

Supplementary Information:
Attosecond timing jitter in millimeter waves via Kerr optical
frequency division

Scott C. Egbert¹, Brendan M. Heffernan¹, James Greenberg¹, William F. McGrew¹, and
Antoine Rolland¹

¹Boulder Research Laboratory, IMRA America, Inc., Longmont, Colorado, USA

February 16, 2026

Contents

S1 Introduction	2
S2 Experimental realization of Kerr optical frequency division	2
S3 System model and Adler-type locking of the repetition rate	3
S3.1 Exact division law and phase-noise transfer functions	4
S3.2 Experimental extraction of the locking bandwidth	5
S3.3 Injection-locking simulations versus experiment	5
S4 Comparison with PLL-based optical frequency division (simulation)	5
S4.1 Idealized PLL model	6
S4.2 Simulated comparison with Kerr optical frequency division	7
S4.3 Implications for high-frequency division and timing jitter	7
S5 Noise contributions and scaling laws	8
S5.1 Low-offset technical noise and $\Delta\nu$ scaling	9
S5.2 Schawlow–Townes limit and quantum noise	9
S5.3 Photodetection noise: thermal and shot-noise limits at 300 GHz	10
S5.4 Comb relative intensity noise and AM-to-PM-induced phase noise	12
S6 Timing jitter and Allan deviation	12
S6.1 Timing jitter	12
S6.2 Allan deviation	13
S7 Summary of closed-form results	14

S1 Introduction

This Supplementary Material provides the experimental, theoretical, and metrological details that underpin and extend the results presented in the main text. Its purpose is to establish a rigorous and transparent connection between the physical principles of Kerr optical frequency division (OFD), the experimental implementation of the system, and the measured phase-noise and timing-jitter performance at millimeter-wave frequencies.

Section S2 presents the complete system-level architecture used to divide a multi-terahertz optical reference down to a 300 GHz millimeter-wave carrier. This section details the optical reference, Kerr soliton microresonator, injection-locking mechanism, spectral conditioning, and photodetection stages. The dual-wavelength Brillouin laser is treated as a subsystem within the overall Kerr OFD architecture; its internal operating principles are not reproduced here and are instead referenced to prior work.

Section S3 develops an oscillator-level description of Kerr OFD based on Adler-type injection locking. This formalism yields the exact division law, defines the finite reference-tracking bandwidth, and provides closed-form transfer functions describing how reference noise, free-running soliton noise, and additive noise sources map onto the divided repetition rate. These results supply the analytical foundation for interpreting the phase-noise spectra reported in the main text.

Section S4 connects this model to experiment and simulation. We describe the experimental extraction of the injection-locking bandwidth, validate the Adler-based description through direct comparison with measured phase-noise spectra, and present numerical simulations contrasting Kerr OFD with an idealized phase-locked-loop (PLL)-based OFD scheme. This comparison isolates the central role of locking bandwidth and highlights the absence of servo-induced resonances in injection-locked division.

Sections S5 and S6 provide a detailed accounting of the noise mechanisms relevant to millimeter-wave operation at 300 GHz and their manifestation in both the frequency and time domains. These include scaling of low-offset technical noise with optical frequency separation, quantum-limited Schawlow–Townes noise, photodetection thermal and shot-noise limits, and excess noise arising from comb relative-intensity noise and amplitude-to-phase conversion. Timing jitter and Allan deviation are derived consistently from the measured phase-noise spectra, providing a time-domain perspective that complements the frequency-domain analysis in the main text.

Finally, Section S7 summarizes the key closed-form relations governing Kerr OFD, serving as a compact reference linking experiment, theory, and noise scaling laws. Together, this Supplementary Material establishes the physical origin, limitations, and generality of Kerr OFD as a route to ultralow-noise millimeter-wave and sub-terahertz oscillators.

S2 Experimental realization of Kerr optical frequency division

This section provides a schematic overview of the complete experimental setup used to implement Kerr optical frequency division and millimeter-wave phase-noise measurement. The diagram in Supplementary Fig. S1 summarizes the optical and electrical signal flow, as well as the interconnections between subsystems.

The purpose of this schematic is to aid interpretation of the measurement configuration and to clarify how the optical reference, Kerr microresonator, and millimeter-wave detection chain are interconnected. Detailed descriptions of the experimental procedures, operating conditions, and measurement techniques are given in the Methods section of the main text.

The diagram highlights the generation of the dual-wavelength Brillouin laser (DWBL) reference,

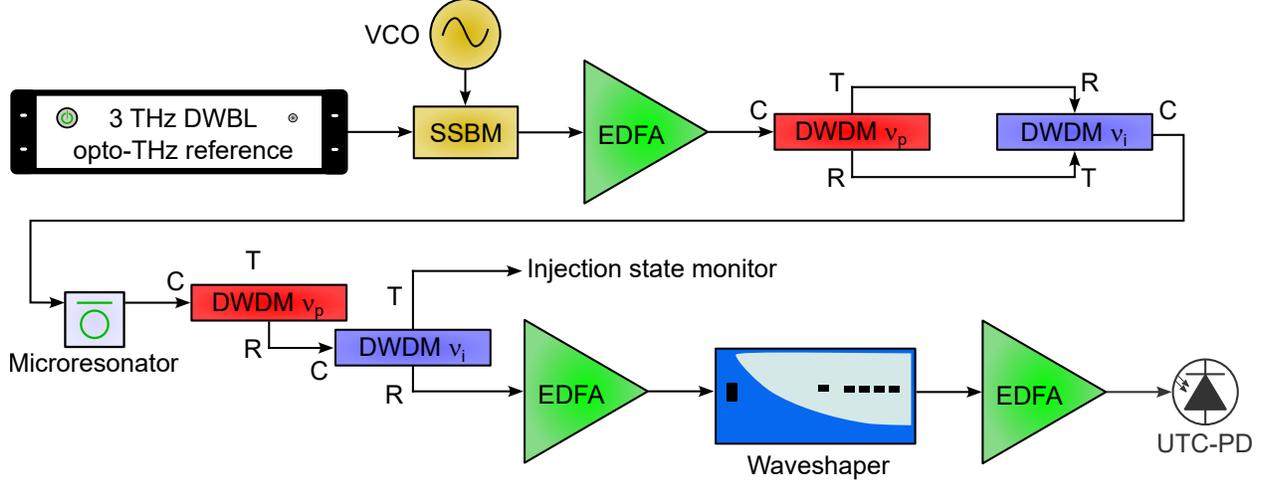


Figure S1: **Detailed experimental schematic of the Kerr optical frequency division (OFD) setup.** Starting from the 3.3 THz dual-wavelength Brillouin laser (DWBL) optical reference, the pump and injection tones, ν_p and ν_i , are frequency shifted using a voltage-controlled oscillator (VCO)-driven modulator to enable Kerr soliton initiation in the microresonator. Both ν_p and ν_i are then amplified using a shared erbium-doped fiber amplifier (EDFA) and spectrally cleaned with a dual bandpass filter implemented by a pair of dense wavelength-division multiplexing (DWDM) filters (ports C: common, T: transmit, R: reject). The two tones then co-propagate to pump and optically inject the microresonator, enabling stable soliton formation. Following the microresonator, residual pump and injection light are rejected, again with DWDM filtering, to prevent saturation of the uni-traveling-carrier photodiode (UTC-PD). The resulting Kerr soliton comb is subsequently amplified, passed through a programmable optical filter (waveshaper) for dispersion compensation, and re-amplified to provide the optical power required to drive the UTC-PD for low-noise terahertz generation.

soliton-based optical frequency division, photodetection at 300 GHz.

