Discrete Real-Time Flight Plan Optimization

by

Francois Le Sellier

Submitted to the Department of Aeronautics and Astronautics
in Partial Fulfillment of the Requirements for the Degree of
Master of Science in Aeronautics and Astronautics

at the

Massachusetts Institute of Technology

September 1999

Signature of Author ....
Department of Aeronautics and Astronautics
August 23, 1999

Certified by ....
James K. Kuchar
Assistant Professor of Aeronautics and Astronautics
Thesis Supervisor

Accepted by ....
Jaime Peraire
Professor of Aeronautics and Astronautics
Chairman, Department Graduate Committee

Discrete Real-Time Flight Plan Optimization

by

Francois Le Sellier

Submitted to the Department of Aeronautics and Astronautics
on August 23, 1999 in Partial Fulfillment of the
requirements for the Degree of Master of Science in
Aeronautics and Astronautics

ABSTRACT

Worldwide, the continuously growing air traffic induces a need for new ATM concepts to be defined. One possibility is using a more decentralized system predicated mainly around free routings (Free Flight), for a more flexible management of airspace.

The present study first highlights the discrepancies and inefficiencies of the current best flight-plan optimizing software that use the Cost Index concept before departure. It then investigates techniques to perform enhanced flight-plan optimizations en-route, with algorithms that are less complex than using the Cost Index. The long-haul flight leg that is considered through the simulations is London (UK) - Boston (MA, USA), flown on a constant flight level. This study shows that running another optimization at the Top of Climb point reduces the average delay at destination from 6.9 minutes to 5.0 minutes. Then, the more futuristic method of considering discrete flight-plan optimizations, while en-route using updated weather forecasts, provides results that are more interesting. If the weather forecasts and the optimizations are done simultaneously every 3-hour or 1.5-hour, the average delay respectively becomes 2.6 minutes or 2.0 minutes.

The second part of this work investigates ways of performing a Linear Program to fly a route close to a 4D-trajectory. This study provides ways of determining the exact weight values for the different state variables used in the cost function to minimize.

Thesis Supervisor: James K. Kuchar

Title: Assistant Professor of Aeronautics and Astronautics

ACKNOWLEDGMENTS

I want to express my gratitude to the following people who have brought a lot of useful information to my work:

- At Airbus Industries: M. Schmidlin (aircraft performance engineer).
- At Air France Airlines: J-Y. Jung (chief, Airline Performance), J. Coquel (contact with the CFMU), G. Cordier (flight planning software), C. Voreux (chief, Repetitive Flight Plans), Mr. Dio and B. Aran (chiefs, Route Charges), P. Fauché (engineer, aircraft performance), G. Tardieu (chief, Operations Control), Mr. Nepveu (chief, Air Inter's Aircraft Operators), P. Gastaud (chief, flight plan document preparation, and chief, dispatchers), J. Mauline, M. Gergaud, M. Allam and M. Antonio (dispatchers), Mr. Epeldé (flight plan document compiler), H. Simatic and B. Serve (aircraft operators), N. Huguenot (chief, aircraft operators), R. Le Ray (chief, flight planning and crew scheduling software), N. Le Gall (engineer, flight planning software), M. Arnoux, J. Bensard, Mr. Charles, Mr. Gussard, Mr. Routier, and others (pilots and copilots).
- At Air Liberté Airlines: S. Depeyre (chief, aircraft operators), H. Julienne and F. Fournier (engineers, aircraft performance), P. Caccia (flight planning software), and B. Hemery (chief, TAT aircraft performance, and Fokker 100 pilot).
- At Air Liberté Industries: Mr. Lemaux (chief, maintenance costs).
- At AOM French Airlines: J-N. Garnier (chief of aircraft operators), L. Tisserand (aircraft operator), O. Fages (aircraft performance engineer), and G. Laforêt (engineer, aircraft performance engineer, and DC10 pilot).
- At the Athis-Mons Area Control Center (ACC) and Upper Airspace Control Center (UAC): C. Couleau (external relations, former air traffic controller), J. and M. Sayous (air traffic controllers), and M. Mathias (chief, ATC room).
- At the EUROCONTROL Experimental Centre: M. Dalichampt (CFMU slots allocating system CASA specialist), C. Magnat (IFPU2 operator), and F. Boulanger (chief, IFPU2 operators).
- At LTU Charter Airline: Mr. Sperber (dispatcher).
- At the Massachusetts Institute of Technology: Professor J. Kuchar (department of Aeronautics and Astronautics), Professor E. Feron (department of Aeronautics and Astronautics), and Professor Illari (department of Earth, Atmospheric and Planetary Science).
- At the Orly Tower: the responsible for external relations.
- At the Roissy-Charles De Gaulle Tower: the chief of the ATC room, S. Sailly (his chief), a few controllers, and S. Ricq (engineer, queuing software).

Table of Contents

Table of Figures

Glossary

ACC	Area Control Center
ADS	Automatic Dependent Surveillance
AO	Airline Operator
AOC	Airline Operating Center
ATC	Air Traffic Control
ATFM	Air Traffic Flow Management
ATM	Air Traffic Ma nagement
ATMS	Air Traffic Management System
CF	Contingency Fuel
CFMU	Central Flow Management Unit
CI	Cost Index
CNS/ATM	Communication Navigation Surveillance for Air Traffic Management
ECAC	European Civil Aviation Community
EP	Entry Point
ETA	Estimated Time of Arrival
ETO	Estimated Time Over
FAA	Federal Aviation Administration
FANS	Future Air Navigation Systems
FFAS	Free Flight Airspace
FIR	Flight Information Region
FL	Flight Level
FMS	Flight Management System
FPL	Flight Plan
GMT	Greenwich Meridian Time
GPS	Global Positioning System
HF	High Frequency
ICAO	International Civil Aviation Organization
INS	Inertial Navigation System
IROPS	Irregular Operations
LAAS	Local Area Augmentation System
lat/long	latitude/longitude
LP	Linear Program
MCT	Minimum Connecting Time
MIT	Massachusetts Institute of Technology
N4D	Near 4D
PFD	Planned Flight Data
RFL	Requested Flight Level
RPL	Repetitive Flight Plan
RTA	Required Time of Arrival
SR	Specific Range
SSR	Secondary Surveillance Radar

SUA	Special Use Airspace
TMA	Terminal Maneuvering Area
TOA	Time Of Arrival
TOC	Top Of Climb
VHF	Very High Frequency
WAAS	Wide Area Augmentation System

Section 1
Introduction

1.1 Air Traffic Management: an Optimization Problem

The objective of air transportation is to carry passengers or cargo between two airports in the world with the maximum safety and in a cost-efficient way. For each flight, the different partners involved in air traffic flow operations (passengers, aircraft operators, authorities) are not interested in optimizing the same criteria. For example, the Air Traffic Control (ATC) systems will try to maximize flow efficiency at a manageable controllers' workload, whereas airlines' Aircraft Operators (AOs) will want to minimize the trip duration and the fuel consumption. The economic efficiency of the system, and thus a large part of its acceptability by the different partners of Air Traffic Management (ATM), comes from the adequacy of the cost function to minimize. Ref.1 studies the issues raised by the definition of a cost function for Air Traffic Flow Management (ATFM) optimization problems. The present study addresses the benefit-driven optimization problem from the airlines' perspective.

Fortunately, the passengers' concerns often match the airlines', as a delay on one flight can affect the costs incurred by further flights chained to it. This phenomenon can lead during a day to a significant disorganization of the entire air traffic system (including the ATFM), particularly around a hub airport. This already suggests that the cost function to consider for the flight path optimization problem should be highly concerned with meeting the scheduled time of arrival.

1.2 Objectives of Present Study

Actually, as seen in the next paragraphs, a look in the field at the different parameters considered by airlines when determining their flight plans shows that meeting the scheduled time of arrival summarizes very well the numerous flight planning concerns. This is useful because it is much easier to consider a single goal rather than several interdependent ones. Moreover, Section 2 shows that the complex Cost Index concept that is currently used in the most advanced flight planning systems has important discrepancies and inefficiencies. Consequently, the first part of the present work simplifies the routing method by focusing on a goal that is more relevant than using a Cost Index: it is to simply concentrate on meeting the time of arrival.

Air traffic is growing worldwide at a rate that imposes new air traffic management concepts to be defined. One such concept, known as Free Flight, is an environment based on free user-preferred routings, permitting dynamic modifications of the flight plan. This requires transitions from the current air traffic system. The present work looks at the efficiency of realistic transition states, first by recalculating the optimized flight plan at the Top of Climb point, then by considering new optimizations run at several discrete times during the flight.

Finally, in the perspective of discrete real-time optimizations, the second part of this work investigates how to cost-efficiently perform a 4D-trajectory.

1.3 Flight Plan Types

1.3.1 Mach Selection

At the flight-planning level, the en-route Mach number is decided in compliance with the strategy chosen by the airline for the flight. This strategy is described by one of the following criteria:

- Minimum fuel
- Minimum time
- Minimum cost
- Minimum fuel within schedule
- Minimum cost within schedule

These are explained in section 1.4.

During the en-route phase, the pilot selects a Mach between the Maxi Range Cruise Mach M_{MRC} and the Long-Range Cruise Mach M_{LRC} (also called the "accelerated Mach"). M_{MRC} is the Mach at which the aircraft flies over its maximal Specific Range (SR), where the Specific Range is the distance flown with one ton of fuel at a given weight, height and temperature. M_{LRC}, typically superior to M_{MRC} by 0.015 or 0.02, is defined by the following relation:

$$SR_{LRC} = 0.99 \times SR_{MRC} \tag{1.2.1}$$

where SR_{LRC} is the Long Range Cruise Specific Range, and SR_{MRC} is the Maxi Range Cruise Specific Range.

Consequently, during a flight, the best possible increase in the Mach M is less than or equal to ΔM=0.02. In the best case scenario, the flight duration is then shortened by a mere 1.5

minute per hour of flight. Indeed, if Δd is the extra distance flown during one hour at Mach $M+\Delta M$, then:

$$\Delta d = \Delta M \times a \times 1hour \tag{1.2.2}$$

where a is the speed of sound. Therefore, the time gain during a one-hour distance flight is:

$$\Delta t = \frac{\Delta d}{Speed} = \frac{\Delta M \times a \times 1hour}{M \times a} = \frac{\Delta M}{M} \times 1hour \tag{1.2.3}$$

If the flight cruise Mach is M=0.8, we conclude that Δt=1.5 minute.

1.3.2 Medium-Haul versus Long-Haul Flights

1.3.2.1 Medium-Haul Flights

The previous result proves that changing the Mach may only slightly reduce the delays at arrival for medium-haul flights. Indeed, combining the use of the "accelerated Mach" M_{LRC} with the decision of flying direct routes (see section 3.5) may reduce the flight plan by only up to five minutes, or ten minutes in extremely rare situations. The time reduction goes down to only two to three minutes if the flight remains on its original trajectory. Even though a few minutes recovered over a delay may be helpful, the airlines agree that a more efficient use of airspace, such as in a Free Flight environment, will be mostly beneficial to long-haul flights. During them, dynamic path optimizations should importantly affect the time of arrival (TOA). The impact will certainly be less obvious with short or medium-haul flights, in particular over Europe where the numerous military restricted areas are strong barriers to a Free Flight environment.

Moreover, medium-haul flights are usually short enough to allow the aircraft to carry a lot of extra fuel, thus enabling them to fly at the maximal Mach and shorten the flight duration. Passengers prefer that option, as well as it matches the airline's interest to lower its operational costs. Indeed, shorter flight duration incurs smaller crew costs, thus resulting in smaller total flight cost, as crew costs are more expensive today than fuel is.

Finally, it must be noticed that airlines generally consider the need for a Cost Index (see Section 2) for medium-haul flights to be questionable. Medium-haul flights' duration may be importantly affected by numerous trajectory constraints, which are too difficult to implement as parameters in the Cost Index model. Therefore, airlines usually simply choose between flying at the Maxi Range Cruise Mach M_{MRC} or at the Long-Range Cruise Mach M_{LRC}.

1.3.2.2 Long-Haul Flights

Long-haul flights divide into two main categories: the limiting and the non-limiting ones.

1.3.2.2.1 Limiting Flights

A flight is said to be limiting when its duration and payload are such that it cannot carry enough fuel to fly at the Long-Range Cruise Mach M_{LRC} during the whole flight. This often causes a flight at Maxi Range Cruise Mach M_{MRC}, mostly with those small airlines that do not have the Cost Index technology. However, a compromise can be found by using an intermediate Mach chosen from a finite set of Mach numbers with the use of a preset Cost Index (see Section 2). As an example, Groupe Air France Airlines makes its choice at the flight planning level between four Machs.

1.3.2.2.2 Non Limiting Flights

When enough fuel can be carried onboard, the aircraft usually flies at the Long-Range Cruise Mach. The reason is that during a flight, the time-related costs (cf. section 2.2.2) are usually - and nowadays - more expensive than fuel is.

1.4 Flight Plan Strategies

1.4.1 Minimum-Fuel Strategy

The minimum-fuel strategy consists in flying at the Maximal Range Cruise Mach M_{MRC} (the lowest Mach possibly in use during the en-route phase) to enable the aircraft to fly over the longest distance. In other words, it provides the lowest fuel burn for a given range. This criterion intervenes when the payload is so heavy that the legal amount of fuel needed to fly at the Long-Range Cruise Mach M_{LRC} cannot be carried. This situation primarily happens with long-haul flights, with a higher probability for longer flights. For example, a flight from Paris (France) to the Reunion Island lasts 10.5 hours and cannot be flown at maximal speed.

The minimum-fuel strategy is also used when the aircraft has been sent into a holding pattern.

This criterion corresponds to the use of a Cost Index (cf. Section 2) equal to zero.

1.4.2 Minimum-Time Strategy

The minimum-time strategy corresponds to flying at the Long-Range Cruise Mach M_{LRC} (the upper bound of the Mach range). In other words, it corresponds to flying at the highest speed with "reasonable" fuel consumption. Some airlines prefer to apply this criterion instead of the minimum cost strategy to save on maintenance costs. This decision varies from one airline to another, depending on the adopted economic strategy.

When an airliner is late on schedule and when payload allows it, the aircraft is usually flown at the "accelerated Mach". This has the most interest for long haul flights, which can this way recover up to twenty or thirty minutes, possibly bringing important useful consequences under circumstances like heading towards a hub airport. The pilot has full authority for deciding at Block Time the strategy applied during the flight - after acknowledgment of recommendations by the Airline Operating Center (AOC).

Moreover, the pilot is responsible for deciding at any moment during a flight whether to continue with the current flight regime or to switch to another one, which generally implies using the accelerated Mach to minimize the flight duration. This decision is up to the pilot because even a Flight Management System (FMS) does not send to the ground the dynamic times-over-future-waypoints calculations, hence no one but the pilot knows better if the schedule will be met. Anyway, it should be remarked that increasing the Mach during the last half-hour of a flight is rarely an option as it unreasonably raises the fuel consumption for less than a thirty-second delay recovery.

Finally, it must be stressed out that an airliner should not arrive at destination ahead of schedule as this may cause problems like not having the gate available yet. In this case, the consequences may include a holding pattern or a long taxiing time, both inducing higher fuel costs and ruining the purpose of flying faster.

This criterion corresponds to setting a Cost Index to its maximal value.

1.4.3 Minimum-Cost Strategy

A priori, the minimum-cost strategy looks the most attractive, as it makes use of the Cost Index (see Section 2) to find a trade-off between fuel and time-related costs to attempt and really minimize the total trip cost. This strategy can be applied to properly FMS-equipped aircraft only. However, a genuine optimization should dynamically take into account the often-unpredictable real-time events that affect the time of arrival, while the Cost Index is a static optimization method. It is entered in the FMS before take-off and is supposedly fixed for the entire flight. In practice, pilots fly in accordance with the chosen Cost Index as long as no real-time event causes a modification of the strategy for the remainder of the flight.

Most airlines are interested in taking the route charges into account (cf. section 3.7), in particular in Europe where the yearly expenses they incur justify trying to reduce them. However, for now, few companies try to reduce the route charges bill because this is a problem relatively difficult to implement in the optimizing tool of a flight planning system. This is bound to evolve; in the meantime, alternatives have sometimes been found. For example, Air France chooses a cheaper route (considering route charges) only if:

- the money benefit over the best route without route charges is higher than a constant threshold,
- and the flight duration exceeds the other's by less than a fixed percentage.

1.4.4 Minimum-Fuel-within-Schedule or Minimum-Cost-within-Schedule Strategies

The minimum-fuel-within-schedule strategy is self-explanatory. The minimum-cost-within-schedule strategy consists in minimizing the trip cost with the constraint of meeting the scheduled arrival time at destination (this is typically of much less interest to charters or cargo flights). During the present study, no existing flight-planning system was found using either one of these criteria.

The minimum-cost-within-schedule strategy has an alternative: fix the Estimated Time of Arrival (ETA) according to the average flight duration known from former years' experience and/or according to the known aircraft performance. This theoretically ensures that the flight will stay close to schedule if its departure is not too much delayed. This is why the existing flight-planning systems simply assume that the time of departure is met to ensure staying close to schedule. The path is calculated from the planned time of departure and uses extrapolations of the weather forecasts to obtain the time at destination. The actual time of arrival is hence different than the scheduled one, notably because time translations occur with the allocation of a departure slot by the Air Traffic Management System (ATMS). The difference is often small, yet often non-negligible in terms of consequence.

It quickly comes to mind that a way of meeting the scheduled time of arrival would be to run the path optimization program with the simple constraint of meeting the time of arrival. Doing the path optimization backwards, using extrapolations of the weather forecasts known before departure, would provide the exact path and time of departure. In fact, the following reasons explain why this method is never and should not be used:

- the time of departure is a constraint, not a parameter,
- the weather forecasts are more likely to be accurate near the time of departure rather than near arrival, causing the basis of the path optimization to be unpredictably obsolete.

Therefore, it should not be assumed in the optimizing process that the time of arrival is known from the beginning. One alternative is to run several flight path optimizations with different Mach values, then to choose the one that brings the aircraft the closest to the scheduled time of arrival (TOA). For example, Air France chooses between four different Machs (the reason is in fact different than trying to meet TOA: the Mach is chosen such that it incurs the minimal flight cost in terms of fuel and crew costs).

1.4.5 Non Universality of Strategies

In practice, an airline's flight plans' strategy may drastically change from one day to the next.

Moreover, two airlines seemingly working the same way may not apply the same criterion over similar flight legs. For example, this may occur when the crew costs are very different, inducing a greater concern to minimize the flight duration for one airline versus the other. Fifteen-minute delays at destinations sum up to high extra crew costs per year, which importance depends on the crew costs per extra hour of flight, causing a more or less important concern about reducing delays.

Section 2

Cost Index

2.1 Definition

The Cost Index (CI) for a flight is the ratio of the time-related costs to the fuel price, with the time-related costs being defined as those incurred by one extra hour of flight.

$$C_I = \frac{C_t}{C_f} \tag{2.1.1}$$

Along a given route, reducing the flight duration by opting for a higher Mach is done at the expense of fuel. The Cost Index method is an attempt towards an appropriate trade-off. It performs an optimization of the total trip cost by controlling both the fuel burn and the flight duration. This optimization is run according to diverse rules relevant to the airline's economic policy.

The Cost Index can only be used with suitably equipped aircraft: a Flight Management System equipped with the Aircraft Performance option is necessary.

2.2 Components

2.2.1 Fuel Price

In theory, the fuel price used in the Cost Index calculation is the sum of:

- the fuel price at departure,
- the fuel price at destination,

- the trip fuel,
- the tankered fuel (to be used on the next rotation of the aircraft).

In the 1980's, fuel expenses represented 25% of the total operating costs. Now, because fuel is cheaper and because most aircraft burn a significantly lower amount of fuel, they count for about only 10% and are now less important than the other costs induced by a flight. This explains why, nowadays, it is considered more important to reduce the flight duration rather than save on fuel.

2.2.2 Time-Related Costs

Also called 'marginal', the time-related costs are those incurred by one extra hour of flight. Their three major components are the following:

- the cockpit crew costs,
- the cabin crew costs,
- the engine and airframe maintenance costs. These include costs per hour of flight (with the assumption that the flight-time-dependent degradations evolve linearly with time) and the cyclic ones. The latter are analytically interpreted as costs per hour of flight, so that all maintenance costs are used in the optimization scheme as time-linear (more details are described in section 2.2.3).

Other components may intervene in the time-related costs:

- the aircraft depreciation,
- the customized costs raised by late arrivals, when delay reaches a threshold; these are significant of :
    - ◊ the crew overtime compensation,

◊ the passengers' dissatisfaction,

◊ the missed connections (mostly in case of a hub), which usually incur high extra costs to take care of:

  * the passengers' rescheduling on other flights,
  * extra meals and accommodation (catering).

2.2.3 Maintenance Costs

A wide variety of parameters can be considered in the Cost Index calculation as "maintenance costs", including the costs incurred by:

- the carpet, seats, etc. (mainly dependent on the type of aircraft),
- the "actual maintenance" of the aircraft,
- the maintenance staff,
- the maintenance contracting (when the airline does not have technicians at one airport),
- the stocks possession,
- the flight monitoring,
- the hangars,
- the engineering,
- possibly other costs.

