

Formal Verification of Curved Flight Collision Avoidance Maneuvers: A Case Study^{*}

André Platzer and Edmund M. Clarke

Computer Science Department, Carnegie Mellon University, Pittsburgh, PA

Abstract. Aircraft collision avoidance maneuvers are important and complex applications. Curved flight exhibits nontrivial continuous behavior. In combination with the control choices during air traffic maneuvers, this yields hybrid systems with challenging interactions of discrete and continuous dynamics. As a case study illustrating the use of a new proof assistant for a logic for nonlinear hybrid systems, we analyze collision freedom of roundabout maneuvers in air traffic control, where appropriate curved flight, good timing, and compatible maneuvering are crucial for guaranteeing safe spatial separation of aircraft throughout their flight. We show that formal verification of hybrid systems can scale to curved flight maneuvers required in aircraft control applications. We introduce a fully flyable variant of the roundabout collision avoidance maneuver and verify safety properties by compositional verification.

1 Introduction

In air traffic control, collision avoidance maneuvers [1,2,3,4] are used to resolve conflicting flight paths that arise during free flight. See Fig. 1 for a series of increasingly more realistic—yet also more complicated—aircraft collision avoidance maneuvers. Fig. 1c shows a malfunctioning collision avoidance attempt. Collision avoidance maneuvers are a “last resort” for resolving air traffic conflicts that could lead to collisions. They are important whenever conflicts have not been detected by the pilots during free flight or by the flight directors of the Air Route Traffic Control Centers. Consequently, complicated online trajectory prediction or maneuver planning may no longer be feasible in the short time that remains for resolving the conflict. In the tragic 2002 mid-flight collision in Überlingen, the aircraft collided tens of seconds after the on-board traffic alert and collision avoidance system TCAS signalled a traffic alert. Thus, for safe aircraft control we need particularly reliable reactions with maneuvers whose correctness has been established previously by a thorough offline analysis. To ensure correct functioning of aircraft collision avoidance maneuvers under all circumstances, the temporal evolution of the aircraft in space must be analyzed carefully together with the effects that maneuvering control decisions have on

^{*} This research was supported by DFG SFB/TR 14 AVACS, NASA NNG05GF84H, Berkman Faculty Award, CMU-GM CRL GM9100096UMA, NSF CCR-0411152, CCF-0429120, CCF-0541245, SRC 2008TJ1860, and AFRO 18727S3.




Fig. 1. Evolution of collision avoidance maneuvers in air traffic control

their dynamics. This results in complicated superpositions of physical system dynamics with control, which is an example of what is called hybrid system [5].

Several numerical [1,6,7,8,4] or optimization-based [6,7,9,4] approaches have been proposed for air traffic control. It is difficult to give sound formal verification results for these approaches due to errors in numerical computations or implicit definition of maneuvers in terms of complicated optimization processes. Formal verification is important to avoid collisions, see Fig. 1c. Formal results have been given by geometrical reasoning [2,3,10,11] in PVS. Yet, one still has to prove by other techniques that the hybrid dynamics of a flight controller actually follows the geometrical shapes. In contrast, we verify the hybrid system dynamics directly using a formally sound approach (assuming sound elementary decision procedures), consider curved flight, and achieve better automation.

Control Challenges. Because of the complicated spatio-temporal movement of aircraft, their maneuvers are challenging for verification. Unlike in ground transportation, braking and waiting is not an option to resolve conflicts. Consequently, aircraft maneuvers have to be coordinated such that the aircraft always respect minimal and maximal lateral and angular speed constraints yet always remain safely separated. Further, angular velocity for curving is the primary means of control, because changes in thrust and linear speed are less efficient for aircraft.

Technical Challenges. Complexities in analysis of aircraft maneuvers manifest most prominently in difficulties with analysing hybrid systems for flight equations. General solutions of flight equations involve trigonometric functions that depend on the angular velocity ω and the orientation of the aircraft in space. For straight line flight ($\omega = 0$), the movement in space is just linear so that classical analysis techniques can be used [5]. These include pure straight line maneuvers [1,12,2,3,4]; see, e.g., Fig. 1a. They have to assume instant turns for heading changes of the aircraft between multiple straight line segments. Instant turns, however, are impossible in midflight, because they are *not flyable*: Aircraft cannot suddenly change their flight direction from 0 to 45 degrees discontinuously. They need to follow a smooth curve instead, in which they slowly steer towards the desired direction by adjusting the angular velocity ω appropriately. Moreover, the area required by maneuvers for which instant turns could possibly

be understood as adequately close approximations of properly curved flight is huge. Curved flight is thus an inherent part of real aircraft control.

During curved flight, the angular velocity ω is non-zero. For $\omega \neq 0$, flight equations have transcendental solutions, which generally fall into undecidable classes of arithmetics; see [13]. Consequently, maneuvers with curves, like in Fig. 1b–1d, are more realistic but also substantially more complicated for verification than straight line maneuvers like that in Fig. 1a. We have recently developed a *sound* verification algorithm that works with differential invariants [14] instead of solutions of differential equations to address this arithmetic. Now we show how a fully curved maneuver can be verified by extending our work [14].

In this paper, we introduce and verify the *fully flyable tangential roundabout maneuver (FTRM)*. It refines the non-flyable tangential roundabout maneuver (NTRM) from Fig. 1d, which has discontinuities at the entry and exit points of roundabouts, to a fully flyable curved maneuver. Unlike most previously proposed maneuvers [1,7,12,2,15,3,4], FTRM does not have non-flyable instant turns. It is flyable and smoothly curved. Unlike other approaches emphasizing the importance of flyability [6], we give formal verification results.

