

NONLINEAR CONTROL OF MECHANICAL SYSTEMS IN THE PRESENCE OF MAGNITUDE AND RATE SATURATIONS

RICHARD M. MURRAY
Mechanical Engineering 104-44
California Institute of Technology
Pasadena, CA 91125

ONR Young Investigator Program

29 September 1995

Abstract

This project is aimed at developing systematic techniques for control of mechanical systems in the presence of magnitude and rate saturations, with particular emphasis on problems arising in the context of high performance aircraft. Magnitude and rate saturations are a major source of nonlinearity in all flight control systems and are a fundamental mechanism of instability in both automated and piloted flight. Recent theoretical developments in nonlinear control theory as well as increasing computational power in offline and online computation are enabling the use of more powerful techniques for control of these systems. The proposed research builds on an established base of work in nonlinear control of mechanical systems and stabilization of strongly nonlinear systems to explore new approaches to this problem. In addition to developing theoretical tools for analysis of flight control systems with saturations, experimental validation of the techniques will be carried out using a flight control experiment at Caltech that exhibits many of the essential features of aircraft systems while remaining simple enough to allow meaningful testing of fundamental feedback mechanisms. Industrial participation with Honeywell Technology Center and McDonnell Douglas Corporation provides a direct path for successful techniques into industry and provides feedback mechanisms to insure the applicability of the proposed work.

CONTENTS

1. Introduction	1
2. Recent Advances in Nonlinear Control	2
2.1. Nonlinear control of systems with input saturations	2
2.2. Two degree of freedom design	3
2.3. Lagrangian control systems and motion control	5
3. Proposed Research	7
3.1. Real-time feedforward in the presence of saturation	7
3.2. Local feedback control in the presence of saturation	9
3.3. Experimental validation on a flight control experiment	10
3.4. Other research directions	12
4. Facilities Available	13
4.1. Caltech ducted fan	13
4.2. Computational facilities	13
4.3. Merrill Wind Tunnel	14
5. Summary and Implementation Plan	14
5.1. Summary	14
5.2. Yearly goals	14
5.3. Interaction with industry	15
References	15
Biographical Information	18
Budget	22
Appendix A. Other Funded Projects	23
Appendix B. Other Pending Proposals	23
Appendix C. Caltech Statement of Support	24
Appendix D. Coordination with Other Projects	25

Nonlinear Control of Mechanical Systems in the Presence of Magnitude and Rate Saturations

Richard M. Murray
California Institute of Technology

1. INTRODUCTION

Linear control theory has become a highly developed and practical method of controlling dynamical systems. Over the past 30 years, techniques have emerged which allow the practicing engineer to quickly and efficiently generate linear controllers for a given system. These techniques include both analytic and software tools for designing controllers given a linear model of the plant and a structured description of plant uncertainties, external disturbances, and desired performance.

The fundamental assumption in linear control theory is that the dynamical system to be controlled can be approximately modeled as a linear system. This assumption is valid for a large class of systems, including many nonlinear systems operating in a neighborhood of an equilibrium point. Applications of linear control theory are widespread: feedback control of chemical processing systems, jet aircraft, and electro-mechanical devices have all benefited from the use of linear control tools.

For many systems, however, the fundamental assumption of a valid linear control model does not hold. High-performance jet aircraft exhibit strongly nonlinear dynamics due to the basic rigid body and aerodynamic forces active on the aircraft as well as nonlinearities in the systems which are used to actuate control surfaces and generate thrust. One of the most significant sources of nonlinearities is actuator saturation, which occurs in all modern control systems but is widely ignored by the existing analysis and synthesis tools. At present, *there is no systematic means of analyzing and designing nonlinear control systems in the presence of magnitude and rate saturations.*

Actuator saturation has a significant effect on the overall stability of aircraft. The recent YF-22 crash (April 1992) has been blamed on a pilot-induced oscillation (PIO) caused in part by time-delay effects introduced by rate saturation of control surfaces [8, 20]. As the complexity and performance of flight systems increase, stronger theoretical understanding is required to avoid such situations and guarantee performance of the system in the face of noise and unmodeled dynamics. These difficulties are already apparent in modern combat aircraft, which have multiple control surfaces that saturate at different magnitudes and rates and must be operated in a coordinated fashion in order achieve performance objectives over a large flight envelope. Without a strong basis for understanding the fundamental limitations and features of such systems, it will be difficult to exploit the full potential offered by active control of these systems.

Recent techniques in nonlinear control theory have begun to tackle some of the difficult issues involved in analysis of systems with input constraints. Nonlinear techniques by Teel, Sussmann, Sontag, and others have begun to show how nonlinear control laws can be used to achieve semi-global stability of linear and nonlinear systems in the presence of saturation. Recent advances in control of Lagrangian systems are providing new methods of utilizing the strong nonlinear structure which is present in mechanical systems and particularly motion control systems. These results are presented in more detail in Section 2.

We propose to tackle the problem of magnitude and rate saturations by applying and extending some of these recent techniques and developing new methods tuned for flight control systems. There are many special properties of such systems which can be exploited in the search for practical techniques for nonlinear control in the presence of saturation. In addition to theoretical work on some of the fundamental issues involved in this problem, we also plan to carry out experimental work on a small, vectored-thrust, flight control experiment at Caltech. This research will have direct application to vectored thrust aircraft, such as the F22, as well as more general air vehicles,

such as ASTOVL aircraft and high performance missiles, which make use of advanced controls to achieve aggressive maneuvering and operation near the limits of the flight envelope.

The broad goals of this research are as follows:

1. To develop systematic techniques for analyzing and designing nonlinear control laws for mechanical systems with magnitude and rate constraints.
2. To develop synthesis tools which allow the use of local linear designs combined with global nonlinear methods to simultaneously achieve local robust performance and global robust stability.
3. To implement and test controllers on representative experiments which replicate important features of full-scale systems and to use these experiments to validate theoretical results and motivate new research directions.
4. To transition successful techniques into military and industrial applications by working with government and corporate partners to further develop methods initiated under this proposal, with the intent of possible implementation on full-scale aircraft.

2. RECENT ADVANCES IN NONLINEAR CONTROL

We begin by reviewing some of the recent approaches in nonlinear control which are directly relevant to this proposal. This section is not intended to be a complete review of the recent literature in nonlinear control theory, but rather a selected review of a few approaches which form the starting point for the work to be pursued under this proposal.

2.1. Nonlinear control of systems with input saturations. At the practical level, saturation in most systems is handled in an *ad hoc* fashion. Gains are chosen and artificial saturations are inserted such that the system performs well in simulations and experimental tests. Rate saturations are sometimes modeled as equivalent time-delays to allow the use of linear control theory [8]. While these techniques work for systems of low dimension and reasonably straightforward dynamics, as systems get more complicated the difficulties in this *ad hoc* approach become more noticeable (as in the crash of the YF22). Further demands for increased performance and agility will exacerbate this problem.