These entire costs sum up to a non-negligible amount in an airline's operating costs as they typically count for 15% to 20% of the total time-related costs. However in practice, many components of the maintenance costs are too difficult to evaluate, and only few of them are used in the Cost Index calculation - when there are any.

2.2.4 Fixed Costs

Fixed costs do not intervene in the Cost Index calculation because they are normally independent from trip fuel and flight duration, and they would be incurred anyway. They include:

- the cyclic maintenance costs,
- the landing and user costs,
- the basic fixed salaries,
- the basic passenger costs (meals, ...).

2.3 Proof that Cost Index is a Performance Optimization Tool

For the purpose of this proof, it is first assumed that the total trip cost can be written as:

$$TripCost = C_f \cdot \Delta f + C_t \cdot \Delta t + C_c \tag{2.3.1}$$

where C_f is the cost of a ton of fuel - constant, D_f is the trip fuel, C_t is the time-related cost per extra flight hour - constant, D_t is the extra flight duration, and C_c represents the fixed costs.

Hence, the only part of the trip cost that could be modified if changes were applied to the flight path is $C_f \cdot \Delta f + C_t \cdot \Delta t$. This provides the cost for one extra nautical mile:

$$ExtraCost = \frac{C_f \cdot \Delta f + C_t \cdot \Delta t}{V \cdot \Delta t} \tag{2.3.2}$$

where V is the magnitude of $\vec{V}$, the ground speed of the aircraft. We have:

$$\vec{V} = a \cdot \vec{M} + \vec{W} \tag{2.3.3}$$

where a is the speed of sound, $\vec{M}$ is the Mach vector, and $\vec{W}$ is the wind speed vector.

We then introduce the Specific Range S_R, which is the distance flown with one ton of fuel:

$$S_R = \frac{V \cdot \Delta t}{\Delta f} \tag{2.3.4}$$

Equation (2.3.2) can then be written as:

$$ExtraCost = \frac{C_f}{S_R} + \frac{C_t}{V} \tag{2.3.5}$$

Hence, for a given barometric altitude (which allows one to consider a as a constant) the minimum cost cruise verifies:

$$\frac{d}{dM}(ExtraCost) = 0 = C_f \cdot \frac{d}{dM}(1/S_R) + C_t \cdot \frac{d}{dM}(1/V) \tag{2.3.6}$$

Also, we have:

$$\frac{dV}{dM} = \frac{a \cdot \left(\frac{\vec{M}}{M} \bullet \vec{V} \right)}{V} = a \cdot \cos\beta \tag{2.3.7}$$

where β is the drift angle; therefore, we obtain:

$$C_I = \frac{C_t}{C_f} = \frac{V^2}{a \cdot \cos\beta} \cdot \frac{d}{dM}(1/S_R) \tag{2.3.8}$$

This ends the proof of CI being a flight performance optimization tool.

Finally, it can be noticed that since β is small in practice, it is legitimate to approximate $\cos\beta$ as 1, such that (2.3.8) can be written as

$$C_I = \frac{C_t}{C_f} = \frac{|\vec{V}|^2}{a} \cdot \frac{d}{dM}(1/S_R) \tag{2.3.9}$$

2.4 Cost Index-Based Flight Path Optimization

Flight Planning systems typically perform Cost Index-based flight path optimization with the following successive steps:

1) For a particular flight leg, the airline first chooses the strategy of the flight (see paragraph 1.4). This translates into opting between flying the en-route phase at the Maximal Range Cruise Mach or at the Long-Range Cruise Mach. Then, the airline extracts from a database a Cost Index that leads to this Mach number in normal atmospheric conditions.
2) The Flight Planning system uses this Cost Index (CI) with the weather forecasts to compute the "best" routes (assuming the Cost Index truly represents the best strategy for the flight) towards destination. Using databases on the aircraft performance, the flight-planning device elaborates the best paths with the best changes of flight level for the desired Mach values.
3) The actual flight plan is then chosen among these paths to best cope with concerns related to the departure slot allocation. From this step onwards, attention is rarely paid to the time of arrival at destination.
4) Later, at Block Time, the pilot enters a Cost Index value in the Flight Management System (FMS). This value is either the same as the one used in the previous steps, or is different when late changes of flight strategy have been decided (e.g., when delay at destination may cause a landing-prohibiting curfew). Whatever the Cost Index value is, it modifies the flight plan which had been originally accepted by the Air Traffic Management System (ATMS) by changing the Mach and the flight level changes (the route is unchanged).
5) Finally, during the flight, the FMS uses the Cost Index on the ATMS-accepted route with updated weather forecasts to dynamically compute the optimal flight level changes. This is a first approach to a real-time optimization, even more because ATC may allow flying direct

routes (cf. section 3.5) or may impose real-time flight level constraints to ensure the avoidance of collisions. However, these modifications have the limitation of always using the original route.

2.5 The Cost Index in Practice

2.5.1 Components

Whether it is suitably equipped for Cost Index use or not, each airline agrees that the Cost Index is a “good idea”. However, its necessity and more importantly its efficiency are usually questioned.

Depending on the airline, in practice only few of the possible cost parameters listed in section 2.2 are considered for the Cost Index calculation. The time-related costs usually comprise at least the crew and some maintenance costs, but the “customized” ones are very rarely taken into account.

Moreover, the fuel cost used in the Cost Index calculation should always be the one at the airport where the fuel was filled in the tanks, but some airlines pay little attention to this. For example, Air Liberté Airlines resets its Cost Index values for each departure airport once every three months. This may be irrelevant because the fuel is not always entirely filled in at the departure airport, some of it may have been tankered at a previous one. In most cases, airlines erase this issue by, for example, considering the average fuel cost for the whole fleet of the same type of aircraft, on all the flight legs they cover during a year. This questionable statistical method shows already that the Cost Index is just an approximation of what would be an optimizing parameter.

Other reasons cause important Cost Index differences between airlines. A major one is that airlines do not use the same parameters for the maintenance cost calculation. Indeed, maintenance data are more or less easily accessible from one airline to another due to evaluation difficulty or simply confidentiality issues. In Reference 3, Airbus Industrie proposes formulas to compute these maintenance costs, provided manufacturers and repairmen can reveal a few specific parameters. These empirical formulas are highly questionable, and we will not even try to study their relevancy.

2.5.2 Efficiency

The Cost Index method may seem a good way to perform path optimization. In practice, it appears that its efficiency is too unpredictable - and too often questioned - to justify such a complex theory. Indeed, numerous real-time events may drastically modify the flight strategy.

2.5.3 Rounded Cost Index

Airlines use approximate values of what should be their optimum Cost Index. To noticeably modify the speed of an airliner, one needs to modify the Cost Index value (CI) by a factor of about two. This shows that speed is not very sensitive to CI, except when it is drastically changed. This allows each airline to use a Cost Index definition that is not extensive and not thoroughly accurate. Consequently, most airlines use rounded and fixed CIs.

As an example[4], a certain airline was recently using the following CIs:

- For B767-200: $CI = 70$ (reason given by pilots: it approximates the Long Range Cruise)
- For B767-300: $CI = 60$ (same reason)
- For B747-400: $CI = 100$ (reason: it allows low fuel burn with acceptable speed stability).

These differences are mostly due to crew wage differences, which are usually the mean values of all the wages on the considered type of aircraft.

- For each aircraft: $CI = 300$ when flight duration is critical, e.g., to avoid crew overtime.

This example shows that CI can be a very rounded approximation that can drastically change during irregular operations (IROPS).

Another airline[4] was recently using a CI of 80 for its B747-400s (obviously using a different Cost Index definition than the previous airline). Its pilots had received the instruction to switch to 250 when behind schedule for connecting passengers (this airline works as a hub).

Some other airlines[4] use non-rounded CIs by considering the fuel cost as the one at departure for an unnecessary "precise" calculation.

Finally, other airlines[4] even let their crews select CI.

These examples prove that the Cost Index definition is very uncertain. Actually, most airlines' engineering departments agree that computing an "exact" CI would not bring noticeable benefit, even during normal operations. On the contrary, improving the definition is often believed to induce more research costs than operational benefits.

2.5.4 An Irrelevantly Linear Definition

The variety of Cost Index values raises a major issue: is it even possible to find a list of parameters to use in the CI calculation to provide a real optimization from the airline's perspective? The answer is no. The following example shows it: Air Liberté Airlines represents the fatigue of the engines as hourly maintenance costs, while Air France does not. Both have justification for their respective decision:

- Air Liberté Airlines pays attention to taking into account the use of engines,

- whereas Air France rightfully considers that the engines' fatigue is not linearly dependent on the flight duration, therefore taking it into account in the CI calculation would induce a wrong calculation anyway. Indeed, the Cost Index value entered in the FMS directly affects the regime of the engines, causing their inside temperature to be one or another. This induces a certain fatigue speed, which eventually incurs hourly maintenance costs for the engines that are specific to this regime.

This example proves one important limitation of the Cost Index theory: it assumes all the cost parameters to be interpretable linearly while some of the maintenance costs are not.

The theory contains numerous interesting aspects, however the CI calculation and its use seem uselessly complex and impose tools that may be too sophisticated - hence expensive - for the limited efficiency that is obtained.

2.5.5 Criticisms about Proof

Actually, the definition of the total trip cost is questionable in itself. Indeed, the time-related costs are not all linear functions of the extra trip time as was mentioned in the former paragraph.

Moreover, the time-related costs during scheduled flight hours are often incorrectly considered as fixed costs - or at least as Mach independent. Using the same example as in the previous paragraph, different CI entries in the FMS lead to different fatigue rates for the engines, causing them different maintenance costs. In other words, the fuel cost during scheduled hours depends on the flight regime, which may change in real-time. Therefore, the total trip cost during scheduled hours should not be considered as a fixed cost.

Moreover, the proof of Cost Index being a performance optimization tool assumes that the aircraft remains on a single flight level. This suggests that a truly optimized flight probably involves different Cost Index values for the different flight levels encountered during the en-route phase (e.g., for a long-haul flight). In practice, things are different as a Cost Index value is entered in the FMS at Block Time and remains constant for the whole flight (see section 2.4). In other words, the Cost Index theory is based on an assumption that is not applicable to long-haul flights.

The CI theory assumes that D_t is positive, while there are situations when it would be interesting to fly faster than schedule.

Several other descriptions - expected at the time to be more practical - of the Cost Index parameters were attempted for the present work. They tried to stress more the possible discontinuity and non-linearity of some costs (notably by using discretized engine's regimes); unfortunately they were not of much use because their derivatives respective to Mach were unsolvable.

These reasons prove important discrepancies in the Cost Index theory, putting a high random weight on its efficiency. Moreover, this theory should certainly not be used for dynamic optimizations.

Section 3

Optimization Goals Analysis

3.1 Rationale

This section intends to show that the primary goal of each flight is to arrive on time. To clarify this, several points are discussed. First, attention is paid to the importance of taking-off on scheduled time, to eventually understand that needing to depart on time is just a consequence of the importance of arriving at destination on schedule.

3.2 Importance of Meeting Departure Time

3.2.1 Medium-Haul Flights

On medium-haul flights, departing on time is roughly equivalent to arriving on time. Indeed, we saw in section 1.3.2.1 that the duration of a medium-haul trip can only be shortened by five minutes (ten in rare situations) with direct routings and use of the "accelerated Mach". Therefore, it is especially important for medium-haul flights to depart on time to reach destination on schedule.

3.2.2 Rotations

Departing on time is a top priority when the aircraft is scheduled for rotation at destination. The daily schedule of an airliner sometimes asks for turn-times down to almost the airport's Minimum Connecting Time (MCT, the minimum time required by the airport's logistics to rotate an aircraft). For example, Air France's "La Navette" airliners have turn-times equal to Orly airport's MCT (thirty minutes). Therefore, a delay occurring at any moment of a day with a flight subject to rotations can generate delays up to the end of the day's operations, or, at least, a delay at next departure. Aircraft Operators work very carefully at avoiding this kind of "threat".

3.2.3 Curfew

The Aircraft Operators' concern about rotations becomes critical when the chained flights can face a curfew at the end of the day. If the destination airport closes before an aircraft lands, the latter is rerouted towards a nearby airport. For example, the 11:00 p.m. - 6:00 a.m. curfew at Orly airport may reroute aircraft towards Roissy Charles De Gaulle. In such occurrence, most passengers get a bad image of the airline and may decide to avoid flying with it in the future. Furthermore, this situation usually incurs very high extra costs by imposing the need to take care of the passengers' accommodation, transport, and food (catering). A rerouting due to a curfew is probably the worst situation for an airline - during safe operations.

3.2.4 Duty-Time Limits

Usually, it is considered less important for long-haul flights to meet the scheduled time of departure than with medium-haul flights because it is much easier to recover from a delay on a long-haul flight. However, for very long flights, the crew duty-time limit can be the number one constraint. It would impose an "accelerated Mach" strategy, with a departure close to the scheduled time (having to reduce the payload if necessary). There is also a duty-time limit for flight attendants, but it is much less a restriction because flight attendants are allowed to fly extra hours.

3.2.5 Psychological Importance of Departing on time

According to Air Liberté Airlines' aircraft operators, experience has shown that passengers generally notice a delay at departure after about three minutes (more relevantly on medium-haul flights). This explains why airliners try to depart within two extra minutes of the scheduled time. When this is not possible, the pilots try to limit the delay to a certain threshold (e.g., fifteen minutes), provided no constraints like rotations strongly recommend a shorter delay.

In general, the airlines without hubs admit that they pay less attention to reach the scheduled time of arrival. Their priority is usually to depart their flights with very small delay. They keep an eye on rotations, but they achieve short delays at arrival by focusing more on the time of departure. This is true as long as there is no risk of arriving at an airport under curfew, in which case the only priority is to reach the airport before it is closed.

3.3 Importance of Meeting the Time of Arrival (TOA)

For medium-haul flights, we explained in section 3.2.1 that trying to meet the scheduled time of departure is often sufficient to ensure arriving close to the scheduled TOA.

On long-haul flights, when payload allows carrying more fuel, direct routings and use of the "accelerated Mach" make it easier to recover from a delay at departure.

Meeting the time of arrival is very important for a number of purposes, among which are the following:

- please passengers (including the business-class passengers who are a priority),

- avoid a curfew after chained flights when turn-times are close to the Minimum Connecting Time (read section 3.2.3),
- permit adequate aircraft rotations (for the reasons given above),
- favor adequate crew rotations,
- allow adequate connections.

3.3.1 Hubbing

When an airline works as a hub at one airport, its arrivals and departures at this airport are sequenced in such a way that several times a day, a wave of airliners arrives and a wave departs shortly after. For example, this happens five times a day for Air France at Roissy Charles De Gaulle. This attracts passengers, as they naturally prefer fast connections to reduce their total trip time.

In practice, not every airliner reaches its allocated gate on schedule; this may result in missed connections. One common reason for this is that the airlines working as hubs reduce their connecting times as much as possible, close to the Minimum Connecting Time (MCT). MCT is airport-dependent: it represents the time needed to move passengers from one gate to another as well as to transit luggage from one aircraft to the other (which is more likely to be time-consuming). Since scheduled times of arrival are, in fact, averages over an entire IATA season (winter's September-to-March or summer's April-to-August), big differences from these averages occur. For example, a flight from Southeast Asia to France sees its duration differ by thirty minutes from April to August. This induces a fair probability to miss some connections.

If an aircraft delay causes missed connections, the airline takes care of:

- finding another connecting flight from the same airline, or from another if this is not possible,
- offering accommodations (food, hotel, etc.).

Typically, the airline takes care of the passengers who have missed their connection in order of economic importance, making a difference between a businessman and a tourist.

These operations are very expensive to the airline and shall be avoided, such that the next flight could be slightly delayed to allow a connection to happen. This decision is taken at the AOC level after consideration of:

- the number of passengers to connect and their economic importance,
- the possibility to connect them to another flight with a reasonable delay,
- the consequence on the next flights if the one to connect with is delayed.

The flights that are the most likely to be imposed such delays at departure are those with long rotation times at destination, to make sure that the delay on the next flight leg is not affected by the first delay. For example, if for one connection situation, there were another possible flight to connect with thirty minutes later, the airline would probably consider that missing the first connection is acceptable. On the contrary, if the next connecting flight were seven hours later, it would be very important to reach the scheduled connection, otherwise the non-connected passengers would be very unpleased and would possibly become forever-lost customers. These aspects are taken into account to possibly fly at the "accelerated" Mach.

3.3.2 Business-Class Passengers

They are very important to the airline because it sees in them as its best source of profit and because they are more likely to fly with the same airline again than tourists are. The airline does its best not to disappoint these customers. Since businessmen, for business reasons, are very scrupulous about reaching the destination airport on time, it becomes a much bigger concern to meet the time of arrival with a flight carrying a number of business-class passengers.

3.3.3 Charters and Cargo Flights

Naturally, meeting the scheduled time of arrival with charters and most cargo flights is less important than with other flights. In those cases, the main concern becomes minimizing the total trip cost - while paying little attention to delays.

3.4 Conclusion for the Flight Strategy

It was described above that, by far, most of the arguments in section 3.2 that justify trying to meet the time of departure are in fact concerns about arriving on time at destination (as is explained further in section 3.3). Actually, for passenger-carrying flights, only one argument was discussed not mentioning the importance to meet the time of arrival: the psychological impact on passengers of leaving the gate late (see section 3.2.5). This leads to the reasonable conclusion that arriving on scheduled time at destination is the one objective that sums up the arguments in sections 3.2 and 3.3. Moreover, it is much easier in a flight-plan optimization tool to concentrate on this single goal rather than to try to model all the different arguments as parameters. The complexity of the model would otherwise be much too important. It would actually be meaningless because the time of departure is not known before the ATMS has allocated a slot, therefore talking about real optimization before the departure slot is known is bound to errors.

These remarks have led the present study to concentrate on the importance of meeting the time of arrival, and to look at real-time optimizations.

3.5 Direct-Routings

3.5.1 Concept

When the airspace is under ATC control, the pilot can ask during the flight for direct-route allowances. Upon clearance, a direct-route shortens the flight path by deleting one or several waypoints from the scheduled path. Figure 1 shows two different "direct paths".

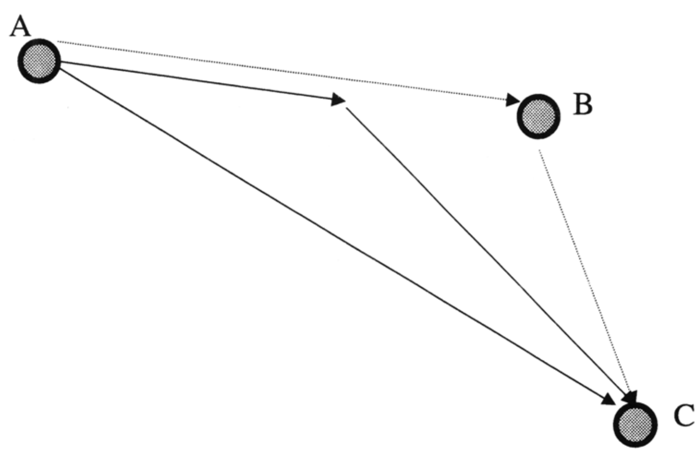

Figure 1. Two direct-routing scenarios
A, B and C are three waypoints. The dashed line represents the original planned path from A to C, via B. The plain lines are two direct-routing alternatives.

Obviously, airlines want their pilots to fly as many direct-routings as possible to reduce the total trip cost by shortening the flight duration and the fuel consumption. In practice, to avoid receiving direct-routing requests too often (thereby risking additional workload), controllers usually propose the direct-routings themselves. These shortcuts match the controllers' interests, as direct-routings reduce the time spent by one aircraft in one controller's surveillance area. However, it remains easier for controllers to handle the aircraft that stay within official airways, where most conflict resolution procedures are known.

3.5.2 Europe

Over Europe, most direct-routings occur at night, when traffic and controllers' workload are lighter (direct-routings are possible in Europe because all aircraft are tracked by radar). For example, a flight in France from Nice to Paris is ensured of receiving direct-path clearances if it takes off around 5:00 a.m., but it is confronted to a peak of traffic if it departs near 8:00 a.m. In the second case, the controller usually prefers the aircraft to stay on its scheduled path because it corresponds to a classic route that he knows how to handle.

Moreover, a number of specific-use-areas block direct-routings; these include the restricted military airspace and the transfer zones. There are numerous military restricted areas over Europe, in particular near Paris, and it is unanimously agreed among air traffic controllers and airlines' crews that they constitute a major constraint against direct-routings.

Finally, direct-routings across borders are not allowed. Therefore, having an aircraft constrained to stop a direct routing at the last waypoint before crossing a border is frequent.

3.5.3 The Atlantic

Waypoints over the Atlantic are defined in "lat/long" coordinates (latitude/longitude). Transatlantic flights are done with a constant Mach along one of four tracks (each defined by a set of "lat/long" waypoints) to ensure safe separation distances. The track concept has been necessary until now because there is no radar coverage over most of the Atlantic. By using tracks, air traffic control centers get reasonable knowledge of each aircraft's position by logging what track is flown, what time it was entered, and by using regular HF-radio position reports from the aircraft's crew. Then, the crew obtains position information from the inertial navigation equipment.

The four tracks are set on a 12-hour basis, switching westbound and eastbound flow. Each track crosses the northern part of the Atlantic and tries to make best use of the weather forecasts for one aircraft category. An aircraft that plans to cross the Atlantic will then be allocated to the track that best fits its cruise performance. Along these tracks, crews are not allowed to ask for direct-routings and aircraft are required to fly in conformance with their announced flight plan (except in case of an emergency or on controller's request).

