Title

A control and signal processing integrated circuit for the JPL-Boeing micromachined gyroscopes

Journal

IEEE Transactions on Control Systems Technology, 13(2)

ISSN

1063-6536

Authors

Chen, YC
M'Closkey, RT
Tran, TA
et al.

Publication Date

2005-03-01

DOI

10.1109/TCST.2004.839558

Peer reviewed

A control and signal processing integrated circuit for the JPL-Boeing micromachined gyroscopes

Yen-Cheng Chen, Robert M'Closkey, *Member, IEEE*, Tuan Tran, Brent Blaes

***Abstract*— A special-purpose integrated circuit that accomplishes the real-time control and filtering tasks for the JPL-Boeing micromachined gyroscopes using a flexible, low-power implementation is presented. Our exposition focuses on the integration of the circuit and a prototype sensor, the synthesis and implementation of the control filters, and the subsequent performance of the closed-loop system. Identified sensor models are also presented because the control approach, and hence, the circuit architecture, is motivated by special features of the sensor dynamics.**

***Index Terms* – gyroscopes, application specific integrated circuits, gain control, microelectromechanical devices**

NOMENCLATURE

ADC	Analog-to-digital converter
AGC	Automatic gain control(ler)
ASIC	Application-specific integrated circuit
DAC	Digital-to-analog converter
FIR	Finite impulse response
FPGA	Field-programmable gate array
RMS	Root-mean-square
V	Volt unit
C_{ij}	ijth channel of compensator C
D_i, $\tilde{D}_i$	ith actuator component
F_{clk}	ASIC operation frequency
K_I	AGC integrator gain
K_P	AGC proportional gain
M,K,C	Mass, stiffness, damping matrices
P_{ij}	ijth channel of plant P
$(PC)_{ij}$	ijth channel of loop gain PC
R	AGC reference amplitude
S_i, $\tilde{S}_i$	ith sensing pick-off component
S_I	Input sensitivity function
S_O	Output sensitivity function
α	Oscillator model gain
c	Oscillator model damping
f_s	Servo rate
Ω	Sensor angular velocity
$\bar{\sigma}$	Maximum singular value
s	Laplace transform variable
ω	Frequency
ω_n	Oscillator model natural frequency

Yen-Cheng Chen is a Research Engineer with the Mechanical and Aerospace Engineering Department at the University of California, Los Angeles.

Robert M'Closkey (corresponding author, rtm@obsidian.seas.ucla.edu) is an Associate Professor with the Mechanical and Aerospace Engineering Department at the University of California, Los Angeles. Funding for this project provided by NSF ECS-9985046, JPL Contract No. 1227052, and UC-MICRO 99-083 and 01-058.

Tuan Tran and Brent Blaes are Staff Engineers with the Jet Propulsion Laboratory, Pasadena, CA.

I. INTRODUCTION

Vibratory gyros are an interesting class of sensor because of the opportunity to implement numerous control loops to improve sensor bandwidth, linearity, and dynamic range and to maintain performance in the presence of perturbations to the sensor dynamics. Devices such as the Jet Propulsion Laboratory-Boeing micromachined-class of rate sensors are designed with degenerate dynamics in which the natural frequencies of two lightly-damped modes are coincident. Fabrication irregularities, however, can produce a wide range of dynamic responses among any batch of supposedly identical devices due to the microscopic scale of key mechanical components. The deviation of the dynamic response from the ideal is caused by the sensitivity of the sensor's eigenvector structure to asymmetries in the mass and stiffness distribution. This sensitivity produces a splitting of the modal frequencies as well as crosschannel coupling.

There are several approaches for contending with the variability of the sensor dynamics. For example, electrostatic tuning can bring the sensor's eigenvectors into more favorable alignment with the sensing pick-off frame, or alternatively, multiple sensing pick-off signals can be blended to essentially align the sensing frame with the sensor's eigenvectors. In practice, both approaches are pursued. Once this task is complete it is necessary to fine-tune the various feedback loops for each individual sensor because, for example, sub-degree perturbations of the pick-offs' phase caused by variations in signal conditioning electronics can have a measurable impact on the sensor's ultimate performance.

As with most prototype sensor development, the JPL-Boeing micromachined gyroscopes were initially custom-built and hand-tuned for optimal performance. Thorough testing can never be avoided, especially with high-performance devices, but it is desirable from a production point of view to devise a systematic and efficient approach for identifying important sensor parameters that define the compensator properties, and then once the compensation is synthesized, to implement it on a flexible computational platform. Pursuing this program has been our focus for several years and in this paper we report one aspect of this work, namely the creation of an application-specific integrated circuit (ASIC) for the JPL-Boeing gyros that performs the requisite control and signal processing tasks. Our objective is to have the ASIC supplant the current analog implementation and its associated time-consuming tuning without sacrificing, and perhaps even increasing, sensor performance.

The computational architecture of the ASIC is fixed, however, there is great flexibility in the individual filters and com-

ponents comprising the architecture that enable its operation with a wide range of sensor dynamics. This adaptability can be achieved with a general purpose digital signal processor (DSP), however, the advantage enjoyed by the ASIC over a DSP is its low power consumption, which makes the ASIC suitable for some applications that preclude the use of a DSP.

The design, fabrication and testing of the ASIC is a herculean task so we will not present the many details related to the engineering choices and trade-offs that we performed. Instead, this paper is focused toward a control audience and so we present 1) interesting features of the JPL-Boeing gyros' dynamics that motivate our control approach, 2) how to synthesize typical compensation, and 3) an assessment of the performance of the ASIC.

The paper is organized as follows. Section II briefly introduces a vibratory gyro model which is useful for understanding the basic operation of the device as a rate sensor. Some metrics for measuring controller performance are also mentioned. Section III presents detailed frequency domain models of a JPL-Boeing microgyro prototype and identifies several challenges for the control loop design. Section IV introduces the ASIC. Section V discusses the synthesis procedure and Section VI presents the closed-loop implementation results.

II. A GENERIC MODEL AND TYPICAL CONTROL LOOPS.

A. Vibratory rate sensor model

A generic linear model for vibratory rate sensors is composed of a two degree-of-freedom (DOF) system with a skew-symmetric matrix that is modulated by the angular rate of rotation of the sensor:

$$M\ddot{\vec{x}} + C\dot{\vec{x}} + \Omega S\dot{\vec{x}} + K\vec{x} = \vec{f}. \tag{1}$$

In this model M, C, and K are the real, positive definite 2×2 mass, damping, and stiffness matrices, respectively, Ω and $\vec{f} = [f_1, f_2]^T$ represent the sensor angular rate of rotation and applied actuation forces, respectively, and S is skew-symmetric. These equations are written in the sensor-fixed coordinates denoted by $\vec{x} = [x_1, x_2]^T$. The reader is referred to [5], [7] for a detailed introduction to vibratory rate sensors.

B. Typical vibratory rate sensor control loops

The mode of operation of generic vibratory rate sensors is now described. Control forces are selected to drive x_1 into a sinusoidal response –this is usually accomplished by a *drive* control loop– then rotation at angular rate Ω about the sense axis transfers momentum from one DOF into the other DOF resulting in a change in amplitude and phase of x_2 from which Ω may be inferred. In practice, a second control loop is used to null the x_2 pick-off signal and in this case the control effort is related to Ω. The second loop has been termed the *force-to-rebalance* loop or simply, rebalance loop, in the literature. These two control loops are common in vibratory rate sensors although their implementation may differ from one technology to another. These loops are depicted in Fig. 1.

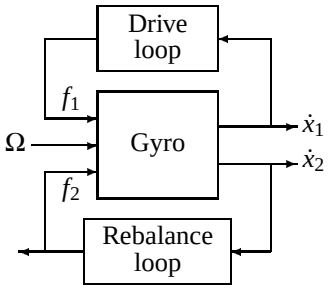

Fig. 1. Sense and drive axis control loops. The rebalance loop control effort, denoted f_2, is demodulated with respect to $\dot{x}_1$ for an estimate of Ω.

Analysis of the closed-loop signals illuminates certain limitations imposed on the detection of Ω. Consider the following simplified version of (1)

$$I\ddot{\vec{x}} + C\dot{\vec{x}} + \alpha\Omega S\dot{\vec{x}} + K\vec{x} = \vec{f}, \tag{2}$$

where the mass matrix is taken to be the identity, the parameter α is the Coriolis coupling coefficient, $S = \begin{bmatrix} 0 & -1; & 1 & 0 \end{bmatrix}$, and the forces $\vec{f} = [f_1, f_2]^T$ are co-located with the sensing pick-offs, and the damping and stiffness matrices are positive definite but otherwise arbitrary. The non-idealities in the sensor dynamics that produce rate-sensing errors in some simple models are itemized in [10], [11] and can be identified according to their influence on the system matrices in (2).

In this analysis we assume that the drive loop has established a constant amplitude sinusoidal response for x_1 and so we take $\dot{x}_1(t) = A\cos\omega t$ for the analysis. The objective is to compute the ideal force co-located with x_2, denoted $\hat{f}_2$, to enforce $x_2(t) \equiv 0$. The control effort of the rebalance loop converges to $\hat{f}_2$ as the loop gain goes to infinity. Define $T(\theta) = \begin{bmatrix} \cos\theta & \sin\theta; & -\sin\theta & \cos\theta \end{bmatrix}$ and let $T(\theta_C)$ and $T(\theta_K)$ be the matrices that diagonalize the damping and stiffness matrices, i.e. $T^T(\theta_C)CT(\theta_C) = \mathrm{diag}(\hat{c}_1, \hat{c}_2)$ and $T^T(\theta_K)KT(\theta_K) = \mathrm{diag}(\omega_1^2, \omega_2^2)$ where $\hat{c}_1$ and $\hat{c}_2$ represent the eigenvalues (principal damping constants) of C and ω_1 and ω_2 are the modal frequencies. The ideal f_2 actuator signal required for $x_2(t) \equiv 0$ is computed to be

$$\begin{aligned} \hat{f}_2 =& A(\alpha\Omega + (\hat{c}_1 - \hat{c}_2)\sin\theta_C\cos\theta_C)\cos\omega t \\ & - A\frac{\omega_1^2 - \omega_2^2}{\omega}\cos\theta_K\sin\theta_K\sin\omega t. \end{aligned} \tag{3}$$

The terms multiplying $\cos\omega t$ are composed of the angular rate Ω and the contributions from the damping. The effects of damping may be reduced by controlling the amount of damping asymmetry ($\hat{c}_1 \approx \hat{c}_2$) or by aligning the eigenvectors of the damping matrix with the x_1–x_2 frame established by the sensing pick-offs ($\theta_C \approx 0$). It is extremely difficult to achieve either of these tasks in practice, hence the sensor is usually designed so as to minimize damping to the greatest extent possible. The second group of terms in (3) are 90° out of phase (quadrature) with the Ω-induced terms and are due to stiffness matrix coupling between the two measurement and drive channels. The sensor's angular rate of rotation is estimated by demodulating

$\hat{f}_2$ with respect to $\dot{x}_1$ and, although this phase-sensitive detection scheme ideally rejects the quadrature terms, any errors in the demodulation signal phase will introduce a spurious component into the estimate of Ω. The quadrature components can be minimized by orienting the stiffness matrix eigenvectors along the x_1-x_2 coordinate frame axes ($\theta_K \approx 0$) or by reducing the frequency split between the modes ($\omega_1 \approx \omega_2$). Both methods are pursued in practice by a variety of tuning schemes that perturb the mass and stiffness matrices. For example, [6], [13] advocate adaptive and closed-loop methods for tuning the sensor dynamics. The JPL-Boeing fabrication process aims to produce sensors with these degenerate dynamics and in fact can consistently produce devices with frequency splits on the order of 0.2%. A complementary method is to change the sensing pick-off frame by creating weighted sums of the pick-off signals to decouple the dynamics. The latter approach is adopted in this paper (see Section III-C).

C. Controller performance metrics

The drive loop controller is selected as a nonlinear filter for reasons to be provided shortly, and the rebalance loop controller is a linear filter, so different experiments are performed to assess their performance. The rebalance loop is designed to null the x_2 pick-off signal, however, the steady-state analysis of the previous section does not illuminate issues such as closed-loop bandwidth, i.e. when Ω is not constant. Linear sensitivity function analysis is an effective means of assessing the rebalance loop performance (closed-loop bandwidth in particular) because the dynamics of the sensor are accurately described by linear equations and the rebalance controller itself is a linear filter. This controller is also used to damp another resonance in the sensor and so the sensitivity function is a useful metric in this case too.

