

Rejecting Oscillatory, Non-Synchronous Mechanical Disturbances in Hard Disk Drives

Rick Ehrlich, Jason Adler, and Haitham Hindi

Abstract—This paper describes a “textbook” technique for rejecting narrow-band disturbances to a hard-disk-drive servo loop, by including a model of the disturbance in the drive’s state-space controller, and then canceling it. Both numerical simulations and experimental evidence are presented which show significant reductions in tracking errors due to the targeted disturbance in steady-state operation. Simulations of the transient behavior of the augmented control loop, however, show some undesirable characteristics.

Index Terms—Disturbance-rejection, HDD, servo.

I. INTRODUCTION

THE TMR (Track Mis-Registration) suffered by a hard-disk-drive while in ontrack mode contains some components that are synchronous with the spindle rotation, as well as some that are not. Typically, a large portion of the nonsynchronous TMR is contained within narrow frequency-bands [1], which invites speculation as to whether it can be identified and rejected by a state-space controller that includes a model of the oscillatory disturbance that caused it [2]. This paper describes such a controller, and shows models of its behavior during simulated disk-drive operation. It is seen that, while the steady-state rejection of the disturbance is very good, the transient behavior of the system is poor.

Most previously published work in the area of HDD disturbance rejection has dealt with cancellation of synchronous disturbances [3]. While there are significant similarities between the problems of rejecting synchronous and nonsynchronous disturbances, the amplitude and phase of the synchronous disturbances typically change very slowly with time, allowing extremely low adaptation gains, with correspondingly small effects on the characteristics of the overall servo loop. Because many nonsynchronous disturbances are driven by white-noise processes, their more rapidly changing magnitude and phase require relatively high adaptation gains which can significantly affect the stability and transient response of the overall servo loop.

Section II describes the model used for simulation of the disturbance controller. The estimator/controller design technique is contained in Section III. The results of simulations of the steady-state behavior of the system are presented in Section IV, followed by experimental data from a prototype HDD in Section V. Section VI uses the model to predict




Fig. 1. Control loop containing a plant with measurement noise and both colored and white disturbances

a number of standard transient responses (to a step in the measurement noise or a pulse in the disturbance torque), as well as the start-up transient that the system experiences when the disturbance cancellation is suddenly turned on. Section VII contains conclusions.

II. MODEL

A simplified model of a discrete-time HDD servo loop, including an oscillatory disturbance, is shown in Fig. 1. The block labeled, “actuator,” represents the power-amplifier and actuator by the transfer-function,

$$P(z) = \frac{1}{2} \frac{z+1}{2(z-1)^2}. \quad (1)$$

In other words, the power-amplifier is assumed to be ideal and the actuator is assumed to be perfectly rigid, suffering no friction or mechanical resonances. The input of $P(z)$ has units of Tracks/Sample² (requested actuator acceleration) and the output has units of tracks.

The block labeled, “disturbance” represents a filter which has a white-noise input and whose output is added to the plant input (“process noise”). The transfer-function, $D(z)$, is the zero-order-hold equivalent of the continuous-time filter, $D_C(s)$,

$$D_C(s) = \frac{\omega_0^2}{s^2 + 2\zeta\omega_0s + \omega_0^2} \quad (2)$$

where ω_0 is the center-frequency of the disturbance filter and ζ is its damping coefficient. The disturbance (produced at the output of that filter) is modeled as an additional (unknown) command to the power-driver.

The block labeled “controller,” with transfer-function, $C(z)$, contains a state-space model of both of the other blocks, $P(z)$ and $D(z)$, with which it estimates the 4 states of the plant and computes a linear state-feedback once per servo sample.

Manuscript received July 7, 2000.

R. Ehrlich and J. Adler are with the Quantum Corporation in Milpitas, CA, USA (e-mail: rick.ehrlich@quantum.com).

H. Hindi is with Stanford University in Palo Alto, CA, USA.

Publisher Item Identifier S 0018-9464(01)02269-5.

The equations which govern the behavior of that current-estimator/controller are well known [2], but are repeated below for later reference in this paper.

$$\hat{x}_k = \bar{x}_k + L(y_k - H\bar{x}_k) \quad (3)$$

$$u_k = -K\hat{x}_k \quad (4)$$

$$\bar{x}_{k+1} = F\hat{x}_k + Gu_k. \quad (5)$$

In (3)–(5), “ k ” represents the sample-index, “ F ,” “ G ,” and “ H ” are the system matrix, input-gain matrix, and output-gain matrix, respectively, “ \bar{x} ” and “ \hat{x} ,” are the predicted and corrected system state-estimates, and “ K ” and “ L ” are the servo’s controller and estimator gains.

