STRATEGIC OPTIONS FOR FUTURE AIR TRAFFIC SYSTEMS

S A N Magill

ATC Systems Research Division
Royal Signals and Radar Establishment
Malvern, Worcs, WR14 3PS, UK

Abstract

This paper explores some possible ways of increasing the air traffic capacity of the busiest parts of en-route airspace in the early years of the next century. Because controller workload is the key capacity-constraining factor, and because there is no prospect of a completely automatic system in the foreseeable future, certain controller-related constraints have a dominant effect on any proposed innovations. Technological developments in navigation, computing and communications are summarized. Options for future air traffic systems are discussed against this background, including those in route structure, airspace sectorization, time control, and flow control. Finally, there is a brief description of a computer simulation which is being built to attempt to quantify the capacity implications of these options.*

1. Introduction

Throughout much of Europe, the demand for air traffic services is rapidly outstripping the capacity of existing Air Traffic Control (ATC) systems. In some areas it is necessary to use Flow Management to restrict traffic flows to manageable levels.[1] Eurocontrol has forecast that the annual demand for IFR movements will double in the years between 1987 and 2000, ICAO.[2] Although much of the increase in annual demand will be accommodated by extending busy periods, the increase in peak-period demand will still be of the order of 60-70%, ICAO.[2] For regulated airspace in the UK, Hunter and Brooker[3] have estimated that peak-period loadings on some en-route ATC sectors will increase by as much as 41% in the years between 1990 and 2000. Although traffic demand forecasting is a notoriously difficult and imprecise business, there can be little doubt that an increase in peak-period demand of the order of 50-70% above today's levels will occur in the not-too-distant future, and that growth will not stop there. If parts of Europe's ATC systems already have insufficient capacity, where is the additional capacity to come from to meet the forecast increased demand?

The capacity problem is more apparent in terminal areas surrounding airports where runway capacity can be quantified fairly precisely, than in en-route airspace where the very concept of capacity is rather elusive. It is possibly for this reason that much more research attention has been devoted to the terminal area problem than to the en-route problem, for example Magill et al,[4] Völckers,[5] Benoît and Swierstra.[6] However, in the long term it may be en-route capacity which will be the greater limitation. Although runways at the busiest airports already operate close to capacity during peak periods, there are many other runways operating at a fraction of their capacity most of the time, and as demand grows, more and more of this unused capacity will be pressed into service. In the UK for example, while Heathrow and Gatwick are already operating close to capacity, Stansted, London City, and the main regional airports have plenty of unused runway capacity, and many minor airports have great potential for developing business traffic. It is also possible that some airfields now used for military purposes will eventually become civil airports. It is time to focus more research attention on the en-route capacity problem.

Against this background, a research project known as 'Air Traffic System Strategy Studies' (ATSTRATS) has been set up and funded by the UK Civil Aviation Authority at the Royal Signals and Radar Establishment. The aim of the project is to find ways of significantly increasing the traffic capacity of the busiest parts of the en-route ATC system in the early years of the 21st century. ATSTRATS will take full account of the contributions which might be made by new technologies – computing, avionics, communications – but will focus more on the ways of organizing and controlling traffic made possible by the new technologies than on the technologies themselves. Some options for future air traffic systems are discussed in this paper.

The paper is structured as follows: Section 2 makes some observations about the nature of traffic capacity and the role of the human air traffic controller. It concludes that a completely automatic ATC system is not an option for the early part of the next century, and that a particular way of looking at possible innovations results from that conclusion. Section 3 attempts to summarize very briefly the contributions which new technology might make. Sections 4 to 7 form the core of the paper. They discuss respectively air route structure, airspace sectorization, time control and flow control, all from the point of view of capacity implications. Section 8 briefly outlines the computer simulation which the ATSTRATS team is building to further study and evaluate the ideas presented here. A fuller discussion of the computer simulation can be found in Magill.[7]

*This research is funded by the UK Civil Aviation Authority. This paper expresses the author's personal views, and is not to be regarded as an official statement by the Civil Aviation Authority or the Royal Signals and Radar Establishment.

