

Design of Guaranteed Safe Maneuvers Using Reachable Sets: Autonomous Quadrotor Aerobatics in Theory and Practice*

Jeremy H. Gillula[†], Haomiao Huang[†], Michael P. Vitus[†], and Claire J. Tomlin

Abstract—For many applications, the control of a complex nonlinear system can be made easier by modeling the system as a collection of simplified hybrid modes, each representing a particular operating regime. An example of this is the decomposition of complex aerobatic flights into sequences of discrete maneuvers, an approach that has proven very successful for both human piloted and autonomously controlled aircraft. However, a critical step when designing such control systems is to ensure the safety and feasibility of transitions between these maneuvers. This work presents a hybrid dynamics framework for the design of guaranteed safe switching regions and is applied to a quadrotor helicopter performing an autonomous backflip. The regions are constructed using reachable sets calculated via a Hamilton-Jacobi differential game formulation, and experimental results are presented from flight tests on the STARMAC quadrotor platform.

I. INTRODUCTION

Modern robotic systems are growing ever more capable and complex. In particular, as UAVs grow in power and maneuverability they require increasingly sophisticated control systems to take advantage of the full range of their capabilities. In the development of these control systems, it is often difficult if not impossible to consider a full nonlinear model of the system. Many approaches to the control of highly maneuverable aircraft have used statistical learning techniques, for example by copying an expert pilot’s example trajectory either through machine learning or via manual creation of approximate trajectories [2], [3]. These methods have been able to push the envelope of what is possible with autonomous control, but since they lack performance guarantees about their stability and robustness, their use may be limited in situations where safety is critical.

An alternative approach that allows more rigorous formal analysis is hybrid decomposition, where the behavior of the system of interest is approximated as a discrete set of simpler modes representing the dynamics in specific regimes or portions of the state space. The decomposed hybrid model (consisting of continuous states and discrete states representing the modes of the system) is then used for analysis and control. This approach has proven successful in a variety of



Fig. 1. STARMAC quadrotor performing an autonomous backflip maneuver.

applications, including manipulator motion planning [4], [5], specifications for mobile robot behaviors [6], and aircraft trajectory planning, where complex trajectories were designed by building up sequences of discrete maneuvers [7].

An important consideration in the design and control of systems with switched dynamics is the safety of transitions between modes. For example, in the case of aircraft maneuver sequences it is necessary to ensure that an aircraft completing one maneuver is able to begin the next maneuver without being in an unsafe or infeasible configuration. In the work described above, this has been accomplished in a variety of ways. The helicopter maneuvering work used “trim states” such as steady flight or hover that the vehicle had to return to after a maneuver before beginning another [7], while the manipulator work has used sequences of specially derived Lyapunov functions to guarantee that a defined sequence could be followed [5] or analytically calculated regions where a given motion was guaranteed to place a part in a desired configuration [4]. There has also been extensive work in the Hybrid Systems literature on construction of switching regions for mode switching [8]–[10], where partitions or manifolds in the state space were found that are regions of attraction for particular modes or controllers. Much of this work, though, has focused on switching under nominal conditions or sensing uncertainty and do not explicitly consider external disturbances.

This paper introduces a novel method for the design of provably safe aerobatic maneuvers using hybrid dynamics and reachability tools in the presence of external disturbances. This method is applied to a quadrotor helicopter performing a backflip, where the backflip maneuver is broken into three main stages: impulse, drift, and recovery. The impulse mode initializes the rotation of the vehicle. Upon reaching the appropriate switching condition the motors are

J. Gillula is a Ph.D. Candidate in Computer Science, Stanford University, Stanford, CA 94305, USA jgillula@cs.stanford.edu

H. Huang is a Ph.D. Candidate in Aeronautics and Astronautics, Stanford University, Stanford, CA 94305, USA haomiao@stanford.edu

M. Vitus is a Ph.D. Candidate in Aeronautics and Astronautics, Stanford University, Stanford, CA 94305, USA vitus@stanford.edu

C. Tomlin is a Professor of Electrical Engineering and Computer Sciences, UC Berkeley, Berkeley, CA 94720, USA tomlin@eecs.berkeley.edu

[†]These authors contributed equally to this work.

