

Clock Synchronization with Weakly Correlated Photons

Justin Yu Xiang Peh,¹ Darren Ming Zhi Koh,¹ Zifang Xu,¹ Xi Jie Yeo,¹ Peng Kian Tan,¹ and Christian Kurtsiefer^{1,2,*}

¹Centre for Quantum Technologies, 3 Science Drive 2, Singapore 117543

²Department of Physics, National University of Singapore, 2 Science Drive 3, Singapore 117542

(Dated: February 24, 2026)

Clock synchronization is necessary for communication and distributed computing tasks. Previous schemes based on photon timing correlations use pulsed light or photon pairs for their strong timing correlations. In this work, we demonstrate successful synchronization of crystal clocks using weakly time-correlated photons of 180 ns coherence time from a bunched light source. A synchronization timing jitter of 10 ns is achieved over symmetric -102 dB optical channel loss between two parties, over a span of 25 hours. We also present a model that gives better estimates to the coherence peak finding success probabilities under low signal.

I. INTRODUCTION

Clock synchronization is used in everyday tasks such as navigation and distributed computing. This is commonly implemented using the Network Time Protocol or global navigation satellite system (e.g., GPS) time synchronization, achieving precision of milliseconds or tens of nanoseconds, respectively [1, 2]. Quantum communication protocols also require clock synchronization, but on the order of nano- to pico-seconds, and are typically achieved either using pulsed laser light [3] or with classical signals on dedicated optical or electronic channels [4, 5].

Modern quantum communication systems can use the resources present in the protocol itself to perform the clock synchronization, such as photons arriving at fixed timing intervals, or photon pairs from spontaneous parametric down-conversion (SPDC) [6, 7]. This has been demonstrated with frequency standards such as Rubidium (Rb) clocks, which provide a long-term frequency stability of <1 ppb/day. In comparison, crystal oscillators without temperature stabilization have frequency stability of only 100 ppb/day.

Recent works have shown that Rb clocks are unnecessary and crystal oscillators are sufficient, by relying on the strong timing correlation (\sim ps) of SPDC photon pairs [8]. Weak coherent photon pulses were also proposed for clock frequency transfer [9]. In both cases, the cross-correlation peak is very strong, i.e., $g^{(2)}(\tau = 0) \gg 1$, with resolution generally limited only by the timing jitter of the generation and detection optoelectronics.

Timing correlations also exist in thermal light, which has been used in applications such as ghost imaging [10] and range finding [11]. These timing correlations arise from temporal photon bunching, also known as the Hanbury-Brown-Twiss effect [12]. In particular, the use of thermal light opens up the potential for distributed clock synchronization due to photon bunching being preserved across arbitrary partitioning, and makes it hard

to manipulate the timing information due to the randomness of the individual photon timing.

However, unlike photon pairs, identification of this photon bunching peak in the correlation $g^{(2)}(\tau)$ is more challenging due to the low signal (i.e., $g^{(2)}(\tau) \leq 2$) as well as fluctuations about the background of $g^{(2)} = 1$. An earlier proposal suggests the use of thermal light for clock synchronization by relying on the use of low efficiency two-photon absorption in single-photon detectors to resolve the bunching characteristic [13, 14]; such a scheme has yet to be demonstrated.

Here, we demonstrate clock synchronization using one such a weakly timing-correlated light source with $g^{(2)}(\tau = 0) = 1.42$ and coherence time $\tau_c = 180$ ns, over two transmission channels with a loss of -102 dB each. We reach an accuracy of 10 ns, and continuously track the frequency drift over a period of 25 hours to demonstrate its stability. Our scheme only requires single-photon detection and crystal oscillators as reference clocks. We additionally derive the probability of achieving clock synchronization over transmission channels, which is applicable across different time-correlated photon sources, including SPDC pair sources and thermal light sources.

II. CLOCK SYNCHRONIZATION WITH BUNCHED LIGHT

The experiment is shown Fig.1(a). Our source of bunched light is formed by laser light of 780 nm wavelength sent into an unbalanced Mach-Zehnder interferometer with optical delay longer than the laser coherence time τ_c [15] (see Appendix A for a more detailed explanation). Phase fluctuations in the laser result in an intensity correlation at the output, that can be measured through the second-order coherence function $g^{(2)}(\tau)$, given by

$$g^{(2)}(\tau) = 1 + \frac{1}{2} \exp\left(-\frac{2|\tau|}{\tau_c}\right). \quad (1)$$

The brightness of the source is 1.55(2) mW, with a measured coherence peak of $g^{(2)}(\tau = 0) = 1.42(1)$ and coherence time of $\tau_c = 180(6)$ ns, as shown in Fig.1(b).

