LASER INTERFEROMETER GRAVITATIONAL WAVE OBSERVATORY
- LIGO -
CALIFORNIA INSTITUTE OF TECHNOLOGY
MASSACHUSETTS INSTITUTE OF TECHNOLOGY

Document Type **LIGO-T970051-00 -** **R** 2-6-97

NPRO frequency stabilization

Rich Abbott, James Mason, and Rick Savage

Distribution of this draft:

This is an internal working note
of the LIGO Project.

California Institute of Technology
LIGO Project - MS 51-33
Pasadena CA 91125
Phone (818) 395-2129
Fax (818) 304-9834
E-mail: info@ligo.caltech.edu

Massachusetts Institute of Technology
LIGO Project - MS 20B-145
Cambridge, MA 01239
Phone (617) 253-4824
Fax (617) 253-7014
E-mail: info@ligo.mit.edu

WWW: http://www.ligo.caltech.edu/

1 ABSTRACT

As a first step in the process in switching to Nd:YAG lasers in LIGO, frequency stabilization techniques were applied to a Lightwave model 126 Nd:YAG 1064 nm 700 mW laser. This has two main motivations. First is that the proposed 10 W laser under development for LIGO by Lightwave uses the Model 126 laser as a master oscillator in a MOPA (master oscillator/power amplifier) configuration, and as such will be the point at which frequency corrections are applied in the 10 W version. The second motivation is simply to develop a 1064 nm light source for various experiments while the 10 W version is still being developed, such as the PNI IR conversion.

In this document, a description of the system is outlined, the design of the servos is detailed, and a procedure for obtaining lock using a fixed length reference cavity is described. This system can readily provide stabilization to the level of 10 mHz/$\sqrt{Hz}$ from 10 Hz to ~10 kHz. Lock is also quite robust, with lock being kept for time periods of at least 24 hours on a regular basis.

2 CONCEPTUAL DESIGN

2.1. Outline of an interferometer

LIGO proposes to measure strains due to gravitational radiation on the order of $10^{-23}/\sqrt{Hz}$. An interferometer is to be used by essentially comparing the phase history of light down one arm of the interferometer with the phase history of the light that travelled down the second arm of the interferometer. If there is any mismatch to the lengths of the arms, it can be seen that phase noise, or equivalently, frequency noise, would be a source of noise in the signal, by the relation $\delta f \ll (\delta l)/l \cdot f \cdot (1 - CMRR)$, CMRR being the common mode rejection ratio. In order for LIGO to measure the proposed strains, frequency fluctuations will need to be kept below 10^{-7} Hz/$\sqrt{Hz}$in the interferometer in the bandwidth of interest, assuming a CMRR of 99%. Since the proposed Nd:YAG laser has a frequency noise level typically about 100 Hz/$\sqrt{Hz}$ at 100 Hz, it's obvious that some frequency stabilization will need to be done to the laser. This level of suppression is a bit beyond the abilities of a single control loop, so one way to do this is to provide suppression in stages. The first stage of this is called the PSL (Pre-stabilized laser). This supplies approximately 80 dB at 100 Hz. Next is the mode cleaner, after which the laser light should be about 10^{-5}Hz/$\sqrt{Hz}$. The common mode servo of the interferometer locks the laser light to the interferometer, and will provide the additional attenuation.

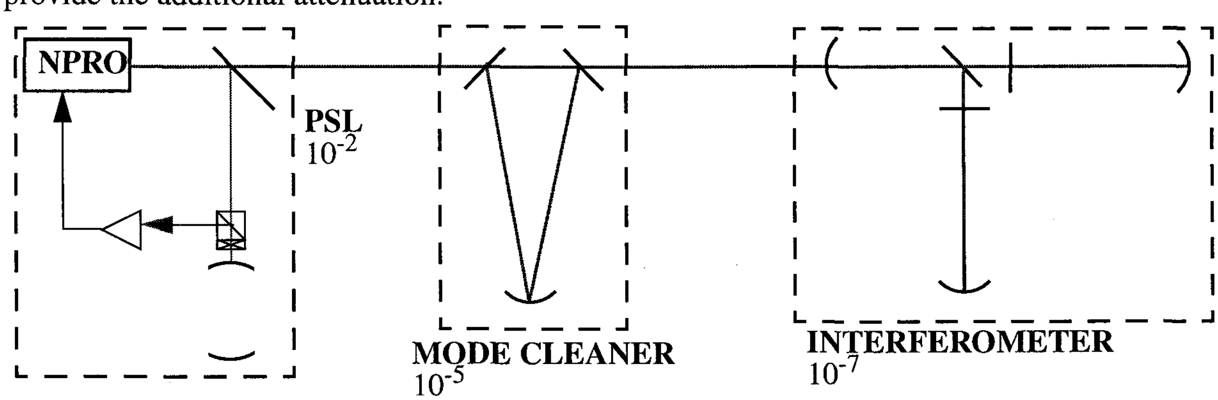

The above diagram gives a rough outline of how a LIGO-like system could be laid out. Not shown in the above diagram are feedback paths from the interferometer and the mode cleaner back to the PSL. The details of how the servo topology for frequency noise in LIGO will be laid out is yet to be finalized.

2.2. Requirements

For this first version of the NPRO PSL, refer to document "NPRO-PSL Design Requirements" (LIGO T960082-00-D) for requirements. Initial requirements of 1 mHz/$\sqrt{Hz}$ were relaxed while testing was going on to 10 mHz/$\sqrt{Hz}$ from approximately 10 Hz to 10 kHz. This calls for a servo gain of at least 46 dB at 10 kHz, and a gain of about 106 dB at 10 Hz. Also two feedaround paths were to be supplied, to allow direct frequency control to be applied from the mode cleaner and the interferometer. The reference cavity to be used would be a fixed length cavity made of two mirrors optically contacted to a ULE fused silica spacer. In order to have some ability to shift the frequency of the output beam from that of the reference cavity, a double-passed AOM would be used in the PSL stabilization path to shift the frequency of the light $2f_{AOM}$ +/- f_{Tuning}, where f_{Tuning} is +/- 5 MHz, determined by the requirement to keep the laser locked to both the reference and some external cavity.

