11th International Symposium on Magnetic Bearings
August 26-29, Nara, Japan

TOWARD AUTOMATED AMB CONTROLLER TUNING: PROGRESS IN IDENTIFICATION AND SYNTHESIS

Jerzy T. Sawicki
Dept. of Mechanical Engineering, Cleveland State University, Cleveland, OH 4411, USA
j.sawicki@csuohio.edu

Eric H. Maslen
Dept. of Mechanical Engineering, University of Virginia, Charlottesville, VA 22904 USA
ehm7s@eservices.virginia.edu

ABSTRACT

A significant impediment to low cost commercial deployment of magnetic bearing systems in many applications is the engineering process of tuning. In applications where the dynamics of the platform contribute significantly to overall system dynamics, AMB tuning in the vendor factory may not be sufficient to provide satisfactory performance once installed. In this case, on-site retuning is done “by hand” by a skilled engineer and the associated cost can be substantial. In addition, later alterations in the plant, the process, or the internal machine dynamics may dictate retuning of the AMB controller, again requiring costly engineering intervention. In this paper, we examine the general characteristics of the plant that govern the requirements of an automated tuning process. The AMB machine tool spindle will serve as an example to illustrate results and outline the approach to the robust control model validation problem.

INTRODUCTION

For many reasons, the idea of automatic tuning of the AMB controller is a recurring theme in the AMB literature and offers an important commercial opportunity. Despite this, relatively little actual progress has been reported. There are many possible causes for this lack of progress, but the central culprit is most likely the combination of the fact that the open loop AMB plant is unstable (this means that many established auto-tuning methods won’t work at all) and that many problems of interest exhibit fairly complex dynamics – right half-plane transmission zeros, gyroscopic-induced sensitivity to rotor speed, and so forth.

The general characteristics of the plant that govern the requirements of an automated tuning will be followed with a discussion of mu-synthesis, which we believe is the most viable technique for automated controller synthesis [1-3], presenting some recent experimental results on a machine tool spindle [4-5]. We then review the literature on system identification in the context of the mu-synthesis control problem, detailing the structure required of a viable identifier, which elements are solved and which remain open. This discussion leads naturally into the problem of ascertaining whether a given model and associated uncertainty description (the components of mu-synthesis) properly *covers* measured dynamics of the plant: we outline a computational scheme for establishing this critical property, a necessary and sufficient condition for plant identification. We conclude with a look to the future: what are the real prospects for commercially robust auto-tuning for AMB systems and what technical challenges remain to attain this valuable goal.

MU-SYNTHESIS

Magnetic bearing systems for rotating machinery represent an archetypal challenge for multi-input, multi-output (MIMO) control: they inherently involve multiple interacting control mechanisms and many conflicting performance objectives. As such, they would appear to be a perfect application of formal MIMO control design techniques such as μ-synthesis, which addresses the stability robustness in a systematic design procedure.

At the most conceptual level, the AMB system may be described by the block diagram in which the control inputs u are signals delivered to the power amplifiers, the measurements y are signals received from position sensors, the loads w are forces or electrical noise acting on the system, and performance measures z are those signals that the engineer will monitor in assessing adequate management of the loads w. Essential to μ-synthesis is an assumption that mapping G is linear time invariant (LTI). Also, it is common to model the controller as LTI for the bulk of the design and analysis work. As such, the controller may be described by a matrix transfer function H and the closed loop system indicated in Fig. 3 maps non-dimensional exogenous signals $\hat{w}$ to non-dimensional performance measures $\hat{z}$ via

$$\hat{z}(s) = P\hat{w}(s) \quad (1)$$

where P is the weighted closed loop performance function:

$$P = W_z^{-1}\left[G_{zw} + G_{zu}H\left(I - G_{yu}H\right)^{-1} G_{yw}\right]W_w \quad (2)$$

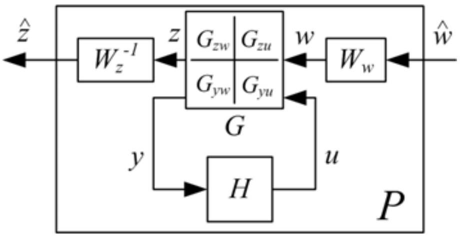

FIGURE 1: AMB system with controller closed loop.

