

This is the accepted manuscript made available via CHORUS. The article has been published as:

Analytically exploiting noise correlations inside the feedback loop to improve locked-oscillator performance

J. Sastrawan, C. Jones, I. Akhalwaya, H. Uys, and M. J. Biercuk

Phys. Rev. E **94**, 022204 — Published 8 August 2016

DOI: [10.1103/PhysRevE.94.022204](https://doi.org/10.1103/PhysRevE.94.022204)

Analytically exploiting noise correlations inside the feedback loop to improve locked-oscillator performance

J. Sastrawan,¹ C. Jones,¹ I. Akhalwaya,² H. Uys,² and M.J. Biercuk^{1,*}

¹*ARC Centre for Engineered Quantum Systems, School of Physics, The University of Sydney, NSW 2006 Australia and National Measurement Institute, West Lindfield, NSW 2070 Australia*

²*National Laser Centre, Council for Scientific and Industrial Research, Pretoria, South Africa*

(Dated: July 14, 2016)

We introduce concepts from optimal estimation to the stabilization of precision frequency standards limited by noisy local oscillators. We develop a theoretical framework casting various measures for frequency standard variance in terms of frequency-domain transfer functions, capturing the effects of feedback stabilization via a time-series of Ramsey measurements. Using this framework we introduce a novel optimized hybrid *predictive feedforward* measurement protocol which employs results from multiple past measurements and transfer-function-based calculations of measurement covariance to improve the accuracy of corrections within the feedback loop. In the presence of common non-Markovian noise processes these measurements will be correlated in a calculable manner, providing a means to capture the stochastic evolution of the LO frequency during the measurement cycle. We present analytic calculations and numerical simulations of oscillator performance under competing feedback schemes and demonstrate benefits in both correction accuracy and long-term oscillator stability using hybrid feedforward. Simulations verify that in the presence of uncompensated dead time and noise with significant spectral weight near the inverse cycle time predictive feedforward outperforms traditional feedback, providing a path towards developing a new class of stabilization “software” routines for frequency standards limited by noisy local oscillators.

High-performance passive frequency standards play a major role in technological applications such as network synchronization and GPS [1] as well as many fields of physical inquiry, including radioastronomy (very-long-baseline interferometry) [2], tests of general relativity [3], and particle physics [4]. Atomic clocks exploiting the stability of Cs [5–8] or other atomic references [9–13] to stabilize an oscillator are known as the most precise timekeeping devices available, but constant performance gains are sought for technical and scientific applications.

In many settings, such as miniaturized deployable frequency standards or in GPS-denied environments, a major performance limitation arises from the quality of the local oscillator (LO) that probes and is locked to the atomic transition. The LO frequency may evolve randomly in time due to intrinsic noise processes in the underlying hardware [10, 11], leading to time-varying deviations of the LO frequency from that of the stable atomic reference. These instabilities are *partially* compensated through use of a feedback protocol designed to transfer the stability of the reference to the LO, but their effects cannot be mitigated completely.

Early work characterizing the so-called Dick effect [14] demonstrated that no matter how good the reference becomes, LO noise will still produce residual instabilities in the locked LO (LLO) through the feedback protocol itself. The dominant mechanism for this is evolution of the LO’s frequency on timescales rapid compared with the shortest measurement and feedback cycle. Major contributors to this phenomenology relate to the presence

of uncompensated LO evolution during initialization and readout stages of the measurement cycle (dead time), as well as aliasing of LO noise at harmonics of the feedback-loop period – the Dick effect [14–16]. Accordingly, significant research focus in the frequency standards community has been placed on improving LO performance, using e.g. ultra-low-phase-noise cryogenic sapphire oscillators or similar [17, 18], with concomitant increases in hardware infrastructure requirements and complexity. Other approaches to mitigating the impact of LO instabilities involve significant modification of the relevant reference hardware, for instance employing multiple atomic references [10, 19].

In this Manuscript we devise and analyze a method by which both the accuracy of the LLO relative to the atomic reference, and the stability of the composite passive frequency standard, can be improved without the need for hardware modification. We develop new analytic tools casting time-domain statistical measures of frequency-standard performance in terms of analytically calculable transfer functions [20], exploiting recent related work in quantum information [21–24]. This approach allows improvement in feedback stabilization by bringing optimal estimation inside the feedback loop of the LO in order to exploit non-Markovianity in the dynamics of LO frequency fluctuations.

Our method expresses the properties of the LLO in terms of the statistics of the unlocked LO at different times as well as correlations between those measurements. We present the relevant transfer functions for time-series measurements of arbitrary-duration Ramsey measurements, and introduce the pair-covariance transfer function explicitly capturing correlations between measurement outcomes at different times. Thus,

* michael.biercuk@sydney.edu.au

given statistical knowledge of the LO noise characteristics, we craft a new form of hybrid feedforward stabilization incorporating the results of an arbitrary number of past measurements with variable duration to calculate an improved *predictive* correction to the LO. This approach shares concepts with techniques of optimal estimation [25] commonly used in engineering to predict the evolution of a dynamical system – here the noisy LO.

In cases where dead time is significant and there is substantial uncompensated LO evolution, we use numerical simulations to show that this approach allows corrections of improved accuracy to be applied to the LO. Simulations also demonstrate that long-term stability of the LLO is improved through a feedforward correction scheme, where corrections are made based on weighting values determined analytically in the same hybrid feedforward approach. The method described here is a technology-independent *software-oriented* approach to improving the performance of frequency standards derived from locked local oscillators. It may be freely used in conjunction with hardware modifications targeted at reducing the same limitations identified, such as interleaving the cycles of two clocks to reduce dead time [10, 19].

The remainder of this manuscript is organized as follows. In Section I we provide an analytic description of the deleterious effects of LO noise on frequency standards, introducing the relevant metrics for performance of interest. This includes presentation of novel analytic expressions explicitly capturing the effects of feedback stabilization on the aggregate system performance through a recursive formulation. Section II demonstrates how to convert these time-domain statistical measures of frequency-standard performance to the Fourier domain, introducing both transfer functions for individual measurements and the pair-covariance transfer function capturing the correlations between arbitrary-duration Ramsey measurements conducted at arbitrary times. We then exploit these tools in Sec. III in order to devise a new hybrid-feedforward correction scheme similar in spirit to concepts from optimal estimation in order to maximize the accuracy of corrections applied to the LO. We demonstrate improvements in correction accuracy and LLO stability via this approach using numerical simulations with realistic LO noise power spectra. Finally, we conclude with a summary and discussion in Sec. IV.

I. THE EFFECT OF LOCAL OSCILLATOR NOISE ON FREQUENCY STANDARD STABILITY

Our primary objective is to suppress the impact of LO frequency noise on the ultimate performance of the *locked* LO, which is stabilized to a reference (e.g. an atomic transition). Accordingly, throughout this analysis we do not consider systematic shifts or uncertainties in the reference and explicitly assume that the reference is per-




FIG. 1. Effect of LO noise on the performance of a locked oscillator. Simulated evolution for a noisy LO, unlocked (black) and locked with traditional feedback (red). The dotted horizontal bars indicate the measurement outcomes (*samples*) over each cycle, \bar{y}_k , which are applied as correction at the end of the cycle, indicated by the bent arrow in the first cycle. Measurement period of duration T_R (white background) is followed by dead time with duration T_D (grey background). Total cycle time $T_c = T_D + T_R$, and here we represent a 50% duty factor, d . Undetected evolution of the LO during the dead time leads corrections to incompletely cancel frequency offsets at the time of correction. The arrows on the far right schematically indicate how locking reduces the variance of $y(t)$ though it does not eliminate it.

fect. This limit provides a reasonable approximation to the performance of many deployable frequency standards where LO stability is far worse than that of the associated atomic reference.

