2009 American Control Conference
Hyatt Regency Riverfront, St. Louis, MO, USA
June 10-12, 2009

On the Controllability of Fixed-Wing Perching

John W. Roberts, Rick Cory and Russ Tedrake

Abstract— **The ability of birds to perch robustly and effectively is a powerful demonstration of the capabilities of nature's control systems. Their apparent robustness to gust disturbances is particularly remarkable because when the airspeed approaches zero just before acquiring a perch, the influence of aerodynamic forces, and therefore potentially the control authority, is severely compromised. In this paper we present a simplified closed-form model for a fixed-wing aircraft which closely agrees with experimental indoor perching data. We then carefully examine the LTV controllability along an optimized perching trajectory for three different actuation scenarios - a glider (no powerplant), a fixed propeller, and a propeller with thrust vectoring. The results reveal that while all three vehicles are LTV controllable along the trajectory, the additional actuators allow the perch to be more easily acquired with less control surface deflections. However, in all three cases, disturbances experienced just before reaching the perch cannot be effectively rejected.**

I. INTRODUCTION

The ability of a bird to land accurately and robustly on a perch represents one of the most striking examples of the high-performance capabilities of nature's aerodynamic control systems—capabilities which are as yet unmatched by man-made controllers. To accomplish perching with the level of performance which birds achieve requires that an aircraft leave the standard flight envelope in which most planes fly, and around which most aircraft controllers are designed. Birds decelerate rapidly to a perch by a fast transient maneuver which exploits the large viscous and pressure drags which occur at extremely high angles-of-attack (AoA). In these high-AoA maneuvers, the sort of control an aircraft possesses over its dynamics is very different from that which it has during standard forward flight. Both this obvious disparity in achieved performance and the growing interest in industry for highly agile aircraft have inspired significant interest in the design of controllers for robust perching.

(Cory & Tedrake, 2008) presents a successful demonstration of a perching glider in a laboratory, where a feedback policy optimized on a coarse state-space discretization was used to execute the perching maneuver in still air conditions. A closely related project at Cornell has examined optimal trajectory generation for a morphing fixed-wing aircraft (Wickenheiser et al., 2005; Wickenheiser & Garcia, 2006; Wickenheiser & Garcia, 2008). Their most recent work focuses on computing optimal perching trajectories in simulation through the solution of a boundary value problem. However, previous work has so far ignored the issue of stabilizing a perching trajectory in the face of aerodynamic disturbances and sensory limitations. As we progress from laboratory to outdoor experiments, stabilization becomes critically important. In this paper, we investigate the controllability and stabilizability of a generated perching trajectory in order to design a feedback control system capable of robustly handling disturbances.

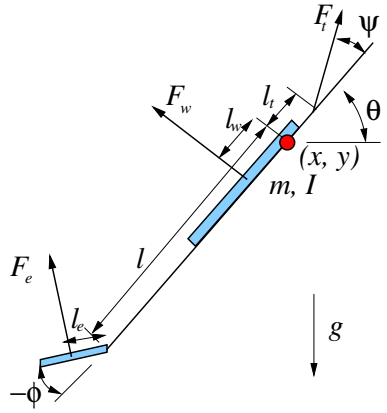

Fig. 1. Simple aircraft model. The blue segments represent the aerodynamic surfaces of the wings and the elevator with surface areas S_w and S_e, respectively, going into the page.

We first present a simplified aircraft model, whose aerodynamics are based upon the work of (Cory & Tedrake, 2008). Using direct collocation, we design nominal control trajectories for this aircraft and examine their controllability and stabilizability, although all forms are very susceptible to disturbances. Our analysis suggests that the ability to generate thrust provides an advantage in practical performance and can improve the robustness of the trajectory. We conclude with a discussion of how these results may be applied to actual perching experiments and of additional techniques which may further improve performance through the use of nonlinear controllers.

