Disruption Recovery in Air Traffic

Disruption Recovery in Air Traffic:

An Airport Perspective

By

Prabhu Manyem

Cambridge Scholars Publishing

Disruption Recovery in Air Traffic: An Airport Perspective

By Prabhu Manyem

This book first published 2021

Cambridge Scholars Publishing

Lady Stephenson Library, Newcastle upon Tyne, NE6 2PA, UK

British Library Cataloguing in Publication Data
A catalogue record for this book is available from the British Library

ISBN (10): 1-5275-6872-5
ISBN (13): 978-1-5275-6872-3

To the air traffic controllers who keep our planes safe in the sky

Contents

List of Figures

List of Tables

Foreword

I first became involved with Prabhu and his co-workers in my capacity as Operations Support Manager in the Air Traffic Control facility at Sydney Airport during the late nineties. I was assigned to assist them in developing familiarity with the Air Traffic Management environment. My interest in their area of research was particularly sparked by my recent experience as an Air Traffic Flow Controller and several years spent developing a software-based ATC flow control information system.

With this experience I had reached the conclusion that in a tactical sense at least, air traffic flow management is frequently if not always a constant process of disruption recovery. The nature and range of the disrupting influences (particularly those relating to changes in airport or network capacity) is discussed in some detail in this book.

This book summarises a significant body of work which, it could be said, seeks to arrive at a “general theory" of air traffic disruption management. In this case a series of objective functions and their constraints which could be applied quickly and efficiently to a disrupted air traffic schedule to produce a new schedule that maximises the use of the changed system capacity while fairly distributing delay costs.

I will confess to a fair degree of cynicism about the concept when first I became involved. This was based on my 25 years plus experience as an air traffic controller. Nevertheless, my more recent studies in industrial mathematics kept me very interested in the research and its progress. At the time I saw practical difficulties in the application of revised schedules. These would be due in no small part to the basic rules of the air as set out in the annexes to the UN conventions on Civil Aviation to which Australia and all of the countries with airspace adjoining Australia’s, are signatories. These among other things set out basic flight priorities which could be seen to be at odds with several of assumptions being made in the research. Developments in the intervening period have shown that the airline industry and the air traffic service providers are open to change in these areas in the interest of the “common good".

Air traffic management in Australia in the late nineties was in a transition from a state run and self-regulated system to one in which the regulatory function was separated from the service provision (and revenue generating) function. This in turn led to better funding to meet the increasing need for an enhanced level of automation and information management to meet the needs of increasing air traffic demand. This was the period when TAAATS (The Australian Advanced Air Traffic System) was implemented, The CDTP (Controlled Departure Time System) was developed and implemented, Slot management was legislated for Sydney Airport and the Maestro tactical flow management system was implemented. Of these, the Slot Management regime and its compliance monitoring led to some serious discussion about how the established system of priorities might be varied to better accommodate compliant flights over those not complying.

Internationally the concept of “Collaborative Decision Making" (CDM) became a hot topic in Eurocontrol and in the ATM panel of the International Civil Aviation Organisation (ICAO). CDM involves the real time sharing of information between airlines and ATM service providers with the objective of optimising the use of available airport and airspace capacity. The Eurocontrol Central Flow Management Unit (CFMU) and the Australian CDTP are good examples of an evolving use of the principals of CDM.

And so, over the past twenty years, the real world of ATM has evolved in a way which shows that the research described in this book was arguably well ahead of its time.

Newcastle (Australia), *Barney Pinney*
September 2019

Preface

Disruptions to commercial airline schedules are frequent and can inflict significant costs. The volume of air traffic has increased considerably over the past few decades, while the capacity of systems such as airports and airways has not kept pace. Demand exceeds capacity at many key airports. Furthermore, unanticipated security related events are also likely to lead to de facto capacity drops.

The high degree of competition among airlines is also noteworthy. This has led to a widespread adoption of Operations Research methodologies and—as a direct result—to schedules that are highly optimised. An unintended consequence of this is that, when there are disturbances to the flight schedules due to unpredictable circumstances such as bad weather, aircraft malfunction, or security checks, the regular (published) schedules are thrown out of gear, sometimes causing chaos at airports and airline operation centres.

The recovery from these schedule perturbations can be assessed by any or all of a number of criteria, some of which may be in conflict. Flights may need to be delayed, diverted or cancelled, causing inconvenience to passengers and reducing airline profits. The goal is to develop techniques facilitating rapid return to normal operations whenever disruptions occur.

