16th IEEE International Conference on Control Applications
Part of IEEE Multi-conference on Systems and Control
Singapore, 1-3 October 2007

Black-Box Identification of a Grating-Stabilized External-Cavity Diode Laser

W. Weyerman†, B. Neyenhuis*, J. Archibald*, M. Washburn*, D. Durfee*, S. Warnick†

†Information and Decision Algorithms Laboratories
Department of Computer Science, Brigham Young University, UT
http://idealabs.byu.edu

* Department of Physics and Astronomy, Brigham Young University, UT

***Abstract*— Diode lasers have many useful properties and have found a variety of uses including CD and DVD players, barcode scanners, laser surgery, water purification, quantum-key cryptography, spectroscopic sensing, etc. Nevertheless, their intrinsic linewidth, or the precision of their emitted wavelengths, is not good enough for many cutting-edge applications such as atom interferometry or high-performance atomic clocks. Using active feedback control we can narrow the linewidth of a diode laser, not allowing the frequency of emitted light to drift away from a reference value. Nevertheless, such feedback designs are challenging because of a lack of first principles models and difficult sensor dynamics. This paper describes our diode laser system and reports our results identifying the system using black-box techniques.**

I. Introduction

Diode lasers have some very unique qualities which have enabled many scientific and technological advances [7]. Compared to other lasers, they are inexpensive, compact, and efficient. They typically require little power and usually do not produce a lot of heat. Their wavelength can be quickly modulated. Diode lasers are also available at many different wavelengths, and a given diode can typically be tuned with temperature and optical feedback over several nanometers.

Despite these benefits, bare laser diodes do not have the stability and narrow linewidth necessary for many applications such as spectroscopy and laser cooling and manipulation of atoms. Placing the diode in an external cavity can reduce the linewidth of a diode laser, making it suitable for many more applications [7]. By additionally locking the laser to a stable reference cavity, the linewidth can be reduced even more. Some of the most stable lasers in the world have been created by actively locking an external-cavity diode laser (ECDL) to an ultra-high finesse optical cavity, [5], [8], [4]. In [4] it is noted that a much smaller linewidth is likely to be achieved through optimized adjustment of the servo amp. Typically, more attention is given to the components that the laser is built with than design of the controller, [5], [8].

We have constructed an external-cavity diode laser to use as the oscillator in a next-generation optical-frequency atomic clock. To achieve the extreme level of stability required for this application, we have mounted the laser in a heavy, mechanically isolated box with a large thermal mass, have developed high-stability, low noise current sources, etc. To cancel out environmental noise and drift, the laser is locked to an ultra-high finesse optical cavity, as mentioned above. Using this cavity, we generate an error signal which we can use for active feedback control of the laser to narrow the linewidth and keep the wavelength from drifting away from its reference value.

In this paper, we first describe our laser system by detailing the operation of the external cavity diode laser, discussing the sensor system that compares the laser's wavelength to a reference to generate an error signal, and then defining notation for the corresponding closed-loop system. We then describe our identification process for the system, which is necessarily black-box and closed-loop. In this process we identify two models for the laser system and characterize the frequency response of an additive noise disturbance.

II. System Description

Our laser system consists of two main components. First is the external cavity diode laser (ECDL) which consists of the laser diode and an external cavity which assists in stabilizing the wavelength of the light from the diode. The second component is the sensor. The wavelength of the light cannot be measured directly, so we use the Pound-Drever-Hall method as discussed in [1] to generate an error signal based on the deviations of the laser's wavelength from a reference wavelength.

A. External Cavity Diode Laser

A diode laser typically consists of a semiconducting junction with two reflecting facets. A current passing through the junction creates electron-hole pairs. When these pairs annihilate, the excess energy is often given off as a photon of light. Light with the correct wavelength passing through the junction can interact with the pairs, causing them to emit more light through stimulated emission. The light produced through stimulated emission is identical to the light

Please direct questions and comments to Sam Weyerman at wsweyerman@gmail.com or Sean Warnick at sean@cs.byu.edu.

which induced the stimulated emission, resulting in coherent amplification of photons.

