Executive Summary

1. Background: United Airlines (UA) operates an average of 5,446 flights a day to more than 370 airports across six continents. In 2012, United Airlines carried more passenger traffic than any other airline in the world and operated nearly two million flights carrying 140 million customers. UA is currently serving 234 airports in the United States with 8 major hubs located throughout the continental United States.

2. Purpose: The purpose of our model is to find the most profitable plane-type with every plane starting at a hub. We then examine the effect of a variety of disruptions to the network and measure the effects. To do this, we create a minimum-cost flow model which is driven by the profit available along the edges of the graph. In addition, we track the plane type moving along each edge. To model our network, we look at a 24-hour period from June 2012 of all flights occurring in mid-to-large airports in the United States by UA.

3. Assumptions:

- Only allowing network to run for two hops
- Network starts with all aircraft at a hub
- All edges between airports are fair game to an aircraft at that airport; network does not force match a particular type of aircraft for a given edge
- UA is incentivized to use certain planes for certain edges based on profit available
- Any edge can be run multiple times, but multiple runs of the same edge are successively less profitable
- Network is time-agnostic
- Not requiring aircraft to return to point of origin
- The network does not include the largest or smallest planes UA uses
- Aircraft are allocated to each hub proportionally to the number of edges coming out of the hub

4. Research questions:

- What percentage of flights are still operating when you remove a hub or multiple hubs?
- What percentage of flights are still operating when you remove a plane-type? What capacity does UA have if you remove a plane-type?
- How many planes should we allocate to each hub?

5. Model: In our simplified version of the network, we have reduced the size but tried to maintain the proportions. For example, the number of airports are reduced from 234 to 63 of the busiest airports in the United States. Additionally, UA normally operates 14 different types of planes, we have reduced that to the five plane-types that constitute the majority of its fleet.

We add a super-start that connects to each airport where a flight begins. Each subsequent node in the model will be hubs, plane-type, possible flight routes, legs of flight routes and costs (see figure). An edge exists if it is possible for a plane to travel the path from Start to End. Each subsequent instance that a leg is used incurs decreasing amounts of revenue, thus penalizing the model for running the same leg multiple times.

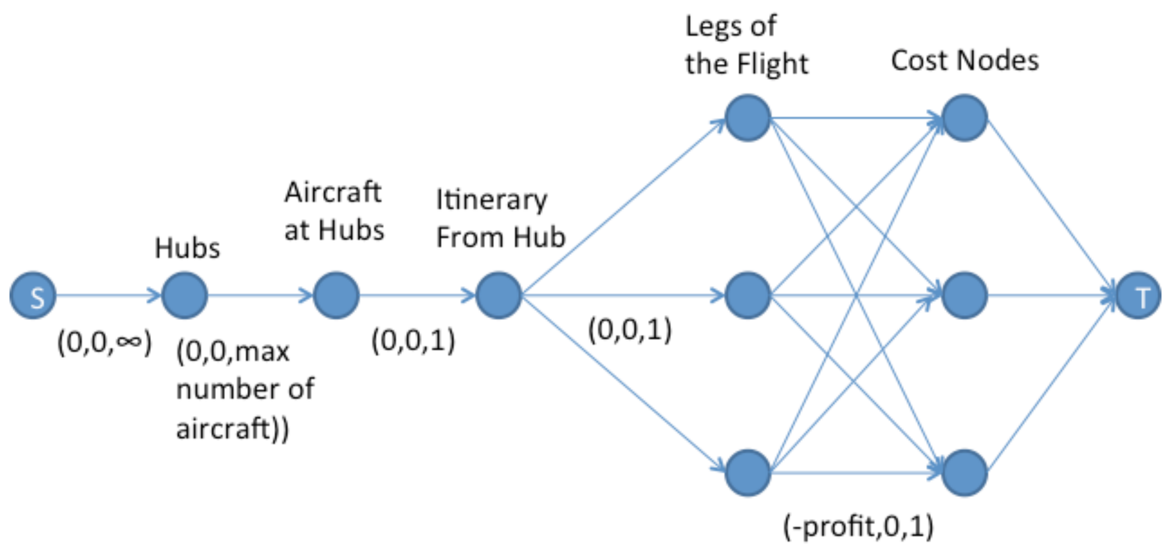

For example, suppose a Boeing 737 starts in San Francisco (SFO), travels to Los Angeles (LAX), then flies back to San Francisco, and further suppose that it was the first plane to travel that leg. Then its path on the network would be: Start -> SFO -> SFO-B737 -> SFO-LAX-SFO -> SFO-LAX and LAX-SFO -> Full Profit -> End. Of note, the Flight node splits the path so that it fulfills both legs—SFO-LAX and LAX-SFO. This is because the same plane is going to be flying each leg. Compare another example with the same plane-type but starting in Los Angeles going to San Francisco and back. The path would be: Start -> LAX -> LAX-B737 -> LAX-SFO-LAX -> LAX-SFO and SFO-LAX -> Less Profit -> End. Notice, this is the second time the legs of LAX-SFO and SFO-LAX is utilized, thus each was sent to the Less Profit node.

