

Active Cancellation of Servo-Induced Noise on Stabilized Lasers via Feedforward

Lintao Li,^{*} William Huie[✉], Neville Chen, Brian DeMarco, and Jacob P. Covey^{✉†}

Department of Physics, The University of Illinois at Urbana-Champaign, Urbana, Illinois 61801, USA



(Received 4 August 2022; accepted 5 October 2022; published 2 December 2022)

Many precision laser applications require active frequency stabilization. However, such stabilization loops operate by pushing noise to frequencies outside their bandwidth, leading to large “servo bumps” that can have deleterious effects for certain applications. The prevailing approach to filtering this noise is to pass the laser through a high-finesse optical cavity, which places constraints on the system design. Here, we propose and demonstrate a different approach where a frequency error signal is derived from a beat note between the laser and the light that passes through the reference cavity. The phase noise derived from this beat note is fed forward to an electro-optic modulator after the laser, carefully accounting for relative delay, for real-time frequency correction. With a hertz-line-width laser, we show ~ 20 dB noise suppression at the peak of the servo bump (approximately 250 kHz) and a noise-suppression bandwidth of approximately 5 MHz—well beyond the servo bump. By simulating the Rabi dynamics of a two-level atom with our measured data, we demonstrate substantial improvements to the pulse fidelity over a wide range of Rabi frequencies. Our approach offers a simple and versatile method for obtaining a clean spectrum of a narrow-line-width laser, as required in many emerging applications of cold atoms, and is readily compatible with commercial systems that may even include wavelength conversion.

DOI: [10.1103/PhysRevApplied.18.064005](https://doi.org/10.1103/PhysRevApplied.18.064005)

I. INTRODUCTION

The use of cold-atom systems in burgeoning quantum science applications is pushing the limits of laser engineering. Many applications now require long-term frequency stability (subhertz) but are also sensitive to the fast frequency noise (approximately megahertz). Nowhere is this more apparent than in the use of long-lived optical metastable states, which, in addition to their rich history in optical atomic clocks [1], are gaining prominence in quantum computing [2–12], quantum simulation [13–15], and quantum networking [16–20]. Additionally, programmable entanglement in atomic clocks [4,21,22] for quantum-enhanced precision [23–27] merges the requirements of optical metrology with quantum computing and networking. Atom-laser coherence at the second scale is required to fully leverage such metastable states, yet short-term stability is crucial for, e.g., gate operations and short-term stability of optical atomic clocks limited by Dick noise [28]. Moreover, the recent progress with the use of Rydberg states [29] for programmable entanglement of neutral atoms [3,4,30–33] enables protocols with similar requirements, albeit with shorter time scales.

In recent years, it has become clear that laser noise at the qubit drive frequency is particularly deleterious [30,34,35].

For stabilized lasers, broad noise peaks that normally range from approximately ± 0.1 to 1 MHz, referred to as “servo bumps,” are inevitable due to the finite loop bandwidth. Therefore, the need for stabilized lasers with spectral filters to remove this high-frequency noise has become ubiquitous, with the pervasive approach based on using an optical cavity as a spectral filter [30,36]—typically the same cavity to which the laser is stabilized. However, additional gain stages such as injection-locked diode lasers and tapered amplifiers are required when using the light transmitted through the cavity, since the transmitted power $P_0 \approx 10 \mu\text{W}$ is limited in practice by the power build-up $F/2\pi \times P_0$, where F is the cavity finesse. This limitation is particularly problematic in the “ultrahigh”-finesse ($F \gtrsim 100\,000$) range that is required for “clock” lasers with fractional stability of $\sim 10^{-15}$ at 1 s [1], where the transmitted power may not even be high enough to injection lock a diode laser. Moreover, the cavity-transmission technique is challenging for laser systems—particularly commercial systems—that involve frequency conversion such as second-harmonic generation, which is common for neutral ytterbium, mercury, and cadmium; ionic aluminum, beryllium, and magnesium; and Rydberg transitions of essentially every species.

Here, we demonstrate an alternative approach to spectral filtering of servo bumps based on active noise cancellation via a feedforward technique. By using the cavity-transmitted light to generate a beat note with the laser

^{*}ltli@illinois.edu

[†]jcovey@illinois.edu




FIG. 1. An overview of the scheme and the main results. (a) An external-cavity diode laser (ECDL) is locked to a high-finesse optical cavity, which produces “servo bumps” on its spectrum. The light transmitted through the cavity serves as a spectral filter to remove these bumps and the real-time frequency and/or phase deviation of the laser with respect to the cavity-transmitted reference is obtained via a beat note between the two. This signal is then fed forward to an electro-optic modulator (EOM), possibly with an optical-fiber delay line added onto the light, to cancel the frequency deviations. For an ideal cavity with infinitesimal line width, the frequency noise of the laser can be denoted as 1ν and the single-frequency laser after the cavity is denoted as ν_0 . (b) Laser spectra with (red and blue) and without (yellow) the feed-forward (FF) applied to the laser. The blue (red) data show the case where a 30-m fiber delay line is used (not used) prior to the EOM. The black data show the background noise level, where a fictitious data point at 200 MHz and approximately 35 dB μ V due to rf leakage is removed. The inset shows the same data over a wider frequency range. The resolution bandwidth is 100 Hz for all the measurements in this plot. (c) Simulations of driven two-level atoms under the three laser noise cases from (b) using a Rabi frequency of $\Omega = 2\pi \times 200$ kHz, which is near the peak of the servo bump.

[37,38], the frequency deviation of the laser from the cavity-filtered reference can be corrected in real time with an active optical device such as an electro-optic modulator (EOM) [39,40] [see Fig. 1(a)]. With optimal signal-delay compensation using an optical fiber, we demonstrate that this technique can suppress noise up to $f_{nc} \approx 5$ MHz (nc means noise canceling) by at least 3 dB below that of the original laser and that the peak of the servo bump ($f_p \approx 250$ kHz) is suppressed by approximately 22 dB. Noise above f_{nc} is not adversely affected and remains at the level of the original laser and f_{nc} is limited only by the electronic circuitry. Without including an optical fiber, $f_{nc} \approx 1$ MHz and the servo-bump peak is suppressed by approximately 15 dB [see Fig. 1(b)]. Moreover, we demonstrate that this technique does not come at the expense of long-term stability, even in the case with a fiber delay line, by showing a Fourier-limited width of ~ 20 mHz for the coherent peak of the beat note. Finally, we simulate the dynamics of a

two-level system under a time-dependent drive based on our data, showing an enormous improvement [Fig. 1(c) shows Rabi frequency $\Omega \approx f_p$]. Our approach is simple and versatile, provides noise suppression out to high enough bandwidths to be relevant for Rydberg excitation, and can readily be applied to complex and commercial laser systems that involve both high-finesse cavities and frequency conversion.

