

# UAV as a Reliable Wingman: A Flight Demonstration

S. Waydo, J. Hauser, R. Bailey, E. Klavins, R.M. Murray

## Abstract

We present the results from a flight experiment demonstrating two significant advances in software enabled control: optimization-based control using real-time trajectory generation and logical programming environments for formal analysis of control software. Our demonstration platform consisted of a human-piloted F-15 jet flying together with an autonomous T-33 jet. We describe the behavior of the system in two scenarios. In the first, nominal state communications were present and the autonomous aircraft maintained formation as the human pilot flew maneuvers. In the second, we imposed the loss of high-rate communications and demonstrated an autonomous safe “lost wingman” procedure to increase separation and re-acquire contact. The flight demonstration included both a nominal formation flight component and an execution of the lost wingman scenario.

## I. INTRODUCTION

In June of 2004 we demonstrated two new control technologies on board two jet aircraft as part of the DARPA Software Enabled Control (SEC) program’s capstone demonstration: optimization-based control using real-time trajectory generation and logical programming environments for formal analysis of control systems. We tested these technologies by implementing a reliable wingman scenario in which an Unmanned Air Vehicle (UAV) performed formation flying with a manned aircraft. The manned aircraft was an F-15 fighter jet, while the UAV surrogate was a T-33 jet trainer outfitted with X-45 UCAV avionics. During formation flight, the vehicles performed a set of coordinated actions using receding horizon control on the UAV to demonstrate real-time trajectory generation. After demonstrating coordinated formation flight, the UAV simulated loss of the high data rate link between the aircraft and begin executing a “lost wingman” protocol designed to (provably) achieve an increased separation between the aircraft. After safe separation was achieved, the high bandwidth link was restored and the UAV requested a rejoin and executed a safe rejoin maneuver. The demonstration scenario is depicted in Figure 1.

S. Waydo (corresponding author: [waydo@cds.caltech.edu](mailto:waydo@cds.caltech.edu)) and R.M. Murray are with the Control and Dynamical Systems department, California Institute of Technology, Pasadena, CA 91125.

J. Hauser and R. Bailey are with the Electrical and Computer Engineering department, University of Colorado, Boulder, CO 80309.

E. Klavins is with the Electrical Engineering department, University of Washington, Seattle, WA 98195.

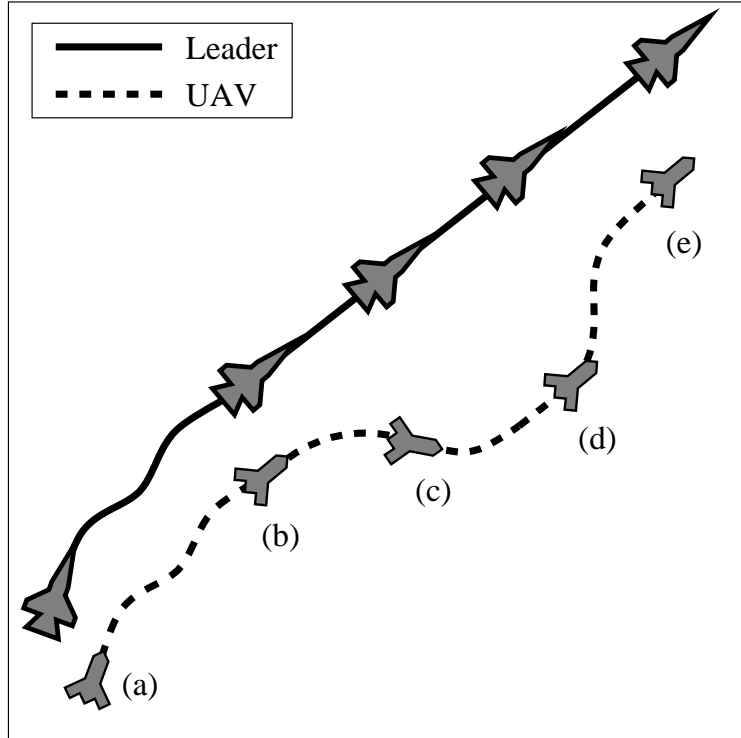


Fig. 1. Lost wingman scenario. At point (a) aircraft are flying in formation. At (b) the data link is lost and the UAV begins executing the lost wingman maneuver. At (c) the UAV receives a confirmation message from the leader and turns to match heading. At (d) the data link is restored and at (e) normal operation resumes.

### A. Research Outline

A nonlinear optimization-based receding horizon controller (RHC) was developed to fly the UAV in formation with the human-piloted aircraft. This controller was implemented using the Nonlinear Trajectory Generation (NTG) software [1], [2], [3], an efficient computational method that combines nonlinear control theory, B-spline basis functions, and nonlinear programming. We describe the design and implementation of this controller in Section II.

A Logical Programming Environment (LPE) framework was used to develop and verify the protocols for executing this scenario, ensuring that the problem specification was satisfied in presence of uncertainty in dynamics, timing and failures. Temporal logic was used to specify the desired performance and our protocol was implemented using the Computation and Control Language (CCL) [4]. In Section III we present the LPE framework and its implementation in CCL and in Section IV we describe its implementation for this flight demonstration.

Both NTG and CCL were interfaced to the Open Control Platform (OCP) [5], [6], a middleware development platform also developed as part of the SEC program to enable researchers to more easily implement their technologies in both simulated and actual embedded system environments. The OCP provided services for controlling and scheduling the execution of components, inter-component communication, and interfacing to the lower-level control system. Also included was an “Experiment Controller” interface through which the test conductor interacted with

the control software. This interface and the results from our simulations and the flight demonstration are presented in Section V.

### *B. Related Work*

Provably safe conflict resolution between aircraft has been extensively studied in the context of air-traffic control [7], [8]. The distinction here is that in addition to verifying the safety of the lost wingman maneuver, we reason about the operation of the underlying decision mechanism (the software) to ensure that the assumptions behind the conflict resolution safety proofs are correct.

Receding horizon control demands significant computing power and can exhibit poor convergence if implemented improperly. For this reason, most prior applications of this technology have been in the process control industry where dynamics are relatively slow. Increasingly inexpensive and powerful computing capabilities coupled with a better understanding of RHC's stability properties have more recently opened up applications to nonlinear systems with fast dynamics and have greatly revived interest in this area [9], [10].

