Proceedings of the 17th World Congress
The International Federation of Automatic Control
Seoul, Korea, July 6-11, 2008

Discrete-Time Exact and Approximate Dynamic Inversion for Settle Performance ⋆

Brian P. Rigney * **Lucy Y. Pao** ** **Dale A. Lawrence** ***

* *Electrical and Computer Engineering Dept., University of Colorado, Boulder, CO 80309-0425, email:* brian.rigney@gmail.com.
** *Electrical and Computer Engineering Dept., University of Colorado, Boulder, CO 80309-0425, email:* pao@colorado.edu.
*** *Aerospace Engineering Sciences Dept., University of Colorado, Boulder, CO 80309-0425, email:* dale.lawrence@colorado.edu.

Abstract:
Single-track hard disk drive (HDD) seek performance is measured by settle time, t_s, defined as the time from the arrival of a seek command until the measured position reaches and stays within an acceptable distance from the target track. In this paper, we show the effective use of feedforward dynamic inversion, coupled with an aggressive desired trajectory y_d, to achieve high performance settle times. It is well known that the exact tracking solution for nonminimum phase (NMP) systems requires noncausal preactuation to maintain bounded internal signals. In the specific HDD operating modes of interest, anticipation of a seek command is unrealistic, and thus preactuation adds to the overall computation of settle time. Unlike many dynamic inversion tracking applications, this negative effect of preactuation leads to interesting trade-offs between preactuation delay, tracking accuracy, and achievable settle performance. We show that, surprisingly, very little preactuation is desirable when truncating the exact tracking solution and applying it to our NMP HDD model. For comparison, we also review the stable Taylor series approximate inverse, and show that a zero-order series' settle performance is comparable to truncated exact inversion while being easier to compute and implement. We experimentally validate this conclusion on a Servo Track Writer (STW).

Keywords: inverse dynamics control, non-minimum phase systems, settling times, storage devices, optimal trajectory.

1. INTRODUCTION

Many applications perform rest-to-rest maneuvers with unsaturated plant inputs, including single-track hard disk drive (HDD) "seeks". Disk drives perform single-track seeks in numerous operating modes, including servo track writing [Takaishi et al., 2003], [Lee, 1991], sequential data transfer [Guttmann et al., 2000], and manufacturing scans to detect media surface defects [Blachek et al., 1999]. Seek performance is measured by settle time, t_s, defined as the time from the arrival of a seek command until the measured position reaches and stays within an acceptable distance from the target track. In these operating modes, minimizing t_s reduces manufacturing time and cost, or increases data throughput.

While we simplify the focus of this paper to discrete-time, single-input single-output (SISO), linear time-invariant (LTI) plant dynamics, there are still many challenges imposed by the HDD applications. HDD dynamics are typically characterized by the presence of nonminimum phase (NMP) zero dynamics. NMP dynamics can arise when the sensors and actuators are noncollocated [Miu, 1993], as in HDDs which have the magnetic reader position sensor and voice-coil actuator on opposite ends of a flexible actuator arm. NMP zeros in discrete-time systems can also result from fast sample rates and high relative degree [Åström et al., 1984]. As we will see, NMP zeros require special treatment in the application of dynamic inversion for aggressive rest-to-rest maneuvers.

HDD computational processing power is limited because high unit volumes dictate the use of low-cost digital signal processors (DSPs). The DSPs are responsible for much more than control tasks, further limiting the amount of processing power available for control. This constrains the single-track seek control design and necessitates the use of computationally efficient algorithms.

Finally, the HDD error feedback compensator C is usually designed for regulation purposes using knowledge of the plant dynamics P, the disturbance and noise spectra, and performance metrics on the regulated state. Typically, the closed-loop natural response is relatively slow to settle, having limited bandwidth due to stability robustness concerns in the presence of unmodeled dynamics. Further, any permanent or temporary change to the feedback compensator for settle performance improvements could negatively affect regulation or require complicated switching to remove transients. Hence, improvements to settle time

⋆ This work has been supported through research grants from Maxtor Corporation and the National Science Foundation (CMS-0201459 and CMMI-0700877).

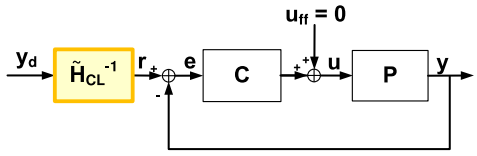

Fig. 1. Block diagram of the feedforward closed-loop inverse (FFCLI) architecture.

are best accomplished through exogenous inputs r and u_{ff} which enter the loop as shown in Fig. 1.

In the presence of these complexities and constraints, we have previously demonstrated the aggressive settle performance capability of dynamic inversion together with reference command generation [Rigney et al., 2006]. In [Rigney et al., 2008], we show inversion-based settle performance compares favorably with minimum energy optimal state transfer techniques [Perez et al., 2003] in our HDD application, while being more amenable to future adaptation. Other researchers have investigated the combined use of dynamic inversion and reference trajectory generation to reduce t_s [de Gelder et al., 2006], [Piazzi et al., 2000], but the complex off-line optimization procedures are not applicable to on-line implementation on low-cost DSPs. [Rigney et al., 2006] and [Rigney et al., 2008] also discuss the settle performance benefits of the feedforward closed-loop inversion (FFCLI) architecture in Fig. 1. FFCLI limits the excitation of the slow closed-loop dynamics and will be used throughout this paper.