S3 System model and Adler-type locking of the repetition rate

The repetition rate of a Kerr soliton frequency comb can be modeled as a self-sustained oscillator subject to injection locking by an external reference. In the present system, the reference is provided by the optical difference between two coherent tones of a dual-wavelength Brillouin laser (DWBL) [1]. Optical fields at frequencies ν_p and ν_i generate a difference frequency $\Delta\nu = \nu_i - \nu_p$ with associated phase $\phi_\Delta(t) = \phi_i(t) - \phi_p(t)$, which is used to coherently drive the microresonator comb.

Injection of the optical-difference tone onto the comb tooth indexed from the pump comb tooth by the integer N constrains the phase of the soliton repetition rate $\theta_{\text{rep}}(t)$ to that of the reference. At the oscillator level, the repetition rate therefore behaves as a slave oscillator driven by a master signal. Introducing the relative phase

$$\psi(t) \equiv N \theta_{\text{rep}}(t) - \phi_\Delta(t), \quad (1)$$

the slow-time, cycle-averaged phase dynamics are described by the Adler equation [2]

$$\dot{\psi}(t) = \Delta\omega - K \sin \psi(t) + \xi(t), \quad (2)$$

where $\Delta\omega = N\omega_{\text{rep}}^{(0)} - \Omega$ is the detuning between the free-running repetition rate $\omega_{\text{rep}}^{(0)}$ and the angular frequency $\Omega = 2\pi\Delta\nu$ of the injected optical-difference signal, K is the effective injection coupling coefficient, and $\xi(t)$ represents intrinsic phase-fluctuation noise of the free-running repetition-rate oscillator, averaging to zero over sufficiently long timescales.

A stationary locked solution exists for $|\Delta\omega| \leq K$, with equilibrium phase $\psi^* = \arcsin(\Delta\omega/K)$. Linearization of Eq. (2) about ψ^* shows that small phase perturbations are tracked only within a finite bandwidth

$$f_c = \frac{K \cos \psi^*}{2\pi}, \quad (3)$$

which defines the frequency range over which the repetition rate follows the optical-difference reference. This tracking bandwidth constitutes the fundamental limit of Kerr OFD and determines the transfer of reference phase noise and intrinsic soliton noise to the output repetition rate, as quantified by the linear transfer functions derived in the following sections.

S3.1 Exact division law and phase-noise transfer functions

Within the injection-locked regime defined previously, the relative phase $\psi(t)$ reaches a stationary value ψ^* , and the repetition-rate phase is constrained by the optical-difference reference. In steady state, setting $\dot{\psi} = 0$ in Eq. (2) yields the exact phase relation

$$N \theta_{\text{rep}}(t) = \phi_{\Delta}(t) + \psi^*, \quad (4)$$

or equivalently,

$$\theta_{\text{rep}}(t) = \frac{1}{N} \phi_{\Delta}(t) + \frac{\psi^*}{N}. \quad (5)$$

This relation shows that, within the locked steady state, the soliton repetition rate realizes an exact division of the optical-difference phase by a factor N .

As a direct consequence, phase fluctuations of the optical-difference reference are reduced by the square of the division factor. For phase noise power spectral densities (PSDs) evaluated at Fourier frequencies well below the tracking bandwidth ($f \ll f_c$),

$$S_{\theta_{\text{rep}}}(f) = \frac{1}{N^2} S_{\phi_{\Delta}}(f), \quad L_{\text{rep}}(f) = L_{\Delta}(f) - 20 \log_{10} N, \quad (6)$$

corresponding to ideal frequency division.

At higher Fourier frequencies, the finite locking bandwidth results in incomplete tracking of the reference. To quantify this behavior, Eq. (2) is linearized about the locked solution ψ^* by writing $\psi(t) = \psi^* + \delta\psi(t)$. In the frequency domain (uppercase letters denote Fourier transforms), the phase fluctuations of the locked repetition rate $\delta\Theta_{\text{rep}}$ are related to fluctuations of the optical-difference reference $\delta\Phi_{\Delta}$, the free-running repetition rate $\delta\Theta_{\text{rep}}^{(0)}$, and additive noise sources Ξ through

$$\delta\Theta_{\text{rep}}(j\omega) = \underbrace{\frac{1}{N} \frac{j\omega}{j\omega + 2\pi f_c}}_{H_m(j\omega)} \delta\Phi_{\Delta}(j\omega) + \underbrace{\frac{j\omega}{j\omega + 2\pi f_c}}_{H_s(j\omega)} \delta\Theta_{\text{rep}}^{(0)}(j\omega) + \underbrace{\frac{1}{N} \frac{1}{j\omega + 2\pi f_c}}_{H_{\xi}(j\omega)} \Xi(j\omega). \quad (7)$$

The resulting repetition-rate phase-noise PSD is therefore

$$S_{\theta_{\text{rep}}}(f) = |H_m|^2 S_{\phi_{\Delta}}(f) + |H_s|^2 S_{\theta_{\text{rep}}^{(0)}}(f) + |H_{\xi}|^2 S_{\xi}(f). \quad (8)$$

In the low-frequency limit $f \ll f_c$, $|H_m|^2 \rightarrow 1$ and $|H_s|^2 \rightarrow 0$, recovering ideal frequency division. For $f \gg f_c$, the reference contribution to the phase-noise PSD rolls off as $(f_c/f)^2$, and the repetition-rate noise approaches that of the free-running soliton oscillator.

S3.2 Experimental extraction of the locking bandwidth

The effective tracking bandwidth f_c defined at the beginning of this section can be determined experimentally from measurements of the repetition rate as a function of the optical-difference injection frequency. In practice, the optical-difference injection frequency is held fixed while the soliton pump–resonance detuning is swept, thereby tuning the free-running repetition rate through the injection point as the locked repetition frequency f_{rep} is monitored.

Within the locking range, the repetition rate remains pinned to the injected reference, producing a flat plateau as a function of detuning (Adler curve). Linearizing the central portion of this plateau yields $\partial f_{\text{rep}}/\partial f_i \approx 0$, confirming rigid phase locking. Near the edges of the plateau, the slope increases sharply as the system approaches the unlocking condition.

For small equilibrium phase offsets ($|\psi^*| \ll 1$), the half-width of the locking plateau in frequency units provides an estimate of the injection coupling strength K . The corresponding small-signal tracking bandwidth is then given by

$$f_c \simeq \frac{1}{2\pi} \times (\text{half-width of locking plateau}), \quad (9)$$

consistent with the Adler model. Outside the locking range, the slope of the pulling curve yields the sensitivity $\partial f_{\text{rep}}^{(0)}/\partial f_i$ of the free-running repetition rate to the injected tone, providing an independent consistency check of the model.