* christian.kurtsiefer@gmail.com

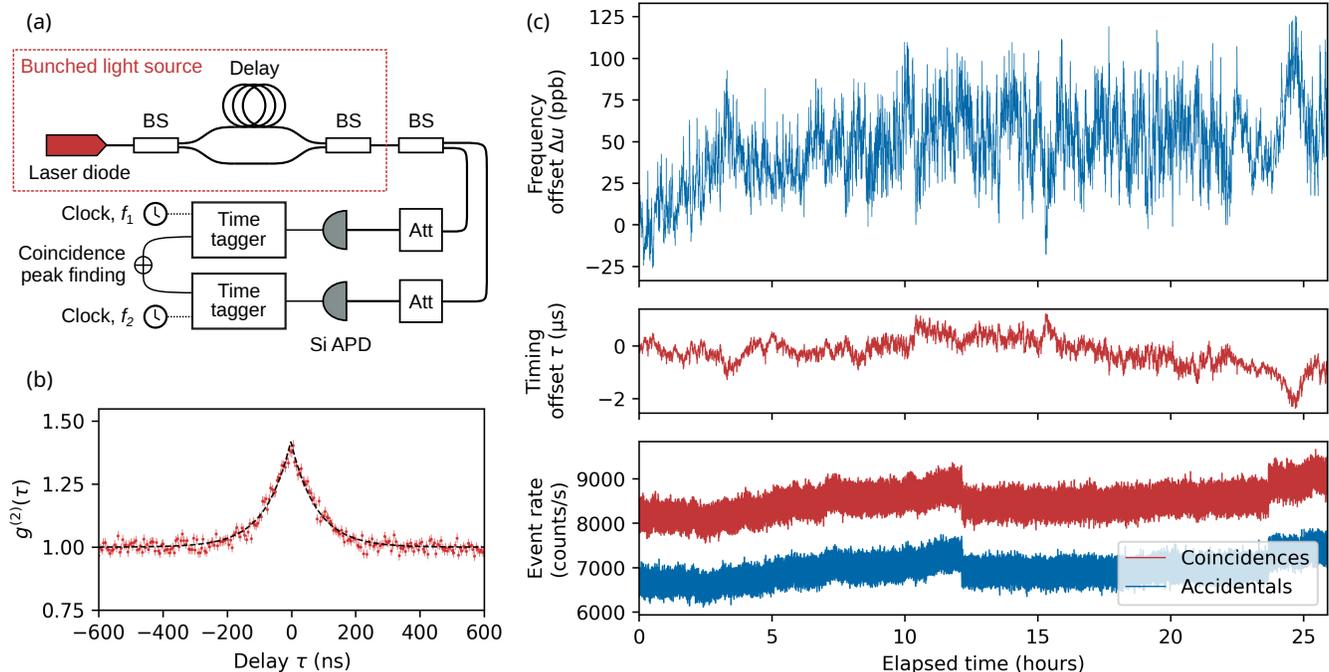


FIG. 1. (a) Simplified experimental setup for clock synchronization. Bunched light is sent to two different parties through separate -102 dB channels, before detection by Silicon avalanche photodetectors (Si APDs) and timestamping by time taggers referenced to independent crystal oscillators of frequencies f_1 and f_2 . Singles count rate s_1 and s_2 of approximately 200 kcounts/s are recorded on each side. BS: beamsplitter; Att: attenuators. (b) Second-order coherence function $g^{(2)}(\tau)$ of the bunched light source, with the curve fit in black corresponding to $g^{(2)}(0) = 1.42(1)$ peak and coherence time $\tau_c = 180(6)$ ns. Error bars correspond to Poissonian errors from counting statistics. (c) Long-term frequency offset ($1 + \Delta u = f_1/f_2$), timing offset τ , and event rates during a 25 hour clock synchronization run between independent time taggers running on separate quartz crystal oscillators, after an initial frequency correction of 4.0 ppm. Correlation peak tracking is performed using a 256 ns coincidence window. The sharp spike and dip in event rates are attributed to laser mode hopping.

The light from the source is shared between two parties with optical fiber channel transmissions of $-101.9(4)$ dB and $-102.2(4)$ dB, respectively, using optical attenuators. Photon arrival events are detected using Silicon avalanche photodetectors (APDs) on each side, with count rates of 192 kcounts/s and 182 kcounts/s without correcting for dark counts and afterpulsing.

The detection events are read by independent time taggers disciplined by different free-running 10 MHz crystal oscillators. The generated timestamps are then exchanged for coincidence peak finding and tracking in real-time. The peak tracking program continuously serves a timing offset between each party by pairing photon detection events between both parties within a 256 ns coincidence window, over the span of the 25 hour measurement. The average coincidence rate of 8500 events/s is consistently higher than the accidental count rate of 7000 events/s, indicative of successful frequency tracking, as shown in Fig. 1(c).

Crystal clocks drift in frequency due to temperature fluctuations and electronic noise. We are able to reconstruct the resulting frequency offset $\Delta u'$ by monitoring the drift in served timings using individual samples span-

ning 10.74 s, with a resolution of 0.537 s (see Fig. 1(c)).

The actual frequency offset Δu was obtained via a heterodyne measurement using identical copies of the clock signals with integration time of 10 s, shown in Fig. 2. A maximum instantaneous frequency difference of 35.4 ppb, corresponding to an average drift of 3.3 ppb/s, was observed.

We estimate the error in the reconstructed signal by performing linear interpolation on the measurement and subsequent differencing. The reconstructed frequency offset is found to be in good agreement with the measured frequency offset, with a root-mean-squared error of 3.2 ppb averaged over the 25 hour span.

III. PEAK FINDING

Clock synchronization can be decomposed into two parts: first identifying the initial clock frequency offset and time difference between two parties, then monitoring the timing drift. The former is performed by distributing photons with time-correlated statistics to each party, and finding the coincidence peak using cross-correlation to identify the time delay.

