Analysis of the Potential for Delay Propagation in Passenger Aviation Flight Networks

Shervin AhmadBeygi, Amy Cohn, and Yihan Guan
University of Michigan
Peter Belobaba,
Massachusetts Institute of Technology

April 1, 2007

Abstract

In this paper, we analyze the potential for the propagation of delays in passenger aviation flight networks. The motivation for this research is to better understand the relationship between the scheduling of aircraft and crew members, and the operational performance of such schedules. In particular, when carriers decide how to schedule these scarce and costly resources, the focus is largely on how to achieve high levels of utilization. The resulting plans, however, often have little slack, limiting the network's ability to absorb disruption; instead, initial flight delays may propagate to delay subsequent flights as well. Understanding the relationship between network plans and delay propagation is a critical precursor to the development of tools for building more robust airline plans. In this paper, we investigate this relationship using flight data provided by two major U.S. carriers, one traditional hub-and-spoke and one "low-fare" carrier.

1 Introduction

In this paper, we present an empirical study to investigate the relationship between aircraft/crew schedules and the potential for delay propagation in passenger aviation flight networks.

Passenger airlines are an essential part of today's society, providing both business and leisure travelers easy access to locations that could not otherwise be reached. In the U.S. alone, hundreds of millions of travelers fly each year. In order to provide this service, airlines face several challenges. One key challenge is the need to efficiently utilize costly resources: For example, the cost of a single aircraft is in the hundreds of millions of dollars, and, in the U.S., senior pilots at major airlines are among the highest-paid workers in the country ([4]).

The operations research (*OR*) community has played a significant role in developing airline planning tools with the aim of utilizing these costly resources efficiently ([1], [2]). In these tools, which typically assume deterministic flight times and other parameters, it is desirable to have little if any slack between flights – by turning crews and aircraft quickly, greater utilization of these resources can be achieved.

In practice, however, flight times are *not* deterministic. Departure delays arise due to mechanical problems, weather delays, ground holds, and other sources. Flights that depart on time can still be delayed in arrival due to causes such as air traffic control issues or re-routings to avoid inclement weather. In isolation, these delays are themselves costly. In a network structure, they can have even greater impact – without adequate slack to absorb the root delay, subsequent flights may also be delayed as they await aircraft and crews from the initially delayed flight. We refer to this as *delay propagation*.

It seems logical that there is an inverse relationship between the efficiency of a network plan and its operational robustness. In a network with high resource utilization, there is limited slack. This in turn limits the opportunity to absorb flight delays, which must instead propagate to subsequent flights. But what is the nature of this relationship? How can it be incorporated in the planning process to produce "better" schedules? And what constitutes "better" – how should the deterministic costs of an airline plan be traded off against the potential costs of delay? These are all challenging questions that are beginning to receive significant attention from both the airline industry and the academic community.

To assist in these efforts, we have undertaken an empirical study of passenger airline flight networks and their potential for delay propagation as a function of their network structure. This study is based on flight data from two U.S. carriers, one traditional "legacy" hub-and-spoke carrier and one niche-based "low-fare" carrier. Using the idea of *propagation trees* as our foundation, we examine the maximum extent that any one given flight delay can propagate through the network (the structure is a tree because one flight uses multiple resources, such as cockpit crews and aircraft, and therefore each flight delay can directly cause multiple subsequent delays, which in turn can continue to branch further). We analyze these trees across the network to gain insight into the distribution of slack throughout the system and the implications of this on delay propagation. We then use this analysis to address commonly-held assumptions about how delays propagate, provide insights into the relationship between network efficiency and operational robustness, and raise questions for further study.

The contributions of this research are in: identifying metrics for helping to assess the potential operational performance of a network plan; providing quantitative analysis of network plans to gain insights as to the relationship between slack and delay propagation; substantiating (in some cases) and disproving (in other cases) commonly-held assumptions about delay propagation; and laying the groundwork for future research on incorporating measures of potential delay propagation in the network planning process.

The paper is organized as follows. In section 2, we outline the details of the study. We present our analysis in section 3. Section 4 contains our conclusions and suggested areas for future research.

2 Analytical Framework

2.1 Motivation

Consider a (hypothetical) airline plan with resource utilization that is maximized in the sense that, for every crew and for every aircraft, their connection times between every pair of consecutive flights are the minimum permitted turn times. Such a schedule is ideal from the perspective that aircraft are being fully utilized and crews are not being paid for excess "sit time" between flights.

