

AFRL-VA-WP-TR-2005-3053

**SOFTWARE ENABLED CONTROL
Design of Hierarchical, Hybrid Systems**

Claire J. Tomlin

Stanford University

Department of Aeronautics and Astronautics

250 Durand

Stanford, CA 94305-4035



JANUARY 2005

Final Report for 15 October 1999 – 31 January 2005

Approved for public release; distribution is unlimited.

STINFO FINAL REPORT

AIR VEHICLES DIRECTORATE

AIR FORCE MATERIEL COMMAND

AIR FORCE RESEARCH LABORATORY

WRIGHT-PATTERSON AIR FORCE BASE, OH 45433-7542

NOTICE

Using Government drawings, specifications, or other data included in this document for any purpose other than Government procurement does not in any way obligate the U.S. Government. The fact that the Government formulated or supplied the drawings, specifications, or other data does not license the holder or any other person or corporation; or convey any rights or permission to manufacture, use, or sell any patented invention that may relate to them.

This report was cleared for public release by the Air Force Research Laboratory Wright Site Public Affairs Office (AFRL/WS) and is releasable to the National Technical Information Service (NTIS). It will be available to the general public, including foreign nationals.

THIS TECHNICAL REPORT IS APPROVED FOR PUBLICATION.

/s/

2Lt David Carlos
Developmental Engineer
Control Design and Analysis Branch

/s/

DEBORAH S. GRISMER
Chief, Control Design and Analysis Branch
Air Force Research Laboratory

/s/

BRIAN W. VAN VLIET, Chief
Control Sciences Division
Air Vehicles Directorate

This report is published in the interest of scientific and technical information exchange and its publication does not constitute the Government's approval or disapproval of its ideas or findings.

REPORT DOCUMENTATION PAGE					Form Approved OMB No. 0704-0188	
The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.						
1. REPORT DATE (DD-MM-YY) January 2005		2. REPORT TYPE Final		3. DATES COVERED (From - To) 10/15/1999 – 01/31/2005		
4. TITLE AND SUBTITLE SOFTWARE ENABLED CONTROL Design of Hierarchical, Hybrid Systems				5a. CONTRACT NUMBER F33615-99-C-3014		
				5b. GRANT NUMBER		
				5c. PROGRAM ELEMENT NUMBER 0602302		
6. AUTHOR(S) Claire J. Tomlin				5d. PROJECT NUMBER A04H		
				5e. TASK NUMBER		
				5f. WORK UNIT NUMBER 0G		
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Stanford University Department of Aeronautics and Astronautics 250 Durand Stanford, CA 94305-4035				8. PERFORMING ORGANIZATION REPORT NUMBER		
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) Air Vehicles Directorate Air Force Research Laboratory Air Force Materiel Command Wright-Patterson Air Force Base, OH 45433-7542				10. SPONSORING/MONITORING AGENCY ACRONYM(S) AFRL/VACA		
				11. SPONSORING/MONITORING AGENCY REPORT NUMBER(S) AFRL-VA-WP-TR-2005-3053		
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution is unlimited.						
13. SUPPLEMENTARY NOTES Report contains color.						
14. ABSTRACT The objective of this effort is to develop a hybrid control theory for multiple UAVs. Specifically, to develop: A) A hybrid interface between discrete and continuous systems, B) A coordination algorithm for UAVs with distributed sensors. Application areas are air traffic control and satellite formation flight. 1) Real time hybrid system analysis and controller design, 2) Distributed sensing systems, and 3) Asynchronous control theory. Hybrid design (1) is based on solving a constrained optimization problem. This is solved using LMIs with ellipsoidal bounding. The discrete modes use a theorem proven to validate that the modes transition correctly. Distributed sensing (2) is based on precision control of formations of UAVs or satellites with SAR or optical interferometry. Coordinated video ad motion is used to estimate position and resolve conflicts. Asynchronous theory (3) is needed for an effective distributed control architecture. Data is time stamped and weighed with value over time.						
15. SUBJECT TERMS Software Enabled Control, SEC, hybrid control, UAV						
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT: SAR	18. NUMBER OF PAGES 40	19a. NAME OF RESPONSIBLE PERSON (Monitor) 2Lt David Carlos 19b. TELEPHONE NUMBER (Include Area Code) (937) 255-8681	
a. REPORT Unclassified	b. ABSTRACT Unclassified	c. THIS PAGE Unclassified				

1 Grant summary

This report summarizes the research completed under the grant “Design of Hierarchical Hybrid Control Systems”, under the DARPA Software Enabled Control (SEC) program, administered by AFRL under contract F33615-99-C-3014, which was awarded to Stanford University over the period 10/15/99 to 01/31/05. The total amount awarded was \$3,048,809. The Principle Investigators for this grant were Professors Claire Tomlin (PI) and Stephen Boyd (Co-PI).

The four pillars of the Stanford SEC program were:

1. design and development of innovative mathematical methods for control of networks of complex systems;
2. design and development of a computational tool for (close to) real-time verification of complex, hybrid systems;
3. design and development of the embedded avionics for each aircraft in a fleet of two small Unmanned Air Vehicles (the Stanford DragonFly UAVs), and the coordination and cooperation of this development with the Boeing OCP developers;
4. the implementation and testing of the hybrid system reachability tool, and the control algorithms for networks of complex systems, on the Stanford DragonFly UAVs, and on the Boeing Capstone Demonstration Platform.

2 Personnel

Over the course of the grant, the personnel supported were:

*Associate Professor, Aeronautics and Astronautics, and, by courtesy, of Electrical Engineering
tomlin@stanford.edu, (phone) 650-723-5164, (fax) 650-723-3738, Durand 250, Stanford CA, 94305-4035.

Faculty: Claire Tomlin (Aero/Astro), Stephen Boyd (EE), David Dill (CS), and Zohar Manna (CS).

Postdoctoral Research Associates: Haitham Hindi (EE, now at PARC, Palo Alto); Dusan Stipanovic (Aero/Astro, now Assistant Professor, University of Illinois at Urbana Champaign); Henny Sipma (CS); Mikael Johansson (EE, now Assistant Professor, KTH Stockholm); Jung Soon Jang (Aero/Astro).

Graduate students: Ian Mitchell (now Assistant Professor, UBC); Michael Colon (CS); Gokhan Inalhan (Aero/Astro, now postdoctoral researcher, MIT); Inseok Hwang (Aero/Astro, now Assistant Professor, Purdue); Jung Soon Jang (Aero/Astro); Meeko Oishi (ME, now postdoctoral researcher, NEON); Lin Xiao (Aero/Astro, now postdoctoral researcher, Caltech); Alexandre Bayen (Aero/Astro, now Assistant Professor, Berkeley).

3 Research Area 1: Controlling systems over networks

The continuing drop in prices for data networks and computers has created a strong drive to let control system components communicate over a shared communication link rather than using separate dedicated communication channels. However, the presence of a communication network in the control loop complicates the control design. Communication over a shared medium introduces varying time delays, occasional loss of data and even the possibility of link failures. Traditional control theory gives little support for dealing with these uncertainties, and an understanding of the issues in networked control systems has just recently started to emerge. So far, the research on networked control systems has been exclusively devoted to single-byte data transfer. Today, however, transfer times are usually a negligible part of the communication delay and most network protocols allow us to send several bytes of data at no extra cost. This makes it interesting to investigate how multiple-byte transfers can be exploited to improve system performance and robustness to varying time-delays, occasional loss of data, and possible link failures. We have derived optimal controllers and estimators for packet-oriented communication, and shown that in situations with varying time delays, they perform better than approaches that use single-byte data transfer, without incurring any additional load on the communication network [1].

Now, consider jointly allocating communication resources and designing a linear control system in order to optimize system performance. This problem is in general not convex, but can be solved heuristically in a way that exploits the special structure of the communication resource allocation problems. Numerical examples show that our approach of jointly optimizing both the system and communication resources can have very significant benefits over the traditional approach [2].

Finally, having established the SRRA framework with theoretical channel capacity functions in previous work, we have moved one step further to solve the SRRA problem with practical fading wireless channels. We showed that the joint optimization of data routing, resource allocation and power control over fading states can be efficiently solved by convex optimization techniques, which opens the door for practical application of the SRRA framework. We also extended the SRRA framework to heterogeneous networks, where wired and wireless networks with different media access schemes are connected together to facilitate ubiquitous data communication. Our SRRA framework can be extended naturally to handle such complex systems.

4 Research Area 2: Computing Reachable Sets for Hybrid Systems

The hybrid systems framework provides an appealing means for verifying the safety of dynamical systems. We address safety verification as a reachability problem: Given an unsafe subset of the system state space and an initial state, is the former reachable from the latter? If it is reachable despite any controllable actions that we can take then the system is provably unsafe.

The continuous-time nonlinear dynamics of the system need to be considered in assessing reachability. Our formulation requires the solution of a Hamilton-Jacobi partial differential equation, and a grid-based numerical solution approach based on level set methods is used for this purpose. Simulation examples are presented for three flight management applications: two-aircraft collision avoidance, the related problem of conflict resolution, and ensuring safety during final landing approach.

The exact reachability computation is prey to the curse of dimensionality: its computational complexity is exponential with respect to the continuous dimension. We also developed alternative approaches, which are based on over-approximating the reachable set of states with a polyhedron. This is also computationally intractable since the propagation of the system's dynamics will result in a potentially unlimited number of constraints (faces of the polyhedron), but we have developed a novel technique for identifying and pruning redundant and irrelevant constraints. This technique promises to be computationally feasible for very high dimensional problems. The basis of the approach is the computation of the maximum volume ellipsoid contained in a polyhedron, a computation that can be formulated as a convex optimization problem (for which global and efficient algorithms are available).

