Ripple delay and its mitigation[1]

James A. Rome
Simon D. Rose
Ronald W. Lee
Oak Ridge National Laboratory, Oak Ridge, TN

James H. Cistone
Lockheed Martin Air Traffic Management, Rockville, MD

George F. Bell
GFB and Associates, Fayetteville, GA

William S.Leber
Northwest Airlines, St. Paul, MN

Abstract

Arrival delays, especially those that occur early in the day, tend to exponentiate with time unless the crew and equipment remain together throughout the day's schedule. Consequently, to achieve optimum results, the airline must account for these downline ripple delay effects when managing its flights. In this study, which uses data from Northwest Airlines, we attempt to minimize the costs of such delays by allowing the airline to swap its landing slots within a ±10 minute window. The study has three parts: (1) calculation of the incremental cost (including affected later flights) as each flight's landing time is moved within the ±10-minute window; (2) determining a near optimum resequencing algorithm for making the swaps and applying it; and (3) extrapolating the results to the industry as a whole. Annual industry-wide savings for airlines at major hubs are over $75 million if airlines are restricted to using only their own slots, and over $100 million if other (available) slots can be used.

---

[1] This study was performed for NASA Ames Research Center by Lockheed Martin Management and Data Systems under Contract No. NAS2-98004 as part of the TO-13 *Benefits Assessment of Airline-Influenced Arrival Scheduling*.

BACKGROUND

Most airlines that use hub-and-spoke operations rely upon a bank of planes arriving at the hub in as short a period as possible (less than an hour). The passengers and crews transfer to their next flight leg and again, all the planes take off and leave the hub as simultaneously as is possible. This sort of operation makes the best use of equipment, gates, and crews so long as everything is working well. However, when a significant number of flights are late, connections will be missed, passenger misconnection costs and ill will escalate, and ripple delays start to build. In a previous study with American Airlines [1], we found that a minute of delay early in the day could cause up to 13 minutes of delay later in the day. This effect was called *ripple delay* in that study.

Ripple delay is caused by the scheduling practices of the airlines. Turn times between flight legs are minimized to fully utilize expensive resources, so there is not much slack in the schedule to absorb delays. In addition, crews are commonly split between flight legs to allow them to live in different places. In addition to the active crew on a flight, deadhead crew (crew being transported to or from their active flights) must also be taken into account if they are assigned to later flights. Because of the crew and equipment splits, one late flight can delay a handful of other flights on the next leg. It is this splitting that causes delays to get worse as the day progresses.

While minutes of delay give an indication of problems, not all minutes are created equal—it is the dollar cost of a delay that ultimately matters to the airline. For Northwest Airlines (NWA), most costs of delay are associated with passenger misconnect costs and ill will, so minimizing the airline costs is also best for the traveling public. Furthermore, the costs of other flights affected by the ripple delay must also be taken into account in determining how to handle the flights in a landing queue; larger cost contributions could arise from later flights.

During both normal and constrained operations, after an aircraft is airborne the estimated time of arrival (ETA) can be accurately computed and adjusted throughout the flight. Although the Federal Aviation Administration (FAA) controls air traffic and assigns landing slots at major airports, it has no knowledge of the economic value of each flight to its customers, the airlines. Once the runway arrival time or landing slots have been computed, it is believed that airlines could further refine and improve their operating efficiency by having the ability to establish the actual landing sequence for their flights at a given airport. In addition to direct cost savings, resequencing could also allow the airline to:

- Control the order in which flights arrive into the terminal complexes thereby reducing taxi time and ramp congestion and improving gate utilization,
- Prioritize arrivals to permit smoother passenger and cargo/mail connections: reducing miss-connections and disrupted passenger expenses, reducing departure delays and the downline effects of these delays and associated disruptions,
- Prioritize arrival for aircraft and flight crew rotations based on next assignment
- Avoid diversion in situations where longer-haul or load-limited flights may not have sufficient fuel to absorb airborne holding delays, by allowing these flights to have landing priority while other company flights, with higher fuel reserves, absorb additional holding delay, and
- Maintain bank integrity. Hub-and-spoke systems work smoothly when all of the flights arrive and depart in groups. While the planes are at the gate, the passengers make their connections. Resequencing tends to reform the bank by delaying early flights and speeding up late ones.

To calculate and mitigate the costs from ripple delays, the first step is to model the airline connectivity. This means integrating the crew, equipment, passenger and flight schedules into a multiply linked list that can be traversed forwards and backwards. The airline cost model and the requirements for equipment and crew turn times must be integrated into this data structure in order to calculate the “ready” time for each flight leg, and the cost associated with moving it forward and backward in time.

Next, the flights must be associated with the landing queue for a bank at a hub airport. If these flights are allowed to swap landing slots within the ±10 minute window, the ripple-delay cost (or savings) associated with each swap must be calculated. Given these costs, a method for selecting the best (or a very good) set of swaps must be devised. Once the swaps are made, the resulting savings can be calculated, and extrapolations can be made for the airline and the industry.

Previous studies [2, 3, 4] of arrival queue resequencing have considered a different problem, namely how to increase the number of landings at a given airport. When cost was introduced in the optimization, only fuel, crew, and aircraft costs per minute were considered; we find these costs to be insignificant compared to passenger costs, especially when the ripple delay effects are included.

The National Aeronautics and Space Administration (NASA) Center Tracon Automation System (CTAS) was designed to sequence an arrival queue so that an airline could specify a preferential arrival order using only their landing slots [5]. This study was undertaken as part of the CTAS effort to quantify the cost benefits that could be derived by an airline from such resequencing. It was also recognized that the downline effects needed to be accounted for in order to make the best use of this resequencing capability.

As is usual, this is a “chicken and egg” problem. The FAA controllers do not ordinarily allow swaps in the air. Planes are often sequenced from takeoff to landing. Such actions require extra work, and there could be safety issues that occur during the process. We have also ignored the problem of aircraft size in the landing sequence separation as was discussed in [4]. However, if the cost savings that can be achieved by these swaps can be shown to be significant, it might motivate the system and in particular Air Traffic Service Providers (ATSPs) to devise safe ways to accommodate these requests. In the meantime, an airline has some control over the landing queue by changing pushback times and air speeds, and occasional requests for queue resequencing can be honored.

Although Northwest Airlines participated in this work, it was not performed for Northwest Airlines and hence the airline has not had any reason to independently verify the results.

DATA AND COST MODELS

This study utilized NWA data from two periods, April 6–8, 1998, and July 30–August 5, 1999. NWA provided us the following data:

- Flight leg data for NWA and its two commuter airlines, Mesaba and Express. The commuters are collectively referred to as Airlink.
- Flight delay causes.
- Passenger ticket data for the three hub airports (incoming and outgoing).
- Passenger connection times.
- Crew patterns for pilots and flight attendants.
- Aircraft schedule and minimum turn time data.
- Flight leg load factors.
- Aircraft tail number, type and model information and seat capacity.
- Aircraft performance data.
- Cost model for passengers, crew, equipment, and fuel.

In addition, we obtained Enhanced Traffic Management System (ETMS) data for the period from Bruce Ware of the FAA's Airspace Redesign Office, and airspace performance data from Carlton Wine's FAA Consolidated Operations and Delay Analysis Systems (CODAS) Web site.

One of the challenges to real-time implementation of resequencing the landing queue is the availability of all the data required to calculate the ripple delay costs. In particular, at the time of our study, NWA had difficulty in assembling the passenger data. In order to get the study finished before our deadlines, we used passenger data from a similar period a month earlier. The data were modified to accommodate schedule changes and flight number changes, but a few flights added to the schedule in the intervening month might have no passengers on them.

Not all data are hard and fast numbers. In particular, the turn times for the crews are scheduled to be 45 minutes, but in fact, crews can generally be ready in 35 minutes if they hustle. We used 35 minutes for crew turn times in this study. The aircraft turn time is a function of the aircraft type and the airport.

Passenger connection times are of course defined by the airline when schedules are created, and are used to schedule connections. However, the actual time needed to make a connection depends on many factors: the agility of the passenger, the distance between gates, and the lateness of the connecting flight. Overseas connections require extra time for customs and immigration. To accommodate this situation, NWA has created a proprietary probabilistic model for passenger connection times. If the effective connection time (ECT = scheduled *out* time – actual *in* time) is large, the chance of making a connecting flight is 100%, but even if the connection time is negative, there is a finite chance of making the connection because the connecting flight might be late or held for connecting passengers or baggage. We used this model when evaluating passenger misconnection costs. There are also costs for passenger ill will, which begins when a flight is 15 minutes late; for interrupted trip expenses, which begins when a flight is 95 minutes late; and for baggage misconnection costs.

Because of the details of crew contracts, the only crew cost that occurs is when the crew exceeds its scheduled block time (scheduled *in* – scheduled *out*). Increased block times can arise from late arrivals and from early departures. We observed some flights pushing off up to 10 minutes ahead of schedule. Although they arrived early, they still incurred a block delay cost.

Ground operation overtime costs occur when a flight takes at least a 30-minute delay. However, by this time, passenger costs are much greater than ground operations costs, so the latter were ignored.

Fuel cost can be broken into gate, taxi, and air delays. We assumed that taxi times, gate costs, and air delays were unchanged by our resequencing, but did calculate the fuel costs (or savings) for resequencing. Fuel costs increase if the plane is speeded up to land earlier and decrease if it slows down to land later.

NWA has maintenance contracts for A320 and DC-10-30 engines based on hours flown. Minor changes in enroute time are assumed to have no impact on either engines or airframes.

Except for baggage misconnection costs, all costs increase monotonically with time. For baggage, the costs increase for a period, and then decrease because the passenger was probably delayed along with his baggage, and then they increase again because the baggage is probably lost. However, when the baggage cost formula is combined with the passenger costs, the result always increases with time. This will be an important factor when determining how to resequence the flight-landing queue.

SCOPE OF THE STUDY

Our study focused on NWA operations at three hubs, Minneapolis (MSP), Detroit (DTW), and Memphis (MEM), over a three-day period (April 6–8) in 1998, and a week period (July 30 –August 5) in 1999. These periods covered a variety of weather and operational conditions. The airline operations for both periods are summarized below in excerpts from the NWA ATC Coordinator daily log highlights.

1998

April 6	Thunderstorms directly impacting MSP terminal operations
April 7	Rain showers at MSP much of the day reducing the Airport Arrival Rate (AAR)
April 8	Thunderstorms at MEM and in the Ohio Valley

1999

July 30	Thunderstorms at MSP causing diversions and single runway ops much of the day
July 31	Thunderstorms across Florida
August 1	Thunderstorms impacting the Washington DC. Metro airports
August 2	Thunderstorms impacting central Florida
August 3	Thunderstorms impacting MIA
August 4	Thunderstorms impacting the southern Great Lakes & Florida. Ground stops at DTW
August 5	Thunderstorms impacting PA through New England in the afternoon

Airline performance for the period April 6 through April 8 was influenced by the following operational factors:

- Labor actions causing high numbers of cancellations on both April 7 and 8
- Runway constructions at MSP reducing the AAR from 60 to 52
- Thunderstorms directly impacting the MSP terminal area causing 15 MSP diversions

Airline Performance for the period July 30 through August 5 was influenced by the following operational factors:

- Runway construction at MSP and MEM
- Runway closure at DTW
- Taxiway construction at all three airports
- Thunderstorms directly impacting MSP and DTW terminal areas
- Thunderstorm outside the HUB airport, terminal areas causing reroutes and a lower AAR
- Thunderstorms outside the DTW terminal area creating enroute congestion impacting DTW
- Ground stops at DTW
- Single Runway operations at MSP
- High levels of enroute congestion and ground stops in the Northeast

The daily percentage of flights that arrive less than 15 minutes late (excluding cancellations) is called the DOT A14%, and is a good metric of performance. Table 1 shows the A14% for NWA for the 1998 and 1999 study periods. A0% is the percentage of flights that arrive on schedule, and A30% is the percentage that arrives within 30 minutes.

Table 1. NWA Arrival Performance.

Date	Destination	A0%	A14%	A30%	Arrivals
4/6/98	DTW	71.9%	92.1%	96.7%	366
4/6/98	MEM	84.2%	93.0%	98.2%	114
4/6/98	MSP	50.6%	76.5%	86.6%	328
4/7/98	DTW	61.3%	83.8%	93.8%	357
4/7/98	MEM	58.4%	85.8%	92.0%	113
4/7/98	MSP	55.8%	81.1%	90.2%	317
4/8/98	DTW	37.7%	69.6%	84.1%	358
4/8/98	MEM	35.7%	51.8%	74.1%	112
4/8/98	MSP	53.3%	79.7%	88.3%	315
7/30/99	DTW	61.4%	82.1%	89.5%	363
7/30/99	MEM	67.0%	83.5%	89.9%	109
7/30/99	MSP	40.4%	56.8%	68.6%	354
7/31/99	DTW	44.0%	63.2%	70.2%	302
7/31/99	MEM	53.8%	76.9%	80.8%	104
7/31/99	MSP	56.9%	74.8%	83.1%	325
8/1/99	DTW	63.7%	83.4%	92.1%	331
8/1/99	MEM	69.4%	92.6%	96.3%	108
8/1/99	MSP	68.5%	84.9%	93.2%	352
8/2/99	DTW	66.6%	90.4%	95.9%	365
8/2/99	MEM	67.3%	85.5%	95.5%	110
8/2/99	MSP	73.0%	89.8%	95.3%	363
8/3/99	DTW	73.9%	90.8%	96.5%	368
8/3/99	MEM	64.9%	91.0%	98.2%	111
8/3/99	MSP	60.7%	87.3%	94.2%	361
8/4/99	DTW	43.4%	61.9%	73.4%	357
8/4/99	MEM	63.6%	83.6%	88.2%	110
8/4/99	MSP	61.6%	87.0%	93.9%	362
8/5/99	DTW	62.9%	84.6%	91.9%	369
8/5/99	MEM	49.1%	79.1%	90.0%	110
8/5/99	MSP	68.9%	89.8%	93.7%	363

From NWA's point of view, August 2, 3, and 5 were good operational days. We have similar data for the two NWA commuter airlines and for departure performance.

AIRLINE CONNECTIVITY

Efficient resequencing cannot be performed by only considering the arriving flights independently of their connectivity to the remainder of the airline schedule, because this connectivity causes downline delay effects that may have more cost impact than the initiating flight. In the delay model, the equipment and the crew cause this connectivity. In order for a flight to take off, the crew and equipment must be available and have had time for their required procedures. We also assumed, in our connectivity calculations, that takeoff could not occur before the time scheduled in the Official Airline Guide (OAG) (although this can happen in practice).

