The 2009 IEEE/RSJ International Conference on
Intelligent Robots and Systems
October 11-15, 2009 St. Louis, USA

Automated Manipulation of Spherical Objects in Three Dimensions Using A Gimbaled Air Jet

Aaron Becker, Robert Sandheinrich, and Timothy Bretl

Abstract— **This paper presents a mechanism and a control strategy that enables automated non-contact manipulation of spherical objects in three dimensions using air flow, and demonstrates several tasks that can be performed with such a system. The mechanism is a 2-DOF gimbaled air jet with a variable flow rate. The control strategy is feedback linearization based on a classical fluid dynamics model with state estimates from stereo vision data. The tasks include palletizing, sorting, and ballistics. All results are verified with hardware experiments.**

I. INTRODUCTION

Our long-term goal is to enable automated, parallel manipulation of multiple objects with air flow. Two key control challenges are presented by this type of manipulation, in contrast to traditional robotic manipulation with a mechanical gripper. First, the dynamics of the flow field itself are difficult to model. These dynamics are typically governed by systems of partial differential equations and may exhibit behavior that is both uncertain and chaotic. Second, the dynamics of the manipulated objects are strongly coupled, since the presence of an object in a flow field changes the structure of that field for other objects.

To make progress, this paper considers the particular example system shown in Fig. 1, for which it is possible to simplify the above two control challenges. In this system, the objects are spheres and the air flow is generated by a single axisymmetric air jet. This air jet has a variable flow rate and is mounted on an actuated 2-DOF rotary motion stage. Our control inputs are the angles θ_1, θ_2 of the stage and the velocity u of the nozzle flow.

The steep velocity gradient outward from the air jet's axis of symmetry creates a stable equilibrium point at a distance that depends on the nozzle velocity and on the physical characteristics of the sphere. By changing the orientation and flow rate of the jet, we can move spherical objects to any point within a three-dimensional workspace.

Although the underlying physics of this equilibrium point are well known for a vertically mounted jet and make for a classic demonstration in the classroom [1], transient behavior is less well understood. Being able to model and control this transient behavior is necessary for automated point-to-point manipulation. In particular, adjusting the flow rate excites low-frequency, high-amplitude oscillations along the axis of symmetry of the jet. These oscillations take a significant amount of time to settle, tend to have a destabilizing effect on the system, and would preclude rapid manipulation.

University of Illinois at Urbana-Champaign
Aaron Becker is a student in the Electrical and Computer Engineering Department, Robert Sandheinrich is a student in the Mechanical Engineering Department, and Timothy Bretl is an Assistant Professor in the Aerospace Engineering Department.

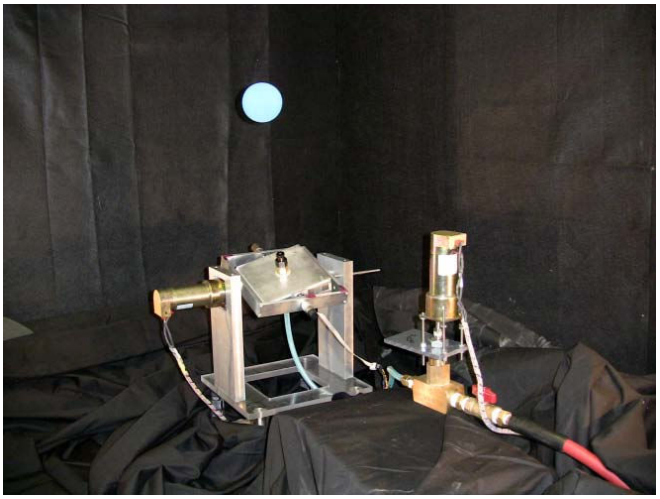

Fig. 1. A spherical object hovering in stable equilibrium above a 2-DOF gimbaled air jet with a variable flow rate. Our control strategy enables automated manipulation of this object in three dimensions.

In this paper we apply feedback linearization based on a classical fluid dynamics model in order to dampen these axial oscillations more quickly. Our approach depends on having a good state estimate, in this case provided by stereo visual feedback from a pair of low-cost cameras.

This control strategy enables a number of manipulation tasks. For example, we can palletize spheres, lifting them on and off a perch and moving them through obstacles in 3D. We can sort spheres without sensors according to their physical characteristics, either stacking several of them in the same flow field or depositing them in bins on the ground. Long-range ballistic positioning is also possible, using a rapid increase in the flow velocity to fire an object to a remote location.