4.1 Introduction

For about the past ten years, researchers in the traditionally distinct fields of control theory and computer science verification have proposed models, and verification and controller synthesis techniques for complex, safety critical systems. The area of *hybrid systems theory* studies systems which involve the interaction of discrete event and continuous time dynamics, and provides an effective modeling framework for complex continuous systems with large numbers of operating modes. Examples include continuous systems controlled by a discrete logic such as the autopilot modes for controlling an aircraft, systems of many interacting processes such as air or ground transportation systems in which discrete dynamics are used to model the coordination protocols among processes, or continuous systems which have a phased operation, such as biological cell growth and division.

The current and potential impact of hybrid systems lies in the confluence of computational methods from control theory and from formal methods in computer science verification. In the examples mentioned above, the system dynamics are complex enough that traditional analysis and control methods based solely on differential equations are not computationally feasible; analysis based solely on discrete event dynamics ignores critical system behavior. Our interest lies in developing computational tools for analyzing and controlling the behavior of hybrid systems; our eventual goal is to develop a real-time tool to provide online verification that a hybrid system satisfies its specified behavior. This goal addresses one of the central themes of Software Enabled Control, in which we seek a systematic methodology for the design, validation, and implementation of control software.

In our work, the system specification that we are most interested in is that of system safety, which asks the question: Is a potentially unsafe configuration of the system reachable from an initial configuration? More important for control theory, given a set of desired configurations, it is crucial to be able to design the hybrid system control inputs to achieve these configurations. Previous work in this area had focused on hybrid systems with very simple continuous dynamic equations (such as clocks, or linear decoupled systems). Our research has three components: the problem formulation for computing *exact reachable sets* of hybrid systems and the design of a software tool to perform such calculations; the development of an algorithm for computing overapproximations of reachable sets which works efficiently in high dimensions; the design of a real-time aircraft testbed for these algorithms. We will focus on the first component in this report, we briefly outline our work in the second.

4.2 Hybrid System Model

The main difference between our hybrid system model and leading models in the literature (see for example the work of Alur and Henzinger [3] on linear hybrid automata) is that we include accurate, nonlinear models of the hybrid system continuous dynamics. The model and algorithm have been developed jointly with Lygeros and Sastry, full details are presented in [4].

A hybrid automaton is a finite state machine with discrete states $\{q_1, q_2, \dots, q_k\}$, in which each discrete state has associated continuous dynamics $\dot{x} = f(q_i, x, \nu)$, with $x \in \mathbb{R}^n$, and continuous inputs and disturbances $\nu = (u, d)$, where $u \in U$ and $d \in D$.

Definition 1 (Hybrid Automaton). *A hybrid automaton H is a collection*

$$H = (Q, X, \Sigma, V, \text{Init}, f, \text{Inv}, R)$$

where

- $Q \cup X$ is the set of state variables, with Q a finite set of discrete states, and $X = \mathbb{R}^n$;
- $\Sigma = \Sigma_1 \cup \Sigma_2$ is a finite collection of discrete input variables, where Σ_1 is the set of discrete control inputs, and Σ_2 is the set of discrete disturbance inputs;
- $V = U \cup D$ is the set of continuous input variables, where U is the set of continuous control inputs, and D is the set of continuous disturbance or uncontrollable inputs; we denote the spaces of input (disturbance) trajectories as the sets of piecewise continuous functions $\mathcal{U}$ ($\mathcal{D}$ respectively), which take values in U (D respectively);
- $\text{Init} \subseteq Q \times X$ is a set of initial states;
- $f : Q \times X \times V \rightarrow TX$ is a vector field describing the evolution of x for each $q \in Q$; f is assumed to be globally Lipschitz in X (for fixed $q \in Q$) and continuous in V ;
- $\text{Inv} \subseteq Q \times X \times \Sigma \times V$ is called an invariant, and defines combinations of states and inputs for which continuous evolution is allowed;
- $R : Q \times X \times \Sigma \times V \rightarrow 2^{Q \times X}$ is a reset relation, which encodes the discrete transitions of the hybrid automaton.

We refer to $(q, x) \in Q \times X$ as the *state* of H and to $((\sigma_1, \sigma_2), (u, d)) \in \Sigma \times V$ as the *input* of H . The control actions (σ_1, u) model those inputs over which the designer has control, and the disturbance actions (σ_2, d) model inputs over which the designer has no control, such as uncertainties in the actions or behaviors of the other aircraft in the system. We assume that the designer has complete knowledge of the *bounds* on these disturbance actions (Σ_2, D) .

Associated to the hybrid automaton H is a *specification* which describes the condition one would like the system to satisfy. Our work has been motivated by verification and synthesis for safety critical applications, and as such we have been primarily interested in safety specifications. These specifications are encoded as subsets of the state space of the hybrid system: the set $G_0 \subseteq Q \times X$ is that subset which is defined *a priori* to be unsafe. The problem specification is therefore to keep the system state outside of the unsafe set.

4.3 Exact Reach Set Computation using Level Sets

In this section, we describe the development of a general purpose tool for exact reachable set computation—the core of which is a new variant of a “local level set” algorithm that efficiently computes an accurate representation of the reachable set boundary. We demonstrate on an example the numerical convergence of our computation by analyzing the results as the continuous state space grid is made finer, a standard method of validation for scientific computing codes. In this way, we show that high accuracy can be achieved at the cost of increased computational time and space. We illustrate our tool on a single mode aircraft conflict resolution example [5, 6], a three mode conflict resolution example, as well as an example of a six mode commercial aircraft auto-lander [7], which exhibits nondeterminism and cycles in its discrete behavior.

Our motivation for this component of the research stems from the belief that for many applications of hybrid systems, it is important to be able to accurately represent the reachable set. We have dealt primarily in the safety verification of avionics systems, where accurate representation of the safe region of operation translates into the ability to operate the system closer to the boundaries of that region, at a higher performance level than previously allowed. For very high dimensional state spaces, additional logic (such as projection operators) or new techniques (such as the convex overapproximations of the next section) will be needed; however, our results in this paper show that it is feasible to do exacting computation for hybrid systems with nonlinear continuous dynamics in three continuous state dimensions and six discrete modes, and we believe it will be feasible to extend this up to five continuous dimensions and large numbers of discrete modes.

4.4 Reachability for Hybrid Systems

In this section we summarize the general framework for handling the interaction between discrete and continuous dynamics (following [4]).

Fundamentally, reachability analysis in discrete, continuous or hybrid systems seeks to partition states into two categories: those that can reach a given target set, and those that cannot. We will label these two sets of states G and $E = G^c$ respectively (the symbol G has been used traditionally to represent unsafe sets, E is used as a symbol for “escape” sets). Any inputs to the hybrid automata are assumed to lie in bounded sets and to have the goal of locally maximizing or minimizing the reachable set: at each iteration, the reachability algorithm chooses values for inputs ξ_G that maximize the size of G and values for inputs ξ_E that minimize the size of G (and hence maximize the




Figure 1: Iterative Reachability Algorithm: Showing detail of iteration for discrete mode j at iteration i .

size of E). Any nondeterminism in the transition relation is also utilized to consistently maximize or minimize G , depending on the goal of the reachability computation. For hybrid automata, the discrete inputs σ and continuous inputs ν can be assigned to the two categories $\xi_G = (\sigma_G, \nu_G)$ and $\xi_E = (\sigma_E, \nu_E)$ according to whether they seek to maximize or minimize G . For example, for our hybrid system model with control actions (σ_1, u) and disturbance actions (σ_2, d) , we would have that $\xi_G = (\sigma_G, \nu_G) = (\sigma_2, d)$, and $\xi_E = (\sigma_E, \nu_E) = (\sigma_1, u)$.

The reachability computation follows an iterative, two stage algorithm shown graphically in Figure 1. The outer iteration computes reachability over the discrete switches, producing iterates G_i and E_i at iteration $i = 0, -1, -2, \dots$, where a negative index is used to indicate that the algorithm is initialized with a target set and run backward to compute those states which can reach the target set. The inner iteration runs a separate continuous reachability problem in each of the discrete modes $j = 1, 2, \dots, K$ to compute the estimates G_i^j and E_i^j . We define the “switch” sets

- $\tilde{G}_i^j$ contains all states in mode j from which a discrete transition to a state in G_{i-1} (typically a state in another mode) can be forced to occur through the application of a discrete input σ_G ; these states will be defined by the invariant of mode j and the guards of the transitions from mode j .
- $\tilde{E}_i^j$ contains all states from which a discrete transition to a state in E_{i-1} can be forced to occur through the application of a discrete input σ_E ; these states are also defined by the invariant of mode j and the guards of transitions from mode j .

Then the goal of the continuous reachability tool is to identify the “flow” sets

- $\tilde{G}_i^j(t)$ contains states from which for all ν_E there exists ν_G that will force the resulting trajectory to flow into $G_{i-1}^j \cup \tilde{G}_i^j$ within time t .
- $\tilde{E}_i^j(t)$ contains states from which there exists ν_E that for all ν_G will force the resulting trajectories to flow into $\tilde{E}_{i-1}^j$ within time t or to stay outside of $G_{i-1}^j \cup \tilde{G}_i^j$ for at least time t .

Note that in some problems the order of the existential and universal quantifiers in the definition

above must be reversed. Given these sets,

$$\begin{aligned} G_i^j &= \lim_{t \rightarrow \infty} \tilde{G}_i^j(t), & G_i &= \bigcup_{j=1}^K G_i^j, & G &= \lim_{i \rightarrow \infty} G_i, \\ E_i^j &= \lim_{t \rightarrow \infty} \tilde{E}_i^j(t), & E_i &= \bigcup_{j=1}^K E_i^j, & E &= \lim_{i \rightarrow \infty} E_i, \end{aligned}$$

where G_0^j is the set of initial conditions of the reachability problem and $E_0^j = (G_0^j)^c$. Simple modifications of this algorithm suffice to solve finite time reachability problems.