Equipment

The Out-Off-On-In (OOOI) data contain the tail numbers for NWA flights and commuter advanced regional jets (ARJs), but not for commuter propeller planes. We sorted the NWA Excel spreadsheet data by tail number, date, and *out* time to obtain the flight leg sequence for each aircraft. We use a list of equipment turn times (a function of the aircraft type and airport category) to determine how much turn time is required before the equipment is ready for the next flight leg. Thus, equipment availability time = actual *in* time + turn time.

Cockpit crew and flight attendants

We have the crew rotations for all NWA flights, both scheduled and actual. We used the actual rotations to determine the sequence of flight legs for each member of the crew. We treated deadhead legs just like ordinary legs. Thirty five minutes was used as the minimum turn time for all crew.

Crew availability time = actual *in* time + 35

In the 1998 version of the study, we assumed that the crew connect time was 45 minutes. Reducing this by 10-minutes made a big difference in the number of downline flights delayed due to late crew availability. The delay chains were much shorter and less frequent using the 1999 value of 35 minutes.

Delay chains

It is important to understand how delays can propagate through the airline's schedule. We developed a graphical representation of these delay chains that helps to clarify several issues.

Figure 1 shows an example of a delay chain created using a 45-minute crew connect time to enhance the effects. Three horizontal lines for the *actual*, *scheduled*, and *model* block times represent each flight leg. The lines connecting the flights show the elements that connect the flight legs, equipment, crew, or flight attendants. For example, the first flight (NW735) delays three other flights. The turn times discussed above must be obeyed between flights. Because we should not predict the future with our model, future block times used in the model are always the same as the scheduled block time. The lower figure shows the effect of moving the on time for NW735 forward by 9 minutes. The number of delayed flights is dramatically reduced (from 2 to 17) and the total number of delay minutes is cut from 578 to 405.

Delay Chain for NWA735

21 flights affected

Total departure delay = 576 min

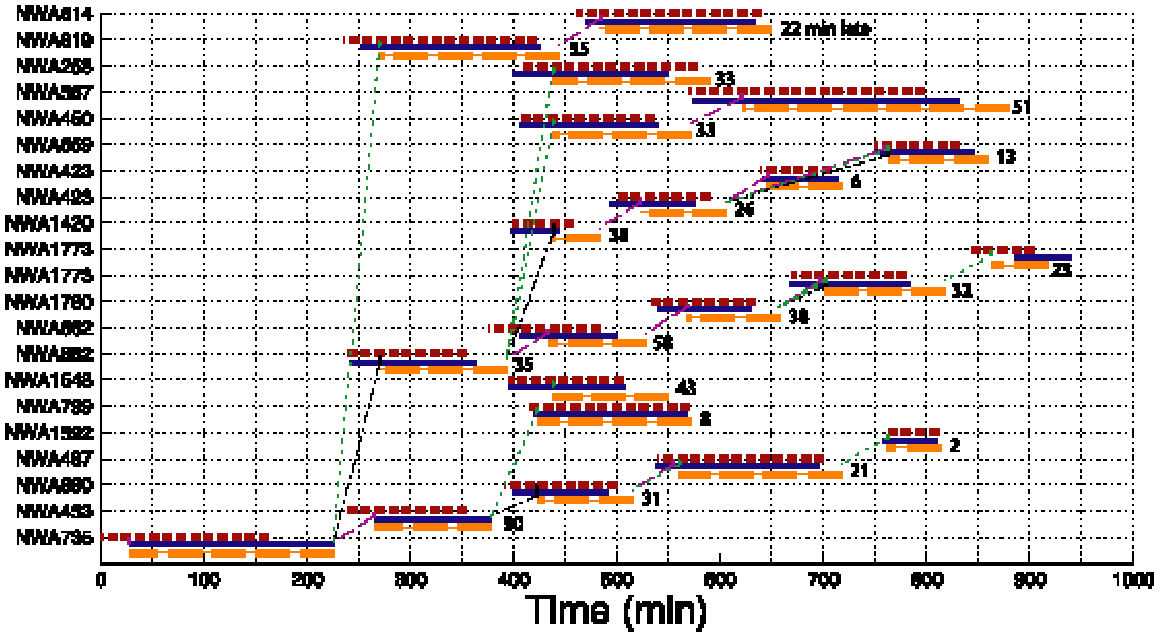

17 flights affected

Total departure delay = 405 min

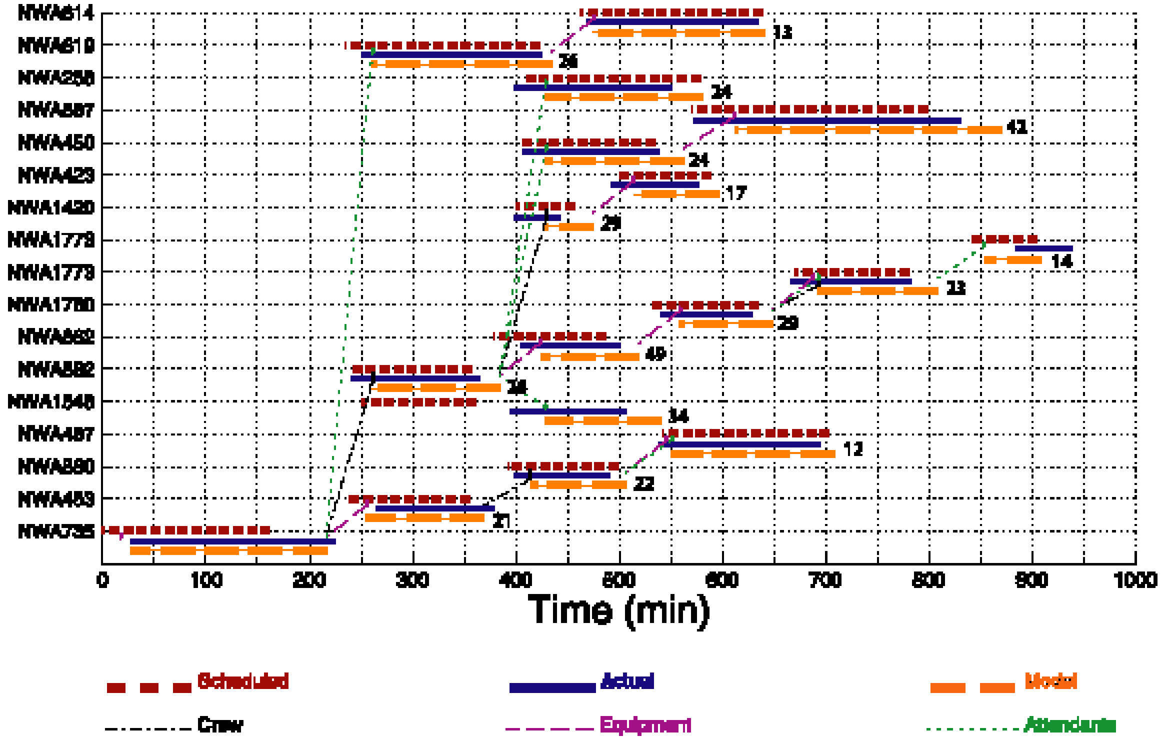

Figure 1. Delay chain for NWA 735. Top figure is as flown which delays 21 flights. Bottom figure shows the effect of moving the initial landing time up by 9 minutes. Only 17 flights are affected and the number of minutes of delay is cut from 578 to 405. A 45-minute crew turn time was used.

Using 35-minute turn times for the crew, we saw many fewer delay chains, and they were of shorter duration. Changes in schedule are evident in the delay chain plots when a leg in the chain is scheduled to depart before its resources were scheduled to arrive. In such cases, there are large delay costs associated with the rescheduled flight that are not associated with the initiator of the delay chain. Nonetheless, moving the head of the chain can still cause large changes to these costs. These changes are calculated correctly and do arise from the resequencing process.

Figure 1 shows that the delay chains can spread, affecting many more flights as time goes on, and also that a given flight can be affected by more than one resource. In this study, we examined resequencing one bank at one hub at a time. It is easy to envision delay chains originating from multiple flights intersecting in the future. If the heads of these chains were not in the same bank at the same hub, we will miss this interaction. Our current model treats resources originating from other banks using their "as flown" times. However, our model does correctly calculate which resource for each future flight leg is "controlling," and recursively traverses the multiply-linked list of flight legs to determine the effects of the crew and equipment on the lateness of a given future flight leg.

The cost model, correctly applied to the flights in each delay chain, determines the cost changes when a flight in the arrival queue is moved forward or backward. Most of the model is concerned with reading the airline data, determining the connectivity of each flight leg, and calculating the cost (including downline effects) for each flight in each landing slot that it can access by reasonable (10-minute) speedup or slowdown.

RESEQUENCING ALGORITHM

Once the costs for landing each flight in each slot are determined, we can try to optimize the landing queue. This optimization is a computationally separate exercise, and depends upon the definition of optimum. For the purposes of this study, "optimum" was defined as minimum cost to the airline. There may be auxiliary constraints that need to be applied. The only one we considered was to avoid swapping flights if the savings were small (e.g., <$100). The effects of this additional constraint will be discussed later.

The model allows us to calculate the ripple delay costs for each landing slot that a flight could use (within ±10 minutes of the actual *on* time). We use two different cost models. The first one we call the *ripple delay cost*. It is the total cost of delay added up for all affected downstream flights. The ripple delay cost represents the potential impact to the airline for each flight in the arrival queue, and we use this cost to make resequencing decisions. However, when we calculate the savings to the airline, we only use the cost savings up to and including the next arrival at an NWA hub. To go further downstream would count cost savings for this flight more than once. We call this the *delay cost*. We will give examples of the differences between these costs in a later section.

Arrival queue

To clarify the algorithm, Figure 2 shows the situation at the end of the first bank at Memphis on April 7, 1998. The left three columns show the NWA slot number, local time, and the flight number that actually landed in that slot. The other columns contain the cost of all flights that can land in each available slot. Slots past the ±10 minute variation contain an asterisk. The flight selected by the optimizing algorithm (see below) for each slot is outlined, and because the algorithm works backwards in time, the earlier slots for that flight are shaded to show they are unavailable. Several slots may occur at the same time because multiple runways are used (e.g., slots 45 and 46), or a given time may not contain an NWA landing slot (8:04). For example, NW1741originally landed at 8:09 in slot 54 for a cost of $160. It could have landed in any slot between 7:59 and 8:19, inclusive. If this flight were moved earlier (left in Figure 2) to 8:06, the cost would be eliminated. If it were delayed more, the cost would increase. Therefore, the task is to determine which flight to land in each available slot to achieve the minimum cost for the airline.

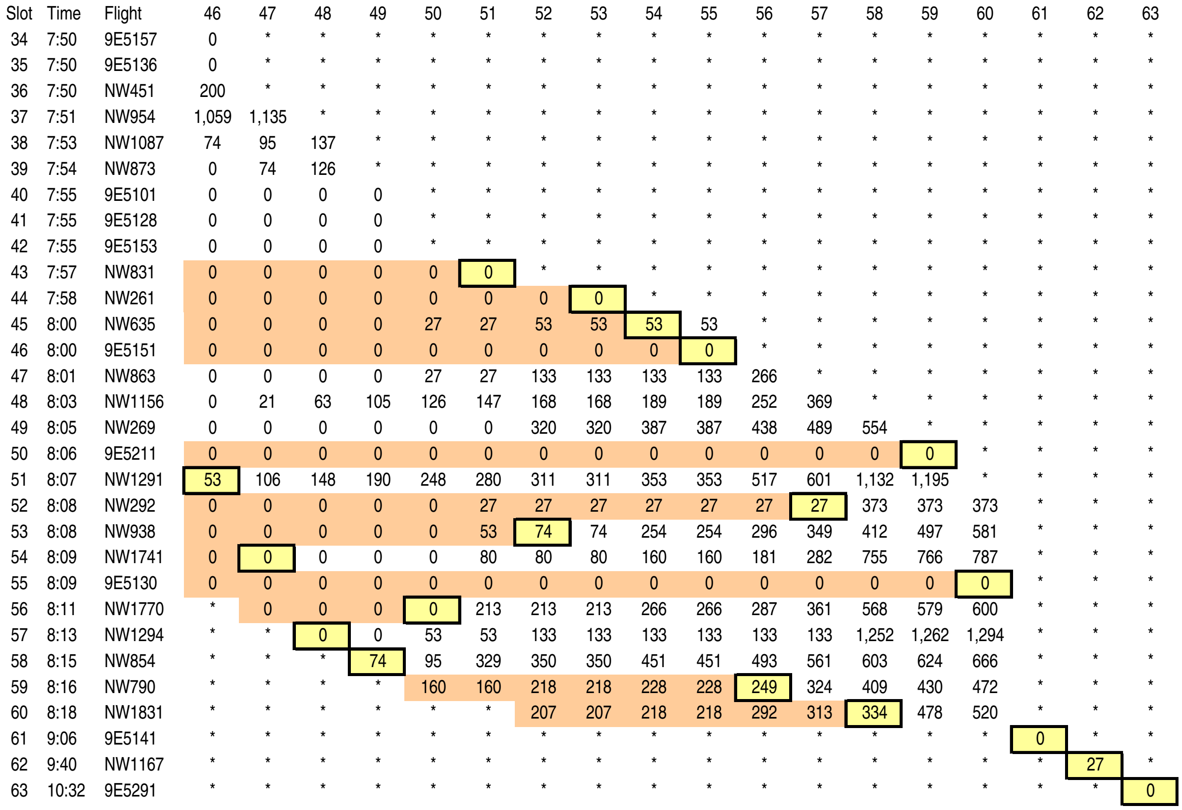

Slot	Time	Flight	46	47	48	49	50	51	52	53	54	55	56	57	58	59	60	61	62	63
34	7:50	9E5157	0	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*
35	7:50	9E5136	0	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*
36	7:50	NW451	200	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*
37	7:51	NW954	1,059	1,135	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*
38	7:53	NW1087	74	95	137	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*
39	7:54	NW873	0	74	126	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*
40	7:55	9E5101	0	0	0	0	*	*	*	*	*	*	*	*	*	*	*	*	*	*
41	7:55	9E5128	0	0	0	0	*	*	*	*	*	*	*	*	*	*	*	*	*	*
42	7:55	9E5153	0	0	0	0	*	*	*	*	*	*	*	*	*	*	*	*	*	*
43	7:57	NW831	0	0	0	0	0	0	*	*	*	*	*	*	*	*	*	*	*	*
44	7:58	NW261	0	0	0	0	0	0	0	0	*	*	*	*	*	*	*	*	*	*
45	8:00	NW635	0	0	0	0	27	27	53	53	53	53	*	*	*	*	*	*	*	*
46	8:00	9E5151	0	0	0	0	0	0	0	0	0	0	*	*	*	*	*	*	*	*
47	8:01	NW863	0	0	0	0	27	27	133	133	133	133	266	*	*	*	*	*	*	*
48	8:03	NW1156	0	21	63	105	126	147	168	168	189	189	252	369	*	*	*	*	*	*
49	8:05	NW269	0	0	0	0	0	0	320	320	387	387	438	489	554	*	*	*	*	*
50	8:06	9E5211	0	0	0	0	0	0	0	0	0	0	0	0	0	0	*	*	*	*
51	8:07	NW1291	53	106	148	190	248	280	311	311	353	353	517	601	1,132	1,195	*	*	*	*
52	8:08	NW292	0	0	0	0	0	27	27	27	27	27	27	27	373	373	373	*	*	*
53	8:08	NW938	0	0	0	0	0	53	74	74	254	254	296	349	412	497	581	*	*	*
54	8:09	NW1741	0	0	0	0	0	80	80	80	160	160	181	282	755	766	787	*	*	*
55	8:09	9E5130	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	*	*	*
56	8:11	NW1770	*	0	0	0	0	213	213	213	266	266	287	361	568	579	600	*	*	*
57	8:13	NW1294	*	*	0	0	53	53	133	133	133	133	133	133	1,252	1,262	1,294	*	*	*
58	8:15	NW854	*	*	*	74	95	329	350	350	451	451	493	561	603	624	666	*	*	*
59	8:16	NW790	*	*	*	*	160	160	218	218	228	228	249	324	409	430	472	*	*	*
60	8:18	NW1831	*	*	*	*	*	*	207	207	218	218	292	313	334	478	520	*	*	*
61	9:06	9E5141	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	0	*	*
62	9:40	NW1167	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	27	*
63	10:32	9E5291	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	0

Figure 2. Queue data for NWA arrivals at the end of the first bank in Memphis on April 7, 1988. The outlined cells designate the flight selected for landing in each slot. The numbers in the cells are the incremental cost in dollars for landing each flight.