Several other research teams participated in the SEC demonstration, including at least two applications of receding horizon control. Keviczky et al. developed an application programming interface (API) for receding horizon control within the OCP and demonstrated reference trajectory tracking with the UAV [11]. Schouwenaars et al. demonstrated the use of Mixed Integer Linear Programming (MILP) implemented in a receding horizon fashion for trajectory generation and task planning [12]. Neither of these demonstrations included a formation flight component; in the first case the UAV followed a pre-planned reference trajectory, while in the second the focus was on dynamic tasking of the UAV from the manned aircraft rather than on formation flight. No other demonstration included an LPE component.

### *C. Research Contributions*

The main contribution of this work was to prove the feasibility of our approach to both nonlinear receding horizon control and formal software modeling and analysis by implementing them on a real flight system operating in an uncertain environment. This flight experiment built on our previous work on applying these tools separately in simulation [13] and on smaller scale experimental testbeds [2], [14], but was the first time they have been joined in a single system. We demonstrated how these two control technologies can work in concert to provide robustness to uncertainty and failures in a dynamic environment and provided a proof-of-concept for how this approach can lead to realizing the goal of human-piloted and autonomous aircraft working together safely.

## II. RECEDING-HORIZON CONTROL WITH NONLINEAR TRAJECTORY GENERATION

### *A. The Nonlinear Trajectory Generation Package*

In receding horizon control (RHC), a finite-horizon constrained optimal control problem is solved starting from the current state. The controls of this trajectory are then applied for some portion of the horizon length, after which the optimization is repeated to provide the feedback necessary for stabilization and error correction. A

major challenge confronting the application of such techniques in systems with fast dynamics such as UAVs is that the algorithms require significant computational power to execute in a timely fashion. Increasingly powerful and affordable computing resources have recently brought the application of RHC techniques to aircraft within reach, but care in the problem formulation and solution is still required to maintain the speed and stability of the algorithms.

Nonlinear Trajectory Generation (NTG) [1], [3], [15] is a software package implementing an efficient computational method combining nonlinear control theory, B-spline basis functions, and nonlinear programming. There are three primary components to NTG. The first is to determine a parameterization such that the equations of motion can be mapped to a lower dimensional space (output space), converting dynamic constraints to algebraic ones. The cost and constraints in the optimal control problem can then also be mapped to the output space. The second is to parameterize each component of the output in terms of an appropriate B-spline polynomial. Finally, nonlinear programming is used to solve for the coefficients of the B-splines.

We here provide an brief overview of the RHC problem and its solution using NTG; a more complete description of this approach is described elsewhere [1], [2]. NTG has previously been used to stabilize and perform high-performance flight maneuvers on an indoor vectored-thrust flight experiment [2] and planar hovercraft-like vehicles [14].

In RHC, the optimal control  $u(\cdot; y_k) \in [0, T]$  for the current state  $y_k$  at time  $t_k$  is the solution to an optimal control problem with a scalar objective and constraints:

$$\begin{aligned} \min_u \int_0^T q(y(\tau), u(\tau)) d\tau + V(y(T)), \\ \dot{y}(t) = f(y(t), u(t)), y(0) = y_k, \\ lb_0 \leq \psi(y(0), u(0)) \leq ub_0, \\ lb_t \leq S(y(t), u(t)) \leq ub_t. \end{aligned} \tag{1}$$

The vectors  $S(y(t), u(t))$  and  $\psi$  are constraints on the trajectory and initial condition, respectively. The optimal control  $u(\cdot; y_k)$  is applied for some portion of the horizon time  $T$ , then the computations are repeated to yield a feedback control law.

To reduce the computational load of the optimization problem, we seek to map the equations of motion to a lower dimensional space. In particular, we want to find an output  $z = z_1, \dots, z_q$  of the form

$$z = \zeta(x, u, u^{(1)}, \dots, u^{(r)})$$

such that  $(x(t), u(t))$  can be determined completely from

$$(x, u) = \psi(z, z^{(1)}, \dots, z^{(l)})$$

for some map  $\psi$ . If such a mapping exists the system is said to be *differentially flat* and trajectories of  $z$  automatically result in dynamically feasible trajectories of  $(x, u)$  [16]. Note that even if a system is not differentially flat we may be able to find a map to an output space that reduces the dimensionality of the optimization problem.

Once a suitable output map is chosen, the outputs are parameterized in terms of B-spline basis functions [17]. Nonlinear programming is then used to solve for the B-spline coefficients that minimize the cost subject to the constraints in output space. NTG implements this strategy, taking as input the cost and constraints in terms of the outputs and their derivatives and the Jacobian of the cost and constraints with respect to the outputs and computing the optimal B-spline coefficients.

### B. Aircraft System Description

A Lockheed T-33A Shooting Star or “T-bird” equipped with a three axis autopilot was used as a UAV surrogate in the flight demonstration. In normal operation, an autopilot is used to reduce pilot workload allowing such flight modes as altitude and heading hold as well as waypoint navigation. The T-33 was not equipped with an actuator to manipulate the throttle but rather our airspeed commands were relayed through the pilot’s display so that the pilot was solely responsible for maintaining safe airspeeds.

For the purpose of the flight demonstration, the T-33 could be directed through low rate (2 Hz) autopilot commands specifying the desired altitude and heading (or heading rate). Airspeed commands could also be sent through the autopilot interface but these *commands* were only *suggestions* to the pilot as to the desired airspeed. Good judgment on the part of the pilot was indeed essential to the safety and success of the flight demonstration.

An interesting question and concern on our part was the extent to which the aircraft could be flown in a closed loop manner using this interface. On the one hand, the use of the autopilot as aircraft control interface presents a conceptually simple control task and, by restricting the class of maneuvers, offers the prospect of a guaranteed level of flight safety. On the other hand, the use of the autopilot interface introduces additional dynamics (including substantial delays) and nonlinear effects limiting the achievable performance of the UAV system.

With the given interface and flight safety requirements, we decided to restrict the desired operations to constant altitude maneuvering with airspeeds within to a reasonable range around a nominal design point. At constant altitude, the throttle manipulations that are needed for safe aircraft maneuvering are intuitive for any qualified pilot and thus easily implemented even when the control of the aircraft attitude is relegated to the autopilot.

