

The Air Bearing Throughput Edge

By Kevin McCarthy, Chief Technology Officer

In addition to the numerous advantages described in the previous section, air bearings have an additional edge over conventional bearing systems that is not generally appreciated. When compared with rolling steel bearing systems, air bearings offer a throughput advantage of up to 10 times, which is a substantial productivity improvement. Photonic alignment systems, with their emphasis on numerous small moves, benefit particularly well from the increased productivity of air bearing ways. This increase in throughput is a direct consequence of the air bearing's freedom from friction, the details of which are covered below.

When large moves are performed by a positioning system, there are benefits to "profiling" the move, with velocity carefully shaped as a function of time so as to minimize higher derivatives and avoid exciting system resonances. A typical motion profile of velocity versus time might resemble that shown in Fig. 1.

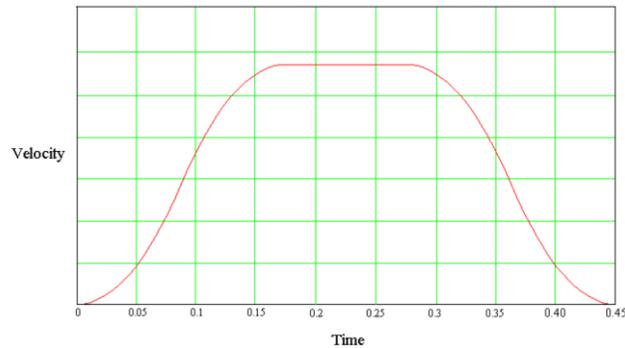


Figure. 1.

SMALL MOVES

As it happens, such careful move profiling is of no utility whatsoever for the numerous small (under 100 micron) moves typical of photonic alignment. In this case, the overall energy is very small, and the servo loop itself is a perfectly adequate trajectory shaper. The position servo loop functions as a low pass filter, allowing us to make the position command a simple step function: at time $t = 0$, we simply command the position loop to be at the destination. Despite the discontinuity of the step command, the actual stage motion follows a smooth curve, as shown in Fig. 2.

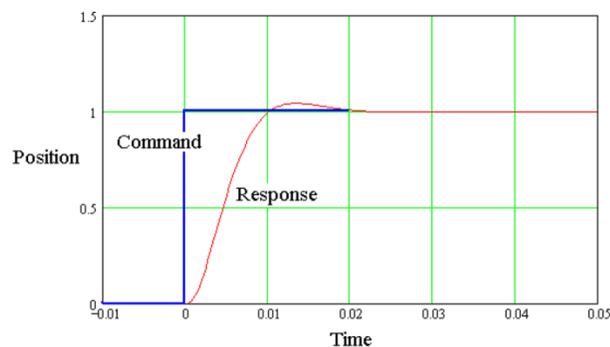


Figure. 2.

In Fig. 2, the commanded and actual position vs. time are shown for a normalized small move. The goal, of course, is to minimize the move and settle time, so as to be able to make as many small moves per second as possible. We may also seek to impose fairly tight values on the settling tolerance. A “move and settle time” is meaningless without a definition of the settling window, which is the acceptable difference between the target position and the actual position. In cutting edge photonic alignment, this settling window may be as small as 10-20 nanometers. Another useful graph of the above move results when the difference between the commanded position and the actual position is plotted. This difference is readily provided by servo controllers, and is called the following error. As shown in Fig. 3, the following error is zero just prior to the move.

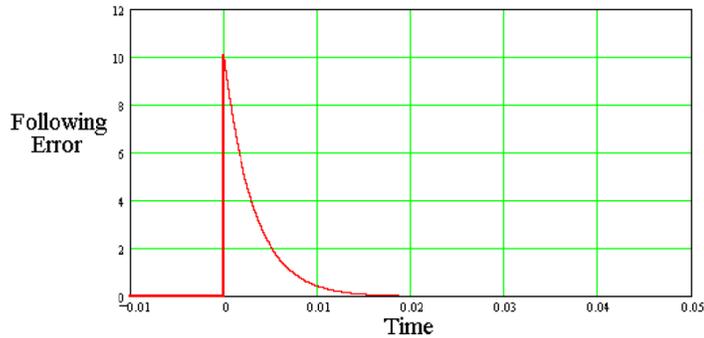


Figure. 3.