The purpose of this paper is the design of the closed-loop dynamic inverse system $\tilde{H}_{CL}^{-1}$ in Fig. 1 for a settle performance objective. This design is complicated by the presence of NMP zeros in the closed-loop dynamics. It is well known that the exact tracking solution for NMP systems requires noncausal preactuation of the closed-loop system to maintain bounded internal signals [Devasia et al., 1996], [Hunt et al., 1996], [Marconi et al., 2001]. Unlike the tracking applications discussed in these references, anticipation of a seek command is unrealistic in the specific HDD operating modes of interest. Thus preactuation adds to the overall computation of settle time and leads to interesting trade-offs between preactuation delay, tracking accuracy, and the achievable settle performance. After precisely defining t_s and describing the experimentally identified example model in Sections 2 and 3, we review the bounded exact inversion technique from an initial condition preload perspective in Section 4. Here, we explore multiple preactuation strategies to preload the initial conditions and evaluate their settle performance. Section 5 reviews an alternative strategy for creating $\tilde{H}_{CL}^{-1}$: stable approximation using a noncausal Taylor series. This technique is capable of producing similar settle performance as the exact inversion algorithms, while being computationally simpler and more amenable to future adaptive inversion approaches. Finally, Section 6 summarizes the conclusions of the paper.

2. SETTLE PERFORMANCE MEASURE

Dynamic inversion algorithms are typically applied to output tracking problems, with performance metrics on the tracking error $y_d - y$, where y is the measured plant output. For settle performance, accurate tracking of a particular y_d is not the objective. In this section, we explicitly define our measure of settle time and present a computationally efficient y_d generation technique suitable for on-line implementation and future adaptation.

2.1 Settle Time Measurement

Settle time t_s is defined as the time from the arrival of a seek command until the measured position reaches and stays within the settle boundary surrounding the desired set-point, as shown in Fig. 2. The settle time, as expressed in number of samples, is k_s, where $t_s = k_s T$ and T is the sample period. In HDD single-track seek applications, there is no penalty on overshoot or undershoot in the plant output y; settle time is the only performance measure of interest, with typical settle boundaries set at $\pm 5\%$ of a track. Further, we do not consider any constraints on plant input magnitude or spectral content. High frequency excitation is of no concern in our application unless it lengthens settle time. [1]

Preactuation, a change in the closed-loop input r before the seek command arrives, and postactuation, a change in r after the seek has settled, are common side-effects of using dynamic inversion algorithms. The example seek in Fig. 2 clearly shows both artifacts, although the postactuation is truncated to save space. As we will discuss in Section 4, both preactuation and postactuation are ultimately a function of the zero dynamics of the system to be inverted. Preactuation requires anticipation of the incoming seek command, which we assume we do not know in the HDD application. While this is not true for all HDD operating modes, we focus on the more generic application which has a wider appeal. Therefore, t_s is lengthened by t_p in order to accommodate the preactuation sequence. Similar to settle time, the preactuation time t_p, expressed

[1] In reality, the acoustical signature of HDD seeks is a concern. While we do not treat this constraint in this work, it would be possible to combine a weighted frequency-domain measure of the plant input u with t_s to form a modified settle performance cost function.

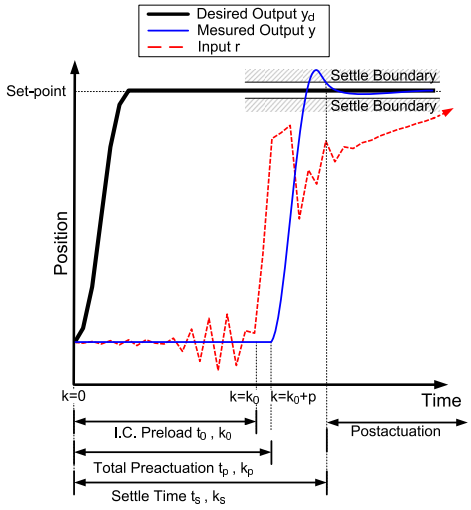

Fig. 2. Example closed-loop input and output trajectories for a single-track seek, with preactuation, postactuation, and total settle time clearly noted.

in number of samples, is k_p, where $t_p = k_p T$. The partitioning of k_p into k_0 and p samples will be fully explained in Section 4. The t_s definition can now be explicitly stated as the time from the initial change in r until y reaches and stays within the settle boundary. Unlike many dynamic inversion tracking applications, this negative effect of preactuation will be critical in the settle performance analysis of NMP inversion algorithms in Sections 4 and 5.

Postactuation is of less concern for all the HDD operating modes of interest. Servo track writing, manufacturing defect scans, and sequential data transfer typically dwell at each track set-point for at least one revolution of the disk. This dwell time is then dependent on the spin speed of the disk, but can be a factor of 10 greater than the settle time. As long as the postactuation is complete within this dwell period, it will not affect performance.

