A frequency stabilization technique for diode lasers based on frequency-shifted beams from an acousto-optic modulator

Mevan Gunawardena, Paul W. Hess,[a] Jared Strait,[b] and P. K. Majumder
Physics Department, Williams College, Williamstown, Massachusetts 01267, USA

(Received 2 September 2008; accepted 6 October 2008; published online 30 October 2008)

We present a simple method for diode laser frequency stabilization that makes use of a Doppler-broadened vapor cell absorption signals of two frequency-shifted laser beams. Using second-order-diffracted, double-passed beams from an acousto-optic modulator, we achieve a frequency separation roughly equal to the Doppler half width. The differential transmission signals of the two beams provide an error signal with a very large linear feature, allowing frequency stabilization over a range of greater than 1 GHz by means of standard proportional-integral-derivative servo feedback to the piezoelectric control of the grating in our external cavity diode laser. We have applied this technique to two different diode laser systems, one used to lock to the 410 nm E1 transition in indium and another for locking to the M1/E2 transition in thallium at 1283 nm. In both cases the technique reduces frequency fluctuation to roughly 1 MHz over time scales from 10^{-3} to 10^{2} s. © *2008 American Institute of Physics.*
[DOI: 10.1063/1.3006386]

I. INTRODUCTION

Diode lasers are ubiquitous in high-precision spectroscopy experiments owing to their convenience, broad frequency coverage, and long lifetime. Extended cavity geometries provide easily tunable single-longitudinal-mode operation, with short-term (millisecond time scale) stability at the 1 MHz level and longer-term drifts (time scales of seconds to minutes) of 100 MHz or more. However, when the precision of the spectroscopic experiment requires longer-term frequency stability at the level of 1 MHz, additional feedback based stabilization is necessary. One of the best established of such locking techniques uses saturated absorption to reveal the Doppler-free spectrum of an atomic transition.[1,2] This technique has limited lock-point tuning range and typically requires an off-center lock. An effective technique not requiring Doppler narrowing uses the dichroism induced by very large magnetic fields.[3] A third technique recently developed in our laboratory utilizes the Faraday effect of an atomic vapor in a very small (few Gauss) magnetic field and uses high-precision polarimetry for stabilization.[4]

Here we present a simple method for frequency stabilization based on the differential vapor cell transmission signal of two frequency-shifted laser beam components derived from an acousto-optic modulator (AOM). Unlike previous laser stabilization work that also makes use of AOMs,[5,6] we have been able to achieve submegahertz long-term stabilization without the use of Doppler-narrowing techniques. In terms of wider applicability, the fact that we do not rely on optical saturation makes the technique more feasible for locking to non-electric-dipole atomic transitions (as will be discussed at the end of this paper). The particular convenience of broad lock-point tunability is a key in the spectroscopy experiment now underway in our laboratory.

Our current experiment requires that we lock a 410 nm external cavity diode laser to drive the $5P_{1/2} \rightarrow 6S_{1/2}$ transition in atomic indium (^{115}In). This is the first step of a two-step excitation experiment, the second being driven by a 1291 nm external cavity diode laser system, which is scanned across the hyperfine structure of the $6S_{1/2} \rightarrow 6P_{3/2}$ transition. The particular frequency of the 410 nm laser determines the velocity class of atoms excited to the intermediate state, which in turn determines the exact resonant frequencies of the 1291 nm transitions. The Doppler-narrowed absorption spectrum of the infrared transition is thus particularly sensitive to the stability of the blue laser. If the blue laser is left unstabilized, the resonant frequency of the second-step transition would drift by more than its full homogeneous width. Frequency noise on the latter translates directly into IR absorption signal noise and drifts of the blue laser over the time scale of our frequency scans can lead to systematic errors in extracted hyperfine splittings. Furthermore, it is very desirable to lock the blue laser to various points in the Doppler-broadened excitation curve since our lineshapes are complicated by the presence of small, unresolved absorption peaks from ^{113}In (4% abundance). Different lock points yield different relative excitation rates and different velocity class selections among the isotopes, allowing us an important systematic error check when we eventually fit our hyperfine spectra and search for the effects of small lineshape asymmetries.

