Proceedings of the 1998 IEEE
International Conference on Robotics & Automation
Leuven, Belgium · May 1998

3-DOF Closed-loop Control for Planar Linear Motors

Arthur E. Quaid and Ralph L. Hollis
The Robotics Institute, Carnegie Mellon University
{aquaid, rhollis}@ri.cmu.edu

Abstract

Planar linear motors (Sawyer motors) have been used in industry as open-loop stepping motors, but their robustness and versatility has been limited. Using a sensor recently integrated into such a motor, a closed-loop 3-DOF controller has been implemented. The software-based control system consists of a commutator for computing amplifier currents from actuator forces, a force resolution function for solving the redundant actuation and saturation problems, and an observer for producing a velocity estimate, together with a PID controller. Experiments are performed using a 2-axis laser interferometer to show that the controller has sub-micron resolution, 2 μm peak-to-peak repeatability, and settling times after trajectories of about 20 ms. Limitations of the PID controller are discussed and ideas for improvements are presented.

1 Introduction

Commercial planar linear motors are available that have micron-level precision over meter-sized planar workspaces. They can move with velocities of several meters per second and accelerations of several g's. However, due to the lack of a suitable position sensor, they have been operated as open-loop stepping motors. It was recognized early that performance could be enhanced through sensing [1], but early attempts to develop sensors were not very successful.

Recently, prototype sensors have been designed and developed. In [2] a 1-DOF magnetic platen sensor was mounted on an outrigger off a commercial planar motor, and a PID controller was implemented, although few results are presented. In [3] a 1-DOF sensor of a different design was similarly mounted, and preliminary results for PD control suggested improved resolution, stiffness, and settling times.

Our group has recently completed development of a compact 3-DOF magnetic platen sensor. It has been integrated into the center of a commercial planar linear motor, in space that was otherwise wasted, and has a linear resolution (1σ) of 0.2 μm and an angular resolution of 0.0014°. With sensing (described in a companion paper [4]) and closed-loop control (the subject of this paper), planar linear motors have the potential for much wider application. Guarded moves, fault detection, automatic registration, compliant motions, and cooperative manipulation are some of the abilities possible with closed-loop control. In [5], an adaptive controller for a 3-DOF planar motor is presented with simulation results, although sensor noise and actuator saturation appear to be neglected.

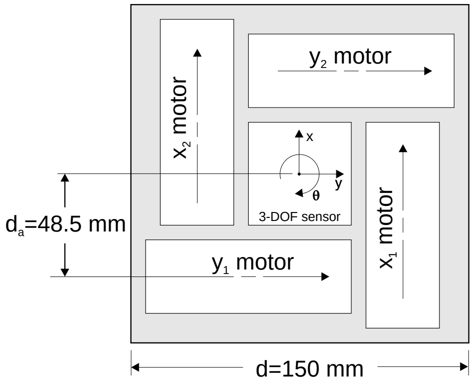

Figure 1: Schematic view of the underside of the forcer. Four linear motors combine to provide a 3-DOF force capability. An integral platen sensor provides 3-DOF sensing for closed-loop control.

In more detail, planar linear motors consist of a moving *forcer* that translates in two directions on a passive steel *platen* stator surface etched with a waffle-iron type pattern. The forcers are supported by a 12-15 μm thick air bearing pre-loaded by permanent magnets, and require a tether to supply air and power. These forcers and platens are available commercially.

The particular forcer examined in this paper is shown schematically in Fig. 1. Two pairs of *motors*[1] mounted orthogonally generate balanced forces about the center of mass. Each of the four motors consists of a stack of laminations and two coils, shown schematically in Fig. 2. The motors operate on a flux-steering principle, with the coil currents acting to switch the permanent magnetic flux from one set of poles to the other. The poles with the most flux tend to align themselves with the platen teeth, so that by activating the poles in the proper order, a stepping motion is achieved. The coil currents can also be *microstepped* by applying a sine wave to one coil and a cosine wave to the other. A more detailed presentation of open-loop operation can be found in [6].

