Benchmark: Flight Envelope Protection in Autonomous Quadrotors

Shahab Kaynama* and Claire J. Tomlin*

v0.1, 2014-3-22

Abstract

The flight envelope represents the safety bounds that arise from aerodynamical and physical limitations of an aerial vehicle. The protection of this envelope consists of (a) analyzing the subset of states, known as the viability or discriminating kernel, for which all *hard* input and state constraints can be satisfied despite external disturbances or model uncertainties, and (b) designing controllers that ensure such constraint satisfaction even in the presence of conflicting performance objectives. Due to the highly complex and under-actuated nature of quadrotors, their flight envelope protection makes for a well-suited benchmark for reachability/viability analysis and safety-preserving controller synthesis.

Category: academic and industrial **Difficulty:** medium to high

1 Context and Origins

The ability to respect hard constraints despite potentially conflicting performance objectives is key in safe operation of many cyberphysical systems. Autonomous quadrotors are a perfect example of such safety-critical systems: On the one hand, the actuators have limited authority in their power throughput, and on the other hand, there are constraints on the state of the system that arise due to flight dynamics and physical-structural limitations. These constraints constitute the *flight envelope* of the system which would have to be analyzed and protected when designing a controller.

The full-dimensional model of a quadrotor consists of twelve states (position, flight angles, and linear and angular velocities) and four inputs (thrust and angular accelerations). The high dimensionality of the model presents a challenge in successfully applying reachability analysis/viability computations, and therefore limits one's ability to synthesize safety-preserving controllers.

In academia, quadrotors are commonly used as a robotics platform for testing of the state-of-the-art control algorithms. But their presence is not limited

*Electrical Engineering and Computer Sciences, University of California at Berkeley

to academic circles: Aside from military and surveillance applications, quadrotors are increasingly employed as a civilian technology to accomplish complex tasks such as monitoring and exploration of hazardous areas [1], in search and rescue missions [2], for delivery of antibiotics in underprivileged countries with inadequate transportation infrastructure [3], and HD filming of sports events [4].

Ensuring that the flight envelope can remain protected is essential in safe and cost effective operation of these systems. This is a challenging problem because the implementation of any relevant control strategy requires the computation of the viability kernel—a task that is particularly hard in such high dimensions. Even for a conventional scheme such as the model predictive control, the maximal controlled-invariant set (infinite-horizon viability kernel) is needed so as to ensure recursive feasibility of the corresponding receding horizon optimization.

The model described here, detailed by Cowling et al. [5], is based on a six degree of freedom Newton-Euler rigid body equations of motion for a particular quadrotor. However, it is generic enough to be readily applicable to other types of quadrotors. The resulting differential equation is normalized so that it is independent of the mass of the vehicle.

2 Description

The states of the system

$$x = \begin{bmatrix} \mathrm{x} & \mathrm{y} & \mathrm{z} & \dot{\mathrm{x}} & \dot{\mathrm{y}} & \dot{\mathrm{z}} & \phi & \theta & \psi & \dot{\phi} & \dot{\theta} & \dot{\psi} \end{bmatrix}^\top \in \mathbb{R}^{12} \tag{1}$$

is comprised of translational positions in [m] with respect to a global origin, their derivatives (linear velocities in x, y, z directions) in [m/s], the Eulerian angles, roll ϕ, pitch θ, and yaw ψ in [rad], and their respective derivatives (angular velocities) in [rad/s]. The control input is the vector $u = [u_1\ u_2\ u_3\ u_4]^\top \in \mathbb{R}^4$ consisting respectively of the total thrust in [m/s^2] normalized with respect to the mass of the vehicle and the second-order derivatives $\ddot{\phi}$, $\ddot{\theta}$, $\ddot{\psi}$ of the Eulerian angles in [rad/s^2]. The system is under-actuated since there are six degrees of freedom but only four actuators.