Copyright © 1990 by ICAS and AIAA. All rights reserved.

2. Traffic Capacity

The traffic capacity of a given volume of en-route airspace is not simply that which would result from packing all aircraft at minimum allowed separation. If it were, then ATC system capacity values in today's systems would be several times larger than those observed. Most informed observers are now agreed that the key constraint on capacity is Air Traffic Controller workload, see for example Parker,[8] Richmond,[9] Miaillier and Planchon.[10] The capacity of a control sector can be defined as the highest mean flow rate which can safely be sustained for a long period. Traffic flow rates can exhibit large random fluctuations, and controllers can safely cope for short periods with flow rates significantly larger than the mean rate, provided there is opportunity to recover from one peak before the next arrives. Thus, sector capacity should be seen as an 'elastic' limit which is frequently exceeded for short periods, rather than a 'hard' limit which must never be exceeded.

ATC systems already make use of computer technology for a number of routine tasks, including radar tracking and display driving, flight plan data processing, and short-term conflict monitoring. It is expected that there will be a substantial increase in the use of computer technology in the near future. Applications will include: electronic storage, display and communication of flight progress data, long-range trajectory prediction to facilitate track and conflict monitoring, and planning tools such as conflict resolvers. With the development of ground/air digital datalinks it is at least in principle conceivable that the whole ATC task might one day be performed completely automatically. This would have the attraction that it would remove the primary capacity constraint identified above. However, there are good reasons why the time at which it might be feasible to build a completely automatic ATC system is so far in the future as to be beyond our timescale:

1. Software technology has not yet developed to the point where it is possible to build a system as large and complex as a fully automatic ATC system, and guarantee that it will perform sufficiently safely under all circumstances.

2. There are legal liability issues. If a fully automatic ATC system should fail, who is to be held responsible?

3. Even when these difficulties have been overcome, there is likely to be resistance from the airline and insurance industries, and from the general public.

A fully automatic ATC system is not an option for the early years of the next century, and controller workload will continue to be the chief capacity constraint. Once this point is accepted, it colours the rest of our thinking about future ATC systems. It follows that all attempts to increase capacity through system design and technological innovation should be directed towards supporting the controllers in their tasks. We must seek ways of reducing the work done per aircraft handled, and ways of subdividing the total ATC task into subtasks with minimal interaction between them, so that more controllers can be applied in parallel to the total task.

When trying to develop new ATC system concepts there are a number of controller-related limiting factors which must be taken into account.

1. When automation is used to support air traffic controllers, this tends to reduce their direct involvement in the details of the traffic situation in hand. On the face of it, reducing controller involvement in details is good for workload and capacity. However, controllers must not be abstracted from detail to the point where they are unable to apply the whole of their human problem-solving skills to specific separation problems which might arise. The idea of a 'high-level' or 'more managerial' role for controllers is an attractive one, but it has not yet been demonstrated, and it is certainly not self-evident, that such a role can be found.

2. The kinds of traffic configuration which arise in the system must be amenable to the kinds of pattern processing which humans are good at, and must at all times be fully comprehensible to controllers. It is likely that simplicity and uniformity of procedures and trajectories will be more successful than complexity and diversity.

3. When automated functions are used to perform some ATC subtasks, because the human air traffic controller is in charge of the total task, communications bandwidth problems can arise between the controller and the automated functions. For example, if a complex set of clearances is generated by a ground-based computer and communicated to the aircraft by digital datalink, and if controller approval of the clearances is required, then the amount of data which the controller can absorb in the time available will limit the complexity of the set of clearances. This limitation seriously constrains the types of system organization which might be feasible.

One contributor to the capacity problem worthy of special note is ATC system uncertainty about predicted aircraft trajectories. Uncertainty contributes to the problem in a number of ways:

- Because of uncertainty, many situations have to be treated as potential conflicts which would not actually develop into real conflicts. This generates extra controller work directly in that conflict-avoiding actions have to be planned and additional monitoring is needed, and indirectly because the complexity level of the total task is increased.