In Chapter 1, we review background knowledge and provide key definitions. Section 1.6 reviews literature from an airport's perspective, as opposed to the vast literature that exists from an airline's point of view.

Navazio and Romanin-Jacur (1998) introduced a basic ground holding model (NRJ). In Chapter 2, we review the NRJ model and extend it to consider cancelling flights and an airport's night curfew restrictions. We report computational results for the extended model.

In Chapter 3, we develop a heuristic based on primal-dual techniques to the disruption recovery model that considers curfew violations for a single airport. We present results associated with its computational testing. Our heuristic produces a solution with approximately the same value as that produced by the Integer Programming solver GLPK, but in much shorter time.

Chapter 4 introduces an adaptive model called MARFE (short for Model to Adaptively Reschedule Flights in Emergencies) which can handle sudden drops in airport capacities. However, when a warning is received, some flights would have already departed, headed for the airport where a capacity drop occurs; such flights may need to be held in the air (*airborne holding*) or diverted to other airports.

Since a linear cost model is too simplistic, in Chapter 5, a new model is developed which takes into account the non-linear nature of the manner in which costs increase with flight delays.

In Chapter 6, we discuss the concept of *fairness* such that all airlines are treated fairly and impartially by airports; not necessarily every day, but over a short period of time, such as a month.

I hope that readers will find the book useful and insightful. You are invited to provide me feedback at prabhu.manyem@gmail.com.

Nanchang, China, *Prabhu Manyem*
March 2021

Acknowledgements

This book wouldn't have been possible without my research collaboration over five years (1999-2004) at the University of South Australia (UniSA) with Jerzy Filar, Kevin White and David Panton. The research carried out during that period serves as a solid foundation for the book.

Valuable domain knowledge in Civil Aviation was obtained from Barney Pinney, Steve Barry and the late David Anderson at Airservices Australia, the Preston Aviation Solutions in Melbourne (now part of the Jeppesen group), Ken Allcott (Sydney Airport Corporation) as well as Andrew Cook and Graham Tanner (University of Westminster, UK).

David Anderson was so kind that he provided us with arrival-departure data at Sydney airport for the year 1999. Barney Pinney has been a constant source of encouragement/motivation in regards to applying optimisation modelling and techniques to the topic.

A grant from the Jiang-Xi Department of Education in China (grant number Grant No. GJJ161113) has been highly valuable. The Yao-Hu Scholar funding from the Nanchang Institute of Technology (China) is very much appreciated.

Research support from the University of South Australia and the Cooperative Research Centre for Sensor Signal and Information Processing (CSSIP) in Adelaide have extensively contributed to the research described in this book.

At Nanchang Institute of Technology, my students Jia-Ting Liu and Shi-Hang Zhang helped me with my research.

I am very grateful to the Sir Ross and Sir Keith Smith Fund (Adelaide, South Australia) which sponsored our project for five years (1999-2004). Preston Aviation Solutions was very generous in allowing us to use their TAAM software for a five-year period from 2000 to 2005.

I would like to thank Adam Rummens and his colleagues at Cambridge Scholars Publishing for their support and patience.

Some of my research was carried out while visiting the Australian National University in Canberra in 2017 and Flinders University in Adelaide in 2019.

Notation

T	set of time periods; $T = \{1, \cdots, t, \cdots, \|T\|\}$
Z	set of airports; $Z = \{1, \cdots, z, \cdots, \|Z\|\}$
F	set of all flights; $F = \{1, \cdots, f, \cdots, \|F\|\}$
F_c	set of flights that violate the curfew.
F_z	set of flights arriving at airport z. For example, F_{syd} = set of flights that arrive at Sydney.
$F_{z,t}$	set of flights scheduled to arrive at airport z in period t
$K_{z,t}$	arrival capacity of airport z for time period t. (arrival capacity is an upper bound on the number of arrivals.)
S_f	set of successors of flight f, $\|S_f\| = n_f$.
r_f	published (or planned) arrival period of flight f ($r_f \in T$).
α_f	flying time of flight f.
Δ_f	delay (in periods) suffered by flight f.
$\Delta_{\max}$	maximum allowed total delay (ground and air combined) for any flight (to be precise, a flight that is not diverted to an airport different from its destination).
A_f	actual arrival period of flight $f = r_f + \Delta_f$
L_f	the last possible arrival period of $f = r_f + \Delta_{\max}$
T_f	set of time intervals in which flight f may land at its destination airport = $[r_f, L_f]$
T_F	set of periods for which the capacity is fully utilised.
c_f	cost of ground delay for flight f per time period.
k_f	cancellation cost of f.
$x_{f,t}$	equal to 1 if f arrives during period t or earlier, 0 otherwise. (Defined differently in Chapter 3. See Page 38.)
w_z^1	first time interval of non-curfew operations at airport z.
w_z^2	last interval of non-curfew operations at airport z. (If the airport name is obvious, we use w_1 and w_2.)