Suppose further that this schedule is implemented, and that an arbitrary flight is delayed in departure by thirty minutes (for example, due to a mechanical problem that must be fixed before take-off). Assuming that this delay is not compensated for by increasing the travel speed, the flight will have a thirty minute arrival delay as well. Because the cockpit crew has a tight turn (i.e. the minimum sit time), this thirty minute delay will propagate to their next flight, causing it to also be delayed in take-off by thirty minutes. If the crew and the aircraft do not stay together, then the aircraft's turn will lead to a second thirty-minute flight delay. The cabin crew could cause a third thirty-minute flight delay if they separate from the aircraft and cockpit crew. Likewise, if flights are held for connecting passengers from this flight, then these flights will be delayed as well.

Now consider this set of flights that have been delayed as a result of the initial flight delay. [We will refer to the initial flight delay as the the *root delay*; it is also sometimes referred to as the *independent delay* in the literature. ([3])] These flights will also arrive late at their destinations, resulting in a second layer of subsequent flight delays. In fact, a delayed aircraft will continue to propagate delay to all of its subsequent flights until the aircraft goes off-rotation into maintenance or enters an overnight phase where there are no longer flights to be covered. Similarly, a crew (cockpit or cabin) will propagate delay to all of its subsequent flights until they go off duty (or, if the overnight rest period between duties is tight, until their pairing ends). If the original crews and aircraft do not stay together for subsequent flights, then each of these resources will ultimately cause other resources (i.e. other crews and aircraft) to enter this stream of delays as well.

The situation we present here considers two improbable extremes – on one hand, a perfectly-efficient schedule and on the other hand a delay cycle that propagates indefinitely. In reality, other factors prevent schedules from being "perfectly efficient" (for example, the market demands that influence flight times and frequencies). Furthermore, recovery alternatives (canceling flights, calling in reserve crews, etc.) frequently prevent delays from propagating fully. Nonetheless, taking into account both objectives – maximizing the profitability of a schedule under ideal conditions; minimizing the propagation of delays in operation – presents an important challenge for airline planners.

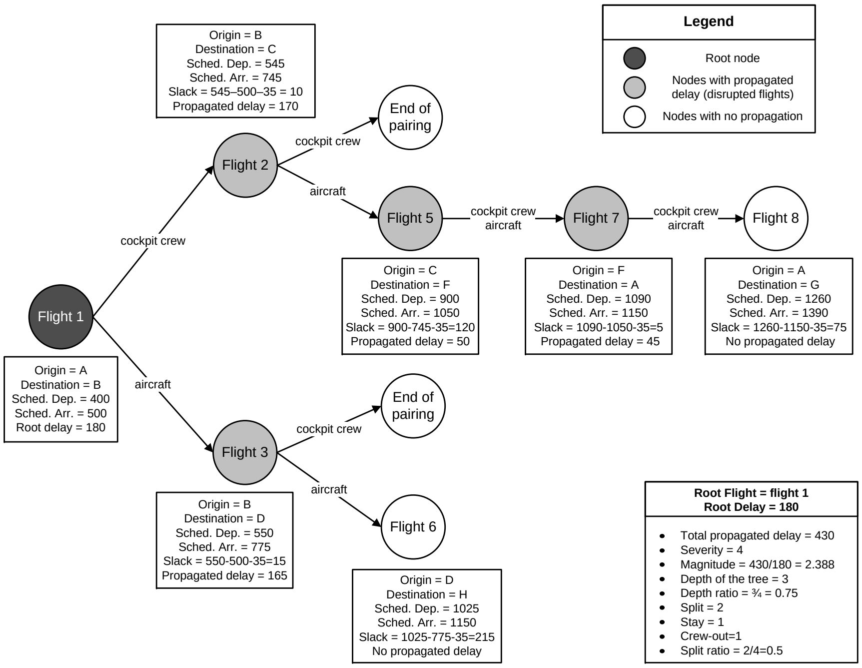

Figure 1: An example of a propagation tree.

It is also a difficult challenge, not only because the planning tools are themselves so difficult and because the real-world problem is highly stochastic, but also because means for quantifying robustness (and the value of this robustness) are not well-defined. In this study, we by no means claim to provide a definitive answer to these daunting questions. Instead, we take an important first step in developing metrics and tools for understanding aviation network structures and their role in propagating flight delays.

2.2 Propagation Trees

A propagation tree is a simple but powerful visual and quantitative device that enables us to better understand how a root delay can, in the absence of other disruptions or schedule modifications, propagate through the network. Figure 1 provides an example of such a tree.