*An earlier version of this paper, with preliminary reachable set calculations and simulation results, was presented at ISRR 2009 [1].

turned off for the drift mode, where the vehicle freely rotates and falls under gravity. Finally, the recovery mode brings the vehicle to a controlled hover condition. Provably safe switching conditions on altitude, altitude rate, attitude, and attitude rate are generated using a Hamilton-Jacobi differential game formulation that guarantees the vehicle will successfully pass through all three modes to arrive at a desired final state. Controlled backflips were successfully performed using these switching rules on the Stanford Testbed of Autonomous Rotorcraft for Multi-Agent Control (STARMAC) [11] and results from these flight tests are presented.

The use of a Hamilton-Jacobi game formulation is a promising tool for the design and verification of safe maneuver sequences. By constructing the reachability problem as a game between a disturbance and control, backwards reachable sets can be calculated where a given controller is guaranteed to either keep out of or arrive into a defined region of the state space within some time horizon [12]. This formulation was first used to derive guaranteed safe switching regions for a collision avoidance controller for manned aircraft, where an evader aircraft could be guaranteed to keep some distance away from a pursuer, but has since been generalized to a wide range of systems [13], [14].

The organization of this paper is as follows. Section II discusses the background for the Hamilton-Jacobi differential game formulation for robust reachability, and the use of reachability for maneuver sequencing and design is presented in Section III. Flight test results are presented in Section IV, with conclusions and future work in Section V.

II. REACHABILITY FOR MODE SWITCHING

The concept of backwards reachability is used in this work as a means of generating guaranteed safe sets for mode transitions. For any mode i , the system evolves under the dynamics $\dot{x} = f_i(t, x, u, d)$, where x is the system state, u is the control input, and d is the disturbance input, where u and d are constrained in some sets U and D , respectively.

A. Hamilton-Jacobi Reachability

A backwards reachable set $Pre_\tau(K)$ is defined as the set of all states where the system can arrive in some set K within time τ , under some appropriate set of assumptions about the disturbance and control. Two conditions are considered here: safety, where the goal of the control is to stay out of an undesired set in the state space which the disturbance is trying to force the system into, and attainability, where the goal is to reach a desired set while the disturbance tries to keep the system out of this set. The reachable sets in this work are calculated using the technique described in Mitchell et al. with some modifications [12]. In this method, the reachability problem is posed as a differential game between the control input and disturbance under which the disturbance chooses the worst-case inputs to either drive the system into the undesired set or away from the desired, set and the control does the opposite.

1) *Unsafe Sets for Safety*: For safety, the goal is to keep the system's state out of some undesired set, and the disturbance is assumed to be attempting to drive the system

into this unsafe set. The unsafe set relative to the keep-out set K for some time τ is denoted $Pre_\tau(K)$, and is defined as the region of the state space where, for any control inputs $u(t) \in U$, there exists some sequence of disturbances $d(t) \in D$ such that $x(t) \in K$ for some $t \in [-\tau, 0]$. That is, if $x(-\tau) \in Pre_\tau(K)$, then no matter what the control does the disturbance can drive the system into K in time less than or equal to τ . Conversely, if the initial state $x(-\tau)$ is outside of $Pre_\tau(K)$, then there exists a control $u(t)$ that keeps the system out of K for up to time τ .

In the game formulation, the boundary of the initial set K is defined as the zero level set of an appropriately selected cost function $l(x)$ that is negative inside K and positive outside (for an example see Fig. 2). To drive the system into the undesired set, the disturbance is modeled as attempting to minimize $l(x)$, while the control attempts to maximize it. To be conservative, the disturbance is allowed to select its input after the control input is known. The reachability calculations use a dynamic programming formulation and require the optimal Hamiltonian $H^*(x, p)$ for this system, which for a particular mode i is defined as

$$H^*(x, p) = \max_{u \in U} \min_{d \in D} p^T f_i(t, x, u, d) \quad (1)$$

where p is the Hamiltonian costate.