- Actual aircraft separations planned and used are often much greater than the permitted minimum separations to allow suitable margins for uncertainty.

- Uncertainty adds to the difficulty of producing effective long-range automatic conflict monitoring aids. Because of uncertainty, it is difficult to achieve a high detection rate and a low false-alarm rate simultaneously.

The considerations of this section pervade most of the ideas in the remainder of this paper.

3. Application Of New Technologies

This section is not concerned with technology applications already incorporated into the design of future ATC systems, but rather with those technology applications which although proved feasible, and possibly in service on board aircraft, have not yet significantly influenced ATC system design. One could focus on the underlying technologies – computing, avionics, communications – or on their manifestation in ATC systems – more precise navigation, improved predictability, time navigation. This section adopts the latter approach.

3.1 Navigation

There have been two important developments in aircraft navigation upon which ATC system designers should be able to capitalize:

1. It is no longer necessary for aircraft to fly from one radio navigation beacon to the next. Area Navigation allows the possibility that aircraft can fly directly from departure airport to destination airport, and raises the question of whether or not an airways system is really necessary.

2. Aircraft can now be navigated so as to pass closer to a target set of points in space than ever before. This capability arises from the combination of several navigation aids (as for example in multiple-DME navigation), and from satellite navigation of the GPS/GLONASS type.

3.2 Transfer Of Information

Until now the main means of communication between ATC and aircraft has been by VHF/UHF radio telephony. The advent of ground/air digital datalinks carried by Mode S secondary radar or satellite brings new possibilities. The information currently exchanged by radio telephony might be exchanged by datalink with much less controller workload. Much more control information might be exchanged, either more complex or frequent clearances or aircraft intentions. Much more ancillary data might be exchanged, including meteorological data, aircraft state data, and trajectory prediction made by flight management systems. With greater use of computers in both ground and airborne systems, ground/air digital datalinks increase the possibilities for completely automating some functions.

3.3 Trajectory Prediction

In section 2, predictability of future aircraft positions was identified as one of the factors contributing to the capacity problem. A number of technology applications can contribute to greater predictability:

1. Use of a suitable computer model of aircraft behaviour. The computations of such a model could be performed on board each aircraft and transmitted to ATC control centres by digital datalink, or could be performed in each control centre. In the latter case, a database of parameters for many different aircraft types would be required. Renteux[11,12] describes a suitable aircraft performance model and parameter database for ground-system use.

2. Improved knowledge of wind conditions aloft. This could come either from wind measurements made by individual aircraft and transmitted by digital datalink, or from wind profiler radar equipment.

3. Transmission by digital datalink of precise details of how an aircraft (crew or avionic system) intends to proceed, for example, the precise place or time where the next turn or height change will begin.

3.4 Time Navigation

Time navigation is the capability to arrive at a particular point in space at (or close to) a specified time. This capability is obviously of great value in the terminal area where it is necessary to arrange aircraft into a strict time sequence prior to arrival at the runway, but it might also have benefits in en-route airspace.

4. Air Route Structure

Airspace organization, including both route structuring and division into control sectors, probably provides the single most powerful tool for increasing the capacity of future air traffic systems. Unfortunately, very little of a general nature is known about the relationship between airspace structures and capacity. Today's ATC system has reached its current state through a long sequence of small evolutionary steps. Attempts to make local improvements have been strongly conditioned by local geography and local traffic demands, and very few general principles have been determined. A more general study of the whole subject is long overdue.

Two important questions must be addressed:

1. Given the new capability of aircraft to navigate directly from departure airport to destination airport, can any capacity benefit be gained by imposing structure which prevents aircraft from flying direct routes with optimal altitude and speed profiles? (There are other kinds of arguments for having an airways system, for example interaction with military traffic, but only capacity arguments are considered here).