Contribution. Our main contribution is to show that reality in model design and coverage in formal verification are no longer incompatible desires even for applications as complex as aircraft maneuvers. As a case study illustrating the use of differential dynamic logic for hybrid systems [16], we demonstrate how tricky and nonlinear dynamics can be verified with our verification algorithm [14] in our verification tool KeYmaera. We introduce a fully curved flight maneuver and verify its hybrid dynamics formally. In contrast to previous approaches, we handle curved flight, hybrid dynamics, and produce formal proofs with almost complete automation. Manual effort is still needed to simplify arithmetical complexity and modularize the proof appropriately. We further illustrate the resulting verification conditions for the respective parts of the maneuver. Finally, we identify the most difficult steps during the verification and present new transformations to handle the enormous computational complexity. To reduce complexity, we still use some of the simplifications assumed in related work, e.g., synchronous maneuvering (i.e. aircraft make simultaneous maneuver choices).

Related Work. Lafferriere et al. [17] gave important decidability results for hybrid systems with some classes of linear continuous dynamics but only random discrete resets. These results do not apply to air traffic maneuvers, because they have non-trivial resets: the aircraft’s position does not just jump randomly when switching modes but, rather, systematically according to the maneuver.

Tomlin et al. [1] analyze competitive aircraft maneuvers game-theoretically using numerical approximations of partial differential equations. As a solution, they propose roundabout maneuvers and give bounded-time verification results for straight-line approximations (Fig. 1a). We verify curved roundabouts with a sound symbolic approach that avoids approximation errors.

Flyability has been identified as one of the major challenges in Košecká et al. [6], where planning based on superposition of potential fields has been used to re-

solve air traffic conflicts. This planning does not guarantee flyability but, rather, defaults to classical vertical altitude changes whenever a nonflyable path is detected. The resulting maneuver has not yet been verified. The planning approach has been pursued by Bicchi and Pallottino [7] with numerical simulations.

Numerical simulation algorithms approximating discrete-time Markov Chain approximations of aircraft behavior have been proposed by Hu et al. [8]. They approximate bounded-time probabilistic reachable sets for one initial state. We consider hybrid systems combining discrete control choices and continuous dynamics instead of uncontrolled, probabilistic continuous dynamics.

Hwang et al. [4] have presented a straight-line aircraft conflict avoidance maneuver that involves optimization over complicated trigonometric computations, and validate it using random numerical simulation and informal arguments.

The work of Doweck et al. [2] and Galdino et al. [3] is probably closest to ours. They consider straight-line maneuvers and formalize geometrical proofs in PVS.

Attempts to Model Check discretizations of roundabout maneuvers [12,15] indicated avoidance of orthogonal collisions (Fig. 1b). Counterexamples found by our Model Checker in previous work show that collision avoidance does not extend to other initial flight paths of the classical roundabout (Fig. 1c).

Pallottino et al. [18] have presented a spatially distributed pattern for multiple roundabout circles at different positions. They reason manually about desirable properties of the system and estimate probabilistic results as in [8]. Pallottino et al. thus take a view that is complementary to ours: they determine the global compatibility of multiple roundabouts while assuming correct functioning within each local roundabout. We verify that the actual hybrid dynamics of each local roundabout is collision free. Generalizing our approach to a spatial pattern of verified local roundabouts could be interesting future work.

Similarly, the work by Umeno and Lynch [11,10] is complementary to ours. They consider real-time properties of airport protocols using Timed I/O Automata. We are interested in proving local properties of the actual hybrid system.

Our approach has a very different focus than other complementary work:

- Our maneuver directly involves curved flight unlike [1,8,2,3,4,11,10]. This makes our maneuver more realistic but much more difficult to analyze.
- Unlike [6,8,4], we do not give results for a finite (sometimes small) number of initial flight positions (simulation). Instead, we verify uncountably many initial states and give unbounded-time horizon verification results.
- Unlike [1,6,7,8,9,4], we use symbolic instead of numerical computation so that numerical and floating point errors cannot cause soundness problems.
- Unlike [7,12,8,2,3,4,11,10], we analyze hybrid system dynamics directly.
- Unlike [6,1,7,8,4,12,18] we produce formal, deductive proofs. Further unlike the formal proofs in [2,3,11,10], our verification is much more automatic.
- In [2,3,4,11,10], it remains to be proven that the hybrid dynamics and flight equations follow the geometrical thoughts. In contrast, our approach directly works for the hybrid flight dynamics. We illustrate verification results graphically to help understand them, but the figures do not prove anything.

- Unlike [19], we consider collision avoidance maneuvers, not just detection.
- Unlike [7,9], we do not guarantee optimality of the resulting maneuver.

2 Background: Differential Dynamic Logic

Hybrid Programs. We use a *hybrid program* (HP) notation [16] for hybrid systems that include hybrid automata (HA) [5]. Each discrete and continuous transition corresponds to a sequence of statements, with a nondeterministic choice (\cup) between these transitions. Line 2 in Fig. 2 represents a continuous transition in a simplistic altitude controller. It tests (denoted by $?q = up$) if the current location q is *up*, and then follows a differential equation $z' = 1$ restricted to invariant region $z \leq 9$ (conjunction $z' = 1 \wedge z \leq 9$). Line 3 tests guard $z \geq 5$ when in state *up*, resets z by a discrete assignment, and then changes location q to *down*. The $*$ at the end indicates that the transitions of a HA repeat indefinitely. We will build HP directly, which gives more natural programs than HA-translation.