[1] Here, *motor* refers to one of the four actuators on the forcer, and *planar linear motor* refers to the entire device.

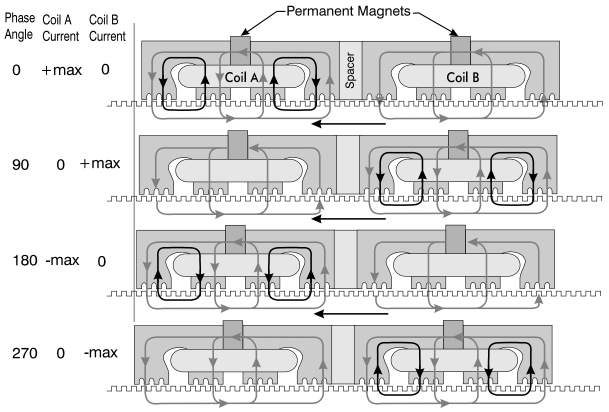

Figure 2: Basic linear motor operation: Currents in the motor coils generate magnetic flux (dark flux path) that sums with the permanent magnet flux (light flux path) to produce forces.

The next section describes the software components of the control system. Section 3 presents the hardware system and the experimental results. The paper concludes with a summary of the results and a discussion of the limitations of the implemented controller.

2 Controller Formulation

A block diagram of the system is shown in Fig. 3. There are four software blocks required. The *estimator* is used to compute a velocity and position estimate based on the sensor position output and the motor dynamics. The *controller* computes a wrench ($\mathbf{w} = [f_x\ f_y\ \tau]^T$) command to be generated by the forcer, based on the estimator outputs and desired trajectory. The *force resolution* block is needed to handle the redundancy of the force generation, as four actuators are used to generate a 3×1 wrench. Finally, the *commutator* computes currents to send to the motor coils based on the forcer position and the desired forces.

It will simplify things to specify wrenches and positions at the center of mass or center of actuation for different blocks. The wrench at the center of actuation, $\mathbf{w}_{ca}$, and the wrench at the center of mass, $\mathbf{w}_{cm}$, are related by:

$$\mathbf{w}_{ca} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ -p_y & p_x & 1 \end{bmatrix} \mathbf{w}_{cm}, \tag{1}$$

where $\mathbf{p} = [p_x\, p_y]^T$ is the location of the center of mass of the forcer expressed in the coordinate system shown in Fig. 1. Similarly, the positions and velocities are related by:

$$\begin{bmatrix} x \\ y \\ \theta \end{bmatrix}_{ca} = \begin{bmatrix} 1 & 0 & p_y \\ 0 & 1 & -p_x \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ \theta \end{bmatrix}_{cm} - \begin{bmatrix} p_x \\ p_y \\ 0 \end{bmatrix}, \text{ and} \tag{2}$$

$$\begin{bmatrix} \dot{x} \\ \dot{y} \\ \dot{\theta} \end{bmatrix}_{ca} = \begin{bmatrix} 1 & 0 & p_y \\ 0 & 1 & -p_x \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} \dot{x} \\ \dot{y} \\ \dot{\theta} \end{bmatrix}_{cm}, \tag{3}$$

where the small rotation range of the motor justifies linearization of the equations about the zero angle.

In this remainder of this section, each software block is examined, with relevant motor modeling introduced as needed.

2.1 Commutation

Derivation of the commutation functions involves first finding a suitable model of the force generated by a single motor segment, given the amplifier inputs. An inversion of this model is then used for commutation.