The edges of the semiconducting junction in a diode laser are usually polished and coated with thin layers of dielectric material to make them reflective. These two reflective surfaces form an optical cavity in which light can oscillate. Light which is spontaneously (or randomly) emitted into one of the modes of the optical cavity can be repeatedly amplified by stimulated emission, resulting in a coherent laser beam. In many diode lasers the semiconductor is doped in such a way as to produce emission in a single transverse mode. In these lasers, the optical cavity can be treated like a one-dimensional oscillator, similar to a guitar string, with resonances at frequencies equal to a constant times the fundamental resonant frequency.

An "ideal" laser generates a beam of light whose oscillating electric field is described by a sine wave. In a real laser, the electric field contains a band of frequencies. We call the width of this band the spectral linewidth. Over short time scales, the linewidth is determined by fundamental properties of the laser, such as the size of the laser cavity and the reflectivity of the mirrors in the cavity. In many lasers, the optical cavity is made of highly reflective mirrors separated by many centimeters. In diode lasers, however, the reflective surfaces of the diode are typically less than a millimeter apart, and the reflective coatings are usually poorer than those on the mirrors used in most other types of lasers. This results in cavity modes which are not as sharply defined, exhibiting a broader linewidth. Furthermore, the semiconducting junction in a diode laser can typically amplify light in a large number of cavity modes, such that the laser will often output light which is a sum of different modes at different wavelengths.

In an ECDL scheme, one or more external mirrors or gratings is used to reflect light back into the laser diode. These mirrors form an external cavity which is longer than the intrinsic cavity of the laser diode. The ECDL configuration which we are using is known as the Littrow configuration. In the Littrow configuration a diffraction grating is placed in front of the diode laser at the Littrow angle — the angle at which the first diffraction order travels exactly back to the diode laser — as shown in Figure 1. A lens between the grating and the diode is used to make the mode of the returning beam match the mode of the beam leaving the diode. This has the effect of extending the size of the laser cavity, resulting in narrower, more sharply defined laser cavity modes. The zeroth diffraction order reflecting from the grating is used as the laser output.

Using a diffraction grating has the added benefit of improved mode selection. Because the angle of the diffracted beam depends on the wavelength of the light, only a particular wavelength will couple correctly back into the laser diode. The different modes of the external cavity will have different losses depending on how far their wavelength is from the wavelength which couples back to the diode correctly. As such, modes near this optimum wavelength will have less loss, resulting in higher light intensities in these modes. Due to the non-linear optical gain of the diode laser, modes with

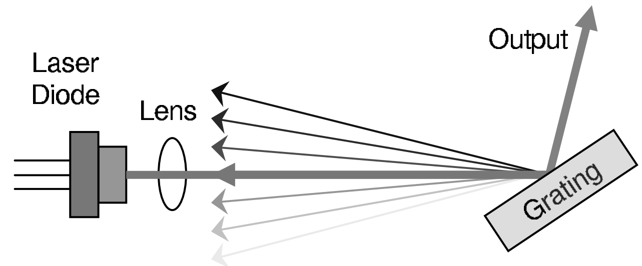

Fig. 1. The Littrow configuration. Changing the wavelength of the laser is accomplished by either rotating the grating, sliding the grating to change the length of the cavity, or changing the effective refractive index of the diode.

more intensity are more efficiently amplified. This results in mode competition which can cause the laser to generate light almost entirely in one cavity mode.

In spite of the external cavity diode laser's ability to generate a strong, sharply defined optical mode, thermal, electrical, and mechanical noise cause the color or wavelength of this mode to drift over time. This makes the ECDL unsuitable for applications such as atom interferometry and high performance atomic clocks. Nevertheless, if the wavelength of the laser could be controlled, then a simple servomechanism could be used to stabilize the wavelength.

There are three ways to change the wavelength of this laser system. First, the wavelength can be changed by rotating the grating to change the resonant cavity mode. When doing this, however, the laser makes abrupt changes in wavelength, known as "mode hops," which are undesirable because of their discontinuous nature.

Another way to change the wavelength of the system is to change the size of the cavity by translating the position of the grating relative to the laser diode. Nevertheless, this approach to effecting wavelength is also undesirable because the mass of the grating makes it difficult to move quickly.