$\phi_{f,g}$	minimum number of time intervals required between arrival of flight f and departure of its successor g, (also known as *f-g turnaround time*)
$\Phi_{f,g}$	the service time between f and g, including the flight time of g, which equals $\phi_{f,g} + \alpha_g$
$\sigma_{f,g}$	slack time between flights f and g, given by $r_g - r_f - \Phi_{f,g}$.
$\mathcal{V}$	cost of violating the curfew per flight.
d_f	1, if flight f violates the curfew, 0 otherwise.
a_f	1, if f violates the morning curfew, 0 otherwise.
b_f	1, if f violates the evening curfew, 0 otherwise.
M	a large positive number.

Chapter 1
Background in Disruption Recovery

Air traffic is considered to be the safest mode of travel. Yet it is highly complex. The available airspace does not change with time. However, with the number of flights increasing steadily, so does the complexity of air traffic management (ATM). Air traffic controllers have a thankless task on their hands.

In this book, we cover *Disruption Recovery* in air traffic management. For instance, bad weather may cause a major airport to suffer significant disruption to its operations. Due to increased separation between aircraft during bad weather, the allowable number of arrivals and departures (for example, in a given 15-minute duration) will reduce. Inevitably flights will be delayed, either on the ground at their origins or in the air enroute. Some flights may need to be cancelled and passengers accommodated in later flights.

To deal with bad weather at flight destinations, apart from (a) delaying flights at their origins (known as *ground holding*), other possible solutions are: (b) holding flights in airborne holding patterns; (b) diversions to other airports; and (c) rerouting flights. Sometimes a mix of these strategies may need to be deployed.

One of the foremost goals in ATM is to develop strategies for the reduction of departure and arrival delays at every airport, preferably using only ground holding and cancellations. Thereby we are able to reduce the need to place incoming planes in airborne *holding patterns*, consequently reducing noise, reducing fuel wastage and enhancing safety.

Disrution Recovery has been studied in other forms of transport, such as Railways (Pender et al. 2013; Fang, Yang, and Yao 2015; Azad, Hassini, and Verma 2016; Ghaemi, Cats, and Goverde 2017).

1.1 Common Good Perspective

Airlines realise that disgruntled passengers may switch to another airline for their next trip. And yet, after giving due consideration to this fact, when airlines optimise disruption recovery, their eventual goal would still be to maximise profits. This also holds true for several airports around the world that have been privatised.

However when a public sector air traffic regulatory authority such as Airservices Australia (the principal air traffic services provider in the country) attempts to optimise disruption recovery, there is a *Common Good Perspective* (CGP) approach. Everyone benefits — all airlines, the flying public, airports and air traffic regulatory agencies. All stakeholders are treated in a fair and equitable manner.

Our belief is that by minimising a suitably constructed cost function of these effects (such as delays, cancellations and diversions), it is possible to significantly reduce recovery costs to all the key participants. CGP has been studied earlier by (Filar et al. 2007).

1.2 Causes of Disruption

The main cause for disruption at airports is bad weather, which affects visibility. In clear weather, aircraft use Visual Flight Rules (VFR) for navigation. In bad weather, they use Instrument Flight Rules (IFR). In bad weather, separation between aircraft should be higher in the airspace near the terminal, and this has an impact on the number of aircraft that can arrive (or depart) within a certain fixed period. This causes a drop in arrival and departure capacities.

Other reasons for disruptions can vary. A runway could be closed for maintenance. Security could be another reason.

1.3 Strategies to Recover from Disruption

A combination of several strategies can be used to minimise the effects of disruption at an airport and to recover efficiently and as quickly as possible. Some of these are:

1. Ground Holding;
2. Airborne Holding;
3. Flight Cancellations; and
4. Diversions to other airports.

1.3.1 Ground Holding

Ground Holding, also known as the *Ground Delay Program (GDP)*, is a vital control strategy suggested by several authors for use in disruption recovery (Bertsimas and Patterson 1998; Filar et al. 2000a; Hoffman and Ball 2000; Richetta 1991; Richetta and Odoni 1993). It consists of delaying the departure of some flights from their airports of origin, in anticipation of bad weather (or *reduced arrival capacity* in general) at destination airports.