This procedure allows direct experimental determination of the tracking bandwidth governing phase-noise transfer, without requiring time-domain locking measurements.

S3.3 Injection-locking simulations versus experiment

To validate the Adler-type description of Kerr OFD, numerical simulations of injection locking were performed and compared directly to experimental phase-noise measurements. Figure S2 shows repetition-rate phase-noise spectra obtained both experimentally and from simulation for two distinct operating points: (i) tuning to the center of the locking range, and (ii) detuning to the boundary of injection locking (boundary locking), where phase locking is still maintained but the equilibrium phase ψ^* approaches the stability limit. Experimentally, this condition is identified by tuning the injected optical-difference frequency until the beat note between the free-running repetition rate and the injected tone collapses from a finite offset (typically ~ 20 MHz) to DC. This abrupt collapse marks the onset of injection locking and indicates an injection bandwidth on the order of the observed collapse frequency.

The measured phase-noise spectra at both the center of lock and the locking boundary are in excellent agreement with simulations based on the Adler phase model. In particular, the simulations reproduce (i) the absolute noise levels, (ii) the frequency at which reference tracking rolls off due to the finite tracking bandwidth, and (iii) the increase in phase noise as the system is tuned toward the locking boundary. These results provide direct experimental validation of the oscillator-level model used throughout this work.

S4 Comparison with PLL-based optical frequency division (simulation)

To illustrate the role of locking bandwidth in high-frequency optical frequency division (OFD), we compare Kerr OFD, based on injection locking, with a representative phase-locked-loop (PLL) based OFD scheme, as can be found experimentally in [3]. The comparison is performed using numerical

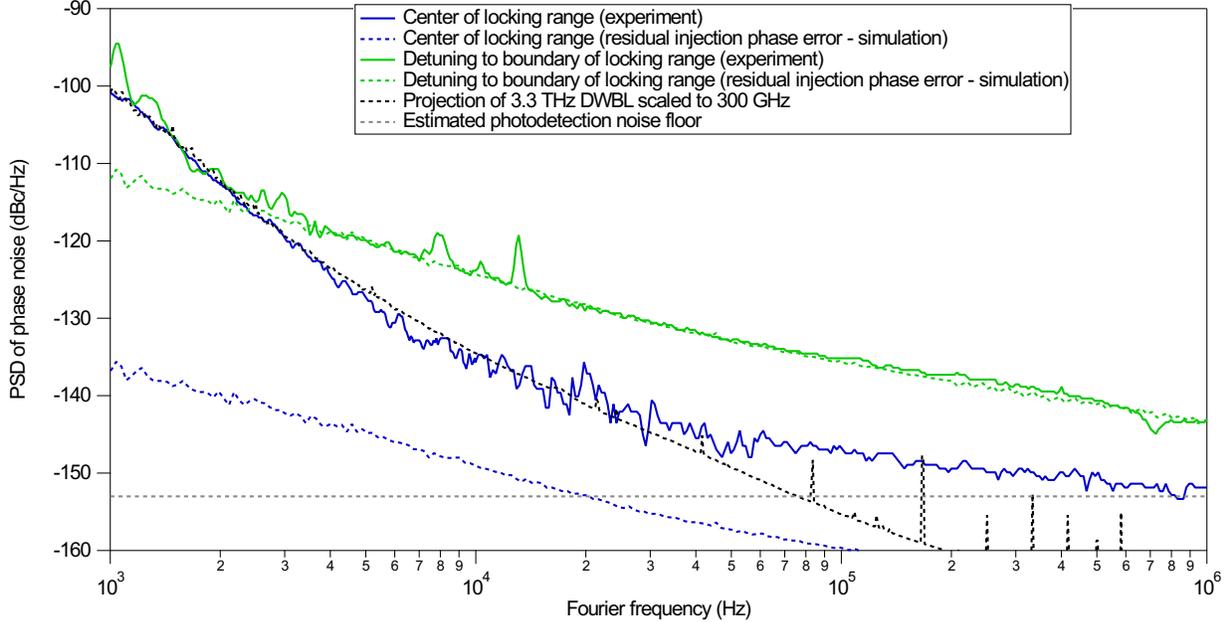


Figure S2: **Injection-locking phase-noise spectra: experiment versus simulation.** Repetition-rate phase-noise power spectral densities are shown for two operating points: tuning to the center of the locking range (blue) and detuning to the boundary of the locking range (green). In the experiment, the boundary-locking condition is identified by the abrupt collapse of the repetition-rate beat note from a finite offset (typically ~ 20 MHz) to DC, marking the onset of injection locking. The agreement between experiment (solid lines) and simulation (dashed lines) confirms bandwidth-limited reference tracking and reproduces the increased phase noise observed near the locking boundary. Projected phase noise of the dual-wavelength Brillouin laser (DWBL) optical reference and photodetection noise floor are included for reference.

simulations that start from the measured free-running phase noise of a 300 GHz Kerr microcomb repetition rate and the measured phase noise of a 3.3 THz optical reference.

S4.1 Idealized PLL model

The PLL-based divider (also called 2P-OFD for two point OFD) is modeled as a standard type-II, second-order closed-loop system, equivalent to a well-tuned proportional–integral–derivative (PID) controller in the phase domain. To allow direct comparison with injection locking, the PLL is parameterized by a single locking bandwidth, while the damping ratio is fixed to

$$\zeta = 0.707, \quad (10)$$

corresponding to a critically damped loop.

The closed-loop phase transfer functions are written as [4]

$$T_{\text{PLL}}(s) = \frac{2\zeta\omega_n s + \omega_n^2}{s^2 + 2\zeta\omega_n s + \omega_n^2}, \quad S_{\text{PLL}}(s) = \frac{s^2}{s^2 + 2\zeta\omega_n s + \omega_n^2}, \quad (11)$$

where T_{PLL} describes the transfer of reference phase noise to the output, S_{PLL} describes the suppression of free-running oscillator noise, and $\omega_n = 2\pi f_n$ is the loop natural frequency. The parameter f_n serves as a practical proxy for the PLL locking bandwidth.

Importantly, this PLL model assumes an ideal phase detector and ignores SNR limitations, detection noise, latency, and quantization effects. As a result, the simulated PLL performance represents a best-case upper bound for practical PLL-based OFD.

S4.2 Simulated comparison with Kerr optical frequency division

Figure S3 compares the simulated repetition-rate phase-noise spectra obtained using Kerr OFD and the idealized PLL-based divider. For Kerr OFD, the reference and free-running noise contributions are combined using the injection-locking transfer functions derived in Sec. S3, with injection bandwidths ranging from 20 MHz to 280 MHz, corresponding to the measured bandwidth of the microresonator used in this work. For the PLL-based divider, type-II loops with locking bandwidths between 10 kHz and 1 MHz are shown, corresponding to the performance of a PLL of exceptional quality and implementation.

Despite its formally type-I nature, Kerr OFD benefits from a substantially larger achievable locking bandwidth, which allows more effective suppression of reference noise over a wide Fourier-frequency range. In contrast, the performance of the PLL-based divider is fundamentally limited by its loop bandwidth, even under the idealized assumptions used here. As a result, Kerr OFD outperforms the type-II PLL in the high-offset-frequency regime most relevant for millimeter-wave and terahertz applications.