A third way to effect the wavelength of the laser system is to impact the effective size of the optical cavity by changing the index of refraction of the diode. This can be accomplished by changing the electrical current powering the laser diode, which has the advantage of enabling both continuous and rapid changes to the resonant wavelength. This current impacts the index of refraction through both joule heating and by changing the density of electron-hole pairs. First, the more current supplied to the device raises its temperature, and since the index of refraction depends on temperature, the wavelength of the modes will increase as the current increases. This effect is particularly important at low frequencies, since at higher frequencies the thermal mass of the diode keeps the temperature from changing significantly over the short period of the modulation. In addition to heating the diode, however, changing the current also effects the index of refraction by changing the density of electron-hole pairs in the diode. This effect decreases the resonant wavelength as current increases. Although this process is not limited by the thermal mass of the diode, the capacitance of the diode reduces the amplitude of the modulation current at higher frequencies. These two effects operate at different time constants. The combination of them, which pulls the resonant wavelength in opposite directions, results in a zero in the laser's response, [3]. Nevertheless, this mechanism for actuating the diode laser enables the possibility of using feedback to control the system and stabilize its drifting wavelength.

B. The Pound-Drever-Hall Mechanism

Although the diode current is an effective actuator for controlling the wavelength of the laser, detecting the laser's wavelength in real time to drive a feedback controller is a challenging issue. One way to effectively accomplish this is through the use of a Pound-Drever-Hall mechanism [2], [3]. This mechanism receives the light from the laser as an input and effectively compares it's wavelength with a desired wavelength to generate an error signal. Although in many control systems a comparison with a reference command is a relatively simple issue, for this system there are substantial challenges associated with the task of generating this type of error signal.

In particular, the Pound-Drever-Hall mechanism behaves linearly only for a tiny band of wavelengths near the reference command. Outside of this region, the error signal goes to zero even though the laser wavelength is not close to the reference command. This nonlinearity is especially pernicious for a feedback system because it represents a loss of observability that effectively "confuses" a feedback controller into thinking things are fine when, in fact, they are not. To understand why the mechanism behaves this way, we must explore the details of how it works.

A diagram of our setup is shown in Figure 2(a). The light from our laser is first sent through an optical isolator. This device only allows light to pass one way, preventing stray light from reflecting back to the laser where it could pull the wavelength or cause instability. Part of the laser beam is then split off with a beamsplitter and sent through an electro-optic modulator (EOM). The EOM modulates the phase of the light and puts frequency sidebands on the nearly-sinusoidal oscillating electric field of the laser beam. The modulation frequency and phase is determined by an electrical signal generated by an rf synthesizer which powers the EOM.

The modulated light passes through a polarizing beamsplitter and strikes the front mirror of an ultra-high finesse optical cavity. This cavity, not to be confused with the laser cavity made up of the laser and grating discussed above, serves as a wavelength reference with which we can measure small changes in the laser's wavelength. The laser has the desired wavelength when it matches, or resonates with, that of this cavity. The size, divergence, direction, and location of the beam is adjusted with mirrors and lenses (not shown in the diagram) to couple into one of the transverse modes of the optical cavity. A birefringent quarter waveplate between the beamsplitter and the cavity causes the light reflected from the cavity back to the beamsplitter to have a polarization which is 90 degrees from the polarization of the incoming light. This