We are thus interested in controlling the aircraft in a two dimensional plane (at constant altitude) for the purpose of formation flying including formation breakup and rejoin maneuvers. Neglecting wind (and using a flat earth model), the kinematics of constant altitude flight are given by

$$\begin{aligned}\dot{x} &= v \cos \chi \\ \dot{y} &= v \sin \chi\end{aligned}\tag{2}$$

where  $x$  and  $y$  are north and east displacements,  $v$  is the velocity, and  $\chi$  is the heading angle. The *velocity coordinate frame* is attached to the aircraft center of mass and oriented so that the  $x$ -axis lines up with the velocity vector. By differentiation, one finds that the acceleration expressed in the velocity frame consists of a longitudinal acceleration  $\dot{v}$  and a lateral acceleration  $v\dot{\chi}$ . For a conventional aircraft, the desired lateral acceleration is obtained by banking the aircraft. Idealizing this situation, we will suppose that the aircraft is flying at constant altitude with zero sideslip

angle (coordinated flight) and that the (normally small) angle of attack is equal to zero. In this case, the lateral acceleration is given by

$$v\dot{\chi} = g \tan \varphi$$

where  $\varphi$  is the roll angle and  $g$  is the acceleration of gravity. Fixing the velocity, we see that the lateral dynamics of this simplified model includes heading and roll angles,  $\chi$  and  $\varphi$ , and turn and roll rates,  $\dot{\chi}$  and  $\dot{\varphi}$ . One can imagine that the autopilot provides tracking of the commanded heading  $\chi$  (or turn rate  $\dot{\chi}$ ) by treating  $\varphi$  or  $\dot{\varphi}$  as a pseudo control and using a higher bandwidth regulator to provide close regulation the roll angle and rate, enforcing limits on the maximum roll rate and roll angle to ensure safe flight operations.

For the purposes of system design and performance evaluation, Boeing provided a simulation environment known as “DemoSim” providing precisely the same interface as used in the flight demonstration. Furthermore, an interface to the DemoSim dynamics was provided for use within the MATLAB/Simulink environment for dynamic analysis and control system design work. Using this environment, we were able to explore the range of maneuvers that would be possible in flight, verifying that the autopilot system (or at least the DemoSim model of it) functioned much as described above. A third order, local linear model describing the mapping from commanded heading to roll, roll rate, and heading was identified from the DemoSim model. This model, with roll rate saturation and time delay included, provided a fast simple model for early testing and evaluation of our control system design. We also identified a second order linear model mapping commanded heading to roll and heading angles for use within the control system. This model was used in developing a simple compensator so that  $\ddot{\chi}$  could be used as a pseudo-control.

### C. Localizer Intercept

A fundamental skill in aircraft navigation is the ability to fly the aircraft to a line in space that goes through a given  $x$ - $y$  point with a desired heading  $\chi$ . When this line is the extension of the runway centerline, it is called the *localizer* in reference to the (now seemingly ancient) navigational aid of the same name and the maneuver is called a *localizer intercept*. We use the localizer intercept (and maintenance) as the basis for formation rejoin and station keeping where the localizer specification varies as a function of the location and heading of the lead aircraft.

For simplicity of exposition, suppose that we would like to intercept a localizer through the point  $(x, y) = (0, 0)$  with heading  $\chi = 0$  and suppose that (through independent regulation) the velocity  $v$  is constant. The lateral dynamics are then described by

$$\begin{aligned} \dot{y} &= v \sin \chi \\ \dot{\chi} &= (g/v) \tan \varphi \\ \dot{\varphi} &= \dot{\varphi}_{cmd} \end{aligned} \tag{3}$$

where  $\dot{\varphi}_{cmd}$  is thought of as a control (or pseudo-control), possibly subject to actuator dynamics. Note that, with  $\dot{\varphi}_{cmd}$  as the control input, this system is differentially flat.

With this setup, the *localizer intercept problem* is to steer the aircraft lateral dynamics to  $(y, \chi, \varphi) = (0, 0, 0)$  with local exponential stability subject to constraints on the roll angle and roll rate. We provide a solution to this problem using a receding horizon control strategy.

To this end, suppose that we desire that the aircraft maneuvers with  $|\varphi| \leq \varphi_{\max}$  and  $|\dot{\varphi}| \leq \dot{\varphi}_{\max}$ . Based on experiments with DemoSim, we chose to use  $\varphi_{\max} = 30\pi/180$  radians and  $\dot{\varphi}_{\max} = 4.5\pi/180$  radians per second. With such limitations, it was much easier to intercept the localizer when we restricted the offset heading angle used to approach it requiring, say,  $|\chi| \leq \chi_{\max}$ . In practice, we found that  $\chi_{\max} = 15\pi/180$  radians worked well, providing a good tradeoff between rapid intercept and requiring excessive speed differential to avoid falling behind the lead aircraft. Working with limitations of this sort has the advantage that small angle approximations to the trigonometric functions will be quite accurate—even the linear approximation for  $\tan \varphi$  will be accurate to within 10%.

As the localizer intercept system (3) is controllable over the intended operating region, local exponential stability may be obtained through the use of receding horizon control with a quadratic cost functional.

We now proceed to specify precisely the optimization problem employed in the receding horizon control. Exploiting differential flatness, our optimal control problem (1) is the constrained calculus of variations problem

$$\begin{aligned} \min \quad & \int_0^T \|(y, \dot{y}, \ddot{y})(\tau)\|_Q^2 + \|y^{(3)}(\tau)\|_R^2 d\tau + \|(y, \dot{y}, \ddot{y})(T)\|_{P_1}^2 \\ & |\dot{y}(t)| \leq v \sin \chi_{\max} \\ & |\ddot{y}(t)| \leq g \varphi_{\max} \\ & |y^{(3)}(t)| \leq g \dot{\varphi}_{\max} \end{aligned} \tag{4}$$

where the constraints provide approximate satisfaction of the physical constraints proposed above. The positive definite symmetric weighting matrices  $Q$  and  $R$  were chosen to obtain, asymptotically, desirable transient dynamics. In fact, these matrices were designed by choosing a desired dynamics for the unconstrained chain of integrators  $y^{(3)} = u$  and solving an inverse optimal control problem by convex optimization. The weight on the terminal state,  $P_1$ , was chosen to be consistent with  $Q$  and  $R$  relative to the chain of integrator dynamics.

In particular,  $P_1$  is the solution of stationary Riccati equation for the corresponding infinite horizon LQR problem. A horizon length of  $T = 30$  seconds was found to work quite well, providing a sufficient preview to allow a smooth intercept of the localizer while satisfying the constraints. The choice of the horizon length and the design of the terminal cost, not to mention the design of the incremental cost, remain somewhat of an art, though theoretical notions that may help guide this process have been developed [18], [19].

