


Measuring Air Bearing Spindle Performance

December 2001





Seagull Solutions, Inc.
Measuring Air Bearing Performance

Seagull Solutions, Inc.



16100 Caputo Drive
Morgan Hill, CA 95037
Phone: 408-778-1127
Fax: 408-779-2806

www.seagullsolutions.net

Report and Tests By:
Peter Polidoro

Testing Support:
Jim Synder



Abstract

In this paper we measure the performance of air bearing spindles with three quantities: the dynamic stiffness, the asynchronous error motion (AEM), and the rotational velocity modulation. The dynamic stiffness is a measure of how well the bearing resists shifts in the rotor axis due to forces on the rotor. AEM also identifies movement of the rotor axis, but is concerned only with motion not synchronized with the spindle rotation. The rotational velocity modulation quantifies how much the RPM varies around the desired RPM value. In our experiments, we found that the dynamic stiffness of the new Seagull Hi-Capacity Biconic air bearing spindle was over double that of the standard Seagull Biconic air bearing spindle and the dynamic stiffness of the standard Seagull Biconic air bearing spindles was higher than the competitor air bearing spindles we tested. We were not able to measure any significant difference in the AEM nor the long-term velocity modulation between the bearings, but we did find that a disk stack mounted on the rotor greatly degrades bearing performance and that fingers placed in the disk stack significantly reduces that performance degradation. We found that signal modulation from an eccentric encoder disk corrupts velocity measurements over the short-term, but that this signal error can be corrected with Seagull's new low modulation encoders.



Table of Contents

Measuring Air Bearing Spindle Performance.....	1
Abstract.....	3
Introduction.....	5
Background.....	6
Dynamic Stiffness.....	7
Asynchronous Error Motion (AEM).....	11
Velocity Modulation.....	13
Test Equipment.....	16
Air Bearing Spindles.....	16
Controller.....	16
Measurement Hardware.....	16
Software.....	16
Accessories.....	16
Test Methodology.....	18
Dynamic Stiffness Measurements.....	18
AEM Measurements.....	24
Velocity Modulation Measurements.....	25
Test Results.....	27
Dynamic Stiffness.....	27
AEM.....	30
Velocity Modulation.....	31
Conclusion.....	35



Introduction

The goal of this paper is to experimentally measure quantitative values of air bearing spindle performance that are as realistic and, hopefully, as meaningful to industry as possible. Often the performance specifications given for air bearing spindles are theoretical and have never been experimentally verified. In addition to never having been verified, many times they are calculated assuming unrealistic or improbable conditions. For example the stiffness is often listed at axial locations inside the bearing, while bearing loads are always placed farther out along the axis where the stiffness is lower. There is not necessarily anything wrong with listing the specifications in this way, it just makes it more difficult for the end user of an air bearing system to know how the air bearing will actually perform under a given set of conditions. With this in mind, we created our performance measurement tests to make the results as useful as possible, even though the results often make the bearings appear to perform more poorly than historical specifications would indicate.

Seagull Solutions, Inc. designs and manufactures air bearing spindles, so of course, we want our products to measure up well against the competition. We made every attempt, however, to keep the tests as objective and as representative as possible to make an equitable comparison between all of the spindles. We put a lot of effort into developing test setups that yield as consistent results as possible. There is no doubt, though, that these tests could be further refined and expanded. This paper is by no means a definitive work on how to measure air bearing spindle performance, but we hope that it provides some insights and perhaps some new perspectives on the subject, for in order to improve the performance of our air bearing spindles, we must first be able to measure it.

We fully encourage others to develop similar tests and experimental setups and verify or question the results we present here. Experienced experimentalists may be able to make many improvements on our testing procedures and other facilities may have more suitable testing equipment or environments. We would be happy to work in collaboration with anyone interested in what they read here to maximize the accuracy and usefulness of these air bearing spindle performance measurement techniques.



Background

For most applications, an ideal air bearing spindle would be one that has a perfectly constant rotational velocity vector. A rotational velocity vector attached to the rotor defines both the orientation of the rotor axis and the magnitude of the angular velocity. In order for the rotational velocity vector to remain constant, the rotor axis must remain stationary with respect to some external reference frame (e.g. no wobbling, tilting, or shifting) and the angular velocity magnitude cannot fluctuate. No real air bearing has a perfectly constant rotational velocity vector, however, and the performance of a bearing can be characterized by how much it deviates from this idealization. A good bearing has a more constant rotational velocity vector under a given set of conditions than a bad one.

In this paper we measure three quantities to characterize the deviation of an air bearing from the idealization. The first quantity is dynamic stiffness, which measures how well the bearing resists radial shifts in the rotor axis due to the centrifugal force created by an imbalance in the rotating spindle. The second quantity is asynchronous error motion (AEM), which measures radial movement of the rotor axis that is not synchronized with the spindle rotation. The third quantity is the rotational velocity modulation or how much the RPM varies around the desired RPM value. These are the primary performance characteristics and in this paper we measure how they change with bearing RPM. In addition, we also keep track of some secondary variables which do not measure the bearing performance explicitly, but do give an indication of overall system performance; the system being composed of the air bearing, air supply, motor, encoder, controller, and amplifier. These secondary variables are the air flow into the bearing and the current through the motor windings and they help indicate problems and inconsistencies in the setup.

One popular air bearing application is hard drive head and gimbal assembly (HGA) testing. During HGA testing, a stack of disks is mounted onto an air bearing spindle and then circular tracks of data bits are written onto and read from the hard disks. Ideally the tracks should be concentric and centered and as perfectly circular as possible so that the track density does not have to be decreased to prevent the tracks from running into each other. Also, the data bits should be spaced as evenly as possible to maximize their density. The disk stack is a very representative real life load condition for an air bearing spindle. Bearings that perform well when spinning freely on their own show very significant performance drops when subjected to the forces of such a load. Even a slight imbalance in the disk stack can create large centrifugal forces on the rotor at high RPM and cause the rotor axis to shift and wobble. An air bearing's dynamic stiffness is a direct indication of how well that bearing can resist those large centrifugal forces. The AEM measurements indicate the potential eccentricity in the data tracks and the velocity modulation indicates how evenly spaced the data bits can be placed.



Dynamic Stiffness

In this paper we use the term “dynamic stiffness” to mean the quantity equal to the centrifugal force acting on the rotor in the radial direction due to an imbalance at a given axial location divided by the radial displacement of the rotor in the same axial location.


$$S_D|_Z \equiv \frac{F_C|_Z}{d|_Z}$$

S_D : DynamicStiffness

F_C : CentrifugalForce

d : RadialAxisDisplacement

Z : GivenAxialLocation



Equation 1 : Dynamic Stiffness Definition

The axial location is a very important parameter when characterizing the dynamic stiffness. Much of the radial displacement at a given axial location is due to rotor axis tilt as opposed to parallel axis shift. For a given amount of tilt, the magnitude of the radial displacement increases as you move away from the tilt fulcrum point. In addition, for a given amount of centrifugal force, the moment on the rotor increases the farther the force acts from the center of the bearing. So for any bearing, the dynamic stiffness measurement will be larger at an axial location closer to the fulcrum point of the bearing than it will be at an axial location that is further away.

The dynamic stiffness is not the same as the stiffness of the air film in the bearing. The dynamic stiffness might be more properly thought of as a mechanical impedance. The dynamic stiffness is a complex number that varies with rotational frequency and is very analogous to electrical impedance, which varies with alternating voltage frequency. At zero rotational frequency, the dynamic stiffness is equal to the air film stiffness just as at zero switching frequency the electrical impedance in a circuit is equal to the electrical resistance. At non-zero rotational frequency, however, the dynamic stiffness is affected by both the inertia of the rotor and the air film damping just as at non-zero switching frequency the electrical impedance is affected by both the capacitance and inductance in the circuit. Dynamic stiffness in an air bearing is a function of rotational velocity, rotor inertia, and the air film stiffness and damping. The air film stiffness and damping are most likely variables that change with rotational velocity and radial rotor displacement as well.



$$S_D = f(\omega, I, k(\omega, d), b(\omega, d))$$

ω : Rotational Velocity

I : Rotor Inertia

k : Air Film Stiffness

b : Air Film Damping

Equation 2 : Dynamic Stiffness Functional Relationship

One way to predict the exact relationship between these variables is to create a dynamics model of the system and solve the equations of motion. As a first approximation we can make an oversimplified model of the system by assuming that the air film stiffness and damping remain constant and by ignoring the rotational inertia and tilting of the rotor. The model becomes a one-dimensional mass-spring-dashpot system with a sinusoidal forcing function.