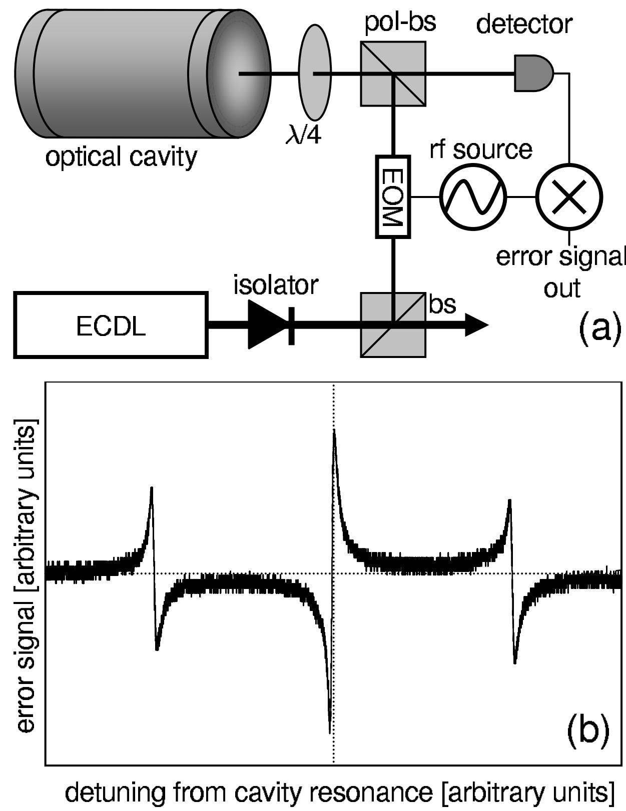

Fig. 2. The Pound-Drever-Hall method of generating an error signal between a laser's emission and a desired wavelength, characterized by a particular optical cavity. In (a) a diagram of our setup is shown. In (b) the magnitude of the error signal generated by our setup is shown as the wavelength of the laser is swept through a cavity resonance. The wavelength of the laser matches the cavity resonance in the center of the plot where the magnitude crosses zero; this is where the laser's output matches the reference. There is a small region centered around this point where the error signal behaves linearly. The slope and width of this section is determined by the fidelity of the cavity: a higher fidelity cavity creates a steeper and more narrow region. Just beyond these wavelengths, error rapidly approaches zero, causing the system to lose observability and effectively go to sleep. If the laser's wavelength is far enough from equilibrium, the error switches sign and the controller will drive the laser away from equilibrium.

causes the light reflected from the cavity into the polarizing beamsplitter to be sent to a photodiode light detector.

If the laser has the desired wavelength, constructive and destructive interference results in a dip in the reflected carrier intensity. The sidebands will be reflected, resulting in a beat note on the detector at twice the modulation frequency. No beat note will occur at the modulation frequency. If the laser is a tiny bit off of the desired wavelength, however, a beatnote will be generated on the detector at the modulation frequency. The amplitude of the beating is a function of the magnitude of the error, or detuning, of the laser's wavelength from the cavity resonance, while the phase of the beating depends on the sign of the error.

The signal from the detector is mixed with the signal used to drive the EOM and then low-pass filtered. This produces an error signal which is zero when the laser is at the desired wavelength, positive if it is just off of the desired wavelength on one side of the cavity mode, and negative on the other. An example of what the magnitude of the error signal might look like as the wavelength of the laser is swept through the cavity resonance is shown in Figure 2. There is a very steep linear regime of the error signal, shown as the steep vertical line in the center of the figure, the width of which is defined by the linewidth of the optical cavity. Outside of this region the Pound-Drever-Hall mechanism behaves nonlinearly. If the laser's wavelength is detuned by more than the cavity linewidth, the "error" signal quickly approaches zero. At a detuning equal to the frequency at which the laser is modulated (32 MHz in our case) the error signal changes sign. If the laser is far enough from the cavity resonance, this will actually cause the controller to drive the system away from the desired wavelength. Beyond these points, the error response of the Pound-Drever-Hall mechanism asymptotically approaches zero until the laser's wavelength approaches the next cavity resonance (roughly 1 GHz away for our cavity). Of particular note is the fact that the error signal quickly approaches zero as true deviations in wavelength begin to leave the linear regime. Once this happens, the system will effectively lose observability since the sensor will report zero error from the desired wavelength when, in fact, the error may be significant.

The steepness of the response in the linear regime derives from the fact that the cavity modes of the ultra-high finesse cavity are very narrow, so that a large error signal will be produced when the laser's wavelength drifts by even a tiny amount. Our best cavity has resonances only 1 kHz wide. It is made of ultra-low expansion (ULE) quartz mirrors contact-bonded to a ULE quartz spacer. This cavity is mounted into a small vacuum chamber to reduce drifts due to changes in the index of refraction of the air between the mirrors. The vacuum chamber is mounted inside of a heavy aluminum box which is lined with leaded foam and mechanically isolated from our optics table. Although such a high fidelity system yields tremendous sensitivity to even tiny errors, the width of the linear regime where the senor works also becomes very small. This can be a problem if nominal system noise perturbs the wavelength outside the linear regime because the resulting loss of observability will prevent a feedback controller from being able to bring the wavelength back to the desired reference value. As a result, we have used a lower fidelity cavity with approximately 10 kHz wide resonances for the studies discussed here. The somewhat lower finesse of this cavity widens the linear region.

C. The Closed-Loop Laser System

A block diagram of the feedback system is given in Figure 3. We will specify five systems, P, P_L, P_S, W_0, and K, with $P = P_S P_L$. P_L is the ECDL, P_S is the Pound-Drever-Hall frequency sensor, as described above. The sensor consists of all the components the light enters after being reflected by the first beamsplitter in Figure 2(a). The system W_0 shapes the noise entering the system and is determined experimentally. The controller is denoted as K.