II. OVERVIEW OF THE TECHNIQUE

Feedforward techniques have been employed to remove laser phase noise for decades, beginning with low-speed fiber-noise cancellation [41], which is still widely used today. More recently, phase-noise feedforward has been extended to higher bandwidth via a self-heterodyne laser phase-noise measurement [40,42,43]. However, this technique suffers from the thermal fluctuation of the fiber delay line, limiting the long-term stability. Meanwhile, researchers working in the microwave-frequency domain have developed a feedforward method by comparing the real-time phase of a noisy oscillator with a clean reference [44]. Similarly, cavity transmission can be used as a phase-noise reference [38] for real-time feedforward to a noisy external-cavity diode laser (ECDL). Our work combines aspects of both the optical and rf approaches to phase-noise cancellation.

We provide a more detailed description of the technique and the circuitry in Fig. 2; additional information can be found in Appendix A. We send the light from the ECDL through an acousto-optic modulator (AOM) and interfere half of it with half of the light collected after the cavity on a 50:50 beam splitter (BS), generating a beat note at 200 MHz that is measured with a photodiode (PD1 in Fig. 2). Note that the amplitude of this beat note scales as the product of the electric field of both sources, so even a small cavity-transmitted power provides sufficient signal (we use approximately 4 μ W for beating with an approximately 40 μ W laser from the ECDL in the experiment). The cavity-transmitted power can be further reduced by scaling up the ECDL laser power by the same factor.

The other half of each source is used to perform the feed-forward for noise cancellation and to measure the result. The laser light after the AOM is sent through an optical-fiber delay line and then into an EOM. We use a free-space EOM rather than a waveguide EOM because it can handle substantially higher optical power and it has lower residual-amplitude modulation (see below). The output is interfered with the cavity-transmitted light on a BS and both ports are monitored with a PD (PD2 and PD3 in Fig. 2).

The signal from PD1 is converted to a time-dependent phase deviation, which is amplified and sent to the EOM with proper polarity. The phase-locked-loop- (PLL) based phase-detection circuit has a working frequency range




FIG. 2. A detailed schematic. The clock ECDL is locked to a high-finesse cavity via the standard Pound-Drever-Hall (PDH) technique (not shown). The ECDL light and the cavity-transmitted light are used in two places for our measurements, so each is split with a 50:50 beam splitter (BS). An acousto-optic modulator (AOM) is used to create a 200-MHz beat note between the ECDL and the cavity-transmitted light on photodiode PD1 by interfering them on a BS. This signal is used to cancel frequency deviations from the cavity-transmitted reference through the use of the feedforward phase-correction circuit shown in the teal box. After the AOM, the laser light is sent through an optical-fiber delay line before going to an electro-optic modulator (EOM). The phase signal from the circuit in the teal box is applied to the EOM. Then, this light is interfered with the cavity-transmitted light on a BS and monitored on two PDs. PD2 is used for measurements and represents what would be sent to the atoms. PD3 is used to measure slow fiber-induced frequency drifts, which are then corrected with the fiber-noise-cancellation circuit (green box). This feedback system corrects the frequency applied to the AOM to compensate these slow drifts. The feedback and feedforward work together to maintain the highly coherent peak of the cavity-stabilized laser while simultaneously removing fast noise from the servo bumps.

from approximately 200 Hz to 4.8 MHz (see Appendix A). The wide-band amplifier that drives the EOM has a bandwidth exceeding 15 MHz. The circuitry, from PD1 to the EOM, introduces a delay of $\tau_e \approx 140$ ns (see Appendix A), so we add an optical delay τ_o using an optical fiber. As discussed below, we identify $L = 30$ m as the optimal fiber length such that $\tau_e \approx \tau_o$. Besides the delay, the gain of the wide-band amplifier is also critical. The coarse adjustment of the gain of the amplifier is achieved by estimating the sensitivity of the phase discriminator and V_π of the EOM. For fine adjustment, we tune the gain of the amplifier by examining the amplitude of the coherent peak (see Appendix B).

The addition of the fiber delay helps to optimally cancel high-frequency noise but it is well known that thermal and acoustic fluctuations introduce low-frequency noise to optical signals passing through fibers, which can be canceled with active feedback [41]. We use PD3 to monitor

this slow noise by mixing the 200-MHz signal with a radio-frequency (rf) source of the same frequency, thereby creating a dc signal that is used in a proportional-integral-differential (PID) servo loop with the frequency-tunable rf source that drives the AOM. Therefore, the frequency of the laser is actively adjusted to compensate for the slow fiber-induced drifts. Due to this fiber-noise cancellation, the rf frequency of the AOM varies slowly with time. Accordingly, the beat signal on PD1 will not only have the high-frequency servo bumps from the ECDL but also the slowly varying phase noise from the AOM. Since the feedforward circuit will only respond to phase noise higher than 200 Hz (for a detailed description, see Appendix A), the slow fiber-noise cancellation does not adversely affect the function of the feedforward circuit. In this way, we can preserve the long-term stability of the laser obtained by locking it (via feedback) to the high-finesse cavity while simultaneously removing high-frequency servo-induced noise via feedforward.

Our specific implementation of the fiber-noise cancellation can only be applied when the fiber is spooled such that the output is proximal to the input and other optical paths. In the case where a fiber is used to span some distance, more standard implementations of fiber-noise cancellation can be employed using the reflection from the fiber output to generate an interferometer [41].

Finally, we emphasize that this feedforward technique, by construction, cannot filter noise better than the cavity does—although it comes close over a wide bandwidth that covers most of the laser phase noise. Specifically, the cavity used in this work has a modest finesse (see Appendix E), so the results discussed below for the feedforward data with the $L = 30$ m delay length probably exceed the quality of the cavity-transmitted spectrum. That is, we assume that the cavity-transmitted signal has no noise. However, our approach is intended for ultrahigh-finesse cavity systems for which the cavity-transmitted noise becomes negligible.