2.2 y_d generation

Motivated by the solution to the time optimal control problem for a rigid body, we derive a family of y_d trajectories generated from the double integral of a bang-bang acceleration pulse, as shown in Fig. 3. The bang-bang acceleration signal has the following $\mathcal{Z}$-transform:

$$Accel(z) = \frac{(z-1)}{d^2T^2} \frac{\left(\sum_{i=0}^{d-1} z^i\right)^2}{z^{2d-1}} \quad , \tag{1}$$

where T is the sample period and $2d$ is the total duration of both the *Accel* signal and y_d. The $\mathcal{Z}$-transform of y_d can then be written as

$$Y_D(z) = \frac{1}{2d^2} \frac{(z+1)}{(z-1)} \frac{\left(\sum_{i=0}^{d-1} z^i\right)^2}{z^{2d-1}} \quad . \tag{2}$$

This simple parameterization provides a single scalar value, d, that can select the aggressiveness of the seek, while also being extremely computationally efficient. Although our future work will adaptively select d on-line for each unit in a population of HDDs [Rigney et al., 2008], this paper uses a fixed extremely aggressive y_d. By fixing y_d, we can focus on the settle performance advantages of the various NMP dynamic inversion algorithms. The analysis in Sections 4 and 5 uses a y_d with duration 0.136 ms, formed from setting $d = 1$ in (2).

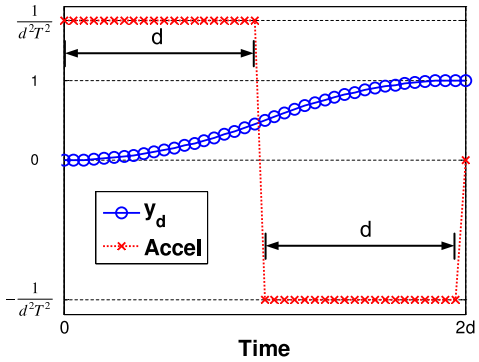

Fig. 3. Desired output trajectory y_d is generated from a bang-bang acceleration pulse. The duration and aggressiveness of y_d is parameterized by d.

3. EXPERIMENTAL SYSTEM IDENTIFICATION

We require a model of the closed-loop system H_{CL} for both the simulation and experimental settle performance analysis in Sections 4 and 5. Our experimental testbed consists of a Servo Track Writer (STW) [Lee, 1991], provided by Maxtor Corporation, and we focus the identification procedure on the STW's closed-loop dynamics. The Servo Track Writer is used to magnetically encode the initial servo position information on the magnetic media during HDD manufacturing. The STW has its own voice-coil motor (VCM) and precision encoder that mechanically interface with the HDD actuator arm and HDD VCM through an opening in the HDD baseplate. Traditionally, the STW moves the HDD actuator one track, magnetically encodes one revolution of servo position information, and repeats this process until all tracks are written. The single-track seek distance is determined by the HDD track density; we use a single-track step size of 1 μrad as a representative angular track width for a modern HDD. The STW has an encoder sensor resolution of 0.5 nanorad, and a compensator sample time of $T = 68$ μs.

We experimentally identify the STW closed-loop dynamics by injecting a pseudo-random sinusoidal sequence as the reference input r in Fig. 1. The input sequence's sinusoidal amplitudes are designed to give adequate signal-to-noise ratio over a wide frequency band of interest, and the frequencies are selected such that each sinusoid has an integer number of periods within the sequence. The frequency-domain weighting vector and model order were selected through trial-and-error to produce a reasonable matching of the experimental step response, shown in Fig. 4. The weighted least squares model for H_{CL} is

$$H_{CL}(z) = \frac{0.10988(z+0.4947)}{(z-0.1541)(z^2-1.859z+0.8695)} \times \frac{(z^2-1.874z+0.8807)(z^2+2.389z+1.574)}{(z^2-1.24z+0.4409)(z^2+1.233z+0.878)} \quad , \tag{3}$$

which matches the step response well. This 7^{th}-order model has a closed-loop bandwidth near 1 kHz, unity DC gain, a relative degree of 2, a high frequency structural mode near 5.3 kHz, and 2 NMP zeros outside the unit circle at $z = -1.1946 \pm j0.3838$.

The step response with 5% settle boundaries in Fig. 4 also highlights requirements for improved settle performance.

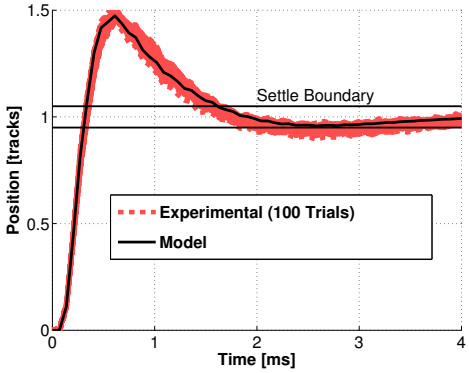

Fig. 4. Experimental and modeled step responses of the STW closed-loop system.

The t_s for the step response of the H_{CL} model is 1.6 ms. While there is some evidence of small high frequency oscillations in the measured output, the step response settle time is dominated by the lower frequency closed-loop dynamics. Any enhancement to settle performance must account for these lower frequency dynamics to produce settle times much shorter than 1.6 ms.