The drive loop controller is designed to regulate the x_1 pick-off signal to a fixed-amplitude sinusoid and in contrast to the linear rebalance controller, it is a nonlinear filter. The reason for choosing a nonlinear controller will become clear when viewing the system identification results of Section III, but to state it briefly, the sensor dynamics, though accurately represented by linear time-invariant equations over intervals of several minutes, are actually slowly time-varying and are well-modeled in the frequency domain by a re-parameterization of the frequency scale. The extremely high quality factors of the sensor's modes can produce enormous changes in gain at a given fixed frequency via a slight shift in the modal frequencies. Hence, a nonlinear controller, the details of which are given in Section V, is necessary to track the slow modal frequency variations. Important performance criteria for the drive loop include drive amplitude stability and modal frequency tracking. In quantifying the nonlinear drive loop bandwidth it is necessary to use local or small signal perturbations, so we will content ourselves with an experimental approach for assessing these performance criteria.

III. JPL-Boeing vibratory gyro models

A. JPL-Boeing micromachined gyro overview

The JPL-Boeing vibratory gyro prototypes (Fig. 2) consist of a silicon micromachined plate suspended above a set of electrodes. The prototype that was tested for this paper was delivered in a high vacuum, hermetically sealed package. The two modes that are exploited in angular rate detection correspond to a two degree-of-freedom rocking motion of the plate that is parameterized by the θ_1 and θ_2 angular coordinates in Fig 2. The central post strongly couples these rocking degrees of freedom via a Coriolis term (the axis which is sensitive to angular rates is the Ω-axis along the post). The electrodes can be used to apply electrostatic forces at points on the plate or, as the plate deflects relative to the electrodes, capacitive sensing can be used to measure the deflection. More details on the design and fabrication of these sensor may be found in [1], [14], [15], [16].

Two electrodes are selected for sensing and another two electrodes are designated as actuators. The actuation electrodes, also commonly referred to as the *drive* electrodes, can be driven directly from a function generator or digital-to-analog converter (DAC). The sensing electrodes use a trans-impedance op-amp configuration to provide a buffered output voltage that is proportional to the average velocity over the electrode surface. Maximum displacements of the sensor's elastic structure, on the order of 1 to 2 μm, correspond to sense electrode potentials of several hundred millivolts in the frequency range of the important modes discussed below. Thus, in the Bode plots that follow, the magnitude scale has Volt/Volt (V/V) units. When defining the equations of motion of this structure we focus on the two closely spaced rocking modes that can be modeled by (1). The angular coordinate parameterization from Fig. 2 is not employed but rather the following perspective is adopted: the generalized coordinates are chosen to be the coordinates established by the sensing electrodes; if we denote the sensing electrode measurements as S_1 and S_2 then $[\dot{x}_1 \quad \dot{x}_2] := [S_1 \quad S_2]$. Furthermore, the electrostatic forces, denoted D_1 and D_2, created by potentials applied to the drive electrodes are different from the generalized forces $[f_1, f_2]$ in (1) since the drive electrodes and the sensing electrodes are not co-located. Thus, we modify (1) to

$$M\dot{\vec{S}} + C\vec{S} + K\int \vec{S} = B\vec{D}, \tag{4}$$

where $\vec{S} := [S_1, \quad S_2]^T$, $\vec{D} := [D_1, \quad D_2]^T$, and B is a real, nonsingular 2×2 matrix that specifies how forces applied by each drive electrode couple into the coordinate frame specified by the sensing electrodes. The Coriolis term has been omitted from this equation. The post is not present in the sensor tested for this paper and so it has reduced sensitivity to rotational motion but nevertheless it retains all of the interesting features of a vibrational gyro. Hence, the controller synthesis, its subsequent implementation on the ASIC, and its performance can be clearly quantified.

B. Sensor empirical frequency response

Fig. 3 shows the wide-band frequency response magnitude of the sensor. In this case $\Omega = 0$ and the device is tested as a 2-input/2-output system. The title of each subplot indicates the sensor/actuator pair. The first large-amplitude, lightly-damped resonance occurs near 2700 Hz. This is a mode in the sensor that corresponds to a linear translation of the sensor's elastic structure along the post direction. This mode is not only easily excited by the control electrodes, but also by linear acceleration in the appropriate direction. The next lightly damped resonance

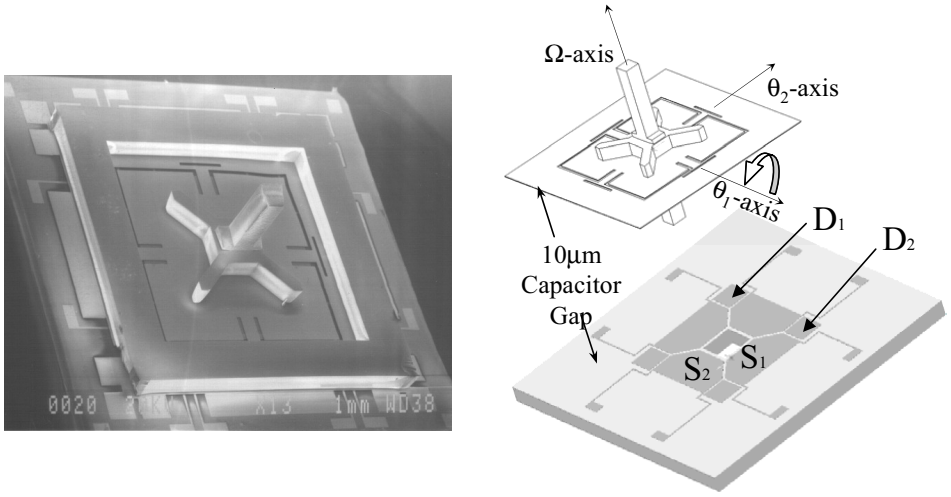

Fig. 2. A photo of the JPL-Boeing vibratory gyro prototype (left) and drawing indicating electrode layout on the baseplate (right). The cloverleaf structure, clearly visible in the photo, has three prominent modes: two angular deflections that can be parameterized by the θ_1 and θ_2 coordinates, and an out-of-plane linear translation in the direction of the Ω-axis. The post along the Ω-axis strongly couples the two angular degrees of freedom with a Coriolis term and it is this coupling that is exploited for detection of the sensor's angular velocity resolved along the Ω-axis. The shaded area on the baseplate drawing reveals the electrode layout. Deflection of the cloverleaf structure changes the capacitance between the cloverleaf and baseplate and the S_1 and S_2 electrodes serve as measurement pick-offs. Electrostatic forces are applied to the cloverleaf by creating a potential difference between the cloverleaf and the D_1 and D_2 electrodes.

near 4420 Hz actually comprises two modes that correspond to the rocking motions of the plate and that have a frequency split of less than 5 Hz (0.11%) and so cannot be individually resolved on this scale. It is these two modes that are strongly coupled via the Coriolis acceleration when the sensor is rotated. Although the sensor was designed to exploit these two modes for angular rate sensing, it is quite evident that the wide-band sensor dynamics are far richer than (4). Nevertheless (4) is an appropriate description if one is interested in a narrow frequency window around the modes. The remaining lightly damped resonances above 5000 Hz are other flexural modes in the sensor's elastic structure. The positive trend in the slope up to 35 kHz is caused by parasitic capacitive coupling between the drive electrodes and sense electrodes. The signal conditioning amplifiers roll-off above 35 kHz and this is responsible for the attention shown in the plots. Lastly, the sharp notches in close proximity to the modes near 2700 Hz and 4420 Hz are created by a charge cancellation on the sensing electrodes due to the opposite phase of parasitic capacitance-induced charge and the charge created by motion of the sensing electrodes. These plots were produced by concatenating results from multiple tests with a dynamic signal analyzer. An alternative approach using ARX models is pursued in [9].

Although each channel in Fig. 3 looks quite similar, there are significant differences in a neighborhood of the two modes near 4420 Hz. It is necessary to explore this region more closely since it is the Coriolis coupling of these modes that render the device sensitive to angular rotation. Fig. 4 reveals the sensor frequency response details in a neighborhood of the 4420 Hz modes with 10 mHz frequency resolution, and demonstrates the high degree of coupling between channels in this device. The process by which we obtain this precise frequency response data is discussed in [2]. The heavy traces in Fig. 4 show the magnitude and phase of the sensor's empirical frequency response taken from a single experiment and are subsequently

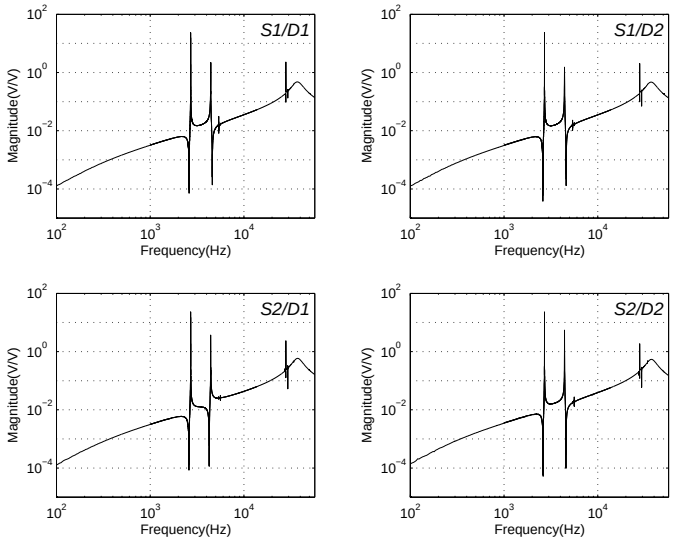

Fig. 3. Wide-band frequency response of the sensor dynamics from the D_1 and D_2 actuation electrodes to the S_1 and S_2 sensing electrodes. The low frequency resonance near 2.7 kHz is the translational mode noted in Fig. 2. The two angular degrees of freedom create resonances near 4.4 kHz but cannot be individually distinguished in this figure due to their approximately 5 Hz separation in frequency. Other flexural modes of the cloverleaf are responsible for the remaining resonances and the upward trend of the frequency response magnitude with increasing frequency is caused by parasitic electrical coupling between the actuation and sensing electrodes.

used in the controller synthesis. The light traces show the magnitude obtained from two additional experiments performed 20 minutes apart (the phase is not included so as to keep the figures uncluttered). It is quite evident that the dynamics exhibit a time-varying characteristic that can be accurately modeled by a re-parameterization of the frequency scale. It is widely accepted that the observed drift is caused by thermally-induced stresses within the sensor [12]. It was pointed out in Section II-

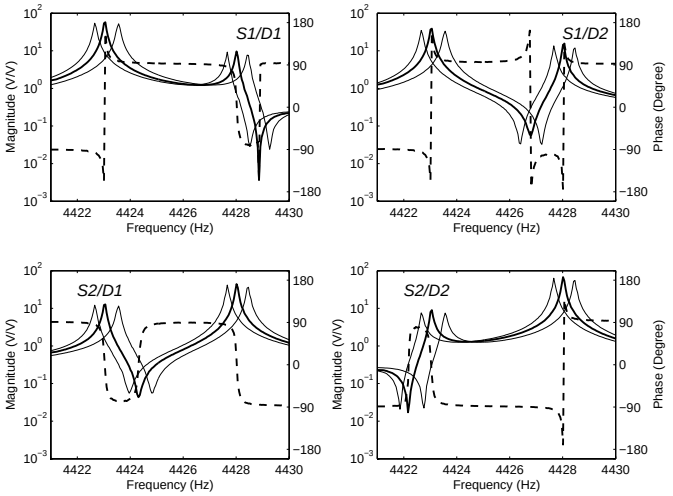

Fig. 4. The sensor's empirical frequency response in a neighborhood of 4420 Hz (magnitude, solid; phase, dash). The heavy traces are used for synthesizing input-output transformations. The light traces, from data sets acquired 20 minutes apart, illustrate the drift phenomenon.

