Temporal Database Technology for Air Traffic Flow Management

S.M.Sripada[1], B.L.Rosser[2], J.M.Bedford[2] and R.A.Kowalski[3]

[1] European Computer-Industry Research Centre
Arabellastrasse 17, 81925 Munich, Germany. sripada@ecrc.de
[2] Ferranti International plc, Simulation and Training
Ty Coch Way, Cwmbran NP44 7XX, U.K.
[3] Imperial College of Science, Technology & Medicine
180 Queen's Gate, London SW7 2BZ, U.K. rak@doc.ic.ac.uk

Abstract. The function of air traffic flow management (ATFM) is to ensure that air traffic operates within adequate margins of safety. Existing ATFM systems are manual which are over-conservative in operation resulting in under-utilisation of available airspace. As well as being costly, such systems are unable to cope with increased demand for air travel in regions such as Europe. Attempts are currently being made to provide computer-based decision support for ATFM. Computerised decision support for ATFM ensures that safety margins are maintained while at the same time increasing the effective capacity of the airspace by more efficient flight scheduling. At the heart of such a system is active temporal database technology which aids the air traffic controllers by keeping track of airspace occupancy (a time-map of spatio-temporal trajectories of aircraft) in controlled regions of airspace, enabling flow managers to process requests for new slots for takeoff and to smoothen and optimise the flow of air traffic. The technology also aids air traffic controllers by alerting them to possible conflicts and by providing tools for re-routing aircraft to avoid mid-air collisions. The paper describes a large scale demonstrator for ATFM that has been developed at Ferranti Simulation and Training.

1 Introduction

In recent years, civil aviation has seen a rapid increase in demand calling for increased air-space utilisation. Existing systems for air traffic flow management (ATFM) are manual and are unable to cope with this increase in demand without compromising safety. Efforts are therefore underway to provide computer-based decision support for ATFM. The function of air traffic flow management is to ensure that air traffic operates within adequate margins of safety. This paper describes a large scale demonstrator for strategic and tactical flow management that has been developed at Ferranti Simulation and Training.

At the heart of the ATFM system is active temporal database technology which aids flow management by supporting the maintenance of a model of airspace occupancy for controlled regions of air space. The ATFM application provides a perfect case for the use of active temporal database technology. Some of the functionalities required by the application which highlight the need for active temporal database technology are described below.

Temporal Databases Flow Managers require a model of airspace occupancy in order to control air traffic. The model comprises a time-map of spatio-temporal trajectories of the various air-borne aircraft in controlled regions of airspace. This requires the ability to store and manipulate spatial and temporal data. In the ATFM application, temporal data is quite complex and requires special database functionalities to support temporal knowledge representation and reasoning.

Real-time updates Space-time coordinates of an aircraft along its route are periodically relayed back to Air Traffic Control units. Based on this input, the airspace occupancy table needs to be continuously and quickly revised to reflect the latest state of the airspace. Information from the airspace occupancy table is then used to allocate slots for new flights as well as to detect and avoid mid-air conflicts.

Integrity constraints Upon every update, the database has to ensure that capacity of a sector of airspace is not exceeded at any time; it should also enforce that a minimum safe separation distance is always maintained between all airborne aircraft.

Triggers Upon every revision, the database has to detect any violations of integrity constraints and automatically initiate actions that alert the flow manager and support remedial actions such as re-routing.

Past records A record of all the past data are to be kept so that analyses may be carried out for post-incident investigations, improving scheduling strategies, optimising the flow of traffic, and discover patterns in air traffic and airspace utilisation for strategic planning purposes. Thus the application requires bi-temporal databases handling both valid time and transaction time [3]. Valid time denotes the time at which information models reality (e.g. an aircraft being at a particular point in space at a certain time) whereas transaction time denotes the time for which the information is believed (or recorded) in the database (e.g. at the time of take-off, the flight is predicted to be at a certain point in space at a certain time; at a later time, this belief may be revised).

