Workload Implications of Free Flight Concepts

John W. Andrews

Jerry D. Welch

ABSTRACT

The removal of constraints upon traffic flow in order to allow more efficient user−preferred routing is a major thrust of current ATM research. Many of the constraints that currently exist arise from the requirement that workload be matched to the capability of a human controller who is responsible for a fixed volume (sector) of airspace. This paper evaluates the average workload, workload peaks, and rates of multi−aircraft conflicts that result from relaxing flow and routing restrictions. The increased variability in sector workload that results from unstructured traffic flow can have a negative effect upon sector productivity even when average traffic levels are unchanged. As traffic density increases, the conflict workload experienced by individual aircraft grows more slowly than the conflict rate experienced by sector controllers. A practical implication is that the accommodation of increasing traffic densities without increased flow restrictions will eventually require the introduction of innovative traffic management architectures in which the conflict resolution workload currently assigned entirely to controllers in static sectors will be distributed among automation, aircrews, and controller teams.

The authors are members of the Air Traffic Automation Group at MIT. Lincoln Laboratory, 244 Wood Street, Lexington, MA 02173−9108, USA. This work was sponsored by the Federal Aviation Administration, Department of Transportation, under Air Force Contract No. F19628−95−C−0002.

Notation

α	Probability that the sector loading will exceed N^{max}.
ε	Sector efficiency (average loading allowed as a fraction of maximum tolerable loading, N^{max}]
G	Sector workload intensity (Schmidt model)
κ	Three−dimensional (volumetric) density of aircraft

λ	Arrival rate of customers at a queue
λ^{AC}	Conflict rate experienced by an individual aircraft (conflicts/hr)
λ^{C}	Conflict rate experienced by a sector controller(conflicts/hr)
λ^{S}	Arrival rate of aircraft entering a sector.(AC/hr)
M^{h}	Half−width of the alert volume for a single aircraft.
M^{v}	Half−height of the alert volume for a single aircraft
N^{max}	Maximum number of aircraft that can enter a sector without overload.
N^{S}	Number of aircraft in a sector
Q	Volume of sector (nmi^{3})
ρ	Traffic intensity in a queueing system.
T	Average sector transit time
τ^{C}	Amount of time sector controller spends to resolve a conflict (including devising solution, communicating solution, monitoring, and recovery)
τ^{R}	Average time required by controller (after detection) to devise and communicate a conflict resolution
τ^{S}	Amount of time sector controller spends on routine tasks for each aircraft that enters the sector.
V^{i}	Ground speed of aircraft i
V^{ij}	Relative speed between two selected aircraft

Workload Implications of Free Flight Concepts

1. Issues Arising from Free Flight

The goal of current free flight initiatives is "a safe and efficient flight operating capability under instrument flight rules (IFR) in which the operators have the freedom to select their path and speed in real time." [Ref. 1] Under this concept, aircraft would fly on user−preferred routes chosen for operating efficiency rather than compatibility with air traffic management (ATM) structures. The attendant relaxing of traffic flow structure has many effects upon the air traffic control process. Of particular interest are the workload and task changes experienced by sector controllers as traffic flow structure is relaxed. Understanding these impacts provides an indication of where the impediments to free flight will arise and where innovations will be needed to ensure the implementation of the free flight vision.

This paper will present a set of simple mathematical models that provide approximate answers to some fundamental ATM performance issues associated with free flight. The effects of increased randomness in traffic flow upon sector loading is analyzed. Conflict rates experienced by both aircrews and sector controllers are estimated and the resulting sector workload intensity is estimated using a simple workload model. The probability of multiple simultaneous conflicts is computed and shown to be a potential constraint upon traffic densities.

1.1 Flow Structure in a Sectorized ATM System

Sectorization is a fundamental architectural feature of the world's air traffic control systems. A sector is a volume of airspace within which one and only one sector controller has responsibility for separation of traffic[1]. As aircraft cross sector boundaries, a reliable and unambiguous hand−off procedure must be executed to transfer control and communication responsibilities.

