Evaluation of a simple timing-based intervention heuristic for trajectory-based Air Traffic Management

Peter Lindsay
ARC Centre for Complex Systems, University of Queensland, Australia
p.lindsay@uq.edu.au

Appeared in: *Proceedings, 5th EUROCONTROL Innovative Research Workshop, Brétigny-sur-Orge, France, 5-7 December 2006.*

Abstract

This paper is concerned with the shift in concept from the current distance-based separation management paradigm of Air Traffic Control towards a timing-based approach to trajectory management. We propose a way of thinking about the sector controller's task, and the interventions they choose, which represents a small change from their current view but which seamlessly integrates many of their key activities into a timing-based trajectory management paradigm. We examine the effectiveness of aspects of the proposed approach using modelling and simulation. The sensitivity of the model to various parameters is examined. Modelling also reveals some areas where the concept will need refinement.

1. Introduction

The move to electronic data management and increased on-ground computing power in Air Traffic Management (ATM) means it is becoming feasible to contemplate management of a flight's 4D trajectory as a whole, rather than today's more fragmented approach to sector-based Air Traffic Control (ATC). Oversimplifying a little to make the contrast clearer, in the current ATC paradigm an aircraft's flight consists of a series of flight segments through individual air sectors, with sector controllers making intervention decisions based on local criteria to ensure that traffic flow in the sector is safe, expeditious and orderly. Controllers have evolved a complicated set of heuristics for deciding how and when to issue instructions or clearances to flight crews (for brevity we refer to these actions collectively as *interventions* in what follows). Decisions about what interventions best solve local problems do not necessarily translate to the best overall system goal, and can have a knock-on effect, causing further disruption in sectors downstream.

In an ideal "optimised ATM" world, on the other hand, airlines would file flight plans (4D proposed trajectories) that are optimal for the particular aircraft and business goals for each flight, the plans would be de-conflicted prior to take-off, and aircraft would fly their planned trajectories and arrive at their destinations on time, without need for intervention [ASTRA03]. In the real world however many factors conspire to thwart the best-laid plans. Operational issues such as delayed push-back or take off times, and unanticipated effects of wind or weather, introduce perturbations into flight trajectories, which in turn means interventions are required in order to avoid conflicts. The challenge is to develop a conceptualisation of trajectory-based ATM that extends down to operational level, so that decisions about how and when to intervene are well aligned with system-wide trajectory management. On the assumption that a role similar to that of the sector controller will persist into the future in some form or other, then it is desirable that current controller best practice be assessed and incorporated into new operational concepts where possible ([PHARE00] p.79).

This paper proposes and examines a simple heuristic for choosing interventions which attempts to integrate tactical (sector controller) decision-making seamlessly into strategic (system-wide) trajectory management. By *heuristic* we mean the "rules of thumb" that a controller would follow when deciding how and when to intervene, in order to resolve conflicts with minimal overall disruption to flight trajectories. The work is part of a larger project which is concerned with investigating the possibility of distributing elements of decision making to other agents in the Air Traffic System (ATS), such as airlines and flight crews. The heuristic proposed here is intended to provide a simple unifying view of the controller's task, to serve as a basic building block in developing complex-systems models of the whole ATS. It is based on the "rules" that controllers have been observed to use in Australian en-route sectors (by which we mean all sectors other than terminal-area sectors).

Structure of the paper:

Section 2 briefly surveys the literature on 4D trajectory management. In Section 3 we start by describing the main kinds of intervention used in the current distance-based separation management approach to ATC in Australia. We describe a simplifying way of thinking about sector design and the controller's task – a conceptualisation that moves closer to the goal of trajectory-based air traffic management. We then describe the proposed intervention heuristic as a simple procedure, based on observations about how expert controllers resolve problems.

Section 4 describes a computer simulation of timing-based aspects of the procedure, which is made more precise for the purposes of modelling and simulation. Some aspects of the procedure are represented by parameters in the model that can be varied, to examine the sensitivity of the model to certain assumptions about controller workload and degree of aversion to risk. We describe the results obtained from an implementation of the procedure when applied to a range of traffic scenarios. Section 5 discusses how we propose to extend the approach in order to investigate proposed new ATM concepts such as User Preferred Trajectories.