The rest of the paper is organised as follows: Section 2 describes the domain of Air Traffic Control (ATC) and specifically the problem of Air Traffic Flow Management (ATFM). The requirements of a computer-based decision support system for ATFM are also described in Section 2. A description of the temporal database technology used in the application and the domain modelling aspects are presented in Section 3. Implementation and performance of the ATFM system are sketched in Section 4. In view of the proprietary nature of the technology, only an outline of the implementation has been included in sections 3 and 4. Plans for commercial exploitation as well as ongoing work and conclusions are presented in Section 5.

2 The ATFM Application

2.1 The ATC Domain

In the early days of commercial air transportation, the skies and runways were relatively empty. Pilots were free to take off from an airport, choose their route to their destination, and land there as soon as they could see that the runaway was free. As both the volume of air traffic and the speed of aircraft increased, the risk of unfettered air travel became too great. First, it became necessary to control the use of runways to avoid problems at take-off and landing. Next, the busy zones in the vicinity of airports needed to be controlled to prevent accidents during approach and ascent. Finally, controls were extended to the enroute section of a flight, i.e. when the aircraft is flying at its cruising altitude and speed towards its destination. In this way, the industry of Air Traffic Control (ATC) evolved.

The means by which ATC functions to control air traffic depend upon

- an internationally agreed set of rules and regulations,
- an infrastructure of equipment, and
- human resources to put the system into effect.

The rules and regulations define safe flying practice. They constrain pilots to follow well-defined routes at definite altitudes (known as flight levels), to inform the authorities of these in advance, and to fly at defined minimum distances one from another. The infrastructure includes ground-based radar for tracking aircraft in flight, beacons which emit direction-finding and position-identifying signals to aircraft, and on-board equipment such as radio, transponders and an increasing array of electronic safety devices. The resources are the Air Traffic Controllers and associated staff. Their main function is to help pilots to follow the rules and regulations by monitoring flights and directing pilots to avoid conflicts (i.e. violations of the standards for safe separation).

The requirements of Air Traffic Control dictate that the controlled airspace is hierarchically divided into distinct three dimensional areas ranging from Flight Information Regions (FIRs) where specific procedures and standards apply to sectors which are the units of operational control.

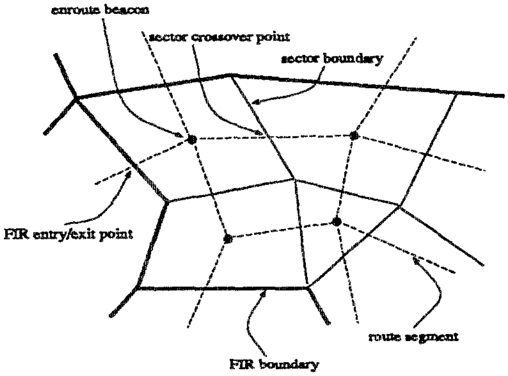

Fig.1. A highly schematic view of a part of an airspace

A highly schematic and simplified plan view of an airspace showing part of an FIRs and its constituent sectors is shown in Fig.1. The dotted lines are the routes that the aircraft actually travel along. They pass over waypoints, shown as heavy dots in the figure (enroute beacons are situated at ground level). Lines joining the waypoints form a grid of imaginary air corridors along which the aircraft travel. An aircraft flying along a route is always under the control of the Controller for the sector it currently occupies. When it passes from one sector to the adjacent one, it is explicitly "handed over" from the first Controller to his counterpart in the next sector.

2.2 The Role of Air Traffic Flow Management in ATC

The over-riding goal of Air Traffic Control (ATC) is maintaining safety in air transportation. It achieves this by a complex of control, management and planning procedures organised into several layers. The details of this organisation vary from country to country according to local conditions. However, a common factor is the fundamental division of the task into Strategic Planning, Radar Control, and Flow Management.

Strategic Planning is carried out often several years in advance. Its object is to ensure that the ATC infrastructure is adequate to meet expected demand. This may involve re-organising the route network, improving equipment levels, or training more Air Traffic Controllers. The Radar Control and Flow Management levels are both concerned with the operation of safety procedures within this infrastructure.

The Radar Controller bears the front line responsibility for maintaining safe separation between aircraft in flight. He does this by monitoring the positions of all aircraft in his sector (i.e. the volume of controlled airspace he is responsible for) on a radar screen, and by communicating instructions to pilots by radio when necessary.