Transatlantic flights should benefit a lot - and should be the most important beneficiary - from new and future CNS systems, as will be discussed in section 4.4.4, saving a lot of money on aircraft operations. This is the reason why improving the path of transatlantic flights is the major point looked at in the present work.

3.5.4 Africa

Since there is very little radar control over most of the African continent, crews rarely ask for direct-path clearances. Actually, over large areas of Africa where the air traffic control efficiency is sometimes dubious, pilots do the air traffic control on their own by keeping in touch with one another on one specific VHF frequency. Free Flight (see section 4.4.6) would have a lot of potential over Africa, however implementing it could require ground equipment that is for now unrealistically expensive for these countries. Therefore, optimizing flight paths over Africa is for the moment a minor priority.

3.5.5 RNAV

RNAV-restricted air segments are for direct-routings only, and an aircraft must be RNAV-equipped to be eligible for clearance to use them.

Furthermore, according to airlines, experience shows that a direct-routing is not possible without RNAV-equipment when the next VOR station is at more than a hundred nautical miles distance. Since direct paths are commonly considered of interest over distances superior to 200 NM, the RNAV-equipment is seen as a necessity for direct-routings.

Nowadays, the European Civil Aviation Community (ECAC) airspace is on continuous rearrangement to increase the number of direct-routing opportunities. Since January 1998, every aircraft above flight level FL245 in Europe has to be RNAV-equipped.

The RNAV-equipment includes the FMS for which the INS, GPS and DME positions are possible inputs, such that a minimal accuracy of 5 NM is achieved. This allows using imaginary waypoints (RNAV waypoints) that are defined in "lat/long" coordinates.

3.6 Optimization Parameters

The entire section has described a number of parameters that should be taken into account in a real flight-plan optimization. As it was concluded in section 3.4, a much more handy way of considering most of these parameters is to set the goal to meeting the time of arrival. Moreover, other constraints such as the regional ones listed in section 3.5 are extra evidence to prove that an exact model that would consider every parameter individually is much too complicated to elaborate. Actually, considering each of these parameters would certainly lead to very elevated optimization computing times, while the efficiency would be very questionable because nobody knows how to correctly implement each of these parameters individually.

3.7 Route Charges

The route charge concern must be mentioned before ending this section. Route charges are the fees for flying over Flight Information Regions (FIR's). They may differ considerably from one FIR to another. They are not easy to take into account because of the difficulty of implementing the contour of FIR's in the flight-planning tool. They are paid according to a fee per kilometer over each FIR (this fee is the same over an entire country, except in the USA).

These charges sum up over a year to a very big amount, inducing increasingly more airlines to consider the route charges in their Flight Planning system. However, the route charges do not interact in the choice of the catalogue of routes that feeds the Flight Planning system. Indeed, the cost for flying over one FIR may drastically change from one day to the next.

The flights considered in our study are transatlantic, so no attention will be paid on the route charge concern.

Section 4
The Future of Flight Plan Optimization

4.1 Current Non-Feasibility of Dynamic Optimizations

With the current ATM system, the airline usually fixes the flight path optimization criterion (see section 1.4) several hours before departure. It uses the latest weather forecasts and assumes that there is no delay at departure. Once the slot has been allocated by the ATMS and accepted by the airline, the flight plan is set. Later, real-time events sometimes cause the Airline Operating Center to ask the pilot during the en-route phase to change the strategy for the flight - while flying the same route. Such possible real-time occurrences show first-hand that a flight cannot be confirmed to be optimized before departure. Furthermore, a delayed departure causes the airliner to encounter other winds than those used for the route calculation, even if the forecasts were correct (while there always are errors, as is discussed in section 6.3). Consequently, the route is known to be frequently not the optimized one, and, as it was asserted in section 3.4, talking trustingly about flight plan optimization is possible only after the departure slot has been allocated.

Improvements of this pseudo-optimization are difficult to achieve with the current air traffic system. The safety of the flight relies a lot on the controller, and all in-flight route changes occur only upon ATC clearance. A dynamic optimization is therefore not applicable in the current system, and the route is fixed before departure. A few changes can be cleared during the flight, but they are limited in number and frequency so that the pilot and mostly the air traffic controller keep good situation awareness.

4.2 Evolution in Air Transport

Since 1983, the International Civil Aviation Organization (ICAO) has launched a group to define what can be used in the future to improve the existing Air Traffic Control system. The future system, called CNS/ATM (Communication Navigation Surveillance for Air traffic Management), will be deployed by taking into account a transition period for the ground equipment, and the necessity to adapt the existing aircraft fleet.

4.3 The Needs of Air Traffic Management

Air Traffic demand is increasing in all parts of the world. Although rates of growth differ between regions, significant increases in air traffic demand are expected to continue into the next century at a rate of about 6% more aircraft per year. The current demands have already increased the pressure on air traffic service providers and users, straining airspace and airport resources. Without change, the result will be further congestion and delays due to the capacity limitations of today's system, which together with environmental considerations could have significant economic consequences.

This continuously growing air traffic induces the need for new ATM concepts to be defined. One possibility is an extensive automated support, a more decentralized system predicated mainly around free routings and, eventually, autonomous aircraft operations. This is the concept of Free Flight, which leads to the introduction of dynamic airspace structures, in particular Free Flight Airspace (FFAS), and the provision of greater flexibility and free routings to suitably equipped aircraft wherever and whenever possible. The concept requirements involve fundamental changes to current roles such as a more flexible management of airspace.

Consequently, the gradual implementation of the Future Air Navigation Systems (FANS), also called the CNS/ATM (Communication Navigation Surveillance for ATM), will provide the capability for closer interaction between ground-based and airborne systems, as well as between several aircraft.

The International Civil Aviation Organization (ICAO) states approved the global CNS/ATM concept aimed at surpassing the current system limits by making better use of the new technologies that are available.

4.4 The Future Environment

4.4.1 Navigation

The CNS/ATM concept takes advantage of the best of line-of-sight systems as well as satellite communication and Global Positioning System (GPS, used for 4D-positioning and speed determination). GPS can provide a high-integrity, highly accurate navigation service suitable to en-route (and terminal) operations. Airlines will hence benefit from both GPS and improved performance of Inertial Navigation Systems (INS), thereby obtaining a better performance than either would alone. Several FAA-sponsored projects are currently under implementation (e.g., the Wide Area Augmentation System, WAAS) or study (e.g., the Local Area Augmentation System, LAAS), to make full use of GPS.

One non-obvious but useful GPS characteristic is its ability to deliver a common and highly accurate time. This allows perfect synchronization between ground and air elements of air traffic systems. This is very useful notably because the Automatic Dependent Surveillance (ADS) position reports only make sense if they can be referenced to a common time. Time provides the means to move from 3D to 4D-navigation. This makes it possible for ground and air systems to cooperate in the real-time management of flight profiles, for example in ensuring that aircraft arrive at reporting points exactly on schedule, or in handling different approach streams to an airport.

GPS time, positioning, and speed accuracy could lead to a reduction in separation distance minimums between aircraft, thus allowing a more efficient use of airspace. Ultimately, GPS will be the tool that will enable Free Flight (see section 4.4.6).

4.4.2 Surveillance

4.4.2.1 Today

In areas of high traffic density, Secondary Surveillance Radar (SSR) Modes A and C currently provide the main method for surveillance and control of air movements, backed up by primary radar and voice reports on VHF. These are line-of-sight systems. Consequently, for oceanic operations, remote land areas, and airspace where primary and secondary radar are not economically justified, voice reports on HF are used for a procedural service that demands wide separation standards to ensure safety.

4.4.2.2 Future

The key feature of the FANS surveillance concept is Automatic Dependent Surveillance (ADS), a means of extending surveillance service to oceanic airspace, remote land areas, and other areas where radar coverage is not available. Instead of relying on voice position reports, an aircraft operating in non-radar areas will automatically transmit its position (and other relevant data such as aircraft intent, speed, and weather) to the air traffic center via satellite or other communication links. This is the basis for potentially significant enhancements to flight safety by reducing position report errors. Use of ADS, supported by direct two-way pilot-controller communications, will allow non-radar areas to evolve to the point where air traffic services are provided the same way as in today's radar airspace. ADS will support reductions in separation minimums in non-radar airspace, thus alleviating delays, minimizing necessary diversions from preferred flight plans, and reducing flight operating costs. Hence, ADS will support increased ATC flexibility, enabling controllers to be more responsive to aircraft flight preferences. With or without reductions in separation minimums, this flexibility will contribute to cost savings for flight operations.

SSR technology will continue to be used for surveillance in terminal areas and high-density continental airspace.

4.4.3 Flow Management

Benefiting from the new CNS systems, Flow Management is to be based on sophisticated models and databases describing the current and projected levels of demand and resources. This new automation should make it possible to predict the possible sources of congestion and delay and to formulate real-time flow management strategies to cope with demand.

4.4.4 Oceanic Operations

International air traffic is growing much more rapidly than domestic operations. This area of Air Traffic Management should benefit significantly from the new technologies and should experience significant improvements through the next decade. Extensive use will be made of ADS, satellite communications, GPS, weather system improvements, etc., to integrate ground-based ATM into the airborne FMS. The goal is to develop flexible oceanic operations that accommodate the users' preferred trajectories to the maximum extent.

4.4.5 En-Route and Terminal Operations

It is expected that the automated flow management will monitor available capacity and demand at airports throughout the en-route (and terminal) airspace, and will implement strategies to prevent the development of congestion.

4.4.6 Free Flight

4.4.6.1 Concept

Using GPS performance, aircraft should be able in a "near" future to navigate directly from departure to destination while obviating the need for rigid, less efficient airway systems and ground-based navigation infrastructure. This is the idea behind "Free Flight". Under Free Flight rules, aircraft operators - or pilots - would be free to choose their own direct routes, speeds and altitudes. Controllers would only intervene to ensure safe separation, prevent unauthorized entry of Special Use Airspace (SUA) and preclude overloading of airport or system traffic capacity. This concept would allow pilots to operate without specific route, speed, or altitude clearances.

4.4.6.2 Savings

The push to Free Flight is coming directly from the system's heaviest users, primarily the air carriers and large fleet operators who have understood the potential for tremendous time and fuel savings that it could make possible. According to industry analysts, the airlines lose between $3.5 and $5 billion a year because of built-in system inefficiencies. As an example, United Airlines estimated in 1997 that its annual costs for traffic-delayed arrivals, missed connections, long gate-holds and take-off queues, inefficient altitude assignments, and circuitous airway routings added $670 million to its cost ledger alone. Similarly, Delta Airlines calculated in 1997 that it would save $16.8 million a year in fuel if using computer programs that would simultaneously optimize the aircraft's route, speed, altitude for current wind, weather and traffic conditions. The transition to Free Flight will not happen overnight, but the evolutionary process has begun.

The trend to free routing is led by the current use of RNAV capabilities (as discussed in section 3.5.5). Table 1 was realized in 1996 by British Airways to show the potential of RNAV-navigation for fuel savings.

Table 1. RNAV-driven fuel savings

	Distance	**Fuel**
London - Hong Kong		
Flight corridor on 11/1/96	5,922 NM	127.6 tons
RNAV (Great Circle)	5,205 NM	112.8 tons
Saving	12%	12%

London -Los Angeles		
Flight corridor	4,905 NM	116 tons
RNAV (Great Circle)	4,727 NM	111.4 tons
Saving	4%	4%

4.4.7 Dynamic Flight Path Optimization

The analysis of optimization goals made in Section 3 led to the conclusion that the one concern that stood for all of these goals was to meet the scheduled time of arrival (TOA). Virtually, this way, connections, rotations, passengers' mood and other time-dependent flight consequences are all optimized at the same time. In other words, virtually all time-dependent flight consequences rely on the accuracy to meet TOA, which is much more important to the airline than saving fuel - as long as there is enough fuel carried by the aircraft. Moreover, it must be recalled that fuel counts for only about 10% of the flight costs: the most avoidable costs are time-related and incurred by extra flight hours. Minimizing these costs is the main concern that defines the optimization of the flight path.

As explained in section 3.4, taking care of arriving on schedule is not the main concern nowadays, nor is it feasible. Indeed, flight plans are set before departure, while dynamic events such as the weather forecasting errors (which count for the most part) affect the efficiency of the flight plan, hence the ability to reach TOA. Therefore, today airlines mostly concentrate on the effort to meet the time of departure, then try to optimize the path by using, for example, a Cost Index. The Cost Index surely helps the pilot by telling him when cost-efficient changes of Flight Level should occur, however it remains a very complicated notion for a limited efficiency, as discussed in section 2.5.2. This induces unneeded extra costs for a result that is often questioned about its benefits by airlines. Indeed, even if all airlines were to agree on one Cost Index definition, the problem of dynamic events affecting the flight and the time of arrival would still be the same.

Consequently, to ensure meeting TOA with a minor error, it is a necessity to consider real-time optimizations during the flight to dynamically adjust the flight characteristics.

Furthermore, the Cost Index concept should not be used because of the numerous irrelevancies that were discussed in section 2.5. The time-related costs cannot be considered as a linear notion summing different independent costs (see section 2.5.4). Instead, the optimization should simply concentrate on reducing the extra flight hours, which would turn out to be similar to reducing the time-related costs - whatever they are.

Section 5

Discrete Real-Time Flight Path Optimization

5.1 Introduction

The work that is described in the next sections concentrates on meeting the time of arrival (when fuel limitation due to a heavy payload is not an issue) through real-time optimizations. The flight path is chosen with the smallest cruise Mach that satisfies this time of arrival. This ensures the lowest flight cost as the time of arrival is met, and for that flight duration, the fuel consumption is minimized.

As mentioned in section 4.4.6.2, some major airlines (e.g., Delta Airlines) have already studied the potential of dynamic optimization, concluding that this could save them up to billions of US dollars a year. However, these studies have considered the very futuristic case of a perfect Free Flight environment where continuous or very frequent path optimizations are possible in real-time. Such frequent optimizations would highly increase the air traffic controllers' and pilots' workload, as they would find it difficult to keep good situation awareness - unless they have a huge automated back-up system that is not realistic now. Therefore, the ideal situation of frequent optimizations is too futuristic to be paid attention to for the moment.

The present work looks at the efficiency of a realistic transition to Free Flight that could be implemented soon; it consists of using a limited number of flight path optimizations during the flight. Interestingly, this method could prove more useful than continuous optimization for many years. Indeed, the errors on the weather forecasts are today still so important that insisting on the accuracy of the flight plan from departure is not very useful. This point raises interesting issues studied in Section 7, such as how to use the better efficiency of flight path optimization on short paths over long flying distances.

The optimizations performed for the present study use a Dynamic Programming-like algorithm, as opposed to the Great Circle technique that a priori does not provide the best flight plan. Indeed, the second technique consists in defining the best path as the set of waypoints closest to the direct trajectory from departure until destination, but such parameters as the weather forecasts are ignored. Section 6 describes the specific algorithm used for the present work.

5.2 Used Data and Assumptions

5.2.1 Transatlantic Flight Leg

As was explained through sections 1.3.2.1, 3.5.3, and 4.4.4, flight path optimization is mostly interesting for long-haul flights, and one major area where the largest benefits can be obtained is transatlantic flights. Therefore, the present work performs simulations on a typical transatlantic flight, which leaves the south of the United Kingdom in late morning to reach the United States near Boston (MA) in early evening (GMT time).

In fact, little room exists for choosing the Top of Climb (TOC, corresponding to the end of the climb where the flight reaches its Requested cruise Flight Level (RFL)) and the Entry Point (EP) of the Terminal Maneuvering Area (TMA). Among several reasons for this are the numerous constraints inherent to the vicinity of airports and urban areas. Therefore, optimizing the flight path is virtually only significant between TOC and EP. Consequently, the present work performs path optimizations between TOC and EP only. TOC is located in "lat/long" coordinates at $[lat_{TOC}, long_{TOC}]$ = [+49°,0°] and is passed over at about 11:00 GMT (about the same time and position as the real flight); EP is defined by $[lat_{EP}, long_{EP}]$ = [+40°,-70°] and its "requested time over" is set at 18:45 GMT. This corresponds to a little longer flight duration than in the real situation, but it is acceptable because the average speed is slower when the aircraft stays at a constant flight level.

5.2.2 Flight Level

Whether the optimizations are performed for the path, Mach, and Flight Level choices, or simply for the first two, considerably varies the complexity of the model as well as the simulation running time. Besides, the results that may be obtained by acting on the path and Mach choices only are sufficient to understand the consequence on the time of arrival of discretized real-time optimizations - the purpose of this work. This is why the following important assumption is afforded in this study: the flight remains from TOC to EP on a single flight level. Most simulations will be run on flight level FL300 - a typical value.

5.2.3 Weather Data

5.2.3.1 Extreme Error Analysis

Nair and Forrester[13] found that gross errors in wind models are more common than what would be expected from a normal distribution. On a three-month winter period, they made approximately fifty reports of errors in the wind speed magnitude that exceeded 30 meters per second. Most of these reports were due to unrealistically large wind speed measurements reported from a couple of aircraft on just a few flights. Moreover, errors in the headwinds above the Atlantic in wintertime (when forecasting errors are at their peak) range up to about 7 meters per second. It should be noted however that Nair and Forrester agree that the error distributions in both wind speed and wind direction conform closely to standard normal distributions.

To be more realistic, it was decided in the present work to use real weather forecasts rather than normal distributions.

5.2.3.2 Weather Forecasts, True Weather

Professor Illari from MIT's department of Earth, Atmospheric and Planetary Science provided weather forecasts and "true weather" grids between longitudes -90° and 0° with a 2.5°-step, and between latitudes +20° and +70° with a 2.5°-step. This area covers a large part of the Northern Atlantic where most transatlantic flights occur. The grids contain information about the horizontal wind speed components and the temperature at different pressure levels (corresponding to flight levels). The "true weather" data are the analysis performed by MIT's department of Earth, Atmospheric, and Planetary Science, providing the actual weather a few hours after they have occurred.

It should be noticed that it is reasonable to use these grids for our simulations because the 2.5°-steps provide a good approximation to the 1.5°-steps actually used by air traffic systems.

For each simulated flight, the following weather data were used:

- the 6-hour weather forecasts released at 0:00 GMT,
- the 12-hour weather forecasts released at 0:00 GMT,
- the 6-hour forecasts released at 12:00 GMT,
- the 12-hour forecasts released at 12:00 GMT,
- the "true weather" for 0:00 GMT,

- the "true weather" for 12:00 GMT,
- the "true weather" for next-day's 0:00 GMT.

5.2.4 Waypoints

The set of waypoints considered in the simulations is disposed similarly to the one used nowadays for actual transatlantic flights. They are defined by their "lat/long" coordinates every 1° of latitude and 10° of longitude as they are represented in Figure 2. For real flights, more waypoints are located near coasts, but we will not consider them in the simulations because - a priori - they affect less the time at destination than the choice of waypoints over the Atlantic does.

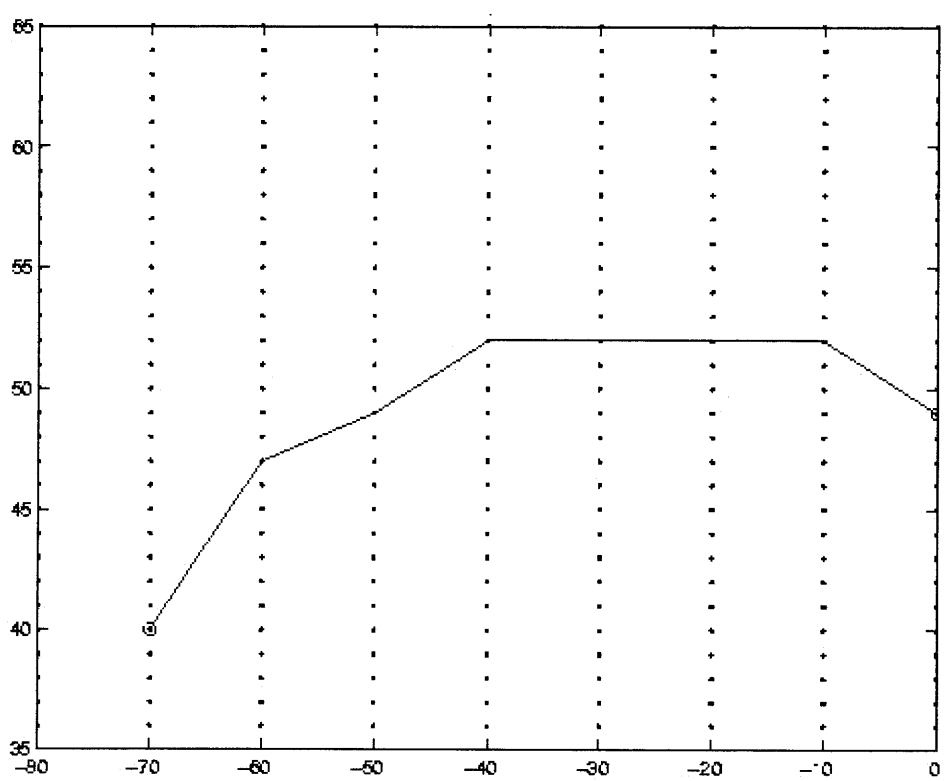

Figure 2. Waypoints database

Each dot represents a waypoint; the coordinates are degrees of latitude (horizontally) and degrees of longitude (vertically). The circles refer to the Top of Climb (right circle, located south of the UK) and the Entry Point (left circle, located near Boston, MA, USA). The plain line is the computed optimal flight path from TOC to EP (determined at departure) on one particular day.