This comparison highlights the central role of locking bandwidth in optical frequency division at very high carrier frequencies, and illustrates why wide-bandwidth injection locking provides a favorable scaling path beyond the capabilities of conventional feedback-based division schemes.

S4.3 Implications for high-frequency division and timing jitter

A key qualitative difference between injection-locked Kerr OFD and feedback-based PLL division lies in the spectral structure introduced by the control loop itself. In a type-II PLL, aggressive noise suppression at low Fourier frequencies is achieved at the expense of pronounced servo-induced features, commonly referred to as servo bumps, near the loop bandwidth. These features arise from the resonant dynamics of the closed-loop transfer function and are an intrinsic consequence of feedback control.

As illustrated in Fig. S3, the presence of servo bumps significantly elevates the phase-noise spectrum in the vicinity of the loop bandwidth, leading to excess noise that directly degrades integrated timing jitter. In the highest-bandwidth iteration (1 MHz) of the simulated best-case PLL scenarios considered here, the phase noise at high Fourier frequencies is more than 20 dB higher than that obtained using the worst-case, boundary injection locking bandwidth (20 MHz), even when the PLL is assumed to be ideal and free of detector noise or latency.

In practice, pushing PLL bandwidths beyond the megahertz range is extremely challenging due to electronic delay, actuator limitations, and stability margins. As a result, the servo-induced noise amplification cannot be easily eliminated and typically becomes more severe as loop bandwidth is increased. Injection locking avoids these limitations entirely by relying on intrinsic oscillator dynamics rather than explicit feedback, enabling smooth, resonance-free noise suppression over bandwidths that are inaccessible to conventional PLL implementations.

More generally, the relatively higher free-running phase noise of Kerr microcombs, which arises from their lower effective resonator quality factor compared to ultrastable mode-locked lasers, can be advantageous for injection-locked optical frequency division. The broader intrinsic linewidth of the free-running soliton oscillator enables exceptionally large injection-locking bandwidths, whereas

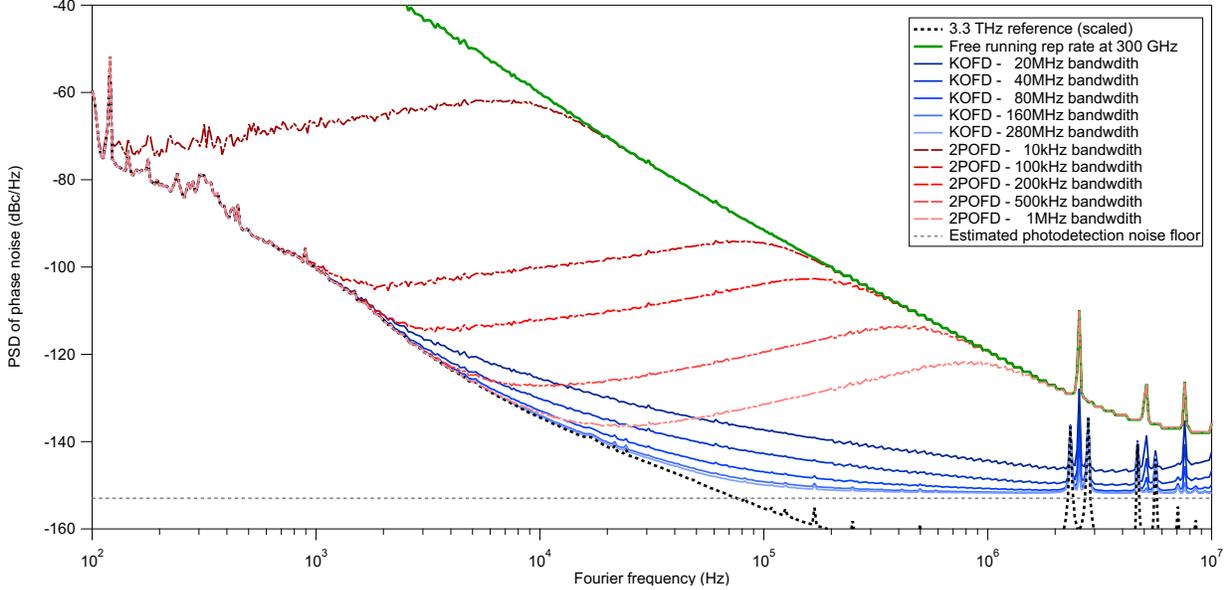


Figure S3: **Simulated comparison of Kerr optical frequency division (OFD) and PLL-based OFD.** Simulated repetition-rate phase-noise spectra starting from the measured free-running 300 GHz microcomb phase noise (green) and a 3.3 THz optical reference, divided to 300 GHz (black, dashed). Kerr OFD is shown for injection bandwidths between 20 MHz and 280 MHz (shades of blue) with an estimated shot-noise floor of -153 dBc/Hz, while a type-II, second-order PLL (2P-OFD) is shown for locking bandwidths from 10 kHz to 1 MHz (shades of red, dashed). The PLL model assumes an ideal phase detector and neglects SNR and latency limitations, representing a best-case performance scenario. Despite its type-I character, the large achievable locking bandwidth of Kerr OFD enables superior noise suppression at high Fourier frequencies compared with bandwidth-limited PLL-based schemes.

feedback-based PLLs and two-point OFD architectures are better suited to ultrastable combs with intrinsically narrow linewidths but correspondingly limited locking or tracking bandwidth.

S5 Noise contributions and scaling laws

The phase noise of the Kerr OFD repetition rate originates from several physically distinct noise sources whose contributions are transferred through the injection-locking dynamics described in Sec. S3. These include low-frequency technical noise associated with environmental perturbations, fundamental phase noise of the optical-difference reference, noise introduced by photodetection and electronic readout, and residual intensity fluctuations of the soliton comb that may couple to phase through amplitude-to-phase conversion. Each contribution arises from a different physical mechanism and follows a characteristic scaling law, allowing their effects on the repetition-rate phase noise to be identified and analyzed separately.

S5.1 Low-offset technical noise and $\Delta\nu$ scaling

Environmental perturbations, such as fluctuations of optical path length in fiber or free-space links, are converted into optical phase noise according to

$$\delta\phi(t) = \omega \delta\tau(t) = \omega \frac{n_g}{c} \delta L(t), \quad (12)$$

where $\delta L(t)$ represents path-length fluctuations and n_g is the group index. For the optical-difference phase $\phi_\Delta(t) = \phi_i(t) - \phi_p(t)$, common-mode contributions cancel to first order, yielding

$$\delta\phi_\Delta(t) = (\omega_i - \omega_p) \frac{n_g}{c} \delta L(t) = \Omega \frac{n_g}{c} \delta L(t), \quad (13)$$

with $\Omega = 2\pi\Delta\nu$.

As a result, the phase-noise power spectral density of the optical-difference reference exhibits the scaling

$$S_{\phi_\Delta}(f) \propto \Delta\nu^2 S_{\delta L}(f), \quad (14)$$

demonstrating that low-frequency technical noise increases quadratically with the optical frequency separation.