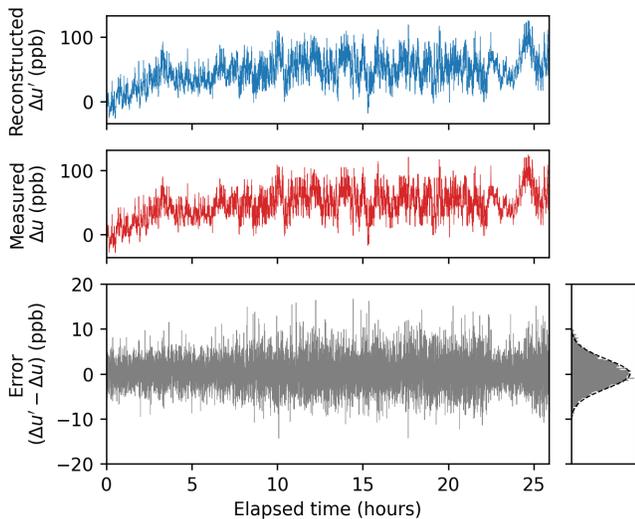


FIG. 2. A concurrent measurement of the actual frequency offset Δu between the two external 10 MHz clocks, as well as the corresponding error in the reconstructed frequency offset $\Delta u'$. The accuracy of the frequency served by the peak tracking algorithm is 3.2 ppb on average.

Efficient peak identification relies on the circular convolution theorem to compute the cross-correlation between two sets of reconstructed detection time traces $a[k]$ and $b[k]$,

$$g^{(2)}(\tau) \sim (a \star b)[k] = \mathcal{F}^{-1} \left\{ \overline{\mathcal{F}\{a\}} \cdot \mathcal{F}\{b\} \right\} [k], \quad (2)$$

under the discretization $\tau = k\delta t$ with $k \in \mathbb{Z}$ and time resolution δt , using the Fourier Transform \mathcal{F} and its inverse¹ [7]. The success of peak identification is constrained by singles rate, rate of true coincidences, and choice of optimal parameters for the Fast Fourier Transform (FFT) to compute the time delay: specifically, the number of time bins $N = 2^q$ ($q \in \mathbb{Z}^+$), and initial bin width δt .

We model the peak finding probability — in both cases, with and without a reliable frequency reference between both parties — and perform an exhaustive parameter scan to identify appropriate FFT parameters.

A. No frequency offset

The minimum acquisition time required for the cross-correlation to obtain a flat cross-correlation noise floor is $T = N\delta t$. This noise floor arises from accidental coincidences — attributed to noise sources, such as coincidences with uncorrelated detection events and dark

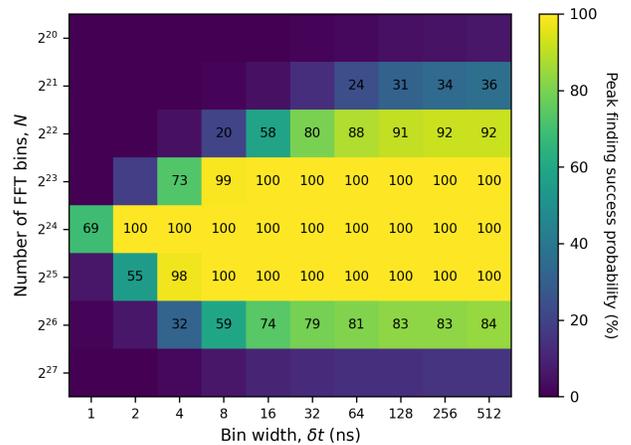


FIG. 3. Success probabilities of finding the correct peak position, numerically obtained by performing $>10^7$ Monte Carlo trials for each $(N, \delta t)$ pair to estimate and interpolate the cumulative distribution in Eqn. 5. The parameters are singles detection rate $s_1 = s_2 = 100$ kcounts/s, coincidence rate $c = 650$ counts/s, bin overlap $\nu = 0.5$, and frequency offset error of $\Delta u = 50$ ppb (precompensation step size of 100 ppb), corresponding to the setup in Fig. 1 with an additional 3 dB attenuation per channel.

counts — with an expected value of $C_a = (s_1 s_2 \delta t) T$, where s_1 and s_2 are the detection event rates on each channel. The detection events are well-approximated by a Poisson distribution after binning, so the observed accidental coincidences X_k in each time bin k also follows a Poisson distribution \mathcal{P} with mean $\lambda = C_a$, i.e., $X_{k \in \{1, \dots, N\}} \sim \mathcal{P}(X_k = x_k | \lambda = C_a)$.

The maximum observed value across all (identically distributed) time bins $X_{(N)} \equiv \max\{X_k\}$ is therefore given by the max-order Poisson distribution (see Appendix B) whose probability distribution is

$$f_{X_{(N)}}(x) = [F_X(x|\lambda)]^N - [F_X(x|\lambda) - f_X(x|\lambda)]^N,$$

where f_X and F_X correspond to the probability mass and cumulative distribution functions of a single bin.

The coincidence rate above background accidentals in a single time bin (denoted c_e) required to be identified as the highest peak in the cross-correlation is thus

$$c_e > \frac{1}{T} (X_{(N)} - X). \quad (3)$$

The coincidence rate per time bin can be maximized by setting the timing resolution δt to be of the same scale as the coherence time τ_c , so that most of the coincidence events fall within the same time bin, i.e., $\delta t \sim \tau_c$. Some of these events may fall into an adjacent time bin instead due to off-centered bins (since the exact time offset is not known *a priori*): this introduces an adjustment factor into Eqn. 3 representing the degree of bin overlap $\nu \in [0.5, 1]$, yielding

$$c_e > \frac{1}{\nu T} (X_{(N)} - X). \quad (4)$$

¹ Algorithms faster than FFT's $\mathcal{O}(n \log n)$ time complexity are also possible [16, 17], but require additional encoding of time-synchronization strings within the light pulses.