Several signals are also represented in Figure 3. $\lambda(t)$ is the wavelength of the light exiting the laser and is represented

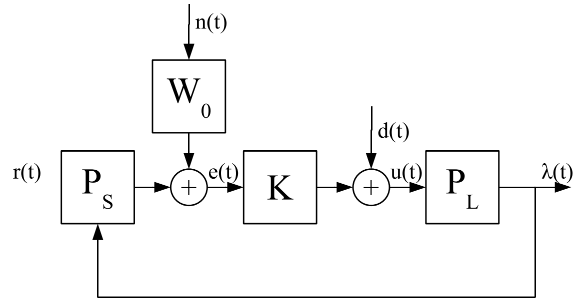

Fig. 3. Diagram of the closed loop system

in Figure 2(a) as the light that is not reflected by the first beamsplitter. We want this signal to stay as close as possible to the reference wavelength, $r(t)$, which is determined by the length of the optical cavity in the sensor. The signal $e(t)$ is the error as measured by the sensor. The signal $n(t)$ is white noise entering the system, $d(t)$ is an artificial disturbance signal intentionally added by us for system identification, and $u(t)$ is the input to the laser. The signals that we can measure are $d(t)$, $u(t)$, and $e(t)$.

III. IDENTIFICATION OF THE LASER SYSTEM

One of the aspects of this system that makes identification difficult is the absence of first principle models for each component that we can easily interconnect and parameterize to create a complete model. Nevertheless, effective controller design demands some understanding and model of the system to guide its development. Thus we turn to black-box identification methods to generate a coarse, linear time invariant control-oriented model of the system. This model will be described in two pieces: the ECDL/PDH system and the noise model.

A. Identification of the ECDL/PDH system

Black-box identification methods are challenging for the ECDL/PDH system because the Pound-Drever-Hall detection technique, which effectively compares the system output with a desired response and generates an error (see Figure 3), operates linearly in only a very narrow region near equilibrium. Outside this region, the error appears to go to zero when, in fact, it remains decidedly non-zero, and the system effectively loses observability. Moreover, since the amplitude of typical noise in the system easily perturbs it outside this linear regime, it is essential to use feedback to attenuate noise and control the system when collecting measurements. As a result, closed-loop identification is essential, even though we have no models to guide the design of this initial stabilizing controller. Nevertheless, through workbench trial-and-error, we obtained such a controller and were able to subsequently collect data for system identification.

The idea behind our identification experiments was to excite the stabilized, closed-loop system with sinusoids of different frequencies and measure the resulting magnitude and phase of the response. We would then fit these points in the frequency domain with the response of a rational transfer

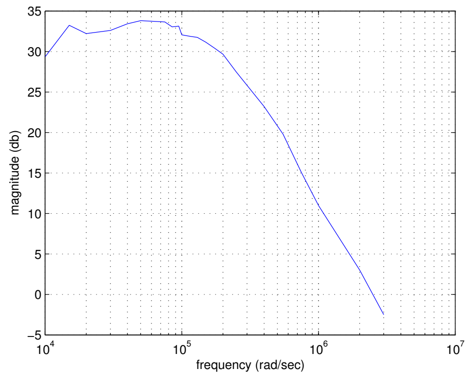

Fig. 4. Magnitude measurements from a single frequency scan to characterize the laser system. At about 10^4 rad/sec we notice some dynamics where the magnitude dips and comes back up. After 10^5 rad/sec the magnitude begins to roll off at a rate of 40 dB/dec.

function, and use this transfer function as our model of the system.

We began by selecting roughly twenty sample frequencies ranging from 1KHz to 3MHz. For each frequency, ω, we introduced our sinusoidal disturbance $d(t)$ as an additive perturbation on the current signal, $u(t)$, driving the external-cavity diode laser. We then measured the resulting error signal, $e(t)$, produced by the Pound-Drever-Hall detection mechanism. Using a Fast Fourier Transform on these signals, we were thus able to measure $|U(j\omega)|$ and $|E(j\omega)|$, and calculate $|P_0(j\omega)| = \frac{|E(j\omega)|}{|U(j\omega)|}$. To determine the phase of the laser system we measured the time delay of the sinusoid, $\delta(\omega)$ from u to e. We then calculated the phase in degrees as $\angle P_0(j\omega) = 360\omega * \delta(\omega)$. Having completed a scan over the entire range of frequencies, we then repeated the experiments 25 times to ensure we captured sufficient data to make meaningful fits and reduce the impact of experimental error.