As will be discussed below, the key component of this locking scheme is an AOM, which produces multiple frequency-shifted beams with which we employ a differential transmission scheme. The ultimate limit of the

[a]Present address: Department of Physics, Harvard University, Cambridge, MA 01238.
[b]Present address: Department of Electrical Engineering, Cornell University, Ithaca NY 14853.

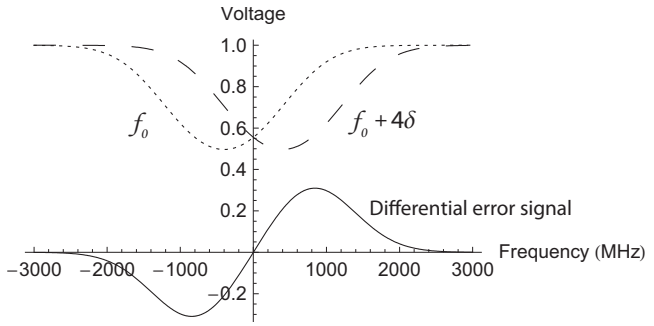

FIG. 1. A model of the Doppler-broadened input beam lineshapes and the differential (error) signal. As described in the text, the frequency separation, 4δ, between beams is provided by the AOM. The Doppler widths correspond to a temperature of 830 °C, the optical depth is 0.7 (typical of our experimental operating conditions), and the total frequency shift is assumed to be 800 MHz. Note the ~1 GHz linear section of the differential lineshape that facilitates locking over a very wide range.

frequency-locking resolution of this method is determined both by the atomic linewidth and the achievable frequency separation of the two laser beam components. The "blue" transition in indium features a very large 1.6 GHz Doppler full width at 800 °C. The AOM crystal used here produces diffracted beams with high efficiency, allowing us to make use of both second-order diffraction and double passing. Typically, AOMs are much less efficient at longer wavelength, making the use of second-order diffracted output infeasible. Yet the Doppler width at these longer wavelengths is correspondingly smaller. In the end, we have achieved roughly comparable final results in terms of frequency stability applying this method at blue and infrared wavelengths.

II. THE DIFFERENTIAL LINESHAPE

Doppler broadening is dominant among lineshape broadening mechanisms for our hot atomic samples. At temperatures near 800 °C, our vapor cell is not optically thin. The hyperfine splittings for this 410 nm transition are on the order of 10 GHz, so that, ignoring the small asymmetry produced by the presence of ^{113}In, we can define the normalized transmission lineshape, $T(\omega,\alpha)$, describing an individual hyperfine transition as

$$T(\omega,\alpha)=\exp[-\alpha A(\omega-\omega_0,\Delta)]. \quad (1)$$

Here α is the measured optical depth, and $A(\omega-\omega_0,\Delta)$ is the normalized Gaussian absorption profile, whose full width at half maximum can be expressed as $\Delta=7.16\times10^{-7}\omega_0\sqrt{T/A}$, where ω_0 is the resonant frequency, T is the Kelvin temperature, and A is the mass in a.u. By producing two frequency-shifted laser beams and detecting the differential transmission signal, we produce an "error" signal as indicated in Fig. 1, where we have assumed a total frequency shift corresponding to roughly the half width of the atomic lineshape. In this case, the linear region of the differential signal extends over more than 1 GHz.