Since our command trajectory is a step function, the following error then jumps to equal the move size, as the stage cannot respond instantly to the step command. The key performance issue then relates to the decay of following error over time. As the kids say when we’re on a long drive: “Are we there yet”? Before we tackle the issues that control settling time, it’s worthwhile to briefly consider the topic of servo bandwidth.

SERVO BANDWIDTH

Since the servo bandwidth is critical to an analysis of stage dynamics, it’s worth taking a moment to explore its meaning. Consider the case where we command a servo system with a small amplitude sine wave, and vary the frequency. The resulting amplitude vs. frequency graph (Fig.4) has a constant value from D.C out through a certain frequency, where the amplitude will peak slightly (if properly tuned), and then decline as $1/f^2$.

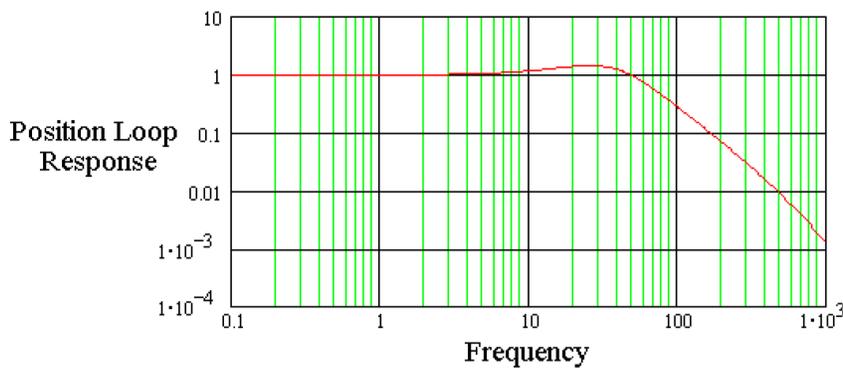


Figure. 4.

The point at which the amplitude has fallen by 3 dB (also the point where there is a 90° phase relationship between the command and the response) is termed the servo bandwidth of the servo system, and is the most important parameter in predicting the dynamic response. We normally wish to have as high a servo bandwidth as possible; this offers the highest rejection of outside perturbations, and provides the greatest dynamic performance. The servo bandwidth is limited by a number of factors, but usually it is the phase lag resulting from the first structural resonance that sets a practical limit. Attempts to increase the servo bandwidth beyond the limit set by structural

resonances turn our servo positioner into an oscillator, with a poor prognosis for its service life. The “natural” form for the servo bandwidth is ω_0 , expressed in radians per second, but the more familiar term for the servo bandwidth is f_0 , expressed in Hertz. A conservative value for the realistic servo bandwidth of either air or mechanical bearing stages is about 50 Hz.

FRICTIONLESS AIR BEARING STAGES

Unlike conventional stages, which suffer from friction and numerous other mechanical maladies, direct-drive air bearing stages are a nearly perfect “physics package”, without friction or physical contact. Their performance is accordingly very easy to model mathematically, and the fidelity of the models to real-world results is quite good. While space here is inadequate for a detailed model description, the decay of following error after a small step move can be fit very well by a simple exponential decay, whose time constant is a direct function of the servo bandwidth. Since these stages lack friction, the integrator term can be zero or very small, and the time constant is then determined by the servo loop proportional term, (assuming, of course, that the derivative term has been properly set so as to provide adequate damping). The proportional time constant “tau” for this decay is $1/\omega_0$, or $1/2\pi f_0$, and for a typical servo bandwidth of 50 Hz, this is about 3.2 milliseconds. The time behavior of the following error, starting with a step function equal to the move size as shown in Fig. 3, is as follows:

$$\text{Following error} = Xe^{-t/\tau}$$

where X is the move size, and τ is the time constant $1/2\pi f_0$. Accordingly, the following error will drop by a factor of e (2.718) every tau. If we will permit a “close counts” approximation, the following error will fall by a factor of 3 every 3 milliseconds. Let’s take the example of our 10 micron move shown in Fig. 3. At time $t=0$, our command position is at +10 microns, the stage has not yet moved, and our following error is by definition 10 microns. At $t = 3$ msec, the error has fallen by a factor of three to 3 microns (remember, close counts). After 6 msec, the error is 1 micron; in 9 msec. it has dropped to 300 nanometers, and continuing in this manner, we have dropped to within 10 nanometers in a mere 18 milliseconds. Viewed from another perspective, we can now make 50 ten micron moves per second, settling to 10 nm. The real world being as it is, a more realistic settling window with moderate cost encoders would be below 25 nm. That is a dramatic improvement over any other positioning technology, and highlights the very real advantages of frictionless air bearing stages. A real-world example is shown in Fig. 5, with the graph displaying laser interferometer data taken at the tooling point. In this example, the stage stack included three axes (X, Y, and Theta), and the move was performed on the bottom-most axis. The vertical scale is 100 nanometers per division, while the horizontal scale is 50 milliseconds per division; the payload mass was 4 kilograms.

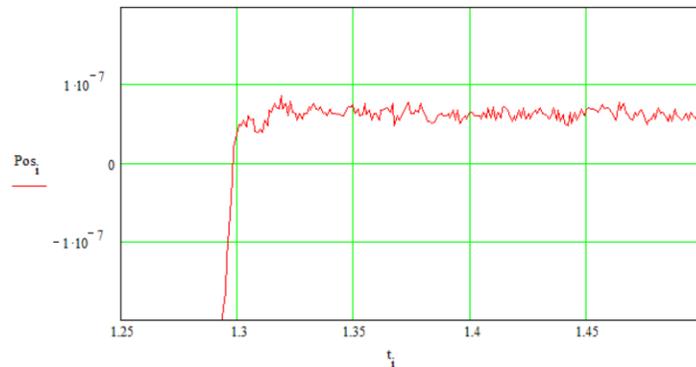


Figure. 5.

FAILURE OF THE PROPORTIONAL TERM

The dominant problem that conventional stages suffer from is friction, which leads to the remarkable fact that the proportional term fails to be of any use whatsoever for small moves with this class of stages. To better understand the problem, refer to the graph in Fig. 6.

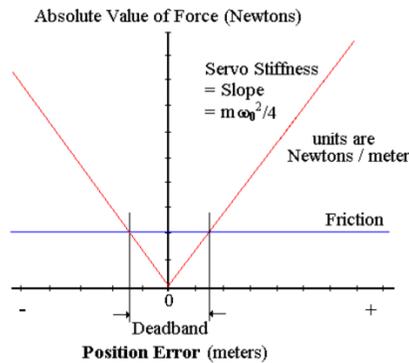


Figure. 6.

In this graph, the vertical axis is the absolute value of force (in a rotary driven system, torque would be substituted), and the horizontal axis is the position error, both positive and negative. If we look at the proportional term of a servo loop, it produces a force which is linearly proportional to the error (hence its name). For example, if an error of 1000 counts results in 100 Newtons of force, then an error of 500 counts would produce 50 Newtons, and so on. The response of the proportional term is shown by the red line in Fig. 6. The slope of this line is the servo stiffness, in Newtons per meter; as it happens, this can be readily calculated – it is equal to $m\omega_0^2/4$, where m is the mass in kilograms, and ω_0 is the servo bandwidth in radians per second. The more familiar servo bandwidth in Hz. is simply $\omega_0/2\pi$. Returning to the graph in Fig. 6, note the horizontal line just above the X axis. This corresponds to the friction in the system, in units of Newtons. The force developed by the proportional term acts to drive the moving element of the stage towards zero position error. A problem arises, however, when the force due to the proportional term is less than or equal to the frictional force. At this point, we're stuck: the stage is, say, 50 microns from the target position. The proportional term responds with 5 Newtons of force, but with a frictional force of 6 Newtons, nothing happens. The stage motion has encountered the "friction boundary", at which the proportional term fails. Were this the only term in the servo loop filter, the following error would remain trapped at this level, never reaching the target position. The proportional term of the servo loop can be thought of as continuously asking the position counter "Where am I"? Upon obtaining the position error, it calculates what it thinks to be the appropriate restoring force, which it then writes to the output DACs. The problem is that this value is less than the system friction, and no motion ensues. Our high performance digital motion controller may have an impressive sample rate of 5 kHz, but the proportional term is pretty simple-minded, and just doesn't "get the picture". Each second it accurately calculates five thousand output values, all equal and inadequate to move the stage. If motion in one direction is considered, this failure of the proportional term will cause the stage to stop well short of its destination; if we include moves in both directions, the error is doubled, as shown by the "deadband" distance in Fig. 6.