After division by the factor N , these low-offset technical noise contributions are transferred to the repetition rate with unity gain at Fourier frequencies well below the locking bandwidth. Consequently, they are only reduced by $1/N^2$ and typically dominate the phase-noise spectrum at small offset frequencies.

S5.2 Schawlow–Townes limit and quantum noise

For an optical tone characterized by a Schawlow–Townes linewidth $\Delta\nu_{\text{ST}}$, the associated fundamental frequency noise is white, with single-sided power spectral density $S_\nu = 2\Delta\nu_{\text{ST}}$. This leads to a phase-noise power spectral density

$$S_\phi(f) = \frac{S_\nu}{(2\pi f)^2}, \quad (15)$$

corresponding to a $1/f^2$ dependence characteristic of quantum-limited frequency noise.

For the optical-difference reference formed by two uncorrelated optical tones, the frequency-noise densities add,

$$S_{\nu,\Delta} = S_{\nu,i} + S_{\nu,p}. \quad (16)$$

Within the locking bandwidth, this contribution is transferred to the soliton repetition rate according to the ideal division law derived in Sec. S3, yielding

$$S_{\theta_{\text{rep}}}(f) = \frac{1}{N^2} \frac{S_{\nu,\Delta}}{(2\pi f)^2}. \quad (17)$$

In terms of single-sideband phase noise, this corresponds to a reduction of $20 \log_{10} N$ relative to the optical-difference reference.

This $1/f^2$ phase-noise region represents the quantum-limited noise floor of the Kerr OFD repetition rate within the locking bandwidth and prior to the onset of photodetection and electronic noise. In Brillouin-based optical references, this limit is set by the fundamental linewidth of the Brillouin lasers and, in general, includes contributions associated with thermal phonon-mediated quantum noise.

S5.3 Photodetection noise: thermal and shot-noise limits at 300 GHz

This subsection estimates the fundamental white phase-noise floor imposed by the photodetection and RF extraction of the 300 GHz carrier. Throughout this subsection, single-sideband (SSB) phase noise is reported as $L(f)$ in dBc/Hz, while independent noise contributions are summed in *linear* units prior to conversion to dBc/Hz.

Let P_μ denote the measured carrier power at $f_\mu = 300$ GHz delivered to a matched 50Ω load at the reference plane of interest. Following standard practice, we define the *linear* SSB phase-noise ratio

$$L_{\text{lin}}(f) \equiv 10^{L(f)/10}, \quad (18)$$

so that independent contributions add as $L_{\text{lin,tot}} = \sum_k L_{\text{lin},k}$ and

$$L_{\text{tot}}(f) [\text{dBc/Hz}] = 10 \log_{10} \left(\sum_k 10^{L_k(f)/10} \right). \quad (19)$$

Thermal (Johnson–Nyquist) noise. A matched resistive environment at effective temperature T_{eff} contributes a broadband thermal noise density. The corresponding SSB phase-noise floor in a 1 Hz bandwidth is [5]

$$L_{\text{th}}(f) = 10 \log_{10} \left(\frac{k_B T_{\text{eff}}}{2P_\mu} \right) [\text{dBc/Hz}], \quad (20)$$

which is white versus Fourier frequency.

Shot noise. The photocurrent shot noise has current PSD $S_i = 2qI_{\text{avg}} = 2\eta \frac{q^2}{h\nu_{\text{opt}}} P_{\text{opt}}$, where I_{avg} is the average photocurrent, η is the quantum efficiency of the detector, h is Planck's constant, ν_{opt} is the optical frequency, and P_{opt} is the optical power. For pulsed photodetection, the resulting shot-noise contribution to the microwave phase noise at frequency f_μ depends on the optical pulse width and the electrical transfer function of the detection chain. While the expressions below are written for Gaussian optical pulses for analytical convenience, sech^2 pulse shapes yield the same dependence on pulse width and detection bandwidth, differing only by order-unity numerical factors. For Gaussian optical pulses, the single-sideband phase-noise contribution can be expressed as [5]:

$$L_{\text{shot}}(f) = 10 \log_{10} \left(\frac{qI_{\text{avg}} |H(f_\mu)|^2 R}{P_\mu} [1 - \exp(-(2\pi f_\mu \tau_{\text{opt}})^2)] \right) [\text{dBc/Hz}], \quad (21)$$

where $R = 50 \Omega$, $|H(f_\mu)|^2$ is the (dimensionless) small-signal microwave transfer function at f_μ , and τ_{opt} is the $1/e$ half-width of the optical pulse intensity profile (including chirp). The bracketed factor satisfies $0 \leq 1 - \exp(-(2\pi f_\mu \tau_{\text{opt}})^2) \leq 1$ and quantifies the extent to which shot noise projects onto the phase quadrature.

Optical amplifier ASE-signal noise. In experiments that make use of an optical amplifier, such as ours, the gain, G , and noise figure, F , of the amplifier lead to a white noise floor, with the dominant effect arising from beating between the broadband ASE spectrum and the coherent comb [5]. For a Gaussian pulsed source, the ASE-signal noise contributes an SSB phase-noise floor of,

$$L_{\text{ASE-sig}}(f) = 10 \log_{10} \left(\frac{\eta^2 \frac{q^2}{h\nu} G^2 F P_{\text{in}} |H(f_\mu)|^2 R}{P_\mu} [1 - \exp(-(2\pi f_\mu \tau_{\text{opt}})^2)] \right) [\text{dBc/Hz}], \quad (22)$$

where P_{in} is the input optical power to the optical amplifier. For our experiment, $G \approx 35$ dB, $F \approx 5$ dB, $\eta \approx 0.15$, $P_{\text{in}} \approx 5 \mu\text{W}$, and $|H(f_\mu)|^2 \approx 0.05$. Assuming a Fourier-transform-limited pulse after dispersion compensation, $\tau_{\text{opt}} \approx 70$ fs denotes the intensity full width at half maximum (FWHM) of the optical pulse.

Numerical evaluation from measured $P_\mu(I)$. Table S1 reports the thermal-noise limit (Eq. 20), the shot-noise (Eq. 21), the amplifier noise (Eq. 22), and their combined contribution computed via Eq. (19) with $T_{\text{eff}} = 290$ K. Here P_μ is the measured 300 GHz power delivered to 50Ω and I_{avg} is the measured average photocurrent.

Table S1: Photodetection noise floors at $f_\mu = 300$ GHz computed from measured carrier power P_μ versus photocurrent. L_{th} uses Eq. (20) with $T_{\text{eff}} = 290$ K. L_{shot} is the shot-noise phase-noise floor of Eq. (21), and $L_{\text{ASE-sig}}$. The total floor is obtained by linear summation as $L_{\text{sum}} = 10 \log_{10}(10^{L_{\text{th}}/10} + 10^{L_{\text{shot}}/10} + 10^{L_{\text{ASE-sig}}/10})$.

