TUTORIAL/ARTICLE DIDACTIQUE

Laser-frequency stabilization using a lock-in amplifier

M. Weel and A. Kumarakrishnan

Abstract: We describe how a lock-in amplifier can be used in a feedback loop to lock the frequency of a laser to an atomic transition in ^{85}Rb. A simple physical explanation is presented to describe the shape of the feedback signal generated by the "lock in" and the dependence of this signal on the phase of the reference signal. We also present a mathematical model of the feedback signal. Using the output of the lock in as feedback, we show that the laser frequency can be stabilized to a few MHz. The experiment is relatively simple to set up and could easily be adapted as an undergraduate laboratory experiment.

PACS No.: 42.60Mi

Résumé : Nous montrons comment un amplificateur synchronisé peut servir dans une boucle de rétroaction pour verrouiller un laser sur une transition atomique dans le ^{85}Rb. Nous expliquons de façon simple la forme du signal de rétroaction généré par la synchronisation et la dépendance de ce signal sur la phase du signal de référence. Nous présentons également un modèle mathématique pour le signal de rétroaction. Utilisant la sortie de l'amplificateur synchronisé comme rétroaction, nous montrons que la fréquence du laser peut être stabilisée à quelques MHz près. L'expérience est relativement simple et peut facilement être adaptée à un laboratoire de 1er cycle.

[Traduit par la Rédaction]

1. Introduction

In many applications of spectroscopy, and experiments involving laser cooling and trapping of atoms [1], the experiment is critically dependent on the stability of the laser frequency. Because the typical linewidth of the principal resonance line in an alkali atom (used widely in trapping experiments) is ∼10 MHz, it is desirable that the uncertainty of the laser frequency be less then this value. Lasers with a linewidth of ∼1 MHz are commercially available and diode lasers with the required linewidth can be made in house [2, 3]. The experimental challenge, however, is to lock the laser frequency to a

Received 6 March 2002. Accepted 9 August 2002. Published on the NRC Research Press Web site at http://cjp.nrc.ca/ on 9 December 2002.

M. Weel[1] and A. Kumarakrishnan. Department of Physics, York University, 4700 Keele Street, Toronto, ON M3J 1P3, Canada.

[1]Corresponding author (e-mail: mweel@yorku.ca).

 DOI: 10.1139/P02-084

stable and reliable reference frequency. In typical laser-cooling experiments, a vapour cell containing a dilute gas of alkali atoms is used as an inexpensive frequency reference. Since it is convenient to maintain the cell at room temperature, the spectral lines are Doppler broadened (linewidth of ∼1 GHz). For this reason, well-known techniques such as Doppler-free spectroscopy [4] are used to isolate narrow spectral features with a linewidth comparable to the atomic linewidth (natural linewidth of the atomic transition). It is then possible to use one of several lock-in techniques to stabilize the laser frequency [5].

One of the easiest lock-in techniques to understand and implement involves frequency modulation using a lock-in amplifier. A tuning element within the laser is used to impose a small modulation on the laser frequency. The modulation causes frequency changes that are small compared to the atomic linewidth (about 10% in our case). A lock-in amplifier is then used to detect a signal at the modulated frequency and produce an error signal that can be used to correct changes in the laser frequency.

Many textbooks describe how a lock-in amplifier can be used for synchronous detection of signals [4, 6]. However, we have not found a simple and satisfactory explanation of how a lock-in amplifier generates the error signal in this application.

In this paper, we provide a simple physical explanation and describe the results based on a suitable mathematical model. Finally, we describe how the stability of the feedback loop can be studied. We have successfully used this technique to lock lasers used in our atom-trapping experiments. We note this technique is very powerful and that it is used widely to stabilize lasers for optical communication. It is possible to adapt this technique as an experiment in advanced undergraduate laboratory experiments. As in several leading institutions, students should be able to use this expertise to trap atoms as part of their undergraduate research experience.