II. LIFT AND DRAG MODELS

Our simple aircraft model can be described by a single rigid body subject to gravitational and aerodynamic forces. The aerodynamic forces are functions of the AoA of both the wing and the elevator control surface, both of which are modeled as thin flat-plates (see Figure 1). Thin flat-plate models are in fact a close approximation to the actual wings and control surfaces found on many aerobatic hobby aircraft,

J. W. Roberts is a PhD student in Mechanical Engineering in the Computer Science and Artificial Intelligence Lab at MIT jwr@mit.edu
R. Cory is a PhD student in Electrical Engineering and Computer Science in the Computer Science and Artificial Intelligence Lab at MIT rcory@csail.mit.edu
R. Tedrake is the X-Consortium Associate Professor of Electrical Engineering and Computer Science, and a member of the Computer Science and Artificial Intelligence Lab at MIT russt@mit.edu

978-1-4244-4524-0/09/$25.00 ©2009 AACC

Fig. 2. A image taken with a high-speed camera from one of the perching trajectories used to obtain data on lift and drag coefficients. The target perch can be seen in the left of the image.

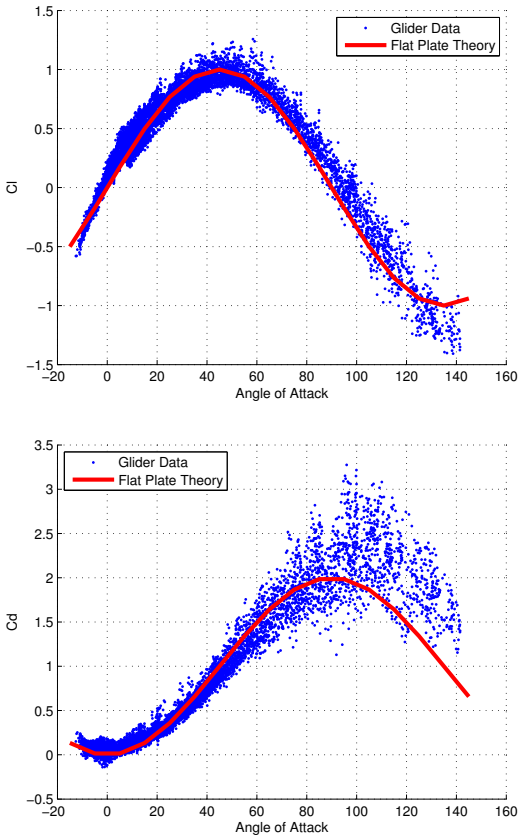

Fig. 3. Scatter plots showing the general trend of the estimated aerodynamic coefficients presented in (Cory & Tedrake, 2008). Shown are lift and drag coefficient data (blue) overlaid with their respective flat plate theory predictions (red).

whose designs have formed the basis of our experimental platforms. Our previous work presented in (Cory & Tedrake, 2008) described a statistical procedure for modeling the aerodynamics of one of our flat-plate experimental gliders (no powerplant). Using tools from nonlinear system identification, lift, drag, and moment coefficients were estimated over a large range of angles-of-attack (including post-stall) using real flight data obtained through the execution of perching trajectories as in Figure 2 (Cory & Tedrake, 2008). Recent work (Hoburg & Tedrake, 2009) shows that the majority of the aerodynamic behavior can be explained using simple trends predicted by flat-plate theory (Tangler & Kocurek, 2005), and thus this minimalist model was used. Figure 3 illustrates the general trend of the estimated coefficients versus their flat-plate theory predictions. Our data agree closely with their respective theory models, which can be written in closed form as:

$$C_L = 2\sin(\alpha)\cos(\alpha) \quad (1)$$

$$C_D = 2\sin^2(\alpha) \quad (2)$$

We suspect that the discrepancy in the drag data at high angles-of-attack are largely due to time-varying vortex shedding effects. For the purposes of this investigation, we use flat-plate theory models which effectively average out the vorticity effects at high angles-of-attack. Also, as the data was taken from a glider, our model does not necessarily capture any of the aerodynamic effects of 'backwash' off a forward-mounted propeller - these effects are thought to be very significant in much of the work on fixed-wing aerobatics(Green & Oh, 2005; How et al., 2008). Future work will focus on investigating higher-fidelity models that include these effects.