Fig. 2. Hybrid automaton vs. hybrid program (simplistic altitude control)

As *terms* we allow polynomials over \mathbb{Q} with variables in a set V . *Hybrid programs* (HP) are built with the statements in Table 1. The effect of $x := \theta$ is an instantaneous discrete jump assigning θ to x . Instead, $x := *$ randomly assigns *any* real value to x by a nondeterministic choice. During a continuous evolution $x'_1 = \theta_1 \wedge \dots \wedge x'_n = \theta_n \wedge \chi$ with terms θ_i , all conjuncts need to hold. Its effect is a continuous transition controlled by the differential equation

Table 1. Statements and (informal) effects of hybrid programs (HP)

notation	statement	effect
$x := \theta$	discrete assignment	assigns term θ to variable $x \in V$
$x := *$	nondet. assignment	assigns any real value to $x \in V$
$x'_1 = \theta_1 \wedge \dots \wedge x'_n = \theta_n \wedge \chi$	continuous evolution	diff. equations for $x_i \in V$ and terms θ_i , with formula χ as evolution domain
$? \chi$	state check	test formula χ at current state
$\alpha; \beta$	seq. composition	HP β starts after HP α finishes
$\alpha \cup \beta$	nondet. choice	choice between alternatives HP α or β
α^*	nondet. repetition	repeats HP α n -times for any $n \in \mathbb{N}$
$do \alpha \text{ until } \chi$	evolve until	evolve HP α until χ holds

$x'_1 = \theta_1, \dots, x'_n = \theta_n$ that always satisfies the arithmetic constraint χ (thus remains in the region described by χ). This directly corresponds to a continuous evolution mode of a HA. The effect of state check $?\chi$ is a *skip* (i.e., no change) if χ is true in the current state and that of *abort*, otherwise. Non-deterministic choice $\alpha \cup \beta$ expresses alternatives in the behavior of the hybrid system. Sequential composition $\alpha; \beta$ expresses a behavior in which β starts after α finishes (β never starts if α continues indefinitely). Non-deterministic repetition α^* , repeats α an arbitrary number of times (≥ 0). The operation *do* α *until* χ expresses that the system follows α exactly until condition χ is true.

Formulas of dL. To express and combine correctness properties of HP, we use a verification logic for HP: The *differential dynamic logic* **dL** [16] is an extension of first-order logic over the reals with modal formulas like $[\alpha]\phi$, which is true iff all states reachable by following the transitions of HP α satisfy property ϕ (*safety*). Reachability properties are expressible using the dual modality $\langle \alpha \rangle \phi$, which is true iff there is a state satisfying ϕ that α can reach from its initial state. *Formulas of dL* are defined by the following grammar, where θ_1, θ_2 are terms, $\sim \in \{=, \leq, <, \geq, >\}$, ϕ, ψ are formulas, $x \in V$, and α is an HP (Table 1):

Formula ::= $\theta_1 \sim \theta_2 \mid \neg \phi \mid \phi \wedge \psi \mid \phi \vee \psi \mid \phi \rightarrow \psi \mid \forall x \phi \mid \exists x \phi \mid [\alpha]\phi \mid \langle \alpha \rangle \phi$.

A Hoare-triple $\{\psi\}\alpha\{\phi\}$ can be expressed as $\psi \rightarrow [\alpha]\phi$, which is true iff all states reachable by HP α satisfy ϕ when starting from an initial state that satisfies ψ .

3 Curved Flight in Roundabout Maneuvers

3.1 Flight Dynamics

The parameters of two aircraft at (planar) position $x = (x_1, x_2)$ and $y = (y_1, y_2)$ in \mathbb{R}^2 flying in directions $d = (d_1, d_2) \in \mathbb{R}^2$ and $e = (e_1, e_2)$ are illustrated in Fig. 3. Their dynamics is determined by their angular speeds $\omega, \rho \in \mathbb{R}$ and linear velocity vectors d and e , which describe both the linear velocity $\|d\| := \sqrt{d_1^2 + d_2^2}$ and orientation of the aircraft in space. Roundabout maneuvers are horizontal collision avoidance maneuvers so that, like [1,12,9,15,18,3,4], we simplify to planar positions.




Fig. 3. Aircraft flight

We denote the flight equations for the aircraft at x and y with angular velocities ω, ρ by $\mathcal{F}(\omega)$ and $\mathcal{G}(\rho)$ respectively, see [1,13]:

$$[x' = d \quad d' = \omega d^\perp] \quad (\mathcal{F}(\omega)) \quad [y' = e \quad e' = \rho e^\perp] \quad (\mathcal{G}(\rho))$$

There $d^\perp := (-d_2, d_1)$ is the *orthogonal complement* of vector d . Differential equations $\mathcal{F}(\omega)$ express that x is moving in direction d , which is rotating with angular velocity ω , i.e., evolves orthogonal to d . Equations $\mathcal{G}(\rho)$ are similar for y, e and ρ . In safe flight configurations, aircraft respect protected zone p . That is, they are separated by at least distance p , i.e., the state satisfies formula $\mathcal{S}(p)$:

$$\mathcal{S}(p) \equiv \|x - y\|^2 \geq p^2 \equiv (x_1 - y_1)^2 + (x_2 - y_2)^2 \geq p^2 \quad \text{for } p \in \mathbb{R} \quad (1)$$




Fig. 4. Protocol cycle and construction of flyable roundabout maneuver

Like all other parameters, we treat p purely symbolically without a specific value. In practice, horizontal separation should be $\geq 5\text{mi}$, vertical separation $\geq 1000\text{ft}$.

3.2 Roundabout Maneuver Overview

FTRM consists of the phases in the protocol cycle in Fig. 4a which correspond to the marked flight phases in Fig. 4b. During free flight, the aircraft move without restriction by repeatedly choosing arbitrary new angular velocities ω and ϱ respectively (as indicated by the self loop at *free* in Fig. 4a). When the aircraft come too close to one another, they agree on a roundabout maneuver by negotiating a compatible roundabout center $c = (c_1, c_2)$ in coordination phase *agree* by communication. Next, the aircraft approach the roundabout circle in a right curve with $\omega < 0$ (*entry* mode) according to Fig. 4b, and reach a tangential position around center c . During the *circ* mode, the aircraft follow the circular roundabout maneuver around the agreed center c with a left curve of common angular velocity $\omega > 0$. Finally, the aircraft leave the roundabout in cruise mode ($\omega = 0$) in their original direction (*exit*) and enter free flight again when they have reached sufficient distance (the protocol cycle repeats as necessary).