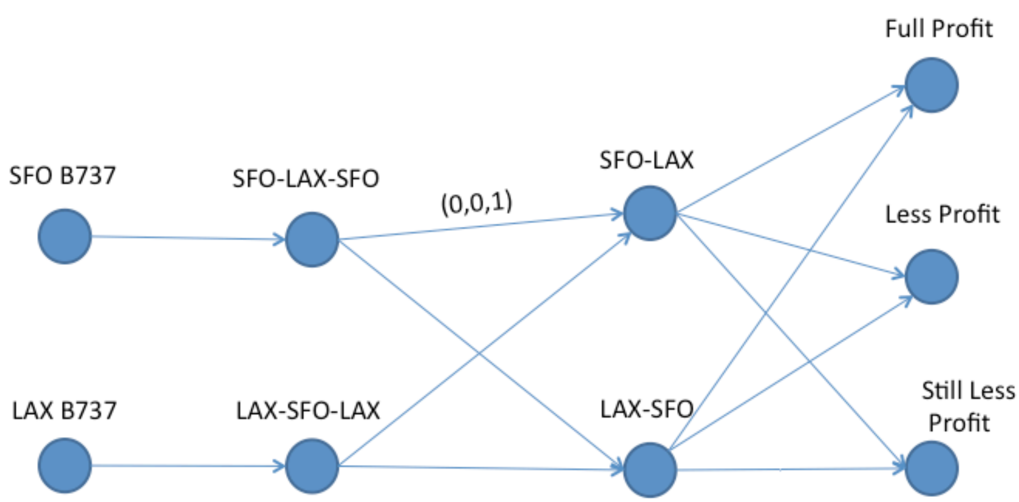

The profit available along the edges is specific to the type of plane. Each plane type in the model has a specific number of seats, is capable of a specific cruising speed, and has a specific cost associated with its operation (variable cost/hour). The cost of running a particular type of plane on a particular edge was determined using these factors, ending in a variable cost/seat for that plane on that edge as shown in Equation 1. A 150% profit was then subtracted from this number in order to drive the flow of the model.

$$\left(\frac{\frac{VariableCost(j)}{hour}}{\frac{PlaneMiles(j)}{hour}}\right)\left(\frac{PlaneMiles(j)}{Seats(j)}\right)=\frac{VariableCost(j)}{Seat(j)}$$

Equation (1)

6. Results: Removing the hubs provides insight into how the network is connected. Removing hubs on either the East or West coasts (SFO, LAX, EWR, CLE, or IAD) does not have as great of an effect as removing the hubs located in the Midwest (IAH, DEN, ORD). This makes sense when looking at how well the hubs in the Midwest are connect to other airports. The three in the Midwest are the most highly connected out of the eight hubs, therefore, you would expect to see more effect on flights-operating by removing them.

Eliminating plane-type also has an intuitive response. If you were to remove the Airbus A319 or the Boeing 767 or 777, then the flight capacity would decrease by 10 percent. However, removing 757 or 737 reduces capacity by 30% and 45%, respectively. This is because the fleet mostly consists of these planes.

To analyze the number of planes that should be allocated to each hub, we switch the nodes Hub and Plane-Type to get the network to "pull" or utilize the planes that it finds most profitable. The analysis shows that the network utilizes all the 757's and 767's and underutilizes everything else. This suggests that our profit function makes the utilization of 757 and 767-type aircraft more appealing to the network than the other types of planes; there is more profit available for them.

7. Conclusion. The model is working as expected, with one exception. There is something limiting the flow into each of the individual cost nodes to 100 flight legs. Aside from that, the model behaves as predicted, and we believe that it is a viable starting point for a deeper analysis of the network. The next avenue in which to improve the model would be to allocate only certain types of planes to certain edges, given certain constraints. This would bring the model more in line with the expectations of real-world demand; i.e., as it sits the model would allow a B777 to fly from the 63rd-busiest airport to the 62nd-busiest airport which was only 100 miles away. The profit function may or may not show that as a profitable venture for the B777 were it in that situation, but in the real world that would probably not be in line with demand. Another interesting challenge would be to drive the model by time rather than by number of legs. The information with which to do this is available, but it would be significantly more intricate network to model.