4. EXACT NMP DYNAMIC INVERSION

Dynamic inversion is complicated by the presence of NMP zeros. The NMP zeros of the original system become unstable poles in the inverse system. For the HDD application using the FFCLI architecture, this causes the system in Fig. 1 to exactly track y_d (with possible delay) while r grows unbounded. Many techniques exist to compute exact and approximate solutions for the inverse dynamics, $\tilde{H}_{CL}^{-1}$, with bounded input signal r. We first investigate the exact inversion procedure and review techniques that use the initial condition responses of $\tilde{H}_{CL}^{-1}$ and H_{CL} to maintain bounded r while exactly tracking y_d with delay. In Section 5, we turn our attention to the most applicable approximation method for the HDD applications.

In order to analyze the exact inverse of the closed-loop system, it is helpful to first partition the plant into minimum phase (MP) and nonminimum phase (NMP) zero polynomials

$$P(z) = \frac{B_m(z)B_n(z)}{A(z)} \ , \tag{4}$$

where B_m is the numerator polynomial including all plant MP zeros, B_n contains all plant NMP zeros, and A is the plant denominator polynomial. The closed-loop system can then be expressed as

$$H_{CL}(z) = \frac{B_{m_{CL}}(z)B_n(z)}{A_{CL}(z)} \ , \tag{5}$$

$$A_{CL}(z) = C_D(z)A(z) + C_N(z)B_m(z)B_n(z) \ , \tag{6}$$

$$B_{m_{CL}}(z) = C_N(z)B_m(z) \ , \tag{7}$$

where C_N and C_D are the Hurwitz numerator and denominator polynomials of the compensator C, which is assumed stable, exactly proper, and minimum phase. The exact causal closed-loop inverse can now be stated as

$$\tilde{H}_{CL}^{-1}(z) = \frac{A_{CL}(z)}{z^p B_{m_{CL}}(z)B_n(z)} \ , \tag{8}$$

where p is the relative degree of H_{CL}. Given our definition of t_s and the inability of our HDD applications to preactuate, we must incorporate p unit delays in $\tilde{H}_{CL}^{-1}$ to maintain causality. Using this exact inverse in a FFCLI architecture produces delayed exact tracking

$$Y(z) = H_{CL}(z)\tilde{H}_{CL}^{-1}(z)Y_d(z) = \frac{1}{z^p}\frac{N(z)}{D(z)} \ , \tag{9}$$

where the $\mathcal{Z}$-transform of y_d has been written as the ratio of two polynomials, N and D.

As previously noted, the system in (9) is internally unstable when H_{CL} is NMP, and the r input signal grows without bound. Many exact inversion algorithms use the initial condition responses of $\tilde{H}_{CL}^{-1}$ and H_{CL} at k_0 (referred to in Fig. 2) to maintain bounded r while still achieving delayed exact tracking. The authors of [Devasia et al., 1996], [Hunt et al., 1996], and [Marconi et al., 2001] derive additional preactuation sequences of length k_0 that preload the required initial conditions of H_{CL}. In this case, the total preactuation is $k_p = k_0 + p$. Similar in spirit to [Marconi et al., 2001], we review a discrete-time polynomial development of the required initial conditions, simplified for SISO LTI system dynamics.

We begin by writing the $\mathcal{Z}$-transform of the closed-loop input r, including both the forced and initial condition responses

$$R(z) = \frac{A_{CL}(z)}{z^p B_{m_{CL}}(z)B_n(z)}\frac{N(z)}{D(z)} + \frac{\widetilde{IC}(z)}{z^p B_n(z)B_{m_{CL}}(z)} \ . \tag{10}$$

Here, $\widetilde{IC}$ is the numerator polynomial of the initial condition response of $\tilde{H}_{CL}^{-1}$, with the $\tilde{}$ modifier used to refer to the inverse system. By choosing to describe $\tilde{H}_{CL}^{-1}$ with an observable canonical state-space realization $(\tilde{\boldsymbol{A}}_{\boldsymbol{obs}}, \tilde{\boldsymbol{b}}_{\boldsymbol{obs}}, \tilde{\boldsymbol{c}}_{\boldsymbol{obs}}, \tilde{\boldsymbol{d}}_{\boldsymbol{obs}})$, the initial condition response numerator polynomial simplifies to

$$\widetilde{IC}(z) = z\,\tilde{\boldsymbol{c}}_{\boldsymbol{obs}}\,\mathrm{adj}(zI - \tilde{\boldsymbol{A}}_{\boldsymbol{obs}})\tilde{x}(k_0) \ , \tag{11}$$

$$= \begin{bmatrix} z^n & z^{n-1} & \cdots & z \end{bmatrix} \tilde{x}(k_0) \ , \tag{12}$$

where n is the order of A_{CL} and the identity used in (12) can be found in [Kailath, 1980]. In order to compute the initial state vector $\tilde{x}(k_0)$ required to cancel the unstable modes in the forced response, we must isolate B_n in (10)

$$R(z) = \frac{U_1(z)}{z^p B_{m_{CL}}(z)D(z)} + \frac{U_2(z)}{B_n(z)} + \frac{\widetilde{IC}(z)}{z^p B_n(z)B_{m_{CL}}(z)} \ . \tag{13}$$