For optimal lock sensitivity, we clearly desire to maximize the slope of the differential signal shown in Fig. 1. While there are practical considerations that restrict us to experimental parameters similar to those indicated in the figure, it is straightforward to model the dependence of the slope on cell temperature and frequency separation of the two laser beam components. Higher temperatures correspond to a greater vapor pressure and higher optical density, which tends to increase the differential slope. Yet higher temperatures increase the Doppler width, broadening the differential curve and decreasing the zero-crossing slope. Similarly, there is an optimal frequency splitting that is set by the experimental Doppler width. Figure 2 shows the results of a simple simulation. Using an analytic form for indium vapor pressure versus temperature,[7] and setting as a fixed point the known optical depth of our cell at one temperature, we then can establish the temperature dependence of both the optical depth and the Doppler width, $\alpha(T)$ and $\Delta(T)$. Figure 2(a) indicates that for the case of our 20-cm-long indium cell, the optimal differential signal slope is found at the readily achievable temperature of 830 °C. On the other hand, it is not surprising that, as seen in Fig. 2(b), the optimal slope occurs when the two-beam frequency splitting is roughly equal to the full width of the Doppler-broadened curve. For our particular AOM, we currently operate at only half of this optimum value, indicating that there is potential for improvement in the performance of this locking scheme. If one could achieve reasonable second-order diffraction efficiency with an AOM model that operates at a rf near 400 MHz, one

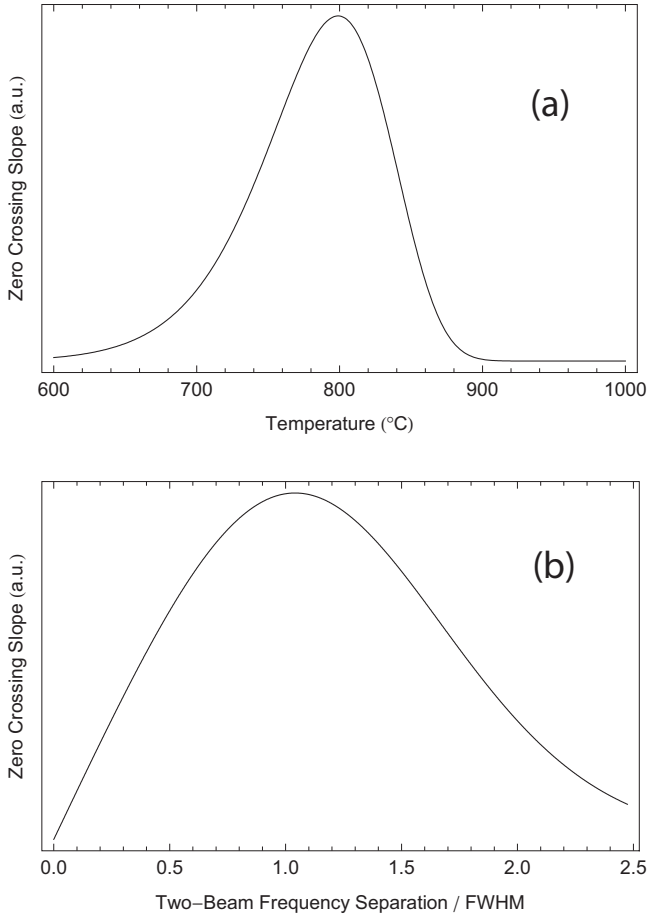

FIG. 2. (Color online) Plot showing magnitude of differential signal slope as a function of cell temperature (a). Below, in plot (b), we show the slope magnitude as a function of two-beam frequency separation, indicating an optimal value near the Doppler FWHM (this is a factor of 2 greater than the value currently used).

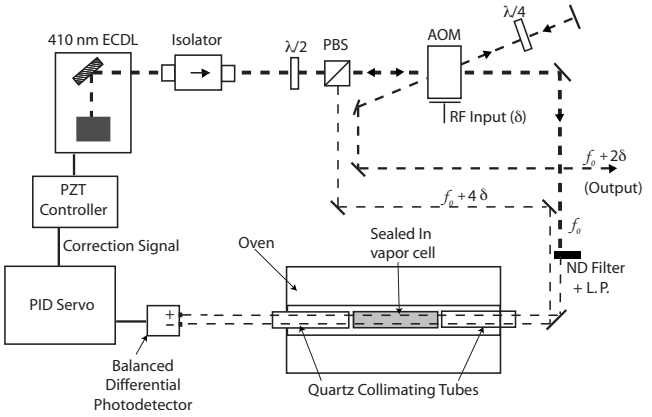

FIG. 3. Layout of the apparatus for the AOM locking scheme. The second-order-diffracted light from the AOM is retroreflected, and the portion of the beam undeflected upon second pass becomes the output beam for the experiment (shown at right). The twice-diffracted beam at $f_0+4\delta$ and the original undiffracted beam become the two inputs to the differential transmission setup. Other diffracted components from the AOM are sent to a scanning Fabry–Pérot cavity, as well as our 0.1 ppm wavemeter for monitoring purposes.

