Clock Synchronization without Frequency References under Weak Timing Correlations

Justin Yu Xiang Peh,[1] Darren Ming Zhi Koh,[1] Zifang Xu,[2] Xi Jie Yeo,[1] Peng Kian Tan,[1] and Christian Kurtsiefer[1,2]

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

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

(Dated: January 8, 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 quartz clocks using weakly time-correlated photons of 180 ns coherence time from a pseudothermal bunched light source. A synchronization timing jitter of 10 ns is achieved over symmetric -102 dB optical loss channels between two parties, over a span of 25 hours. We also derive a model to accurately estimate the coherence peak searching success probabilities.

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 is typically achieved either using pulsed laser light [3] or with classical signals on dedicated optical or electronic channels [4, 5].

Modern systems can use the resources typically in the protocol itself to perform the clock synchronization, such as photon pairs from spontaneous parametric down-conversion (SPDC) [6, 7]. This is possible with the use of 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.

To remove the dependency on Rb frequency standards, it was proposed to find and track the timing difference between photon pairs due to their strong timing correlation on the order of picoseconds [7, 8], by identifying an initial coincidence peak with low timing resolution then iteratively apply frequency corrections to improve resolution. Weak coherent photon pulses were also proposed for clock frequency transfer [9]. In both cases, the cross-correlation peak is very strong, i.e., $g^{(2)}(\tau = 0) \gg 2$, with resolution generally limited only by the timing jitter of the generation and detection optoelectronics.

Other sources of timing-correlated light are thermal light sources which has been used in application 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.

However, unlike photon pairs, identification of this peak is more challenging due to the low signal (i.e., $g^{(2)}(\tau) \leq 2$) as well as fluctuations about the background of $g^{(2)} = 1$. There is an earlier proposal that suggests the use of thermal light for clock synchronization, 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 weak timing correlation light source, i.e., light from pseudothermal source with a $g^{(2)}(\tau) = 1.44$ and coherence time $\tau_c = 180(6)$ ns, over a -102 dB transmission channel with 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 sources.

II. CLOCK SYNCHRONIZATION WITH BUNCHED LIGHT

Timing synchronization was achieved using bunched light between two remote parties over symmetrical loss channels, each with -102 dB average transmission, for more than 25 hours using only simple crystal oscillators, see Fig. 1.

The source of bunched light is a pseudothermal source [15], established using a laser of 780 nm wavelength sent into an unbalanced Mach-Zehnder interferometer with optical delay longer than the laser coherence time τ_c; phase fluctuations in the laser manifests the bunching effect, see Appendix A for a detailed explanation. This results in an intensity correlation that can be measured through the second-order coherence func-