A skilled Radar Controller is well able to perform this function as long as his workload remains within his capacity. If there is too much activity within his sector at a given time, he may become overloaded and risk making an error. Examples of potential errors are failing to observe an impending "conflict" (i.e. a violation of the separation standard applying to the sector), mis-directing a pilot, or failing to effect a smooth hand-over of control when an aircraft leaves his sector.

One measure employed to reduce the risk of such an overload is to predict the Radar Controller's workload for a short period ahead of real time (i.e. up to about 30 minutes), and to schedule his operations so that no two actions become urgent at the same time. (Often, this function is performed by a separate member of the Radar Control team called the Planning Controller.) However, this presupposes that the workload profile averaged over time remains below the limit set by his capacity. In the busier regions of the air transport network, where traffic loads vary widely over time and space, this requires Flow Management.

The need for Air Traffic Flow Management (ATFM) arose to remove all risk that any sector of (controlled) airspace could become overloaded for an extended period, with the consequent threat to safety. ATFM covers a range of activities which can be characterised as strategic, pre-tactical and tactical - according to the time-scales over which they operate. For all of these activities, it is necessary to carry out a prediction of traffic flows.

At the strategic level, these predictions are used to determine appropriate staffing levels for particular sectors and to set restrictions when necessary. Underloaded sectors can often be grouped together under one Radar Controller, thus freeing resources, while more heavily loaded sectors are assigned to the most experienced Radar Control teams in order to maximise their capacity.

When an overload is predicted - i.e. the predicted traffic flow will exceed the maximum safe capacity of the sector - the Flow Manager must plan measures to reduce the load. These measures are designed to re-distribute the traffic in time, in space, or both. When the overload is of relatively short duration (e.g. one hour or so at the peak of the daily traffic cycle), some of the peak traffic can be delayed, thus smoothing out the cycle. In other cases, it may be more effective to divert traffic from an overloaded sector through nearby sectors which are not overloaded.

The restriction is the strategic Flow Manager's main instrument to effect such re-distributions. A basic restriction is applied to a sector and specifies the maximum number of aircraft permitted to enter the sector over a specified period. Usually, restrictions are refined to take account of the distribution of flows through a sector. In this case, separate restrictions may be placed on the individual entry points to the overloaded sector.

A Flow Management policy (including restrictions and guidelines for re-routing or delaying flights) is initially set well in advance of real time (usually at the start of a busy season). This is reviewed and tailored, by adjusting the restrictions if necessary, one day before the day of operation. The restrictions, together with general principles of day-to-day flow management, are then enforced at the tactical level by the procedure of Slot Allocation. This takes place in a time window extending from about thirty minutes to about four hours before real time, when airlines and pilots log their flight plans and request permission to traverse controlled airspace. The Flow Manager will either approve the flight plan as requested, or will negotiate time or route changes in accordance with the restriction plan. It is at this point that air travellers all too frequently become aware of the Flow Management process, when the pilot informs them that the delay in their departure is due to "Air Traffic Control".

2.3 Requirements for Computer-Based Decision Support for ATFM

The ideal decision support system for ATFM must be both integrated and distributed. It must integrate all levels of the ATFM process - from strategic to tactical. It must be accessible to users and information suppliers throughout its geographical range. It must be responsive for interactive use, and reactive to frequent updates and data revisions. It must cater for the special requirements of each user in each subregion, while at the same time integrating traffic predictions and flow management measures over the whole range. It must deal intelligently with the large volumes of data at its disposal, taking account of the reliability, precision, and temporal characteristics of the data.

Above all, the system must be reliable in use. In increasing the effectiveness and efficiency of ATFM, it must not introduce additional hazards which could affect safety. To this end, it must begin by automating the lower-level aspects of the flow management process using well-established technology, leaving the Flow Manager in full control of the task. Nevertheless, it must also have built-in flexibility and capacity for progressive task automation, so that it may be developed and extended incrementally as technology allows. Without this, the system could quickly become out-dated and impede further progress.