5.2.5 Linear Motion

5.2.5.1 Turns

Large angles of turn at cruise speed are unusual. They will become even rarer with the advent of free flight. As it is suggested for long-haul flights in Reference 5, turns were first considered for the simulations to be at a rate of 1° per second.

Let us call θ the accumulated course change by the aircraft between the Top of Climb and the Entry Point, M the Mach, a the speed of sound, and TR = 1°/s the turn rate. We find that the travel distance (aside from wind effects) flown during turns is given by:

$$s = \frac{M \cdot a \cdot \theta}{TR} \qquad (5.2.5.1.1)$$

If we then consider the case where turns are performed over waypoints only, immediately heading to the next waypoint, an easy calculation shows that the distance flown during what should have been turns is:

$$l = 2\frac{M \cdot a}{TR}\tan\left(\frac{\theta}{2}\right) \qquad (5.2.5.1.2)$$

Considering the very bad example where θ = 90° and M is the "accelerated Mach" 0.82, we obtain that the flight duration difference between the two cases is of about 25 seconds (corresponding to about 5,040 meters). Compared to the scheduled flight duration from TOC to EP of 7:45 hours, 25 seconds is a mere 0.09%. Moreover, a 25-second difference from schedule is not noticed. As a result, we will make the approximation for the simulations that turns are performed immediately over waypoints.

5.2.5.2 Between Turns

The reasoning used in the previous paragraph would similarly prove the credibility of the assumption that motion between two consecutive turns can be considered linear for the simulations (and in fact, pilots really try to fly linearly).

5.3 Flight Path Optimization

5.3.1 Waypoints Search Space

5.3.1.1 An Elliptic Search Space

The optimized path is searched among all the contiguous waypoint combinations from TOC to EP. Of course, the computing time evolves exponentially with the number of waypoints. Since it is obviously useless to consider some waypoints in the path-search algorithm, a first concern is to reduce the flight path search space. One analytically practical and natural way of achieving this is to consider only the waypoints within an ellipse, whose foci are the Top of Climb and the Entry Point. Several flight-planning systems used by actual airlines do the same.

5.3.1.2 Pujet's Ellipse

5.3.1.2.1 Proof of Concept

Pujet[6] found a way to analytically justify the reduction of the search space to an ellipse. He explained that a waypoint M cannot be on the optimized path if the direct routes TOC-to-M and M-to-EP at speed ($V_{max}+w_{max}$) bring the aircraft at destination later than the Required Time of Arrival (RTA). V_{max} is the maximum true air speed in the aircraft envelope, and w_{max} is the maximum wind velocity in the entire weather grid; the sum of the two represents the highest ground speed magnitude reachable by the aircraft. In other words, what Pujet means is that any acceptable waypoint M for the search space follows:

$$dist(TOC, M) + dist(M, EP) \leq t_{\max}(V_{\max} + w_{\max}) \qquad (5.3.1.2.1.1)$$

where *dist(A,B)* is defined as the shortest distance around the Earth between points A and B, and t_{max} is the maximal flight duration determined by the Required Time of Arrival (RTA).

Relation (5.2.1.2.1) describes the inside of an ellipse, whose eccentricity e is given by:

$$e = \frac{dist(TOC, EP)}{t_{\max}(V_{\max} + w_{\max})} \qquad (5.3.1.2.1.2)$$

5.3.1.2.2 First Eccentricity Improvement

The search ellipse can be slightly reduced by using a smaller V_{max} to compute the eccentricity value. Since the weather grids that are used in the simulations include temperatures, if T_{max} is the highest temperature in the search space, then we can consider that V_{max} is the maximal speed magnitude relatively to the wind. It is then given:

$$V_{\max} = M_{\max}\sqrt{\gamma R T_{\max}} \qquad (5.3.1.2.2.1)$$

where M_{max} is the maximal acceptable Mach, and γR is classically approximated as the following constant:

$$\gamma R = 223.2 \text{ m}^2/(\text{K.s}^2) \qquad (5.3.1.2.2.2)$$

5.3.1.2.3 Insufficient Search Space Reduction

The assumption was made in section 5.2.2 that the motion of the aircraft was on a constant flight level. Compared to the Earth's radius, there are small altitude differences on a same flight level. This allows the following approximation to be made in this study, giving the distance between two positions A and B on a single flight level:

$$dist(A,B) = \alpha \cdot (R_E + alt_{FL}) \qquad (5.3.1.2.3.1)$$

where α is the angle at the center of the Earth between A and B, R_E is the Earth's radius (R_E = 6,378 km), and alt_{FL} is a typical altitude on the considered flight level.

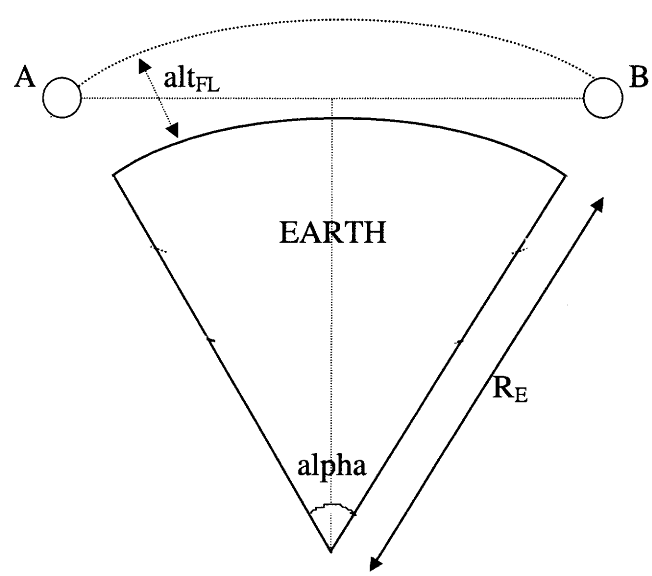

Figure 3. Calculating distances on a single flight level

This figure proves the following relation:

$$\|AB\| = 2 \cdot (R_E + alt_{FL}) \cdot \sin\left(\frac{\alpha}{2}\right) \qquad (5.3.1.2.3.2)$$

On relatively short distances, such as between two contiguous waypoints, the sine function can be approximated by the identity function. However, the approximation is no longer valid when calculating much longer distances such as between TOC and EP.

By working in spherical coordinates with origin the center of the Earth, we obtain the practical approximation:

$$\|AB\| = (R_E + alt_{FL}) \cdot \sqrt{2} \cdot [1 - \cos(L_B - L_A) \cdot \cos l_A \cdot \cos l_B - \sin l_A \cdot \sin l_B]^{1/2} \qquad (5.3.1.2.3.3)$$

where L_M and l_M respectively refer to the latitude and longitude of point M.

Equations (5.3.1.2.3.1), (5.3.1.2.3.2) and (5.3.1.2.3.3) combine into:

$$dist(A,B) = 2(R_E + alt_{FL}) \cdot a\sin\left\{\frac{1}{\sqrt{2}}[1 - \cos(L_B - L_A)\cos l_A \cos l_B - \sin l_A \sin l_B]^{1/2}\right\}$$
(5.3.1.2.3.4)

To take an example, if we assume in our simulations that the flight level is FL300, then $alt_{FL} \sim$ 10,000 m, and (5.3.1.2.3.4) provides:

$$dist(TOC,EP) = 5455 \text{ km}$$

The weather database on a typical day (e.g., March 6, 1998) provides w_{max} = 78.1 m/s, T_{max} = 243.8 K and t_{max} = 8 hr. Then (5.2.1.2.2) gives V_{max} = 191.3 m/s = 371.9 kt, and equation (5.3.1.2.1.2) eventually reveals the value of the elliptic search space eccentricity:

$$e = 0.703$$

This value is possibly useful for medium-haul flights, since these are subject to numerous regulations (e.g., having to avoid Special Use Airspaces), which incur a need for big search spaces. However, using a 0.703-eccentricity is not interesting for transatlantic flights. The corresponding elliptic search space is huge and barely reduces the full original search space, as is seen on Figure 4.

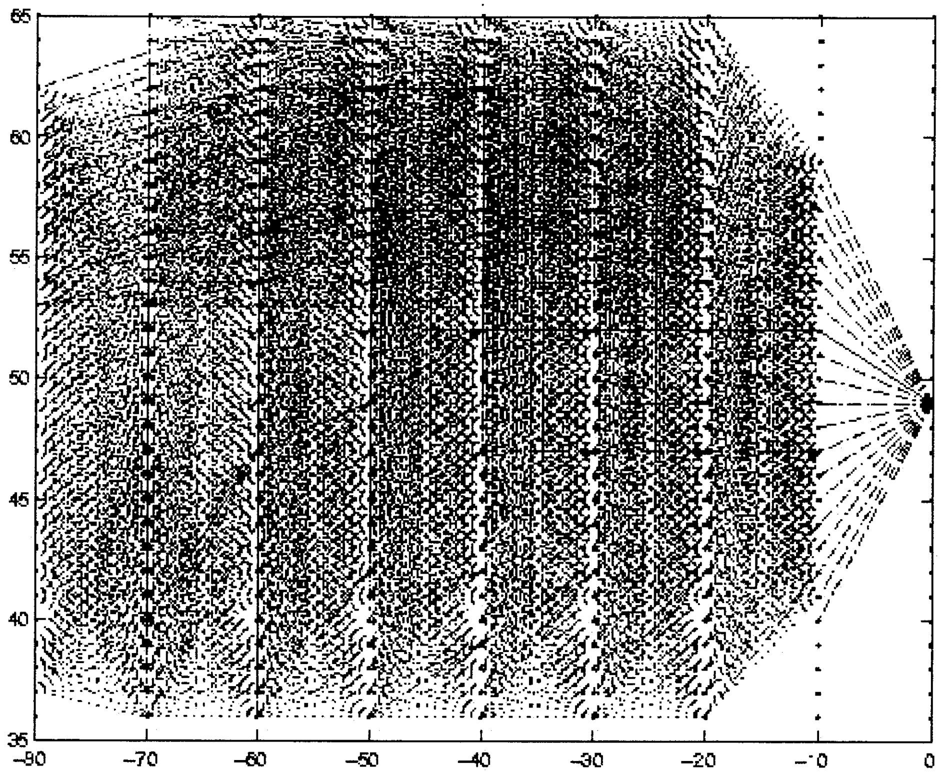

Figure 4. Projection of the search space ellipse in the "lat/long" plane for e = 0.703
All possible paths starting at TOC and verifying the criteria described later are shown. The thick plain trajectory is the optimal one (forecasted at departure).

Even if the 0.703-value was analytically proven to reduce the size of the search space, the ellipse still contains numerous waypoints that would definitely never be part of the optimized route. What this eccentricity value does instead is to increase the optimization running time because it takes more time to validate the presence of waypoints inside the ellipse than to simply consider all of them in the path optimization algorithm.

It should be noted that reducing the search space by simply affording a smaller Required Time of Arrival (RTA) is useless. Even if RTA were equal to the scheduled time of arrival over

EP (this would provide the minimal RTA), the eccentricity of the previous example would be equal to 0.775. This defines an ellipse very similar to the 0.703-eccentricity case.

In conclusion, in practice, when an airline uses an elliptic search space, the eccentricity value is empirically chosen. One method for setting it is to start from a value close to 1, then to decrease it until the optimized path lies well inside the ellipse. For the example given above, e = 0.94 provides an all-the-time acceptable size of search area for long-haul flights (an example is seen on Figure 5). It can be noticed that the obtained elliptic search space is very similar to the one actually used by Air France[6], shown on Figure 6 (corresponding to e = 0.92).

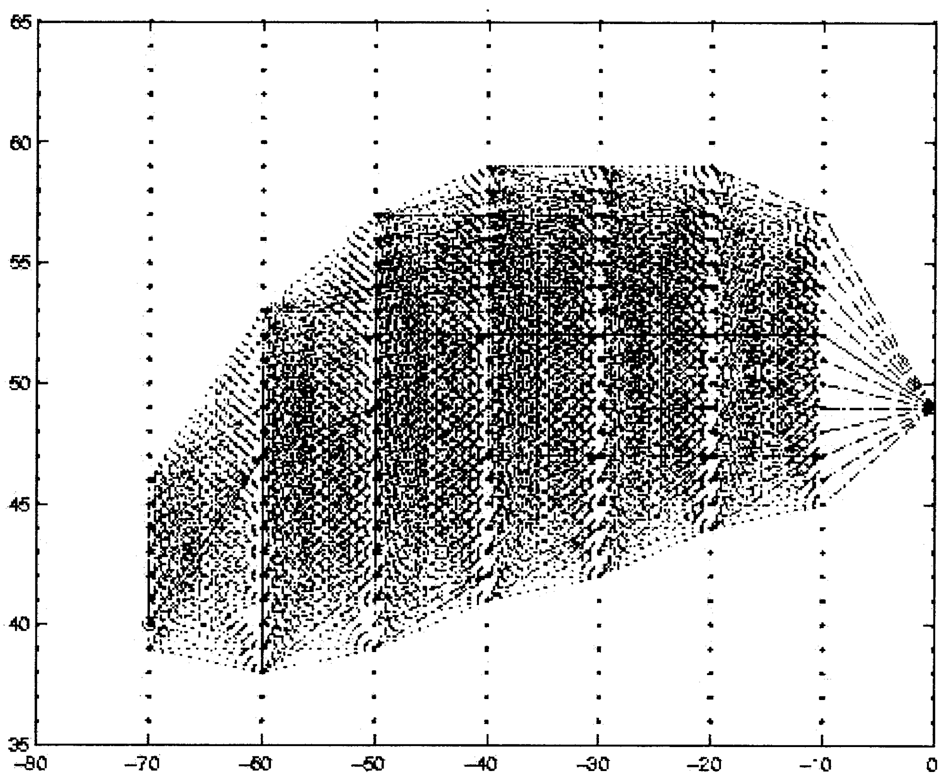

Figure 5. Projection of the search space ellipse in the "lat/long" plane for e = 0.94

All possible paths starting at TOC and verifying the criteria described later are shown. The thick plain trajectory is the optimal one (forecasted at departure).

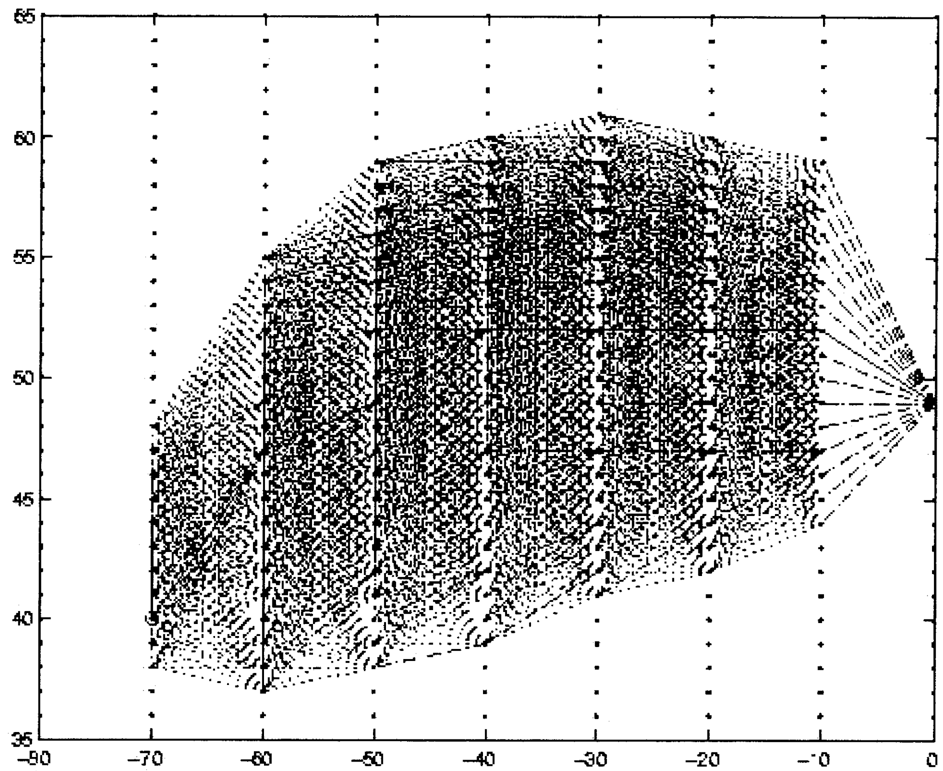

Figure 6. Air France's search space ellipse for the same flight leg as in our study

Actually, even $e = 0.94$ is probably an unnecessarily low eccentricity value, as in our example, the forecasted optimal flight path is still found with $e = 0.994$ (Figure 7).

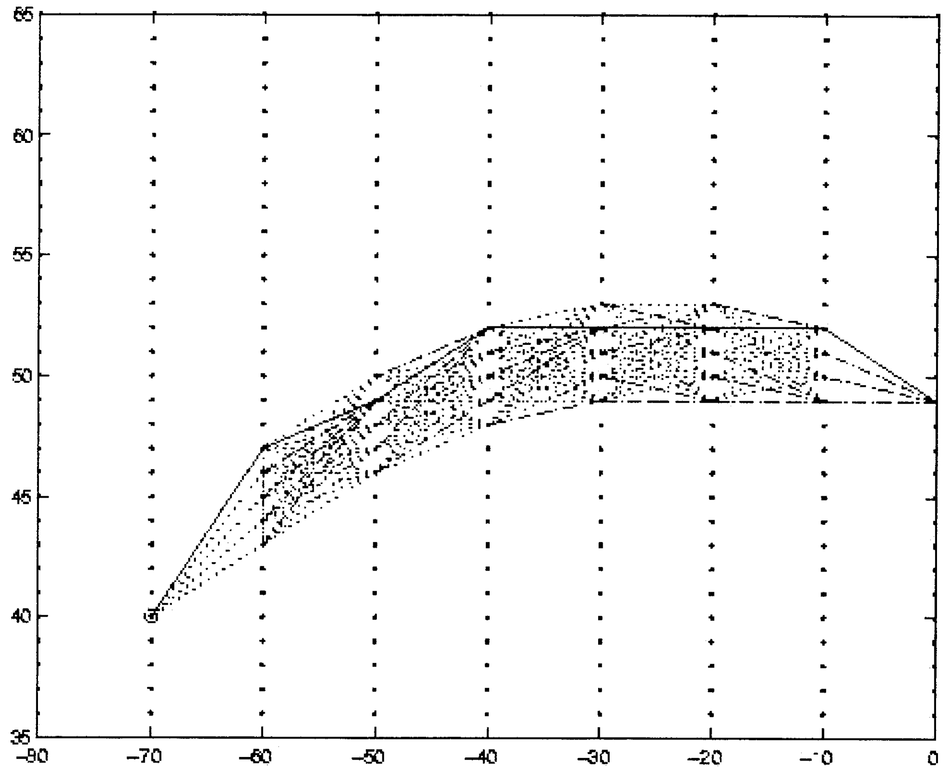

Figure 7. *e* = 0.994
The plain line is the optimal path.

For the present study, several algorithms - inspired from Pujet's idea - were thought of to methodically diminish the search space size better than the $e = 0.703$ case. Unfortunately, none managed to significantly improve the eccentricity value, and the size of the search space remained huge with every attempt. The reason is most probably that the weather data grids are too large to afford being used quickly and easily to reduce the search space dimension.

Further work - not done here - could be performed on analytically proving a useful eccentricity value.

5.3.2 Waypoint Pairs

This paragraph describes the algorithm used in the simulations to detect when waypoint B can follow waypoint A on a route. B is considered "potentially after A" when:

- the angle at the center of the Earth between A and B is less than 12 degrees. This way, there are numerous possible B choices, but B is never - with our set of waypoints - at more than one longitude step from A (2.5° longitude). This allows all pairs of contiguous waypoints used in actual flight-planning systems for calculating transatlantic Random Routes (see paragraph A.1.4) to be comprised. For example, Air France looks at few possibilities for waypoint B and only keeps in its database the four that are inside the search space ellipse and that bring the least cost from A to B. The 12-degree criterion used in the present work leads to more than fifteen B possibilities, and those used by Air France lie well within. This ensures to not reduce this part of the efficiency of the path optimization compared to actual flight-planning systems. Nevertheless, a 12-degree criterion is unlikely to ever be used in a real flight-planning tool because it considers numerous waypoint pairs that would definitely never be flown during safe procedures. Much worse, this method has the huge drawback to considerably increase the computing time, as this one increases exponentially with the number of considered waypoint pairs. Some further work - not done here - could study the efficiency of looking at only 3, 4, 5, or n possible next waypoints.
- B can be reached from A only if B is closer to the Entry Point (EP).
- If A and EP are on the same latitude, then B is necessarily EP.

5.3.3 Reachable Waypoints

Having obtained all the possible "pairs", a waypoint W shall be said to be reachable from TOC if there exists a chain of waypoint pairs from TOC to W. Figures 4, 5, and 6 show all the reachable waypoint pairs for different eccentricity values of the search space ellipse.