As described more fully in [7], each linear motor segment generates forces according to the equation:

$$\begin{aligned} f_{x1} &\triangleq f(i_{x1}, x_1, \psi_{x1}) \\ &= k_a(i_{x1}, x_1, \psi_{x1}) \sin\left(\frac{2\pi}{p}x_1 - \psi_{x1}\right), \end{aligned} \tag{4}$$

where x_1 is the motor position in its direction of force generation, p is the pitch of the motor, and ψ_{x1} is the motor phase commanded by the amplifier. k_a is a proportionality factor which depends largely on amplifier current, but also on the motor position and skew angle. In this work, all experiments are performed with small ($\ll 1°$) rotations and are not overly sensitive to force ripple. Thus, it is sufficient to ignore the dependence on position and angle. Then, k_a is a linear function of only the amplifier current i_{x1}, and (4) becomes:

$$f_{x1} = k i_{x1}\ \sin\left(\frac{2\pi}{p}x_1 - \psi_{x1}\right). \tag{5}$$

The commutator needs to find an i_{x1} and ψ_{x1} that generate a commanded force f_{x1} by inverting (5). There are an infinite number of solutions, but as introduced in [7], two interesting possibilities are a *fixed amplitude* solution, where i_{x1} is set to a constant value and ψ_{x1} is varied according to f_{x1}, and a *fixed phase* solution, where ψ_{x1} is chosen constant relative to the motor position and i_{x1} is varied. In this work, the *fixed phase* approach is used, allowing the amplifier currents to be zero when f_{x1} is zero, which reduces thermal effects. With this approach, the commutator chooses amplifier inputs according to the equations:

$$i_{x1} = \frac{f_{x1}}{k}, \text{and} \tag{6}$$

$$\psi_{x1} = \frac{2\pi}{p}x_1 - \frac{\pi}{2}. \tag{7}$$

The position x_1 is computed based on the sensed forcer position. At high speeds, it is also important to add *phase advance* to compensate for the latency between the sensor reading and the control output [7].

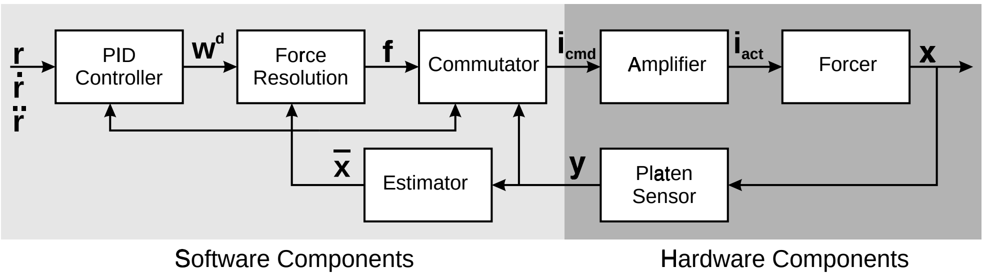

Figure 3: Controller block diagram

This term requires a velocity estimate. The position of each linear motor can then be computed:

$$\begin{bmatrix} x_1 \\ x_2 \\ y_1 \\ y_2 \end{bmatrix} = \begin{bmatrix} 1 & 0 & -d_a \\ 1 & 0 & d_a \\ 0 & 1 & -d_a \\ 0 & 1 & d_a \end{bmatrix} \left(\mathbf{y}_{ca} + t_a \begin{bmatrix} \bar{\dot{x}} \\ \bar{\dot{y}} \\ \bar{\dot{\theta}} \end{bmatrix}_{ca} \right), \tag{8}$$

where t_a is the phase advance time, $[\bar{\dot{x}}\ \bar{\dot{y}}\ \bar{\dot{\theta}}]^T$ is the *estimated* forcer velocity, and $\mathbf{y}$ is the *sensed* forcer position, both expressed at the center of actuation.

2.2 Force resolution

The force resolution function must consider the force kinematics and force saturation properties of the forcer. The force kinematics are given by:

$$\mathbf{w}^d_{ca} = \begin{bmatrix} f^d_x \\ f^d_y \\ \tau^d \end{bmatrix}_{ca} = \begin{bmatrix} 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 \\ -d_a & d_a & -d_a & d_a \end{bmatrix} \begin{bmatrix} f_{x1} \\ f_{x2} \\ f_{y1} \\ f_{y2} \end{bmatrix}, \tag{9}$$

where $\mathbf{w}^d_{ca}$ is a wrench applied at the center of the forcer, and d_a is the distance from the center of the motor to the center of the forcer, as shown in Fig. 1.