Many techniques are available in the linear literature for incorporating actuator saturations into the design process. One example is the use of l_1 analysis and synthesis techniques, which allow specification of the maximum output response as a function of the maximum size of noise and disturbances [17, 16]. Other techniques include the use of convex optimization to design controllers with a variety of input and performance constraints [6]. A major limitation of these approaches is that they only work for linear systems and they generate linear controllers. As a result, the designs can be very conservative (since the gain must be small enough to tolerate the worst case scenario) and it may be difficult to achieve high performance.

Several new nonlinear tools have been introduced in the last three years for analyzing and controlling linear and nonlinear systems with saturation. One of the fundamental techniques is based on the thesis work of Teel [55, 56], who showed how to stabilize a chain of integrators using nested saturation functions. This result is significant since it is known that it is not possible to stabilize a chain of three or more integrators using a linear control law followed by a saturation function. Thus even simple linear systems with simple saturations can give rise to difficult nonlinear problems. Teel's approach generates nonlinear controllers which are locally linear, but become nonlinear as the inputs grow toward the saturation limits. We briefly review these results to indicate some of the basic ideas.

The difficulty with using linear control laws in the presence of saturation is that directional information about the state error is lost. Consider, a simple linear system of the form

$$\dot{x} = Ax + Bu \quad x \in \mathbb{R}^n, u \in \mathbb{R},$$

where x is the state, u is the (single) input, and (A, B) are the state space representation of the dynamics. A saturated, state-space linear controller for this system has the form

$$u = \text{sat}(k_1 x_1 + \cdots + k_n x_n),$$

where $k_i \in \mathbb{R}$ are the feedback gains and $\text{sat} : \mathbb{R} \rightarrow \mathbb{R}$ is a continuous saturation function that is the identity function on some interval containing the origin and constant outside that interval. The essential difficulty with this control action is that it does not distinguish between large values of the state x_1 and those of x_n or the other states. In either case, the control law will eventually saturate and this can lead to instability.

In Teel's work, and later generalizations by Sussman, Sontag, and Yang [52], the saturated linear controller is replaced by a *nested saturation controller* of the form

$$u = \text{sat}(k_1 x_1 + \sigma_1(k_2 x_2 + \cdots + \sigma_{n-1}(k_n x_n))),$$

where the functions σ_i are artificially imposed saturations. With appropriate choice of σ_i and k_i , it is possible to show that the system is semi-globally stable, meaning that given the size of the initial conditions, one can find σ_i , k_i to stabilize all such initial conditions. It is also possible to use sums of saturations functions, as shown in [52]. Some similar results are also available for nonlinear systems, namely those in so-called feedforward form [56]:

$$\begin{aligned}\dot{x}_1 &= f(x_1, x_2, \dots, x_n, u) \\ \dot{x}_2 &= f(x_2, \dots, x_n, u) \\ &\vdots \\ \dot{x}_n &= f(x_n, u)\end{aligned}$$

These results show how a nonlinear controller can be used to modify a linear feedback law to achieve practical stability in the face of saturation.

Despite the success of this approach to stabilization of systems with saturation, there are a number of limitations with this technique and other nonlinear techniques that have appeared in the literature:

1. Existing techniques only apply to linear systems or nonlinear systems in special (non-generic) normal forms. Analytic expressions for the dynamics are required in order to synthesize controllers.
2. Most approaches in the literature focus on the problem of set point stabilization in the presence of saturation, and not aggressive trajectory tracking away from equilibrium points and potentially near the operating envelope of the system.
3. For nonlinear systems, only static state-feedback control laws are considered. The use of local (dynamic) control laws, such as those generated by LQG or H_∞ techniques, has not been developed in a systematic fashion. This leads to weak performance away from saturation.
4. No evaluation of the control laws on experimental hardware has been performed, so it is not known how difficult the controllers are to implement in the presence of noise, uncertainty, and practical concerns.

One of the motivations of the current proposal is to begin to address these difficulties in a systematic fashion.

2.2. Two degree of freedom design. A second technique which is becoming more widespread is the use of two degree of freedom nonlinear control techniques. A large class of industrial and military control problems consist of planning and following a trajectory in the presence of noise and uncertainty. Examples range from unmanned and remotely piloted airplanes and submarines performing surveillance and inspection tasks, to mobile robots moving on factory floors, to multi-fingered robot hands performing inspection and manipulation tasks inside the human body under




FIGURE 1. Two degree of freedom controller design.

the control of a surgeon. All of these systems are highly nonlinear and demand accurate performance.

Modern geometric approaches to nonlinear control often rely on the use of feedback transformations to convert a system into a simplified form which can then be controlled with relatively standard techniques (such as linear feedback). While this approach very effectively exploits the nonlinear nature of the system, it often does so by “converting” the nonlinear system into a linear one. This can be disadvantageous if one is concerned with disturbance rejection and other performance specifications since the nonlinear transformations typically do not preserve many important properties of the system.

One way around this limitation is to make use of the notion of *two degree of freedom* controller design for nonlinear plants. Two degree of freedom controller design is a standard technique in linear control theory that separates a controller into a feedforward compensator and a feedback compensator. The feedforward compensator generates the nominal input required to track a given reference trajectory. The feedback compensator corrects for errors between the desired and actual trajectories. This is shown schematically in Figure 1.

Many modern nonlinear control methodologies can be viewed as synthesizing controllers which fall into this general framework. For example, traditional nonlinear trajectory tracking approaches, such as feedback linearization [25, 27] and nonlinear output regulation [26], are easily viewed as a feedforward piece and a feedback piece. Indeed, when the tracking error is small, the primary difference between the methods is the form of the error correction term: output regulation uses the linearization of the system about a single equilibrium point; feedback linearization uses a linear control law in an appropriate set of coordinates. It is important to note that these approaches rely on the availability of state feedback in order to generate feedforward commands. Thus, they are not traditional open loop feedforward controllers and this difference is crucial to their operation.

This two step approach can be carried one step further by completely decoupling the trajectory generation and asymptotic tracking problems. Given a desired output trajectory, we first construct a state space trajectory x_d and a nominal input u_d that satisfy the equations of motion. The error system can then be written as a time-varying control system in terms of the error, $e = x - x_d$. Under the assumption that that tracking error remains small, we can linearize this time-varying system about $e = 0$ and stabilize the $e = 0$ state. One method of doing this is to solve the linear quadratic optimal control problem to obtain the optimal (time-varying) feedback gain for the path. More advanced techniques include the use of linear-time varying robust synthesis (see, for example,

Shamma [51] for recent results and a survey of the literature) and the use of linear parameter varying synthesis developed by Packard [48] and others.

The use of two degree of freedom techniques has been studied in a variety of applications. It is a relatively standard approach in classical robotics [43] and has also seen recent application in flexible robot systems [1, 19, 49]. Applications to flight control include the work of Meyer et al. [38] and Martin et al. [32]. General theoretical results have been explored by Paden and Chen [13, 49] as well as ourselves [41, 57, 59]. It is also implicit in the output regulation problem studies by Isidori and Byrnes [26]. Finally, some initial experimental results to a simple flight control system have been reported in [58].