#### D. Implementation

NTG relies on NPSOL [20], a numerical optimization library that uses sequential quadratic programming to solve nonlinear optimization problems, and PGS, the companion software to the Practical Guide to Splines [17], both of which are implemented in Fortran 77. These utilities were compiled into libraries that were linked to the main OCP executable. The optimization problem (4) was coded in C++ within the main executable and solved by calls

to these linked libraries. For the demonstration the NTG solution was computed and implemented at 2 Hz with a 30 second horizon length.

### III. FORMAL METHODS FOR SPECIFICATION AND ANALYSIS OF SOFTWARE

In this section we introduce our approach to using formal methods in computer science to specify and verify our reliable wingman algorithms. In particular we discuss the “lost wingman” scenario we implemented for the flight demonstration. The goal of this work is to be able to specify desired safety and correctness properties of our programs and then write our programs in such a way that we can formally prove they meet the specifications given a model of the world, including aircraft and controller dynamics as well as communications link properties, etc.

#### A. The Computation and Control Language

The Computation and Control Language (CCL) is a programming language designed to make it easy to specify programs together with a model of the environment in a way that makes it possible to prove theorems about them [13], [4]. Inspired by the UNITY formalism [21] and adapted to real-time control systems, CCL is a specification and implementation language for the operation of concurrent systems that interact both through dynamics and logical protocols. The advantages of this approach are several:

- The formalism of CCL and the analysis techniques allow us to analyze the dynamical behavior and the logical behavior each in an environment naturally suited to them.
- CCL is both a specification and implementation language, meaning that the protocols we analyze and the code we implement are nearly identical; in some cases they are one and the same.
- Reasoning about CCL specifications is done using standard logical tools amenable to the application of automated theorem proving software.
- We can model and reason about portions of the system using CCL specifications even if they will not be implemented as such.

For this demonstration we used CCL as both a modeling environment for the complete system and as our implementation language for a portion of the software. We summarize here a few important points about CCL; this material has been described in considerably more detail in previous work [13], [4].

#### B. Structure and Semantics of a CCL Specification

A CCL specification, or *program*,  $P$  consists of two parts  $I$  and  $C$ .  $I$  is a predicate on states called the initial condition.  $C$  is a set of *guarded commands* of the form  $g : r$  where  $g$  is a predicate on states and  $r$  is a *relation* on states. In a rule, primed variables (such as  $x'_i$ ) refer to the new state and unprimed variables to the old state. For example,

$$x < 0 : x' > y + 1$$

is a guarded command stating: If  $x$  is less than zero, then set the new value of  $x$  to be greater than the current value of  $y$  plus 1. If  $x$  is not less than zero, do nothing. Note that guarded commands can be non-deterministic, as



is the one above since it does not specify the exact new value for  $x$ , only that the new value must be greater than  $y + 1$ .

Programs are composed in a straightforward manner: If  $P_1 = (I_1, C_1)$  and  $P_2 = (I_2, C_2)$  then  $P_1 \circ P_2 = (I_1 \wedge I_2, C_1 \cup C_2)$ . This type of composition allows us to modularize our specifications into separate programs such as one for a finite state machine and one to model the dynamics. In the real system, then, we can implement only the state machine and as long as the real dynamics satisfy the CCL specification any proofs about the original specification will remain valid.

In addition to the initial condition and the collection of guarded commands, we need to specify the *semantics* we will use to interpret the specification. The semantics determine how commands are picked for execution, i.e. in what order and with what relative frequency. In general commands are picked in a nondeterministic way to model the uncertainty in how they may be interleaved in the actual system; the semantics we choose place restrictions on this nondeterminism. For the purposes of this paper, where the commands are picked at very close to the same rate, the *EPOCH* semantics will be sufficient [4]. Informally EPOCH semantics can be described as follows:

**CCL EPOCH Semantics:** Commands are chosen non-deterministically from  $C$ , but each command must be chosen once before any command is chosen again.

Thus, the execution of a system is divided into *epochs* during which each command is executed exactly once. This is an attempt to capture the small-time interleaving that may occur between processors that are essentially synchronized, as well as the nondeterministic ordering of commands that can occur when some are picked by some event-driven process while others are picked by a deterministic process. We generally view epochs as occurring at some fixed frequency, although that has not yet been explicitly modeled.

### C. Specification of Experimental System

We now turn to the task of specifying the demonstration system. Here we provide a few examples of the CCL specifications; the full details and a partial safety proof have been provided elsewhere [22].

1) *Dynamics and Semantics:* Using techniques described in [22], we specify two programs,  $F_{dyn}$  and  $T_{dyn}$ , to keep track of the dynamics of the F-15 and T-33, respectively, using the subscript 1 to denote the F-15 (so for example the position of the F-15 is  $(x_1, y_1)$ ) and 2 to denote the T-33. These programs are rules that periodically update the aircraft states according to the discrete-time version of Equation 2. We also constrain the execution to obey the EPOCH semantics, with the additional requirement that each epoch begin with the execution of the command describing the dynamics. In the actual system the execution of each epoch is triggered by an accurate software timer every  $\Delta t$  seconds and each epoch will complete before the next trigger, so this is a realistic model.

2) *Controllers:* Each aircraft runs a controller that at each update uses its own data plus what is known about the other aircraft to generate controls  $\mathbf{u}$ . We envision that this is accomplished by some function  $control : \mathcal{Q} \times \mathcal{Q} \rightarrow \mathcal{U}$  (where  $\mathcal{Q}$  is the state space of a single vehicle) that we will leave unspecified aside from basic assumptions required for proofs of safety properties [22]. This allows us to use any number of control techniques in the system to be

implemented while retaining the correctness of the safety proofs, provided the implemented controller satisfies the safety specification. In particular, we can implement a simple model of the controllers in CCL for analysis and test purposes and a more complex receding-horizon controller for the flight test and retain the validity of our analysis as long as both meet the higher-level specifications. On the T-33 we can implement this as a constraint on the optimization-based controller described in Section II.

$$\begin{array}{l} \text{Program } P_c \\ \hline \text{Initial } \mathbf{u} = (0, 0) \\ \text{Commands } \textit{true} : \mathbf{u}' = \textit{control}(\mathbf{x}_1, \mathbf{x}_2) \end{array}$$

As with the dynamics we denote the F-15 controller by  $F_c$  and the T-33 controller by  $T_c$ .

