Learning Agents for Air Traffic Coordination: Key Challenges

Kagan Tumer, Oregon State University
kagan.tumer@oregonstate.edu

Transportation systems are by definition Cyberphysical Systems, because they require a tight connection between the computational algorithms and the physical system to result in safe, reliable and efficient operation. Air traffic management is a particular example of such a system, where problems range from long term optimization (route selection), to scheduling (gate assignments), to human-agent interaction (air traffic controller and automation), to dynamic coordination (re-routing due to weather), to minute by minute control (separation assurance). Though each problem is complex in its own right, their combination leads to a truly congested system, one where over the US airspace alone, 87,000 flights need to be managed daily.

To address this congestion problem, we need to determine how to increase the capacity of our airspace while retaining the stellar safety record of air traffic systems. Indeed, it is critical to note that current regulations and procedures have resulted in air traffic systems to attain an unprecedented safety record. Yet, these regulations are also partly responsible for the congestion and delays that not only cost the industry billions of dollars every year, but get worse every year as the demand for air travel increases. The key observation that makes this problem ripe for a cyberphysical systems solution is that *congestion is at least partially caused by how we use the airspace*, inviting computational tools to increase air space capacity, with little to no physical expansion of the system.

In this position paper, we focus on the use of learning agents in the airspace and determine how such agents can both alleviate congestion and help humans make more informed decisions [4, 13]. But for learning agents to be accepted as part of the national airspace, the following three challenges need to be addressed:

1. The role of computational agents in air traffic;
2. The coordination of a large number of computational agents; and
3. The impact of the highly dynamic (including human decisions) environment on agent performance.

1. Role of Computational Agents in Air Traffic: There are two high level ways in which the roles of learning agents can be viewed in transportation systems. The first, and more 'natural" approach, is to identify system components (aircraft, controller, airport) with computational agents. The agents then learn the same actions that are currently performed by those entities, and use learning methods to develop policies that lead to desirable system behavior. The second, and more abstract approach, is to introduce learning agents in the system to shape the behavior of the "natural" entities in the system. An example of the first approach is to identify each aircraft with an agent [1, 4, 11]. An example of the second approach is to introduce agents that can restrict access to certain parts of the airspace, force ground holds, or set miles in trail for the aircraft.

The air traffic domain is well-suited for the second approach, because the introduction of learning agents does not directly conflict with well-defined human roles. In this respect, such agents are there to help humans make better decisions. But when new agents are introduced into the system, it becomes critical to ensure that those agents' objectives and purpose are selected in a way that does not create unintended conflicts in the agent-agent and agent-human interactions [2, 5, 6]. That topic leads us to the second challenge.

2. Coordination of a large number of computational agents: Ensuring coordination among thousands of agents is a critical research issue that arises when many agents are inserted into the system. Not only do the decisions of these agents need to be understandable by a human operator (pilot, controller), they also need to support system level objectives (increase safety, reduce congestions, reduce delay) [3, 7]. This is a particularly key step when the agents are adaptive (such as reinforcement learning agents). It is imperative to derive objective functions for the agents in such a way that when each agent achieves its own objective, the system behaves in the desired manner [2, 8, 9].

Though there have been many recent results in coordination [3, 11, 13], acceptance of learning in the air traffic domain requires those results to scale to thousands of agents, to handle heterogeneous agents, and to operate when agents have limited information about the state of the system. Finally, this coordination must occur in a highly dynamic environment where the "best" decisions will vary greatly based on the decisions of other agents (both human and computational), which leads us to the third challenge.

3. Operation in highly dynamic environments: Handling highly dynamic environments where the best action not only depends on events occurring outside an agent's purview (weather, equipment problems, human decisions), but also changes rapidly, is the third challenge that arises in air traffic control [3, 10, 12]. The agents' policies and decisions need to be robust both to small variations caused by fluctuations in normal aircraft operation (re-routes, gate assignments) and to large variations caused by unexpected external factors (weather).