5.3.4 Best Paths Computation

An algorithm that finds the best path from the Top of Climb to any reachable waypoint is then run. The path is calculated for a constant Mach as it is done in actual flight-planning tools. A better Mach choice is later (see section 5.3.6) easily obtained by just slightly increasing the computing time and without running the complete optimizing algorithm again with another Mach (which is what real flight-planning software do).

5.3.4.1 Best Path Definitions

Referring to the conclusions in sections 3.4 and 4.4.7, the optimized path shall be considered as the one that allows reaching destination the closest to the scheduled Time of Arrival, with the smallest Mach that affords this.

Moreover, the quickest trajectory that brings the aircraft from the Top of Climb to a reachable waypoint at a particular Mach shall be called the best path to this waypoint at this Mach. For each Mach we are working with, a Dynamic Programming-like algorithm (detailed in Section 6) will find the quickest path from TOC to any "reachable waypoint" (as defined in section 5.3.3) of the search space. Consequently, this path will be the best path towards this waypoint at the used Mach.

5.3.4.2 Best Path Determination

5.3.4.2.1 Two-Step Process

As explained in the previous paragraph, the simulations run for this study consider that the best trajectory up to the Entry Point (EP) can be found in two steps:

1) for the usual Mach on this flight leg, the fastest path up to EP is calculated,

2) then, on the same route, a binary search determines a Mach that allows arriving very close to the scheduled time of arrival.

The simulations confirmed the relevancy of this method. Always, deriving the Mach this way, or directly setting it to this value before calculating the best path, both lead to the same forecasted optimal path. This is even more true when the optimization is run during the flight, because the path that is left to fly is shorter, such that there are fewer route options. However, it must be understood that this method might have a different efficiency with a denser grid of waypoints.

5.3.4.2.2 Algorithm

Once the "reachable waypoints" have been listed (see section 5.3.3), the following algorithm derives the optimal path from the Top of Climb (TOC) to any reachable waypoint:

1) for each reachable waypoint B, the list $(A_i)_{i=1,2,...n}$ of waypoints possibly right before B on the flight route is determined. The A_i's are the waypoints that are paired with (see paragraph 5.3.2), and are "anterior to", B.

2) The minimal flight cost from TOC to B via a specific waypoint A_i - right before B - is the sum of the minimal cost from TOC to A_i and of the cost from A_i to B.

$$BestCost(TOC, A_i, B) = BestCost(TOC, A_i) + Cost(A_i, B) \qquad (5.3.4.2.2.1)$$

Consequently, if the best path towards each A_i right before B is known, then the best trajectory from TOC to B is the one with the lowest sum of costs from TOC to A_i and A_i to B. This technique is inspired by the Dynamic Programming algorithm.

$$BestCost(TOC, B) = \inf_i \{BestCost(TOC, A_i, B)\} \qquad (5.3.4.2.2.2)$$

3) If the best path to one A_i has not been determined yet, then the program applies step 1) to A_i. It will get back to waypoint B only when the best path to A_i has been determined.

An example of the obtained "best paths" is seen on Figure 8.

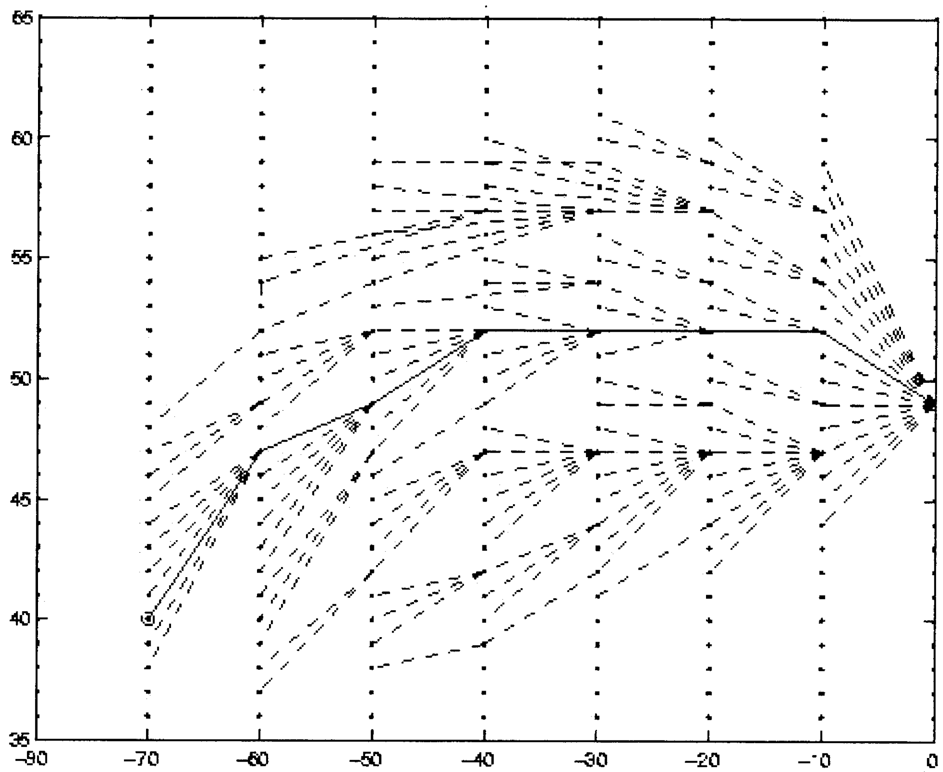

Figure 8. Best paths from TOC to any reachable waypoint. *e*=0.92.

The dashed lines represent the best paths from TOC to each point in the search ellipse. The plain line is the optimal path.

Interestingly, a 0.994-eccentricity provides the "best paths" plot shown on Figure 9. In this case, the optimal path from TOC to EP barely fits inside the search space ellipse. This indicates that one should not rely on this eccentricity value to compute an optimal route.

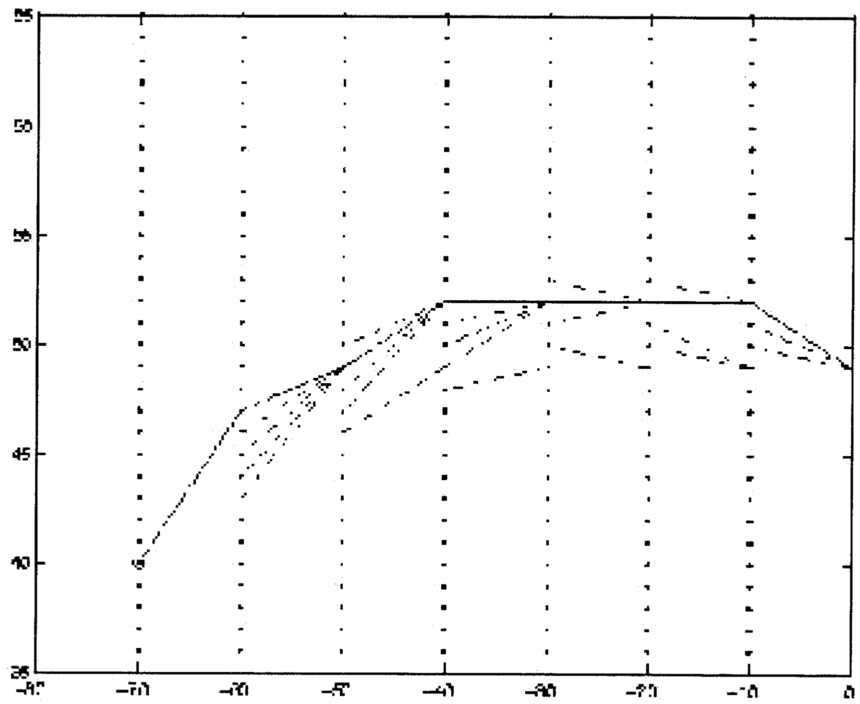

Figure 9. Best paths, *e* = 0.994.
The plain line is the optimal path.

5.3.5 Cost Function

To analytically translate the goal that was explained in section 5.3.4.1, the Cost Function to minimize in the optimization problem is defined as the trip duration. Then, once the best path – according to this cost function and one particular Mach - has been obtained, the Mach is adjusted to meet the Time of Arrival, or at least to arrive with the minimal delay.

The Cost Function having been defined as the flight duration, it is the sum of the flight duration over each route segment. A segment is delimited by two consecutive waypoints (section 5.2.5.1 proved that motion could be assumed from top-of-waypoint to top-of-waypoint).

5.3.5.1 Segment Discretization

Over each segment (whose linear assumption was justified in section 5.2.5.2), a number n of equidistant discretizations, defining n sub-segments, is done to favor a more adequate use of the weather data. The weather information is assumed constant over each sub-segment, equal to the value at one bound. The n value did not affect the results of the simulations by much, yet a safe $n = 10$ was adopted as it seemed a very reasonable choice. However, a smaller value could be chosen to reduce the computing time.

5.3.5.2 Tri-Linear Extrapolation

The weather information at the beginning of a sub-segment is calculated via a tri-linear extrapolation of the weather data grids along the East-West, North-South, and vertical axis. This is in conformance with what actual flight-planning systems do. The tri-linearization is performed with the corners of the smallest paralleloid that contains the sub-segment (the corners are points of the weather grids), not taking into account the influence of the weather external from the paralleloid.

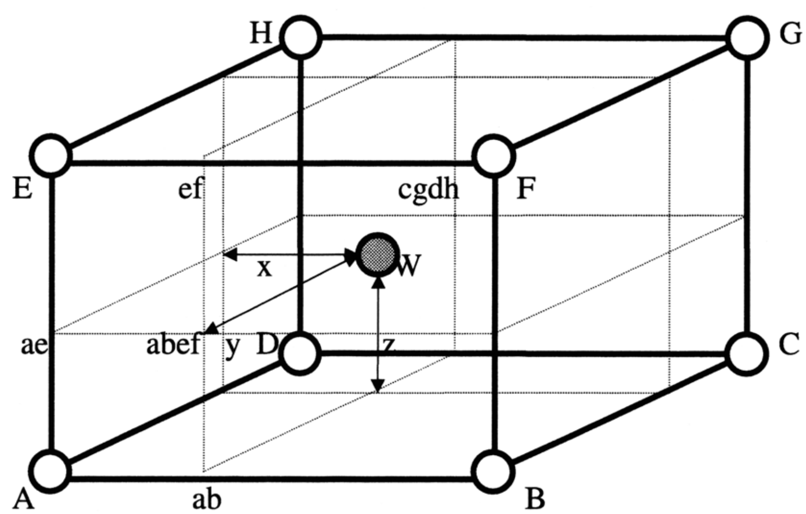

Figure 10. Weather data tri-linearization
x, y and z are normalized.

By linearization, we have (for example):

$$WeatherData(ab) = (1-x)\cdot WeatherData(A) + x\cdot WeatherData(B) \quad (5.3.5.2.1)$$

By bi-linearization, we then obtain (for example):

$$WeatherData(abef) = (1-z)\cdot WeatherData(ab) + z\cdot WeatherData(ef) \quad (5.3.5.2.2)$$

Finally, by tri-linearization, we conclude:

$$WeatherData(W) = (1-y)\cdot WeatherData(abef) + y\cdot WeatherData(cgdh) \quad (5.3.5.2.3)$$

5.3.6 The Optimal Mach

5.3.6.1 Take-Off Weight Limited Knowledge

The weather uncertainties (described in section 6.3) understandably limit the efficiency of the Mach choice to meet the scheduled time of arrival (see the simulations results in Section 8). Next to this, airlines have at least one other important reason for not seeking the exact optimal Mach value. It is that there is limited knowledge about the take-off weight three hours before departure when the flight-planning system tries to compute the "best" route. At this time, the take-off weight (and the fuel to refill the tanks with) is estimated by using:

- the number of currently booked tickets,
- the number of passengers who have already checked in,
- the cargo weight that has already been checked in,
- an estimate of the amount of fuel left onboard from the previous flight,
- prior take-off weight statistics for that route.

These parameters are used two to three hours before the flight, when it is probable that many passengers have not checked in yet. Moreover, it happens often that the pilot demands more fuel to be carried onboard to ensure a safe fuel margin in case of "holding", rerouting, or other important unpredicted occurrence. This happens one hour before the flight, while the passengers and cargo mass is still not known. As a conclusion, the aircraft's take-off mass used for the flight-path optimization is only based on statistics. This goes against what E. Hoffman proved in Reference 7: accurate knowledge of the take-off mass is mandatory to compute a flight-plan that is really optimized.

Therefore, optimizing a flight path can first take its meaning only shortly before departure, when the take-off mass can be guessed with the best accuracy.

This conclusion inspired the present work to first look at the efficiency of running a path optimization during the climb phase (see section 6.1). Some airlines have manifested their interest in this idea: they see in it a first step towards real-time optimizations, and they are curious about the benefits of this technique. However, E. Hoffman insists that the ability to correctly estimate the aircraft's mass, even during the flight, is very much questionable. This explains why the airlines' interest in looking at a new optimization during the climb phase is much more due to an interest in looking at the benefit of using last-minute weather forecasts.

5.3.6.2 Best Mach Determination

The best route from the Top of Climb to the Entry Point was determined in section 5.3.4.2.2 with a certain Mach that was statistically estimated for the flight leg in question, probably close to the optimal Mach value. Once this route has been calculated, a binary search is run until it finds a Mach that allows the aircraft to reach the Entry Point at less than one minute from the scheduled time of arrival, on this same path. We discussed in section 5.3.4.2.1 the relevancy of this method for finding the optimal Mach. The important benefit of making a binary search (rather than running the route optimization for a whole set of Mach values) is a very fast computing time. Moreover, the one-minute-threshold corresponds to a realistic acceptable difference from schedule.

Of course, if the binary search exceeds the Mach range, the Mach is set to the nearest boundary value (the Long-Range Cruise Mach or the Max Range Cruise Mach).

Section 6

Optimizations in Future Context

6.1 Optimization at the Top of Climb

As mentioned in section 5.3.6.1, some major airlines are curious about the benefit of running a new flight-plan optimization during the climb phase with updated weather forecasts.

Therefore the present study first compares the efficiencies of running a single optimization three hours before departure - as is currently done - and of running another one at the Top of Climb (TOC) with last-minute updated weather forecasts. The two strategies may truly differ in their results as the flight may notably have suffered an important delay at departure - for example. These results are shown in section 8.1.4.

6.2 Dynamic Optimization

Farther in the perspective of Free Flight, it is interesting to investigate the effectiveness of real-time flight-plan optimizations - during the en-route phase. The present work focused on the efficiency of optimizations periodically performed in real-time. Another study could look at the potential benefit of realizing optimizations at non-periodic times, if not non-correlated ones.

Each real-time optimization computes the best path from the next waypoint until the Entry Point (the trajectory is still assumed from waypoint to waypoint). Choices had to be made concerning the weather forecast periodicity, which would certainly have in practice an impact on the optimization frequency. The decision held for this work was that each optimization would be run at the weather forecast update rate. Currently, the weather forecasts are given every 12 hours for 6, 12, 18 and 24 hours later (closer forecasts are not possible yet, but will become so in the future with faster computers). The 6-hour and 12-hour forecasts (i.e., for 6 and 12 hours later) were used along all the simulations. Then the efficiency of switching to 3-hour forecasts (i.e., 3, 6, 9…) or 1.5-hour forecasts (i.e., 1.5, 3, 4.5, 6, 7.5…) was investigated.

6.2.1 New Weather Forecasts

To use weather forecasts that are more frequent in the simulations, a realistic degree of uncertainty for the new forecasts had to be modeled. Hereafter is the algorithm that was adopted for the present work. It is based on a multi-linear extrapolation of the current weather status and forecasts, with a realistic amount of uncertainty that is obtained via linear extrapolation.

First of all, let us assume there exists a parameter $a_p \in [0,1]$ that represents the level of accuracy of the new forecasts given at time t, such that:

$$forecast(t) = (1 - a_p) \cdot forecast_0(t) + a_p \cdot actual(t) \qquad (6.2.1.1)$$

where:

- *forecast(t)* is the weather forecast matrix at time t,
- *forecast$_0$(t)* is the linear extrapolation at time t of the 6 and 12-hour weather forecast matrices that correspond to the boundaries of the 6-hour time window containing t ([6:00GMT,12:00GMT], [12:00GMT,18:00GMT], or [18:00GMT, 24:00GMT]),
- and *actual(t)* is the actual weather data matrix.

In other words, *forecast()* is barycenter of *forecast$_0$()* and *actual()*. The 6-hour and 12-hour forecasts are given for times t_0 and t_1 ($t \in [t_0, t_1]$). a_p is equal to 1 when the forecast and the actual weather are identical, and is equal to 0 when the forecast is the one at time t_0.

Similarly, let us assume that there exists a function a_f of time t, with $a_f(t) \in [0,1]$, such that the new forecast $forecast(t)(t_1)$ given at time t of the weather at time t_1 is defined by:

$$forecast(t)(t_1) = (1 - a_f(t)) \cdot forecast_0(t_1) + a_f(t) \cdot actual(t_1) \qquad (6.2.1.2)$$

where $forecast_0(t_1)$ is the actual 6-hour (or 12-hour) forecast at time t_1.

It seems reasonable to first-degree that $a_f(t)$ is a linear function of t. Therefore, by using the following initial conditions:

$$a_f(t=t_0) = 0 \qquad (6.2.1.3)$$

$$a_f(t_1) = a_p \qquad (6.2.1.4)$$

we obtain:

$$a_f(t) = \frac{(t - t_0)}{t_1 - t_0} \cdot a_p \qquad (6.2.1.5)$$

Then, relations (6.2.1.2) and (6.2.1.5) combine together to provide the forecast at time t of the weather at time t_1. The result is a function of t and a_p.

Moreover, the forecast at time t of the weather at any time in $[t,t_1]$ is given by linear extrapolation between the forecast at time t of the weather at time t, and the forecast at time t of the weather at time t_1.

6.2.2 Flight Path Stability

In the context of dynamic optimizations, the flight path is considered stable when:

- it is subject to only small modifications if a new optimization reduces only slightly the cost function value,
- and a route modification must occur only if it brings substantial benefit (a substantial decrease of the cost function).

Stability is very important to ensure comfortable situation awareness for both the pilot and the controller.

The cost function used in the present work is the flight duration (see section 5.3.4.1). This has turned out through the simulations to bring a very good stability. More precisely, almost no real-time optimization modified the remaining flight route (instead, these optimizations adjusted the Mach value to meet the scheduled time of arrival). A few restrictions, such as allowing modification of the route only if it brings more benefit than a certain threshold, can be used to ensure stability; the simulations of the present study did not need them. Nevertheless, in the context of denser grid of waypoints, it is expected that restrictions would become of interest.

6.3 Wind Uncertainties

6.3.1 Impact of Weather Turbulence on Time of Arrival

Certain scientists showed their interest in investigating the impact of turbulence on the efficiency of a Flight Plan optimization. Actually, the mean value of the turbulence wind speed being near zero, it is expected that the effect on the time at destination is not important (assuming the aircraft does not reduce Mach). This was verified throughout the simulations, as we will see in section 8.1.2.

6.3.2 Turbulence Model

6.3.2.1 Dryden Model

The turbulence models that are frequently used to evaluate aircraft performance are the Dryden and von Karman models. The Dryden model, which is used much more frequently, was the one retained for the present work; it has indeed a simpler form and can be easily implemented by passing white noise through linear filters. TR. Beal[9]'s algorithm for generating turbulence based on the Dryden model is described in the next paragraph.

Note that the vertical turbulence effects are not considered, as the present work is limited to a two-dimensional study (on a single flight level).

6.3.2.2 Beal's Algorithm

- Let u and v be the turbulence velocities along the x and y axes of the aircraft, supposed to remain on a single flight level,
- let σ_u and σ_v be the standard deviations of u and v,
- and let L_u and L_v be the scale lengths for power spectra.

The en-route flight levels that are considered in the simulations are above 2,000 ft. As it is recommended in Ref.10, this allows the following typical assumptions to be made: $\sigma_u = \sigma_v$ and $L_u = L_v = 1{,}750$ ft.

For the spectral densities of the filtered signals to correspond to those of the Dryden spectra, the standard deviations of the unfiltered signals must be as follows:

$$\sigma_{wn} = \sigma_u \cdot \sqrt{\frac{2L_u}{Dx}} \qquad (6.3.3.1)$$

for the u-velocity filter, and:

$$\sigma_{wn} = \sigma_v \cdot \sqrt{\frac{2L_v}{Dx}} \qquad (6.3.3.2)$$

for the v-velocity filter, where Dx is the integration step.

With Matlab ®, the uncorrelated random discrete functions x_u and x_v are generated such that their standard deviations respectively correspond to (6.3.3.1) and (6.3.3.2). Then, Beal's algorithm provides discrete sets for u and v:

$$u_i = A_u u_{i-1} + 2B_u (x_u)_{i-1} \tag{6.3.3.3}$$

with:

$$R_u = \frac{Dx}{2L_u} \qquad A_u = \frac{1-R_u}{1+R_u} \qquad B_u = \frac{R_u}{1+R_u} \tag{6.3.3.4}$$

and:

$$y_i = A_v y_{i-1} + 2B_v (x_v)_{i-1} \tag{6.3.3.5}$$

$$v_i = A_v v_{i-1} + B_v (y_i + y_{i-1}) + C_v (y_i - y_{i-1}) \tag{6.3.3.6}$$

with:

$$R_v = \frac{Dx}{2L_v} \qquad A_v = \frac{1-R_v}{1+R_v} \qquad B_v = \frac{R_v}{1+R_v} \qquad C_v = \frac{\sqrt{3}}{1+R_v} \tag{6.3.3.7}$$

Using these relations, the turbulence effects are incorporated in the weather data by appropriately adding the random u and v to the average wind speed coordinates.