Our hope is that some of these manipulation tasks can be transitioned out of the laboratory and into real-world situations. For instance, because air flow avoids the need for mechanical contact, it is particularly appropriate for applications in the textile, printing, and foodstuffs industries that involve the conveyance or rearrangement of flexible, porous, or delicate objects. Examples include the handling of clothes [2], [3], paper [4], sliced fruit and vegetables [5], and biscuits [6]. Similarly, this type of manipulation can move many objects at the same time, and may increase the throughput of systems for industrial parts handling.

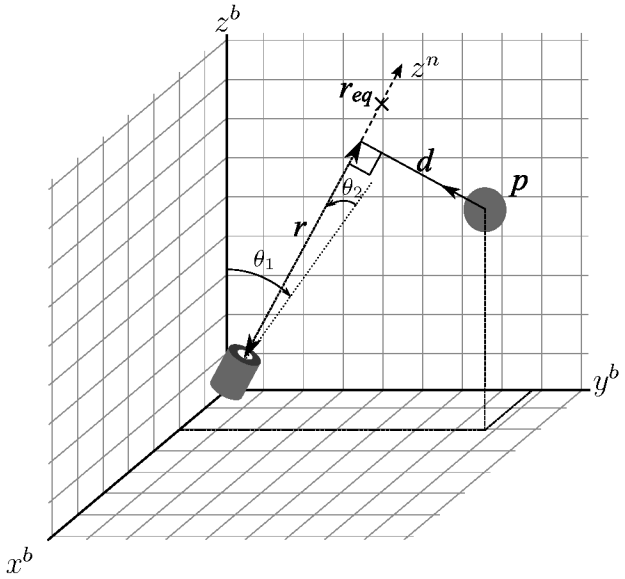

Fig. 2. Schematic of our hardware system. The air nozzle is located at the origin and the sphere position, nozzle direction and equilibrium position are shown. The sphere's location is p, the distance from origin to the sphere equilibrium position is r_{eq}. The perpendicular distance from the sphere to the nozzle centerline is d. The component of p along e_3 of the nozzle frame is r.

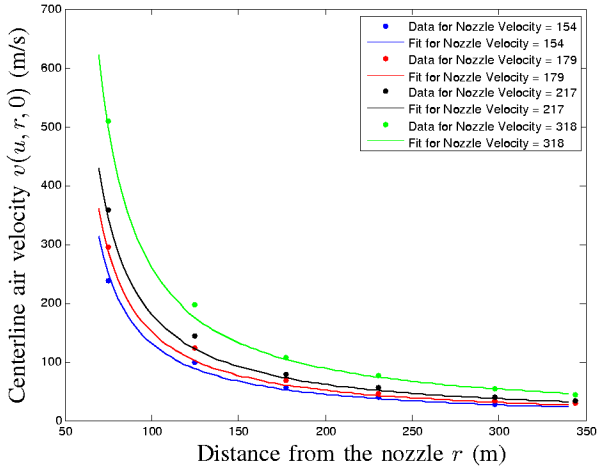

Fig. 3. Measured air velocity along the jet's centerline $d = 0$ for increasing distance from the nozzle r. Each line represents a different nozzle velocity u. The solid lines are fits of the form $v(u, r, 0) = \frac{u}{r+c_1}$ where c_1 is a constant.

The outline of this paper is as follows. In Section II we describe our hardware system in more detail and present our control strategy. In Section III we show how our control strategy can be applied to enable a number of manipulation tasks. Finally, we present some concluding remarks in Section IV.

II. CONTROL STRATEGY

A. Mechanism

Figure 1 shows our mechanism for automated manipulation of objects with air flow. It consists of an air jet mounted to a 2-DOF rotary motion stage. The air jet is supplied up to 620 kPa through a valve that is continously adjustable via a DC motor with encoder feedback. Both the motion stage and the valve are controlled by a digital signal processor (DSP) running a 1 kHz control loop, which allows us to command the angles θ_1, θ_2 of the motion stage and the flow velocity u as inputs to the system. The objects are spheres, with radii from 12 to 97 mm and masses from 2.6 to 188 g. The positions of these spheres are captured with stereo vision from two orthogonally mounted cameras triggered at 55 Hz.