For the example system considered in this paper, $\omega_0 = 2 * \pi * 500 \text{sec}^{-1}$, and $\zeta = 0.016$. The noise input, v , and the disturbance inputs, w_1 , and w_2 , were all assumed to be white Gaussian random variables with variances, 0.01 Tracks, 0.001 Tracks/Sample², and 0.002 Tracks/Sample², respectively. The sample-rate of the system was assumed to be 10 kHz, giving a sample-time, T , of 100 μsec . The authors chose to scale the internal states of the estimator such that the first state represented the position of the actuator (the signal labeled “ e ” in Fig. 1), the second state represented the actuator’s velocity (\dot{e}), the third state represented $1/(\omega_0 T)$ times the derivative of the disturbance¹, and the fourth state represented the effective disturbance current (the output of the disturbance filter in Fig. 1). With that scaling, and the parameters given above, the values of F , G , and H , were:

$$F \cong \begin{bmatrix} 1 & 1 & 0 & 0.5 \\ 0 & 1 & 0 & 1 \\ 0 & 0 & 0.9414 & -0.3075 \\ 0 & 0 & 0.3075 & 0.9512 \end{bmatrix} \quad (6)$$

$$G = \begin{bmatrix} 0.5 \\ 1 \\ 0 \\ 0 \end{bmatrix} \quad (7)$$

$$H = [1 \ 0 \ 0 \ 0]. \quad (8)$$

The simulations described below assumed perfect knowledge of the disturbance center frequency, but the authors found that the results were not extremely sensitive to the accuracy of the estimated frequency. A practical application of this technique might need to include either prior identification or on-line adaptation of the compensation to the actual disturbance.

III. ESTIMATOR/CONTROLLER DESIGN

All of the estimator/controller designs, as well as all of the simulations, were performed using MATLAB Version 5.3.0.101 83, along with Matlab’s Control System Toolbox, Version 4.2 [4]. The estimators were designed using the Control Toolbox function, “Kalman,” which determines the steady-state Kalman filter estimator gains for a given plant and set of noise/disturbances. The controllers were designed using

¹The scale factor of the third state was chosen so as to keep the third state’s magnitude similar to that of the fourth state.

the Control Toolbox function, “dlqr,” which determines the controller gains which minimize the quadratic cost function, $J(u)$:

$$J(u) \equiv \sum_{k=0}^{\infty} (x_k^T Q x_k + u_k^T R u_k) \quad (9)$$

where Q is a state-weighting matrix and R is a control-weighting matrix supplied by the user. For this paper, the control weighting matrix was set to unity, and Q was set to

$$Q = \rho H^T H. \quad (10)$$

In other words, the dlqr function was asked to minimize a weighted sum of the squares of the control effort and plant output. The remaining scalar constant, ρ , was adjusted for each case to give an overall system open-loop crossover frequency of 900 Hz.

IV. SIMULATIONS OF STEADY-STATE BEHAVIOR

The authors compared the modeled TMR obtained with two different LQG (**L**inear-**Q**uadratic **G**aussian) controllers. The first (referred to later as the “simple” controller) was made “aware” of only the two rigid-body states of the actuator, while the second controller (later referred to as the “enhanced” controller) was designed with all 4 plant states “in mind.” For the first system, the F , G , and H matrices used in the estimator/controller design were upper/left portions of the matrices shown in (6)–(8). With ρ set equal to 60^2 , the open-loop crossover of the system was roughly 900 Hz. The estimator and controller gains for that system were:

$$L_{Simple} \cong \begin{bmatrix} 0.4673 \\ 0.1460 \end{bmatrix} \quad (11)$$

$$K_{Simple} \cong [1.3595 \ 1.6490]. \quad (12)$$

For the second controller, a 900 Hz open-loop crossover was obtained by setting $\rho = 5.3$. The estimator and controller gains for that system were:

$$L_{Enhanced} \cong \begin{bmatrix} 0.5307 \\ 0.1984 \\ -0.0164 \\ 0.0108 \end{bmatrix} \quad (13)$$

$$K_{Enhanced} \cong [0.6335 \ 1.1256 \ 0.0127 \ 1.0270]. \quad (14)$$

Because the fourth estimator state is the value of the oscillatory disturbance, one might expect that the third and fourth controller gains in (14) would be 0 and 1, respectively, so as to exactly cancel the (estimated) disturbance. The third gain is small, and the fourth is close to, but not exactly equal to, 1. This is due to the fact that the LQR design optimizes a weighted sum of (the squares of) the control effort and tracking error. If the value of ρ chosen for the enhanced controller was very large (so as to emphasize reduction of tracking-error over reduction of control effort), then the third gain would tend toward zero and the fourth would approach 1.