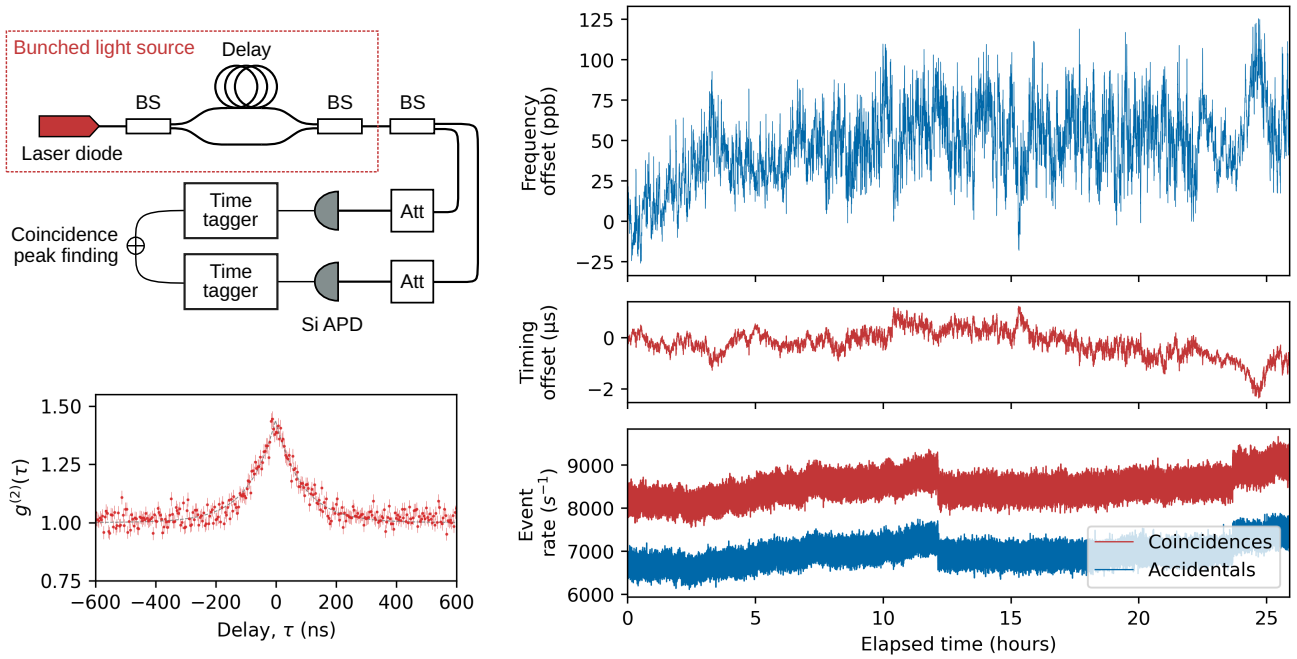

FIG. 1. **Top left:** Simplified experimental setup for clock synchronization. Bunched (pseudothermal) light is sent to two different parties through separate -102 dB channels, before detection by Silicon avalanche photodetectors (Si APDs) and timestamping by time taggers running on independent quartz crystal oscillators. Singles count rate s_1 and s_2 of approximately 200 kcounts/s are recorded on each side. BS: beamsplitter; Att: attenuators.
Bottom left: Plot of $g^{(2)}(\tau)$ of the pseudothermal light source, with the curve fit in grey corresponding to $g^{(2)}(0) = 1.44(1)$ peak and coherence time $\tau_c = 180(6)$ ns. Error bars correspond to Poissonian errors from counting statistics.
Right: Long-term trace of frequency offset, timing offset, 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.

tion $g^{(2)}(\tau)$, given by

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

due to the Lorentzian spectral profile of the laser, where the coherence time τ_c is inversely proportional to the frequency spectral width Δf of the source [16]. The brightness of the source is 1.55(2) mW, with a measured coherence peak of $g^{(2)}(\tau = 0) = 1.44(1)$ and coherence time of $\tau_c = 180(6)$ ns.

The light from the source is shared between $n = 2$ parties with optical fiber channel tranmissions of $-101.9(4)$ dB and $-102.2(4)$ dB 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 quartz 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 8,500 events/s is consistently higher than the accidental count rate of 7,000 events/s, indicative of successful frequency tracking.

Due to the presence of clock frequency drift, typically from temperature fluctuations and electronic noise, the frequency offset between the clocks changes over time as well. We are able to reconstruct this offset, by monitoring the drift in served timings using individual samples spanning 10.74 s, with a resolution of 0.537 s, shown in Fig. 1.

The actual frequency offset 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.