2) *Capture Sets for Attainability*: Symmetrically, for attainability the desired goal of the controller is to drive the system into some desired set D by minimizing the cost function, while the disturbance is assumed to be attempting to drive the system away, thereby maximizing the cost function. Thus a capture set $Pre_\tau(D)$ can be defined for the system which guarantees that there exists a control input that drives the system into D within time τ no matter what the disturbance does. This reverses the role of the control and disturbance, although again for robustness the disturbance is allowed knowledge of the control input. The optimal Hamiltonian in this case is then

$$H^*(x, p) = \min_{u \in U} \max_{d \in D} p^T f_i(t, x, u, d) \quad (2)$$

3) *Reachable Set Calculations*: In many cases (such as the quadrotor maneuvers described below) it can be inconvenient to use the optimal control input described by the differential game, and instead a particular pre-determined controller may be used. This is easily accounted for by defining a controller $u = C_i(t, x)$ specified for the currently active mode. The controller can be subsumed into the modified system dynamics as $\dot{x} = \hat{f}_i(t, x, d)$ and the optimal Hamiltonians for safety and reachability can be modified by removing the max and min, respectively, over u , and leaving only the optimization over d .

Once the dynamics, initial set, and defining cost function are selected, the appropriate Hamiltonian of the system can be formulated and the backwards reachable sets calculated using the Level Set Toolbox developed at the University of British Columbia [15]. Using the given cost function and Hamiltonian, the toolbox propagates the boundary of the level set describing the initial set of interest (either K or D)

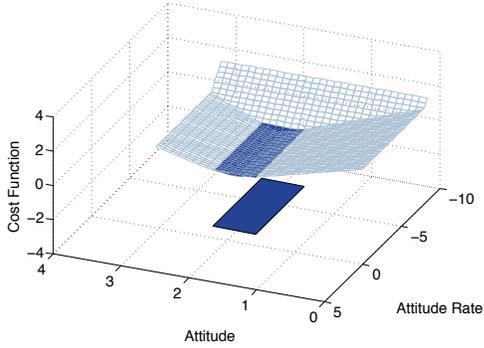


Fig. 2. The zero sub-level set of an appropriate cost function (light blue) is used to define capture or unsafe regions (dark blue) in the state space (in this case, the capture set for attitude in the recovery mode).

backwards in time to find the sets $Pre_\tau(K)$ and $Pre_\tau(D)$. This toolbox has been used in a variety of reachability applications, and was used to compute the results of this paper. The specific implementation details are omitted here, but details on computing reachable sets using the optimal Hamiltonian can be found in the Toolbox documentation.

B. Mode Switching

The application of backwards reachable sets to the design of mode switching control is pursued in the following manner. Given a sequence of n modes and a final desired set D_n , the capture set for the n^{th} mode $Pre_\tau(D_n)$ can be computed for a time horizon τ . Then, for the $(n-1)^{\text{st}}$ mode, a desired set D_{n-1} is chosen where $D_{n-1} \in Pre_\tau(D_n)$, and a capture set $Pre_\tau(D_{n-1})$ can be computed that guarantees that mode $n-1$ will take the state into $Pre_\tau(D_n)$, where mode n can be activated to take the state into D_n . This can be repeated until an initial capture set $Pre_\tau(D_1)$ is found, where if $x(0) \in Pre_\tau(D_1)$, then x is guaranteed to be taken to D_n through the defined modes where the transition from mode i to $i+1$ occurs when $x \in D_i$. A similar process can be conducted for safety calculations, where an initial unsafe set $Pre_\tau(K_1)$ can be found that keeps the system out of K_n through all of the mode transitions.

III. APPLICATION TO QUADROTOR MANEUVER SEQUENCING

A. Experimental Platform

To demonstrate the validity of this approach to maneuver sequencing, this method was used to develop a backflip maneuver for the STARMAC quadrotor helicopter. The vehicle has a total max gross thrust of 3.6kg and a total weight of 1.1kg, which leaves 2.5kg of thrust left for actuation. The vehicle is equipped with a Microstrain 3DMG-X1 inertial measurement unit (IMU) which provides three-axis attitude, attitude rate and acceleration. The resulting attitude estimates are accurate to $\pm 2^\circ$, so long as sustained accelerations are not maintained. Height above the ground is determined using the Senscomp Mini-AE (10m range) sonic ranging sensor which has an accuracy of 3–5cm. Computation and control are managed at two separate levels. The low level control,

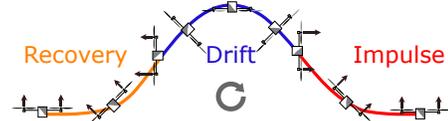


Fig. 3. The backflip maneuver, broken down into three modes. The vehicle travels from right to left, spinning clockwise as it does so. The size of each arrow indicates the relative thrust from each rotor.

which performs real-time control loop execution and outputs PWM motor commands, occurs on an Atmega 128 processor. The high level planning, estimation and control occurs on a lightweight Gumstix Verdex, a PXA270 based single board computer running embedded Linux.