The polynomials U_1 and U_2 result from the partial fraction expansion of the forced response and must satisfy

$$U_1(z)B_n(z) + U_2(z)z^p B_{m_{CL}}(z)D(z) = A_{CL}(z)N(z) \ . \tag{14}$$

Hence, if $\tilde{x}(k_0)$ satisfies the following equation

$$\begin{bmatrix} z^n & z^{n-1} & \cdots & z \end{bmatrix} \tilde{x}(k_0) = z^p B_{m_{CL}}(z)\left(\gamma B_n(z) - U_2(z)\right) \ , \tag{15}$$

where γ is an unknown scalar constant, then the r response is bounded and has the following $\mathcal{Z}$-transform

$$R(z) = \frac{U_1(z)}{z^p B_{m_{CL}}(z)D(z)} + \gamma \ . \tag{16}$$

We postpone discussion of the existence of $\tilde{x}(k_0)$ and the computation of γ in (15) until we develop the required initial conditions on H_{CL}.

An unfortunate consequence of removing the unstable components from r is that we no longer achieve delayed exact tracking. The $\mathcal{Z}$-transform of the output y, again including both forced and initial condition responses, is

$$Y(z) = \frac{B_n(z)}{A_{CL}(z)} \frac{U_1(z)}{z^p D(z)} + \gamma \frac{B_{m_{CL}}(z) B_n(z)}{A_{CL}(z)} + \frac{IC(z)}{A_{CL}(z)} \quad , \tag{17}$$

where $IC(z)$ is the numerator polynomial of the initial condition response of H_{CL}. Similar to the inverse dynamics, an observable canonical state-space realization for H_{CL}, $(\boldsymbol{A_{obs}}, \boldsymbol{b_{obs}}, \boldsymbol{c_{obs}}, \boldsymbol{d_{obs}})$, leads to

$$IC(z) = \begin{bmatrix} z^n & z^{n-1} & \cdots & z \end{bmatrix} x(k_0) \quad , \tag{18}$$

with $x(k_0)$ representing the initial condition vector of H_{CL}. Through the use of (14), we can isolate the delayed exact tracking portion of the forced response in (17)

$$Y(z) = \frac{N(z)}{z^p D(z)} - \frac{U_2(z) B_{m_{CL}}(z)}{A_{CL}(z)} + \gamma \frac{B_{m_{CL}}(z) B_n(z)}{A_{CL}(z)} + \frac{IC(z)}{A_{CL}(z)} \quad . \tag{19}$$

We can then use $x(k_0)$ to cancel the unwanted components

$$\begin{bmatrix} z^n & z^{n-1} & \cdots & z \end{bmatrix} x(k_0) = B_{m_{CL}}(z) \left(-\gamma B_n(z) + U_2(z) \right) \quad . \tag{20}$$

Equations (15) and (20) have solutions $\tilde{x}(k_0)$ and $x(k_0)$, respectively, if

(1) The order of the right-hand side (RHS) of each equation is less than or equal to n. By definition, the order of $z^p B_{m_{CL}} B_n$ is n and the order of $z^p B_{m_{CL}} U_2$ is less than or equal to n. Therefore the order constraint is easily satisfied.

(2) The RHS of each equation has at least one root at $z = 0$. If $B_{m_{CL}}$ does not have a root at $z = 0$, then γ must satisfy

$$\gamma = \frac{U_2(0)}{B_n(0)} \quad , \tag{21}$$

where $B_n(0)$ is nonzero by definition.

The unstable mode cancellation in the inverse system can be analytically computed and exactly enforced in the feedforward controller software, i.e., by only implementing (16); precisely setting the initial conditions of the physical closed-loop H_{CL} system is much more difficult. We discuss three techniques to address this difficulty.

4.1 H_{CL} State Transfer to Desired $x(k_0)$

As a byproduct of discrete-time controllability, there exists a sequence $r(k)$ that transfers the state of H_{CL} from 0 to $x(k_0)$ in at most n samples [Kailath, 1980]. While it may be possible to control the state of H_{CL} from 0 to $x(k_0)$ in ℓ samples, where $\ell < n$, this requires the very restrictive condition that $x(k_0)$ lies in the range space of the ℓ-step input-to-state map. For our example H_{CL} in (3) and the required $x(k_0)$, this does not apply and we are left with $k_p = n + p = 9$ samples of preactuation. This method of preloading initial conditions is not typically used in standard dynamic inversion tracking applications because y is non-zero during preactuation and causes increased tracking errors. In contrast, it is a perfectly legitimate technique in this settle time application where preactuation is already included in t_s. Unfortunately, as we will see, 9 samples of preactuation is a very long delay compared to other techniques.