Figure 1 : Simplified Air Bearing Spindle Dynamics Model

The figure above is a graphical representation of the very simplified air bearing spindle system mechanical model. The total mass “M” represents the rotor mass and the imbalance mass. The imbalance mass “m” is located a distance “r” from the axis of rotation and spins at a rotational velocity equal to “ ω ”. The variable “x” in the figure corresponds to an arbitrary radial direction. The centrifugal force on the rotor caused by the rotating imbalance mass acts as a sinusoidal driving force on the rotor in the x direction at a frequency equal to the rotating speed. The variables “k” and “b” are the stiffness and damping of the air film respectively. If we solve the equations of motion for this system and divide the centrifugal force by the maximum displacement, we come up with the following relationship between the system variables:



$$S_D = \sqrt{(k - M\omega^2)^2 + b^2\omega^2}$$

$$\theta = \tan^{-1}\left(\frac{b\omega}{k - M\omega^2}\right)$$

θ : PhaseAngle

Equation 3: Solutions to the Simplified Model Equations of Motion

Note that at zero rotational frequency ($\omega=0$), S_D simply equals k and the phase angle is zero. The phase angle is the phase shift of the output sinusoid with respect to the input sinusoid. When the phase angle is zero, the maximum positive displacement occurs when the input force is in the same direction as the displacement. When the phase angle is 180 degrees, the maximum positive displacement occurs when the input force is the opposite direction as the displacement. In the figure above, the air film thickness “ t ” will be smallest when the mass “ m ” is at the very bottom of its revolution and the phase angle is zero. When the phase angle is 180 degrees, “ t ” will be smallest when the mass is at the very top of its revolution.

In dynamics analysis, bode diagrams are often used to represent the frequency-response characteristics of dynamic systems. The frequency-response is the steady-state response of a system to a sinusoidal input. Bode diagrams consist of two separate plots, one giving the magnitude of the response versus the frequency and the other the phase angle of the response versus frequency. The dynamic stiffness is inversely proportional to the bode magnitude plot. Another way of determining the relationship between system variables would be to directly measure the frequency-response of the system, create a set of bode plots, then back out the transfer functions from the shape of the plots.




Figure 2: Example Bode Plot of the Simplified Model

The figure above is a typical bode diagram of the simplified air bearing mechanical model. The magnitude plot is shown above the phase plot. Note that the magnitude is plotted on a log-log scale. The magnitude plot is proportional to the reciprocal of the dynamic stiffness. So the simplified model predicts that for low RPM the air bearing dynamic stiffness will remain constant and the phase shift will remain zero. It then predicts that the dynamic stiffness will undergo a short drop as the phase angle begins to change, then increase indefinitely as the RPM rises. If we knew the exact proportions of the plot, we could estimate values for “M”, “b”, and “k”.

The simplified one-dimensional mass-spring-dashpot model is most likely too simplified to predict any real trends in an air bearing. Another thing to keep in mind is that the simplified model considers movement of the rotor with respect to the stator. Often the real concern is how the rotor moves with respect to some other reference point, such as a capacitance measurement probe or write head that is attached to some support structure. Since no support structure is perfectly rigid, a better model would need to take into account the transfer function that describes the motion relationship between the stator and the real reference point. The lack of perfect rigidity in the support structure makes analysis and testing more difficult, but does allow for the interesting possibility of custom designing the support structure to tune the transfer function between the rotor and the external reference point to better meet the designers needs. Instead of simply trying to make the structure as rigid as possible, the designer might be able to tune it to minimize motion between rotor and reference point or to shift natural frequencies in the system.



Knowing the details of the relationship between the system variables, whether the relationship was found theoretically, experimentally, or by a combination of the two methods, would be very useful for predicting how the bearing would react under a set of operating conditions. A detailed bode plot or set of transfer functions would indicate resonant modes in the bearing and stability issues so the designer of an air bearing system could work around them. In this paper we make the first step in the frequency-response characterization of an air bearing by measuring the dynamic stiffness across a range of rotational speeds.

Asynchronous Error Motion (AEM)

Asynchronous error motion is movement of the rotational axis that does not occur at frequencies that are multiples of the rotational frequency. While the source of the error motions may be well defined and repetitive, they appear random relative to time, or they repeat at a frequency that is not a multiple of the rotational frequency.¹ AEM can be caused by any number of disturbing forces acting on the rotor. Similar to the dynamic stiffness, much of the AEM in the rotational axis is tilt, so the farther the measurements are made from the tilt axis, the larger the motion magnitude. NRR or non-repeatable runout refers to the same type of motion. AEM is often measured both axially and radially and the radial measurements can either be from a fixed direction or one that rotates with the spindle. For the purposes of this paper we are most concerned with fixed radial AEM and when we use the term AEM this is what we imply.

Most air bearings easily measure less than one micro-inch AEM under relatively controlled conditions, with no load, and when the measurement is made close to the end of the bearing. The motion magnitude increases significantly, though, when a load is placed on the bearing and the measurements are taken at an axial location where activity normally occurs. Since the numbers are on such a small scale, even minor-seeming environmental disturbances can increase the motion dramatically. The vibrations from an air compressor running on the roof of a building can double the AEM of an air bearing spinning within that building. The first step to reducing AEM is to reduce the disturbing forces on and within the system. Vibration isolators on massive system foundations are often used to try to minimize the effects of outside disturbances.

The air bearing load itself can be a source of disturbing forces acting on the rotor. For example, a stack of disks in the HGA testing application imparts random forces in many directions on the rotor as it spins around. The disks in the stack are separated by layers of air. When the disks spin around at high speed, packets of air are ejected creating low pressure voids in the gaps between the disks. These voids are continually replenished by the surrounding air to fuel further packet loss. The effect is similar to tiny rocket blasts going off in the disk stack pushing on the rotor. These little pushes can create enormous AEM in the spindle.

¹ Lion Precision Advanced Spindle Error Analyzer Version 7 Instruction Manual page 51



Figure 3 : Disk Stack

The picture above shows a typical disk stack that might be used in an HGA testing application. Seagull has developed a way to minimize the random air packet loss (“puffing” as we refer to it) from the disk stack as it is spinning around by inserting “fingers” in between the disks that comb off the layers of air.²




Figure 4 : Disk Stack with Fingers

The picture above shows the fingers inserted into the disk stack to comb off the air layers before they have a chance to exhibit the puffing phenomenon. In previous tests we have seen very significant AEM reduction in a disk stack when using the fingers and in this paper we subject them to more formal tests and compare the finger effects on a variety of bearings.

² Patent Number 6,097,568



Velocity Modulation

The third performance quantity that we measure in this paper is the velocity modulation of the bearing as it is spinning. The velocity of the bearing is controlled by a negative feedback loop. The controller in the system senses the velocity of the rotor from the encoder signal, compares it to the desired velocity, calculates a torque response, and sends that information to the amplifier as a current command. The amplifier sends current through the motor windings based on the current command in its own little negative feedback loop and the motor applies torque to the rotor.

