Research Statement of Hamsa Balakrishnan

The air transportation system is a large, complex, global system that transports over 2.1 billion passengers each year. Air traffic delays have become a huge problem for passengers and airlines; they even make the headlines in the popular press and lead to new legislation. Aircraft are also the fastest growing contributor to man-made greenhouse gas emissions. According to the Joint Economic Committee of the US Senate, domestic air traffic delays in 2007 cost airlines over $19 billion and the US economy over $41 billion, wasted 740 million gallons of jet fuel, and released an additional 7.1 billion kilograms of CO_2 into the atmosphere.

My research is in the design, analysis, implementation, and evaluation of *practical algorithms for air transportation systems* that will help air traffic controllers and system operators make better decisions in the face of increasing traffic. This research is important because of the high costs of delays and pollution today, as well as the projected 2x increase in air traffic over the next 15 years.

To prevent large-scale, cascading delays and congestive collapse, and to mitigate pollution, we need new techniques and strategies, perhaps even a radical redesign of certain aspects of the system. My research at MIT seeks to develop these methods.

My research style is to develop new algorithms that are grounded in real-world operational data, implement them, and test them in both simulation and in field trials to gain a fundamental understanding of which techniques work well and why. By analyzing large amounts of weather, flight, and operational data, I have (together with collaborators) developed robust optimization algorithms for scheduling and routing air traffic, accounting for multiple objectives and stakeholders, coping with the intrinsic uncertainty in the operating conditions (e.g., weather), and incorporating environmental objectives into decision-making.

The rest of this statement summarizes my contributions in these areas, and future plans. References to papers are given in parentheses; detailed citations are listed in the complete list of publications.

Research activities and contributions

Terminal-area scheduling for multiple objectives: The efficient operation of airports, and runways in particular, is critical to the throughput of the air transportation system. Scheduling arrivals and departures at runways is challenging because it needs to address the diverse and competing considerations of efficiency, safety, and equity among airlines. A natural approach to runway scheduling that arises from operational and fairness considerations is Constrained Position Shifting (CPS), [Dear and Odoni, 1976], which requires that an aircraft's position in the scheduled sequence not deviate significantly from its position in the first-come-first-served sequence. In collaboration with Bala Chandran and later my PhD student Hanbong Lee, I developed a new family of scalable dynamic programming algorithms for runway scheduling under CPS and other operational constraints [AIAA-GNC 2006; ATM-R&D-Seminar 2007; ACC 2008; *Operations Research* 2010]. The key insight that we had was that even though the space of all feasible solutions is exponential in size, we can represent the solution space as a directed acyclic graph whose size is *linear* in the number of aircraft being scheduled. This insight enabled us to reduce scheduling under CPS to a shortest path problem on that directed acyclic graph. We have developed a prototype implementation, which is fast enough for real-time use. We have also shown how this framework can be extended to many practical problems, including the simultaneous scheduling of arrival and departure operations at an airport (which may involve merging multiple queues) and the optimization of more general cost functions, including fuel burn and robustness metrics [ACC 2007; *Proc. of the IEEE* 2008]. We released our implementation to researchers at the NASA Ames Research Center. It has been integrated with NASA's Stochastic Terminal Area Scheduling Simulation (STASS) for evaluating future terminal-area operating concepts. Our techniques have, for the first time, made scheduling under CPS a practical way to increase terminal-area throughput.

Reducing airport emissions: Aircraft taxiing on the surface contribute significantly to the fuel burn and emissions at airports. My research (with my students Ioannis Simaiakis and Harshad Khadilkar, as well as John Hansman and Tom Reynolds) identifies opportunities to reduce airport emissions, estimates the potential benefits of these strategies, and assesses the implementation barriers that need to be overcome before these approaches can be deployed in the field [*Transp. Res. Record* 2010]. With Ioannis Simaiakis, I have developed and validated a new queuing network model of departure processes at airports, and used this model to develop advanced queue management policies to decrease fuel burn and emissions [AIAA-GNC-2009].

With the help of the Massachusetts Port Authority and the FAA, we conducted a field trial at Boston's Logan International Airport (BOS) to show that a significant portion of these impacts can be reduced through measures to limit airport surface congestion. The central idea is to control the rate at which aircraft push back from their gates during times when the airport is congested. Aircraft therefore spend time at their gates with their engines off, instead of adding to an already congested taxiway system with their engines on.

By designing and testing a simple but novel pushback rate control strategy, we achieved significant reductions in taxi-out delays, fuel burn, and emissions. Results from field tests conducted between August 23 and September 24, 2010, showed that during eight four-hour demonstration periods, more than 15,000 kg of fuel were saved, at the rate of 50-60 kg per gate-held flight. Moreover, these savings were achieved with average gate-hold times of only 4.3 minutes [ATM R&D Seminar 2011 (best paper award)]. The stakeholder feedback and detailed surface surveillance data collected during the trials are being used to investigate the potential for wider adoption of the methods at other airports.

An important aspect of reducing surface inefficiencies is the use of surface surveillance data to evaluate and improve the performance of an airport. We have developed algorithms for processing ASDE-X data from DFW and BOS airports to identify locations where aircraft queues form, the effect of congestion on these queues, and the resultant wait times. At BOS, we have used these analyses to develop tools that provide air traffic controllers feedback on their performance. We provide these results to the BOS Operations Manager in the form of daily operational efficiency reports. In addition, we have been analyzing taxi profiles from Flight Data Recorder information to assess the impact of congestion on fuel burn and emissions (with Aerodyne Research).