B. Dynamic model

1) Coordinates: We assume the object to be manipulated is a perfect sphere and define its configuration by $p \in \mathbb{R}^3$ as shown in Fig. 2. It will be useful for us to express p in two different Cartesian coordinate systems, one fixed in the workspace and one aligned with the rotary motion stage. We call the former the *base frame* b and denote the corresponding coordinates by p^b; we call the latter the *nozzle frame* n and denote the corresponding coordinates by p^n. We assume the air jet is aligned with the unit vector $e_3 = \begin{bmatrix} 0 & 0 & 1 \end{bmatrix}^T$ in the nozzle frame. Given the angles θ_1 and θ_2 of the stage, we define a rotation matrix

$$R_n^b = \begin{bmatrix} 1 & 0 & 0 \\ 0 & \cos\theta_1 & -\sin\theta_1 \\ 0 & \sin\theta_1 & \cos\theta_1 \end{bmatrix} \begin{bmatrix} \cos\theta_2 & 0 & \sin\theta_2 \\ 0 & 1 & 0 \\ -\sin\theta_2 & 0 & \cos\theta_2 \end{bmatrix}$$
$$= \begin{bmatrix} \cos\theta_2 & 0 & \sin\theta_2 \\ \sin\theta_1\sin\theta_2 & \cos\theta_1 & -\sin\theta_1\cos\theta_2 \\ -\cos\theta_1\sin\theta_2 & \sin\theta_1 & \cos\theta_1\cos\theta_2 \end{bmatrix}$$

between these two coordinate frames, so we can say

$$p^b = R_n^b p^n \qquad \text{and} \qquad p^n = R_b^n p^b$$

where

$$R_b^n = \left(R_b^n\right)^T.$$

We define the *axial distance* r and the *perpendicular distance* d to the object as follows:

$$r = e_3^T R_b^n p^b \qquad \text{(axial distance)}$$
$$d = \left\| \begin{bmatrix} e_1 & e_2 \end{bmatrix}^T R_b^n p^b \right\| \qquad \text{(perpendicular distance).}$$

In what follows, we will characterize the dynamics of the spherical object in terms of these variables r and d.

2) Axial flow velocity: We make the common assumption that the flow field of the air jet is axisymmetric and is dominated by the velocity component in the axial direction, i.e., the direction e_3 as expressed in the nozzle frame [7]. For a given velocity u at the nozzle outlet, the axial velocity v varies as a function of the axial distance r and the perpendicular distance d from the nozzle. In particular, Pitot tube measurements, as shown in Figs. 3- 4, confirm the relationship

$$v(u, r, d) = \frac{u}{r + c_1} \operatorname{sech}^2\left(c_2 \frac{d}{r}\right) \qquad (1)$$

where c_1 and c_2 are fixed constants [8], [9].

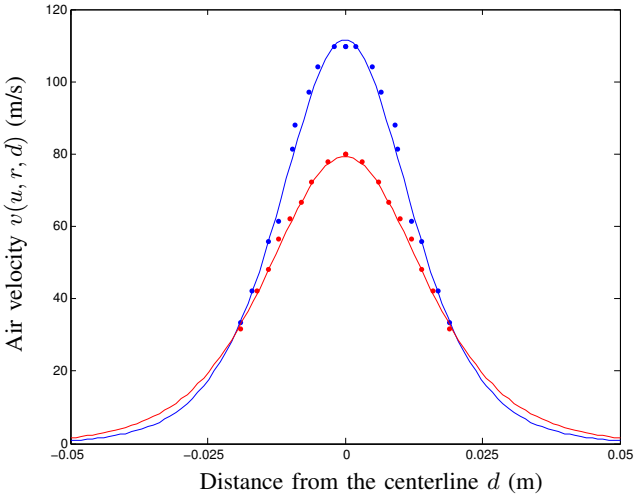

Fig. 4. Measured free air velocity v in the flow field for varying distance from the jet's centerline d. Each line represents a different axial distance r. The red line corresponds to $r = 0.232$ m and the blue line $r = 0.187$ m. The solid lines are fits of the form $v(u, r, d) = v(u, r, 0)\,\mathrm{sech}^2\left(c_2 \frac{d}{r}\right)$ where c_2 is a constant. Notice that the nozzle velocity decays significantly for $d > 0.02$ m.

