2005).
[143] D. D. Hudson, S. M. Foreman, S. T. Cundiff, and J. Ye, Opt. Lett. **31**, 1951 (2006).
[144] H. Tsuchida, Opt. Lett. **24**, 1434 (1999).
[145] L.-P. Chen, Y. Wang, and J.-M. Liu, IEEE J. Quantum Electron. **32**, 1817 (1996).
[146] E. N. Ivanov, S. A. Diddams, and L. Hollberg, IEEE Trans. Ultrason. Ferroelectr. Freq. Control **50**, 355 (2003).
[147] D. D. Hudson, K. W. Holman, R. J. Jones, S. T. Cundiff, J. Ye, and D. J. Jones, Opt. Lett. **30**, 2948 (2005).
[148] D. Kleppner, Phys. Today **59**, 10 (2006).
[149] J. Ye, Opt. Lett. **29**, 1153 (2004).