The procedure described above, developed in [4, 6], was motivated by the work of [8, 9] for reachability computation and controller synthesis on timed automata, and that of [10] for controller synthesis on linear hybrid automata. In that development the reachability problem's objective was to determine E —the largest controllable invariant subset of the state space—by computing the set of states G which were reachable in backwards time from the set of predefined unsafe states. In terms of the definitions above, control inputs from this problem lie in ξ_E and disturbance inputs in ξ_G . For safety, any model nondeterminism would be used to maximize the unsafe set G .

4.5 Continuous Reachability using Level Sets

While practical algorithms for computing discrete reachability over many thousands of states have been designed and implemented, determination of continuous reachability for even low dimensional systems is still an open problem. The continuous portion of a hybrid reachability problem requires methods of performing four key operations on sets: unions, intersections, tests of equality, and evolution according to the discrete mode's continuous flow field. The choice of representation for sets dictates the complexity and accuracy of these operations; consequently, continuous reachability algorithms can be classified according to how they represent sets.

For our exact representation scheme, we characterize the set being tracked implicitly by defining a “level set function” $J(x, t)$ throughout the continuous state space which is negative inside the set, zero on its boundary, and positive outside, and which encodes the initial data in $J(x, 0)$. The intersection of two such sets is simply the maximum of their level set functions at each point in state space, and the union is the minimum; a variety of easily implemented equality tests are possible. Evolution of a level set under a nonlinear flow field is governed by the Hamilton-Jacobi (HJ) partial differential equation (PDE) (see, for example, [5])

$$\begin{aligned} -\frac{\partial J(x, t)}{\partial t} &= \min\{0, \max_{\nu_{\min}} \min_{\nu_{\max}} f(x, \nu_{\min}, \nu_{\max})^T \nabla J(x, t)\}, \\ &= \min\{0, H(x, \nabla J(x, t))\} \end{aligned} \tag{1}$$

where $\nu_{\min}$ are those continuous inputs trying to minimize the size of the set being tracked, and $\nu_{\max}$ are those inputs trying to maximize its size. The order of the optimization must be chosen appropriately for the situation. The first minimization on the right hand side of (1) ensures that the set can only grow (states that are unsafe cannot become safe at a future time). That the subzero level sets of the viscosity solution of (1) yields the correct reachable set, is proved in [11]. The implicit representation has a number of advantages when compared with the explicit representations that other researchers are pursuing, including a conceptually simple representation

of very general sets and a size which is independent of the complexity of the set (although it grows exponentially with dimension). In addition, a set of sophisticated numerical techniques to accurately solve PDEs may be drawn upon for computation. In the remainder of this section, we focus on the representation (1), and assume that the modeler can compute the appropriate optimization over inputs in (1) if given x and $\nabla J(x, t)$.

The HJ PDE (1) is well known to have complex behavior. Even with smooth initial data $J(x, 0)$ and continuous Hamiltonian $H(x, \nabla J)$, the solution $J(x, t)$ can develop discontinuous derivatives in finite time; consequently, classical infinite time solutions to the PDE are generally not possible. In the quest for a unique weak solution Crandall and Lions introduced the concept of the viscosity solution [12]. For most problems of interest, finding the analytic viscosity solution is not possible (see [13] for cases in which it is possible), and so we seek a numerical solution. We approximate the solution of (1) on a Cartesian grid of nodes. Three terms in the equation must be approximated at each node, based on the values of the level set function at that node and its neighbors: the gradient ∇J , the Hamiltonian H , and the time derivative $\frac{\partial J(x, t)}{\partial t}$. We discuss each of these separately.

In each dimension at each grid point there exist both left and right approximations of the gradient ∇J , depending on which neighboring grid points' values are used in the finite difference calculation. We label the vector of left approximations ∇J^- , the vector of right approximations ∇J^+ , and will see below that ∇J^- , ∇J^+ or some combination of the two will be used to compute the numerical Hamiltonian $\hat{H}$. The accuracy of a derivative approximation is measured in terms of the order of its local truncation error; an order p method has error $\|\nabla J - \nabla J^\pm\| = \mathcal{O}(\Delta x^p)$. At the current time, we have implemented the basic first order accurate approximation for speed and a weighted, essentially non-oscillatory fifth order accurate approximation for high fidelity. We have chosen to use the well studied Lax-Friedrichs numerical Hamiltonian approximation $\hat{H}$ [14]

$$\hat{H}(x, \nabla J^-, \nabla J^+) = H(x, \frac{\nabla J^- + \nabla J^+}{2}) - \frac{1}{2}\alpha^T(\nabla J^+ - \nabla J^-), \quad (2)$$

where $H(x, \nabla J)$ is given by (1) and the term containing the vector coefficient α is a high order numerical dissipation added to damp out spurious oscillations in the solution. The time derivative of the PDE is handled by the method of lines: the value of the level set function J at each node is treated as an ODE $\frac{dJ}{dt} = \hat{H}$, with $\hat{H}$ given by (2). General ODE solvers, such as Runge-Kutta (RK) schemes, can then be applied.

The Hamilton-Jacobi equation (1) describes the evolution of the level set function over all of space. But we are only interested in its zero level set; thus, we can restrict our computational updates to nodes near the boundary between positive and negative $J(x, t)$ —an idea variously called “local level sets” [15] or “narrowbanding” [16]. We have implemented a new variant of this method in our code. Because the boundary is of one dimension less than the state space, considerable savings are available for two and three dimensional problems. If the number of nodes in each dimension is n (proportional to Δx^{-1}) and the dimension d , the total number of nodes is $\mathcal{O}(n^d)$; with the restriction on the timestep necessary for numerical stability, the total computational cost is $\mathcal{O}(n^{d+1})$. With local level sets, we reduce computational costs back down to $\mathcal{O}(n^d)$.

4.6 Examples

In this section we present three examples: a validation of our numerical implementation on a single mode, three dimensional aircraft collision avoidance example (see [6, 5] for details), and two multimodal examples. Once a method of determining continuous reachability is available, the

discrete iteration of the algorithm described in Section 4.4 is relatively straightforward. In fact, for discrete transition graphs with no cycles it is possible to order the continuous reachability problems such that no discrete iteration is required (the three mode example below). In order to examine the complications induced by discrete cycles—such as how to avoid zenoness, in what order to execute the continuous reachability problems, and how to determine which switches are active— we present an example representing the landing of a civilian airliner.

Numerical Validation of Aircraft Collision Avoidance

This example features a control aircraft trying to avoid collision with a disturbance aircraft, where both aircraft have fixed and equal altitude, speed and turning radius—they may only choose which direction they will turn:

$$\dot{x}_r = -v_u + v_d \cos \psi_r + u y_r, \quad \dot{y}_r = v_d \sin \psi_r - u x_r, \quad \dot{\psi}_r = d - u,$$

where $v_u = v_d = 5$ are the aircraft speeds, x_r and y_r are the relative planar location of the aircraft and ψ_r is their relative heading. The inputs $|u| \leq 1$ and $|d| \leq 1$ are the control’s and disturbance’s respective turn rates. The initial unsafe set $J(x, 0)$ is the interior of the radius five cylinder centered on the ψ_r axis. Choosing optimal inputs according to (1) with $\nu_G = \nu_{\max} = d$ and $\nu_E = \nu_{\min} = u$, we get the optimal Hamiltonian:

$$H(x, p) = -p_1 v_u + p_1 v_d \cos \psi_r + p_2 v_d \sin \psi_r + |p_1 y_r - p_2 x_r - p_3| - |p_3|.$$

Using our C++ implementation, grid sizes corresponding to 50, 70, 100, 140, and 200 nodes in each dimension were tried with a low order accurate scheme (first order space and time, hereafter referred to as the “(1,1)” scheme) and a high resolution scheme (fifth order space and second order time, hereafter the “(5,2)” scheme). On the eight million node finest grid—only around 10% of which is being actively updated on any one timestep by the local level set algorithm—execution time for the (5,2) scheme was about eighteen hours on a Sun UltraSparc II with lots of memory. Reducing the grid size in half results in the expected eightfold savings in memory and time; hence, the coarsest grid takes only fifteen minutes with the (5,2) scheme.

Results are visualized¹ by the zero level isosurface of the unsafe reachable set G , shown in Figure 2. On the left is a head-on view of the (5,2) solution. On the right is a zoomed overhead view of the point of the bulge computed by the (1,1) scheme for several grid sizes.

We compare solutions on the four coarser grids to the solution on the finest grid, using linear interpolation on the finest grid if necessary. Figure 3 demonstrates that the scheme is converging to the finest grid’s solution of (1) at approximately a linear rate in both average error and pointwise maximum error. We cannot expect to show a higher order convergence rate because of the linear interpolation used to evaluate the error.

Two conclusions can be drawn from Figures 2 and 3. First, for this example, low order schemes are not at all competitive in terms of accuracy with the (5,2) scheme: our (5,2) implementation can produce more accurate results in about fifteen minutes using only the coarsest grid. Second, the pointwise maximum error of the (5,2) scheme is always less than the grid spacing, so if a $50^{-1} = 2\%$ error is tolerable for this application, only this fastest, coarsest grid need ever be run.

¹Figure 2 and Figure 8 visualize some level set surfaces as triangular meshes; these are not the meshes on which the Hamilton-Jacobi PDE was solved, but rather an artifact of three dimensional Matlab visualization techniques.