#### D. Communication

We suppose there are two communications links between the two aircraft, a “high-speed” data connection for state information and a less frequently used “low-speed” connection. The high-speed link is implemented by the lower-level command and control software, and because the CCL program on the T-33 interacts with this subsystem only to check if new data have arrived we abstract it into a command  $c_{data}$  that updates  $T_s$ , the next time data will be sent by the F-15, and  $T$ , the next time data will be received by the T-33. The send time  $T_s$  is incremented by  $\Delta T$  whenever data are sent, reflecting the fact that the F-15 sends data at a regular rate, while the receive time  $T$  can be anything in the interval  $(T_s, T_s + \tau_d]$ , reflecting the uncertain time delay in the system. We keep track of the status of the data link using the boolean variable  $data\_on$ .

We model the low-speed communications link using a mailbox with a queue and a nondeterministic time delay of up to  $\tau_c$ . When a message is sent, a record is added to the end of the queue containing a scheduled arrival time  $t_a \in (t, t + \tau_c]$ . We write  $send(i, y)$  as shorthand for the process of sending data  $y$  to vehicle  $i$ . We then have the predicate  $in(i)$  which is  $true$  if a message has arrived at vehicle  $i$ .

As will be seen below, the nature of the messages we send is such that we are only interested in the most recent message received. We use  $msg' = recv(i)$  to denote setting the new value of  $msg$  to the  $data$  field of the most recent record in  $queue_i$  for which  $t \geq t_a$  and then deleting from  $queue_i$  all messages for which  $t \geq t_a$ .

All of these communications are specified by the program  $P_{comm}$ :

$$\begin{array}{l} \text{Program } P_{comm} \\ \hline \text{Initial } T_s = t_0 \wedge T \in (T_s, T_s + \Delta T] \\ \text{Commands } c_{data} \equiv t > T \wedge data\_on : \\ \quad T' \in (T_s + \Delta T, T_s + \Delta T + \tau_d] \\ \quad \wedge T'_s = T_s + \Delta T \\ c_{msg,1} \equiv in(1) : msg'_1 = recv(1) \\ c_{msg,2} \equiv in(2) : msg'_2 = recv(2) \end{array}$$

#### IV. LOST WINGMAN PROTOCOL

The lost wingman protocol consists of two parts: a CCL state machine managing the T-33 and a detailed flight procedure for the pilot of the F-15. The flight procedure can also be written as a state machine, and so for purposes of analysis we can model the pilot's behavior using CCL as well.

##### A. T-33

The state machine running on the T-33 is depicted in Figure 2. The system has four modes denoted by the variable  $m_2$  in the programs:

- *normal*, for normal formation flight, abbreviated  $n$
- *lost<sub>1</sub>*, for the case where the T-33 has ceased receiving state data from the F-15 and has not yet received a confirmation of lost status, abbreviated  $l_1$ ,
- *lost<sub>2</sub>*, for the case where lost confirmation has been received from the F-15, abbreviated  $l_2$ , and
- *found*, for the case where the state data link has been re-acquired but normal formation flight has not resumed, abbreviated  $f$ .

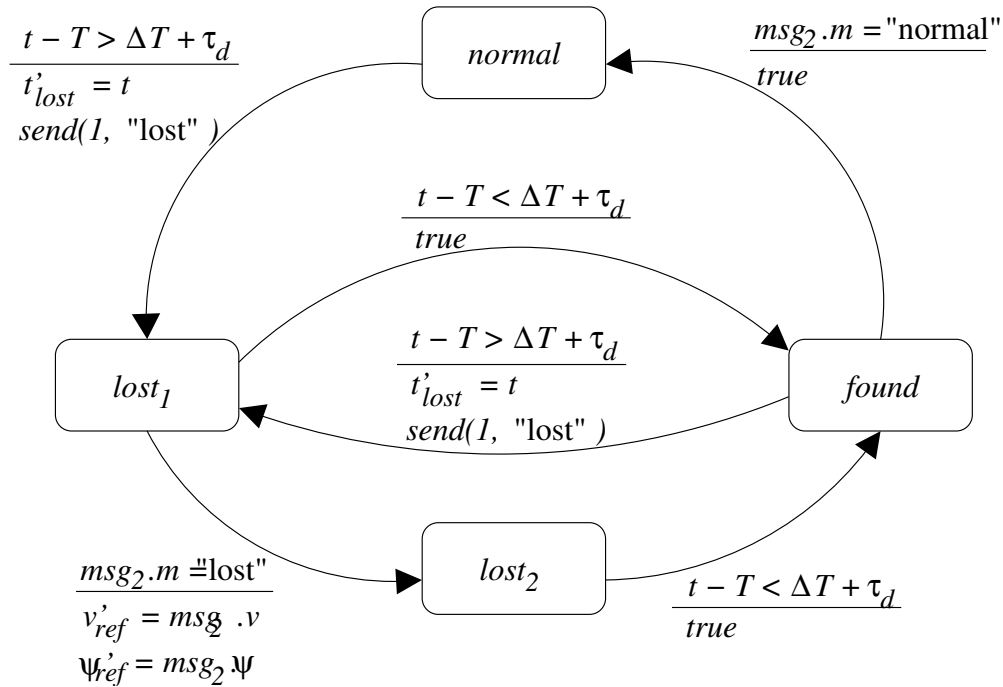


Fig. 2. T-33 state machine.

This flight demonstration was designed to test what happens when the T-33 enters  $lost_1$  mode from *normal* mode. The portion of the state machine specification involving this portion of operation is:

Program $T_{sm}$	
Initial	$m_2 = n$
Commands	$c_{lost} \equiv m_2 \in \{n, f\} \wedge t - T > \Delta T + \tau_d :$ $m'_2 = l_1 \wedge t'_{lost} = t$ $\wedge send(1, "lost")$
	$c_{found} \equiv m_2 \in \{l_1, l_2\} \wedge t - T < \Delta T + \tau_d :$ $m'_2 = f$
	$c_{lost2} \equiv m_2 = l_1 \wedge msg_2.m = "lost" :$ $m'_2 = l_2 \wedge v'_{ref} = msg_2.v$ $\wedge \psi'_{ref} = msg_2.\psi$
	...

### B. F-15

The F-15 state machine consists of just two modes (denoted by  $m_1$  in the specifications), *normal*, for normal formation flight, and *lost*, for the lost wingman scenario. In *normal* mode the pilot is free to fly at will within a performance envelope that ensures safe formation flight is possible, that is, an envelope defined by the T-33's flight capabilities. If the pilot receives a lost message from the T-33, the procedure is to transition to straight and level flight and transmit to the T-33 the resulting speed and heading (using a button on the rear-seat test conductor's control interface to do so automatically). Upon receiving and acknowledging a rejoin request and observing that the T-33 has rejoined formation safely, the pilot can resume *normal* operation.