3.3 Compositional Verification Plan

For verifying safety properties and collision avoidance of FTRM, we decompose the verification problem and pursue the following overall verification plan:

- AC1 *Tangential roundabout maneuver cycle*: We prove that the protected zones of aircraft are safely separated at all times during the whole maneuver (including repetitive collision avoidance maneuver initiation and including multiple aircraft) with a simplified but not yet flyable entry operation $entry_n$. Subsequently, we refine this verification result to a flyable maneuver by verifying that we can replace $entry_n$ with its flyable variant *entry*.
- AC2 *Bounded control choices for aircraft velocities*: We show that linear speeds remain unchanged during the whole maneuver (the aircraft do not stall).
- AC3 *Flyable entry*: We prove that the simplified $entry_n$ procedure can be replaced by a flyable curve *entry* reaching the same position as $entry_n$.

- AC4** *Bounded entry duration:* Flyable entry procedure succeeds in bounded time, i.e., aircraft reach the roundabout circle in some bounded time $\leq T$.
- AC5** *Safe entry separation:* Most importantly, we prove that the protected zones of aircraft are still respected during the flyable entry procedure.
- AC6** *Successful negotiation:* We prove that the negotiation phase (*agree*) satisfies the respective requirements of multiple aircraft simultaneously.
- AC7** *Safe exit separation:* We show that, for its bounded duration, the exit procedure cannot produce collisions and that the initial *far separation* for free flight is reached again so that the FTRM cycle repeats safely.

This plan modularizes the proof and allows us to identify the respective safety constraints imposed by the various maneuver phases successively. We present details of these verification tasks in the sequel and summarize the respective verification results into a joint safety property of FTRM in Section 5. The proof and formulation for **AC2** is a simple variation of **AC1** and will not be discussed.

3.4 Tangential Roundabout Maneuver Cycles (AC1)

First, we analyze roundabouts with a simplified instant entry procedure and without an exit procedure (**AC1**), i.e., the non-flyable NTRM depicted in Fig. 1d. We refine this maneuver and its verification to the flyable FTRM afterwards.

Modular Correctness of Tangential Roundabout Cycles. We verify that NTRM safely avoids collisions, i.e., the aircraft always maintain a safe distance $\geq p$ during the curved flight in roundabout. In addition, these results show that arbitrary repetitions of the protocol cycle are always safe when, as a first step, we simplify the entry maneuver. The NTRM model and property are summarized in Fig. 5. The simplified flight controller in Fig. 5 performs collision avoidance maneuvers by tangential roundabouts and repeats these maneuvers any number of times as needed. During each cycle of the loop of *NTRM*, the aircraft first perform arbitrary free flight (*free*) by choosing arbitrary new angular velocities ω and ϱ (repeatedly as indicated by the loop in *free*).

$$\begin{aligned}
 \psi &\equiv \mathcal{S}(p) \rightarrow [NTRM] \mathcal{S}(p) \\
 NTRM &\equiv (free; agree; entry_n; circ)^* \\
 free &\equiv (\omega := *; \varrho := *; \mathcal{F}(\omega) \wedge \mathcal{G}(\varrho) \wedge \mathcal{S}(p))^* \\
 agree &\equiv \omega := *; c := * \\
 entry_n &\equiv d := \omega(x - c)^\perp; e := \omega(y - c)^\perp \\
 circ &\equiv \mathcal{F}(\omega) \wedge \mathcal{G}(\omega)
 \end{aligned}$$

Fig. 5. Nonflyable tangential roundabout collision avoidance maneuver NTRM

Aircraft only fly freely while they are safely separated, which is expressed by constraint $\mathcal{S}(p)$ in the differential equation for *free*. Then the aircraft agree on an arbitrary roundabout center c and angular velocity ω (*agree*). We model this communication by nondeterministic assignments to the shared variables ω, c . Refinements include all negotiation processes that reach an agreement on common ω, c in bounded time. Next, they perform the simplified non-flyable entry

procedure ($entry_n$) with instant turns (Fig. 1d). This operation identifies the goal state that $entry$ needs to reach:

$$\mathcal{R} \equiv d = \omega(x - c)^\perp \wedge e = \omega(y - c)^\perp \tag{2}$$

It expresses that, at the positions x and y , respectively, the directions d and e are tangential to the roundabout circle at center c and angular velocity ω ; see Fig. 6. Finally, the roundabout maneuver itself is carried out in $circ$. The collision avoidance roundabouts can be left again by repeating the loop and entering arbitrary free flight at any time. When further conflicts occur during free flight, the controller in Fig. 5 again enters roundabout conflict resolution maneuvers.




Fig. 6. \mathcal{R}

Multiple Aircraft. We prove separation for up to 5 aircraft participating in the roundabout at the same time. There, the safety property is mutual collision avoidance, i.e., each aircraft has a safe distance $\geq p$ to every other aircraft, which yields a quadratic number of separation properties that have to be verified. This quadratic increase in the size of the property that actually needs to be proven for a safe roundabout of n aircraft and the increased dimension of the underlying continuous state space increase verification times. Also see [13].

3.5 Flyable Entry Procedures (AC3)

For property **AC3** in Section 3.3, we generalize the verification results about NTRM with simplified entry procedures (Fig. 1d) to FTRM (Fig. 4b) by replacing the non-flyable $entry_n$ procedure with flyable curves (called $entry$). This turns the non-flyable NTRM into the flyable FTRM maneuver.