III. RESULTS

To quantify our results, we compare at each frequency the spectral density of the feedforward signal $I_{w/f}(f)$ to the case without engaging the feedforward $I_{w/o}(f)$ (a detailed calculation can be found in Appendix C). That is, $I(f) = I_{w/f}(f) - I_{w/o}(f)$. Hence, $I < 0$ in the spectral region where the feedforward reduces the noise and $I > 0$ in regimes where the feedforward increases the noise. The bandwidth in which the feedforward reduces the noise depends on the relative delay between the optical and electrical signals. Figure 3(a) shows $I(f)$ for the case of $L = 0, 15$, and 30 m, corresponding to $\tau_o \approx 0, 70$, and 140 ns, respectively. While all three cases show $I < 0$ below approximately 1 MHz, the $L = 0$ and $L = 15$ m cases have $I > 0$ beyond this frequency before approaching zero. In




FIG. 3. The results. (a) Noise-difference spectra $I(f)$ between the case without feedforward engaged and cases with it engaged versus the frequency for several optical-fiber delay lengths L : red, 0 m; purple, 15 m; blue, 30 m. Negative values show reduced noise and positive values show increased noise. The points are measured data while the lines are simulations based on the circuit transfer function (see Appendix A). (b) The integrated noise difference $T(f^0)$ between $f = 0$ and f^0 for several values of $f^0 \in \{1, \dots, 9\}$ MHz versus the optical-fiber delay length L , showing a clear minimum at $L = 30$ m. (c) The noise spectrum over a narrow range with $L = 30$ m, with and without feedforward engaged. The main figure shows the case with a slow feedback loop to the AOM to cancel fiber noise, giving a Fourier-limited coherent peak. This indicates that our technique maintains the intrinsic laser line width. The inset shows the case without canceling fiber noise. (d) The laser-intensity noise with and without applying the feedforward, and the background. We see that the feedforward makes the intensity noise higher due to residual-amplitude modulation from the EOM. We show below that this effect is negligible.

the case of $L = 0$, our numerical model based on the circuit transfer function (see Appendix B) predicts that there will be a dip below zero near 6 MHz where the relative phase delay is roughly 2π , such that the feedforward can successfully cancel noise. Our data does not corroborate this dip, which we believe is because there is limited noise to cancel at 6 MHz due to the profile of the servo bump being primarily below 1 MHz. For $L = 30$ m, our data are in agreement with our numerical modeling to suggest that the relative delay is small, offering optimal performance of the feedforward method. The region where $I < 0$ is maximized under this condition and the signal asymptotes to zero rather than first going positive.

To identify the optimal fiber delay length L , we define a figure of merit $T(f^0)$ as the integrated noise under $I(f)$ from $f = 0$ up to a chosen bandwidth $f = f^0$. Figure 3(b) shows $T(f^0)$ for several values of $f^0 \in \{1, \dots, 9\}$ MHz, where there are minima at $L = 30$ m for all values of f^0 . This is in good agreement with our understanding based on the transfer function of our system (see Appendix B),

suggesting that there is $\tau_e \approx 140$ ns of electronic delay. For $L = 30$ m, $\tau_o = nL/c \approx 140$ ns for optical fiber with a refractive index of $n = 1.44$.

As noted above, the fiber delay line can introduce noise due to acoustic and thermal path-length fluctuations. Thus, the use of a fiber to compensate electronic delay introduces low-frequency noise that must be removed with active feedback. Figure 3(c) shows beat-note spectra similar to those in Fig. 1(b), except over a much narrower frequency range. We consider the optimal case of $L = 30$ m. Figure 3(c, inset) shows the situation without removing the fiber-induced noise for two cases: where the feedforward that removes the high-frequency noise (servo bumps) is and is not engaged. There is no apparent difference on this scale. Figure 3(c, main figure) shows the same two cases, where now the fiber-induced noise is removed by feeding back to the AOM as described above. Here, the spectra are Fourier limited and there is no apparent broadening due to the fiber. We see that the peak corresponding to the case where the feedforward is engaged has a higher amplitude

than the peak where feedforward is not engaged. This is due to the noise-cancellation process that actively transfers the spectral weight of the laser noise back into the coherent peak. Indeed, the increased amplitude shown here is comparable with the spectral weight under the servo bumps. The laser power is unchanged during the measurement.

Note that this spectrum is not a true measure of the laser line width since the coherence length is much longer than the path-length difference in the interferometer arms. We believe from other data that our laser has ~ 1 Hz line width (for details about the laser system, see Appendix E).

We also consider the intensity noise of the laser with and without applying the feedforward. Figure 3(d) shows the intensity noise of the background as well as the laser light in both cases. We see that the case with the feedforward applied has increased noise within approximately 4 MHz compared to the case without the feedforward applied. We attribute this to residual-amplitude modulation (RAM) on the EOM, in which frequency modulation is partially mapped onto amplitude modulation, typically due to an imperfect input polarization. This problem is more prevalent with waveguide EOMs and our findings underline our decision to use a free-space EOM. We believe that improvements to our setup would further reduce the measured RAM. To quantify the significance of this added intensity noise, we turn to numerical modeling.

We simulate a driven two-level atomic system under the present of the measured noise to gauge the effect of our noise-cancellation system on atomic dynamics (for details, see Appendix D). First, we consider only frequency noise. We specifically focus on the cases of the feedforward engaged with $L = 0$ and $L = 30$ m and the feedforward not engaged. Figure 4 shows the 9π pulse fidelity, which is the fifth maximum of the population-transfer fraction of the Rabi evolution, versus the Rabi frequency used in the simulation. We observe that Rabi frequencies in regions of high noise give particularly poor fidelity and thus these results are qualitatively similar to the inverse of the noise spectra shown in Fig. 1(b). Figure 1(c) shows time traces of these Rabi evolutions for a Rabi frequency of $\nu = 2\pi \times 200$ kHz. Correspondingly, we see in Fig. 4 that the case without feedforward has poor fidelity between 0.2 MHz and approximately 1 MHz, which then asymptotically approaches unity well beyond the servo-bump bandwidth. The case of feedforward with $L = 0$ m is significantly improved at low frequency but this improvement is lost beyond approximately 1.3 MHz, where the relative delay time causes the feedforward to degrade the fidelity. The case of feedforward with $L = 30$ m has excellent fidelity over the entire frequency range, since noise is efficiently canceled within the servo-bump bandwidth and noise is not added outside this bandwidth.