4.2 Infinite Preactuation

The authors of [Devasia et al., 1996] and [Hunt et al., 1996] derive a continuous-time noncausal input $r(t)$ over the preactuation interval $t \in (-\infty, t_0]$ which drives the state of H_{CL} from rest to the required $x(t_0)$ while maintaining $y(t) = 0$. A similar procedure is detailed in [Marconi et al., 2001] for a noncausal input $r(k)$ in discrete-time, which we briefly review here. Consider (13), without the initial condition response of $\tilde{H}_{CL}^{-1}$ included

$$R(z) = \frac{U_1(z)}{z^p B_{m_{CL}}(z) D(z)} + \frac{U_2(z)}{B_n(z)} \quad . \tag{22}$$

In order to compute a bounded $r(k)$ from $R(z)$, we use the standard right-sided inverse $\mathcal{Z}$-transform for the stable first term in (22) and a left-sided inverse $\mathcal{Z}$-transform for the unstable second term [Haykin et al., 2003]. The left-sided inverse transform for $\frac{U_2(z)}{B_n(z)}$ has a region of convergence (ROC) that includes the unit circle and produces a bounded noncausal sequence which is nonzero on $k \in (-\infty, k_0 - 1]$. This noncausal portion of $r(k)$ drives the state of H_{CL} to the required initial conditions $x(k_0)$ such that the undesired forced response terms in (19) are cancelled and y achieves delayed exact tracking. A portion of the $r(k)$ preactuation sequence for the example system was previously shown in Fig. 2. An infinite k_p causes an obviously unrealistic infinite t_s, and we thus investigate truncation strategies.

4.3 Truncated Preactuation

The example system's $r(k)$ in Fig. 2 converges quickly to zero as $k \to -\infty$, allowing for practical implementation through truncation. While truncated preactuation is no longer an exact inversion approach, it is directly motivated by the exact inversion development and we thus include it with the other exact algorithms. The authors of [Zou et al., 1999] and [Marconi et al., 2001] discuss truncation strategies when y_d is only known over a finite window ahead of the current sample. Instead, we take advantage of full knowledge of y_d in our application, which allows for the partition in (22). We then use the first k_0 samples of the left-sided inverse $\mathcal{Z}$-transform of the unstable second term in (22) as the truncated preactuation sequence applied on $k \in [0, k_0 - 1]$. Large values for k_0, and consequently t_p, result in small errors in $x(k_0)$, very little tracking error (disregarding delay), but large offsets in the calculation of t_s. There is thus an interesting trade-off between the length of t_p and the achievable settle performance, explored further for the example system in Fig. 5. Unlike tracking performance, which generally deteriorates as k_p decreases, settle performance improves for smaller k_p until $k_p = 8$. At this point, the tracking errors induced by the unwanted forced response terms in (19) cause excursions in y that exceed the settle boundary. The tracking error is dependent on the size of the undesired terms in the y response, which in turn is determined by the degree to which the closed-loop dynamics H_{CL} are excited by the desired reference trajectory y_d. For our example system, as shown in Fig. 5, $k_p = 3$ yields the best trade-off between $x(k_0)$ accuracy and k_p delay, and achieves a t_s of 0.476 ms.

We further explore the results of Fig. 5 by plotting the y output responses for the $k_p = 3$ truncated preactuation

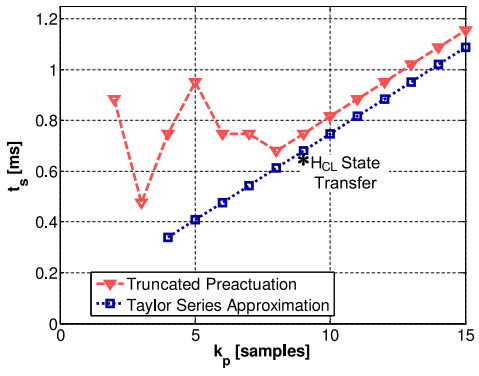

Fig. 5. Settle time t_s versus number of preactuation samples k_p for exact inversion with truncated preactuation and noncausal Taylor series approximation.

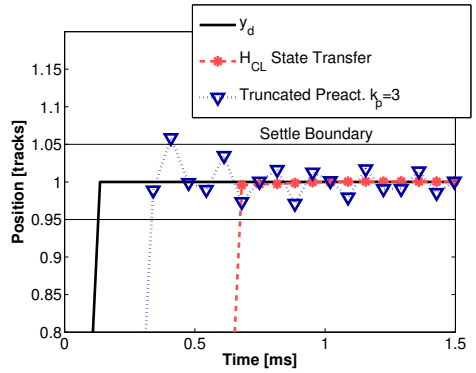

Fig. 6. Output y response comparison between H_{CL} state transfer technique to preload $x(k_0)$ and truncated preactuation with $k_p = 3$.

method and the H_{CL} state transfer technique for our example system in Fig. 6. While the H_{CL} state transfer method exactly tracks a delayed version of y_d, the $k_p = 9$ samples of preactuation yield a very slow $t_s = 0.680$ ms. Alternatively, while high frequency resonance and slow closed-loop mode excitation in y initially cause excursions outside the settle boundary, the $k_p = 3$ truncated preactuation method's achievable settle time is significantly faster: $t_s = 0.476$ ms. Although the truncated preactuation technique with $k_p = 3$ achieves a faster t_s, the response in Fig. 6 is of concern considering the HDD population has large unstructured uncertainty at high frequency. Also, even after t_s, the y trajectory is close to exceeding the settle boundary. Disturbances and sensor noise on y can lead to very inconsistent t_s's for a given y_d with this truncated preactuation technique. The next section discusses an alternative stable approximate NMP inversion technique that is capable of producing better settle performance than these exact techniques with less high frequency excitation.

