LIGO-T010113-00-W 10/05/01

Core Optics Auto-Alignment Sequencer

Keely Willis

Distribution of this document:
LIGO Science Collaboration

This is an internal working note
of the LIGO Project.

California Institute of Technology
LIGO Project – MS 18-34
1200 E. California Blvd.
Pasadena, CA 91125
Phone (626) 395-2129
Fax (626) 304-9834
E-mail: info@ligo.caltech.edu

Massachusetts Institute of Technology
LIGO Project – NW17-161
175 Albany St
Cambridge, MA 02139
Phone (617) 253-4824
Fax (617) 253-7014
E-mail: info@ligo.mit.edu

LIGO Hanford Observatory
P.O. Box 1970
Mail Stop S9-02
Richland WA 99352
Phone 509-372-8106
Fax 509-372-8137

LIGO Livingston Observatory
P.O. Box 940
Livingston, LA 70754
Phone 225-686-3100
Fax 225-686-7189

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

Core Optics Auto-Alignment Sequencer

Keely Willis

Mentor: Michael Landry

Abstract:

The fine-tuning of the alignment of LIGO's suspended optics is lengthy and tedious. To automate this process, we have written Align, a sequencing program in Expect that interacts with existing processes to find the optimal alignment of the interferometer, one cavity at a time. After it is initiated, Align contacts the Experimental Physics and Industrial Control System (EPICS) to find the initial locked positions of the mirrors. Align then uses EPICS to repeatedly reposition the mirrors one at a time around this initial position, adjusting both yaw and pitch. In each new position, Align interacts with Diagnostic Test Tools (DTT) to run a sine response test on the cavity currently being aligned, which applies a sine wave of frequency f to dither the mirror, measures the power output of the cavity, and runs a fast Fourier transform on the results. From this Align extracts the fundamental frequency and first harmonic coefficients and combines them in a ratio, p(f)/p(2f). The intention of Align is to reposition the mirrors until this ratio is minimized, the minimal ratio corresponding to the optimal alignment for that cavity.

Introduction:

In the LIGO interferometer (IFO) alignment process, the core optics go through several different stages of adjustment before optimal alignment is attained. In the first part of this process the mirrors are adjusted in yaw and pitch, and the cavities in length, until resonance is attained within the cavities. At this point, the optics undergo a much finer alignment, again in yaw and pitch. The intent of this fine-adjustment is to find the maximal intensity of light inside the Fabry-Perot cavity, which increases the sensitivity of the IFO as a detector. This is a crucial step in the alignment of the IFO; however, it is extremely time-consuming, sometimes requiring hours of manual adjustment. The intent of this project is to automate this fine-alignment process, so that it can be completed quickly and efficiently, necessitating a minimum of human interaction.

To this end, we have written a program, Align. Align is written in the scripting languages Tcl and Expect. The basic concept governing the program is for Align to emulate the process the operator would follow. To proceed with the alignment, Align interacts with several different processes, including the Experimental Physics and Industrial Control System (EPICS) software, and the Diagnostic Test Tools (DTT) software. Align controls the optics via EPICS, and measures the response of the cavities with DTT.

External Processes:

The functions of DTT that Align uses are the sine response test and the triggered time series test. In the sine response test, DTT applies a sinusoidal waveform of frequency f to dither the mirror's angular offset in pitch. DTT records the values of the cavity power through a predefined measurement channel. DTT then runs a fast Fourier transform (FFT) on the recorded results, producing a power spectrum and thus a series of harmonic coefficients, which are the data of interest. In the triggered time series test, DTT simply records the power output over a specified period of time as communicated through a predefined measurement channel. For both of these tests, the results can be saved to a file in xml format, and using the C command `xmlconv`, converted to ASCII format and accessed by Align. This allows Align to extract the necessary information.

Align uses EPICS for a more direct interaction with the mirrors. Given a file containing a list of channels corresponding to optic angular offsets, the EPICS command `caGet` allows Align to find the angular offsets of the mirrors, and save these values to a file. Align also uses the EPICS command `caPut`, which takes values from a file along with their corresponding angular offset channels and set the mirrors to those angular offsets. Two other EPICS commands, `caget` and `caput`, allow Align to interact with its Motif-based Editor and Display Manager (MEDM) screen, and thus with the user.