Many learning algorithms rely on the stationarity of the environment to provide guarantees of performance that do not apply to multiagent learning. Yet, in many multiagent settings, the environment changes slowly enough that single agent learning algorithms are effective [13]. But in highly dynamic environments, it is critical to explore new methods that can not only ensure the agents are learning the right behaviors, but also provide guarantees that they will do so while remaining within the specified safety bounds (possibly set to ensure understanding by human operators).

Conclusion: Each of these directions is critical to the success of cyberphysical systems providing implementable solutions to the air traffic management problem. But from a broader perspective, they are critical challenges to any transportation system. For example, though the specifics are different for the acceptance of autonomous vehicles in our roadways, and the role of autonomous vehicles more directly overlap with those of drivers, the remaining challenges are similar: those systems must also address the coordination of a large number of computational agents, the operation in highly dynamic environments, and the integration of human drivers and autonomous vehicles.

These critical challenges highlight that what is critical for the adoption of learning agents in transportation systems is not just designing better algorithms. It is to find the right insertion points of the agents into the system, and to make sure we focus on "what" the agents learn, rather than "how" they learn. Only when computational agents provide safer and more efficient transportation systems, can we expect them to be deployed in everyday operation.

References

[1] A. Agogino and K. Tumer. Regulating air traffic flow with coupled agents. In *Proceedings of the Seventh International Joint Conference on Autonomous Agents and Multi-Agent Systems*, Estoril,Portugal, May 2008.

[2] A. K. Agogino and K. Tumer. Analyzing and visualizing multiagent rewards in dynamic and stochastic environments. *Journal of Autonomous Agents and Multi Agent Systems*, 17(2):320–338, 2008.

[3] A. K. Agogino and K. Tumer. Learning indirect actions in complex domains: Action suggestions for air traffic control. *Advances in Complex Systems*, 12:493–512, 2009.

[4] A. K. Agogino and K. Tumer. A multiagent approach to managing air traffic flow. *Journal of Autonomous Agents and Multi-Agent Systems*, 24:1–25, 2012. DOI: 10.1007/s10458-010-9142-5.

[5] M. Colby and K. Tumer. Shaping fitness functions for coevolving cooperative multiagent systems. In *Proceedings of the Eleventh International Joint Conference on Autonomous Agents and Multiagent Systems*, Valencia, Spain, June 2012.

[6] M. Colby and K. Tumer. Multiagent reinforcement learning in a distributed sensor network with indirect feedback. In *Proceedings of the Twelveth International Joint Conference on Autonomous Agents and Multiagent Systems*, Minneapolis, MN, May 2013.

[7] W. Curran, A. Agogino, and K. Tumer. Addressing hard constraints in the air traffic problem through partitioning and difference rewards (extended abstract). In *Proceedings of the Twelveth International Joint Conference on Autonomous Agents and Multiagent Systems*, Minneapolis, MN, May 2013.

[8] A. Iscen and K. Tumer. Decentralized coordination via task decomposition and reward shaping (extended abstract). In *Proceedings of the Twelveth International Joint Conference on Autonomous Agents and Multiagent Systems*, Minneapolis, MN, May 2013.

[9] S. Proper and K. Tumer. Coordinating actions in congestion problems: Impact of top-down and bottom-up utilities. *Autonomous Agents and MultiAgent Systems*, 27(3), 2013.

[10] S. Proper and K. Tumer. Multiagent learning with a noisy global reward signal. In *Proceedings of the Twenty-Seventh AAAI Conference on Artificial Intelligence (AAAI-13)*, Bellevue, WA, July 2013.

[11] K. Tumer and A. Agogino. Distributed agent-based air traffic flow management. In *Proceedings of the Sixth International Joint Conference on Autonomous Agents and Multi-Agent Systems*, pages 330–337, Honolulu, HI, May 2007.

[12] K. Tumer and A. K. Agogino. Multiagent learning for black box system reward functions. *Advances in Complex Systems*, 12:475–492, 2009.

[13] K. Tuyls and K. Tumer. Multiagent learning. In G. Weiss, editor, *Multiagent Systems*. Springer, 2013.