The rest of the paper is organized as follows. In Sect. 2, we describe the basic elements of the feedback loop and explain how a lock-in amplifier produces an output that can be used to stabilize the laser frequency. In Sect. 3, we present a physical model that explains how the "lock in" produces the desired feedback signal. In Sect. 4, we describe the shape of the feedback signal using a simple mathematical model. In Sect. 5, we describe the experimental results and tests of the frequency stability.

2. Using a lock in in a feedback loop

In typical applications, the input to the lock-in amplifier is a signal from a photodiode that detects a laser beam. A modulation may be imposed on the input by chopping a laser beam at the modulation frequency. The modulation frequency or reference frequency is usually derived from an internal oscillator within the lock in. The input signal will, therefore, have a frequency component corresponding to the modulation frequency. The lock in multiplies the input with the reference frequency and integrates the product over several periods. The lock in is, therefore, sensitive to the component of the signal corresponding to the reference frequency. By choosing the reference to be higher than typical sources of noise, the signal-to-noise ratio can be greatly improved at the reference frequency.

For our application, which involves atom trapping, we modulate the input by modulating the laser frequency. Since we use diode lasers, this is achieved by modulating the laser–diode current. The modulated laser beam passes through a reference cell containing a dilute gas. A photodiode detects the absorption of the laser beam through the cell. When the laser is scanned across the Doppler-free line shape detected by the photodiode, the lock-in output has the desired shape for stabilizing the laser frequency (Figs. 1 and 3*b*). The output of the lock in is fed back to the laser current to complete the loop.

To maintain the laser frequency at the desired value, we require the feedback signal generated by the lock in to have a specific shape. If the laser frequency drifts above the desired value, we would like the feedback signal to be negative. This will change the laser current so as to correct the laser frequency. Similarly, if the frequency drifts below the desired value, we would like the feedback signal to be positive. In our case, the desired frequency corresponds to the peak of an atomic resonance. The atomic resonance line shape can be represented by a Lorentzian. As a result, if the feedback signal

Fig. 1. Block diagram of feedback loop. The input to the lock in is generated by subtracting the signals from the two photodiodes.

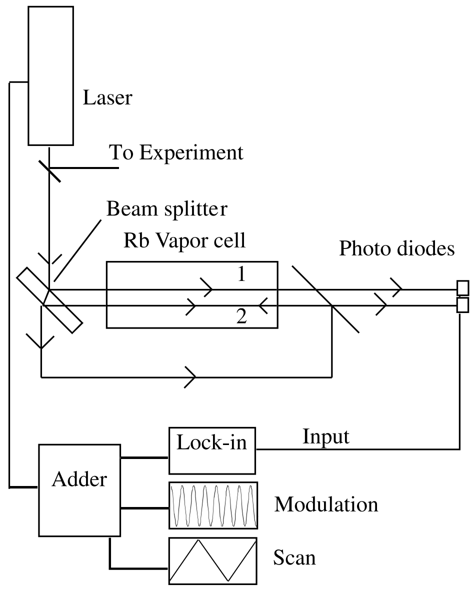

approximates the derivative of the absorption line shape, it will correct for small changes in the laser frequency.

3. Physical description

In this section we provide a physical description of the shape of the feedback signal. We first consider a frequency-modulated laser that is tuned to an atomic resonance in a vapour cell. The frequency modulation is provided by a sinusoidal voltage derived from the lock in. The Doppler-free absorption line shape is detected on a photodiode when the laser is scanned across resonance (Fig. 1). This signal shape is a broad Doppler-broadened dip with narrow spectral features in the middle. These features correspond to regions of decreased absorption. A second photodiode detects the Doppler-broadened absorption without these spectral features. These signals are subtracted to generate the Doppler-free absorption signal. The subtracted signal consists of spectral peaks that correspond to regions of decreased absorption. Hereafter, we refer to the subtracted signal as the absorption signal. To lock the laser frequency to an atomic resonance, the scan is centred on a Doppler-free absorption peak and the scan amplitude is zeroed. The feedback loop is engaged by sending the sum of the lock-in output and the modulation to the laser. The details of the experimental setup and the observed spectrum are described in Sect. 5, where we describe the experimental results.