The magnitude of the problem can be readily calculated. To do so we simply divide the friction (in Newtons) by the stiffness, in Newtons per meter, to obtain the error in meters. This result can then be doubled if we want to consider bi-directional motion. Taking advantage of the stiffness formula provided above, and using the more familiar Hz. value for the servo bandwidth, the resulting formula becomes:

$$\text{Error (in meters)} = F/m\pi^2f_0^2$$

Where F is the friction in Newtons, m is the mass in kilograms, and f_0 is the servo bandwidth in Hz. If we plug in values of 2 Newtons friction, 1 kilo moving mass, and a servo bandwidth of 50 Hz, the resulting value for the friction boundary is a whopping 81 microns! For moves of this size or smaller, the proportional term might as well be turned off. These values are pretty typical of mechanical bearing stages; the friction would be a little larger for a recirculating bearing stage, and a little less for a crossed roller stage, but none of this affects the basic conclusion. For the micron and sub-micron

sized moves we hope to use during alignment operations, our servo loop simply doesn't work. This is clearly not an acceptable situation.

In the above example, we were assuming a direct drive for our mechanical bearing stage. It's worth asking if the use of a leadscrew can improve the situation. As pointed out in the section "Limitations of Leadscrews", there are numerous reasons to avoid leadscrews in high precision mechanisms, but let's see how they do in addressing the issue of friction. We can reformulate the above equation to reflect the case of a leadscrew based system. In this case, the angular error, in meters, is as follows:

$$\text{Error (in meters)} = LT/2J\pi^3f_0^2$$

Where L is the screw lead (advance per revolution) in meters, T is the torque in Newton-meters, J is the total rotary moment of inertia in kilogram-meters squared, and f_0 is again the servo bandwidth in Hertz. The rotary inertia is dominated by the motor rotor, followed by the leadscrew, and, in a distant third place, the reflected payload inertia. If we plug in typical values (a screw lead of .002 meters, leadscrew torque of 0.05 Newton-meters, total rotary inertia of $5 \times 10^{-5} \text{ kg-m}^2$, and a servo bandwidth of 50 Hz.), we arrive at a friction boundary value of 13 microns. Well, it appears that the mechanical advantage of leadscrews helps a little bit here, but the inability to make moves smaller than 13 microns is of little value.

CAN THE INTEGRATOR TERM HELP?

In a typical servo controller PID loop, however, there are two additional terms present. The "D" term supplies a force or torque which opposes motion, and which is proportional to velocity. While this term provides the damping necessary to ensure stability, it is of no use once motion has ceased at the friction boundary, and the force it produces was of the wrong sign to begin with. If conventional stages with friction are to close to final position at all, they must turn to the last of the three terms in the PID loop: the "I", or integrator term. The good news is that the integrator term, unlike the proportional term, "gets the picture", and will slowly sum the errors of past samples to produce a growing output command that will eventually get us to zero steady-state position error. The bad news is that the introduction of the integrator term degrades stability, and that for stable systems (a reasonable expectation, after all), the integrator "tau", or system time constant, will be five to ten times that of the proportional tau. Unlike the proportional tau (τ_p) of 3.2 milliseconds, the integrator tau (τ_i) will be on the order of 25 milliseconds. The ten micron move in our frictionless air bearing example, which took a mere 18 milliseconds, will be extended to ~150 milliseconds when a conventional stage with friction is chosen. While there are a series of tricks (gain scheduling, friction bias, backlash compensation, etc.) that can be used to try to patch over this fundamental problem, the issue remains, and looms large when competitive comparisons are made.

SUMMARY

The whole thrust of this section can be summarized as follows: our miniature, frictionless, air bearing stages can follow the rapid exponential decay of the proportional servo down "into the noise", producing dramatic benefits in system throughput and accuracy. The ultimate limits on performance for these stages are set only by encoder resolution, amplifier linearity, and external vibration. Conventional stages, beset by friction and other mechanical nonlinearities, are dependent on the action of a sluggish integrator term in the servo loop filter, and this sets a fundamental limit on their dynamics. As a result, they are considerably more limited in performance, and will have substantially lower throughput and precision.

ABOUT THE AUTHOR...

Kevin McCarthy is the chief technology officer of Dover Motion.