B. Quadrotor Dynamics

To simplify the reachability calculations the quadrotor's dynamics were modeled in a 2D plane, on the assumption (later verified) that the out-of-plane dynamics could be stabilized without affecting the vehicle's performance or safety during the backflip. The planar dynamics are given as:

$$\frac{d}{dt} \begin{bmatrix} x \\ \dot{x} \\ y \\ \dot{y} \\ \phi \\ \dot{\phi} \end{bmatrix} = \begin{bmatrix} \dot{x} \\ -\frac{1}{m} C_D^v \dot{x} \\ \dot{y} \\ -\frac{1}{m} (mg + C_D^v \dot{y}) \\ \dot{\phi} \\ -\frac{1}{I_{yy}} C_D^\phi \dot{\phi} \end{bmatrix} + \begin{bmatrix} 0 \\ D_x \\ 0 \\ D_y \\ 0 \\ D_\phi \end{bmatrix} + \begin{bmatrix} 0 & -\frac{1}{m} \sin \phi & 0 & \frac{1}{m} \cos \phi & 0 & -\frac{l}{I_{yy}} \\ 0 & -\frac{1}{m} \sin \phi & 0 & \frac{1}{m} \cos \phi & 0 & \frac{l}{I_{yy}} \end{bmatrix}^T \begin{bmatrix} T_1 \\ T_2 \end{bmatrix} \quad (3)$$

where the state variables x , y , and ϕ represent the vehicle's lateral, vertical, and rotational motion, respectively; D_x , D_y and D_ϕ are disturbances; and constant system parameters are m for the vehicle's mass, g for gravity, C_D^v for linear drag¹, C_D^ϕ for rotational drag, and I_{yy} for the moment of inertia. Six-dimensional problems are currently not tractable using the Level Set Toolbox, so the system's states were divided into three sets for independent analysis: the rotational dynamics were analyzed to ensure the attainability of the backflip; the vertical dynamics were analyzed to ensure safety (i.e. the vehicle remained above some minimum altitude); and the horizontal dynamics were ignored for simplicity.

C. Backflip Attainability

For the purpose of guaranteeing attainability the backflip was divided into three modes, as shown in Fig. 3: impulse, which initializes the rotation of the vehicle; drift, where the vehicle rotates freely and falls under gravity; and recovery, which brings the vehicle to a controlled hover condition.²

¹For simplicity the vehicle's drag was modeled as linear w.r.t. velocity, an assumption that was later shown via experiment to be sufficiently accurate.

²This division was driven by the fact that unlike a standard helicopter, a quadrotor's blades have a fixed pitch, which means that a quadrotor is only capable of generating thrust in one direction. Thus, to successfully complete a backflip maneuver with a slow rotational rate (e.g. around $400^\circ/\text{sec}$), it was necessary to turn off the motors while the vehicle was inverted to prevent the vehicle from propelling itself into the ground.

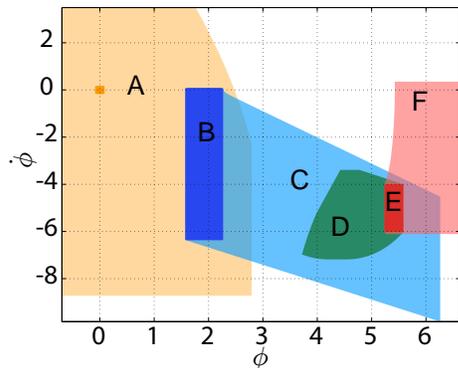


Fig. 4. Composite attainable sets of the backflip maneuver are plotted in the ϕ (radians) vs $\dot{\phi}$ (radians/second) plane. The backflip maneuver starts in the region labeled F and ends in the region labeled A.

Each mode was designed using the method described in Section II-B. To provide further detail on this process, each mode is described in reverse order below.