Section 7

Formulation of a Linear Program to Fly 4D

7.1 Introduction

A lot of research (notably done by Eurocontrol) has been performed on trying to build "4D-trajectories" for approaches in Terminal Maneuvering Areas (TMA's). This concept literally consists in seeking that the aircraft passes over a given set of waypoints at precise given times. This enables the ATC to deal more easily with air traffic management in TMA's. One subsequent issue is whether this concept would be of interest outside TMA's. Notably, 4D-trajectories for the entire route could provide airlines with a more comprehensible reaction to airspace congestion, connections to get at a destination airport, and other cost-inducing factors to take into account in real-time. However, an obvious consequence of using a 4D-strategy is that the Mach can no longer be considered a constant. Moreover, due to the weather uncertainties (see paragraph 6.3), the 4D-strategy is likely to burn more fuel than would be the case with less stringent time constraints.

The previous sections were focused on the scheduled time of arrival only. Now, it is interesting to study the consequences of imposing several intermediary times to meet. The current section is a first approach towards a strategy that affords 4D-routing while minimizing the difference between the forecasted fuel-cost for the trip and the actual one.

7.2 Assumptions and Goals

7.2.1 Optimizations Goal

The 4D-trajectory to follow is assumed to be determined before departure. The simulation goal is to find how to minimize the fuel consumption while passing on top of each waypoint at a time very close to the corresponding Estimated Time Over (ETO, time constraint that is set before departure).

More exactly, to match real needs, the concern of this study is to minimize the impact of wind uncertainties on the optimized fuel cost (defined as the fuel cost for the whole flight, as calculated before departure). Indeed, airlines want the actual costs of their flights to remain reasonably close to the predicted costs. The contrary would give little meaning to running flight path 'optimizations'; it could even turn out as a drawback if the pilot wants to meet the Estimated Times Over "at any cost", possibly causing a lot of extra fuel burn. Actually, the simulations showed that different definitions for the cost function of the 4D-optimization problem resulted in very different errors between the actual and the forecasted fuel costs. This proved that finding a correct definition was a major concern to solve.

7.2.2 4D-Trajectory Choice

The 4D-trajectory corresponds to the route calculated before departure by a flight-planning optimizing software. For each waypoint, the Estimated Time Over is set to the time at which the waypoint is calculated in this 'optimal path' to be flown over.

Similarly to previous sections, this study is restricted to the case of a flight on a single Flight Level (cf. section 5.2.2), and motion between two consecutive waypoints is legitimately considered linear (cf. section 5.2.5). Also, only the components of the wind in the horizontal plane are considered. With these assumptions, the 4D-trajectory can be set to a typical optimal path obtained with the algorithms of Section 5. It can be noted that by considering 'reasonable' weather errors, we are then ensured that the path is stable (cf. section 6.2.2).

7.2.3 Mono-Dimensional Study

Actually, because the set of waypoints identifying the trajectory has already been determined, the models used in the simulations are simplified by considering that the entire en-route phase is in a single dimension. In other words, the flight shall be considered linear on one flight level from the Top of Climb until the Entry Point but with the weather data of a real path (a typical optimal path obtained in Section 5). This mono-dimensional study provides results that can be extrapolated to the 4D-problem (which is much more difficult to model, and which uses much more computing time) when trying to define a useful cost function to minimize.

7.2.4 Weather Data

The 4D-trajectory goal is calculated before departure as the optimized path according to the weather forecasts (and the optimization criteria). This path, obtained through the algorithms of Section 5, assumes that the Mach is constant (this is consistent with reality). Consequently, if the aircraft flies at this Mach, meeting the Estimated Time Over each waypoint is only dependent on the accuracy of the weather forecasts. More specifically, whatever these weather forecasts were before departure, the only relevant wind information for our work is the error in the forecasts. In other words, once we have the 4D-trajectory goal, it is useless in our study to consider the original forecasted winds. Only the impact of the error of forecasts and (of the next forecasts if any) are of interest to us.

Two different ways of considering the forecast errors were used in the simulations:

- either by using the real differences between the forecasts and the actual weather obtained via the weather grids provided by MIT's department of Earth, Atmospheric and Planetary Science (cf. section 5.2.3.2),
- or by using Normal distributions of wind errors (for both the North-South and the East-West components), as Nair and Forrester showed it was relevant to do (cf. section 5.2.3.1). This simple technique is used in the simulations when different magnitudes of the wind forecasts errors are considered.

7.3 Linear Programming

Our problem is subject to a number of constraints that, it will be seen, can be formulated as linear. The classic, natural and easy choice is then to use Linear Programming (LP) to perform the fuel flow optimization.

7.3.1 Notations

7.3.1.1 Linear Motion

To summarize the previous paragraphs, the goal of this study is to look at an optimal use of the fuel flow for a mono-dimensional motion along what will be called the x-axis. The flight is considered between the Top of Climb at position $x0$ and time 0, and the Entry Point at position X, originally scheduled to be reached at time T.

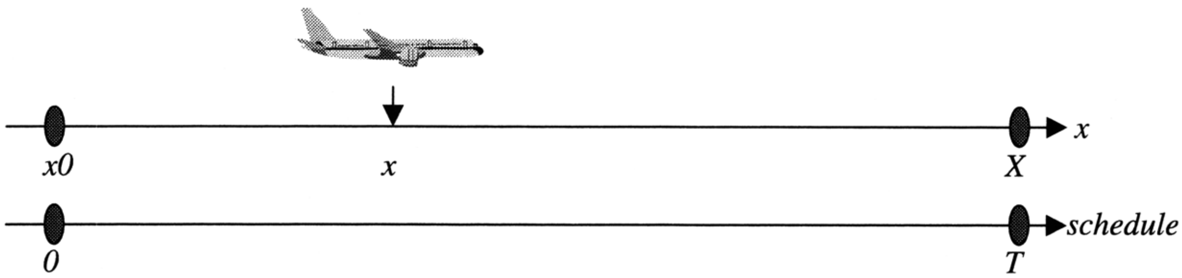

Figure 11. Linear motion from the Top of Climb until the Entry Point

7.3.1.2 Intermediate Waypoints

For the general case study, let us call N_{iwp} the number of intermediate waypoints before the Entry Point. Waypoint number i is located at $x=x_{wp}(i)$, and is expected to be flown over at the exact Estimated-Time-Over $T_{wp}(i)$. The actual location of the aircraft at time $T_{wp}(i)$ is noted $x_{wp}i$.

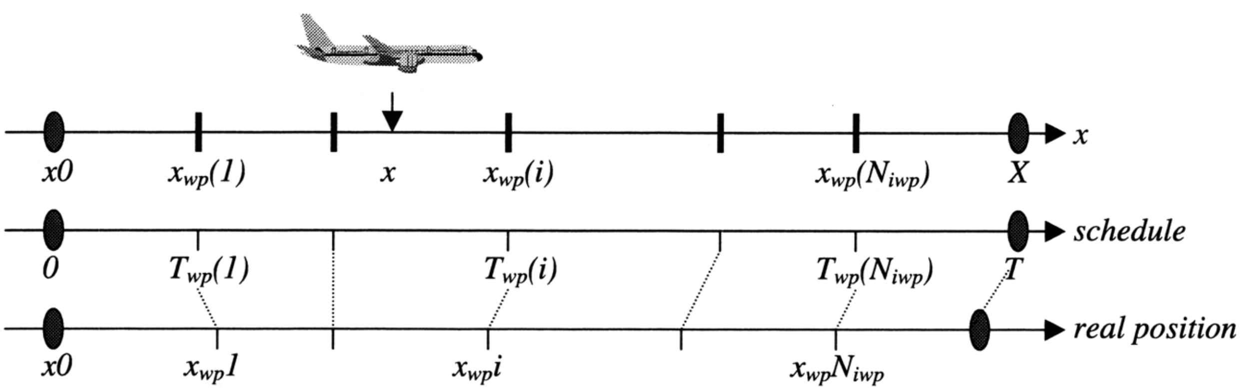

Figure 12. Intermediate Waypoints Notations

7.3.1.3 Discretized Real-Time Optimization

The classic way for a problem like ours to use Linear Programming (LP) is to discretize time between the starting position (at time 0) and the destination one (at time T). The period of time that separates two discretization steps is chosen as a constant called dt. If N_{disc} is the number of time steps, then we have:

$$dt = \frac{T}{N_{disc}} \qquad (7.3.1.3.1)$$

The higher N_{disc}, the more "continuous" the optimization. However, similar to the work in Section 5, the higher N_{disc}, and the much higher the computing time. Indeed, we will see later that the dimensions of the matrices involved in the Linear Program are almost proportional to N_{disc}, causing the optimization running time to increase exponentially with N_{disc}. In actual flight situations, cockpit software functions are wanted to provide their results after a few seconds (at worst), this explains why N_{disc} cannot be chosen too big if onboard real-time flight optimizations are to be run. Consequently, the computing time dictates a N_{disc} threshold. However, even if a big N_{disc} would seem to favor more continuous variations of Mach, in fact, it is typical to such LP problems that accelerations will be done only at a few discretization steps.

The LP program is run one first time at the Top of Climb, then is - for instance - run at each update of data sent from the AOC. For the simulations, it makes sense to run a new optimization at each "discretization step" following an update of the weather forecasts. Considering that two consecutive weather forecasts are always separated by the same amount of time, two consecutive optimizations are almost separated the same. However, some further work could investigate when optimizations during the flight should be computed.

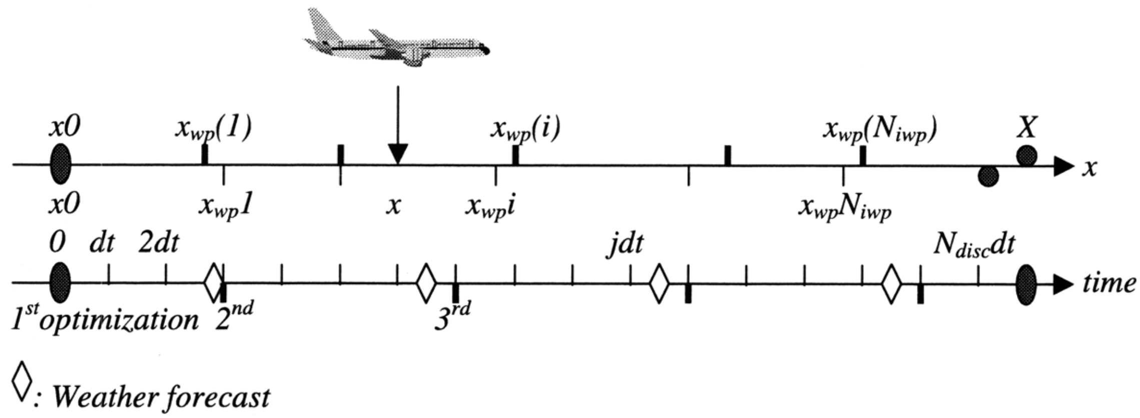

Figure 13. Time discretizations, weather forecasts, and optimizations steps

Parameters at the top of climb or where and when a new optimization occurs are indexed *0* (e.g., time is t_0). Parameters at the n^{th} time discretization step after the optimization location are indexed n (at time $t_n = t_0+n.dt$):

- the aircraft's position is x_n,
- the speed relative to the ground is dx_{rn},
- the acceleration command is u_n,
- and the wind horizontal speed forecast is U_n.

7.3.2 Acceleration Command u_n

Having explained the linear motion assumptions, we can consider that only one command acts on the 'aircraft' system: the acceleration $(u_n)_{n=0..(N-1)}$, with N the number of time discretization steps left before the scheduled time over the Entry Point (T). This useful property of a single command for the system highly encouraged working in one dimension. It is easy to implement, while the results of the optimizations provide exactly what this work is looking at:

knowing the fuel consumption for the whole trip. For that, we do the first-degree approximation that the acceleration command corresponds to the extra fuel consumption (extra to the flight-with-constant-Mach case). The optimization will act on the command to reduce the extra fuel consumption.

7.3.3 Equality Constraints

7.3.3.1 First Relations on Positions and Speeds

Since by construction, the LP program provides a constant acceleration on every discretization step, we have:

$$\dot{x}_r(t) = \dot{x}_{r_n} + (t - t_n) \cdot u_n \qquad \forall n \in [0, N-1], \forall t \in [t\ , t_{n+1}] \quad (7.3.3.1.1)$$

where $\dot{x}_r(t)$ is the speed of the aircraft relative to the wind, at time *t*.

We also have:

$$\dot{x}(t) = \dot{x}_r(t) + U(t) \qquad \forall t \qquad (7.3.3.1.2)$$

Notably, this provides at time $t = t_n$ the following relation:

$$\dot{x}_{r_n} = \dot{x}_n - U_n \qquad \forall n \in [0, N-1] \qquad (7.3.3.1.3)$$

(7.3.3.1.1) and (7.3.3.1.2) combine into the following relation:

$$\dot{x}(t) = \dot{x}_{r_n} + U(t) + (t - t_n) \cdot u_n \qquad \forall n \in [0, N-1], \forall t \in [t_n, t_{n+1}] \qquad (7.3.3.1.4)$$

In particular, the two last equations provide at time $t = t_{n+1}$:

$$\dot{x}_{n+1} = \dot{x}_n + dt \cdot u_n + U_{n+1} - U_n \qquad \forall n \in [0, N-1] \qquad (7.3.3.1.5)$$

This relation is used as such in the linear program.

Finally, (7.3.3.1.4) can be integrated between t_n and t as:

$$x(t) = x_n + \int_{t_n}^{t} \dot{x}_n(t)dt = x_n + (t-t_n)\dot{x}_n + \frac{(t-t_n)^2}{2}u_n + \int_{t_n}^{t}(U(t)-U_n)dt \qquad \forall n \in [0, N-1], \forall t \in [t_n, t_{n+1}]$$

(7.3.3.1.6)

7.3.3.2 Wind Speed Linear Approximation, Positions Incrementation

The wind forecast at time t is used as a first-degree linear extrapolation:

$$U(t) = U_n + \frac{(t-t_n)}{dt}\cdot(U_{n+1}-U_n) \qquad \forall n \in [0, N-1], \forall t \in [t_n, t_{n+1}] \quad (7.3.3.2.1)$$

which allows us to write (7.3.3.1.6) in the following easier form:

$$x(t) = x_n + \int_{t_n}^{t} \dot{x}_n(t)dt = x_n + (t-t_n)\dot{x}_n + \frac{(t-t_n)^2}{2}u_n + \frac{(t-t_n)^2}{2dt}(U_{n+1}-U_n) \qquad \forall n \in [0, N-1], \forall t \in [t_n, t_{n+1}]$$

(7.3.3.2.2)

In particular, this provides at time $t = t_{n+1}$:

$$x_{n+1} = x_n + dt\cdot\dot{x}_n + \frac{dt^2}{2}\cdot u_n + \frac{dt}{2}\cdot(U_{n+1}-U_n) \qquad \forall n \in [0, N-1] \quad (7.3.3.2.3)$$

This equation is used as such in the linear program.

7.3.3.3 Aircraft's Position at an Estimated Time Over

Let us call $n_{\mathrm{Twp}}(i)$ the discretization step number which corresponds to the discretization time directly preceding $T_{wp}(i)$ (waypoint i's Estimated Time Over). The time difference $dt_{\mathrm{wp}}(i)$ between $T_{\mathrm{wp}}(i)$ and the preceding discretization step is:

$$dt_{wp}(i) = T_{wp}(i) + \left[N_{disc} - N - 1 + n_{Twp}(i)\right]\cdot dt \qquad \forall i \in \left[1, N_{iwp}\right] \qquad (7.3.3.3.1)$$

With (7.3.3.2.2), we obtain:

$$x\left(T_{wp}(i)\right) = x_{n_{Twp}(i)} + dt_{wp}(i)\dot{x}_{n_{Twp}(i)} + \frac{dt_{wp}(i)^2}{2}u_{n_{Twp}(i)} + \frac{dt_{wp}(i)^2}{2\cdot dt}\left(U_{n_{Twp}(i)+1} - U_{n_{Twp}(i)}\right) \qquad \forall i \in \left[1, N_{iwp}\right]$$

(7.3.3.3.2)

This value will be called xwp_i.

7.3.3.4 Summary

The equality constraints for the LP formulation are equations (7.3.3.1.5), (7.3.3.2.3), (7.3.3.3.2), as well as the three initial and final conditions:

$$x_0 = x0 \qquad (7.3.3.4.1)$$

$$\dot{x}_0 = dx0 \qquad (7.3.3.4.2)$$

$$x(T) = X \qquad (7.3.3.4.3)$$

with *x0* and *X* the respective Toc of Climb and Entry Point position, and *dx0* the initial ground speed.

7.3.4 N4D Formulation

7.3.4.1 Definition

As it will prove useful, we define the notion of Near-4D trajectory (N4D). It shall be similar to the 4D notion, except that the strong time equality constraints over each intermediate waypoint become the following goal: at each original Estimated Time Over, keep the miss distance to (or from) the corresponding waypoint reasonably small.

7.3.4.2 4D versus N4D

The N_{iwp} time equality constraints imposed by the 4D formulation (versus N4D) are:

$$xwp_i = x_{wp}(i) \qquad \forall i \in [1, N_{iwp}] \qquad (7.3.4.2.1)$$

To compare 4D and N4D efficiencies, there is no need to make two different linear programs. 4D is just the particular case of N4D with infinite weights - in the cost function - on the miss distance with intermediate waypoints (these infinite weights will simply be simulated with very big values). The notion of weight on miss distance will be explained in paragraph 7.3.6.2.1.

For our study, two different LP's were originally written with one being specific to 4D. The results confirmed the statement that 4D is just the particular case of N4D with infinite miss distance weights. This is why only the N4D formulation is described hereafter.

Finally, it should be remarked that a 4D formulation has an advantage: the involved matrices are smaller, so that the optimization computing time is reduced. However, experience showed that a 4D program does not run faster than twice the speed of a N4D algorithm. This time advantage is outweighed by the drawbacks of "4D versus N4D" that will be explained later.

7.3.5 First Inequality Constraints

Here follows the list of "obvious" inequality constraints acting on the system:

- the acceleration (relative to the wind) limitations:

$$u_{\min} \leq u_n \leq u_{\max} \qquad \forall n \in [0, N-1] \tag{7.3.5.1}$$

which can be notably relevant of a desired passengers' comfort;

- the Mach number bounds:

$$M_{\min} \leq M \leq M_{\max} \tag{7.3.5.2}$$

The speed relative to the wind is:

$$\dot{x} = M \cdot \sqrt{\gamma R \cdot Temperature} \tag{7.3.5.3}$$

(7.3.5.2) and (7.3.5.3) combine into the following speed inequality constraints to implement in the linear program:

$$M_{\min} \le \dot{x}_n \cdot \sqrt{\gamma R \cdot Temperature_n} \le M_{\max} \qquad \forall n \in [1, N] \qquad (7.3.5.4)$$

where $Temperature_n$ is the temperature at waypoint number n when the aircraft is on top of it. In fact, $Temperature_n$ should depend on the result of the optimization, as it is a function of the time at which waypoint n is predicted reached. This problematic situation is resolved by defining $Temperature_n$ as the temperature over waypoint number n at the originally expected time there. This first-degree approximation is very reasonable.

7.3.6 Cost Function

7.3.6.1 4D Study

A first idea is to look at minimizing the positive acceleration increment *u(t)* along the path of the aircraft. Indeed, to recall a fuel cost minimization, we can consider that if *u(t)* is positive, the value of the cost function augments, while if it is negative the cost function remains unchanged. This is mathematically the same as saying that we want to minimize the integral over the flight of the quantity $\frac{u(t)+|u(t)|}{2}$.

The cost function to minimize would then be written as:

$$C_1 = \int_{t_0}^{t} \frac{u(t)+|u(t)|}{2} dt \qquad (7.3.6.1.1)$$

Since time is discretized, we obtain:

$$C_1 = \sum_{n=0}^{N-1} \frac{u_n + |u_n|}{2} \cdot dt \qquad (7.3.6.1.2)$$

The absolute value is typically taken care of in the LP formulation by adding new decision variables: $(v_n)_{n=0,..,N-1}$. They verify:

$$-v_n \leq u_n \leq v_n \qquad \forall n \in [0, N-1] \tag{7.3.6.1.3}$$

and the cost function to minimize in the 4D study finally becomes:

$$C_2 = \sum_{n=0}^{N-1} \frac{u_n + v_n}{2} \cdot dt \tag{7.3.6.1.4}$$

7.3.6.2 N4D Study

7.3.6.2.1 Weights on Miss Distances with Waypoints

We choose an arbitrary weight $W_{md}(i)$ for the miss distance at time $T_{wp}(i)$ with waypoint i. The new objective function to minimize is thus:

$$C_3 = C_2 + \sum_{i=1}^{N_{iwp}} W_{md}(i) \cdot \left| x_{wp}(i) - xwp_i \right| \tag{7.3.6.2.1.1}$$

A constant value was chosen for these weights. This value influences the results of the optimization, as will be shown in section 8.2.2.