We first consider a situation where the laser has drifted below resonance. The applied modulation causes the laser frequency to periodically shift closer and further from resonance. Therefore, when the laser frequency increases, the absorption signal increases, and when the laser frequency decreases the absorption signal decreases. We also note that the photodiode signal is in phase with the frequency modulation (Fig. 2*a*).

The lock in multiplies the input from the photodiode by a square wave, which has the same frequency as the modulation. The product of these two wave forms is then integrated over several cycles of modulation. From Fig. 2*a*, we can see that the result of the integration is positive.

When the laser is on resonance, any change in the laser frequency due to the modulation takes the frequency away from resonance and produces a decrease in the absorption signal. As a result the

Fig. 2. Absorption signal with frequency modulation (*a*) below resonance, (*b*) at resonance, and (*c*) above resonance. In each plot the top curve is the absorption signal, the bottom continuous curve is the modulation, and the broken curve is the product.

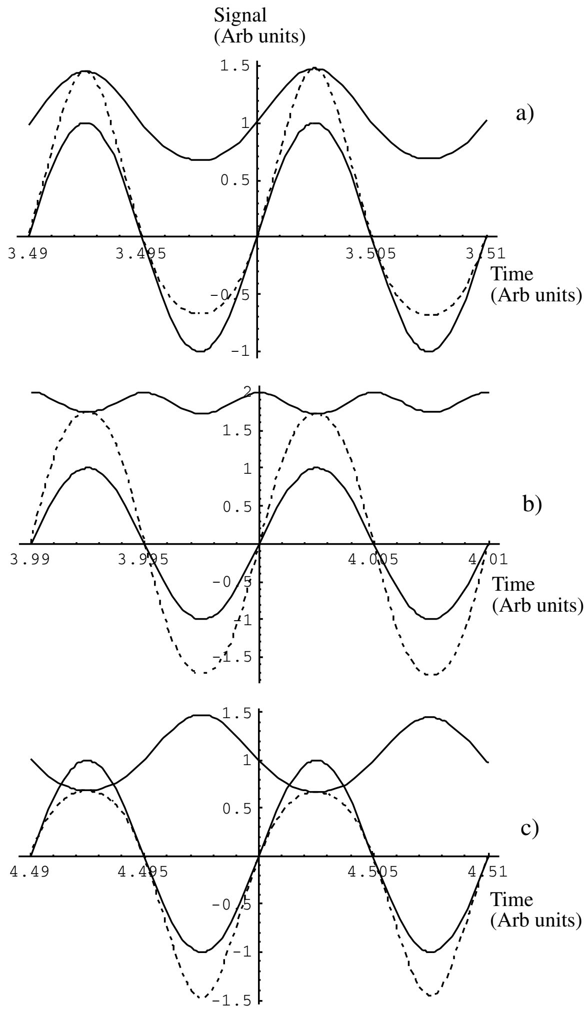

photodiode signal will have a frequency component twice the frequency of modulation (Fig. 2*b*). In this case the integrated signal averages to zero.

Finally, we consider the drift of the laser frequency above resonance. The photodiode signal once again has the same frequency as the modulation. However, the phase of the photodiode signal is 180° out of phase with respect to the modulation frequency. It is clear from the Fig. 2*c* that the integrated signal will be negative in this case.

When the frequency modulated laser is scanned across the absorption line shape without engaging the feedback, the output of the lock in will have the desired shape near the resonance. As a result, when the feedback is engaged the lock in can correct small changes in the laser frequency. The lock in detects the signal at the modulation frequency. Since the modulation frequency can be chosen to be higher than typical sources of noise (for example, mechanical vibrations, 120 Hz modulation due to room lights, or frequency components associated with pumps and motors) the technique is very effective and is widely used.