Flyable Entry Properties. A flyable entry maneuver that follows the smooth entry curve from Fig. 4b is constructed according to Fig. 7a and specified formally as:

$$\begin{aligned} (r\omega)^2 = \|d\|^2 \wedge \|x - c\| = \sqrt{3}r \wedge \exists \lambda \geq 0 (x + \lambda d = c) \wedge \|h - c\| = 2r \wedge d = -\omega(x - h)^\perp \\ \rightarrow [\mathcal{F}(-\omega) \wedge \|x - c\| \geq r] (\|x - c\| \leq r \rightarrow d = \omega(x - c)^\perp) \end{aligned} \tag{3}$$

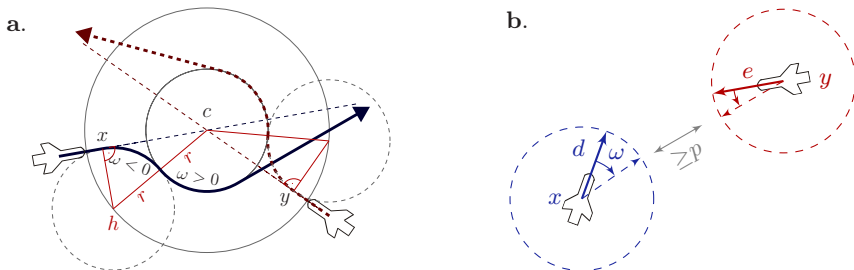


Fig. 7. Flyable entry maneuver: characteristics and separation

The assumptions in (3) express that r is the radius corresponding to speed $\|d\|$ and angular velocity ω ($(r\omega)^2 = \|d\|^2$) and that *entry* starts with distance $\sqrt{3}r$ to c heading towards c ($\exists \lambda \geq 0 (x + \lambda d = c)$). For the construction of the maneuver and positioning in space, we use the auxiliary anchor point $h \in \mathbb{R}^2$ identified in Fig. 7a and line 1 of (3). It is positioned relative to the roundabout center c and the x position at the start of the entry curve (i.e., with x at the right angle indicated in Fig. 7a). The entry curve around h is similar to the roundabout curve around c . Formally, h is characterized by distance r to x , distance $2r$ to c ($\|h - c\| = 2r$) and, further, vector $x - h$ is orthogonal to d and obeys the relative orientation of the curve belonging to $-\omega$ (hence $d = -\omega(x - h)^\perp$). The property in (3) specifies that the tangential goal configuration (2) around c is reached by a flyable curve when waiting until aircraft x and center c have distance r , because the domain restriction of the dynamics is $\|x - c\| \geq r$ (line 2) and the postcondition assumes $\|x - c\| \leq r$, which imply $\|x - c\| = r$. The feasibility of choosing anchor point h can be shown by proving an existence property; see [13].

Spatial Symmetry Reduction. The property in (3) can be verified in a simplified version. We use a new *spatial symmetry reduction* to simplify property (3) computationally. We exploit symmetries to reduce the spatial dimension by fixing variables. Without loss of generality, we recenter the coordinate system with c at position 0. Further, we can assume aircraft x comes from the left by changing the orientation of the coordinate system. Finally, we assume, without loss of generality, linear speed 1 (by rescaling units appropriately). Observe that we *cannot* fix a value for both the linear speed and the angular velocity, because the units are interdependent. In other words, if we fix the linear speed, we need to consider all angular velocities in order to verify the maneuver for each possible radius r of the roundabout maneuver (and corresponding ω). The x position resulting from these symmetry reductions can be determined easily by Pythagoras theorem (i.e., $(2r)^2 = r^2 + x_1^2$ for the triangle enclosed by h, x, c in Fig. 7a):

$$x = (\sqrt{(2r)^2 - r^2}, 0) = (\sqrt{3}r, 0) . \tag{4}$$

3.6 Bounded Entry Duration (AC4)

As the first step for showing that the entry procedure finally succeeds at goal (2) and maintains a safe distance all the time, we show that *entry* succeeds in bounded time and cannot take arbitrarily long to succeed (**AC4** in Section 3.3).

By a simple consequence of (3), the entry procedure follows a circular motion around anchor point h , see Fig. 7a. That is, when r is the radius belonging to angular velocity ω and linear speed $\|d\|$, the property $\|x - h\| = r$ is an invariant of *entry*; see [13]. By **AC2**, which can be proven easily, the speed $\|d\|$ is constant during the *entry* procedure. Thus, the aircraft proceeds with nonzero minimum progress rate $\|d\|$ around the circle. The flight duration for a full circle of radius r around h at constant linear speed $\|d\|$ is $\frac{2\pi r}{\|d\|}$, because its arc length is $2\pi r$. From the trigonometric identities underlying equation (4), we can read off that the aircraft completes a $\frac{\pi}{3} = 60^\circ$ arc, see Fig. 7a. Hence, the maximum duration T

of the *entry* procedure is: $T := \frac{1}{6} \cdot \frac{2\pi r}{\|d\|} = \frac{\pi r}{3\|d\|}$ Instead of π , which is not definable in first-order real arithmetic, we can use any overapproximation, e.g., 3.15.

3.7 Safe Entry Separation (AC5)

In Section 3.5, we have shown that the simplified *entry_n* procedure from NTRM can be replaced by a flyable *entry* maneuver that meets the requirements of approaching tangentially for each aircraft. Unlike in instant turns (*entry_n*), we have to show that the flyable entry maneuvers of multiple aircraft do not produce mutually conflicting flight paths, i.e., spatial separation of all aircraft is maintained during the entry of multiple aircraft (AC5). See Fig. 8 for multiple aircraft FTRM where separation is important.