The algorithm

The first task is to determine the start and end of an arrival bank. For the purposes of our study, any gap of more than 10 minutes between landing slots will delineate a bank because swaps cannot span a longer time gap. At the start of a bank, a flight can only land in its assigned slot or in later slots. At the end of a bank, the flight can only land in earlier slots or in its assigned slot. In Figure 2, slot 60 is the end of the first bank

The key to the resequencing algorithm is to start from the end of each bank and to work forward. The "cost" to optimize is the cost in a slot minus the cost in the first available slot for the flight. It may occur that a flight incurs a cost no matter in which slot it lands. Only the cost savings count when performing the resequencing. The rules of the algorithm are as follows:

- Don't do any resequencing unless the number of available slots exceeds the number of planes that must land before they run out of slot times on the early side. In slot 49, we must take flight 854 because there are no earlier slots, later slots are already taken, and it must land. This situation also occurs for flight 1294 in slot 48. We discuss this in more detail in the section *Results of the Analysis*.
- In each slot, select the minimum cost flight and land it in the slot (9E5130 in slot 60). In slot 58, the lowest cost flight is NW1831 ($334 – $207 = $127). In slot 56 it is NW790 ($249 –$160 = $89)
- In case of a tie for minimum cost, land the flight with the latest original slot time, i.e., the bottom flight on the spreadsheet. (Pick NW1770 in slot 50)
- Because we select the lowest cost flight in each slot, and costs increase with time, this algorithm yields results that are close to optimum. Unaccommodated flights cost less in an earlier time slot.

- This resequencing algorithm has the ability to sacrifice one flight for the common good. NW635 was delayed, thus incurring an extra $27 in cost because landing any other flight in this slot would cost at least $133.

You might think that the same results could be achieved by starting at the beginning of a bank and working forward in time. Indeed, this method was the first approach tried. In this case, you must look ahead to find the flight that is going to yield the biggest cost and land it first. Flights must be considered that would land more that 10 minutes ahead, and so are not listed in the column of the slot being optimized. By starting at the back and working backward in time, we obtain an easier to code, and easier to understand algorithm that works better in all cases.

It is not trivial to determine whether resequencing is allowed at a given slot, because we must look backward in time for all the flights that could land in this slot, and make sure that enough slots are available to land them all. The algorithm for this is as follows:

- Create a *needs* array of dimension at least as big is the maximum number of planes that could land in any slot, and set all elements to zero.
- For each plane that could land in the slot (and has not already landed), find the number of available slots, *ns*, between the slot time and the earliest possible landing time for the flight. If there are multiple flights landing at this slot time, do not count slots with a higher slot number.
- Add 1 to each element of the *needs* array from *ns* upward, indicating that a flight could need to land in that slot.
- If any element of the array is equal to the index of that element, it means that there are just enough slots available to land the flights under consideration and hence, swapping is not allowed.

It may occur that a subgroup of *n* contiguous flights all must land in the same range of slots. In that case, we determine the flights that can swap among themselves and allow resequencing among these *n* flights.

In general, determining an optimum algorithm for resequencing is a very complicated Operations Research problem. However, in our case, several things make this easier. First, as discussed in the cost model section, the cost of delaying a flight increases monotonically in time. So landing a flight earlier will never cost more if fuel costs (for speedup) are ignored. Second, all planes are subject to the same cost model, so the change in costs between slots for one flight is generally similar to the change in costs for another flight. Finally, the cost slopes for later flights are usually higher because more costs start to kick in. If the cost slope was the same for all flights, we could prove that our resequencing algorithm was optimum. Because the cost slopes do change from flight to flight, we can merely say that this algorithm is easy and fast to apply and yields excellent results. A counter example is shown in Figure 3. The figure caption explains why this case is unusual.

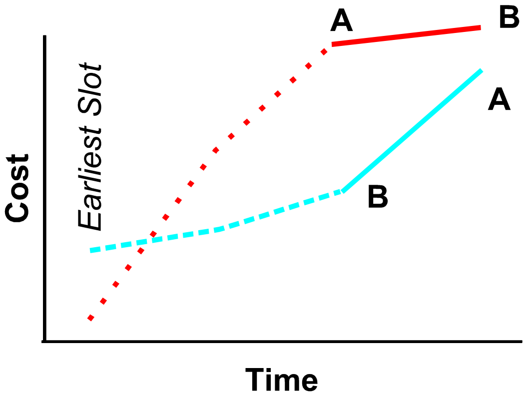

Figure 3. A counter example for the resequencing algorithm. The algorithm would select the two slots labeled *A*, whereas the selection labeled *B* would yield more cost savings. (Remember it is the difference between the current cost and the earliest possible cost that counts.) However, in practice, most cost curves are concave up because the cost rate increases as time goes on (up to a maximum incurred cost). In addition, the curve with the highest cost usually has the greatest slope because more passengers are affected and additional costs arise as the flight becomes later. In this case, the curve with the higher slope would be above the curve with the lower slope and the algorithm is optimum.

Costs increase with time

The resequencing algorithm assumes that costs increase monotonically with increasing delay. This is true for everything except the fuel costs for speedup or slowdown. Fuel costs generally decrease when a plane is slowed down and increase when it is speeded up. Accordingly, we ignored the fuel speedup costs (which are usually much smaller than the passenger delay costs) when performing the optimization. Then, after the resequencing, we adjusted the costs for the resequenced flights to reflect the speedup fuel costs. As a result, some flights that were speeded up to reduce passenger costs actually cost a little more due to these increased fuel costs. On the other hand, some delayed flights incurred a negative cost (savings) due to the slowdown. If flights with small cost savings are not swapped, the increased fuel costs will become negligible. We examine the effects of such actions in the analysis section.

The cost of misconnected luggage is also non-monotonic. It increases with time, then decreases because the owner is delayed also, and then increases again. However, the increases in other passenger costs always make the sum of the passenger costs increase monotonically.

Airplane Performance

Using Aircraft Operating Manual data provided by NWA Performance Engineering a series of flight path computations were performed for each of the Northwest Airlines turbo jet aircraft:

- B747 (-200, -400)
- DC10 (-40 -30)
- B757 –200
- DC9 (MD80, -30, -40, -50)
- A310, A320
- B727-200
- ARJ
- Turbo Props (SF series)

To simplify the model, aircraft fleet types were combined where the performance and cost deltas were similar.

Flight profiles were computed over a wide range of speeds and altitudes, and a cost model was built for the variance of time around the nominal operating speed for each airplane type. Time-fuel-speed (Cost Index) profiles were integrated to resolve the minimum cost for a given flight time for each aircraft. The cost model included the price of fuel and an allocation for the operating cost per minute for aircraft itself. The strategic value for crew and airplane utilization costs was not included because the crew block time is included in the main model. In general, it is believed that resequencing will in fact keep airlines operating closer to their published schedule and therefore within their business allocations for aircraft and crew utilization costs.

The data in Table 2 for the B747-400 are typical of the cost delta for flight time variance around nominal.

•

Table 2. B747-400 flight data.

Cost Index		Time Variance	Fuel Cost ($)	A/C Cost ($)	Cost Delta ($)
300		–12	351	(104)	247
200		–7	170	(61)	109
107	Nominal	0	0	0	0
100		1	(12)	9	(4)
0		15	(145)	130	(15)

The descent profiles were based on Idle Thrust Descents from the Top of Descent Point (TOD) to FL100, deceleration to KCAS250, continued descent and deceleration to the final landing runway at Vth+10 knots. These data produce a general view of airplane performance that can be applied to all turbo jet aircraft.

Specific engineering data for the turbo prop aircraft was not available so an assignment of cost was made using the basic turbo jet cost profile and a best fit achieved for the turbo prop aircraft from knowledge of their fuel consumption and operating costs.

The model was further adjusted in recognition that not all the speed change would be achieved by changing the cruise profile for the aircraft. Other techniques can be used to reduce or increase the time to landing such as:

- Reducing Distance to Landing (more direct routing to the airport)
- Increasing Distance to Landing (such as extending the downwind-base-final)
- Reducing time to touch down by utilizing drag devices to enable a steeper approach profile allowing for higher cruise speed to be held to a closer TOD point.

The airplane flight profile data demonstrated that for both turbo jet and turboprop airplanes, a time window for speed up/slow down of –10 (earlier) to +10 (later) around the nominal landing time is practical and achievable. As a result of these considerations, we arrived at the fleet cost model for NWA equipment that was used to calculate fuel costs for speedup or slowdown during the landing phase of the flight as shown in Table 3.

Table 3. Fleet cost model. Costs are in Dollars.

Time Δ	B727	DC10	B747	B757	DC9	RJ	SB340	A319/320
-10	90	180	210	110	70	45	30	75
-9	80	160	187	98	62	40	27	67
-8	70	140	163	86	54	35	23	58
-7	60	120	140	73	47	30	20	50
-6	50	100	117	61	39	25	17	42

-5	40	80	93	49	31	20	13	33
-4	30	60	70	37	23	15	10	25
-3	20	40	47	24	16	10	7	17
-2	10	20	23	12	8	5	3	8
-1	5	10	12	6	4	3	2	4
0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
1	(3)	(6)	(7)	(4)	(2)	(2)	(1)	(3)
2	(5)	(10)	(12)	(6)	(4)	(3)	(2)	(4)
3	(7)	(14)	(16)	(9)	(5)	(4)	(2)	(6)
4	(9)	(18)	(21)	(11)	(7)	(5)	(3)	(8)
5	(11)	(22)	(26)	(13)	(9)	(6)	(4)	(9)
6	(13)	(26)	(30)	(16)	(10)	(7)	(4)	(11)
7	(15)	(29)	(34)	(18)	(11)	(7)	(5)	(12)
8	(16)	(32)	(37)	(20)	(12)	(8)	(5)	(13)
9	(18)	(35)	(41)	(21)	(14)	(9)	(6)	(15)
10	(19)	(38)	(44)	(23)	(15)	(10)	(6)	(16)

Equipment connectivity

- Flight leg data for NWA and its two commuter airlines were based on actual tail number utilization, (*actual* rather than *planned* tail number determined equipment connectivity).
- There was no tail number information available for the two commuter airlines except for the ARJs in 1999.

Crew patterns for pilots and flight attendants

- Patterns were based on actual flown rather than scheduled.
- The deadhead crews (crews being transported to or from their active flights) were treated the same as regular crews for the connection matrix. But often they are not on a tight schedule to meet a flight.
- Crew information for the two commuter airlines was not available.

Commuter flights are treated with prejudice

Aside from ARJs in 1999, there is no downline connectivity for the two commuter carriers. Consequently, there are no downline delay costs. In addition, commuters carry fewer passengers so that the ill-will costs are always lower than on an NWA flight delayed by the same amount. Because the cost of delaying commuter flights is small, the resequencing algorithm will delay commuter flights and advance NWA flights, all things being equal.

UNDERSTANDING THE MODEL OUTPUT

A few words are in order to explain the assumptions and limitations of the model, to show examples of the output data, and to explain where the final numbers come from.

Internal Consistency

The primary "errors" in the model concern the quality of the input data—its consistency and its correctness. The model reads, parses, and stores hundreds of megabytes of data for a week's run. These data sets are all preprocessed using Excel. The commuter data are not always in the same format (or contain the same quantities) as the NWA data, so these must be converted and reconciled. Most problems in this process occur at date boundaries, conversions to and from Zulu (GMT), or in associating a time with the correct date. However, the airline data are collected on a daily basis, and we want the "daily" runs to be associated with a local day, not a Zulu day. These issues lead to problems when a flight crosses a date boundary between *out* and *off* or between *on* and *in*. The code and the Excel preprocessing must look for and correct these problems. The recursive nature of the delay calculation and the complicated interconnectedness of the airline are very efficient at pointing out any such problems. Nonetheless, a few such cases remain, and the ones we have examined manually all have faulty or inconsistent input data.

The code prints out numerous diagnostics for items such as large delay times, large numbers of passengers, and passenger connections that cannot be found.

Limitations of the Model

Given the duration of this study and the complexity of the data being analyzed, it was not possible to build a model that could simulate the total airline operation.

The current model does not deal with:

- Determining if moving flights forward could reduce the number of overnight passengers.
- Calculating if moving up the last flight of the day could avoid or reduce a crew rest delay before the next day's departure.
- Costs associated with gate occupancy impacts for early flight arrivals or the delays and additional taxi in times caused by late departing flights.
- Complete simulation of the resequencing of flights in the arrival stream for the entire airline.

The model calculates each delay chain separately. In general, delay chains from multiple flights may intersect a few airports downstream, and the downstream flight must await all its resources before departure. The model has the capability of following all resequencing downline for the entire day. However, to simulate the rest of the day we would have to reorganize all of the data so that the entire NWA system was advanced simultaneously rather than an airport at a time. There was insufficient time to perform this effort. To resequence the whole airline schedule, it would also be necessary to resolve the availability of landing slots at the future times that the model decides they are needed. In an extreme case, the demand for arrival slots could bunch up so that dozens of flights might want to land at the same time.