2. If it turns out to be beneficial to impose structure, then which kinds of structure bring the greatest capacity benefits?

Quantitative answers to these questions can really only be obtained from computer simulation (see Section 8), but there are various qualitative and intuitive arguments which can cast some light.

4.1 Direct Routes

The airways system in use today originally grew up because, in the early days of radio navigation, it was necessary for aircraft to fly from beacon to beacon. ATC procedures grew up in parallel with the airways system in response to the increase in air traffic, and so the 'airways way of looking at things' became part of the 'culture' of ATC. Today, if there were just one aircraft in the sky (and it had modern avionic equipment on board), the only sensible course of action for its crew would be to choose the Great Circle route from departure airport to destination airport (with possible deviations for weather systems), and to choose the altitude and speed profile to give the cost-optimal flight. But ATC is not about the single aircraft, it is all about handling many aircraft simultaneously. Since controller workload is seen as the chief capacity-limiting factor, is less controller workload per aircraft implied by a direct-routes system, or by a system based on airways?

One way of approaching this question is to consider the occurrence of conflicts (in the absence of avoiding action) in systems based on direct routes and on airways. Since a great deal of controller workload is concerned with detecting potential conflicts, and planning and implementing conflict-avoiding manoeuvres, this is not an unreasonable approach. Superficially it appears that an airways system concentrates all aircraft into restricted regions of space, increasing traffic densities and thus conflict probabilities. However, with a carefully designed network of airways it might be possible to segregate the various traffic flows in such a way as to greatly reduce interaction between them, and thus reduce the conflict probability below the direct-routes value. Assuming that, for a given geographical distribution of airports and traffic demands, a direct-routes system does generate fewer conflicts per unit time than a structured system, this in itself does not necessarily mean that the former system has a higher capacity than the latter. It may be that the human pattern-processing abilities function much better in a system with airways, so that although potential conflicts develop more frequently, the workload involved in conflict-avoidance is less. There is very little published research in this area.

So far it has been assumed that, in a direct-routes system, very few of the direct routes would in fact coincide with one another. But this might not be the case. The busiest regions of en-route airspace occur where there are several important cities situated fairly close together, with each connected to each of the others by a busy air route. An informal analysis of London FIR flight plan data done by the ATSTRATS project team indicates that a large majority of flights involve a surprisingly small number of city pairs. This suggests that, for the London FIR at least, if there was no formal airways system there would still be an informal one. Thus traffic would be concentrated into a small region of space without the benefit of measures to segregate the various flows – the worst of all possible worlds.

4.2 Structure In The Horizontal Dimensions

Two horizontal structuring options suggest themselves: unidirectional airways, and airways with multiple parallel lanes. Unidirectional airways essentially segregate the two flows in opposite directions between two end-points, and thus eliminate a whole class of potential conflicts. They are already used today, and will no doubt be more extensively used in future. Airways with several parallel lanes (inspired by the road traffic analogy) are the more innovative option, and here we will confine ourselves to considering that option. There are two reasons for using such a structure:

1. Airways with multiple lanes provide a means of distributing traffic over more space, thus reducing densities and conflict probabilities.

2. Multiple parallel lanes provide a new means of resolving conflicts which is much more predictable and requires less controller work than the method known as 'vectoring' used today.

Perhaps the second reason requires some elaboration. In current ATC practice, when an aircraft is 'vectored', it is asked to suspend its own navigation along its route and adopt a heading specified by ATC instead. Some time later it may be asked to adopt another heading, and eventually it will be asked to resume own navigation. This procedure generates extra work for ATC because the controller must give special attention to aircraft being vectored so as to determine when to return them to their own navigation, and because at least two control instructions are needed per vectoring operation. Vectoring also increases uncertainty in the total ATC system about such aircraft because their future trajectories are known only to the controller involved – other controllers and automated tools do not share this information. The implementation of conflict avoidance by a single very predictable 'change lane' instruction has much to recommend it.