Fig. 8. Multiple aircraft

Bounded Overapproximation. We show that entry separation is a consequence of the bounded speed (AC2) and bounded duration (AC4) of the flyable entry procedure when initiating the negotiation phase *agree* with sufficient distance. We prove that, when following bounded speed for a bounded duration, aircraft only come closer by a bounded distance. Let b denote the overall speed bound during FTRM according to AC2 and let T be the time bound for the duration of the entry procedure due to AC4. We overapproximate the actual behavior during the *entry* phase by arbitrary curved flight (see Fig. 7b). When the *entry* procedure is initiated with sufficient distance $\sqrt{2}(p + 2bT)$, the protected zone $p \geq 0$ will still be respected after the 2 aircraft follow *any* curved flight (including the actual choices during the *entry* phase and subsequent *circ* phase) with speed $\|d\| \leq b$ and $\|e\| \leq b$ up to $T \geq 0$ time units (see Fig. 7b):

$$\|x - y\| \geq \sqrt{2}(p + 2bT) \wedge p \geq 0 \wedge \|d\|^2 \leq \|e\|^2 \leq b^2 \wedge b \geq 0 \wedge T \geq 0 \rightarrow [entry] (\|x - y\| \geq p) \quad (5)$$

In [13], we show that this property follows from the more general fact that aircraft only make limited progress in bounded time from some initial point z when starting with bounded speeds (even when changing ω arbitrarily):

$$x = z \wedge \|d\|^2 \leq b^2 \wedge b \geq 0 \rightarrow [\tau := 0; \mathcal{F}(\omega) \wedge \tau' = 1] (\|x - z\|_\infty \leq \tau b) \quad (6)$$

The maximum distance $\|x - z\|_\infty$ from z depends on clock τ and bound b . To reduce the polynomial degree and the verification complexity, we overapproximate distances from quadratic Euclidean norm $\|\cdot\|$ in terms of linearly definable supremum norm $\|\cdot\|_\infty$, instead, which is $\|x\|_\infty \leq c \equiv -c \leq x_1 \leq c \wedge -c \leq x_2 \leq c$.

Far Separation. By combining the estimation of the entry duration (3.6) at speed $\|d\| = b$ with the entry separation property (5), we determine the following magnitude as the *far separation* f , i.e., the initial distance guaranteeing that the FTRM protocol can be repeated safely in case new collision avoidance is needed:

$$f := \sqrt{2}(p + 2bT) = \sqrt{2} \left(p + \frac{2}{3}\pi r \right) \quad (7)$$

4 Synchronization of Roundabout Maneuvers

Following our verification plan in Section 3.3, we show that the various actions of multiple aircraft can be synchronized appropriately to ensure safety of the maneuver. We analyze the negotiation phase and compatible exit procedures.

4.1 Successful Negotiation (AC6)

For negotiation to succeed (AC6), we have to show that there is a common choice of the roundabout center c and angular velocity ω (or radius r) so that multiple participating aircraft can satisfy the local requirements of their respective entry procedures simultaneously, i.e., of the property (3) for AC3.

We prove that all corresponding choices of *agree* satisfy the mutual requirements of multiple aircraft simultaneously. As one possible option among others: when choosing roundabout center c as the simultaneous intersection (intersection $x + \lambda d = y + \lambda e$ after time λ) of the flight paths of the aircraft at x and y , the choices for c, r, ω are compatible for multiple aircraft; see Fig. 9a:

$$\begin{aligned} &\lambda > 0 \wedge x + \lambda d = y + \lambda e \wedge \|d\| = \|e\| \rightarrow \\ &[c := x + \lambda d; r := *; ?\|x - c\| = \sqrt{3}r; ?\|y - c\| = \sqrt{3}r; \omega := *; ?(r\omega)^2 = \|d\|^2] \\ &(\|x - c\| = \sqrt{3}r \wedge \lambda \geq 0 \wedge x + \lambda d = c \wedge \|y - c\| = \sqrt{3}r \wedge y + \lambda e = c) \quad (8) \end{aligned}$$

The tests in the dynamics ensure that the *entry* curve starts when x, y and c have appropriate distance $\sqrt{3}r$ identified in Section 3 and that r is the radius belonging to angular velocity ω and linear speed $\|d\|$. This property expresses that, for aircraft heading towards the simultaneous intersection of their flight paths with speed $\|d\| = \|e\|$ (line 1), the intersection of the linear flight paths (line 2) is a safe choice for c satisfying the joint requirements (line 3) identified in Section 3. For an analysis of far separation during negotiation and of the feasibility of these choices, see[13]. Other choices of c, ω than Fig. 9a are possible for asymmetric initial positions of aircraft, but computationally more involved.

4.2 Safe Exit Separation (AC7)

NTRM (Fig. 1d) does not need an exit procedure for safety, because the maneuver repeats when further air traffic conflicts arise. For FTRM, instead, we




Fig. 9. Separation of negotiation and good and bad exit procedure separation

need to show that the exit procedure produces safe flight paths until the aircraft are sufficiently separated: When repeating the FTRM maneuver, the *entry* procedure needs far separation (7) not just distance p for safety, see Fig. 4b.

Safe Separation. If the aircraft enter simultaneously, they can exit simultaneously. For **AC7**, we first show that aircraft that exit simultaneously (from tangential positions of the roundabout circle) always respect their protected zones:

$$\mathcal{R} \wedge \|x - y\|^2 \geq p^2 \rightarrow [x' = d \wedge y' = e] (\|x - y\|^2 \geq p^2) \quad (9)$$

Thus, safely separated aircraft exiting simultaneously along straight lines from tangential positions (\mathcal{R} by eqn. 2) of a roundabout always remain safely separated. We prove an overapproximation: exit rays (Fig. 9b–9c) are separated [13].

Far Separation. Aircraft reach arbitrary separation when following the exit procedure long enough. Using overapproximation Fig. 9b, we prove that—due to different exit directions $d \neq e$ —the exit procedure will finally separate the aircraft arbitrarily far (starting from tangential configuration (2) of the roundabout):

$$\mathcal{R} \wedge d \neq e \rightarrow \forall a \langle x' = d \wedge y' = e \rangle (\|x - y\|^2 > a^2) \quad (10)$$

5 Flyable Tangential Roundabout Maneuver

We combine the results about the individual phases of flyable roundabouts into a full model of FTRM that inherits safety modularly. We collect the maneuver phases according to the protocol cycle of Fig. 4 and take care to ensure that the safety prerequisites are met, as identified for the respective phases in Section 3-4.