For nonlinear systems, there are several advantages to separating the real-time trajectory generation and local trajectory stabilization portions of the controller:

1. The feedforward controller can make use of the geometric nature of the nominal system for generating feasible trajectories and appropriate nominal forces for following aggressive trajectories.
2. Local (scheduled) designs can be used to achieve good local performance given the nominal state space trajectory and the appropriate nominal forces.
3. Strong nonlinearities such as input saturation can initially be treated separately from the problem of robust stabilization and disturbance rejection along a reference trajectory.

Of course, eventually one must understand the interaction between the feedforward and feedback designs and the individual control designs must exploit knowledge of the behavior of the complementary controller.

2.3. Lagrangian control systems and motion control. The likelihood of building a general nonlinear theory which can be used for all nonlinear dynamical control systems is extremely small due to the diverse nature of behaviors that can be present in nonlinear systems. In this work, we plan to concentrate on one particularly well-structured class of nonlinear control systems: second order systems whose unforced motion is described by Lagrangian mechanics. The special features of this class of systems must be exploited if we wish to design good nonlinear controllers which can extend and eventually outperform existing design techniques for flight control.

There are several reasons why this class of systems is a good candidate for new results in nonlinear control. On the practical end, mechanical systems are often quite well identified and very accurate models exist for specific systems, such as F18-HARV [9, 21]. Furthermore, instrumentation of mechanical systems is relatively easy to achieve and hence modern nonlinear techniques (which often rely on full state feedback) can be readily applied. On the theoretical end, new results in Lagrangian mechanics, motivated in large part by problems in control theory, have generated new insights into the control problem and have produced several new techniques. These techniques exploit the symmetries, constraints, and energy properties of Lagrangian systems to understand the underlying behavior of the system.

One of the main classes of Lagrangian control systems is *motion control systems*, where the control task is to regulate the position and orientation of a rigid body, or set of rigid bodies, in Euclidean space. The dynamics for these systems have a tremendous amount of structure and new techniques are being developed to make use of this structure. The dynamics of motion control systems on the Euclidean group can all be written in terms of three fundamental equations: a *group equation*, which describes the motion of the position and orientation of the system; a *momentum equation*, which describes the evolution of the momentum (conserved in the absence of drag and applied forces); and the *base equations* which describe the dynamics of the internal shape of the system. A complete description of these equations for systems with symmetries and/or constraints is derived in [2] and outlined in a controls context in [41]. We briefly survey those results here.

Consider a mechanical system whose configuration consists of a set of “position” variables $g \in SE(3)$ which represents the position and orientation of the system and a set of “shape” variables $r \in$

M which represents the internal shape of the system (for example, the angles of the control surfaces). Here $SE(3)$ is the special Euclidean group and M is a smooth, finite dimensional manifold. The dynamics are described by a system of equations having the form of a reconstruction equation for a group element g , an equation for the momentum p (no longer conserved in when external forces and/or constraints are present), and the equations of motion for the variables r . In terms of these variables, the equations of motion have the functional form

$$\begin{aligned}\dot{g} &= g(-A(r)\dot{r} + I^{-1}(r)p) \\ \dot{p} &= \alpha(r, \dot{r}, p) + F_p \\ M(r)\ddot{r} + C(r, \dot{r})\dot{r} + N(r, \dot{r}, p) &= F_r\end{aligned}$$

The first equation describes the motion in the group variables as the flow of a left invariant vector field determined by the internal shape r , the velocity $\dot{r}$, as well as the generalized momentum p . The momentum equation describes the evolution of p and can be shown to be bilinear in $(\dot{r}, p)$. Finally, the bottom (second-order) equation describes the motion of the variables r which describe the configuration up to a symmetry (*i.e.*, the shape). The term $M(r)$ is the mass matrix of the system, C is the Coriolis term which is quadratic in $\dot{r}$ and N is quadratic in $\dot{r}$ and $\dot{p}$. The variable F_r represents the potential forces and external forces applied to the system that only affect the shape variables while F_p represents the forces which affect the momentum directly. Note that in the absence of external forces, the evolution of the momentum p and the shape r *decouple* from the group variables.

The utility of this form of the equations is that it separates the dynamics into pieces consistent with the overall geometry of the system. This can be quite powerful in the context of control theory. For example, recent results by Kelly, Ostrowski, and Murray in controllability for locomotion systems (for which the base dynamics are fully actuated) have made explicit use of this structure [29, 47]. Other work which has made use of the special structure of mechanical systems include the work of Lewis and Murray on configuration controllability [31], the work of Bloch and Marsden on satellite stabilization [3], the work of Leonard on stability of underwater vehicles [30], and the work of Bullo and Murray on control of fully actuated systems on Lie groups [10, 12], to name a few.

This form of the equations is also very relevant to flight control problems. Indeed, the standard equations of motion for an aircraft (see, for example, Etkin [22]) are a local coordinate version of these equations with lift, drag, and pitching moment considered as external forces. However, traditionally work in flight control has not made explicit use of the geometry of the Euclidean group $SE(3)$, but rather relied on the use of local coordinates (such as Euler angles), which give complicated coordinate expressions and destroy some of the natural structure present in motion control problems on $SE(3)$. While this is appropriate for local control of low-acceleration maneuvers, the use of such local constructions for aggressive maneuvers can be problematic and often leads to controllers which are extremely complicated and difficult to interpret. By making use of the global geometry of the problem, it is possible to design nonlinear controllers which are easily understood in terms of the problem data and the desired performance.

One example of this type of global control law for mechanical systems is the “proportional derivative” control laws defined in [10, 11] for fully actuated mechanical systems on Lie groups and Riemannian manifolds. For that class of systems, there is a naturally defined set of feasible trajectories for connecting two configurations: one parameter subgroups in the case of Lie groups and geodesics in the case of Riemannian manifolds. Furthermore, there is a naturally defined transport map which allows desired trajectories to be mapped from desired configurations to current configurations, allowing a meaningful computation of the feedforward forces and position/velocity errors that should be used when the system has tracking errors. These ideas have not yet been extended to underactuated mechanical systems (such as aircraft), but are indicative of the types of results which are available by exploiting some of the structure present in motion control problems for mechanical systems.

3. PROPOSED RESEARCH

There are many open problems in control of mechanical systems with magnitude and rate saturations. In this section we indicate some specific problems which we plan to pursue.

3.1. Real-time feedforward in the presence of saturation. Accurate and aggressive trajectory tracking requires good calculation of the nominal trajectory to be regulated as well as the dynamic trim conditions which provide the feedforward forces. It is important to note that in the context of the two degree of freedom design paradigm discussed above, the feedforward portion of the controller is *not* open loop. Rather, it is free to use current information about the state as well as a description of the desired trajectory up to the present time. The entire desired trajectory is typically not available in flight control examples (unlike the situation in robotics), so it is important that the feedforward portion of the controller operate causally, using only current state and past information about the desired trajectory.