When the plant model satisfies these assumptions (LTI) and the performance objective is to ensure that the peak gain of P is less than some target threshold, γ, then the problem is one of $\mathcal{H}_\infty$ control and synthesis is accomplished by the sole additional step of delivering the weighted model to the machinery of $\mathcal{H}_\infty$ synthesis.

However, in reality, there is an additional consideration in this controller design: the model provided by G may not exactly match the actual plant. That is, the plant G_p may be better described by G_0 in combination with some specific perturbation, Δ. This means that the model structure of Fig. 1 should be amended as in Fig. 2 to include this perturbation.

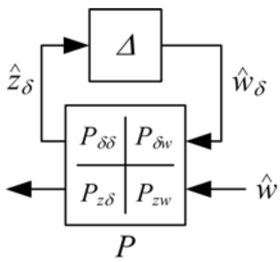

FIGURE 2: Model with uncertainty.

It is assumed that the actual values of the elements of are not known (if they were, then we would simply modify G_0 accordingly) but that they are bounded so that we require that $\Delta \in \mathbf{\Delta}$ where we know the structure of $\mathbf{\Delta}$ and can bound its elements. In this case, the control problem becomes: ensure that the closed loop system is stable for all $\Delta \in \mathbf{\Delta}$ and that $|P|_\infty < \gamma$, again for all $\Delta \in \mathbf{\Delta}$. In this case, the problem becomes one of μ-synthesis. As in the previous case, the solution strategy is to simply specify the nominal plant G_0, the target gain bound γ, and the uncertainty set $\mathbf{\Delta}$. When these three elements are submitted to the computational machinery of μ-synthesis, the required controller is produced automatically, if one can be formulated to meet the specifications.

OPEN-LOOP TRANSFER FUNCTION MEASUREMENT

Consider the control loop shown in Figure 3. Assume that we can simultaneously measure the full set of signals $\mathbf{y}(s)$ and $\mathbf{u}(s)$. In this case,

$$\mathbf{y}(s) = \mathbf{P}(s)\mathbf{u}(s) \quad (3)$$

but it is not possible to extract the elements of $\mathbf{P}(s)$ because there is insufficient information. However, suppose that we construct a series of tests $i = 1...n$ for which

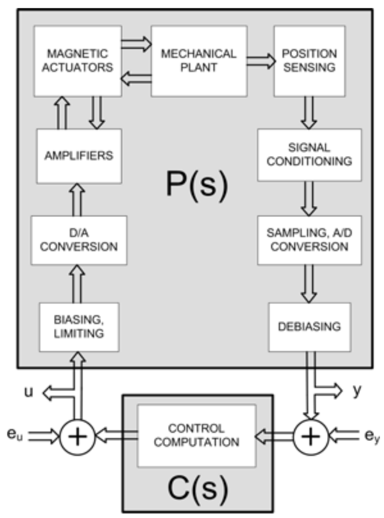

FIGURE 3: AMB control with all components indicated.

$$\mathbf{y}_i(s) = \mathbf{P}(s)\left[\mathbf{I} - \mathbf{C}(s)\mathbf{P}(s)\right]^{-1}\hat{\mathbf{v}}_i e_{u,i}(s) \quad (4a)$$

and

$$\mathbf{u}_i(s) = \left[\mathbf{I} - \mathbf{C}(s)\mathbf{P}(s)\right]^{-1}\hat{\mathbf{v}}_i e_{u,i}(s) \quad (4b)$$

in which $\hat{\mathbf{v}}_i$ is a vector of zeros except for the ith element, which is 1.0.

In this case, we may construct the matrices

$$\mathbf{Y}(s) \equiv \left[\mathbf{y}_1(s) \quad \mathbf{y}_2(s) \quad \cdots \quad \mathbf{y}_n(s)\right] \quad (5a)$$

and

$$\mathbf{U}(s) \equiv \left[\mathbf{u}_1(s) \quad \mathbf{u}_2(s) \quad \cdots \quad \mathbf{u}_n(s)\right] \quad (5b)$$

for which it is trivially established that

$$\mathbf{Y}(s)=\mathbf{P}(s)\left[\mathbf{I}-\mathbf{C}(s)\mathbf{P}(s)\right]^{-1} diag\left[e_{u,i}(s)\right] \quad (6)$$