### C. Implementation

While we may reason about specifications in which the rules are arbitrary relations on states, any code we write will be deterministic, meaning that all rules will be assignments. In this case there exists a CCL interpreter known as CCLi [4], [23] that can be used to execute CCL code. The implementation code for all of the programs described here is available online [24].

Because we have modularized the different aspects of our CCL specification into different programs, we are free to implement these specifications in whatever manner is most appropriate for each. For this demonstration, the T-33 state machine and low-speed communications protocols were implemented directly in CCL programs updated by the primary executable, while the high-speed communications and receding-horizon control were implemented in the main C++ source code.

For development and testing of our algorithms, we implemented a simulation of the system and our controllers in CCL in addition to the state management code used for the flight test. This allowed us to quickly iterate and test our designs in a simple desktop environment prior to integrating with the complex system required to fly the actual

aircraft. When it came time to integrate CCL with the flight system, the CCLi library was linked with the flight code so our CCL protocols could be periodically updated from within the demonstration executable. Implementing the CCL code in this way rather than recoding it in C++ (the language of the flight code) had two significant advantages. First, we could have 100% confidence that our flight protocols would execute just as those we analyzed and tested because the code was identical. Secondly, because the version of CCL we used is interpreted rather than compiled, this strategy allowed us to modify our protocols and re-test without having to go through the time consuming process of recompiling the main executable.

The F-15 state machine was implemented by the detailed flight procedures provided to the test crew specifying what actions to take upon receiving messages from the autonomous wingman. Messages from the F-15 to the T-33 were either automatic (in the case of the detailed state information) or sent by the test conductor pressing buttons on a laptop computer interface (described in Section V below). For the testing of our algorithms a keyboard controllable model of the F-15 was written in CCL and used to demonstrate the lost wingman scenario.

## V. FLIGHT DEMONSTRATION

### A. *Demonstration Overview*

The main objectives of the flight experiment were to

- 1) demonstrate the use of NTG to perform coordinated formation flight, and
- 2) demonstrate the use of CCL to manage the lost wingman scenario.

Satisfying these objectives required us to also successfully integrate both of these software packages with the OCP, which provided the inputs to the lower-level controllers.

The UAV test platform was a T-33 jet aircraft outfitted with the avionics package of the X-45 UCAV unmanned aircraft. The OCP comprised the main structure of the control code, collecting the aircraft's state information and providing hooks through which our software could process this data and generate control inputs. The available control inputs were ground speed, altitude, and turn rate. Altitude and turn rate were passed to an autopilot that performed the low-level control, while ground speed was controlled by the software instructing the pilot to make manual adjustments to the throttle. For our tests the altitude set point was held fixed and only turn rate and ground speed were controlled. The OCP also handled communications between the two aircraft, automatically providing the transfer of state information needed for formation flight and allowing us to pass other information such as that required to execute the lost wingman scenario.

Our test scenario proceeded in three stages. In the first stage, the F-15 pilot was instructed to fly a nondeterministic path within the test range to test the station-keeping performance of the NTG-based controller. In the second stage, the test conductor eliminated the high-rate data link between the aircraft and the lost wingman behavior was observed. Finally, the high-rate link was restored and the T-33 was allowed to execute a rejoin maneuver and resume station-keeping. At all times the aircraft maintained at least a 1-mile separation for safety reasons, and a test pilot on board the T-33 handled flying the aircraft to and from the test area and was available to intervene in the case of problems with the autonomous system.

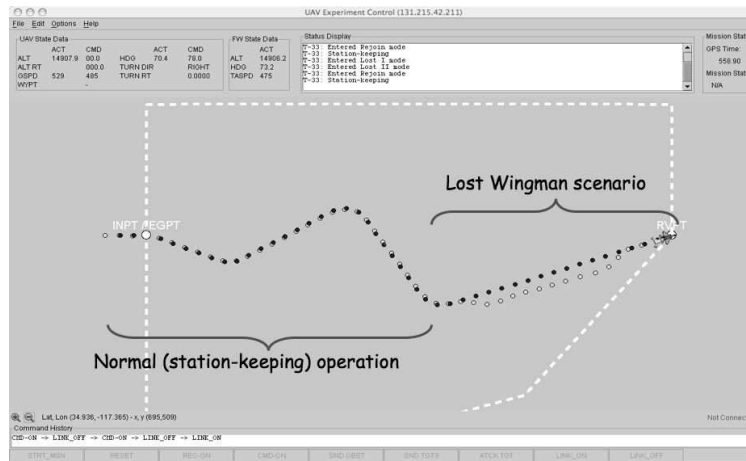


Fig. 3. Screen capture depicting the Experiment Controller interface on the UAV computer. Results from a full demonstration simulation including the lost wingman behavior are shown. The path of the F-15 is marked with dark circles while the path of the T-33 is marked with lighter circles.

On both aircraft the test engineer (riding in the aircraft rear seat) interacted with the control software using a Java-based “Experiment Controller” interface. This interface, shown with simulation data in Figures 3 and 4, was common between all of the teams participating in the flight demonstration, each of which provided an xml file defining the particular buttons needed for their test in addition to the built-in functions to start and stop the test software execution. For our demonstration, on the T-33 side the test engineer had buttons to cut off and restore the high-speed data link to test the lost wingman performance. On the F-15 side, the test engineer had a button to confirm the lost wingman status upon receipt of a lost message from the T-33 and another button to approve a formation rejoin request from the T-33 upon restoration of normal state communications. In addition to these buttons the interface provided the test engineer with a view of both aircraft’s states and any messages being sent between the aircraft in the course of the test.

### B. Simulation Results

To test the system performance prior to the flight demonstration we used the DemoSim environment described above. In this way we were able to test our software in an integrated environment with the T-33 and F-15 flight software running on separate laptop computers networked via Ethernet. To test the formation flight capabilities of the T-33, an interface was provided to have the F-15 follow a pre-planned trajectory while the T-33 followed. Figures 3 and 4 give an example of typical simulation results obtained for a full demonstration including the lost wingman scenario.