Theory:

In order to find the optimal alignment we need a measure of the intensity of the light within the cavity. If we assume that there is a single, optimal position for the mirror, which results in the maximal intensity of light within this cavity, we can define the angle between the actual position of this mirror and the ideal position as α_{off}. In Figure 1, the normal to the mirror in its ideal alignment is shown as a dashed line between the two mirrors, and α_{off} is shown as the angle between the dashed line and the normal to the mirror in an imperfect alignment. Using these definitions, we can define the power spectrum as a function of α, where α is as defined below, approximating the spectrum as an inverted parabola, as shown in Figure 2. The equation describing this is

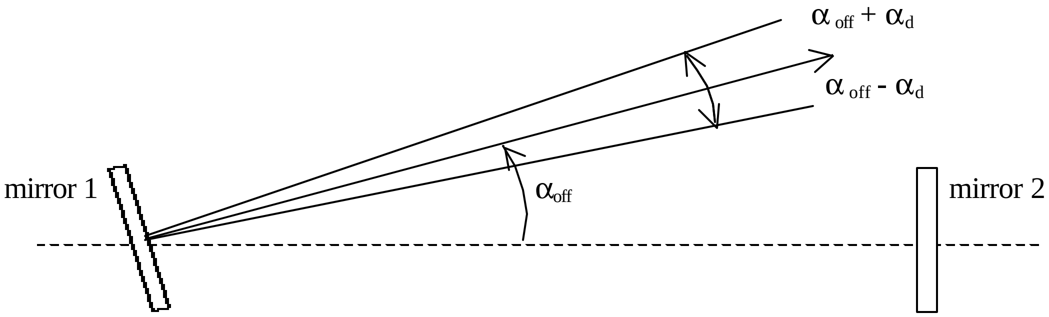

Figure 1 – Optic Angular Offsets (exaggerated)

$$P(\alpha) = P_0 - P_2\alpha^2 \tag{1}$$

where P_0 is the maximal power, and P_2 is a scaling coefficient. We modulate the cavity mirror angle, or set

$$\alpha = \alpha_d \cos(\omega t) + \alpha_{off} \tag{2}$$

where α_d is the angular dithering amplitude, and ω is the angular frequency of the dither, $\omega = 2\pi f$. Figure 1 depicts the angles α and α_{off}. If we combine Equations 1 and 2, we get the following:

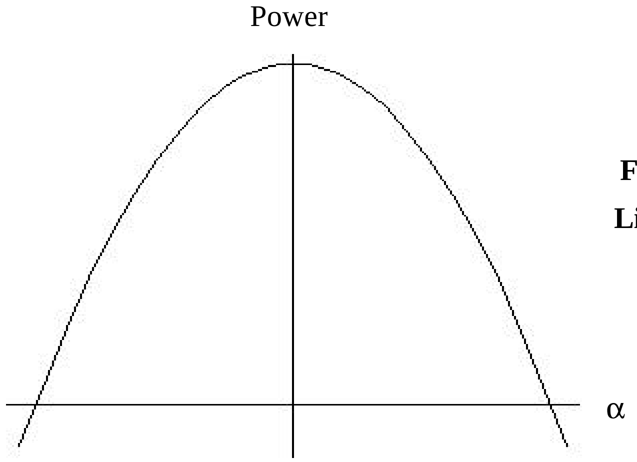

Figure 2 – Approximation of the Power of Light Inside Cavity vs. Angular Offset a of Core Optic

$$P(\alpha) = P_0 - P_2\left[\frac{\alpha_d^2}{2}(\cos(2\omega t) + 1) + 2\alpha_{off}\alpha_d\cos(\omega t) + \alpha_{off}^2\right]. \tag{3}$$

From Equation 3 we can see that a ratio of the f and 2f components of the power spectrum will be minimized for the optimal alignment, i.e. as α_{off} goes to zero. To find the values of these components, we need to access information from DTT. Figure 3 shows graphical results provided by DTT regarding the sine response test. In the top half of the window we see the harmonic coefficients of the power spectrum. The red bars in the graph represent the harmonic coefficients of the drive. The blue bars correspond to the harmonic coefficients of the measurement, these coefficients being the data of interest. Align extracts the coefficients for fundamental frequency and first harmonic from DTT and combines them in a ratio, p(f)/p(2f). The position that produces the minimal ratio is the optimal position of those measured.