The force saturation properties of the motor sections are determined by their maximum rated current: $f_{max} = ki_{max}$. If each of the four motor sections can generate up to a maximum force f_{max}, the wrench saturation constraints for the forcer are:

$$|f^d_x| \leq 2f_{max}, \tag{10}$$

$$|f^d_y| \leq 2f_{max}, \text{and} \tag{11}$$

$$|\tau^d| + d_a\left(|f^d_x| + |f^d_y|\right) \leq 4f_{max}d_a. \tag{12}$$

The *force resolution problem* is to find a suitable solution to (9) subject to these constraints.

It is helpful to consider this problem geometrically. By eliminating the absolute values, these equations can be expanded into 12 inequalities linear in the wrench parameters f_x, f_y, and τ. These constraints combine to form a *wrench envelope* $\mathcal{E}$ that can be represented in $I\!R^3$ as a rhombic dodecahedron, depicted in Fig. 4. For example, at point $\mathbf{p}_1$, the forcer is generating the maximum possible force in both the $+x$ and $-y$ directions and cannot generate any torque. Alternatively, at point $\mathbf{p}_2$, the forcer is using its full capabilities to generate torque, and cannot generate any forces.

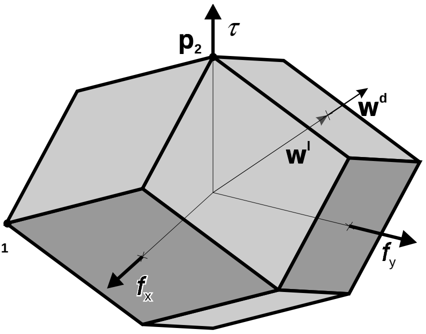

Figure 4: A *wrench envelope* represents the force/torque limits of the forcer. The darker faces result from constraint equations (10,11), and the lighter faces result from constraint equation (12).

The input to the force resolution function is the desired wrench $\mathbf{w}^d_{ca}$, which may be outside, inside or on $\mathcal{E}$ leading to over-, under-, and uniquely-constrained cases.

In the under-constrained case, an infinite number of solutions can be found that satisfy (9)-(12). Here, we choose the solution:

$$\begin{bmatrix} f_{x1} \\ f_{x2} \\ f_{y1} \\ f_{y2} \end{bmatrix} = \begin{bmatrix} \frac{1}{2} & 0 & -\frac{a}{a+b}\frac{1}{2d_a} \\ \frac{1}{2} & 0 & \frac{a}{a+b}\frac{1}{2d_a} \\ 0 & \frac{1}{2} & -\frac{b}{a+b}\frac{1}{2d_a} \\ 0 & \frac{1}{2} & \frac{b}{a+b}\frac{1}{2d_a} \end{bmatrix} \begin{bmatrix} f^d_x \\ f^d_y \\ \tau^d \end{bmatrix}_{ca}, \tag{13}$$

where $a = 2f_{max} - |f^d_x|$ and $b = 2f_{max} - |f^d_y|$. Note that a and b are a measure of the remaining force capability of the forcer in the x and y directions. Note that a or b are both non-negative, and because the wrench is inside the force envelope, at least one is non-zero. Therefore, (13) will always be defined, and (by examining the derivative) can also be shown to be continuous.

If the desired wrench lies on envelope $\mathcal{E}$, there is a single solution. This solution is identical to the under-constrained case except that it becomes undefined when $a = b = 0$. In this special case the desired wrench is at one of the corners in the f_x, f_y plane (*e.g.* $\mathbf{p}_1$) in Fig. 4, so that $\tau = 0$, and (13) is replaced by:

$$\begin{bmatrix} f_{x1} \\ f_{x2} \\ f_{y1} \\ f_{y2} \end{bmatrix}^d = \begin{bmatrix} 1/2 & 0 & 0 \\ 1/2 & 0 & 0 \\ 0 & 1/2 & 0 \\ 0 & 1/2 & 0 \end{bmatrix} \begin{bmatrix} f_x^d \\ f_y^d \\ \tau^d \end{bmatrix}_{ca} . \qquad (14)$$