I_{avg} (mA)	P_μ (μW)	L_{shot} (dBc/Hz)	L_{th} (dBc/Hz)	$L_{\text{ASE-sig}}$ (dBc/Hz)	L_{sum} (dBc/Hz)
2	18.9	-181.43	-159.75	-149.67	-149.26
2.5	23.75	-181.50	-160.74	-149.74	-149.40
3	30.2	-181.73	-161.78	-149.96	-149.69
3.5	38.1	-182.07	-162.78	-150.30	-150.06
4	50	-182.58	-163.91	-150.82	-150.61
5	70.7	-183.16	-165.48	-151.40	-151.23
6	100.1	-183.88	-166.99	-152.12	-151.98
7	129.37	-184.25	-168.11	-152.49	-152.37
8	167.7	-184.78	-169.23	-153.02	-152.92

The values in Table S1 show that, given the measured carrier power at 300 GHz, the white noise floor originating from the optical amplifier dominates the fundamental additive noise floor over the explored operating range, while the thermal contribution is at least 7 dB lower, and the shot noise is tens of decibels lower yet. We note that, since the white phase noise levels are significantly different from each other, cross-spectral collapse is not present in the measurement [6]. Any measured white phase-noise level significantly above L_{sum} therefore indicates excess noise mechanisms beyond fundamental detection noise (e.g., ASE-related intensity noise and/or AM-to-PM conversion), which are addressed separately.

Experimental limitation and operating point selection. Despite the favorable detection-noise limits estimated above, we did not observe a clean, photocurrent-independent white phase-noise floor that could be unambiguously attributed to the fundamental thermal or shot-noise limits. In practice, the measured high-frequency phase noise exhibited a residual dependence on photocurrent, indicating the presence of excess noise associated with amplitude-to-phase (AM-to-PM) conversion in the UTC photodiode and subsequent RF chain. This excess conversion noise, which depends on bias point and optical operating conditions, inevitably limits the experimental rigor with which the fundamental detection floor can be identified.

For this reason, all Kerr OFD measurements reported in the main text were performed at a fixed average photocurrent of approximately $I_{\text{avg}} = 4$ mA, which was experimentally identified as an operating point where AM-to-PM conversion was strongly minimized while maintaining high 300 GHz carrier power. This choice represents a practical trade-off between carrier power and excess detection noise, and ensures the best achievable phase-noise performance in the present system.

S5.4 Comb relative intensity noise and AM-to-PM-induced phase noise

Relative intensity noise (RIN) of the optical comb lines can contribute to the measured repetition-rate phase noise through amplitude-to-phase (AM-to-PM) conversion in the photodetector and downstream RF electronics. This mechanism becomes particularly relevant when operating at photocurrents that do not minimize AM-to-PM conversion, or in the presence of excess optical intensity noise, and typically manifests over an intermediate range of Fourier frequencies.

To quantify this effect, we consider the phase-noise contribution expected from RIN mapped through AM-to-PM conversion during photodetection. If $S_{\text{RIN}}(f)$ denotes the double-sideband fractional intensity-noise power spectral density and α is the effective AM-to-PM conversion coefficient (in radians per unit fractional power), assumed here to be weakly frequency dependent over the band of interest, the corresponding single-sideband phase-noise contribution is given by

$$L_{\phi, \text{AM-to-PM}}(f) = \frac{\alpha^2}{4} S_{\text{RIN}}(f). \quad (23)$$

Figure S4 shows the measured Kerr OFD phase-noise spectrum at 300 GHz together with the phase-noise contribution expected from AM-to-PM conversion of the measured pump RIN for filtered and unfiltered pump amplified spontaneous emission (ASE). The Kerr OFD trace is obtained at a photocurrent (8 mA) chosen to deliberately enhance AM-to-PM sensitivity, thereby directly revealing the RIN-induced phase-noise contribution.

Using a single heuristic value of α , the RIN mapped through Eq. (23) quantitatively reproduces the observed phase-noise plateau between approximately 10 kHz and 100 kHz for the filtered-pump condition, indicating that comb-line intensity noise dominates the phase noise in this Fourier frequency range. At Fourier frequencies above ~ 100 kHz, the dark gray trace becomes limited by the measurement floor of the single-comb-line RIN characterization, which explains why the measured Kerr OFD phase noise at 300 GHz does not follow this trend at higher offset frequencies. Applying the same conversion coefficient to the unfiltered-pump RIN predicts a substantially larger phase-noise contribution (see light gray trace), demonstrating that pump ASE filtering is essential to suppress comb-line RIN and avoid AM-to-PM-limited millimeter-wave performance.

While the AM-to-PM conversion coefficient is not independently calibrated here, the agreement between the mapped RIN contribution and the measured phase-noise plateau confirms the physical origin of the excess noise and validates the noise attribution.

S6 Timing jitter and Allan deviation

Phase-noise spectra provide a frequency-domain description of oscillator stability, while their impact in the time domain is commonly quantified through timing jitter and Allan deviation. For the Kerr OFD repetition rate, timing jitter directly reflects fluctuations of the zero crossings of the millimeter-wave signal derived from the soliton pulse train.

S6.1 Timing jitter

The root-mean-square (RMS) timing jitter σ_τ is obtained by integrating the repetition-rate phase-noise PSD according to

$$\sigma_\tau^2 = \frac{1}{(2\pi f_{\text{rep}})^2} \int_{f_1}^{f_2} S_{\theta_{\text{rep}}}(f) df, \quad (24)$$

where f_1 and f_2 define the lower and upper integration limits. Using the transfer functions derived in Sec. S3, the total jitter can be decomposed into independent contributions associated with reference noise, free-running soliton noise, and additive white noise floors.