Similarly to what was done in the previous paragraph, to keep an LP-like formulation, we deal with the absolute values by introducing new decision variables: $(\alpha_i)_{i=1,..,Niwp}$. They verify:

$$-\alpha_i \leq x_{wp}(i) - xwp_i \leq \alpha_i \qquad \forall i \in \left[1, N_{iwp}\right] \tag{7.3.6.2.1.2}$$

and the objective function becomes:

$$C_4 = C_2 + \sum_{i=1}^{N_{iwp}} W_{md}(i) \cdot \alpha_i \tag{7.3.6.2.1.3}$$

7.3.6.2.2 Weights on Positive Acceleration Increments

The good-but-limited efficiency of the simulations using C_4 as the cost function showed that forecasting late in-flight acceleration increments is not very useful for precise arrival over the Entry Point (because of weather uncertainties during the flight). The main consequence is a waste of fuel consumption. This has led to the consideration of weights for the positive acceleration increments, such that eventually we obtain the following cost function:

$$C_5 = \left(\sum_{n=0}^{N-1} W_a(n) \cdot \frac{(u_n + v_n)}{2} \cdot dt \right) + \left(\sum_{i=1}^{N_{iwp}} W_{md}(i) \cdot \alpha_i \right) \qquad (7.3.6.2.2.1)$$

where $W_a(n)$ is the weight on the n^{th} positive acceleration increment.

The intuitive idea, which will be confirmed with the results of the simulations, is that we should choose $(W_a(n))_{n=0,..,N-1}$ as a monotonically decreasing series. Indeed, this way, early positive acceleration increments occur only when necessary. The series shall also depend on the flight length left to travel to make sure that the first big parenthesis in (7.3.6.2.4) remains on the same order of magnitude as the second.

Section 8

Simulation Results and Analysis

8.1 Objective: Meeting the Time of Arrival

8.1.1 Stability of Real-Time Dynamic-Programming Optimizations

To show the stability of the optimal paths obtained with the Section 5 and Section 6 algorithms, the next page shows a typical example of the best paths these algorithms would provide. The optimizations are run before the Top of Climb, and in real-time during the en-route phase. The best path towards each reachable waypoint is represented (cf. Section 5), and the optimal trajectory is the plain line. These plots show the very good stability of the optimized trajectory. However, it must be kept in mind that this may be true only for the case of waypoints separated by 10-degree latitudes; the impact of a denser grid of waypoints was not studied in the present work.

Although the route is very stable, at each optimization, a new optimal Mach is calculated and it sometimes varies a lot. This is acceptable because the pilot and the controller keep good situation awareness if the route is stable.

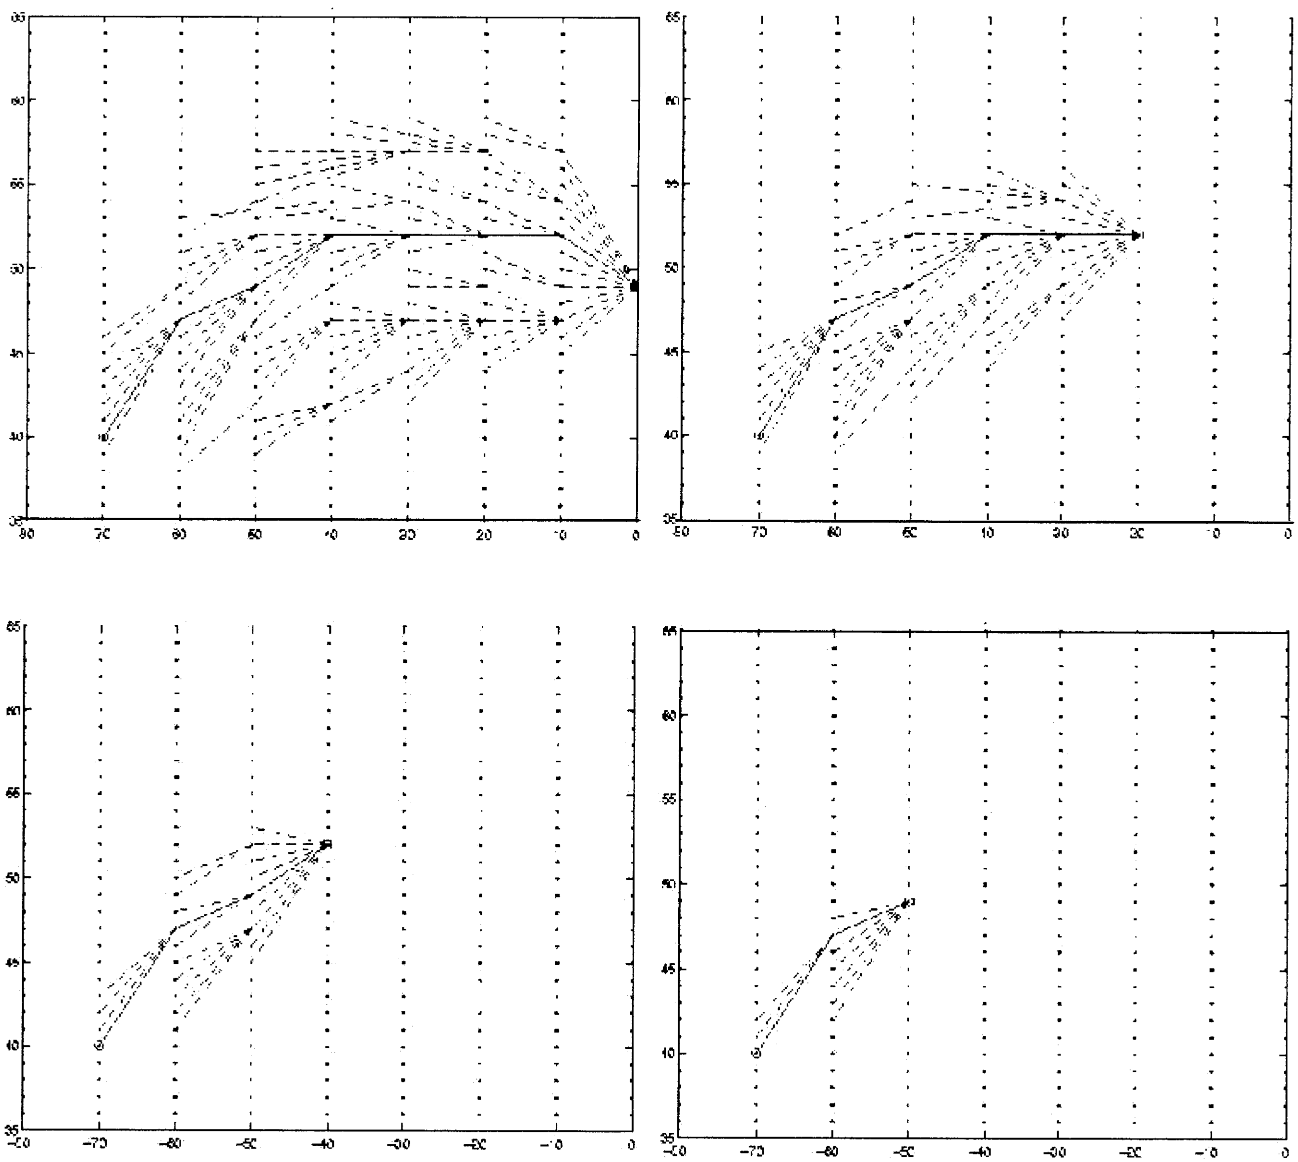

Figure 14. Real-Time Dynamic-Programming Optimizations Sample

Best paths computed at optimization steps distant from each other by time t = 1.5 hour. The actual route is in this case the one that was computed three hours before departure. Other simulations rarely show a real-time change of the optimal path (but the Mach is changed after each optimization). Here, the first optimization gives an optimal Mach of 0.8025. After the second optimization, it is set to 0.82 (maximal value), then to 0.8178, and finally to 0.8132.

8.1.2 Impact of Weather Turbulence on Time at Destination

Using the notations defined in section 6.3.2.2, the standard deviations of wind turbulence components were varied to see their impact on the delay at arrival. On the following figure, each plotted value is the average of eight simulations:

- four different days,
- the optimization is run twice on each day for the same flight leg.

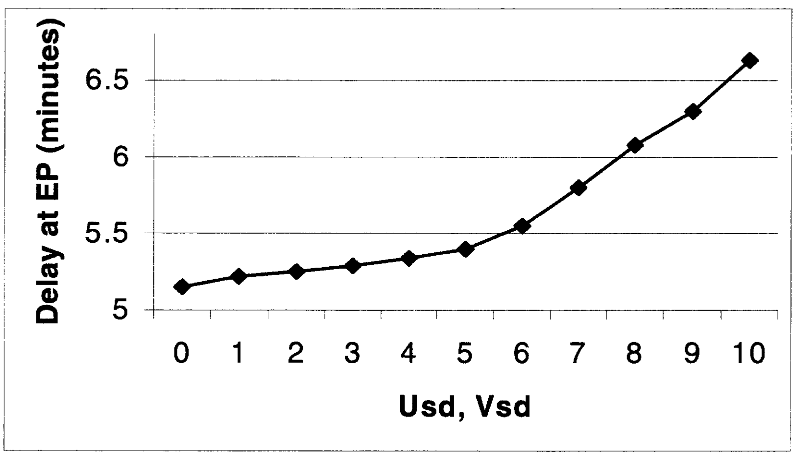

Figure 15. Impact of weather turbulence on time at destination

Optimizations are run every 1.5-hour; at every attempt, they try to compensate the error on schedule. Consequently, it is the last optimization that generally affects the delay the most at destination.
The delay when there is no turbulence is seen in more detail in section 8.1.4.

Figure 15 shows the obvious result that the delay at destination increases with the importance of the weather turbulence. Nevertheless, the delay values remain confined very close to the non-turbulence situation. Consequently, the weather turbulence is no longer considered in the rest of the study. This saves computing time.

8.1.3 Comparison with Actual Flight-Planning Systems

The most famous flight-planning system is probably PHOENIX (see Ref.10). PHOENIX uses 12, 18 and 24-hour weather forecasts, not the 6-hour forecasts that are more likely to be accurate. Air France has added many features to the PHOENIX software by creating the OCTAVE (Ref.7). This advanced flight-planning tool works in a way that is very similar to the algorithms described in Section 5 (with, of course, the extra possibility to change the flight level). Consequently, the present work is a study into the consequences on the time at destination of running other flight-plan optimizations rather than only one before departure, with the best current flight-planning systems.

The flight leg that is considered for the simulations was chosen because it corresponds to a flight leg flown by Air France, for which a typical optimal path is known from Ref.7. Air France's "best path" is very similar to the one obtained through the simulations of the present work (as seen on Figure 2). Naturally, a few differences exist because the waypoint grids are not exactly similar, and because real flights fly in 3D and not 2D, but these differences are very limited.

In conclusion, it is legitimate to compare the results of the real-time optimizations with the situation of a single optimization run before departure, and to extrapolate the conclusions to the real case with the best real flight-planning tools.

8.1.4 Optimization Before Departure

The two next paragraphs show the consequence on the delay at destination of running real-time flight path optimizations. The results need baseline delays to compare with; these are the delays over the Entry Point when the simulations are run similarly to the current practice (only once, and before departure). Figure 16 shows the distribution of these delays over 50

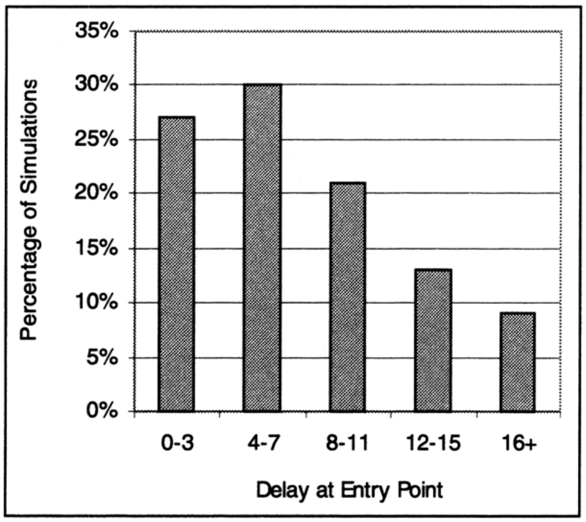

simulations.

Figure 16. Baseline delays

The values involve 50 simulations on flights for which a single optimization before departure induces delay at destination. The average delay is 6.88 minutes. The average delay when also considering the non-delayed flights was 5.17 minutes (over 50 simulations too).

8.1.5 Optimization at the Top of Climb

The first step towards real-time optimizations is to run a flight-plan optimization at the Top of Climb with the latest updates of weather forecasts (cf. section 6.1). Figure 17 illustrates the improvement of such method over the flight delay due to an optimization done three hours before reaching the Top of Climb (which roughly corresponds to the current real situations). The assumption made is that the weather forecasting system has improved from 6-hour to 3-hour forecasts, allowing a forecast update to be obtained right before the Top of Climb.

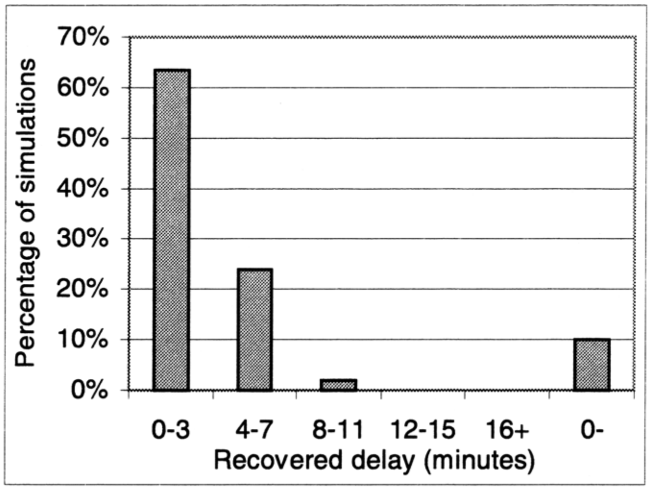

Figure 17. Delay reduction through optimization at the Top of Climb

The values involve 50 simulations on flights for which a single optimization before departure induces delay at destination. The average recovered delay is 1.92 minute.

This graphic shows that about 90% of flights shortened their delay over the Entry Point by performing a flight path optimization over the Top of Climb with updated 3-hour weather forecasts. However, the improvement was always very limited.

10% of flights worsened their delay over the Entry Point after running an optimization at the Top of Climb. Fortunately, the lengthening of the flight duration never exceed 4 minutes (out of five simulations with increased delays, only one augmented delay by more than 2 minutes).

In conclusion, running a flight-plan optimization with updated weather forecasts at the Top of Climb is beneficial to the airline, but this improvement is very limited.

8.1.6 Dynamic Optimization

A farther look in the future is then investigated by considering a discrete set of real-time flight path optimizations (cf. section 6.2). The logical evolution in the complex ATC-world context is that the first step towards real-time flight path optimizations will probably be to re-optimize the path over the Top of Climb (read previous paragraph). Therefore, the current 6-hour weather forecasts are considered enhanced to 3-hour forecasts, or even 1.5-hour forecasts. Figures 18 and 19 show the results for the corresponding sets of simulations. Optimizations are run at each weather forecasts update.

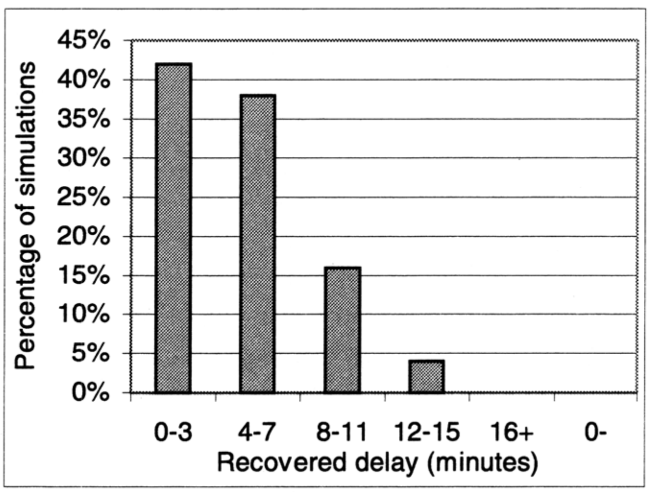

Figure 18. Delay reduction through real-time optimizations every 3-hour

The values involve 50 simulations on flights for which a single optimization before departure induces delay at destination. The average recovered delay is 4.28 minutes.

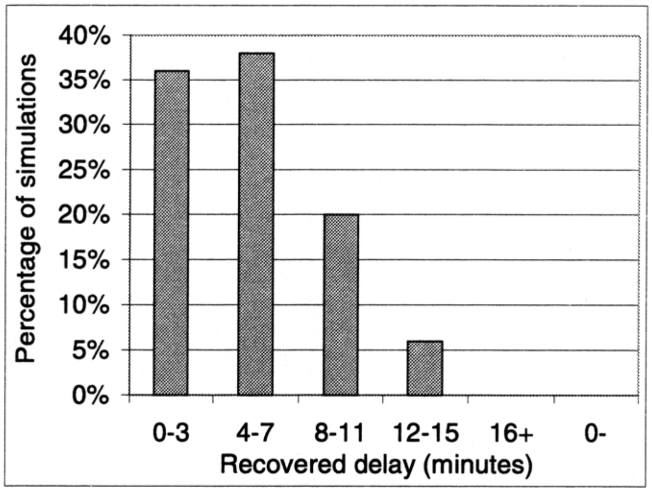

Figure 19. Delay reduction through real-time optimizations every 1.5-hour

The values involve 50 simulations on flights for which a single optimization before departure induces delay at destination. The average recovered delay is 4.84 minutes.

Figures 18 and 19 prove the efficiency of considering real-time discrete optimizations. For both cases, 100% of flights decreased their delay at destination, and the time gain is often relatively significant. Moreover, Figure 19 shows results slightly better than on Figure 18, which tends to conclude that more frequent real-time optimizations can reduce even more delay. However, the differences are not big, so it seems reasonable to stay with 3-hour weather forecasts (which will be much more affordable in terms of weather system computing time than 1.5-hour forecasts).

8.2 4D and N4D Trajectories

The following paragraphs refer to the work described in Section 7.

8.2.1 Linear Program Outputs

Figure 20 shows an example of aircraft velocity results obtained with the linear program described in Section 7.

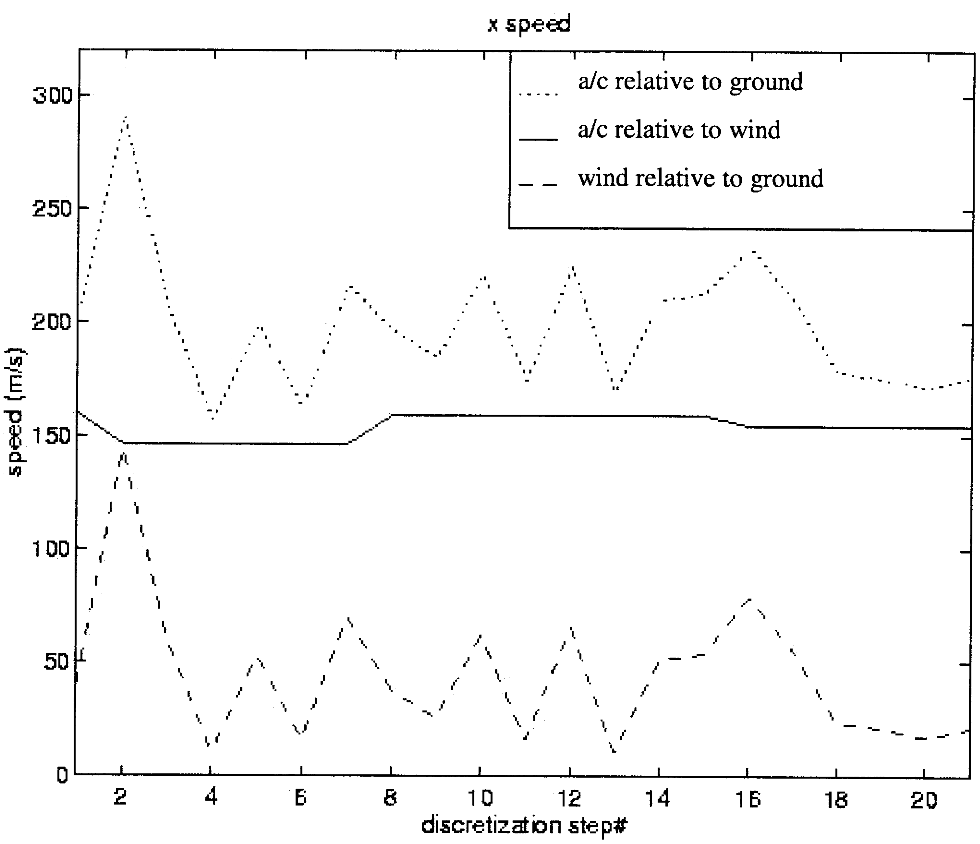

Figure 20. Typical N4D linear program outputs (W_{md}=0.05)

The weather data is here chosen such that it varies a lot. Five intermediate waypoints are considered, and three optimizations are performed. The acceleration can be seen from the variations of aircraft speed relative to the wind.

