VIBRATION CONTROL FEEDBACK R&D
AT UNIVERSITY OF BRITISH COLUMBIA

T. Mattison, R. Greenall, T. Downs
Dept. of Physics & Astronomy, University of British Columbia, Vancouver BC V6T 1Z1 CANADA

Abstract

The final quadrupoles of a linear collider must be stable vertically to the nanometer (nm) level, despite being cantilevered inside a particle physics detector. One approach to this problem is to measure the position of the quads with a laser interferometer, and reposition them with piezoelectric actuators. We have constructed a test platform, interferometer, and software to develop this technology. The platform vibration can be reduced from 90 nm to 5 nm on the ground, and 4.5 nm to 1.5 nm when isolated. This performance seems to be limited by additional modes of the platform supports. The interferometer precision is up to 0.01 nm, and an interferometer end mirror can be stabilized to 0.06 nm with piezoelectric feedback.

1 INTRODUCTION

The beam spot of a TeV-scale linear collider must be small to produce adequate luminosity at affordable beam power. The vertical dimension is of order nanometers (nm), while the horizontal dimension is of order 100 nm. If the final quadrupoles vibrate by a comparable amount, the luminosity will be greatly reduced. Beam-beam deflection can be used to monitor the beam offset, and for steering feedback. This feedback can compensate for motions at frequencies much less than the linac pulse rate (and also within a sufficiently long bunch train). Natural ground motion is high at low frequencies, but it is highly correlated between the quadrupoles. At a good site, at higher frequencies the motion is tolerable, so mounting the quadrupoles rigidly to “bedrock” would probably work. However, in a particle physics application, the quadrupoles will probably be cantilevered into the detector on supports whose resonances will amplify the ground motion. Some form of real-time measurement and compensation of the quadrupole positions will be required.

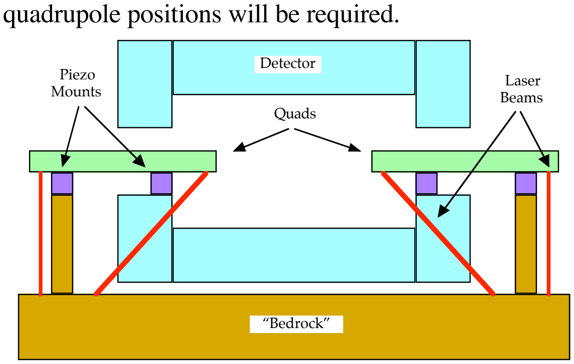

Figure 1: Optical Anchor concept

One approach to this problem is the “Optical Anchor” concept. Laser interferometry is used to measure the positions of the quadrupoles, and piezoelectric actuators are used to correct the position. Several light paths are required, some requiring holes through the detector, as illustrated schematically in Figure 1. Mike Woods demonstrated 0.2 nm resolution and 20 nm/hour drifts in a 10 meter baseline interferometer at SLAC.

2 UBC PROGRAM

Our goal at the University of British Columbia (UBC) is to demonstrate sub-nanometer position stability in one dimension for a 100 kg object over a 10 meter baseline. The SLAC laser interferometer equipment was moved to UBC and installed in the basement of the Hennings Physics Building. The vertical ground motion (Figure 2) is several hundred nm at a few Hz, with a continuum that is about 10 nm at 30 Hz and about 1 nm at 100 Hz, with several large narrow peaks from machinery in the building, some of which come and go.

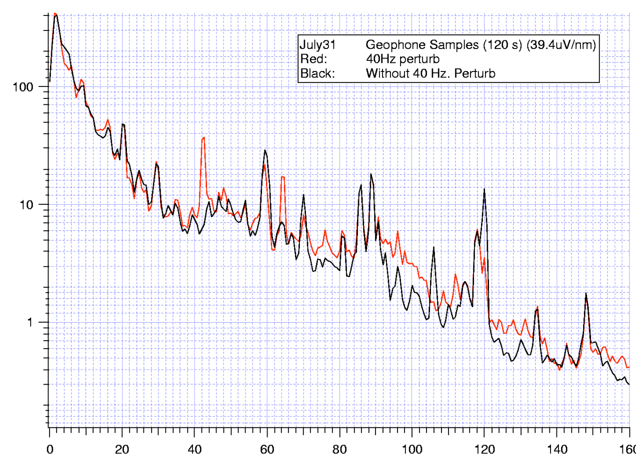

Figure 2: Vertical Ground Motion Amplitude (nm) vs Frequency in UBC Physics Basement