One possible instance of FTRM is the HP in Fig. 10, which is composed of previously illustrated parts of the maneuver. The technical construction and protocol cycle of the entry procedure have already been illustrated in Fig. 4. In FTRM, Π denotes the synchronous parallel product. By communication, FTRM operates synchronously, i.e., all aircraft make simultaneous mode changes [4].

$$\psi \equiv \|d\| = \|e\| \wedge r > 0 \wedge \mathcal{S}(f) \rightarrow [FTRM^*] \mathcal{S}(p)$$

$$\mathcal{C} \equiv \|x - c\| = \sqrt{3}r \wedge \exists \lambda \geq 0 (x + \lambda d = c) \wedge \|y - c\| = \sqrt{3}r \wedge \exists \lambda \geq 0 (y + \lambda e = c)$$

$$FTRM \equiv free^*; agree; \Pi(entry; circ; exit)$$

$$free \equiv \omega := *; \varrho := *; \mathcal{F}(\omega) \wedge \mathcal{G}(\varrho) \wedge \mathcal{S}(f)$$

$$agree \equiv c := *; r := *; ?(\mathcal{C} \wedge r > 0); ?\mathcal{S}(f);$$


$$\omega := *; ?(r\omega)^2 = \|d\|^2; x_0 := x; d_0 := d; y_0 := y; e_0 := e$$

$$entry \equiv do \mathcal{F}(-\omega) \text{ until } \|x - c\|^2 = r^2$$

$$circ \equiv do \mathcal{F}(\omega) \text{ until } \exists \lambda \geq 0 \exists \mu > 0 (x + \lambda d = x_0 + \mu d_0)$$

$$exit \equiv \mathcal{F}(0); ?\mathcal{S}(f)$$

Fig. 10. Flight control with flyable tangential roundabout collision avoidance



Decomposed property of system dynamics	See
$S(f) \rightarrow [free] S(f)$	Fig. 5
$S(f) \rightarrow [agree](S(f) \wedge C)$	(8), [13]
$C \wedge S(f) \rightarrow [entry] S(p)$	(5)
$C \wedge S(f) \rightarrow [entry] \mathcal{R}$	(3)
$\mathcal{R} \wedge S(p) \rightarrow [circ](S(p) \wedge \mathcal{R})$	Fig. 5
$\mathcal{R} \wedge S(p) \rightarrow [exit] S(p)$	(9)
$\mathcal{R} \wedge S(p) \rightarrow [exit] S(f)$	(9), (10)

Fig. 11. Composing verification for flyable tangential roundabout maneuvers

Consequently, the parallel product $\Pi(entry; circ; exit)$ of HP simplifies to the conjunction of the respective differential equations in the various modes and can be defined easily: $(entry_x \wedge entry_y)$; $(circ_x \wedge circ_y)$; $(exit_x \wedge exit_y)$ where $entry_x$ is the entry procedure of the aircraft at position x (likewise for more aircraft).

To verify this maneuver, we split the proof into the modular properties that we have already shown previously following the verification plan from Section 3.3. Formally, we split the system at its sequential compositions, giving the subproperties depicted in Fig. 11. Formula \mathcal{R} is due to equation (2) and $S(p)$ by (1).

By combining the results about the FTRM flight phases as summarized in Fig. 11, we conclude that FTRM avoids collisions safely. The modular proof structure in Fig. 11 still holds when replacing any part of the maneuver with a different choice that still satisfies the specification, e.g., for different entry procedures that still succeed in tangential configuration \mathcal{R} within bounded time. This includes roundabouts with *asymmetric positions*, i.e., where the initial distance to c can be different, and with *near conflicts*, where the flight paths do not intersect in one point but in a larger critical region [4]. Most notably, the separation proof in Section 3.7 tolerates asymmetric distances to c (Fig. 7b).

Theorem 1 (Safety property of flyable tangential roundabouts). *FTRM is collision free, i.e., the collision avoidance property ψ in Fig. 10 is valid. Furthermore any variation of FTRM with a modified entry procedure that safely reaches tangential configuration \mathcal{R} in some bounded time T is safe. That is if the following formula holds, saying that, until time T , the aircraft have safe distance p and will have reached configuration \mathcal{R} at time T , where τ is a clock:*

$$S(f) \rightarrow [\tau := 0; entry \wedge \tau' = 1]((\tau \leq T \rightarrow S(p)) \wedge (\tau = T \rightarrow \mathcal{R})) .$$

6 Experimental Results

Table 2 summarizes experimental results obtained using the tool KeYmaera on a 2.6GHz AMD Opteron with 4GB memory; we use different proof search settings than in [14]. Rows marked with * indicate a property where simplifications like symmetry reduction have been used to reduce the computational complexity. Table 2 shows that even aircraft maneuvers with challenging hybrid curve

Table 2. Experimental results for air traffic control (see [13] for details)

Case study	See	Time(s)	Memory(MB)	Steps	Dimension
tangential roundabout	2 aircraft	10.4	6.8	197	13
tangential roundabout	3 aircraft	253.6	7.2	342	18
tangential roundabout	4 aircraft	382.9	10.2	520	23
tangential roundabout	5 aircraft	1882.9	39.1	735	28
bounded maneuver speed	AC2	0.5	6.3	14	4
flyable roundabout entry*	(3)	10.1	9.6	132	8
flyable entry feasible*	[13]	104.5	87.9	16	10
flyable entry circular	[13]	3.2	7.6	81	5
limited entry progress	(6)	1.9	6.5	60	8
entry separation	[13]	140.1	20.1	512	16
mutual negotiation successful	(8)	0.8	6.4	60	12
mutual negotiation feasible*	[13]	7.5	23.8	21	11
mutual far negotiation	[13]	2.4	8.1	67	14
simultaneous exit separation*	[13]	4.3	12.9	44	9
different exit directions	[13]	3.1	11.1	42	11

dynamics can be verified formally. Memory consumption of quantifier elimination is shown in Table 2, excluding the front-end. The dimension of the continuous state space and number of automatic proof steps are indicated. Except for simple help in the proof of one property, the proofs for Table 2 are automatic.

