Distributed Agent-Based Air Traffic Flow Management

Kagan Tumer and Adrian Agogino

November 6, 2007

1 Introduction

The efficient, safe and reliable management of our ever increasing air traffic is one of the fundamental challenges facing the aerospace industry today. On a typical day, more than 40,000 commercial flights operate within the US airspace [12]. In order to efficiently and safely route this air traffic, current traffic flow control relies on a centralized, hierarchical routing strategy that performs flow projections ranging from one to six hours. As a consequence, the system is slow to respond to developing weather or airport conditions leading potentially minor local delays to cascade into large regional congestions. In 2005, weather, routing decisions and airport conditions caused 437,667 delays, accounting for 322,272 hours of delays. The total cost of these delays was estimated to exceed three billion dollars by industry [6].

Furthermore, as the traffic flow increases, the current procedures increase the load on the system, the airports, and the air traffic controllers (more aircraft per region) without providing any of them with means to shape the traffic patterns beyond minor reroutes. The Next Generation Air Transportation Systems (NGATS) initiative aims to address this issues and, not only account for a threefold increase in traffic, but also for the increasing heterogeneity of aircraft and decreasing restrictions on flight paths. Unlike many other flow problems where the increasing traffic is to some extent absorbed by improved hardware (e.g., more servers with larger memories and faster CPUs for internet routing) the air traffic domain needs to find mainly algorithmic solutions, as the infrastructure (e.g., number of the airports) will not change significantly to impact the flow problem. There is therefore a strong need to explore new, distributed and adaptive solutions to the air flow control problem.

An adaptive, multi-agent approach is an ideal fit to this naturally distributed problem where the complex interaction among the aircraft, airports and traffic controllers renders a pre-determined centralized solution severely suboptimal at the first deviation from the expected plan. Though a truly distributed and adaptive solution (e.g., free flight where aircraft can choose almost any path) offers the most potential in terms of optimizing flow, it also provides the most radical departure from the current system. As a consequence, a shift to such a system presents tremendous difficulties both in terms of implementation (e.g., scheduling and airport capacity) and political fallout (e.g., impact on air traffic controllers). In this paper, we focus on agent based system that can be implemented readily. In this approach, we assign an agent to a "fix," a specific location in 2D. Because aircraft flight plans consist of a sequence of fixes, this representation allows localized fixes (or agents) to have direct impact on the flow of air traffic[1]. In this approach, the agents' actions are to set the separation that approaching aircraft are required to keep. This simple agent-action pair allows the agents to slow down or speed up local traffic and allows agents to a have significant impact on the overall air traffic flow. Agents learn the most appropriate separation for their location using a reinforcement learning (RL) algorithm [13].

In a reinforcement learning approach, the selection of the agent reward has a large impact on the performance of the system. In this work, we explore four different agent reward functions, and compare them to simulating various changes to the system and selecting the best solution (e.g, equivalent to a Monte-Carlo search). The first explored reward consisted of the system reward. The second reward was a personalized agent reward based on collectives [2, 16, 17]. The last two rewards were personalized rewards based on estimations to lower the computational burden of the reward computation. All three personalized rewards aim to align agent rewards with the system reward and ensure that the rewards remain sensitive to the agents' actions.

2 Air Traffic Flow Management

With over 40,000 flights operating within the United States airspace on an average day, the management of traffic flow is a complex and demanding problem. Not only are there concerns for the efficiency of the system, but also for fairness (e.g., different airlines), adaptability (e.g., developing weather patterns), reliability and safety (e.g., airport management) [3]. In order to address such issues, the management of this traffic flow occurs over four hierarchical levels:

1. Separation assurance (2-30 minute decisions);
2. Regional flow (20 minutes to 2 hours);
3. National flow (1-8 hours); and
4. Dynamic airspace configuration (6 hours to 1 year).

Because of the strict guidelines and safety concerns surrounding aircraft separation, we will not address that control level in this paper. Similarly, because of the business and political impact of dynamic airspace configuration, we will not address the outermost flow control level either. Instead, we will focus on the regional and national flow management problems, restricting our impact to

[1] We discuss how flight plans with few fixes can be handled in more detail in Section 2.

decisions with time horizons between twenty minutes and eight hours. The proposed algorithm will fit between long term planning by the FAA and the very short term decisions by air traffic controllers.