3 SYSTEM LAYOUT

The proposed topology for frequency stabilization of the PSL is specified in the document, "NPRO-PSL Conceptual Design" (LIGO T960089-00-D).

The Model 126 NPRO laser comes with two frequency correction actuators built in. The "slow" controller consists of a thermoelectric cooler, which sets the temperature of the NPRO crystal. This has two effects, to change the physical size of the resonator cavity and to change the index of refraction of the crystal. Lightwave reports that the change Δn of the crystal has a much larger effect than the change in the size Δl. This input has a usable bandwidth of about 0.1 to 1 Hz. The gain is approximately 3-4 GHz/V, with a range of +/- 10 V. The "fast" controller consists of a piezo bonded to the resonator crystal. Voltage applied to it stresses the crystal, inducing changes in frequency. The bandwidth of this path is nominally limited by the internal resonances of the piezo, which begin around 200 kHz. The gain here is approximately 4 MHz/V, with a range of +/- 50 V. Both inputs give a positive change in frequency for a positive voltage. The free running frequency noise of the laser is approximately ($2x10^4$/f) Hz/$\sqrt{Hz}$ from 10 to 10 kHz.

The reference cavity is a ULE fused silica spacer, 20 cm in length, with mirrors optically contacted to each end. The mirror transmissions are 300 ppm, with losses < 30 ppm. This gives a finesse of ~ 10000, and a bandwidth of about 75 kHz. The temperature induced resonant frequency change is ~ 150 MHz/° C. The light is to be locked to this cavity using the Pound-Drever-Hall reflection locking technique, using sidebands at 12.3 MHz. The cavity is suspended in a vacuum chamber at 10^{-7} torr by two loops of wire from a 3 layer seismic isolation stack.

Since the cavity has no length adjustment, a double passed AOM is used to shift the frequency of the output laser light from that of the reference cavity. Double passing the 1st order diffracted beam output of the AOM, the light is shifted by twice the drive frequency of the AOM. This frequency is a nominal 80 MHz +/- 5 MHz. This gives a tuning range of +/- 10 MHz, which was chosen after analysis of data from the 40m interferometer and the 12m triangular mode cleaner.

The stabilization servo would need about 46 dB of gain at 10 kHz, but the fast piezos probably don't have a usable bandwidth much more than 20 kHz due to their resonances. This makes it necessary to utilize an external phase correcting Pockels cell to extend the bandwidth of the servo so it can be stable. The Pockels cell used is a New Focus model 4004 broadband Pockels cell, with a specified modulation depth of 15 mrad/V.