Data acquisition and feedback control is through a Linux PC with analog-digital interface card, or through an AdWin Gold digitizer-DSP interfaced to a Windows PC. The required feedback bandwidth is about 1 kHz, so sampling and control at 5-10 kHz is adequate. To ensure continuous control, the Linux version is implemented as a kernel driver module. Kernel programming proved to be not much more difficult than normal C programming. Interrupt response is adequate for this application, and the full double-precision power of a Pentium IV CPU is available. The AdWin DSP system has zero operating system overhead, but is limited to relatively slow single-precision arithmetic, and is programmed in a proprietary BASIC. The Windows PC compiles and downloads the code, presents a GUI interface, and generates time and frequency domain data plots.

3 TEST PLATFORM

The test fixture (Figure 3) has a 10 kg instrumented platform mounted on flexures to a baseplate with two end-posts. A feedback piezo with variable preload springs on one end of the platform pushes against a variable-stiffness spring to an excitation piezo on the end-post. The other end has one interferometer mirror on the platform, a reference mirror on the end-post, and a capacitive position sensor to measure the platform to end-post distance directly.

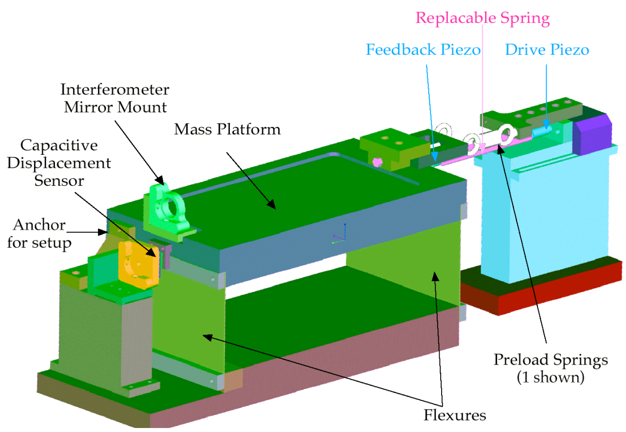

Figure 3: Test Platform

Theoretically, if the feedback applied with a gain of K microns of piezo displacement per micron of mass displacement, the resonant frequency would increase by a factor of $\sqrt{1+K}$. If the ground motion were tolerable above 1000 Hz and the natural frequency were 10 Hz, a gain of K=10000 would be required. Feedback proportional to the velocity of the mass (derivative term) can be used to critically damp the natural resonance, and larger velocity feedback can critically damp the increased frequency due to proportional feedback.

In principle, it should be possible to correct the position of the platform in three sampling time steps, independent of the sampling rate. On the first step, the piezo is moved far enough beyond the target postion that the mass moves half of the desired distance during one clock tick. On the next step, the piezo is moved an equal distance in the opposite direction, to kill off the velocity acquired during the first step. The mass will come to a stop cur at the desired target point. On the third step, the piezo is moved to the point that exerts no force on the mass, so it stays at the target point.

Our first experiments used a rigid rod between the piezo and end post and simple proportional-differential feedback, but we could not make this stable at any useful gain. We replaced the rigid rod by a soft spring (and softer flexures) to lower the resonant frequency. This allowed derivative (velocity) feedback to damp large resonant motions, but increasing the gain caused high frequency oscillations of the spring itself. Simulations that modelled the spring resonance by a lumped mass in the middle of the spring qualitatatively matched the behavior, and showed that the velocity filter time constant and the current limit of the piezo amplifier were contributing to the instability. We built a high-current piezo driver, and better springs with fewer internal modes, but were still not able to apply enough gain to do much more than damp the resonance without inducing high-frequency oscillations.