Figure 4 also includes simulations with the measured intensity noise from Fig. 3(d) (see Appendix D for details) and we consider both frequency noise and intensity noise




FIG. 4. The simulation of a driven two-level atom. The fidelity after a 9π -pulse (the fifth maximum of the population-transfer fraction) of Rabi dynamics with time-dependent drives based on measured noise spectra (see Appendix D) is shown for the cases without feedforward (solid yellow) and with feedforward for $L = 0$ and $L = 30$ m (solid red and blue). We see that the $L = 30$ m case (blue) has excellent fidelity over the entire bandwidth. The $L = 0$ m case (red) is improved compared to the no-feedforward case below approximately 1 MHz but degrades substantially at higher frequency. The case without feedforward (solid yellow) returns to high fidelity for frequencies well above the servo-bump regime. We also consider the case where intensity noise is included (dashed lines). We see that the intensity noise is negligible compared to the frequency noise.

for the two cases with feedforward applied. In practice, we believe that the frequency noise may be synchronized with the intensity noise but the two are uncorrelated in our simulations. We find that the intensity noise is irrelevant compared to the frequency noise over the entire range of • considered. This can be understood by considering the similarity between the amplitude and the phase modulation for small index. We can express a phase-modulated signal with the phase variation defined by $1\varphi \rightarrow 0$ as $\sin(\omega t + 1\varphi) \approx \sin(\omega t) + 1\varphi \times \cos(\omega t)$, where ω is the carrier frequency and the second term is amplitude modulated by the phase variation 1φ . By converting the curves in Fig. 1(b) into phase noise (see Appendix D), we find that the phase noise is always higher than -90 dBc/Hz even for the best configuration (with feedforward and a 30-m fiber delay). This phase-noise level is still much higher than the intensity noise in Fig. 3(d).

IV. DISCUSSION AND CONCLUSIONS

We present a simple and versatile alternative to the cavity-filtration-based approach to removing servo-induced noise on stabilized lasers. Our technique is based on measuring the real-time frequency deviation from the cavity resonance via a beat note between the cavity-transmitted light and the laser, which is then fed forward

to an active device. We find that an EOM after a $L = 30$ m optical-fiber delay line to compensate the $\tau_e \approx 140$ ns electronic delay provides noise suppression of approximately 22 dB within the servo bump and can suppress noise with a bandwidth as high as approximately 5 MHz. We also show that this technique can be applied without compromising the long-term stability of the laser by demonstrating a Fourier-limited ~ 20 mHz line-width beat note for the coherent peak. This is accomplished by adding a fiber-noise-cancellation feedback loop and our technique seamlessly combines the feedforward for high-frequency noise removal with the feedback for low-frequency noise removal.

In cases where the laser system includes frequency doubling or quadrupling from the wavelength referenced to the optical cavity, the only required change to our approach—if it is to be applied at the final output—is to apply a commensurate factor to the frequency modulation applied to the EOM downstream. More generally, any other transfer function can be added to the electrical circuit to match a transfer function in the laser system such as, for instance, sum- or difference-frequency generation that could be combined with second-harmonic generation (SHG). Given the large bandwidths of single-pass nonlinear optics and resonant SHG cavities, we do not anticipate significant effects on the feedforward bandwidth even in such complex laser systems. Finally, for UV systems where it may be impractical to have an optical fiber, we point out that the fiber delay line could come before a frequency-doubling stage, for instance. In particular, the fiber and EOM could be before the doubling stage(s) and even before an amplification stage with saturable gain, obviating the already negligible added intensity noise due to the EOM RAM.

Additionally, we suggest a nonoptical variant of our feedforward technique specifically for the case of driving transitions between a ground state and a Rydberg state. Instead of using the real-time frequency deviation to apply a correction to the light via the EOM, one could instead apply the correction to the atomic transition frequency. In particular, a microwave field detuned from a nearby opposite-parity Rydberg state could be used to Stark shift the target Rydberg state [45]. Amplitude modulation of this microwave field would thereby modulate the optical frequency of the ground-Rydberg transition. This technique could, in principle, have higher bandwidth than the EOM.

In conclusion, we believe that our fast feedforward technique offers opportunities for state-of-the-art laser engineering and may find application in intensity noise mitigation as well. Our approach provides increased access to ultranarrow lasers with a clean spectrum free of servo bumps, approaching the true delta-function limit without the complexity associated with the cavity-filtration technique—especially for ultrahigh-finesse cavity systems.

Therefore, we believe that our techniques will find application in myriad emerging quantum science directions that rely on optically metastable states and Rydberg states.

ACKNOWLEDGMENTS

We thank Brett Merriman for helpful discussions and Jun Ye for critical reading of the manuscript. We acknowledge funding from the National Science Foundation (NSF) Quantum Leap Challenge Institutes (QLCI) for Hybrid Quantum Architectures and Networks (NSF OMA-2016136); the NSF Division of Physics (PHY) (NSF Award No. 2112663); the NSF Quantum Interconnects Challenge for Transformational Advances in Quantum Systems (NSF Award No. 2137642); and the Office of Naval Research (ONR) Young Investigator Program (ONR Award No. N00014-22-1-2311). This material is based upon work supported by the U.S. Department of Energy Office of Science National Quantum Information Science Research Centers as part of the Q-NEXT center.

APPENDIX A: DRIVE CIRCUITRY AND SIMULATION

The drive circuitry is divided into two parts, the phase-measurement circuit and the driver for the free-space EOM.

As described in the main text, the beat signal on PD1 includes both the high-frequency phase noise from the servo bump and the slowly varying phase noise from the AOM. The working principle of the phase-measurement circuit is based on using a low-bandwidth PLL that tracks the slowly varying phase noise but ignores the high-frequency phase noise. Since the voltage-controlled oscillator tracks the slowly varying phase noise from the AOM, only the servo-bump phase noise will appear after the phase detector.

The real-time phase-measurement circuit is shown in Fig. 5(a). The phase comparator U2 compares the phase difference between the input beat signal J1 and the output of the voltage controlled crystal oscillator X1. The oscillator X1 has very low phase noise and serves as the reference oscillator for the phase comparison. Thus the output of U2 is mainly dependent on the phase noise of the input beat signal. Since the phase comparator only works for a phase difference ranging from $-\pi$ to π , the output of the phase comparator is fed back to X1 via a slow loop filter formed by U1B. Consequently, the slow fiber-noise-cancellation servo, the long-term phase drifts, and the unwanted phase wraps will not affect the measurement of the phase noise. Finally, the wide-band dc amplifier U1A converts the differential output of U2 to a single-ended phase-noise voltage signal, which is sent to the EOM driver circuit.




