In the Proceedings of the 1996 IFAC World Congress in San Francisco, CA, July 1996.

REJECTING ROTATIONAL DISTURBANCES ON SMALL DISK DRIVES USING ROTATIONAL ACCELEROMETERS

Daniel Y. Abramovitch *

Storage Technologies Department, Hewlett-Packard Laboratories, 1501 Page Mill Road, M/S 4U-12,Palo Alto, CA 94304-1126, USA, Phone: (650) 857-3806, E-mail: danny@hpldya.hpl.hp.com

Abstract: Small disk drives are inherently designed for portable applications and thus must be able to reject external shock and vibration. This paper expands on previous efforts at using the signal from a rotational accelerometer to minimize the the effects of these disturbances by dealing with several issues that come up: accelerometer beam resonances, low sample rate of the embedded servo on the disk drive, and widely varying accelerometer gains. The resulting algorithm is both simple and effective, making it practical for in-drive use. Experimental data is provided.

Key Words: Accelerometers, adaptive control, disk memory, disturbance rejection, feedforward compensation, multirate, vibration.

1. INTRODUCTION

1.1 Motivation

Small disk drives face several problems that have yet to become major issues for large disk drives. Previous papers have discussed the effects of friction in the actuator arm pivot (Abramovitch *et al.*, 1994; Wang *et al.*, 1994). Another issue, that was mentioned in these papers but not dealt with, was that of external shock and vibration. Small disk drives are inherently designed for portable applications. In the mobile environment, the drive must tolerate much more severe shock and vibration than typical in traditional disk drives. This translates to a requirement for additional gain at relatively low frequencies, where the shock and vibration play a more significant role.

In order to desensitize disk drives to translational disturbances, a balanced mechanical actuator has traditionally been used. However, since the actuator must pivot freely in order to access the data (as free as friction will allow) the effects of rotary disturbances about the axis normal to the disk surface can be considerable.

While there have been numerous publications on accelerometer feedforward algorithms (as will be described below), current use of accelerometers is limited to the role of a threshold detector for stopping writes of data to the disk (HP, 1993; Smith, 1993). Using a rotational accelerometer to sense rotational shock and vibration and using the accelerometer signal as the input to an auxiliary feedforward controller has been proposed as a method of making these drives more robust to rotary shocks and vibrations. This work investigates the practicality of such a scheme using an existing low-cost drive accelerometer.

1.2 Some History

The idea of using accelerometer signals to compensate for external shock and vibration of a disk drive is not new. In fact as far back as 1977, White proposed a scheme to use accelerometer signals to sense vertical shock and vibrations, and then take action to minimize the possibility of the magnetic heads slapping against the magnetic media. While improved mechanics and stiffer air bearings have minimized the need for such a system, it is interesting to note that White anticipates the two primary modes of using accelerometers in disk drives: as a simple shock-protection device and as part of a control loop. In the former mode, when the accelerometer detects a large enough shock, the magnetic heads are moved away from the disk so as to avoid possible head crashes. In the latter mode, the effect of the shock on the head-to-disk spacing is actively minimized by feeding the accelerometer signal into a control loop. (The vertical position control can be accomplished by either changing the internal pressure of the drive and thus the air bearing stiffness, or by using a servomotor on the drive arm in the vertical direction.)

More recently, the use of accelerometers for minimizing the effects of both seek reaction torque and external excitation has been studied by Davies and Sidman (1993). This work calculated analytically the filter needed to make the effect of both of these disturbances 0. Adding some practical constraints leads to a workable solution. However, knowledge of both the drive and accelerometer parameters was necessary to achieve this result. In fact, the paper tends to imply that a high-quality accelerometer is being used, as no discussion is given of accelerometer resonances in the servo bandwidth or noise. The low-pass filtering that was done was motivated to limit the accelerometer loop's gain at high frequency and thus to prevent unmodeled head disk assembly (HDA) dynamics from destabilizing the system.

One of the main practical issues in disk drives is a continual push towards lowering the manufacturing cost. Thus, it is impractical to use expensive laboratory-grade accelerometers. In the work of Knowles and Hanks (1987), a linear accelerometer was used to minimize the effect of translational shock on the position error signal. The accelerometer was mounted directly on the head disk assembly (HDA) so that both internally and externally produced disturbances could be sensed. However, each of these accelerometers needed to be calibrated in the drive during manufacture, which raised the production costs. This motivated more recent work by Hanks (1994), which has shown how accelerometers can be calibrated while the drive is in operation. This allows allow less-expensive accelerometers to be used, and shortens the manufacture time. This paper builds on these previous results, making use of both multirate control and adaptation thresholding, to improve the achievable performance.