Figure 4 shows the measurements obtained for one of the 25 scans. Noting that the system rolls off at 40db/dec, the simplest model to fit the data would be a second order system. Fitting such a system, we obtain

$$P_2 = \frac{3.278e12}{(s + 2.53e5)^2}$$

and observe its fit to all of the scan data in Figure 5(a).

To obtain a higher fidelity model, we note in Figure 4 the slight dip and hump around 10-20 kilorad/sec and 50-60 kilorad/sec, respectively. Using a fourth order model to capture this feature, we find

$$P_4 = \frac{1.071e13(s^2 + 3.197e4s + 3.044e8)}{(s + 1.674e6)(s + 1.105e5)(s + 3.228e4)(s + 9931)}$$

and observe its fit to all of the scan data in Figure 5(b).

These models appear to capture much of the dynamic information about the laser system obtained through our black-box identification experiments. In particular, however, we note that the quality of fit of both models appears to deteriorate significantly at frequencies higher than 10^6 rad/sec. The second order model exhibits strong deviations

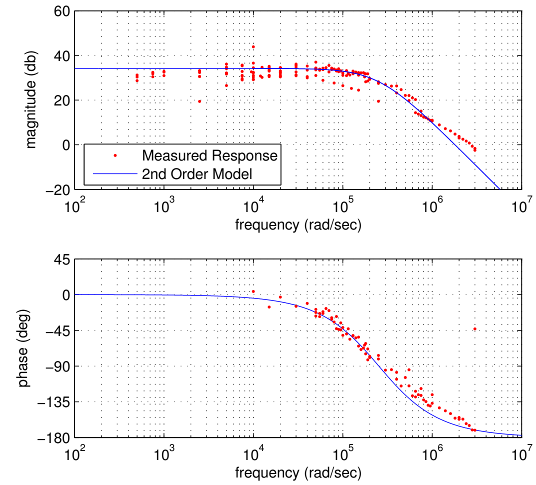

(a)

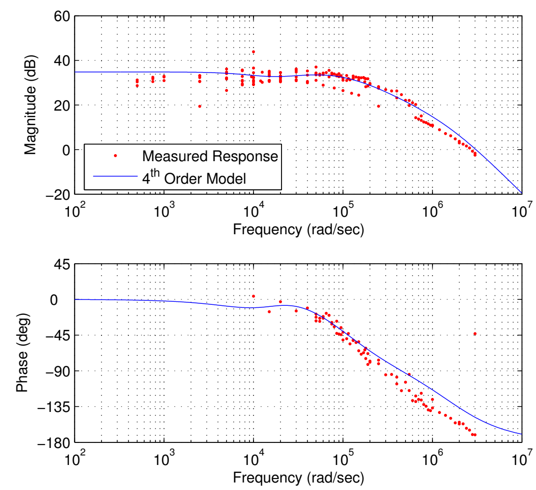

(b)

Fig. 5. Open loop models of the ECDL/PDH system plotted against the measured response of the laser from u to e. In (a) is the 2nd order model, and in (b) is the 4th order model.

in it's magnitude response, while the fourth order model exhibits strong deviations in its phase response.

B. Identification of the Noise Model

Considering discrepancies in the model to be explained by additive noise on the error signal generated by the Pound-Drever-Hall system, we developed a frequency weight model of this discrepancy. Given a stabilizing controller, the idea was that any deviation from equilibrium would be explained by this model. Thus, measuring the error signal when the closed-loop system should be in equilibrium allows us to factor out the impact of the controller, K, and the laser system, P, yielding $W_0 = (1 + PK)\frac{E}{N}$. Here we assume that n is unit intensity white noise. We repeated this process 30 times to compare the results of a single experiment with the experimental average. Figure 6 demonstrates the results of these measurements, along with the frequency response

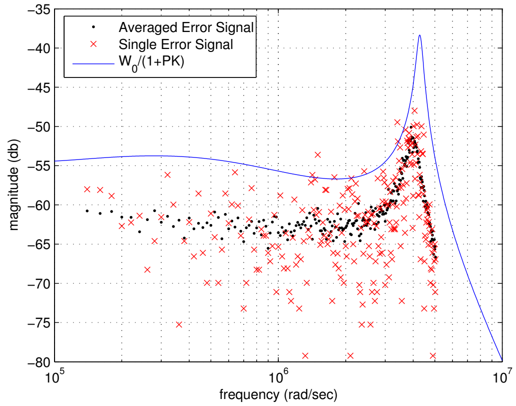

Fig. 6. Frequency weight on noise, W_0, chosen to conservatively cover experimental values.