Figure 2: Reachable Set for Aircraft Collision Avoidance Example

Three Mode Conflict Resolution

Consider the following conflict resolution problem between two aircraft ($i \in \{1, 2\}$), which is presented in [17]. Let $(x_r, y_r, \psi_r) \in \mathbb{R}^2 \times [-\pi, \pi)$ represent the relative position and orientation of aircraft 2 with respect to aircraft 1. In terms of the absolute positions and orientations of the two aircraft, (x_i, y_i, ψ_i) for $i = 1, 2$, it may be verified that $x_r = \cos \psi_1(x_2 - x_1) + \sin \psi_1(y_2 - y_1)$, $y_r = -\sin \psi_1(x_2 - x_1) + \cos \psi_1(y_2 - y_1)$, $\psi_r = \psi_2 - \psi_1$, and thus

$$\begin{aligned}\dot{x}_r &= -v_1 + v_2 \cos \psi_r + \omega_1 y_r \\ \dot{y}_r &= v_2 \sin \psi_r - \omega_1 x_r \\ \dot{\psi}_r &= \omega_2 - \omega_1\end{aligned}\tag{3}$$




Figure 3: Convergence of (5,2) Scheme to Finest Grid Solution (J_n is the solution $J(x, t)$ on a grid size of n)




Figure 4: Two aircraft in two modes of operation: in the first and third figures, the aircraft follow a constant heading and velocity (in *Mode 1*); in the second figure, the aircraft follow a half circle (in *Mode 2*).




Figure 5: In q_1 both aircraft follow a straight course, in q_2 a half circle, and in q_3 both aircraft return to a straight course.

where v_i is the velocity along the body axis of aircraft i , and ω_i is the rate of change of heading. A conflict occurs when a pair of aircraft incur a lateral spacing of less than 5 nautical miles; for safety of the maneuver, the relative position (x_r, y_r) must remain outside of the interior of the 5 mile radius disk centered at the origin $\{(x_r, y_r, \psi_r) : x_r^2 + y_r^2 \leq 5^2\}$.

Consider the maneuver illustrated in Figure 4, with the protocol of the maneuver defined as follows: the aircraft are nominally in *Mode 1*, in which the aircraft fly straight and level with given constant velocities v_i and constant heading ψ_i ; at a certain relative separation distance each aircraft turns to its right, follows an arc of a circle in *Mode 2*, in which the velocity is held at a prescribed constant value v_i , until it intersects its original trajectory, then turns to its right and returns to its desired trajectory. The model allows instantaneous changes in heading which is an obvious abstraction of the true aircraft dynamics and is treated here in order to clearly illustrate the algorithm; a model with heading capture dynamics is presented in [17]. In each mode, the continuous dynamics may be expressed in terms of the relative motion of the two aircraft (3): in *Mode 1*, $\omega_i = 0$ for $i = 1, 2$ and in *Mode 2*, $\omega_i = 1$ for $i = 1, 2$. We assume that both aircraft switch modes simultaneously, so that the relative orientation ψ_r is constant. This assumption simply allows us to display the state space in two dimensions, making the results easier to present. The problem statement is therefore: generate the relative distance between aircraft at which the aircraft may switch safely from *Mode 1* to *Mode 2*, and a turning radius r in *Mode 2*, to ensure that the 5 nautical mile separation is maintained.

The dynamics of the maneuver can be encoded by the hybrid automaton of Figure 5, where q_1 corresponds to *Mode 1* before the maneuver, q_2 corresponds to *Mode 2* “avoid mode”, and q_3

corresponds to *Mode 1* after the avoid maneuver has been completed. There is one discrete control input σ_1 , such that the switch from $\sigma_1 = 0$ to $\sigma_1 = 1$ triggers the transition from q_1 to q_2 . The transition from q_2 to q_3 is required to take place after the aircraft have completed a half circle: note that with $\omega_i = 1$, for $i = 1, 2$, it takes π time units to complete a half circle. The continuous state space is augmented with a timer $z \in \mathbb{R}^+$ to force this transition. Let $x = (x_r, y_r, \psi_r, z)^T$. At each transition, both aircraft change heading instantaneously by $\pi/2$ radians; we represent this with the standard rotation matrix $Rot(\frac{\pi}{2})$. As discussed in the previous section, we assume that v_1 is controllable, and v_2 is the disturbance input with known bounds. Safety is defined in terms of the unsafe set:

$$G_0 = \{q_1, q_2, q_3\} \times \{x \in X : x_r^2 + y_r^2 \leq 5^2\}$$

which encodes the fact that the two aircraft should never come within 5 nautical miles of each other.

The solution of the three mode conflict resolution example is presented in Figure 6. Each of the twelve graphs in Figure 6 is a plot in the (x_r, y_r) axes, with each of the four rows representing an iteration of the algorithm, and the three columns corresponding to discrete states q_1, q_2 , and q_3 . The computation is initialized with the set G_0 , which is shown as the dark shaded disks. A fixed point is reached after three iterations. The dark shaded region in the fourth row, first column of Figure 6 is the most interesting: it is a plot of the unsafe set of states in q_1 . As long as the relative position of aircraft 2 with respect to aircraft 1 is not inside this region, then there exists a control policy such that the conflict is resolved. If the relative position is inside the unshaded (white) region, then switching instantaneously will lead to a conflict; the aircraft must remain in q_1 until the relative state enters the light shaded region, then either switch to q_2 , or remain in q_1 .

Aircraft landing example

The autopilots of modern jets are highly automated systems which assist the pilot in constructing and flying four-dimensional trajectories, as well as altering these trajectories online in response to air traffic control directives. The autopilot typically controls the throttle input and the vertical and lateral trajectories of the aircraft to automatically perform such functions as: acquiring a specified altitude and then leveling, holding a specified altitude, acquiring a specified vertical climb or descend rate, automatic vertical or lateral navigation between specified way points, or holding a specified throttle value. The combination of these throttle-vertical-lateral modes is referred to as the *flight mode* of the aircraft. A typical autopilot has several hundred flight modes – it is interesting to note that these flight modes were designed to automate the way pilots fly aircraft manually: by controlling the lateral and vertical states of the aircraft to set points for fixed periods of time, pilots simplify the complex task of flying an aircraft. Those autopilot functions which are specific to aircraft landing are among the most safety critical, as it is extremely difficult, if not impossible, for the pilot to safely take over control of the aircraft when the aircraft is close to the ground. Thus, the need for automation designs which guarantee safe operation of the aircraft has become paramount. Testing and simulation may overlook trajectories to unsafe states: “automation surprises” have been extensively studied [18] *after* the unsafe situation occurs, and “band-aids” are added to the design to ensure the same problem does not occur again. We believe that the ability to compute accurate reachable sets inside the aerodynamic flight envelope under flight mode switching, may help prevent the occurrence of automation surprises.

A simple point mass model for aircraft vertical navigation is used, which accounts for lift L , drag

D , thrust T , and gravity mg (see [6] and references therein). State variables are aircraft height z , horizontal position x , velocity $V = \sqrt{\dot{x}^2 + \dot{z}^2}$ and flight path angle $\gamma = \tan^{-1}(\frac{\dot{z}}{\dot{x}})$. Inputs are thrust T and angle of attack α , where aircraft pitch $\theta = \gamma + \alpha$. The equations of motion can be expressed as follows:

$$\frac{d}{dt} \begin{bmatrix} V \\ \gamma \\ x \\ z \end{bmatrix} = \begin{bmatrix} \frac{1}{m}[T \cos \alpha - D(\alpha, V) - mg \sin \gamma] \\ \frac{1}{mV}[T \sin \alpha + L(\alpha, V) - mg \cos \gamma] \\ V \cos \gamma \\ V \sin \gamma \end{bmatrix} \quad (4)$$

The functions $L(\alpha, V)$ and $D(\alpha, V)$ are modeled based on empirical data [19] and Prandtl's lifting line theory [20]:

$$L(\alpha, V) = \frac{1}{2}\rho S V^2 C_L(\alpha), \quad D(\alpha, V) = \frac{1}{2}\rho S V^2 C_D(\alpha),$$

where ρ is the density of air, S is wing area, and $C_L(\alpha)$ and $C_D(\alpha)$ are the dimensionless lift and drag coefficients.

In determining $C_L(\alpha)$ we will follow standard auto-lander design and assume that the aircraft switches between three fixed flap deflections $\delta = 0^\circ$, $\delta = 25^\circ$ and $\delta = 50^\circ$ (with slats either extended or retracted), thus constituting a hybrid system with different nonlinear dynamics in each mode. This model is representative of current aircraft technology; for example, in Airbus cockpits the pilot uses a lever to select among four predefined flap deflection settings. We assume a linear form for the lift coefficient $C_L(\alpha) = h_\delta + 4.2\alpha$, where parameters $h_{0^\circ} = 0.2$, $h_{25^\circ} = 0.8$ and $h_{50^\circ} = 1.2$ are determined from experimental data for a DC9-30 [19]. The value of α at which the vehicle stalls decreases with increasing flap deflection: $\alpha_{0^\circ}^{\max} = 16^\circ$, $\alpha_{25^\circ}^{\max} = 13^\circ$, $\alpha_{50^\circ}^{\max} = 11^\circ$; slat deflection adds 7° to the $\alpha^{\max}$ in each mode. The left side of Figure 7 gives a graphical summary of the possible configurations. The drag coefficient is computed from the lift coefficient as [20] $C_D(\alpha) = 0.041 + 0.045C_L^2(\alpha)$ and includes flap deflection, slat extension and gear deployment corrections. So for a DC9-30 landing at sea level and for all $\alpha \in [-5^\circ, \alpha_\delta^{\max}]$, the lift and drag terms in (4) are given by

$$L(\alpha, V) = 68.6 (h_\delta + 4.2\alpha)V^2 \quad D(\alpha, V) = (2.81 + 3.09 (h_\delta + 4.2\alpha)^2)V^2$$

Flap deflection dynamics model: In reality, the decision to move from one deflection setting to another can occur at any time, but approximately 10 seconds are required for a 25° degree change in flap deflection. For our preliminary implementation, we have chosen to ignore the continuous dynamics associated with discrete mode switching, allowing the flaps and slats to move instantly to their commanded positions. However, if such instantaneous controlled switches were always enabled then the system would be zeno; therefore, we introduce Delay, or *transition* modes $0t$, $25t$ and $50t$, which use the envelopes and flight dynamics of the regular modes $0u$, $25d$ and $50d$ and include a timer to model the actual delay in flap change (the discrete automaton is shown on the right side of Figure 7). A regular mode may make a controlled switch to a transition mode, so flight dynamics can be changed instantly. Transition modes have only a timed switch at $t = t_{\text{delay}}$, so controlled switches will be separated by at least t_{delay} time units and the system is nonzeno. For the executions shown below, $t_{\text{delay}} = 0.5$ seconds.