B. Variable frequency offset

In the case of two separate clocks with slightly different clock frequencies, the timing delay between photon arrivals between each party can drift by $\Delta\tau_{\text{drift}} \approx T\Delta u(t)$ due to the non-zero clock-frequency offset $\Delta u(t)$ after elapsed time T . For a sufficiently small T , the frequency offset can be approximated as a constant, i.e., $\Delta u(t) = \Delta u$.

In order to maximize the number of coincidences in a single time bin, the time bin should ideally be as wide as the timing drift, i.e., $\delta t \geq \Delta\tau_{\text{drift}}$, or in other words, $N \leq 1/\Delta u$. We model this as an additional correction factor $\xi \equiv \max\{1, N\Delta u\}$, and together with Eqn. 4 yields the minimum required true coincidence rate for successful peak finding given by

$$c_e > \frac{\xi}{\nu T} (X_{(N)} - X). \quad (5)$$

The surface of Eqn. 5 is shown in Fig. 3, by performing Monte Carlo simulations across different N and δt . Our model provides better estimations of the peak finding probability compared to previous works [7, 8] by avoiding the normal approximation to the noise in the FFT (see Appendix C).

The reduction in coincidence counts due to the presence of frequency offset can be mitigated either by choosing a smaller N , or by performing a frequency precompensation sweep [8] on the set of timestamps $\{t_i\}$ to reduce the apparent Δu between the two clocks and allow for larger N values. The compensation is given by the mapping

$$t_i \rightarrow \Delta t_i(1 + \Delta u) + t_{i-1} = t_i + \Delta t_i \Delta u, \quad (6)$$

where $\Delta t_i \equiv t_i - t_{i-1}$ is the separation between consecutive timestamps.

In practice, while the frequency offset between crystal clocks can be high (~ 10 ppm), the short-term stability of the clocks themselves is much better (< 10 ppb/s). This means the full precompensation sweep need not be repeated even if clock synchronization were lost; using the last estimated frequency offset Δu_0 as well as $\Delta u_0 \pm 100$ ppb is typically sufficient for peak recovery.

Once an initial peak has been found, the frequency and timing resolution can be further improved by repeating the respective corrections with progressively smaller time bins, until the desired resolution has been reached [7]. Our optimized implementations of peak finding and frequency compensation are open-sourced [18].

IV. PEAK TRACKING

The use of crystal oscillators as frequency references results in changes in frequency offset over time, due to their long-term stability in $\Delta u(t)$ of up to 100 ppb/day. This causes a drift in the correlation peak position over time which needs to be tracked.

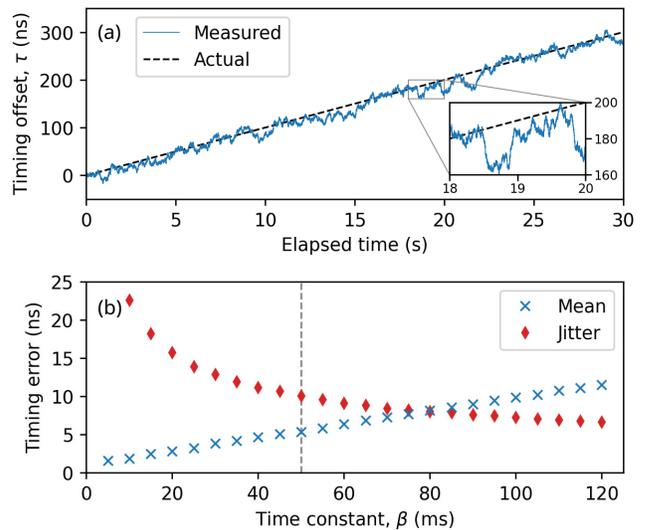


FIG. 4. Characterization of peak tracking accuracy under a constant frequency offset $\Delta u = 10$ ppb, with a time constant $\beta = 50$ ms chosen for the exponential moving average filter used in the experiment. (a) Measurement of reported timing offsets (blue) against the actual offset (black) over a period of 30 seconds. The inset highlights a ≈ 30 ns timing error over a 0.5-second window, which suggests that a histogram fit does not improve accuracy. (b) Measurements of average offset error and jitter across different time constants β , over a measurement period of 10 minutes. Timing jitter is equivalent to the root-mean-squared error, and corresponds to 10 ns with $\beta = 50$ ms (marked in gray dashed line). Peak tracking fails with time constants less than 5 ms.

Tracking can be performed by looking for coincidences within a sufficiently wide coincidence window, such that drifts in the peak are captured. However, when using bunched light for peak tracking, coincidences events are effectively dominated by background accidentals due to the low $g^{(2)}(\tau)$ of the light source. Directly returning each time difference found will result in the peak being quickly lost.

We apply a smoothing operation using an exponential moving average filter to introduce damping against events far from the current estimated peak position. This effectively minimizes noise fluctuations and allows for the peak to be tracked.

In order to quantify the accuracy of the peak tracking, we supply a constant 10 ppb frequency offset between the two timestamps using a function generator and derive the true timing offset. Using a time constant $\beta = 50$ ms for the moving average filter, a timing jitter of 10 ns can be obtained, as shown in Fig. 4, which is an order of magnitude better than the $\tau_c = 180(6)$ ns coherence time of the source. The jitter was tested to remain the same even under larger frequency offsets of 50 ppb as well.

A tracking lag is also present and can be seen in Fig. 4, where the measurement best-fit line is slightly displaced from the true timing offset. This is because the frequency offset is non-zero and constant, and manifests as a mean

timing error of around 6 ns at $\beta = 50$ ns. In the case of free-running clocks after frequency compensation, this lag is expected to be insignificant due to the frequency offset fluctuating about zero.