We measured the amplitude and phase response to piezo excitation vs frequency, and also the response to square-wave excitation by the piezo (Figure 4). In addition to the fundamental mass-spring resonance, there were a large number of other modes in the springs and the platform base. These matched the frequencies of oscillations of the system when feedback gain was too high.

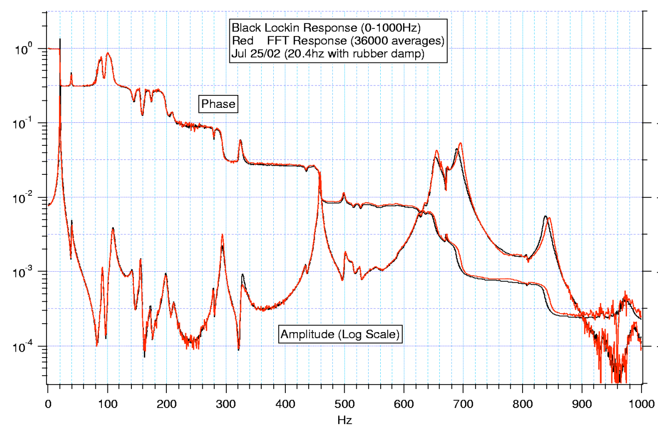

Figure 4: Platform Response to Piezo vs Frequency (log amplitude, cycles of phase vs Hz)

We implemented a bank of digital filters in the AdWin software that could be tuned in frequency and width, and used either to notch out frequencies in the proportional-derivative feedback input, or to apply narrow-band feedback with variable gain and phase.

The tuning algorithm was to start with all filters turned off, and increase the feedback gain until an oscillation appeared. A filter was tuned to this frequency and used to notch out or damp the mode. This allowed the feedback gain to be increased until an oscillation appeared at another frequency. Another filter was tuned to this frequency, and the feedback gain increased more. It was sometimes necessary to change the frequency or width of some of the early filters as more were added. Gain was increased until it was no longer possible to control the oscillations by adding filters.

Even with this more sophisticated algorithm, it was still not possible to apply enough gain to substantially increase the resonant frequency. The RMS motion with feedback on is still dominated by narrow peaks of ground motion due to local vibration sources (pumps, etc) at frequencies higher than the natural resonance.

Even though it was not possible to apply broadband feedback to control these peaks, it was possible to reduce many of them by narrowband feedback. We tuned software resonators to the frequencies of the vibration peaks,, which caused their amplitudes and phases to track the ground motion at those frequencies. Then using the measured response of the platform at each frequency, we we set the control amplitude and phase to drive an equal motion, to make the platform track the ground motion.

The best feedback results that we achieved were and RMS motion integrated down to 5 Hz of 1.5 nm. This was done with the test platform isolated from the floor by two small innertubes, which passively reduced the platform motion down to 4.5 nm. With the platform on the ground, the motion was 90 nm without feedback, and it was reduced to 5 nm with feedback (Figure 5).

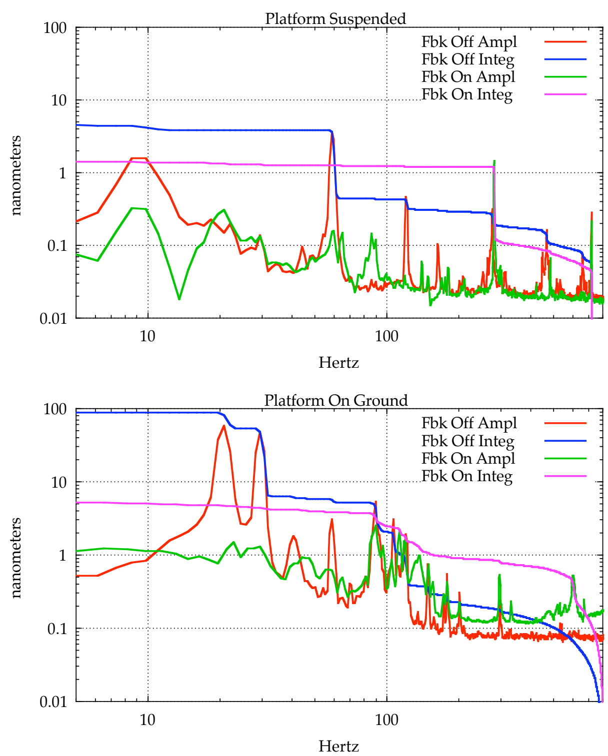

Figure 5: Feedback Performance

Our explanation for the poor broadband performance is that the piezo excites not just the fundamental resonance, but also the other mechanical resonances of the platform and baseplate. While the mass can be moved quickly by opposite kicks on successive time steps calculated to cancel for the fundamental resonance, they cannot be made to cancel for all the other resonances.