1) *Recovery*: The final target set for the recovery mode was chosen to be $\phi = 0 \pm 5^\circ$, $\dot{\phi} = 0 \pm 10^\circ/\text{sec}$, essentially a stable hover configuration. As described in Section II-A.3, a fixed controller (in this case a standard PD controller on ϕ) was used to drive the vehicle to this configuration. (The control input u was divided between the two motors as $T_1 = T_{nom} - u$, $T_2 = T_{nom} + u$, where T_{nom} was the nominal total thrust necessary to counteract gravity.) This target set was then propagated backwards using the Hamilton-Jacobi framework, taking into account the worst-case disturbances due to motor noise and wind. It should be noted that the magnitude of these disturbances has a significant impact on the resulting reachable sets; as one would expect, if the potential worst-case disturbances are too large, the vehicle may not be able to reach the target set. For this calculation (and those following) the worst-case disturbances were determined from previous experience, and were on the order of 5% of the total nominal thrust necessary to keep the vehicle aloft. The resulting level set (region A in Fig. 4) represents the capture set for this maneuver.

2) *Drift*: For the drift mode the target set was chosen as $\phi = 110 \pm 20^\circ$, $\dot{\phi} = -180 \pm 185^\circ/\text{sec}$ (region B in Fig. 4), and was again propagated back (this time with no control input, and thus reduced worst-case disturbances due to the lack of motor noise) to produce the capture set for the drift mode (region C of Fig. 4).

3) *Impulse*: Finally, for the impulse mode, the target set was $\phi = 310 \pm 10^\circ$, $\dot{\phi} = -287 \pm 58^\circ/\text{sec}$ (region E of Fig. 4). Once again a fixed controller was used, and the worst-case disturbances were chosen so as to account for motor noise and wind. The resulting capture set is pictured in region F of Fig. 4.

4) *Motor Turnoff*: Originally, it was assumed that the motors would turn off instantaneously when the vehicle entered the drift mode; that assumption proved to be false after some initial experiments. As a result an additional mode was added for the purpose of analysis. In this mode, the

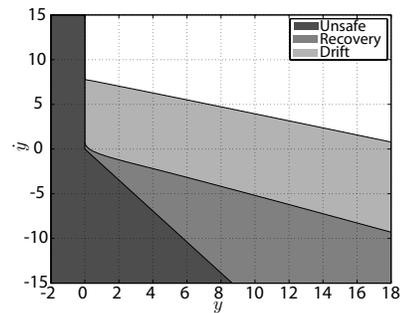


Fig. 5. Unsafe vertical sets of the backflip maneuver are plotted in the y (meters) vs \dot{y} (meters/second) plane. As long as the vehicle begins a given mode outside that mode's unsafe set, safety is guaranteed.

motor turn off was modeled as a linear decay in the vehicle's angular acceleration, i.e.: $\tau = f_{sat}(\alpha t + \ddot{\phi})/I_{yy}$ where $f_{sat}(y) = \{y, \text{if } y < 0; 0, \text{otherwise}\}$, and the parameter α was found using linear regression. These dynamics were then propagated forward from the target set of the impulse mode; the resulting level set (pictured in region D of Fig. 4) contains all possible states the vehicle could be in while the motors are turning off. Thus, as long as this set was contained in the drift set, attainability of the backflip was guaranteed.

D. Backflip Safety

To ensure the vehicle would perform the backflip safely, a similar procedure to that described for attainability was used. First, a final unsafe set was chosen to represent all configurations the vehicle would need to avoid during the recovery mode. Because the vehicle's rotational and vertical dynamics are coupled during powered thrust, however, it was first necessary to find a way to decouple them so that safety could be analyzed solely in the vertical state space. This decoupling was accomplished by taking advantage of the fact that the recovery mode was designed to use a fixed control law. As a result, a nominal trajectory could be generated that could then be plugged into the system dynamics, allowing the backwards reachable set to be computed as usual by propagating it backward for a fixed time T_r , based on the maximum time that the recovery mode could take. The resulting level set indicates all the configurations in which it would be unsafe for the vehicle to enter the recovery mode.

In the drift mode the rotational and vertical dynamics decouple, and so the unsafe set for the drift mode was generated by propagating backward the unsafe recovery set under the vertical dynamics. Once again, this was done for a fixed time T_d , based on the maximum length of the maneuver as calculated from the rotational dynamics. The resulting level set represents all the configurations in which it would be unsafe for the vehicle to enter the drift mode.