Assuming then, that the inverses of $\mathbf{U}(s)$ and $diag\left[e_{u,i}(s)\right]$ exist, one can construct

$$\mathbf{U}^{-1}(s)=diag\left[e_{u,i}(s)\right]^{-1}\left[\mathbf{I}-\mathbf{C}(s)\mathbf{P}(s)\right] \quad (7)$$

so that

$$\mathbf{Y}(s)\mathbf{U}^{-1}(s)\equiv\mathbf{P}(s) \quad (8)$$

Obviously, the choices of $e_{u,i}$ and number of exemplars, n, must be made in order to ensure invertibility of these matrices (or, rather, to minimize their condition numbers). Further, all elements of each column of $\mathbf{Y}(s)$ and $\mathbf{U}(s)$ must be acquired simultaneously in order to ensure that $e_{u,i}$ is identical for every element of each column. The various columns, must of course, be measured independently. Hence, n measurements must be made where n is at least equal to the dimension of **u**.

AMB MACHINING SPINDLE

The platform for this study is an AMB supported machine tool spindle with the cross-section shown in Figure 4. The spindle was originally developed by Revolve Magnetic Bearings, a subsidiary of SKF, Inc., and adapted to permit control using a dSPACE digital controller. The spindle rotor is supported by two radial bearings and one thrust bearing. The maximum static radial load capacities are approximately 1400 and 600 N for the front and rear bearings, respectively, and the maximum axial capacity for thrust bearing is 500 N. The spindle reaches a rotational speed of 50,000 rpm at 10 kW. The AC motor acts on the rotor between the thrust and rear radial bearing [4-5].

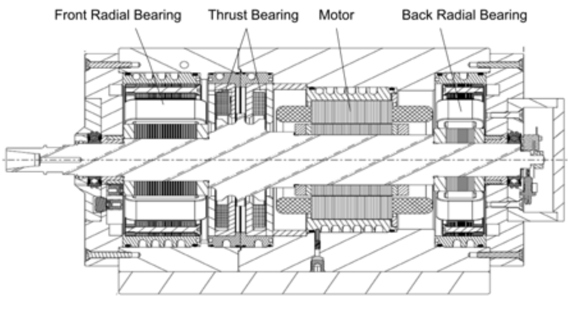

FIGURE 4: Cross section of AMB machining spindle without tool holder.

The rotor is modeled using finite Timoshenko beam elements to produce rotor mass (M), gyroscopic (G), and stiffness (K) matrices of fairly high dimension:

$$\mathbf{M}\begin{Bmatrix}\ddot{x}\\ \ddot{y}\end{Bmatrix}+\omega\mathbf{G}\begin{Bmatrix}\dot{x}\\ \dot{y}\end{Bmatrix}+\mathbf{K}\begin{Bmatrix}x\\ y\end{Bmatrix}=\begin{Bmatrix}f_x\\ f_y\end{Bmatrix} \quad (x)$$

where ω is the rotor speed. The original finite element model of this rotor had 64 mass stations in each plane. This model was transformed to modal using modal matrix (Φ):

$$\Phi^T M\Phi=I,\ \Phi^T K\Phi=\Lambda,\ \begin{Bmatrix}x\\ y\end{Bmatrix}=\begin{bmatrix}\Phi & 0\\ 0 & \Phi\end{bmatrix}\begin{Bmatrix}z_x\\ z_y\end{Bmatrix} \quad (x)$$

and

$$\begin{Bmatrix}\ddot{z}_x\\ \ddot{z}_y\end{Bmatrix}+\omega\Phi^T\mathbf{G}\Phi\begin{Bmatrix}\dot{z}_x\\ \dot{z}_y\end{Bmatrix}+\mathbf{\Lambda}\begin{Bmatrix}z_x\\ z_y\end{Bmatrix}=\Phi^T\begin{Bmatrix}f_x\\ f_y\end{Bmatrix} \quad (x)$$

The standard modal truncation was then applied wherein it was assumed that the modes at frequencies beyond some cutoff have negligible effects. In this case modal truncation retained only four modes in each plane: two rigid body modes and two flexible modes. At this point, modal damping ratio were selected and inserted in the model.