III. AIRCRAFT MODEL

First, let us define unit vectors normal to the control surfaces (in the directions of the illustrated force vectors):

$$\mathbf{n}_w = \begin{bmatrix} -s_\theta \\ c_\theta \end{bmatrix}, \quad \mathbf{n}_e = \begin{bmatrix} -s_{\theta+\phi} \\ c_{\theta+\phi} \end{bmatrix},$$

where we introduce the notation $s_\theta = \sin\theta$ and $c_\theta = \cos\theta$ which will be used throughout this section. Also, let us define a unit vector along the primary fuselage axis (tail to nose) of the vehicle (and thus the thrust force):

$$\mathbf{n}_p = \begin{bmatrix} c_\theta \\ s_\theta \end{bmatrix}.$$

Next, solving for the kinematics of the geometric centroid of aerodynamic surfaces (here equivalent to the the mean aerodynamic chord, which is the center of pressure for flat plate theory) for each aerodynamic surface, we have:

$$\mathbf{x}_w = \begin{bmatrix} x - l_w c_\theta \\ y - l_w s_\theta \end{bmatrix}, \quad \mathbf{x}_e = \begin{bmatrix} x - l c_\theta - l_e c_{\theta+\phi} \\ y - l s_\theta - l_e s_{\theta+\phi} \end{bmatrix}, \quad (3)$$

$$\dot{\mathbf{x}}_w = \begin{bmatrix} \dot{x} + l_w\dot{\theta}s_\theta \\ \dot{y} - l_w\dot{\theta}c_\theta \end{bmatrix}, \quad \dot{\mathbf{x}}_e = \begin{bmatrix} \dot{x} + l\dot{\theta}s_\theta + l_e(\dot{\theta}+\dot{\phi})s_{\theta+\phi} \\ \dot{y} - l\dot{\theta}c_\theta - l_e(\dot{\theta}+\dot{\phi})c_{\theta+\phi} \end{bmatrix}. \quad (4)$$

The determination of the aerodynamic forces is of utmost importance to the fidelity of the model. Equations (1) and (2) provide expressions for lift and drag coefficients that model the observed dynamics very well. Using these coefficients, the forces on the wing and elevator can be written in closed form. The force on the wing can be written as components in lift and drag:

$$\mathbf{F}_{wL} = \frac{1}{2}\rho|\dot{\mathbf{x}}_w|^2 C_L S_w, \quad \mathbf{F}_{wD} = \frac{1}{2}\rho|\dot{\mathbf{x}}_w|^2 C_D S_w. \quad (5)$$

where ρ is the density of air, and S_w and S_e are the surface areas of the wing and tail control surfaces, respectively. Using C_L and C_D from Eqns. (1) and (2), these can be written:

$$\mathbf{F}_{wL} = \rho|\dot{\mathbf{x}}_w|^2 \cos(\alpha_w)\sin(\alpha_w)S_w, \quad (6)$$

$$\mathbf{F}_{wD} = \rho|\dot{\mathbf{x}}_w|^2 \sin^2(\alpha_w)S_w. \quad (7)$$

where α_w and α_e are the angles-of-attack of the wing and elevator (computed using the atan2 function in Matlab), respectively. Noting that the lift force is perpendicular to velocity in the positive y direction and the drag force is anti-parallel to the velocity, this can be written (using the normal of the wing) as:

$$\mathbf{F}_w = \rho|\dot{\mathbf{x}}_w|^2 \sin(\alpha_w)S_w\mathbf{n}_w \quad (8)$$

Doing the same for the elevator we can represent the aerodynamic forces as:

$$\alpha_w = \theta - \tan^{-1}\frac{\dot{y}_w}{\dot{x}_w}, \quad \alpha_e = \theta + \phi - \tan^{-1}\frac{\dot{y}_e}{\dot{x}_e} \quad (9)$$

$$\mathbf{F}_w = \rho S_w|\dot{\mathbf{x}}_w|^2 \sin\alpha_w\mathbf{n}_w = f_w\mathbf{n_w}, \quad (10)$$

$$\mathbf{F}_e = \rho S_e|\dot{\mathbf{x}}_e|^2 \sin\alpha_e\mathbf{n}_e = f_e\mathbf{n_e}, \quad (11)$$