We have already begun to investigate the role of feedforward control in the context of a small flight control experiment at Caltech [57, 58]. The experiment itself is described in more detail below. However, these results confirm that good feedforward forces are essential for accurate tracking of non-equilibrium maneuvers. In past and current work, saturation has not been explicitly incorporated into the controller design and we propose to extend our work in this direction, including experimental validation.

For our current work, we have been making use of the fact that the pitch axis dynamics of an aircraft are approximately “differentially flat”. A system is said to be differentially flat if all of the feasible trajectories for the system can be written as functions of a flat output $z(\cdot)$ and its derivatives. In other words, given a nonlinear control system

$$(1) \quad \begin{aligned} \dot{x} &= f(x, u) \\ y &= h(x) \end{aligned}$$

we say the system is *differentially flat* if there exists a function $z(x, u, \dot{u}, \dots, u^{(p)})$ such that all feasible solutions of the underdetermined differential equation (1) can be written as

$$\begin{aligned} x &= \alpha(z, \dot{z}, \dots, z^{(q)}) \\ u &= \beta(z, \dot{z}, \dots, z^{(q)}). \end{aligned}$$

Differentially flat systems were originally studied by Fliess et al. in the context of differential algebra [23] and later using Lie-Backlund transformations [24]. In [59] we reinterpreted flatness in a differential geometric setting. We made extensive use of the tools offered by exterior differential systems and the ideas of Cartan. Using this framework we were able to recover most of the results currently available using the differential algebraic formulation and achieve a deeper geometric understanding of flatness. We also showed that differential flatness is more general than feedback linearization in the multi-input case. More importantly, the point of view is quite different, focusing on trajectories rather than feedback transformations. See [41] for a description of the role of flatness in control of mechanical systems.

As a simple example, consider the dynamics of a planar, vectored thrust flight control system as shown in Figure 2. This system consists of a rigid body with body fixed forces and is a simplified model for the Caltech ducted fan described in Section 3.3. Let (x, y, θ) denote the position and orientation of the center of mass of the fan. We assume that the forces acting on the fan consist of a force f_1 perpendicular to the axis of the fan acting at a distance r from the center of mass, and a force f_2 parallel to the axis of the fan. Let m be the mass of the fan, J the moment of inertia, and g the gravitational constant. We ignore aerodynamic forces for the purpose of this example.




FIGURE 2. Planar ducted fan engine. Thrust is vectored by moving the flaps at the end of the duct.

The dynamics for the system are

$$\begin{aligned}
 m\ddot{x} &= f_1 \cos \theta - f_2 \sin \theta \\
 m\ddot{y} &= f_1 \sin \theta + f_2 \cos \theta - mg \\
 J\ddot{\theta} &= rf_1.
 \end{aligned}
 \tag{2}$$

Martin et al. [32] showed that this system is flat and that one set of flat outputs is given by

$$\begin{aligned}
 z_1 &= x - (J/mr) \sin \theta \\
 z_2 &= y + (J/mr) \cos \theta.
 \end{aligned}
 \tag{3}$$

Using the system dynamics, it can be shown that

$$\ddot{z}_1 \cos \theta + (\ddot{z}_2 + g) \sin \theta = 0.
 \tag{4}$$

and thus given $z_1(t)$ and $z_2(t)$ we can find $\theta(t)$ except for an ambiguity of π and away from the singularity $\ddot{z}_1 = \ddot{z}_2 + g = 0$. The remaining states and the forces $f_1(t)$ and $f_2(t)$ can then be obtained from the dynamic equations, all in terms of z_1 , z_2 , and their higher order derivatives.

Having determined that the system is flat, it follows that all feasible trajectories for the system are characterized by the evolution of the flat outputs. Using this fact, we can convert the problem of point to point motion generation to one of finding a curve $z(\cdot)$ which joins an initial $z(0), \dot{z}(0), \dots, \dot{z}^{(q)}(0)$, corresponding to the initial state, to a final set of values for the same quantities, corresponding to the final state. In this way, we reduce the problem of generating a feasible trajectory for the system to a classical algebraic problem in interpolation (Bezier splines provide one particularly nice solution technique). Similarly, problems in trajectory generation can also be converted to problems involving curves $z(\cdot)$ and algebraic methods can be used to provide real-time solutions [57, 58].

Thus, for differentially flat systems, trajectory generation can be reduced from a dynamic problem to an algebraic one. Specifically, one can parameterize the flat outputs using basis functions $\phi_i(t)$,

$$z = \sum a_i \phi_i(t),$$

and then write the feasible trajectories as functions of the coefficients a :

$$\begin{aligned}x_d &= \alpha(z, \dot{z}, \dots, z^{(q)}) = x_d(a) \\ u_d &= \beta(z, \dot{z}, \dots, z^{(q)}) = u_d(a).\end{aligned}$$

Note that no ODEs need to be integrated in order to compute the feasible trajectories (unlike optimal control methods, which involve parameterizing the *input* and then solving the ODEs). This is the defining feature of differentially flat systems. The practical implication is that nominal trajectories and inputs which satisfy the equations of motion for a differentially flat system can be computed in a computationally efficient way (solution of algebraic equations).

One of the immediate directions that we plan to explore is the application of similar ideas to systems with input constraints. In particular, since we have an *explicit* description of the inputs and states as a function of the flat outputs, it is possible to examine tradeoffs between stability and performance in the presence of input constraints. In essence, one gets a nonlinear function of the coefficients a which must satisfy some bound while at the same time a must be chosen to track a given trajectory. When both conditions cannot be simultaneously met, a tradeoff must be established to resolve the conflict. The choice of basis functions which simplify this tradeoff is an open problem with significant consequences.

Another issue with respect to choice of basis functions is the allocation of control surfaces to achieve a given input. Consider a flight control system with multiple control surfaces that can be used to generate a desired force and moment (a simple example is the vectored thrust system of Figure 2 with a wing and controllable flap added). In general, the different control surfaces will have different magnitude and rate saturations which must be respected. For linear systems, one might separate the control action by assigning the high-frequency portion of the control signal to the faster actuator and use the slow actuator for the low-frequency portion of the control signal. When the actuators experience both magnitude and rate saturations, these simple strategies are not guaranteed to give good performance.

If the system is differentially flat, one potential solution is to choose basis functions which are matched to the performance characteristics of the actuators. For example, we might parameterize the flat output as

$$z(t) = \sum q_i \phi_i(t) + \sum b_j \psi_j(t),$$

where the basis functions $\{\phi_i\}$ are associated with one actuator and $\{\psi_j\}$ are associated with the other. By choosing basis functions which generate appropriate commands for the actuators, we allow the feedforward controller to properly trade off performance with actuator bandwidth (now interpreted in a nonlinear context).