2. Literature survey

There is growing industry consensus that 4D trajectories (2D routes plus altitude plus time) will form the basis for more efficient ATM operations in the future [Mohleji03, ASTRA03, Euro03, FAA06, ICAO03]. Benefits expected from a shift to 4D trajectories include:

- Airlines will have a much greater say in optimising flight trajectories to suit their business needs [ASTRA03]. Trajectories can be tailored to take advantage of expected winds and aircraft load conditions, and to suit different operational imperatives such as fuel use or time of arrival. Following the shortest route is not necessarily optimal, especially for long-haul flights in the presence of jet streams.
- 4D trajectories provide a precise description of the aircraft's intent. Distributing trajectory information will thus enhance situation awareness of the different agents involved in ATM decision making [Euro03].
- Full gate-to-gate management of trajectories is expected to lead to better system load balancing and thus increase the capacity of airspace [FAA06].

A number of research programs over the last decade have studied how an ATM system based on 4D trajectories might work in practice. Simulations within the DAG-TM framework [DAGTM99] have shown that, in an en-route operational environment where trajectories are negotiated between pilots and controllers, aircraft can meet their estimated time over a feeder fix more accurately than in the current operational environment [Lee03]. In a 4D trajectory-based ATS, aircraft fly at or close to their requested cruise flight level for longer periods of time than in the current system [PHARE97]. These last two studies both reported concerns about potential controller workload when switching to a trajectory based ATM system. There is need for further research into the ATM concepts based on 4D trajectories [NGATS06].

Most of the studies mentioned above treat conflict resolution as separate from trajectory planning. Some authors have proposed frameworks that encompass the two concepts (eg [Haraldsdottir03, Post99, Leiden00, Prevot03]), but do not try to integrate conflict resolution and 4D trajectories as a controller procedure. The concept proposed here tries to do so, in a manner consistent with current controller best practice.

3. Integrating intervention decisions with timing-based trajectory management

This section outlines the proposed concept for integrating tactical (sector controller) decision making into strategic (system-wide) trajectory management. Section 4 describes the results of some modelling and simulation experiments to evaluate the concept.

3.1. Background: the sector controller's task

This section is included by way of contextual background, to outline the kinds of intervention that en-route sector controllers make in the distance-based separation management approach to ATC as currently practised in Australia and to explain the difficulties in attempting to model operator behaviour in a top-down reductionist manner.

For the purposes of the current discussion, the main reasons interventions are required are:

- to resolve conflicts: ie, to ensure that aircraft are separated by at least 5NM whenever they are flying within the same Flight Level (FL).
- to authorise an aircraft to climb or descend to another Flight Level, in order to reach preferred cruise level or to descend into an airport. Such FL clearances are often limited in extent, for example because of the geometry of the sector (eg a requirement to stay above a certain altitude until a certain point is reached) or because of inter-sector protocols that limit the range of altitudes at which an aircraft can be handed off from one sector controller to the next.
- to divert around weather – but this won't be treated here.
- to meet downstream flow management requirements, eg by holding a flight until a landing slot is available for it.

A wide range of interventions is available to controllers, but for current purposes the main possibilities are as follows:

- Change Flight Level – either to resolve a conflict by moving one of the aircraft out of the way of the other, or by causing a climbing or descending aircraft to level out to prevent them from violating separation; or simply to clear them to continue their next stage of climb or descent.

- Change speed: typically used when a only small correction is required in order to ensure separation or to improve flow (eg spacing).
- Vector off course temporarily: typically used when a larger correction is required to ensure separation or to improve flow.
- Put in a holding pattern: typically used when flow management indicates that a large delay (more than 5 mins) is required to sort out problems further downstream.
- Track parallel: move off to the left or right of the nominal route by 5NM in order to overtake a slower aircraft travelling on the same route, or to let a faster aircraft pass.