From weather forecasts to capacity forecasts: Convective weather (thunderstorms) is responsible for large delays and disruptions in many parts of the world, particularly in the US during the summer. Current algorithms require reliable forecasts of whether or not routes are blocked by weather to schedule and route traffic. In work with PhD student Diana Michalek, I showed how raw convective weather forecasts, which provide deterministic predictions of the Vertically Integrated Liquid, VIL (the moisture content of a region of airspace), can be translated into probabilistic forecasts of whether or not a route into or out of an airport will be blocked.

The problem can be described as follows: meteorologists predict the VIL on a scale of 0-256 around the airport, for each pixel on a 1 km x 1 km grid. Because aircraft are not supposed to fly through a VIL of more that 133, routes must remain on pixels with VIL less than 133. Aircraft headed for the airport (typically when they are 30-60 minutes from the terminal-area, or 60-90 minutes from the airport) need to be told which route they should take inside the terminal-area. If an aircraft is asked to fly a route that ends up being blocked by weather, it will have to be diverted significantly and rescheduled. Predicting whether a route will be open for use is tricky because the VIL forecasts are inaccurate, and arrival and departure routes are much longer than the granularity of VIL forecasts.

Using techniques from machine learning, we developed and validated classification algorithms that predict whether or not a given route is likely to be open in actual weather [AMS-Annual-Meeting 2009, ATM-R&D-Seminar 2009]. Our approach uses historical forecasts and the characteristic features of the route. This is the first algorithm that combines different features of the route to predict the probability of blockage, and provides several insights into the relationship between VIL forecasts and route blockage. Surprisingly, we found that the theoretical capacity (a measure of how many routes into the airport do not pass through forecast weather obstacles) was a poor predictor of route blockage, despite it being a frequently cited metric. The reason is that although the theoretical capacity is a prediction of how many routes will be open, it does not give any indication of which ones they would be. Knowing just the number of open routes is not enough for terminal-area route planning.

We have also used our forecasts of route blockage to dynamically modify routes to optimize the expected capacity of the terminal-area. Experiments using real weather scenarios show that our algorithms recommend that a terminal-area route be modified 30% of the time, opening up 11% of available routes that would have otherwise been closed. We also found that 97% of routes predicted by our method as being “open with probability greater than 95%” are in fact open in the weather that actually materializes [CDC 2010, *Transportation Science* 2011].

Researchers at MIT Lincoln Laboratory have adopted our methods to develop new metrics for evaluating en route aviation convective weather forecasts. A key implication of our work is that improved metrics that assess the skill in predicting route blockage are a better measure of the ability to support air traffic management than traditional metrics, which focus on the accuracy of predicting convective activity in a 1km x 1km pixel. We are currently developing an implementation of these techniques as part of a prototype tool for arrival/departure gate management in the Chicago-area terminal airspace.

Distributed feedback control of the National Airspace System: Today’s airspace is partitioned into sectors, and each air traffic controller is responsible for managing traffic within his or her sector. Controllers communicate just locally with their neighboring sectors; the control of flows between sectors is done by negotiating handoffs in an ad hoc manner. Air traffic controllers do not really prioritize flows. Because of this lack of prioritization, local weather disruptions (say, in the New York area) can lead to holding patterns in the mid-West, affecting all flows through the mid-West, even those not bound for New York.

With Jerome Le Ny, I developed a queuing network representation of traffic flows in the NAS that accurately models current operations. We developed distributed feedback control techniques that guarantee that the aircraft queues in each airspace sector, which are an indicator of controller workload, are kept small. We showed that under realistic conditions, our feedback control policy for scheduling and routing aircraft stabilizes the system (i.e., all queues remain bounded). Our approach provides the first distributed feedback control strategy for realistic, multi-airport settings. We also showed that our model generalizes existing flow-based models, and that our methods can mitigate the impact of weather disruptions [ACC 2009, CDC 2010, *AIAA Journal of Guidance, Control and Dynamics* 2011].

Future Plans

In the long term, I am interested in the development of algorithmic approaches that enable robust and sustainable infrastructure systems. In the near term, we are currently discussing the next steps of our airport surface congestion management research with the FAA, including follow-on field tests at BOS in Summer 2011, as well as other possible sites for similar efforts. On the topic of reducing surface emissions, I believe that we are hindered by the absence of a reliable model of aircraft fuel burn on the airport surface. I am therefore interested in developing an accurate model that can translate surface trajectories into estimates of fuel burn and emissions. I would like to use this model to develop airport optimization algorithms that balance the tradeoffs between multiple objectives. The airport environment involves complex interactions among its different components; I therefore also intend to investigate the integration of scheduling algorithms for different airport elements, such as the runways, taxiways, and ramps. The development of architectures for such schedule integration will be essential for enabling efficient surface operations. On the topic of weather forecasts, I am interested in developing classification algorithms that can predict whether pilots will deviate from a particular region of airspace or fly though it, and using them to optimally schedule traffic. In addition, I am looking forward to the prototyping and testing of our approaches in the Chicago area airspace.

In summary, I intend to continue using operational data to design, implement, and evaluate practical algorithms for a more efficient, robust, and environment-friendly air traffic system.