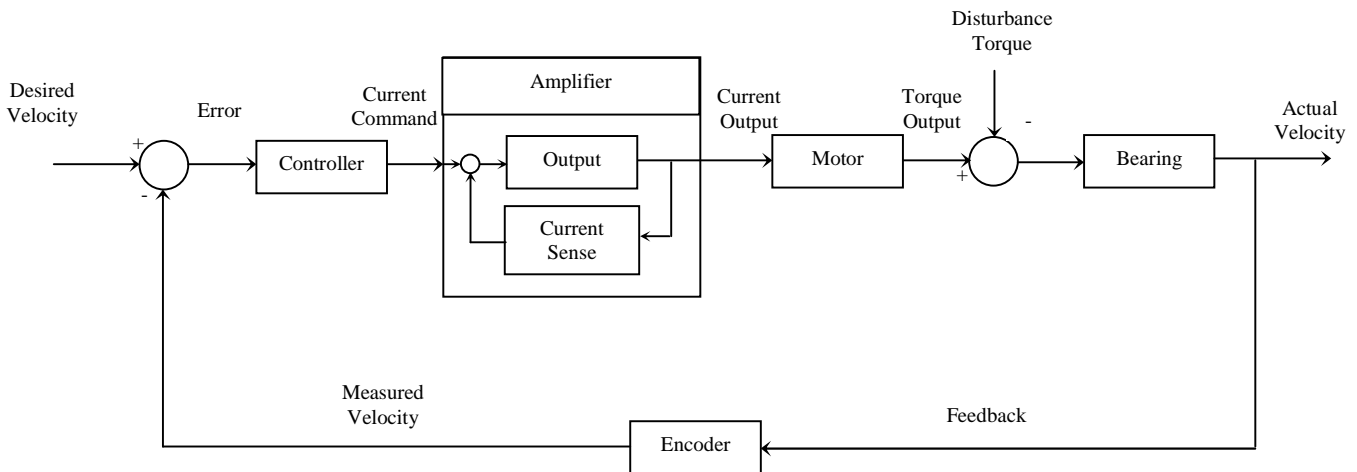


Figure 5 : Air Bearing Spindle System Feedback Loop

The figure above is a block diagram of the system. In order to improve the actual velocity of the bearing, you need to both optimize the design of the feedback loop (and the individual system components) and reduce the disturbance torque acting on the bearing.

Disturbance torque on the bearing can come directly from the bearing load. The same puffing phenomenon in a disk stack that affects the AEM also acts as a source of input forces on the bearing that fight the controller's goal of maintaining constant velocity. The packets of air that continually move in and out of the disk stack and air layer system vary the momentum of the load, in effect acting as disturbance torques on the bearing rotor. The disturbance torques combine with the motor torque output to accelerate and decelerate the rotor. Unless the controller can sense the changes in velocity quickly enough to help reduce them, the bearing will speed up and slow down, speed up and slow down, and the velocity will fluctuate randomly around the desired value. The fingers placed in between the disks to stop the puffing phenomenon helps to reduce the velocity modulation of the bearing as well.

There are several ways in which the velocity modulation might be defined and measured. It is often expressed as a percentage of variation about the desired value. For example you might see a performance specification of 0.001% velocity error. One issue



with expressing the error as a percentage is that for a given amount of fluctuation the percent error changes with the magnitude of the desired value. A percent velocity error at 10,000 RPM for example may not be straightforwardly comparable to a percent velocity error found at 20,000 RPM. If the variation remained constant at ± 1 RPM say, then the percent error measurement will be twice as large at 10,000 RPM than at 20,000 RPM.

The time scale of the test is also an important variable in velocity modulation measurements. The velocity variation measured with a high sample rate over a very short period of time might be very different than the variation measured with a low sample rate over a long time period. Depending on the update frequency of the controller and the input frequency of the disturbance forces, the two measurements might be quite independent of each other. The spindle could be speeding up and slowing down at very high frequency, but the average velocity could seem very constant if it was examined infrequently over a long period of time. Conversely the spindle could be going through long cycles of slowly picking up speed, then slowly slowing down, which would not be noticed if the velocity was examined only for a short burst.

In an attempt to reconcile these issues with the velocity modulation measurements, we examine the spindle velocity on both a short-term time scale and a long-term time scale. On the short-term scale we measure intra-revolution velocity error using the individual encoder pulses over a relatively short sample time. On the long-term scale we measure inter-revolution velocity error using only the encoder index pulse over a relatively long sample time. In addition, instead of taking data over a given amount of time, we take it over a given number of spindle rotations, so that the data will be more comparable across a range of rotational velocities.

Another issue complicating the velocity error is frequency modulation of the encoder feedback signal. Encoders output a digital pulse stream at a frequency proportional to the rotational speed. A popular type of encoder is the incremental optical encoder, which consists of a rotating disk, a light source, and a photodetector (light sensor). The disk, which is mounted on the rotating shaft, has coded patterns of opaque and transparent sectors. As the disk rotates, these patterns interrupt the light emitted onto the photodetector, generating a digital or pulse signal output.




Figure 6 : Optical Encoder



Figure 6 shows the basic principle of an optical encoder. If the shaft rotates at a constant angular velocity, the markings on the disk pass by the light stream at a constant rate. The encoder outputs the digital pulse signal at the rate that the markings pass by the light stream. The controller then measures the frequency of the signal and calculates the angular velocity, knowing how many marks are on the disk. A problem arises, though, when the encoder disk is not perfectly centered on the shaft. When the disk is not centered, some of the markings are farther from the center of the shaft than they are supposed to be and some are closer. For a given shaft rotational velocity, the markings that are too far away will pass by the light stream at a higher rate than the markings that are too close. The result is that the frequency of the encoder output signal is no longer constant, but modulates around the mean frequency. The off-center encoder disk causes the output signal to vary sinusoidally.

The sinusoidal frequency modulation in the encoder output creates several problems. The encoder introduces error in the feedback signal. A very tight control loop will cause the bearing to track the incorrect feedback signal, introducing a physical velocity modulation that matches the error signal modulation. Even if a control loop is loose enough to be effectively blind to the feedback error, the error still needs to be taken into account when using the encoder output for velocity modulation testing. It can be tough, though, to separate the true velocity error from the encoder signal error.

As a solution to the encoder frequency modulation problem, Seagull has developed a special low modulation encoder with an adjustable encoder disk. With a centered disk, the encoder no longer introduces the sinusoidal error into the feedback signal and the frequency modulation is reduced. This allows for a much more accurate velocity measurement that can be used by the controller for tighter tracking or by a tester for measuring more accurate performance parameters.

Error motion in the spindle and masking errors on the encoder disk also introduce errors in the encoder output. Since the encoder disk is connected rigidly to the rotor, it feels the effects of both synchronous and asynchronous error motion of the rotor. This motion corrupts the encoder output signal just as the eccentricity of the encoder disk does. Again, since most of the error motion is rotor axis tilt, it is very important that the encoder disk is located axially as close as possible to the center of the bearing to minimize this corruption. Even if the encoder disk is perfectly centered and undergoes no unwanted error motion, though, slight variations in the transparent and opaque sectors of the encoder mask can still introduce error in the encoder output signal. It is important for the photodetector to average the signal over several of the mask sectors to minimize the effects of their variations or potential contamination. Some encoders average over more sectors than others. Seagull's new low modulation encoders average over 9 sectors while the old Seagull encoders only average over 5.



Test Equipment

Air Bearing Spindles

Seagull:

2 Hi-Capacity Biconics (biconic geometry)

2 Standard Biconics (biconic geometry)

Competitors:

2 Brand A (bispheric geometry)

2 Brand D (planar cylindrical geometry)

Controller

Seagull P2 Controller for all bearings except for Brand A.

(Brand A bearings use a 512 count encoder not compatible with the P2 controller)

Measurement Hardware

Lion Precision Spindle Error Analyzer:

Capacitance Probe C4-D

Capacitance Probe Driver DMT-22

Data Acquisition:

National Instruments Data Acquisition board PCI-MIO-16E-1

Modulation Domain Analyzer:

Hewlett Packard 53310A

Software

National Instruments Labview 6.0.2

Mathworks Matlab 6.1.0

Lion Spindle Error Analyzer 7.1.0

Accessories

Stiffness Adapters:



Figure 7 : Stiffness Adapter



Disk Stack Adapters:




Figure 8 : Disk Stack Adapter

Disks and Spacers:

- 9 Aluminum Hard Drive Disks 95mm Diameter x 1 mm thick
- 8 Disk Spacers 3mm thick

Fingers:




Figure 9 : Fingers



Test Methodology

Dynamic Stiffness Measurements

The concept behind the dynamic stiffness measurements is relatively straightforward. In order to determine how much force it takes to displace the rotor axis, we simply introduce a known centrifugal force on the rotor and measure how far the rotor axis shifts off-center.