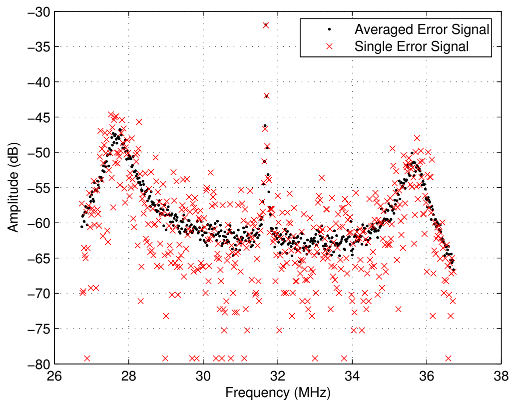

Fig. 7. FFT of the power spectrum of the signal measured at the photodetector while the laser is locked. Note the center spike corresponds to our modulation frequency of about 32 MHz.

of a conservative bound we used as our noise model, given by

$$W_0 = \frac{9.107e016}{s^3 + 2872e5s^2 + 1.84e13s}.$$

Figure 7 demonstrates the corresponding Pound-Drever-Hall detector (output) signal for one experiment, as well as for the ensemble average over all 30 experiments. Note the resonant spike corresponding to the equilibrium frequency of the laser locked to the high fidelity optical cavity of the Pound-Drever-Hall mechanism. Equipped with these models of the laser system and its corresponding frequency-weighted noise distribution, we are now prepared to begin systematic validation experiments and controller design, which is explored in [6].

IV. CONCLUSION

In this study we have explored the identification of a grating-stabilized external-cavity diode laser using a Pound-Drever-Hall mechanism to generate an error signal comparing the actual laser output with a desired reference wavelength. We used black-box identification techniques due to the complexity of our system and the absence of first-principle models. The identification process was necessarily closed-loop because of nonlinear effects associated with the Pound-Drever-Hall mechanism as an error sensor. Nevertheless, we obtained two linear models of the laser system that captured the prominent dynamics up to about 10^6 rad/sec. Moreover, we characterized the frequency response of the resulting model discrepancy to use as a noise model to help guide feedback design. Having identified the system, we are prepared to systematically design experiments to validate our models and design controllers to narrow the linewidth and stabilize the drifting wavelength of the laser.

REFERENCES

[1] E.D. Black. An introduction to Pound-Drever-Hall laser frequency stabilization. *Am. J. Phys.*, 69(1):79–87, January 2001.

[2] 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. *App. Phys. B*, 31:97–105, 1983.

[3] Richard W. Fox, Chris W. Oates, and Leo W. Hollberg. *Stabilizing diode lasers to high-finesse cavities*, volume 40 of *Experimental Methods in the Physical Sciences*, chapter 1, pages 1–46. Elsevier Science, San Diego, California, 2002.

[4] A. Schoof, J. Grünert, S. Ritter, and A. Hemmerich. Reducing the linewidth of a diode laser below 30 Hz by stabilization to a reference cavity with a finesse above 10^5. *Opt. Lett.*, 26(20):1562–1564, October 2001.

[5] H. Stoehr, F. Mensing, J. Helmcke, and U. Sterr. Diode laser with 1 Hz linewidth. *Opt. Lett.*, 31:736–738, March 2006.

[6] W. Weyerman, B. Neyenhuis, J. Archibald, M. Washburn, D. Durfee, and S. Warnick. Identification and control of a grating-stabilized external-cavity diode laser. *accepted for publication in IEEE Transcations on Control Systems Technology*, 2007.

[7] Carl E. Wieman and Leo Hollberg. Using diode lasers for atomic physics. *Rev. Sci. Instrum.*, 62(1):1–20, January 1991.

[8] B. C. Young, F. C. Cruz, W. M. Itano, and J. C. Bergquist. Visible lasers with subhertz linewidths. *Phys. Rev. Lett.*, 82(19):3799–3802, May 1999.