2.4 Characteristics of ATFM Data

An effective ATFM decision support system will inevitably centre on the prediction of air traffic in terms of its volume and its detailed spatio-temporal distribution. This depends upon the availability of suitable data, and also upon appropriate methods of handling this data. The data available to an ATFM system can be described in terms of three basic characteristics, namely: its level of abstraction, its accuracy and precision, its lead time, and its source.

Air traffic data presented to ATFM is at three main levels of abstraction:

- individual flights: flight plans describing routes and associated timing information,
- traffic flow rates: the number of aircraft crossing a point (e.g. a beacon) per unit time as a function of time, and
- traffic load rates: the number of aircraft in an airspace entity (e.g. a sector or a route segment) as a function of time.

These are genuine levels of abstraction since information is progressively lost at each level. (Thus, reasoning from individual flights to either flow or load rates is straightforward, but the inverse is not.) This temporal data makes reference to a structural description of the airspace, i.e. the route network, its sectorisation, and the various kinds of "fixes" - instrumented beacons, reporting points, etc.

Accuracy and precision are generally inversely related when applied to predictive information. For example, the prediction that a flight will take off at 10.19am exactly will prove inaccurate if the flight actually takes off at 10.21am, whereas "about 10.20am" would have been true. Of course, neither prediction would be accurate if the flight was cancelled. In fact, ATFM data is generally presented at a constant high precision, even if this means that the accuracy is low. For example, every flight plan includes a take-off time expressed in hours and minutes whether this time is confirmed or is just a pious wish. The main determinants of accuracy are the data's lead time and its source.

In general, the earlier an item of ATFM information is available, the less accurate it is likely to be. The range of lead times is wide - from about one year ahead of real time down to minutes or hours after real time. In the context of current ATFM practice, however, information received later than about 30 minutes before real time tends to be useless for flow management purposes because it is too late to act on it.

Sources of Information There are five main types of information source used in ATFM:

- Intentions: flights intending to use controlled airspace are required to log their flight plans well before scheduled take-off.
- Historical data: seasonal, weekly and daily patterns in the level of utilisation of airspace due to holiday and commuter traffic.
- Trends: longer term trends such as the growing demands on controlled airspace.
- Special factors: factors such as major sporting events, state visits which produce localised effects.

- Actual events: Ultimately, all predictions are confirmed, refined or corrected by what happens in the real world. Many real time events are relevant to tactical ATFM because they have enduring consequences which fall within the tactical ATFM time window (i.e. 4 hours to 0.5 hours in advance). Examples are take-offs, position updates, temporary runway or route segment closures, weather fronts.

2.5 Prediction and Data Manipulation

An essential requirement for the Decision Support System posed by tactical ATFM is the ability to maintain a consistent picture of the traffic over a moving time window (of four hours) in the face of a constant stream of updates and revisions. As well as providing tools for dynamically managing the traffic at the level of individual flights, it must actively warn of any impending overloads. It must also be able to recover the exact context in which a Flow Management decision was made for the purpose of Post-Incident Analysis.

An important attribute of the Decision Support System for all levels of ATFM is the ability to support hypothetical reasoning. The Flow Manager will frequently wish to evaluate a number of alternative solutions to a problem before committing to a single solution. For example, a Strategic Flow Manager may wish to compare different restrictions and choose the one whose effects are closest to the ones desired. Similarly, a Tactical Flow Manager may wish to allocate a take-off slot tentatively, and move it forwards or backwards in time to clear any restriction violations that it might otherwise induce. The integrity of the database must not be affected by such tentative operations especially in view of the fact that it will be accessed by several users simultaneously - and the alternatives tried should not interfere with each other.

3 Temporal Database Technology for ATFM

The temporal database technology that has been developed for the ATFM application is based on the research solutions developed by Sripada in [7]. In particular, Sripada proposes a declarative representation of temporal knowledge in a database and the pre-computation of appropriate derived relations (materialised views) so as to improve query-time performance. To improve performance at update-time, an efficient method for the maintenance of these derived views has been proposed [8]. The pre-computed relations are then stored in an extended relational database for efficient access.