In the over-constrained case, $\mathbf{w}_{ca}^d$ lies outside the wrench envelope. In this case, the forcer is saturated, and there are no solutions to (9) that satisfy the constraints. However, if we redefine the problem for this case to be mapping $\mathbf{w}_{ca}^d$ back onto the wrench envelope, there are once again infinite solutions. One simple solution is to linearly scale the desired wrench vector back to the point where it pierces the wrench envelope. For example, $\mathbf{w}_{ca}^d$ in Fig. 4 is mapped to $\mathbf{w}_{ca}^l$. They are related by a scale factor s_l:

$$s_l = \max\left(\frac{|f_x^d|}{2f_{max}}, \frac{|f_y^d|}{2f_{max}}, \frac{|f_x^d| + |f_y^d| + |\tau^d|/d_a}{4f_{max}} \right), \qquad (15)$$

such that $\mathbf{w}_{ca}^l = \mathbf{w}_{ca}^d / s_l$.

Although this solution has been implemented, the experiments are designed to avoid saturation cases. Saturation in general is a difficult non-linearity to deal with effectively. In this case it also acts to couple the axes, requiring a more complicated controller design.

2.3 Control

Given the commutator and force resolution functions, the controller can be designed around a simple linear model. The mechanical dynamics of the forcer are simply those of a mass moving in the plane:

$$\mathbf{w}_{cm} = \begin{bmatrix} m & 0 & 0 \\ 0 & m & 0 \\ 0 & 0 & I_z \end{bmatrix} \begin{bmatrix} \ddot{x} \\ \ddot{y} \\ \ddot{\theta} \end{bmatrix} + \begin{bmatrix} b_v(\dot{x}) & 0 & 0 \\ 0 & b_v(\dot{y}) & 0 \\ 0 & 0 & b_\omega(\dot{\theta}) \end{bmatrix}, \qquad (16)$$

where m is the forcer mass and I_z is the rotational inertia. The b_v and b_ω functions model the eddy-current damping, but are neglected for simplicity and because the damping is negligible relative to the amount of added controller damping. To express the dynamics in this decoupled form, I_z and $\mathbf{w}_{cm}$ are defined relative to the center of mass of the forcer. The gains of a PID control law,

$$u = K\left(e + T_d \dot{e} + \int \frac{1}{T_i} e \right), \qquad (17)$$

can then be chosen independently for each axis. Feedforward acceleration terms from Eq. 16 are also added to the controller forces.

2.4 State estimation

The estimator exploits the linear dynamic model of the forcer to produce a filtered position and velocity signal without excessive lag. The discrete time estimator takes the form:

$$\bar{\mathbf{x}}_{cm}(k+1) = \begin{bmatrix} \mathbf{\Phi} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{\Phi} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{\Phi} \end{bmatrix} \bar{\mathbf{x}}_{cm}(k) +$$

$$\begin{bmatrix} \mathbf{L} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{L} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{L} \end{bmatrix} \tilde{\mathbf{y}}_{cm}(k) + \begin{bmatrix} \frac{T^2}{2m} & 0 & 0 \\ \frac{T}{m} & 0 & 0 \\ 0 & \frac{T^2}{2m} & 0 \\ 0 & \frac{T}{m} & 0 \\ 0 & 0 & \frac{T^2}{2I_z} \\ 0 & 0 & \frac{T}{I_z} \end{bmatrix} \mathbf{w}_{cm}^l (18)$$

with

$$\mathbf{\Phi} = \begin{bmatrix} 1 & T \\ 0 & 1 \end{bmatrix}, \quad \mathbf{L} = \begin{bmatrix} l_1 & 0 \\ 0 & l_2 \end{bmatrix}, \text{ and}$$

$$\tilde{\mathbf{y}}_{cm}(k) = \mathbf{y}_{cm}(k) - \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 \end{bmatrix} \bar{\mathbf{x}}_{cm}(k).$$

Here, $\mathbf{y}_{cm}(k) = [x\ y\ \theta]^T$ is the sensor position output at time k, $\bar{\mathbf{x}}_{cm}(k) = [\bar{x}\ \bar{\dot{x}}\ \bar{y}\ \bar{\dot{y}}\ \bar{\theta}\ \bar{\dot{\theta}}]^T$ is the state estimate at time k, and T is the sample time.