A. Time-domain description of Ramsey measurements and feedback stabilization

We represent the fractional frequency offset of the LO relative to an ontologically perfect reference at time t , $y(t) \equiv (\nu(t) - \nu_0)/\nu_0$, where ν_0 is the reference frequency and $\nu(t)$ is the LO frequency. In such a setting, Ramsey spectroscopy provides a means to determine the average value of $y(t)$ over a period T_R . Point-like realisations of the stochastic process $y(t)$ cannot be obtained experimentally; instead, the LO frequency error produces integrated *samples*, denoted \bar{y}_k and indexed in time by k :

$$\bar{y}_k \equiv \frac{1}{T_R^{(k)}} \int_{t_k^s}^{t_k^e} y(t) g(t - t_k^s) dt \quad (1)$$

where $T_R^{(k)} \equiv t_k^e - t_k^s$, $[t_k^s, t_k^e]$ is the time interval over which the k th sample is taken (“s” and “e” indicating “start” and “end” respectively), and $g(t)$ is a *sensitivity function* capturing the extent to which LO fluctuations at some instant t contribute to the measured outcome for that sample [26]. The range of $g(t)$ is $[0, 1]$ and its domain over the k interval is $t \in [t_k^s, t_k^e]$. The ideal case is the rectangular window case, where

$$g(t) = \begin{cases} 1 & \text{for } t \in [t_k^s, t_k^e] \\ 0 & \text{otherwise} \end{cases} \quad (2)$$

in which case \bar{y}_k reduces to the time-average of $y(t)$ over the interval $[t_k^s, t_k^e]$.

In traditional feedback stabilization, the samples, \bar{y}_k , are used to determine corrections to be applied to the LO in order to reduce frequency differences from the reference (Fig. 1). Consider the trajectory of the same frequency noise realisation $y(t)$ in the cases of no correction, denoted $y^{LO}(t)$ and correction, $y^{LLO}(t)$. The relation between these two cases of $y(t)$ is

$$y^{LLO}(t) = y^{LO}(t) + \sum_{k=1}^n C_k \quad (3)$$

where C_k refers to the value of the k th frequency correction applied to the LO, n of which have occurred before time t .

Under traditional feedback stabilization, each correction is directly proportional to the immediately preceding measurement outcome: $C_k = w_k \bar{y}_k^{LLO}$, where w_k is correction gain. Since \bar{y}_k^{LLO} is calculated by convolving $y^{LLO}(t)$ with a sensitivity function pertaining to the measurement parameters, (3) is a recursive equation in general. It is possible to cancel all but one of the recursive terms by setting the correction gain equal to the inverse of the average sensitivity $\bar{g}_k \equiv \int_0^{T_R^{(k)}} g(t)/T_R^{(k)} dt$ of the preceding measurement, i.e. $w_k = -\bar{g}_k^{-1}$, where the minus sign indicates negative feedback. With this constraint we can write

$$\bar{y}_k^{LLO} = \bar{y}_k^{LO} - \frac{\bar{g}_k}{\bar{g}_{k-1}} \bar{y}_{k-1}^{LO} \quad (4)$$

and for a Ramsey interrogation and measurement with negligibly short pulses, $\bar{g}_k = 1$. Applying feedback corrections sequentially after each measurement is able to effectively reduce the fractional frequency offset of the locked oscillator, $y^{LLO}(t)$, over many cycles, thus improving long-term stability.

We may now consider the limitations of this general approach. In the limit of a static frequency offset in time, a single (perfect) correction will set the frequency offset error of the LLO to zero; however, such perfect correction is in general not achieved. The primary reason for this in the limit of perfect measurements and corrections is *dynamic evolution* of the LO on timescales rapid compared to the measurements which cannot be fully compensated by the feedback loop possessing cycle time T_c .

In Fig. 1 we demonstrate how evolution of the LO frequency during T_R leads the feedback protocol to incompletely correct the offset $y(t)$. From the formalism presented above we see that incomplete feedback arises because the corrections are based only on the *average*

value of the frequency offset as measured over the k th period, \bar{y}_k (horizontal solid lines in Fig. 1), rather than the instantaneous value of the LO frequency offset at the time of correction *which cannot be known*. The difference between these two values leads to incomplete compensation of time-varying frequency offsets, and hence residual fractional instability in the quantity $y^{(LLO)}(t)$. In fact, LO evolution due to noise components evolving rapidly on the timescale of T_c contribute to effective aliasing of the noise in the measurement-feedback routine, and ultimately giving rise to the Dick effect. The impact of these effects on the ultimate stability of the LLO is exacerbated in circumstances where there is nonzero *dead time*, T_D , during which the LO may evolve, but this evolution is not captured by a measurement. Dead time arises due to e.g. the need to reinitialize the reference between measurements, or perform classical processing of the measurement outcome before a correction can be applied.

The net impact of this uncompensated evolution is a reduction in the long-term stability of the locked local oscillator. We now move on to describe the relevant quantitative metrics for LLO *variance* in both free-running and feedback-locked settings.

B. Measures of frequency standard stability for unlocked and locked LOs

The performance of the frequency standard is statistically characterized by various time-domain measures capturing the evolution of LO frequency as a function of time.

The variance of \bar{y}_k , denoted $\sigma_y^2(k)$ and often called *true variance* [26] is,

$$\sigma_y^2(k) = (\langle \bar{y}_k^2 \rangle - \langle \bar{y}_k \rangle^2) \rightarrow \text{E}[\bar{y}_k^2] \quad (5)$$

$$= \text{E} \left[\left(\frac{1}{T_R^{(k)}} \int_{t_k^s}^{t_k^e} y(t) g(t - t_k^s) dt \right)^2 \right] \quad (6)$$

where in the first line we assume that the true variance is simply equal to the expected value, E, of \bar{y}_k^2 , since $y(t)$ is assumed to be a zero-mean process. The true variance captures the spread of measurement outcomes due to different noise realizations in an infinite ensemble, over a single timestep. However, in a measurement context one does not have immediate access to an infinite ensemble of noise realizations, but rather a single time-series of measurement outcomes accumulated sequentially over a single noise realization. As a result we rely on a measure more conducive to this setting, the *sample variance* for N sequential finite-duration measurements $\{\bar{y}_k\}$ [26]

$$\sigma_y^2[N] = \frac{1}{N-1} \sum_{k=1}^N (\bar{y}_k - \frac{1}{N} \sum_{l=1}^N \bar{y}_l)^2. \quad (7)$$

In this work we will rely on such measures of frequency stability, rather than the more commonly employed Allan variance, in line with recent experiments [27]. Our decision to avoid the Allan variance is deliberate, as its form – effectively a moving average – specifically *masks* the effect of LO noise components with long correlation times. In fact the Allan variance is employed by the community in part because it does not diverge at long integration times τ due to LO drifts, as would the sample or true variance [26, 28–30]. In the limit where the stability of a frequency reference is dominated by LO noise (and the reference can be treated as perfect) this approach gives physically meaningful results. For completeness we introduce the formal definition of the Allan variance, calculated by finding the variance of the difference between consecutive pairs of measurement outcomes:

$$^A\sigma_y^2(y) = \frac{1}{2} \langle (\bar{y}_{k+1} - \bar{y}_k)^2 \rangle \quad (8)$$

where \bar{y}_k is the k th measurement outcome and $\langle \dots \rangle$ may indicate a time average or an ensemble average, depending on whether $y(t)$ is assumed to be ergodic.