Landing: The aircraft enters its final stage of landing close to 50 feet above ground level ([19, 21]). Restrictions on the flight path angle, aircraft velocity and touchdown (TD) speed are used to

determine the initial safe set E_0 :

$$\left\{ \begin{array}{ll} z \leq 0 & \text{landing or has landed} \\ V > V_{\delta}^{\text{stall}} & \text{faster than stall speed} \\ V < V^{\text{max}} & \text{slower than limit speed} \\ V \sin \gamma \geq \dot{z}_0 & \text{limited TD speed} \\ \gamma \leq 0 & \text{monotonic descent} \end{array} \right\} \cup \left\{ \begin{array}{ll} z > 0 & \text{aircraft in the air} \\ V > V_{\delta}^{\text{stall}} & \text{faster than stall speed} \\ V < V^{\text{max}} & \text{slower than limit speed} \\ \gamma > -3^\circ & \text{limited descent flight path} \\ \gamma \leq 0 & \text{monotonic descent} \end{array} \right\} \quad (5)$$

We again draw on numerical values for a DC9-30 [19]: stall speeds $V_{0u}^{\text{stall}} = 78$ m/s, $V_{25d}^{\text{stall}} = 61$ m/s, $V_{50d}^{\text{stall}} = 58$ m/s, maximal touchdown speed $\dot{z}_0 = 0.9144$ m/s, and maximal velocity $V^{\text{max}} = 83$ m/s. For passenger comfort, the aircraft's input range is restricted to $T \in [0 \text{ kN}, 160 \text{ kN}]$ and $\alpha \in [0^\circ, 10^\circ]$.

The interior of the surface shown in the first row of Figure 8 represents E_0 for each of the $0u$, $25d$ and $50d$ modes. The second row of the figure shows the safe envelope E when there is no mode switching. Portions of E_0 are excluded from E for two reasons. States near $z = 0$ correspond to low altitudes and are too close to the ground at steep flight path angles to allow control inputs time to prevent the plane from crashing. States close to the stall velocity correspond to low speeds where there is insufficient lift and the flight path angle becomes steeper than that allowed by the flight envelope. This latter condition holds throughout the very narrow range of speeds allowed in mode $0u$, with the result that only post-touchdown states ($z \leq 0$) are controllable in this mode. The third row shows how E can be increased if switches are permitted (for example, mode $0u$ becomes completely controllable). Mode $50d$ is the best to be in for landing and there is no difference in E with or without switching enabled. The fourth row shows slices of the set in the third row, taken at $z = 3$ meters. The light grey regions are unsafe G and the dark grey are safe E . The figure shows that modes $0u$ and $25d$ are safe only because there exists a discrete switch to a safe state in another mode.

4.7 Overapproximations of Reachable Sets

In the reachable set computation of the previous section, computational complexity introduced by increasing the dimension of the continuous state is a limiting factor: the technique suffers from exponential growth with respect to the continuous dimension. To challenge this curse of dimensionality, we are investigating techniques for computing overapproximations to the reachable set of states.

In [22], we devise and implement a method based on projecting the reachable set of a high dimensional system into a collection of lower dimensional subspaces where computational costs are much lower. We formulate a method to evolve the lower dimensional reachable sets such that they are each an overapproximation of the full reachable set, and thus their intersection will also be an overapproximation of the reachable set: the method uses a set of lower dimensional Hamilton-Jacobi PDEs for which, in any projection, the set of disturbance inputs is augmented with the unmodeled dimensions. For the three dimensional aircraft collision avoidance problem of the previous section, the result of performing the computation in the (x_r, y_r) space, and then back-projecting this set to form a cylindrical overapproximation of the reachable set in three dimensions, is computed in about 40 seconds.

In addition, we have studied polytopic overapproximations. Here, the main idea is to propagate a set of supporting hyperplanes of the initial set such that, at any time instant, the bounded

intersection of these hyperplanes (known as a polytope) includes the reachable set. The faces of the polytope constitute a level set of the special polytopic level set function. This approach can be successfully applied to feedback linearizable nonlinear systems, linear dynamic games, and systems that can be represented as linear systems with norm bounded perturbations [23]. Reduces the three dimensional collision avoidance problem computation time from 15 minutes (time to compute exact set) to about 1 second [23]. To the best of our knowledge, the proposed method computes an over-approximation of the unsafe set faster than any other known method [24].

A second technique is to overapproximate nonlinear dynamics by linear dynamics with bounded disturbance, and to overapproximate reachable sets by polyhedra. Propagation of sets is then simply described by the propagation, through linear dynamics, of linear inequalities describing the polyhedral set. With this technique however, comes the problem of the potentially unlimited number of inequalities needed to describe the polyhedral estimate. This problem has been addressed in the literature, yet previous solutions fall short in several ways: exact pruning of polytopes [25] requires complex data structures to be maintained and the number of “active” hyperplanes could still be unlimited; the use of simple polyhedral shapes for bounding the polyhedral set [26] is itself computationally intensive; the ellipsoidal overapproximation to the reachable set [27, 28, 29], while an elegant and efficient representation of the set, is generally too conservative an overapproximation.

We have developed a technique which combines the efficiency of the ellipsoidal data structure with the tightness of polyhedral overapproximation. The heart of our technique is a novel pruning procedure based on the Maximum Volume Ellipsoid (MVE) contained in the polyhedron. The basic idea of the method is to rank the constraints by their distance (measured in the norm induced by the ellipsoid) to the center of the MVE, and delete the ones that appear to be least relevant. Conceptually, constraints that are far away from the center are less relevant than those that are close. This approach will allow us to perform guaranteed pruning (where only provably redundant constraints are deleted).

For example, consider a discrete-time linear system (which represents a linear overapproximation of the nonlinear systems of the previous section):

$$x(t+1) = Ax(t) + Bw(t)$$

where $x(t) \in \mathbb{R}^n$ is the state vector, $w(t) \in \mathbb{R}^d$ is a bounded disturbance, and A and B are matrices with consistent dimensions. We assume that $w(t) \in \mathcal{W}(t)$ for some bounded convex set $\mathcal{W}(t)$. Assume that $x(t)$ is known to lie in a polyhedron $\mathcal{P}(t)$, where $\mathcal{P} \subseteq \mathbb{R}^n$ is the intersection of a finite number of half-spaces,

$$\mathcal{P} = \bigcap_{i=1}^m \mathcal{H}_i$$

Using the parameterization of half-spaces introduced above, we can represent $\mathcal{P}$ in the form

$$\mathcal{P} = \{x \mid f_i^T x \leq g_i, i = 1, \dots, m\} \equiv \{x \mid Fx \preceq g\} \quad (6)$$

where $\preceq$ denotes componentwise inequality. Now, an ellipsoid is defined as the image of a unit ball under an affine transformation

$$\mathcal{E} = \{Pu + q \mid \|u\| \leq 1\} \quad (7)$$

We assume that A is invertible, which is always the case if the dynamics is discretized from some

continuous-time system, and we propagate the set $\mathcal{P}(t)$ through the dynamic matrix A :

$$\begin{aligned} A\mathcal{P}(t) &= \{x \mid f_i(t)^T A^{-1}x \leq g_i(t), \ i = 1, \dots, m(t)\} \\ &= \{x \mid F(t)A^{-1}x \leq g(t)\}. \end{aligned}$$

To account for the bounded disturbance, we need to find the polytope $A\mathcal{P}(t) + B\mathcal{W}(t)$. An easily computed outer approximation of it is given by

$$\begin{aligned} \tilde{\mathcal{P}}_{t+1} &= \{z \mid f_i(t)^T A^{-1}z \leq g_i(t) + \|f_i(t)^T A^{-1}B\|_1, \ i = 1, \dots, m(t)\} \\ &= \{z \mid \tilde{F}(t)z \leq \tilde{g}(t)\}. \end{aligned}$$

As $\mathcal{P}$ evolves through hybrid dynamics, its representation grows in dimension. To limit the computational requirements for storing and operating on $\mathcal{P}$, it is therefore necessary to occasionally simplify the description of $\mathcal{P}$ by dropping, or discarding, constraints. We will call this procedure polyhedral pruning.