The model currently gives a valid representation of the activities for one bank. As was mentioned earlier, rescheduled equipment or crew can lead to downline costs that may not be attributable (solely) to the landing time of a given flight. Because of this, the "cost" of landings at a hub on a given day must be taken as indicative rather than absolute. However, the cost savings achieved by resequencing are independent of these issues and are accurate.

The Landing Slot Viewpoint

Because we are resequencing the landing queue to save money for the airline, one way of looking at the model output is on a slot-by-slot basis. Unless extra slots are introduced (as will be discussed later), NWA has n landing slots for n flights, so the optimization algorithm is trying to find the rearrangement of these slots that will save the most money. Other constraints can also be considered.

If a flight originally lands in a slot s at time t, it incurs a certain cost to the airline. By speeding up or slowing the flight, it could be possible land in earlier or later slots. For this study, we chose a ±10-minute landing window from its actual landing time. This may encompass a range of slots that depends upon the number of nearby NWA flights. For example, a flight scheduled to land at the start of a bank can only land in later slots, and one at the end of the bank can only land in earlier slots. Given the available slots for a flight, the cost model determines the cost of landing the flight in each slot. Then the optimizer resequences all the flights

There are many costs associated with a given flight. For the purposes of resequencing the landing queue, we use the costs for the incoming flight plus the costs for all the downstream flights that will be affected by this flight. We call this the *ripple delay* cost because it gives the airline scheduler the best idea of the eventual impact of moving this flight. However, after the flights from one bank are resequenced, we cut off the downline costs the next time a flight returns to a hub when we calculate the savings to the airline (to avoid double counting), and call this the *delay cost*.

We have developed a three-dimensional plot that illustrates the day at a hub from the slot viewpoint. We illustrate this in Figures 4 and 5 with a plot of arrivals at MSP on 7/30/1999 — a very bad day for NWA. The *Original Slot* axis represents the slot for each flight. The time of the slot is printed in the front, and the number of original flight that landed in the slot is printed in the rear. Colored boxes are plotted for each other slot that the flight could land in (the *Δ Slot* axis). In fact, the horizontal plane is just Figure 2 replotted along the diagonal axis. The color of the box represents the ripple delay cost if the flight were to land in the slot. Earlier slots are in front and later slots are behind the *Δ Slot* axis. Banks start and end when the number of early or late slots for a flight goes to zero.

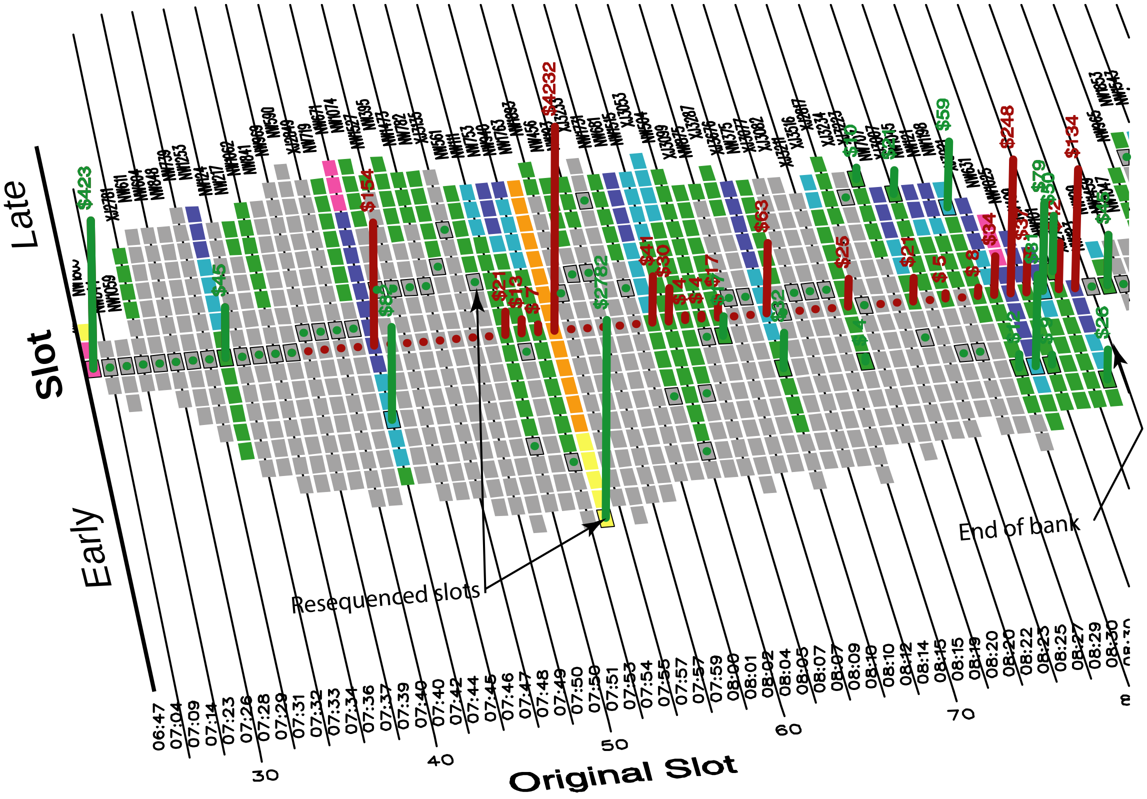

Figure 4. An enlargement of Figure 5 to display the details. For example, NW1763 originally was to land in slot 53 at 7:55 AM for a ripple delay cost of $4232. It was swapped with another flight to land 10 minutes earlier, so that ripple delays cost $2782, a potential savings of about $1500. The outlined rectangles represent the resequenced landing slot. Before resequencing, all planes landed on the *Δ Slot* = 0 axis.

A gray patch in a slot represents no cost to NWA for landing in that slot. Colored patches in front of the Δ Slot = 0 line represent cost-savings opportunities for the airline. Gray slots behind the axis represent flights that can be delayed without incurring a cost. In this run, we did no resequencing if the cost for all flights that could land in the slot was below $100. Accordingly, the first flight that needs a swap originally landed in the slot at 7:42 am.

The vertical bars represent the pre (red) and post (green) costs for the flight that was originally in the slot. A bar that moves forward (Δ Slot < 0) represents a flight that was moved earlier, and a bar that moves backward represents a flight that moves later. The vertical (Cost) axis is a nonlinear scale that represents the cumulative *delay cost* savings achieved by resequencing.

As the day progresses, things get steadily worse due to a thunderstorm line that moved across the airport. On the last two plots, it is evident that bank integrity has been destroyed; almost every flight arrived late. The gray patches have almost vanished. The front-to-back extent of the possible arrival patches gets smaller, indicating a reduced AAR.

Unfortunately, these 3-D plots get rather crowded when many flights are delayed. It is then useful to examine the raw data by slot. In Table 4 we show the time corresponding to the slots from 293–320.

Table 4. Slot occupancy before and after resequencing.

Slot Number	Slot Time	Local Time	Original Flight	Minutes Late	Original Cost	Rescheduled Flight	Minutes Late	Optimum Cost	Rescheduled Cost
293	*15556165*	*16:25*	*NW1021*	*53*	*1280.22*	*XJ2767*	*20*	*45.98*	*39.65*
294	*15556172*	*16:32*	*NW725*	*32*	*15.08*	*XJ3611*	*31*	*219.83*	*236.5*
295	*15556175*	*16:35*	*XJ3611*	*34*	*232.25*	*NW751*	*39*	*1577.01*	*1600.34*
296	*15556176*	*16:36*	*XJ2724*	*50*	*243.64*	*NW137*	*14*	*384.38*	*384.38*
297	*15556176*	*16:36*	*NW137*	*14*	*384.38*	*XJ2724*	*50*	*243.64*	*243.64*
298	*15556176*	*16:36*	*XJ2791*	*22*	*13.01*	*NW725*	*36*	*30.16*	*23.16*
299	*15556179*	*16:39*	*NW751*	*43*	*2307.44*	*NW861*	*37*	*6222.65*	*6320.43*
300	*15556184*	*16:44*	*NW1269*	*42*	*2403.92*	*XJ2791*	*30*	*21.6*	*16.27*
301	*15556188*	*16:48*	*NW861*	*46*	*9664.77*	*NW1269*	*46*	*2858.98*	*2851.98*
302	*15556190*	*16:50*	*NW44*	*-9*	*0*	*NW1056*	*67*	*4096.21*	*4106.21*
303	*15556192*	*16:52*	*NW1056*	*69*	*4339.71*	*XJ3128*	*71*	*1184.1*	*1184.1*
304	*15556192*	*16:52*	*XJ3128*	*71*	*1184.1*	*XJ3158*	*63*	*1460.85*	*1484.18*
305	*15556194*	*16:54*	*NW614*	*10*	*48.75*	*NW102*	*10*	*52.5*	*62.5*
306	*15556196*	*16:56*	*NW102*	*12*	*60*	*NW218*	*-11*	*5076.93*	*5126.93*
307	*15556200*	*17:00*	*XJ3095*	*7*	*0*	*NW614*	*16*	*145.83*	*132.83*
308	*15556200*	*17:00*	*XJ3158*	*71*	*1592.1*	*NW44*	*1*	*14.2*	*-23.8*
309	*15556203*	*17:03*	*NW316*	*13*	*113.6*	*NW316*	*13*	*113.6*	*113.6*
310	*15556203*	*17:03*	*NW218*	*-4*	*5931.45*	*NW678*	*34*	*628.63*	*668.63*
311	*15556205*	*17:05*	*NW1298*	*-7*	*0*	*NW1046*	*17*	*478.98*	*478.98*
312	*15556205*	*17:05*	*NW1046*	*17*	*478.98*	*NW960*	*16*	*233.53*	*256.86*
313	*15556208*	*17:08*	*NW678*	*39*	*876.34*	*NW714*	*3*	*20.96*	*33.18*
314	*15556209*	*17:09*	*NW960*	*20*	*248.61*	*XJ3095*	*16*	*16.74*	*10.91*
315	*15556210*	*17:10*	*NW714*	*5*	*31.44*	*NW1298*	*-2*	*8.28*	*-0.89*
316	*15556211*	*17:11*	*NW1842*	*-9*	*0*	*NW327*	*-2*	*1373.66*	*1420.33*
317	*15556215*	*17:15*	*XJ2877*	*15*	*0*	*NW1842*	*-5*	*15.08*	*8.08*
318	*15556217*	*17:17*	*NW114*	*2*	*20.56*	*NW1252*	*9*	*33.75*	*63.75*
319	*15556218*	*17:18*	*NW327*	*5*	*1906.05*	*NW114*	*3*	*24.7*	*22.2*
320	*15556220*	*17:20*	*XJ3166*	*70*	*777.14*	*NW860*	*-2*	*103.44*	*127.88*

In this table, the Rescheduled Cost column includes the fuel cost for speedup or slowdown. To help illustrate the swaps, we have inserted lines that connect the flights before and after resequencing. Because we display the start of a fractured bank, few slots are available. Thus, in order to move up NW861 (from slot 301 to slot 299) to obtain a $3300 savings, we must move several other flights. Because there are no extra slots, it is not obvious how to do these swaps manually. Furthermore, the aggregate of swaps involves almost all of the flights, although many of the swaps are very small. NWA feels that the cost of a swap of less than five minutes is in the noise level. There is no real penalty to small moves, but unless the move produces significant savings, perhaps it should be avoided. We discuss the effect of requiring a minimum savings level in the *Results of the Analysis* section.

