2018 18th International Conference on Control, Automation and Systems (ICCAS 2018)
Oct. 17–20, 2018; YongPyong Resort, PyeongChang, GangWon, Korea

Simulated Physics for High Speed Aerial Systems

MinKu Kang[1] and Kee-Eung Kim[2*]

[1]School of Computing, KAIST (tominku@kaist.ac.kr)
[2]School of Computing, KAIST (kekim@cs.kaist.ac.kr) * Corresponding author

Abstract: In this work, we introduce a model of an aerial system based on a physics-based simulation engine. We investigate some basic properties of the proposed model, showing its potential benefit for autonomous control.

Keywords: Aerial System, Physics-based Simulation Environment, Reinforcement Learning, Self Stability, Simulation-based Optimization

1. INTRODUCTION

Building an accurate aerial system simulator requires considerable expert knowledge on the related fields (e.g. aerodynamics [3]), and sometimes significant amount of computation time (e.g. Computational Fluid Dynamics). In this work, we propose a simple aerial system simulation environment, which is built on top of a computationally efficient 3D physics engine called MuJoCo [1]. We introduce some basic properties of our proposed environment, which renders our environment as a testbed for simulation-based optimization of control, such as reinforcement learning [2].

2. ENVIRONMENT

We adopt a simple phenomenological model of fluid dyanmics implemented in MuJoCo physics engine [1]. According to the model, the aerodynamic force and torque exerted on i-th component of the body frame are:

$$F_i \propto -\rho s_j s_k |v_i|\, v_i \tag{1}$$

$$T_i \propto -\rho s_i \left(s_j^4 + s_k^4\right)|\omega_i|\, \omega_i, \tag{2}$$

where v_i is the velocity of the system in the i-th component of the body frame (local frame) axes, s_i is the area of the surface of the equivalent inertia box (the right side of Figure 1) corresponding to the i-th component of the body frame axes. Note that, in the model, the magnitude of the aerodynamic force is proportional to the square of the velocity, highly dissipating the energy of a fast moving aerial system. This simple phenomenological model is fit for the case when inertial effects are dominant over viscous effects [1] (e.g. a high speed aerial system moving across the air with high Reynolds number). At every simulation time step, the physics engine computes the acceleration of the rigid body based on various forces acting on it, including external control forces, gravitational and aerodynamics forces. The simulator then generates the trajectory of the system by numerically integrating (e.g. via RK4) the forward dynamics of the system over time. The resulting trajectory (data) of the system can then be fed as trajectory samples into data-driven control optimization methods such as reinforcement learning.

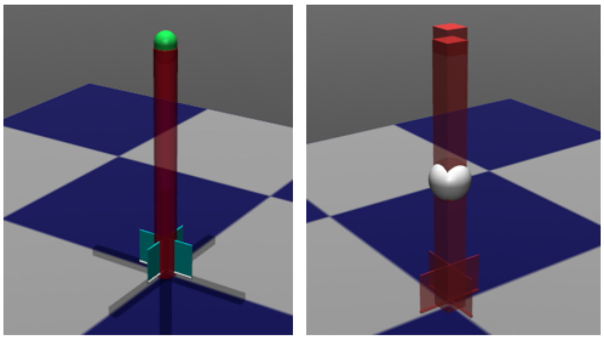

Fig. 1 The aerial system model with static fins attached at the bottom for self-stability. The longitudinal length of it is 5.2 m, and the mass is 81 kg, with a gravitational acceleration of 9.8 m/s^2 and air density of 1.0 kg/m^3 being applied. An equivalent inertia box (right) is shown as red boxes with the center of mass (COM) indicated as a white sphere. The transparent black rods represent three axes for thrusts (one for axial thrust and the other two for lateral thrusts which additionally induce torques due to the moment arms). Our work rather focuses on the topology of the model rather than its physical specifications.

The MuJoCo physics engine also allows one to define a callback function through which customized passive forces (e.g. aerodynamic forces) can be applied to the target system at every simulation time step. Note that the data-driven control methods, which rely on forward-trajectory samples from the physics-based simulator, are different to classical control methods for aerial system where originally nonlinear aerial system dynamics models (e.g. ODE) are linearized around equilibrium points, rather than integrated over time [5].