Note that the estimator is decoupled and the estimator gains l_1 and l_2 can be computed based on the desired estimator pole locations using well-known pole-placement techniques (*i.e.* MATLAB's[2] `PLACE` command).

3 Experimental results

The software blocks in Fig. 3 are implemented on a Motorola PowerPC 133 MHz computer running the LynxOS real-time operating system. I/O hardware consist of a number of Industry Pack (IP) credit-card sized modules on an ISA bus carrier. The computer and I/O hardware fit in a standard mini-tower PC case, and connect to the planar linear motor through a tether. The planar linear motor consists of a Normag platen and a modified Normag 2-phase forcer with a 1.016 mm pitch, 60 N nominal static force, 1.4 Kg mass, and 4 A peak operating current. The test setup includes a Zygo 2-axis laser interferometer, which can measure the differential skew angle and one translational axis of the forcer, providing a position measurement independent of the magnetic platen sensor.

The software is structured with a single high-priority thread running at 3500 Hz that includes the commutation, sensor I/O, force resolution, controller,

[2] Product of MathWorks, Inc.

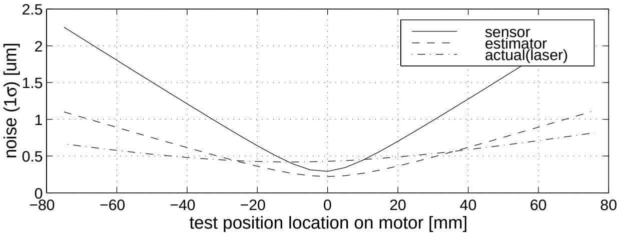

Figure 5: The controller regulates points near the sensor better than points on the periphery of the forcer.

observer, and trajectory generator functions. This set of functions takes approximately 150 μs to complete, leaving some time for lower-priority user interface and network communications threads to execute. The fast update rate is necessary to provide a reasonable number of updates per pitch. With a forcer pitch of 1.016 mm, and a peak speed above 1 m/s, even with 3500 Hz updates there may may be less than four updates per pitch. However, a more complicated controller can still be implemented by separating the commutator and controller into separate threads, with the controller running at a slower rate than the commutator.

3.1 Resolution tests:

To characterize the position resolution under closed-loop control, the PID controller was used to regulate the forcer to zero position and angle. The gains of the controller were $K = 220$ N/mm, $T_d = 0.0053$ s, and $T_i = 0.028$ s. The controller poles are underdamped with natural frequency of 40 Hz. Estimator poles were placed at 80 Hz, which was near the upper limit for the controller rate of 3500 Hz. Above this rate, an unacceptable level of noise from the sensor was passed into the controller, causing the forcer to be *audibly* noisy.

Readings for the x translation and θ skew angle from the platen sensor, estimator, and laser interferometer were recorded at 3500 Hz for 1000 samples. The amount of motion in the x direction of a particular point on the forcer was computed using the simple differential kinematic equation:

$$\delta x_t = \delta x - \delta\theta \; y_t, \qquad (19)$$

where δx_t is the x differential motion of a test point (x_t, y_t) on the forcer given a differential motion at the middle of the forcer of δx and $\delta\theta$. Figure 5 shows the standard deviation of δx_t as recorded by the sensor, estimator, and laser interferometer as y_t is varied. Note that points with $|y_t| = 75$ mm correspond to the edge of the forcer. Because the sensor measures angle by differencing two parallel position measurements that are close together (see [4]), there is a low-noise *sweet spot* in the middle of the forcer, where the sensor is located. However, Fig. 5 indicates that even at the edge of the forcer, the controller maintains micron-level resolution (1σ), which is sufficient for many applications.