For systems which are not known to be differentially flat (necessary and sufficient condition are not yet available), there are several directions to pursue. One is to generate a low-dimensional “normal system” which captures the essential structure of the problem while at the same time reducing the complexity of the system description. For differential flat systems, the normal system is zero dimensional, meaning that the solutions of the system do not have to satisfy any differential equations. For more general systems, non-empty minimal systems are necessary (see [41] for some initial thoughts in this direction). Another approach is the use of a limited set of feasible trajectories which are sufficient for generating paths between any two points but provide a simplified description of the systems. This is the basic idea behind the work on PD control on Lie groups and Riemannian manifolds reported in [10, 11, 12], where one-parameter subgroups and geodesics form the set of sufficient trajectories. Finally, one can always approximate a system by a differentially flat one, hoping to retain more of the global geometry than simple Jacobian linearization about an equilibrium point. This technique in fact works remarkably well for the full dynamic model of the ducted fan (including aerodynamics and stand dynamics), which fails all currently known tests for differential flatness.

3.2. Local feedback control in the presence of saturation. While the use of feedforward trajectories will be used to insure that the *nominal* trajectory satisfies the nominal dynamics and constraints, in the presence of noise and uncertainty it is necessary to also insure that the local feedback controller, which is correcting for (hopefully) small deviations, is robust with respect to input constraints. Methods for analyzing and designing these local feedbacks must be applicable to systems which are not necessarily in a special normal form (beyond that which is available for all mechanical systems) if they are to be applied to standard flight control models. Furthermore, it is important that nonlinear controllers be capable of making use of advanced linear synthesis techniques for providing good local performance away from saturation.

One of the initial techniques that we plan to explore is the use of nonlinear gain scheduling to modify the gains as we near the limits of the operating envelope. While this is a standard *ad hoc* technique for controller design, we expect to be able to make substantial progress on systematic techniques for gain scheduling for the specific case of actuator saturation, in part by exploiting both the mechanical nature of the system as well as knowledge of the design of the feedforward portion of the controller.

A guiding principle in the synthesis of local control laws is to attempt to find techniques which give both local robust performance and global robust stability. That is, when the actual trajectory and the desired trajectory are close, we wish to utilize controllers that give guaranteed performance in the presence of uncertainty. This is possible, for example, by the use of scheduled or linear parameter varying (LPV) robust controllers when operating near the desired trajectory [5, 28, 48]. At the same time, for large deviations from the desired trajectory or expected operating condition, it is important that the controller not lead to instability of the overall system. However, performance issues are now secondary and can be relaxed to insure stability. This point of view has been explored in the context of flight control by Morton et al. [40].

To illustrate more specifically the approach which we propose to pursue, consider the problem of local stabilization of the following strongly nonlinear system:

$$\begin{aligned}\dot{x}_1 &= u_1 \\ \dot{x}_2 &= u_2 \\ \dot{x}_3 &= x_2 u_1.\end{aligned}$$

One can readily check that the linearization of this system about the origin is not controllable and hence most existing techniques for nonlinear control are not applicable. It can further be shown that this system cannot be stabilized to the origin using any static state feedback which is continuous on an open neighborhood of the origin [7].

Coron [15] showed that this system could be smoothly stabilized using time-varying feedback. For example, the feedback law

$$\begin{aligned}u_1 &= -x_1 + x_3 \cos t \\ u_2 &= -x_2 + x_3^2 \sin t\end{aligned}$$

can be shown to asymptotically stabilize the origin [54]. However, the convergence rate for this feedback is extremely slow and goes only as $1/\sqrt{t}$. M'Closkey and Murray developed a technique for improving the convergence of the system by applying a *nonlinear wrapper* to the original control law [44]. Given the original control law $u = \alpha(x, t)$, a new control law is computed of the form

$$u' = \rho(x, t) \cdot \alpha(\gamma(x, t), t).$$

The map $\gamma : \mathbb{R}^{n+1} \rightarrow \mathbb{R}^n$ modifies the argument passed to the original control while the function $\rho : \mathbb{R}^{n+1} \rightarrow \mathbb{R}^{m \times m}$ scales the output from the control law. The techniques used to design the wrapper functions γ and ρ involve the use of homogeneous vector fields relative to a nonstandard dilation and generate a control law which gives provably exponential convergence and is everywhere continuous and smooth except at the origin (see [33, 34, 35, 36, 37, 44] for a full development as well as experimental results).




FIGURE 3. Caltech ducted fan with support stand. The figures at the right indicate different thrust modes of the fan unit.

Similar techniques for treating saturation in nonlinear systems have been developed by Praly [50], also using tools from homogeneous vector fields. In Praly's work, he constructs nonlinear wrapper functions which scale the gains of a linear controller so as to avoid instability near the input limits of the system. We intend to explore the application of these specific techniques to flight control systems, but more importantly we hope to make use of nonlinear wrappers to simultaneously provide local performance and global stability in the presence of saturation. This would provide substantial improvements to some of the techniques described in Section 2.1 for using nonlinear control to achieve semi-global stability. The challenge is to design nonlinear wrappers which can modify a dynamic linear controller so as to properly account for saturation and insure global stability without sacrificing local performance.

3.3. Experimental validation on a flight control experiment. In addition to pursuing research in the basic theory and algorithms for controller analysis and synthesis, we will implement successful control techniques on a small flight control experiment at Caltech. In this section we give a brief description of the experiment and indicate the role of experimental validation in the research to be pursued under this proposal.

As part of an experimental project funded by the F18-HARV project at NASA's Dryden Flight Research Center, we have constructed a small vectored-thrust flight control experiment for use in robust nonlinear controls research. The current system consists of a high-efficiency electric motor with a 6-inch diameter blade, capable of generating up to 9 Newtons of thrust. Flaps on the fan allow the thrust to be vectored from side to side and even reversed. The engine is mounted on a three degree of freedom stand which allows horizontal and vertical translation as well as unrestricted pitch angle. A diagram of the system is shown in Figure 3.

One of our goals for this system is to allow real-time input of the desired xy position or velocity of the fan and to design controllers that can track any trajectory within the performance limits of the fan and that respond well to commands outside of those limits. Even in the absence of constraints, standard approaches to this problem, such as I/O linearization, cannot be used here for several reasons: 1) the system has non-minimum phase zero dynamics with respect to the xy position, 2) the system is not full state linearizable with static feedback (nor obvious dynamic feedbacks), and




FIGURE 4. New ducted fan with attached wing.

3) we do not have access to higher derivatives of the desired trajectory. For these reasons we must study the application of more advanced techniques to generate controllers for the system.

The system has many characteristics that are found on real aircraft. Aerodynamic forces are significant at high velocities and actuator saturations are common. In particular, some controllers have gotten “locked in” to high velocity, forward flight modes because the control authority of the vectored thrust was not sufficient to overcome the aerodynamic pitching moment. Rate saturations also play a role since the flaps are controlled by R/C servos which rate limit for large excursions of the flaps. Uncertainties, noise, and computational limits are also present and must be accounted for in order to get working control designs.