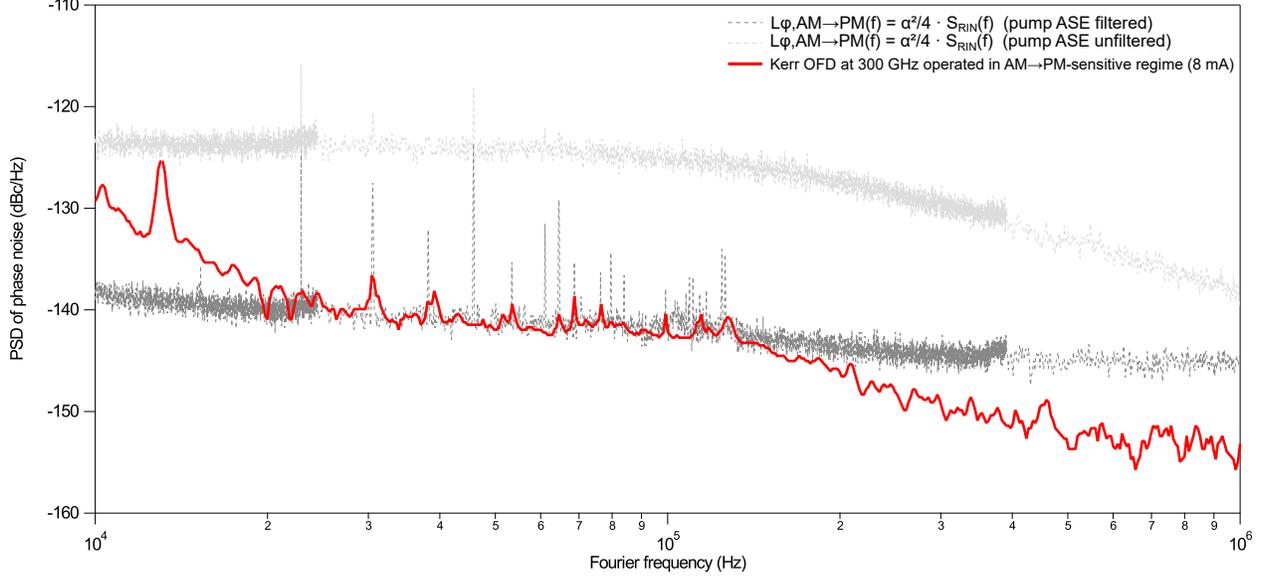


Figure S4: **AM-to-PM-limited phase noise in Kerr OFD.** Phase-noise spectra at 300 GHz showing the contribution expected from pump relative intensity noise (RIN) mapped into phase noise via AM-to-PM conversion, $L_{\phi,AM\text{-to-}PM}(f) = \alpha^2 S_{RIN}(f)/4$, for filtered and unfiltered pump amplified spontaneous emission (dark and light gray traces, respectively). The red trace corresponds to Kerr optical frequency division operated in an AM-to-PM-sensitive regime (8 mA photocurrent), intentionally enhancing AM-to-PM coupling to directly reveal the RIN-limited phase-noise contribution.

When ideal division places the quantum-limited reference noise below the photodetection shot-noise floor at higher Fourier frequencies, the jitter integral is typically dominated by low-offset technical noise scaling with $\Delta\nu$. In this regime, the contribution from high-frequency noise is reduced to the attosecond level, consistent with the measured phase-noise spectra.

S6.2 Allan deviation

While phase-noise spectra provide a detailed frequency-domain characterization of the Kerr OFD repetition rate, long-term frequency stability is more naturally quantified using Allan deviation. Allan deviation measurements of the divided millimeter-wave signal were therefore performed to independently verify the low-frequency behavior inferred from the phase-noise analysis.

Figure S5 shows the measured fractional frequency stability as a function of averaging time. The observed Allan deviation exhibits a monotonic increase with averaging time, consistent with the presence of low-frequency technical noise dominating the spectrum at small Fourier frequencies. This behavior is expected from the $\Delta\nu$ -scaled environmental noise discussed in Sec. S3.1 and is consistent with the absence of additional long-term stabilization loops in the Kerr OFD system.

Importantly, the Allan deviation measurements do not reveal excess instability beyond that predicted from the measured phase-noise spectra. This confirms that the phase-noise data presented in the main text and Supplementary Material provide an accurate description of the system stability over the full range of relevant timescales. The Allan deviation therefore serves as an independent consistency check linking the frequency-domain noise analysis to time-domain stability.

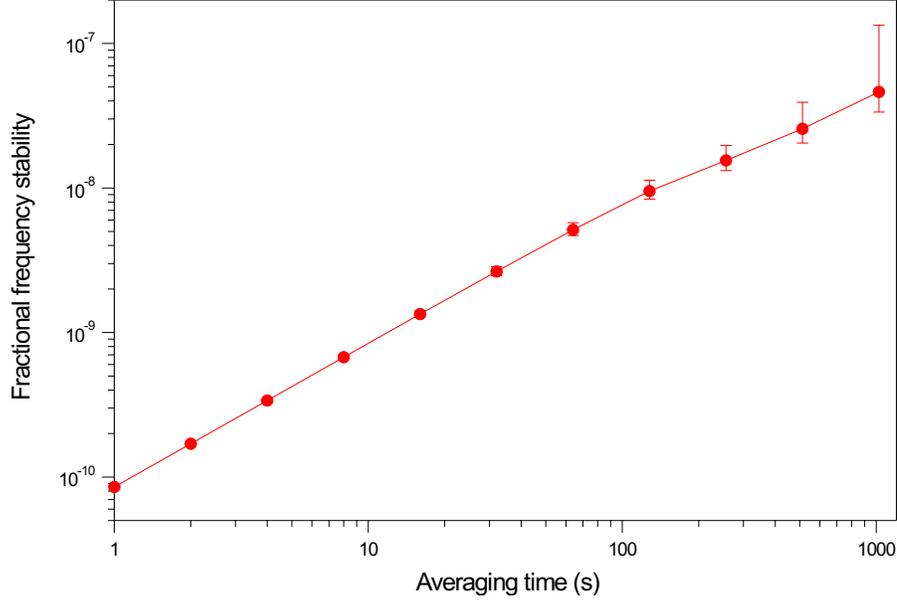


Figure S5: **Allan deviation of the Kerr OFD repetition rate.** Measured fractional frequency stability as a function of averaging time for the millimeter-wave signal. The monotonic increase with averaging time is consistent with low-frequency technical noise dominating the spectrum at small Fourier frequencies, in agreement with the phase-noise analysis presented in the main text and Supplementary Material.

S7 Summary of closed-form results

For Kerr optical frequency division of an optical-difference reference $\Delta\nu$ to a millimeter-wave repetition rate f_{rep} , the key closed-form relations are summarized below.

1. **Injection locking and exact division law (inside the locking band).** With the relative phase $\psi(t) = N\theta_{\text{rep}}(t) - \phi_{\Delta}(t)$ governed by an Adler-type dynamics (Sec. S3), a phase-locked solution exists within the injection range $|\Delta\omega| \leq K$. In this regime,

$$f_{\text{rep}} = \frac{\Delta\nu}{N}, \quad \theta_{\text{rep}}(t) = \frac{\phi_{\Delta}(t)}{N} + \frac{\psi^*}{N}, \quad (25)$$

where $\psi^* = \arcsin(\Delta\omega/K)$ is the equilibrium phase offset. Equation (25) is the exact division law: within the tracking bandwidth, the repetition-rate phase is rigidly constrained to the optical-difference phase, reduced by the factor N .

2. **Finite tracking bandwidth.** Small perturbations about ψ^* are tracked only within the effective injection-locking bandwidth

$$f_c = \frac{K \cos \psi^*}{2\pi}, \quad (26)$$

which sets the Fourier-frequency range over which ideal division is enforced. Near the boundary of the locking range, $|\psi^*| \rightarrow \pi/2$ and $\cos \psi^* \rightarrow 0$, so f_c collapses even though phase locking may still be maintained.

3. **Phase-noise transfer functions (reference vs. free-running noise).** Linearizing the Adler dynamics about ψ^* yields the frequency-domain relation (Sec. S3)

$$\delta\Theta_{\text{rep}}(j\omega) = \underbrace{\frac{1}{N} \frac{j\omega}{j\omega + 2\pi f_c}}_{H_m(j\omega)} \delta\Phi_{\Delta}(j\omega) + \underbrace{\frac{j\omega}{j\omega + 2\pi f_c}}_{H_s(j\omega)} \delta\Theta_{\text{rep}}^{(0)}(j\omega) + H_{\xi}(j\omega) \Xi(j\omega), \quad (27)$$

where $\delta\Theta_{\text{rep}}^{(0)}$ is the free-running repetition-rate phase fluctuation, and Ξ represents additive noise associated with the injection process and readout.