would produce a near-optimum peak separation for this application. Finally, simulations also show that the extent of the linear locking region maximizes near this same temperature. As we heat further, and α becomes large, saturation of the transmission dip leads to a flat portion at the center of the differential lineshape so that locking becomes infeasible.

III. EXPERIMENTAL DETAILS

The laser we use for our primary application is a 410 nm GaN diode laser in the "Littrow" external cavity configuration (Toptica Photonics, Model DL 100), which produces roughly 10 mW of light at its output. After passing through an optical isolator, the laser is incident on the acousto-optic modulator (Isomet 1250C) arranged in a standard double pass configuration, as indicated in Fig. 3. We drive the AOM with roughly 1 W of rf power at $f_{\mathrm{AOM}} \equiv \delta = 200$ MHz. When the AOM is aligned carefully, we diffract roughly 15% of the incident laser beam (frequency f_0) into the second order ($f \equiv f_0+2\delta$). The retroreflected, second-order-diffracted beam passes twice through a quarter-wave plate, producing a beam at $f_0+4\delta$. This counterpropagating double-passed beam is extracted via a polarizing beam splitter. The undiffracted portion of the original incident laser beam, at $f=f_0$, as well as the four times shifted beam, is then directed into our locking vapor cell for differential transmission analysis. Another very convenient feature of our method is the presence of a laser beam component at $f=f_0+2\delta$, which naturally appears as the undiffracted portion of the second-pass beam. We direct this laser component (roughly 0.3 mW) into the vapor cell oven used for our eventual two-step spectroscopy experiment. It is by definition tuned to the lock point of our servo system, which corresponds to the center of the absorption dip for this frequency component. If more power were required for a particular application, one could pick off a portion of the undeflected laser beam, although in this case the "output" beam to the experiment would be shifted away from the center of the absorption peak. Other AOM output beams are typically used for wavelength monitoring, or to monitor the stability and mode structure of the 410 nm laser via a low-finesse Fabry–Pérot cavity. After the double-pass procedure is complete, we typically obtain 75 μW of power at $f_0+4\delta$, as compared to slightly more than 1 mW at f_0.

The two frequency-shifted laser beams are directed in parallel paths roughly 0.5 cm apart into a 20-cm-long, 1-in.-diameter sealed quartz cell containing indium vapor. The cell is placed inside a split-element commercial oven (Applied Test Systems Series 3210). We supply roughly 300 W of ac power to the oven to maintain a temperature of 800 °C. We also feed a stainless steel tube into the center of the oven, allowing us to supply a very small amount of air near the stem of our cell, maintaining a slightly lower temperature in the center, preventing indium condensation from accumulating on the windows. The open design of this oven allows easy optical access to the cell but has the drawback of creating convection currents external to the cell windows. Inserting hollow quartz tubes of matching 1 in. diameter up against the cell windows and protruding several inches out of the oven effectively eliminate the laser beam transmission instability caused by the convective air flow.

After passing through the oven, the two laser beams are directed into the two input channels of a balanced differential photodetector (Thorlabs model PDB150A). Because of the large power asymmetry of the two beams, we insert a neutral density filter and a rotatable polarizer in the path of the stronger beam upstream of the vapor cell. We reduce its power until the error signal from the differential signal output of the detector is precisely zeroed, either by tuning to the symmetric point between absorption dips or by temporarily tuning to a point well off-resonance. An attractive feature of our differential locking scheme is its insensitivity to overall common-mode changes in laser power level. However, changes in AOM diffraction efficiency or other differential or uncorrelated changes to the individual laser beam transmission signals can limit our frequency-locking precision, as will be discussed below.