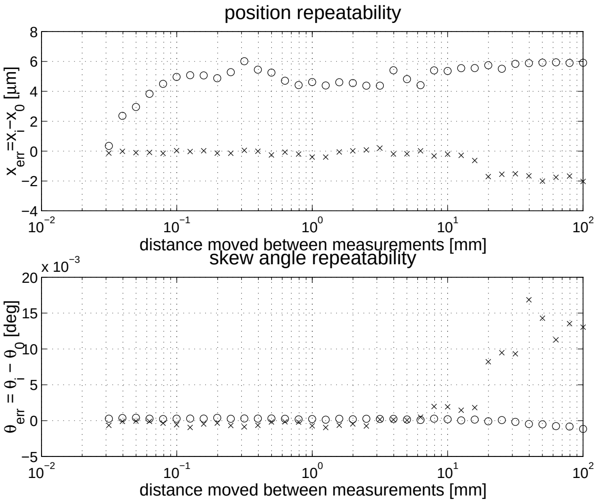

Figure 6: Bi-directional repeatability: errors in moving to a reference position from two approach directions are shown for open-loop (circles) and PID control (crosses) for a range of motion distances.

3.2 Repeatability

Repeatability is the ability of an actuator to return to the same position. The difference in the forcer positions after moving to a reference location from two different directions was measured with the laser interferometer. This process was repeated with varying move distances. The controller was started several minutes before testing began, and the test was designed to be completed in under a minute to minimize thermal effects. The crosses in Fig. 6 show this bi-directional position repeatability for the PID controller to be under 2.5 μm peak-to-peak, and the skew angle repeatability to be under 0.02° peak-to-peak over 36 motions. Note that the error increases at a travel distance of about $20mm$, which is when the linear motors start to overlap the reference position, so it appears likely that the motors are leaving a residual magnetic field in the platen that is slightly affecting the sensor operation. The exact mechanism for this interaction is under investigation. For comparison, the bi-directional repeatability tests are repeated under open-loop control. The error here is probably due to a combination of tether disturbance forces and the same magnetic hysteresis that affects the sensor.

3.3 Trajectory commands

The closed-loop controller was used to track a trajectory with a bang-bang acceleration of 10 m/s^2, maximum velocity of 0.8 m/s, and a position change of 0.1 m. Integral gains were disabled for this experiment to prevent integral windup during the motion.[3]

As shown in the dark traces in Fig. 7, the PD controller tracks the trajectory to within 50 μm, and settles to 1 μm within 20 ms. The tracking error comes

[3] Gain-scheduling could be used to re-enable the integral gain near the end of the trajectory to correct the steady-state errors.

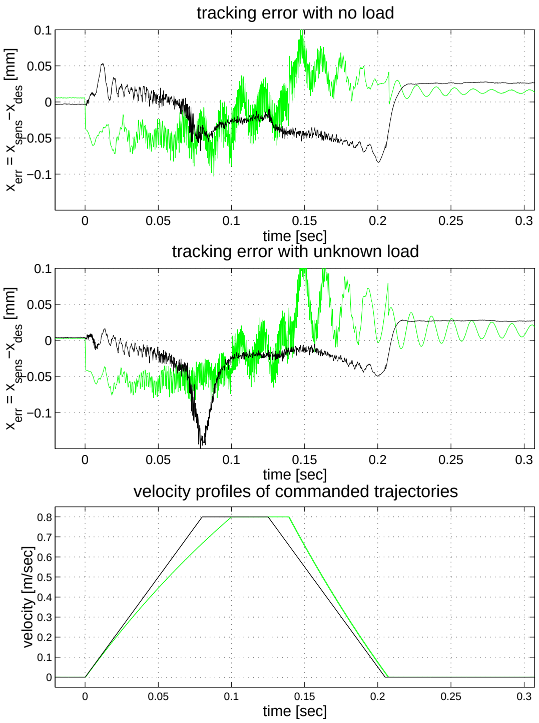

Figure 7: Trajectory tracking: tracking errors are shown for PD (dark trace) and open-loop (light trace) control given no load (upper plot) and an "unknown" load of 240 g (middle plot). The lower plot shows the velocity profile of the trajectories.

from motor modeling errors (i.e. eddy-current damping, errors in the mass or peak force model parameters, or unmodeled actuator nonlinearities.). The tracking error is also shown when a mass of 240 g is attached to the edge of the forcer, causing errors in in the mass, inertia, and center of mass of the model. Even in this case, the controller settles just as well, although the tracking error increases significantly during one part of the trajectory.