The normalized equations of motion that describe the dynamics are

$$\ddot{\mathrm{x}} = -u_1 \cos(\phi) \sin(\theta), \tag{2a}$$
$$\ddot{\mathrm{y}} = u_1 \sin(\phi), \tag{2b}$$
$$\ddot{\mathrm{z}} = -g + u_1 \cos(\phi) \cos(\theta), \tag{2c}$$
$$\ddot{\phi} = u_2, \tag{2d}$$
$$\ddot{\theta} = u_3, \tag{2e}$$
$$\ddot{\psi} = u_4, \tag{2f}$$

with $g \approx 9.81$ being the acceleration of gravity.

The viability of the system could either be established globally for the nonlinear case using the above dynamics, or locally around the hover condition, for example, by linearizing these differential equations. The constraints on the state and input vectors will be slightly different for each case:

1. **Nonlinear case:** The angles ϕ and θ and the speed profile $V := \|[\dot{\mathrm{x}}\ \dot{\mathrm{y}}\ \dot{\mathrm{z}}]\|$ are bounded as $\phi, \theta \in [-\frac{\pi}{4}, \frac{\pi}{4}]$ and $V \leq 5$. The angular velocities are constrained as $\dot{\phi}, \dot{\theta}, \dot{\psi} \in [-3, 3]$. We further assume that the vehicle must safely fly within the range of 1 to 7 m above the ground in z direction in an environment that stretches 6 m in each direction in the x-y plane. These constraints form the flight envelope $\mathcal{K}$. The input vector u is constrained by the hyper-rectangle $\mathcal{U} := [0, g + 2.38] \times [-0.5, 0.5]^3$.

2. **Linear case (hover mode):** By linearizing the equations of motion (2) about the hover condition $\phi = 0$, $\theta = 0$, and $u_1 = g$, we obtain the linear time-invariant (LTI) dynamics $\dot{\bar{x}} = A\bar{x} + B\bar{u}$ where the bar notation is used for the deviations from the equilibrium. The system matrices A and B can be found in the Appendix. The flight envelope $\mathcal{K}$ is essentially the same as in the nonlinear case. The new input vector $\bar{u}$ is constrained by the hyper-rectangle $\mathcal{U} := [-g, 2.38] \times [-0.5, 0.5]^3$.

Note that due to the agility of the quadrotor, it can travel a significant distance in a short time interval. The safety problem can be tackled in continuous-time or sampled-data[1] frameworks (see Section 4 for arguments about the discrete-time case). In describing these two frameworks, we ignore the bar notation and denote the state and the control input respectively as x and u for both linear and nonlinear cases.

1. **Continuous-time framework:** For a given time horizon τ in [s] (possibly infinite) we wish to compute the set of initial states $x_0 := x(0)$ for which there exists a *measurable* control signal $u(\cdot)$ satisfying $u(t) \in \mathcal{U}$ almost everywhere such that the resulting trajectory $x^u_{x_0}(\cdot)$ satisfies $x^u_{x_0}(t) \in \mathcal{K}$ over at least $\mathbb{T} := [0, \tau]$. This set is the continuous-time viability kernel

$$\mathrm{Viab}_{\mathbb{T}}(\mathcal{K}) := \left\{x_0 \in \mathcal{K} \mid \exists u(\cdot) \in \mathscr{U}_{\mathbb{T}},\, \forall t \in \mathbb{T},\, x^u_{x_0}(t) \in \mathcal{K}\right\} \tag{3}$$

   with

$$\mathscr{U}_{\mathbb{T}} := \{u \colon \mathbb{T} \to \mathbb{R}^4 \text{ measurable},\ \ u(t) \in \mathcal{U} \text{ a.e. } t \in \mathbb{T}\}. \tag{4}$$

   A Lebesgue measurable control may not be implementable in practice.