1.3 Technical Issues

Several technical issues have to be dealt with. First, the accelerometer itself has limited bandwidth and thus the accelerometer resonances come into play. Second, the limited sample rate of the embedded servo on the drive limits the effectiveness of any feedforward compensation. Next, the manufacturing variations in the accelerometer mean that there is a large swing in the device gain from drive to drive. Finally, the need to implement any scheme on a low-cost disk drive DSP rules out the more complicated schemes that one might try.

This paper will show how each of these issues is dealt with, and shows a dramatic improvement in the drive's ability to reject rotational disturbances. The final algorithm is quite simple, so it should be easily implementable on most disk drive DSPs. Laboratory verification of the algorithm was done using the Banshee Multivariable Workstation (BMW), described previously by the author (Abramovitch, 1993).

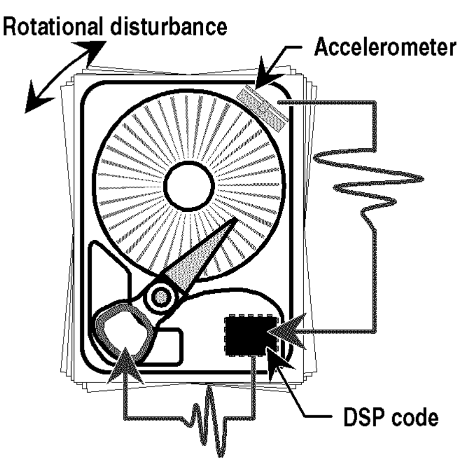

Fig. 1. Pictorial view of rotational shaking of a 1.3 inch disk drive.

2. EFFECTS OF ROTARY EXCITATION ON PES

Figure 1 shows a conceptual view of a 1.3 inch disk drive being shaken rotationally. Figure 2 shows the results of an actual experiment on a drive being shaken with a random rotary vibration in the frequency range of 50 to 500 Hz. The vibration clearly has a significant effect on the position error signal (PES) of the drive.

The goal of this work is to use an existing 1.3 inch disk drive rotational accelerometer to substantially diminish the effects of rotary shock and vibration on the disk drive control loop. The block diagram for this is shown in Figure 3.

As shown in Figure 3, the main control loop of the disk drive includes the Electro-Mechanical System (plant) which includes both the actuator components (power

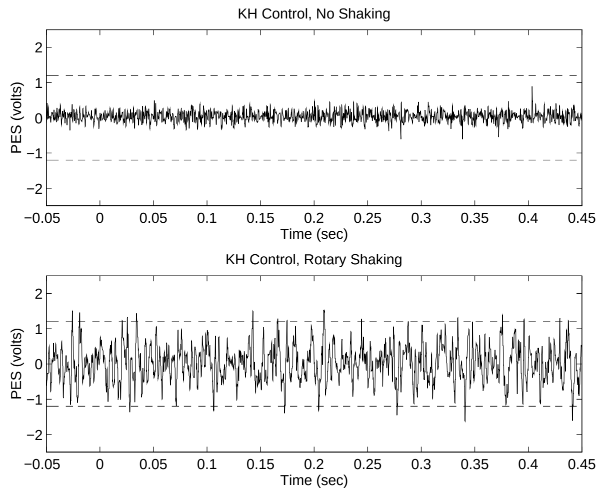

Fig. 2. The Position Error Signal is considerably affected by rotational vibration. On this plot 1.2 volts equals 12% of a track. PES is the position error signal. The frequency range of the excitation is 50 to 500 Hz. The vibration has an rms value of 86.5 rad/s^2.