The architecture of the ATFM temporal database system is shown in Fig.2. The ATFM system has been developed around a Temporal Database kernel: a proprietary database system for real-time applications (SoftRP of Ferranti International plc, denoted by DMS in Fig.2.) has been extended and optimised to handle temporal data with suitable indexing techniques for efficient storage and retrieval of temporal data [2]. A number of interfaces are supported by the temporal relational database including TSQL, a temporal extension of SQL (as illustrated in Fig.2). A Deductive Temporal Database is then developed by tightly coupling Quintus Prolog and the Temporal Relational Database kernel with various optimisations for temporal database access. Data modelling of the application domain has been done in the deductive framework of GRF (General Representation Formalism) based on an extended version of the Event Calculus [4], a logic based framework for representing and reasoning with time and change. Mechanisms for materialisation, triggers and efficient belief revision have also been integrated into the tightly coupled system. Various decision support modules have been developed in Prolog which are served by the deductive temporal database system.

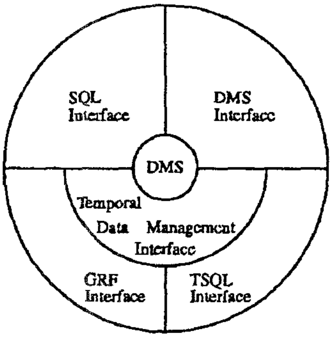

Fig.2. Architecture of the Deductive Temporal Database Management System. The DMS is the relational database kernel with SQL and TSQL interfaces. The DMS and GRF interfaces are deductive interfaces for the Prolog language.

3.1 Modelling Based on the Event Calculus

The deductive level of a DTDBMS calls for a suitable temporal knowledge representation formalism, for which there are many candidates (e.g. see [5]). The requirement here is for a temporal logic which can meet the requirements of "temporal projection" for industrial domains, and which can be provided with an efficient computational basis (temporal projection is model-based temporal reasoning which often equates to prediction in practice, but can equally well work backwards in time). The one chosen for the ATFM application was the Event Calculus because it enables the representation of events and processes (which abound in the ATFM application) plus its ability to deal with default persistence and continuous change and facilitate tasks such as abductive planning. Although the Event Calculus presents some computational problems for temporal reasoning in general, it proved possible to develop a very efficient implementation for temporal projection. Event Calculus was extended in the project with more powerful facilities (GRF) for modelling process-based domains such as air traffic.

The basic element of the ATFM domain is the *flight* process. As a physical process, flight is a continuous function that maps time to the position of the aircraft in space (altitude, latitude and longitude). This is modelled in the extended Event Calculus at a qualitative level, i.e. as a sequence of events (instantaneous changes in the qualitative state of the flight) and properties (persistent states of the flight). Typical events are *at_waypoint(flight,waypoint)* and *enter_sector(flight,sector)*. Typical properties are *enroute_between(flight,waypoint1,waypoint2)* and *in_sector(flight,sector)*.

The main model of the flight process model includes a unique identifier for the flight, an unambiguous representation of the flight path in terms of *take_off*, *at_waypoint* and *landing* events, and values for cruising altitude and cruising speed. From this, all other information concerning the flight can be recovered, sometimes with the aid of default information. For example, the *enter_sector* events can be found by solving for the intersections between the flight path and the sector boundaries, while the altitude at a point on the climb to cruising altitude can be found using a default climb profile.

Air traffic rates are deduced from the total set of known flights. If all of the *in_sector(flight,sector)* properties have been explicitly deduced, the sector load at any instant t can be easily recovered by the Event Calculus query *holds_at(in_sector(_,sector_name),⟨t⟩)* and counting the set of properties returned. To plot the sector load as a function of time, a simple mechanism can be defined which initialises the function with the above query, obtains all *enter_sector* and *exit_sector* events over the required time range (in time order), and increments or decrements the function at each such event.

Calculating flow rates at beacons or other waypoints can be done in a similar way. However, since flow is integrated over time, it is necessary to define a function for this integration. The simplest form of this is a time window. Thus, if the flow is to be expressed as number of aircraft per 10 minutes at time T, the time window will extend from T-5 minutes to T+5 minutes. For each event

happens(at_waypoint(flight,waypoint),⟨t⟩),

the events

happens(enter_time_window(flight,waypoint),⟨t-5⟩) and
happens(exit_time_window(flight,waypoint),⟨t+5⟩)

are generated. These are then used exactly as the *enter_sector* and *exit_sector* events were to compute the required function.