The great difficulty with implementing a system with multiple parallel lanes is how to ensure with sufficient confidence that streams of traffic on neighbouring lanes will always be safely separated. For straight portions of route this is not particularly difficult. It is sometimes suggested that turns can be avoided altogether in an airways system, but this is not the case. There are various regions of airspace which must be avoided, including terminal areas around airports, danger areas, and military exercise areas. For turns on airways with multiple lanes, it will be necessary to define aircraft turning procedures very rigorously. It will be necessary to specify the points at which turns must begin and end, and the trajectories between these points.

Another problem occurs where two multi-lane airways merge into one, or one airway branches into two multi-lane parts. It is essential that each lane before the merge or branch should have a defined successor after the merge or branch, so that controller workload is not generated by the need to give each aircraft a clearance at such a point. But this implies that lanes will necessarily cross one another at merging and branching points. This in turn reduces the capacity of individual lanes, and produces a new source of potential conflicts. The problem of lane crossings at merging and branching points certainly requires further study, but it does not seem to be a great enough problem to invalidate the whole multiple-lane concept.

4.3 Structure In The Vertical Dimension

Two kinds of vertical structuring rule can be distinguished:

1. Global rules where allowed altitude is independent of position, but is related to some other aspect of flight such as heading or speed. The semi-circular and quadrantal rules used today are examples.

2. Local rules where allowed altitude is a function of horizontal position. This kind of rule is already in use at some busy crossing points near terminal areas.

The first kind of rule is only helpful for segregating traffic flows in level flight. In the busiest parts of Europe a large proportion of traffic is climbing or descending, so that rules of the second kind are more helpful. Although such rules are already used today, they would appear to have potential for much wider use for separating flows at busy crossing points. However they do have some problems. Their use will inevitably prevent some aircraft from flying optimal profiles. Also, if one route passes through several crossing points where altitude constraints are imposed, then there will be a strong interaction between the constraints at successive crossing points. This is so because an aircraft can gain or lose only a limited amount of altitude in the horizontal distance between two successive crossing points. Climb rates are very variable, and wind variations contribute to the variability of climb and descent gradients, so when planning how much altitude can be gained or lost in a given horizontal distance, conservative assumptions must be used. In some cases it might be possible to allow alternatives for both low-performance and high-performance trajectories.

In current ATC practice, aircraft are free to navigate along their routes in the horizontal dimension, but (except in the case of Standard Instrument Departures) they must be given explicit clearance to change altitude. In the long term, it might be possible to define three-dimensional routes in such a way that aircraft could follow them in all three dimensions without the need to obtain separate clearances for each altitude clearance. This could make a useful contribution to reducing controller workload and increasing traffic capacity.

5. Airspace Sectorization

A sector is a contiguous volume of airspace controlled by a single air traffic controller or by a group of controllers working together as a tightly-knit team. The way in which a volume of airspace is divided into sectors is intimately bound up with its horizontal and vertical route structure, and has a significant effect on its traffic capacity. There are two important sectorization issues: sector size, and the nature of inter-sector interfaces.

If a large sector size can be used, this has the advantage that each aircraft experiences fewer transfers of control from sector to sector (which generate controller workload), and the advantage that the sector's controllers have more room for conflict-avoiding manoeuvres. The maximum sector size is set by the number of aircraft which can be handled simultaneously. As traffic densities rise, sector size is forced downwards. As sector size becomes smaller, controllers spend more and more of their time co-ordinating and transferring control between sectors, and the lower limit is the point where they spend all their time doing this. Where possible, it is best to sectorize the airspace in such a way that neighbouring sectors work in parallel on traffic streams with minimal interaction between them, so that an individual aircraft passes through only a small subset of the total set of sectors. This principle is already seen in today's systems in the form of vertical sectorization, but its full potential has not yet been exploited.