The standard measures for oscillator performance consider either a free-running LO or provide a means only to statistically characterize measurement outcomes under black-box conditions. We may derive explicit analytic forms for different measurements of variance in the presence of feedback locking in order to provide insights into opportunities to improve net LLO performance through modification of the stabilization protocol.

We write time-domain expressions for variance using the relevant definitions provided above and the link between corrections in feedback and the history of the LLO's evolution. For the true variance we substitute Eq. 4 to find

$$\sigma_{yLLO}^2(k) = \text{Var}[\bar{y}_k^{LLO}] = \sigma_{yLO}^2(k) + \left(\frac{\bar{g}_k}{\bar{g}_{k-1}} \right)^2 \sigma_{yLO}^2(k-1) - \frac{2\bar{g}_k}{\bar{g}_{k-1}} \sigma(\bar{y}_{k-1}^{LO}, \bar{y}_k^{LO}) \quad (9)$$

and calculate the expected value of the LLO sample variance in a similar manner using Eq. 3

$$\begin{aligned} \mathbb{E}[\sigma_{yLLO}^2[N]] = & \frac{1}{N-1} \sum_{k=1}^N \left\{ \left(\sigma_{yLO}^2(k) + \bar{g}_k^2 \sum_{r=1}^{k-1} \sum_{s=1}^{k-1} \sigma(C_r, C_s) - 2\bar{g}_k \sum_{u=1}^{k-1} \sigma(\bar{y}_k^{LO}, C_u) \right) \right. \\ & \left. + \frac{1}{N^2} \sum_{p=1}^N \sum_{q=1}^N \left(\sigma(\bar{y}_p^{LO}, \bar{y}_q^{LO}) + \bar{g}_p \bar{g}_q \sum_{w=1}^{p-1} \sum_{x=1}^{q-1} \sigma(C_x, C_y) \right) - \frac{2}{N} \sum_{l=1}^N \left(\sigma(\bar{y}_k^{LO}, \bar{y}_l^{LO}) + \bar{g}_k \bar{g}_l \sum_{y=1}^{k-1} \sum_{z=1}^{l-1} \sigma(C_y, C_z) \right) \right\} \quad (10) \end{aligned}$$

We see that the characteristics of the locked LO can be expressed in terms of the unlocked LO and the *co-variance* between two quantities, $\sigma(x, y)$, capturing correlations between them. This may include the covariance of different measurement outcomes on the LO (e.g. $\sigma(\bar{y}_k^{LO}, \bar{y}_l^{LO})$), or different corrections applied to the LO (e.g. $\sigma(C_r, C_s)$). It is this observation – that we may express relevant statistical quantities surrounding the performance of locked local oscillators in terms of measurement covariances – that will provide a path towards the development of new stabilization routines exploiting temporal correlations in the LO noise (and hence measurement outcomes).

II. PERFORMANCE MEASURES FOR FREQUENCY STANDARDS IN THE FOURIER DOMAIN

We require an efficient theoretical framework in which to capture these effects, and hence transition to the frequency domain, making use of the power spectral density of the LO's fractional frequency error, $S_y(\omega)$, in order to characterize average performance over a hypothetical statistical ensemble. In this description residual LLO instability persists because the feedback is insensitive to LO noise at high frequencies relative to the inverse measurement time. Additional instability due to the Dick effect comes from aliasing of noise at harmonics of, T_c^{-1} , the effective loop bandwidth.

We may analytically calculate the effects of measurement, dead time, and the feedback protocol itself on frequency standard performance in the frequency domain as follows. Defining a normalised, time-reversed sensitivity

function $\bar{g}(t_k^m - t) = g(t - t_k^s)/T_R^{(k)}$, where $g(t)$ is assumed to be time-reversal symmetric about t_k^m , the midpoint of the k th interval $[t_k^s, t_k^e]$, we can express the true variance as a convolution $\sigma_y^2(k) = \mathbb{E} \left[\left(\int_{-\infty}^{\infty} y(t) \bar{g}(t_k^m - t) dt \right)^2 \right]$. Expanding this expression gives

$$\sigma_y^2(k) = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \mathbb{E}[y(t)y(t')] \bar{g}(t_k^m - t) \bar{g}(t_k^m - t') dt' dt \quad (11)$$

$$= \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} R_y^{TS}(\Delta t) \bar{g}(t_k^m - t) \bar{g}(t_k^m - t') dt' dt \quad (12)$$

where $R_y^{TS}(\Delta t)$ is the two-sided (TS) autocorrelation function and $\Delta t \equiv t' - t$. Using the Wiener-Khinchin theorem we write $R_y^{TS}(\Delta t) = \mathcal{F}^{-1}\{S_y^{TS}(\omega)\}$, relating the autocorrelation function to the Fourier transform of the power spectral density of the LO noise. Defining the Fourier transform of $\bar{g}(t_k^m - t)$:

$$G_k(\omega) \equiv \int_{-\infty}^{\infty} \bar{g}(t_k^m - t) e^{i\omega t} dt \quad (13)$$

Here $|G_k(\omega)|^2$ is called the *transfer function* for the k th sample, describing the spectral properties of the measurement protocol itself. For measurements performed using Ramsey interrogation with $\pi/2$ pulses of negligible duration and zero dead time, the transfer function has a sinc-squared analytic form $|G_k(\omega)|^2 = (\sin(\omega T_R^{(k)}/2)/(\omega T_R^{(k)}/2))^2$.

We may then express the true variance

$$\sigma_y^2(k) = \frac{1}{2\pi} \int_0^{\infty} S_y(\omega) |G_k(\omega)|^2 d\omega \quad (14)$$

where the substitution of the one-sided PSD $S_y(\omega)$ is possible because $|G_k(\omega)|^2$ is even. This result is similar to the convolution theorem, which states that $\mathcal{F}\{f \star g\} = \mathcal{F}\{f\} \cdot \mathcal{F}\{g\}$, where \star denotes a convolution and f and g are Fourier-invertible functions. This framework, which expresses the true variance as the overlap integral of the noise power spectrum and the transfer function for the measurement protocol has recently seen broad adoption in the quantum information community where time-varying dephasing noise is a major concern for the stability of quantum bits [21–24, 31–34].