Arrivals to MSP on 7/30/1999

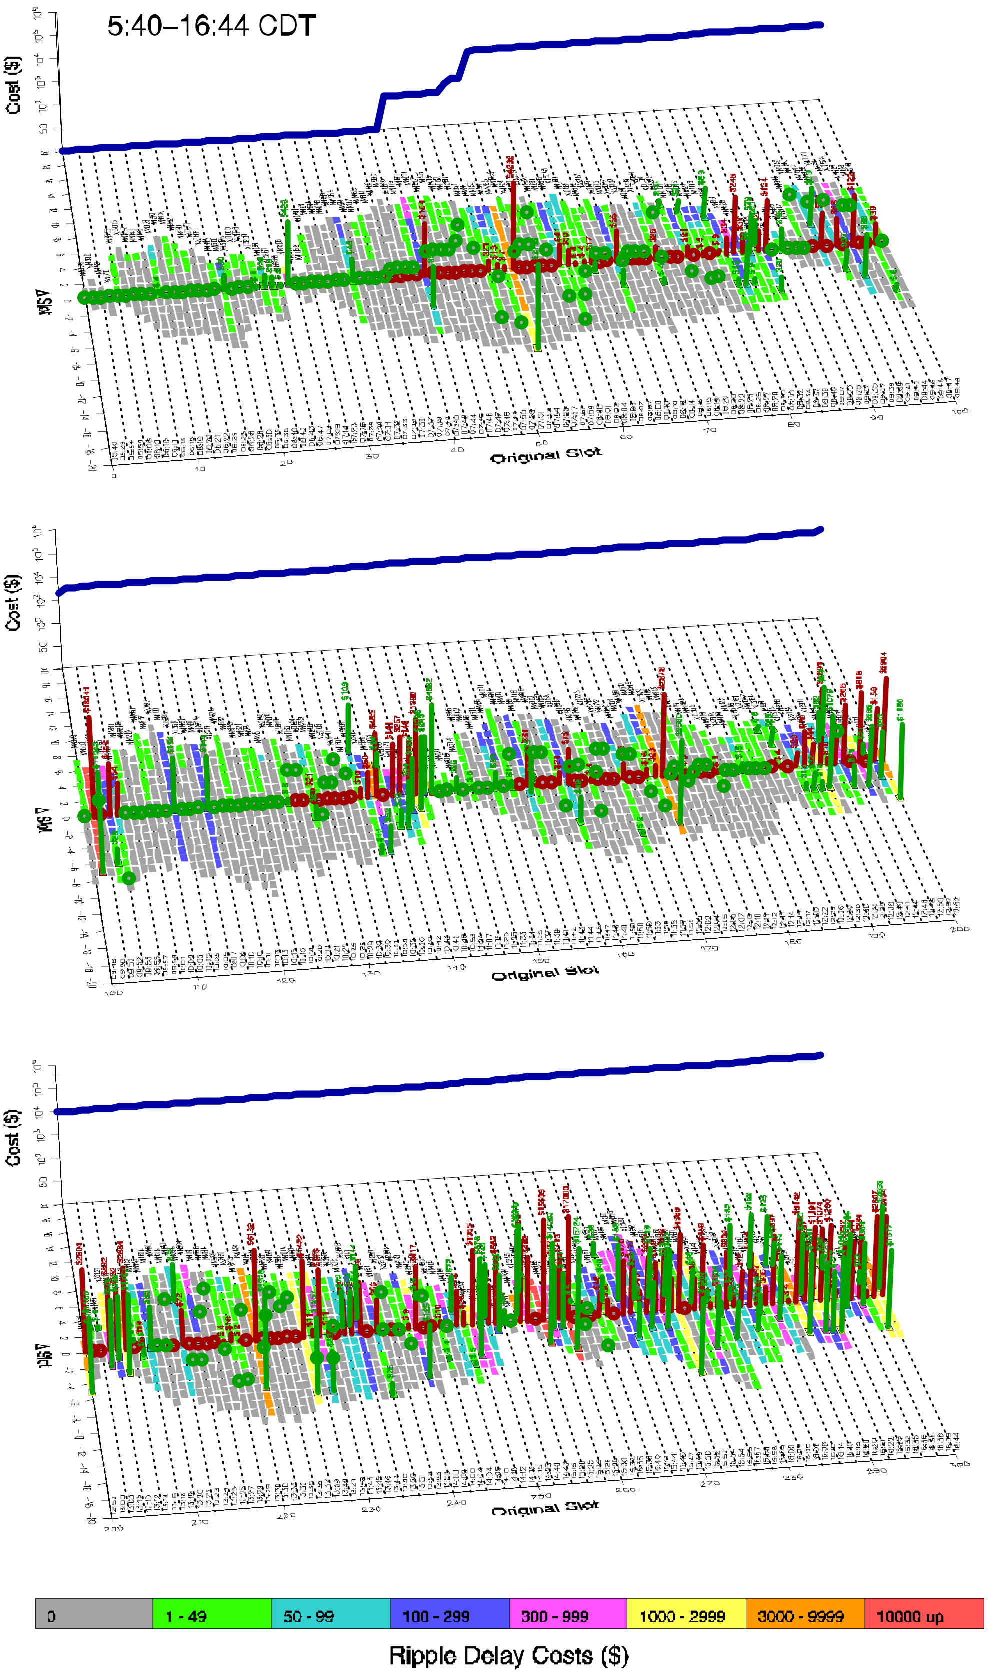

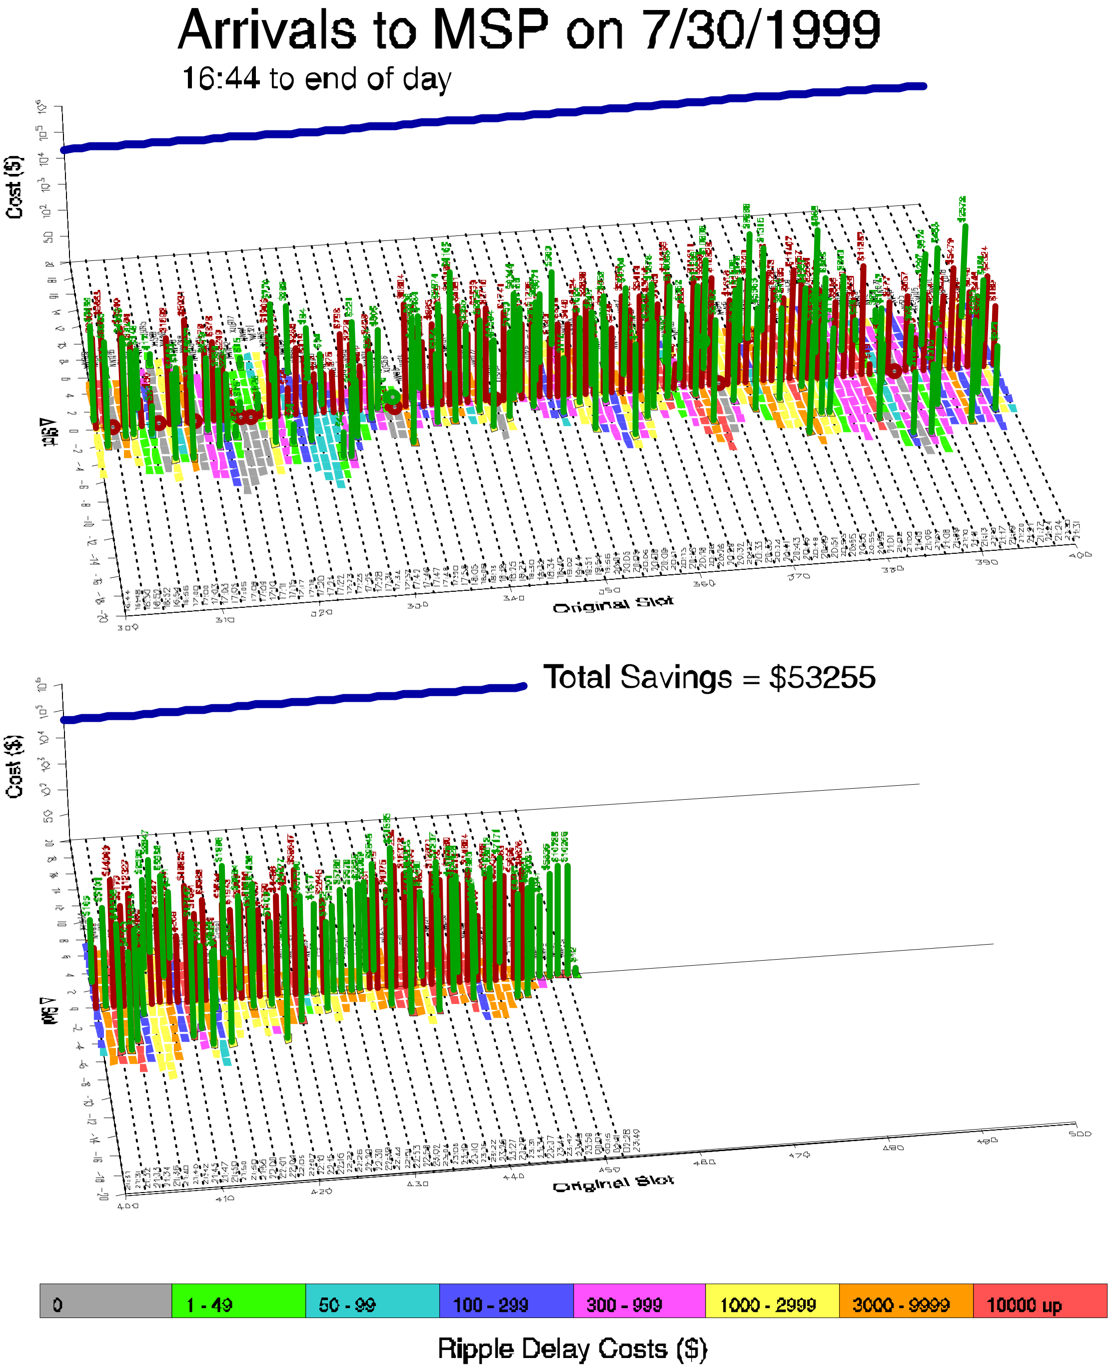

Figure 5. NWA arrival queue for Minneapolis on 7/30/1999 when the airport was impacted by thunderstorms. Resequencing the arrival queue could save over $50,000 for the day.

Extra slots

It is an unnecessary (and usually inaccurate) constraint to restrict NWA to only use its slots for swaps. Although there may be no extra slots available in the middle of a bank, there are certainly many available between banks at the NWA hubs. We decided to add extra slots to the model to see how the optimizer would use them.

For DTW, we determined the weather conditions and hence the airport acceptance and departure rates. We also obtained the OOOI data for all other airlines landing at the airport, although there were some discrepancies in this data, and unscheduled traffic is not included. We also traded off departure and arrival slots when necessary to come up with a list of times when extra landing slots would be available.

To insert these extra slots into the model, we created a fictitious airline (JR) that carries no passengers, flies planes that have no fuel charges, with crews that incur no block delay. Consequently, these flights can be moved at will with no cost. We inserted these flights into the input data and reran the DTW cases for three days. We can illustrate some of the results with excerpts from the DTW run for 7/30/1999

Table 5. Slot usage when extra landing slots (JR flight numbers) are added.

Slot Number	Local Time	Original	Minutes Late	Original Cost	Resched Flight	Minutes Late	Optimum Cost	Resched Cost	Swap?	Empty Slot?
Start of a bank										
740	*19:00*	*JR10259*	*-19*	*0.00*	*JR10259*	*-19*	*0.00*	*0.00*	*0*	*1*
741	*19:01*	*JR10260*	*-19*	*0.00*	*JR10260*	*-19*	*0.00*	*0.00*	*0*	*1*
742	*19:02*	*JR10261*	*-19*	*0.00*	*JR10261*	*-19*	*0.00*	*0.00*	*0*	*1*
743	*19:03*	*JR10262*	*-19*	*0.00*	*JR10262*	*-19*	*0.00*	*0.00*	*0*	*1*
744	*19:04*	*JR10263*	*-19*	*0.00*	*NW1192*	*-27*	*244.86*	*311.53*	*1*	*1*
745	*19:05*	*JR10264*	*-19*	*0.00*	*JR10263*	*-18*	*0.00*	*0.00*	*1*	*1*
746	*19:06*	*JR10265*	*-19*	*0.00*	*NW330*	*77*	*8,812.94*	*8,922.94*	*1*	*1*
747	*19:07*	*JR10266*	*-19*	*0.00*	*JR10264*	*-17*	*0.00*	*0.00*	*1*	*1*
748	*19:08*	*JR10267*	*-19*	*0.00*	*JR10265*	*-17*	*0.00*	*0.00*	*1*	*1*
749	*19:09*	*JR10268*	*-19*	*0.00*	*NW1417*	*87*	*5,290.45*	*5,360.45*	*1*	*1*
750	*19:10*	*NW268*	*-27*	*0.00*	*JR10266*	*-16*	*0.00*	*0.00*	*1*	*0*
751	*19:11*	*JR10269*	*-19*	*0.00*	*JR10267*	*-16*	*0.00*	*0.00*	*1*	*1*
752	*19:13*	*NW1192*	*-18*	*844.69*	*JR10268*	*-15*	*0.00*	*0.00*	*1*	*0*
753	*19:15*	*NW642*	*-14*	*648.00*	*NW268*	*-22*	*0.00*	*-22.00*	*1*	*0*
754	*19:15*	*XJ2975*	*-10*	*0.00*	*JR10269*	*-15*	*0.00*	*0.00*	*1*	*0*
755	*19:16*	*NW330*	*87*	*10,662.73*	*NW642*	*-13*	*648.00*	*645.67*	*1*	*0*
756	*19:17*	*NW696*	*-9*	*72.00*	*XJ2975*	*-8*	*0.00*	*-1.67*	*1*	*0*
757	*19:19*	*NW1417*	*97*	*6,383.39*	*NW696*	*-7*	*72.00*	*68.11*	*1*	*0*
758	*19:20*	*XJ3097*	*-10*	*0.00*	*XJ3097*	*-10*	*0.00*	*0.00*	*0*	*0*
759	*19:21*	*XJ3580*	*-9*	*0.00*	*XJ3580*	*-9*	*0.00*	*0.00*	*0*	*0*
End of a bank										
916	*21:50*	*XJ3326*	*-10*	*0.00*	*XJ3326*	*-10*	*0.00*	*0.00*	*0*	*0*
917	*21:50*	*XJ3362*	*-1*	*0.00*	*XJ3362*	*-1*	*0.00*	*0.00*	*0*	*0*
918	*21:50*	*NW749*	*2*	*72.00*	*NW749*	*2*	*72.00*	*72.00*	*0*	*0*
919	*21:53*	*NW1571*	*-2*	*288.00*	*NW1721*	*0*	*22.62*	*22.62*	*1*	*0*
920	*21:53*	*NW1721*	*0*	*22.62*	*XJ3561*	*4*	*24.84*	*58.17*	*1*	*0*
921	*21:54*	*NW443*	*-11*	*0.00*	*NW231*	*5*	*30.16*	*61.27*	*1*	*0*
922	*21:55*	*XJ3347*	*4*	*0.00*	*XJ3588*	*27*	*0.00*	*4.17*	*1*	*0*
923	*21:56*	*XJ3588*	*28*	*4.14*	*NW1571*	*1*	*288.00*	*282.56*	*1*	*0*
924	*21:58*	*XJ3561*	*9*	*45.54*	*NW443*	*-7*	*0.00*	*-7.00*	*1*	*0*
925	*21:59*	*NW231*	*10*	*193.01*	*XJ3347*	*8*	*0.00*	*-3.00*	*1*	*0*
926	*22:00*	*JR10338*	*-19*	*0.00*	*NW1441*	*20*	*59.06*	*121.28*	*1*	*1*
927	*22:01*	*JR10339*	*-19*	*0.00*	*NW990*	*6*	*0.00*	*40.00*	*1*	*1*
928	*22:02*	*JR10340*	*-19*	*0.00*	*NW1189*	*14*	*1,834.56*	*1,892.89*	*1*	*1*
929	*22:03*	*JR10341*	*-19*	*0.00*	*NW1754*	*123*	*5,409.07*	*5,471.29*	*1*	*1*

930	*22:06*	*NW990*	*11*	*14.00*	*NW339*	*13*	*0.00*	*73.33*	*1*	*0*
931	*22:07*	*JR10342*	*-19*	*0.00*	*NW1410*	*5*	*98.62*	*168.62*	*1*	*1*
932	*22:08*	*JR10343*	*-19*	*0.00*	*NW694*	*5*	*73.00*	*111.89*	*1*	*1*
933	*22:09*	*NW1441*	*29*	*247.88*	*JR10338*	*-10*	*0.00*	*0.00*	*1*	*0*
934	*22:10*	*JR10344*	*-19*	*0.00*	*JR10339*	*-10*	*0.00*	*0.00*	*1*	*1*
935	*22:10*	*NW1189*	*22*	*2,133.69*	*JR10340*	*-11*	*0.00*	*0.00*	*1*	*0*
936	*22:11*	*JR10345*	*-19*	*0.00*	*JR10341*	*-11*	*0.00*	*0.00*	*1*	*1*
937	*22:12*	*NW1754*	*132*	*5,945.72*	*JR10342*	*-14*	*0.00*	*0.00*	*1*	*0*
938	*22:13*	*NW339*	*20*	*120.26*	*JR10343*	*-14*	*0.00*	*0.00*	*1*	*0*
939	*22:14*	*NW694*	*11*	*82.00*	*JR10344*	*-15*	*0.00*	*0.00*	*1*	*0*
940	*22:15*	*JR10346*	*-19*	*0.00*	*JR10345*	*-15*	*0.00*	*0.00*	*1*	*1*
941	*22:16*	*JR10347*	*-19*	*0.00*	*JR10346*	*-18*	*0.00*	*0.00*	*1*	*1*
942	*22:17*	*JR10348*	*-19*	*0.00*	*JR10347*	*-18*	*0.00*	*0.00*	*1*	*1*

The first set of flights is from the start of a bank. The empty slots are shaded in gray and are originally occupied by JR flight numbers. The usage of the extra slots is indicated in the last two columns of the table. Without the extra slots, late flights at the start of a bank have no earlier slots available. Adding these extra slots at the start of the bank enabled a large savings for NW330 and NW1417 that were actually scheduled to arrive in an earlier bank. At the end of a bank, late flights "straggle in," and extra slots among these stragglers enable flights such as NW1754 to move forward from slot 937to slot 929. Extra slots at the end of a bank also enable early flights to move back (beyond the bank end) to make earlier slots available for late NWA flights.