Once the system has been aligned and the two input channels equalized, the procedure for establishing a lock is straightforward. Using an external Fabry–Pérot cavity, we ensure that the 410 nm laser tunes smoothly in a single-mode fashion over several gigahertz surrounding the absorption dip. The laser sweep is turned off, and we manually tune the laser to a point between the absorption dips of the frequency-shifted beams. The differential transmission error signal is processed by a custom-built proportional-integral-derivative (PID) servo circuit of standard design,[1] and the correction signal is returned to the piezoelectric transducer (PZT) control of the laser external cavity grating. The gain of the servo system is set just below the level that induces oscillation. The time constant of the feedback loop is in our case limited by the response time of the photodetector and locking circuitry. Its value can be determined experimentally by inserting a small electronic square wave component of variable frequency directly into the summing amplifier from which the error signal of the servo circuit is derived. Monitoring the response of the correction signal to this perturbation allows straightforward determination of the time constant, which we

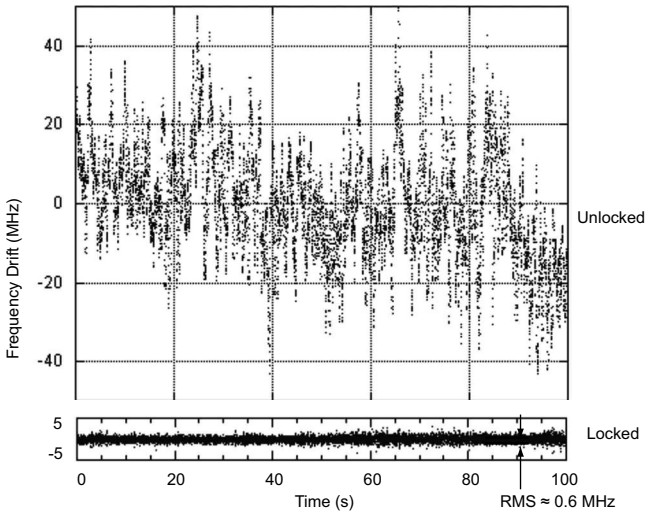

FIG. 4. Plotted on an equivalent vertical scale, the frequency fluctuations of the free-running diode laser are displayed, as well as the residual noise of the locked laser (below).

determine to be close to 1 ms, corresponding to an effective bandwidth of roughly 150 Hz. This is sufficient to remove frequency fluctuations and drifts on time scales from a few milliseconds to a few seconds.

IV. SIGNAL CALIBRATION AND LOCKING SENSITIVITY

The first step in determining the frequency stability of our locked laser is to calibrate the error signal. We do this in three ways. First, we model the differential lineshape theoretically and predict the slope in the zero-crossing region for a given peak transmission voltage and differential signal gain. Alternatively, we can use our 0.1 ppm wavemeter, which has a frequency resolution of roughly 30 MHz. We simply monitor the frequency on the wavemeter as we manually tune the laser center frequency through the broad linear range of our differential signal, and relate the measured differential signal voltage to the observed frequency change. The third method uses a superimposed set of Fabry–Pérot fringes of known frequency on the differential signal allowing the time axis to be converted directly to frequency. These three methods give consistent results for the frequency calibration. The frequency conversion factor is dependent on the amplification of the differential signal. However, the noise in the locked signal is relatively unaffected by a broad range of amplification factors. In the current setup we use an amplification that results in a conversion ratio of 180 MHz/V.

Figure 4 shows a 100 s trace of the differential signal for both the locked and unlocked cases. The vertical frequency axis was derived from the observed voltage fluctuations as described above. One can see evidence of both the short-term frequency jitter on time scales of 10–100 ms and the longer-term frequency drift in the unlocked signal. The rms fluctuation of the locked signal here is 0.6 MHz (150 Hz bandwidth). We note that this quoted frequency stability assumes the absence of other sources of fluctuations at bandwidths beyond what this servo system corrects for. The baseline noise level over time scales of several seconds, measured by monitoring the error signal when we are well off-resonance, shows noise comparable to that in our "locked" configuration, indicating that we are at or near the limit of the frequency-locking resolution for the current configuration.