²For the normalization chosen by the authors, the units of ρ are Samples⁻⁴.




Fig. 2. Plots of open-loop gain and open-loop phase-Margin for the two servo loops modeled. (a) is for the loop that was designed without knowledge of the oscillatory disturbance, and (b) is for the loop whose controller was “aware” of the disturbance.




Fig. 3. Comparison of modeled TMR spectra obtained with second-order state-space controller (a) with that obtained from the fourth-order one (b).

Plots of the open-loop gain and open-loop phase-margin (open-loop phase, plus 180 degrees) for the two systems are shown in Fig. 2. Examination of the plots in Fig. 2 shows that the phase-margin of the enhanced control loop is about 31 degrees; considerably lower than that of the original system’s margin of about 45 degrees. The TMR spectra for the two systems are shown in Fig. 3. Each spectrum in that figure is the RMS of 100 simulations, and is scaled by the square-root of the number of time-points per trace (so that the RMS of the plotted spectrum is equal to the RMS of the measured position-error). Each plot also contains a “baseline” TMR spectrum, which is the root-sum-square of the TMR due to the white torque disturbance and the PES noise. The RMS measured position (the signal, y , in Fig. 1) for the loop that contained a state-space model of the oscillatory disturbance was about 21% less than that of the loop that had the simpler controller. The RMS value of the actual position (the signal, e , in Fig. 1) was about 29% less with the enhanced controller than it was with the simple one.

V. EXPERIMENTAL VERIFICATION

The authors augmented the estimator/controller of a laboratory prototype drive to include the ability to simultaneously

estimate and reject two oscillatory disturbances of the sort described above. For practical reasons, the experimental code was limited to model only un-damped oscillatory disturbances ($\zeta = 0$), and to cancel them exactly [in other words, for each target disturbance the state-feedback gains for the two resonant states, $K(3)$ and $K(4)$, were 0 and 1, respectively]. The track-density of the drive was about 780 000 tracks/m ($\sim 20\,000$ TPI, or “Tracks-Per-Inch”). Both the actuator and the original estimator/controller were of higher order than the second-order actuator modeled above, but the results were similar to those obtained with the simplified model. Fig. 4 shows a portion of the measured PES (Position-Error-Signal) spectrum both with and without the resonant disturbance cancellation turned on. As was the case with the model, the modified controller greatly reduced the TMR contribution of the oscillatory disturbances, and amplified the contributions of the other sources at nearby frequencies. The net reduction in nonsynchronous TMR for the experiment was about 18% (from 1.59% of a track RMS TMR to 1.31%).

VI. SIMULATIONS OF TRANSIENT BEHAVIOR

The results of the previous sections show that the controller is able to reject an oscillatory disturbance very well in steady-state.




Fig. 4. Comparison of experimentally measured TMR spectra obtained with low-order state-space controller (a) with that obtained from the higher-order one, which estimates and attempts to cancel two oscillatory disturbances (b).




Fig. 5. Transient in RMS value of measured position-error as the servo controller is switched from the simple one to the enhanced one. The settle-transient is essentially complete in about 5 msec.

For most purposes, however, disk drive servos must function well immediately after transient events, such as seeks from one track and head to another.

Some oscillatory disturbances (disk-vibrations, in particular) may be effectively un-correlated from one head to another. In that case, any state estimates pertaining to the disturbance that were obtained before the head-switch are useless afterwards, and the best initial value for the disturbance state-estimates is zero.

It is of interest to see how quickly the enhanced controller reaches its improved TMR (relative to the simple controller), once it is switched on. Fig. 5 shows the predicted RMS value of y versus time for the situation in which the servo control is switched from the simple controller to the enhanced one. That prediction was made by simulating the evolution of the state covariance matrix for the entire 8-state closed-loop system (4 states for the plant, 4 more for the estimator), with all three inputs (w_1 , w_2 , and v) [5]. The simulation was begun with a covariance matrix of all zeros (with the system at rest), and the

simple controller³, and was run for 2000 samples, at which point all startup transients had died out. Then, the estimator and controller gains were switched to those of the enhanced controller [shown in (13) and (14)], and the simulation was run for another 100 samples. As shown in Fig. 5, the expected RMS tracking error actually gets worse for awhile, and then settles down to a lower value. The transient has an effective time-constant of about 2 msec.