Fig. 3. (*a*) Absorption signal with frequency modulation during a scan period as a function of time (modelled by (1) and (2)). (*b*) Integrated signal of the absorption signal from Fig. 2*a* as a function of time (modelled by (3)); the zero crossing in (*b*) corresponds to the peak in (*a*).

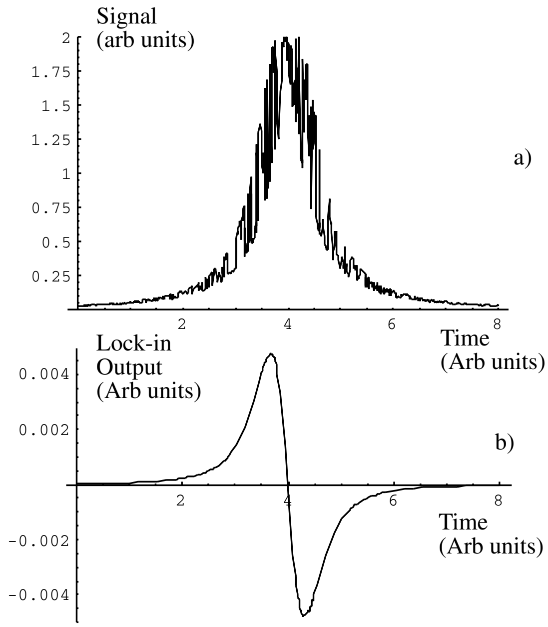

4. Mathematical model

The physical picture presented in the previous section can be described by a simple mathematical model. We model the input of the lock-in amplifier by describing the absorption signal of the laser beam through the reference cell as a function of frequency. The Doppler-free absorption spectrum generated using saturated absorption spectroscopy (see Sect. 5) is a Lorentzian function, $A(f)$, and can be described by

$$A = \frac{2\alpha}{(\alpha^2 + (\omega_0 - f)^2)} \tag{1}$$

where ω_0 is the resonance frequency, α is the width of the absorption line shape, and f is the laser frequency. In our experiment, the absorption line shape is detected by the photodiode. To observe the absorption spectrum, we vary the laser frequency as a function of time. The frequency will have a linear component (or scan) and a modulation that can be described by

$$f = at + b\sin(f_{\text{mod}}t) \tag{2}$$

Here, a is the slope of the scan, b is the amplitude of the modulation, t is time, and f_{mod} is the frequency of the modulation. A plot of the detected signal as a function of time can be generated using (1) and (2) and is shown in Fig. 3*a*.

Several interesting features of the plot shown in Fig. 3*a* appear when we look closely at various sections of the curve as shown in Figs. 2*a*–2*c*. In Fig. 2*a*, we see that for frequencies below resonance, the signal and the modulation are in phase. When we look at the signal at the resonant frequency, we see that it has twice the frequency of the oscillation, as shown in Fig. 2*b*. Finally, we see that above resonance, the signal is 180° out of phase with the modulation, as shown in Fig. 2*c*.

The output of the lock in is produced by multiplying the absorption signal and the modulation and integrating the product over several oscillations. In a lock-in amplifier, it is possible to vary the phase

Fig. 4. Lock-in output during a scan period as a function of time (*a*) for a $\phi = 90°$ and (*b*) for $\phi = 180°$.