In this example flight 1 is the root flight and is delayed by 180 minutes. Note that this an "independent" delay which can be caused by a mechanical problem, weather conditions, etc. After landing, the cockpit crew of flight 1 is assigned to flight 2 and the aircraft is turned to flight 3. This is shown in Figure 1 by two arcs coming out of flight 1. Flight 1 is scheduled to arrive at 500 (the time unit is *minutes after midnight*); however, as a consequence of the 180 minute delay, the actual arrival time is now 680. Considering a minimum turn time of 35 minutes, the cockpit crew will be ready for flight 2 at 500+180+35=715. This will translate to a 170 minute delay for flight 2; in other words, the actual departure and arrival times of flight two are 715 and 915 respectively. After flight 2, the cockpit crew ends their pairing, but the aircraft which is assigned to flight 2 turns to flight 5. Assuming a minimum turn time of 35 minutes, the aircraft will be ready at 915+35 = 950. This means that flight 5 will be delayed by 50 minutes because of the delay in the upstream flight. Both the crew and the aircraft of flight 5 are also assigned to flight 7 which is scheduled to depart at 1090. Assuming that flight 5 is scheduled to arrive at 1050 and considering a minimum required turn time of 35 minutes, the slack time between flights 5 and 7 is only 1090-1050-35 = 5 minutes. This slack will dampen the 50 minute delay of flight 5 by 5 minutes. Therefore flight 7 is delayed by 45 minutes. This means that the actual landing time for flight 7 will be 1195. Considering 35 minutes for turn time, the crew and the aircraft will be ready for flight 8 at 1230. Flight 8 is scheduled to depart at 1260, therefore the 45 minute delay does not propagate from flight 7 to flight 8.

Similarly, consider the aircraft which was assigned to flight 1. This aircraft is turned to flight 3, but as a result of the 180 minute root delay, the aircraft will not be ready until 715. This means that flight 3 will be delayed by 715-550=165 minutes. After flight 3 lands, the cockpit crew ends its pairing and the aircraft turns to flight 6. However, between flights 6 and 3 there is slack of 1025-775-35=215. Therefore the delay does not propagate from flight 3 to flight 6.

As we can see, the 180 minute root delay in flight 1 causes 430 additional delay minutes in the downstream flights, disrupting a total of 4 additional flights.

In this propagation tree, we observe two different paths: $1 \rightarrow 2 \rightarrow 5 \rightarrow 7$ and $1 \rightarrow 3$. Therefore the longest path ($1 \rightarrow 2 \rightarrow 5 \rightarrow 7$) includes 3 flights beyond the original delayed flight.

Observe that in this tree, we only consider the impact of delay on two resources – the aircraft and the cockpit crews. We do this for a number of reasons. The first is data availability – these are the two resources provided by the supporting carriers in our preliminary study. The second is importance – these are the two most costly resources, with cabin crews and passenger delays having less significant impact. The third is complexity – we felt that initially limiting our study to only two resources would enable us to develop an understanding of some of the critical drivers of delay propagation without being overwhelmed by the network interactions. We hope that the results of this first study will facilitate a second study that incorporates additional network resources (particularly, cabin crews and connecting passengers) as well.

In this process of constructing propagation trees, we track a number of metrics, defined here:

- **Total propagated delay**:The sum of the delays (in minutes) imposed on downstream flights in a propagation tree; note that the root delay is not included. In figure 1, the total propagated delay is 430 minutes.
- **Magnitude**: The ratio of total propagated delay to root delay. In figure 1, the magnitude is 2.388.
- **Severity**: The total number of disrupted flights, excluding the root flight. In this analysis, we consider a flight as disrupted only if the value of the propagated delay is greater than or equal to 5 minutes. In figure 1, the severity is 4.
- **Depth**: The number of nodes corresponding to the longest path in a propagation tree (not counting the root delay). In figure 1, the depth is 3.
- **Depth ratio**: The ratio of depth to severity in a propagation tree. In figure 1, the depth ratio is 0.75.
- **Stay**: The total number of nodes (disrupted flights) in which both the crew and the aircraft are the same as in the preceding node. In figure 1, the stay is 1 (flight 7).
- **Crew-out**: The total number of nodes (disrupted flights) in which the crew is not the same as the preceding node, because the crew in the preceding node has ended their pairing. In figure 1, the crew-out is 1 (flight 5).
- **Split**: The total number of nodes (disrupted flights) in which either the crew or the aircraft is not the same as the preceding flight, because these resources split to serve two different subsequent flights. In figure 1, the split is 2 (flights 2 and 3).
- **Split ratio**: The ratio of split to severity in a propagation tree. In figure 1, the split ratio is 0.5.

Through the use of these metrics, we are able to quantitatively evaluate the relationship between an individual root delay and the rest of the network. This provides insights into the relationship between airline schedules and their robustness. In particular, we use these metrics to address "conventional wisdom" from the airline and academic communities about these relationships. For example, we consider the following commonly-held beliefs:

- Propagated delays create significantly more impact than the original root delays themselves.
- A single delay can "snowball" through the entire network.
- Keeping aircraft and crews together can help to mitigate the impact of disruption.