A second major architectural feature of conventional air traffic control is standardized routing in which aircraft follow specified routes which require them to enter and exit sectors at defined points. Although controllers often approve non−standard routing, there is generally an understanding that when traffic loading makes it advisable, aircraft will be required to follow standard routes.

A variety of operational considerations are taken into account in designing sectors and specifying standard routes. For example, the shape of sectors is usually such that the predominant traffic flow is parallel to the longest dimension of the sector. Routes are usually specified so that they remain clear of boundaries and well within radar coverage, and avoid high−speed (head−on) encounters between opposing streams of traffic.

The resulting traffic flow structure, while helpful to the sector controller, results in a number of inefficiencies from the point of view of aircraft operators. One of the principal goals of the free flight concept is to remove these restrictions and thus allow more efficient user−preferred routing. However allowing unstructured traffic flow has potentially significant implications for sector controller workload. For instance:

- The probability of simultaneous hand−offs may increase due to aircraft flying on parallel paths.
- Aircraft may fly near the sector boundary and thus interact with traffic not under sector control.
- Sector penetration may occur more frequently through the floor or ceiling of the sector.
- The likelihood of multi−aircraft conflicts may increase.
- The likelihood of simultaneous conflicts may increase.
- The likelihood of conflicts occurring immediately after hand−off may increase.
- Conventional flow control tools, such as imposing miles−in−trail restrictions on major routes, may be unavailable.
- The generally more complex and varied nature of the flow may degrade situation awareness. In order to ensure a continue progress toward the ultimate free flight objective, these potential stresses upon the sector controller must be understood and an adequate set of compensating procedures and automation must be provided in a timely manner.

[1] In some cases, a team of two or more controllers assigned to a sector may work in a closely coordinated way to carry out the air traffic control functions. However, even in such cases one can usually identify a single "radar controller" who bears most of the mental workload associated with conflict resolution

2. Sector Loading Analysis

The term sector loading will refer to the number of aircraft, NS, that are simultaneously present within a given air traffic control sector. This section will examine the statistics of this quantity under structured and unstructured scenarios.

2.1 Sector Loading (Poisson Model)

When aircraft fly independent user−preferred routes, their locations with respect to sector boundaries will be effectively randomized[2]. If aircraft positions are independently distributed in the airspace of interest with volumetric density κ, then the probability that there will be NS aircraft in a sector of volume Q at a given time is given by the Poisson distribution. HenceThe mean and standard deviation of NS are

$$P\left[N_S = n\right] = \frac{(\kappa Q)^n}{n!} \quad e^{-\kappa Q} \qquad [2.1]$$

The mean and standard deviation of N_S are

$$E\left[N_S\right] = \kappa Q \text{ and } \quad \sigma_{N_S} = \sqrt{\kappa Q} \qquad [2.2]$$

2.2 Flow Rate Relationships

Sector loading can be related to the sector throughput (aircraft/hour entering the sector).
Suppose that there are several routes feeding the sector, each contributing an arrival rate λi. The total arrival rate for the sector is then

$$\lambda_S = \sum_i \lambda i \qquad [2.3]$$

If the traffic is well−distributed among several independent routes, then the contributions of each route to NS will be independent. If Ti is the average time that aircraft on route i are in the sector, then the expected number of aircraft in the sector due to route i is λiTi and the expected number of aircraft in the sector is

$$E\left[N_S\right] = \sum_i \lambda i \quad Ti \quad . \qquad [2.4]$$

If all routes have the same average sector transition time, E[Ti] = T, then this simplifies further to

$$E\left[N_S\right] = \sum_i \lambda_i \quad T = T\lambda_S \quad . \; [2.5]$$

Combining equation [2.2] and [2.5], we see that

$$\lambda_S = \frac{\kappa Q}{T} \quad \text{when } T_i = T \text{ for all i.} \qquad [2.6]$$