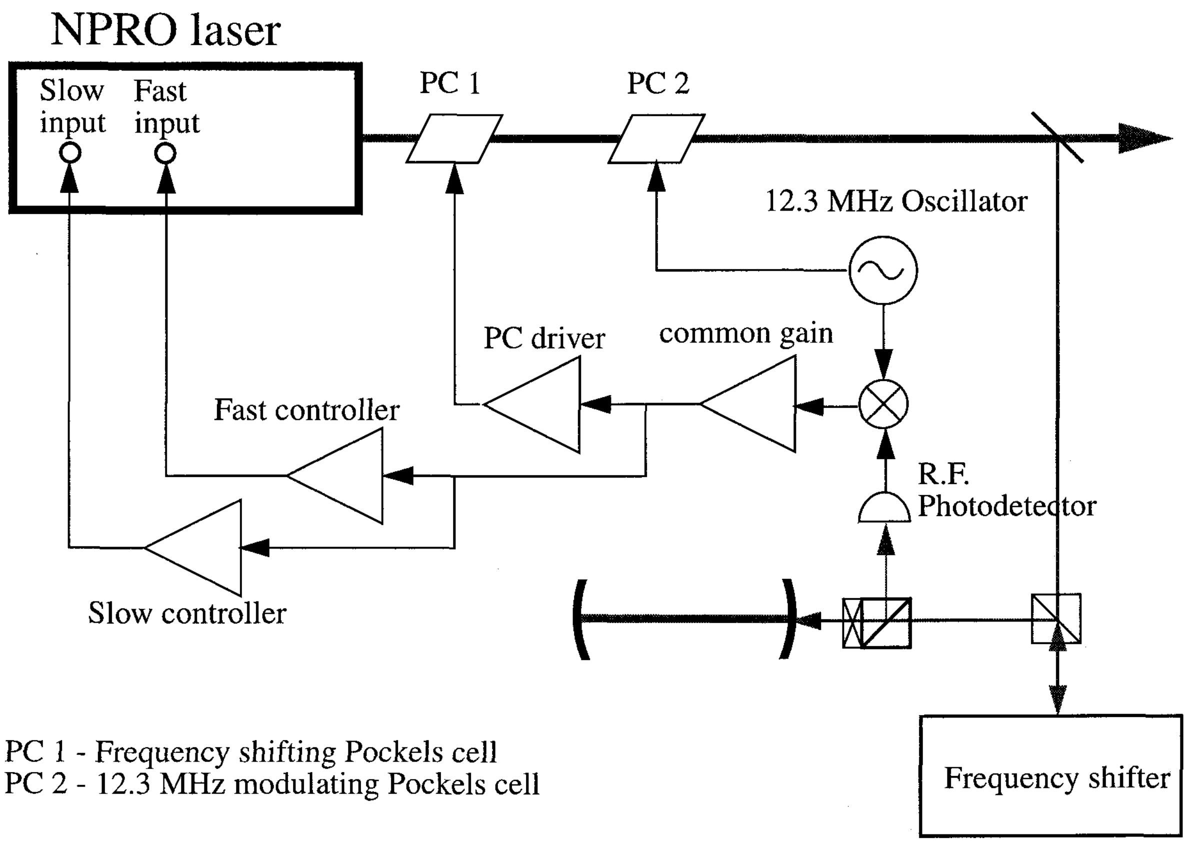

PC 1 - Frequency shifting Pockels cell
PC 2 - 12.3 MHz modulating Pockels cell

The diagram above lays out the general design of the NPRO PSL. The other Pockels cell in the diagram is the 12.3 MHz Pockels cell, a New Focus 4003 resonant Pockels cell. The 12.3 MHz oscillator and amplifier are existing modules. The photodiode is an existing RF photodetector, using a tuned tank circuit at 12.3 MHz, modified for 1064 nm with either a YAG444 or a 220A photodiode.

4 SERVO DESIGN

Servo design requires some modeling, into which physical parameters need to be put. Measurements were made of the various actuators to build the model.

4.1. Actuators and existing gains

4.1.1. Slow actuator

The slow actuator frequency response was measured by locking the laser to a Coherent optical spectrum analyzer in transmission. The slow input was driven and the transfer function from the slow input to the locking error signal was taken. The closed loop gain was divided out, giving the frequency response of the actuator. This is shown in Figure 1 (All figures are collected at the end of this document). A simple model for this response is a 3-pole roll off at 0.2 Hz, with a DC gain of 3 GHz/V.

4.1.2. Fast actuator

Nominally, the piezo of the fast actuator should have a relatively flat frequency response out to the mechanical resonances of the piezo. A measurement of this was made by locking the laser to the reference cavity with low gain and bandwidth (< 1 kHz). Above the servo bandwidth, then, the signal out of the demodulator is essentially an open loop measurement of the frequency noise. The fast piezo was driven above the unity gain of the servo, and a transfer function was measured from this input to the demodulator out. Figure 2 shows a 10-100 kHz span of this measurement. Included in the dynamics of this measurement is the cavity pole, which shows up at 35 kHz. Subtracting the cavity pole results in a flat response at least to 100 kHz. Of note are the features around 30 kHz, 60-70 kHz, and above 90 kHz. These are consistent with parallel, or parasitic resonances, most likely of the structure the piezo is mounted on. A measurement was made to higher frequencies, also. Piezo resonances were found around 250 kHz, above which the response was not coherent. The gain of the piezo was determined by driving the piezo open loop with a triangle wave generator. The amplitude was high enough to scan through the carrier and both sidebands. A photodiode monitored the transmitted light through the cavity, and an oscilloscope was used to determine the amount of voltage needed to scan from sideband to carrier, then carrier to the next sideband. Several measurements were averaged to give 4.1 MHz/V.

4.1.3. Pockels cell

The Pockels cell response was measured in the same fashion. The laser was locked with low gain and bandwidth (~2 kHz). The Pockels cell was driven, and the transfer function from the Pockels cell to the demodulator output was measured. Figure 3 shows this transfer function, with the cavity pole, the demodulator gain, and the zero of the Pockels cell divided out. The measured gain of the Pockels cell is 19 mrad/V.

4.1.4. Demodulator gain

The "demodulator gain" refers to the voltage out of the mixer due to a frequency fluctuation about the point of resonance. This can be estimated using a relatively simple analytical model. The power incident on the photodiode which contains the signal proportional to frequency fluctuations goes as $Re\{E_0 \cdot (E_{+1} + E_{-1})^*\}$, where E_0, E_{+1}, and E_{-1} are the carrier and sideband fields reflected from the cavity. Using 85% photodiode efficiency, a known impedance of 8000 Ω in the tank circuit of the R.F. photodiodes, and a 6 dB insertion loss into the mixers, a demodulator gain of about 1 V/2.5 kHz is arrived at. In reality, the actual gain will probably be lower, due to imperfect modematching, etc. However, this number is sufficient to use in a model.

4.2. Controller design

A block diagram of the controller is shown below. Each of the actuator loops needs to be used in parallel with the other 2. The rule for parallel open loop gains is that the open loop looks like the frequency response of the particular loop which dominates in gain. In this case, we expect to see the slow loop dominate below .1 Hz, the fast loop to dominate up to about 20 kHz, and the Pockels cell to dominate out to the bandwidth of the servo.