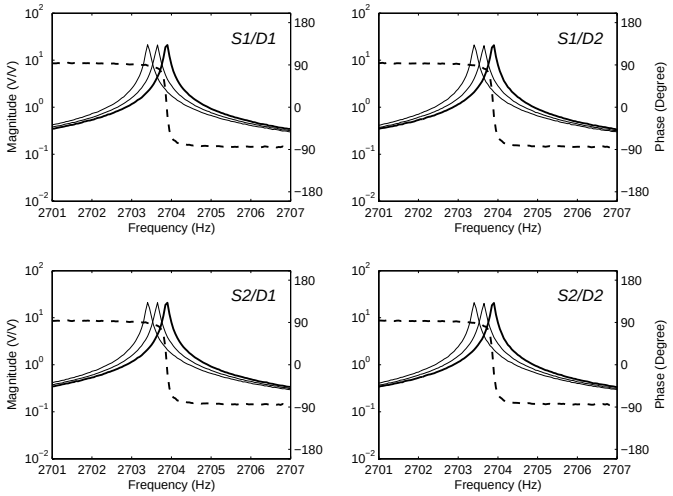

Fig. 5. The sensor's empirical frequency response in a neighborhood of 2700 Hz (magnitude –solid; phase –dash).

C that it is necessary for the drive control loop to excite one of the modes at the frequency corresponding to its peak response. This is evidently a challenging task due to the time varying nature of the plant.

The scale factors associated with each electrode establish the electrode's sensitivity (if used for sensing) or authority (if used for actuation). The prominent mode near 2700 Hz is not used in sensing angular rate, however, it is desired to regulate this mode about a given fixed position in order to hold constant the scale factors because they are dependent upon the nominal gap between the vibrating plate and electrodes. As previously noted though, the sense electrode signal conditioning electronics are configured so that the velocity is measured at points on the vibrating structure. Hence, maintaining a constant average plate position is not possible so instead this mode is damped, using feedback, about the equilibrium established in a constant acceleration field (reorienting the sensor in the gravitational field will cause a small shift in the plate's equilibrium position even though the mode at 2700 Hz is effectively damped). Detail of the sensor's frequency response in a neighborhood of 2700 Hz is shown in Fig. 5 along with several additional data sets illustrating the time-dependent frequency response.

C. Synthesis models

The JPL-Boeing sensors are designed to operate in a "degenerate" condition (when $\omega_1 = \omega_2$), however, the 5 Hertz difference between these modes in Fig. 4 is the native split due to manufacturing imperfections and packaging-induced stresses. There are post-fabrication methods that can be employed to further reduce the difference between modal frequencies, however, we will not elaborate on these but instead adopt an approach of decoupling the sensor's closely-spaced modes by performing input and output transformations. Recall that decoupling the modes near 4420 Hz into separate channels is a viable approach to reducing the quadrature error in (3) in lieu of achieving $\omega_1 = \omega_2$. Thus, it is desired to find input and output transformations that render the sensor dynamics diagonally dominant such that each diagonal channel has a different mode in the 4420 Hz region. With this approach the sensor's mass and stiffness matrices in (4) are not perturbed so that their generalized eigenvectors align with the native sensing electrodes frame but rather a set of *virtual* sense and drive electrodes are created to align with the generalized eigenvectors so that in these new coordinates the modes near 4420 Hz are decoupled to the greatest degree possible (see Section II-B). Although this diagonalizing process should be considered a part of the control synthesis procedure outlined in Section V it is convenient to address it while discussing sensor modeling.

Denoting the sensor frequency response as H_{sens}, we compute 2×2 constant-gain input and output transformations, denoted T_{in} and T_{out} respectively, such that $T_{out}H_{sens}T_{in}$ is diagonally dominant with a separate mode isolated to each channel. The individual gains in the transformations are rounded to signed, 8-bit constants (the largest magnitude being $\frac{255}{256}$) and then subsequently realized with an analog signal summing network possessing 8-bit programmable coefficients. The transformations are nearly orthogonal matrices (maximum condition number of 1.2) and hence do not exacerbate any modeling uncertainty associated with the frequency response data. The algorithm that produces T_{in} and T_{out} uses the frequency response data (heavy traces) from Fig. 4 to fit a model of the form (4) and then from this model we compute T_{in} and T_{out}. Presentation of the algorithm would take us too far from the paper's focus so it is our hope that the reader will be content with the final result and the fact that these modeling-oriented details will appear in the literature at some future date.

The virtual sense electrode potentials are denoted by $\tilde{S}$ and are related to the native sense pick-off potentials by $\tilde{S} = T_{out}\vec{S}$, where the notation of (4) is used to represent the native pick-off signal vector $\vec{S}$. Similarly, the new virtual drive electrode vector, denoted $\tilde{D}$, is related to the native drive electrode potentials by $\vec{D} = T_{in}\tilde{D}$. The wide-band *identified* frequency response of $T_{out}H_{sens}T_{in}$ is shown in Fig. 6. The efficacy of the transformations in decoupling the modes near 4420 Hz into separate chan-

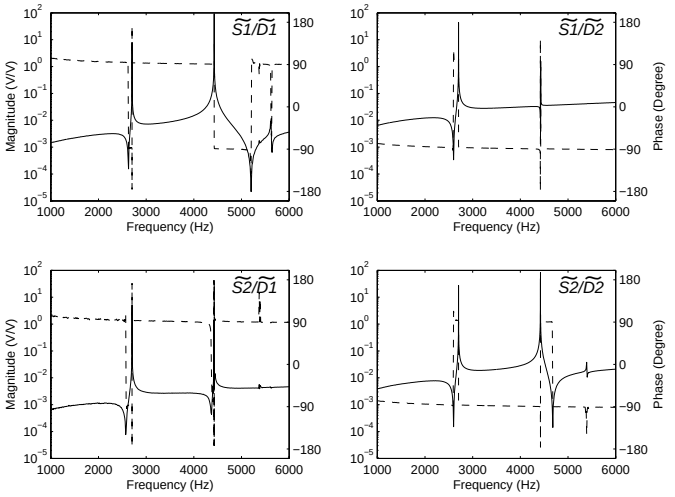

Fig. 6. Identified frequency response of the system $T_{out}H_{sens}T_{in}$ (magnitude, solid; phase, dash). Programmable analog networks implement T_{in} and T_{out} with 8-bit coefficient accuracy.

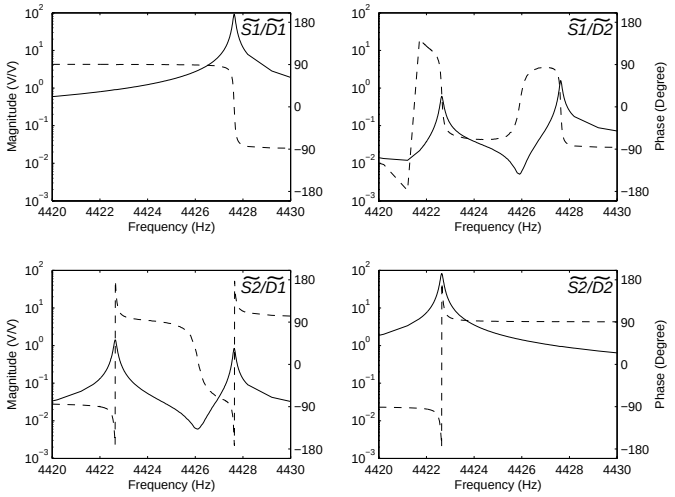

Fig. 7. Data from Fig. 6 near 4420 Hz illustrating the diagonally dominant transfer matrix with a different mode in each diagonal channel (magnitude, solid; phase, dash). The $\tilde{S}_1/\tilde{D}_1$ and channel $\tilde{S}_2/\tilde{D}_2$ are fit with the frequency domain models (5) and (6), respectively.

nels is evident in Fig. 7, which is a zoomed version of Fig. 6. In this frequency range the system is very much diagonally dominant with a different mode in each of the diagonal channels. In fact, the peak gain in the diagonal channels now exceeds the peak gain in the off diagonal channels by almost 40 dB. Fig. 8 displays the detail of the mode near 2700 Hz.

The sensor with the virtual actuation and sensing pick-offs is ideal for controller synthesis. We will have occasion to use simple single degree-of-freedom resonator models of the form $\alpha s/(s^2+cs+\omega^2)$ for the diagonal channels in Fig. 7. These transfer functions are

$$\begin{aligned} h_1 := \tilde{S}_1/\tilde{D}_1 &= \frac{\alpha_1 s}{s^2+c_1 s+(\omega_1 * 2 * \pi)^2} \\ &= \frac{56.4847s}{s^2+0.614618s+(4427.63*2*\pi)^2} \end{aligned} \tag{5}$$

$$\begin{aligned} h_2 := \tilde{S}_2/\tilde{D}_2 &= \frac{\alpha_2 s}{s^2+c_2 s+(\omega_2 * 2 * \pi)^2} \\ &= \frac{-60.1777s}{s^2+0.727505s+(4422.64*2*\pi)^2}. \end{aligned} \tag{6}$$

The time constants $(2/c)$ of h_1 and h_2 are approximately 3.25 seconds and 2.75 seconds, respectively. The off-diagonal transfer functions are not modeled since they are ignored in the synthesis procedure, although their effect is fully accounted for in the post-synthesis analysis. Similarly, a resonator model for the sensor's frequency response near 2700 Hz in the 2-2 channel of Fig. 8 is,

$$\begin{aligned} h_3 := \tilde{S}_2/\tilde{D}_2 &= \frac{\alpha_3 s}{s^2+c_3 s+(\omega_3 * 2 * \pi)^2} \\ &= \frac{16.3403s}{s^2+0.619545s+(2704.67*2*\pi)^2}. \end{aligned} \tag{7}$$

The time constant of this mode is 3.22 seconds. Only the 2-2 channel is modeled since it will be used in the synthesis of the rebalance control filter.

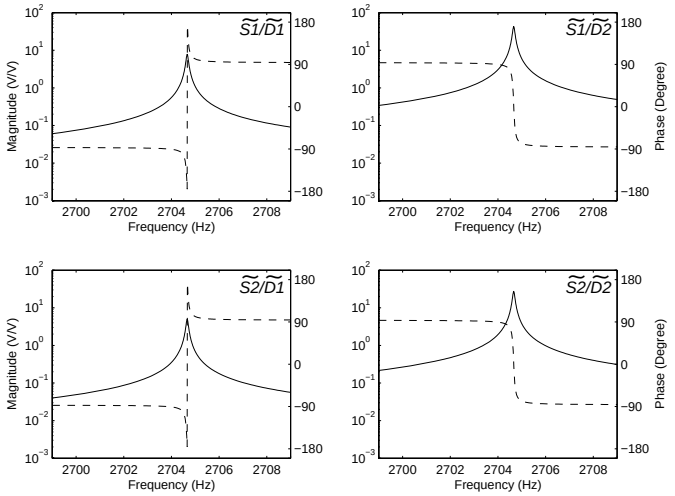

Fig. 8. Data from Fig. 6 near 2700 Hz (magnitude, solid; phase, dash).

IV. A CONTROL ASIC FOR VIBRATORY GYROSCOPES

This section describes the high-level functionality, architecture, and features of the special-purpose digital integrated circuit that implements the control and signal processing tasks for the family of JPL-Boeing gyroscopes. The ASIC architecture, shown in Fig. 9, employs a nonlinear *automatic gain control* for the drive loop, a linear filter for the rebalance loop and also provides filtering and demodulation of the rebalance signal for detection of its in-phase and quadrature components. The linear filters are realized with an FIR implementation. A feature that makes this ASIC attractive for a variety of rate gyro applications is its low power consumption because the ASIC core consumes only 0.37μW/tap/f_s, where f_s is the servo rate in kHz. The efficient hardware implementation that produces these low power requirements is described in [3]. Our exposition, however, will focus on specifying the filters and gains in Fig. 9 for a successful integration with a prototype sensor.