The thrust force can then be given simply as:

$$\mathbf{F}_t = f_t\left[c_{\theta+\psi}\ s_{\theta+\psi}\right], \quad (12)$$

where f_t is the signed magnitude of the thrust and ψ is the angle of the thrust vector with respect to the body. Finally, these forces can be integrated into the dynamics as follows:

$$m\ddot{x} = -f_w s_\theta - f_e s_{\theta+\phi} + f_t c_{\theta+\psi} \quad (13)$$

$$m\ddot{z} = f_w c_\theta + f_e c_{\theta+\phi} + f_t s_{\theta+\psi} - mg \quad (14)$$

$$I\ddot{\theta} = -f_w l_w - f_e(l c_\phi + l_e) + f_t l_t s_\psi \quad (15)$$

We can then write the state of the system as $\mathbf{x} = [x\ y\ \theta\ \phi\ \dot{x}\ \dot{y}\ \dot{\theta}\ \dot{\phi}]^T$, and the actuation as $\mathbf{u} = [\ddot{\phi}\ f_t\ \psi]^T$. The parameters for this system were then chosen based upon the properties of the foam glider used for the data collection in §II, and are given below in Table I.

TABLE I
MODEL PARAMETERS

Parameter		Value
m	$[kg]$	0.05
g	$[m/s^2]$	9.81
ρ	$[kg/m^3]$	1.292
S_w	$[m^2]$	0.1
S_e	$[m^2]$	0.025
I	$[kg \cdot m^2]$	$6 \cdot 10^{-3}$
l	$[m]$	0.35
l_w	$[m]$	-0.03
l_e	$[m]$	0.04
l_t	$[m]$	0.05

While this is an extremely simplified model of an aircraft, it has sufficient richness to capture a number of the relevant properties of the system. It allows for stable gliding without thrust, as well as a number of trim conditions for steady level flight, steady climb and so forth. Most critically, however, is the fact that the aerodynamic force models agree well with those seen on the experimental glider, even up to very high angles of attack. As high-AoA maneuvers are a critical aspect of the perching problem to be solved, the fidelity of the model in this regime is of utmost importance.

IV. PROBLEM SPECIFICATION

As previously mentioned, the state of the aircraft consists of seven state variables $\mathbf{x} = [x\ y\ \theta\ \phi\ \dot{x}\ \dot{y}\ \dot{\theta}\ \dot{\phi}]^T$, where $\mathbf{x}$ denotes the state vector, x and y denote the horizontal and vertical distance to the perch, respectively, θ denotes the pitch of the aircraft, and ϕ is the elevator deflection angle. The perching task is described as follows: the aircraft approaches a perch that is 4m away with an initial forward velocity of 6m/s at zero angle-of-attack, (i.e., $\mathbf{x}_0 = [0\ 1\ 0\ 0\ 6\ 0\ 0\ 0]^T$). Given these initial conditions, the goal is to land the aircraft on the perch with a negligible final forward velocity ($\mathbf{x}_d = [4\ \ 0.75\ \ \pi/4\ \ 0\ \ 0\ \ -0.5\ \ -0.5\ \ 0]^T$) in one second using the control inputs: elevator torque, thrust (propulsive) force and thrust direction (i.e., thrust vectoring). Thrust is saturated between -0.03 and 0.1 Newtons, the elevator angle is constrained to be between $-40°$ and $40°$ and ψ was set between $-15°$ and $15°$. A nominal perching trajectory is illustrated in Figure 4.

In the following sections, we divide the control problem into two steps: 1) Generate a feasible open-loop perching trajectory using direct collocation optimization(von Stryk, 1993; Betts, 2001) and 2) Stabilize this nominal perching trajectory by computing a local feedback policy and analyzing the controllability of the closed loop system. The constraints on our trajectory generation differ from those used in Wickenheiser's (Wickenheiser & Garcia, 2008) gradient descent approach in that we allow for some deviations around the desired state (e.g., in pitch velocity).

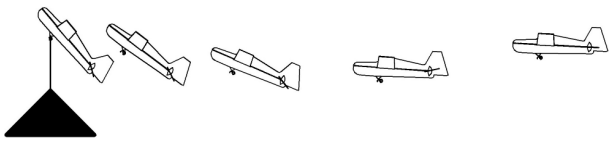

Fig. 4. A nominal perching trajectory. The trajectory of the glider illustrated here is derived in (Cory & Tedrake, 2008) from real flight data of a successful perch.