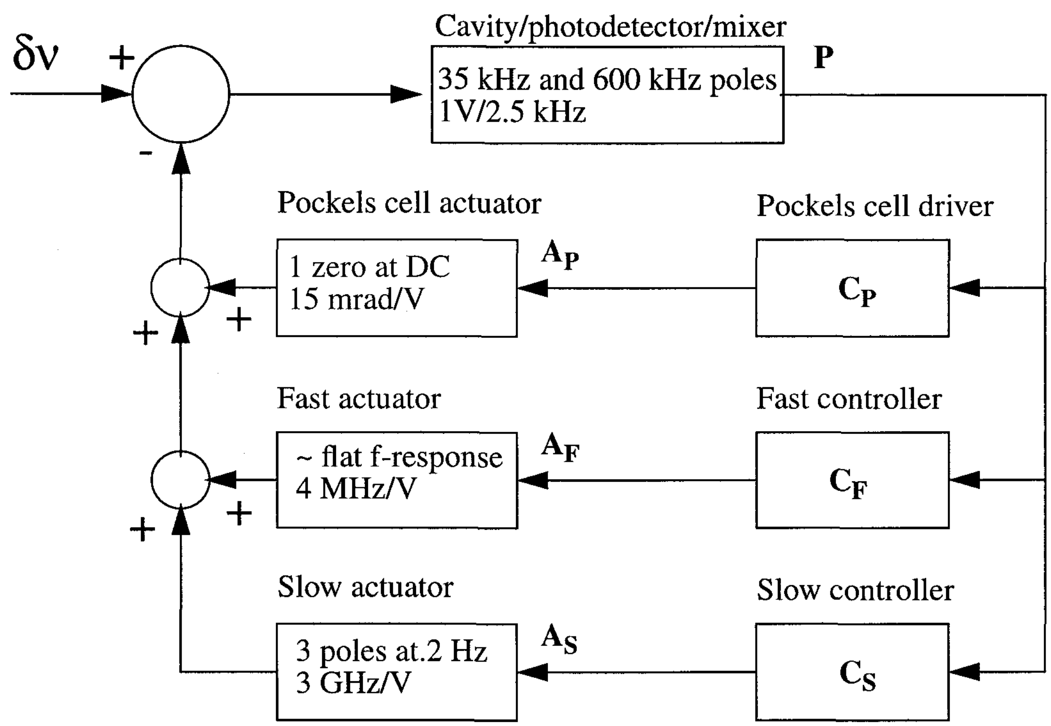

4.2.1. Slow controller

The main purpose of the slow controller is to provide DC control of the laser. That is, for long term drifts due to variations in temperature of the laser, the reference cavity, etc., the slow actuator, with a possible tuning range of about 600 GHz, is used. Since the reference cavity has no length control, and a free spectral range of 750 MHz, DC fluctuations are most easily controlled using the slow actuator. This is most effectively done by using an integrator in the slow controller. For practical purposes, this controller needs a "switch" that turns the integrator off, which is done by moving the pole from DC to some finite frequency, in our case 0.05 Hz. This is useful to short the integrator if any offsets have been integrated, and allows for faster time responses of the slow controller in this mode. Since acquiring lock would also be very difficult with the integrator on, the shorted integrator is referred to as "acquisition mode". This loop should have a bandwidth of approximately 0.1 Hz, due to the poles at 0.2 Hz in the actuator.

The range of the slow controller would in principle need to be within one free spectral range of the cavity, or 750 MHz, which corresponds to about 200 mV in the slow actuator. Standard op-amps will have plenty voltage output to tune over many free spectral ranges, so the dynamic range of this controller would not be a problem.

At the upper end of the slow bandwidth, 0.1 Hz, we have no specific frequency noise requirement, but the requirement at 10 Hz of 10 mHz/$\sqrt{Hz}$ might suggest a level of 100 mHz/$\sqrt{Hz}$ at 0.1 Hz. Based on a conservative 1V/10 kHz mixer output, this corresponds to an input referred electronics noise of ~ 1μV/$\sqrt{Hz}$, which is pretty trivial.

4.2.2. Fast controller

The fast path will be responsible for the frequency noise suppression in the bandwidth of interest. The main design requirement is that the servo has enough gain, and has low enough noise. Given a 1/f free running noise spectrum, then the controller needs to be at least 1/f. A 2 pole controller was designed, mostly because it was simple enough to do. A third pole was added to help roll off the response at higher frequencies, in response to the presence of piezo resonances in the 200's of kHz. A zero was added to improve the phase of the fast loop at the point where control is handed over to the Pockels cell. So, the fast controller has poles at 10, 100, and 10000 Hz, and a zero at 500 Hz. The nominal gain at DC is about 10^6.

Given the free running frequency noise quoted in section 3.1, and assuming this is reasonably good down to .1 Hz, integrating this power gives about 2×10^6 Hz_{RMS} from .1 to 10000 Hz. Working back from the fast actuator, this corresponds to about 0.5 V_{RMS}, which is not a problem for controller or actuator.

The noise of this servo must be small, at least a factor of 10 below the frequency noise goal. Using 1 mHz/$\sqrt{Hz}$, and the demodulator gain, we specified an input referred noise of 300 nV/$\sqrt{Hz}$, which is also fairly simple, as long as the gain of this loop is kept in the early stages of the electronics.

4.2.3. Pockels cell driver

The main purpose of the Pockels cell path is to extend the bandwidth of the overall loop gain, so that there can be ~46 dB of gain at 10 kHz. There are several factors that motivate the shape of this loop. Drive voltages to a frequency correcting Pockels cell typically are rather high. In order to keep this number as low as possible, the loop is ac-coupled with 2 zeros at DC. This puts the phase of the loop +270 degrees. At the point where the Pockels cell loop gain equals the fast loop gain, however, stability considerations dictate that the difference in phase of the two loops needs to be less than 180 degrees. Since the fast loop has phase close to -180, several poles need to be incorporated in the Pockels cell path to bring the phase to an acceptable level. Zeros are also incorporated the controller to account for the cavity pole and the photodiode pole. To this end, the Pockels cell controller has two zeros at 0 Hz, one at 50kHz, and one at 1MHz. A pole is put at 1 kHz, and 3 more at 5 kHz. Gain is adjusted to cross the fast loop gain around 20 kHz.

As mentioned before, the two zeros in the Pockels cell path are included to reduce the amount of rms contribution from frequency noise out of the band dominated by the Pockels cell. A model of the loop was put together in Matlab, the transfer function was calculated from frequency noise to the voltage into the Pockels cell. Multiplied by the frequency noise of the laser, and integrated, the total rms to the Pockels cell is about 1-2 V_{RMS}. This allows the Pockels cell controller to be made using standard op-amps, which simplifies the problem somewhat.

Noise consideration are similar to the slow controller in that the bandwidth dominated by this controller is out of the specified frequency range for control.