5. NONCAUSAL TAYLOR SERIES STABLE APPROXIMATION

The noncausal Taylor series approximation seeks to approximate $\frac{1}{B_n}$ with a noncausal polynomial of chosen order [Gross et al., 1994]. We review the algorithm by first starting with the NMP zero polynomial

$$B_n(z) = b_{m_n} z^{m_n} + b_{m_n-1} z^{m_n-1} + \cdots + b_0 \quad , \qquad (23)$$

where m_n is the number of NMP zeros. The resulting $\tilde{H}_{CL}^{-1}$ system with the Taylor series approximation is then

$$\tilde{H}_{CL}^{-1}(z)\Big|_{Taylor} = \frac{A_{CL}}{z^{p+m_T+m_n} B_{m_{CL}}} \frac{\sum_{i=0}^{m_T} \alpha_i z^i}{B_n(1) \sum_{i=0}^{m_T} \alpha_i} \quad , \qquad (24)$$

where m_T is the order of the series approximation and the α_i sequence is derived from the Taylor series expansion of $\frac{1}{B_n}$, as in [Gross et al., 1994]. In this case, $k_p = p + m_T + m_n$ samples of delay have been incorporated into $\tilde{H}_{CL}^{-1}$ to maintain causality. The denominator gain $B_n(1) \sum_{i=0}^{m_T} \alpha_i$ ensures the resulting FFCLI system from y_d to y has unity DC gain

$$\frac{Y(z)}{Y_d(z)}\Big|_{Taylor} = \frac{B_n(z) \sum_{i=0}^{m_T} \alpha_i z^i}{z^{p+m_T+m_n} B_n(1) \sum_{i=0}^{m_T} \alpha_i} \quad . \qquad (25)$$

Surprisingly, the y_d to y system has a finite impulse response (FIR). This has significant settle performance advantages because y perfectly achieves the desired set-point after a fixed number of samples. Unfortunately, the FIR order is $k_p = p + m_n + m_T$ and can be large for high-order series approximations. Similar to the trade-offs between settle performance and tracking accuracy with the truncated preactuation method, increasing m_T improves the approximation to $\frac{1}{B_n}$ and yields better tracking accuracy, while limiting the achievable settle performance. Referring back to Fig. 5, we plot t_s over a range of series orders, from 0 to 11. This yields a k_p range from 4 to 15 for our system in (3) due to the extra delay in (25) from the relative degree p and NMP polynomial order of $B_n(z)$. In this case, the shortest series, a zero-order approximation, provides the best settle performance. A zero-order series is the equivalent of ignoring the NMP zeros in H_{CL} and assuming $B_n(z) = 1$. The tracking improvements that come from higher-order approximations are negated by the accompanying delay in the calculation of t_s.

We further explore the settle performance of the zero-order series approximation by plotting the y output response in Fig. 7. While the simulated response does not achieve delayed exact tracking, it does exactly reach the set-point k_p samples after y_d, with $t_s = 0.34$ ms or 6 samples. Due to our definition of t_s in Section 2.1, we count $k = 0$ as the first sample because the r input changes at this instant. This is 2 samples faster than the truncated preactuation exact method, and 1 sample faster than the minimum energy optimal state transfer techniques which are limited by the order of the system [Rigney et al., 2008]. Further, there is no evidence of the high-frequency excitation present in the truncated preactuation output response. The zero-order series inverse also requires less computational complexity than the exact inversion methods, and is more amenable to future adaptive inverse approaches.

Fig. 7 also plots the average of 100 experimental seeks using the zero-order Taylor series approximation to vali-

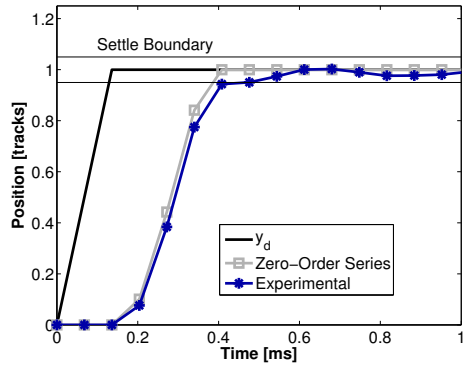

Fig. 7. Simulated and average experimental output y responses for the zero-order Taylor series stable approximate inversion algorithm.

date the simulation results. The average experimental t_s is 0.404 ms, which is 1 sample slower than the simulated response. The average experimental y trajectory shows increased undershoot and phase lag relative to the simulation results. This is indicative of lower frequency modeling errors in the identified model for H_{CL}, rather than excitation of high frequency resonance dynamics. Conceivably, these lower frequency parametric errors in H_{CL} could be adaptively identified on-line to further improve the settle performance.