Other forms of vectoring are sometimes used to resolve conflicts (such as "pointing" one aircraft towards another aircraft's current position, in order to let the first aircraft pass safely behind the second aircraft) but these will not be treated here.

Controllers get trained in some basic heuristics for selecting how and when to intervene in order to ensure that flow in their sector is safe and orderly. Over time, as they learn the common patterns of traffic and the traffic configurations that crop up repeatedly, they develop their own heuristics, some of which are almost subconscious, or at least are hard to articulate. Attempts to formalise these heuristics (eg [Spaeth01]) typically follow a classical reductionist approach, roughly along the following lines:

1. Determine the trigger conditions (ie the conditions that determine when an intervention might be necessary, such as when a conflict is detected, or when an aircraft is approaching Top Of Descent and requires clearance to start descending from cruise level).
2. Determine the range of interventions that might achieve the desired goal.
3. Eliminate the options that are not feasible (this step may be interlaced with the previous step).
4. Select the one which best fits the current circumstances (which may involve other consideration such as how it will affect sector traffic patterns later in time).

While this approach has had some success in cases where the task is highly constrained (eg [Ljundberg92]), in more general cases the number degrees of freedom available to the controller and the difficulty of defining a suitable fitness function makes it very difficult to come up with an agent-based model that satisfies this task [Callantine02].

Note that there are other aspects of the controller's task that can contribute to intervention decisions but which will not be treated here, such as accepting and handing off aircraft and issuing weather advisories.

3.2. The intervention heuristic

The heuristic we propose is a generalisation based on solutions generated by expert controllers conducted in the ATC Workload collaboration between the University of Queensland and AirServices Australia. A series of traffic scenarios – both static and dynamic – were prepared by AirServices and presented to individual controllers, and their solutions and reasons for choice of intervention were recorded. (The results will be published elsewhere.) A wide variety of solutions were employed, which confirms that controllers are not trained to follow a rigid set of rules, but develop a whole set of different tactics and implement them in an individual way. For example, some controllers appeared to be more "risk averse" than others and applied greater separation standards than the prescribed 5NM standard. Some intervened early to "prevent situations developing" and others let the scenarios play out much longer, to see whether intervention was really required. The wide variety indicates that is it likely to be very difficult to develop a "one size fits all" model of controller behaviour.

The heuristic described below embodies many of the rules observed and emulates aspects of the different controller solutions. While it does not emulate any one particular controller, our experiments to date indicate that it generates credible solutions.

As explained, we first make some simplifying assumptions about sector design and how flow management requirements are expressed, in part to make subsequent modelling simpler, but also to "cut down the trees to better see the forest":

- Each flight has a predetermined route, given as part of its Flight Plan.
- Somewhere on the route, close to where the route would cross the sector boundary (which may be a vertical boundary or a horizontal one, such as the floor of a feeder sector) we will assume there is a waypoint which we'll call the *Handoff Fix* for that flight in that sector.[1]
- A *Target Flight Level (TFL)* is assigned to the Handoff Fix for the flight: this represents the last Cleared Flight Level that the flight would normally get assigned in that sector.[2] If the

---

[1] It is relatively straightforward to extend this notion to cover multiple waypoints, both inside and outside the sector, each with targeted times and/or flight levels. We use a single waypoint near the sector exit boundary for simplicity of modelling.

[2] Again, this notion could be extended to cover a range of flight levels if desired.

aircraft is en-route, the TFL would thus normally be the preferred cruise level. If the aircraft is climbing or on descent, the TFL will be determined by the inter-sector protocol on Standard Assignable Flight Levels for the two sectors involved (eg a feeder sector controller might give clearance for the flight to descend to 8000ft prior to it entering the terminal approach sector).

- A target *Time Over Handoff Fix (TOHF)* is assigned to the flight by an external agent (typically the Flow Controller). This notion is already in operation in many parts of the world. For example, in Australia a tool called Maestro is used to calculate aircraft time targets for terminal-approach feeder fixes, using knowledge of the aircraft's likely descent profile and landing-slot time.