We introduce the centrifugal force on the rotor by imbalancing it.

$$F_C = mr\omega^2$$

Equation 4 : Centrifugal Force Due to Rotating Imbalance

The centrifugal force on the rotor is equal to the imbalance on the rotor times the rotational velocity “ ω ” squared. The imbalance is equal to the mass of the imbalance “ m ” times the distance “ r ” the imbalance is located from the axis of rotation. The dimensions of the imbalance are mass-length and the dimensions of “ ω ” are angle/time. The centrifugal force on the rotor is linearly proportional to the imbalance magnitude. If we double the imbalance in the rotor by either adding more mass or by moving the mass farther from the axis of rotation, the centrifugal force also doubles. The centrifugal force, however, grows geometrically with the rotational velocity. If we double the angular velocity of the rotor, the centrifugal force quadruples. So even small imbalances can create quite large forces on the rotor at high RPM.

We use the stiffness adapter to apply the imbalance to the rotor for the stiffness tests. The adapter bolts onto the end of the rotor and has a setscrew hole drilled through it radially. As mentioned in the background, the axial location is a very important variable when measuring and specifying the dynamic stiffness. We chose to use an axial location 2 inches from the top of the bolt heads holding the adapter onto the rotor, since this is a typical load location for many air bearing spindle applications. The setscrew can be moved radially to change the amount of imbalance in the adapter. The setscrews in the adapters (we need a different adapter for every incompatible bolt hole pattern in the bearings) have mass of around 2.5 grams. The setscrews can be moved as far as 0.18 inches from their balanced positions, so we can introduce as much as 0.45 gram-inches of imbalance (mixing English and metric units). It would not be difficult to have this much imbalance in a haphazardly placed hard drive disk stack in an HGA tester. At 30,000 RPM that relatively small amount of imbalance applies over 25 lbs of force on the rotor, so it is very important that the bearing is stiff enough to resist those high forces if you want to run at high speed. One approximation we make is that the imbalance remains constant during the tests, even though it changes slightly with the rotor axis displacement. We ignore this change, though, since it is several orders of magnitude less than the initial imbalance distance from the axis of rotation.



So after we have introduced a known amount of imbalance into the rotor, the next step is to measure how far that imbalance displaces the rotor axis as the rotor spins at speed. We find the location of the rotor axis by examining the runout of the adapter with a non-contact capacitance probe. The probe measures the distance between the probe tip and the side of the adapter.



Figure 10: Dynamic Stiffness Test Setup

The picture above shows a typical setup for the dynamic stiffness tests. Note that because the setscrew hole punctures both sides of the adapter, the probe cannot be placed at the exact same axial location as the imbalance because the holes would interfere with the runout measurements. So we chose to place the probe at a slightly greater axial distance from the center of the bearing than the imbalance. This makes all of the stiffness measurements slightly lower than they would be if the probe measured the same location as the imbalance. We made sure to keep this distance the same on all of the bearings, though, so the results could be compared properly. The Lion Precision capacitance probes can be used in two ranges: HI and LO. The HI range measures at a higher resolution over a shorter distance, while the LO range measures a lower resolution over a larger distance. We use the Lion Precision capacitance probes in LO range for the dynamic stiffness tests to give ourselves plenty of space between the probe tip and the adapter. We do not want the adapter to hit the probe when the runout is large.

To measure the runout of the adapter and compute the location of the rotor axis, we wrote several programs in Labview and Matlab to do the data acquisition and manipulation. The Labview data acquisition program first finds the rotational velocity from the encoder, then determines the sample rate necessary to measure the runout in 40 locations around the circumference of the adapter. We chose 40 samples per revolution because that was the maximum number that could be taken by the data acquisition card while the spindle was spinning at 30,000 RPM. The program then waits until it sees the index pulse from the encoder, then takes the 40 runout measurements and stores them into an array. It does this for 20 revolutions, so afterwards it has 20 runout measurements at each of the 40 circumferential locations around the adapter. The Labview program



then sends this information to the Matlab program. The Matlab program takes the median value of the 20 samples at each of the 40 locations and calls that the runout for that spot. (We make multiple measurements at each point to minimize the corrupting effects that the AEM and the velocity modulation have on the tests.) So now the program has 40 runout measurements equally spaced around the circumference of the adapter. It then scales the data by subtracting the smallest runout measurement from all of the runout measurements. The smallest runout measurement is the smallest gap size from the probe to the adapter, which is not necessarily constant between bearings. So after all that we are left with 40 scaled runout measurements equally spaced around the circumference of the adapter, with the first measurement corresponding to the index pulse.

A perfectly balanced, perfectly cylindrical, perfectly centered adapter would leave the programs above with 40 scaled runout measurements all equal to zero. If we plotted these data in polar coordinates, we would simply see a single point at the origin.

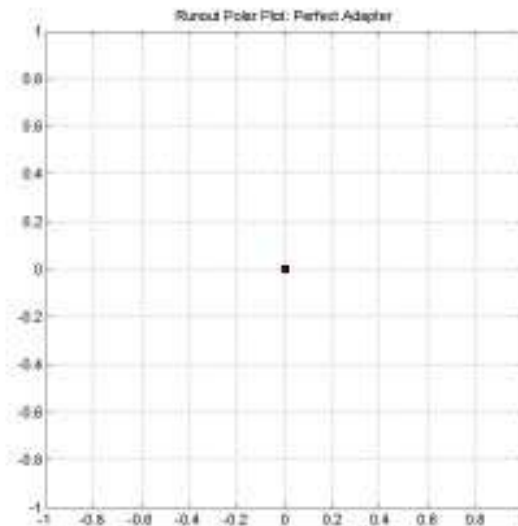


Figure 11 : Runout Polar Plot for a Perfect Adapter

The runout polar plot for the perfect adapter is not very exciting. All of the points are on the origin. If the cross-section of the adapter was square shaped, however, instead of circular, but still perfectly balanced and centered, the runout polar plot would look a little more interesting.




Figure 12 : Runout Polar Plot for a Square Adapter

The four points at the origin of the runout polar plot with the square adapter correspond to the four corners of the square. Not all of the points are on the origin this time, but the plot is still centered about the origin. Now if we take the square adapter and simply move it off center from the axis of rotation of the bearing, the runout polar plot changes again.

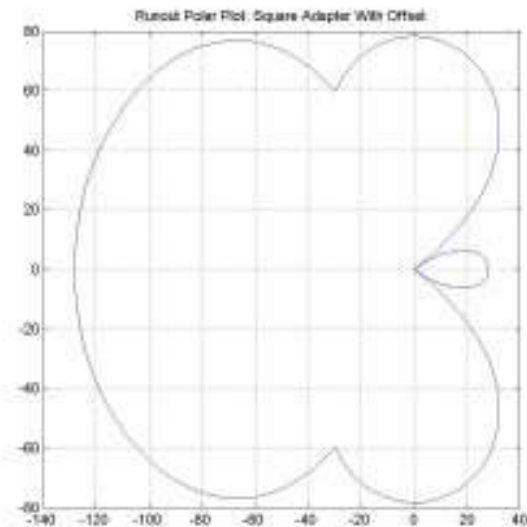


Figure 13 : Runout Polar Plot of a Square Adapter with Offset

The plot above shows the runout polar plot of a square adapter with an offset of magnitude 50. Now the plot is no longer centered about the origin. To find the approximate center of the data, we can do a least squares circle fit to the data and find the location of the center of the best-fit circle. If we do least squares circle fits on both of the



runout polar plots and take the difference between the two centers of the circles, we can, in effect, measure the offset of the adapter.

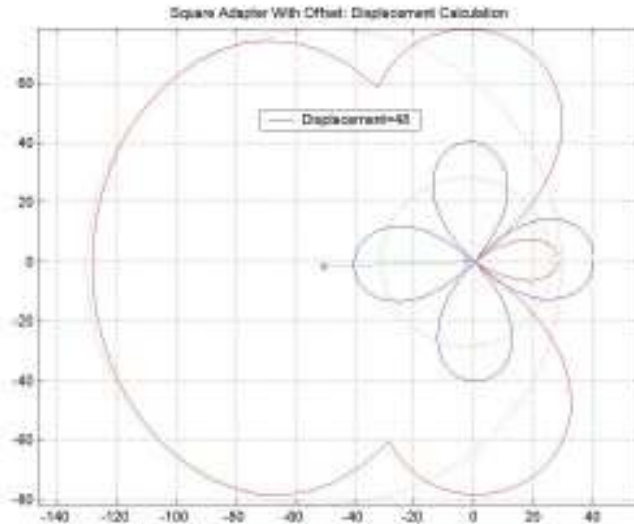


Figure 14 : Offset Measurement for the Square Adapter

The barely visible circular dotted lines in the plot above are the best-fit circles to the two data sets. The centers of both of the best-fit circles are shown on the plot along with a line that connects them. The length of that connecting line is approximately the magnitude of the adapter offset. In this simulation, the offset of 50 was measured as 48.