4.2.4. Additional features

In addition to the poles and zeros and voltage requirements for each of the three loops above, other features of note are included in the design. First, a feedaround input, as required in the conceptual design document, is incorporated into the early stages of the amplifier. This is simply to be a unity gain buffer. Also, to simplify the amount of cabling, the mixer is included on board, along with notch filters for the modulation frequency and its first two harmonics. The slow path also includes a buffered input at the end of its path in order to sum an external DC voltage offset. As mentioned in the conceptual design document, this is necessary for lock acquisition. A DC bias is applied to the slow input in order to coarsely tune the laser to the resonant frequency of the reference cavity. Once the laser is close to the cavity frequency, lock can be acquired and the integrator in the slow path can be used to maintain the proper bias to the slow controller to keep the laser at that frequency. Also included are test inputs and outputs to measure transfer functions, and monitor outputs are on each of the servo path outputs. Since this is a 3 degree of freedom system, 3 gain controls are included. Since the Pockels cell path goes to high bandwidth where phase delays can have significant effect, the placement of the 3 gain stages are an overall gain stage, and one each in the slow and fast paths, leaving out any direct gain control in the Pockels cell path. Switches are included to switch off the slow controller, to switch between “acquisition” and “integration” mode in the slow controller, and to switch to a test input after the mixer output for diagnostic purposes.

A schematic of the controller along with a list of the front panel features and labels is included at the end of this document.

4.3. Modeling results and predictions

4.3.1. Loop gain and expected residual frequency noise

As mentioned in section 4.2.3, a model of this control loop was built in Matlab, using the measured parameters of the system and the proposed control electronics, as laid out in the diagram above. The paradigm consisted of *.m files which contained the frequency response of each of the blocks in the diagram in section 4.2. A master file called these transfer functions and multiplied them in the appropriate fashion to generate whatever transfer function was required. The resulting open loop gain is given by the following equation, and shown in the following figure. The C’s, P’s and A’s are defined as in the block diagram in section 4.2. Calculating the residual frequency

$$L_{OL} = P \cdot (C_S A_S + C_F A_F + C_P A_P)$$

noise requires the closed loop gain, shown below, multiplied by the open loop frequency noise spectrum. Based on this transfer function, and assuming a frequency noise spectrum like that

$$L_{CL} = \frac{1}{1 + L_{OL}}$$

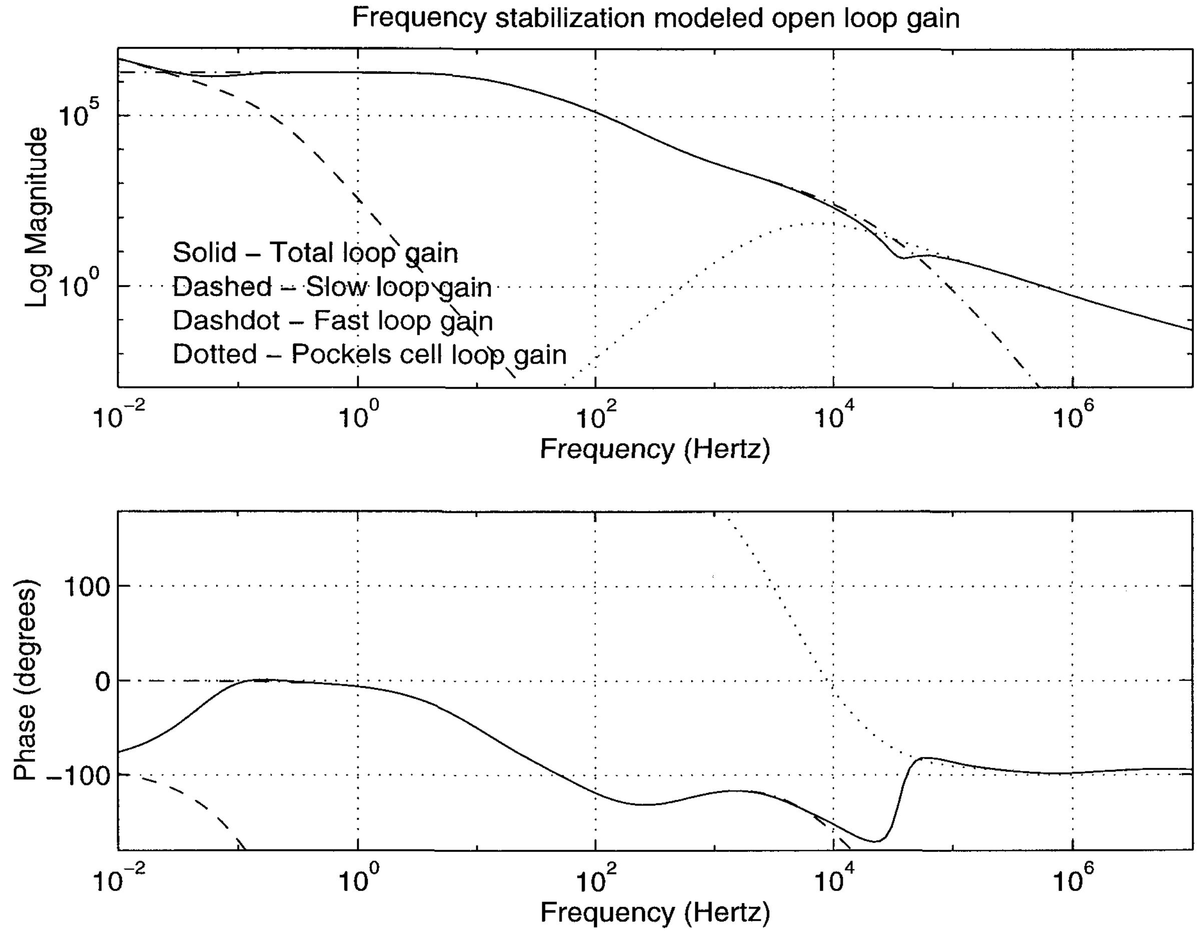

stated in section 3.1, this results in a residual frequency noise spectrum shown in the following figure.