Once generated in this way from the basic flight information, the rate functions can be themselves manipulated by the Flow Manager. For example, if the Flow Manager restricts the flow across a particular waypoint to 50% of its previous value, then 50% of the flights contributing to the value of the function at that point must be delayed to the end of the restriction period. Because of the heterogeneity (in respect of origin, destination and route) of the set of contributing flights, the system must select a representative subset to delay. This involves a degree of intelligent reasoning on the part of the system. In order to fine-tune the restriction, the Flow Manager must then be given the tools to modify the system's selection, i.e. by setting other restrictions or by explicitly applying the restriction to flights with specified destinations or origins. This is where the essential skill of strategic flow management lies. The problem is not yet sufficiently understood to be able to automate it, but the use of the ATFM system in manual mode will contribute to this understanding and permit automation in the longer term.

Integrity Enforcement through Triggers A flow management system must be able to maintain a consistent model of airspace usage in a changing world. To achieve this, the TDBMS must support consistency maintenance using integrity constraints. In the ATFM domain, the temporal database systems is required to enforce the following principal integrity constraints:

1. that the projected number of aircraft in a sector does not exceed the capacity of the sector at any point in time
2. that no two aircraft do not arrive at the same waypoint within a specified time window

Contrary to normal integrity constraints such as "an aircraft cannot be at two places in space the same time" which can be satisfied by deleting old conflicting information or rejecting the update, the above integrity constraints need to be maintained by the database system actively. When any revision to the database occurs that violates the above integrity constraints, the database should initiate a set of actions such as alerting the air traffic controller to the impending danger, invoking conflict resolution modules to present re-routing options to the controller.

The integrity enforcement mechanism is based on active database technology and is invoked by triggers from event-condition-action rules [1].

Real-time updates The projected trajectories of the aircraft need to be constantly revised with updates to a flights route or take-off time. The implications of the new time-space co-ordinates of the aircraft are then used to project the new (revised) trajectory of the flight, and integrity maintenance has to be carried out. All of this activity requires real-time belief revision capabilities. The belief revision not only involves deletion of old, incorrect data, but also the computation of new implicit data. Real-time updates are achieved through a novel technique for efficient belief revision based on transaction time stamping developed by Sripada [8].

4 Implementation

4.1 The Architecture of the ATFM system

The ATFM system that has been developed is called TIMS-ATM (Temporal Information Management System - Air Traffic Management). The TIMS-ATM architecture contains two distinct databases (see Fig.3). A "static" database stores the ATC infrastructure model. This needs to be optimised for fast retrieval, but updates are generally done off-line and infrequently. The other database stores the dynamically changing traffic predictions, and needs to be optimised for both retrieval and updates. The deductive level is provided by

1. a rule base and a data-driven inference engine for update time deduction, and
2. a goal-driven query processor for deduction at query time.

The nature of the application as essentially a monitoring system places a premium on update-time deduction. This means that as much information as possible is explicitly

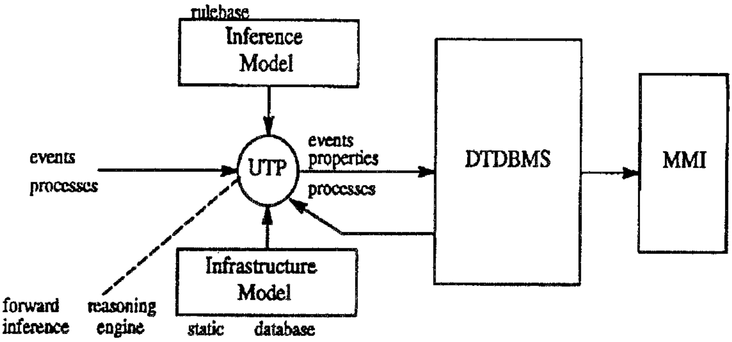

Fig.3. The TIMS-ATM system architecture

deduced from input messages by the Update Transaction Processor (UTP) and is stored in the temporal database as a time-map. This is then ready for access by the Flow Manager (or other reasoning agents) with minimal latency.