We use the same technique for measuring the displacement of the rotor axis during the dynamic stiffness tests. In the tests we use a cylindrical adapter, centered with a dial indicator, and take runout measurements on it both when it is balanced and when it is unbalanced. The difference between the centers of both data sets is the displacement of the rotor axis of the bearing due to the centrifugal force caused by the imbalance.

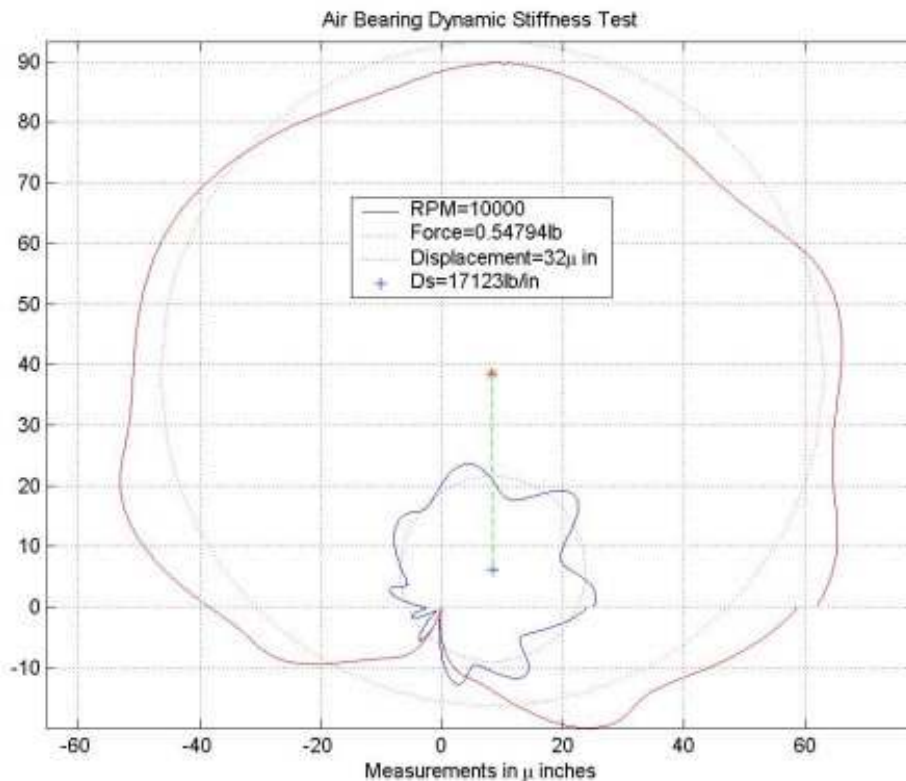


Figure 15 : Typical Dynamic Stiffness Measurement Polar Plot

The figure above shows a typical dynamic stiffness measurement for a spindle running at 10,000 RPM. The small data set and best-fit circle are the runout measurements when the adapter was balanced. As you can see on the plot, the adapter is not perfectly round, nor perfectly centered, otherwise all of the points would be on the origin. If the adapter was perfectly round and centered, then we would not need to take the initial balanced runout measurements every time, we could simply assume the initial adapter center was at the origin. Taking the initial measurements improves the accuracy of the tests in light of the imperfect conditions. The large data set and best-fit circle correspond to the runout on the adapter when it is unbalanced. The center of the adapter did not change with respect to the center of the rotor, rather the rotor axis shifts with respect to the axis of rotation. The program calculates the centrifugal force on the rotor knowing the imbalance magnitude and the rotational velocity, then divides the force by the rotor axis displacement to find the dynamic stiffness.

In addition to finding the distance between the centers of the data sets, we might also find the direction. The direction of the rotor axis shift corresponds to the phase angle. If we knew the location of the imbalance with respect to the index pulse of the encoder, plus we knew of all the time delays in the data acquisition system, we could use the direction information to map the phase angle across a range of bearing speeds to get the second half of the bode plot of the frequency response for the system. We noticed, by matching polar runout plots with physical markings on the adapters, that there is indeed a



significant time delay from the index pulse signal to the time it takes the data acquisition system to take the first measurement. This means that the zero degree angle does not necessarily match up with the index pulse. For example if there was a 1 ms time delay in the data acquisition system, then at 20,000 RPM (3 ms/rotation), the bearing will have spun 1/3 of a rotation between the time the encoder output the index pulse and the time the software was able to take the first runout measurement with the probe. The software continues to take the forty measurements, so it measures over an entire revolution of the spindle, but the first measurement does not exactly coincide with the index pulse. With a constant time delay, the angle between the index pulse and the first measurement will change with RPM. We found this to indeed be the case. This angle change would make a phase angle calculation more difficult, but does not necessarily corrupt the dynamic stiffness measurements. It is another reason, though, why we need to take initial runout measurements of the adapter while it is balanced at every RPM increment, before we take the imbalanced runout measurements at the same RPM increments.

Another assumption we make with these dynamic stiffness measurements is that the dynamic stiffness does not change with radial direction. We arbitrarily choose a radial location when we place the probe. It is possible for the dynamic stiffness to be slightly direction dependent, though. For example a slightly out of round stator might affect the air film properties creating changes in the dynamic stiffness. Along these same lines, we also assume that the air properties remain constant between the capacitance probe and the adapter. An out of round adapter may create air flow patterns that change the air capacitance and slightly corrupt the capacitance probe readings, but we assume these effects are negligible.

AEM Measurements

We use the Lion Precision Spindle Error Analyzer software to make the asynchronous error motion measurements. We use the same Lion capacitance probes and drivers as for the dynamic stiffness tests, but this time we use the probes on the HI range setting, since the magnitude of the AEM is much smaller and we need more accurate readings. The Lion software automatically subtracts out the synchronous error motion when making the asynchronous error measurements, which should minimize the effects of any imbalance in the target. The software also triggers its measurements from the encoder pulses, so that should minimize the corrupting effects of the velocity modulation. The measurements are on such a small scale, however, that the condition of the probe target is very critical for accuracy. The target mounted on the bearing spindle needs to be as round and centered and have the best surface finish possible to sense very delicate motions of the rotor.

For these tests we use the disk stack adapter attachment. Again, the axial location of the measurements is a very important variable, so we chose a distance of one inch from the top of the bolt heads holding the adapter onto the rotor. This spot gives us enough room to leave the probe in place and put the stack of disks and spacers on the adapter. Since the probe is closer to the center of the bearing than the disks, the AEM measurements will be best-case readings. We would have preferred to have the probe in



a location that would have given us worst-case readings so we could estimate the motion of the top disk, but we found that difficult to pull off procedurally.




Figure 16 : AEM Test Setup (Probe and Adapter with Disk Stack and Fingers)

The picture above shows a typical disk stack setup for making the AEM measurements. We take the measurements over a range of spindle speeds with disks and fingers in place, with only the disks and no fingers, and then without disks or fingers. We use nine disks in the disk stack since that is the maximum number of disks that the model of the disk stack adapter could hold.

The Lion Spindle Error Analyzer software measures and computes the AEM measurement by taking readings over 25 spindle revolutions. The software continues to take measurements every 25 revolutions or so and displays the latest AEM value. The AEM values vary slightly from measurement to measurement, so to find a representative value we sample 20 AEM measurements in a row and take the mean.

Velocity Modulation Measurements

For the velocity modulation measurements we again use the disk stack adapters just as we did for the AEM measurements. Again we take the measurements over a range of bearing speeds with disks and fingers in place, with only the disks and no fingers, and then without disks or fingers.

Table 1: Velocity Modulation Test Setups

		
Fingers	Disks	Bare



This time, though, instead of using a capacitance probe to look at the rotor runout, we use a modulation domain analyzer to look at the output of the encoder to figure out how the velocity of the spindle changes with time.