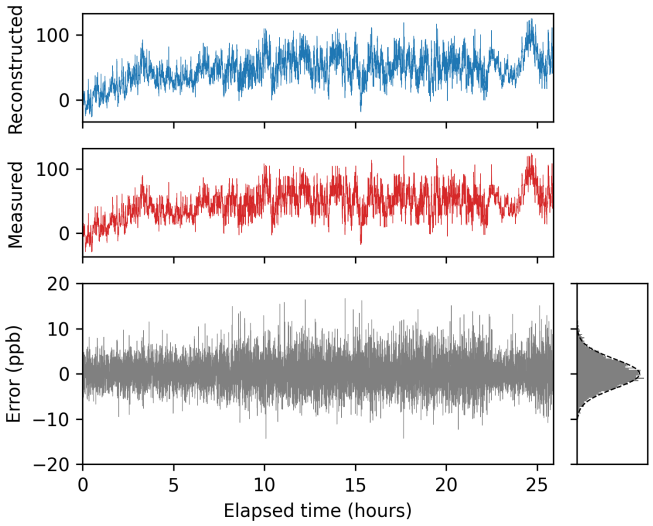

FIG. 2. A concurrent measurement of the actual frequency offset between the two external 10 MHz clocks, as well as the corresponding error in the reconstructed frequency offset. The accuracy of the frequency servoed by the peak tracking algorithm is 3.2 ppb on average.

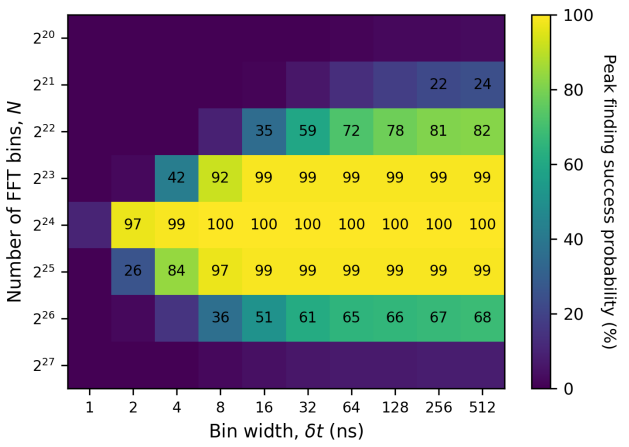

FIG. 3. Simulated probabilities to find the correct peak position by solving the required true coincidence equation found in Eqn. 5, given the singles detection rate $s_1 = s_2 = 100$ kcounts/s, coincidence rate $c = 650$ counts/s, bin overlap $\nu = 0.5$, and frequency offset error of up to $\Delta u = 100$ ppb. These parameters correspond to the setup in Fig. 1 with an additional 3 dB attenuation per channel, and shows good agreement with measurements.

III. PEAK SEARCHING

Clock synchronization can be decomposed into two parts: first identifying the initial clock frequency and time difference between two parties, then tracking of said peak to monitor 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.

Efficient peak identification relies on the circular convolution theorem to compute the cross-correlation between two sets of detection timestamps $a[k]$ 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 inverse [7][17]. The success of peak identification is constrained by the singles rate and the rate of true coincidences, and additionally requires the choice of optimal parameters for the Fast Fourier Transform (FFT) parameters for computing 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 in Fig. 3. Our model provides better estimations of the peak finding probability compared to previous works, by avoiding the normal approximation to the noise in the FFT (see Appendix B).

A. No frequency offset

The minimum acquisition time required for the cross-correlation is $T = N\delta t$ to obtain a flat cross-correlation noise floor. This noise floor arises from accidental coincidences — attributed to noise sources, such as coincidences with uncorrelated detection events and dark 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 accidental coincidences in each time bin also follows a Poisson distribution P, i.e., $X_{k\in\{1,\ldots,N\}} \sim \mathrm{P}(\lambda = C_a)$.

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

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

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

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

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 a smudge factor into Eqn. 3 representing the degree of bin overlap $\nu \in [0.5, 1]$, yielding

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

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 some elapsed time T. Under 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., $\delta t > \Delta\tau_{\text{drift}}$, or in other words, $N < 1/\Delta u$. We model this as an additional smudge factor $\mu \equiv \max\{1, N\Delta u\}$, and together with Eqn. 4 yields the minimum required true coincidence rate for successful peak finding given by

$$c_e > \frac{1}{\mu\nu T}\left(X_{(N)} - X\right). \tag{5}$$

The surface of this equation is plotted in Fig. 3, by performing Monte Carlo simulations for a specific set of parameters and across different N and δt.