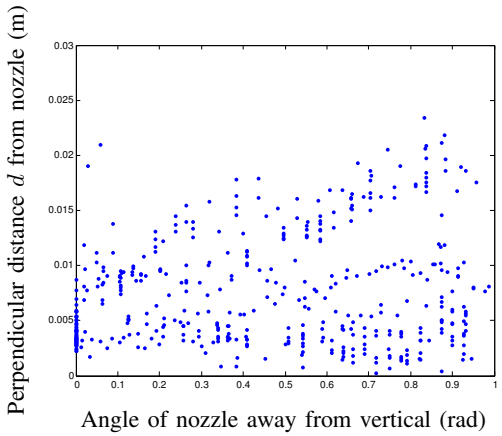

Fig. 5. Dependence of equilibrium perpendicular distance d on the angle of the nozzle away from vertical for several nozzle velocities u. This raw position data shows that for any configuration, the equilibrium position of a 6.7 g 40 mm diameter sphere is within 25 mm of the nozzle axis. By assuming that the perpendicular position is stable, the modeling is greatly simplified.

3) Flow field: It is important to note that the dynamics of flow itself can have a significant effect on the system. Changes to the nozzle pressure take time to propagate to objects in the flow. In addition, the volume of air in the system can store and output energy over time. These effects are part of the complex dynamics of the overall system that make a precise model impractical. For our model, we assume changes to the fluid are instantaneous and memoryless, so changes to u have an immediate and time invariant effect on objects in the flow.

Flow is also assumed to be in the turbulent regime with Reynolds number $R_e = \frac{v\beta}{\nu}$ in the range of $10^4 - 10^6$, where β is the radius of the sphere and ν is the kinematic viscosity of the fluid. In this regime, the relationship between air velocity and drag has a nonlinear dependence on R_e [10]. By assuming the sphere is always near the equilibrium point, this relationship can be neglected. While flow through and near the nozzle may be supersonic, the models here assume that flow around the sphere is subsonic and incompressible.

4) Perpendicular motion: Because of the steep velocity gradient at small distances from the nozzle axis (Fig. 4), the position of a spherical object in the direction perpendicular to this axis is stable about $d = 0$. A position offset from $d = 0$ causes a velocity, and therefore pressure, difference across the cylinder. This pressure difference makes the sphere stable in the direction perpendicular to the flow, as shown in Fig. 5.

5) Axial motion: The axial dynamics r of a spherical object are governed by the standard drag equation and gravity.

$$\ddot{r} = \frac{1}{2}\frac{C_d}{m}\rho A(v(u, r, d) - \dot{r})^2 - g e_3^T R_b^n e_3 \qquad (2)$$

Here, C_d is the coefficient of drag, m the mass of the sphere, ρ the density of the fluid and A the sphere cross sectional area. In writing this equation of motion, we assume that the angular velocity of the nozzle frame is small relative to the dynamics of axial motion. For the applications we consider in this paper, this assumption is reasonable. Substituting equation 1 into equation 2 and solving for zero acceleration and velocity gives the axial equilibrium position

$$r_{\text{eq}}(\theta_1, \theta_2, u) = u\sqrt{\frac{C_d}{m}\frac{\rho A}{2g\cos\theta_1\cos\theta_2}} + c_3 \qquad (3)$$

where c_3 is a fixed constant. In other words, for a given θ_1, θ_2, and u, the following configuration is stable:

$$p^b = R_n^b \begin{bmatrix} 0 \\ 0 \\ r_{\text{eq}} \end{bmatrix}.$$

It is easy to invert this relationship to find the values of θ_1, θ_2, u required to achieve a given configuration p^b.

6) Frequency response: Figure 6 shows the helical motion of the sphere in three dimensions about this equilibrium configuration. It traces elliptical patterns around the centerline of the nozzle, but the dominant motion is in the axial direction.

To characterize the stability of this equilibrium configuration, we measured the frequency response of the system from the nozzle velocity u to sphere position, in both the axial and perpendicular directions (Fig. 7). In each case, there are two resonant frequencies, and the corresponding low-frequency oscillations are lightly damped.

Notice that the amplitude of these oscillations is much larger in the axial direction than in the perpendicular direction. The large-amplitude axial oscillations make the settling time for a step response large, precluding rapid manipulation. It is this problem that we will correct with our control design.