The plot in Fig. 5 implies that for a drive sitting on a track in steady-state, with no prior knowledge about the phase or magnitude of an oscillatory disturbance (but with perfect knowledge of its frequency), it would take about 5 msec for the enhanced controller to significantly reduce the TMR due to that disturbance. This actually represents an optimistic lower bound on the time interval between the end of a seek and full reduction of the expected TMR. Because of effects of slight errors in the mechanical model combined with the very large actuator commands involved in seeks, any adaptation of phase and magnitude of an oscillatory disturbance would have to be turned off during a seek. In fact, because it can take some time for the drive to “settle-out” such estimation errors after a seek, it is likely that cancellation of the oscillatory disturbance would have to be postponed for some time after the end of each seek. Turning off the correction updates can be easily accomplished by simply setting $L(3)$ and $L(4)$ equal to zero.

To give a general idea of the level of sensitivity of the enhanced controller to transient disturbances, the authors examined the transient response of the system to two standard excitations. Fig. 6 compares the responses of the simple and enhanced controllers to a step-change in the measurement noise, v , and Fig. 7 makes a similar comparison for a single-sample pulse in the disturbance, w_2 .

When subjected to a step in the measurement-noise input, v , of 1 track, the simple controller overshoots by about 0.3 tracks and settles in about 1 msec. The “enhanced” controller

³Even though the simple controller could be implemented with only two estimated states, for the purposes of this simulation, it was implemented with all four states, but with the corresponding estimator and controller gains set to zero. That allowed simulation of both the simple and enhanced controller with the same structure, but different estimator/controller gains.




Fig. 6. Comparison of simulated responses of the enhanced and simple controllers to a unit step at the measurement-noise input, v .




Fig. 7. Comparison of simulated responses of the enhanced and simple controllers to a pulse (of amplitude 0.2 Tracks/Sample²) at the disturbance input, w_2 .

overshoots by almost 0.5 tracks, and rings (at the frequency of the disturbance that it was designed to reject) for several msec. When subjected to a (single-sample) pulse of magnitude 0.2 track/Sample², the simple system moves about 0.3 tracks, and again recovers in about 1 msec, while the “enhanced” system rings for a long time.

VII. CONCLUSIONS

In a sense, the authors have done little more than to verify, both numerically and experimentally, that the principle of internal model control [2] works. That is, a HDD subjected to a narrow-band disturbance can in steady-state reject that disturbance very well if the servo estimator/controller contains an accurate model of the disturbance source. Because of the rapid rate of track-density growth in the HDD industry today, and the relative importance of narrow-band TMR sources [1], however, this “textbook” control technique deserves careful examination.

One should expect that an improvement in the control loop’s ability to reject an additional disturbance would come at some cost. The enhanced servo loop does a very good job of rejecting narrow-band disturbances because of its high loop-gain at the expected frequency of the disturbance, as shown in Fig. 2 (lower plot). That same peak in the open-loop gain causes ringing in the system’s transient response to almost any “surprises.” Because of this poor transient response, it is likely that an actual HDD would need to wait a long time after the end of each seek that it does before enabling its enhanced disturbance-rejection, and then wait for several more msec for the TMR-reduction transient shown in Fig. 5. This would place significant restrictions upon the performance of such a drive. Designers using this control technique would need to trade performance for TMR.

The enhanced servo loop presented here was designed specifically for steady-state operation. It is likely that a different estimator/controller, designed to deal with both steady-state and transient conditions, would have better transient behavior. Such a servo loop would still need to have a gain peak at the expected frequency of the disturbance, and thus would still likely suffer similar (perhaps smaller) transient problems.

REFERENCES

- [1] R. Ehrlich and D. Curran, “Major HDD TMR sources, and projected scaling with TPI,” *IEEE Trans. Magn.*, vol. 35, no. 2, pp. 892–897, Mar. 1999.
- [2] G. F. Franklin, J. D. Powell, and A. Emami-Naeini, *Feedback Control of Dynamic Systems*, 2nd ed: Addison-Wesley.
- [3] A. H. Sacks, M. Bodson, and W. Messner, “Advanced methods for repeatable runout compensation,” *IEEE Trans. Magn.*, vol. 31, no. 2, pp. 1031–1036, Mar. 1995.
- [4] “Matlab control system toolbox users guide,” in *The Mathworks*, Jan. 1998.
- [5] G. F. Franklin, J. D. Powell, and M. Workman, *Digital Control of Dynamic Systems*, 3rd ed: Addison-Wesley, pp. 719–723.