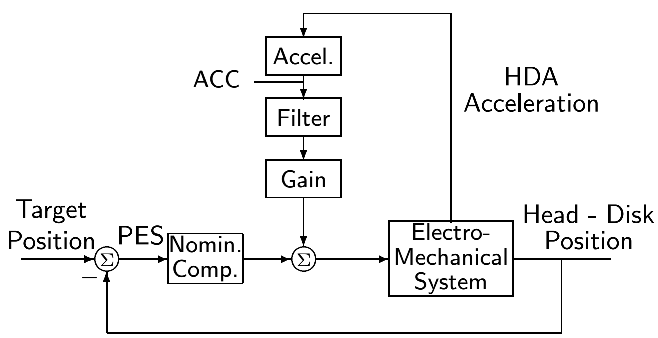

Fig. 3. Block diagram of feedforward cancellation

amplifier, voice coil motor, actuator arm, suspension, and read/write head) and the HDA components (disk platters, spindle, spindle motor, and base plate). The Electro-Mechanical System determines the relative position difference between the read/write head and the center of the target track. This difference is subtracted from the Target Track Position (or simply Target Position) to form the drive's Position Error Signal (PES). PES enters the Nominal Compensator to produce a command input to the Electro-Mechanical System.

An auxiliary loop is formed by sensing the acceleration of the Head Disk Assembly (HDA). In particular, for this paper, the rotational acceleration is sensed. This acceleration signal (ACC) is filtered and passed through a gain stage before being summed into the command signal. Thus, the aforementioned goal of this paper can be restated as designing this auxiliary loop so that rotational HDA acceleration is decoupled from PES.

The rotational accelerometer is mounted on the drive base plate, rather than the actuator arm, so it senses motion of the drive's head disk assembly (HDA). Modulo actuator pivot friction, the actuator can be considered to be floating free of the HDA; thus when the HDA is bumped rotationally in the plane of the disk[1], the actuator will stay still in inertial space and a position error will result. The drive feedback loop can reject some of the disturbance, but may lack the gain to reject large amounts of it. Furthermore, the phase lag of having to go through the feedback loop diminishes the cancellation ability of the drive. However, by feeding forward the accelerometer signal — at the proper gain and phase — into the position control loop, the actuator can be made to follow the rotational disturbance.

The natural reaction might be to "Just do it", *i.e.* to simply use the accelerometer signal in a feedforward fashion. This is not possible, as there are technical issues to be overcome:

1) The accelerometer has limited bandwidth due to a resonance.
2) The disk drive has an embedded servo, *i.e.* the position information is interleaved with the user data. While this itself does not guarantee a low sample rate, the small geometry of the drive, combined with the need to preserve space for user data, forces the drive to have a relatively low sample rate (3717 Hz).
3) The built in accelerometer has a large gain variation (± 50%) from unit to unit.

3. DEALING WITH THE ISSUES

3.1 Filtering and sample rate

The rotary accelerometer in question consists of two cantilever beams, sensed differentially, as shown in Figure 4 (perspective view) and Figure 5 (top view). The differential sensing of the signal outputs is used to cancel the individual responses to translational motion, while boosting the signal output during rotational motion.

The first bending mode of the accelerometer beam puts a limit on the frequency that can be followed. The beam

[1] Or if the bump has a rotational component to it.

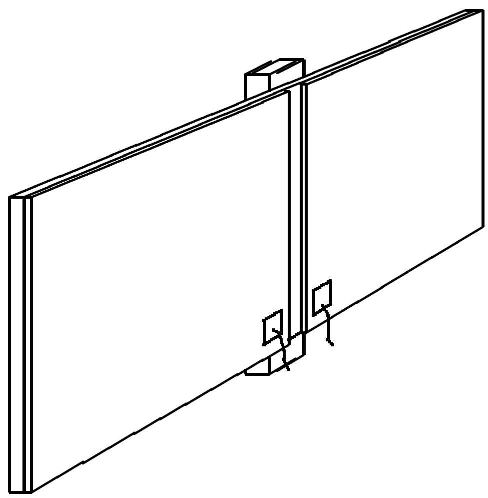

Fig. 4. Rotary accelerometer. The beam has two separate pieces of piezo-electric material bonded to it. Flexibility in the beam results in a voltage at either pad where the wires are attatched.