Recalling that statistical measures of LLO variance rely not only on expressions for the true variance over noise ensembles, but also of covariances between measurements or corrections, we must equivalently express the covariance in terms of transfer functions. Using the identity $\sigma^2(A \pm B) = \sigma^2(A) + \sigma^2(B) \pm 2\sigma(A, B)$, we define a sum and a difference sensitivity function: $g_{k,l}^+(t)$ and $g_{k,l}^-(t)$, with respect to two measurements indexed k

and l . These expressions are general functions of time with two regions of high sensitivity corresponding to the individual measurement periods.

$$g_{k,l}^{\pm}(t) \equiv \begin{cases} g(t - t_k^s), & \text{for } t \in [t_k^s, t_k^e] \\ \pm g(t - t_l^s), & \text{for } t \in [t_l^s, t_l^e] \\ 0, & \text{otherwise} \end{cases} \quad (15)$$

These time-domain sum and difference sensitivity functions have their corresponding frequency-domain transfer functions, defined as their Fourier transforms normalised by $T_R^{(k,l)}$:

$$G_{k,l}^{\pm}(\omega) \equiv \int_{-\infty}^{\infty} \left(\frac{g(t_k^m - t)}{T_R^{(k)}} \pm \frac{g(t_l^m - t)}{T_R^{(l)}} \right) e^{i\omega t} dt \quad (16)$$

Substituting this and the form of the true variance (14) into the variance identity above and rearranging terms gives the covariance of the two measurement outcomes

$$\sigma(\bar{y}_k, \bar{y}_l) = \frac{1}{2\pi} \int_0^{\infty} \frac{S_y(\omega)}{4} \left(|G_{k,l}^+(\omega)|^2 - |G_{k,l}^-(\omega)|^2 \right) d\omega \quad (17)$$

$$\equiv \frac{1}{2\pi} \int_0^{\infty} S_y(\omega) G_{k,l}^2(\omega) d\omega \quad (18)$$

whereby $G_{k,l}^2(\omega)$ is defined to be the *pair covariance transfer function*. For the case of flat-top Ramsey measurements over the intervals $[t_{k,l}^s, t_{k,l}^e]$ this term takes the form

$$G_{k,l}^2(\omega) = (\omega^2 T_R^{(k)} T_R^{(l)})^{-1} \left[\cos(\omega(t_l^s - t_k^s)) + \cos(\omega(t_l^e - t_k^e)) - \cos(\omega(t_l^e - t_k^s)) - \cos(\omega(t_l^s - t_k^e)) \right]. \quad (19)$$

This is a generalization of the transfer function previously derived for the special case of periodic, equal-duration Ramsey interrogations [26, 29], and allows effective estimation of $y(t)$ for any t and for any set of measured samples $\bar{\mathbf{y}}_k$. Now, with the expressions above we may calculate all relevant terms in Equations 9 and 9 in the Fourier domain.

We thus see that this approach allows expression of time-domain LO variances as overlap integrals between $S_y(\omega)$ and the transfer functions capturing the effects of the measurement and feedback protocol, including correlations between measurements or corrections in time. Through this formalism we may incorporate arbitrary measurement protocols (e.g. arbitrary and dynamic Ramsey periods and dead times): the underlying physics of e.g. changing linewidth of the measurement is explicitly captured through the form and implicit time-dependence of the transfer function used to characterize the measurement protocol.

III. EXPLOITING NOISE CORRELATIONS TO IMPROVE FEEDBACK STABILIZATION

Recasting variance metrics for the stability of LOs in terms of transfer functions is particularly powerful because it provides a path to craft new measurement feedback protocols designed to reduce residual variance measures for the LLO by modifying the protocol's spectral response. Our key insight is that the non-Markovianity of dominant noise processes in typical LOs – captured through the low-frequency bias in $S_y(\omega)$ [26, 29] – implies the presence of temporal correlations in $y(t)$ that may be exploited to improve feedback stabilization. These correlations are captured in the set of n past measurement outcomes $\bar{\mathbf{y}}_k = \{\bar{y}_{k,1}, \dots, \bar{y}_{k,n}\}$; accordingly *future evolution* of $y(t)$ may be predicted based on a past subset of measurements within $\bar{\mathbf{y}}_k$, so long as the past measurements and point of prediction fall within the characteristic *correlation time* for the LO noise given by $S_y(\omega)$. This approach provides a direct means to account for LO evolution that is normally not compensated during *dead time* in the measurement process and is missed by the averaging process over T_R .

A. Optimal estimator for corrections

The formal basis of our analytic approach, in summary, is to calculate a covariance matrix in the frequency domain via transfer functions to capture the relative correlations between sequential measurement outcomes of an LLO, and use this matrix to derive a linear *predictor* of the LLO frequency offset at the moment of correction. Under appropriate conditions this predictor provides a correction with higher accuracy than that derived from a single measurement, allowing us to improve the ultimate performance of the LLO. Since the predictor is found using information from previous measurements (feedback) and a priori statistical knowledge of the LO noise to *predict* the evolution of the LO (feedforward), we call the scheme *hybrid feedforward*.

This approach shares common objectives with application of optimal control techniques such as Kalman filtering [25] in the production of composite frequency standards from an ensemble of physical clocks [35], or in compensating for deterministic frequency shifts due to e.g. aging or changes in the ambient temperature of a clock [36, 37]. The primary advance of this work is the insight that *stochastic* evolution of the LO can be predicted and compensated using optimal control protocols *inside the feedback loop*.

In hybrid feedforward, results from a set of n past measurements are linearly combined with weighting coefficients \mathbf{c}_k optimized such that the k th correction, C_k , provides maximum correlation to $y(t_k^c)$ at the instant of correction t_k^c (Fig. 1c). Assuming that the LO noise is Gaussian, the optimal least minimum mean squares estimator (MMSE) is linear, and the optimal value of the

correction is given by $C_k = \mathbf{c}_k \cdot \bar{\mathbf{y}}_k$: the dot product of a set of correlation coefficients \mathbf{c}_k derived from knowledge of $S_y(\omega)$ and a set of n past measured samples, $\bar{\mathbf{y}}_k = \{\bar{y}_{k,1}, \dots, \bar{y}_{k,n}\}$. We define an $(n+1) \times (n+1)$ covariance matrix where the $(n+1)$ th term represents an ideal zero-duration sample at t_k^c and in the second line we write the covariance matrix in block form:

$$\Sigma_k \equiv \begin{bmatrix} \sigma(\bar{y}_{k,1}, \bar{y}_{k,1}) & \cdots & \sigma(\bar{y}_{k,1}, y(t_k^c)) \\ \sigma(\bar{y}_{k,2}, \bar{y}_{k,1}) & \cdots & \sigma(\bar{y}_{k,2}, y(t_k^c)) \\ \vdots & \ddots & \vdots \\ \sigma(y(t_k^c), \bar{y}_{k,1}) & \cdots & \sigma(y(t_k^c), y(t_k^c)) \end{bmatrix} \quad (20)$$

$$\equiv \begin{bmatrix} \mathbf{M}_k & \mathbf{F}_k \\ \mathbf{F}_k^T & \sigma(y(t_k^c), y(t_k^c)) \end{bmatrix}. \quad (21)$$

In this form the matrix \mathbf{M}_k describes correlations between measurement outcomes while the vector \mathbf{F}_k describes correlations between each measurement and the LLO at the time of correction. The MMSE optimality condition is then fulfilled for

$$\mathbf{c}_k = \frac{\mathbf{F}_k}{\sqrt{\mathbf{F}_k^T \mathbf{M}_k \mathbf{F}_k}} \frac{w_k}{2\pi} \int_0^\infty S_y(\omega) d\omega \quad (22)$$

where w_k is an overall correction gain. The covariance matrix elements are calculated as defined above in terms of the LO noise power spectrum.

In the practical setting of a frequency standard experiment, we wish to improve two metrics simultaneously: the accuracy of each correction, and the long-term stability of the LLO output. The former is achieved by maximising the correlation between C_k and $y(t_k^c)$, while the latter is captured by the metrics of frequency variance, sample variance, and Allan variance at long averaging times.