V. TRAJECTORY GENERATION

A. Trajectory Generation Method

We computed an open-loop perching trajectory for our model using direct collocation: a standard policy optimization method for finite horizon problems. While solutions from this method can be susceptible to local minima, repeated runs from several initial trajectories converged to the same result.

This optimization was performed with a finite horizon of one second, and a quadratic cost function with costs imposed

on actuation and final position error, taking the form:

$$J = \int_{t_i}^{t_f} \left(\mathbf{u}^T(t)R\mathbf{u}(t) + \mathbf{x}^T(t)Q\mathbf{x}(t)\right) dt + \mathbf{x}^T(t_f)Q_f\mathbf{x}(t_f), \tag{16}$$

where $R = R^T > 0$ is cost on actuation, $Q = Q^T \geq 0$ is cost on state error and $Q_f = Q_f^T \geq 0$ is the cost on final state error. For this work all were taken to be diagonal matrices, with R set to $10^{-6} \cdot I$ (and I as the identity matrix), Q as the zero matrix, while Q_f's diagonal was set as $\text{diag}(Q_f) = [100\ 100\ 25\ 0\ 10\ 10\ 0]$ (corresponding to gains of 100 on x and y, 25 on θ and 10 on both $\dot{x}$ and $\dot{y}$).

B. Resulting Trajectory

The optimization was performed for all three variants of the aircraft—a glider, an aircraft with thrust and an aircraft with thrust-vectoring. The open-loop trajectory found for each of these systems was slightly different, although they are qualitatively similar (see Figure 5). Saturations were imposed as hard constraints in the optimization.

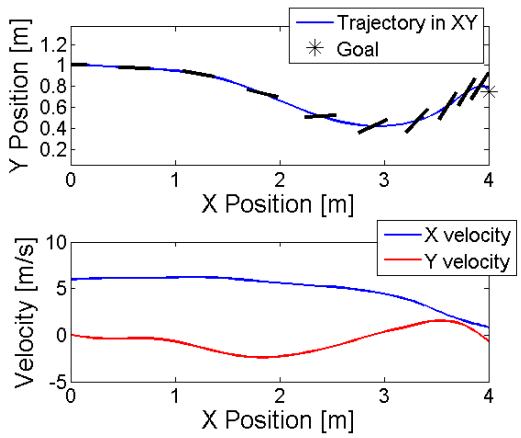

Fig. 5. Nominal trajectory for the airplane with thrust vectoring as found by direct collocation. The black bars show the plane's pitch at that point in the trajectory.

These trajectories were in accord with the optimal trajectories previously proposed by (Wickenheiser & Garcia, 2008) in a qualitative sense. The dip in altitude before the perch that was observed in that work is seen here as well, demonstrating that the proposed simplified model captures many of the relevant dynamical aspects for the perching task. Different initial conditions will result in different nominal trajectories, but qualitatively many of the found trajectories are similar over a reasonable range of initial conditions.

VI. CONTROLLABILITY

The controllability of the aircraft around the generated open-loop trajectory is of great interest due to its bearing on the robustness of the perching trajectory. In this section the linear time-varying (LTV) controllability of the system is investigated, a process which revealed several important properties of the system which must be dealt with by any proposed controller.

A. LTV Analysis

The full non-linear system is converted to an LTV systems by linearizing the nonlinear dynamics around the nominal trajectory. The behavior of the system along the trajectory can then be written:

$$\dot{\mathbf{x}} - \dot{\mathbf{x}}_t = A(t)(\mathbf{x} - \mathbf{x}_t) + B(t)(\mathbf{u} - \mathbf{u}_t) \tag{17}$$

where A and B are the linearizations of the dynamics at time t of the trajectory, while $\mathbf{x}_t$ and $\mathbf{u}_t$ are the nominal trajectory state and action vectors, respectively.

To investigate the controllability of this LTV system, we evaluate the rank of the controllability Gramian $G(t_0, t_f)$, which, for LTV systems is a function of both the initial and final time. It is a well-known result that for LTV systems the cost can be written $J = \mathbf{x}^T P^{-1}(t)\mathbf{x}$, and that if the Gramian is full rank, the system is controllable. A common method to compute $G(t_0, t_f)$ is by integrating $P(t)$ backwards in time from the final time t_f to the initial time t_i, according to the equation:

$$\dot{P}(t) = A(t)P(t) + P(t)A^T(t) - B(t)R^{-1}(t)B^T(t) \tag{18}$$

with $P(t_f)$ equaling the zero matrix as the boundary value (Lewis, 1992). $P(t)$ being full rank then indicates that errors in the trajectory following at time t can be rejected completely (with unbounded control actions), and therefore the system is controllable.