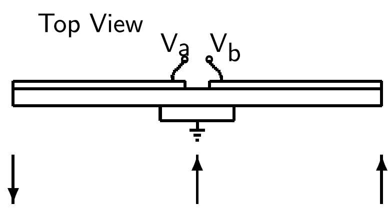

Fig. 5. Top view of rotary accelerometer. Rotary acceleration is sensed by subtracting $V_a - V_b$. Translational acceleration causes the two sides to move in-phase and is rejected by $V_a - V_b$.

resonance can be handled by filtering. However, at the low sample rate of the sectored servo, a significant phase lag accompanies most filter designs, as is shown in Figure 6. It becomes apparent from experimental observation that having a very small phase angle about 0 is as critical as having the proper gain setting. Thus, the above phase lag limits the benefit of the accelerometer feedforward to relatively low frequencies.

Noting that the sample rate of position sensing is limited by the sectored servo, but that the accelerometer has no such limit, a multi-rate scheme can be employed. As shown in Figure 7, the accelerometer signal is sampled not only when the position error signal is sampled, but at several time instants in between. This extra sampling

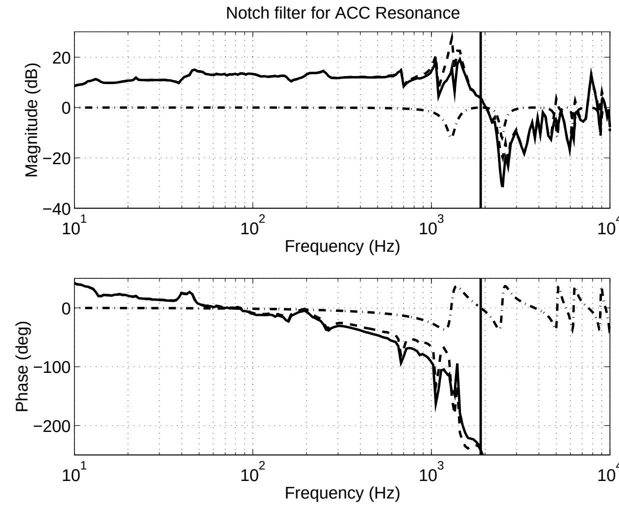

Fig. 6. Frequency response of accelerometer and notch filter. The accelerometer signal is sampled at the nominal rate. The dashed line represents the accelerometer frequency response, the dashed-dotted line is the frequency response of the filter, and the combined response is shown by the solid line. The vertical solid line is at the Nyquist frequency of 1858.5 Hz.

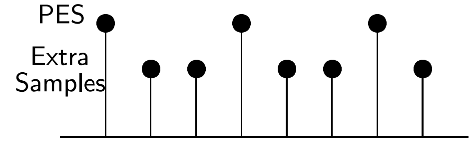

Fig. 7. Extra accelerometer samples.

of the accelerometer has no effect on the user overhead of the drive. All that is required is the CPU bandwidth (in the DSP) to do it. In the case shown in Figure 8, sample rate of 4 times the nominal (position) sample rate is used on the accelerometer. The higher sample rate on the accelerometer:

- broadens the usable bandwidth,
- reduces the phase lag of the filters used, and
- gives more freedom in filter design.

The filter design shown in Figure 8 includes a notch to damp the resonance, a low-pass filter to roll off the filter gain above the Nyquist rate, and a lead to restore some of the phase in the 100–600 Hz range. This results in some significantly improved filter designs and performance. A comparison of frequency response functions from disturbance to position error signal (PES) is shown in Figure 9. In this plot, the further down the plot, the more suppressed the disturbances are. Note that with

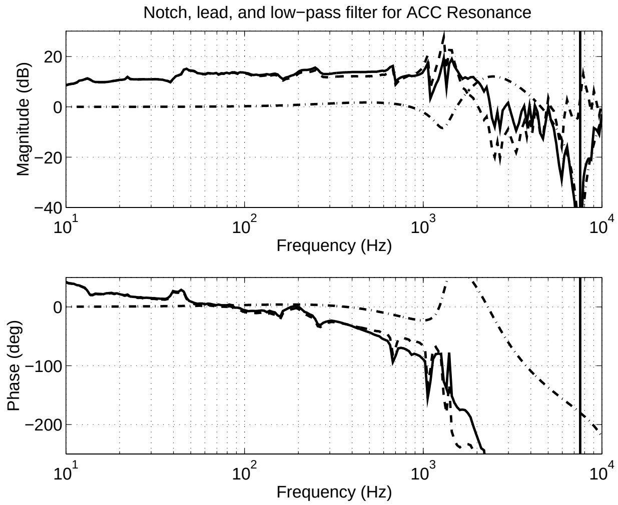

Fig. 8. Frequency response of accelerometer plus lead, notch and low-pass filter. The accelerometer signal is sampled at 4 times the nominal rate. The dashed line represents the accelerometer frequency response, the dashed-dotted line is the frequency response of the filter, and the combined response is shown by the solid line. The vertical solid line is at the Nyquist frequency of 7434 Hz.