Consideration of a single cycle of hybrid feedforward correction calculated over a time-sequence \mathbf{y}_k can provide a value for the ensemble averaged $\langle y^{LLO}(t_k^c)^2 \rangle$ in terms of covariance matrix elements, despite the fact the LLO frequency variance under hybrid feedforward cannot be expressed in a closed non-recursive form for more than a single cycle. This in turn provides a metric for the *correction accuracy* for hybrid feedforward, defined as the inverse of the LLO frequency variance at time t_k^c normalized by that for the free-running LO. This captures the extent to which a correction brings $y^{LLO}(t) \rightarrow 0$ at the instant of correction, $t = t_k^c$

$$\begin{aligned} A_k &\equiv \frac{\langle y^{LO}(t_k^c)^2 \rangle}{\langle y^{LLO}(t_k^c)^2 \rangle} \\ &= \left(1 + w_k^2 - w_k \frac{|\mathbf{F}_k|^2}{\sqrt{\mathbf{F}_k^T \mathbf{M}_k \mathbf{F}_k}} \right)^{-1} \end{aligned} \quad (23)$$

We can gain insights into the performance of the correction protocol by considering limiting cases. In the

limit of white noise with negligible correlations, $\mathbf{M}_k \rightarrow \mathbb{I}$, the identity matrix possessing only diagonal elements. In this limit the rightmost term in Eq. 23 reduces to $w_k |\mathbf{F}_k|$, which is small (there are negligible correlations between measurement outcomes and $y(t_k^c)$). Here accuracy $A_k \rightarrow 1/(1 + w_k^2)$, and is maximized by setting $w_k = 0$; this is interpreted as performing no feedback at all, as measurements (and hence corrections) are uncorrelated with $y(t_k^c)$. By contrast with perfect correlations all elements of the covariance matrix takes value unity. Standard feedback works perfectly by selecting unity gain and selecting the number of measurements to be combined, $n = 1$, to correct based on a single measurement.

In intermediate regimes induced by LLO noise with finite correlation times (e.g. colored power spectra $S_y(\omega)$) the ensemble-averaged accuracy of the hybrid feedforward correction is maximized (for $w_k > 0$) by minimizing the term $\sqrt{\mathbf{F}_k^T \mathbf{M}_k \mathbf{F}_k}$. We can interpret the effect of \mathbf{M}_k as an effective rotation matrix, reducing the magnitude of this expression by effectively maximizing the “angle” between $\mathbf{M}_k \mathbf{F}_k$ and \mathbf{F}_k . While it is unphysical to reduce this to zero based on the limiting cases discussed above, it is possible to appropriately select k , based on characteristics of $S_y(\omega)$, in order to improve correction accuracy.

Beyond the accuracy of single corrections, in all slaved frequency standards we rely on repeated measurements and corrections to provide long-term *stability*, a measure of how the output frequency of the LLO deviates from its mean value over time. We study this by calculating the sample variance of a time-sequence of measurement outcomes averaged over an ensemble of noise realizations, $\langle \sigma_y^2[N] \rangle$. A “moving average” style of hybrid feedfor-




FIG. 2. Schematic diagram of hybrid feedforward with an example protocol using $n = 3$. Start and end times of measurements are defined arbitrarily permitting non-uniform-duration measurements, although measurements are illustrated as uniform for clarity. Corrections $C_k^{(n=3)}$ are applied in either non-overlapping blocks of three measurements or as a moving average (depicted here). In the latter case, the covariance matrix must be recalculated to correctly account for any variations in measurement duration. Dashed red arrows indicate the first corrections performed without full calculation of the covariance matrix. This effect vanishes for $k > n$.

ward provides improved long-term stability, as the correction C_k will depend on the set of measurement outcomes $\bar{\mathbf{y}}_k = \{\bar{y}_{k-n+1}, \dots, \bar{y}_k\}$, among which previous corrections have been interleaved, as illustrated in Fig. 2. In this case the covariance matrix must be updated to reflect the action of each correction. See Appendix for a detailed form of the Sample Variance in the case of this form of stabilization.

B. Numerical Simulations

In order to test the general performance of hybrid feedforward in different regimes we perform numerical simulations of noisy LOs with user-defined statistical properties, characterized by $S_y(\omega)$. We produce a fixed number of LO realizations in the time domain [38] and then use these to calculate measures such as the sample variance over a sequence of simulated measurement outcomes with user-defined Ramsey measurement times, dead times, and the like. In these calculations we may assume that the LO is free running, experiencing standard feedback, or employing hybrid feedforward, and then take an ensemble average over LO noise realizations in order to produce the expectation value of the relevant variance metric. Our calculations include various noise power spectra, with tunable high-frequency cutoffs, including common ‘flicker frequency’ ($S_y(\omega) \propto 1/\omega$), and ‘random walk frequency’ ($S_y(\omega) \propto 1/\omega^2$) noise, as appropriate for experiments incorporating realistic LOs.

We begin by exploring the value of varying the Ramsey periods, $T_R^{(k)}$ in an n measurement sequence in order to vary sampling of the LO frequency noise power spectrum $S_y(\omega)$. As an example, we fix $n = 2$ and permit $T_R^{(1)}$ and $T_R^{(2)}$ to be varied as optimization parameters in order to maximize correction accuracy under hybrid-feedforward correction. We employ a multidimensional optimization in the form of a Nelder-Mead simplex over the measurement durations, finding that a measurement protocol consisting of a long measurement period followed by a short period maximizes correction accuracy (inset Fig. 3). Intuitively this structure ensures that the measurement routine samples both high and low-frequency regimes of $S_y(\omega)$. The ordering of a long followed by a short Ramsey period ensures minimum delay (and hence maximum correlation) between the most recent measurement sampling the high-frequency regime and the noise at time of correction, $y(t_k^c)$.

With $S_y(\omega) \propto 1/\omega$ and $S_y(\omega) \propto 1/\omega^2$ we observe increased correction accuracy under hybrid feedforward relative to traditional feedback in Fig. 3. For a ratio of $T_R^{(1)}/T_R^{(2)} \rightarrow 100$ we see improvements in accuracy up to 50% through use of our protocol. This indicates clearly that the use of statistical information about the noise can be used to improve feedback corrections via incorporation of prediction of the LO’s evolution. In the presence of rapid fluctuations in $y(t)$ arising from a broadband




FIG. 3. Calculated correction accuracy of the first correction for hybrid feedforward normalized to feedback (accuracy = 1), under different forms of $S_y(\omega)$ as a function of the ratio of Ramsey periods between the two measurements employed in constructing $C_k^{(2)}$. Correction accuracy for feedback is calculated assuming the minimum Ramsey time; thus for the ratio of Ramsey measurements taking value unity on the x -axis, the hybrid feedforward scheme takes twice as long as feedback. Inset: depiction of the form of $C_k^{(2)}$ used in hybrid feedforward, depicting the “slower” measurement being performed first.

white power spectrum, the benefits of hybrid feedforward are mitigated, as expected. Under such circumstances the LO noise evolution is delta-correlated in time, meaning there is no temporal correlation to exploit via hybrid feedforward. In the parameter ranges we have studied numerically we find that correction accuracy is maximized for $n = 2 - 3$, with diminishing performance for larger n . Again, this is determined by the relevant correlation time of the LO noise.