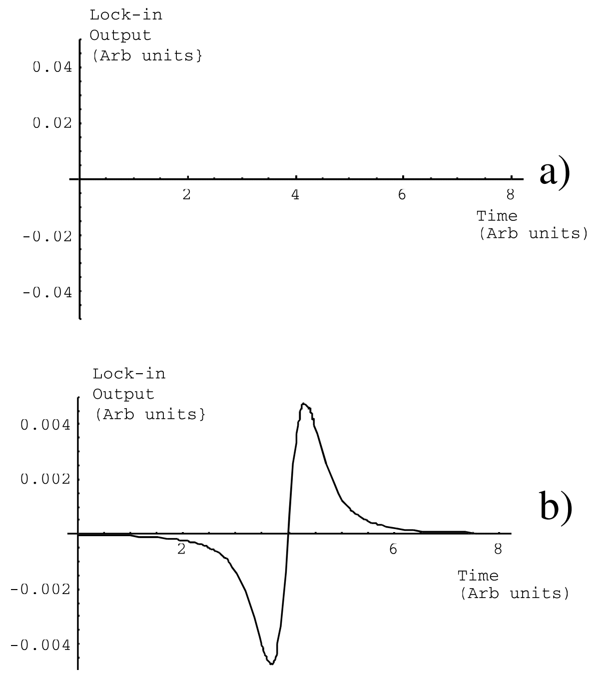

of the modulation signal. As a result the output of the lock in can be modelled as

$$\int_{t=T}^{t=T+N/f_{\text{mod}}} \frac{2\alpha}{\alpha^2 - (\omega_0 - at + \sin(f_{\text{mod}}t))^2} \sin(f_{\text{mod}}t + \phi)\,\mathrm{d}t \tag{3}$$

where $t = T$ is the start time of the integration, $t = T + N/f_{\text{mod}}$ corresponds to N modulation periods after the start time, ϕ is the phase difference between the signal and the modulation. We can plot (3) over a range of T. From Fig. 3*b*, we see that for the input in Fig. 3*a*, the output of the lock in will have the shape predicted by the physical model described in Sect. 3.

The output of the lock in is critically dependent on the phase difference between the two sine terms in (3). Figure 4*a* shows the output of the lock in predicted by the model for $\phi = 90°$.

The output is essentially zero under these conditions. Figure 4*b* shows the output when $\phi = 180°$. In this case the output is the negative of the desired output.

The output of the lock in is used as feedback to correct changes in the laser frequency. This is achieved by electronically adding the output of the lock-in amplifier described in (3) to the sum of the scan voltage and the modulation described by (2) and feeding the resultant signal to the laser current as shown in Fig. 1. Figure 5 shows the sum of the linear scan and the lock-in output.

It can be seen that the signals will add to give a total signal that is constant near resonance. If the signal is flat near resonance, this implies that change in the laser frequency due to the scan are exactly cancelled by changes in the feedback. As a result, the total signal sent to the laser is constant. Therefore, the laser frequency remains unchanged. If the scan amplitude is lowered, so that the entire scan is inside the flat region in Fig. 5, the laser frequency will remain on resonance for the entire scan period. If the scan amplitude is lowered to zero, the laser will remain on resonance for small variations in laser current, laser cavity length, or other factors, such as, temperature that may cause the unlocked laser frequency to drift.

It is interesting to note that the input signal can be multiplied by a variety of reference wave forms

Fig. 5. Sum of scan voltage and lock-in output as a function of time. Resonance corresponds to $t = 4$.

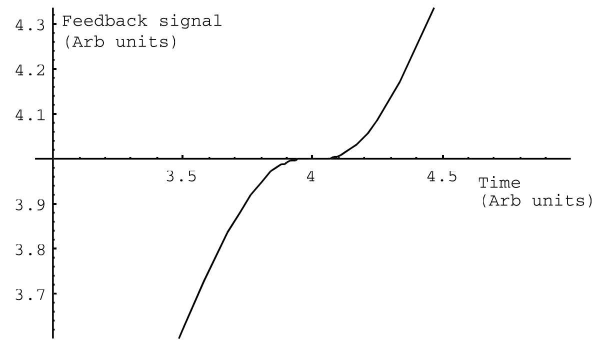

(sinusoidal, triangular, square, etc.) in commercially available lock-in amplifiers. We have compared the output signal generated using different reference wave forms of the same amplitude and period. This is done by replacing the sinusoidal component in the numerator of (3), with triangular and square wave forms (since the modulation remains sinusoidal, the denominator of (3) is unchanged). We have calculated the height of the peak of the output signal in Fig. 3*b*. We note that for a square wave, the output is 22% higher then for a sine wave. The triangular wave form produces an output 4% lower than for a sine wave. This suggests that the feedback signal can be produced with a smaller amplitude laser-frequency modulation with a square wave. A lower amplitude would lower the fluctuations in the laser frequency due to the modulation. This means that a square wave gives the highest signal-to-noise ratio.