We divide the interventions into two different categories, each with an internal ranking:

Flight Level (FL) based intervention: Instruct a change of FL to X say, where X is a value between the current FL and TFL which is *safe* – ie, would resolve the conflicts (if any) and introduce no new conflicts. If the flight is already at TFL, choose a value of X which is safe. In both cases choose X to be as close to TFL as possible.

Timing based intervention: This category consists of all the other interventions we are considering: ie, changing speed, vectoring, parallel tracking and holding. These interventions have no effect on FL but can significantly affect TOHF; by contrast, changing FL has relatively little effect on TOHF. So if TOHF is more important than TFL (see below for more details), choose the intervention which is safe and which minimises the effect on TOHF.

The heuristic can be briefly described as follows (see Section 4 for a more detailed explanation of the timing-based part of it): The controller scans each flight in turn and checks to see if it is in conflict and/or is not going to meet its TFL and TOHF. If so, the controller chooses whether to consider a FL-based or timing-based intervention as follows:

- If the flight is climbing towards cruise level, or is but approaching its nominal Top of Descent or has started descent, use the best safe FL-based intervention.
- If the flight is at cruise level and a long way away from its nominal Top of Descent, then use the best safe timing-based intervention.

If the flight is in conflict with one or more other flights, the controller may decide to intervene on the other flights instead, if the first flight is on track to achieve its TFL and TOHF.

It is of course possible to generate scenarios where the above heuristic will fail, but we postulate that if it is applied early enough and often enough, such cases will not arise in practice, providing of course the TOHF targets are set appropriately.

4. Modelling and simulation of the proposed heuristic

This section describes the results of some preliminary modelling and simulation of timing-based aspects of the concept. The TFL part of the heuristic is not treated here, but Flight Level changes may be used to resolve conflicts.

4.1. Implementation of the model

In order to test the heuristic and to explore its robustness, an agent-based model has been developed which implements a proceduralised version of the heuristic. The model simulates agents for the controller of the sector in question, for a controller of the rest of the airspace, and for each airborne aircraft. The model gets run on *scenarios* which consist of sector configuration data, aircraft performance data and flight plans. For each aircraft in a scenario, one of the way-points in its flight-plan gets tagged as the Handoff Fix and assigned a (fixed) target Time Over Handoff Fix.[3]

The aircraft agent simply starts by flying according to its flight plan; it responds immediately to any instructions issued by controller agents. The sector controller agent can intervene if they detect that an aircraft under their jurisdiction has a conflict or is off time. The look-ahead time and distance separation standard used are parameters that can be varied to simulate controllers with different characteristics.[4] The non-sector controller is there simply to take part in the handoff/accept interaction with the sector controller.

The controller agent maintains a queue of aircraft to be scanned, with aircraft being added at the head of the queue at, e.g., take-off or sector entry and removed from the queue when they are handed off. The agent is activated at regular intervals (the scan interval), and

---

[3] The current results are thus of interest primarily for upper-airspace air sectors, where intermediate flight levels are less of a concern, since aircraft are typically flying at cruise level, or in the final stages of climbing towards it, or need to be given Top Of Descent clearance (down to a pre-determined level). This situation contrasts with feeder sectors and "intermediate level" airspaces, where a significant number of aircraft need to climb or descend through intermediate flight levels, and handing them off at the appropriate flight levels and/or with the appropriate flight level clearances is an important part of the controller's task.

[4] For example, some controllers report that they use a "soft standard" of say 8NM or even 10NM as a kind of added safety margin, rather than simply use the 5NM mandated minimum.

checks the aircraft at the head of the queue to see if it is in conflict or is not going to meet its timing target (TOHF). Parameters are used for key values, so we can test sensitivity of the heuristic to different factors: eg the *scan interval* (the time the elapses between scans of successive aircraft – a fixed amount here), *distance tolerance* (the separation standard applied), and *timing tolerance* (how much predicted TOHF needs to differ from target TOHF before an intervention is deemed necessary).