FIG. 5. The drive circuitry used in the feedforward-noise-canceling setup. (a) The PLL-based phase-detection circuit, which converts the real-time phase deviation of the input beat signal into voltage. (b) The high-bandwidth and high-voltage wide-band amplifier, which amplifies the phase-deviation signal to high voltage that matches the V_{π} of the free-space EOM.

The free-space EOM is preferred in this setup because of its ability to handle high power and maintain good polarization stability. However, this type of EOM needs a relatively high driving voltage compared with a waveguide EOM. A high-speed and high-voltage EOM driver is required to decrease the electronics delay and increase the modulation bandwidth. Figure 5(b) shows the schematic of our homemade EOM driver that meets these requirements. The key component is a cathode-ray-tube (CRT) driver (Philips Semiconductors, TDA6120Q) that, in past decades, has been used to drive the cathodes of a CRT in High Definition televisions or monitors. The amplifier is tested up to a peak-to-peak voltage output of 160 V with a slew rate around 10 000 V/ μ s. A typical low-capacitance free-space EOM (Thorlabs, EO-AM-NR-C4) with smaller than 20 pF can be used since it is comparable to the capacitance of a CRT cathode. However, any

long-distance coax cable should be avoided due to the high parasitic capacitance.

The simulation of the phase-calculation circuit is based on the TINA-TI software [46] and the voltage-controlled crystal oscillator and the phase detector are replaced with an integrator that gives the same overall transfer function. The simulation indicates a lower-frequency response of around 200 Hz and a high-frequency response up to 4.8 MHz. For the step response, the simulation gives a bandwidth-limited delay of around 130 ns; this delay together with the 10 ns propagation delay in the high-voltage amplifier contributes the overall electronics delay $\tau_e \approx 140$ ns. The circuit simulation matches the measurement of the noise-canceling factor calculated in Figure 3(a).

The physical concept underlying our phase-noise cancellation is the fact that all the noise that we want to remove

comes from the feedback system, being the intrinsic laser phase noise multiplied by the noise gain $[1 + A(s)]^{-1}$ of the feedback loop, where $A(s)$ is the open-loop gain of the system. The function $[1 + A(s)]^{-1}$ is also known as the noise-shaping function, which moves the laser phase noise at low frequency up to higher frequency [47]. We note that this noise-shaping picture is valid only for broadband noise or when the loop bandwidth is significantly smaller than the bandwidth of the noise. The information of the phase noise can either be derived from the error signal that modulates the laser frequency or by interfering the laser output with the cavity transmission [38].

APPENDIX B: FEEDFORWARD NOISE CANCELLATION AND THE MODELING

With the knowledge of the phase noise, we use the feedforward system following a high-finesse cavity-locked laser to remove the servo bump. Compared with a feedback-based approach, a feedforward system gives us more flexibility and advantages for removing the phase noise. First, in a feedforward system, the control variable is not error based and thus there is no stability issue that is pervasive for a closed-loop system. Second, the delay τ_e of the phase-noise measurement and frequency-modulation device can be compensated by introducing the same amount of delay in the optical path. In a closed-loop system, the loop bandwidth of a feedback system is always limited by $1/4\tau_e$ [48].

Considering the finite bandwidth and the delay in the system, the transfer function of the feedforward phase-noise canceling system can be written as [49]

$$H(s) = 1 - \frac{e^{-(\tau_e - \tau_o)s}}{1 + \frac{s}{2\pi f_3 \text{ dB}}}, \quad (\text{B1})$$

where $f_3 \text{ dB}$ is the bandwidth of the feedforward system. The factor $e^{-(\tau_e - \tau_o)s}$ describes the phase delay introduced by the fixed optical and electronics time delay. When $\tau_o = \tau_e$, the transfer function $H(s)$ gives a noise attenuation factor $f_3^2 \text{ dB} / (f_3^2 \text{ dB} + f^2)$, which is 0.5 when $f = f_3 \text{ dB}$ and gradually approaches 0 for higher frequency. For a system with infinite bandwidth, the feedforward can be understood as a self-heterodyne measurement with the fiber delay τ_o replaced by the overall delay of $\tau_e - \tau_o$ [50,51]. For more general cases, the noise attenuation is determined by the Bode magnitude $|H(s)|$ of the transfer function and the frequency response is given by $20 \text{ Log}_{10}[|H(s)|]$ (in decibels) [49], where s is the complex frequency. The frequency response of the transfer function $H(s)$ under different $\tau_e - \tau_o$ is shown in Fig. 3(a) and the dip in the red curve can be understood as the laser phase noise around 6 MHz being canceled by the measurement of the same phase noise around one period (approximately 140 ns) before.




FIG. 6. The circuit simulation of the phase-detection circuit. (a) The Bode plot of the output of the phase-detection circuit versus the phase-variation signal input. (b) The step response of the output of the phase-detection circuit.

The transfer function B1 is based on the ideal match between the phase disturbance and the amplitude of the feedforward, which is not exactly true for our experiment. However, for noise suppression up to 20 dB (1% in power), this only requires that an amplitude error is less than 10%, which can readily be achieved. In fact, the performance of the feedforward system is mostly limited by the bandwidth of the phase-noise measurement and correction, which is fully described by the transfer function. Thus, this transfer function is a good approximation for our system and adjustment of the feedforward gain is straightforward and robust.

Finally, we note that a similar feedforward method is also employed in phase-noise removal for a microwave-band oscillator. A subsampling phase detector (SSPD) is used for measuring the phase noise and feeding it forward to a variable delay line for canceling the phase noise [44].

APPENDIX C: PHASE-NOISE CALCULATION

For the spectral density of the feedforward signal $I_{w/f}(f)$ and $I_{w/o}(f)$, we normalize the noise spectra by the signal power at dc, which is the laser power in the coherent peak. For both $I_{w/f}(f)$ and $I_{w/o}(f)$, the dc part of the spectra is removed after the normalization. As shown in Fig. 3(c), the feedforward actively transfers the laser noise into the coherent peak and thus the normalization will introduce an error to the calculated noise suppression. However, this difference in power is less than 14% and has no significant effect on the results shown in Figs. 3(a) and 3(b).