Figure 17: Modulation Domain Analyzer HP 53310A

The modulation domain analyzer has two inputs, the frequency signal input and an optional trigger. The modulation analyzer measures how the frequency of the input signal changes with time. It can display the information in two ways: 1) Cartesian coordinates with frequency on the dependent axis versus time on the independent and 2) a histogram of the distribution of measured frequencies. The analyzer computes the mean (μ) and standard deviation (σ) of the distribution of frequencies. The units of these numbers are hertz and we use them to compute the percent velocity error. We approximate the range of the distribution as 6σ , so the total variation about the mean is plus or minus 3σ .

$$\% \text{ Error} = \frac{3\sigma}{\mu} \times 100\%$$

Equation 5: Velocity Modulation Percent Error

We examine the standard encoder output for the short-term intra-revolution velocity modulation and the index pulse from the encoder for the long-term inter-revolution velocity modulation. As mentioned in the background, the frequency variance depends heavily on the time scale that the frequency is measured over. Somewhat arbitrarily we chose to examine the intra-revolution velocity error over 3 revolutions and the inter-revolution velocity error over 1000 revolutions. So at each RPM, we had to set the modulation analyzer time scale so that the range corresponded to the number of revolutions we wanted to measure over. The standard deviation also depends on the number of samples, so we fixed the number of samples per test at 20,000. The intra-revolution measurements are somewhat flawed in that they depend on the encoder count, but all of the encoders used in these tests had 1024 count encoders, so at least they were consistent. We did not include the Brand A spindles in the velocity modulation tests since they use encoders with 512 counts and since we had to use another controller to spin them.



Test Results

The first step of the testing was to design setups that yielded repeatable results. There were many sources of potential corruption and they all needed to be controlled or at least minimized. The air flow and motor current measurements were valuable for diagnosing system problems. We found that some of the most influential variables were the air supply pressure, the electrical ground paths throughout the system, the bearing support structure, and the probe holders. Some of the bearings with very high flow caused the air supply pressure to fluctuate, so we added air tanks to act as air supply capacitors. This allowed the spindles to all run steadily at their rated air supply pressures. Good electrical grounding in the system was crucial for reliable capacitance probe readings. It was especially important that the rotors stay properly grounded. At high bearing speeds the imbalances used in the dynamic stiffness tests created very large forces throughout the entire support structure. Everything in the support structure had to be as rigid and held as securely as possible to minimize unwanted motion. Since we were trying to measure the motion of the rotor with respect to the stator, the probe holders needed to be as rigid and as lightweight as possible to minimize the amount of motion between the probe tip and the stator. Any probe holder has natural frequencies and resonant modes, however, so the transfer function between stator and probe tip might always affect the frequency response tests, depending on the frequency ranges.

Dynamic Stiffness

The first relationship between variables that we wanted to check with the dynamic stiffness tests was how the dynamic stiffness varied with the radial displacement of the rotor axis. We assumed that the dynamic stiffness might both be a function of bearing RPM and the radial displacement of the rotor axis. This could mean potentially having to examine the data three-dimensionally, with dynamic stiffness as the dependent variable and both the RPM and the radial displacement as independent variables. Since it is more convenient to examine the data two-dimensionally, we decided to first check how the dynamic stiffness varied with radial displacement at a constant RPM value, on the chance that the variation would be small and we could assume it to be negligible.

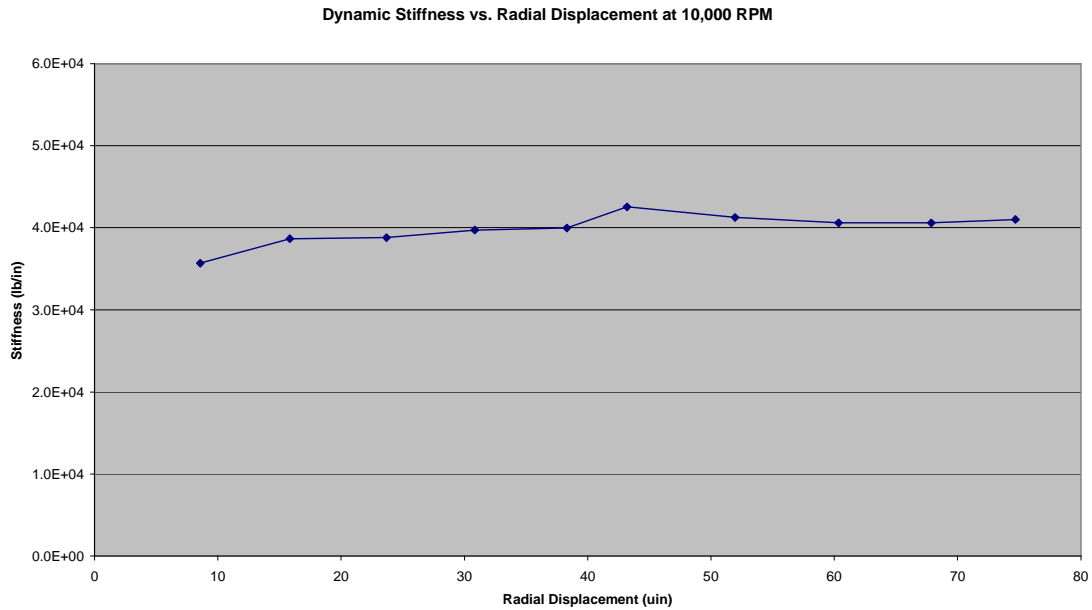


Figure 18 : Dynamic Stiffness vs. Radial Displacement at 10,000 RPM

The plot above shows how the dynamic stiffness varies with radial displacement for one of the bearings at 10,000 RPM. Each data point represents a unique imbalance introduced into the stiffness adapter. The dynamic stiffness seemed to remain very constant with radial displacement, so for the rest of the dynamic stiffness tests we only measured the dynamic stiffness with a single imbalance magnitude (0.3 gram-inches). As an additional check, we also examined how the direction of the displacement changed with the various imbalance amounts at 10,000 RPM. Assuming that the phase angle remains constant at a given rotational velocity (not necessarily true if the air film properties change with rotor radial displacement), the displacement should always be in the same radial direction, regardless of the amount of imbalance.

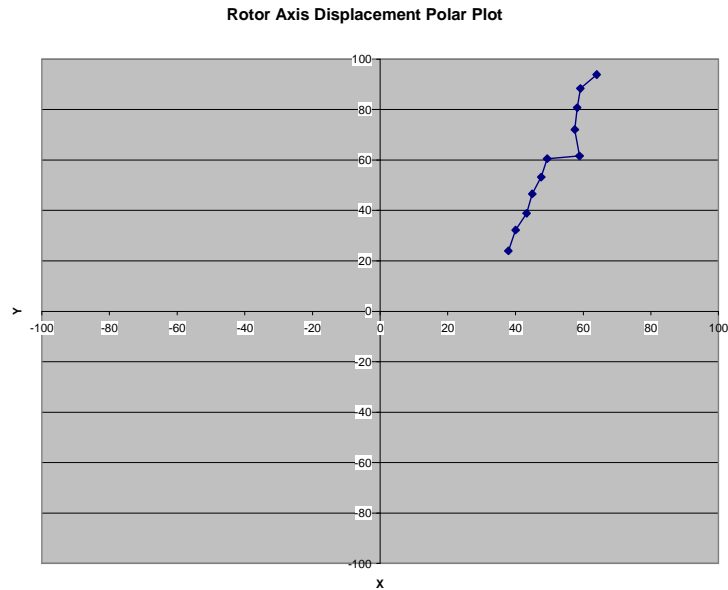


Figure 19 : Rotor Axis Displacement Polar Plot

The plot above is a polar plot of the rotor axis and how it moves with the varying magnitudes of centrifugal force acting on the rotor. Except for one outlying point, the rotor axis seems to move repeatedly in the same radial direction. The point closest to the origin is the location of the adapter axis when the adapter is balanced. The points move away from the origin as the imbalance, and therefore the centrifugal force, increases. Note that the balanced location does not coincide with the origin. The distance between the balanced location and the origin is the amount that the stiffness adapter axis is off center from the rotor axis.

The next step in the stiffness testing was to perform the stiffness tests for each of the eight bearings over a range of speeds.