The intervention logic is given below. Note that this function has to cope correctly with a wide variety of situations. Not only may the scanned aircraft be in conflict or off time (or both), but it may not have a target time-over, or it may conflict with more than one aircraft, or it may conflict with aircraft not in the controller agent’s jurisdiction. Note also that the agent can choose to intervene on the scanned aircraft *or* on the aircraft it is in conflict with, provided they are all under the controller’s jurisdiction.

```
set SS, the current solution set, to empty
set BB, the current solution badness, to infinity
if AA has a target way-point in its future flight-plan
  set WW, the active way-point, to the target way-point
  set TT, the required time over, as in flight-plan
else
  set WW to the last way-point in AA's flight-plan
  set TT to the predicted time over WW
calculate BB*, the current badness of AA with respect to WW/TT
set MM, the set of manipulable aircraft, to {AA}
if AA is in conflict
  add those conflicting aircraft controlled by Agent to MM
  set WWi/TTi, as for AA, for each of the added aircraft
generate LL, the list of all intervention sets involving subsets of MM
for each II in LL, in order
  if II is not safe then next II
  calculate BB', the badness of II with respect to WW/TT and the WWi/TTi
  if BB' < BB
    set SS to II and set BB to BB'
if (SS is empty) or (AA not in conflict and BB* <= BB)
  report/log no new intervention
else
  implement SS
return
```

Intervention logic

The intervention list for each aircraft consists of a short selection of speed changes, off-track vectorings and height changes, with each being stepped over a range centred on its nominal/current value. Each list is headed by the speed change required to meet the timing target, if this is possible, and entries from the lists are combined to generate all possible intervention sets, with unsafe sets being rejected. Interventions are ranked by their *badness*, which is the absolute value of the amount by which an aircraft is off time or, if there is no recorded timing target, the amount by which the time over the last way-point would be altered. The badness for a set of interventions is the sum of the badnesses for the individual interventions, and the set with smallest badness is used. Note that there may not be a solution (in which case no intervention is taken), and that more than one intervention may be taken in the same scan step (eg if interventions are taken on multiple aircraft that conflict with the scanned aircraft).

4.2. Results

Our simulations are driven by scenario files which give begin and end times, the control sector and the aircraft flight plans (call-sign, type, entry time, cruising altitude, way-points, time over). These files can be created manually for test/illustration purposes or can be derived from recordings of actual traffic. We have tested the model on a wide variety of scenarios, to check the correctness of our implementation and to observe the model’s behaviour. Generally the model accurately and safely “solves” the scenarios, and its performance is robust in the face of parameter variations. Of course, if the scenarios are unrealistic or the parameter values extreme then reasonable solutions may not be possible. The model also displays some idiosyncrasies that need to be addressed.

To illustrate our discussion we will use a test scenario based on Figure 1: the scenario is completely hypothetical and kept small for the purposes of illustration.[5] The sector is a 100NM radius circle. Flight OAF142 enters at time 9h55m00s, the others at time 10h00m00s, and they all fly towards the handoff fix HF and then off to the right. The aircraft are of different types, and they are all at the same height throughout. At their nominal speeds, they pass over HF in the order UAL, BIZ, OAF, AAL, with approximate timings 10h27m10s, 10h29m35s, 10h32m30s, 10h34m40s. The UAL/BIZ and OAF/AAL pairs come into conflict after this. To test the intervention heuristic, TOHF times are set for the four aircraft as follows: BIZ

---

[5] Note in particular that this scenario is not at all representative of the complexity that arises in intermediate sectors when aircraft are climbing and descending on different routes.

(10h30m), AAL (10h35m), UAL (10h38m), OAF (10h41m).