2.3 Sector Loading Simulation Results

When traffic flow is conditioned by requiring aircraft to fly in−trail on standard routes, fluctuations in sector loading are reduced. Freedom in routing results in greater independence between aircraft trajectories and produces greater variation in the traffic loading experienced by any given sector, even though the average sector loading is unchanged. This is evident in Figure 2.1which provides a histogram of the number of aircraft in a sector resulting from simulation of 50 hours of operation under two different flow conditions. In the first case aircraft enter on a single route with a 60 second in−trail requirement. In the second, aircraft enter on separate routes with no in−trail requirement.

The average number of aircraft in the sector is the same for both cases. But a significant difference can be seen in the tails of the distribution. In the in−trail case, there are never more than 15 aircraft in the sector. For the unstructured case, there is a significant probability of having 16 or more aircraft in the sector. In fact, as many as 22 aircraft were observed at one point.

2.4 Efficiency in Sector Loading

In a practical ATM system, the design of sectors and the control of traffic flow must ensure that individual sectors are not often presented with a sector loading that exceeds the level that the sector controller can safely handle. Let this constraint be expressed as a limit, Nmax, to the number of aircraft that can be accommodated in a sector.

If the loading for each sector in a system were precisely regulated, we could envision a system in which the average sector loading is very close to Nmax. Such a system would have a high level of ATM productivity in the sense that it would minimize the number of controllers required to handle the total traffic load. However, when the flow is random there will be fluctuations in the sector loading. The average loading of sectors must then be kept well below the maximum loading so that random fluctuations do not produce frequent sector overloads.

Suppose that a design criterion were imposed upon the ATM system to require that the probability of overload (NS > Nmax) be less than a specified probability α. In actual practice, the process for sector design will be much more complicated, but the basic constraints resulting from this criterion should nevertheless provide insight into a major consideration affecting the actual result. We can define the sector loading efficiency, ε, as the fraction of the maximum loading that can be tolerated on average, i.e.,