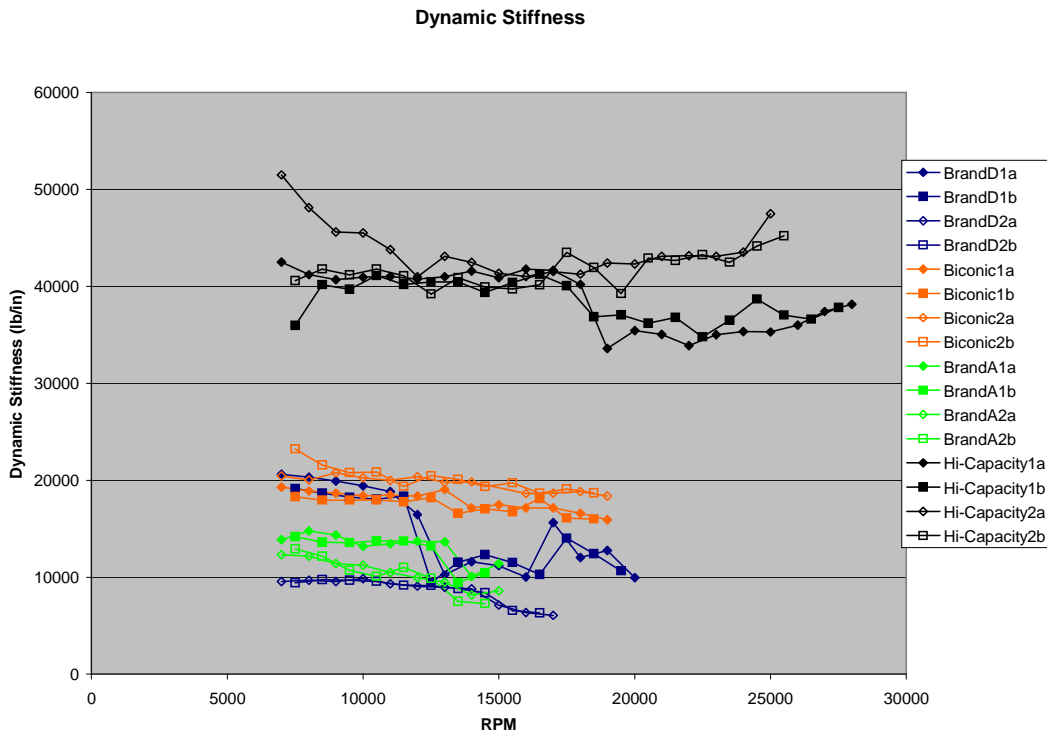


Figure 20 : Dynamic Stiffness vs. RPM for All Air Bearings Tested

The plot above shows the results of the dynamic stiffness tests for the bearings. The dynamic stiffness is plotted on the dependent axis versus the RPM on the independent axis. Each color represents one type of bearing (4 types in total), there were two bearings of each type (8 bearings), and each bearing was tested twice (16 tests in total). Brand A measured between 10,000 and 15,000 lb/in and the Seagull standard Biconic measured between 15,000 and 20,000 lb/in. The two Brand D bearings showed very different results. The one with the higher stiffness had almost twice the air flow of the one with the lower stiffness, so perhaps they were built differently or one was damaged in some way. The new Seagull Hi-Capacity Biconic air bearing spindle was designed for high stiffness and high speed and the results show this to indeed be true. The stiffness of the Hi-Capacity Biconic was over twice that of the standard Biconic and was measured all the way up to 28,000 RPM, limited only by the amplifier supply voltage. We had to limit the maximum speed of the other bearings because we either reached their maximum rated RPM or because the rotor displacement was too large to measure under the load.

AEM

We found that the AEM measurements are very sensitive to external conditions. Any vibrations transmitted through the foundation from the outside world showed up in tests and increased AEM. With our setup, we were not able to measure any significant differences in AEM magnitude between the bearings. They all measured comparably with similar external conditions to the accuracy of our experimental setup. We were,



however, able to measure significant differences in the AEM with changes in the disk stack for all of the bearings.

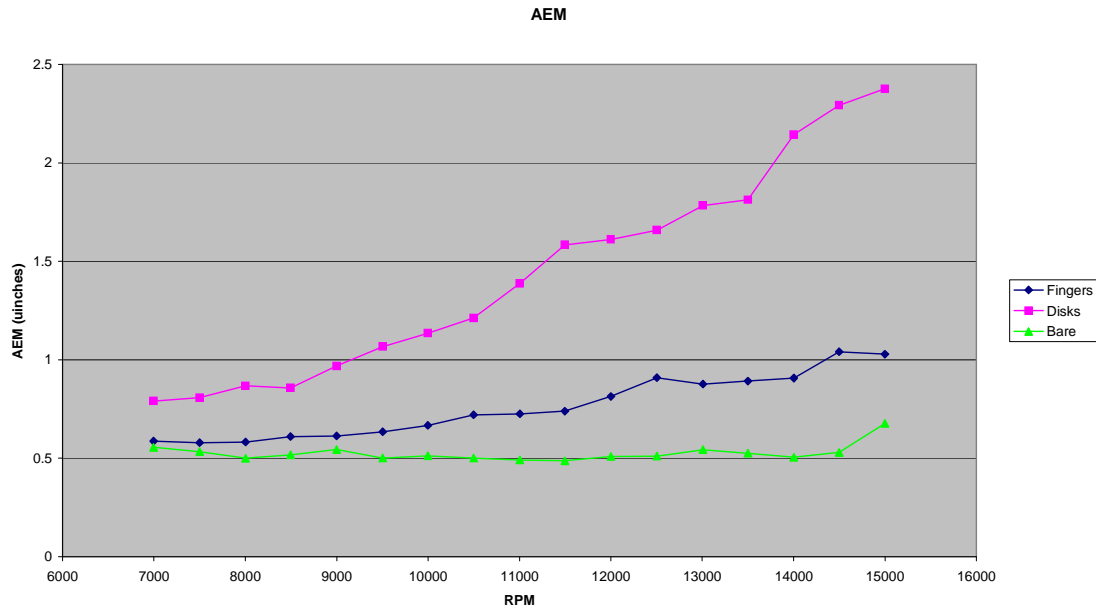


Figure 21 : Typical AEM vs. RPM Test Results

The chart above shows typical AEM results for the bearings. With just the disk stack mounted on the rotor, all of the bearings showed around half a micro-inch AEM. In reality this number is probably somewhat lower due to probe noise and external vibrations. The probe noise was measured around 0.27 micro-inches. The numbers shown in the chart are the raw measurements without any sort of noise correction. The AEM increased dramatically when the disks were placed on the adapter and seemed to increase linearly or perhaps quadratically with RPM. When the fingers were placed in the disk stack, the AEM reduced significantly. This suggests that the puffing phenomenon in the air layers between the disks in the stack is a major source of disturbance motion in the rotors. Another observation we made was that the motor current doubled with the disks on the adapter. The motor current increased by another 20% with the fingers in the stack. This extra current can create thermal issues in the amplifier and the bearing, so Seagull’s Hi-Capacity air bearing spindle uses an ironless-coreless low induction motor, to reduce the motor current consumption.

Velocity Modulation

Intra-Revolution

We found that the frequency modulation of the typical encoder output was so large that we were unable to distinguish the velocity variation from the encoder distortion for the short-term intra-revolution tests. The following table shows our results for these tests, both with the special Low-Modulation (LM) encoder with the centerable encoder disk used on the Hi-Capacity Bionics and the typical encoders used on the others.