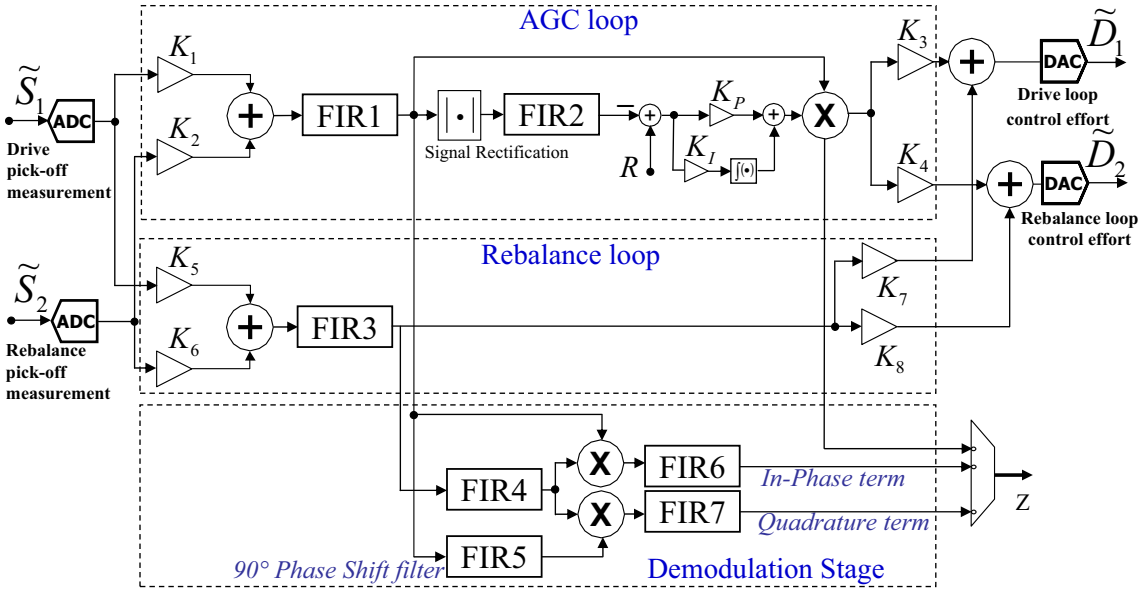

Fig. 9. ASIC architecture: each FIR filter possess 128 taps with 18-bit precision coefficients; gains K_1 through K_8 are programmable with 18-bit precision; K_P and K_I are programmable and represent the PI compensation for the AGC loop.

It is necessary at this point to discuss some details of the drive control loop in order to interpret the architecture in Fig. 9. Section II-C alluded to the fact that the drive control loop is a nonlinear filter that selectively excites and maintains a stable amplitude of a modal resonance in the sensor's dynamics even though these dynamics are slowly time-dependent (see Figs. 4 and 5). A classical approach would employ linear compensation with an external sinusoidal reference signal to drive the appropriate mode. This, of course, does not perform well unless the reference signal frequency tracks the resonance shift. Scheduling the reference frequency (with respect to a temperature measurement, for example) is not practical since it requires exhaustive testing of the sensor to precisely quantify the temperature dependence of its dynamics. The automatic gain control filter, or AGC, overcomes these challenges in an elegant and simple manner.

The AGC may be described using a simple harmonic oscillator model with a velocity measurement that determines the feedback force as shown in Fig. 10. The magnitude and sign of the feedback gain are determined by the AGC which employs an amplitude detection scheme, a comparator to generate the amplitude error signal and some compensation, commonly a PI filter, to set the feedback gain based on the amplitude error. The amplitude detection scheme is depicted as a rectifier and low-pass filter in Fig. 10 and R is the reference amplitude. The AGC functions in the following manner: if the oscillation amplitude is smaller than the desired reference amplitude, the AGC essentially destabilizes the harmonic oscillator with positive feedback; on the other hand, if the oscillation amplitude is too large, the AGC adds damping to the loop to reduce the oscillator amplitude. During steady state operation the AGC adds just enough energy to the oscillator to counter its intrinsic damping thereby maintaining a stable amplitude. The primary advantage is that the oscillator is always excited at its natural

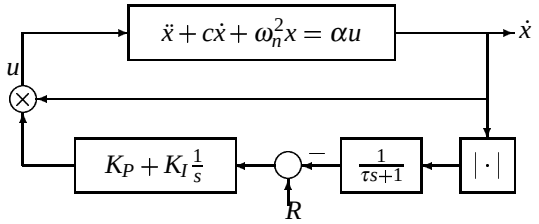

Fig. 10. Automatic gain control loop for regulating the oscillator amplitude. The AGC gains are K_I and K_P, R is the reference amplitude, and $\tau\omega_n >> 1$ is necessary for effective amplitude detection.

frequency, denoted ω_n in the figure, even when ω_n is slowly time varying. This will be clearly demonstrated by the experimental results in Section V.

The closed-loop system is nonlinear but the distinct time scales associated with the dynamics, corresponding to a fast oscillator frequency and a slower time scale for the modulated amplitude, can be exploited to analyze the closed-loop system response as demonstrated in [4], [8]. These references analyze the continuous-time case shown in Fig. 10 and reveal how the response of the oscillator depends upon the various controller parameters. In contrast to the thoroughly analyzed continuous-time models, however, the ASIC filters are purely discrete-time and FIR. The discrete-time implementation is pursued because of the flexibility it imparts to the ASIC for adapting to a wide range of sensor dynamics. The choice of FIR filters is due to the desire to avoid limit cycling because all computations use fixed-point arithmetic. Consequently, rigorous analysis of the closed-loop system consisting of a sampled-data sensor model and a controller employing fixed-point implementations of FIR filters is beyond the scope of this paper.

We now return to the ASIC architecture in Fig. 9 and indicate its main features:

1) **Filters.** There are seven, 128-tap, fully programmable

FIR filters and a proportional-integral section for the automatic gain control.

2) **Precision.** The computations are fixed-point with 18-bit input data precision, 18-bit output data precision, 18-bit coefficient precision, and 20-bit internal data precision (24-bits for the integrator).
3) **I/O Signals.** The $\tilde{S}_1$ and $\tilde{S}_2$ input channels are the signals from the analog network that decouples the two modes near 4420 Hz. Similarly, the $\tilde{D}_1$ and $\tilde{D}_2$ output channels are the analog virtual drive electrode potentials. The A-D and D-A converters are not integrated into the ASIC but are shown in order to clarify the input and output signals. The Z channel reads out the AGC gain, and in-phase and quadrature components of the rebalance-loop signal.
4) **I/O gains.** Gains K_1 through K_8 are programmable and can implement the T_{in} and T_{out} transformations discussed in Section III-C.
5) **AGC loop.** The AGC consists of FIR1, the rectifier, FIR2, the proportional-integral (PI) compensation and programmable limiters on the integrator and AGC gain.
6) **Rebalance loop.** The rebalance loop compensation is implemented in FIR3.
7) **Demodulation.** The rebalance signal may be demodulated with respect to the input signal to the AGC stage. FIR4 and FIR5 are typically used to shift signal phases and FIR6 and FIR7 are low-pass filters located after the multipliers.

Evidently the ASIC implements a diagonal controller if one considers the signal flow *between* the gains K_i, $i = 1,\ldots,8$, in Fig. 9. Our experiments will demonstrate that there is some crosstalk between the channels but that it can be effectively ignored. The ASIC servo rate for real-time control is constrained by the filter complexity. For example, with a core voltage of 2.5 V the ASIC operating frequency, denoted F_{clk}, is 37 MHz, and because one multiply can be performed per cycle, the maximal I/O rate of the ASIC is determined by the longest FIR path length through the AGC (in which case the path would be {FIR1,FIR2}) or the rebalance-demodulation path (in which case the path would be {FIR3,FIR4,FIR6/7} or {FIR1,FIR5,FIR7}). Hence, the maximal I/O rate is then $F_{clk}/\max(N_1+N_2, N_3+N_4+N_6, N_3+N_4+N_7, N_1+N_5+N_7)$, where N_i is the length of FIRi. If we assume that all of the FIRs have maximum length (128 taps each), then the servo rate limit of the ASIC is $F_{clk}/(N_3+N_4+N_6) = 96.4$ kHz and the total core power consumption is 31.9 mW. If the complexity of the filters is reduced, then a corresponding increase in the maximum servo rate, or a reduction in power requirements, can be realized.

A. ASIC Brassboard

The ASIC is a digital processor and requires interfacing to the analog sensor signals. A brassboard was constructed to accomplish this task. It hosts the ASIC, an audio codec chip with stereo analog-to-digital converters (ADC) and digital-to-analog converters (DAC), analog antialiasing and smoothing filters, six 12-bit ADCs for gyro biasing, three with a high-voltage (-10V to +50V) output range, and a Xilinx SRAM-based field-programmable gate array (FPGA). The FPGA contains digital logic for a personal computer interface via the IEEE 1284 enhanced parallel port (EPP), logic for setting the bias DACs, configuring the codec, and configuring, controlling, and monitoring the digital ASIC. The FPGA also includes circuitry to synchronize data flows between the ASIC and the codec and to measure the gyroscope rocking frequency with mHz resolution. The FPGA may be replaced with more special purpose circuitry, however, it was included in this brassboard version as a back-up processor in the event that the ASIC did not perform to expectations. The hardware and signal relationships are shown in Fig. 11.

The computer controls and monitors the brassboard through Enhanced Parallel Port (EPP) interface. In EPP mode, two independent 8-bit I/O channels are available, one designated as address and the other designated as data. Both share 8-bit bidirectional bus. The data/address transfer handshake for EPP is handled by hardware allowing transfer rates on the order of 500 KB/s to 2 MB/s. Two independent strobes differentiate input/output transfers to the data/address channels. The EPP interface is implemented as part of the FPGA and an addressed I/O approach is employed where the contents of an address register, updated by address write operations, is used to select the source or destination of data read/write operations.

The codec chip is a Crystal CS4224 and is a highly integrated, high performance, 24-bit, audio codec providing stereo analog-to-digital and stereo digital-to-analog converters using delta-sigma conversion techniques. One channel of the codec is used for the AGC loop and the other for the rebalance loop. The servo rate is set by the codec at 48 kHz so all ASIC filters are designed with this value. The conversion technique provides 24-bit resolution at the expense of an appreciable transport lag. In fact, the total transport lag associated with the control hardware on the brassboard is 980 μs and this is almost entirely due to the ADC and DAC conversion process. A lag of such long duration would under typical circumstances vastly limit the bandwidth of the closed-loop system however we will demonstrate that effective regulation of the lightly damped modes is still possible.

V. Controller synthesis and implementation

The decoupled sensor dynamics in a neighborhood of 4420 Hz permits a greatly simplified design process because the off-diagonal terms in Figure 7 can be ignored. The wide-band response in Figure 6 demonstrates that large loop gain can be achieved in a neighborhood of 2705Hz and 4420Hz with relatively modest controller gain.

Some notation is now introduced. The sensor is designated by P, the control hardware with ASIC filters (programmed according to Figs. 12 and 15) is designated by C, and the ijth channel of P is denoted P_{ij}, with the same convention applying to C. Our design process ignores the off-diagonal coupling in both C and P but the analysis of the designs in Section V-C includes the coupling effects. Hence, we will have occasion to discuss properties of the loop gains PC and CP and in this case the ijth channel of PC is denoted $(PC)_{ij}$ with the same convention applying to CP.

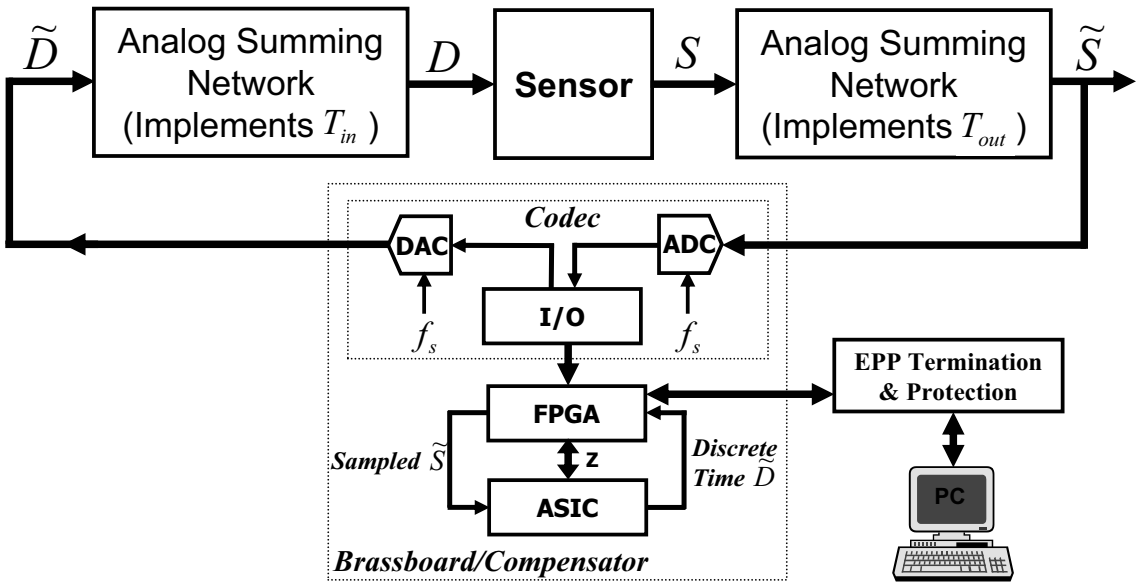

Fig. 11. This figure clarifies the signal flow in the closed-loop system. The sensor dynamics shown in Figs. 3, 4, and 5 reveal the system response with direct access to the sensor's sensing and drive electrodes (the S/D transfer functions in the diagram). The analog networks create virtual drive and sense electrodes that decouple the gyro dynamics in a neighborhood of the 4420 Hz modes. The system dynamics from the perspective of the new measurements and actuators ($\tilde{S}/\tilde{D}$ transfer functions) are shown in Figs. 6, 7, and 8. The ASIC has access to the sampled $\tilde{S}$ signals and specifies the discrete-time $\tilde{D}$ sequence. The servo rate, denoted f_s, is 48 kHz. The implemented controllers' frequency responses, including the brassboard I-O dynamics, are shown in Figs. 12 and 15. These filters are discussed in detail in Section V.