We used the measured square-wave response of the platform to find a theoretical optimum piezo command sequence. We averaged thousands of cycles of motion with the piezo driven by a 0.25 Hz square wave. We fit the first half of the data to a sum of a few decaying sinusoids at the frequencies of low resonances. We extrapolated this to the second half of the data file and subtracted. We took this as the piezo step response of the platform. We then solved for the step sequence that would produce the fastest response (minimum RMS deviation from unity during the sequence), no residual ringing (minimum RMS deviation from unity after the end of the sequence with high weight), and reasonable control power (minimum RMS command deviation from linear ramp with low weight). The resulting command sequence and the calculated platform motion are shown in Figure 6. The result indicates that we should be able to position the platform in about 5 milliseconds, about 1/4 of the period of the fundamental resonance, but much worse than the 0.5 milliseconds that would be possible at our sampling rate for a system that had only a single mode .

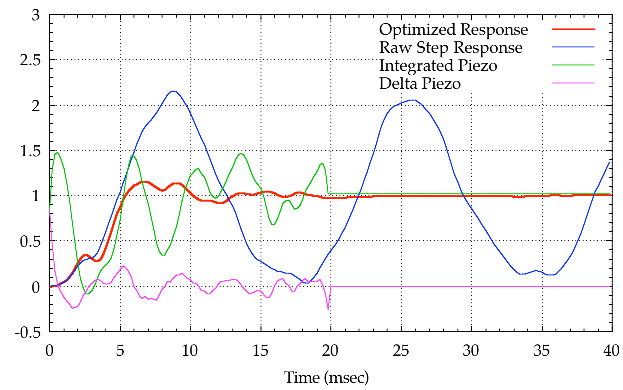

Figure 6: Calculated Optimal Piezo Sequence and Calculated Platform Response

We plan to try to improve the system response by putting the piezo between the test mass and a large reaction mass with no spring in series (there will still be a variable spring, but between the reaction mass and the end post). This should allow the test mass to be repositioned very rapidly (at the expense of moving the reaction mass the other direction). The feedback will then have to compensate for the large (but slow) motion of the reaction mass on the variable spring. We will also stiffen the base plate and grout it to the floor to reduce the effect of some of the other resonances.

4 INTERFEROMETER

The above experiments with the test platform were done using a capacitive position sensor, which has the advantage of giving an analog voltage that is directly proportional to position, and has adequate resolution. We have achieved far better position resolution by using optical interferometry. The classical Michelson configuration is shown in Figure 7. The HeNe laser light goes through a beam splitter, down two perpendicular arms with mirrors at the ends, and back to the beam splitter which recombines the beams onto a linear photodiode array. If the mirrors are adjusted so the two beams hit the photodiode array at a slightly different angles, interference fringes form. If an interferometer mirror changes position, the ringe pattern shifts. Fitting the photodiode signals allows the motion to be measured. The wavelength of HeNe light is 632 nm, so one milliradian of phase is 0.1 nm.

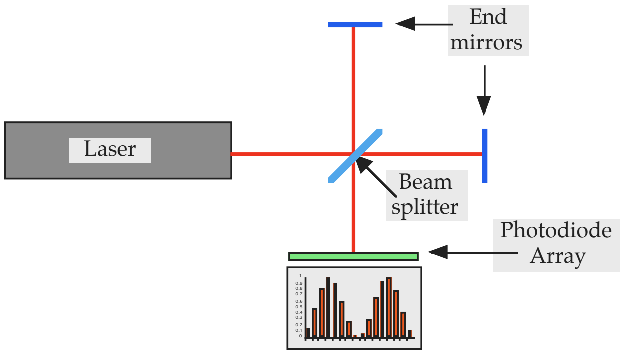

Figure 7: Michelson Interferometer

To calibrate, we use a piezo to slowly move one of the mirrors while rapidly digitizing the photodiodes. We fit each photodiode voltage vs piezo setting to an offset sinusoid. These fits (Figure 8) have noticable residuals from mirror vibrations from ground motion. The parameters from these fits are used to calculate a phase deviation that is common to all photodiodes due to vibration, for each