2.1 Stability under feed-forward control

In many cases it is desirable for a physical system to have some degree of damping mechanism which dissipates the energy of the system over time [4]. This dissipative behavior could be highly desirable especially for an RL (reinforcement learning) agent where boundedness of the trajectories could be important for the stability of

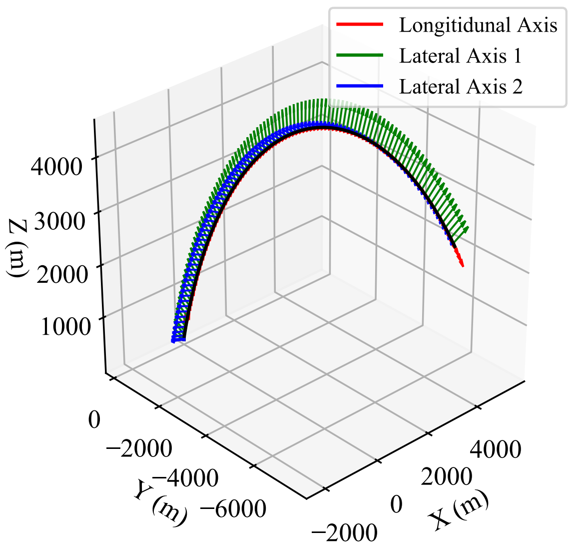

Fig. 2 A nominal trajectory of the aerial system only under feed-forward thrusts. A stable parabolic trajectory is maintained over a long flight time (50 seconds). The arrows along the trajectory indicate the axes of the body frame (local frame) of the aerial system, with the solid black curve being the position trajectory.

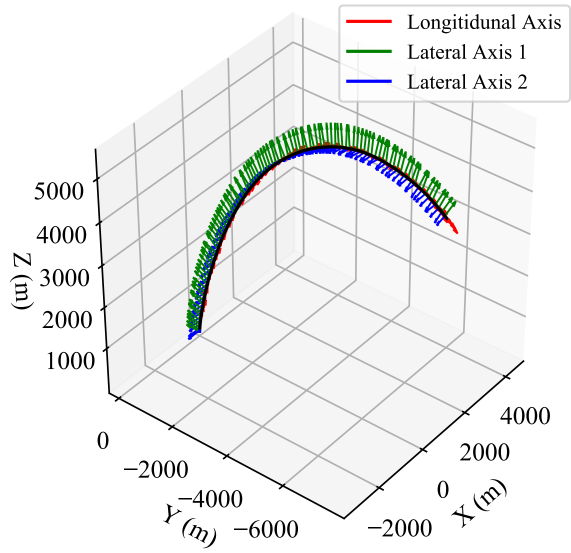

Fig. 3 A trajectory of the aerial system under control disturbances. The trajectory is kept stable overall due to the stabilizing characteristics of the attached fins even under the disturbances.

optimization process. While excessive damping forces can prevent an agent from exploring the state space sufficiently, we argue that at least some degree of damping mechanism should be equipped in a physical system to make the trajectories bounded in some sense, thus securing some degree of learning stability for the RL agent.

While there exist various types of stability [4], we focus on the statistics of angle of attacks along a trajectory as a measure of open-loop stability. We apply feed-forward control (i.e. open-loop control) signals to the aerial system through the axial (i.e. longitudinal) thrust channel along with small magnitude of lateral thrusts for 50 seconds. The axial control signal is equivalent to the longitudinal thrust of 15,000 N, while the lateral thrust is of 75 N which results in a slight yaw rate (a slight turn). The aerial system is launched in a straight upright attitude, resulting in a parabolic trajectory (Figure 2). Note that even under the open-loop control (no feedback), the attitude of the aerial system is well aligned with the velocity direction (i.e. low angle of attack) in overall. More specifically, the longitudinal axes of the system are almost in consistence with the tangent vectors of the position trajectory (Figure 2). This self-stabilizing characteristic could be highly desirable especially in simulation-based control optimization such as reinforcement learning, where radical trajectory changes due to controls can significantly deteriorate the stability of optimization. Considering the fact that the aerial systems move across larger spaces compared to ground systems, unboundedness of the trajectories would be more severe if no self-stabilizing mechanism is equipped in the system.