the proper filtering and a high enough sample rate, the in-drive accelerometer can nearly match the laboratory grade accelerometer over a large freqency range in its usefulness for disturbance rejection. The time-domain performance of the filter from Figure 8, used in feedforward of the accelerometer signal (ACC), is shown in Figure 10.

Implementing multirate in this case is not conceptually difficult. As shown in Figure 11, the filters can be designed separately. While this simple procedure may not take full advantage of the interrelationship between the PES and accelerometer signal, it is quite effective and straightforward to program on a DSP. In the structure in Figure 11 the nominal PES loop is clocked only by the disk sector. The accelerometer loop is clocked both by the disk sector and by an extra clock for the inter-sector samples. In this case, the sample rate multiplier for the accelerometer is designated by M.

3.2 Adapting the accelerometer gain

Finally, there is the issue of the varying gain of the accelerometer (± 50% from drive to drive). It is natural to want to adapt the filter gain to compensate for the accelerometer gain. Note that one thing about the accelerometer does make the problem simpler. That is,

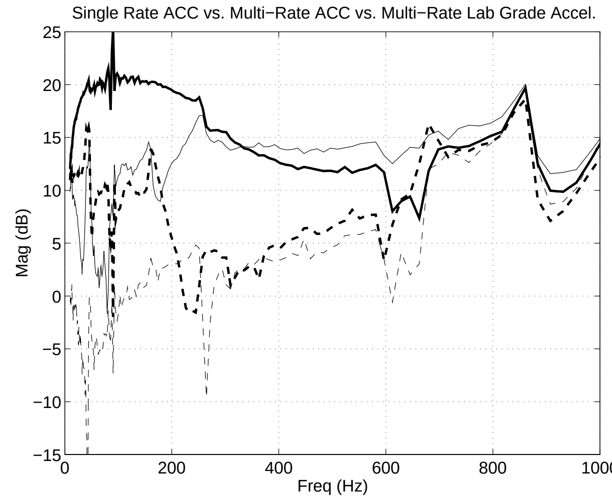

Fig. 9. Comparison of frequency response functions from disturbance to PES. In this plot, the further down the plot, the more suppressed the disturbances are. The thick solid line is without feedforward. The thin solid line is with the on drive accelerometer (ACC) being fed forward with single rate sampling. The thick dashed line is with ACC being fed forwad with 4 X multi-rate sampling. The thin dashed line is with the signal from a laboratory grade accelerometer being fed forward with 4X multi-rate sampling. Note that with the proper filtering and a high enough sample rate, the in-drive accelerometer can nearly match the laboratory-grade accelerometer over a large freqency range in its usefulness for disturbance rejection.

the shape of the accelerometer response is approximately constant. Only the accelerometer gain is highly variable. This allows one to design one filter and merely adapt the gain. In the structure shown in Figure 11 that means that the filter, $\mathsf{C}_{\mathsf{ACC}}$, can remain constant while only the gain, $\mathsf{K}_{\mathsf{ACC}}$, is adapted.

Returning to Figure 3, there are two choices for how to proceed with adaptation. The first method, which is probably more intellectually appealing, is to identify the entire accelerometer response. However, given that the available signals are the PES, the accelerometer signal (ACC), and the signal being sent to the DAC, this approach would require estimating the head-disk-assembly (HDA) acceleration from the PES signal and the compensator output. This could then be used to form an error signal with the measured ACC signal and thus, the accelerometer could be adapted. Note that a major challenge with this approach is that estimating the HDA acceleration from PES and the output to the DAC