For the impulse mode it was assumed that there would be no loss in altitude, due to the fact that the impulse mode was designed so that the vehicle's thrust would always be upward during this mode. The resulting unsafe sets are pictured in Fig. 5; as long as the vehicle began each mode outside that mode's unsafe set, the overall safety of the system was guaranteed. To ensure that the vehicle began the

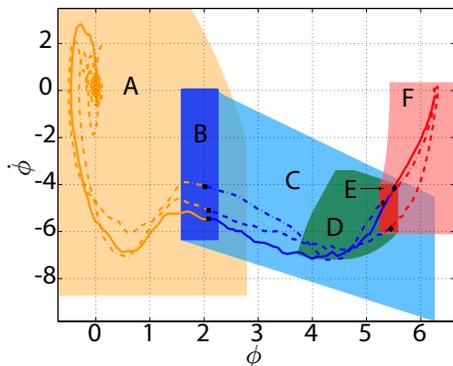


Fig. 6. Three experimental validations (solid, dash and dash-dot lines) of the backflip maneuver overlaid on the composite reach sets. The transitions from the impulse to drift mode are shown as black diamonds which are contained in region E, and the transitions from the drift to the recovery mode are indicated by the black squares that are confined to region B.

entire maneuver outside of these unsafe sets, an additional preliminary climb mode was added before the impulse mode, in which the vehicle would accelerate upward until it reached a safe altitude and velocity.

IV. RESULTS

A mosaic of one of the demonstrations of the backflip maneuver is shown in Fig. 7. Fig. 7(a) depicts the quadrotor after the initial climb mode which is the start of the impulse mode, and Fig. 7(b) is at the end of the impulse mode and at the beginning of the drift mode. Figs. 7(b)-(f) display the entire drift portion of the maneuver and Fig. 7(e) shows the quadrotor inverted. Finally, Figs. 7(f)-(j) display the recovery mode of the backflip maneuver which successfully returns the quadrotor to a safe condition of $\phi = 0^\circ$ and $\dot{\phi} = 0^\circ/\text{sec}$. Video of the backflip maneuver can be viewed at <http://hybrid.eecs.berkeley.edu/aerobatics.html>.

Fig. 6 shows the $(\phi, \dot{\phi})$ trajectory of three experimental validations through the designed attainable sets for the backflip maneuver. As the figure illustrates, the trajectories are contained within the capture sets for each maneuver. The transition between the impulse and drift modes is denoted by a diamond, and the transition between the drift and recovery modes is indicated by a square. The switch between the maneuvers are contained within each of their goal regions, E and B, respectively. When the quadrotor switches into the drift mode, it takes approximately 0.2 seconds for the motors to spin down which explains why the quadrotor is still accelerating at the beginning of the drift maneuver. Table I displays the time spent in each mode for each trial.

TABLE I

THE AMOUNT OF TIME SPENT IN EACH MODE FOR EACH EXPERIMENTAL TRIAL OF THE BACKFLIP MANEUVER.

	Impulse (sec.)	Drift (sec.)	Recovery (sec.)
Trial 1	0.55	0.55	2.56
Trial 2	0.50	0.55	5.54
Trial 3	0.55	0.61	4.70

A. Backflip Safety

Fig. 8 displays the unsafe vertical reachable sets and the switching points for the three trials of the maneuver. The blue points correspond to entries into the drift mode and the orange points correspond to entry into the recovery mode. As the figure illustrates, all the points are outside of their respective unsafe set and therefore the vehicle can safely perform the maneuver without hitting the ground. Fig. 9 displays that the pitch of the quadrotor is within $\pm 5^\circ$ for almost the entire maneuver. This validates the assumption that the backflip maneuver can be modeled in the 2D $(\phi, \dot{\phi})$ plane.

Finally, it should be noted that while the results of only three trials are presented here, several additional trials were also conducted, with varying levels of success. However, these other trials were unsuccessful solely due to factors outside the scope of the reachable set analysis. For example, some trials failed because of human error, in the form of bugs in the code running on the vehicle. Others were unsuccessful due to hardware malfunctions (e.g. a broken sonic ranger, or saturation of the IMU's accelerometers). As every roboticist knows, these sorts of failures are typical when making the transition from theory to "real" engineering.

V. CONCLUSIONS AND FUTURE WORK

Reachable set analysis was successfully used to examine and design an autonomous backflip maneuver. The analysis provided a theoretical proof that under the worst case bounded disturbances, which account for modeling errors and external disturbances such as wind, the quadrotor would safely complete the complicated maneuver. The backflip maneuver was split into three main stages: impulse, drift and recovery in which switching conditions were provided to ensure a safe transition between the three different modes. The aerobatic maneuver was successfully demonstrated on the STARMAC quadrotor platform.