5. Experimental setup

In this section, we describe the experimental setup, consisting of a saturated absorption spectrometer, experimental results, and tests of the frequency stability.

The laser used in the experiments is a grating-stabilized external cavity laser [2, 3] with a linewidth of ~1 MHz. The laser frequency can be adjusted by tilting the grating and by changing the diode current. To lock the laser to an atomic transition it is necessary to obtain a Doppler-free line shape using saturated absorption.

The experimental setup involves a Rb vapour cell at room temperature. The saturated absorption setup, shown in Fig. 1, uses a beam splitter to produce two parallel weak beams, labelled 1 and 2, which are reflections off the front and back surfaces of the beams splitter. These beams traverse the vapour cell and are incident on two photodiodes. When the laser is scanned, both photodiodes produce a temperature-dependent Doppler-broadened Gaussian absorption spectrum (full width half maximum of ~500 MHz). Most of the laser light passes through the beam splitter and is directed to counter propagate along the path of one of the weaker beams (beam 2). The absorption spectrum detected by the photodiode placed in beam 2 now has additional features due to the interaction of the atoms with the intense beam.

A sample of room temperature atoms contains a Maxwell–Boltzmann distribution of velocities. At a given frequency, the CW (continuous wave) laser will only interact with a narrow range of velocities. The absorption spectrum of beam 2 will not be affected when the two counter-propagating beams excite different velocity classes. When the laser is on resonance, however, both beams excite the same velocity class, namely, the zero-velocity class. In this case, the more intense beam saturates the atomic transition, so that at resonance the photodiode observes a decrease in the absorption signal. This results in a spectral feature called a Lamb dip. The width of this spectral feature has no contribution due to Doppler broadening. The width of the peak is determined only by the natural linewidth of the atomic

Fig. 6. Atomic structure of $5S_{1/2}$–$5P_{3/2}$ transition in ^{85}Rb; transition wavelength is ~780 nm.

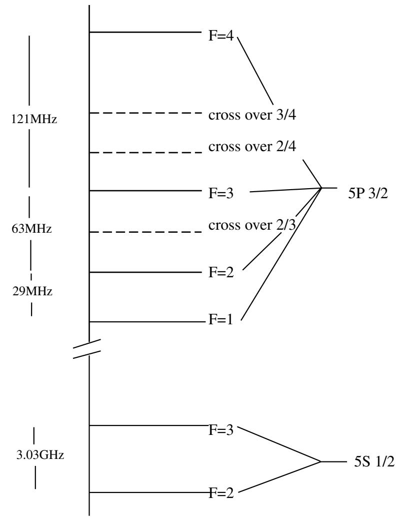

transition and power broadening due to the intensity of the beam. In Rb, the natural linewidth is ~6 MHz and the Lamb dip typically has a width of about 10 MHz.

Rb has the atomic structure shown in Fig. 6. At room temperature, a dilute gas containing two isotopes of Rb (^{85}Rb and ^{87}Rb) in the ratio ~72:28 are present in the cell. Therefore, both isotopes contribute to the observed spectrum. When the laser is scanned, four Doppler-broadened transitions are observed. These correspond to transitions from the two ground-state hyperfine levels of each isotope. We will now discuss the saturated absorption spectrum corresponding to absorption out of the ^{85}Rb $F = 3$ $5S^{1/2}$ ground state. Quantum mechanical selection rules permit transitions from the $F = 3$ ground state to $F' = 4$, 3, 2 excited states. The excited states fine-structure splitting is shown in Fig. 6.