6. CONCLUSIONS

Dynamic inversion, coupled with reference trajectory generation, is an effective and practical means to achieve high performance settle times. Unlike many NMP dynamic inversion *tracking* applications, which freely use noncausal preactuation to minimize the $y_d - y$ tracking error, preactuation in the HDD settle application adds to the overall settle time. This leads to interesting trade-offs between the amount of preactuation delay, tracking accuracy, and the achievable settle performance.

In this paper, we review the development of the noncausal exact tracking solution in discrete-time and investigate truncation strategies for implementation. We use simulation analysis for an example HDD system to show that the adverse effects of preactuation on settle time overshadow improvements in y_d tracking performance. While our example system performs best with $k_p = 3$ samples of preactuation, the measured output y has increased high frequency content. There are implementation concerns with this method because the plant dynamics are not well known at high frequency.

For comparison, we also review a stable approximate NMP inversion algorithm using noncausal Taylor series approximations. The noncausal series order is a design variable, and can be selected to match the NMP dynamics with desired accuracy. Because the noncausal approximation adds to the computation of settle time, the best settle performance comes from a zero-order approximation. This approximate inversion technique is capable of producing better settle performance than the truncated exact inversion method, with the added benefits of being computationally simpler to implement and producing a finite impulse response (FIR) between y_d and the measured output y. We experimentally validate this conclusion by implementing the approximate inversion technique on an STW. Future work will attempt to address the experimental performance degradation through the use of *adaptive* inverse systems.

REFERENCES

K. Åström, P. Hagander, and J. Sternby. Zeros of sampled-data systems. *Automatica*, vol. 20, no. 1, pp. 31-38, 1984.

M. Blachek, M. Neumann, G. Smith, and P. Wachowiak. Method and apparatus for detecting handling damage in a disk drive. U.S. Patent 5935261, Aug. 10, 1999.

E. de Gelder, M. van de Wal, C. Scherer, C. Hol, and O. Bosgra. Nominal and robust feedforward design with time domain constraints applied to a wafer stage. *ASME J. Dyn. Sys., Meas., and Contr.*, vol. 128, pp. 204-215, Jun. 2006.

S. Devasia, D. Chen, and B. Paden. Nonlinear inversion-based output tracking. *IEEE Trans. Automat. Contr.*, vol. 41, no. 7, pp. 930-942, Jul. 1996.

E. Gross and M. Tomizuka. Experimental beam tip tracking control with a truncated series approximation to uncancelable inverse dynamics. *IEEE Trans. Ctrl. Sys. Tech.*, vol. 2, no. 4, pp. 382-391, Dec. 1994.

J. Guttmann, M. Heminger, M. Hicken, S. Howe, and T. Swatosh. Method and apparatus for optimizing the data transfer rate to and from a plurality of disk surfaces. U.S. Patent 6105104, Aug. 15, 2000.

S. Haykin and B. Van Veen. Inversion of the $\mathcal{Z}$-Transform. *Signals and Systems*. New York, NY: John Wiley and Sons, Inc., 2003, pp. 572-579.

L. Hunt, G. Meyer, and R. Su. Noncausal inverses for linear systems. *IEEE Trans. Automat. Contr.*, vol. 41, no. 4, pp. 608-611, Apr. 1996.

T. Kailath. *Linear Systems*. Englewood Cliffs, NJ: Prentice-Hall, 1980.

C. Lee. Servowriters: A critical tool in hard disk manufacturing. *Solid State Tech.*, vol. 34, no. 5, pp. 207-211, 1991.

L. Marconi, G. Marro, and C. Melchiorri. A solution technique for almost perfect tracking of non-minimum-phase, discrete-time linear systems. *Int. J. Control*, vol. 74, no. 5, pp. 496-506, 2001.

D. Miu. *Mechatronics*. New York, NY: Springer-Verlag, 1993, pp. 114-155.

H. Perez and S. Devasia. Optimal output-transitions for linear systems. *Automatica*, vol. 39, no. 2, pp. 181-192, 2003.

A. Piazzi and A. Visioli. Minimum-time system inversion for residual vibration reduction. *IEEE Trans. Ctrl. Sys. Tech.*, vol. 5, no. 1, pp. 12-22, Mar. 2000.

B. Rigney, L. Pao, and D. Lawrence. Model inversion architectures for settle time applications with uncertainty. *Proc.* 45^{th} *IEEE Conf. Decision Control*, Dec. 2006, pp. 6518-6524.

B. Rigney, L. Pao, and D. Lawrence. Nonminimum phase dynamic inversion for settle time applications. *IEEE Trans. Ctrl. Sys. Tech.* Provisionally accepted in Jan. 2008. Pre-print available at `http://brian.rigney.googlepages.com/pubs`

K. Takaishi, Y. Uematsu, T. Yamada, M. Kamimura, M. Fukushi, and Y. Kuroba. Hard disk drive servo technology for media-level servo track writing. *IEEE Trans. Magn.*, vol. 39, pp. 851-856, Mar. 2003.

Q. Zou and S. Devasia. Preview-based stable-inversion for output tracking. *Proc. American Control Conf.*, Jun. 1999, pp. 3544-3548.