The phase noise is calculated by the beat signal in Fig. 1(b), following the standard procedure [52]. From the normalized spectra $I_{w/}(f)$ and $I_{w/o}(f)$ (in dBc/Hz), we can also calculate the phase noise:

$$L_{w/,\text{w/o}}(f) = I_{w/,\text{w/o}}(f + f_0) - 10\text{Log}_{10}(R), \quad (\text{C1})$$

where $L(f)$ is the single-sideband (SSB) phase noise, f_0 is the frequency of the coherent peak, and $f > 0$. R is the resolution bandwidth used in the spectra measurement.

For the phase-noise simulation, we need convert this phase noise from dBc/Hz to rad^2/Hz . This relation is as follows:

$$S_\phi(f) = [\arcsin(10^{L(f)/20})]^2, \quad (\text{C2})$$

where $S_\phi(f)$ is the single-sided phase noise in rad^2/Hz .

APPENDIX D: MODELING A DRIVEN TWO-LEVEL ATOM

In order to quantitatively evaluate the influence of the laser phase noise on the atomic state dynamics, we numerically solve the Schrödinger equation of a two-level system driven by the time-dependent phase factor $\varphi(t)$ and coupling fluctuation factor (t) . For zero laser detuning, the Hamiltonian of the system is given by

$$\hat{H}(t) = \frac{[1 + (t)] \bullet}{2} e^{-i\varphi(t)} |e\rangle\langle g| + \text{H.c.}, \quad (\text{D1})$$

where $\bullet$ is the Rabi frequency and $\varphi(t)$ and (t) are the time traces of the laser phase and the field-amplitude variation, respectively.

From measured noise spectra, we calculate the phase noise in Appendix C. Using the measured single-sided phase noise $S_\phi(f)$ (in rad^2/Hz) [38], we can calculate sets of the time trace $\varphi^i(t)$ for the given laser phase noise [53,54]:

$$\varphi^i(t) = \sum_f \frac{\bullet}{2} \sqrt{\frac{S_\phi(f)}{1f}} \cos(2\pi ft + \varphi_f^i), \quad (\text{D2})$$

where φ_f^i is the random phase offset of frequency f and phase-noise trace i and $1f$ is the frequency resolution of $S_\phi(f)$. The phase offset is sampled uniformly, $\varphi_f^i \in [0, 2\pi)$.

By the same method, we can calculate the laser-intensity noise trace $1I^i(t)$ with the intensity noise spectrum $S_I(f)$ shown in Fig. 3(d). Since the Rabi frequency $\bullet \ll \bar{I}$, (t) is determined by $\frac{[I_0 + 1I^i(t)]/I_0 - 1}{\bar{I}}$, where I_0 is the average laser intensity. Considering that the noise is wide band and small ($1I \ll I_0$), we can use the approximation $\frac{[I_0 + 1I^i(t)]/I_0 - 1}{\bar{I}} = \frac{1}{2} 1I^i(t)/I_0$.

With the calculated time traces $\varphi^i(t)$ and $1I^i(t)$, we numerically simulate the time-dependent Hamiltonians

using the Runge-Kutta method [55]. We average the time evolution of the state vector over a series of $N = 50$ trials of the phase-noise trace. The Rabi dynamics and the corresponding fidelity under different Rabi frequencies are shown in Figs. 1(c) and 4. We note that the intensity noise and phase noise are uncorrelated in our simulation, which is not correct for the experiment. However, we believe that this discrepancy introduces only a minor effect, since the intensity noise is significantly smaller than the phase noise.

APPENDIX E: DETAILS OF THE LASER SYSTEM

The laser we use in this experiment is an ECDL from Toptica Photonics, working at 674 nm for Sr^+ ions. The laser is locked to a high-finesse ($F \approx 35\,000$) cavity provided by Stable Laser Systems using the Pound-Drever-Hall (PDH) technique and the cavity line width is around 45 kHz. For a state-of-the-art high-finesse cavity with $F \approx 100\,000$, the laser line width of the cavity is usually smaller than 10 kHz. For the frequency stability of our system, the measured Allan deviation is 8×10^{-15} for $\tau = 1$ s, according to the documentation provided by Stable Laser Systems.

Finally, we note that some PDH locking systems may involve the sideband locking technique, which gives a frequency difference between the laser and the cavity-transmitted light of up to half of the free spectral range of the cavity. Our method is also suitable for this tunable sideband technique by frequency down-converting (i.e., using an rf down converter or frequency divider) the beat signal to a lower rf frequency. In fact, the ability to easily handle this offset is another advantage of our approach: if the cavity-transmitted light is used to interrogate atoms, the frequency offset must be addressed with modulators.

-
- [1] A. D. Ludlow, M. M. Boyd, J. Ye, E. Peik, and P. O. Schmidt, Optical atomic clocks, *Rev. Mod. Phys.* **87**, 637 (2015).
 - [2] T. Monz, D. Nigg, E. A. Martinez, M. F. Brandl, P. Schindler, R. Rines, S. X. Wang, I. L. Chuang, and R. Blatt, Realization of a scalable Shor algorithm, *Science* **351**, 1068 (2016).
 - [3] I. S. Madjarov, J. P. Covey, A. L. Shaw, J. Choi, A. Kale, A. Cooper, H. Pichler, V. Schkolnik, J. R. Williams, and M. Endres, High-fidelity entanglement and detection of alkaline-earth Rydberg atoms, *Nat. Phys.* **16**, 857 (2020).
 - [4] N. Schine, A. W. Young, W. J. Eckner, M. J. Martin, and A. M. Kaufman, Long-lived Bell states in an array of optical clock qubits, *Nat. Phys.* **18**, 1067 (2022).
 - [5] H. X. Yang, J. Y. Ma, Y. K. Wu, Y. Wang, M. M. Cao, W. X. Guo, Y. Y. Huang, L. Feng, Z. C. Zhou, and L. M. Duan, Realizing coherently convertible dual-type qubits with the same ion species, *Nat. Phys.* **18**, 1058 (2022).
 - [6] D. T. C. Allcock, W. C. Campbell, J. Chiaverini, I. L. Chuang, E. R. Hudson, I. D. Moore, A. Ransford, C.