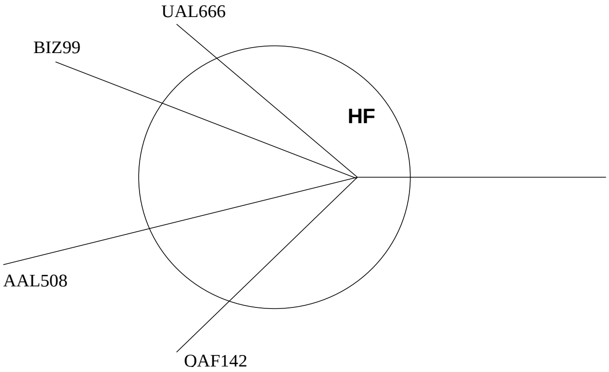

Figure 1: The 4-aircraft scenario used for testing

The first point to note about the model's behaviour is that it is a discrete system, and its behaviour is very sensitive to the parameters and the scenario used. In order to mitigate the effect of this, we incorporated into our parameter sweep code the ability to make multiple runs for each setting, where each run is over a slightly different scenario. This was achieved by simply offsetting the aircraft entry times by a value drawn from a normal distribution of standard deviation 30 seconds and mean zero. We did 100 runs for each setting and report the averaged results. Our input parameters were the interval between the agent's scans of its queue, the timing tolerance threshold that triggers intervention, and the separation standard enforced. The model's behaviour for this particular set of scenarios was found not to be very sensitive to the distance tolerance applied, so it was set at 8 NM for the remainder of the experiments.

Figure 2 illustrates how the model's performance degrades abruptly if the controller's scan interval is too long. The *aggregate badness* is the sum of the badnesses for the four aircraft; it should be compared with the pre-intervention badness of 1216 seconds. Note how the aggregate badness increases as the timing tolerance increases. In a run in which TOHF is achieved for all 4 aircraft, the aggregate badness will be less than four times the tolerance. For larger tolerances intervention is not triggered for BIZ or AAL, since they are off time by small amounts. An abrupt degradation in ability to resolve conflicts also occurs around the same point: until the agent's scan time exceeds 100 seconds the aircraft are on time and there are no conflicts.

Figure 3 shows the average number of interventions that are triggered as the model tries to achieve the TOHF targets safely. In principal, four interventions should suffice for small tolerances and two for large tolerances. However the model may not be able to find a safe intervention which reduces the badness to within its tolerance or, indeed, which reduces the badness by any amount. Thus, in some circumstances, the model may intervene many times for a single aircraft, reducing its badness in many small steps until it is within tolerance. In part this behaviour is due to the fact that discrete speed values were considered, and so interventions typically overshoot or undershoot. But uncertainty of this nature will undoubtedly arise in practice, so this is one place where the heuristic needs improvement.