As mentioned earlier, another important feature of our frequency stabilization scheme is the wide capture range of the error signal. Referring to Fig. 1, we can see that the linear capture range extends for about 1000 MHz around the zero crossing. Adjusting the servo loop lock point electronically allows us to lock to either side of the Doppler-broadened output beam absorption profile in a straightforward way.

V. POTENTIAL SYSTEMATIC ERRORS AND ACCURACY LIMITATIONS

We now consider several potential systematic errors that can limit the accuracy of our frequency stabilization scheme. First, we note that any fluctuations in the differential error signal which are *not* associated with frequency changes will nevertheless produce a response from the servo system, effectively altering the frequency lock point. Such spurious changes to the error signal can occur due to slow drifts in the power level of one (but not both) laser beam component or from electronic drifts in the offset. To explore the typical size of such systematic drifts, we sampled the differential signal with the servo lock turned off but with the vapor cell heated to a typical operating point. By tuning well off the atomic resonance, we eliminate differential transmission changes due to actual laser frequency fluctuations. Using the error signal voltage-to-frequency calibration already established, we can express observed drifts in the differential signal in terms of effective frequency change. Typically, for 15–30 min after optimizing the optics, we see offset drifts equivalent to 5–10 MHz due either to mechanical settling of the optics or to slow temperature stabilization of the AOM. As shown in Fig. 5, after this initial period, we typically observed smaller and more random drifts, whose maximum drift rate corresponds to an effective frequency change of less than 1 MHz/min. As the typical frequency scans in our spectroscopy experiment require 4–5 s, such offset drifts (and corresponding frequency instability of the lock) do not cause significant errors in fitting spectra and extracting hyperfine splittings in our spectroscopy experiment. Clearly, further attention to the issue of drift of the differential offset would be required in an application demanding megahertz-level absolute frequency stability over time scales of minutes to hours.

A second concern is the effect of vapor cell temperature fluctuations on the lock point. Because the temperature affects both optical depth and linewidth, a small change in temperature would clearly affect the slope of the differential signal and thus slightly alter the calibration of the slope in terms of frequency. It is easy to see that if we choose as our lock point the frequency at the symmetric point (i.e., halfway between the two frequency-shifted absorption dips), a temperature change will not change the frequency of the lock point and the zero crossing of the differential signal is immune to such common-mode changes to the lineshapes. However, if we lock away from the center of resonance, the particular frequency corresponding to a given non-zero dif-

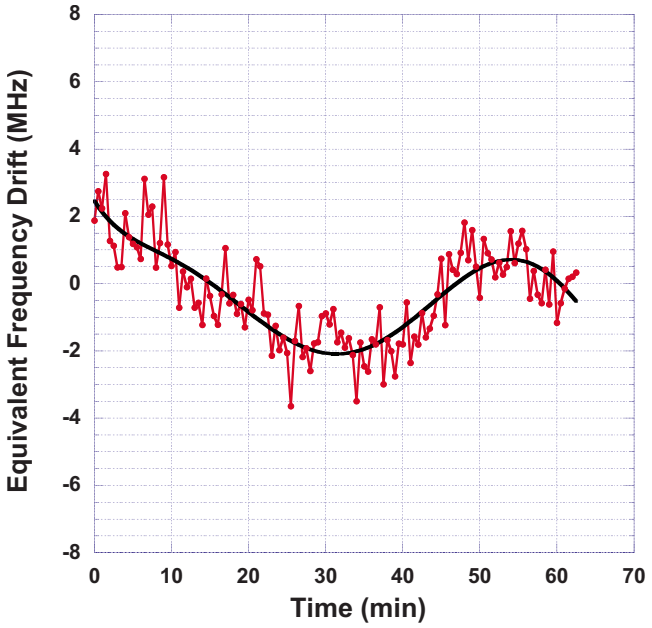

FIG. 5. A 1 h record of the differential signal taken with the laser unlocked and tuned well off-resonance. The differential signal is sampled every 30 s. If the system were in locked configuration, drifts such as this would lead to a systematic frequency change. In our spectroscopy experiment, an individual frequency scan requires only a few seconds so that typical long-term background drift rates such as shown here do not contribute significantly to noise or systematic errors in our observed spectra. The solid line is drawn simply to indicate the long-term trend.