2.2 Stability under disturbances

We further investigate the stability (in the sense of AOAs) characteristics of the aerial system by injecting disturbances (i.e. forces) into the thrust channels (i.e. control channels). An additive axial disturbance is sampled from a normal distribution $\mathcal{N}\left(0, 1500^2\right)$ N, and it is applied to the axial thrust channel every time step. Lateral thrust disturbances are sampled from $\mathcal{N}\left(0, 150^2\right)$ N and applied to the lateral thrust channels every time step. The resultant trajectory is represented in Figure 3. It shows that the system still maintains a stable (i.e. low AOAs) parabolic trajectory even under the disturbances, with a slight deviation from the nominal one (i.e. disturbance-free trajectory) in Figure 2.

Furthermore, to see the effect of the static fins as a stabilizer, we conduct an additional experiment with an aerial system with the fins removed. Figure 5 shows the resultant trajectory of the system with no fins under the same disturbance setting. The trajectory certainly indicates that the system is quite unstable, meaning that it cannot maintain a consistent parabolic trajectory. The system quickly falls down and wanders around the ground. The results seemingly imply that the fins attached at the bottom of the aerial system serve as a trajectory stabilizer.

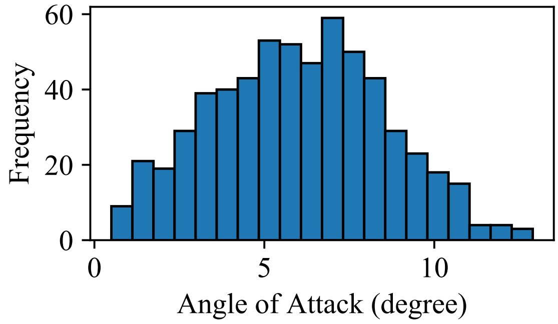

Fig. 4 A statistics of angle of attacks along the trajectory of an aerial system with fins. The mean is 5.5 deg. and the standard deviation is 2.2 deg. The AOAs are maintained at a low level indicating some degree of stability of trajectory.

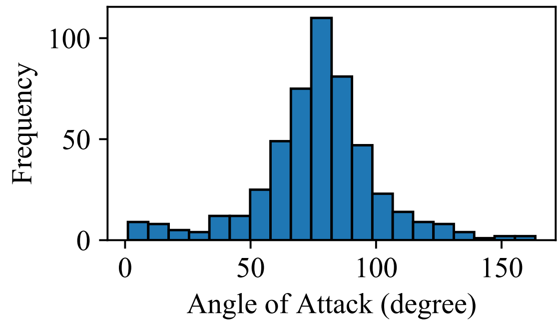

Fig. 6 A statistics of angle of attacks along the trajectory of an aerial system without fins. The mean is 79 deg. and the standard deviation is 24 deg. The AOAs are spread over a wide range with the high mean value, indicating an instability of the trajectory.

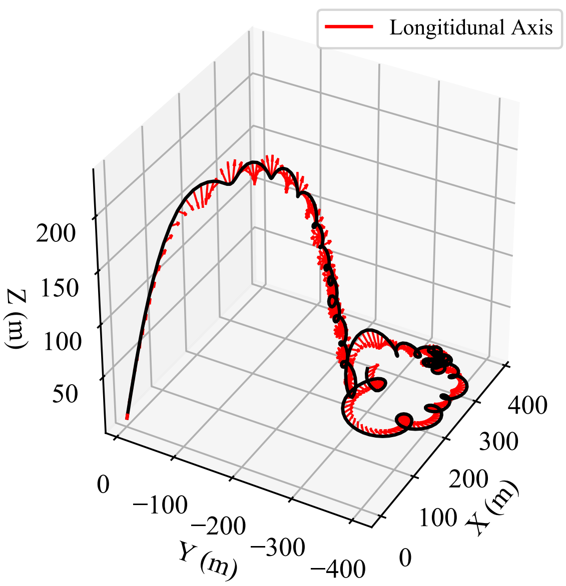

Fig. 5 A trajectory of an aerial system with the fins removed with the other settings unchanged. The longitudinal axis of the system is not aligned well with the tangent vector of the position trajectory (the black solid curve). The aerial system hits the ground at some point and then wanders around it.