However, due to the high-speeds of certain modes of the system, an evaluation of rank is difficult to do numerically. These difficulties can be overcome by periodically saturating the singular values of the matrix P. This was accomplished by performing the singular value decomposition of $P = U\Sigma V^T$ every 0.1 seconds, saturating all singular values to be less than or equal to 1, then replacing P with $P_{sat} = U\Sigma_{sat}V^T$. After this process the integration can continue for another 0.1 seconds. The span of this saturated matrix will be the same as that of the original matrix, thus the saturation will not change the rank. This periodic saturation resolves the numerical issues resulting from the disparity in mode speeds, and the controllability of the system can now be evaluated numerically.

This analysis indicated that the system is controllable over virtually the entire trajectory, with or without the use of thrust. However, achieving the perch may require very strong responses from the actuators, and as the physical actuators will saturate relatively quickly, the actual envelope of control around the nominal trajectory may be very small. To achieve robustness, the trajectory should be controllable with as little actuator effort as possible, and in this metric the glider and airplanes with thrust perform very differently. This difference is analyzed in detail in §VI-C.

B. Stabilizing Controller

Due to the fact that controllability exists (as seen in the previous section), there is the possibility of developing a simple linear controller to stabilize the trajectory. For a system with a quadratic cost on state and action as given in

Eqn. (16), LTV analysis can give an optimal LTV controller for the system. However, to avoid the numerical difficulties associated with computing P and to allow some acceptable deviations from the desired final state, $S = P^{-1}$ was used to derive the LTV controller. This was obtained by integrating the Riccati equation backwards in time according to:

$$-\dot{S}(t) = A^T(t)S(t)+S(t)A(t)-S(t)B(t)R^{-1}(t)B^T(t)S(t), \tag{19}$$

using a large but finite final cost. S can then be used to derive the optimal LTV controller of the form (Kwakernaak & Sivan, 1972):

$$\mathbf{u}(t) = -F(t)(\mathbf{x}(t) - \mathbf{x}_t(t)), \tag{20}$$

$$F(t) = R^{-1}(t)B^T(t)S(t), \tag{21}$$

where $\mathbf{x}_t(t)$ is the nominal trajectory of the system. The terminal value $S(t_f)$ for this integration is then a final cost equivalent to the Q_f of Eqn. (16). While the cost function structure to develop the controller was the same as that used in §V for trajectory generation, the gains were changed. For this LTV controller, R was set to a diagonal matrix with entries $[.1\ 20\ 5]$, Q as $10I$ (I again the identity matrix), and Q_f the same as in Section V-A.

C. Advantages of Thrust

The inclusion of thrust in this model (distinguishing it from the glider system upon which the aerodynamic data is based) provides the critical ability of being able to add energy to the system. This allows the plane to respond to larger perturbations than would otherwise be possible. A glider can respond to small perturbations by achieving the final state via a path with less losses than the nominal, but in the case of large perturbations there can simply be insufficient energy to reach the final state. Furthermore, even in the case of small perturbations from the nominal trajectory, thrust provides a number of advantages, including smaller required actuator responses (and thus less risk of saturation) and faster convergence to the nominal trajectory (in practice reducing the errors and thus allowing a closer approach to the perch).

A clear indication of the advantages of thrust can be seen in Figure 6, where the value of the greatest eigenvalue of S is plotted as a function of time along the trajectory. This eigenvalue is representative of the cost of a perturbation in the most expensive direction. A high cost represents a combination of larger actuator exertions and greater error at time $t = t_f$ (as the terminal cost is finite, a final error is allowed). The figure shows that the addition of thrust results in less-expensive responses to perturbations over much of the trajectory, with thrust-vectoring providing still greater performance. This is indicative of the fact that a thrust-capable plane can respond well even near the perch, while the glider finds it very difficult to exert control as its airspeed drops and the remaining time decreases.