In addition to developing new theory for use on the existing experiment, we intend to continue to develop this experiment into a more realistic flight test experiment. We have recently built a new engine unit and added a small wing with actuated control surfaces for use in forward flight (see Figure 4). This modification adds more characteristics to the system that are found in full-scale aircraft and will enable us to study control issues which arise on many vectored thrust aircraft. In particular, the thrust to weight ratio of the new fan is approximately 0.5, which is much more realistic for the types of maneuvers we wish to execute, and the addition of a wing flap introduces multiple surfaces which produce similar effects but have different saturation limits (in both magnitude and rate). Detailed models of the existing system are currently being developed and work in this direction will continue using wind tunnels and other facilities already available at Caltech.

We will use this experiment to validate the control approaches pursued as part of this proposal and to help better understand and narrow the gap between theory and implementation. The ducted fan is an ideal platform to help bridge this gap since it is simple enough to be used in a university environment while being complicated enough to provide a challenge for existing and future control techniques. Our experience with the ducted fan has shown that industry is much more interested in devoting resources to implementing control methods developed in academia if they see a working demo which has enough of the characteristics of a real system to indicate that the control methods have the potential of leading to improved performance on practical systems (and not just in simulation). We have actively sought input from the government and industry researchers about additional modifications to the system which can be made to enhance its use as a realistic experiment for flight control systems. For example, the addition of the wing is due in part to suggestions from personnel at NASA Dryden.

3.4. Other research directions. There are many other interesting problems which we hope to pursue if our initial objectives are quickly achieved. One of the most interesting is the problem of modeling and predicting pilot induced oscillations in nonlinear flight control systems. Very little work has been done on nonlinear characterizations of pilot induced oscillations, even though it appears that nonlinear effects can play a significant role [20]. There is considerable interest in industry in better characterizing the sources of pilot induced oscillations so that susceptible flight modes can be identified and avoided.

A related problem to actuator saturation is dead zones. Like rate saturations, dead zones are present in a large variety of nonlinear systems and, analyzed linearly, can lead to time-delays which can reduce the phase margin of the control system. Furthermore, dead zones boundaries are often unknown or change during operation. Some nonlinear approaches to (adaptive) control of systems with dead zones have been studied by Tao and Kokotovic [53]. A new model for control of systems with friction (a common source of dead zones due to break-away forces) has recently been proposed by Canudas de Wit et al. [18] and may provide a starting point for work on dead zone nonlinearities.

Finally, there are a number of special cases of flight control systems for which some of the techniques described above may be easier to apply than others. For example, it is possible to show that the rigid body dynamics of a nonrolling missile are differentially flat and certain choices of actuator locations for rolling missiles also result in differential flatness [45]. Thus the geometry of these systems is considerably simpler than more general aircraft systems. Other specific systems which have been studied in some detail and may lead to special results include underwater vehicles [30] and aerial towed cable systems [42].

4. FACILITIES AVAILABLE

Caltech is an ideal environment for carrying out the types of studies proposed here. Graduate students studying controls all take a strong basic set of courses in operator theory, differential geometry, dynamical systems, and modern control theory as part of the required courses for a degree in Control and Dynamical Systems, and are therefore equipped with a strong set of mathematical tools that is necessary for developing new results in nonlinear control. In addition, we have a well developed experimental infrastructure which encourages students to get involved with experiments in flight controls, propulsion systems, flexible structures and robotics. A few of the specific resources to be used for the research presented in this proposal are detailed below.

4.1. Caltech ducted fan. The Caltech ducted fan, described briefly above, is a small flight control experiment which has been in use for the past two years as part of research in robust nonlinear control of aerospace vehicles [14, 28]. It roughly mimics the pitch axis dynamics of a vectored-thrust aircraft and is under complete computer control. The goals of the experiment are to study the use of active control for rapid transition between hover and forward flight as well as transition between forward and reverse flight (the ducted fan equivalent of a Herbst maneuver). A broad overview of the project and current progress is available on the World Wide Web at URL <http://avalon.caltech.edu/~dfan>.

The system which we have built is shown in Figure 3. It consists of a high-efficiency electric motor with a 6-inch diameter blade, capable of generating up to 9 Newtons of thrust. Flaps on the fan allow the thrust to be vectored from side to side and even reversed. The engine is mounted on a three degree of freedom stand which allows horizontal and vertical translation as well as unrestricted pitch angle. Recent modifications (some still in progress) include design and construction of a new fan unit (capable of providing 20 Newtons of thrust), addition of a wing with actuated flap, redesign of the stand to eliminate unwanted dynamics, improved aerodynamic models using wind tunnel data, and synthesis of more aggressive trajectories.

Previous work on this system has compared the application of a number of different control paradigms to this system and generated new insights into the practical requirements of a good

nonlinear control method [28]. This system has also been used for experimental application of linear parameter varying (LPV) control methods [4, 5], system identification [5, 39], and model validation [39, 46].

4.2. Computational facilities. Computing for this project will be directly supported by a cluster of Sun workstations. The cluster currently consists of a SPARCstation 20 dual processor computation server and approximately 10 smaller clients. Software includes a number commercially available packages, including Matlab (for controller synthesis) and Mathematica.

Real-time control is performed using PC-based data acquisition and control systems and custom-designed software. The basic hardware and software infrastructure has been in place for over 18 months and is currently being used on approximately 8 different research and instructional experiments. Servo rates of up to 4000 Hz are possible using this hardware, which is acceptable for most low frequency experiments. Software and documentation are in place for rapid synthesis and testing of Matlab control designs, including a module for implementing linear parameter varying (LPV) controllers.

4.3. Merrill Wind Tunnel. The Graduate Aeronautical Laboratory at Caltech (GALCIT) has a number of experimental facilities which are available for use by faculty and students. The Merrill Wind Tunnel is a small wind tunnel which is easily operated by a single person and is currently being used for identification of aerodynamic coefficients for the Caltech ducted fan. It has a three-foot diameter working section and is capable of operating from very low speeds through about 30 m/s (65 mph).

5. SUMMARY AND IMPLEMENTATION PLAN

5.1. Summary. Magnitude and rate saturations are a major source of nonlinearity in flight control and other mechanical systems. To date, there exist no systematic methods of analyzing and synthesizing controllers which simultaneously give global stability and local performance for nonlinear systems. The research described in this proposal will explore fundamental theoretical issues as well as develop tools and algorithms which can be implemented on realistic experiments. This is consistent with other research projects that have been undertaken by the PI and there is considerable expertise and infrastructure available at Caltech for pursuing this type of research.

Initially, the main techniques that will be explored involve the separation of the overall problem into real-time generation of feasible trajectories, followed by local control onto those trajectories. Current work (funded by the Air Force and the National Science Foundation) is concentrating on some of the theoretical and practical issues associated with trajectory generation and control of mechanical systems without explicitly taking into account magnitude and rate saturations. The work proposed here will complement that effort since it will be primarily focused on the role of actuator saturations in both the feedforward and feedback sections of the controller.