A. AGC design

The automatic gain controller is defined by specifying FIR1, FIR2, the reference amplitude R, and the AGC gains K_P and K_I. The plant model for this loop is the 1-1 channel in Fig. 7 which is described by (5). The performance criteria outlined in Section II-C are amplitude stability with a constant reference, modal frequency tracking capability, and closed-loop bandwidth. We can add one more criteria, namely excitation rate from zero amplitude (this determines how rapidly the sensor can reach steady-state). The first two criteria must be evaluated through experimentation because the sources that induce the modal frequency drift include the complicated dependence of the sensor dynamics on temperature and thermistor effects in the signal conditioning electronics, both of which are difficult to analytically capture. If the AGC is implemented using the analog filters shown in Fig. 10 then the last two criteria may be accurately predicted from the parameters in the simple parametric plant model (5) and those in the controller, as illustrated in [8], however, the fact that the AGC uses FIR filters complicates the analysis of the closed-loop system.

Although the closed-loop system is easily simulated with FIR filters, we forgo presentation of simulation results in the interest of minimizing the paper length. A useful heuristic, however, emerged from [8] that is somewhat independent of the filter implementation details and that may be applied here: the two-time scale behavior of the closed-loop system allows us to design the AGC filters by assuming the compensator is a fixed, linear filter. That is, from the sensor's perspective, the compensation can be treated as a fixed filter for any given sensor response because the resonances in the sensor are much higher in frequency than the bandwidth of the AGC gain. Here we refer to the *AGC gain* as the subsystem that modulates the feedback signal in Fig. 9 to distinguish it from the *compensator* that is defined as the AGC gain times FIR1. Suppose the resonance with AGC loop closed has reached a steady state amplitude. Let a denote the amplitude of the post-FIR1 signal and assuming $K_I = 0$ the steady-state AGC gain is given by $(R-a)\cdot K_P$. We set $K_I = 0$ for convenience because this reduces the AGC gain to an affine function of a and whose bandwidth is determined by FIR2. This gain modulates the output of FIR1 to determine the control effort.

The details of the compensator implemented in Section VI are now presented. The measured frequency response of compensator when the input amplitude is zero (that is, FIR1 times the AGC gain $R\cdot K_P$) is shown in Fig. 12. This figure represents the controller magnitude and phase when the amplitude of the AGC input ($\tilde{S}_1$) is much smaller than the reference, i.e. $R \gg a$. In fact, the frequency response was measured by setting FIR2$\equiv 0$ thereby yielding the compensator $K_P \cdot R \cdot$ FIR1. Under normal operation FIR2 is a 51-coefficient linear phase, low-pass filter with 50 Hz cut-off frequency. Note, however, that the magnitude shape in Fig. 12 is defined solely by FIR1 since R and K_P are real gains (also, $K_1 = K_3 = K_6 = K_8 = 1$ and $K_2 = K_4 = K_5 = K_7 = 0$ since the plant has been decoupled via T_{in} and T_{out}; furthermore, the compensator servo rate is 48kHz for these tests). The corresponding loop gain is shown in Fig. 13 and confirms that only the response of the ω_1-mode is actually regulated because outside of a very small band around ω_1 the loop gain magnitude is less than one. It is especially important to filter out the ω_3-mode so that it cannot be excited by the AGC nor corrupt the amplitude measurement of the ω_1-mode. The notch at ω_3 in Fig. 12 is quite effective in this regard.

The reader will observe in Fig. 12 that there is significant group delay in the controller phase. This is caused by the 980

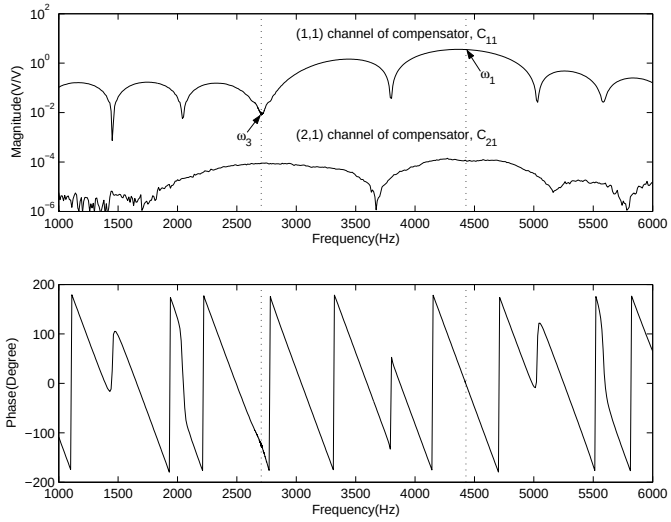

Fig. 12. Measurement of the AGC filter gain (top) and phase (bottom). The AGC is implemented in the (1,1) channel of the compensator, i.e. $\tilde{D}_1/\tilde{S}_1$ transfer function of the compensator in Fig. 11, and is denoted C_{11}. The FIR2 filter used in the amplitude detection scheme is set to zero for the measurement of this transfer function to demonstrate the initial controller gain when the loop is closed about a quiescent sensor. FIR1 is a 121-tap filter that achieves attenuation at 2700Hz and the target phase at the ω_1-mode frequency. The compensator cross-channel response $C_{21} = \tilde{D}_2/\tilde{S}_1$ is also shown (magnitude only) and demonstrates that the extremely small magnitudes of the compensator's off-diagonal terms justifies our treating it as a diagonal system.

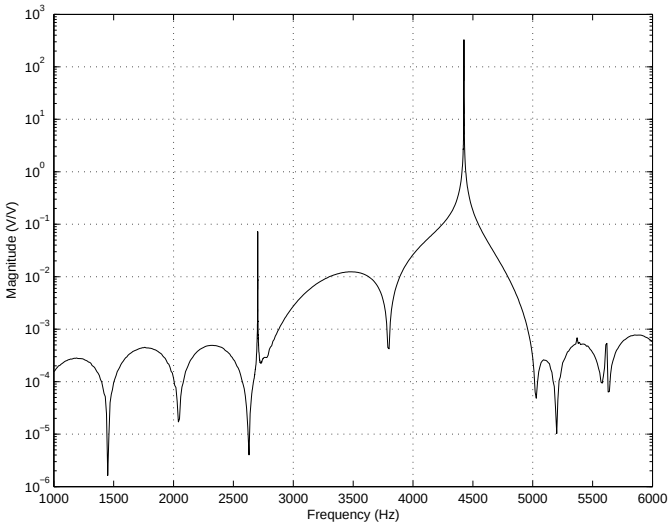

Fig. 13. This figure shows the magnitude of the loop gain of the AGC channel $P_{11}C_{11}$. The plot is constructed from the open-loop frequency response data of the sensor and control hardware. The compensator C_{11} is designed to attenuate the ω_3 resonance in the loop gain, and to excite the ω_1 resonance by appropriate shaping of loop phase. P_{11} represents the decoupled sensor transfer function $\tilde{S}_1/\tilde{D}_1$ whose dynamics are approximated by (5) in a neighborhood of ω_1.

μs transport lag associated with the 24-bit sigma-delta ADC and DAC converters. Additional lag is introduced by the anti-aliasing and smoothing filters on the brassboard and by the FIRs themselves. By bookkeeping all of these contributions it is possible to shape the phase of FIR1 to achieve the target loop phase of 0 (modulo 360) degrees at ω_1. Thus, the AGC loop feedsback a signal whose phase is 0 degrees at the model frequency and by our positive feedback convention, this implies the mode is destabilized when the loop is initially closed around a quiescent sensor. Nyquist plots readily demonstrate that the desired phase is achieved at the modal frequency, however, these plot are deferred until Section V-C.

The picture revealed thus far demonstrates the compensator response for very small oscillator amplitudes. When the loop is closed about a quiescent sensor the small-input condition is satisfied and the large loop gain at ω_1 in conjunction with the loop phase lead to exponential growth in the amplitude of the ω_1-mode. The initial growth rate of the amplitude is computed to be $\exp\left(\frac{1}{2}(\bar{K}\alpha_1 - c_1)t\right) = \exp(100.5t)$, where $\bar{K} := K_P \cdot R \cdot \mathrm{FIR1}(j2\pi\omega_1) = 3.57$ is the compensator magnitude at ω_1, and α and c are the modal parameters from (5). The growth in amplitude is detected by the AGC and, accordingly, a reduction in the AGC gain occurs because the amplitude error signal $R - a$ has been reduced. This has the effect of scaling the magnitudes in Figs. 12 and 13 to smaller values than those shown in the plots. In fact, when the steady-state amplitude of the ω_1-mode is achieved, the loop gain has been reduced to one. The results of an empirical test of the compensator are shown in Fig. 14 and reveal the dependence of its steady-state output amplitude for a given input amplitude. This figure was created by driving the compensator with $a\sin\omega_1 t$ over a range of amplitudes to emulate the sensor's resonance at ω_1. Note that the compensator gain at zero input amplitude obtained from plot shown at the bottom of Fig. 14 matches the compensator magnitude at ω_1 in Fig. 12. The figure demonstrates that the compensator gain is reduced as the input amplitude increases.

A limit on the maximum compensator magnitude is imposed by the 980 μs transport lag. In the case that the closed-loop system can be treated as a linear system, excessive gain will lead to the formation of large lobes in the closed-loop sensitivity function. In fact, these lobes are present in the sensitivity function of the rebalance channel shown in Fig. 22. If we restrict the controller magnitude to be less than five at the sensor's resonances then this sensitivity function is less than two at all frequencies. Although these arguments are not directly applicable to the nonlinear AGC, we make a similar restriction on its maximum magnitude. With regard to the parameters comprising the controller magnitude, $R = 0.405$ (volts) is chosen to achieve a given peak mechanical displacement in the sensor, and K_P and the FIR1 coefficients are scaled to optimize fixed-point computation performance while respecting the maximum gain limit noted above.

B. Rebalance loop control design

The linear rebalance loop controller implemented by FIR3 uses $\tilde{S}_2$ as its measurement and $\tilde{D}_2$ is its output. One objective of this loop is to achieve disturbance rejection in a neighborhood of 2700Hz, which corresponds to damping an elas-

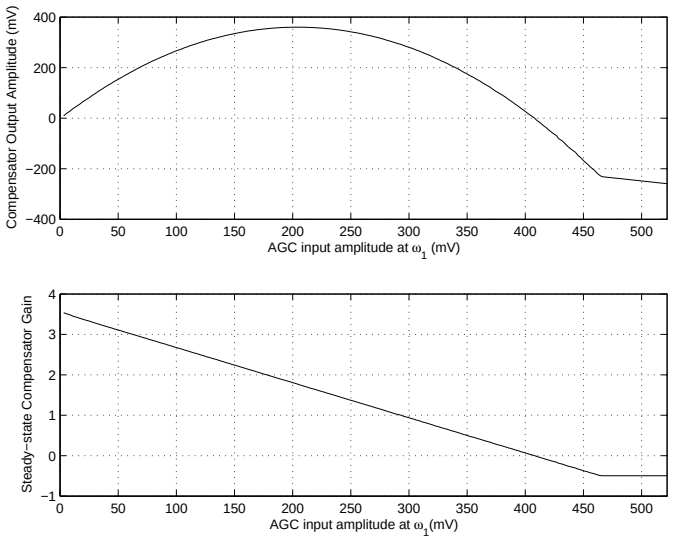

Fig. 14. Even though the AGC is a nonlinear compensator it is still possible to define its steady-state input-output properties. This figure displays the steady-state AGC compensator behavior in two different formats. The top plot shows the compensator output amplitude as a function of the input amplitude when the input is a sinusoid with frequency ω_1. Since very little excitation amplitude is required to maintain a constant response amplitude of the sensor at ω_1, the sensor's final amplitude is very near where the graph crosses zero. Alternatively, the bottom plot is what we defined in Section V-A as the AGC gain. It is the top graph normalized by the input amplitude. The top representation is useful for determining the peak actuator effort (cf. Fig. 23) while the bottom representation permits rapid identification of the steady-state amplitude (which differs slightly from R since only proportional compensation is used on the amplitude error).