The corresponding repetition-rate phase PSD is

$$S_{\theta_{\text{rep}}}(f) = |H_m|^2 S_{\phi_{\Delta}}(f) + |H_s|^2 S_{\theta_{\text{rep}}}^{(0)}(f) + |H_{\xi}|^2 S_{\xi}(f). \quad (28)$$

4. **Asymptotes and practical interpretation.** From Eqs. (27)–(28):

$$f \ll f_c: \quad |H_m|^2 \rightarrow 1, \quad |H_s|^2 \rightarrow 0 \quad \Rightarrow \quad S_{\theta_{\text{rep}}}(f) \simeq \frac{1}{N^2} S_{\phi_{\Delta}}(f), \quad (29)$$

$$f \gg f_c: \quad |H_m|^2 \sim \left(\frac{f_c}{f}\right)^2, \quad |H_s|^2 \rightarrow 1 \quad \Rightarrow \quad S_{\theta_{\text{rep}}}(f) \rightarrow S_{\theta_{\text{rep}}}^{(0)}(f) \text{ (plus white floors)}. \quad (30)$$

Thus, f_c directly determines where the spectrum transitions from division-limited (reference-dominated) to soliton-limited (free-running).

5. **Noise scalings of the optical-difference reference.**

- *Low-offset technical noise and $\Delta\nu$ scaling (Sec. S5.1).* Path-length fluctuations $\delta L(t)$ imprint on the optical-difference phase as $\delta\phi_{\Delta}(t) = \Omega(n_g/c)\delta L(t)$ with $\Omega = 2\pi\Delta\nu$, so

$$S_{\phi_{\Delta}}(f) \propto \Delta\nu^2 S_{\delta L}(f), \quad (31)$$

and after division the corresponding term scales as $S_{\theta_{\text{rep}}} \propto (\Delta\nu^2/N^2)$ at Fourier frequencies where it dominates.

- *Schawlow–Townes / quantum-limited $1/f^2$ phase noise (Sec. S5.2).* White frequency noise S_{ν} produces

$$S_{\phi_{\Delta}}(f) = \frac{S_{\nu,\Delta}}{(2\pi f)^2}, \quad S_{\theta_{\text{rep}}}(f) = \frac{1}{N^2} \frac{S_{\nu,\Delta}}{(2\pi f)^2} \quad (f \ll f_c), \quad (32)$$

i.e., a $20 \log_{10} N$ reduction in single-sideband phase noise inside the tracking bandwidth.

6. **Photodetection white floors (shot and thermal) expressed in dBc/Hz.** At sufficiently large Fourier frequencies (where the divided reference noise is below the detector/electronics floor), the measured single-sideband phase noise approaches a white floor set by additive noise in the RF readout chain.

For an RF tone at the carrier frequency f_{rep} with carrier power P_{RF} delivered to $R = 50 \Omega$, additive white *power* noise spectral density at the output, N_{add} (W/Hz), produces a phase-noise floor

$$L_{\text{add}} [\text{dBc/Hz}] = 10 \log_{10} \left(\frac{N_{\text{add}}}{P_{\text{RF}}} \right), \quad (33)$$

where N_{add} is the *sum in linear units* of independent white-noise contributions.

For the two dominant white terms:

- *Thermal noise (Johnson–Nyquist)* of a matched $50\text{-}\Omega$ system at temperature T :

$$N_{\text{th}} = k_B T \quad \Rightarrow \quad L_{\text{th}} = 10 \log_{10} \left(\frac{k_B T}{P_{\text{RF}}} \right). \quad (34)$$

- *Shot-noise-limited additive power noise* referred to the RF tone can be written in the same form,

$$L_{\text{shot}} = 10 \log_{10} \left(\frac{N_{\text{shot}}}{P_{\text{RF}}} \right), \quad (35)$$

where N_{shot} depends on the photocurrent I_{dc} and the photodetection transfer to the RF harmonic being used. In practice, for photonic-microwave generation it is common to estimate L_{shot} using validated photodetection floor models (e.g. the pulse-train/comb-based shot-noise floor treatment) rather than a CW approximation when the pulse shape/duty cycle is unknown or not independently measured.

Independent white floors add *linearly* in N_{add} :

$$N_{\text{add}} = N_{\text{th}} + N_{\text{shot}} + \dots, \quad L_{\text{floor}} = 10 \log_{10} \left(\frac{N_{\text{add}}}{P_{\text{RF}}} \right). \quad (36)$$

7. **AM-to-PM conversion and RIN-driven plateaus.** Comb-line RIN, $S_{\text{RIN}}(f)$, converts to phase noise via AM-to-PM with coefficient α (rad per fractional power),

$$L_{\phi, \text{AM-to-PM}}(f) = 10 \log_{10} \left(\frac{\alpha^2}{4} S_{\text{RIN}}(f) \right), \quad (37)$$

and can dominate over the white shot/thermal floors over intermediate Fourier frequencies when operating away from an AM-to-PM null.

8. **Experimental practical note (floor identification).** Although the closed-form relations above predict a clean transition to a white photodetection floor at high Fourier frequencies, experimentally the extraction of an unambiguous shot/thermal-limited floor can be hindered by photocurrent-dependent excess noise introduced by AM-to-PM in the UTC and downstream electronics. In this work, the best Kerr OFD phase-noise data were therefore taken at a heuristic operating point (photocurrent ~ 4 mA) where AM-to-PM was found to be strongly minimized, enabling the most faithful comparison to the injection-locking transfer-function framework.

Together, Eqs. (25)–(28) provide the compact “system law” for Kerr OFD: (i) exact division by N inside the tracking bandwidth f_c , (ii) a first-order roll-off of reference tracking beyond f_c , and (iii) a transparent decomposition of the measured phase-noise spectrum into reference noise, free-running soliton noise, and additive readout floors.

References

- [1] Brendan M Heffernan, James Greenberg, Takashi Hori, Tatsuya Tanigawa, and Antoine Rolland. Brillouin laser-driven terahertz oscillator up to 3 thz with femtosecond-level timing jitter. *Nature Photonics*, 18(12):1263–1268, 2024.
- [2] Robert Adler. A study of locking phenomena in oscillators. *Proceedings of the IRE*, 34(6):351–357, 1946.

- [3] Tomohiro Tetsumoto, Tadao Nagatsuma, Martin E Fermann, Gabriele Navickaite, Michael Geiselmann, and Antoine Rolland. Optically referenced 300 ghz millimetre-wave oscillator. *Nature Photonics*, 15(7):516–522, 2021.
- [4] Floyd M. Gardner. *Phaselock Techniques*. Wiley, 3 edition, 2005.
- [5] Franklyn Quinlan. The photodetection of ultrashort optical pulse trains for low noise microwave signal generation. *Laser & photonics reviews*, 17(12):2200773, 2023.
- [6] C. W. Nelson, A. Hati, and D. A. Howe. A collapse of the cross-spectral function in phase noise metrology. *Review of Scientific Instruments*, 85:024705, 2014.