OPEN-LOOP TRANSFER FUNCTIONS: EXPERIMENTAL RESULTS

The controller was implemented using dSPACE based on differential control, with 10 output channels required and provided by the two D/A boards [3, 4]. The hardware consisted of the DS1005 PPC Board featuring the PowerPC 750GX running at 1GHz. The controller sampling time was 10 kHz. To measure the open loop transfer function in a MIMO system like an AMB supported rotor, some care must be taken in the signal processing. In this case, the system has four inputs (amplifier perturbations for the x- and y- axes of each bearing plane) and four outputs (sensor signals for the x- and y- axes of each sensing plane). Thus, the open loop transfer function has the form:

$$\{\mathbf{y}\}=G(s)\{\mathbf{u}\} \quad (9)$$

in which $G(s)$ is a 4×4 matrix of transfer functions.

To measure $G(s)$, we conducted four experiments in which each of the input signals $e_{u,i}$ is perturbed individually. That is, for the first experiment, we inject a sinusoidal perturbation to $e_{u,1}$, for the second, to $e_{u,2}$ and so on. For each conducted experiment, a vector of *all* inputs to the plant and a vector of all outputs from the plant were recorded. That is, at each frequency, four sets of Fourier coefficients were measured for the four amplifier inputs and four sensor outputs.

In this manner, the signals are expected to be related by

$$Y(\omega_i)=G(j\omega_i)U(\omega_i) \quad (10)$$

and the transfer function may be obtain by the simple arithmetic

$$G(j\omega_i) = U^{-1}(\omega_i)Y(\omega_i) \tag{11}$$

If the perturbations for the four experiments have been chosen well, then the required inverse $U^{-1}(\omega_i)$ will exist for all of the test frequencies.

Figure 5 shows a typical measurement of the open loop transfer function: in this case, from excitation of the amplifiers driving the tool end bearing to the output of the tool end position sensor. Notice, in particular, that the magnitude levels to a constant value at low frequency: this is because the magnetic stiffness of the magnetic bearing asymptotically approaches a constant value at low frequency. If the rotor were truly "free-free", then the magnitude would converge at low frequency to a constant slope of -2. The asymptotic gain is the ratio $k_a k_i k_s / k_x$, so it can be used as an independent measure of k_x if the other parameters are known (usually, k_x is the most difficult parameter to identify); here k_a is the composite gain of the D/A and power amplifier, k_i is the actuator gain, k_s is the composite gain of sensor and A/D, and k_x is the magnetic stiffness.

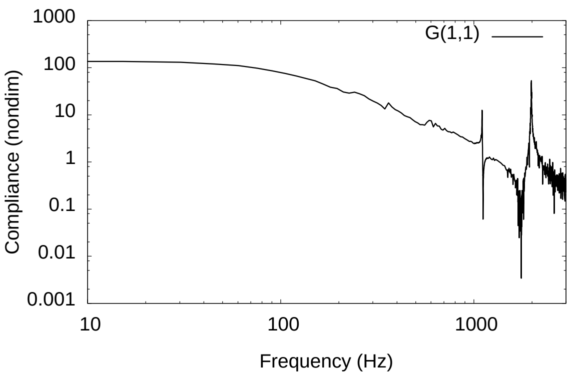

FIGURE 5: Open loop transfer function G(1,1).

In Figure 6, the region between 800 Hz and 2100 Hz has been expanded to permit inspection of the bending modes of the rotor. First, note that the first and second bending modes are clearly identified at roughly 1100 Hz and 1950 Hz. Second, note that there is a transmission zero immediately following the first bending mode and another preceding the second bending mode. The interpretation of this is that the first bending mode has a node between the sensor and actuator used in this measurement while the second bending mode does not [6]. This observation is crucial in validating a computational model of the rotor.

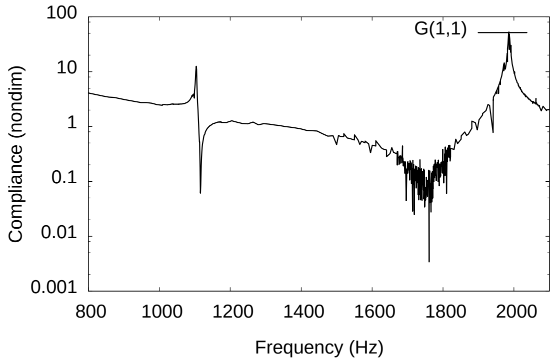

FIGURE 6: A detail view of Figure 5.