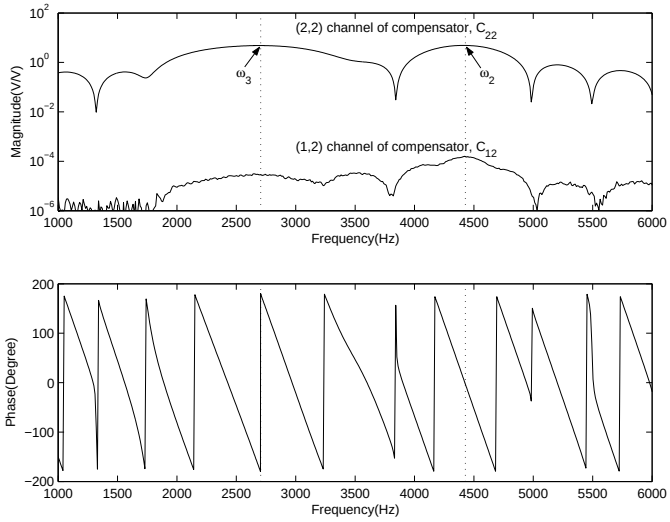

Fig. 15. The measured rebalance controller frequency response, denoted C_{22}, is the (2,2) channel of the compensator. Note that the cross-channel coupling transfer function C_{12} is zero for all practical purposes. The controller magnitude at ω_2 and ω_3 is approximately 4.6 implying that the closed-loop bandwidth associated with the ω_2-mode exceeds 20 Hz. This channel's sensitivity function magnitude is also less than two (Fig. 22).

tic mode whose excitation can lead to degraded sensor performance (see Section III-B). Since this mode is easily excited by linear acceleration of the sensor case we can ascertain the effectiveness of the control by analysis of the sensitivity function. The primary control objective of this loop, however, is to reject the disturbance injected by the Coriolis term in (1) into the second degree of freedom, i.e. $\tilde{S}_2$, and in this case it is not the magnitude of the sensitivity *per se* but rather the rapidity with which the signal reaches steady-state after the disturbance is imposed. In other words, the time constant of the closed-loop system is the pertinent metric for controller performance.

As in the AGC design, velocity-to-force feedback for the rebalance loop is used to achieve the objectives. The controller phase must be shaped to produce -180 (modulo 360) degrees loop phase at ω_2 and ω_3 in order to realize optimum damping of these modes under the action of positive feedback. One such filter, shown Fig. 15, achieves the target loop phases at ω_2 and ω_3. This figure is an empirical frequency response of the control hardware with the test signal injected at $\tilde{S}_2$ and measurements made at $\tilde{D}_2$ and $\tilde{D}_1$, with the latter measurement representing the cross-channel coupling. The loop gain magnitude of the rebalance loop is shown in Fig. 16 to demonstrate that the compensator does indeed regulate the ω_2 mode and ω_3 mode but unlike the AGC, the phase of this loop is shaped to achieve increased damping of the resonances.

The relationship between closed-loop sensitivity and time constant can be derived using the simple harmonic oscillator model in Fig. 17. Although the controller is much more complex than the constant gain K in this model, the fact that the loop gain can only be made larger than one in a narrow band about each mode suggests that the large lag discussed in Section V-A, and also present in the rebalance controller channel, can be effectively ignored if the controller gain is not excessive. This view is also supported by Nyquist plots of the loop gains shown in Section V-C, where one cannot discern the effect of the lag.

The loop gain magnitude from the model in Fig. 17 at ω_n, which we denote L, and closed-loop time constant are given by $K\alpha/c$ and $2/(K\alpha + c) \approx 2/(K\alpha)$, respectively. Hence, if the closed-loop time constant is desired to be smaller than some minimum time constant $\tau_{\min}$, then we have the loop-gain constraint $L > 2/(c\tau_{\min})$ which in turn implies that the sensitivity function magnitude $|u/d|$ at ω_n satisfy $|u/d| \approx 1/L < c\tau_{\min}/2$. This is achieved with a controller gain satisfying $K > 2/(\alpha\tau_{\min})$. For example, a minimum bandwidth of 20Hz corresponds to $\tau_{\min} = 1/(20 \cdot 2 \cdot \pi)$ and the lower bound on the controller gain is computed to be $K > 4.18$ using the plant parameters from (6). The controller magnitudes at ω_2 and ω_3 are 4.87 and 4.82, respectively, and respect the maximum value noted in Section V-A that ensures the peak sensitivity is less than two.

C. Multi-input/Multi-output analysis

The ASIC filters were designed without regard to cross-coupling in the controller and subsequent tests of the hardware confirmed that the off-diagonal terms are negligible. We also chose, for the sake of design expediency, to ignore the small cross-coupling in the sensor in a neighborhood of 4420 Hz while synthesizing the filters. Since the loop gains PC and CP have two distinct regions where their norms are greater than one, we can present their Nyquist plots in separate figures from 1kHz to 3.5kHz and from 3.5kHz to 6kHz. To wit,

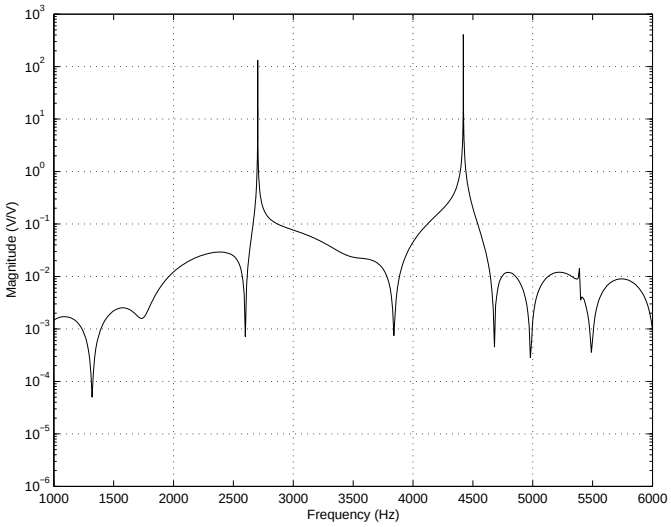

Fig. 16. This figure shows the magnitude of the loop gain of the rebalance channel, $P_{22}C_{22}$. Open-loop frequency response data were used to construct the loop gain. The design plant dynamics are approximated by (6) and (7) in a neighborhood of ω_2 and ω_3.

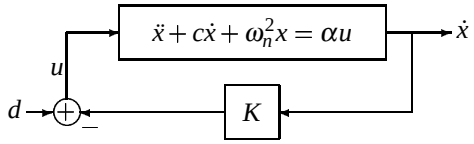

Fig. 17. Oscillator model for deriving the relation between the closed-loop time constant and sensitivity function.

Figs. 18 and 19 show Nyquist plots of both loop gains over these frequency bands. Nyquist plots complement Figs. 12, 15, 13 and 16 because they clearly demonstrate that the target loop phases are achieved at the resonances. Since the AGC is nonlinear, these plots were created by fixing the AGC gain at its maximum value as we described for Fig. 12. Note that $(PC)_{ii} = (CP)_{ii} = P_{ii}C_{ii}$, for $i = 1,2$, as a consequence of the diagonal controller.

Let us first consider the loop gains in the lower frequency range encompassing the ω_3-mode (see Fig. 18). Recall that the rebalance controller C_{22} has the ancillary task of damping the resonance at ω_3 and that the magnitude of C_{11} (the AGC) is greatly reduced so as not to regulate this mode. The ability of the ASIC to shape the loop phase to the desired values is evident in $(PC)_{22}$ because the phase is almost -180 degrees at the maximum loop magnitude in this channel. Furthermore, since C_{11} possess very low gain in this region, it is reasonable to treat the closed-loop as linear and time-invariant and thereby employ singular value analysis tools for performance assessment. It is important, however, to first select the appropriate closed-loop functions for performance analysis. For example, in Fig. 18 it is evident that $(PC)_{i1} \ll 1$, $i = 1,2$, leading one to conclude that the closed-loop system should suffer from very poor disturbance rejection if we are considering the output sensitivity function as the best function to capture performance. A more realistic model of the disturbance, and in particular linear acceleration of the sensor case, is effectively modeled as a disturbance at the plant input (summed with $\tilde{D}$ in Fig. 11) because

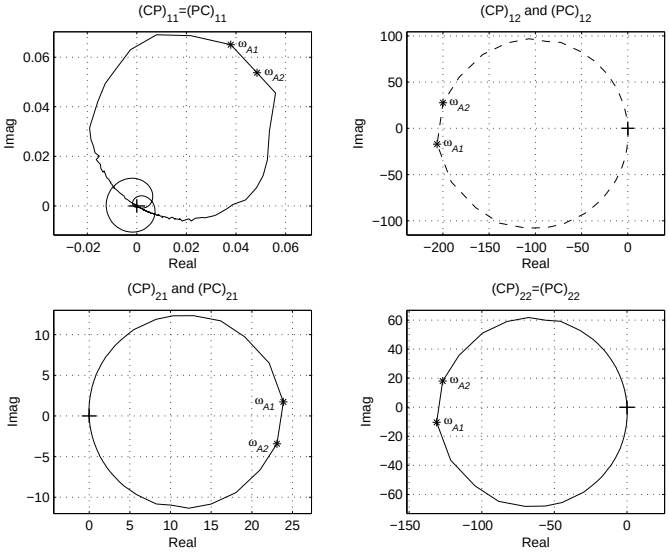

Fig. 18. *CP* (solid) and *PC* (dash) from 1 kHz to 3.5 kHz constructed from empirical frequency responses of the sensor and controller. The direction of increasing frequency can be determined from $\omega_{A1} = 2704.60$ Hz and $\omega_{A2} = 2704.61$ Hz. Both are shown with largest controller gain determined by the AGC. Positive feedback is used when the loop is closed.

the acceleration-induced forces applied through the supports to the sensor's vibrating structure are similar to the electrostatic forces created by the actuation electrodes. The output, however, is taken at the plant output since we are interested in how these disturbances affect the sensor response. Hence, we are interested in the magnitude of the closed-loop transfer function PS_I, where $S_I = (I - CP)^{-1}$ is the input sensitivity function, for assessing the performance of the controller with regard to linear acceleration disturbance over the band 1 kHz to 3.5 kHz. A plot of the maximum singular values of PS_I (using *measured* closed-loop transfer functions) is compared with the maximum singular values of P in Fig. 21 and shows a 40 dB reduction in the peak sensitivity. We elaborate on this result in Section VI.

The loop gains in the higher frequency band extending from 3.5 kHz to 6 kHz are shown in Fig. 19. Note that the phase and gain of $(PC)_{11}$ will lead to destabilization of the ω_1-mode but that the ω_2-mode is damped as is evident from $(PC)_{22}$. The variable AGC gain scales the magnitudes of $(PC)_{11}$, $(PC)_{21}$, $(CP)_{11}$, and $(CP)_{12}$ to smaller values from maximums shown in the plots. For example, when the loop is initially closed about the sensor and the amplitude estimate in the AGC is close to zero, the $(PC)_{11}$ and $(PC)_{21}$ loci (alternatively $(CP)_{11}$ and $(CP)_{12}$), are given as shown in the plot, however, when the amplitude estimate is half that of the reference amplitude then the magnitude of these plots is halved which leads to continued growth in the amplitude of the ω_1-mode but at a reduced rate. This reduction in AGC gain continues to a critical value at which the AGC has excited the ω_1-mode to a steady-state amplitude and the locus of $(PC)_{11}$ essentially passes through $+1 + j0$ on the Nyquist plot.