This distinction is further clarified in Figure 7, in which the performance of the systems are examined using the optimal LTV controllers stabilizing the trajectory optimized for the glider, operating on the full nonlinear dynamics with actuator saturations. The airplanes which are capable of producing thrust respond to disturbances in their initial velocity much better than the glider, particularly as the initial velocity is decreased (thrust and thrust vectoring are equivalent around the glider trajectory due to the linearization about zero thrust). Table II shows what this means in more practical terms, comparing the errors in position and velocity at the end of the trajectory for the glider and thrust capable plane.

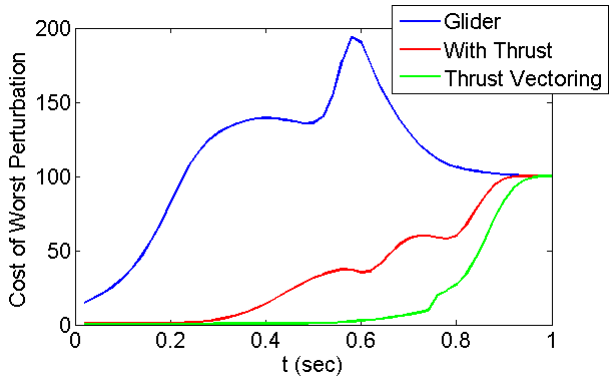

Fig. 6. The largest eigenvalue of S plotted against time t along the trajectory. The cost of a perturbation in the most expensive direction quickly falls for the thrust-vectoring plane, followed shortly thereafter by the thrust-capable plane. Additional actuation clearly increases the amount of the trajectory over which disturbances can be effectively and cheaply rejected, however all three do poorly during the last 0.1 seconds of the trajectory.

TABLE II
COMPARISON OF MAXIMUM ERROR AT FINAL STATE FOR INITIAL VELOCITIES WITHIN 1 m/s OF NOMINAL

Error Type		Glider	Thrust
Position	[m]	0.4306	0.3339
Velocity	[m/s]	0.4949	0.2806
Pitch	[rad]	0.5330	0.4472

The picture that emerges is that while all three systems are controllable for the full length of the trajectory, additional actuators offer what is in some sense greater control authority, as they are able to respond to disturbances more quickly, more easily, and with lower efforts on the individual actuators. The addition of these actuators have the practical effect demonstrated in Table II of making the goal of robust perching much more attainable.

VII. DISCUSSION

The ability of a plane capable of generating thrust to execute a highly successful perching trajectory, and to stabilize this trajectory over a reasonable range of initial conditions, is a significant step forward in the development of man-made perching controllers. Flyers in nature are able to perch in a wide range of conditions, dealing with steady wind, intermittent gusts and irregular shaped perches, all while maintaining a high degree of robustness. Stabilizing a nominal trajectory, even with a simple linear controller, is the first step toward achieving this goal.

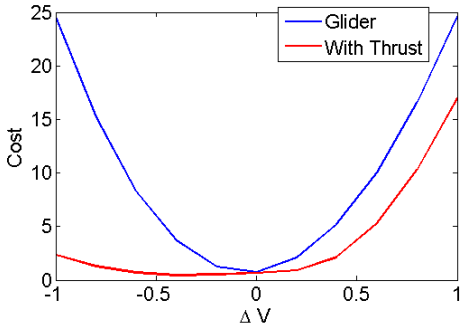

Fig. 7. Cost for initial velocities of $V_{nominal} + \Delta V_{init}$ executing on the glider trajectory. As is clear from the figure, the addition of thrust improves robustness, particularly for low initial velocities. Thrust vectoring in this context is equivalent to thrust when linearized around the glider trajectory.

Eventually, however, richer controllers will be required to truly match the abilities of birds and insects, particularly as control for different parameters such as larger inertias or smaller aerodynamic surfaces can become much more challenging. Nonlinear control laws could be capable of perching over a much wider range of initial conditions while rejecting more significant disturbances during the trajectory. A more advanced controller could identify when a trajectory has become impossible or inefficient to stabilize, and follow a more appropriate trajectory as a result. This ability would make gliders more competitive, but would also make planes capable of thrust even more robust. How to obtain these nonlinear controllers remains an open question, particularly as even the heavily simplified dynamics of the plane model presented here are rather complex. The authors believe that machine learning techniques, both value methods such a value iteration and policy gradient methods such as weight-perturbation, offer the possibility of obtaining these high-performance controllers.