We also make a small note regarding the use of timing offset histograms to improve timing resolution. For photon pair and pulsed sources, higher timing resolution can be achieved by collecting a histogram within a small timing window and extracting the peak from a normal distribution curve fit, due to the narrow peak coherence signature of $\tau_c \ll 1$ ns [8]. However, for our bunched light source with a significantly noisier coherence peak signature, this technique does not improve the timing accuracy, as can be seen in the served timing offset in the inset of Fig. 4 being consistently far from the mean.

Active frequency compensation is also performed online to keep the frequency offset small, to avoid peak tracking loss since the correlation peak is less likely to drift out of the coincidence window. We do this by estimating the frequency offset Δu from the rate of change in timing drift (see Appendix D), then performing timing compensation for each timestamp using the estimated frequency offset, identical to that of Eqn. 6.

V. CONCLUSION

We demonstrated successful clock synchronization between two parties over a symmetrical -102 dB transmission channel, using a bunched light source of coherence time $\tau_c = 180(6)$ ns. Peak tracking was performed online over a span of 25 hours with active frequency compensation and an exponential moving average filter of time constant $\beta = 50$ ns, achieving an overall timing jitter of 10 ns.

Previous papers on clock synchronization performed using photon pair sources rely on a peak significance metric for peak finding [7, 8], used as a threshold for quantifying the probability of the cross-correlation peak being attributed to noise. While it remains a useful metric for estimating peak location, this underestimates the peak finding probabilities under low signal conditions with small time bins due to the use of normal approximation. We develop a model that accounts for the Poisson nature of the coincidences, and calculate instead the true probability of the cross-correlation peak being the

signal, as well as derive optimal FFT parameters for a given true coincidence and singles rate. This remains applicable even when using any other sources of timing correlations, including photon pairs and thermal light.

The clock synchronization scheme in this work can be directly applied to protocols that use photon bunching as a resource, or indirectly by software timing compensation using the served timing and frequency offsets. Direct frequency compensation can also be performed by actively tuning the crystal oscillator frequencies, so that the clock signals themselves can be utilized as part of a clock distribution network.

We can also take advantage of the fact that the second-order coherence is preserved across arbitrary partitioning of the light, to distribute the signal amongst multiple parties in a star topology. Since splitting light into two separate channels incurs an additional insertion loss of 3 dB per channel, clock distribution to 2^n parties can be achieved with only $-3n$ dB of additional loss per channel, e.g., a 128-party setup with this source incurs about -80 dB.

This work paves the way towards clock synchronization using telecommunication O-band and C-band bunched light, which will be able to propagate with minimal chromatic dispersion over longer distances, e.g., 500 km of G.652 telecommunication fiber with 0.2 dB/km optical loss at 1550 nm. This additionally opens up the possibility of using erbium-doped fiber amplifiers to amplify the correlation signal. While timing offset changes as a result of thermal expansion in optical fibers will need to be addressed, this effect is fundamentally indistinguishable from a clock frequency drift in this scheme and can hence be compensated for in the same manner.

VI. ACKNOWLEDGEMENTS

The authors would like to thank Ng Boon Long for helpful discussions in setting up the concurrent frequency measurement. This research is supported by the National Research Foundation, Singapore through the National Quantum Office, hosted in A*STAR, under its Centre for Quantum Technologies Funding Initiative (S24Q2d0009) and the Quantum Engineering Programme 3.0 under grant W25Q3D0004.

-
- [1] D. L. Mills, *Computer network time synchronization: the network time protocol* (CRC press, 2006).
 - [2] B. A. Renfro, M. Stein, E. B. Reed, and E. J. Villalba, An analysis of global positioning system standard positioning service performance for 2020 (2022).
 - [3] F. R. Giorgetta, W. C. Swann, L. C. Sinclair, E. Baumann, I. Coddington, and N. R. Newbury, Optical two-way time and frequency transfer over free space, *Nature Photonics* **7**, 434 (2013).
 - [4] T. Gerrits, I. A. Burenkov, Y. Li-Baboud, A. Rahmouni, D. M. Anand, Hala, O. Slattery, A. Battou, and S. V. Polyakov, White Rabbit-assisted quantum network node synchronization with quantum channel coexistence, in *Conference on Lasers and Electro-Optics* (Optica Publishing Group, San Jose, California, 2022) p. FM1C.2.
 - [5] M. Alshoukan, P. G. Evans, B. P. Williams, N. S. V. Rao, C. E. Marvinney, Y.-Y. Pai, B. J. Lawrie, N. A. Peters, and J. M. Lukens, Synchronizing a quantum local

area network with White Rabbit, in *Conference on Lasers and Electro-Optics* (Optica Publishing Group, San Jose, California, 2022) p. FM1C.4.