A more rigorous argument for the stability condition of the closed-loop system for different AGC gains is desired, however, because the off-diagonal elements in Figs. 18 and 19 have appreciable magnitudes (> 1) that heretofore have been ignored.

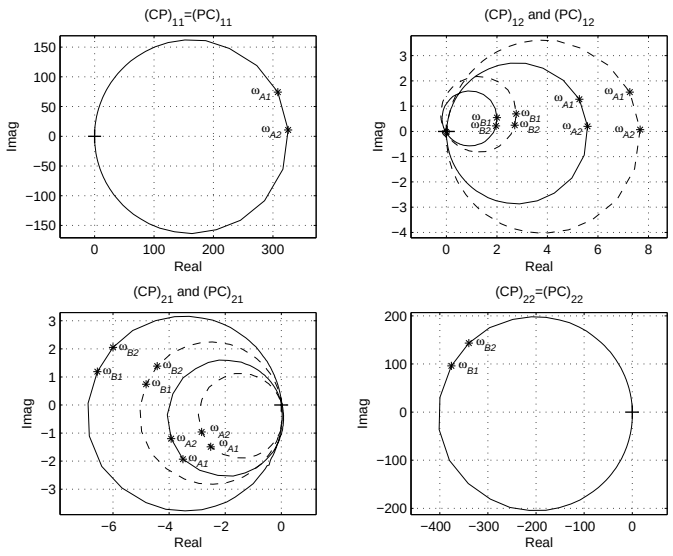

Fig. 19. *CP* (solid) and *PC* (dash) from 3.5kHz to 6kHz. The following frequencies are shown on the plots $\omega_{A1} = 4427.60$ Hz, $\omega_{A2} = 4427.61$ Hz, $\omega_{B1} = 4422.60$ Hz, and $\omega_{B2} = 4422.61$ Hz. Both are shown with largest controller gain determined by the AGC. The variable AGC gain affects the magnitude of $(PC)_{11}$, $(PC)_{21}$, $(CP)_{11}$, and $(CP)_{12}$ by scaling. Positive feedback is used when the loop is closed.

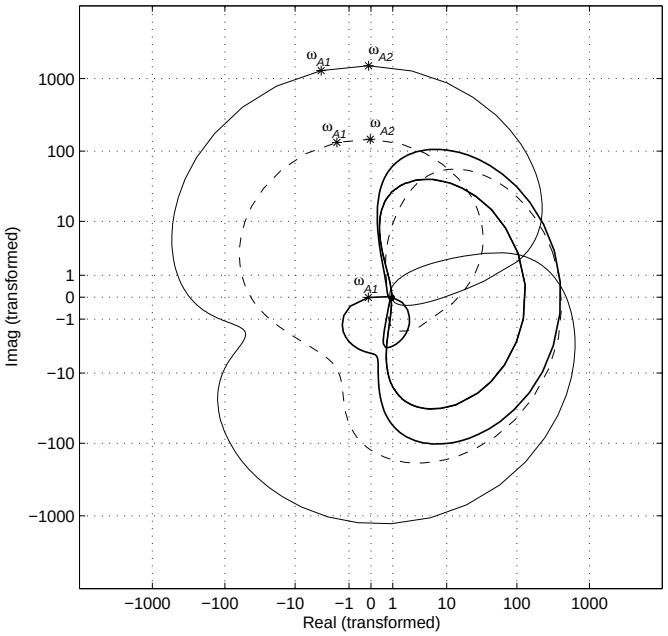

Fig. 20. Nyquist plots of $\det(I - P\Delta C)$ for $\Delta = [\delta, 0; 0, 1]$, $\delta = \{1, 1/10, 1/328\}$, with logarithmically compressed magnitude. The frequency scale spans 1 kHz to 6 kHz. The loci corresponding to the larger AGC gains imply that the closed-loop system has two unstable closed-loop poles. When the AGC gain is reduced to the point where the locus passes through the origin it is revealed by the frequency marker that it is the ω_1 mode that is excited by the AGC ($\omega_{A1} = 4427.63$ Hz and $\omega_{A2} = 4427.64$ Hz).

Thus, the stability question is more intricate than the single loop arguments we employed for designing C. We will apply the multivariable Nyquist criterion for several AGC gains to resolve this question. A Nyquist plot of $\det(I - P\Delta C)$, where $\Delta := [\delta, 0; 0, 1]$, $\delta = \{1, 1/10, 1/328\}$, is shown in Fig. 20 for $\omega \in [1\text{kHz}, 6\text{kHz}]$. This plot employs the plant and controller empirical frequency response data over this frequency band. Note that the δ parameter models the reduction of the (1,1) channel loop gain (due to the reduction in C_{11} magnitude via the AGC) from its maximum value of 328. The Nyquist plot reveals that for $\delta = \{1, 1/10\}$ the closed-loop system possess one pair of closed-loop poles with real part strictly greater than zero since there is one clockwise encirclement of the origin. When $\delta = 1/328$ the locus of $\det(I - P\Delta C)$ essentially passes through the origin and hence a pair of closed-loop poles lie on the $j\omega$-axis. The location of these poles is revealed by the ω_{A1} marker to be in extremely close proximity to ω_1, i.e. the resonance that was selected for excitation by the AGC. The axes of the Nyquist plot have been scaled to more clearly reveal the details near the origin: if L represents a complex quantity to be graphed on a Nyquist plot we scale it according to $\log_{10}(1 + |L|)e^{j\angle L}$ so its phase does not change but its amplitude is logarithmically compressed.

VI. Closed-loop results

The closed-loop results are presented in this section. We first discuss the disturbance rejection properties in the lower frequency region where the ω_3-mode can be excited by linear acceleration of the sensor case. The AGC controller has very small gain in this region and so we can perform a standard closed-loop frequency response experiment even though the mode at ω_1 is excited to a stable amplitude during the course of the experiment. We discussed in Section V-C the motivation for considering the maximum gain of PS_I as the appropriate measure of linear acceleration disturbance rejection. A plot of $\bar{\sigma}(PS_I)$, where $\bar{\sigma}$ denotes the maximum singular value, over the interval 2.4 kHz to 3 kHz, is shown in Fig. 21 and demonstrates a 40dB reduction in peak gain compared to the open-loop case $\bar{\sigma}(P)$. This result appears counterintuitive because the loop gains in Fig. 18 have either a row or column of low gain in the neighborhood of ω_3 and hence both sensitivity functions possess input directions that provide essentially no signal attenuation. If a disturbance near ω_3 is introduced into the low-gain direction at the plant input, however, the (2,2) controller element generates a signal in a different input direction which has the effect of canceling the disturbance through the plant due to the highly-coupled nature of its dynamics near ω_3 (cf. Fig 8). This provides the excellent performance observed in Fig. 21 for all input directions.

The sensitivity function can also be measured for the (2,2) channel due to the decoupled plant dynamics in a neighborhood of ω_1 and ω_2 and because of the low gain of the AGC loop at ω_3. Furthermore $S_{O,22} = S_{I,22}$ (the controller is diagonal) so only $S_{I,22}$ is displayed in Fig. 22. The primary objective of the rebalance control is to regulate ω_2-mode amplitude to zero and the $S_{I,22}$ function can be used to quantify the closed-loop time constant associated with the ω_2-mode. The sensitivity function exhibits 50 dB of attenuation at ω_2 and this corresponds to a closed-loop bandwidth of nearly 20 Hz.

We now turn our attention to the drive loop regulated by the AGC. Both the rebalance and AGC loops are closed simultaneously about the quiescent sensor to obtain the response shown in Fig. 23. Ideally, the $\tilde{D}_2$ and $\tilde{S}_2$ signals associated with the

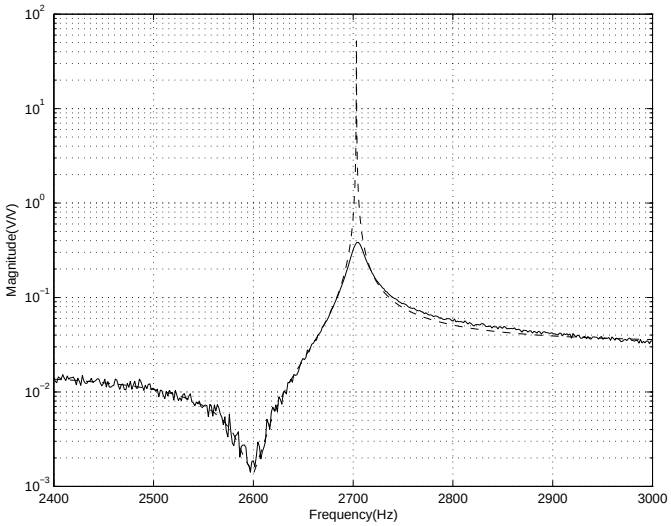

Fig. 21. The ability to reject linear acceleration (vibration) disturbances is quantified by $\overline{\sigma}(PS_I)$ (solid) in relation to $\overline{\sigma}(P)$ (dash). The closed-loop transfer function magnitude is reduced by over 40 dB compared to the open-loop sensor. The closed-loop transfer function PS_I was measured for this plot.

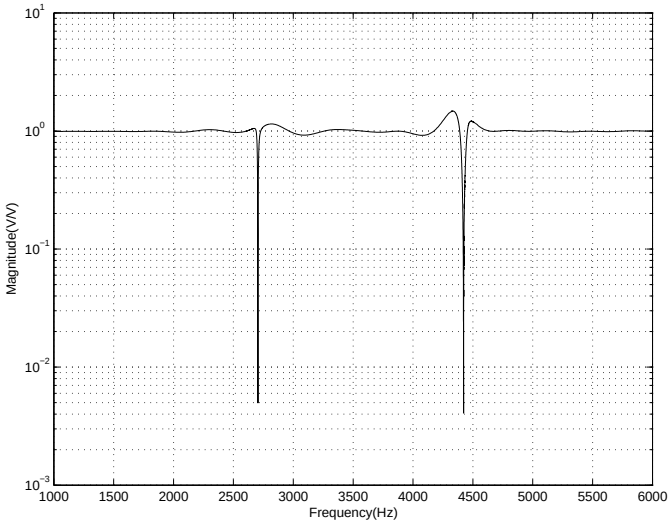

Fig. 22. The measured (2,2) channel of S_I shows a 50 dB reduction at ω_2. This corresponds to a 20 Hz bandwidth for the ω_2-mode.

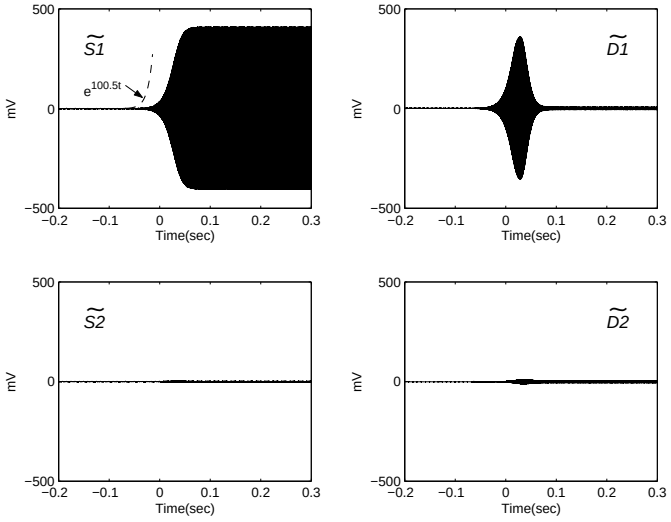

Fig. 23. Start-up transients in the closed-loop system. The $\tilde{S}_1$-$\tilde{D}_1$ pair represent the AGC response which is largely decoupled from $\tilde{S}_2$ and $\tilde{D}_2$, which represent the rebalance loop pick-off signal and rebalance loop control command, respectively. A plot of $e^{100.5t}$ is shown in the $\tilde{S}_1$ plot to demonstrate that the amplitude envelope predicted by the initial AGC loop gain does indeed match extremely well with the closed-loop response.