The maximum volume inscribed ellipsoid (MVE)

We are interested in finding the maximum volume ellipsoid that is contained in a polyhedron. Let $\mathcal{E}$ be an ellipsoid of the form (7), and $\mathcal{P}$ be a polyhedron parameterized as in (6). Then, $\mathcal{E} \subset \mathcal{P}$ if and only if

$$\sup_{x \in \mathcal{E}} f_i^T x \leq g_i, \quad i = 1, \dots, m$$

Using the parameterization of the ellipsoid, this condition can be expressed as

$$\sup_{\|u\|_2 \leq 1} f_i^T Pu + f_i^T q \leq g_i, \quad i = 1, \dots, m$$

Maximization of the left hand side expression gives the equivalent containment condition

$$\|Pf_i\|_2 + f_i^T q \leq g_i, \quad i = 1, \dots, m$$

Hence, we can find the MVE inside a polyhedron by solving the convex optimization problem

$$\begin{aligned} &\text{maximize} \quad \log \det P \\ &\text{subject to} \quad P = P^T \succeq 0 \\ &\quad \quad \quad \|Pf_i\|_2 + f_i^T q \leq g_i, \quad i = 1, \dots, m \end{aligned}$$

This problem follows in the general class of maxdet problems, which can be solved globally and efficiently using the algorithms described in [30]. Moreover, specialized algorithms for computing the MVE have been proposed in [31, 32]. The MVE gives an optimal (in terms of volume) ellipsoidal approximation of the polyhedron. Clearly, the MVE also provides an inner bound on the polyhedron, since $\mathcal{E}_{\text{mve}} \subseteq \mathcal{P}$. More surprisingly, perhaps, is that the MVE also provides an *outer* bound of the polyhedron: the ellipsoid obtained by scaling the MVE a factor n around its center is guaranteed to contain the polyhedron:

$$\mathcal{E}_{\text{mve}} \subseteq \mathcal{P} \subseteq q_{\text{mve}} + n(\mathcal{E}_{\text{mve}} - q_{\text{mve}})$$

This is a dual result to the celebrated Löwner-John ellipsoid [33, 34] (which treats the minimum volume ellipsoid that contains a convex set). Thus, the MVE gives a good approximation of a

polyhedron, and provides simultaneous inner and outer bounds.

Polyhedral pruning using the MVE

The MVE also gives a natural measure for ranking the relevance of the constraints that define $\mathcal{P}$. For a halfspace $\mathcal{H}_i = \{x \mid f_i^T x \leq g_i\}$, we define the *relevance ranking* η_i as the largest factor by which the ellipsoid can be enlarged and yet still lie within the halfspace:

$$\eta_i = \max\{\alpha \mid q_{\text{mve}} + \alpha(\mathcal{E}_{\text{mve}} - q_{\text{mve}}) \subseteq \mathcal{H}_i\}$$

It is easy to verify that the relevance ranking is given by

$$\eta_i = (g_i - f_i^T q) / \|Pf_i\|_2$$

This is, in fact, nothing more than the minimum distance from the halfspace to the center of the ellipsoid in the P^2 metric. The relevance ranking has the following properties:

1. $\eta_i \geq 1$ for all i
2. $\eta_i = 1 \Rightarrow f_i^T x \leq g_i$ active
3. $\eta_i \geq n \Rightarrow f_i^T x \leq g_i$ redundant

We define a constraint to be *redundant* if the polyhedron $\mathcal{P}$ does not change when the constraint is removed. We will say that the polyhedron is *minimal* if it does not contain any redundant constraints, and we will say that a constraint is *active* if $f_i^T x^* = g_i$ for some $x^* \in \mathcal{P}$. A constraint is (strongly) redundant if and only if it is inactive for all $x \in \mathcal{P}$. The last property allows us to discard redundant constraints. There is a gray zone $\eta_i \in (1, n)$ for which the relevance ranking neither proves that a constraint is active nor that it is redundant. In practice, guaranteed pruning can often be too conservative in that it fails to discard many redundant constraints. We therefore use the ranking in a heuristic pruning procedure: given $\mathcal{P} = \{x \mid f_i^T x \leq g_i \mid i = 1, \dots, m\}$, we compute $\mathcal{E}_{\text{mve}}$, sort the inequalities by their relevance ranking, η_i , and keep the k most relevant (the ones with smallest η_i). This yields an approximate simplified polyhedron $\hat{\mathcal{P}}$, with $\mathcal{P} \subseteq \hat{\mathcal{P}}$. We believe this algorithm to be an effective method for computing fast polyhedral overapproximations of reachable sets, in very high dimensions.

We have also developed two methods for computing fast overapproximations to reachable sets for verifying safety properties of continuous and hybrid systems. These approximative methods will reduce computation times by orders of magnitude, and thus allow us to treat problems of higher dimension than before.

4.8 Summary

We have presented and numerically validated a tool for determining accurate approximations of reachable sets for hybrid systems with nonlinear continuous dynamics and adversarial continuous and discrete inputs. By developing convergent approximations of such complex systems, we will be better able to synthesize aggressive but safe controllers. As an example, the six mode auto-lander shows that for envelope protection purposes the safest control decisions are to switch directly to full flap deflection, but to maintain airspeed until touchdown. With the summary data from the reachability analysis, such decisions can be made based on local state information; without it the

auto-lander may not detect that low speeds—while still within the flight envelope—lead inevitably to unsafe flight path angles. In this arena, we have made real progress on problems in a few dimensions, where we have close to exact solutions. We show that the computation may be made faster as we loosen the numerical approximation, however, even with such approximation, these methods will not scale beyond five or six continuous dimensions.

In response to this problem, we are developing methods that gracefully extend to problems with a much higher dimension of the continuous state (these methods have been shown to work well in dimension up to 100). While these are approximate methods, they are, however, non-heuristic: when they determine that a reachable set does not intersect some other set, the result is certain.

Finally, in order to test our reachability tools, we are designing and building a system of two Unmanned Aerial Vehicle (UAVs), with on-board sensing, navigation, and control systems enabling automatic flight. The Stanford DragonFly UAV Project currently consists of two 10-foot wingspan, fixed wing, unmanned aircraft, each equipped with GPS, inertial sensors, and on-board computers running QNX Neutrino. These will be described in an ensuing section.

5 Research Areas 3 and 4: Design, Implementation, and Validation on Stanford DragonFly testbed and on Boeing SEC Platform (F-15, T-33)

In order to validate the algorithms described above, we computed the reachable set for a two-aircraft conflict, in which one aircraft, the “blunderer”, is assumed to fly into the path of the other aircraft, the “evader”. Here, we assume that the evader can perform one of only three “emergency escape maneuvers (EEMs)”: this gives a more conservative reachable set than if the evader were given full range of maneuvers, yet it provides a realistic and practical test for collision avoidance. The reachable set, called the “blunder zone”, is computed for the three EEMS: the maneuvers are shown in Figure 9. This example was motivated by the problems of Closely Spaced Parallel Approaches (CSPA) of civilian aircraft into airports with closely spaced runways.

5.1 Implementation and Validation

The schematic for the implementation of the new algorithm is shown in Figure 10. The algorithm runs on the ‘own aircraft’. It takes speed and heading information from both aircraft and computes the danger zone.

With the position information from both aircraft, it then performs the check of whether the boundary of the danger zone encroaches the position of the own aircraft. If so, it activates the necessary EEM. During the EEM, speed and heading information of the blundering aircraft will not be required unless they change beyond the limits expected.

Some issues and considerations regarding the implementation of the algorithm are as follows.

- **Communication dropouts:** It is reasonable to expect communication dropouts over a wireless link. The algorithm takes this into account and guarantees safety despite dropouts.
- **Simplification:** As it is, the danger zone boundary, for the most part, is given in the form of a series of parameterized explicit analytical expressions (only one case involves a bisection

search algorithm). To simplify the computation further, a slightly conservative approximation by a polygon is used, as shown in Figure 11.

Using Microsoft Visual C++ ver 6.0 on a Dell Pentium 3 machine, each computation of this approximation of the danger zone can run in 0.05 sec.

- **Sensor inaccuracy:** Sensor inaccuracy is accounted for assuming WAAS-like characteristics. This is achieved by ‘pushing out’ the sides of the polygon as shown in Figure 11.
- **Cockpit display:** A cockpit display of the danger zones to the pilot have been developed by Jennings [35], as shown in Figure 12.

The danger zone computation is based on a 2-dimensional kinematic model which is engineered to contain realistic aircraft trajectories. These constitute approximate aircraft models. A validation of the implementation of the algorithm is thus required. Also, it is useful in verifying the algorithm’s strengths.

The implementation of the algorithm has been validated through 1,106,091 simulations using high dimensional dynamic aircraft models representing a Piper Saratoga as blunderer and a Cessna Caravan as evader. This choice of aircraft is because of the collaboration work with Jennings [35] who performed piloted simulation studies of various displays (including these danger zones) using these aircraft, although the motivation of the algorithm is towards commercial aircraft. The simulation is performed in 3-D using non-linear dynamic models whose basic control inputs are desired bank angle, lift coefficient and thrust. The algorithm is run on-line on the Cessna Caravan. It computes the danger zone boundary around the Piper Saratoga and checks the Cessna Caravan’s position with respect to the boundary. Upon encroachment, the Cessna Caravan is made to evade according to the part of the boundary that encroaches it. The horizontal separation between the two aircraft at the closest point of approach, together with the result of each run, is recorded.

The objective of the simulation runs is to validate the algorithm with realistic aircraft models under realistic conditions involving turbulence, sensor noise and ADS(B) drop outs. While in the simulation it is necessary to choose only a finite number of blunder trajectories, recall that the analysis considered all blunder trajectories permitted by the model and bounds on control inputs. The validation approach is adapted from Honeywell’s work [36]. As presented in the Honeywell approach [36, 37], three metrics are used to evaluate the validation results. They are the probabilities of

1. Successful Alerts (SA): this refers to events where alerts are given, the EEMs are executed, and no collisions occur; and if the EEMs were not executed, collisions would have occurred.
2. Unnecessary Alerts (UA): this refers to events where alerts are given, the EEMs are executed, and no collisions occur; and if the EEMs were not executed, no collision would have occurred either. If the alert were given because the other aircraft is blundering, the alert will not be deemed as unnecessary. However, such cases are still categorized as UAs in order to be consistent with the literature [36, 37].
3. Collision, Given an Alert (CGA): this refers to events where alerts are given, the EEMs are executed, and yet, collisions occur; and if the EEMs were not executed, no collisions would have occurred.