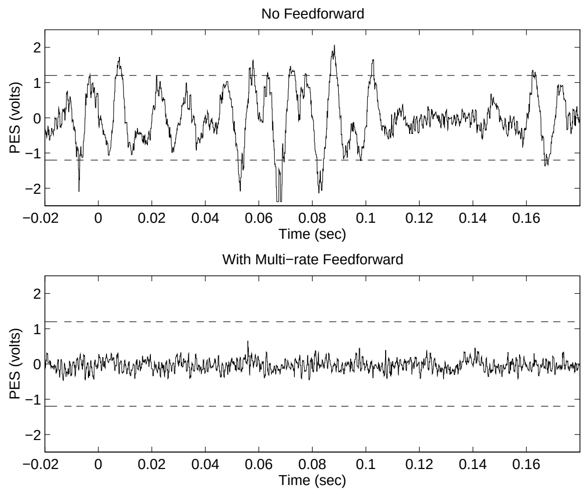

Fig. 10. Multi-rate feedforward of in-drive rotational accelerometer signal (ACC). ACC is sampled at 4X the nominal sample rate. This is with a tuned filter gain.

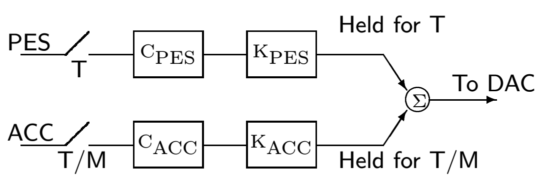

Fig. 11. Compensator structure. $C_{(\cdot)}$ is the compensator filter, $K_{(\cdot)}$ is the gain.

involves estimating second derivatives. There are often issues with noise sensitivity of such a technique.

A far simpler method involves simply understanding what the requirements for the system are. What is really desired is to remove the effects of HDA acceleration from PES. Of course, the ever popular Widrow-Hoff LMS (Least Mean Squares) algorithm (Widrow and Stearns, 1985) essentially decorrelates signals, *i.e.* it removes the effects of one from another. This is the method suggested in Hanks' work (1994). Going slightly beyond this, the algorithm below adds some upper and lower bounds on the values that the gain can adapt to, and

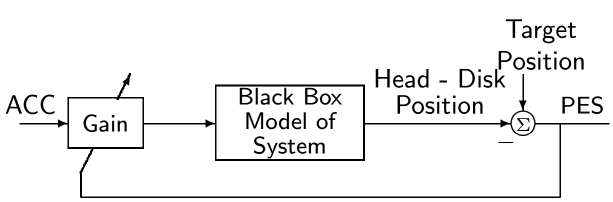

Fig. 12. Block diagram of adaptation.

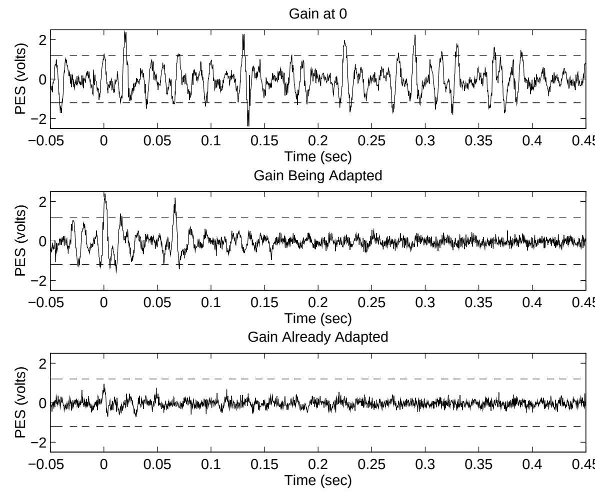

Fig. 13. Adaptation of accelerometer filter gain. Note the rapid decrease in the magnitude of the position error. The disturbance is random rotary vibration between 50 and 150 Hz with a level of 57.7 rad/s^2 rms.

includes logic to adapt the accelerometer gain only when there is a large enough signal at the accelerometer to guarantee that the algorithm is not adapting to noise. It turns out that this rather simple method, based on an understanding of what the control system is trying to achieve, works extremely well. Thus, the moral might be that occasionally when one has thought about the problem long enough, one can "Just do it."