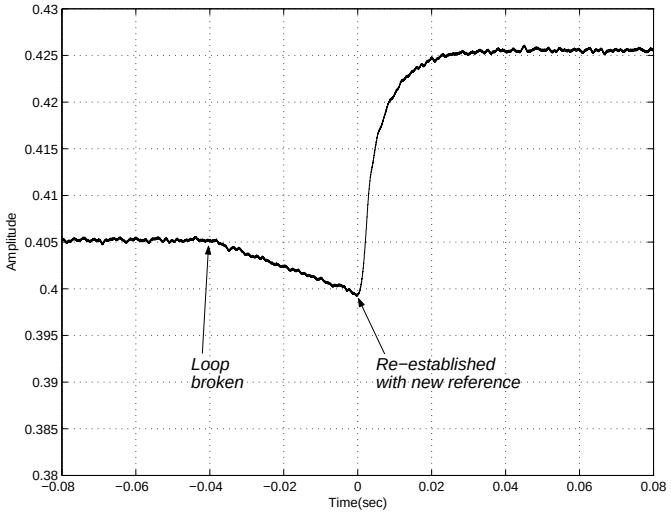

Fig. 24. Amplitude of $\tilde{S}_1$ with positive perturbation about the nominal reference amplitude. Loop is broken at $t = -0.04$s and then reestablished with a new reference at $t = 0$s. This step response test on the amplitude of $\tilde{S}_1$ establishes the drive loop bandwidth to be 26Hz in a neighborhood of the nominal reference amplitude.

rebalance loop are isolated from the $\tilde{D}_1$ and $\tilde{S}_1$ signals and this is largely the case although some cross-coupling is to be expected through the plant. A point to note is that the ω_1-mode is excited from rest with an amplitude envelope initially given by the exponential $e^{100.5t}$. This growth rate was predicted from the sensor and controller parameters in Section V-A. Of course as the amplitude of $\tilde{S}_1$ increases, the AGC gain is reduced until a steady-state amplitude is reached. Due to the extremely light damping in the sensor's structural modes, very little drive amplitude is required to maintain the steady-state condition as is evident from the later half of the $\tilde{D}_1$ trace.

The bandwidth of the closed loop is determined by applying a small perturbation to the reference amplitude while the sensor is operating in steady-state. The result of this experiment is shown in Fig. 24 where the RMS value of $\tilde{S}_1$ (over a short time window) is graphed. This effectively recovers the signal's amplitude and is more convenient to graph than the original signal. The initial reference amplitude is 0.405 V and the perturbed reference 5% larger. As the new parameters load into the ASIC the loop is temporarily opened and this accounts for the droop in the amplitude starting at $t = -0.04$ seconds. At $t = 0$ loop closure is reestablished with the slightly larger reference amplitude. The ensuing response indicates an AGC bandwidth of 26Hz.

The *raison d'être* for employing the AGC to excite the selected mode within the sensor is its ability to track the slow modal frequency drift. The tracking capability of the loop

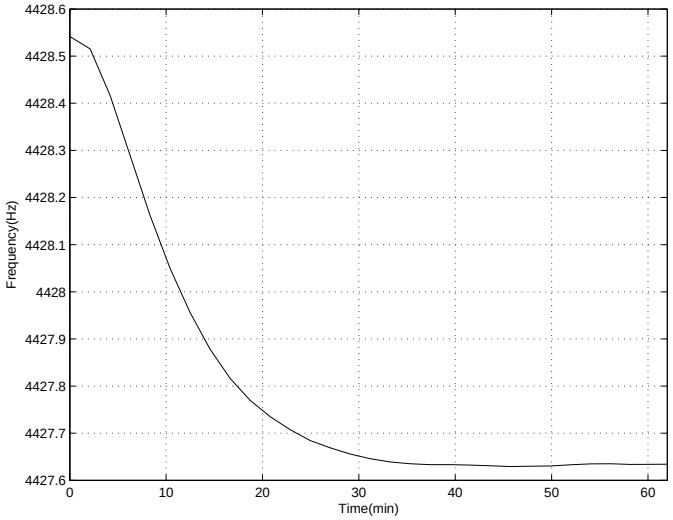

Fig. 25. Steady-state AGC excitation frequency versus time. Two seconds of $\tilde{S}$ and $\tilde{D}$ data were acquired at a 50 kHz sampling frequency every two minutes. The frequency and amplitude of $\tilde{S}_1$ were extracted from a model fit. The frequency is plotted in this figure and the amplitude is plotted in Fig. 26. The drift in modal frequency is believed to be caused by thermal transients in the sensor.

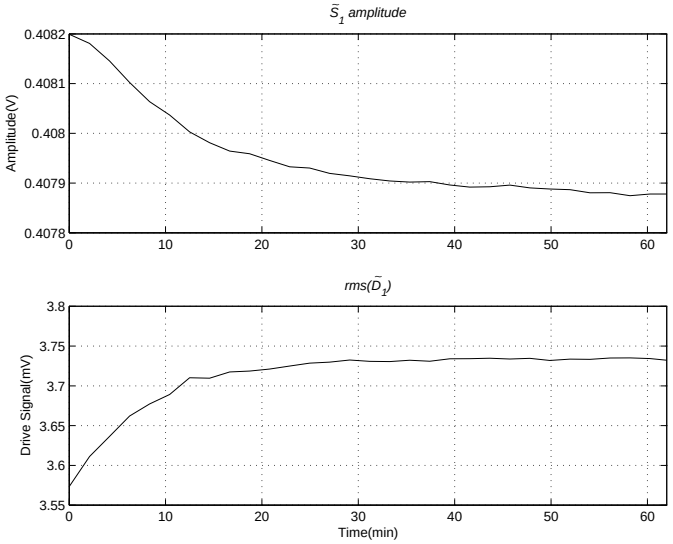

Fig. 26. Top plot: stability of the steady-state $\tilde{S}_1$ amplitude versus time. The amplitude information is extracted from the same data sets that were used to quantify the frequency drift in Fig. 25. Bottom plot: the RMS value of $\tilde{D}_1$ for this experiment exhibits a 4% change thus implying that the AGC is driving the the sensor at its resonant frequency even though this frequency slowly drifts with time.

is shown in Fig. 25. This experiment was conducted over a one hour period during which several seconds of the $\tilde{S}$ and $\tilde{D}$ signals were acquired every two minutes. The frequency of the steady-state $\tilde{S}_1$ response was obtained by fitting a sinusoid model with amplitude, frequency and phase parameters to the sampled data. The fitting procedure was benchmarked by adjusting the frequency and amplitude of a precision function generator in steps of 1 mHz and 0.5 mV, respectively. The algorithm easily tracked these minute changes. Fig. 25 shows that over a one hour interval the frequency of the sustained oscillation changes by almost 1 Hz. The fact that the AGC successfully drives the mode at its natural frequency is supported by the RMS measurement the AGC effort, i.e. RMS of $\tilde{D}_1$, shown in Fig. 26. The change in RMS control effort is approximately 4% and indicates a similar change in plant gain at the oscillation frequency from which it can be inferred that the AGC is driving the mode at its natural frequency. The stability of the $\tilde{S}_1$ amplitude is also shown in Fig. 26. As with the excitation frequency and control effort values, the amplitude seems to have stabilized after one hour. The change in amplitude is quite small though, approximately 0.07%, and this complicates efforts to identify the sources that contribute to the drift. We included the AGC integrator ($K_I \neq 0$) in a separate experiment and still observed roughly the same amplitude change so this indicates that its cause is located primarily in the sensing mechanism which includes the signal conditioning circuit, the analog network, and the FIR filter ripple in the AGC. We will report at a future date our efforts to eliminate the drift.

VII. Conclusion

The ASIC shows great promise as a technology for supplanting the traditional but cumbersome analog control electronics for the JPL-Boeing micromachined rate sensors. We demonstrated a process by which detailed frequency domain models of the sensor can be used to generate the control filters via a set of interpolation constraints on the loop phase at the sensor's resonances. Our current effort is producing a detailed comparison between the ASIC implementation and the analog electronics and these results will be made available at a future date.

Acknowledgment

We are indebted to Dorian Challoner and Boeing Satellite Systems for supporting our research program.

References

[1] Bae, S.Y., Hayworth, K.J., Yee, K.Y., Shcheglov, K., and Wiberg, D.V., "High performance MEMS micro-gyroscope," *SPIE-Int. Soc. Opt. Eng. Proceedings of SPIE - the International Society for Optical Engineering*, Vol. 4755, pp. 316–24, 2002.

[2] Chen, Y.-C., Hui, J., and M'Closkey, R.T., "Closed-loop Identification of a Micro-sensor," *Proc. 42nd IEEE Conf. Decision and Control*, pp. 2632–7, 2003.

[3] Grayver, E. and Daneshrad, B., "Word-serial Architectures for Filtering and Variable Rate Decimation," *VLSI Design*, Vol. 14, No. 4, pp. 363–72, 2002.

[4] Grayver, E., and M'Closkey, R.T., "Automatic Gain Control ASIC for MEMS Gyro Applications," *Proc. 2001 American Control Conference*, Vol. 2, pp. 1219–22, 2001.

[5] Lawrence, A., *Modern Inertial Technology: Navigation, Guidance, and Control*. Springer-Verlag, New York, 1998.

[6] Leland, R.P., "Adaptive mode tuning for vibrational gyroscopes," *IEEE Trans. Control Systems Tech.*, Vol. 11, No. 2, pp. 242–7, 2003.

[7] Lynch, D.D., "Vibratory Gyro Analysis by the Method of Averaging," *The 2nd Saint Petersburg Intl. Conf. on Gyroscope Technology and Navigation*, pp. 26-34, 1995.

[8] M'Closkey, R. T., Vakakis, A., and Gutierrez, R., "Mode Localization Induced by a Nonlinear Control Loop," *Nonlinear Dynamics*, Vol. 25, No. 1, pp. 221–36, 2001.

[9] M'Closkey, R.T., Gibson, J.S., and Hui, J., "System identification of a MEMS gyroscope," *Journal of Dynamic Systems, Measurement, and Control, Transactions of the ASME*, Vol. 123, No. 2, pp. 201–10., 2001.

[10] Morrow, C.T., "Zero Signals in the Sperry Tuning Fork Gyrotron," *The Journal of Acoustic Society of America*, Vol. 27, No. 3, pp. 581–5, 1955.

[11] Newton, G.C. Jr., “Thoery and Pratice in Vibratory Rate Sensor,” *Control Engineering*, pp. 95–9, 1963.

[12] Shcheglov, K., Evans, C., Gutierrez, R., and Tang, T.K., “Temperature depedent characteristics of the JPL silicon MEMS gyroscope,” *Proc. 2000 IEEE Aerospace Conf.*, Vol. 1, pp. 403–11, 2000.

[13] Painter C.C., Shkel A.M., “Active structural error suppression in MEMS vibratory rate integrating gyroscopes,” *IEEE Sensors Journal*, vol.3, no.5, pp.595–606, Oct. 2003.

[14] Tang, T.K., Gutierrez, R.C., Wilcox, J.Z., Stell, C., Vorperian, V., Dickerson, M., Goldstein, B., Savino, J.L., Li, W.J., Calvet, R., Charkaborty, I., Bartman, R., and Kaiser, W.J., “Silicon Bulk Micromachined Vibratory Gyroscope for Microspacecraft,” *SPIE-Int. Soc. Opt. Eng. Proceedings of SPIE - the International Society for Optical Engineering*, vol. 2810, pp.101–15., 1996.

[15] Tang, T.K., Gutierrez, R.C., Wilcox, J.Z., Stell, C., Vorperian, V., Calvet, R., Li, W.J., Charkaborty, I., Bartman, R., and Kaiser, W.J., “Silicon Bulk Micromachined Vibratory Gyroscope,” *Solid-State Sensor and Actuator Workshop, Hilton Head, SC*, pp. 288–93, 1996.

[16] Tang, T.K., Gutierrez, R.C., Stell, C.B., Vorperian, V., Arakaki, G.A., Rice, J.T., Li, W.J., Chakraborty, I., Shcheglov, K., Wilcox, J.Z., and Kaiser, W.J., “A Packaged Silicon MEMS Vibratory Gyroscope for Microspacecraft,” *Proc. IEEE, The Tenth Annual International Workshop on Micro Electro Mechanical Systems, Nagoya, Japan*, pp. 500–5, 1997.