In other words, flights are held on the ground at their origin due to capacity reductions at their destination; once such flights are cleared for take-off, they arrive at their destination without any additional delay.

Ground holding is practised at Australian airports such as Sydney, Brisbane, Melbourne and Perth. It is also widely practised by the Federal Aviation Administration (FAA) in the United States. Given that the demand for arrivals at an airport will, at times, exceed capacity, it is beneficial to delay some of these flights at their originating airport. This is because delays of aircraft on the ground are cheaper and safer than equivalent delays in the air.

However, appropriate ground holding of flights requires an accurate prediction (perhaps hours in advance) of the amount of congestion that will occur at the destination airport. An inaccurate weather forecast, combined with ground holding, can result in under-utilisation at the destination airport.

Hence if the flight duration is large (as in long-haul flights), flights are *not* subjected to ground holding; this is because, by the time the aircraft gets close to its destination, the weather (at the destination) could significantly change. Furthermore, the aircraft could adjust its speed or take a different route more easily than short-haul flights. Long-haul flights can also be diverted to other airports close to their destinations; this option is not meaningful for the short-hauls (they might as well be ground held at their origins).

1.3.2 Airborne holding

Airborne holding is the practice of holding aircraft in holding patterns in the airspace near the airports until they are able to land. It is more expensive and perceivably less safe than ground holding.

1.3.3 Cancellations and Diversions

The other two strategies (flight cancellations and diversions) are self-explanatory. When an airline cancels a flight, obviously it needs to accommodate those passengers on other flights. If a flight bound for Airport *A* is diverted Airport *B*, naturally it needs to be flown from *B* to *A* at some future time.

1.4 Importance of Arrival Capacities

As in (Navazio and Romanin-Jacur 1998), in several models presented in this book, only arrival capacities at airports have been implemented, *not* departure capacities. Generally, arrival capacities are more restrictive than departure capacities. In bad weather, an airport can handle fewer arrivals than departures.

Gilbo has studied the interaction between arrival and departure capacities at airports in (Gilbo 1993, 1997; Gilbo and West 2013). An example for Chicago airport is given below in Fig. 1.1 by (Gilbo 1997).

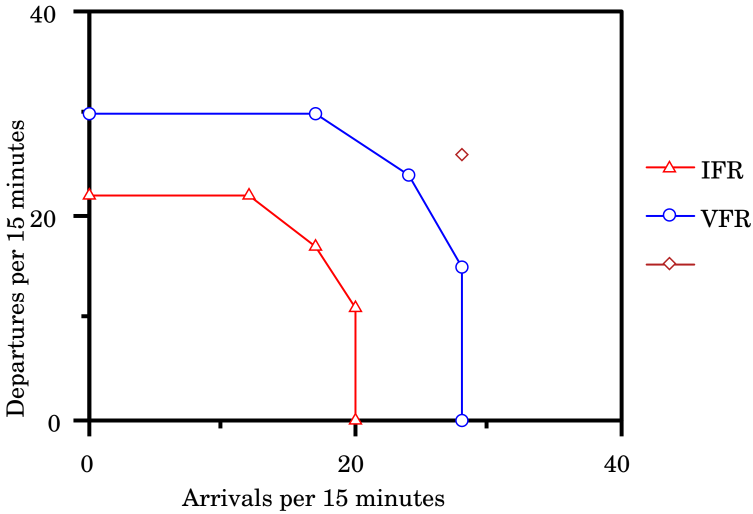

Fig. 1.1 Interaction between arrival capacity and departure capacity

1.5 Discretisation of Time

A key feature of the models presented here is the discretisation of time into periods (or intervals). There could be any fixed number of these periods, and they could be uniform or varying in length throughout the day. Most commonly we will deal with a 24-hour day divided into 96 periods of 15 minutes each or 48 periods of 30 minutes each.

In discrete optimisation, an integer programming (I.P.) framework is frequently used to state the problem.

1.6 Literature Review

A detailed survey of airport recovery literature up to 1999 appeared in the work of Filar, Manyem, and White (2001).

1.6.1 Integer Programming Models

We now discuss some of the results on the Integer Programming (I.P.) modelling and testing in disruption recovery at airports.

Vranas et al. (Vranas, Bertsimas, and Odoni 1994) provide the first model that considers this problem in a dynamic environment, that is, as the weather changes, and as aircraft fleet become available (or unavailable), the model provides updated ground holding decisions. The objective minimises the sum of airborne delay and ground delay. They provide 0-1 integer programming formulations that can be extended to include flight cancellations, arrival-departure capacity curves, hub and spoke network, and enroute speed adjustments.