The LMS algorithm involves one equation, namely:

$$w_{k+1} = w_k + 2\mu\varepsilon_k x_k. \qquad (1)$$

In terms of the simplified block diagram of Figure 12 LMS is implemented as:

$$(\mathrm{Gain})_{k+1} = (\mathrm{Gain})_k + 2\mu\,(\mathrm{PES})_k\,(\mathrm{ACC})_k\,, \qquad (2)$$

where ACC is the signal coming off the uncalibrated accelerometer (replacing x in Equation 1) and PES is the drive's Position Error Signal (replacing ε in Equation 1). Note that Figure 12 is the block diagram for the LMS algorithm in this system, and is simpler than the system block diagram of Figure 3. In fact, this algorithm is so simple that it involves only 6 instructions in a DSP to code LMS. Another 4 instructions are used for limiting the upper and lower bounds of the gain. Finally, a thresholding routine is added which only allows adaptation when the system is actually being shaken. This merely checks to make certain that the drive is being shaken before doing any adaptation of the accelerometer gain. This consumes another 6 instructions. Thus,

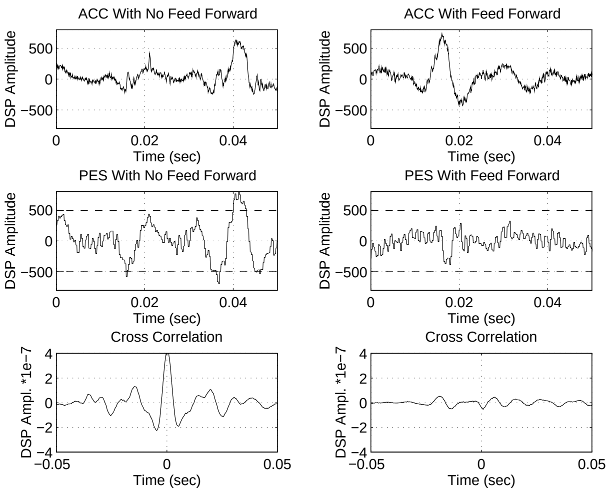

Fig. 14. Comparison of signal correlations. The top two plots have the accelerometer signals, the center two plots have the PES signals, and the bottom two have the cross-correlation between the two. On the left-hand side it is clear that with no feedforward ACC and PES are highly correlated. On the right-hand side, it is clear that once the algorithm has adjusted the gain, the correlation is considerably reduced.

16 DSP instructions yields an effective and robust adaptation algorithm. Because this simple adaptive scheme is effective, it is reasonable to implement this on a disk drive DSP. An example that shows the effectiveness of this algorithm is shown in Figure 13. One can also see that the LMS algorithm does in fact decorrelate ACC and PES by looking at Figure 14. The top two plots have the accelerometer signals, the center two plots have the PES signals, and the bottom two have the cross-correlation between the two[2]. On the left-hand side it is clear that with no feedforward, ACC and PES are highly correlated. On the right-hand side, it is clear that once the algorithm has adjusted the gain, the correlation is considerably reduced.

3.3 Shocks and intermittent excitation

It was mentioned earlier that a thresholding routine was added to the adaptation mechanism. It turned out that a fairly simple routine, merely checking the magnitude of the ACC signal against some minimum level, was enough to keep the accelerometer filter gain from drifting away from its "converged" value. Thus, in Figure 15 the intermittent shocks are rejected when adaptation is allowed. The net effect of the thresholding routine is that

[2] Calculated with the Matlab *xcorr* function.

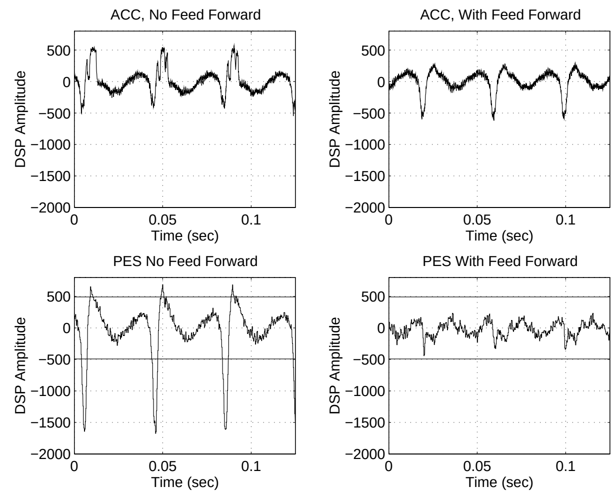

Fig. 15. Adaptation during intermittent 3 mS shock pulses. Note that even though the magnitude of the pulses does not change, their effect on the system does. The thresholding routine keeps the parameters from drifting in between shock pulses.

as these intermittent shocks are spaced further and further apart, the thresholding will keep the accelerometer filter gain at a value that was converged to during the rotary shocks. If the thresholding were not present, then the accelerometer gain would adapt to whatever signal noise was present on the line. Eventually, the accelerometer filter gain could be detuned so that when another rotary shock came along it would have a severe effect on PES.