- Roman, J. M. Sage, and D. J. Wineland, *omg* blueprint for trapped ion quantum computing with metastable states, *Appl. Phys. Lett.* **119**, 214002 (2021).
- [7] A. Erhard, H. Poulsen Nautrup, M. Meth, L. Postler, R. Stricker, M. Stadler, V. Negnevitsky, M. Ringbauer, P. Schindler, H. J. Briegel, R. Blatt, N. Friis, and T. Monz, and Entangling logical qubits with lattice surgery, *Nature* **589**, 220 (2021).
- [8] N. Chen, L. Li, W. Huie, M. Zhao, I. Vetter, C. H. Greene, and J. P. Covey, Analyzing the Rydberg-based optical-metastable-ground architecture for ^{171}Yb nuclear spins, *Phys. Rev. A* **105**, 052438 (2022).
- [9] Y. Wu, S. Kolkowitz, S. Puri, and J. D. Thompson, Erasure conversion for fault-tolerant quantum computing in alkaline earth Rydberg atom arrays, *Nat. Commun.* **13**, 4657 (2022).
- [10] D. Okuno, Y. Nakamura, T. Kusano, Y. Takasu, N. Takei, H. Konishi, and Y. Takahashi, High-resolution spectroscopy and single-photon Rydberg excitation of reconfigurable ytterbium atom tweezer arrays utilizing a metastable state, *J. Phys. Soc. Jpn.* **91**, 084301 (2022).
- [11] A. Pagano, S. Weber, D. Jaschke, T. Pfau, F. Meinert, S. Montangero, and H. P. Büchler, Error budgeting for a controlled-phase gate with strontium-88 Rydberg atoms, *Phys. Rev. Res.* **4**, 033019 (2022).
- [12] M. Ringbauer, M. Meth, L. Postler, R. Stricker, R. Blatt, P. Schindler, and T. Monz, A universal qudit quantum processor with trapped ions, *Nat. Phys.* **18**, 1053 (2022).
- [13] C. Kokail, C. Maier, R. van Bijnen, T. Brydges, M. K. Joshi, P. Jurcevic, C. A. Muschik, P. Silvi, R. Blatt, C. F. Roos, and P. Zoller, Self-verifying variational quantum simulation of lattice models, *Nature* **569**, 355 (2019).
- [14] M. K. Joshi, A. Elben, B. Vermersch, T. Brydges, C. Maier, P. Zoller, R. Blatt, and C. F. Roos, Quantum Information Scrambling in a Trapped-Ion Quantum Simulator with Tunable Range Interactions, *Phys. Rev. Lett.* **124**, 240505 (2020).
- [15] J. Choi, A. L. Shaw, I. S. Madjarov, X. Xie, J. P. Covey, J. S. Cotler, D. K. Mark, H.-Y. Huang, A. Kale, H. Pichler, F. G. S. L. Brandão, S. Choi, and M. Endres, Emergent randomness and benchmarking from many-body quantum chaos (2021), arXiv:2103.03535.
- [16] B. Casabone, A. Stute, K. Friebe, B. Brandstätter, K. Schüppert, R. Blatt, and T. E. Northup, Heralded Entanglement of Two Ions in an Optical Cavity, *Phys. Rev. Lett.* **111**, 100505 (2013).
- [17] M. Lee, K. Friebe, D. A. Fioretto, K. Schüppert, F. R. Ong, D. Plankensteiner, V. Torggler, H. Ritsch, R. Blatt, and T. E. Northup, Ion-Based Quantum Sensor for Optical Cavity Photon Numbers, *Phys. Rev. Lett.* **122**, 153603 (2019).
- [18] J. P. Covey, A. Sipahigil, S. Szoke, N. Sinclair, M. Endres, and O. Painter, Telecom-Band Quantum Optics with Ytterbium Atoms and Silicon Nanophotonics, *Phys. Rev. Appl.* **11**, 034044 (2019).
- [19] J. P. Covey, A. Sipahigil, and M. Saffman, Microwave-to-optical conversion via four-wave mixing in a cold ytterbium ensemble, *Phys. Rev. A* **100**, 012307 (2019).
- [20] W. Huie, S. G. Menon, H. Bernien, and J. P. Covey, Multiplexed telecommunication-band quantum networking with atom arrays in optical cavities, *Phys. Rev. Res.* **3**, 043154 (2021).
- [21] B. C. Nichol, R. Srinivas, D. P. Nadlinger, P. Drmota, D. Main, G. Araneda, C. J. Ballance, and D. M. Lucas, An elementary quantum network of entangled optical atomic clocks, *Nature* **609**, 689 (2022).
- [22] C. D. Marciniak, T. Feldker, I. Pogorelov, R. Kaubruegger, D. V. Vasilyev, R. van Bijnen, P. Schindler, P. Zoller, R. Blatt, and T. Monz, Optimal metrology with programmable quantum sensors, *Nature* **603**, 604 (2022).
- [23] L. I. R. Gil, R. Mukherjee, E. M. Bridge, M. P. A. Jones, and T. Pohl, Spin Squeezing in a Rydberg Lattice Clock, *Phys. Rev. Lett.* **112**, 103601 (2014).
- [24] E. M. Kessler, P. Kómár, M. Bishof, L. Jiang, A. S. Sørensen, J. Ye, and M. D. Lukin, Heisenberg-Limited Atom Clocks Based on Entangled Qubits, *Phys. Rev. Lett.* **112**, 190403 (2014).
- [25] P. Kómár, E. M. Kessler, M. Bishof, L. Jiang, A. S. Sørensen, J. Ye, and M. D. Lukin, A quantum network of clocks, *Nat. Phys.* **10**, 582 (2014).
- [26] L. Pezzè, A. Smerzi, M. K. Oberthaler, R. Schmied, and P. Treutlein, Quantum metrology with nonclassical states of atomic ensembles, *Rev. Mod. Phys.* **90**, 035005 (2018).
- [27] R. Kaubruegger, P. Silvi, C. Kokail, R. van Bijnen, A. M. Rey, J. Ye, A. M. Kaufman, and P. Zoller, Variational Spin-Squeezing Algorithms on Programmable Quantum Sensors, *Phys. Rev. Lett.* **123**, 260505 (2019).
- [28] E. Oelker, R. B. Hutson, C. J. Kennedy, L. Sonderhouse, T. Bothwell, A. Goban, D. Kedar, C. Sanner, J. M. Robinson, G. E. Marti, D. G. Matei, T. Legero, M. Giunta, R. Holzwarth, F. Riehle, U. Sterr, and J. Ye, Demonstration of 4.8×10^{-17} stability at 1 s for two independent optical clocks, *Nat. Photonics* **13**, 714 (2019).
- [29] M. Saffman, T. G. Walker, and K. Mølmer, Quantum information with Rydberg atoms, *Rev. Mod. Phys.* **82**, 2313 (2010).
- [30] H. Levine, A. Keesling, A. Omran, H. Bernien, S. Schwartz, A. S. Zibrov, M. Endres, M. Greiner, V. Vuletić, and M. D. Lukin, High-Fidelity Control and Entanglement of Rydberg-Atom Qubits, *Phys. Rev. Lett.* **121**, 123603 (2018).
- [31] H. Levine, A. Keesling, G. Semeghini, A. Omran, T. T. Wang, S. Ebadi, H. Bernien, M. Greiner, V. Vuletić, H. Pichler, and M. D. Lukin, Parallel Implementation of High-Fidelity Multiqubit Gates with Neutral Atoms, *Phys. Rev. Lett.* **123**, 170503 (2019).
- [32] T. M. Graham, M. Kwon, B. Grinkemeyer, Z. Marra, X. Jiang, M. T. Lichtman, Y. Sun, M. Ebert, and M. Saffman, Rydberg-Mediated Entanglement in a Two-Dimensional Neutral Atom Qubit Array, *Phys. Rev. Lett.* **123**, 230501 (2019).
- [33] S. Ma, A. P. Burgers, G. Liu, J. Wilson, B. Zhang, and J. D. Thompson, Universal Gate Operations on Nuclear Spin Qubits in an Optical Tweezer Array of ^{171}Yb Atoms, *Phys. Rev. X* **12**, 021028 (2022).
- [34] S. de Léséleuc, D. Barredo, V. Lienhard, A. Browaeys, and T. Lahaye, Analysis of imperfections in the coherent optical excitation of single atoms to Rydberg states, *Phys. Rev. A* **97**, 053803 (2018).
- [35] M. L. Day, P. J. Low, B. White, R. Islam, and C. Senko, Limits on atomic qubit control from laser noise, *npj Quantum Inf.* **8**, 72 (2022).