Bertsimas and Patterson (1998) consider the Traffic Flow Management Problem (TFMP) caused by disturbances to flight schedules. This includes determination of aircraft release times at airports (ground holding) and optimal aircraft speeds while airborne. Their model is an integer program that considers the capacities of the en route airspace, airport capacities at different time intervals, and cost per unit time of holding aircraft on the ground and in the air. The objective function minimises the total delay cost.

Navazio and Romanin-Jacur (1998) construct an integer programming model to analyse multi-airport ground holding. Henceforth, we call this the *NRJ model*. Their model minimises the overall delay cost subject to airport arrival capacity, connections, and time constraints imposed by airlines. It assumes that the flying time of each flight is fixed, and hence any delay suffered by a flight is due to ground holding at the origin airport. The model distributes the ground holding delays among a set of flights originating at a set of airports. We describe this model in greater detail in Sec. 2.2.

Hoffman and Ball (2000) present five different models of the Single Airport Ground Holding problem in the presence of *banking* constraints to accommodate the hubbing operations of major airlines in the United States. A *bank* is a group of flights. The banking constraints at an airport connect flights in an *arrival bank* with those in a *departure bank*. For two of these models, they show that these constraints induce facets of the convex hull of integer solutions.

The MARFE model in Chapter 4 can be viewed as descending from the models of (Bertsimas and Patterson 1998) and (Navazio and Romanin-Jacur 1998), and as an extension of an intermediate ground holding model reported in (Filar et al. 2003). Its main distinguishing features are the multistage, adaptive structure and the incorporation of the many realistic constraints that apply at major airports currently.

Several authors have studied disruptions from an airline's perspective. See (Ball et al. 2007) and (Clausen et al. 2010) for surveys. A comprehensive survey of air traffic flow management was undertaken in (Vossen, Hoffman, and Mukherjee 2012). However, as in (Filar et al. 2003) and (Filar et al. 2007), we focus on a *common good perspective* (Sec. 1.1) that serves the interests of everyone: the airports, the airlines, the air service providers such as air traffic control, and last but certainly not the least, the travelling public.

In Chapter 3 of his Ph.D. thesis, (White 2011) discusses conditions under which various models (in particular, the linear relaxations of these integer programming models) in air traffic disruption recovery are guaranteed to return optimal solutions in polynomial time.

More recently, (Manyem 2018) developed polynomial time algorithms for (a) the NRJ model, and (b) an extension of the NRJ model that allows flight cancellations. These algorithms are based on the primal-dual algorithms used in Integer Programming. The author proved that the algorithms are guaranteed to return optimal solutions as long as the problem instances are feasible.

A version of the NRJ mathematical model for disruption recovery that implements non-linear delay costing with an accompanying heuristic was presented in (Manyem 2021a). In (Manyem 2021b), the author models curfew violations at airports, proves that the problem is NP-hard and presents an approximation algorithm.

1.6.2 Stochastic Models

In (Terrab and Paulose 1992), a stochastic programming model for a single airport is considered. In this model, capacity profiles are considered to be random and it is assumed that a probability distribution on these capacity scenarios is known.

(Terrab and Odoni 1993) formulate two models and analyse them. The first, an integer programming model, is deterministic. The objective is to minimise the total ground holding costs of all flights scheduled to arrive in the time horizon selected by the study. It assumes that airport capacities and flight duration times are known in advance with certainty and that airborne holding is always more expensive than ground holding. Under these conditions the ground holding solution will ensure that the number of flights reaching the airport never exceeds capacity and that no airborne holding is necessary. This model can be solved by standard minimum cost network flow algorithms. However, the authors also present a faster algorithm that gives priority to the flights with the highest marginal cost of delay.

The second model in (Terrab and Odoni 1993) is stochastic, and some airborne holding may occur because of uncertainty of airport capacity over the subsequent few hours. A number of airport capacity scenarios are possible, and the probabilities of the scenarios are assumed to be static. A dynamic ground holding policy would be more realistic since capacity forecasts become more certain as time passes. A dynamic policy would favour long-range flights over short-range flights since the ground holding of short flights can take advantage of more accurate forecasts at the scheduled takeoff time. In the current system in use in Australia and the US, long-range flights are generally exempt from ground holding.

1.6.3 Computational results

In (Navazio and Romanin-Jacur 1998), the authors provide a heuristic (the NRJ heuristic), based on the limited-resource Critical Path Method, that provides suboptimal results but has a much shorter running time on average. Among other data, the input to the heuristic includes the published schedules, airport capacities at different time intervals, and the number of delay periods