$$\varepsilon = \frac{E[N]\alpha}{N_{max}} \qquad [2.7]$$

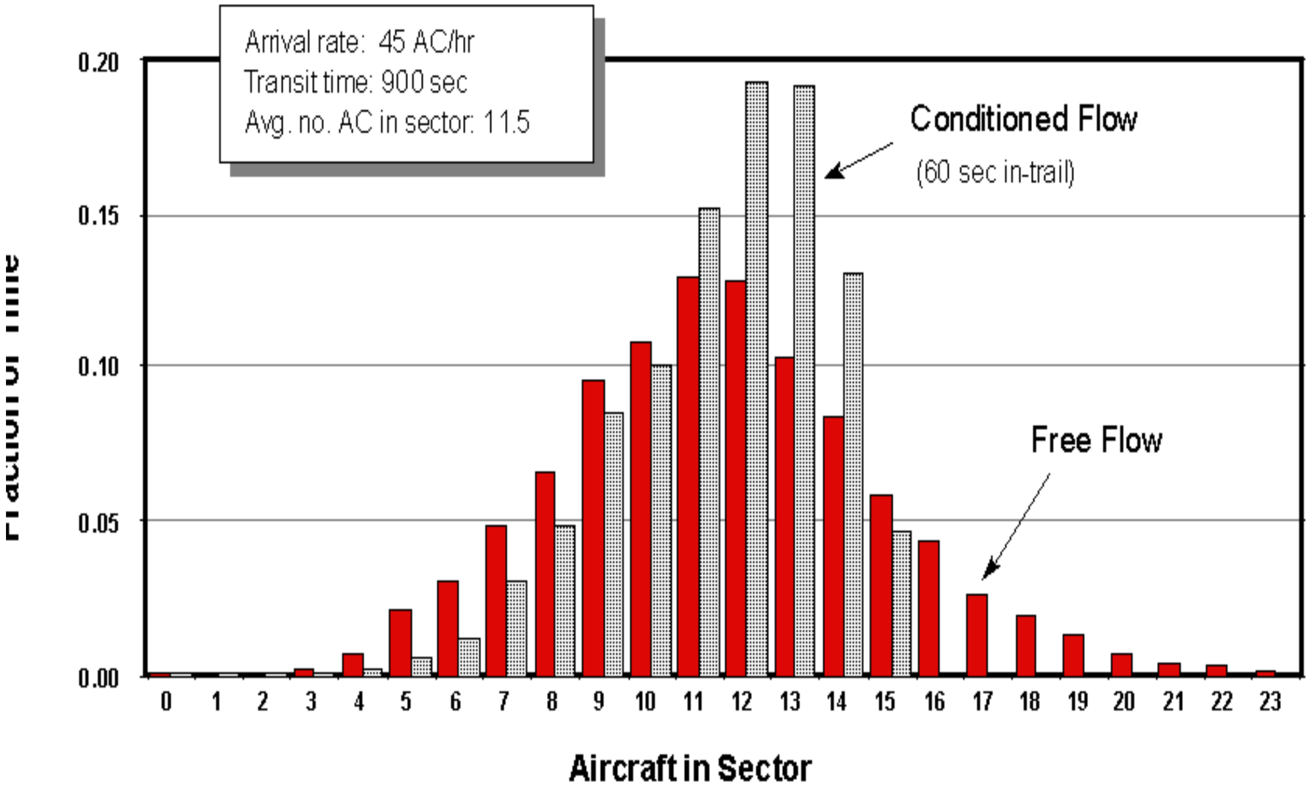

Figure 2.1 Free flow of traffic into a sector results in greater probability of transient peaks in section loading.

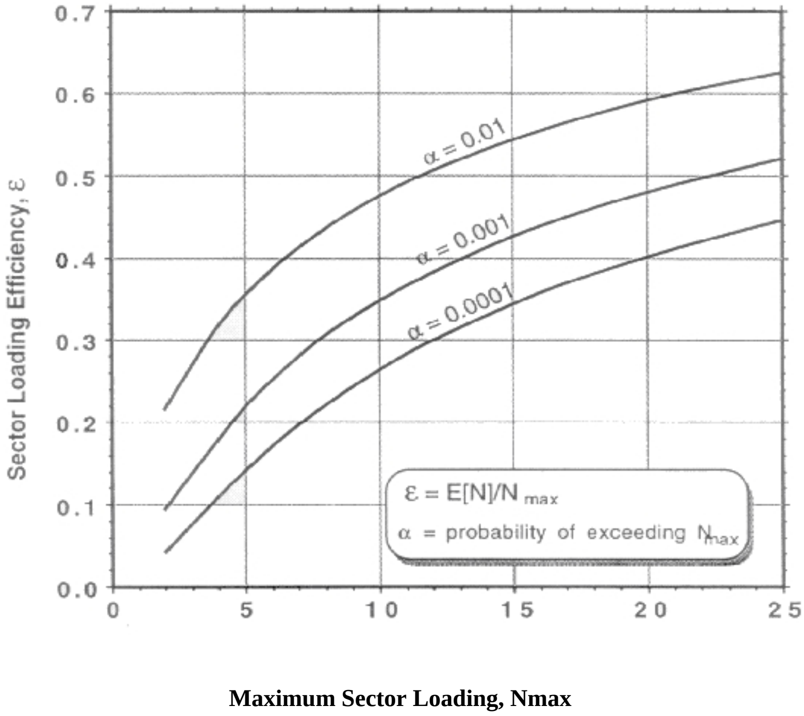

Figure 2.2. Efficiency in sector loading for random flow conditions.

Efficiency is the average number of aircraft in a sector divided

by the maximum acceptable traffic loading for that sector.