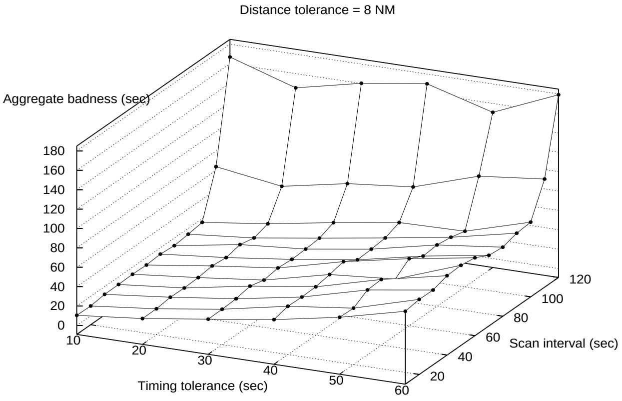

Figure 2. Sensitivity of aggregate badness to scan interval and timing tolerance

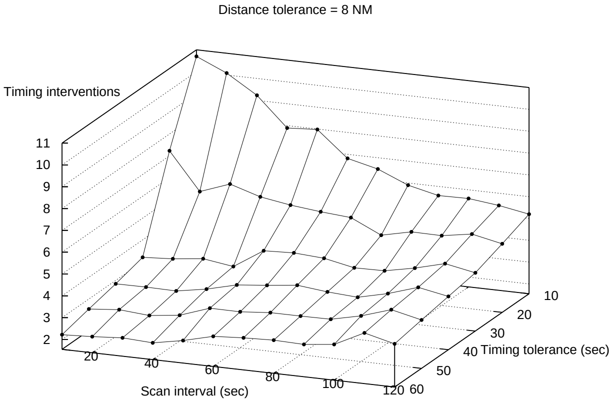

Figure 3. Sensitivity of number of interventions taken

The last of our results reported here concerns the amount of computational processing involved in the procedure. In practice this would relate to some combination of workload imposed on the controller and/or the required performance of the tools involved, but here we use running time of our simulation model as a proxy measure. We would expect that, the more frequently the agent scans aircraft and the tighter the tolerances they work to, the longer the model takes to run. That this is indeed the case is illustrated in Figure 4. The running times here are short enough to support real-time use. However in more complex situations, particularly when trying to resolve multi-aircraft conflicts, the running time can increase dramatically as we generate and test all intervention sets.

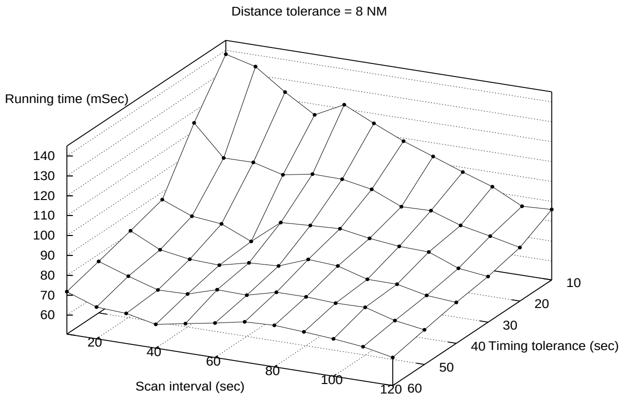

Figure 4. Sensitivity of total computation time

5. Summary and conclusions

In summary, the paper reports early results in developing an approach that integrates sector controller interventions into 4D trajectory management. The approach is expressed as a heuristic, or combination of rules for how and when controllers make interventions that combine three important aspects of en-route sector control: conflict resolution, issuing of flight level clearances, and achievement of time-over-waypoint requirement (although the experiments reported here do not cover the flight level clearance aspect except in as much as it relates to conflict resolution). The heuristic was derived from observations about how expert controllers resolve problems in radar-controlled airspace in Australia, but it seems to be reasonably generic, at least for sectors that have timing requirements at waypoints.

For the purposes and modelling and simulation the heuristic was implemented as a parameterised procedure and tested on a wide range of scenarios, including traffic and sector configurations based on real-world data (not reported here). The results indicate that the procedure comes up with credible solutions when its key parameters are kept within reasonable bounds, but begins to break down when those bounds are crossed. We think the approach has significant potential as a test bed for investigating the complexity that may arise in new sector designs and/or traffic configurations, or new operational concepts such as User Preferred Trajectories (UPTs) [ASTRA03].

Note that we are *not* proposing that our heuristic become standard procedure for sector controllers. We are simply trying to develop a model of “suitably realistic” controller behaviour to form the basis for our subsequent studies into how trajectory-based ATM concepts can be accommodated within (an extension of) the current ATC environment. Our working hypothesis is that, while our model does not generate all of the possible outcomes that human controllers would generate, it is sufficiently representative of controller behaviour that we can use it to test new concepts. For example, we are using it – and the trajectory-manipulation environment that supports it – to investigate what happens if some of the responsibility for detecting problems and proposing solutions is delegated to airlines, either at tactical level (by proposing interventions) or at strategic level (by negotiating timing requirements).

It could of course be said that we have only tackled one part of the problem, in that we have not discussed how timing requirements would be set at waypoints. In particular, if they are not set appropriately, they could result in unsafe traffic patterns, or undue workload for the controller. This is an area we propose to consider more closely in the future, together with the issue of how airlines can best be brought into the decision making process. Our working hypothesis is that UPTs will require airline input at both levels: at indicating a preference for the kind of intervention that gets taken if one is needed (eg depending on whether they want to optimize for time or fuel use), and at renegotiating trajectories at “macro level”, as represented by timing requirements at important waypoints (assuming that the user preferred 3D route is followed fairly closely).