- [6] J. Lee, L. Shen, A. Cerè, J. Troupe, A. Lamas-Linares, and C. Kurtsiefer, Symmetrical clock synchronization with time-correlated photon pairs, *Applied Physics Letters* **114** (2019).
- [7] C. Ho, A. Lamas-Linares, and C. Kurtsiefer, Clock synchronization by remote detection of correlated photon pairs, *New Journal of Physics* **11**, 045011 (2009).
- [8] C. Spiess, S. Töpfer, S. Sharma, A. Kržič, M. Cabrejo-Ponce, U. Chandrashekhara, N. L. Döll, D. Rieländer, and F. Steinlechner, Clock Synchronization with Correlated Photons, *Physical Review Applied* **19**, 054082 (2023).
- [9] C. Spiess and F. Steinlechner, Clock synchronization with pulsed single photon sources, *Quantum Science and Technology* **9**, 015019 (2023).
- [10] R. S. Bennink, S. J. Bentley, and R. W. Boyd, Two-photon coincidence imaging with a classical source, *Physical Review Letters* **89**, 113601 (2002).
- [11] P. K. Tan, X. J. Yeo, A. Z. W. Leow, L. Shen, and C. Kurtsiefer, Practical range sensing with thermal light, *Physical Review Applied* **20**, 014060 (2023).
- [12] R. H. Brown and R. Q. Twiss, Correlation between photons in two coherent beams of light, *Nature* **177**, 27 (1956).
- [13] J. Zhu, P. Huang, X. Xiao, and G. Zeng, A new clock synchronization scheme based on the second-order coherence of thermal light, *Optica Applicata* **43** (2013).
- [14] G. Xu, X. Ren, Q. Miao, M. Yan, H. Pan, X. Chen, G. Wu, and E. Wu, Determination of the contributions of degenerate and non-degenerate two-photon absorption response in silicon avalanche photodiode for infrared photon detection, *Journal of Modern Optics* **67**, 1321 (2020).
- [15] X. J. Yeo, D. M. Z. Koh, J. Y. X. Peh, C. Kurtsiefer, and P. K. Tan, Near-lossless method for generating thermal photon-bunched light (2026), arXiv:2602.00633 [quant-ph].
- [16] L. Calderaro, A. Stanco, C. Agnesi, M. Avesani, D. Dequal, P. Villoresi, and G. Vallone, Fast and Simple Qubit-Based Synchronization for Quantum Key Distribution, *Physical Review Applied* **13**, 054041 (2020).
- [17] J. Krause, N. Walenta, J. Hilt, and R. Freund, Clock-offset recovery with sublinear complexity enables synchronization on low-level hardware for quantum key distribution, *Physical Review Applied* **23**, 044015 (2025).
- [18] J. Y. X. Peh, Z. Xu, and C. Kurtsiefer, fpfind, <https://github.com/s-fifteen-instruments/fpfind> (2024).
- [19] X. J. Yeo, E. Ernst, A. Leow, J. Hwang, L. Shen, C. Kurtsiefer, and P. K. Tan, Direct measurement of the coherent light proportion from a practical laser source, *Physical Review A* **109**, 013706 (2024).

Appendix A: Bunched light source and experimental setup

The experimental setup in this work uses light from a distributed-feedback (DFB) laser diode of wavelength 780 nm, coupled into 780-HP fiber to project into a single optical mode. This light is sent into an unbalanced fiber-based Mach-Zehnder interferometer with 400 m delay fiber. This corresponds to a $\Delta = 2 \mu\text{s}$ delay, which is

longer than the coherence time $\tau_c = 180 \text{ ns}$.

Assuming that the power and polarization in both arms match, the output field from one of the interferometer ports is given by

$$E(t) = \frac{1}{2}[E_0(t) + E_0(t + \Delta)],$$

where $E_0(t)$ is the electric field of the laser light. The overall second-order coherence function thus gives

$$\begin{aligned} g^{(2)}(\tau) &= \frac{1}{4} \left[g_0^{(2)}(\tau + \Delta) + g_0^{(2)}(\tau - \Delta) \right] \\ &\quad + \frac{1}{2} \left[g_0^{(2)}(\tau) + |g_0^{(1)}(\tau)|^2 \right] \\ &= 1 + \frac{1}{2} \exp\left(-\frac{2|\tau|}{\tau_c}\right), \end{aligned}$$

where $g_0^{(2)}(\tau) = 1$ is the second-order coherence function of the laser light and $g_0^{(1)}(\tau) = \exp(-|\tau|/\tau_c)$ is the first-order coherence function of the laser light, assuming a Lorentzian spectral profile. The peak at $g^{(2)}(\tau = 0)$ is a manifestation of the photon bunching effect that saturates at 1.5 in our setup. Our measured peak value of 1.42(1) is attributed to the fact that the laser is not fully coherent [19] and other experimental imperfections. A more in-depth explanation and discussion about the source can be found in [15].

A second BS is used to further split the light into two symmetrical channels for downstream detection using fiber-pigtailed active-quenched Si avalanche photodiodes (Excelitas SPCM-800-10-FC). Attenuation in each channel is achieved by cascading fiber beamsplitters as well as fiber decoupling at the mating sleeves, and is measured to be relatively stable over multiple days.

Time tagging is subsequently performed by timestamp devices (S-Fifteen Instruments TDC2) with a $1\text{-}\sigma$ timing jitter of 20 ps. The clock to each timestamp device is supplied by external 10 MHz crystal oscillators without any temperature stabilization.

The actual frequency offset between the two clocks was measured concurrently with the clock synchronization experiment. The 10 MHz signal was duplicated and combined in a frequency mixer (Mini-Circuits ZFM-2+) followed by a low-pass filter (Mini-Circuits SLP-2.5+, 2.5 MHz cut-off). The mixed signal was then sampled by an oscilloscope at a rate of 2.5 ksamples/s over 10 s, which performed a 2^{14} -bin FFT (nominal resolution of approximately 0.15 Hz) with a Hann window.

The frequency offset error is then calculated by measuring the difference between the served frequency offset (from the frequency estimation step in the clock synchronization) and the measured frequency offset with linear interpolation. The histogram of offset errors is fitted using a Gaussian probability density function, obtaining a standard deviation of 3.2 ppb as seen in Fig. 2.