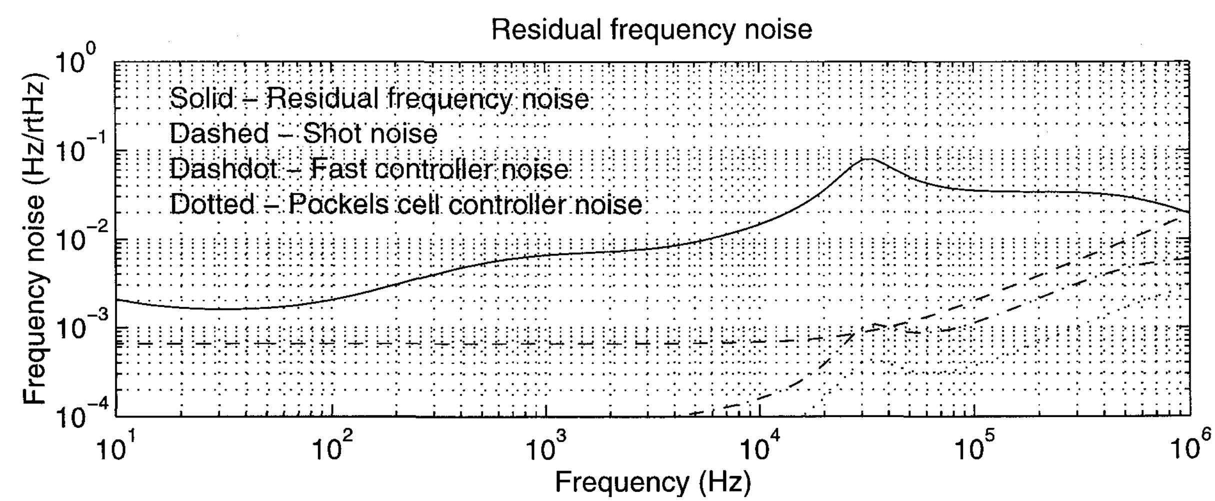

4.3.2. Limiting noise sources

4.3.2.1 Shot noise

The frequency detection method is a process dependent on sensitivity to optical power, in other words, it counts photons. A fundamental noise source with this sort of detection is shot noise, which essentially goes as $\sqrt{N}$, where N is the number of photons detected. Detailed calculations have derived formulas for the shot noise sensitivity of this detection method, given as

$$\frac{S}{N}(f) = \vartheta_{FSR}\sqrt{\frac{e}{\sigma}} \cdot \frac{\sqrt{3\left|E_+\right|^2 + E^2{}_{DC}}}{2\pi E_2 E_+} \cdot \frac{(1 - r_a r_b)^2}{T_a r_b} \cdot \sqrt{1 + \left(\frac{2\pi f}{\omega_c}\right)^2}$$

Definitions for the various parameters are found by referring to either "Shot Noise in a Recycled Unbalanced LIGO" by Torrey Lyons and Martin Regehr, or "Calculations for the Shot Noise in the Recycled 40m" by Malik Rahkmanov. Given the appropriate parameters for the optical configuration, the level of shot noise is shown on the previous figure.

4.3.2.2 Electronic noise

The designed electronics were modeled in Cadence to predict their transfer functions, and estimate their noise outputs. Cadence is a sophisticated program which takes into account real properties of op-amps, phase delays in circuits, etc. From the output referred noise predicted by the Cadence model, transfer functions were derived from the output of the controllers to the frequency error point to estimate the contribution to frequency noise due to noise in the electronics. The figure above containing the residual frequency noise also contains the level of noise contributions from both the fast and Pockels cell controllers.

5 PERFORMANCE

5.1. Measured transfer functions and noise

Figures 4 and 5 show measured transfer functions of the fast and pockels cell controllers. These all agree with the modeled transfer functions to reasonable levels of accuracy. Figures 6 to 8 show measured output referred noises of each of the controller electronics. These also agree remarkably well with the Cadence model prediction.

5.2. Locked laser measurements

Figure 9 shows an in-loop (and hence lower bound) measurement of the residual frequency noise. The following figures, 10 and 11, shows measurements of the open loop gain for this particular measurement, over different bandwidths. Of note is that the frequency noise requirements have mostly been met within the bandwidth of interest as measured inside the loop. Although the figure shows 10 mHz/$\sqrt{Hz}$ only to 7 kHz, obtaining this level at 10 kHz was not difficult, however, it entailed increasing the gain of the fast loop, which causes the bump at 30 kHz to grow. This is a result of the relative phase of the fast and Pockels cell paths at the crossover point, which gets worse at higher frequency. Agreement between this and the predicted residual frequency noise is very good. However, some problems show up in the measurement of the loop gain, most notably the amount of phase available at unity gain. The model prediction shows that there should be plenty of phase at unity gain, and bandwidth shouldn't be limited at all by this consideration. The measurement on the other hand indicates that the phase drops considerably in the 100 kHz band, and that phase margin goes to 0 at about 600 kHz. The fact that the phase is decreasing in the fashion that it does tends to suggest that phase delays are accumulating in the Pockels cell path. The measurement of the pockels cell controller supports this, and modeling phase delay produces a loop gain which agrees with the measurement.

6 OPERATION

6.1. Setup

The figure below lays out our optical table and elements.