The saturated absorption spectrum consists of a sequence of Lamb dips superimposed on the Doppler-broadened spectrum. The Lamb dips occur when the counter-propagating beams are resonant with each of the three transitions. Additional Lamb dips are produced when the two counter-propagating beams are each resonant with a different transition. These additional Lamb dips are called "cross-over resonances" and occur when the laser frequency (in the laboratory frame) is exactly halfway between a pair of atomic levels. Cross-over peaks occur exactly halfway between transitions because the beams are counter propagating and the atoms "see" the two beams Doppler shifted by the same amount but in opposite directions. The cross-over resonances are stronger than the Lamb dips because two velocity classes are addressed, that is, velocity classes corresponding to atoms moving along the direction of the intense beam, and those moving along the direction of the weak beam.

The second weak beam (beam 1 in Fig. 1) is used to provide just the Doppler-broadened absorption spectrum. The Doppler-free spectrum shown in Fig. 7 is produced by subtracting the photodiode signals due to beams 1 and 2.

6. Results

When the range of the frequency scan is lowered so that the laser scans over a single-absorption peak ($F' = 4$), we observe the signal shown in Fig. 8*a*.

Figure 8*b* shows the corresponding output of the lock-in amplifier. It is clear that the output has the desired shape described by the theoretical model (Fig. 3*b*). This output is used as feedback to correct the laser frequency.

Our discussion of the lock-in output in Sect. 3 assumed that the frequency modulation and the

Fig. 7. Observed Doppler-free absorption spectrum of 5S 1/2–5P3/2 transition in ^{85}Rb.

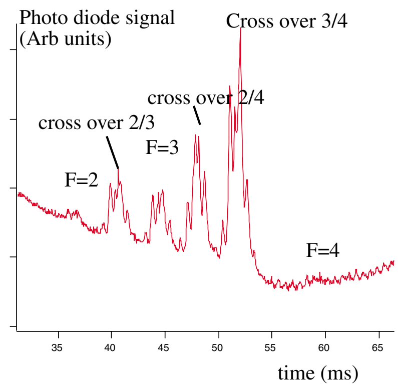

Fig. 8. (*a*) Observed saturated absorption spectrum of a single transition and (*b*) lock-in output corresponding to the signal in (*a*); the phase shift between the signal and modulation, $\phi = 0$.

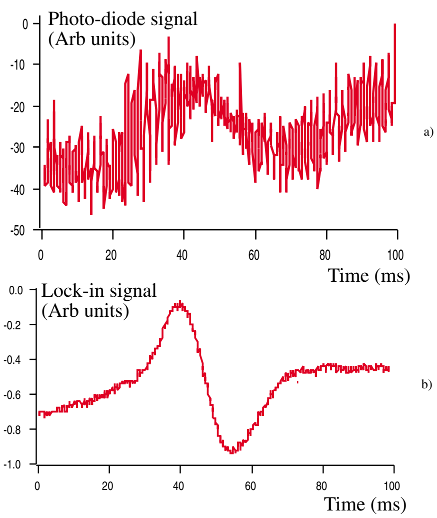

Fig. 9. Lock-in output as a function of phase shift between signal and modulation (*a*) $\phi = 90°$ and (*b*) 180°.

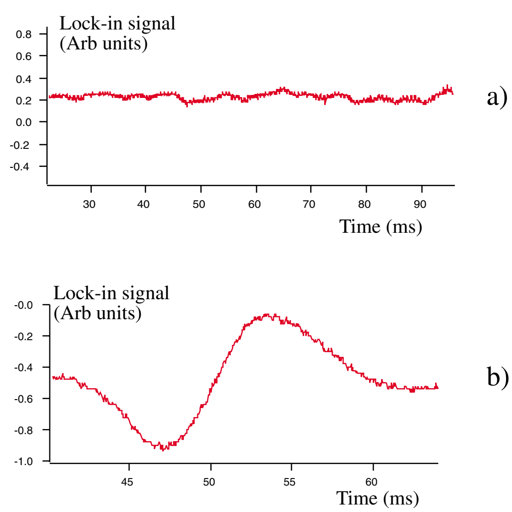

absorption signal are in phase. In Sect. 4, we discussed the dependence of the output on the phase. Figure 9*a* shows the variation of the lock-in output with the phase of the modulation shifted by 90° and Fig. 9*b* shows the output for a phase shift of 180°.