C. Control design

Based on our fluid dynamic model, we applied feedback linearization to control the axial position of the sphere. Given

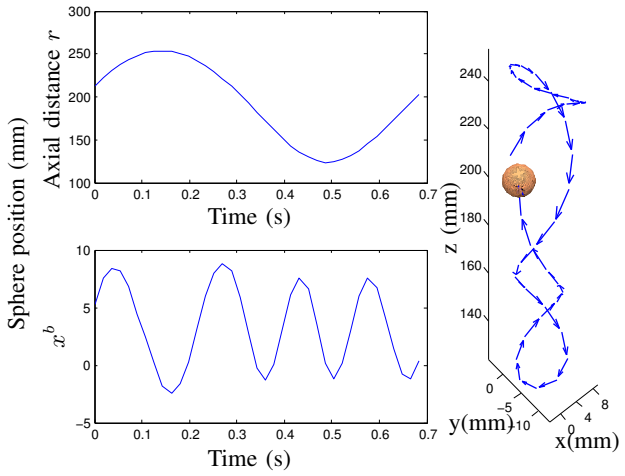

Fig. 6. Recorded 0.7 s trajectory of a sphere in Cartesian coordinates. Top left, r oscillations vs. time. Bottom left, x-axis oscillations over time. The perpendicular amplitude is an order of magnitude less than the r amplitude and oscillates at roughly four times the frequency. Right, 3D plot of sphere trajectory.

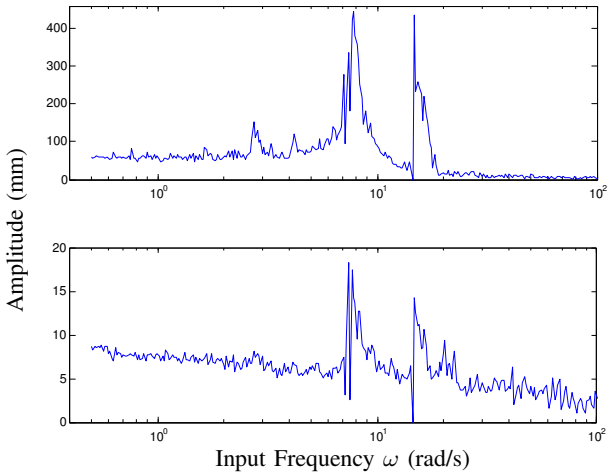

Fig. 7. Input to output Bode analysis with sinusoidal nozzle velocity input and the amplitude of oscillations in r (top) and d (bottom) recorded for a vertically oriented air jet. Note that the oscillations in the axial direction are an order of magnitude larger than in the perpendicular plane.

a desired axial acceleration a of the sphere (provided by an outer PID loop), we compute the desired nozzle velocity of our controller

$$u_{des} = \frac{r + c_1}{\operatorname{sech}^2\left(c_2 \frac{d}{r}\right)} \left(\sqrt{\frac{2a}{\frac{C_d}{m}\rho A}} + \dot{r} \right), \qquad (4)$$

where c_1 and c_2 are constants determined by the free fluid flow. Through this choice of u_{des}, we try to eliminate the nonlinear dependence of $\ddot{r}$ on r and $\dot{r}$. For this approach to work, we need an accurate state estimate. Our cameras only sample at 55 Hz, so we use an extended Kalman filter to estimate the state between each image capture (Fig. 8). Occasionally the camera provides spurious measurements. The Kalman filter provides a robust method to handle sensor noise.

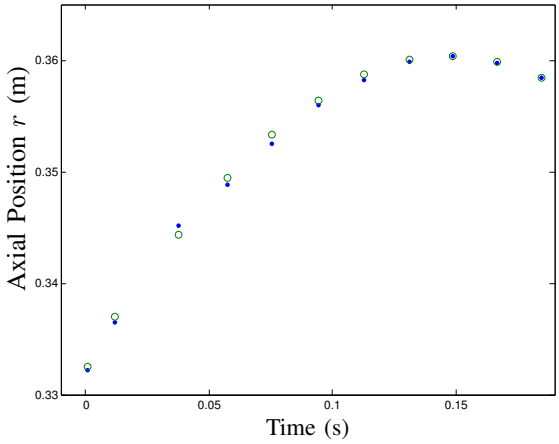

Fig. 8. State tracking using a Kalman filter. The Kalman filter propagates the past state according to the system model, and combines the model with imperfect sensor data to produce a robust state estimate. With these estimates the controller can use the predicted position of the sphere (blue dot) instead of the measured position of the sphere (green circle) from the previous time step.