As described above we also consider the improvement in long-term stability achievable for a locked LO under different feedback stabilization protocols. We calculate $\langle \sigma_y^2[N] \rangle$ up to $N = 100$ measurements, calculated using feedback and hybrid feedforward with $n = 2$ assuming the simple case of uniform T_R and overlapping corrections (Fig. 4a-b). We present the resulting *normalized improvement* in $\langle \sigma_y^2[N] \rangle$ relative to the LLO under standard feedback in Fig. 4c, observing clear improvement (reduction) in $\langle \sigma_y^2[N] \rangle$ through the hybrid feedforward approach. Benefits vary with the details of the selected noise power spectrum but vary 5 – 25% of $\langle \sigma_y^2[N] \rangle$ relative to standard measurement feedback. We present data for different functional forms of $S_y(\omega)$, including low-frequency dominated flicker noise ($\propto 1/\omega$), and power spectra with more significant noise near T_c^{-1} ($\propto 1/\omega^{1/2}$). The benefits of our approach appear most significant in the long term when high-frequency noise reduces the efficacy of standard feedback. Notably, because of well known relationships between LO *phase noise* and LO *frequency noise* [28], significant high-frequency weight in $S_y(\omega)$ is commonly encountered.

As another example we explicitly explore how HFF




FIG. 4. (a, b) Calculated sample variance for an unlocked LO, feedback, and hybrid feedforward, as a function of measurement number N , for different power spectra (indicated on graphs). Calculations assume $S_y(\omega) \propto 1/\omega$, with a high-frequency cutoff $\omega_c/2\pi = 100/T_c$ and $S_y(\omega) \propto 1/\omega^{1/2}$ with a cutoff frequency $\omega_c/2\pi = 1/T_c$, demonstrating the importance of high-f noise near $\omega/2\pi = T_c^{-1}$. PSDs with different ω -dependences are normalised to have the same value at $\omega_{low} = 1/100T_c$. (c) Normalized sample variance data from panels (a) and (b) presented as the ratio of $\langle \sigma_y^2[N] \rangle^{(HFF)} / \langle \sigma_y^2[N] \rangle^{(FB)}$ in order to demonstrate improvement due to hybrid feedforward (numbers less than unity indicate smaller sample variance under hybrid feedforward). (d) Calculated $\langle \sigma_y^2[N] \rangle$ for $N = 20$ as a function of duty factor, normalized to the sample variance for the free-running LO. Data above red dashed line indicate that the standard feedback approach produces instability *larger* than that for the free-running oscillator. Both data sets assume $S_y(\omega) \propto 1/\omega$, with $\omega_c/2\pi = 100/T_c$. Crosses represent data with ten noise spurs superimposed on $S_y(\omega)$, starting at $\omega/2\pi = 1.15T_c^{-1}$, and increasing linearly with step size $0.15T_c^{-1}$.

can mitigate Dick-effect-related aliasing by performing calculations using noise with strong contributions near T_c^{-1} . In Fig. 4c we calculate the expectation value of the sample variance at a fixed value of $N = 20$ for a LLO stabilized using either traditional feedback or hybrid feedforward. The sample variances are normalized by that for the free-running LO, meaning that values of this metric less than unity demonstrate improvement due to stabilization, and smaller values indicate better stabilization. On the horizontal axis we vary the duty factor d , defined as the ratio of the interrogation time to total cycle time: $d \equiv T_R/T_c$ from 1% to unity (no dead time).

Knowledge of correlations in the noise allows hybrid feedforward to provide metrologically significant gains in stability relative to traditional feedback. These differences arise because even though the noise processes are random, knowledge of the statistical properties of the noise provides a means to effectively model the average dynamical evolution of the system, and accurately predict how the system will evolve in the future. Exact performance depends sensitively on the form and magnitude of $S_y(\omega)$, but results demonstrate that systems with high-frequency noise content around $\omega/2\pi \approx T_c^{-1}$ benefit significantly from hybrid feedforward.

Results presented in Fig. 4d indicate an improvement provided by hybrid feedforward that is most marked for low duty factor d where the Dick effect is known to be most severe, although overall sample variance grows for both stabilization protocols in this limit. As $d \rightarrow 1$ the performance of traditional feedback and hybrid feedforward converge, as standard feedback corrections become most effective when dead time is shortest. We even find that in certain regimes where noise near T_c^{-1} is so strong that standard feedback makes long-term stability *worse* than applying no feedback at all, hybrid feedforward can provide useful stabilization.

For these calculations we have selected noise power spectra that, while conservative, are inspired by typical LO phase noise specifications weighted to enhanced high-frequency content due to the conversion between phase and frequency instability [28]. We specifically compare $S_y(\omega) \propto 1/\omega$ and $S_y(\omega) \propto 1/\omega$ with added noise spurs near the inverse cycle time, T_c^{-1} . The latter condition mimics LO noise in real experimental situations with tight-SWAP clocks such as the POP atomic clock [39]. The POP clock has a duty cycle $d \approx 0.7$, and a dominant LO noise term $S_y(\omega) \propto 1/\omega$ with sharp spurs near the inverse cycle time, matching the conditions treated in Fig. 4. From these data our results suggest a potential improvement in sample variance $\sim 5 - 10\%$ using this extremely simple (and un-optimized) protocol. We note that the noise strength treated in this figure is very strong; we have generally observed that relative gains are determined by the form of $S_y(\omega)$ rather than its absolute magnitude.

IV. CONCLUSION

In summary, we have presented a set of analytical tools describing LLO performance in the frequency domain for arbitrary measurement times, durations, and duty cycles. We have employed this formalism, based on generalized transfer functions, to develop a new software-only approach to LO feedback stabilization in slaved passive frequency standards, bringing optimal estimation techniques inside the feedback loop. The techniques we propose incorporate a series of past measurements and statistical knowledge of the noise to improve the accuracy of feedback corrections and ultimately improve the stabil-

ity of the slaved LO. We have validated these theoretical insights using numerical simulations of noisy local oscillators and calculations of relevant stability metrics.

The results we have presented represent only the first steps which we believe may be possible focused on the integration of optimal estimation techniques to improve clock stabilization without the need for hardware modification. For instance we have numerically demonstrated improved correction accuracy using nonuniform-duration T_R over a cycle, as well as long-term stability improvement using only the simplest case of uniform T_R . These approaches may be combined to produce LLOs with improved accuracy relative to the reference at the time of correction and improved long-term stability. In cases where the penalty associated with increasing T_R is modest (captured in the statistical properties of the noise), such composite schemes can provide substantial benefits as well, improving both accuracy of correction to the LLO and overall frequency standard stability.

Other expansions may leverage the basic analytic formalism we have introduced; we have introduced the transfer functions, $|G(\omega)|^2$ and $G_{k,l}^2(\omega)$, but have assumed only the simplest form for the time-domain sensitivity function and fixed overall gain. However, it is possible to craft a measurement protocol to yield $|G(\omega)|^2$ that suppresses the dominant spectral features of the LO noise. We have observed that through such an approach one may reduce the impact of aliasing on clock stabilization, indicating a path for future work on reducing of the so-called Dick limit in precision frequency references. Finally, recent experiments from our group [40] have demonstrated that similar concepts in optimal estimation and predictive correction may be integrated using machine learning techniques to gain information about $S_y(\omega)$ on the fly, and experimentally demonstrating that similar hybrid feedforward stabilization can yield experimental gains in long-term LLO stability. It is exciting to explore combinations of analytic and numerical techniques to find tradeoffs in performance vs computational efficiency.