The results, together with those of the Honeywell AILS Segmented [36] and MIT Probabilistic [38] Alerting Algorithms, and the corresponding formulae for the metrics are given in Table 1. *In the*

Table 1: Metrics for three algorithms: % of runs.

		750ft	1700ft		2500ft		
	Formula	CSPA	CSPA	MIT	CSPA	AILS	MIT
SA	$\frac{CD+FA}{CD+FA+MD+IC+LA}$	100%	100%	99.5%	100%	99.9%	99.3%
UA	$\frac{IC+FA}{CD+FA+MD+IC+LA}$	98.78%	77.14%	55.5%	69.02%	93.7%	48.7%
CGA	$\frac{IC+LA}{CD+FA+IC+LA}$	0%	0%	0.5%	0%	0.1%	0.7%

table, the acronym CSPA refers to the new algorithm demonstrated in this paper, and is referred to as the CSPA algorithm, while AILS refers to the Honeywell algorithm and MIT refers to the MIT algorithm. To compare the results with the other two algorithms, some background is first given. The Honeywell algorithm is designed for independent approaches with runway spacings of 2500 ft or greater. Thus, only the results for a runway spacing of 2500 ft are available. The thresholds used in the Honeywell algorithm to trigger alerts are obtained empirically through simulating a large of number of parallel approaches with blunders. The thresholds picked are the ones that satisfactorily prevent collision. It is not stated in the Honeywell paper [36] what ‘satisfactory’ quantitatively corresponds to. The MIT algorithm, on the other hand, is designed so that the probability of collision in the event of a blunder is 0.001. This criterion is considered to be the acceptable level of risk in the Precision Runway Monitor System studies [39].

The CSPA algorithm has 100% SA and 0% CGA whereas both the Honeywell and MIT algorithms have some Collisions, Given Alerts (CGAs). With regard to Unnecessary Alerts (UA), the CSPA algorithm has a lower UA rate than the Honeywell algorithm but a higher one than the MIT algorithm. This latter comparison may appear to be a disadvantage. However, as explained earlier, alerts in the UA category would be considered necessary when they are raised by blunders that have the potential to later cause a collision should their trajectories be modified mid-course. This viewpoint is also expressed by Honeywell in their paper [36]. The blunders that were correctly rejected (CR) by the CSPA algorithm are those that did not have the potential *as yet* to cause a collision. In CSPA, it is more important to ensure that False Alarms (FAs) do not occur during normal aircraft approaches with normal flight technical error (FTE). With the CSPA algorithm, no False Alarms will be raised so long as the aircraft maintain their longitudinal separation according to the MLS.

The MIT algorithm might have produced better SA and CGA results if it were designed for a lower collision-probability value. However, it must be noted that it still relies on a specified blunder trajectory set, whereas the CSPA algorithm does not.

In simulation, the new algorithm has been successfully demonstrated to work according to expectations and is able to guarantee safety as long as the blunder trajectories are captured by the model and are within the bounds on the control inputs assumed.

5.2 Flight Demonstration

This section describes the flight demonstration of the implementation of the danger zone computation algorithm on the Stanford DragonFly UAVs and the Boeing testbed of F-15 and T33 aircraft and discusses the results.

5.2.1 Dual UAV Demonstration

In this flight demonstrations, DF2 is designated the Own Ship which will run the CSPA algorithm and conduct the evasive maneuver when required. DF3 is designated the Other Ship which will blunder in its ‘final approach’. Segment 1 of the Danger Zone is not used. Segment 2 is used in its place. In the demonstration, both aircraft will conduct the closely spaced parallel approach autonomously. DF3 will be programmed to blunder in its approach at a certain point. The algorithm will run on board DF2 and will trigger an EEM according to the danger zone computation. When triggered to conduct the EEM, DF2 will fly the EEM autonomously.

Prior to the flight demonstration, a ground taxiing demonstration is conducted. For the ground taxiing demonstration, ground taxiing controllers had to be designed. In addition, a navigation algorithm which blends GPS measurements with IMU measurements through an Extended Kalman Filter is designed. The ground taxiing phase is conducted prior to the in-flight phase so that valuable experience in conducting the demonstration is gained without risking the UAVs in the flight phase. The benefits of the ground taxiing phase are as follows:

- **Testing Software Readiness.** The ground taxiing phase tests that the CSPA algorithm, the server and the client processes can communicate with each other and run simultaneously at the desired rates. It also shows that the CSPA algorithm has been coded correctly to generate the results that match the Matlab version of the algorithm.
- **Partially Testing System Readiness.** This phase also partially checks that the dual UAVs are ready for operations. This includes checking that the necessary logistics are available for engine starting, battery charging, refuelling, etc, that the UAVs can be switched between manual and flight computer control, that the installed sensors work, and that the procedures for operations run smoothly.

After the ground taxiing demonstration, parameter identification of the flight dynamics of the aircraft are obtained experimentally and the flight controllers for airspeed hold, heading hold, bank angle hold, altitude hold, and straight line tracking are designed and flight tested. Thereafter, the flight demonstration could commence. For the demonstrations, each UAV has a top level architecture that is shown in Figure 13. Below, the ground taxiing results are first presented followed by the flight results.

Ground Taxiing Demonstration

Figure 14 shows the demonstration conducted at Moffett Federal Airfield. About 10 runs are made, each with a different relative start position. In all of the runs, the distance between the two UAVs at the closest point of approach is greater than the threshold of 8m. The results of two runs are shown in Figures 15 and 16. In Figure 15, the EEM turns out to be the one associated with Segment 2 of the danger zone, i.e., acceleration and turn to 45 deg. In Figure 16, the EEM is the one associated with Segment 3 of the danger zone, i.e., coast and turn to 60 deg. For the former case, the separation distance between the two UAVs almost exceeded the threshold. This is because DF3 turned at almost 10 deg/s in the experiment whereas a lower value of 5 deg/s is assumed to be its maximum turn rate in computing the danger zone. In the latter case, the separation distance between the two UAVs is above the threshold.

The completion of the ground demonstration shows that the CSPA algorithm can run on this dual UAV test bed in real-time.

In-Flight Demonstration

To reiterate, the purpose of the in-flight demonstration is to show that the algorithm can be implemented and in particular, to demonstrate that it can run in real-time.

Figure 17 shows the demonstration conducted over Moffett Federal Airfield. Over twenty runs are made with DF3 blundering and DF2 performing the EEM as triggered by the danger zone computation algorithm. All runs are successful. Sample results are shown in Figure 18. The former emulates an independent approach scenario where there is no restriction in the longitudinal spacing between the two aircraft during normal approaches. The latter emulates a dependent approach scenario.

Note that the separation threshold is not violated. The flight demonstration shows the implementation of the algorithm and in particular, demonstrates that it can run in real-time.

5.3 SEC Capstone Demonstration

The Stanford algorithm of “dual vehicle conflict resolution” for the Software Enabled Control (SEC) Capstone demonstration sponsored by DARPA and Boeing was evaluated successfully on the Boeing Open Control Platform (OCP). The technology is based on the Reachable Set Computation done by Ian Mitchell and Rodney Teo. The experiment involved two aircraft (F-15 and T-33), where one aircraft is assigned as pursuer (F-15) and the other as evader (T-33). The flight tests were performed at Edward Airforce Base in June 2004, and Stanford had six sorties during three flight experiments. The algorithm demonstrated the guaranteed safety (maintaining a minimum separation distance) for every single flight test performed (there was a single failure of the algorithm due to an improper initialization of the OCP).

Figure 19 shows snapshots of a flight experiment displayed on the Experimental Controller Display installed on both pursuer and evader. As seen in **Figure 19(a)**, as pursuer blunders into the path of the evader, the algorithm computes the corresponding danger zone based on position and velocity information of two aircraft, yet the evader does not commence an emergence escape maneuver (EEM). As the pursuer enters a caution zone (which is 3,000 ft away from the danger zone) as shown in **Figure 19(b)**, the algorithm alerts the pilot by displaying a warning message as well as changing the color of the danger zone. Finally, once pursuer touches the danger zone, the algorithm commences an EEM immediately based on which part of the danger zone the pursuer touches. **Figure 19(d)** shows the resulting separation distance between aircraft during this experiment and it is seen that the minimum separation distance (1 mile) was well maintained.