D. Performance

We compared an open-loop strategy to our control strategy both with and without predictive estimation. Figure 9 shows the results for an open-loop strategy. As expected, the axial position of the sphere is stable when r_{ref} is changed. However, this position differs significantly from the reference position, due to imperfections in our model. In addition, there is significant steady-state oscillation. The second plot shows the results for feedback linearization, and the third the results for feedback linearization with predictive estimation using an extended Kalman filter. The last controller shows the fastest settling time and damping (Table I). Moreover, feedback linearization exhibits these responses over a larger range of r than the PID controller, because of the nonlinear dependence of the control effort on r.

III. APPLICATION TO MANIPULATION TASKS

A. Sphere sorting

When spheres with differing drag to mass ratios are introduced to the same fluid jet, the spheres quickly arrange themselves in order of increasing drag to mass, barring

TABLE I
QUANTITATIVE RESPONSE COMPARISON

Control	Estimation	t_r(s)	t_s(s)	e_{ss}(mm)	M_p%
OL	No	1.31	11.18	-10	34.0
u	No	1.29	> 20	8	35.5
u	Yes	1.21	7.95	-2	32.2
FL	No	1.23	11.99	0	28.7
FL	Yes	1.34	1.21	0	18.5

Comparison of controllers: OL, open loop; u, PID control on nozzle air velocity u; and FL, feedback linearization on force. Here t_r is the 10-90% rise time, t_s is the time to steady state within 25% of the step input, e_{ss} is the mean steady state error and M_p is the maximum overshoot as a percentage of the step size.

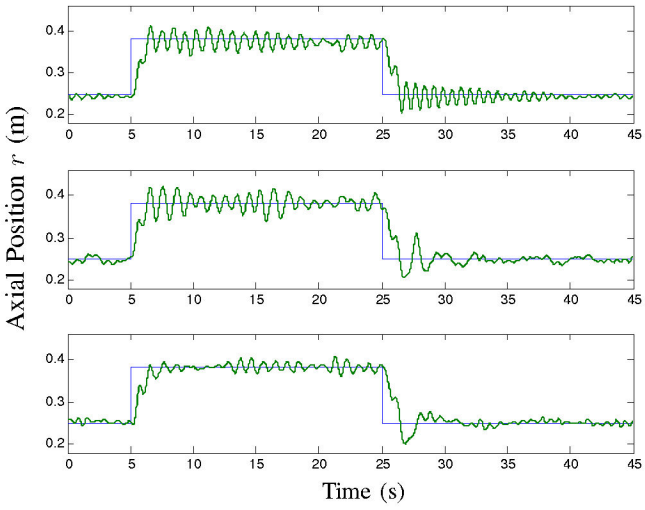

Fig. 9. Three control policies on a 10.8 g 40 mm diameter sphere. Open loop (top) uses no feedback and simply chooses the u that corresponds with r_{eq}. Feedback linearization (center) noticeably reduces the oscillation and removes steady state error. Adding estimation (bottom) using an extended Kalman filter further improves the transient response.

Fig. 10. Sorting multiple spheres simultaneously without camera feedback according to their ratio of drag to mass. The spheres with the highest C_d/m have equilibria higher in the jet. Left: At a flow velocity of 220 m/s, a sphere weighted to 12.9 g has an equilibrium position of 0.34 m and a sphere weighted to 6.7 g has an equilibrium of 0.49 m. By tilting the air jet, the spheres can be sorted into bins. Right: Three spheres of equal shape but differing weights are shown in equilibrium.

collisions. This behavior provides a simple, non-contact, sensorless sorting solution that arranges dissimilar spheres into a column. This column can be tilted to guide spheres into corresponding bins (Fig 10, left). Spheres can also be classified individually, using a constant velocity jet. The height above the nozzle is determined by AC_d/m (Eq. 3), and again, a single actuator can sort spheres into bins according to their drag to mass ratios. Our experiment successfully sorted spheres by gradually increasing u to 220 m/s, tilting the nozzle from vertical to $\theta_1 = \pi/4$ in 4 seconds, then rapidly zeroing the flow rate to drop the spheres into their respective containers. Reliably sorting stacked spheres into bins takes 6 seconds from the time they are introduced to the flow field, with up to three spheres per operation.