Appendix B: Derivation of max-order distribution

The FFT for peak finding according to Eqn. 2 is performed with N bins and accidental coincidences in each of these bins follow the Poisson distribution \mathcal{P} with mean λ . The maximum value across all bins increases with the number of bins. The distribution of the max-order statistic (i.e., the maximum of all bins) with respect to the number of coincidences x and number of bins N of mean value λ is denoted $f_{(N)}(x|\lambda)$.

Since each bin is assumed to be independent and identically distributed, the corresponding max-order cumulative distribution function (CDF) can be expressed as the product of the CDF of individual bins,

$$F_{(N)}(x|\lambda) = \underbrace{F(x|\lambda) \times \dots \times F(x|\lambda)}_{N \text{ times}} = [F(x|\lambda)]^N.$$

We therefore express the max-order probability mass function (PMF) in terms of the Poisson PMF and CDF of a single bin,

$$\begin{aligned} f_{(N)}(x|\lambda) &= F_{(N)}(x|\lambda) - F_{(N)}(x-1|\lambda) \\ &= [F(x|\lambda)]^N - [F(x-1|\lambda)]^N \\ &= [F(x|\lambda)]^N - [F(x|\lambda) - f(x|\lambda)]^N. \end{aligned}$$

Notably, this form is already tenable for direct computation without further simplification, even though computing the difference of n^{th} -powers generally causes catastrophic cancellation due to floating-point rounding errors. The max-order Poisson PMF distribution width heuristically scales roughly with $\sqrt{\lambda}$, which requires the PMF to be accurate to at least $1/\sqrt{\lambda}$, e.g., $\lambda \leq 10^4$ needs at least $\sim 10^{-2}$ accuracy. Since $F(x|\lambda) \in [0, 1]$ and the fact that exponentiation can be easily applied for large N using numerical techniques (such as exponentiation-by-squaring), the rounding errors can be minimized to near the floating-point precision, e.g., $\sim 10^{-16}$ for 64-bit floats.

For larger $\lambda > 10^4$ (i.e., high accidental coincidence rates per time bin), the normal approximation for each bin remains appropriate,

$$f(x) \sim \mathcal{P}(x|\lambda) \approx \mathcal{N}(x|\mu = \lambda, \sigma = \sqrt{\lambda}), \quad (\text{B1})$$

and the corresponding max-order probability distribution function follows a more tractable form for numerical computation,

$$\begin{aligned} f_{(N)}(x|\mu, \sigma) &= \frac{d}{dx} F_{(N)}(x|\mu, \sigma) \\ &= \frac{d}{dx} [F(x|\mu, \sigma)]^N \\ &= N f(x|\mu, \sigma) [F(x|\mu, \sigma)]^{N-1}. \end{aligned}$$

This in fact corresponds to the first-order term for the

discrete case after a binomial expansion,

$$\begin{aligned} f_{(N)}(x|\lambda) &= [F(x|\lambda)]^N - [F(x|\lambda) - f(x|\lambda)]^N \\ &= [F(x|\lambda)]^N - \left[[F(x|\lambda)]^N \right. \\ &\quad \left. - N f(x|\lambda) [F(x|\lambda)]^{N-1} + \dots \right] \\ &= N f(x|\lambda) [F(x|\lambda)]^{N-1} - \mathcal{O}(f^2 F^{N-2}), \end{aligned}$$

noting that $f(x|\lambda)/F(x|\lambda) \ll 1$ at large λ .

Appendix C: Comparison of peak finding probabilities with Poisson / normal approximation

To understand the effect of normal approximation on peak finding success probabilities, it is useful to first understand the use of the max-order distribution.

Previous models derive a statistical significance S representing the number of standard deviations σ_X the coincidence peak value deviates from the mean number of accidental coincidences X (of a *single* time bin). Ignoring any corrections for frequency difference, the significance shares the form

$$S = \frac{c_e T}{\sigma_X} = \frac{c_e T}{\sqrt{s_1 s_2 T \delta t}} = c_e \sqrt{\frac{N}{s_1 s_2}} \quad (\text{C1})$$

using notation consistent with this paper (corresponding to Eqn. 8 of [7] and Eqn. A3 of [8]). Notably, this equation only depends on the number of time bins N and has no dependence on time bin width δt .

The search for the true coincidence peak is more concerned with distinguishing the true coincidence peak from the *next highest* peak arising from noise, and less to do with the average fluctuations of a single time bin. Since the number of accidental coincidences X in each time bin is stochastic, the maximum observed noise peak across all time bins intuitively increases with increasing number of observations (i.e., number of bins N). This maximum is distributed according to the order statistics $X_{(N)}$ of all time bins.

In contrast, the derivation in Eqn. C1 implicitly assumes that the noise peak is given only by the accidental coincidences X from a single time bin. The use of statistical significance is thus more accurately described as an indication of the *peak misidentification rate*: the probability that the true coincidence peak is misidentified as noise, given that the location of said peak is *already known*. Our derivation of Eqn. 5 in the main text, which accounts for $X_{(N)}$, yields the actual peak identification rate.