ferential signal voltage becomes temperature dependent—and increasingly so the further from line center we choose the lock point. Our simulations indicate that for the extreme case of a 600 MHz off-center lock, we predict a lock-point dependence on temperature of roughly 0.7 MHz/°C. At present we do not employ any active temperature control to our oven, relying on the large thermal mass to limit observed temperature drifts to less than 1 °C/h. While this is a sufficient stability for our present needs, it is clear that even fairly modest temperature control would reduce this potential systematic error well below the level of the residual statistical fluctuations discussed here.

Our locking apparatus exists in a space absent of magnetic shielding. One can consider the effects of ambient field fluctuations on the performance of our stabilization scheme. Similar to the case of temperature fluctuations, the (unresolved) Zeeman splittings, the indium levels in the ambient ~0.5 G field cause slight lineshape broadening (as well as possible small asymmetries). Given the g-factors of the relevant states, we expect that Zeeman splittings cause additional broadening on the order of 1 MHz/G. Thus we would only be concerned about a potentially significant systematic error for the case of significant off-center locking, and then only if ambient field changes on the order of 1 G were to occur over the time scale of our spectral scans. Even for our unshielded apparatus, we have measured the typical field fluctuations to be one to two orders of magnitude smaller than this over this time scale. One can also consider the effects of the potentially larger 60 Hz ac magnetic fields induced by the 3–4 A (rms) current driving the heating elements of the oven. First, we note that these currents do not

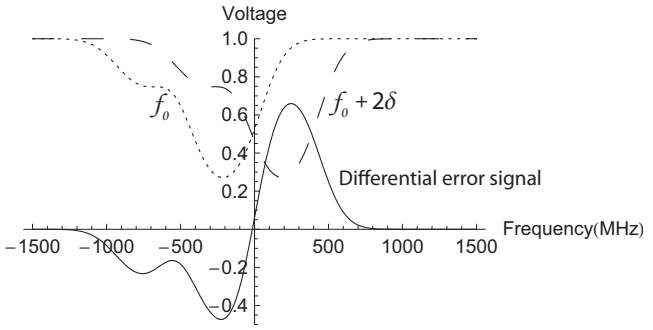

FIG. 6. (Color online) A theoretical plot of the composite thallium $6P_{1/2} \rightarrow 6P_{3/2}(F=1 \rightarrow F'=1,2)$ hyperfine transition, containing contributions from both ^{203}Tl and ^{205}Tl. Plotted below is the corresponding differential lineshape for frequency separation of 400 MHz.

change significantly over the time scale of our measurements. Also, the winding pattern of the heating elements is chosen to minimize the longitudinal field created by the current. Finally, as a test, we studied the correction signal when the system was locked upon switching on and off the oven currents. At the level of our statistical resolution, no systematic change was detected.

VI. APPLICATION OF METHOD TO NON-ELECTRIC-DIPOLE TRANSITION IN THALLIUM

As discussed, a desirable feature of our frequency stabilization technique is its broad applicability to other wavelengths and to transitions that are not necessarily of the electric dipole variety. As one considers other frequencies, particularly longer wavelengths in the red and infrared that are common in atomic spectroscopy, we note that the Doppler widths are correspondingly smaller. This is beneficial here, as the differential signal slope would be roughly three times larger for an analogous locking setup in the infrared region. At the same time, the optimal frequency separation of the two beams would be correspondingly smaller. This in turn eases the requirement on the performance of the AOM. Indeed, since typical AOM diffraction efficiency is far lower in the infrared as compared to the blue/visible, one can realistically only make use of the first-order diffracted peak.