### C. Flight Test Results

The flight control experiment was executed 7 times over 5 flights between June 17 and June 25, 2004. We here describe the results from Tests 6 and 7 on June 25, in which the aircraft had to be diverted away from other traffic,

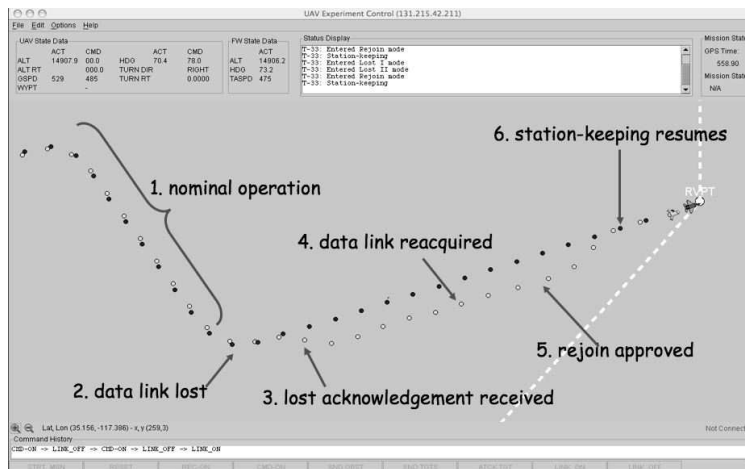


Fig. 4. Expanded view of the lost wingman portion of the simulation depicted in Figure 3.

resulting in a nondeterministic set of maneuvers that nicely stressed the formation keeping capabilities of the T-33. Test 6 included an execution of the lost wingman protocol. Figures 5 and 6 provide an overview of the flight paths of both aircraft taken over the experiments, with the position of the F-15 marked at 1-minute intervals.

#### D. Control Performance

As can be seen from Figures 5 and 6, these flight tests included several maneuvers induced by the avoidance of other traffic in the test airspace. The execution of the lost wingman scenario occurred 4–6 minutes into Test 6 (in the dashed box in Figure 5) and was the only occasion in which the T-33 deviated significantly from the flight path of the F-15 (by design of the protocol).

Figures 7 and 8 depict the heading of the two aircraft during Tests 6 and 7, respectively. The T-33's heading matched that of the F-15, subject to a significant time delay incurred by the control and communications software. The significant deviations at about 300 and 380 seconds in Test 6 correspond to the Lost Wingman divert and rejoin maneuvers. The other small deviations were caused by the F-15 making tighter turns than the autopilot of the T-33 would allow the controller to follow. This behavior was also the report of the flight test conductors.

Naturally the physical autopilot was developed to reduce pilot workload, especially during routine flight modes including navigation tasks such as waypoint following and altitude regulation. For the purpose of safe and intuitive (to the pilot) operation, the autopilot restricts the aggressiveness of the maneuvers that it attempts. For example, experiments with the supplied simulation model DemoSim indicated that the autopilot would restrict the roll rate to approximately 4 or 5 degrees per second. In flight, we observed typical roll rates of 5 to 6 degrees per second with a maximum of approximately 9 degrees per second.




Fig. 5. Flight path of the two vehicles for the entire Flight Test 6 on June 25, 2004. Points of interest: (1) Experiment start, (2) Divert maneuvers, (3) Lost Wingman start, (4) Lost Wingman complete, (5) Experiment finish. Diamonds mark the position of the F-15 at 60 second intervals. Dashed box is area of interest for Figure 9.

### E. Lost Wingman Performance

Figure 9 depicts the portion of Test 6 in which the Lost Wingman protocol was executed. At point (1) the high-rate state communications were cut off and the T-33 automatically transitioned to its lost state upon detecting this condition. From point (1) to point (2) the T-33 executed a tight turn to ensure it would stay clear of the F-15 while broadcasting its lost status via the low speed link. At point (2) the lost confirmation message was received by the T-33, having been sent by the F-15 test engineer pressing a confirmation button on the laptop interface after straight and level flight was achieved. At this time the T-33 turned to match the heading commanded by the F-15. At point (3) the state communications were restored and the T-33 automatically requested a formation rejoin, which was granted by the F-15. At point (4) the T-33 completed its rejoin maneuver and nominal operation resumed. As can be seen from the similarity between this plot and the schematic of Figure 1 and the simulation results of Figure 4 the lost wingman procedure executed exactly as designed.

## VI. LESSONS LEARNED

Transitioning the technologies described here from idealized simulations and high precision laboratory experiments to a real flight system with significant uncertainties and performance limitations proved to be an immensely challenging task. Many lessons were learned about such practical implementation.

Much of the success of this flight demonstration can be attributed to our control design properly respecting the performance constraints on the flight system imposed by having to command the aircraft through the autopilot. The RHC framework allowed us to explicitly capture these constraints in our control design, resulting in achieving the






Fig. 6. Flight path of the two vehicles for the entire Flight Test 7 on June 25, 2004. Points of interest: (1) Experiment start, (2) Experiment finish. Diamonds mark the position of the F-15 at 60 second intervals.

best performance that could be expected given the capabilities of our UAV surrogate. We found that we obtained the best performance when we ensured that the constraints were appropriately conservative, that is, when we made sure our controllers would not attempt to push the flight performance too close to the allowable limits.

As described above, the choice of parameters for the RHC problem (horizon length, costs, etc.) is still in large part an art requiring significant expertise and extensive simulation. The tactic we used here, namely choosing desired dynamics for the simplified unconstrained system and solving an inverse optimal control problem, was found to be a very effective approach to this problem.

These successes depended heavily on the accuracy of the DemoSim environment provided by Boeing for experiment development. The entire flight test program took place over about two weeks after extensive testing of our control software in simulation. Because of the required testing and the flight demonstration schedule, no revision of our control software after the start of flight tests was feasible. This meant that our control design was entirely based on the system identification we were able to do on the DemoSim model. Much more than in laboratory environments where iteration based on “flight” tests is the norm, a highly accurate simulation environment was critically important.