In the TIMS-ATM, a DTDBMS (Deductive Temporal Database Management System) forms the basis of a subsystem for inferring and storing temporal information (in this case, air traffic predictions). This subsystem acts as a server to a set of reasoning agents, which perform the ATFM tasks. These agents are initially human Flow Managers, interacting via specially designed MMI visualisation and input tools. However, the architecture supports the progressive automation of ATFM tasks by allowing additional task modules to be easily integrated. A few simple examples of these are demonstrated in the TIMS-ATM.

4.2 MMIs for Strategic and Tactical Flow Management

The TIMS-ATM demonstrator for Air Traffic Flow Management comprises a set of tools that aid the controller in performing both Strategic and Tactical Flow Management. These tools allow the controller to experiment with possible flow management actions before committing to a preferred decision. The flow management tools were developed to allow the flow manager to visualise and modify the information contained within the underlying temporal database via a Motif/X Window graphical user interface.

Strategic and Pre-Tactical Flow Management Given a projected airspace occupancy model for some time in the future, the TIMS-ATM tools for strategic and pre-tactical flow management allow the flow manager to visualise the impact of flow on air traffic controller's workload and to modify flows to alleviate this.

The two main approaches for displaying information contained within the temporal database are:

- plots of waypoint flow and sector load against time, and
- plots of Radar Controller's workload against time.

Radar Controller's workload is an additional level of abstraction which is implemented experimentally in the TIMS-ATM. It is calculated by summing the elements of workload generated by the several aircraft in the sector. As well as giving rise to a general workload increment for as long as it is in the sector, an aircraft generates additional workload according to its specific behaviour within the sector - e.g. as it is entering, leaving, or reaching a reporting point, if it is changing altitude, or if there is a prospect of conflict with another aircraft. The precise specification of these workload increments will require expert knowledge.

An additional display feature is the colour-coded analysis of a sector load by route. That is, the routes passing through a sector are coloured to indicate the volume of traffic using them. Saturated colours indicate heavy flows, while unsaturated colours indicate light flows. Numerical annotations are also used. This display gives the flow manager an immediate visual impression of the traffic pattern, which can help him identify the critical bottlenecks in the sector.

By clicking on the waypoints and selecting from a menu of operations, the Flow Manager can invoke a restriction setting tool. This enables him to alter the flow at the selected waypoint by a variable percentage. When he does this, a percentage of the traffic is automatically delayed. The tool also shows him the loads on all other sectors which might be affected by his action, so that he can check any knock-on effects produced by the delay.

Tactical Flow Management The Tactical Flow Management Tools provide support for the flow manager in the activity of slot allocation. The flow manager can select a flight from a list of flights awaiting slot allocation and a consistency check of the plan associated with the flight is performed. If the plan is valid then the flight is hypothetically projected into the airspace occupancy model of the temporal database. If no conflicts, congestion or restriction violations are detected by the system, the flow manager can commit to the slot allocation. Otherwise, the flow manager is provided with a set of tools to repair the conflict interactively. These tools include:

- a colour annotated display of the route on the geographical display, annotated with times at key points enroute,
- waypoint flow and sector load plots for the relevant parts of the route,
- a time map display showing all of the flights passing a selected waypoint, together with their altitudes,

The flow manager is able to reschedule a slot manually by using either mouse or keyboard interaction. An experimental automatic slot allocator is also provided. This works by moving the flight forward in time until a clear flight path is found (i.e. one free of conflicts, congestion and restriction violations).

4.3 Performance and Optimisation

The TDBMS has been developed for efficiency in updating and accessing temporal data. It has been implemented using a proprietary real time data management system that provides constant time access, query optimisation and high update performance. The constant time access means that the time to search for a record using a single unique index is independent of the number of entries in the database. The temporal database has been further optimised for manipulating temporal attributes by transparently incorporating a multi-index intersection operation and an optimised join.

The benefit of using this database technology for ATFM is that it is possible to project an average of 2000 flights per hour into the airspace usage model that covers a large section of European airspace.