Figure 7 presents the same detail but for the right end of the rotor. In this case, the bending modes clearly occur at the same frequencies as for the left end of the rotor but the zeros are quite different. Although the zeros are not as easily identified as for the left end data (roughly 950 Hz and 1750 Hz), it is clear that both zeros precede the associated poles which indicates that neither bending mode has a node between this sensor/actuator pair.

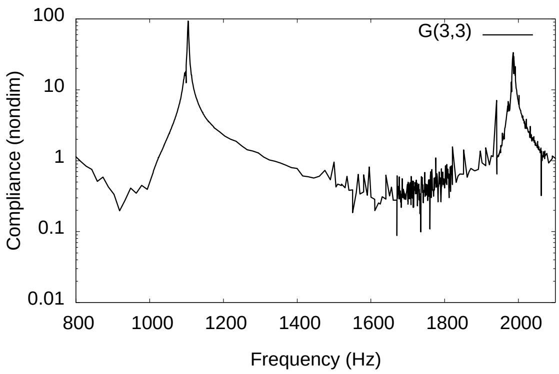

FIGURE 7: Detail of the magnitude of transfer function G(3,3).

The immediate importance of these observations to the problem of constructing an adequate model for mu-synthesis of the controller is that the finite element model of the rotor, in conjunction with electrical models of the surrounding components has these modes at the right frequencies, but the transmission zeros are not correct: for the left end (G_{11}), the model indicates a zero prior to the first pole, rather than after it. From a modal point of view, this amounts to an apparently rather minor error in the mode shape, but if this discrepancy between model and measurement is represented as an uncertainty for the purposes of μ-synthesis, then the magnitude of this uncertainty may be very large: indeed, if the uncertainty is modeled as additive to the plant, it will be extremely large in the immediate vicinity of this mode. Figure

8 presents the error between the measured G_{11} and the modeled G_{11}, which is a lower bound to the additive uncertainty required for μ-synthesis.

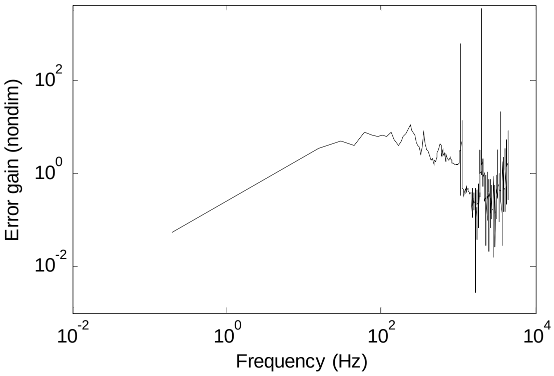

FIGURE 8: The error between the model and the measured plant transfer function.

When using any controller synthesis tool, there is always the possibility that the plant model used for controller synthesis will differ from the dynamics of the physical plant and that, as a result, the performance or relative stability of the combination of controller and physical plant will be poor relative to the design goals. This concern is a primary basis for introducing design objectives such as gain and phase margin: to permit deviation between model and physical plant without suffering excessive loss of relative stability or performance. This issue is accentuated when using automated synthesis tools like μ- or $\mathcal{H}_\infty$, and is the primary reason for the emergence of μ- from its progenitor, H_∞. An important uncertainty analysis problem is associated with the output sensitivity function. Output sensitivity is specifically used in the ISO-14839-3 [7] standard as a measure of AMB system robustness, requiring that the peak gain of the diagonal elements of the output sensitivity function is less than about 4 for commercial systems.

To use μ-synthesis, the analyst needs three components:

1. an analytic model of the physical plant
2. a set of performance objectives written in terms of acceptable gains from various input channels to various output channels
3. a description of the potential sources of deviation between the physical plant and its analytic model.

What μ-synthesis provides, in turn, is a controller which is guaranteed to stabilize any physical plant *covered* by the combination of the analytic model and the uncertainty description. The term "cover" means, essentially, that if you were to generate all possible Bode plots characterizing the analytic model in combination with all possible values of the uncertainty, then the actual physical plant's Bode plot lies within the resulting gain/phase *envelope*.