The continental US airspace consists of 20 regional centers (handling 200-300 flights on a given day) and 830 sectors (handling 10-40 flights). The flow control problem has to address the integration of policies across these sectors and centers, account for the complexity of the system (e.g., over 5200 public use airports and 16,000 air traffic controllers) and handle changes to the policies caused by weather patterns. Two of the fundamental problems in addressing the flow problem are: (i) modeling and simulating such a large complex system as the fidelity required to provide reliable results is difficult to achieve; and (ii) establishing the method by which the flow management is evaluated, as directly minimizing the total delay may lead to inequities towards particular regions or commercial entities.

Previous work in this domain fell into one of two distinct categories: The first principles based modeling approaches used by domain experts [3, 7, 9, 11] and the algorithmic approaches explored by the learning and/or agents community [5, 8, 10]. In [3] "geometric optimization" was proposed, where the geometry of a particular air space pattern was utilized to create policies that reduced conflict. In [7] dynamic programming was used over a model of air traffic routes. In [8] a multi-agent approach to the "free flight" problem was proposed, where agents used utilities that balanced their local needs with the system goals. Though our approach comes from the second category, we aim to bridge the gap by using FACET to test our algorithms, a simulator introduced and widely used (i.e., over 40 organizations and 5000 users) by work in the first category [4, 1].

References

[1] 2006 NASA Software of the Year Award. FACET: Future ATM concepts evaluation tool. Case no. ARC-14653-1, 2006.

[2] A. K. Agogino and K. Tumer. Handling communication restrictions and team formation in congestion games. *Journal of Autonomous Agents and Multi Agent Systems*, 13(1):97–115, 2006.

[3] K. D. Bilimoria. A geometric optimization approach to aircraft conflict resolution. In *AIAA Guidance, Navigation, and Control Conference and Exhibit*, Denver, CO, 2000.

[4] K. D. Bilimoria, B. Sridhar, G. B. Chatterji, K. S. Shethand, and S. R. Grabbe. Facet: Future atm concepts evaluation tool. *Air Traffic Control Quarterly*, 9(1), 2001.

[5] M. S. Eby and W. E. Kelly III. Free flight separation assurance using distributed algorithms. In *Proceedings of Aerospace Conference, 1999*, Aspen, CO, 1999.

[6] FAA OPSNET data Jan-Dec 2005. US Department of Transportation website. (http://www.faa.gov/data_statistics/), 2005.

[7] S. Grabbe and B. Sridhar. Central East pacific flight routing. In *AIAA Guidance, Navigation, and Control Conference and Exhibit*, Keystone, CO, 2006.

[8] J. C. Hill, F. R. Johnson, J. K. Archibald, R. L. Frost, and W. C. Stirling. A cooperative multi-agent approach to free flight. In *AAMAS '05: Proceedings of the fourth international joint conference on Autonomous agents and multiagent systems*, pages 1083–1090, New York, NY, USA, 2005. ACM Press.

[9] P. K. Menon, G. D. Sweriduk, and B. Sridhar. Optimal strategies for free flight air traffic conflict resolution. *Journal of Guidance, Control, and Dynamics*, 22(2):202–211, 1999.

[10] M. Pechoucek, D. Sislak, D. Pavlicek, and M. Uller. Autonomous agents for air-traffic deconfliction. In *Proceedings of the Fifth International Joint Conference on Autonomous Agents and Multi-Agent Systems*, Hakodate, Japan, May 2006.

[11] B. Sridhar, G. Chatterji, and S. Grabbe. Benefits of direct-to in national airspace system. In *AIAA Guidance, Navigation, and Control Conference and Exhibit*, Denver, CO, 2000.

[12] B. Sridhar, T. Soni, K. Sheth, and G. B. Chatterji. Aggregate flow model for air-traffic management. *Journal of Guidance, Control, and Dynamics*, 29(4):992–997, 2006.

[13] R. S. Sutton and A. G. Barto. *Reinforcement Learning: An Introduction*. MIT Press, Cambridge, MA, 1998.

[14] C. Tomlin, G. Pappas, and S. Sastry. Conflict resolution for air traffic management: A study in multiagent hybrid systems. *IEEE Transaction on Automatic Control*, 43(4):509–521, 1998.

[15] K. Tumer and D. Wolpert, editors. *Collectives and the Design of Complex Systems*. Springer, New York, 2004.

[16] D. H. Wolpert and K. Tumer. Optimal payoff functions for members of collectives. *Advances in Complex Systems*, 4(2/3):265–279, 2001.