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 [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

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

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

In practice, while the frequency offset between quartz clocks can be high (~ 10 ppm), the short-term stability of the clocks themselves are much higher (< 10 ppb/s), so after an initial precompensation reference, a small frequency precompensation step size of 100 ppb is sufficient to identify most peaks within 3 peak searching attempts [18].

Once a 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 in Python, and frequency compensation in C, are available under GPLv2 license in GitHub [19].

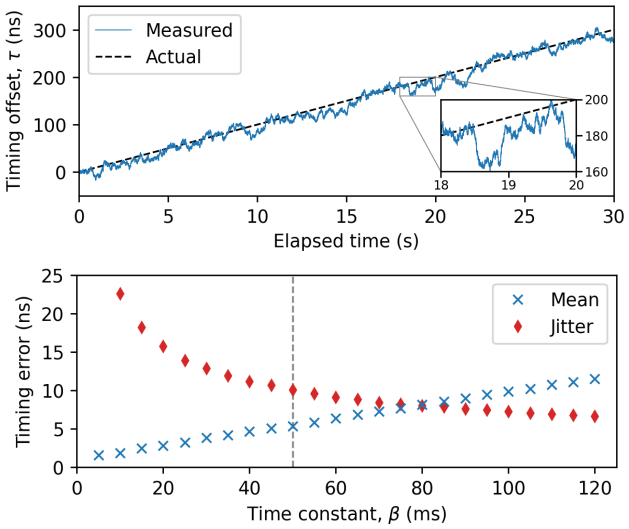

FIG. 4. Characterization of peak tracking accuracy under a constant frequency offset $\Delta u = 10$ ppb, with a time constant $\beta = 50$ ms chosen for the exponential moving average filter used in the experiment.
Top: Measurement of reported timing offsets (blue) against the actual offset (black) over a period of 30 seconds. The inset highlights a $\approx$30 ns timing error over a 0.5-second window, which suggests that a histogram fit does not improve accuracy.
Bottom: 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 grey dashed line). Peak tracking fails with time constants less than 5 ms.

IV. PEAK TRACKING

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

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 τ_i 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 attributed to the frequency offset being non-zero and constant, and manifests as a mean timing error of around 6 ns at $\beta = 50$ ms. While this lag scales proportionally to the frequency offset, this lag is expected to not contribute significantly due to the random fluctuations of the frequency offset about its mean, in the case of free-running clocks.

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. However, for the bunched source with a significantly noisier coherence peak signature, this technique does not improve the timing jitter, 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 pseudothermal 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. 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 the photon bunching as a resource, or indirectly by compensating the reported timings with the measured timing and frequency offsets. Direct frequency compensation can also be performed by actively correcting the quartz oscillator frequencies, e.g., by means of a voltage error signal, 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 author would like to thank Ng Boon Long for helpful discussions in setting up the concurrent frequency measurement.

[1] David L Mills. *Computer network time synchronization: the network time protocol*. CRC press, 2006.

[2] Brent A Renfro, MS Stein, Emery B Reed, and Eduardo J Villalba. An analysis of global positioning system standard positioning service performance for 2020, 2022.

[3] Fabrizio R Giorgetta, William C Swann, Laura C Sinclair, Esther Baumann, Ian Coddington, and Nathan R Newbury. Optical two-way time and frequency transfer over free space. *Nature Photonics*, 7(6):434–438, 2013.

[4] T. Gerrits, I. A. Burenkov, Y.S. 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*, page FM1C.2, San Jose, California, 2022. Optica Publishing Group.