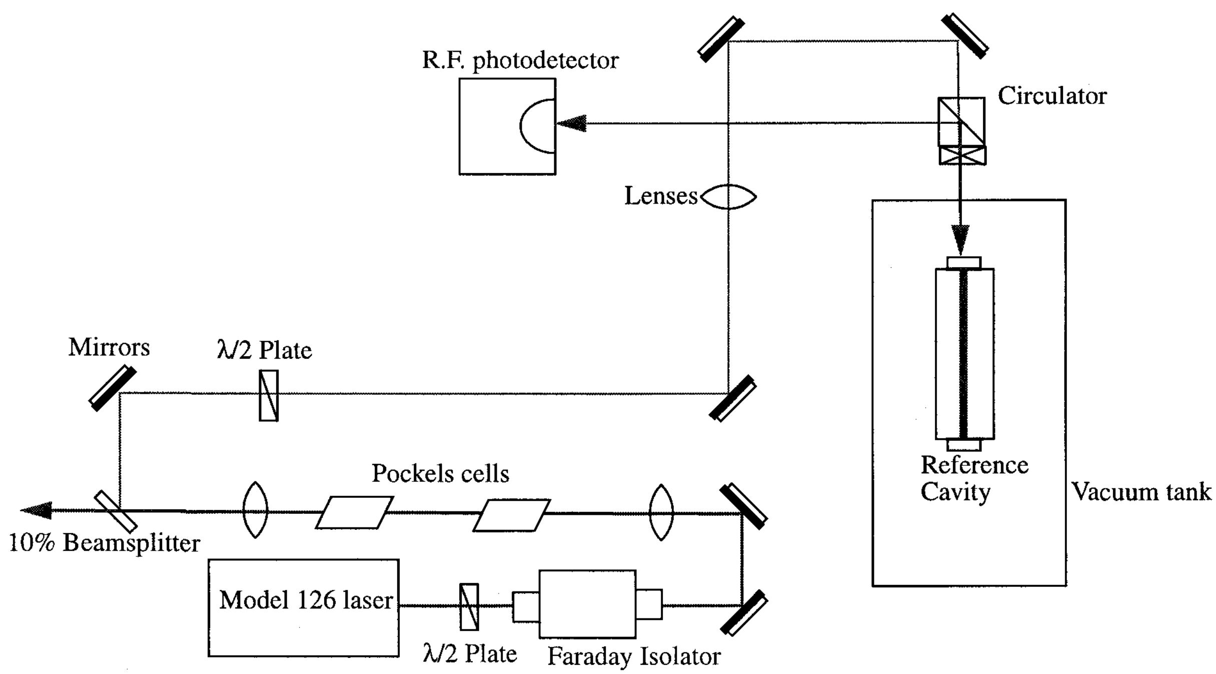

The first λ/2 plate is used to rotate the polarization of the light from vertical to 45 degrees for insertion into the Faraday isolator (measured insertion of 93%). The light leaving the isolator is then polarized perpendicular to the table. The first lens is used to focus the light to a waist position between the two Pockels cells. The measured spot size at 5 cm from the laser is ~.2 mm, and the waist positioned between the Pockels cells is nominally.1 mm. The first Pockels cell is the broadband frequency correcting Pockels cell, and the second is used to impose the 12.3 MHz modulation. This was driven at ~ 6 V_{P-P}, resulting in a modulation of ~ Γ =.75. The next lens is used in conjunction with the 3rd lens for modematching. A 10% window is used to pick off a portion of the light for frequency stabilization. Since the laser is outputting about 600 mW, and we really only would like about 10 mW input to the cavity, the second λ/2 plate is used in conjunction with the circulator's polarizing beamsplitter to dump the rest of the light. This obviously is not optimal for use in an interferometer, however it suited this particular setup. The circulator is a polarizing beamsplitter with an optically contacted $\lambda/4$ plate. This causes the light reflected from the cavity to be reflected by the polarizing beamsplitter, not transmitted. This light is then detected by an R.F. photodetector, which has a tuned resonant circuit to maximally transmit power at 12.3 MHz.

The vacuum tank was kept at 10^{-7} torr using a Vac-Ion pump. Inside the vacuum, the reference cavity was hung from small springs mounted on posts. Two small copper vanes were hung from the cavity close to two sets of 4 magnets in a quadrupolar configuration for eddy current damping of the two swinging modes of the suspended cavity. The supports for the cavity are mounted on an isolation system made up of 3 plates with RTV silicone springs between them. Not shown but also used was an infrared camera, placed at the far end of the vacuum tank in order to see the light transmitted when the cavity was on resonance. Also not shown is an optical spectrum analyzer, which has another 10% pick off just before the RF photodiode. Note that the frequency shifter was not incorporated into this layout, because time did not allow.

6.2. Alignment

Typical alignment procedure involved first passing the light through the Faraday isolator. This was done using an infrared photo card, and eyeballing the position of the beam approximately to the center of the F.I. apertures. The alignment was done by shifting the position and tilt of the F.I.. The next alignment was the first lens, which simply involved positioning the lens such that the beam passed through the center of the lens, by marking the position of the beam without the lens at the other end of the table, and bringing the position of the beam back to the same place once the lens was installed. Next, in order to align the beam through the 2 mm apertures of the Pockels cell, a negative lens was used to blow the beam up so it was approximately 2 cm in diameter on a beam block. The position of the front aperture for each Pockels cell was adjusted by finding where the beam clipped the aperture on each side, then centering the aperture in between these points. Then the positions of the rear apertures were dithered to minimize the distortion to the beam as viewed in the expanded spot. The next two lenses were also centered, and placed according to calculations to mode match into the cavity. Alignment into the cavity was performed first by roughly eyeballing the light to the center of the input cavity mirror, using the final two mirrors in the optical path. Then a function generator was used to dither the frequency of the laser over a large range in the slow actuator (by at least one free spectral range), while the camera which looked at the output of the cavity was monitored. Once modes began flashing in the cavity, adjustments were made based on the strength of the modes, to begin optimizing for lower order modes. Once a TEM_{00} mode was found, a slow DC offset was applied to the slow actuator to bring the laser close to the frequency of the 00 mode, and the function generator drove the fast actuator (the time scales of the slow actuator were irritatingly slow). Then, the alignment attempted to maximize the output of the 00 mode through the cavity. For reference, irises were placed in the optical path and centered on the beam for future alignment.