A factor that must be repeatedly checked is the accuracy of the results from the sine response test. With the harmonic coefficients comes a coherence coefficient, shown in the bottom half of the window in Figure 3. Here, the red bar represents the coherence of the driver, and the blue bar represents the coherence of the measurement. If we run a sine response test, taking multiple averages, and use both a driver and a measurement channel as is the case with our tests, “the coherence between these two channels can be calculated as follows:

$$\frac{|\langle c_1 \cdot d_1^* \rangle|^2}{\langle |c_1|^2 \rangle \langle |d_1|^2 \rangle} \tag{4}$$

where $\{c_1\}$ and $\{d_1\}$ are the sets of complex amplitudes resulting from n sine detections on each of the two channels, and the averaging is performed on these n values.(1)”

Figure 3 – Results of a Sine Response Test

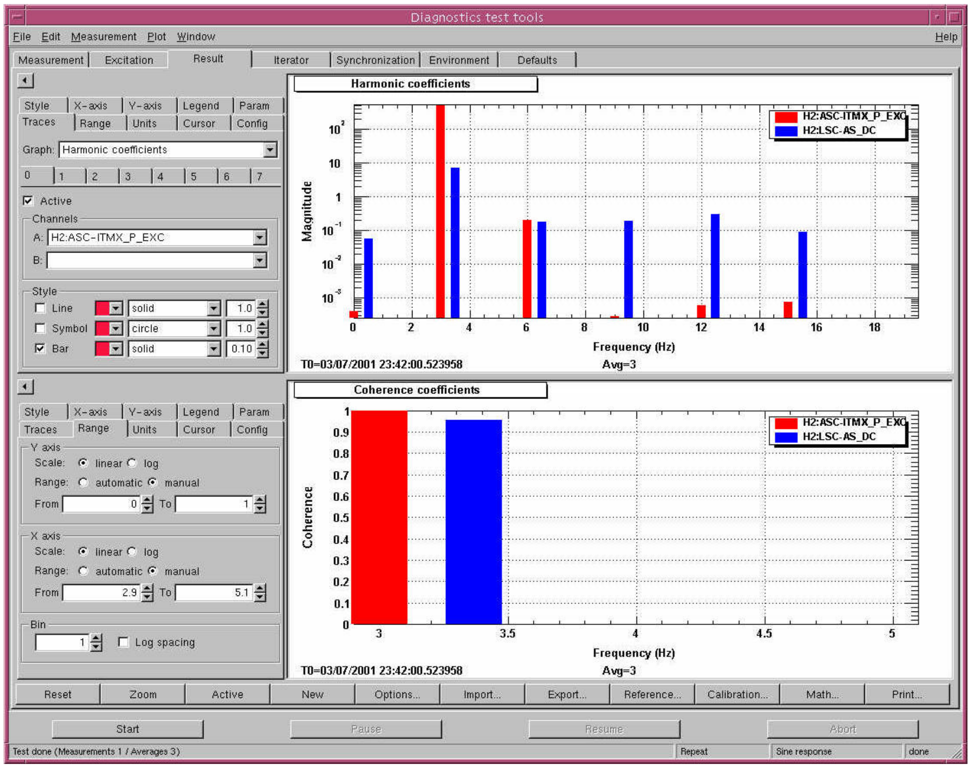

Graph: Harmonic coefficients
☑ Active
Channels
A: H2:ASC-ITMX_P_EXC
B:
Style
☐ Line solid 1.0
☐ Symbol circle 1.0
☑ Bar solid 0.10

Y axis
Scale: ● linear ○ log
Range: ○ automatic ● manual
From 0 To 1
X axis
Scale: ● linear ○ log
Range: ○ automatic ● manual
From 2.9 To 5.1
Bin
1 ☐ Log spacing

T0=03/07/2001 23:42:00.523958 Avg=3

Test done (Measurements 1 / Averages 3)
Repeat
Sine response
done

In simple terms, the coherence is a value between zero and one, indicating the response of the measurement channel to the driver (dithering of mirror.) A high coherence indicates the drive is affecting the cavity, whereas a low coherence means the drive has had little or no impact on cavity power. After each sine response test, Align checks the coherence of the measurement DTT has taken, and if the coherence is below a certain threshold value (0.7) the sine response test is repeated.