To show the result more quantitatively, we measure the angle of attacks along trajectories. The angle of attack, a bit informally, is the angle between the velocity direction and the longitudinal direction of a system (refer [3][5] for a more formal definition). While a large discrepancy between the two vectors (i.e. high angle of attack) could be desirable in some cases, this quantity is usually kept in a small bounded range ($<$ 90 deg.) [3][5]. The two histograms (Figure 4, Figure 6) show the statistics of angle of attacks measured along each trajectory of the system with fins and without fins respectively. Overall, the angle of attack is maintained at a desirable low level in case the fins are attached. In case the fins are removed, however, the angle of attacks are spread over a more wider range with a higher mean value indicating that the system has some degree of instability (Figure 5).

Especially for data-driven control optimization methods such as reinforcement learning whose optimization performance highly relies on the quality of the trajectories, the seemingly unpredictable and unstable trajectory (Figure 5) could make the optimization procedure quite unstable. Even worse, since an RL agent interacts with the environment with a stochastic policy during the learning process, where random thrusts (actions) sampled from the policy are exerted on the system, the trajectory can be more unpredictable and can change radically resulting in high degree of learning instability.

2.3 Trajectories under Varying Air Densities

Eqs. (1) $\sim$ (2) imply that the trajectory of an aerial system might depend on the density of the medium (i.e. air density, ρ) the system travels across. To see the effect of varying air densities on trajectories, we collect trajectories under various air density settings (Figure 7). Note that, since an aerial system under high air density experiences high damping forces according to eqs. (1) $\sim$ (2), the overall travel distance of the system is short compared to the one under low air density (Figure 7). In other words, these resultant trajectories are consistent with our intuition based on the phenomenological model.

3. CONCLUSION AND FUTURE WORK

In this work, we proposed a simple aerial system agent simulated in an environment built on top of an efficient physics-based engine where a simple phenomenological aerodynamics model was adopted. We investigated some

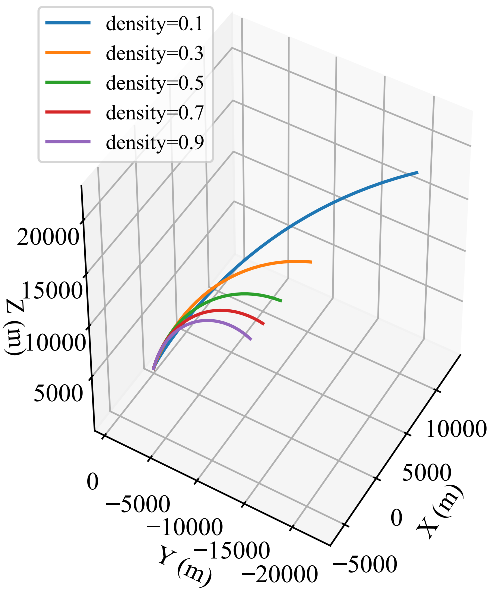

Fig. 7 Trajectories of an aerial system under different air densities (unit in kg/m^3). The travel distances of systems in dense medium are short due to the dissipative aerodynamic forces which depend on density of medium.

basic properties of the proposed model such as self stability which is exhibited in most aerial systems. We leave simulation-based control optimization of the aerial system as a future work.

4. ACKNOWLEDGMENTS

This work was conducted at High-Speed system Research Center of KAIST with the support of the Defense Acquisition Program Administration and the Agency for Defense Development under Contract UD170018CD.

REFERENCES

[1] Todorov, Emanuel, Tom Erez, and Yuval Tassa. ”Mujoco: A physics engine for model-based control.” Intelligent Robots and Systems (IROS), 2012 IEEE/RSJ International Conference on. IEEE, 2012.

[2] Sutton, Richard S., and Andrew G. Barto. Reinforcement learning: An introduction. Vol. 1. No. 1. Cambridge: MIT press, 1998.

[3] Munson, Bruce R., et al. Fundamentals of fluid mechanics. John Wiley & Sons, 2014.

[4] Slotine, Jean-Jacques E., and Weiping Li. Applied nonlinear control. Vol. 199. No. 1. Englewood Cliffs, NJ: Prentice hall, 1991.

[5] Cook, Michael V. Flight dynamics principles: a linear systems approach to aircraft stability and control. Butterworth-Heinemann, 2012.