[5] Muneer Alshowkan, Philip G. Evans, Brian P. Williams, Nageswara S. V. Rao, Claire E. Marvinney, Yun-Yi Pai, Benjamin J. Lawrie, Nicholas A. Peters, and Joseph M. Lukens. Synchronizing a quantum local area network with White Rabbit. In *Conference on Lasers and Electro-Optics*, page FM1C.4, San Jose, California, 2022. Optica Publishing Group.

[6] Jianwei Lee, Lijiong Shen, Alessandro Cerè, James Troupe, Antia Lamas-Linares, and Christian Kurtsiefer. Symmetrical clock synchronization with time-correlated photon pairs. *Applied Physics Letters*, 114(10), 2019.

[7] Caleb Ho, Antia Lamas-Linares, and Christian Kurtsiefer. Clock synchronization by remote detection of correlated photon pairs. *New Journal of Physics*, 11(4):045011, 2009.

[8] Christopher Spiess, Sebastian Töpfer, Sakshi Sharma, Andrej Kržič, Meritxell Cabrejo-Ponce, Uday Chandrashekara, Nico Lennart Döll, Daniel Rieländer, and Fabian Steinlechner. Clock Synchronization with Correlated Photons. 19(5):054082.

[9] Christopher Spiess and Fabian Steinlechner. Clock synchronization with pulsed single photon sources. 9(1):015019.

[10] Ryan S Bennink, Sean J Bentley, and Robert W Boyd. "two-photon" coincidence imaging with a classical source. *Physical review letters*, 89(11):113601, 2002.

[11] Peng Kian Tan, Xi Jie Yeo, Alvin Zhen Wei Leow, Lijiong Shen, and Christian Kurtsiefer. Practical range sensing with thermal light. *Physical Review Applied*, 20(1):014060, 2023.

[12] R Hanbury Brown and Richard Q Twiss. Correlation between photons in two coherent beams of light. *Nature*, 177(4497):27–29, 1956.

[13] Jun Zhu, Peng Huang, Xiaoqi Xiao, and Guihua Zeng. A new clock synchronization scheme based on the second-order coherence of thermal light. *Optica Applicata*, 43(2), 2013.

[14] Guangjian Xu, Xinyi Ren, Qucheng Miao, Ming Yan, Haifeng Pan, Xiuliang Chen, Guang 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(16):1321–1326, 2020.

[15] Xi Jie Yeo. Preprint. *Pending publication*, 2024.

[16] R. Loudon. *The Quantum Theory of Light*. Oxford University Press, 2000.

[17] We additionally note that algorithms faster than FFT's $\mathcal{O}(n \log n)$ time complexity are also possible [20, 21], but require additional encoding of time-synchronization strings within the light pulses.

[18] Using offsets of Δu_0 and $\Delta u_0 \pm 100\,\mathrm{ppb}$.

[19] Justin Yu Xiang Peh, Zifang Xu, and Christian Kurtsiefer. fpfind. `https://github.com/s-fifteen-instruments/fpfind`, 2024.

[20] Luca Calderaro, Andrea Stanco, Costantino Agnesi, Marco Avesani, Daniele Dequal, Paolo Villoresi, and Giuseppe Vallone. Fast and Simple Qubit-Based Synchronization for Quantum Key Distribution. *Physical Review Applied*, 13(5):054041, May 2020.

[21] Jan Krause, Nino Walenta, Jonas Hilt, and Ronald Freund. Clock offset recovery with sublinear complexity enables synchronization on low-level hardware for quantum key distribution, April 2024.

[22] A. Khintchine. Korrelationstheorie der stationären stochastischen prozesse. *Mathematische Annalen*, 109(1):604–615, Dec 1934.

[23] AJF Siegert. *On the fluctuations in signals returned by many independently moving scatterers*. Radiation Laboratory, Massachusetts Institute of Technology, 1943.

[24] Xi Jie Yeo, Eva Ernst, Alvin Leow, Jaesuk Hwang, Lijiong Shen, Christian Kurtsiefer, and Peng Kian Tan. Direct measurement of the coherent light proportion from a practical laser source. *Phys. Rev. A*, 109:013706, Jan 2024.