The Flight Leg Viewpoint

The slot viewpoint displays *how* savings can occur by swapping slots, but not *why*. To understand where costs and savings occur, we must examine the effect of a swap upon the flight involved.

Consider NW560 landing at MSP at 9:40 am. Although it only lands 3 minutes late, it is 20 minutes late by the time it reaches the gate and kicks off a delay chain:

> NW561 MSP to DEN uses equipment from NW560
>
> NW580 DEN to MSP uses equipment and crew from NW561
>
> (only "ripple delay" costs past here)
>
> NW580 MSP to MCO uses equipment from NW580
>
> NW577 from MCO to MSP uses equipment and flight attendants from NW580

This chain can be represented graphically in Figure 6. In this plot, the actual, scheduled, and modeled block times for each leg are displayed. It appears that NW561 was scheduled to depart before its equipment was scheduled to arrive. This is an obvious indicator that there was a change in the original equipment schedule. The equipment for NW561 didn't arrive and NW560's equipment was substituted. So, the whole cost of delay for flight 561 is not attributable to 560's late arrival. Nonetheless, if flight 560 were to arrive earlier, 561 could leave earlier and save costs for the airline, so the savings due to a move up will be correct. Figure 6 also shows the situation after NW560 was allowed to land 9 minutes earlier. The four downline flights are each delayed 9 minutes less for a time savings of 45 minutes.

In our model, landings at hubs are treated as independent events. In this example for NW560, although NW580 from DEN to MSP was moved up 9 minutes because of the resequencing of NW560, when we resequence the NW580 arrival bank later in the day, we treat NW580 as arriving at its actual *on* time.

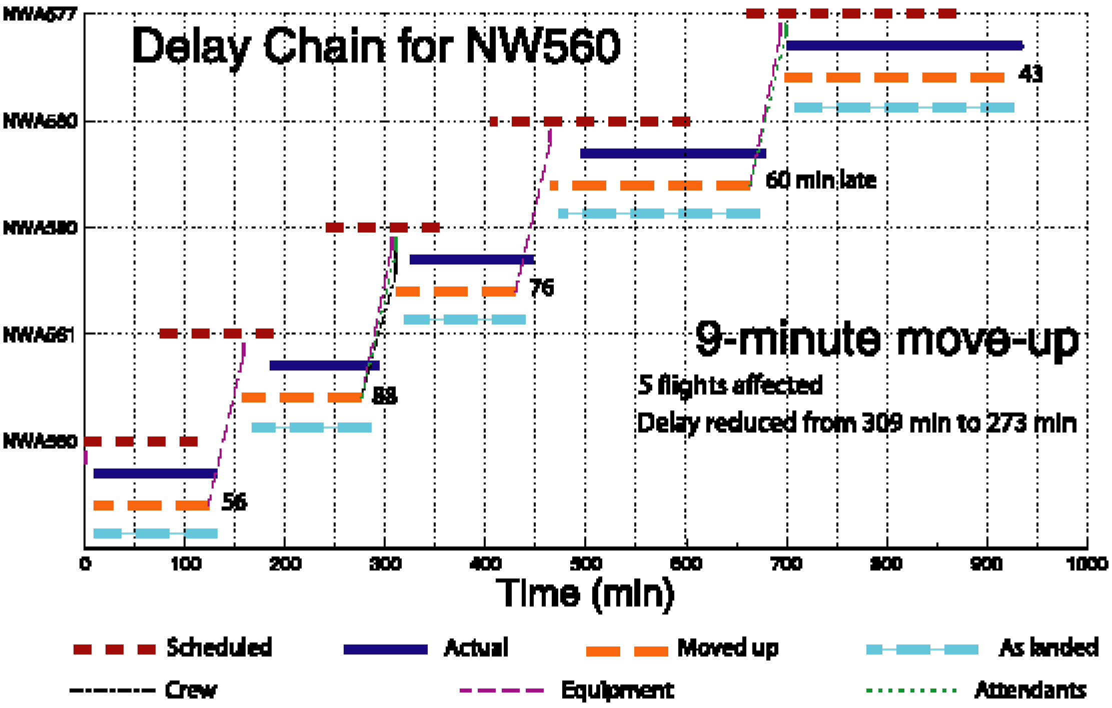

Figure 6. Delay chain before and after resequencing for flight NW560.

Table 6. Costs accumulated in the NW560 delay chain.

Flight Leg	**Cost item**	**Original**	**Resequenced**
NW 561 from MSP to DEN delayed by NW560 equipment available 97 min. late	Ill will terminating passenger	5328.62	4726.54
NW 580 from DEN to MSP delayed by NW561 equipment available 88 min. late NW561 cockpit available 93 min. late NW561 attendants available 93 min. late	Ill will terminating passenger Passenger connect cost	2165.46 5385.84	2057.34 5149.78
-------- End of non-ripple delay costs ------			
NW 580 from MSP to MCO delayed by NW580 equipment available 78 min. late	Ill will terminating passenger	2959.38	2408.34
NW 577 from MCO to MSP Delayed by NW580 equipment available 59 min. late	Ill will terminating passenger Passenger connect cost (The connect cost went up because luggage costs do not increase monotonically.)	3289.44 233.76	2259.04 236.16

Table 6 shows the costs that were attributable to each of the legs in this delay chain. The passenger connect costs have three components: an interrupted trip expense cost (ITE) for long delays, a luggage misconnect cost which increases, then decreases, and finally increases again, and a passenger misconnection cost. In this model, we assume that passengers on legs that leave a hub terminate. If they arrive late, they incur an ill-will cost for terminating passengers. Because we apply resequencing a bank at a time (rather than for the whole schedule), we terminate cost benefits from the resequencing after the NW580 leg. Terminating any cost savings when a plane arrives at the next hub should prevent double counting of the cost benefits that we claim for resequencing.

The downline costs are added to the costs for the passengers on flight 560 to obtain the total costs as follows:

Table 7. Summary for NW560.

Item	**Pre Swap**	**Post Swap**
Slot number	101	94
Local slot time	9:50 am	9:41 am
Flight number	NW560	
Passengers	143	
Scheduled *in* time	9:52 am	
Minutes late	20	11
Block delay $	45.54	8.28
Passenger delay $	0.00	0.00
Passenger ill will	115.25	0.00
Downline effects		
Delay $	12,879.92	11,933.66
Ripple delay $	19,362.50	16,837.20
Fuel costs for resequencing	0.00	66.67
Total cost	13,040.71	12,008.61
Savings		1,032.10
Delay minutes	182	164
Ripple delay minutes	309	273
DOT A14%	Late	On time

For NW560, most of the savings were due to downline effects. The only direct costs attributable to this flight were a block time overage for the crew, and a small passenger ill-will cost. After resequencing, this flight had the further advantage of being less than 15 minutes late and therefore improved the DOT A14% (the percentage of flights arriving within 15 minutes of scheduled in time) on-time performance. However, without utilizing extra slots, other flights are delayed, so the overall improvement in the DOT14 percentage is small.

RESULTS OF THE ANALYSIS

The model was run to determine the baseline cost of operation for the day, before any resequencing was applied. Then, our (close to) optimum-resequencing algorithm was applied and an "optimum" arrival sequence was determined. The savings in delay minutes and actual dollar cost were computed for each flight and then summed for the day. In the case of resequencing, the optimization algorithm was applied using only NWA slots. In a separate run, NWA flights were permitted to swap with added empty slots at DTW. Because it was not feasible to determine empty slot availability for all days of the study, we extrapolated the results from 3 days in DTW.

Based upon the rationale described in previous sections, the results of the study were annualized for NWA operations at the three hubs studied. Subsequently, the results were extrapolated to other airline hub operations based upon DOT delay figures.

Financial summary of resequencing

We looked at the dependency of the savings achieved and the number of flights that were resequenced as a function of the minimum resequencing cost (MRC). The results are shown in Figure 7. The maximum savings seem to be achieved for an MRC of $100. The reason for this is that it does not pay to swap flights with a small cost saving because of the fuel costs for speedup and slowdown . On the other hand, the savings do not decrease very much as the MRC is increased to $300, but the number of resequenced flights decrease by nearly a factor of two (from about 300 to 150 flights). For the purposes of this study, we picked an MRC of $100 to achieve the maximum cost savings, but it would be interesting to repeat the extrapolations for the $300 case.

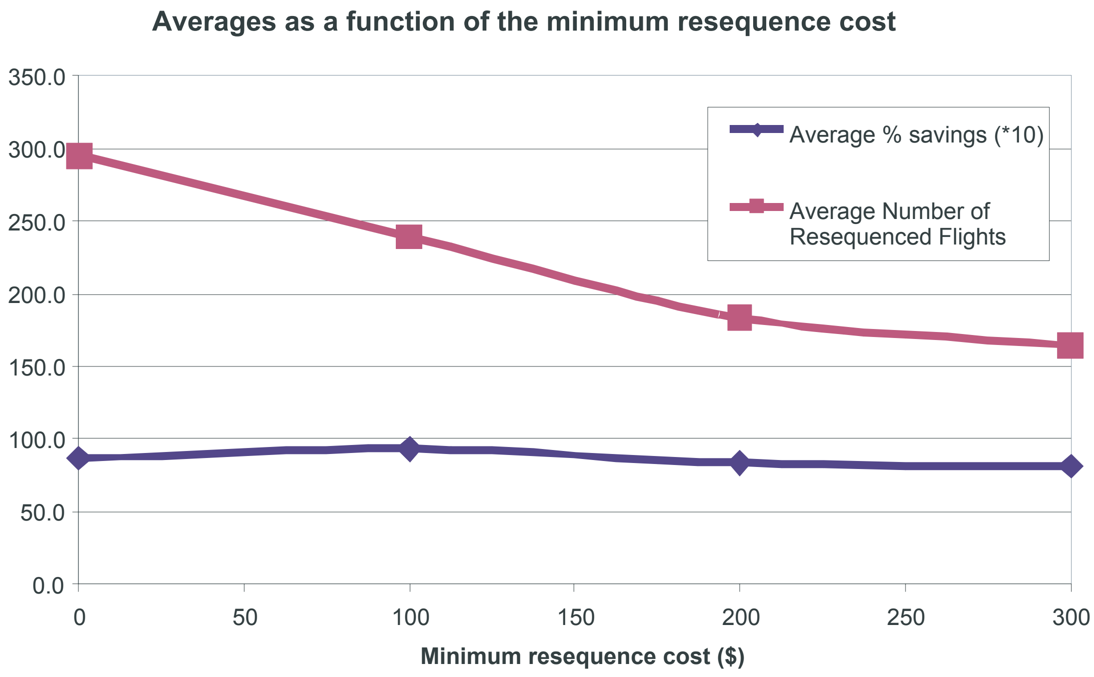

Figure 7. Minimum resequencing cost dependence averaged over all NWA hubs and all days.

The model run for an MRC of $100 produced the results shown in Table 8 for cases where only NWA and NWA commuter slots were used for resequencing. The entries are ordered according to their pre-swap DOT A14% .

Table 8. Resequencing savings for NWA and its commuters if no extra slots are used.

Date	Total Passengers	NWA Flights	Commuter Flights	Pre Swap A14%	Savings	Post Swap A14%	Reseq. Flights	Reseq. Passengers	Savings Per Flight
Detroit									
8/4/1999	30558	354	150	62.4	$42,941.44	67.5	343	21157	$85.20
7/31/1999	33667	294	131	64.6	$36,248.06	66.7	283	22270	$85.29
4/8/1998	33310	358	126	69.8	$28,190.19	72.6	385	27384	$58.24
7/30/1999	42753	363	172	82.6	$30,182.78	85.1	285	24981	$56.42
8/1/1999	42673	330	140	83	$27,444.17	87	324	29084	$58.39
4/7/1998	31201	356	127	84.3	$9,538.84	88.8	307	21063	$19.75
8/5/1999	30230	369	169	84.6	$14,209.63	86.4	310	16904	$26.41
8/2/1999	45580	364	165	90.7	$10,943.97	92.6	293	23476	$20.69
8/3/1999	31039	368	169	91	$7,155.47	93.2	259	14274	$13.32
4/6/1998	33297	354	128	92.4	$5,648.43	93.8	252	19506	$11.72
Memphis									
4/8/1998	10620	110	98	56.3	$13,995.20	57.2	140	7475	$67.28
7/31/1999	11558	104	87	79.1	$2,652.16	81.2	67	4352	$13.89
8/5/1999	13684	110	92	80.2	$6,430.75	81.7	108	8022	$31.84
7/30/1999	13923	109	96	86.8	$2,629.61	87.8	61	4220	$12.83
8/4/1999	13029	110	92	88.6	$1,584.50	87.6	59	3784	$7.84
8/2/1999	13996	110	95	88.8	$1,631.73	91.2	52	3839	$7.96
4/7/1998	9459	112	98	90.5	$1,269.85	91.4	107	5489	$6.05
8/3/1999	12324	111	94	91.2	$930.48	92.7	62	4154	$4.54
8/1/1999	13615	108	85	94.3	$533.69	92.2	20	1879	$2.77
4/6/1998	11011	114	98	94.3	$497.83	94.8	45	2958	$2.35
Minneapolis									
7/30/1999	41682	333	121	60.4	$53,255.69	62.2	326	32480	$117.30
7/31/1999	35711	325	125	74.8	$44,444.47	76.6	297	24514	$98.77
4/6/1998	29425	311	141	79.7	$25,981.49	81	342	21683	$57.48
4/8/1998	33119	315	145	79.7	$23,614.97	85.4	339	24552	$51.34
4/7/1998	28506	316	142	81.6	$19,013.58	83.5	324	20662	$41.51
8/1/1999	42698	350	138	85.4	$27,865.47	88.9	261	24388	$57.10
8/4/1999	46348	362	150	86.7	$23,264.44	90.1	341	32361	$45.44
8/3/1999	40149	360	155	87.8	$22,857.69	90.8	344	27950	$44.38
8/2/1999	44889	364	152	89.6	$26,362.17	91.8	372	33899	$51.09
8/5/1999	47909	363	153	89.8	$28,394.92	91.5	290	27953	$55.03

Using CODAS data, the total number of operations at DTW was plotted against the actual landing time for NWA and its commuter flights. The total arrival rate was then compared with the weather conditions and the AAR from the Air Traffic Control System Command Center (ATCSCC) logs to identify where additional landing slots were available. The additional landing slots were then injected into the model for use within the resequencing algorithm as described previously. Open slots were added into the model for the days of July 30, 31, and August 1, 1999 as shown in Table 9.