In traditional ATC practice, when an aircraft flies through sector A followed by sector B, a controller in sector B must explicitly agree to accept control of the aircraft before control is transferred. This process of getting agreement is known as 'co-ordination'. It is usually carried out by telephone, and it can account for a great deal of controller workload. On some inter-sector boundaries near busy terminal areas another method known as 'agreed levels' has evolved. According to this method, sector B agrees in advance to accept all aircraft transferred to it by sector A provided they satisfy a specified set of conditions (altitude, direction, speed), and provided that sector B can ask sector A to stop the stream if the need arises. Thus the emphasis is shifted from individual aircraft to streams of aircraft. If the 'agreed levels' method could be generalized to work on all inter-sector boundaries, this could make an important contribution to increasing capacity. As practised at present, the method is limited by the repertoire of agreed conditions which can be used on any one boundary. Use of electronic storage and display of flight progress data in future systems will allow the possibility of a much larger repertoire of agreed conditions on any one inter-sector boundary, and generally facilitate this kind of inter-sector interface.

6. Time Control

The term 'Time Control' is used here to mean control of the time at which an aircraft arrives at a specified point in space. It may be achieved by varying the aircraft's path to the point, or by varying its speed, or by a combination of both. In current systems, Time Control is used by approach controllers to adjust the longitudinal position of each aircraft in an arriving stream relative to the rest of the stream. It is not used in en-route airspace except occasionally as an extension of the approach control process. The question addressed here is, could an en-route ATC system obtain significant capacity benefits by making use of Time Control?

The trajectory calculations needed for Time Control could be performed by computers on board aircraft (as in the case of the most modern flight management systems), or they could be performed by computers forming part of the ATC system. If Time Control is implemented by path-variation, then as argued by Magill[13] in another context, this must necessarily involve communication between the air traffic controller and the aircraft. Either the aircraft proposes and seeks clearance for an alternative path, or the controller asks the aircraft to fly an alternative path. If such manoeuvres are to be well-structured and predictable (as opposed to ad hoc vectoring manoeuvres), provision must be made for them in the route structure. This could be done by having a series of alternative but standard paths at turning points, or by having path-stretching loops alongside straight portions of route. If Time Control is implemented by the speed-variation method, this could be done entirely on board aircraft, and need not involve communication with the controller. However the speed-variation method does have another serious problem. At cruising altitudes the interval between the minimum and maximum normal operating speeds is very restricted. Assuming that an aircraft is operated somewhere near the middle of this interval, the amount of speed variation available for Time Control purposes might be of the order of ±5% of true air speed.

There appear to be two possible applications for Time Control in en-route ATC:

1. as a new member of the armoury of conflict-avoidance techniques;

2. as a means of increasing the predictability of aircraft trajectories.

Speed variation is not very attractive as a general conflict-avoidance tool because it is so slow-acting; for example, to obtain a 1-minute separation it might be necessary to begin speed variation 20 minutes in advance. Path-length variation is not very attractive as a general conflict-avoiding mechanism because the necessary path-stretching areas will not be available at all points on each route. However there are two special situations where Time Control might be appropriate for conflict avoidance: overtaking situations, and situations where traffic streams merge. These two situations are special in that a segregation method can not be used for them; they occur because the aircraft involved are part of the same stream. Overtaking conflicts arise because of speed differentials. They develop slowly, and so speed variation is an appropriate tool for dealing with them. The ATC system knows about potential merging conflicts a long time in advance, and so both path variation and speed variation can be used for avoiding them.

Time Control by means of speed variation could be used to improve the predictability of aircraft trajectories. If there are differences between the conditions assumed by the ATC system's trajectory prediction calculations and the conditions experienced by aircraft – wind conditions for example - these differences will lead to prediction errors. If aircraft can monitor their own along-track progress, and adjust their own speeds to achieve the desired rate of progress, such prediction errors will be greatly reduced. However this will only work if the discrepancies between the ATC system's assumptions and reality are small enough to be accommodated within the aircraft's range of operating speeds. Also, this method of improving predictability has a price tag attached to it: it does not permit aircraft to operate at their most economic speeds.