Appendix A: Pseudothermal source and experimental setup

The primary source used in the frequency tracking measurement is a pseudo-thermal light source, where light from a laser running above the lasing threshold is split into two separate paths, one of which is delayed beyond the coherence time of the laser. If the two paths are perfectly indistinguishable (i.e., in polarization, spatial mode, and brightness), they act as independent light emitters and exhibit bunching effects.

From the Wiener-Khintchine theorem [22], the Lorentzian spectral profile from a coherent laser light source corresponds to a $g^{(1)}$ with a Laplacian timing profile of the form,

$$g^{(1)}(\tau) := \frac{\langle E^*(t)E(t+\tau)\rangle}{\langle E^*(t)E(t)\rangle} = \exp\left(-\frac{|\tau|}{\tau_c}\right)$$

as a function of coherence time τ_c.

Thermal light can be modelled as a large collection of independent emitters of light [16]. In the case of ν emitters, we have

$$\langle E^*(t)E(t+\tau)\rangle = \nu\langle E_i^*(t)E_i(t+\tau)\rangle$$

and

$$\begin{aligned}\langle E^*(t)E^*(t+\tau)E(t+\tau)E(t)\rangle &= \nu\langle E_i^*(t)E_i^*(t+\tau)E_i(t+\tau)E_i(t)\rangle \\ &+ \nu(\nu-1)\left[\langle E_i^*(t)E_i(t)\rangle + |\langle E_i^*(t)E_i(t+\tau)\rangle|\right],\end{aligned}$$

and we can recover the Siegert relation [23] and determine the $g^{(2)}$ of the light source

$$\begin{aligned}g^{(2)}(\tau) &:= \frac{\langle E^*(t)E^*(t+\tau)E(t+\tau)E(t)\rangle}{\langle E^*(t)E(t)\rangle^2} \\ &= 1 + \frac{\nu-1}{\nu}\left|g^{(1)}(\tau)\right|^2 \\ &= 1 + \frac{1}{2}\exp\left(-\frac{2|\tau|}{\tau_c}\right), \quad \nu = 2,\end{aligned}$$

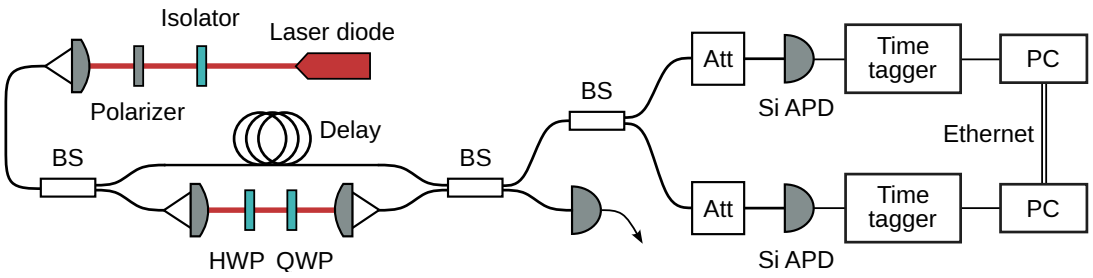

FIG. 5. Detailed experimental setup using a pseudothermal bunched light source. BS: 50:50 beamsplitter; Att: attenuator; HWP: half-wave plate; QWP: quarter-wave plate; APD: avalanche photodiode.

which saturates at a value of $g^{(2)}(\tau = 0) = 1.5$.

Our experimental setup, per Fig. 5, 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. The coherence length of the laser is approximately 200 m; a delay fiber of length 400 m was used in the delay arm, corresponding to a 2 µs delay. Polarization rotation in the delay arm is compensated using a free-space link with a HWP and QWP, while attenuation is achieved by slight decoupling of the collection mode, before recombination using a fused-coupler beamsplitter (BS). The $g^{(2)}(\tau = 0) = 1.44(3)$ achieved in the setup is attributed to the fact that the laser is not fully coherent [24] and other experimental imperfections.