These are consistent with the theoretical predictions shown in Fig. 4. When the laser is first locked, it is convenient to find the desired phase by studying the dependence of the lock-in output on the phase of the modulation frequency. The phase of the modulation is varied until the lock-in output is a minimum. The phase is then shifted by plus or minus 90° so the signal is a maximum. One of these settings will correspond to the desired phase. At the other setting, the feedback signal will push the laser frequency away from resonance.

The known level splittings can be used to calibrate the horizontal scale of Fig. 7 in units of frequency. To determine the frequency jitter associated with the feedback loop, the scan amplitude was turned off, leaving the laser on the desired peak. The variation in the absorption signal at this frequency was then studied, with and without engaging the feedback loop. Without the feedback, the typical variation in the laser frequency over 10 min is about 10 MHz. With the feedback correcting the frequency the typical variation is about 2 MHz over the same period. Using the feedback loop discussed in this paper, it was possible to maintain the laser frequency on a resonance peak for a period of more than 1 h. Our stability is limited by the time constant of the lock in (∼10 ms). This stability can be improved by increasing the modulation frequency and decreasing the time constant if the tuning element of the laser has sufficient bandwidth.

An alternative way to show the effectiveness of the feedback loop is to scan the laser frequency with the feedback loop engaged. The correction produced by the lock in can then be observed. It is possible to observe that these two signals will sum to give no net contribution to the laser frequency as shown in Fig. 5 and discussed in Sect. 4.

7. Conclusions

We have presented a discussion of how a lock-in amplifier can be used to stabilize the laser frequency and lock it to an atomic transition. The feedback loop involved a commercial lock-in amplifier, a home-built phase-shifter circuit, two inexpensive function generators to produce the modulation and the scan, and a home-built electronic adder to sum the lock-in output modulation and scan signals. The total cost of these elements was about $1500. It is, therefore, possible to introduce this experiment into advanced student laboratory experiments that involve spectroscopy with tunable diode lasers.

An obvious disadvantage of this technique is that the frequency of the laser is modulated even when the laser frequency is locked. In atom-trapping experiments, the modulation can be detected in the fluorescence of the atom trap even for small modulation amplitudes. The modulation acts as an additional heating mechanism for the trapped atoms. To prevent this problem, it is possible to apply the modulation to a tuning element such as an acousto-optic modulator placed in the saturated absorption arm of the optical setup. In this way, the laser beam used for trapping atoms is not modulated. Alternatively, other methods such as FM spectroscopy [5] produce the desired feedback signal using an electro-optic modulator placed in the saturation absorption setup.

Acknowledgements

We acknowledge the contributions of A. Vorozcovs who wrote the software for automated data acquisition. AV was partially supported by a York University International scholarship. MW was partially supported by an undergraduate summer research scholarship. This research was supported by the Canada Foundation for Innovation, Ontario Innovation Trust, York University, and the Natural Sciences and Engineering Research Council of Canada.

References

1. W.D. Phillips. Rev. Mod. Phys. **70**, 721 (1998).
2. C. Wieman and L. Hollberg. Rev. Sci. Instrum. **62**, 1 (1991).
3. K.B. MacAdam, A. Steinback, and C. Wieman. Am. J. Phys. **60**, 1098 (1992).
4. W. Demtroder. Laser spectroscopy: basic concepts and instrumentation. Springer Verlag, Berlin. 1981.
5. M.D. Levenson. Introduction to nonlinear laser spectroscopy. Boston Academic Press, Boston. 1988.
6. A. Corney. Atomic and laser spectroscopy. Clarendon Press, Oxford, New York. 1977.