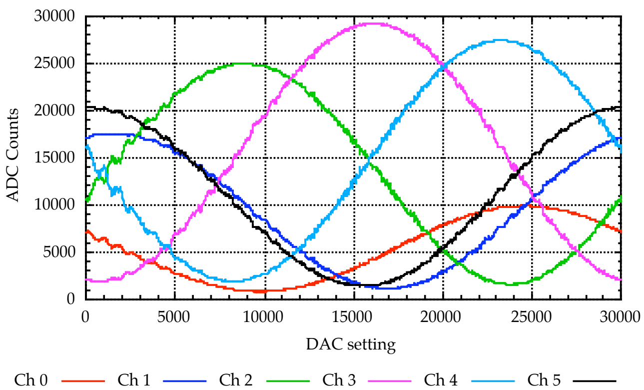

Figure 8: Photodiode Voltage vs Piezo Setting

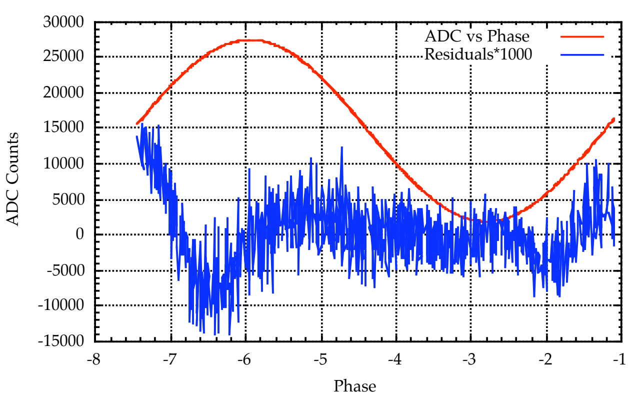

Figure 9: Photodiode Voltage vs Fitted Phase and Residuals multiplied by 1000

piezo setting. Then the raw photodiode voltages vs this phase deviation are fit to offset sinusoids. These fits have very small residuals (Figure 9).

The parameters from these fits are then used for position reconstruction. The ADC card is set up to digitize at 100 kHz, DMA the results to memory, and interrupt the CPU when the 6 photodiodes have been read. The interrupt handler does a nonlinear fit for the laser amplitude and the interferometer phase (which measures position). This can be repeated at over 5 kHz.

The position error can be estimated from the residuals of the interferometer fit. The error on the fit phase parameter, translated into a position error, is less than 1 nanometer with fresh calibration constants (Figure 10). The resolution can be improved by adjusting the calibration constants by a small fraction of the residual after each measurement. After several hundred updates (which takes only a small fraction of a second), the fit residuals indicate that the position error is less than 0.01 nanometers!

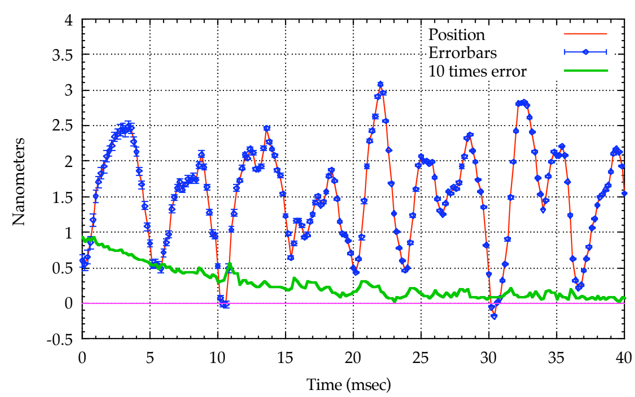

Figure 10: Interferometer Position and Error vs Time

It should be pointed out that this the resolution on the light phase, rather than the mirror position alone. The light phase is a combination of the position of all the optical elements, any air-density dependent difference in the optical path length of the arms, and any frequency drift in the laser (times the absolute difference in the arm lengths). One must take care that the optics and paths are engineered so the light phase only depends on the mirror-postion of interest. The data presented here were taken with the interferometer not connected to the test platform.