It will be important to understand the system-level effects of interactions between strategic and tactical decision making, to avoid undesirable consequences. An example of an unintended event cascade could be the following: a sector controller slows down a flight temporarily to avoid a conflict; this results in a Terminal Arrivals System thinking that the aircraft will not achieve its current landing slot and allocating a later slot to it; this in turn results in the flow tool allocating a later TOHF to it; this in turn results in the controller delaying the aircraft yet further; and so on.

Acknowledgements: This research was supported by the Australian Research Council’s Centre of Excellence scheme funding. The author gratefully acknowledges the contributions of Colin Ramsay, Martijn Mooij, Ariel Liebman and colleagues in the ATC Workload project. In particular, Colin did most of the design and implementation of the trajectory simulator, and generated the results reported here, and Martijn has been invaluable in monitoring the literature and providing domain expertise.

6. References

[ASTRA03] Australian Strategic Air Traffic Management Group. *Air Traffic Management: A Strategic Plan for Australia. Vol 1: Optimising Safety, Efficiency, Capacity and the Environment*, 2003.

[Callantine02] T. Callantine. *CATS-based air traffic controller agents*. NASA Contractor Report 2002-211856. Moffatt Field, CA: NASA Ames Research Center.

[DAGTM99] NASA. *Concept Definition for Distributed Air/Ground Traffic Management (DAG-TM)*, Version 1.0, 1999.

[Euro03] Eurocontrol. *Air Traffic Management Strategy for the Years 2000+* (Vol 2), 2003.

[FAA06] FAA, *National Airspace System (NAS) Architecture*, V6.0, 2006.

[Haraldsdottir03] A. Haraldsdottir et al. *Boeing capacity-increasing ATM concept for 2020*. 5th Eurocontrol / FAA ATM R&D Seminar, Budapest, Hungary, 2003.

[ICAO03] International Civil Aviation Organization. *ATM Operational Concept Document*, ICAO AN-Conf/11-WP/4, 2003.

[Lee03] P. Lee et al. *Free maneuvering, trajectory negotiation and self-spacing concepts in distributed air-ground traffic management*, ATM2003, FAA/Eurocontrol R&D Seminar, Budapest, Hungary, 2003.

[Leiden00] K. Leiden & M. Green. *Trajectory orientation: a technology-enabled concept requiring a shift in controller roles and responsibilities*. 3rd USA/Europe ATM R&D seminar, Napoli, Italy, June 2000.

[Ljungberg92] M. Ljungberg & A. Lucas. *The OASIS air-traffic management system*. In Proceedings of the Second Pacific Rim International Conference on Artificial Intelligence, PRICAI '92, Seoul, Korea, 1992. http://citeseer.ist.psu.edu/ljungberg92oasis.html

[NGATS06] NASA. *Next Generation Air Transport System (NGATS): Research Opportunities in Aeronautics*. NASA Research Announcement NNH06ZNH001, 2006.

[PHARE97] Eurocontrol. *PHARE PD/1 Final Report*, Annex C, DOC 96-70-24, 1997.

[PHARE00] Eurocontrol. *PHARE Advanced Tools Project Final Report*, PHARE/EQH/PAT-8.2/FR1.1, March 2000.

[Post99] W. Post & I. Wilson. *Conflict detection & resolution: lessons learned*. Paper presented at the Technical Interchange Meeting. Memphis, TN, USA, 1999.

[Prevot03] T. Prevot et al. *Air traffic concept utilizing 4D trajectories and airborne separation assistance*. Proceedings of the AIAA Guidance, Navigation, and Control Conference, AIAA-2003-5770, Austin, TX, USA, 2003.

[Spaeth01] O. Spaeth & K. Eyferth. *Conflict resolution in en route traffic – a draft concept for an assistance system compatible with solutions of air traffic controllers*. MMI-Interaktiv, Nr. 5, July 2001, ISSN 1439-7854.