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. Fiber spools for delay was not used to avoid potential complications from timing delays attributed to fiber length changes.

Time tagging is subsequently performed by timestamp devices (S-Fifteen Instruments TDC2) with 4 ps nominal timing resolution and 20 ps 1-σ timing jitter. The clock to each timestamp device is supplied by external 10 MHz clock distribution boards with quartz crystal oscillators and without any onboard temperature stabilization.

The actual frequency offset between the two clocks was measured concurrently with the clock synchronization experiment by mixing separate copies of the 10 MHz signals and passing through a low pass filter with 4 GHz cut-off frequency. This mixed signal was sampled by an oscilloscope at a rate of 2.5 ksamples/s over 10 s, before performing 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 servoed 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.

Appendix B: Peak finding probabilities using Poisson statistics vs Gaussian assumption

The corresponding probabilities for Poisson model and the corresponding Gaussian approximation model are presented in Fig. 6. We find that the previous models assuming Gaussian-distributed bin values incorrectly predict the success probability of peak finding: in particular, it generally underestimates the probability for low time bin widths, and overestimates for small number of FFT bins.

FFT can be efficiently implemented if the number of bins N are either in powers of 2, or has factorizations with small primes. Here we use $N = 2^q, q \in \mathbb{Z}^+$ for its logarithmic scale.

Appendix C: Derivation of max-order distribution

Given that the FFT is performed with N bins, the accidental coincidences in each of these bins follow the Poisson distribution with mean λ. The maximum value across all these bins increases with an increasing number of bins. The distribution of the max-order statistic (i.e., the maximum of all bins) with respect to the number of coincidences k and number of bins N of mean value λ is denoted $f_{(N)}(k|\lambda)$.

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,

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

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

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

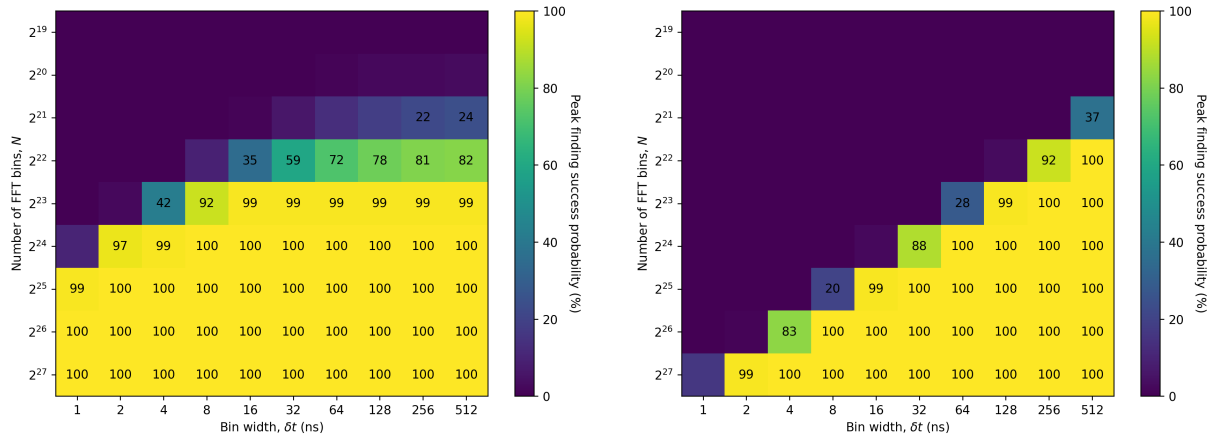

FIG. 6. **Left:** Probabilities to find the correct peak position by solving Eqn. 4, given the singles detection rate $s_1 = s_2 = 100\,\text{kcounts/s}$, coincidence rate $c = 650\,\text{counts/s}$, and bin overlap of $\nu = 0.5$. Since there is no frequency offset, the peak finding probabilities reaches 100% with increasing number of FFT bins.
Right: Probabilities to find the correct peak position using the normal approximation for X_i in Eqn. 4, with the same parameters.