More specifically, if the plant and associated uncertainty are defined by

$$G(s;\Delta) = \mathcal{F}_u\left(G_0(s),\Delta\right): \quad \Delta \in \mathbf{\Delta}$$

then this model covers an actual plant $G_p(s)$ if and only if there exists some $\Delta_p \in \mathbf{\Delta}$ such that $\left\| G(s;\Delta_p) - G(s) \right\|_\infty = 0$. The symbol $\mathcal{F}_u(\cdot,\cdot)$ denotes the upper linear fractional transformation (LFT) of the first argument by the second argument. It is a general notation for any feedback connection of the nominal plant, $G_0(s)$, with the uncertainty, Δ, and permits uncertainty descriptions of a very broad range including simple additive uncertainty, multiplicative uncertainty, or nearly any kind of structured uncertainty such as parametric, modal, and so forth. The distinction between Δ and $\mathbf{\Delta}$ is that is a specific instance of an uncertainty operator that lives in the bounded class specified by $\mathbf{\Delta}$. Generally, specification of $\mathbf{\Delta}$ is a matter of describing its block structure (which elements of the operator are zero, which are equal one to the other, which are real-valued) and providing a bound for the blocks in terms of maximum singular value, potentially as a function of frequency.

The importance of this notion that the components G_0 and $\mathbf{\Delta}$ must cover the actual plant G_p cannot be over-emphasized: if they do not, then the resulting controller *may* stabilize the physical combination, but it is equally likely that it will not. Our experience in developing controllers for the machine tool spindle described in this paper using μ-synthesis has been that our model has not covered the actual plant (primarily because of the misplaced zero prior to the first flexible mode in G_{11}) and the result is that many of the controllers we have synthesized have not stabilized the system.

Surprisingly, this problem of determining whether or not a given combination of G_0 and $\mathbf{\Delta}$ covers a physical plant has received very little attention in the open literature. Part of the reason for this is fundamental: as elaborated by Newlin and Smith [8], it is actually (and perhaps obviously) not possible to determine whether (G_0,$\mathbf{\Delta}$) covers G_p for, at the least, the simple reason that G_p is never known precisely (hence the uncertainty description!). However, it is possible to determine that (G_0,$\mathbf{\Delta}$) *does not cover* G_p, a piece of information of nearly the same utility. So the approach advocated in [8] is to determine the smallest $\mathbf{\Delta}$ which cannot be shown not to cover G_p and then assume that, at least with some measure of confidence, that the resulting pair (G_0,$\mathbf{\Delta}$) does cover G_p. Obviously, this approach is far from perfect, but at least it screens out uncertainty models which can be demonstrated not to cover.

Computing the condition examined in [8] is nontrivial and there is no software in the public domain to carry out this computation. The authors are presently working to produce a validated implementation of the algorithm described in [8] and will make this available to the larger technical community when this work is completed.

Reflecting on the three components needed to properly manage μ-synthesis, it should be apparent that the first component (analytic model) has been examined exhaustively in the literature. Models are, of course, not perfect, but can certainly exhibit very high levels of sophistication and concomitant fidelity. The second component is important only to ensuring suitable performance: the act of specifying "suitable performance" is entirely equivalent to producing these performance objectives so it is nearly axiomatic that producing this component is completely tractable.

The problematic component is obviously the uncertainty description. The literature on robust analysis and control provides many suggestions for structuring the uncertainty description, but ultimately, the hard requirement remains that (G_0,$\mathbf{\Delta}$) must cover G_p. It may be tempting to simply be conservative: choose a very large Δ to ensure that the cover requirement is met. However, the closed loop system always presents a trade-off between performance and robustness. Hence, choosing an unnecessarily large $\mathbf{\Delta}$ will generally imply significantly worse closed loop system performance than can ideally be obtained. Indeed, an excessively large class $\mathbf{\Delta}$ will frequently lead to the result that no controller can be found that stabilizes this excessively uncertain plant description.