In Figure 2.1, the conditioned flow resulted in an average sector loading of 11.5 aircraft while the sector loading never exceeded 15 aircraft. This would correspond to a sector loading efficiency of 76% assuming that Nmax=15. Efficiencies for the free flow case can be obtained by using equations [2.1] and [2.2] to note that if P[NS > Nmax] < α then

$$\sum_{i=0}^{N_{max}} \left[\frac{E[N]^n}{n!} \; e^{-E[N]} \right] \geq 1 - \alpha \quad . \qquad [2.8]$$

E[N]α is found by solving [2.8].

Figure 2.2 shows the resulting efficiency for several values of α. For practical values of Nmax and α, the efficiency tends to be less than 50 percent. There appears to be a significant penalty in efficiency when the sector loading is conditioned to accommodate the randomness of traffic flow. The value of Nmax may also decrease with randomness due to greater workload being expended per aircraft. (This effect will be discussed in Section 4 below).

[2]. This is not to imply that aircraft densities are uniformly distributed over the entire airspace. This assumption requires only that the dispersion produced by removing the structure produce the appearance of randomness with respect to a single sector (which is typically less than 100 miles in extent).

3.0 Basic Conflict Rate Calculations

3.1 Alert Rate Model for Aircraft

For purposes of calculating conflict rates, it will be assumed that a conflict occurs whenever two aircraft approach within a distance Mh horizontally and Mv vertically simultaneously. If we define a cylindrical alert volume of diameter 2Mh and height 2Mv, then a conflict occurs whenever an intruding aircraft penetrates this volume [Ref. 2]. If Vij is the relative speed between the subject aircraft and some set of randomly located intruders, then the relative motion of the alert volume sweeps out new volume at the rate 4 Mh Mv Vij (see Figure 3.1). The expected rate at which aircraft appear within the swept volume is simply the volumetric density of traffic multiplied by the volume rate. Let κ be the volumetric density of all traffic. Then κ fV(Vij) dVij is the density of traffic in a relative velocity interval dVij at Vij. The incremental contribution to the conflict rate from this subset of aircraft is

$$d\lambda_{AC} = 4\,\kappa\,M_h\,M_v V_{ij}\,f_V\left(V_{ij}\right)dV_{ij}\,. \qquad [3.1]$$

The total conflict rate, found by integrating this over the entire range of relative speeds, is

$${}^{\lambda}\mathrm{AC} = 4\,\kappa\,M_h\,M_v\,E\left[V_{ij}\right] \qquad [3.2]$$

3.2 Alert Rate Model for Sector Controller

At what rate do conflicts arise within a single air traffic control sector? If there are NS aircraft within a sector, each experiencing conflicts at an average rate of λAC, then the total rate of conflicts seen by all aircraft is NS λAC. At practical traffic densities, the great majority of these conflicts are pairwise conflicts (involving only two aircraft). Thus the rate of pair conflicts seen by the sector controller is approximately

$${}^{\lambda}\mathrm{C} = 4\,\frac{N_S\,{}^{\lambda}\mathrm{AC}}{2} \qquad [3.3]$$

The ratio between the conflict rate for a pilot and that for a controller is NS /2. Since the aircraft conflict rate increases as NS, the rate experienced by the controller increases as NS 2. Today, en route controllers are sometimes expected to handle 10−15 aircraft simultaneously [Ref. 7]. At this loading, controllers must be able to handle a conflict rate that is 6 to 10 times larger than the conflict rate that experienced by an individual aircraft flying in the same sector. This has significant workload implications as traffic densities increase. The workload imposed by conflict resolution can appear to the sector controller to be increasing to intolerable levels while individual pilots are observing only moderate increases in interactions with other aircraft.

