Learning to fly more efficiently in groups[1]

Esteban A. L. Hufstedler, James Riehl, Julien M. Hendrickx, Philippe Chatelain
UCLouvain, 1348 Louvain-la-Neuve, Belgium
{esteban.hufstedler, james.riehl, julien.hendrickx, philippe.chatelain}@uclouvain.be

1 Introduction

Inspired by the impressive capabilities of migratory bird flocks, we investigate how groups of fixed-wing aircraft can efficiently fly long distances in the presence of aerodynamic wake effects. Winged agents in flight generate a pair of vortices containing regions of upward airflow (upwash) and downward airflow (downwash) as shown in Fig. 1. These wake effects offer the opportunity to save energy by flying in the upwash region of a preceding agent. Indeed, there is evidence that migratory bird flocks are able to reduce their energy use by 12% by doing exactly this [1]. While simplified wake models can predict the approximate locations of these regions, complex aerodynamic effects render the task of tracking these locations quite challenging in practice.

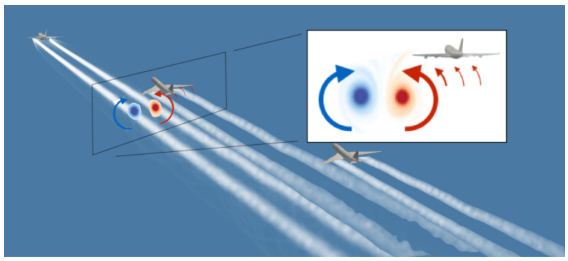


Figure 1: Vortices generated by aircraft in flight

In this work, we explore the use of reinforcement learning (RL) to optimize the efficiency of formation flight. The RL algorithm prescribes relative offset vectors between leading and following aircraft, which are then stabilized using standard formation control methods. We also consider the stability of longer sequences of aircraft flying in this type of formation.

2 Aircraft model, dynamics, low-level control

We use models of fixed-wing aircraft and wake velocity based on steady potential theory at various levels of fidelity. We then apply proportional-derivative control to stabilize the relative positions and velocities between each agent and the agent immediately ahead in the formation, the set points of which are degrees of freedom in the learning approach.

[1]This project has received funding from the European Research Council (ERC) under the European Unions Horizon 2020 research and innovation program (grant agreement no. 725627) and from the French community of Belgium in the Joint Research Activity RevealFlight (convention no. 17/22-080).

3 Reinforcement learning

RL is a type of embodied machine learning that learns to maximize a reward function through interacting with the environment. With the given low-level controller and aircraft dynamics, the RL controller for each agent attempts to learn and choose the spanwise position behind its leader that minimizes its thrust. RL is suitable for this task because it is a model-free technique, and can learn regardless of the precise physical parameters involved (such as flight speed, weight). The spanwise positions over time of two learning aircraft (following Plane 1) are shown in Fig. 2. This shows their exploration of the environment and choice of most rewarding position. Compared to isolated flight, this formation reduced the required thrust by 8%.

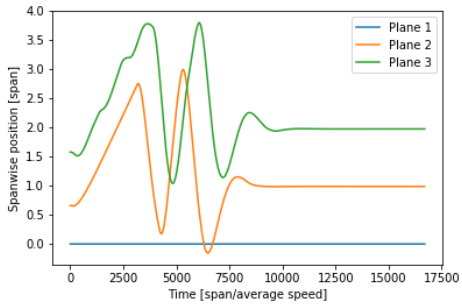


Figure 2: Relative spanwise positions of the aircraft.

4 String stability of aircraft formations

An important result in cascaded multi-agent systems such as vehicle platoons is that even when each leader-follower position error is asymptotically stable, disturbances may amplify from front to back resulting in instability as the length of the formation increases [2]. Indeed, we have observed this phenomenon, called “string instability”, in the aircraft formations. We are investigating known remedies for string instabilities such as communicating the control signals from leader to follower or knowing the position of the absolute leader, and whether these may correspond to observations of real migratory bird flocks. While these methods seem to work in simpler models, wake-induced nonlinearities add complexity to the picture.

References

[1] B. Voelkl, S. J. Portugal, M. Unsöld, J. R. Usherwood, A. M. Wilson, and J. Fritz. ”Matching times of leading and following suggest cooperation through direct reciprocity during V-formation flight in ibis.” *Proc. Natl. Acad. Sci.* 112, no. 7 (2015): 2115-2120.

[2] Swaroop, D., and J. Karl Hedrick. ”String stability of interconnected systems.” *IEEE transactions on automatic control* 41, no. 3 (1996): 349-357.