This problem is further confounded by the fact that the uncertainty description is not unique. A good choice of the structure of the interaction between the plant and its uncertainty may lead to cover with an uncertainty of very modest bound, resulting in easy controller synthesis and excellent closed loop performance. By contrast, some alternately structured description of the same level of plant uncertainty may require a very large uncertainty bound and result in very poor closed loop behavior.

To illustrate this, consider the zero location problem illustrated by the machine tool spindle discussed above. If the uncertainty is treated as essentially unstructured so that

$$G(s) = G_0(s) + \Delta$$

then a lower bound on one element of Δ is described by Figure 8. To use this in a synthesis model, we must find a reasonably low ordered transfer function $G_d(s)$ whose gain mimics this additive error and then describe Δ as the product of the deterministic description G_d and some entirely unknown but bounded complex number $\delta : |\delta| < 1$. To cast the result in a form amenable to μ-synthesis, an augmented plant is constructed which contains both G_0 and G_d, with G_d scaled by δ. The more closely G_d cleaves to the data of Figure 8, the less conservative the description becomes, but such tuning of G_d generally raises its order, resulting in a high order nominal augmented plant, which high order is inherited by the synthesized controller. In addition, because this uncertainty is now attributed, in some sense, to the sensing or actuation of the plant, it is a non-physical description of the cause of the uncertainty and, generally, is conservative (δ is larger than it needs to be).

Conceptually, this uncertainty description does, in fact, cover the actual plant G_p, but it also covers many other plants which are entirely infeasible from the underlying physics of the problem. Since the synthesized controller must stabilize the entire implied class of possible plants, including those which are infeasible, it must be unnecessarily conservative.

A better approach is to examine the structure of the rotor model to determine what might cause the displaced transmission zero. Some care must be exercised here because it is undesirable to introduce a causal attribution which produces uncertainty not only in the zero location but also in the associated pole location: the pole location is not uncertain, only the zero. One relatively simple approach to this is to examine the modal description of the rotor and look for mode shape perturbations that could shift the location of the transmission zero. With a bit of care, the uncertainty can then be cast entirely in modal coordinates as a mode shape uncertainty with the result that the analyst is provided with an uncertainty knob that affects only the location of this particular transmission zero, leaving the remaining significant features of the plant transfer function unaffected. By focusing on phenomena with a more nearly causal influence on the observed uncertainty, the nominal uncertainty model becomes simpler (in this case, the uncertainty in modal coordinates becomes a simple scalar with no attributed dynamics - G_d) and can generally be less conservative. Such an apparently deeply embedded uncertainty description is readily described using an LFT structure, so it is entirely amenable to the base formulation of the μ-synthesis problem.

CONCLUSIONS

There is no identification or design problem where physical system can be described by a nominal model. Robust control models account for perturbations and unknown signals; however, it is still a matter of judgment whether or not the model is proper to describe the system.

REFERENCES

1. Zhou, K., Doyle, J.C., and Glover, K., *Robust and Optimal Control*, Prentice-Hall, Inc., 1996.
2. Green, M., and Limebeer, D. J. N., *Linear Robust Control*. Prentice-Hall, Inc., 1995.
3. E.H. Maslen and J.T. Sawicki, "μ-Synthesis for Magnetic Bearings: Why Use Such a Complicated Tool?," ASME Int. Mechanical Engineering Congress and Exposition, Seattle, Washington, USA, November 12-15, 2007.
4. Sawicki, J.T. and Maslen, E.H., "Rotordynamic Response and Identification of AMB Machining Spindle," Paper GT2007-28018, Turbo ASME Turbo Expo Conference, May 14-17, 2007, Montreal, Canada.
5. Sawicki, J.T. and Maslen, E.H., "AMB Controller Design for a Machining Spindle using μ-Synthesis," The Tenth International Symposium on Magnetic Bearings (ISMB-10), Martigny, Switzerland, August 21-23, 2006.
6. Maslen, E.H.: Positive Real Zeros in Flexible Beams, Shock and Vibration, **2**(6), 429-435, 1995.
7. ISO TC108/SC2 Working Group 7, 2006. Active magnetic bearings - evaluation of stability margin. ISO-14839-3, September 2006.
8. Newline, M. P. and Smith, R. S., "A Generalization of the Structured Singular Value and Its Application to Model Validation," IEEE *Transactions on Automatic Control*, **43**(7), pp. 901-907, July 1998.