Table 2: Encoder Frequency Modulation vs. Time

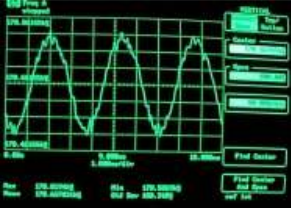
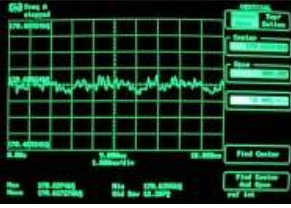
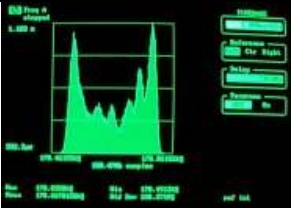
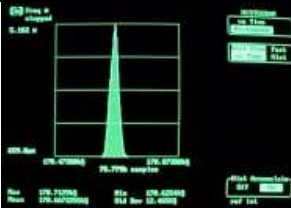
Encoder Frequency Modulation Versus Time RPM: 10,000 Encoder Count: 1024 Sample Time: 3 Revolutions	
 <p>Typical Encoder 170.666kHz +/-300Hz +/- 0.18% error</p>	 <p>Adjusted LM Encoder 170.666kHz +/-36Hz +/- 0.02% error</p>
<p>Ideally the graphs above should be straight horizontal lines. For a given RPM, the encoder pulse output frequency should be a single constant value. The typical encoder output varies like a sine wave, though, because the optical disk is off center. Within the revolution, some of the encoder pulses are too fast and others are too slow. The optical disk on the LM encoder can be centered, however, giving a much more constant frequency output. The frequency modulation of the LM encoder is over 8 times less than that of a typical unadjusted encoder.</p>	

Table 3 : Encoder Frequency Modulation Histograms

Frequency Histogram	
 <p>Unadjusted Encoder</p>	 <p>Adjusted LM Encoder</p>
<p>These two graphs show the distribution of the frequency output from both encoders. The unadjusted encoder has a large frequency range with most of the values jumping back and forth between the two frequency extremes. A double-peaked histogram is typical of a sinusoidal variation. The frequency distribution of the LM encoder, however, is very tight and centered about the correct value, with most of the variation due to random noise.</p>	

Note: All of the images from the modulation domain analyzer are shown on equal scales for comparison.



Inter-Revolution

The long-term inter-revolution velocity modulation tests were not affected by the encoder modulation, since only the index pulse from the encoder was examined. The controller used for these experiments calculates the rotational velocity by averaging the encoder output over several bearing revolutions, so it is effectively blind to the encoder modulation within each revolution. This means that the encoder modulation does not show up as velocity error, but unfortunately this also means that the controller cannot correct short-term velocity error either. The rotor essentially coasts under open loop control over the time it takes the encoder to determine the rotational velocity. Ideally controllers should be able to make short-term corrections in the velocity error, but with such a tight control loop, the encoder modulation will affect velocity error unless a low modulation encoder is used.

As with the AEM experiments, we were unable to measure much variation in the long-term velocity modulation between the spindles. Also, the velocity modulation is very controller dependent, so it is not practical to compare the bearings in this test since the controller we used was optimized for the motors in the Seagull air bearing spindles. Again, though, just as with the AEM experiments, we were able to measure a very significant difference when the disks were placed on the adapter for all of the bearings.

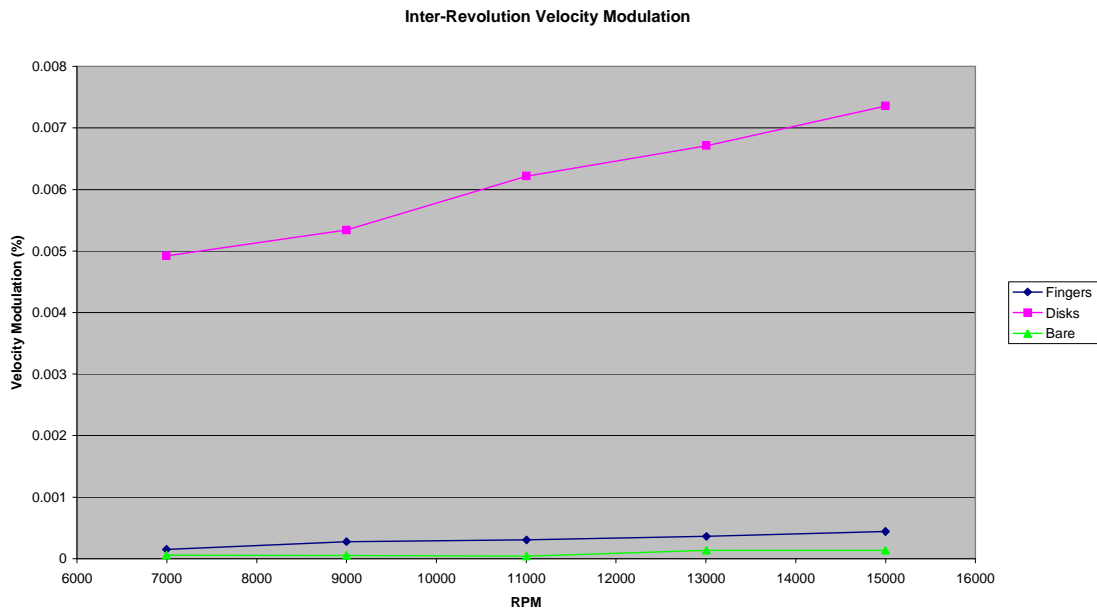


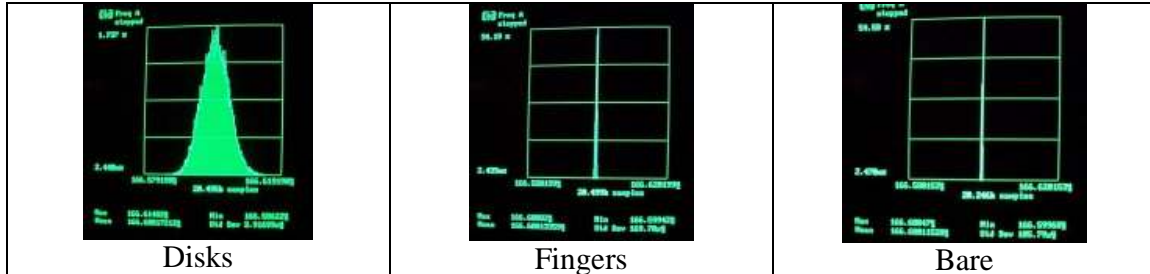
Figure 22 : Typical Inter-Revolution Velocity Modulation vs. RPM Test Results

With the adapter alone, the inter-revolution velocity error was well under 0.001%. With the disks on the adapter, however, the velocity modulation increased to over 0.005% for all of the bearings tested and increased with RPM. With the fingers installed, though, the velocity modulation decreased to almost the same level as just the bare adapter. Again the puffing phenomenon seems to be having a very significant influence



on the spindle performance. The histograms of the velocity distribution are also very telling as to the effect of the disks and the fingers.

Table 4: Velocity Distribution Histograms



The histograms are screen shots from the modulation domain analyzer. All of the histograms are shown on the same scale for proper comparison. The bare adapter shows almost no variation in the velocity and most of the measurements occur at the correct velocity value. When the disks were placed on the adapter, on the other hand, the velocity measurements show a large amount of random variation. The distribution is Gaussian, as typical for random phenomenon, and wide showing that the values occur over a broad range. With the fingers in the disk stack, however, the distribution returns to almost as narrow of a peak as the bare adapter alone.



Conclusion

The three parameters tested in these experiments seem to be very valuable for quantitatively measuring the performance of air bearings and their support systems. We found Seagull's new Hi-Capacity Biconic air bearing to be over twice as stiff as Seagull's standard Biconic air bearing, which itself had higher stiffness than the competitor bearings we tested. We were not able to measure any significant difference in the AEM and the velocity modulation between the spindles, but we did find that a stack of disks mounted on the rotor greatly reduces spindle performance and that fingers placed in between the disks in the stack almost negates the negative effects of the disks. We found that we were not able to measure the short-term velocity modulation of the spindles because the frequency modulation of a typical encoder is too large. We were, however, able to reduce the encoder modulation almost entirely with Seagull's new low modulation encoder, which will be critical for control loop improvements in the future. The dynamic stiffness tests were valuable for measuring bearing stiffness, but they also have potential to provide a great deal of information as to the frequency response characteristics of the air bearing spindle and the surrounding support structure. Detailed frequency response tests would indicate natural frequencies and resonant modes in the system and give clues about the relationship between variables in the system and the approximate magnitude of various system parameters. We here at Seagull intend to continue to improve the performance of our spindles and the support system using these tests, and improvements of these tests, as a guide. We encourage others to try similar experiments and to work to refine them so the entire industry can benefit from improvement through testing and experimentation.