Table 9. Savings when extra slots were used at DTW.

Date	Total Passengers	NWA Flights	Commuter Flights	Pre Swap A14 %	Savings	Post Swap A14 %	Reseq. Flights	Reseq. Passengers	Savings Per Flight
7/31/1999	33667	294	131	64.6	$75,774.31	68.4	300	23548	$178.29
8/1/1999	42673	330	140	83	$39,902.78	87.3	325	29470	$84.90
7/30/1999	42753	363	172	82.6	$44,347.53	86.2	283	24900	$82.89

The tables above rank each day of operation by the pre-resequencing A14% arrival performance and compute an average saving per flight operation for each of the hub airports. We note here that in the model, we exclude cancelled flights, which are included in the usual DOT On-Time Percentage figures, hence the A14% notation.

We originally plotted the savings per flight versus the A14% On-Time data and noticed that the data were fit well by a straight line at each hub. However, the Memphis result differed from the MSP and DTW results. By analysis, we determined that there is also a dependence on the number of flights at a hub. Multiplying by the number of flights that day to the –0.8 power brought the three straight lines closer together. This value was normalized by the number of flights at DTW on a nominal day (530). The results of this analysis are shown in Figure 8.

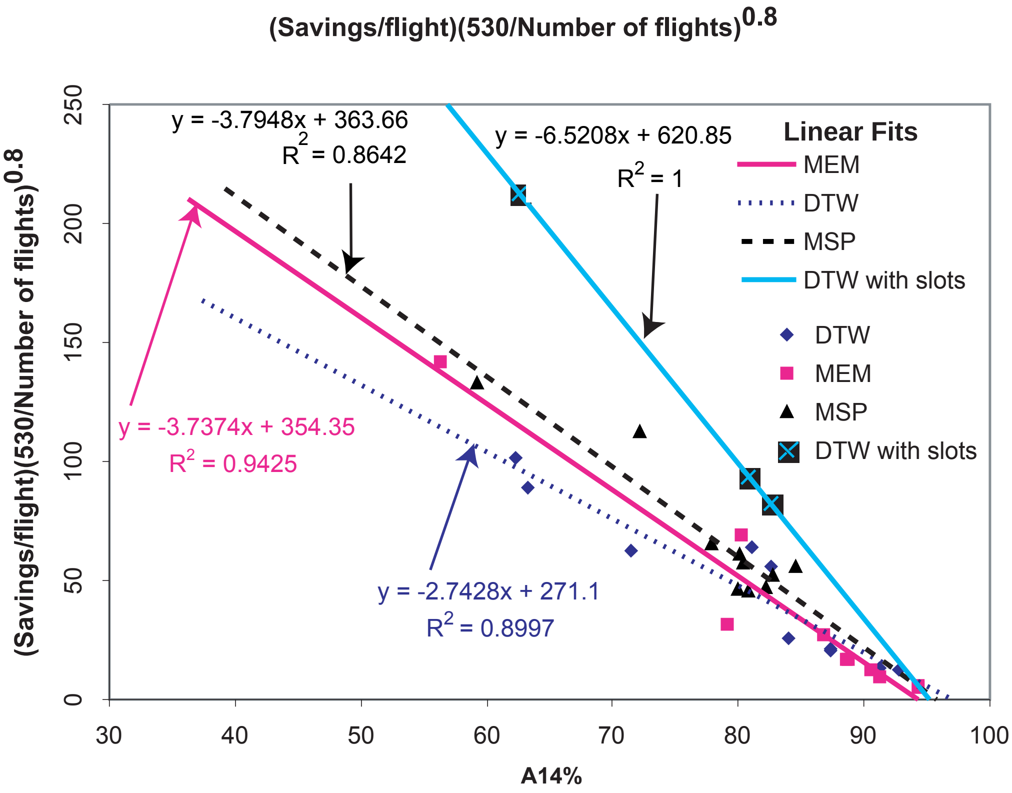

Figure 8. The fit for the savings at each of the three NWA hubs. The results are also shown when extra slots are used at DTW.

From this analysis, we decided to use the MSP results for extrapolation when no extra slots are available. Accordingly, we can predict the savings/flight by the following formulas:

$$\text{savings/flight} = \left(\frac{\#\text{flights}}{530}\right)^{0.8}\left(354.95 - 3.74 * \text{A14\%}\right)$$

for the case in which only NWA – commuter swaps were used, and

$$\text{savings/flight} = \left(\frac{\#\text{flights}}{530}\right)^{0.8}\left(620.85 - 6.52 * \text{A14\%}\right)$$

for the case in which all open slots were used. When the DOT On-Time percentage is above about 95%, no savings can be achieved with the $100 MRC, so the negative values for savings/flight predicted by these formulas do not occur.

Plotting the predicted savings/flight versus the actual savings/flight without extra slots can assess the goodness of this prediction. (The formula fits the cases with extra slots almost exactly.) The results of this analysis are shown in Figure 9.

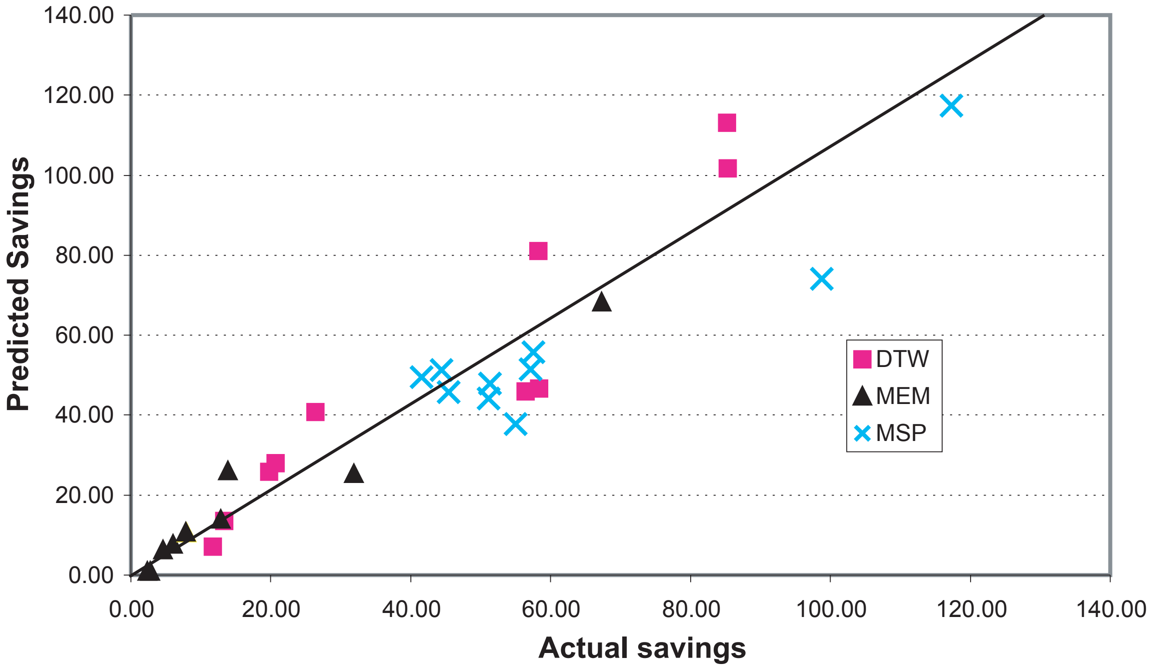

Figure 9. A test of the formula for predicted savings per flight.

We used these formulas and the DOT Statistics for NWA Hubs for the period 01/01/98 through 11/20/99 to extrapolate our results to a longer operational period. The results for Detroit are plotted in Figure 10, which shows the A14% and the savings per flight with and without extra slots.

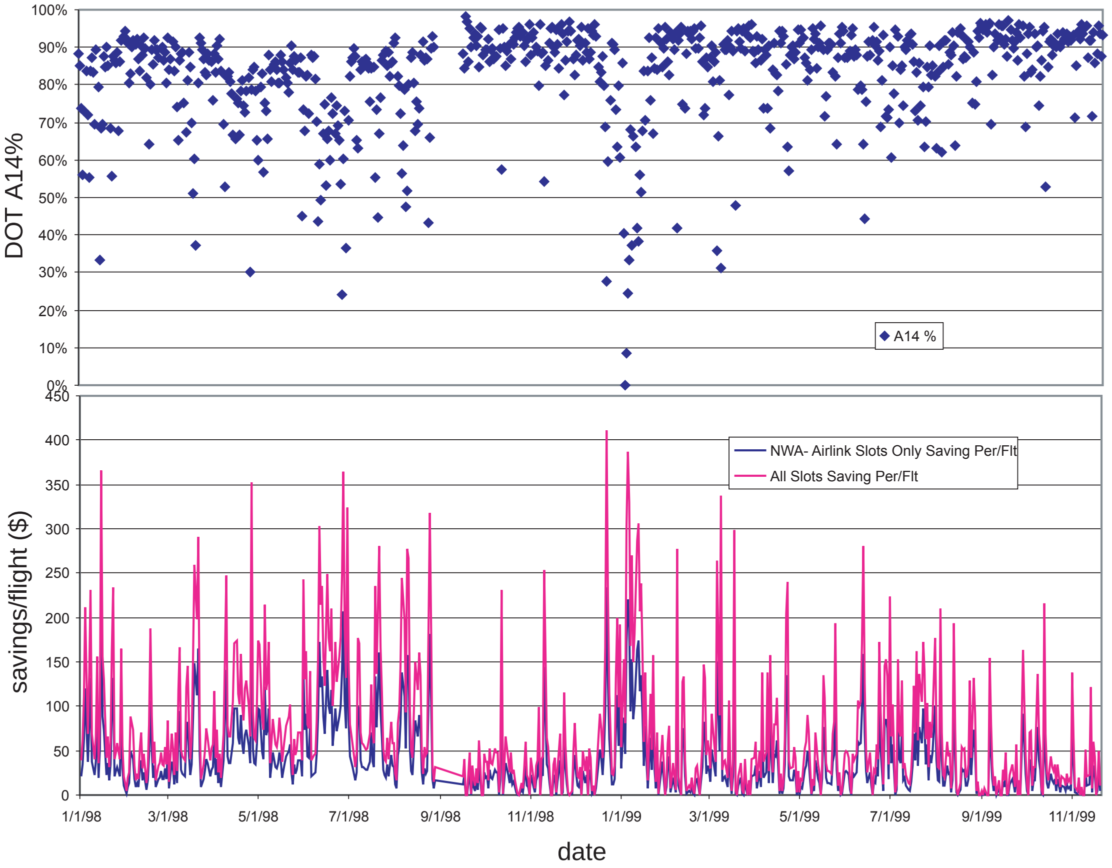

Figure 10. A14% (top) and the savings per flight (bottom) for NWA at Detroit.

Using these data for the three NWA hubs for this period of almost 2 years, we add up the savings achievable by NWA for each hub. Then we calculate the daily total, and annualize the results for the cases where only NWA and Airlink slots are used for swaps and when extra free slots are used. The results are shown in Table 10.

Table 10. Summary of projected savings for NWA hubs.

	NWA + Airlink Slots Only			Extra Unused Slots		
	Period of Figure 10	Daily	Annual	Period of Figure 10	Daily	Annual
DTW	$ 12,657,373	$ 18,892	**$ 6,895,434**	$ 22,821,259	$34,062	**$ 12,432,477**
MEM	$ 1,722,540	$ 2,571	**$ 938,399**	$ 3,136,150	$4,681	**$ 1,708,499**
MSP	$ 7,846,138	$ 11,711	**$ 4,274,389**	$ 14,148,033	$21,116	**$ 7,707,510**
Total	**$ 22,226,051**		**$ 12,108,222**			**$ 21,848,487**

Using the DOT On-Time Performance Reports for January through September 1999, airline operations at the major hub airports were analyzed. At medium density airports the formula for both the "intra-airline slot swapping" and "all open slot" values were applied to project potential savings, given an assumption that the NWA cost model generally applies to these operations. At high density hubs (ATL, DFW & ORD) only the "intra airline swapping" formula was applied as an airport-specific determination of how many additional landing slots might exist at these airports during peak arrival time would need to be undertaken before any credit for additional slots could be taken. The results for all these airline hub operations can be summarized to estimate an industry-wide annualized savings, and are shown in Table 11.

Table 11. Summary of Annualized Savings.

		Savings	Savings
Airport	**Airline**	**Airline only Slots**	**All Open Slots**
ATL*	**DAL**	$ 15,855,173	$ 15,855,173
CLT	**USA**	$ 3,610,138	$ 6,474,808
CVG	**DAL**	$ 1,471,161	$ 2,652,239
DEN	**UAL**	$ 2,740,567	$ 4,937,415
DFW*	**AAL**	$ 6,693,129	$ 6,693,129
DFW*	**DAL**	$ 436,946	$ 436,946
EWR	**COA**	$ 2,650,660	$ 4,707,271
IAH	**COA**	$ 2,458,417	$ 4,428,237
LAX	**UAL**	$ 1,518,791	$ 2,708,316
ORD*	**AAL**	$ 4,571,899	$ 4,571,899
ORD*	**UAL**	$ 7,671,878	$ 7,671,878
PHL	**USA**	$ 2,993,007	$ 5,308,449
PHX	**AWA**	$ 2,122,502	$ 3,784,348
PIT	**USA**	$ 2,391,137	$ 4,283,630
SFO	**UAL**	$ 3,112,731	$ 5,529,491
SLC	**DAL**	$ 609,090	$ 1,111,329
STL	**TWA**	$ 2,830,283	$ 5,142,910
DTW	**NWA**	$ 6,895,434	$ 12,432,477
MEM	**NWA**	$ 938,399	$ 1,708,499
MSP	**NWA**	$ 4,274,389	$ 7,707,510
	Total	**$ 75,845,729**	**$ 108,145,953**

*Intra-airline swapping only

Operational summary of resequencing

The model results were analyzed to determine if any patterns existed that would allow the CTAS to resequence flights based upon some predetermined criteria. The following discusses potential criteria:

Sequence by Aircraft Type

The model results for cost savings due to aircraft resequencing do not indicate any correlation between resequenced aircraft and aircraft type. There are instances when it is advantageous to move a larger aircraft, such as an A-320, back to give way to a DC-9 that is operating late. In other cases, swaps between DC-9s are very cost effective.