VIII. CONCLUSION

If one desires to build a perching plane, this work suggests that thrust aids in achieving robustness to reasonably large disturbances (e.g., 1m/s error on a nominal speed of 6m/s). Furthermore, while thrust vectoring allows for superior performance, in practice the additional complexity may not be warranted for the marginal increase in performance. Perhaps most important, however, is the fact that none of the systems displayed the level of robustness seen in nature (all three effectively losing authority during the last 0.1 seconds of the trajectory), suggesting richer actuation and control may be required to match the performance of the humblest bird.

IX. ACKNOWLEDGEMENTS

The authors would like to thank Dr. Alexander Megretski for fruitful discussions related to this work. The work was supported by the MIT Lincoln Laboratory Advanced Concepts Committee, by the National Science Foundation graduate fellowship program, and the Microsoft Research New Faculty Fellowship program.

REFERENCES

Bayraktar, S., & Feron, E. (2007). Experiments with small helicopter automated landings at unusual attitudes. *arXiv*.

Betts, J. T. (2001). *Practical methods for optimal control using nonlinear programming*. SIAM Advances in Design and Control. Society for Industrial and Applied Mathematics.

Blauwe, H. D., Bayraktar, S., Feron, E., & Lokumcu, F. (2007). Flight modeling and experimental autonomous hover control of a fixed wing mini-uav at high angle of attack. *AIAA Guidance, Navigation and Control Conference and Exhibit*.

Cory, R., & Tedrake, R. (2007). On the controllability of agile fixed-wing flight. *Proceedings of the 2007 Symposium on Flying Insects and Robots (FIR)*.

Cory, R., & Tedrake, R. (2008). Experiments in fixed-wing UAV perching. *Proceedings of the AIAA Guidance, Navigation, and Control Conference*.

Frank, A., McGrew, J. S., Valenti, M., Levine, D., & How, J. (2007). Hover, transition, and level flight control design for a single-propeller indoor airplane. *AIAA Guidance, Navigation and Control Conference and Exhibit*.

Green, W. E., & Oh, P. Y. (2005). A mav that flies like an airplane and hovers like a helicopter. *Proceedings of the IEEE/ASME International Conference on Advanced Intelligent Mechatronics*.

Green, W. E., & Oh, P. Y. (2006). A fixed-wing aircraft for hovering in caves, tunnels, and buildings. .

Hoburg, W., & Tedrake, R. (2009). System identification of post stall aerodynamics for uav perching. *AIAA Infotech@Aerospace Conference, Seattle, WA*.

How, J. P., Teo, J., & Michini, B. (2008). Adaptive flight control experiments using raven. *Proceedings of the 14th Yale Workshop on Adaptive and Learning Systems*.

Kwakernaak, H., & Sivan, R. (1972). *Linear optimal control systems*. John Wiley & Sons, Inc.

Lewis, F. L. (1992). *Applied optimal control and estimation*. Digital Signal Processing Series. Prentice Hall and Texas Instruments.

Tangler, J., & Kocurek, J. D. (2005). Wind turbine post-stall airfoil performance characteristics guidelines for blade-element momentum methods. *43rd AIAA Aerospace Sciences Meeting and Exhibit*.

von Stryk, O. (1993). Numerical solution of optimal control problems by direct collocation. *Optimal Control, (International Series in Numerical Mathematics 111)* (pp. 129–143).

Wickenheiser, A., Garcia, E., & Waszak, M. (2005). Longitudinal dynamics of a perching aircraft concept. *Proc. SPIE - Int. Soc. Opt. Eng. (USA)*, *5764*, 192 – 202.

Wickenheiser, A. M., & Garcia, E. (2006). Longitudinal dynamics of a perching aircraft. *Journal of Aircraft*, *43*, 1386–1392.

Wickenheiser, A. M., & Garcia, E. (2008). Optimization of perching maneuvers through vehicle morphing. *Journal of Guidance, Control, and Dynamics*, *31*, 815–824.