3.4 System requirements

In a competitive market product, a major issue for any algorithm is what it costs in terms of system complexity (and therefore production cost). The breakdown can be applied as follows:

- Accelerometer feedforward requires a rotational accelerometer.
- Multi-rate compensation requires 1 clock, direct DSP access to ADC and DAC, and CPU bandwidth.
- Adaptation requires 16 DSP instructions and excitation,*i.e.* shaking.

None of these requirements is extreme. Given that a rotational accelerometer is on the drive already, feedforward control using that signal is a possibility, as stressed by other authors (White, 1977; Knowles and Hanks, 1987; Davies and Sidman, 1993). The addition of multi-rate sampling of the accelerometer is quite important when the the base sample rate of the control system is limited. The requirements for this are rather modest by modern disk drive standards. Furthermore, the self-calibration of the accelerometer, as proposed by Hanks (1994) and expanded upon here, removes the added cost of having to calibrate each accelerometer in the drive at manufacture time. The fact that the algorithm presented here requires a relatively small number of DSP instructions makes it extremely practical for actual production drives. Thus, there is a strong belief by the author that this method is readily implementable in most disk drives.

4. CONCLUSIONS

This paper has shown how the signal from a rotational accelerometer can be used to substantially diminish the effects of rotational shocks and vibrations from a small disk drive. Several practical issues have been dealt with, such as accelerometer beam resonances, slow base sample rate of the servo system, large variation in the accelerometer gain, and ensuring excitation during adaptation. The resulting algorithm is quite simple and readily implementable in most disk drive DSPs. The effectiveness of this algorithm has been displayed in laboratory trials on 1.3 inch disk drives.

References

Abramovitch, Daniel, Feei Wang and Gene Franklin (1994). Disk drive pivot nonlinearity modeling Part I: Frequency Domain. In: *Proceedings of the 1994 American Control Conference*. AACC. IEEE. Baltimore, MD. pp. 2600–2603.

Abramovitch, Daniel Y. (1993). The Banshee Multivariable Workstation: A tool for disk drive research. In: *Advances in Information Storage Systems, Vol. 5* (Bharat Bhushan, Ed.). pp. 59–72. ASME Press. ISBN 0-7918-0031-8.

Davies, David B. and Michael D. Sidman (1993). Active compensation of shock, vibration, and wind-up in disk drives. In: *Advances in Information Storage Systems, Vol. 5* (Bharat Bhushan, Ed.). pp. 5–20. ASME Press. ISBN 0-7918-0031-8.

Hanks, D. Mitchel (1994). Apparatus and method for shock attenuation in a disk recording and/or reproduction system using variable gain acceleration sensor. United States Patent 5,299,075. Hewlett-Packard Co.. Corporate Patent Department, M/S 20B-O, 3000 Hanover Street, Palo Alto, CA 94304 USA.

HP (1993). *HP Kittyhawk Personal Storage Modules Product Brief.*

Knowles, Vernon L. and Darwin M. Hanks (1987). Shock and vibration disturbance compensation system for disc drives. European Patent Application 871065555.3. Hewlett-Packard Co.. Corporate Patent Department, M/S 20B-O, 3000 Hanover Street, Palo Alto, CA 94304 USA.

Smith, Robert F. (1993). Apparatus for sensing operating shock on a disk drive. United States Patent 5,235,472. Seagate Technology, Inc.. Scotts Valley, CA.

Wang, Feei, Terril Hurst, Daniel Abramovitch and Gene Franklin (1994). Disk drive pivot nonlinearity modeling Part II: Time Domain. In: *Proceedings of the 1994 American Control Conference*. AACC. IEEE. Baltimore, MD. pp. 2604–2607.

White, Robert A. (1977). Shock force compensating system. United States Patent 4,040,103. Sperry Rand Corporation. New York, NY USA.

Widrow, B. and S. D. Stearns (1985). *Adaptive Signal Processing*. Prentice-Hall. Englewood Cliffs, New Jersey.