2. **Sampled-data framework:** If the state of the quadrotor can only be measured at fixed sampling intervals $\delta \in \mathbb{R}_{>0}$, a sampled-data framework represents the problem setup more adequately. In that case, let $N_\delta := \lceil \tau/\delta \rceil$ denote the number of sampling intervals in $\mathbb{T}$ and consider the case where the input is applied at the beginning of each sampling interval; during the interval the input is kept constant until the next sampling

[1] In a sampled-data system, the dynamics are allowed to evolved continuously in time, but the control is restricted to the class of piecewise constant functions; e.g. an analog system controlled by a digital platform.

instant (zero-order hold). We now wish to compute the set of initial states x_0 for which there exists a *piecewise constant* control signal $u(\cdot)$ satisfying $u(t) \in \mathcal{U}$ such that $x^u_{x_0}(t) \in \mathcal{K}$ over at least $\mathbb{T} := [0, \tau]$. This is the sampled-data viability kernel

$$\mathrm{Viab}^{\mathrm{sd}}_{\mathbb{T}}(\mathcal{K}) := \big\{x_0 \in \mathcal{K} \mid \exists u(\cdot) \in \mathscr{U}^{\mathrm{pwc}}_{\mathbb{T}},\, \forall t \in \mathbb{T},\, x^u_{x_0}(t) \in \mathcal{K}\big\}, \tag{5}$$

where

$$\mathscr{U}^{\mathrm{pwc}}_{\mathbb{T}} := \{u \colon \mathbb{T} \to \mathbb{R}^4 \text{ piecewise const.}, \\ u(t_k) \in \mathcal{U}\ \forall k \in \{0, \ldots, N_\delta\},\ u(t) = u(t_k)\ \forall t \in [t_k, t_{k+1})\}. \tag{6}$$

3 Key Observations

For safety, an approximation of the viability kernel must be conservative—i.e. in form of an *under-approximation*; an over-approximation includes states for which one or more of the constraints would inevitably be violated.

The presented benchmark problem can be primarily used to assess the performance of *any algorithm* that attempts to automatically generate an under-approximation of the viability kernel. Additionally, the algorithm may synthesize the associated safety-preserving control laws (inputs that enforce the desired constraint satisfaction) and test them in protecting the flight envelope.

Several properties can serve as quality indicators for approximations generated by the algorithm:

1. **Conservatism:** Check if every state that belongs to the approximation, also belongs to the true (unknown) viability kernel. In other words, for every state in the approximate set, does there exists an admissible control signal $u(\cdot)$ that keeps the system contained in $\mathcal{K}$ over at least $\mathbb{T}$?

2. **Accuracy:** Either a quantitative measure (via some metric of choice, e.g. volume, Hausdorff distance) or a qualitative measure (by also computing an over-approximation to compare against the under-approximation) of accuracy can be provided. When there is a tradeoff between accuracy and computational complexity, it can be explicitly stated in terms of the inputs to the algorithm.

3. **Convergence:** The rate of convergence of the under-approximation to the true kernel can be a useful indication of the performance of the algorithm. Whether the algorithm can asymptotically recover the true kernel must be discussed. If not, to what subset does the algorithm converge?

4. **Termination:** Even when the horizon τ is finite, the algorithm should be able to terminate once the under-approximate set is *controlled-invariant* (meaning there exists an admissible control that make the set positively invariant under the closed-loop dynamics). Alternatively, the algorithm should be able to test and verify that a generated set is controlled-invariant.

4 Outlook and Possible Variants

Robustness: The unmodeled dynamics that may arise due to the ground effect, air drag, the turbulence due to the flow of air within the propellers, and other complicated or otherwise significant nonlinearities are neglected here. However, these uncertainties can be treated as an additive, bounded but known disturbance to the system if necessary. Exogenous perturbations can also be taken into account by augmenting the additive disturbance input (for example, the effect of wind can be treated as perturbations to ẋ, ẏ, and ż).

Robustifying the analysis against an unknown, Lebesgue measurable additive disturbance that draws values from a bounded set $\mathcal{V}$ results in a zero-sum differential game in the continuous-time case. The winning domain of the control input, i.e. the set of initial states for which there exists an admissible safety-preserving control despite the worst-case actions of the disturbance, is known as the *discriminating kernel* of $\mathcal{K}$:

$$\mathrm{Disc}_{\mathbb{T}}(\mathcal{K},\mathcal{U},\mathcal{V}) := \left\{x_0 \in \mathcal{K} \mid \forall v(\cdot) \in \mathscr{V}_{\mathbb{T}}^{\mathrm{na}},\, \exists u(\cdot) \in \mathscr{U}_{\mathbb{T}}^{\mathrm{fb}},\, \forall t \in \mathbb{T},\, x_{x_0}^{u,v}(t) \in \mathcal{K}\right\} \tag{7}$$

with $\mathscr{U}_{\mathbb{T}}^{\mathrm{fb}}$ denoting the class of possibly time-dependent state feedback control signals and $\mathscr{V}_{\mathbb{T}}^{\mathrm{na}}$ denoting the class of non-anticipative disturbance signals taking values pointwise in $\mathcal{V}$. For more discussions surrounding this formulation as well as the rationale behind the order of quantifiers in (8) please refer to [6].