Some “electrical” alignment is also done. First, the phase of the local oscillator applied to the mixer needs to be set. Since the phase shifter has, at the smallest, 10 degree divisions, this was done by eye. The 00 mode was found, and the function generator used to dither the fast input of the NPRO. The output of the mixer was viewed on an oscilloscope set to trigger as the laser went through resonance. The phase was adjusted to maximize the symmetry of the demodulator output, verified by setting the phase to the “wrong” phase and confirming the symmetry in that output.

Also, the electronics comes with the ability to tune out any voltage offset out of the mixer. This was done by looking at the MIXROUT output through a SR560 and a gain of about 100, low passed at .03 Hz. The pot, located inside the NIM module, was tuned such that the DMM read about 1 mV, after the bias of the SR560 was tuned out, which seemed to be about as good as could be done.

6.3. Nominal settings

A set of parameters was developed as indications of the state of the system. Nominal values for these parameters were worked out to insure repeatability of results. These are listed in the table below.

Modulation voltage	2.5 V_{Peak}
Modulation depth	.75
Cavity input power	10 mW
Visibility	80-90%
Laser power	600 mW
Vacuum pressure	10^{-7} torr
RFPD V_{DC} (out of lock)	-130 mV
RFPD V_{DC} (locked)	-48 mV

Table 1: Nominal parameters for NPRO laser

6.4. Lock acquisition

The process used to acquire lock in these experiments was not automated. Gains in all loops are turned to 0, the slow loop is left open, and the integrator is switched off to keep from integrating up any offsets in the path. The slow actuator is ramped in the slow DC input using the Calibrators DC voltage supply until the 00 mode is found, usually by watching the camera which is looking at the transmitted cavity light. Once the mode is found, the gain in the fast path is turned up slightly. If the laser is very close to resonating, the laser will usually lock right away. If not, the slow actuator must be used to tune the laser closer to the right frequency, with the fast gain small, maybe about 0.1, and the common gain at minimum. The reason that the fast gain must be kept low is that when the laser needs to tune through the point where the sidebands are resonant, the servo has the wrong sign, and a large voltage builds up in the fast path as the slow DC tries to push against the fast gain of the servo. When the slow finally manages to exceed the ability of the fast loop to keep away from the sideband resonance, the fast voltage drops to zero and the laser frequency shifts very rapidly through the carrier and the other sideband, to the tuned point of the slow DC. However, with low fast gain, it's easy to tune via the slow DC close to resonance, then turning up the fast gain usually will lock the laser immediately. At this point, the fast gain needs to be increased to about 0.5 to 1.0 before the common gain can be turned up. This sets the crossover between the fast loop and Pockels cell at approximately the right place in the 20 kHz region. Then the common gain can be turned up to approximately 0.6 or 0.7. This brings the loop very close to the maximum gain possible, limited by the loss of phase margin at 600 kHz. If the MIXROUT output is being monitored on an oscilloscope, this is evident by the output voltage beginning to grow as the gain is increased. This is not due to a noisier laser, but rather unity gain oscillations at around 500 kHz. At this point, the slow loop is turned on, and the slow mode is switched from acquisition to integration. The gain of the slow loop can be turned up to about 0.3 or 0.4 before oscillations begin to develop.

Re-acquisition once lock is lost is a similar process. The first thing that needs to be done fairly quickly is the integrator must be switched off. Also the gains must be turned down in the same way as acquiring lock the first time. If the slow DC bias is still hooked up and supplying the DC voltage which brought the laser roughly to resonance, frequently all that's needed is to wait until the laser relaxes and returns to equilibrium. Loosing lock typically causes the laser to shift its frequency somewhat, and requires about a half a minute to relax. Once close to resonance, again, the fast gain is turned up, then the common gain. However, if the laser frequency set by the slow DC bias has drifted relative to the cavity sufficiently, the laser will not return to the point where the fast loop can acquire. At this point, the slow DC bias voltage must be scanned again to find the resonance. This typically is not far away, though, so radical shifts in voltage should not be required.

6.5. Problems

Below is a list of problems and other notes concerning frequency stabilization of the NPRO, both understood and not understood.

- When the common gain is turned up too high, although not high enough to lose lock, occasionally 2 spikes in the frequency spectrum show up around 3 and 5 kHz. The origin of this is unknown, however it's suspected that a stage in the pockels cell path may be saturating. The pockels cell path has very high gain, which peaks at about 4 kHz, and the first pass at these electronics had bad saturation problems in the pockels cell path (not on the output, however).
- Concerning the use of the power adjust input for power stabilization. There is a very strong coupling between the power adjust input and frequency noise. Figure 12 shows a transfer function between the power adjust input and frequency noise output. This measurement was made with the built in "noise eater" off, however with the noise eater on, this trace goes down maybe only 10 dB. It was made by driving the power adjust and looking at the voltage out of the demodulator. The cavity pole and the loop gain have been divided out, as well as the frequency to voltage gain of the demodulator. This coupling is both good and bad. The good is that stabilizing the power output of the laser diodes does help the frequency stabilization, which tends to suggest that a large amount of frequency noise comes from intensity fluctuations of the laser diodes. The bad is that stabilizing the laser light at some point far down the optical path may introduce noise into the laser diode in order to correct for artificial intensity fluctuations (i.e., fluctuations due to beam jitter through a mode cleaner, parasitic interferometers). This in turn will make the frequency stabilization worse.
- Changing the common gain of the servo while the integrator is on seems to occasionally cause problems, that is the slow path seems to want to begin oscillating. Not understood.