In summary then, en-route Time Control does not hold out the hope of dramatic capacity gains. However it does appear to offer some capacity gains through conflict avoidance in merging and overtaking situations, and through improved predictability. Further assessment of en-route Time Control must await quantitative results from computer simulation.

7. Flow Control

In the absence of Flow Control measures, the number of aircraft passing through a volume of airspace per unit time is not a constant or smoothly varying function of time; rather, it exhibits large random fluctuations. Randomness is an unavoidable feature of air traffic for several reasons. The variation in along-track component of wind velocity from day to day is greater than the range of cruising speeds available to most aircraft. The preparation of an airliner for departure involves many separate processes such as the control of passengers and their baggage, each with some associated randomness, so take-off times inevitably deviate from schedules. Competition for limited resources increases these effects. Randomness causes traffic streams to have bunches and gaps, and as a result, Flow Control procedures have evolved to limit the number of aircraft in a region of airspace at any one time. Flow Control is implemented by allocating time slots for take-off to individual aircraft, and to a lesser extent by using alternative routes. The question to be addressed here is, can greater use of Flow Control or use of more sophisticated forms of Flow Control significantly affect en-route traffic capacity?

There are two ways in which Flow Control might help in a busy region of airspace:

1. It can limit the maximum number of aircraft to be handled at any one time, and thus prevent any part of the system from becoming overloaded.

2. It can smooth out the bunches and gaps in traffic streams and so reduce the probability of conflict.

Preventing overload is important for the safety of the system, and for reducing the occasional long delays associated with recovering from overload, but it has little effect on capacity in non-overload situations. If traffic capacity were a 'hard' limit beyond which the system could not function, then by limiting flow rate peaks, average flow rates could be increased. However, as explained in Section 2, capacity is an 'elastic' limit which can safely be exceeded for short periods. As a result, the limiting of peaks has a much smaller effect on the mean flow rates which can be handled safely.

Reducing conflict probabilities will tend to increase capacity. However the only type of conflict whose occurrence is significantly reduced by removing bunching is that where an aircraft conflicts with another in the same traffic stream – overtaking conflicts, or conflicts in the main climb or descent phase. These are only a proportion of the total set of potential conflicts causing controller workload, but in some sectors they could be an important proportion.

The relationship between smoothness of flow and system capacity is well-worth exploring in a computer simulation study, but we should view Flow Control more as a means of ensuring the safety of the ATC system than as a means of significantly increasing capacity.

8. A Computer Simulation Study

At a number of points in this paper it has been noted that very little information exists about the precise effects of particular system-structuring concepts on traffic capacity, and that the only really satisfactory way to obtain such information is by computer simulation. Accordingly, the ATSTRATS project team is now engaged in building a simulation known as ASIM which is capable of modelling a wide range of ATC en-route structuring concepts. This will enable the quantification of the capacity implications of the various structuring options, and it is expected that the activity of building and experimenting with the simulation will add greatly to the team's insight into the possibilities for future systems. The simulation technology aspects of ASIM are discussed by Magill.(7) The initial list of topics for study includes the following:

- The relationship between the ATC system's uncertainty about future aircraft positions (that is, how good the trajectory prediction is), and traffic capacity.

- The capacity implications of the various route and airspace structuring concepts discussed in Sections 4 and 5.
- The capacity benefits of en-route Time Control.
- The effect on capacity of the smoothness or randomness of traffic flows.

ATC simulations may be either of the real-time man-in-the-loop kind, or of the non-real-time kind where human actions are modelled by the simulation program. ASIM is of the latter kind because of cost and flexibility considerations. However, as a consequence, any methods or concepts recommended by ASIM results will have to be validated and refined by air traffic controllers working in a real-time simulation environment before being considered for operational service; ASIM should be seen as a 'first filter' for new ATC ideas.