For comparison, the tracking experiment was repeated under open-loop control. However, it would be unfair to have the open-loop controller attempt to track the same bang-bang acceleration trajectories. Instead, the trajectory described in [8] was implemented. This trajectory is identical to the bang-bang acceleration trajectory, except for *burst* (step changes applied to the desired position) and *acceleration rolloff* as the velocity increases. These changes are designed so that the open-loop forcer dynamics will not be excited and, in the ideal case, there will be a constant tracking error during each phase (acceleration, slew, and deceleration) of the trajectory.

After a reasonable amount of tuning of the trajectory, tracking errors were as shown in the light traces of the top plot of Fig. 7. Although the tracking error levels are reasonably good, there is a significant oscillation, indicating that some dynamic parameters had errors or there were unmodeled nonlinearities. This oscillation grew worse when the extra load was added, as shown in the middle plot. Furthermore, in both cases, the open-loop controller takes much longer to settle at the end of the trajectory.

4 Conclusions

This work has demonstrated 3-DOF closed-loop control of a planar linear motor using an integrated platen position sensor. Experiments indicate submicron resolution and repeatability, and improved tracking and settling time relative to open-loop control. The *force resolution* problem was identified as a consequence of the actuator redundancy and a solution was presented.

Limitations were found with the PID controller. Trajectories with accelerations or velocities near the limits of the forcer could not be reliably performed due to the non-linear and coupling effects of actuator saturation. Also, at high velocities, disturbance torques (in practice, disturbances from the tether were most common) may momentarily saturate the torque capability of the motor. To reject this disturbance, a PD or PID controller would need an unrealistically high bandwidth. We are presently exploring the use of nonlinear switching controllers to address these problems.

Acknowledgements

This work is supported in part by NSF grants DMI-9523156 and DMI-9527190, and the CMU Engineering Design Research Center. Quaid was supported by an AT&T Fellowship. We acknowledge Yangsheng Xu's early participation in this work, and Alfred Rizzi for work on the real-time computing hardware and operating system and for many helpful discussions. Finally, this work would not be possible without the efforts of Zack Butler in development of the platen sensor.

References

[1] B. A. Sawyer, "Linear magnetic drive system." U. S. Patent 3,735,231, May 22 1973.

[2] F. Y. Wong, H. Schulze-Lauen, and K. Youcef-Toumi, "Modelling and digital servo control of a two-axis linear motor," in *Proc. American Control Conference*, pp. 3659–3663, June 1995.

[3] J. Ish-Shalom, "Sub-micron large motion multi-robot planar motion system," in *Proc. 7th Int'l Symp. on Robotics Research*, pp. 582–595, October 1995.

[4] Z. Butler, A. Rizzi, and R. Hollis, "Integrated precision 3-DOF position sensor for planar linear motors," in *Proc. IEEE Int'l Conf. on Robotics and Automation*, May 1998.

[5] F. Khorrami, H. Melkote, and J. Ish-Shalom, "Advanced control system design for high speed ultra accurate manufacturing systems," in *Proc. of the American Control Conference*, pp. 164–165, June 1997.

[6] E. R. Pelta, "Precise positioning without geartrains," *Machine Design*, pp. 79–83, April 1987.

[7] A. E. Quaid, Y. Xu, and R. L. Hollis, "Force characterization and commutation of planar linear motors," in *Proc. IEEE Int'l Conf. on Robotics and Automation*, April 1997.

[8] J. I. Nordquist and P. M. Smit, "A motion-control system for (linear) stepper motors," in *Proceedings of the Fourteenth Annual Symposium on Incremental Motion Control Systems and Devices*, pp. 215–231, 1985.