As an indication of the versatility of our method, we used the technique presented here to stabilize a diode laser operating at 1283 nm. We locked the InGaAs diode laser to the "forbidden" M1/E2 $6P_{1/2} \rightarrow 6P_{3/2}$ transition in thallium using a quartz vapor cell heated to a comparable temperature to that of our indium cell (800 °C). Because thallium's vapor pressure at a given temperature is so much larger than that of indium, this temperature achieves a thallium optical depth for this non-E1 transition comparable to that of the indium 410 nm E1 transition. Performing a standard double-pass experiment with an AOM optimized for this wavelength (and using only first-order diffraction), we generate two beams separated in frequency by 400 MHz (Fig. 6). The thallium absorption lineshape is more complicated than that of indium due to the presence of both an unresolved hyperfine splitting and the existence of two isotopes whose small transition isotope shift here is also unresolved. Nevertheless, as shown in Fig. 7, this composite transmission spectrum still yields a

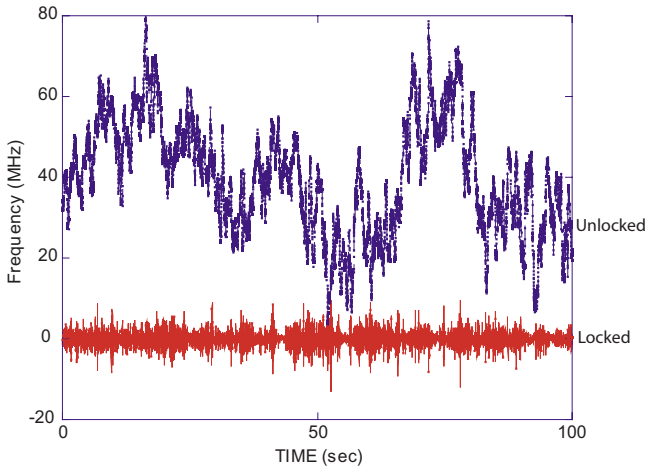

FIG. 7. Shows the laser lock to the M1/E2 "forbidden" transition in Tl, indicating a residual frequency fluctuation below 2 MHz. Effective locking bandwidth here is approximately 200 Hz.

differential signal with a significant linear feature, making laser locking feasible. Figure 7 shows a successful lock to this M1/E2 transition with a residual frequency noise of less than 2 MHz.

VII. CONCLUSION

In conclusion, we have developed a simple, robust, and widely applicable technique for frequency-stabilizing external cavity diode lasers based on multipass, frequency-shifted beams from an acousto-optic modulator. The method offers submegahertz frequency lock capability in a 150 Hz bandwidth, extending to time scales of several seconds, and features a broad capture range and the ability to tune the lock point over an ~1 GHz frequency range. We rely on all of these features in our current experiment, which involves two-step, two-color laser excitation of indium atoms. Further improvement in the performance of the system would result from an AOM operating at higher frequency and from improved mechanical stability. We are currently employing the same technique in a two-step excitation experiment in thallium where the initial transition at 378 nm will be driven by an UV external cavity laser diode system.

ACKNOWLEDGMENTS

This work was supported by the National Science Foundation Grant No. 055552. We gratefully acknowledge the contributions from Owen Simpson and Huajie Cao at various stages of this work.

[1] K. B. MacAdam, A. Steinbach, and C. Wieman, Am. J. Phys. **60**, 1098 (1992).
[2] Y. Yoshikawa, T. Umeki, T. Mukae, Y. Torii, and T. Kuga, Appl. Opt. **42**, 6645 (2003).
[3] K. L. Corwin, Z.-T. Lu, C. F. Hand, R. J. Epstein, and C. E. Wieman, Appl. Opt. **37**, 3295 (1998).
[4] J. A. Kerckhoff, C. D. Bruzewicz, R. Uhl, and P. K. Majumder, Rev. Sci. Instrum. **76**, 093108 (2005).
[5] E. D. Van Ooijen, G. Katgert, and P. Van Der Straten, Appl. Phys. A: Mater. Sci. Process. **79**, 57 (2004).
[6] C. I. Sukenik, H. C. Busch, and M. Shiddiq, Opt. Commun. **203**, 133 (2002).
[7] D. Stull and D. E. Gray, *American Institute of Physics Handbook* (McGraw-Hill, New York, 1972).