5.2. Yearly goals. The broad goals of this project are to develop new theoretical and computational tools for control of mechanical systems in the presence of magnitude and rate saturations, to implement and test these tools on a university flight control experiment, and to transition successful techniques into industry for further development. This research builds on several of the PI's previous research areas and complements existing projects at Caltech in which the PI is involved. The yearly goals for the first three years of the project are as follows:

Year 1:

1. Extend real-time trajectory generation techniques which are currently being developed to work in the presence of magnitude and rate saturations, for example by specific choice of basis functions.

2. Develop techniques for local feedback control of systems with input constraints to allow the use of existing linear (and nonlinear) control methodologies by designing nonlinear control “wrappers” which suitably modify the inputs and outputs to prespecified linear control laws.
3. Demonstrate initial theoretical paradigms on the Caltech ducted fan and compare to standard control techniques which have already been implemented.

Year 2:

1. Develop real-time trajectory generation techniques which account for magnitude and rate saturation as well as multiple control surfaces with different actuator bandwidths and performance limits.
2. Using results of experimental implementations from the first year, explore the use of more complicated nonlinear controllers for local control (if necessary) or more complicated interactions between the feedforward and feedback portions of the two degree of freedom controllers.
3. Implement and test techniques developed in the first year of the proposal on the Caltech ducted fan and design trajectories and test procedures that are relevant to full-scale flight.

Year 3:

1. Demonstrate and test a nonlinear controller capable of providing local performance and global stability on the Caltech flight control experiment while performing aggressive maneuvers.
2. Continue research in new directions consistent with the goals of this proposal.

5.3. Interaction with industry. One of the measures of ultimate success of this project is the degree to which it provides new insight and tools for use in industrial and military applications. In order to maximize the potential benefit of this research to the Navy and to industry, contacts will be established during the course of the project with the goal of transferring successful techniques from university labs to military and industrial research programs as well as transferring experience and realistic problem specifications from industry back into the university.

Two specific contacts have already been identified and have agreed to work closely with the PI on this project: Blaise Morton at the Honeywell Technology Center and Kevin Wise of McDonnell Douglas Corporation. Dr. Morton is a staff scientist at Honeywell and has had substantial interaction with a number of different flight control projects. He already has a close working relationship with Caltech as part of an AFOSR program on Partnerships for Research Excellence and Transition (PRET) which has recently been established at Caltech. Kevin Wise will provide expertise in aircraft and missile control systems, as well as facilitate the transfer of successfully demonstrated techniques into MDC research programs. Dr. Wise has expressed a strong interest in working with Caltech and in particular the problem of pilot induced oscillations. Letters of support from Honeywell and McDonnell Douglas are attached at the end of the proposal.

REFERENCES

- [1] E. Bayo and B. Paden. On trajectory generation for flexible robots. *Journal of Robotic Systems*, 4(2):229–235, 1991.
- [2] A. M. Bloch, P. S. Krishnaprasad, J. E. Marsden, and R. M. Murray. Nonholonomic mechanical systems with symmetry. Technical Report CIT/CDS 94-013, California Institute of Technology, 1994. To appear *Archive for Rational Mechanics*. Available electronically from <http://avalon.caltech.edu/cds/reports>.
- [3] A. M. Bloch and J. E. Marsden. Stabilization of rigid body dynamics by the energy-Casimir method. *Systems and Control Letters*, 14:341–346, 1990.
- [4] B. Bodenheimer, P. Bendotti, and M. Kantner. Linear parameter-varying control of a ducted fan engine. *International Journal of Robust and Nonlinear Control*, 1995. (submitted).

- [5] R. Bodenheimer. *The Whirling Blad and the Steaming Cauldron*. PhD thesis, California Institute of Technology, 1995.
- [6] S. P. Boyd and C. H. Barratt. *Linear Controller Design: Limits of Performance*. Prentice-Hall, 1991.
- [7] R. W. Brockett. Asymptotic stability and feedback stabilization. In R. W. Brockett, R. S. Millman, and H. J. Sussman, editors, *Differential Geometric Control Theory*, pages 181–191. Birkhauser, 1983.
- [8] J. J. Buchholz. Time delay induced by control surface rate saturation. *Z. Flugwiss. Weltraumforsch.*, 17:287–293, 1993.
- [9] D. J. Bugajski and D. F. Enns. Nonlinear control law with application to high angle-of-attack flight. *Journal of Guidance, Control, and Dynamics*, 15(3):761–767, 1992.
- [10] F. Bullo and R. M. Murray. Proportional derivative (PD) control on the Euclidean group. In *Proc. European Control Conference*, 1995.
- [11] F. Bullo and R. M. Murray. Proportional derivative (PD) control on the Euclidean group. CDS Technical Report CIT/CDS 95-010, California Institute of Technology, 1995. Available electronically from <http://avalon.caltech.edu/cds/reports/>.
- [12] F. Bullo, R. M. Murray, and A. Sarti. Control on the sphere and reduced attitude stabilization. In *IFAC Symposium on Nonlinear Control Systems Design (NOLCOS)*, 1995.
- [13] D. Chen. An iterative solution to stable inversion of nonminimum phase systems. In *Proc. American Control Conference*, pages 2960–2964, 1993.
- [14] H. Choi, P. Sturdza, and R. M. Murray. Design and construction of a small ducted fan engine for nonlinear control experiments. In *Proc. American Control Conference*, pages 2618–2622, 1994.
- [15] J-M. Coron. Global asymptotic stabilization for controllable systems without drift. *Mathematics of Control, Signals, and Systems*, 5:295–312, 1991.
- [16] M. A. Dahleh and I. J. Diaz-Bobillo. *Control of Uncertain Systems: A Linear Programming Approach*. Prentice Hall, 1995.
- [17] M. A. Dahleh and J. B. Pearson. L_1 -optimal compensators for continuous-time systems. *IEEE Transactions on Automatic Control*, 32(10):889–895, 1987.
- [18] C. Canudas de Wit, H. Olsson, K. J. Astrom, and P. Lischinsky. A new model for control of systems with friction. *IEEE Transactions on Automatic Control*, 40(3):419–425, 1995.
- [19] S. Devasia and E. Bayo. Inverse dynamics of articulated flexible structures—Simultaneous tracking and vibration reduction. *Dynamics and Control*, 4(3):299–309, 1994.
- [20] M. A. Dornheim. Report pinpoints factors leading to YF-22 crash. *Aviation Week and Space Technology*, November 1992.
- [21] M. A. Dornheim. X-31, F-16 MATV, F/A-18 HARV explore diverse missions. *Aviation Week and Space Technology*, 140(16):46–47, 1994.
- [22] B. Etkin. *Dynamics of Flight: Stability and Control*. Wiley, 1982.
- [23] M. Fliess, J. Levine, P. Martin, and P. Rouchon. On differentially flat nonlinear systems. *Comptes Rendus des Séances de l'Académie des Sciences*, 315:619–624, 1992. Serie I.
- [24] M. Fliess, J. Levine, P. Martin, and P. Rouchon. Linéarisation par bouclage dynamique et transformations de Lie-Bäcklund. *Comptes Rendus des Séances de l'Académie des Sciences*, 317(10):981–986, 1993. Serie I.
- [25] L. R. Hunt, R. Su, and G. Meyer. Global transformations of nonlinear systems. *IEEE Transactions on Automatic Control*, AC-28(1):24–31, 1983.
- [26] A. Isidori and C. I. Byrnes. Output regulation of nonlinear systems. *IEEE Transactions on Automatic Control*, 35(2):131–140, 1990.
- [27] B. Jakubczyk and W. Respondek. On linearization of control systems. *Bulletin de L'Academie Polonaise des Sciences, Série des sciences mathématiques*, XXVIII:517–522, 1980.
- [28] M. Kantner, B. Bodenheimer, P. Bendotti, and R. M. Murray. An experimental comparison of controllers for a vectored thrust, ducted fan engine. In *Proc. American Control Conference*, 1995.
- [29] S. D. Kelly and R. M. Murray. Geometric phases and robotic locomotion. *Journal of Robotic Systems*, 12(6):417–431, 1995.
- [30] N. E. Leonard. Underwater vehicle dynamics and stability. Technical report, Princeton University, 1995. Preprint.
- [31] A. D. Lewis and R. M. Murray. Equilibrium controllability for a class of mechanical systems. In *Proc. IEEE Control and Decision Conference*, 1995.
- [32] P. Martin, S. Devasia, and B. Paden. A different look at output tracking: Control of a VTOL aircraft. In *Proc. IEEE Control and Decision Conference*, pages 2376–2381, 1994.
- [33] R. T. M'Closkey and R. M. Murray. Nonholonomic systems and exponential convergence: Some analysis tools. In *Proc. IEEE Control and Decision Conference*, pages 943–948, 1993.
- [34] R. T. M'Closkey and R. M. Murray. Experiments in exponential stabilization of a mobile robot towing a trailer. In *Proc. American Control Conference*, pages 988–993, 1994.