The sampled-data case is simpler since the inputs need not react to one another instantaneously. In that case, the kernel to be approximated is

$$\mathrm{Disc}_{\mathbb{T}}^{\mathrm{sd}}(\mathcal{K},\mathcal{U},\mathcal{V}) := \left\{x_0 \in \mathcal{K} \mid \exists u(\cdot) \in \mathscr{U}_{\mathbb{T}}^{\mathrm{pwc}},\, \forall v(\cdot) \in \mathscr{V}_{\mathbb{T}},\, \forall t \in \mathbb{T},\, x_{x_0}^{u,v}(t) \in \mathcal{K}\right\}, \tag{8}$$

where $\mathscr{V}_{\mathbb{T}}$ denotes the class of measurable signals from $\mathbb{T}$ to $\mathcal{V}$. (cf. [7])

Simplified dynamics: It is possible to simplify the problem and deal with a lower dimensional model by restricting the movement of the quadrotor and suppressing the degrees of freedom. For example, limiting the vehicle to fly only along the z-axis without turning or twisting yields a double-integrator dynamics.

Discrete time: Discretizing the dynamics and performing the analysis in discrete time may not be ideal due to the agility and safety-critical nature of the quadrotor. For instance, with a sampling frequency of 10 Hz the vehicle could travel a distance of half a meter in between each two consecutive time instants. That said, with a sufficiently high sampling frequency one can attempt to approximate the discrete-time viability kernel and use that approximation for safety analysis (and synthesis).

A Appendix

Linearizing the equations of motion (2) about the hover condition yields

$$A = \begin{bmatrix} 0&0&0&1&0&0&0&0&0&0&0&0\\ 0&0&0&0&1&0&0&0&0&0&0&0\\ 0&0&0&0&0&1&0&0&0&0&0&0\\ 0&0&0&0&0&0&0&-g&0&0&0&0\\ 0&0&0&0&0&0&g&0&0&0&0&0\\ 0&0&0&0&0&0&0&0&0&0&0&0\\ 0&0&0&0&0&0&0&0&0&1&0&0\\ 0&0&0&0&0&0&0&0&0&0&1&0\\ 0&0&0&0&0&0&0&0&0&0&0&1\\ 0&0&0&0&0&0&0&0&0&0&0&0\\ 0&0&0&0&0&0&0&0&0&0&0&0\\ 0&0&0&0&0&0&0&0&0&0&0&0 \end{bmatrix}, \quad B = \begin{bmatrix} 0&0&0&0\\ 0&0&0&0\\ 0&0&0&0\\ 0&0&0&0\\ 0&0&0&0\\ 1&0&0&0\\ 0&0&0&0\\ 0&0&0&0\\ 0&0&0&0\\ 0&1&0&0\\ 0&0&1&0\\ 0&0&0&1 \end{bmatrix}.$$

As mentioned in Section 3, the capabilities of an algorithm can be tested by computing the viability kernel (or the discriminating kernel in the presence of a disturbance input) and by possibly synthesizing the corresponding safety-preserving control laws. We refrain from enlisting all existing algorithms, but only mention in passing that for the continuous-time LTI case above, a permissive, scalable, robust safety-preserving switched control strategy has been proposed in [6]. Instead, as an example, we will discuss the performance of a new algorithm [8] for approximation of the *sampled-data* viability kernel. Figure 1 shows selected 2D projections of a polytopic approximation of $\mathrm{Viab}^{\mathrm{sd}}_{[0,2]}(\mathcal{K})$ (with sampling interval $\delta = 0.1$). The algorithm, which strikes a direct tradeoff between accuracy and computational complexity, requires about 3 seconds (on a common laptop computer, in Matlab, without optimizing the code for speed) to generate a new vertex of the polytopic under-approximation and 5 seconds to generate a facet of the over-approximation.