In some cases, it was observed that swapping a commuter flight for a larger jet was often cost-effective; however, more analysis on this result is required. The model discriminates against commuter aircraft because we have no crew rotations, and only tail numbers for the ARJs. Therefore, there are almost no downline delay chains, and the number of passengers (and hence their delay costs) is small. The results may also be anomalous due to the way that NWA schedules commuter aircraft. These results are not conclusive.

Sequence Aircraft by Schedule

The cost savings that were determined from the analysis of the model generally resulted from moving up a late operating flight to a slot occupied by an early flight. Thus, a rule that resequences flights such that the operation approaches its original schedule might be a general rule that could produce cost savings. It is certainly in accord with the airline's connectivity plans. In this case, CTAS would need to know how each flight is operating relative to its original schedule, not relative to an amended flight plan ETA. This could be accomplished by accessing the original arrival schedule intent from ETMS data and having CTAS create a sequence that brings the fleet closest to that original arrival schedule.

As the resequenced data were analyzed, we found that while a general movement of aircraft toward their original schedule often resulted in an optimum sequence, there is no obvious simple rule that achieves this result. In some cases, we observed flights that departed early and arrived early, but had positive costs due to exceeding the flight crew’s block time. If these flights were moved closer to their original schedules, the costs of these flights would actually increase. If no extra slots are available, for each flight that is moved forward, another flight must move backward; often, an on-time flight can be delayed with no cost penalty to achieve this swap. Passenger misconnection and ill will heavily influence cost savings. Therefore, a flight with many potential misconnecting passengers would have the potential to save more delay costs than a flight where all passengers terminate at the hub, regardless of which flight was operating closer to the original schedule.

It follows that a key ingredient to implementing resequencing is the ability to calculate the cost for each flight to land in each available landing slot (including downline effects). If this is known, correct resequencing decisions can be made. Operationally, airlines have all the required data, but the passenger component is not yet integrated into the flight operations data stream. Because most of the costs are attributable to passenger ill will and misconnection costs, the ability to calculate the “worth” of a given flight in real time is currently unavailable.

Interaction with Airline Operational Centers (AOCs)

The cost savings presented here resulted from detailed analysis of the data and an optimized resequencing algorithm. The implementation of this methodology represents a real-time collaboration between an AOC and an ATSP; however, decision support tools to achieve this collaboration in real-time do not exist and are not under development at this time. In order to create an operational environment that emulates the current analysis, interactive communication between AOCs and ATSPs is required.

The determination of optimum cost savings appears to be a function of many variables. These variables are ultimately known to the airline; however, all the information to determine the lowest delay cost sequence is not presently available in real-time to the airline. It is clear that the airline business rules and cost structures are the most significant element in the cost equation, and the delay costs are strongly influenced by the amount of operational delay for each aircraft. At present, CTAS does not have any information to determine the optimum sequence, and a massive upload of AOC operational information, including highly proprietary airline business rules would be required for CTAS to perform these computations. A more natural distribution of this functionality would be for CTAS to present the AOC with updated arrival sequences, including current sequence and projected *on* times, and allow the AOC to determine the optimum situation based upon internal information and the CTAS-provided sequence. This interchange would continue as the real-life situation evolved and the airline AOC tools would react to the sequence reality. In this scenario, CTAS would need the following additional capabilities:

- Accept the AOC's desired arrival sequence,
- Determine the optimum method for ATC implementation of the desired sequence, and
- Integrate, arbitrate, and resolve conflicts between the requests from multiple AOCs at a single hub.

These conclusions are based upon the study of a single airline operation over a 10-day window within a two-year period. It is highly likely that other airline operations, such as Federal Express at MEM, will have considerably different definition for what is optimum in their arrival sequence. These additional rules may provide for an interaction between CTAS and an airport surface management system, allowing the opportunity to develop a complete arrival management system.

In addition to CTAS, MITRE's Self-Managed Arrival Sequencing Tool [6] and Sabre's Dependability Predictor Model [7] provide the tools that allow airlines to influence their arrival queues. The results from this study should provide airlines the mechanism to correctly assess the "worth" of the planes in their arrival queues so that they can use such tools to maximum benefit.

CONCLUSIONS

The summary of bank operations was derived from NWA hub operations during the ten-day assessment period. Our model of NWA operations and costs was applied to determine the cost differences that can be obtained by swapping arrival slots, either among just NWA flights or by also using open slots. Annual savings resulting from preferential arrival ordering for NWA operations at DTW, MSP, and MEM range from approximately $12 million when only NWA slots are used, to $22 million, when empty arrival slots are available and utilized for resequencing. These results assume that the 10 days studied were a representative sample, and are subject to the caveats mentioned in this paper. If we extrapolate to the major airline's hubs, the annual savings that could be achieved with only intra-airline swaps would be over $75 million, and over $100 million if other available slots can be used.

Further conclusions are as follows:

- Ripple delay costs are often dominant. Therefore, the airline connectivity must be known and the downline costs evaluated, to determine the effect of a slot change.
- There are significantly fewer delay chains when the crew turn time is reduced from 45 minutes to the 35 minutes used in this study.
- We have developed a method for near-optimum resequencing that works on a bank-by-bank basis.
- There is a tradeoff between the cost savings and the number of resequenced flights. The cost savings is optimized with an MRC of $100, but the number of flights is almost halved, with just a little less savings, if the MRC is $300.
- Moving late flights forward certainly yields savings. But, there may be no cost penalty for making an on-time flight a little later to create a landing slot.
- In general, the cost savings increase linearly with decreased DOT A14%.
- The savings seem to also scale as the number of flights to the 0.8 power.
- Resequencing has little effect on the overall A14% if extra slots are not added.
- Adding extra (available) slots almost doubles the savings.
- The extra slots are used to move late flights at the start of a bank forward or to find a slot for stragglers at the end of a bank. Extra slots would help greatly in the middle of a bank but few are available.
- In general, moving late flights forward towards their original schedule reduces costs. However, for each flight that moves forwards, another must move backwards unless extra unused slots are available.

Acknowledgements

The authors wish to thank Nageswara Rao of Oak Ridge National Laboratory for his helpful analysis of the resequencing algorithm. Andrew W. Danz, Jay A. Friedman , and Todd S. Ekstrom from Northwest Airlines provided the airline data used in this study, answered questions about it, and provided general advice and insight into Northwest's operations.

Acronyms

A14%	daily percentage of flights that arrive less than 15 minutes late (excluding cancellations)
AAR	Airport arrival (or acceptance) rate
AOC	Airline operational center
ARJ	Advanced regional jet
ATCSCC	Air Traffic Control System Command Center
ATSP	Air traffic service providers
CODAS	Consolidated Operations and Delay Analysis Systems
CTAS	Center Tracon Automation System
DTW	Detroit Metropolitan Wayne County Airport
ECT	Effective connection time
ETA	Estimated time of arrival
ETMS	Enhanced Traffic Management System
FAA	Federal Aviation Administration
FL	Flight level. (Aircraft fly with altimeters set to the ISA Standard Sea Level Pressure setting of 29.92 inches or 1013.2 HP)
FL100	Flight level in hundreds of feet. 100 = 10,000 feet pressure level.
GMT	Greenwich mean time (ZULU)
ITE	Interrupted trip expense cost
KCAS	Knot calibrated airspeed (this is indicated airspeed corrected for probe position, compression effects and indicating error)
KCAS250	250 knots calibrated airspeed In the US (and most other countries) aircraft are required to limit speed to 250 knots or slower below 10,000 feet
MEM	Memphis International Airport
MRC	Minimum resequencing cost
MSP	Minneapolis-St. Paul International Airport
NASA	National Aeronautics and Space Administration
NWA	Northwest Airlines
OAG	Official Airline Guide
OOOI	Out-Off-On-In
Vth	V-speed threshold. This is the target runway threshold crossing speed for the aircraft at its landing weight
Vth+10	This is the typical "final approach target speed" (sometimes referred to as Vprog). For a large transport category aircraft, this represents the Vth speed landing flaps + 10 knots KCAS.

Biographies

James A. Rome

Jim Rome received all his degrees in Electrical Engineering from the Massachusetts Institute of Technology. After he graduated with the Sc.D degree in 1971, he joined Oak Ridge National Laboratory. He spent 25 years doing theoretical Plasma Physics and is a Fellow of the American Physical Society. He has been Editor of the international newsletter *Stellarator News* since its inception 11 years ago. Since 1990, Dr. Rome has been involved with air traffic modeling. He studied a month's worth of data to try to determine the capacity of the airspace in 1991. More recently, he has been modeling the economic impact of allowing airlines to reorder arrival queues at major hubs. Dr. Rome is also working on computer security for a Department of Energy project that allows electron microscopes and beamlines to be operated remotely across the Internet. He is President and co-founder of Scientific Endeavors Corp., producer of the *GraphiC* Scientific Graphics Library (used for many of the plots in this paper). More details of Dr. Rome's current activities are found on his Web page http://www.ornl.gov/~jar

George F. (Rusty) Bell.

Rusty Bell is an aviation professional with over 30 years experience in commercial air operations. He founded GFB & Associates in 1997, which offers a wide range of aviation consulting services in both the science and business of air operations. He started his career in 1970 with Pan American World Airway where he held a variety of dispatcher and management positions serving in Europe, Middle East, Asia, and Africa. Following transfer to headquarters, he was appointed Director of Flight Control where he was responsible for all day-to-day flight operations. In 1991, he joined Delta Air Lines as a senior manager in flight operations where he worked on the development of advanced flight and operations control systems. Experienced in the worldwide theatre of flight operations, he is recognized as an expert in the field of Airline Operations Control and Air Traffic Management. Qualified as an aircraft dispatcher, with formal training as a navigator he is an active general aviation pilot.

James H. Cistone

Jim Cistone has been manager of Advanced Programs for Lockheed Martin Management & Data Systems in King of Prussia, Pennsylvania. In this capacity he lead several tasks in aviation research for NASA. He is Program Manager for Lockheed Martin's Advanced Air Transportation Technologies (AATT) Research Project, which was awarded in November 1997. Cistone's background includes software and systems engineering for government and commercial automation systems, the development of automation systems for aviation and air traffic managament, and program and project management.

Jim is an active commercial and instrument-rated pilot working toward his flight instructor's rating. He was graduated from Muhlenberg College with a B.S. in Physics and received an M.S. in Physics from Drexel University. He also completed the Ph.D. course requirements in Physics and Atmospheric Science at Drexel University.

Cistone has just moved to Rockville, MD where he assumes new duties with Lockheed Martin Air Traffic Management.

William (Bill) S. Leber Jr.

Bill Leber, as a Chief Flight Dispatcher with Northwest Airlines, oversees the Air Traffic Coordinator position at NWA's System Operations Control Center. He was a key developer of NWA's industry-leading strategic planning team and advocate of its extension to the FAA's Air Traffic Control System Command Center. He is currently a qualified Aircraft Dispatcher, and has served as a check airman and an international Chief Flight Dispatcher.

Bill has been a participant in the Collaborative Decision Making efforts since the early 1990's where he was a key contributor to the implementation of the Oakland and Anchorage Center Track Advisory systems. He is presently a participant in several areas of the Air Transport Association's Collaborative Decision Making efforts including Co-Chairmanship of the Collaborative Routing – Long-term Working Group He has been a member of NASA's ATM ESC since 1999

He has 10 years of air traffic management experience coordinating with FAA, JCAB, Euro-control and other air traffic service providers in Europe, Russia and the Pacific Basin. He has been with Northwest Airlines since 1983. Before joining Northwest Airlines he was employed by McDonnel Douglas Astronautics division (now Boeing) in St. Louis as an Engineering Planner.

Bill is the former President and Co-founder of the Airline Dispatchers Federation a non-union professional association. He earned his B.S. in Aeronautical Administration at Parks College of St. Louis University. He holds both an Aircraft Dispatcher and Private Pilot certificate..

Simon D. Rose

Mr. Rose is a technical project manager for Oak Ridge National Laboratory and is involved in transportation and safety initiatives that involve several federal agencies including the Federal Aviation Administration (FAA) and National Aeronautics and Space Administration (NASA). Mr. Rose has a B.S. Degree in Physics from the University of Kent, England and joined ORNL in 1980. He has extensive experience in system safety analyses, modeling and simulation and information systems development.

Mr. Rose has a wide experience within the Department of Transportation and several Department of Defense programs. Work including safety information systems work in support of the U.S. Coast Guard and technical analyses and regulatory assessments concerning oil pollution law for pipelines and other transportation related facilities for the Office of Pipeline Safety. His aviation safety and analyses work includes support to the Collaborative Decision Making program, the Global Aviation and Information Network program, and the Advanced Air Transportation Technologies program.

Ronald W. Lee

Ron Lee received a B.S. in Information and Computer Science from Georgia Tech in 1985 and an M.S. in Computer Systems from the Air Force Institute of Technology in 1986. As an intelligence systems staff officer, Mr. Lee was involved in the design and implementation of several intelligence support applications. From 1989 until 1994 he worked in the Data Systems Research and Development Division of Lockheed-Martin Energy Systems on a variety of tasks including military decision support tools and flight message processing and distribution systems for the FAA. He was one of the principle developers of the prototype CONUS data access tool (CONDAT) and participated in a Congressionally-mandated study of U.S. Airspace capacity. Since joining ORNL in 1994, Mr. Lee has architected and implemented distributed, platform-independent, Web-accessible tools and applications supporting the Defense Intelligence Agency and Defense Threat Reduction Agency. He has a broad range of experience in object-oriented and distributed architectures, software engineering, data analysis, and graphic data representation.

[1]. Roger Beatty, Rose Hsu, Lee Berry, and James Rome, *Preliminary Evaluation of Flight Delay Propagation Through an Airline Schedule*, Air Traffic Control Quarterly, **74** (4) 259-270 (1999).

[2]. R. Dear, *Dynamic Scheduling of Aircraft in the Near Terminal Area*, Technical Report R76-9, Flight Transportation Laboratory, M.I.T. (September 1976).

[3]. H. Psaraftis, *A Dynamic Programming Approach to the Aircraft Sequencing Problem*, Technical Report R78-4, Flight Transportation Laboratory, M.I.T. (October 1978).

[4]. C. S. Venkatakrishnan, Arnold Barnett, and Amedeo R. Odoni, *Landings at Logan Airport: Describing and Increasing Airport Capacity*, Transportation Science, **27** (3) 211-227 (August 1993).

[5]. Gregory C. Carr, Heinz Erzberger, and Frank Neuman, *Airline Arrival Prioritization in Sequencing and Scheduling*, in 2nd USA/Europe Air Traffic Management R&D Seminar, Orlando FL (December 1988).

[6]. http://www.mitre.org/news/articles_99/msr.html

[7] http://www.sabre.com/airline/airline_planning.html