There is another aspect to the information provided by DTT that is useful to Align. The information returned contains magnitudes both for the coherence and harmonic coefficients already described; it also provides phase information. The value given represents the difference in phase between the driver and the measured results. As we adjust a mirror in angle, this phase remains largely constant, until the mirror is moved from one side of the optimal position to the other. At this point, the phase changes by 180°, resulting in a positive phase value on one side of the optimal position, and a negative phase value on the other side.

Program:

The first function of Align is to find the initial positions of the optic to be aligned, which has been specified to it through its MEDM screen. It does this by using the EPICS command caGet. Without moving the mirror’s angular offset, Align accesses DTT and uses it to run a sine response test on the mirror, dithering in pitch. Align finds the sign of the phase information and uses that to choose the direction to move the mirror, in order to set it closer to its optimal alignment. Align then uses caPut to move the mirror. This process is repeated until the sign of the phase changes, at which point the mirror is set to a point midway between the last two positions. This process is then repeated for yaw. After the adjustment in yaw is completed, the mirror should be fairly close to it’s optimal alignment.

Next, Align runs the mirror through a series of positions close to this new initial position. To begin, Align adjusts the mirror's angular offset in pitch alone. In each new pitch position, Align commands DTT to run a sine response test on the cavity. Align extracts the fundamental frequency and first harmonic coefficients and combines them in the power ratio p(f)/p(2f). This is stored as the ratio for this position of the mirrors. Align iterates through a series of positions varied in pitch, finding this ratio for each. At each position, before it moves on to the next, Align checks the coherence of the measurement just made. If the coherence is below the threshold value of 0.7, the test is repeated. The coherence of the new result is also tested. This test can be repeated up to six times; after that, Align moves on to the next position, and doesn't use the ratio value for the position with consistently bad coherences. Align then finds the minimum of these ratios, and sets the position corresponding to that minimum as the new pitch value. With the mirror now set to this pitch, Align repeats the procedure described above, except now varying the position and dithering in yaw. This scanning of angular positions defines a three dimensional hyperspace in which the x- and y- axes are pitch and yaw, and the z-axis is the ratio value. Align repeats this procedure, until it finds a position that has the minimum ratio for the

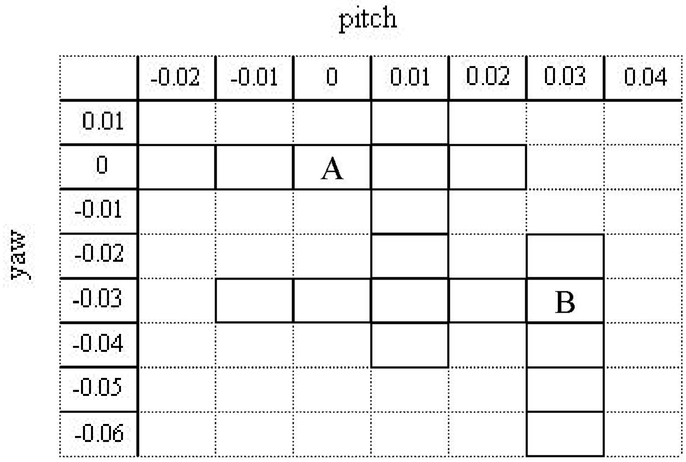

Figure 4 – The Hyperspace

positions in both its pitch and yaw range, corresponding to the minimal value of the hyperspace. This is the optimal alignment for this mirror. In Figure 4, above, the mirror was initially at the position marked ‘A’ in the hyperspace. The spaces with the dark outlines represent the range of positions the mirror was moved through, acquiring data for each position and using that data to decide the next range of positions. The end position in this simple example was position ‘B’. Align sets this mirror to it’s optimal position, and exits.