There are also resolution degradations from velocity. The photodiodes are digitized sequentially instead of simultaneously, so the resonstructed is an average over the time window. A larger problem is that the photodiode amplifiers have a time constant of 10 microseconds, so the voltage measured for any single diode is an average over that interval. If the fringes are moving, this reduces the maximum and increases the minimum voltage that will be observed. This results in estimated fit errors of several percent of the vibration magnitude for moderate frequencies.

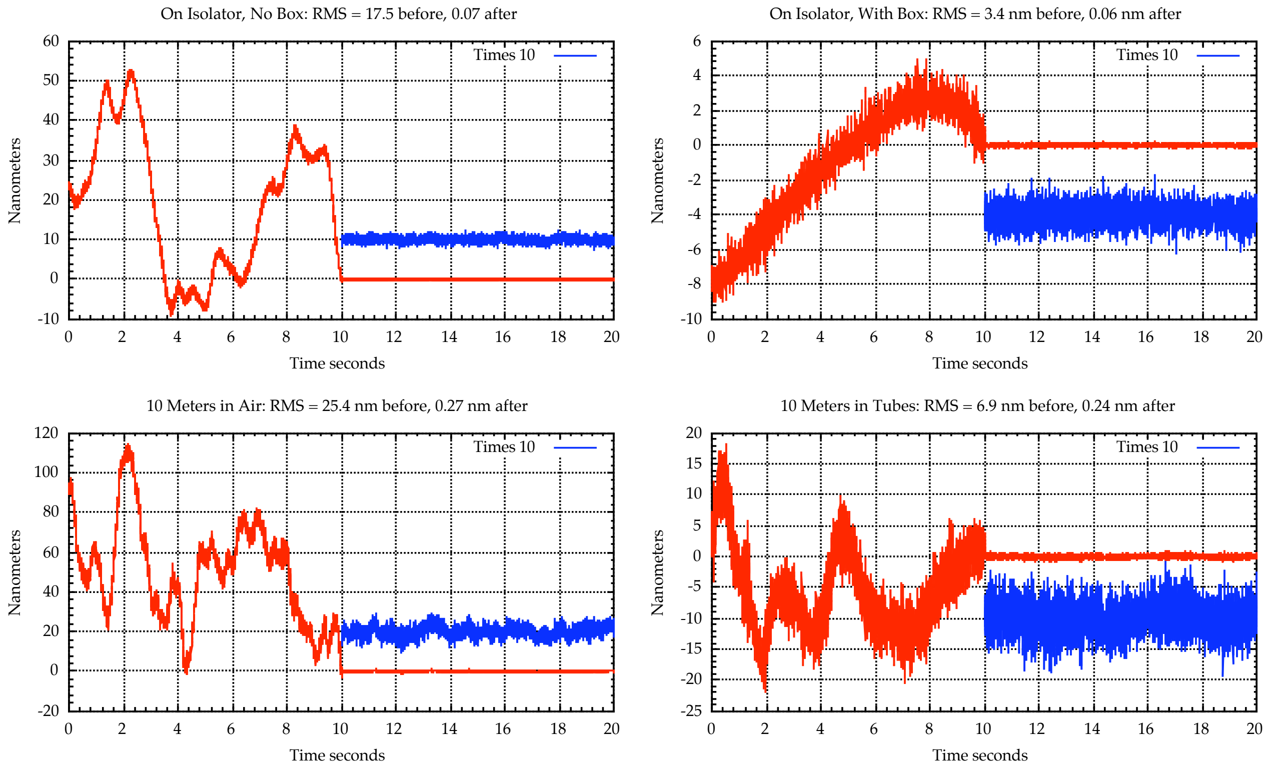

Figure 11: Interferometer Measurement vs Time
Feedback turned on after 10 seconds, feedback-on data also shown magn ified 10 times and offset