The success of several trials coupled with the failure of others (due to bugs or hardware malfunctions) emphasize how useful the presented methodology can be. Because practical robotics can be fraught with engineering challenges that can cause robots to fail, it is all the more important that as many guarantees as possible about safety be made. While

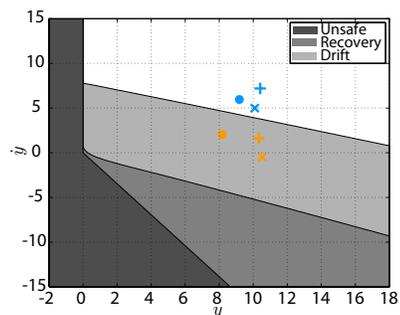


Fig. 8. Three experimental validations (\times , $+$, and \bullet) of the backflip maneuver. The light blue symbols correspond to when the vehicle entered into the drift mode and the orange symbols correspond to when the vehicle entered into the recovery mode. Since all points are outside of their respective reachable set, the vehicle was safe to execute the maneuver.

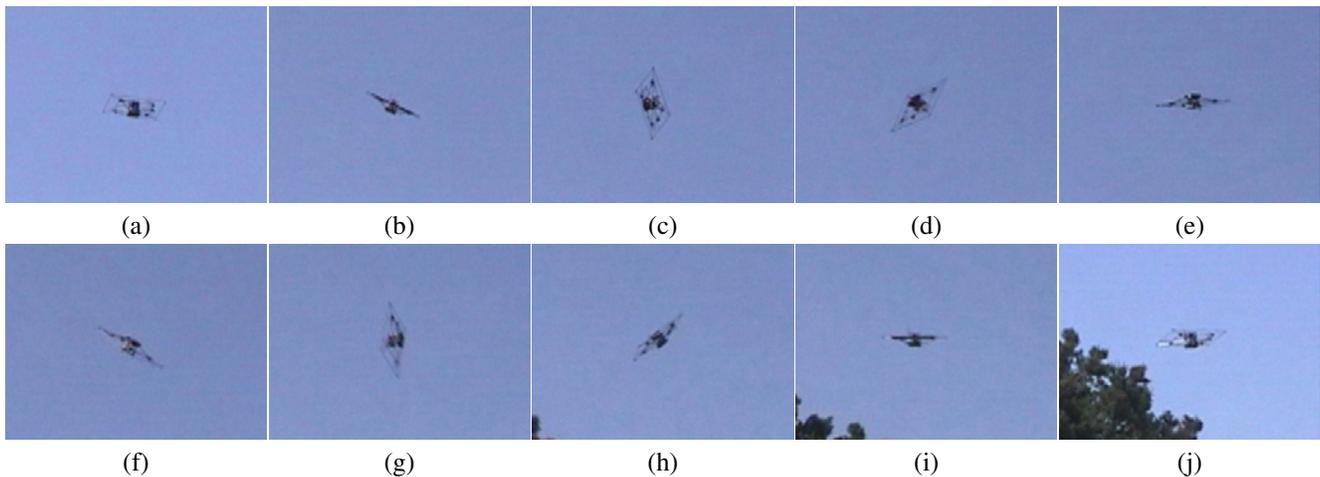


Fig. 7. A mosaic of the successful demonstration of the backflip maneuver. (a): The quadrotor has finished the climb portion of the backflip and is starting the impulse mode. (b): The quadrotor has finished the impulse stage and is entering into the drift portion. (b)-(f): Display of the drift stage of the backflip. (f): The drift mode is concluding and the recovery mode has started. (f)-(j): The recovery mode is safely returning the quadrotor to its hovering position.

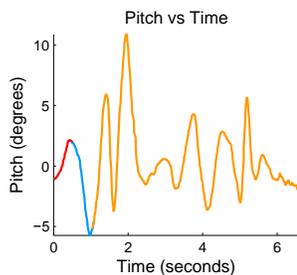


Fig. 9. Experimental data from a single run of the backflip maneuver showing the out-of-plane pitch (degrees) of the quadrotor.

the reachable set framework presented here cannot eliminate all possible failure modes, it can help to ensure that under reasonable circumstances a robot will perform in a provably safe manner.