- [36] M. J. Martin, Ph.D. thesis, JILA/NIST-CU, University of Colorado Boulder (2013).
- [37] E. Ivanov and M. Tobar, Low phase-noise microwave oscillators with interferometric signal processing, *IEEE Trans. Microw. Theory Tech.* **54**, 3284 (2006).
- [38] F. Schmid, J. Weitenberg, T. W. Hänsch, T. Udem, and A. Ozawa, Simple phase noise measurement scheme for cavity-stabilized laser systems, *Opt. Lett.* **44**, 2709 (2019).
- [39] J. L. Hall and T. W. Hänsch, External dye-laser frequency stabilizer, *Opt. Lett.* **9**, 502 (1984).
- [40] F. Aflatouni and H. Hashemi, Wideband tunable laser phase noise reduction using single sideband modulation in an electro-optical feed-forward scheme, *Opt. Lett.* **37**, 196 (2012).
- [41] L.-S. Ma, P. Jungner, J. Ye, and J. L. Hall, Delivering the same optical frequency at two places: Accurate cancellation of phase noise introduced by an optical fiber or other time-varying path, *Opt. Lett.* **19**, 1777 (1994).
- [42] M. Bagheri, F. Aflatouni, A. Imani, A. Goel, and H. Hashemi, Semiconductor laser phase-noise cancellation using an electrical feed-forward scheme, *Opt. Lett.* **34**, 2979 (2009).
- [43] M. Parniak, I. Galinskiy, T. Zwettler, and E. S. Polzik, High-frequency broadband laser phase noise cancellation using a delay line, *Opt. Express* **29**, 6935 (2021).
- [44] S. S. Nagam and P. R. Kinget, in *2018 IEEE Cust. Integr. Circuits Conf.* (IEEE, 2018).
- [45] A. Kumar, A. Suleymanzade, M. Stone, L. Taneja, A. Anferov, D. I. Schuster, and J. Simon, Quantum-limited millimeter wave to optical transduction (2022), arXiv Prepr. [ArXiv:2207.10121](https://arxiv.org/abs/2207.10121).
- [46] T. Instruments, Getting started with TINA-TI2008.
- [47] S. Pavan, R. Schreier, and G. C. Temes, *Understanding Delta-Sigma Data Converters*, 2nd ed. (Wiley-IEEE Press, 2017).
- [48] R. W. P. Drever, J. L. Hall, F. V. Kowalski, J. Hough, G. M. Ford, A. J. Munley, and H. Ward, Laser phase and frequency stabilization using an optical resonator, *Appl. Phys. B Photophys. Laser Chem.* **31**, 97 (1983).
- [49] A. V. Oppenheim, A. S. Willsky, and S. H. Nawab, *Signals and Systems* (Pearson, London, 1996), 2nd ed.
- [50] T. Okoshi, K. Kikuchi, and A. Nakayama, Novel method for high resolution measurement of laser output spectrum, *Electron. Lett.* **16**, 630 (1980).
- [51] L. Richter, H. Mandelberg, M. Kruger, and P. McGrath, Linewidth determination from self-heterodyne measurements with subcoherence delay times, *IEEE J. Quantum Electron.* **22**, 2070 (1986).
- [52] E. Rubiola, Basics of phase noise (2005 IEEE Frequency Control Symposium (FCS) and PTTI Joint Conference, Vancouver, Canada, <http://rubiola.org/pdf-slides/2005T-ifcs-basics.pdf>, 2005).
- [53] P. Cladé, Ph.D. thesis, Laboratoire Kastler Brossel, Université Pierre et Marie Curie-Paris VI (2005).
- [54] Equation D2 in Appendix D is an inverse Fourier transform that converts the phase-noise spectrum into the time domain. The factor of $1/f$ denotes the frequency resolution of the noise spectra. Because the spectrum of phase noise spans frequencies above and below the coherent peak, the factor 2 accounts for the lost noise power in a single-sided phase-noise spectra. We note that this coefficient is 2 in the equation given by Ref. [53] rather than $\sqrt{2}$, which we believe is an error.
- [55] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery, *Numerical Recipes 3rd Edition: The Art of Scientific Computing* (Cambridge University Press, New York, 2007), 3rd ed.