The piezo can then be used in feedback to keep the interferometer phase constant. Figure 11 shows the measured position over 10 seconds with feedback off, then 10 seconds with a simple proportional-intergral-differential feedback on. With feedback turned off, there are slow drifts that can be greatly reduced by covering the interferometer and the light paths with cardboard boxes and tubes. This is presumably air density changes due to air currents, and could probably be reduced even further by more attention to air currents. The RMS position with feedback on is as good as 0.06 nanometers with short arms, and 0.25 nanometers with 10 meter arms (folded to be parallel). The good results with feedback are a further indication that the position resolution is quite good, because feedbacks inherently amplify any noise in the signal.

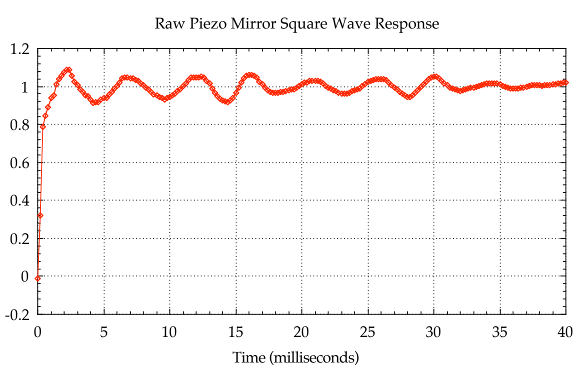

Figure 12: Measured Step Response of Piezo Mirror

The reason that the piezo-mirror feedback performs so much better than the test-platform feedback can be seen in Figure 12, which shows the response of the piezo mirror position to a step in the piezo voltage. The mirror moves 80 percent of the way in just 2 ticks (400 microseconds), and the magnitude of the ringing is less than 10 percent of the step. The time scale is the same as for Figure 6.

While the basic interferometer resolution appears to be more than adequate, there are a number of improvements possible. A non-multiplexed ADC would remove the systematic due to non-simultaneous samples, and faster photodiode amplifiers would reduce the error due to the exponentially weighted average of the recent signal. Both effects could also be incorporated into a more elaborate fit that took velocity explicitly into account, which would be natural in a more sophisticated Kalman filter in any case. More attention could be paid to enclosing the optical paths. The beam splitter and other optical elements could use more rigid mounts rather than the convenient laboratory mounts presently used.

5 CONCLUSION

We have been developing the "Optical Anchor" concept for control of vibrations in final focusing magnets for future linear colliders. Our test platform was intended to be one variable mass and one variable spring moving in one degree of freedom. Inevitably it has more vibration modes, and we did not initially appreciate how much that complicates the problem of controlling its position. We have demonstrated theoreticaly that broadband control of our particular platform is not possible at the required level. However, we have also shown that it is possible to apply narrow-band feedback to cancel motions due to coherent ground motion from cultural sources, even above the broadband control frequency limit. We plan to add a large reaction mass to the system, which should allow us to move the payload mass quickly despite a low resonant frequency of the main spring (which is necessary to fairly simulate the stiffness of a practical linear collider quadrupole mount).

We have developed calibration techniques that allow off the shelf laser, mirrors, and photodiode arrays with simple electronics, read by a Linux PC with an ADC card to make interferometric measurements with a resolution as low as 0.01 nanometers at 5 kHz. When the only thing being moved by the piezo is a small mirror, simple feedback algorithms applied to the data stream can control the mirror position to 0.06 nanometers with short interferometer arms in air, and 0.24 nanometers with 10 meter arms in air.

We are optimistic that the combination of adding a reaction mass to the test platform, and some further control software developments (incorporating standard modern control theory) will bring us to the goal of demonstrating sub-nanometer control of the position of a 100 kg mass with a 10 meter baseline.

6 ACKNOWLEDGEMENTS

A large amount of this work has been done by students at the University of British Columbia Department of Physics and Astronomy. In addition to Travis Downs who produced the interferometer feedback results and Russ Greenall who produced the test platform feedback results, previous undergraduate students Ken Yau, Jason Thompson, and Parry Fung did inportant development work along the way. This work is supported by the Canadian National Science and Engineering Research Council (NSERC).