Objects can be introduced into the flow field in many ways, including being delivered from another non-contact manipulator, caught, manually placed, or lifted from a perch. Sphere stacking has last in, first out (LIFO) properties. Since objects with lower drag to mass ratios are always below objects with larger ratios, moving a stack of spheres is an exact implementation of the tower of Hanoi problem [11].

For three air jets performing this operation, any stack of n spheres that can be stably stacked in the air flow can be moved in $2^n - 1$ operations. We have stably stacked three spheres of uniform size (Fig 10, right).

Fig. 11. Long-range ballistic positioning of spheres. Note the sphere falling towards the target beaker on the right.

B. *Long-range ballistic positioning*

Because supersonic air velocities can be generated close to the nozzle, a nearby sphere can experience very large accelerations due to drag. Also, at low nozzle velocities, it is possible to hold the sphere in a nearby equilibrium even if the air jet is tilted far from vertical. This allows spheres to be aligned precisely along the barrel of the nozzle. The accuracy can be improved by waiting for oscillations to settle before firing the sphere. By rapidly increasing the flow rate to a calibrated set point, the sphere becomes a predictable projectile (Fig.11). While the impulse on the ball is difficult to calculate analytically for a rapidly accelerating u, the results are reliable and can be determined empirically to calibrate the ballistic system. Our platform repeatedly tilted and held a 40 mm ping pong ball, weighted to 12.6 g, to any angle within $\pi/4$ radians from the vertical axis. By increasing u from 77 to 312 m/s, the sphere landed in a 74 mm target located 1.75 m from the airjet. Ballistic positioning extends the reachable space of the robot from 0.7 m to 2.5 m for a 12.6 g ping pong ball. This could allow teams of similar manipulators to exchange spheres over longer distances.

C. *3D trajectory tracking*

The gimbaled air jet can place a weighted ping pong ball at any point within $\pi/4$ radians from the vertical axis and 0.7 m from the tip of the air nozzle. Two loops of 3 mm wire, each twice the ball diameter in size, and one loop fashioned from 10 mm steel strap that was 1.5 ball diameters wide were used as obstacles to test if the controller could pass spheres though loops without losing stability. The controller was tested with the loops at arbitrary angles to the flow (Fig. 12). Currently, these trajectories are learned from an operator, streamlined, sped up and implemented by the controller. A future controller could visually identify obstacles and plan an appropriate path.

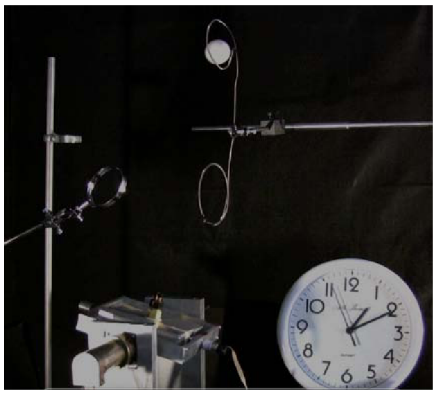

Fig. 12. Navigating a sphere among obstacles in three dimensions. Two wire obstacles are each twice the sphere diameter. The wire strap obstacle is 1.5 sphere diameters in width and is aligned with the flow.

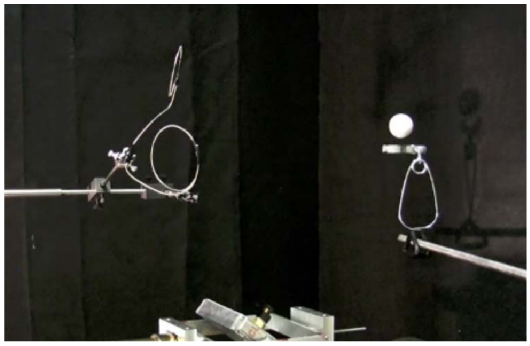

Fig. 13. Placing a sphere on a perch. The air jet lifted the sphere from a perch on the left and navigated through both obstacles before landing on the 35 mm perch shown.