4. Sector Workload Modeling

4.1 Sector Workload Intensity

Schmidt [Ref. 8] found that a useful predictor of the workload experienced by the sector controller could be obtained by simply adding up the total work time imposed by all the various tasks given to the controller. Schmidt found that when this sum equaled approximately 80% of the total time available, controllers reported that the sector had reached its maximum loading. For purposes of this analysis, we will employ the Schmidt approach by defining a workload intensity measure, G, as

$$G = {}^{\tau}S^{\lambda}S + {}^{\tau}C^{\lambda}C\ . \qquad [4.1]$$

where τS is the seconds of routine work per aircraft handled by the sector and τS is the seconds of work associated with each conflict resolved. Schmidt suggested that τS= 60 seconds and τC=50 seconds were

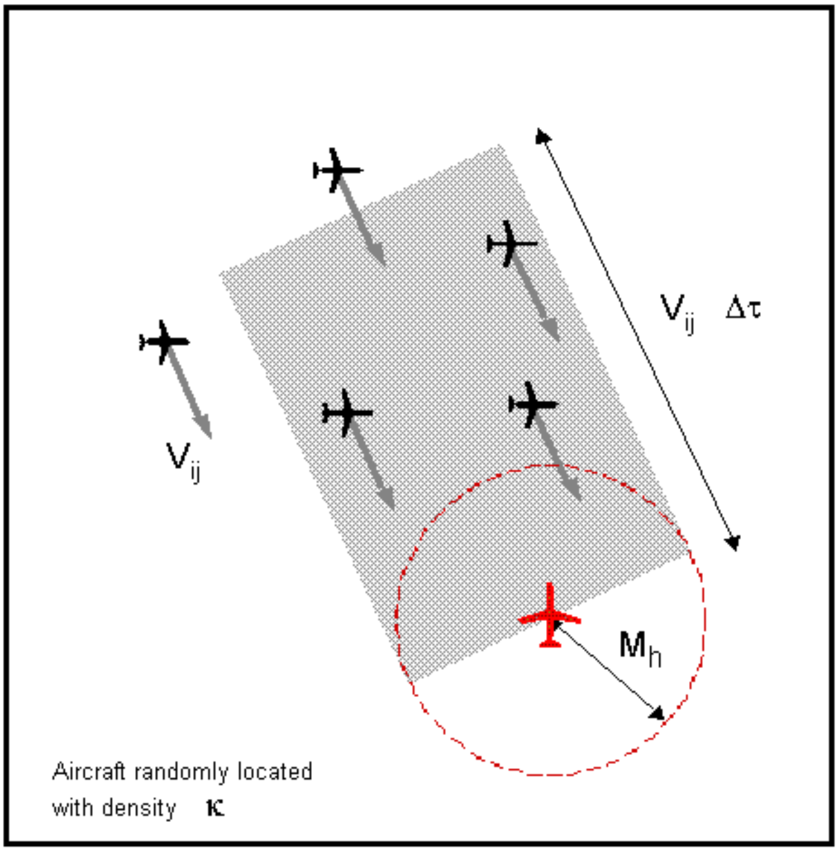

Figure 3.1 Alert rate model (horizontal view). Any aircraft with relative velocity Vij will alarm if it is within the shaded area.