There are several possible interesting areas of future work. First, a method of describing the reachable sets in a parametric way could be explored. In the current framework, the resulting reachable set is greatly dependent on the parameters used when generating it. Consequently, if the parameters change then the entire reachable set needs to be recalculated offline. Second, the reachable set analysis was performed using a continuous time model, but the implementation of the maneuver was performed on a discrete computing platform. Therefore, the vehicle may pass through the switching regions before its controller can change modes. Even though this was not an issue in the experiments that were conducted, it is still an important area of future work.

ACKNOWLEDGMENTS

The authors would like to thank Tony Mercer for his assistance in performing the flight tests.

REFERENCES

[1] J. H. Gillula, H. Huang, M. P. Vius, and C. J. Tomlin, "Design and analysis of hybrid systems, with applications to robotic aerial vehicles," in *Proceedings of the International Symposium of Robotics Research (ISRR)*, Lucerne, Switzerland, August 2009.

[2] A. Coates, P. Abbeel, and A. Y. Ng, "Apprenticeship learning for helicopter control," *Commun. ACM*, vol. 52, no. 7, pp. 97–105, 2009.

[3] V. Gavrillets, I. Martinos, B. Mettler, and E. Feron, "Flight test and simulation results for an autonomous aerobatic helicopter," in *Proceedings of the 21st Digital Avionics Systems Conference*, 2002, pp. 8.C.3/1–6, doi: 10.1109/DASC.2002.1052943.

[4] T. Lozano-Perez, M. T. Mason, and R. H. Taylor, "Automatic Synthesis of Fine-Motion Strategies for Robots," *The International Journal of Robotics Research*, vol. 3, no. 1, pp. 3–24, 1984.

[5] R. Burridge, A. Rizzi, and D. Koditschek, "Sequential composition of dynamically dexterous robot behaviors," *International Journal of Robotics Research*, vol. 18, no. 6, pp. 534–555, 1999.

[6] H. Kress-Gazit, G. E. Fainekos, and G. J. Pappas, "Translating structured english to robot controllers," *Advanced Robotics Special Issue on Selected Papers from IROS 2007*, vol. 22, no. 12, pp. 1343–1359, 2008.

[7] E. Frazzoli, M. A. Dahleh, and E. Feron, "Maneuver-based motion planning for nonlinear systems with symmetries," *IEEE Transactions on Robotics*, vol. 21, no. 6, pp. 1077–1091, 2005.

[8] M. Zefran and J. W. Burdick, "Stabilization of systems with changing dynamics," in *Hybrid Systems: Computation and Control Conference*, Berg en Dal, the Netherlands, 1999.

[9] M. Lazar and A. Jokic, "Synthesis of trajectory-dependent control lyapunov functions by a single linear program," in *Hybrid Systems: Computation and Control Conference*, San Francisco, California, 2009.

[10] M. Egerstedt, T. J. Koo, F. Hoffmann, and S. Sastry, "Path planning and flight controller scheduling for an autonomous helicopter," in *Hybrid Systems: Computation and Control Conference*, Berkeley, California, 1998.

[11] G. M. Hoffmann, H. Huang, S. L. Waslander, and C. J. Tomlin, "Quadrotor helicopter flight dynamics and control: Theory and experiment," in *Proceedings of the AIAA Guidance, Navigation, and Control Conference and Exhibit*, Hilton Head, SC, August 2007.

[12] I. M. Mitchell, A. M. Bayen, and C. J. Tomlin, "A time-dependent hamilton-jacobi formulation of reachable sets for continuous dynamic games," *IEEE Transactions on Automatic Control*, vol. 50, no. 7, pp. 947–957, 2005.

[13] M. Oishi, I. M. Mitchell, A. M. Bayen, and C. J. Tomlin, "Invariance-preserving abstractions of hybrid systems: Application to user interface design," *IEEE Transactions on Control Systems Technology*, vol. 16, no. 2, pp. 229–244, 2008.

[14] J. Ding, J. Sprinkle, S. S. Sastry, and C. J. Tomlin, "Reachability calculations for automated aerial refueling," in *Proceedings of the 47th IEEE Conference on Decision and Control*, Cancun, Mexico, 2008.

[15] I. Mitchell, *A Toolbox of Level Set Methods*, 2009, <http://people.cs.ubc.ca/~mitchell/ToolboxLS/index.html>.