- Delays that occur early in the day cause greater propagation than delays later in the day.
- It is most important to prevent delay propagation early in the day (in other words, slack should be more pronounced in the early parts of the schedule).

2.3 Study Parameters

In our analysis, we consider the impact of root delays for each flight on the rest of the network. Furthermore, we evaluate the impact of the length of the root delay, considering values ranging from 15 minutes to 180 minutes, in 15 minute intervals, thus providing 12 trees for each root flight. [Beyond 180 minutes, it is highly likely that recovery actions, such as a flight cancelation or a crew swap, would be imposed, rather than allowing the delay to propagate. Therefore, we do not extend our analysis beyond this point.] In the first data set, there are 1719 flights; the second data set contains 410 flights. We assume a 35-minute turn time. We also consider a minimum rest of 570 minutes (9.5 hours) for cockpit crews between two consecutive duty periods. In addition to computing the individual metrics (as defined in section 2.3) for each of these flight/delay pairs, we also compile aggregate statistics, looking at flights not only individually, but grouped by origin, and by time of day. The sub-routine we use in order to generate the propagation trees is shown in Table 1.

```
for each value of root delay
for each flight
{
    initialize the propagation tree
    initialize the list of nodes
    create the root node and add it to the list
    while there is a node in the list
    {
        calculate the slack between the flight in the current node and the succeeding flights
        if there is not enough slack
        {
            create a new node and add it to the list
            update the propagation tree statistics
        }
        delete the current node
    }
}
```

Table 1: the sub-routine for constructing a tree

We analyze these results and discuss their implications in section 3. Before doing so, we conclude this section by noting limitations of our initial study. First, we only take into account aircraft and cockpit crews. Other resources (in particular, cabin crews and connecting passengers) can cause additional delay propagation. Second, we do not take into account interactions between delays. In reality, there is often correlation between delays (in particular, due to weather conditions) and thus propagation trees will impact one another. Third, we do not consider recovery options (canceling flights, calling in reserve crews, etc.). One of the challenges of doing so is that the means for making these decisions are rarely codified by the airlines, but instead conducted in an ad hoc manner by individual operations personnel, based on their experience and intuition. Fourth, we do not weight the probability of root delays. In our aggregate data, all root delays are treated equally. Finally, our data sets are restricted to only two carriers, and one specific day in each carrier's schedule. As a result, the analysis is by no means intended to make universal statements but rather to gain preliminary insights and to develop metrics and tools for analysis. We hope that this initial research will motivate future investigations taking a more complex network viewpoint.

3 Empirical Analysis

PENDING APPROVAL FROM CARRIERS TO RELEASE RESULTS

4 Conclusions and Future Research

In this paper, we have presented an empirical study investigating the relationship between aircraft/crew schedules and the potential for delay propagation in passenger aviation flight networks. Although this is by no means intended as a definitive and exhaustive study, it nonetheless provides a starting point both to increase understanding of how delay propagates in a flight network and also to motivate future research in several directions.

First, we see value in extending the complexity of this analysis – taking into account the probabilities of the occurrence of different root delays, recognizing correlations (for example, weather-based) between groups of root delays, adding in the propagation due to cabin crews and connecting passengers, and incorporating recovery decisions, such as crew swaps and flight cancelations.

Second, our research can be useful in helping to identify mechanisms for strategically using slack in the system to mitigate the impact of disruption, by recognizing where this slack can provide the greatest benefit.

Third, there is important work left to be done in quantifying the *value* of increased robustness – what is the trade-off between improved robustness (i.e. decreasing the likelihood of delay propagation) and scheduled costs (i.e. the cost of a plan if it operates without disruption)?

Finally, once the value of robustness can be quantified, it is possible to begin incorporating metrics of robustness within the planning process itself.

References

[1] C. Barnhart, A. Cohn, E. Johnson, D. Klabjan, G. Nemhauser, and P. Vance (2003). “Airline Crew Scheduling.” *Handbook of Transportation Science*, second edition, Kluwer’s International Series, 2003.

[2] Barnhart, C., K. Talluri. (1997) “Airline operations research.“Arthur McGarity, Charles ReVelle, eds. **Design and Operation of Civil and Environmental Engineering Systems**. John Wiley and Sons, New York, 435–469.

[3] Lan, S., Clarke, J., and Barnhart, C. (2003) “Robust aircraft routing”, Technical report, Massachusetts Institute of Technology.

[4] http://www.bls.gov/oco/cg/cgs016.htm.