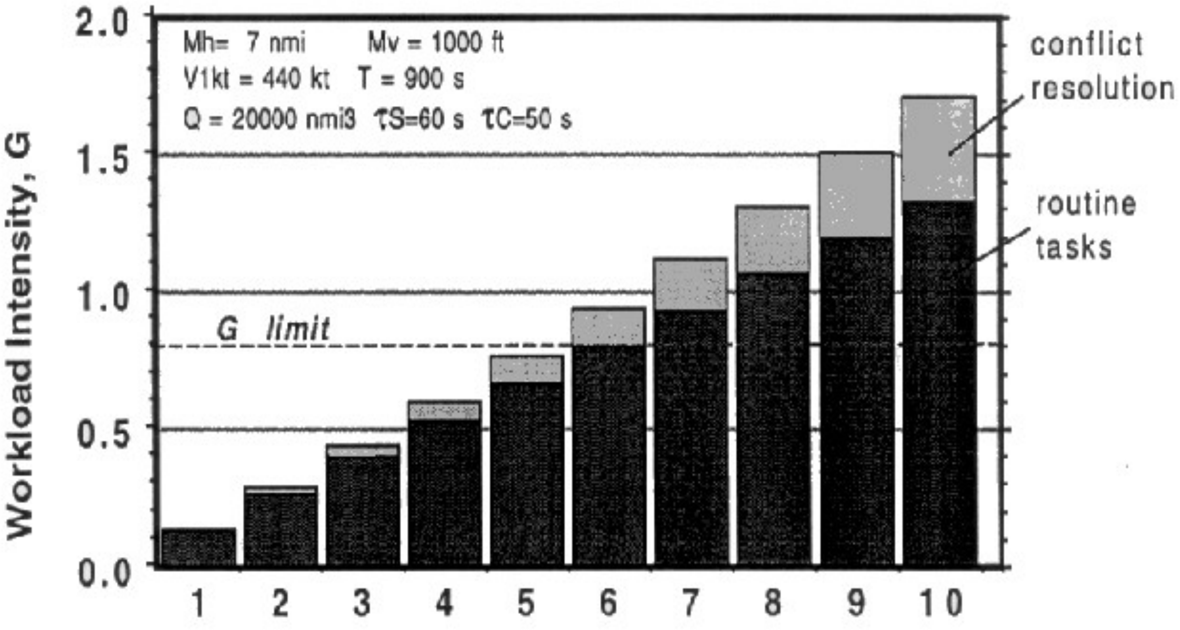

Traffic Density (AC per 10000 nmi3)

Figure 4.1. The amount of workload devoted to conflict resolution increases with traffic density.

appropriate for en route systems that he observed. When G=80%, the sector will be considered to be loaded to capacity. Note that as defined above G is an average value and does not take into account fluctuations in loading such as those we discussed earlier in regard to sector loading. Hence, the 80% value may have to be reduced in practice.

Using the sector loading and alert rate equations previously derived, the G value for a particular value of NS can be written

$$\tau_S = \frac{N_S}{T} + \tau_C \frac{2\, M_h M_v\, E\left[V_{ij}\right]}{Q}\, N_S\, 2 \quad [4.2]$$

Since NS has a Poisson distribution for unstructured flow, the expected value of this quantity can be written as

$$G = \tau_S \frac{E\left[N_S\right]}{T} + \tau_C \frac{2\, M_h M_v\, E\left[V_{ij}\right]}{Q} \left(E[N_S]^2 + E[N_S]\right) [4.3]$$

Substituting κQ for E[NS] then yields

$$G = \tau_S \frac{\kappa Q}{T} + 2\tau_C \; M_h M_v E\left[V_{ij}\right] \kappa \left(\kappa Q + 1\right) \qquad [4.4]$$

Because κQ >>1 for sectors of reasonable size and loading, the τC term is roughly proportional to κ2 while the τS term is proportional to κ.

Equation [4.4] is a key result that provides insight into the pressure placed upon the sector controller as traffic grows. Because the part of the workload intensity that results from conflict resolution increases roughly as the square of the traffic density, the sector workload grows at an increasingly rapid pace as traffic density increases. The fraction of the total workload associated with conflict resolution also grows. Figure 4.1 illustrates this effect.

4.2 Sector Loading Summary

Figure 4.2 provides a more complete summary of the sector loading characteristics for the case shown in Figure 4.1. For this case, the limit of supportable workload (G =80%) is reached at a traffic density of approximately 0.0005 AC/nmi3. At this traffic density, the sector acceptance rate is λS=40 AC/hr and there are an average of NS=10 aircraft in the sector simultaneously. The aircraft experience conflicts at a rate of λAC=1.3/hr, and the controller is handling λC=6.5 conflicts/hr. If τS and τC are reduced by increased automation, the workload curve (G) will be reduced, but the other curves will not change. This figure clearly shows that increasing traffic density has a more dramatic effect on the controller conflict rate (λC) than upon the aircraft conflict rate (λAC).