At any time during this process, the cavity can fall out of lock, so that the cavity is uncontrolled and nonresonant. This can result from any one of many factors, such as an initially unstable lock, noise from another part of the interferometer, or the fact that Align is moving the optics comprising the cavity. If the cavity is no longer in lock, the results being read by Align are invalid, which would in turn, cause Align to find a bogus optimal position for the mirror, if it found any final position at all. In order to avoid this, Align tests repeatedly to verify that the cavity is still in resonance. If the cavity is in lock, even if poorly aligned, the light inside it is considerably more intense than it is when the cavity is out of lock. Therefore, to test for lock, Align commands DTT to run a time-series test on a specified measurement channel. Align averages the values from the test and compares the resulting single value to a preset threshold value. This value can be set from the MEDM screen, or Align will use its default values. This threshold value is set cavity by cavity. If the measured power is less than the threshold value for a Fabry-Perot cavity, or greater for the Michelson cavity locked on a dark fringe, Align assumes the cavity has fallen out of lock. In this case, Align immediately ends any sequence of tests it had been doing, and sets the mirror back to the original position left by the operator. From here, Align commands DTT to run time-series tests repeatedly over several minutes. Align checks each of the resulting power averages from these tests to see if the cavity has gone back into lock. If, after ten repeated tests the cavity has not gone back into lock, Align prints a comment to the screen and exits. If, in this time, the cavity does go back into lock, the position at which the cavity lost lock is recorded, and the entire alignment procedure is begun again. From now on, Align will check each new position of the mirror to verify that it is not a position that let the cavity fall out of lock. If it is such a position, Align moves on to the next position without ever having put the mirror back to the position that caused it to lose lock.

Results:

Because of time constraints and the difficulty in scheduling periods to run such invasive tests on a cavity of the IFO, little conclusive testing of the final form of Align was possible. Table 1, below, lists the final yaw and pitch values of twelve successive tests that were possible. These tests were run on the Michelson cavity of the 2-kilometer IFO, adjusting the position of ITMY. Before each test, the mirror was set to some arbitrary position that allowed the cavity to remain in lock, but was different from previously tested positions. The data in the table shows the consistency of the results. The average pitch value of those found is 1.154 counts, and the average yaw value is 0.770 counts. The standard deviation σ for pitch is 0.002, and for yaw, 0.001.

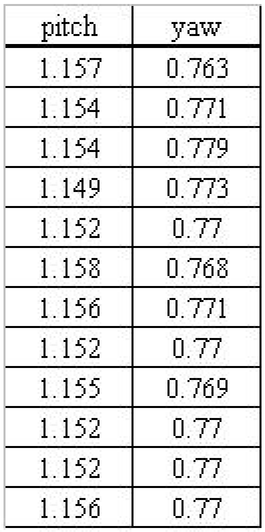

pitch	yaw
1.157	0.763
1.154	0.771
1.154	0.779
1.149	0.773
1.152	0.77
1.158	0.768
1.156	0.771
1.152	0.77
1.155	0.769
1.152	0.77
1.152	0.77
1.156	0.77

Table 1 - Test Data

Discussion:

In this way, Align can be used to find the optimal alignment for a single mirror. With the exception of the first cavity which requires the alignment of two mirrors simultaneously and thus is not currently within the ability of Align to align, the operator can align the core optics in the entire interferometer much more quickly and easily than before. Given a well-aligned cavity, the operator then aligns the adjacent cavities simply by running Align on the mirrors in those cavities, not comprising the original cavity. The amount of time the alignment takes is, of course, largely dependant on how well aligned the cavities were initially.

Conclusions:

The purpose of the core optics auto-alignment sequencer project was to write a program to reduce the work and time needed for the operator to align the IFO. Align is able to find the optimal angular alignment of a single mirror with a consistency within a standard deviation of 0.001 counts in pitch or 0.002 counts in yaw. To continue this project, it would be desirable to expand the range of Align's abilities, to make it able to align an entire cavity instead of a single mirror. This could be achieved by adding a small subprocess that would move both mirrors in the cavity together. The idea behind this would be much the same as that behind Align's current procedure. The mirrors would be iterated through a series of positions, testing the intensity of light within the cavity for each new position. The only way the new subprocess would be different would be in that two mirrors would be moved, between measurements, instead of one.

As it stands, Align can successfully align the majority of the cavities of a LIGO interferometer. The only cavity Align cannot align is the initial Fabry-Perot; however, a simple modification will allow Align to optimize this first cavity.

Acknowledgements:

First and foremost, I would like to thank my mentor, Dr. Michael Landry, for the many hours spent explaining concepts and working through problems with me, and for all the guidance along the way. Also, thanks go to Dr. Daniel Sigg, for answering all my questions about DTT, to Dr. Stan Whitcomb for always being available with advice in the control room, and to Dr. Greg Mendell for introducing me to Expect and letting me use Exploring Expect for the entire summer.

References:

D. Sigg and P. Fritschel. Diagnostic Test Software, LIGO Internal Document T990013-A-D.