References

- [1] M. Johansson and L. Xiao, “On optimal control over packet-oriented communication links,” in *38th Allerton Conference on Communication, Control and Computing*, 2000.
- [2] L. Xiao, M. Johansson, H. Hindi, S. Boyd, and A. Goldsmith, “Joint optimization of communication rates and linear systems,” *IEEE Transactions on Automatic Control*, vol. 48, pp. 148–153, January 2003.
- [3] R. Alur, C. Courcoubetis, T. A. Henzinger, and P.-H. Ho, “Hybrid automata: An algorithmic approach to the specification and verification of hybrid systems,” in *Hybrid Systems* (R. L. Grossman, A. Nerode, A. P. Ravn, and H. Rischel, eds.), LNCS, pp. 366–392, New York: Springer Verlag, 1993.
- [4] C. J. Tomlin, J. Lygeros, and S. Sastry, “A game theoretic approach to controller design for hybrid systems,” *Proceedings of the IEEE*, vol. 88, no. 7, pp. 949–970, July 2000.
- [5] I. Mitchell and C. Tomlin, “Level set methods for computation in hybrid systems,” in *Hybrid Systems: Computation and Control* (B. Krogh and N. Lynch, eds.), LNCS 1790, pp. 310–323, Springer Verlag, 2000.
- [6] C. J. Tomlin, *Hybrid Control of Air Traffic Management Systems*. PhD thesis, Department of Electrical Engineering, University of California, Berkeley, 1998.
- [7] I. Mitchell, A. M. Bayen, and C. J. Tomlin, “Validating a Hamilton-Jacobi approximation to hybrid system reachable sets,” in *Hybrid Systems: Computation and Control* (M. D. D. Benedetto and A. Sangiovanni-Vincentelli, eds.), LNCS 2034, pp. 418–432, Springer Verlag, 2001.
- [8] O. Maler, A. Pnueli, and J. Sifakis, “On the synthesis of discrete controllers for timed systems,” in *STACS 95: Theoretical Aspects of Computer Science* (E. W. Mayr and C. Puech, eds.), no. 900 in LNCS, pp. 229–242, Munich: Springer Verlag, 1995.
- [9] E. Asarin, O. Maler, and A. Pnueli, “Symbolic controller synthesis for discrete and timed systems,” in *Proceedings of Hybrid Systems II, Volume 999 of LNCS* (P. Antsaklis, W. Kohn, A. Nerode, and S. Sastry, eds.), Cambridge: Springer Verlag, 1995.
- [10] H. Wong-Toi, “The synthesis of controllers for linear hybrid automata,” in *Proceedings of the IEEE Conference on Decision and Control*, (San Diego, CA), 1997.
- [11] I. Mitchell, A. M. Bayen, and C. J. Tomlin, “Computing reachable sets for continuous dynamic games using level set methods,” *IEEE Transactions on Automatic Control*, January 2002. Submitted.
- [12] M. G. Crandall, L. C. Evans, and P.-L. Lions, “Some properties of viscosity solutions of Hamilton-Jacobi equations,” *Transactions of the American Mathematical Society*, vol. 282, no. 2, pp. 487–502, 1984.
- [13] A. M. Bayen and C. J. Tomlin, “A construction procedure using characteristics for viscosity solutions of the Hamilton-Jacobi equation,” in *Proceedings of the IEEE Conference on Decision and Control*, (Orlando, FL), pp. 1657–1662, 2001.

- [14] M. G. Crandall and P.-L. Lions, “Two approximations of solutions of Hamilton-Jacobi equations,” *Mathematics of Computation*, vol. 43, no. 167, pp. 1–19, 1984.
- [15] D. Peng, B. Merriman, S. Osher, H. Zhao, and M. Kang, “A PDE based fast local level set method,” *Journal of Computational Physics*, vol. 165, pp. 410–438, 1999.
- [16] J. A. Sethian, *Level Set Methods and Fast Marching Methods*. New York: Cambridge University Press, 1999.
- [17] C. J. Tomlin, I. Mitchell, and R. Ghosh, “Safety verification of conflict resolution maneuvers,” *IEEE Transactions on Intelligent Transportation Systems*, vol. 2, no. 2, pp. 110–120, 2001. June.
- [18] A. Degani, *Modeling Human-Machine Systems: On Modes, Error, and Patterns of Interaction*. PhD thesis, Department of Industrial and Systems Engineering, Georgia Institute of Technology, 1996.
- [19] I. M. Kroo, *Aircraft Design: Synthesis and Analysis*. Stanford, California: Desktop Aeronautics Inc., 1999.
- [20] J. Anderson, *Fundamentals of Aerodynamics*. New York: McGraw Hill Inc., 1991.
- [21] United States Federal Aviation Administration, *Federal Aviation Regulations*, 1990. Section 25.125 (landing).
- [22] I. Mitchell and C. J. Tomlin, “Overapproximating reachable sets by Hamilton-Jacobi projections,” *Journal of Scientific Computing*, vol. 19, pp. 323–346, December 2003.
- [23] I. Hwang, D. Stipanovic, and C. Tomlin, “Polytopic approximations of reachable sets to linear dynamic games and a class of nonlinear systems,” in *In the Proceedings of the AACC American Control Conference*.
- [24] I. H. D. M. Stipanović and C. J. Tomlin, “Computation of an over-approximation of the backward reachable set using subsystem level set functions,” *Dynamics of Continuous, Discrete & Impulsive Systems Series A: Mathematical Analysis*, vol. 11, pp. 399–411, 2004.
- [25] N. A. Bruinsma and M. Steinbuch, “A fast algorithm to compute the $\mathbf{H}_\infty$ -norm of a transfer function matrix,” *Systems and Control Letters*, vol. 14, no. 5, pp. 287–293, 1990.
- [26] M. Milanese and G. Belforte, “Estimation theory and uncertainty intervals evaluation in presence of unknown but bounded errors – linear families of models and estimators,” *IEEE Transactions on Automatic Control*, vol. 27, no. 2, pp. 408–414, 1982.
- [27] F. C. Schweppe, “Recursive state estimation: Unknown but bounded errors and system inputs,” *IEEE Transactions on Automatic Control*, vol. AC-13, no. 1, pp. 22–28, 1968.
- [28] H. S. Witsenhausen, “Sets of possible states of linear systems given perturbed observations,” *IEEE Transactions on Automatic Control*, vol. 13, pp. 556–558, 1968.
- [29] N. Shishido and C. J. Tomlin, “Ellipsoidal approximations of reachable sets for linear games,” in *Proceedings of the IEEE Conference on Decision and Control*, (Sydney, Australia), 2000.

- [30] S.-P. Wu and S. Boyd, *SDPSOL: A Parser/Solver for Semidefinite Programming and Determinant Maximization Problems with Matrix Structure. User's Guide, Version Beta*. Stanford University, June 1996.
- [31] Y. Nesterov and A. Nemirovsky, *Interior-point polynomial methods in convex programming*, vol. 13 of *Studies in Applied Mathematics*. Philadelphia, PA: SIAM, 1994.
- [32] Y. Zhang, "An interior-point algorithm for the maximum-volume ellipsoid problem," Technical Report TR98-15, Rice University, Houston, Texas, 1998.
- [33] F. John, "Extremum problems with inequalities as subsidiary conditions," in *Fritz John, Collected Papers* (J. Moser, ed.), pp. 543–560, Birkhauser, Boston, Massachusetts, 1985.
- [34] M. Grötschel, L. Lovász, and A. Schrijver, *Geometric Algorithms and Combinatorial Optimization*. Algorithms and Combinatorics, New York: Springer-Verlag, 1988.
- [35] C. Jennings, *Threat Displays for Final Approach*. PhD thesis, Department of Aeronautics and Astronautics, Stanford University, 2003.
- [36] M. Jackson, P. Samanant, and C. Haissig, "Analysis of airborne alerting algorithms for closely spaced parallel approaches," *Air Traffic Control Quarterly*, vol. 9, December 2001.
- [37] J. K. Kuchar, "Methodology for alerting-system performance evaluation," *Journal of Guidance, Control and Dynamics*, vol. 19, pp. 438–444, March-April 1996.
- [38] B. Carpenter and J. Kuchar, "Probability-based collision alerting logic for closely-spaced parallel approach," in *Proceedings of the AIAA 35th Aerospace Sciences Meeting and Exhibit*, (Reno NV), January 1997. AIAA 97-0222.
- [39] E. M. Shank and K. M. Hollister, "A statistical risk assessment model for the precision runway monitor system," in *Proceedings of the ATCA Conference*, 1992.




Figure 6: The rows indicate the iterations of the algorithm, the columns indicate the discrete state. Each figure displays a projection of part of the unsafe set onto the (x_r, y_r) axes. A fixed point is reached after 3 iterations. Angular velocity $\omega_i = 1$ in q_2 .




Figure 7: Left: lift coefficient $C_L(\alpha)$ model for the DC9-30 [19]. Circles located at $(\alpha_\delta^{\max}, C_L(\alpha_\delta^{\max}))$ indicate the stall angle and the corresponding lift coefficient in each mode. Right: Discrete transition graph of slat and flap settings. Solid lines are controlled switches (σ_E in this version of the reachability problem) and dashed lines are uncontrolled switches (σ_G).




Figure 8: Maximally controllable safe envelopes for the multimode landing example. From left to right the columns represent modes $0u$, $25d$ and $50d$.




Figure 9: A sample danger zone result generated using 3 different EEMs.




Figure 10: Schematic showing Algorithm Implementation.




Figure 11: Conservative Approximation of Danger Zone for On-line Computation.




Figure 12: Cockpit Displays (for aircraft on the right) showing a normal approach (Courtesy of Chad W. Jennings [35]).




Figure 13: Top Level Architecture of each UAV.



Figure 14: Ground Taxiing Demonstration.




Figure 15: Ground Taxiing Demonstration: Accelerate, Turn 45 deg EEM.




Figure 16: Ground Taxiing Demonstration: Coast, Turn 60 deg EEM.



Figure 17: In-Flight Demonstration.






Figure 18: Flight Demonstration: Independent Approach Scenario.




(a) Pursuer blunders into the path of evader and the algorithm computes corresponding danger zone displayed as blue line segments.



(b) Pursuer enters a caution zone and a warning message is displayed and the color of the danger zone turns into yellow.



(c) Pursuer touches the danger zone and the algorithm commences an EEM.



(d) Separation distance between aircraft during the experiment.

Figure 19: Snap shots of the flight experiment for the Stanford dual vehicle conflict resolution algorithm on the Boeing SEC Platform (T-33, F-15).