To illustrate the performance of the database on multiple time attributes, Fig.4. gives the result of a set of range queries on a database containing 16758 records. Time is typically represented as an integer in the TDBMS and the data contains two integer time attributes labelled X and Y. The results of the benchmark show the superior and predictable performance of the intersection operation independent of the data distribution.

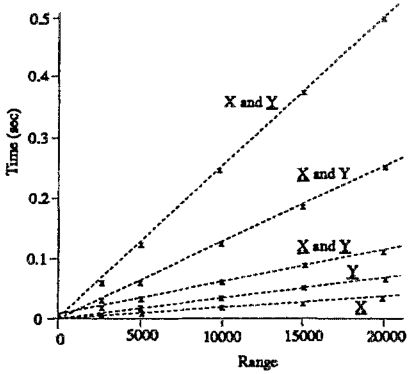

Range	Records returned		
	X	Y	X and Y
2500	424	852	23
5000	749	1696	52
10000	1670	3679	227
15000	2652	5700	972
20000	3662	7331	1827

(a) Time taken

(b) Records retrieved

Fig.4. Range Select benchmark on one or two integer attributes. The select operations were performed on the X, Y and X and Y attributes. An underscore indicates that an index was used for that attribute. Timings were taken from a Sun 4/330(17.8 mips).

5 Conclusions

We have described a computerised decision support system for Air Traffic Flow Management developed at Ferranti International plc (UK). The potential benefits of such a system is higher capacity utilisation without a decrease in air safety, resulting in lower operational costs and lower air travel costs. One of the crucial enabling technologies in the development of the demonstrator was active temporal database technology.

The demonstrator has been successfully tested in simulated real-time with real flight information from the UK airspace. Negotiations are currently in progress for the commercial deployment of the ATFM system in Europe. The temporal database technology is to be further extended and exploited in another project involving several members of European aircraft manufacturers and software vendors to develop a commercial product for pilot procedures training.

Application of database technology in advanced applications such as ATFM also has potential research gains. Firstly, it validates the research results such as those in [7] as providing practicable solutions to real world problems. Secondly, it provides necessary application oriented feedback towards establishing appropriate database standards (e.g. [6]). Lastly, it stimulates further research in the area towards the development of advanced temporal database technology for more sophisticated applications such as Intelligent Vehicle Highway Systems [9].

Acknowledgements

The work of the first author was supported by a Doctoral Fellowship from the Jawaharlal Nehru Memorial Trust(UK). The development of the ATFM system at Ferranti International plc was supported by the ESPRIT project EQUATOR. Many thanks are due to Todd Ferguson for his invaluable help in integrating figures into the camera-ready version of this paper.

References

1. Bayer, P.: State-Of-The-Art Report on Reactive Processing in Databases and Artificial Intelligence. The Knowledge Engineering Review, Vol.8, No.2, 1993
2. Bedford, J.M.: EQUATOR Temporal Database Management System - User Guide and Report. Technical Report AMS&T/EQ/REP/33, Ferranti International plc, May 1992
3. Jensen, C.S. et al.: A Concensus Glossary of Temporal Database Concepts. ACM SIGMOD Record, Vol.23, No.1, March 1994
4. Kowalski,R.A., Sergot,M.J.: A logic-based calculus of events. In New Generation Computing. 1986
5. Maiocchi, R., Pernici, B., Barbic, F.: Automatic Deduction of Temporal Information. ACM Transactions on Database Systems. Vol.17, No.4, December 1992
6. Snodgrass, R.T. et al.: TSQL2 Language Specification. ACM SIGMOD Record, Vol.23, No.1, March 1994
7. Sripada, S.M.: Temporal Reasoning in Deductive Databases. Doctoral Thesis. Department of Computing, Imperial College of Science, Technology & Medicine, University of London, January 1991
8. Sripada, S.M.: A Temporal Approach to Belief Revision in Knowledge Bases. In Proc. 9th IEEE Conference on Artificial Intelligence for Applications, CAIA'93, Orlando, Florida, March 1993
9. Sripada, S.M.: The Design of the ChronoBase Temporal Deductive Database System. In Proc. International Workshop on Infrastructure for Temporal Databases (NSF Workshop), Arlington, Texas, June 1993