On top of the accuracy of the simulation module, the ability to implement our controllers and interact with DemoSim first in a MATLAB environment and later in a full OCP software environment was crucial. The MATLAB interface to DemoSim enabled the system identification upon which our controller design was based, a process which would have been massively more cumbersome had we needed to operate the full flight software to do experiments for identification. This interface also allowed us to more quickly iterate on our control design without the overhead involved in the flight software implementation. Because the MATLAB simulation was one and the same as the

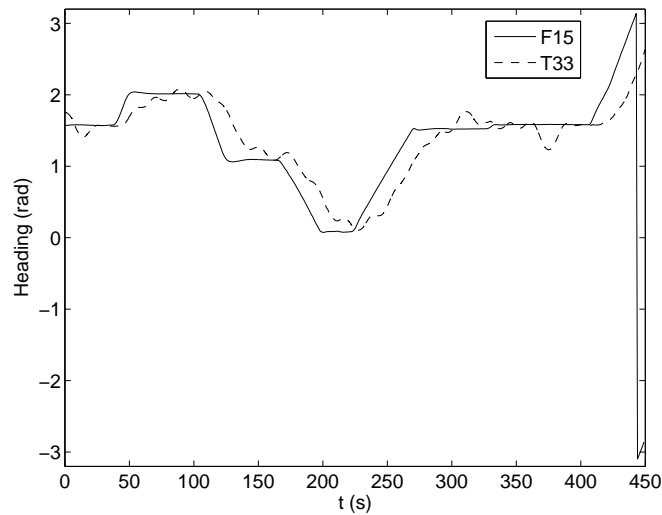


Fig. 7. Aircraft heading during flight demonstration 6. Time lag between F-15 and T-33 trajectories is the result of both inherent time delay in the flight system (communications delay, etc.) and the large ( $\approx 1$  mi) separation distance between the aircraft.

simulation used with the flight software, we were able to implement our final control design in the OCP with high confidence in its success.

#### ACKNOWLEDGMENTS

We would like to thank Brian Mendel and Jared Rossen at the Boeing Company for their extensive help with the integration of our software with the Open Control Platform.

This work was supported by the DARPA SEC program, F33615-98-C-3613, as well as a Fannie and John Hertz Foundation Fellowship (to S.W.).

#### REFERENCES

- [1] M. Milam, "Real-time optimal trajectory generation for constrained dynamical systems," Ph.D. dissertation, California Institute of Technology, 2003.
- [2] M. Milam, R. Franz, J. Hauser, and R. Murray, "Receding horizon control of a vectored thrust flight experiment," *IEE Proceedings on Control Theory and Applications*, vol. 152, no. 3, pp. 340–348, 2005.
- [3] R. Murray, J. Hauser, A. Jadbabaie, M. Milam, N. Petit, W. Dunbar, and R. Franz, Eds., *Software-Enabled Control: Information Technology for Dynamical Systems*. Wiley - IEEE Press, 2003, ch. Online Control Customization via Optimization-based Control, pp. 149–174.
- [4] E. Klavins, "A language for modeling and programming cooperative control systems," in *Proceedings of the International Conference on Robotics and Automation*, 2004.
- [5] J. Paunicka, B. Mendel, and D. Corman, "The OCP - an open middleware solution for embedded systems," in *Proceedings of the American Control Conference*, 2001.
- [6] J. Paunicka, B. Bendel, and D. Corman, Eds., *Software-Enabled Control: Information Technology for Dynamical Systems*. Wiley - IEEE Press, 2003, ch. Open Control Platform: a software platform supporting advances in UAV control technology, pp. 38–62.
- [7] C. 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.




Fig. 8. Aircraft heading during flight demonstration 7.

- [8] Z.-H. Mau, E. Feron, and K. Billimoria, "Stability of intersecting aircraft flows under decentralized conflict avoidance rules," in *AIAA Conference on Guidance, Navigation, and Control*, 2000.
- [9] D. Mayne, J. Rawlings, C. Rao, and P. Scokaert, "Constrained model predictive control: Stability and optimality," *Automatica*, vol. 36, pp. 789–814, 2000.
- [10] R. Findeisen and F. Allgöwer, "An introduction to nonlinear model predictive control," in *21st Benelux Meeting on Systems and Control*, 2002.
- [11] T. Keviczky, A. Packard, O. Natale, and G. Balas, "Application programming interface for real-time receding horizon control," in *Conference on Decision and Control and European Control Conference*, 2005.
- [12] T. Schouwenaars, M. Valenti, E. Feron, and J. How, "Implementation and flight test results of milp-based uav guidance," in *IEEE Aerospace Conference*, 2005.
- [13] E. Klavins, "A formal model of a multi-robot control and communication task," in *42nd IEEE Conference on Decision and Control*, December 2003.
- [14] J. Chauvin, L. Sinigra, and R. Murray, "Nonlinear trajectory generation for the caltech multi-vehicle wireless testbed," in *European Control Conference*, 2003.
- [15] "Nonlinear trajectory generation." [Online]. Available: <http://www.cds.caltech.edu/~ntg/>
- [16] M. Fliess, J. Lévine, P. Martin, and P. Rouchon, "Flatness and defect of nonlinear systems: introductory theory and examples," *International Journal of Control*, vol. 6, pp. 1327–1360, 1995.
- [17] C. D. Boor, Ed., *A Practical Guide to Splines*. Springer, 2001.
- [18] A. Jadbabie, J. Yu, and J. Hauser, "Unconstrained receding-horizon control of nonlinear systems," *IEEE Transactions on Automatic Control*, vol. 46, no. 5, pp. 776–783, 2001.
- [19] A. Jadbabie and J. Hauser, "On the stability of receding horizon control with a general terminal cost," *IEEE Transactions on Automatic Control*, vol. 50, no. 5, pp. 674–678, 2005.
- [20] "Npsol." [Online]. Available: [http://www.sbsi-sol-optimize.com/asp/sol\\_product\\_npsol.htm](http://www.sbsi-sol-optimize.com/asp/sol_product_npsol.htm)
- [21] K. Chandy and J. Misra, *Parallel Program Design: A Foundation*. Addison-Wesley, 1988.
- [22] S. Waydo and E. Klavins, "Verification of an Autonomous Reliable Wingman Using CCL," California Institute of Technology, Tech. Rep. CaltechCDSTR:2006.001, 2006.
- [23] "The Computation and Control Language (CCL)." [Online]. Available: <http://sveiks.ee.washington.edu/ccl/>




Fig. 9. Flight path of the two vehicles for the Lost Wingman portion of Flight Test 6 on June 25, 2004. Points of interest: (1) Lost maneuver start (Lost 1 mode), (2) Lost heading established (Lost 2 mode), (3) Rejoin maneuver start (Found mode), (4) Rejoin maneuver complete (Normal mode).

[24] "Verification of an Autonomous Reliable Wingman Using CCL." [Online]. Available: [http://www.cds.caltech.edu/~waydo/lost\\_wingman/](http://www.cds.caltech.edu/~waydo/lost_wingman/)