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 $F(k|\lambda) \in [0, 1]$ 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,

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

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

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

This in fact corresponds to the first order term for the discrete case after a binomial expansion, noting that $f(k|\lambda) \leq F(k|\lambda)$, with

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

Appendix D: Peak tracking derivation

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

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

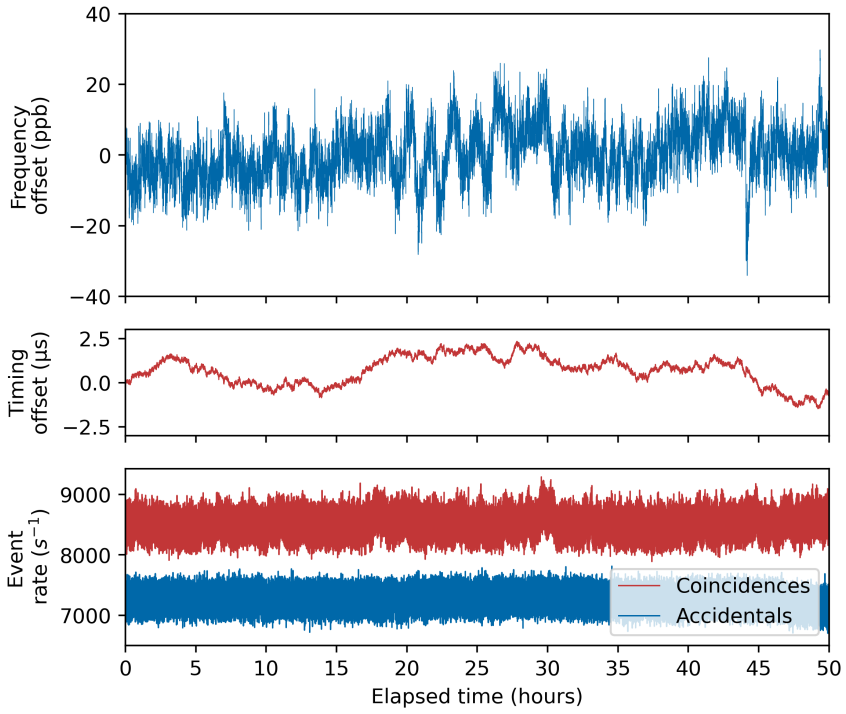

FIG. 7. Tracking measurement over 50 hours, similar to that of Fig. 1.

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

$$t_i \to \Delta t_i(1+\Delta u) + t_{i-1} = t_i + \Delta t_i \Delta u, \tag{D1}$$

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

The peak tracking is intrinsically noisy due to the signal being dominated by accidental coincidences. 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 cofficient 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,

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

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 $\overline{\Delta t} = 1/s = 5\,\mu$s. We set the time constant $\beta = 50$ ms, which is 4 orders of magnitude longer than $\overline{\Delta t}$ for averaging.

Appendix E: Longer time scale measurement

Fig. 7 shows a longer experimental measurement spanning over 50 hours, with a similar drift behaviour observed in the reconstructed frequency offset. The frequency offsets were not measured concurrently, so no tracking accuracy estimation is available for this dataset.

Measurement was performed using a different pair of Si APDs instead (Micro Photon Devices SPAD), which had an order of magnitude lower system efficiency and resulted in a lower -90 dB attenuation channel possible at the same detection rate parameters. The measurement was also performed using suboptimal parameters for peak tracking that dropped off too many coincidences.

The timestamps were disciplined with onboard voltage-controlled temperature-compensated quartz crystal oscillators, so the frequency offset between the two timestamps generally remained the same, as well as nearly an order of magnitude better frequency stability.