In the parameter regimes we have studied the relative performance benefits of the hybrid feedforward approach are of metrological significance - especially considering they may be gained using only “software” modification without the need for wholesale changes to the clock hardware or measurement procedure. We believe the approach may find special significance in tight-SWAP (size, weight, and power) applications such as space-based clocks where significantly augmenting LO quality is generally impossible due to system-level limitations. Future work will focus on detailed studies of measurement routines optimized for specific clocks such as HORACE. Overall, we believe that this work indicates clear potential to improve passive frequency standards by incorporation of optimal estimation techniques in the feedback loop itself. *Note:* While preparing this manuscript we became aware of related work seeking to employ covariance techniques to improve measurements of quantum

clocks [41].

Acknowledgements: The authors thank H. Ball, D. Hayes, and J. Bergquist for useful discussions. This work partially supported by Australian Research Council Discovery Project DP130103823, US Army Research Office under Contract Number W911NF-11-1-0068, and the Lockheed Martin Corporation.

APPENDIX

Variances for locked local oscillators with hybrid feedforward

The standard measures for oscillator performance consider either a free-running LO or provide a means only

to statistically characterize measurement outcomes under black-box conditions. Here we present explicit analytic forms for different measurements of variance in the presence of feedback locking.

The expected value of the LLO sample variance can be found by substituting (3) into the definition of the sample variance, producing a generic expression for traditional feedback (one measurement per correction cycle) and hybrid feedforward (multiple measurements per cycle):

$$\begin{aligned}
 E[\sigma_{y_{LLO}}^2[N]] &= \frac{1}{N-1} \sum_{k'=1}^N \left\{ \sigma_{y_{LLO}}^2(k') + \frac{1}{N^2} \sum_{p'=1}^N \sum_{q'=1}^N \sigma(\bar{y}_{p'}^{LLO}, \bar{y}_{q'}^{LLO}) - \frac{2}{N} \sum_{l'=1}^N \sigma(\bar{y}_{k'}^{LLO}, \bar{y}_{l'}^{LLO}) \right\} \\
 &= \frac{1}{N-1} \sum_{k'=1}^N \left\{ \left(\sigma_{y_{LLO}}^2(k') + \bar{g}_{k'}^2 \sum_{r=1}^{\lfloor k'/n \rfloor} \sum_{s=1}^{\lfloor k'/n \rfloor} \sigma(C_r, C_s) - 2\bar{g}_{k'} \sum_{u=1}^{\lfloor k'/n \rfloor} \sigma(\bar{y}_{k'}^{LO}, C_u) \right) \right. \\
 &\quad \left. + \frac{1}{N^2} \sum_{p'=1}^N \sum_{q'=1}^N \sigma(\bar{y}_{p'}^{LO} + \bar{g}_{p'} \sum_{p=1}^{\lfloor p'/n \rfloor} C_p, \bar{y}_{q'}^{LO} + \bar{g}_{q'} \sum_{q=1}^{\lfloor q'/n \rfloor} C_q) - \frac{2}{N} \sum_{l'=1}^N \sigma(\bar{y}_{k'}^{LO} + \bar{g}_{k'} \sum_{u=1}^{\lfloor k'/n \rfloor} C_u, \bar{y}_{l'}^{LO} + \bar{g}_{l'} \sum_{v=1}^{\lfloor l'/n \rfloor} C_v) \right\} \\
 &= \frac{1}{N-1} \sum_{k'=1}^N \left\{ \left(\sigma_{y_{LLO}}^2(k') + \bar{g}_{k'}^2 \sum_{r=1}^{\lfloor k'/n \rfloor} \sum_{s=1}^{\lfloor k'/n \rfloor} \sigma(C_r, C_s) - 2\bar{g}_{k'} \sum_{u=1}^{\lfloor k'/n \rfloor} \sigma(\bar{y}_{k'}^{LO}, C_u) \right) \right. \\
 &\quad \left. + \frac{1}{N^2} \sum_{p'=1}^N \sum_{q'=1}^N \left(\sigma(\bar{y}_{p'}^{LO}, \bar{y}_{q'}^{LO}) + \bar{g}_{p'} \bar{g}_{q'} \sum_{p=1}^{\lfloor p'/n \rfloor} \sum_{q=1}^{\lfloor q'/n \rfloor} \sigma(C_p, C_q) \right) \right. \\
 &\quad \left. - \frac{2}{N} \sum_{l'=1}^N \left(\sigma(\bar{y}_{k'}^{LO}, \bar{y}_{l'}^{LO}) + \bar{g}_{k'} \bar{g}_{l'} \sum_{k=1}^{\lfloor k'/n \rfloor} \sum_{l=1}^{\lfloor l'/n \rfloor} \sigma(C_k, C_l) \right) \right\}
 \end{aligned} \tag{24}$$

$$\tag{25}$$

$$\tag{26}$$

where in the case of hybrid feedback, N is defined to be total number of measurements and n is the number of measurements per cycle. The summation signs with unprimed indices are sums over whole cycles (of which there are $\lfloor N/n \rfloor$) and the primed indices are sums over all N measurements. In general, $E[\sigma_{y_{LLO}}^2[N]]$ contains recursive terms that cannot be concisely expressed in terms of the LO PSD $S_y(\omega)$ and covariance transfer function $G^2(\omega)$.

The Allan variance, the conventional measure of frequency standard instability, can be expressed analogously

$${}^A\sigma_y^2(y) = \frac{1}{2\pi} \int_0^\infty S_y(\omega) |{}^AG(\omega)|^2 d\omega \tag{27}$$

where the transfer function, for ideal Ramsey interrogation, is

$$|{}^AG(\omega)|^2 = \frac{2 \sin^4(\omega T_R/2)}{(\omega T_R/2)^2} \tag{28}$$

where T_R lacks an index because the definition of the Allan variance assumes equal-duration interrogation bins [26]. The Allan variance calculated via this frequency-domain approach can be compared to its value via the time-domain approach, which consists of finding the variance of the difference between consecutive pairs of measurement outcomes:

$$^A\sigma_y^2(y) = \frac{1}{2}\langle(\bar{y}_{k+1} - \bar{y}_k)^2\rangle \quad (29)$$

where \bar{y}_k is the k th measurement outcome and $\langle\cdots\rangle$ may

indicate a time average or an ensemble average, depending on whether $y(t)$ is assumed to be ergodic.

The LLO Allan variance can be found by substituting (4) into the definition of the Allan variance (29):

$$^A\sigma_{yLLO}^2(k) = \frac{1}{2}\mathbb{E}[(\bar{y}_{k+1}^{LLO} - \bar{y}_k^{LLO})^2] \quad (30)$$

$$= \frac{1}{2}\mathbb{E}\left[\left(\bar{y}_{k+1}^{LO} - \frac{\bar{g}_{k+1}}{\bar{g}_k}\bar{y}_k^{LO} - \bar{y}_k^{LO} + \frac{\bar{g}_k}{\bar{g}_{k-1}}\bar{y}_{k-1}^{LO}\right)^2\right] \quad (31)$$

$$= \frac{1}{2}\left(\sigma_{yLO}^2(k+1) + \left(1 + \frac{\bar{g}_{k+1}}{\bar{g}_k}\right)^2\sigma_{yLO}^2(k) + \left(\frac{\bar{g}_k}{\bar{g}_{k-1}}\right)^2\sigma_{yLO}^2(k-1) + \frac{2\bar{g}_k}{\bar{g}_{k-1}}\sigma(\bar{y}_{k+1}^{LO}, \bar{y}_{k-1}^{LO}) - 2\left(1 + \frac{\bar{g}_{k+1}}{\bar{g}_k}\right)\sigma(\bar{y}_k^{LO}, \bar{y}_{k+1}^{LO}) - \frac{2(\bar{g}_k + \bar{g}_{k+1})}{\bar{g}_{k-1}}\sigma(\bar{y}_k^{LO}, \bar{y}_{k-1}^{LO})\right) \quad (32)$$