D. *Liftoff and landing*

The ability to do pick-and-place tasks (a mainstay of robotic manipulation) adds a great deal of versatility to the air jet controller (Fig. 13). Two perch styles were tested, one with rubber jaws and the other with metal rails. Each had an open front to reduce flow disturbances, were cupped to passively stabilize the sphere, and were 35 mm wide. Perches were tested at five distances from the nozzle, ranging from 0.1 to 0.6 m with equal success, but were never successful for a starting angle more than $\pi/4$ radians from vertical. The most successful liftoff method was to slowly lower the jet over the sphere until it rocked free, using a u corresponding to $r_{eq} = r$ (Eq. 4). Larger nozzle velocities push the sphere away, while lower nozzle velocities give no response. Landing is similar to the liftoff operation, but requires greater precision.

IV. CONCLUSION

In this paper we showed one way to manipulate spherical objects in three dimensions using air flow. Our mechanism was a 2-DOF gimbaled air jet with a variable flow rate. The position of the spherical object was stable along the axis of the jet at a distance that depended on the flow rate. However, it was still necessary to apply feedback control in order to damp the residual oscillations about the equilibrium point. In particular, we focused on the axial oscillations, which were of lower frequency and higher amplitude than the perpendicular oscillations. Our control strategy was feedback linearization based on a classical fluid dynamics model with state estimates from stereo vision data. We applied our control strategy to enable several different manipulation tasks. These tasks included sorting based on physical characteristics, repeatable long-range ballistic positioning, 3D trajectory tracking among obstacles, and liftoff and landing for pick-and-place operations.

There are many opportunities for future work. First, our control strategy did not completely eliminate axial oscillations. Doing so, and in particular eliminating perpendicular oscillations, is a challenging control problem that may require online learning and adaptation or a more complete model of the system dynamics. Second, we could explore in more detail the manipulation tasks suggested in this paper. For example, 3D trajectory tracking poses an interesting optimal control problem—for our model, maximum lift on a sphere with radius β is achieved when $d \approx \min\left(\arctan(\beta/r), 0.06\right)$. Maintaining this distance creates maximum acceleration, but also increases the probability the system will become unstable. Third, there are many other manipulation tasks that could be done, in particular ones involving the manipulation of more than one object at a time. For example, it is possible to catch thrown objects as well as to "juggle" several objects simultaneously—we are currently working on enabling these tasks. Finally, as a long-term goal, we would like to extend our work to a consideration of non-spherical and non-rigid objects.

REFERENCES

[1] E. H. Brandt, "Levitation in physics," *Science*, vol. 243, no. 4889, pp. 349–355, 1989. [Online]. Available: http://www.sciencemag.org/cgi/content/abstract/243/4889/349

[2] B. Ozcelik and F. Erzincanli, "A non-contact end-effector for the handling of garments," *Robotica*, vol. 20, pp. 447–450, 2002.

[3] ——, "Examination of the movement of a woven fabric in the horizontal direction using a non-contact end-effector," *International Journal of Advanced Manufacturing Technology*, vol. 25, pp. 527–532, 2005.

[4] D. Biegelsen, A. Berlin, P. Cheung, M. Fromherz, and D. Goldberg, "Airjet paper mover," in *SPIE Int. Symposium on Micromachining and Microfabrication*, Sep 2000, pp. 4176–11.

[5] S. Davis, J. Gray, and D. G. Caldwell, "An end effector based on the bernoulli principle for handling sliced fruit and vegetables," *Robotics and Computer-Integrated Manufacturing*, vol. 24, no. 1, pp. 249–257, 2008.

[6] J. N. Reed and S. J. Miles, "High-speed conveyor junction based on an air-jet floatation technique," *Mechatronics*, vol. 14, no. 6, pp. 685–699, 2004.

[7] G. Sharpe, *Fluid Dynamics*. Essex, England: Longman Scientific and Technical, 1994.

[8] W. Niessen, *Combustion and incineration processes: Applications in Environmental Engineering*. CRC Press, 2002.

[9] F. White, *Viscous Fluid Flow*. New York, NY: McGraw Hill, 2006.

[10] T. Kawamura, H. Takami, and K. Kuwahara, "Computation of high reynolds number flow around a circular cylinder with surface roughness," *Fluid Dynamics Research*, vol. 1, no. 2, pp. 145–162, 1986/12.

[11] J. Bode and A. Hinz, "Results and open problems on the tower of hanoi," *30th Southeastern International Conference on Combinatorics, Graph Theory, and Computing*, vol. 139, pp. 112 – 122, 1999.