The under-approximation is tight in the sense that each vertex of the polytope belongs to the boundary of the true viability kernel with some *a priori* known accuracy. The over-approximation is also tight in that each facet of the polytope touches the boundary of the true kernel in at least one point. The two approximating sets sandwich the boundary of the viability kernel to within a certain precision in at least half of the number of vertices of the under-approximation, providing an added layer of confidence about the precise location of the kernel. The tightness of the sets are unfortunately unobservable in the projection plots shown in the figure.

References

[1] J. Bradley. (2013, May) (CNN) Tiny flying robots! Meet the quadrotor. [Online]. Available: http://goo.gl/Tx0LT4

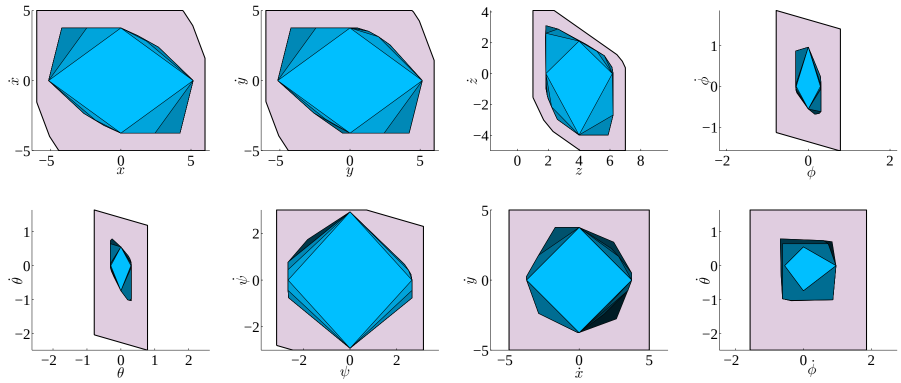

Figure 1: Selected 2D projections of a polytopic approximation [8] of the sampled-data viability kernel for the linearized version of the benchmark example. Under-approximations for $N = 24, 48, 96, 500, 1000, 2000, 3000$ vertices are shown with $N = 24$ in the lightest shade of blue (innermost set), and $N = 3000$ in the darkest shade of blue. An over-approximation (outermost set) with 1500 facets is also shown in lavender. The approximations are tight and touch each other to within a constant accuracy in $N/2$ points—a fact that is unrecognizable due to projections.

[2] Royal Canadian Mounted Police. (2013, May) Single vehicle rollover – Saskatoon RCMP search for injured driver with unmanned aerial vehicle. [Online]. Available: http://goo.gl/TvVxA

[3] T. Lee. (2013, Dec.) (The Washington Post) Low-income countries might get drone deliveries before the U.S. Here's why. [Online]. Available: http://wapo.st/1g3YLK8

[4] DediCam. [Online]. Available: http://dedicam.tv

[5] I. Cowling, O. Yakimenko, J. Whidborne, and A. Cooke, "Direct method based control system for an autonomous quadrotor," *Journal of Intelligent & Robotic Systems*, vol. 60, pp. 285–316, 2010.

[6] S. Kaynama, I. Mitchell, M. Oishi, and G. Dumont, "Scalable safety-preserving robust control synthesis for continuous-time linear systems," *IEEE Transactions on Automatic Control*, (accepted for publication; preprint at arXiv:1312.3399).

[7] I. Mitchell, S. Kaynama, M. Chen, and M. Oishi, "Safety preserving control synthesis for sampled data systems," *Nonlinear Analysis: Hybrid Systems*, vol. 10, pp. 63–82, 2013.

[8] J. Gillula, S. Kaynama, and C. Tomlin, "Sampling-based approximation of the viability kernel for high-dimensional linear sampled-data systems," in *Hybrid Systems: Computation and Control*, Berlin, Germany, April 2014.