-
- [1] H. Katori, *Nat. Photon.* **5**, 203 (2011).
 - [2] N. Nand, J. Hartnett, E. N. Ivanov, and G. Santarelli, *IEEE Trans. Microwave Thy and Tech.* **59** (2011).
 - [3] C. W. Chou, D. B. Hume, T. Rosenband, and D. J. Wineland, *Science* **329**, 1630 (2010).
 - [4] S. Blatt, A. D. Ludlow, G. K. Campbell, J. W. Thomsen, T. Zelevinsky, M. M. Boyd, J. Ye, X. Baillard, M. Fouché, R. Le Targat, A. Brusch, P. Lemonde, M. Takamoto, F.-L. Hong, H. Katori, and V. V. Flambaum, *Phys. Rev. Lett.* **100**, 140801 (2008).
 - [5] J. Guéna, M. Abgrall, D. Rovera, P. Laurent, B. Chupin, M. Lours, G. Santarelli, P. Rosenbusch, M. Tobar, R. Li, K. Gibble, A. Clairon, and S. Bize, *IEEE Trans. Ultrason., Ferroelec., and Freq. Control* **59** (2012).
 - [6] O. I. de la Convention du Mètre, “The international system of units (si),” (2006).
 - [7] D. B. Sullivan, J. C. Bergquist, J. J. Bollinger, R. E. Drullinger, W. M. Itano, S. R. Jefferts, W. D. Lee, D. Meekhof, T. E. Parker, F. L. Walls, and D. J. Wineland, *J. Res. Natl. Inst. Stand. Technol.* **106**, 47 (2001).
 - [8] C. Audoin and B. Guinot, *The Measurement of Time*, 1st ed. (Cambridge University Press, 2001).
 - [9] P. T. H. Fisk, *Rep. Prog. Phys.* **60**, 761 (1997).
 - [10] N. Hinkley, J. Sherman, N. Phillips, M. Schioppo, N. Lemke, K. Beloy, M. Pizzocaro, C. Oates, and A. Ludlow, *Science* **341** (2013).
 - [11] N. Huntemann, M. Okhapkin, B. Lipphardt, S. Weyers, C. Tamm, and E. Peik, *phys. rev. lett.* **108** (2012).
 - [12] C. W. Chou, D. B. Hume, J. C. J. Koelemeij, D. J. Wineland, and T. Rosenband, *Phys. Rev. Lett.* **104**, 070802 (2010).
 - [13] T. Rosenband, D. B. Hume, P. O. Schmidt, C. W. Chou, A. Brusch, L. Lorini, W. H. Oskay, R. E. Drullinger, T. M. Fortier, J. E. Stalnaker, S. A. Diddams, W. C. Swann, N. R. Newbury, W. M. Itano, D. J. Wineland, and J. C. Bergquist, *Science* **319** (2008).
 - [14] G. Dick, *Proc. 19th Annual Precise Time and Time Interval (PTTI) Appls. and Planning Meeting* (1987).
 - [15] Y. Y. Jiang, A. D. Ludlow, N. D. Lemke, R. W. Fox, J. A. Sherman, L.-S. Ma, and C. W. Oates, *Nat. Photon.* **5**, 158 (2011).
 - [16] M. Takamoto, T. Takano, and H. Katori, *Nature Photonics* **5**, 288 (2011).
 - [17] J. Hartnett and N. Nand, *IEEE Trans. Microwave Thy and Tech.* **58**, 3580 (2010).
 - [18] C. Hagemann, C. Grebing, T. Kessler, S. Falke, N. Lemke, C. Lisdat, H. Schnatz, F. Riehle, and U. Sterr, *IEEE Trans. Instr. and Meas.* **62**, 1556 (2013).
 - [19] G. Biedermann, K. Takase, X. Wu, L. Deslauriers, S. Roy, and M. Kasevich, *Phys. Rev. Lett.* **111** (2013).
 - [20] B. Girod, R. Rabenstein, and A. Stenger, *Signals and Systems* (Wiley, New York, 2001).
 - [21] A. Kofman and G. Kurizki, *Phys. Rev. Lett.* **87**, 270405 (2001).
 - [22] T. J. Green, H. Uys, and M. J. Biercuk, *Phys. Rev. Lett.* **109**, 020501 (2012).
 - [23] T. J. Green, J. Sastrawan, H. Uys, and M. J. Biercuk, *New J. Phys.* **15**, 095004 (2013).
 - [24] A. Soare, H. Ball, D. Hayes, J. Sastrawan, M. C. Jarratt, J. J. McLoughlin, X. Zhen, T. J. Green, and M. J. Biercuk, *Nat. Phys.* **10** (2014).
 - [25] R. F. Stengel, *Optimal Control and Estimation* (Dover, Mineola, New York, 1994).
 - [26] J. Rutman, *Proceedings of the IEEE* **66**, 1048 (1978).
 - [27] W. H. Oskay, S. A. Diddams, E. A. Donley, T. M. Fortier, T. P. Heavner, L. Hollberg, W. M. Itano, S. R. Jefferts, M. J. Delaney, K. Kim, F. Levi, T. E. Parker, and J. C. Bergquist, *Phys. Rev. Lett.* **97**, 020801 (2006).
 - [28] *Characterization of Clocks and Oscillators*, Tech. Rep. 1337 (National Institute of Standards and Technology, 1990).
 - [29] J. Barnes, A. Chi, L. Cutler, D. Healey, D. Leeson, T. McGunigal, J. Mullen, W. Smith, R. Sydnor, R. C. Vessot,

- and G. Winkler, *IEEE Trans. Instr. and Meas.* **IM-20**, 105 (1971).
- [30] C. Greenhall, D. Howe, and D. Percival, *IEEE Trans. Ultrason., Ferroelec., and Freq. Control* **46**, 1183 (1999).
- [31] G. Uhrig, *Phys. Rev. Lett.* **98**, 100504 (2007).
- [32] L. Cywinski, R. M. Lutchyn, C. P. Nave, and S. D. Sarma, *Phys. Rev. B* **77**, 174509 (2008).
- [33] M. J. Biercuk, H. Uys, A. P. VanDevender, N. Shiga, W. M. Itano, and J. J. Bollinger, *Nature* **458**, 996 (2009).
- [34] M. J. Biercuk, A. C. Doherty, and H. Uys, *J. Phys. B* **44**, 154002 (2011).
- [35] C. A. Greenhall, *Metrologia* **40**, S335 (2003).
- [36] B. Penrod, in *Frequency Control Symposium, 1996. 50th., Proceedings of the 1996 IEEE International*. (1996) pp. 980–987.
- [37] C. Nichols, P. Wu, G. Carleton, and S. Beaudin, (2006).
- [38] A. Soare, H. Ball, D. Hayes, X. Zhen, M. C. Jarratt, J. Sastrawan, H. Uys, and M. J. Biercuk, *Phys. Rev. A* **89**, 042329 (2014).
- [39] S. Micalizio, C. E. Calosso, A. Godone, and F. Levi, *Metrologia* **49**, 425 (2012).
- [40] S. Mavadia, V. Frey, J. Sastrawan, S. Dona, and M. J. Biercuk, *arXiv:1604.03991* (2016).
- [41] M. Mullan and E. Knill, *arXiv* **1404.3810** (2014).