- [35] R. T. M'Closkey and R. M. Murray. Exponential stabilization of driftless nonlinear control systems via time-varying, homogeneous feedback. In *Proc. IEEE Control and Decision Conference*, pages 1317–1322, 1994.
- [36] R. T. M'Closkey and R. M. Murray. Extending exponential stabilizers for nonholonomic systems from kinematic controllers to dynamic controllers. In *Proc. IFAC Symposium on Robot Control (SYROCO)*, pages 243–248, 1994.
- [37] R. T. M'Closkey and R. M. Murray. Exponential stabilization of driftless nonlinear control systems using homogeneous feedback. Technical Report CIT/CDS 95-012, California Institute of Technology, 1995. Submitted, *IEEE T. Automatic Control*. Available electronically from <http://avalon.caltech.edu/cds/reports>.
- [38] G. Meyer, L. R. Hunt, and R. Su. Nonlinear system guidance in the presence of transmission zero dynamics. Technical report, NASA, 1994. Preprint.
- [39] J. Morris. *Experimental Control: A Helicopter Case Study*. PhD thesis, California Institute of Technology, 1995.
- [40] B. Morton, D. Enns, and B.-Y. Zhang. Stability of dynamic inversion control laws applied to nonlinear aircraft pitch-axis models. *International Journal of Control*, 1995. (to appear).
- [41] R. M. Murray. Nonlinear control of mechanical systems: A Lagrangian perspective. In *IFAC Symposium on Nonlinear Control Systems Design (NOLCOS)*, 1995. Available electronically from <http://avalon.caltech.edu/cds/~murray>.
- [42] R. M. Murray. Trajectory generation for a towed cable flight control system with input constraints using differential flatness. In *IFAC World Congress*, 1996. (submitted).
- [43] R. M. Murray, Z. Li, and S. S. Sastry. *A Mathematical Introduction to Robotic Manipulation*. CRC Press, 1994.
- [44] R. M. Murray and R. T. M'Closkey. Converting smooth, time-varying, asymptotic stabilizers for driftless systems to homogeneous, exponential stabilizers. In *Proc. European Control Conference*, 1995.
- [45] R. M. Murray, M. Rathinam, and W. Sluis. Differential flatness of mechanical control systems: A catalog of prototype systems. In *ASME International Mechanical Engineering Congress and Exposition*, 1995.
- [46] M. Newlin. *Model Validation, Computation, and Control*. PhD thesis, California Institute of Technology, 1995.
- [47] J. P. Ostrowski. *The Mechanics and Control of Undulatory Robotic Locomotion*. PhD thesis, California Institute of Technology, 1995.
- [48] A. Packard. Gain scheduling via linear fractional transformations. *Systems and Control Letters*, 22(2):79–92, 1994.
- [49] B. Paden, D. Chen, R. Ledesma, and E. Bayo. Exponentially stable tracking control for multi-joint flexible-link manipulators. *ASME Journal of Dynamic Systems, Measurement and Control*, pages 1–17, 1991.
- [50] L. Praly. Robustesse des systèmes non linéaires. Technical Report CAS E161, Ecole Nationale Supérieure des Mines de Paris, 1994.
- [51] J. S. Shamma. Robust stability with time-varying structured uncertainty. *IEEE Transactions on Automatic Control*, 39(4):714–724, 1994.
- [52] H. J. Sussmann, E. D. Sontag, and Y. D. Yang. A general result on the stabilization of linear systems using bounded controls. *IEEE Transactions on Automatic Control*, 39(12):2411–2425, 1994.
- [53] G. Tao and P. V. Kokotovic. Adaptive control of plants with unknown dead zones. *IEEE Transactions on Automatic Control*, 39(1):59–68, 1994.
- [54] A. Teel, R. M. Murray, and G. Walsh. Nonholonomic control systems: From steering to stabilization with sinusoids. In *Proc. IEEE Control and Decision Conference*, pages 1603–1609, 1992. To appear, *Int'l. J. Control*.
- [55] A. R. Teel. *Feedback Stabilization: Nonlinear Solutions to Inherently Nonlinear Problems*. PhD thesis, University of California, Berkeley, 1992.
- [56] A. R. Teel. Semiglobal stabilization of minimum phase nonlinear systems in special normal forms. *Systems and Control Letters*, 19(3):187–192, 1992.
- [57] M. van Nieuwstadt and R. M. Murray. Approximate trajectory generation for differentially flat systems with zero dynamics. In *Proc. IEEE Control and Decision Conference*, 1995. (to appear).
- [58] M. van Nieuwstadt and R. M. Murray. Real time trajectory generation for differentially flat systems with unstable zero dynamics. In *IFAC World Congress*, 1996. (submitted).
- [59] M. van Nieuwstadt, M. Rathinam, and R. M. Murray. Differential flatness and absolute equivalence. In *Proc. IEEE Control and Decision Conference*, 1994. Also available as Caltech technical report CIT/CDS 94-006. Submitted, *SIAM J. Control and Optimization*.