8.2.2 4D versus N4D efficiency

The difference between 4D and N4D is set by the choice of the weight on miss distances with the intermediate waypoints (see section 7.3.6.2.1). If $W_{md} < 10^{-2}$, it means that no effort is paid on trying to pass by the intermediate waypoints on schedule, such that the miss distances become sometimes significant, causing delays (positive or negative) over the intermediate waypoints of several minutes. If $W_{md} > 10^{2}$, it means that the goal is set to fly 4D. The consequence of the effective trip fuel cost of varying the value of W_{md} is shown on Figure 21.

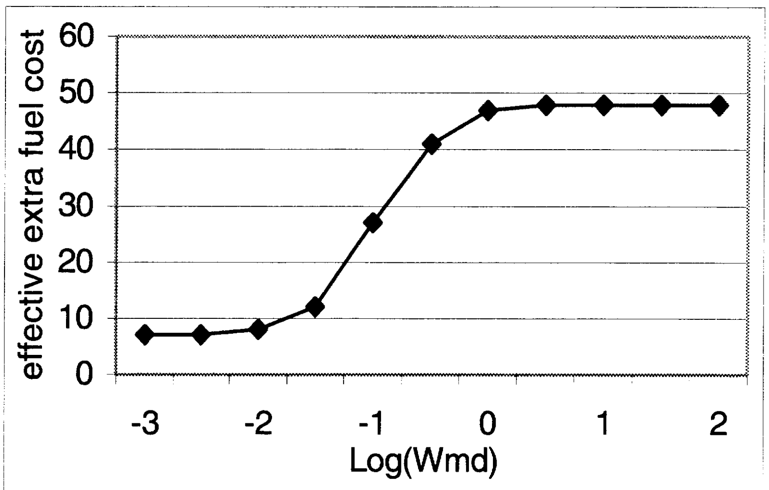

Figure 21. Impact of W_{md} on trip fuel consumption

Each value is the average of eight simulations, with a wind speed variance U_{var} = 5m/s, five intermediate waypoints, and three time-equidistant optimizations. The fuel cost represents the extra fuel cost over the flight at a constant Mach (calculated to fly over the Entry Point on schedule). When W_{md} is very small, the extra fuel cost is not zero because there still exists the constraint of having to adjust the Mach to try to arrive at the Entry Point on schedule.
The extra fuel cost was assimilated in the simulations as the positive acceleration increments, such that the unit for the extra fuel cost is m/s^2.

Figure 21 clearly shows that imposing to fly a 4D-trajectory considerably increases the extra fuel consumption. A trade-off must be found if the goal is to stay close to a 4D-path. It is found by adequately setting the W_{md} value in the N4D problem. Figure 22 provides another reason for imposing to adequately choose W_{md}.

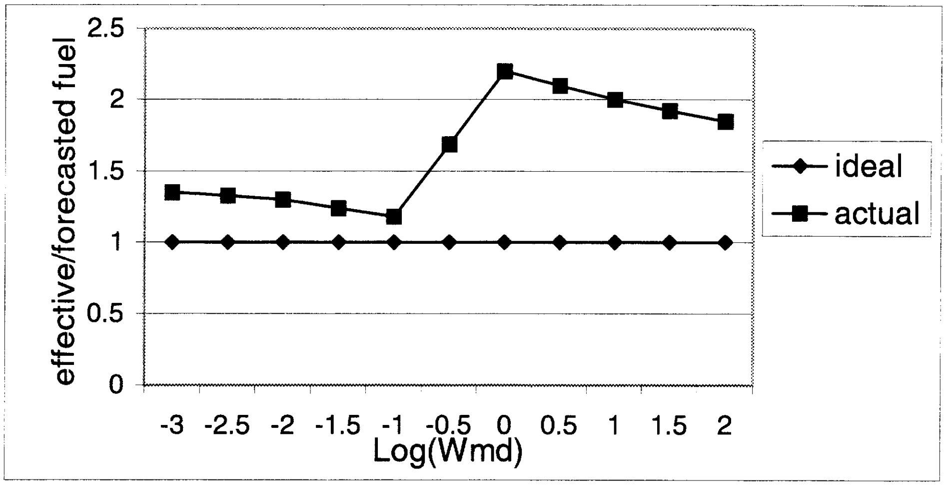

Figure 22. Impact of W_{md} on reliability of fuel consumption forecast

Each plotted value is the average of eight simulations, with a wind speed variance U_{var} = 5m/s, five intermediate waypoints, and three time-equidistant optimizations. The forecasted fuel represents the extra fuel calculated before departure to be needed compared to the flight at a constant Mach. The effective fuel is the extra fuel that has actually been burnt during the flight, different from the forecasted one because of the weather forecasts errors.
The ideal result - in terms of reliability of the flight-path optimization run before departure - is to have the effective fuel equal to the forecasted one.

Figure 22 shows that the 4D concept would result in poor reliability in the flight-path optimization run before departure. More exactly, wanting to fly 4D causes the actual fuel consumption to be possibly very different from the one predicted. Again, flying near a 4D-trajectory requires a trade-off with the value of W_{md}.

It is possible to guess a good value for W_{md} by looking at the cost function involved (cf. section 7.3.6.2.1) and wanting to roughly equal the fuel consumption terms with the miss distance ones. This balances the importance of flying N4D and of minimizing the fuel consumption. Let us consider the examples used for Figures 20, 21 and 22. Figure 20 tells us that the fuel term should be less than 20. By considering for example that a miss distance is wanted less than 2000 meters, since there are five intermediate waypoints, this suggests a W_{md} value of 20/(2000*5) = 0.002. This method seems efficient, as Figures 21 and 22 suggest.

8.2.3 Improved Strategy

As mentioned in section 7.3.6.2.2, introducing weights on the positive acceleration increments in the cost function allows us to reduce the fuel consumption. The idea is to consider that the constraints near the end of the flight should be less stringent than near where the optimization is run. Indeed, strong constraints at the end of the flight could cause unnecessarily elevated extra fuel consumption, while dynamic events such as the weather forecast error could make useless the previous acceleration increments. This idea was verified through the simulations. No results are displayed here, as they are not based on analytical proof, but the reader could find it interesting to learn that choosing $W_{fc}(i)$ (see section 7.3.6.2.2) as a decreasing function of i, where i is the next i^{th} waypoint, sometimes worked remarkably well to reduce the extra fuel consumption. For example, it seemed that choosing

$$W_{fc}(i) = N_{disc} - i + discretization\_step\_number$$

provided promising results.

8.2.4 Other Results

Finally, Figure 23 confirms the previous comparison between 4D and N4D. It also shows the intuitive result that more errors on the wind forecast bring more errors in the fuel consumption forecast. It should be noticed that the reason for the local minimum in fuel cost ratio at five waypoints was not investigated.

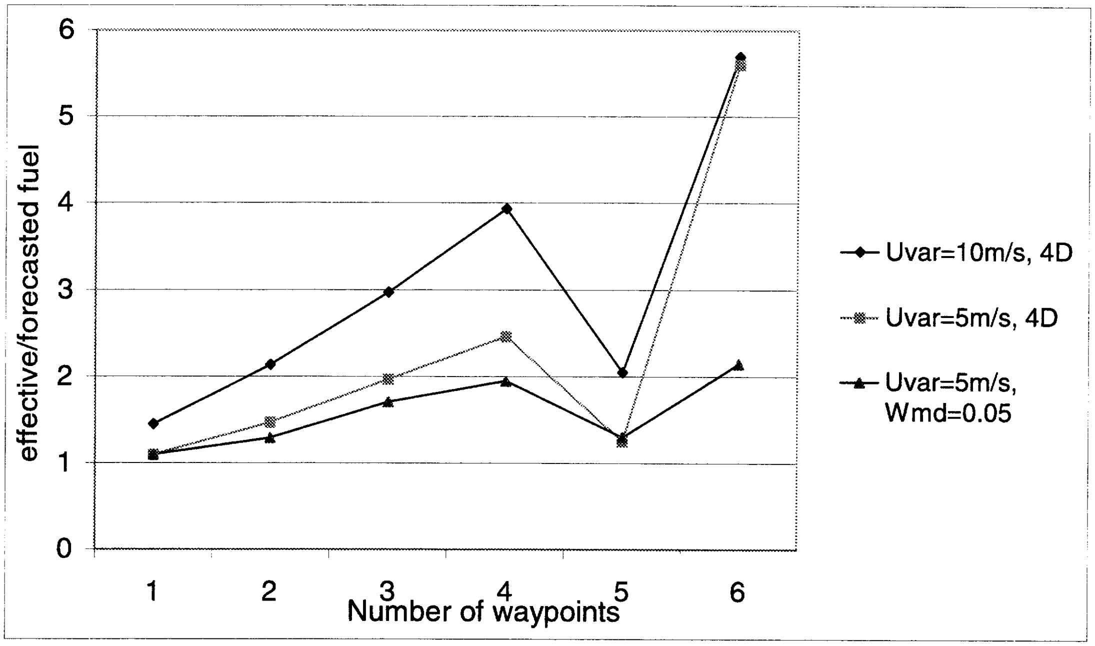

Figure 23. Impact of number of waypoints and of weather forecasts errors on reliability of forecasted trip fuel

U_{var} is the variance of the wind speed. The number of waypoints includes the waypoints of the path without the Top of Climb.

Section 9

Conclusion

Weather uncertainties prove to be such that calculating an optimal Mach number at departure results in significant errors in arrival time and fuel burn. The strategy used by airlines to run several optimizations for several constant Mach numbers is then justifiable. It provides them with a fast method to obtain different optimized routes for different Mach numbers, then to choose the one whose characteristics fit best with their goals (e.g., arriving at the Entry Point very close to the scheduled time). Their choice is made between about four Mach numbers only because they know additional resolution is inefficient when confronted with weather uncertainties.

We described a study into the effect of wind forecast error on flight plan accuracy. We saw that the time over the Entry Point, hence at the destination airport too, can be very much different from the one computed before take-off. This results in flight cost uncertainties that make irrelevant many criteria that could be thought of for a Flight Plan optimization. The only way for now of dealing with wind forecasts errors is to carry contingency fuel (CF) to cover deviations from the planned route. Deviations may be due to ATC intervention, weather, or equipment failure. As for now, CF typically counts for 5% of the total trip fuel, the quantity of which is calculated according to the forecasted headwinds. The European regulation will soon go down to 3%. CF can occasionally be used in taxiing before take-off.

Providing guidance on the forecast wind error would have definite benefits for long-haul flights, where approximately one third of the fuel is burnt to carry CF. A small decrease in CF could produce large financial savings. For short-haul flights, there is less scope for making financial savings by reducing CF.

There is another area of benefit, operating on time-scales of minutes rather than hours, where uncertainties in forecast winds produce uncertainties in aircraft position. A reduction in these uncertainties might result in a reduction in ATC conflicts and hence a reduction in controller workload.

The only wind forecast accuracy information currently available is monthly or seasonal and RMS error statistics. This information covers areas such as NW Europe (Ref. 14) and the North Atlantic, and it is stratified as a function of flight level. A more sophisticated error model would lead to higher confidence in the forecasts thus allowing airlines to safely reduce the CF.

The present study investigated ways of performing enhanced flight-plan optimizations compared to what is currently done. We first saw the benefit of running another optimization at the Top of Climb, yet we realized that the benefits were limited. On the contrary, the more futuristic method of considering discrete flight-plan optimizations certainly seemed promising to ensure reducing delays at destination.

Finally, we investigated ways of performing a Linear Program to fly a route close to a 4D-trajectory. It is not sure whether people will further study the 4D concept. Nevertheless, the results described in paragraph 8.2 bring precious information to help creating the cost function to minimize in this optimizing problem.

Appendix

A.1 Different Flight Plan Preparation for Different Flight Legs

A.1.1 Repetitive Flight Plan (RPL)

Repetitive Flight Plans are planned months in advance. They represent the statistically "best" routes for each flight leg and are often the shortest path. An RPL is the route frequently chosen by the airline for one flight leg. The ATMS gathers all RPLs months prior to flights, thereby roughly forecasting the traffic density.

RPLs exist for medium-haul flight plans only, as the weather will not affect the duration of a medium-haul flight a lot, contrarily to what may happen to a long-haul flight.

A.1.2 Planned Flight Data (PFD)

Aside from RPLs, months before the corresponding flights, the airline sends PFDs to the CFMU that only contain the scheduled flight legs (for long-haul flights). Using complete flight plans (RPLs) ready for such flights would not be very useful.

A.1.3 Flight Plan (FPL)

On the actual day of the flight, the Flight Planning software uses the weather forecasts to determine the route that fits best the airline's strategy. The route is selected from a catalogue of routes, oceanic tracks (for transatlantic flights), computed random routes (if the airline has the corresponding optimizing tools), or a combination of all these.

A.1.4 Random Routes

The best Flight-Planning tools work in a way similar to Dynamic Programming to determine on long-haul flights the set of waypoints that will constitute a so-called "Random Route". This route is supposedly the best one to fly accordingly to the weather forecasts.

Other flight-planning tools consider that the best path is given by the Great Circle technique, which looks at the path the closest to the direct route from departure to destination. However, a better path requires the use of weather forecasts.

Random Routes are computed from a set of imaginary (given by their latitude/longitude coordinates) and real waypoints. Therefore they require that the aircraft be RNAV-equipped.

A.2 Flight Levels

A.2.1 Introduction

At each moment, for each Mach, there is an optimal barometric altitude where the aircraft burns minimal fuel. The lighter the aircraft is, the higher this altitude is. While the aircraft burns fuel, its weight decreases; consequently, the optimal path would be a constant climb. In practice, except Concorde (only commercial aircraft allowed above FL410) and over certain areas of the Pacific Ocean where air traffic is almost non-existent, this is not permitted. Instead, aircraft fly on Flight Levels (FL's) to ensure vertical separation.

A.2.2 Optimal Flight Level Determination

It is the responsibility of the crew to find an optimal flight level - as long as it has been cleared by the ATC. Some onboard equipment helps the crew for this purpose, e.g., the FMS with use of a Cost Index.

Let us recall briefly how changes of Flight Levels are decided by the pilot when done manually (i.e., when he replaces the FMS). He uses plots of the optimum Flight Levels versus the aircraft weight, according to the chosen flight strategy (see section 1.4). A limit of 1% extra fuel consumption than on the optimal Flight Level is used to determine when the Flight Level change should occur. On a DC10 for instance, this corresponds to a climb every around two hours.

Currently, the experimental RVSM zone between FL330 and FL370 only requires 1,000ft vertical separation (as opposed to 2,000ft), this allows flying closer to the optimal climb.

It should be noticed that it is economically a very bad choice to climb on too small slopes. The climb rates are given on the plots and the Captain reacts in compliance with them while staying within the 1% extra fuel margin. These rates have been determined such that the aircraft reaches the desired Flight Plan with the maximal mass to remain in the 1% extra fuel margin.

Actually, pilots always try to climb as long as they respect this margin, except when "close" to destination (two hours for a DC10 not on a low Flight Level, one hour for an A320), when it would not be profitable. The reasons for not climbing vary:

- for DC10s, it would cost too much fuel for little gain;
- for A320s, the climbing over-consumption is low (as long as the aircraft remains under the optimal FL, this is an improvement on older aircraft). However, the few kilos of fuel that can

be saved by climbing during the last hour of flight are much less profitable than the possible rerouting to reduce trip duration.

Moreover, in case of an airline working as a hub, the airliner would not increase the trip duration by climbing because the primary concern is to arrive on time (more important than minimizing fuel consumption).

Even when the load of the aircraft allows flying at the "accelerated" Mach, climbs are often done with the "slow" Mach, as it brings more lift and improved vertical speed (the ATC wants fast Flight Level transitions). For example, a DC10 doing his en-route phase at the accelerated Mach 0.84 would do his climbs at Mach 0.82.

A.2.3 Irregular Constraints

The ATC may ask the Captain to stay on a certain Flight Level for a few minutes to ensure safe aircraft separation. In addition, the controller may ask for a small (e.g., 10 degrees) change in heading followed by the climb to a desired FL to ensure safe separation. The aircraft would then re-join the planned route.

By the way, the one-percent margin (see the previous paragraph) corresponds to regular flight procedures. The pilot, confronted to real-time weather events like turbulence, favorable winds, or unfavorable ones, can ask to move to another Flight Level. Avoiding dangerous weather (storms, cumulonimbus clouds...) is a first priority for safety. The pilot could therefore choose to fly out from the one percent-margin, but this decision remains dependent on ATC-clearance.

A noticeable case is the flights from Europe to Asia, which are very ATC-dependent. Indeed, the traffic on this "route" is heavy and entirely done in airways (even when flying to

Reunion). Towards Asia, there exists only one airway, and all aircraft fly through it. This may induce for example a four-hour wait on the flight level taken when leaving the Paris TMA (e.g., FL290 for a DC10) before being cleared for a climb to FL330 - the next allowed FL. This is because when it exits the TMA, the DC10 flying to Asia is too heavy to afford going directly to FL330. It then enters the Turkish airspace on FL290, where traffic is very heavy, forcing the aircraft to stay on this FL.

Finally, it must be remembered that FL's over a few countries are fixed (e.g., over Great Britain and Netherlands), as well as certain routes are (e.g., Great Britain to Paris).

A.3 Holding Pattern in Area Control Center (ACC)

A Holding Pattern consists of one - rarely several - lap on one Flight Level along a fixed track. The flight routes do not cross the tracks (hence constraining the routes to avoid the tracks). To ensure safety, a "held" aircraft flies at a speed imposed by the ATC such that all aircraft in a holding track at the same Flight Level fly at the same assigned speed. This speed (and/or flight level) depends on the type of aircraft, such that the airliner flies close to its Maximal Range Cruise Mach (lowest speed for lowest fuel consumption).

Two different types of holding patterns exist: inside the Terminal Maneuvering Area (TMA), or right before entrance into the TMA. The second one (the one of interest to us) is regulated by an Area Control Center (ACC). On arrival to Paris airports for example, the Athis-Mons ACC takes care of sequencing at an "adequate" rate the aircraft scheduled to enter the Paris TMA. The access to a TMA is done through one of several (three around Paris) Entry Points, and one role of the ACC is to ensure that the aircraft rate through the Entry Point corresponds to the current runways capacity. For that purpose, an ACC controller may request a pilot to fly along a holding track.

In France, the ACC controller knows the delays to apply to aircraft with the help of the sequencing tool MAESTRO. Other systems are used in other countries. MAESTRO knows the updated runways capacities and assigns each flight aiming at an Entry Point with the number of minutes of delay it should be imposed to permit a good sequencing. It should be noted that all sequencing tools, as for the CFMU slot allocation system CASA, are based on a "First come, First served" basis.

Holding Patterns outside TMA's are increasingly seldom because the sequencing tools do often well at predicting enough time in advance the delays to apply to aircraft, allowing the controller to find early solutions to those delays. When needed, the controller asks the pilot to reduce speed, or to deviate to slightly lengthen or shorten the path.

References

[1] Sofréavia / INRETS, "New Optimisation Approaches to Air Traffic Flow Management. FINAL REPORT. Sequencing, speed control, multi-sector planning and their interfaces. Cost function in ATFM optimization problems." NOAA/WP3.3a/V1.1, October 1996

[2] L. Chan Kan Fai, "Cost Index Calculation", training report, Airbus Industrie, May 1992

[3] Airbus Industrie, "Cost index - Calculation, operational use, fuel tankering", AI/EV-O 472.1413/93, April 1993

[4] Airbus Industrie Engineering Directorate (Flight Division), "Cost Index Determination and Utilisation", paper presented at the fifth Performance and Operations Conference, May 1988

[5] Survey by Eurocontrol giving numerous airlines' pilots feedback on Cost Index, 1997

[6] EATCHIP, "Trajectory Prediction for EATCHIP Phase III concept and principles", Edition 0.2, Eurocontrol Working Draft OPR.ET1.ST03.1000-ORD-01-00, 21 April 1997

[7] Slides presented by Air France airlines on their Flight Planning System OCTAVE, 1997

[8] N. Pujet and E. Feron, "Flight Plan Optimization in Flexible Air Traffic Environments", MIT department of Aeronautics and Astronautics, draft, 1997

[9] E. Hoffman and A. Bossu, "Basic Statistical Analysis of Aircraft Mass", Eurocontrol Experimental Centre Report No. 302, August 1996

[10] EDS - Air Transport Services, "PHOENIX Flight Planning System - Technical Overview", 1997

[11] Several Air France airlines' Flight Plan documents for transatlantic flights, August 8, 1997

[12] Mr. Sperber (LTU dispatcher), e-mail correspondence, August 1997

[13] T.R. Beal, "Digital Simulation of Atmospheric Turbulence for Dryden and von Karman Models", Journal of Guidance, Control, and Dynamics, Vol. 16, No. 1, January-February 1993

[14] I.J. McNair and D.A. Forrester, "Wind Forecast Accuracy Quantisation", Meteorological Office, Bracknell, UK, Version 1.0 B, 12 February 1997

[15] D.A. Forrester, A.K. Davis, J. Coiffier, G. Hello and D.C. Underwood, "The Improvement of Meteorological Data for Air Traffic Management Purpose", Stage 2b Report for Eurocontrol Contract No. C/1.184/HQ/PW/94, PHARE Program Task TM02, December 1994

[16] E. Hoffman, "Performance Modeling for ATC Use", Task BM14, Eurocontrol Experimental Centre Report No. 258, November 1993

[17] W. Austyn Mair and D.L. Birdsall, "Aircraft Performance", Number 5 in Cambridge Aerospace Series, Cambridge University press, 1992