Due to the positive excess kurtosis (i.e., “fatter tail”) of the Poisson distribution relative to the normal distribution, the corresponding max-order distribution $f_{(N)}(x|\lambda)$ is skewed more towards higher peak values. At low coincidence rates per time bin ($\lambda \lesssim 10^4$) occurring at small

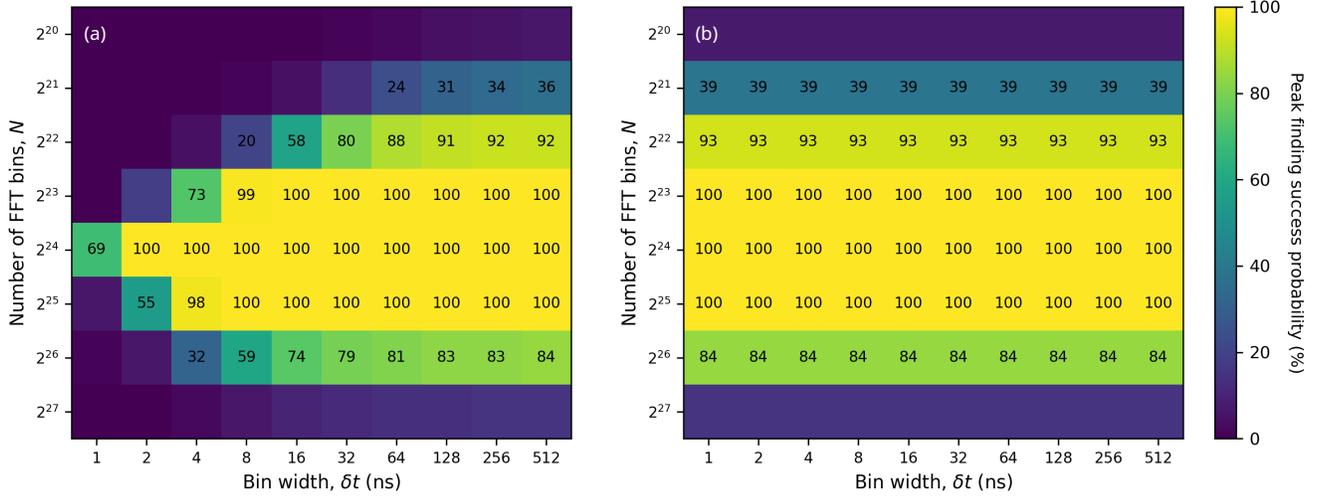


FIG. 5. (a) Success probabilities of finding the correct peak position by solving Eqn. 5, given singles detection rate $s_1 = s_2 = 100$ kcounts/s, coincidence rate $c = 650$ counts/s, and bin overlap of $\nu = 0.5$, and frequency offset error of $\Delta u = 50$ ppb (precompensation step size of 100 ppb). (b) Corresponding probabilities using the normal approximation for observed accidental coincidences X in Eqn. 5, with the same parameters. The probabilities are strongly dependent on N and minimally with δt , similar to the estimations provided using the statistical significance framework.

N and δt values, the normal approximation $f_{(N)}(x|\mu, \sigma)$ underestimates the noise peak value, and thus overestimates the peak finding success probabilities.

The success probabilities for the Poisson model and the normal approximation model are shown in Fig. 5 by numerically solving Eqn. 5 using Monte Carlo simulations of $> 10^7$ trials per $(N, \delta t)$ pair. The normal approximation indeed overestimates the peak finding probabilities at small time bin width δt . We use $N = 2^q$ ($q \in \mathbb{Z}^+$) for its logarithmic scale.

Appendix D: Derivation of peak tracking equation and parameters

We perform active frequency compensation by estimating the clock frequency offset from the set of timestamps. The timing difference $\tau_i := t_i - t'_i$, between a timestamp pair $\{t_i, t'_i\}$ from both parties, is continuously served by searching for photon pair detection events within a prescribed coincidence tracking window.

The frequency offset Δu is given by the ratio between measured elapsed time $\Delta t_i := t_i - t_{i-1}$ with respect to some reference elapsed time, in this case the elapsed time measured by the other peer $\Delta t'_i$,

$$\frac{\Delta t_i}{\Delta t'_i} = 1 + \Delta u.$$

Rewriting in terms of the measured successive timing difference, we can estimate the frequency offset by measur-

ing the rate of change in the timing difference, i.e.,

$$\Delta u_i = \frac{\tau_i - \tau_{i-1}}{\Delta t'_i}.$$

Active frequency compensation is therefore achieved by performing a timing correction for each timestamp t_i using the estimated frequency offset,

$$t_i \rightarrow \Delta t_i(1 + \Delta u) + t_{i-1} = t_i + \Delta t_i \Delta u,$$

where the overall frequency offset Δu accumulated is also given by $\Delta u = \prod_0^i (1 + \Delta u_i) - 1$.

Peak tracking with thermal light is intrinsically noisy due to the signal being dominated by accidental coincidences. Hence, we use an exponential moving average filter given by,

$$\tau'_i = \alpha \tau_i + (1 - \alpha) \tau'_{i-1}, \quad \tau'_0 = \tau_0,$$

which behaves like a low-pass filter to smooth the signal. The coefficient α represents the relative weight of timing offset τ_i with respect to the accumulated average τ'_{i-1} , and can be adjusted depending on the observed signal rate.

Given a unit step impulse, the time it takes to reach $1 - 1/e$ of the signal is associated with a time constant β related to the coefficient α ,

$$\alpha = 1 - \exp\left(-\frac{\Delta t}{\beta}\right) \approx \frac{\Delta t}{\beta}, \quad \beta \gg \Delta t,$$

after Taylor expansion of the exponential, where $\overline{\Delta t}$ is the average separation time between consecutive timestamp events. The singles rate s in our experiment is 200 kcounts/s, so $\overline{\Delta t} = 1/s = 5 \mu\text{s}$. We set the time constant $\beta = 50$ ms, which is 4 orders of magnitude longer than $\overline{\Delta t}$ for averaging.