7 Summary

We have analyzed complex air traffic control applications. Real aircraft can only follow sufficiently smooth flyable curves. Hence, mathematical maneuvers that require instant turns give physically impossible conflict resolution advice. We have developed a new collision avoidance maneuver with smooth, fully flyable curves. Despite its complicated dynamics and maneuvering, we have verified collision avoidance in this flyable tangential roundabout maneuver formally using our verification algorithm for a logic of hybrid systems. Because of the intricate spatio-temporal movement of aircraft in curved roundabouts, some of the properties require intricate arithmetic, which we handled by symmetry reduction and degree-based reductions. The proof is automatic except for modularization and arithmetical simplifications to overcome the computational complexity.

While the flyable roundabout maneuver is a highly nontrivial and challenging study, we still use modeling assumptions that should be relaxed in future work, e.g., synchronous, symmetric conflict resolution. Further generalizations include different varying cruise speeds, disturbances or new aircraft. The proof structure behind Theorem 1 is already sufficiently general, but the computational complexity high. It would be interesting future work to see if the informal robustness studies of Hwang et al. [4] can be carried over to a formal verification result.

Acknowledgements. We would like to thank the anonymous referees for their helpful comments and César Muñoz for his feedback.

References

1. Tomlin, C., Pappas, G.J., Sastry, S.: Conflict resolution for air traffic management. *IEEE T. Automat. Contr.* 43(4), 509–521 (1998)
2. Dowek, G., Muñoz, C., Carreño, V.A.: Provably safe coordinated strategy for distributed conflict resolution. In: *AIAA-2005-6047* (2005)
3. Galdino, A.L., Muñoz, C., Ayala-Rincón, M.: Formal verification of an optimal air traffic conflict resolution and recovery algorithm. In: Leivant, D., de Queiroz, R. (eds.) *WoLLIC 2007*. LNCS, vol. 4576, pp. 177–188. Springer, Heidelberg (2007)
4. Hwang, I., Kim, J., Tomlin, C.: Protocol-based conflict resolution for air traffic control. *Air Traffic Control Quarterly* 15(1), 1–34 (2007)
5. Henzinger, T.A.: The theory of hybrid automata. In: *LICS*, pp. 278–292. IEEE, Los Alamitos (1996)
6. Košecká, J., Tomlin, C., Pappas, G., Sastry, S.: 2-1/2D conflict resolution maneuvers for ATMS. In: *CDC*, Tampa, FL, USA, vol. 3, pp. 2650–2655 (1998)
7. Bicchi, A., Pallottino, L.: On optimal cooperative conflict resolution for air traffic management systems. *IEEE Trans. ITS* 1(4), 221–231 (2000)
8. Hu, J., Prandini, M., Sastry, S.: Probabilistic safety analysis in three-dimensional aircraft flight. In: *CDC*, vol. 5, pp. 5335–5340 (2003)
9. Hu, J., Prandini, M., Sastry, S.: Optimal coordinated motions of multiple agents moving on a plane. *SIAM Journal on Control and Optimization* 42, 637–668 (2003)
10. Umeno, S., Lynch, N.A.: Proving safety properties of an aircraft landing protocol using I/O automata and the PVS theorem prover. In: Misra, J., Nipkow, T., Sekerinski, E. (eds.) *FM 2006*. LNCS, vol. 4085, pp. 64–80. Springer, Heidelberg (2006)
11. Umeno, S., Lynch, N.A.: Safety verification of an aircraft landing protocol: A refinement approach. In: Bemporad, A., Bicchi, A., Buttazzo, G. (eds.) *HSCC 2007*. LNCS, vol. 4416, pp. 557–572. Springer, Heidelberg (2007)
12. Massink, M., Francesco, N.D.: Modelling free flight with collision avoidance. In: Andler, S.F., Offutt, J. (eds.) *ICECCS*, pp. 270–280. IEEE, Los Alamitos (2001)
13. Platzer, A., Clarke, E.M.: Formal verification of curved flight collision avoidance maneuvers: A case study. Technical Report CMU-CS-09-147, School of Computer Science, Carnegie Mellon University, Pittsburgh, PA (2009)
14. Platzer, A., Clarke, E.M.: Computing differential invariants of hybrid systems as fixedpoints. *Form. Methods Syst. Des.* (2009); In: Gupta, A., Malik, S. (eds.) *CAV 2008*. LNCS, vol. 5123, pp. 176–189. Springer, Heidelberg (2008)
15. Damm, W., Pinto, G., Ratschan, S.: Guaranteed termination in the verification of LTL properties of non-linear robust discrete time hybrid systems. In: Peled, D.A., Tsay, Y.-K. (eds.) *ATVA 2005*. LNCS, vol. 3707, pp. 99–113. Springer, Heidelberg (2005)
16. Platzer, A.: Differential dynamic logic for hybrid systems. *J. Autom. Reasoning* 41(2), 143–189 (2008)
17. Lafferriere, G., Pappas, G.J., Yovine, S.: A new class of decidable hybrid systems. In: Vaandrager, F.W., van Schuppen, J.H. (eds.) *HSCC 1999*. LNCS, vol. 1569, pp. 137–151. Springer, Heidelberg (1999)
18. Pallottino, L., Scordio, V.G., Frazzoli, E., Bicchi, A.: Decentralized cooperative policy for conflict resolution in multi-vehicle systems. *IEEE Trans. on Robotics* 23(6), 1170–1183 (2007)
19. Muñoz, C., Carreño, V., Dowek, G., Butler, R.W.: Formal verification of conflict detection algorithms. *STTT* 4(3), 371–380 (2003)