It would be possible to base such a simulation on a purely abstract volume of airspace with abstract traffic sources and sinks. Instead, ASIM is based on a volume of airspace which corresponds approximately to the London FIR, and models accurately the positions of airports in and near the UK, (although the program is written in such a way that a different geographical region could be modelled if required). Details of terminal area operations and traffic below 8000 feet are excluded. The simulation will generate pseudo-random traffic demand, but the mean flow rates and aircraft type mixes between the various airports involved will be based on the best available forecasts for the year 2000. Since the primary capacity-limiting factor is seen as controller workload, it is important for the simulation program to model this carefully. It is extremely difficult to model workload satisfactorily over a wide range of possible ATC system scenarios. For this reason the workload estimates produced will be fairly crude, but they should be sufficiently precise to compare the main system structuring options. The workload modelling method being used is based on the DORATASK method described by Richmond.[9]

9. Conclusions

There can be little doubt that the peak-period air traffic demand in the busiest parts of European airspace is going to grow dramatically over the next 10 to 20 years, and that research is needed to find new ways to cope with the resulting traffic levels. Although controller workload is seen as the chief capacity-limiting factor in today's systems, a completely automatic ATC system is not an option for the foreseeable future. It is more helpful to seek ATC system innovations which will reduce the amount of controller workload per aircraft handled, or will partition the total task in such a way that more controllers can be applied in parallel to it. New technology – computing, avionics, communications – makes possible some new ways of organizing and controlling air traffic, and this paper has considered four of them. From the qualitative and rather intuitive arguments presented, it would appear that new route structuring and airspace sectorization concepts hold out the greatest hope of capacity increase; time control offers some prospect of capacity increase in specific circumstances, and flow control is best seen as a means of preventing system overload. At the present time very little is known about the capacity implications of these new possibilities, and the only satisfactory way of rectifying this situation is by computer simulation. The construction of such a simulation is under way, and we look forward to its results.

References

1. Civil Aviation Authority, 'NATS Operational Strategy Plan', Civil Aviation Authority, London, Sept 88.
2. International Civil Aviation Organization, 'Description of the future Air Traffic Management System in the European region', European Office of ICAO, June 1989.
3. R D Hunter and P Brooker, 'The challenge facing Air Traffic Management in 2000 ad', ATC-2000 Seminar, Eurocontrol Institute, Luxembourg, Feb 1988.
4. S A N Magill, A C F Tyler, E T Wilkinson and W Finch, 'Computer-based tools for assisting air traffic controllers with arrivals flow management', RSRE Report 88001, Royal Signals and Radar Establishment, Malvern, Worcs, UK, Sept 1988.
5. U Völckers, 'Computer assisted arrival sequencing and scheduling with the COMPAS system', AGARD Guidance and Control Conference, Brussels, June 1986.
6. A Benoît and S Swierstra, 'Air traffic control in a Zone Of Convergence: Assessment within Belgian airspace', ICAS84 Conference, Toulouse, Sept 1984.
7. S A N Magill, 'Simulation of options for future Air Traffic Systems', UKSC90 Conference on Computer Simulation, Brighton, UK, Sept 1990.
8. I G Parker, 'A review of some sector capacity estimation techniques', DR Report 8703, Civil Aviation Authority, London, August 1987.
9. G C Richmond, 'An interim description of the DORATASK methodology for the assessment of sector capacity', DORA Report 8818, Civil Aviation Authority, London, Dec 1988.
10. B Miaillier and P Planchon, 'The role of the controller in the ATC loop', ATC-2000 Seminar, Eurocontrol Institute, Luxembourg, Feb 1988.
11. J-L Renteux, 'Aircraft modelling standards for future ATC systems', Eurocontrol doc 872003, Eurocontrol, Brussels, 1987.
12. J-L Renteux, 'BASE OF AIRCRAFT DATA Users Manual', Eurocontrol doc 892009, Eurocontrol, Brussels, 1989.
13. S A N Magill, 'Use of 4-D FMS in a complex terminal area in the pre-data-link era', COMPAS 87 conference on 'computer applications in aircraft design and operation', Paris, June 1987, Proceedings published by Springer Verlag, 1987.