4.3 Benefits of Automation

One way of reducing workload, and hence improving sector productivity, is to reduce the task time parameters τS and τC . Figure 4.3 shows the permitted sector traffic density (assuming a limit of G=80%) as a function of the task workload parameters. Note that when the routine workload is high, as in a system with little automation, the value of τC has little impact on capacity. This is because the sector controller reaches the workload limit while most of the workload is still devoted to routine tasks. But as τS decreases, the workload limit is reached at a point at which conflict tasks constitute a greater fraction of the total workload. The permitted traffic density is then more sensitive to τC.

Automation applied to routine traffic management tasks (such as arranging flight strips, accepting hand−offs, and conveying frequency assignments to aircraft) continues to make progress in reducing the time associated with routine tasks. There has not been such obvious success in the reduction of task times for tasks that require evaluation and judgment, such as conflict resolution. Research on conflict probes and other decision−support aids should anticipate the need to operate under conditions where the value of τC is critical to the viability of the system.

5. Multiple Conflict Considerations

The workload analysis presented in the previous section focused upon the average sector loading and average sector workload. However, an additional performance consideration arises with regard to the probability of simultaneous multiple conflicts. Prompt resolution also allows more cost−effective maneuvers to be implemented (requiring less deviation from the user−preferred trajectory). Several considerations make it undesirable for a sector controller to have more than one unresolved conflict at the same time. First, there is the obvious result that the resolution of one conflict must be delayed while the other is being serviced. Delayed resolution is less efficient and may decrease the typical margin of safety. Second, there is the increased probability of distraction, reduction in situation awareness, and human error. This section will examine how the probability of simultaneous multiple conflicts depends upon traffic densities.

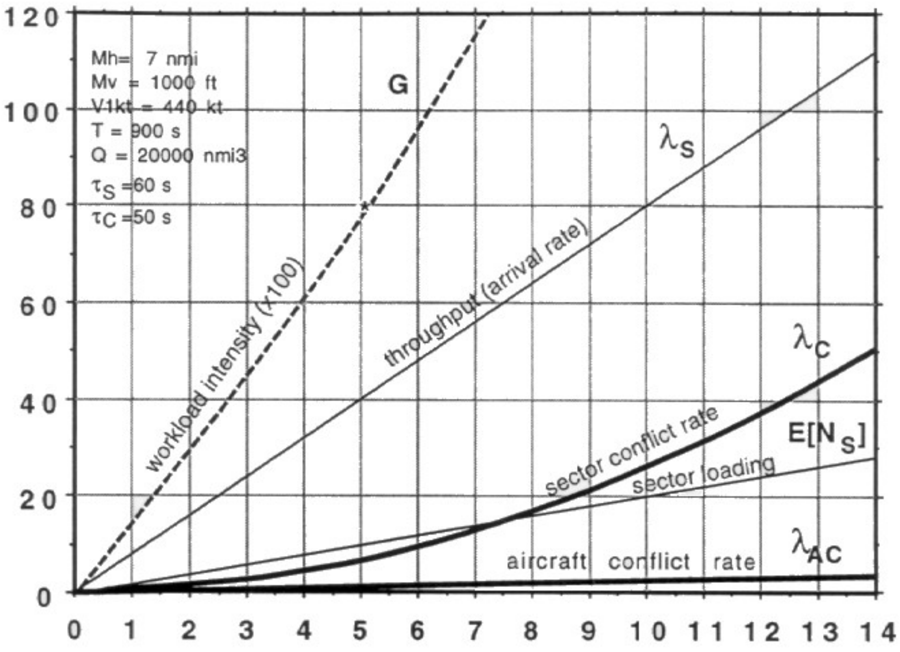

Traffic Density (AC/10000 NMI 3)

Figure 4.2. Effect of traffic density on sector loading variables and workload

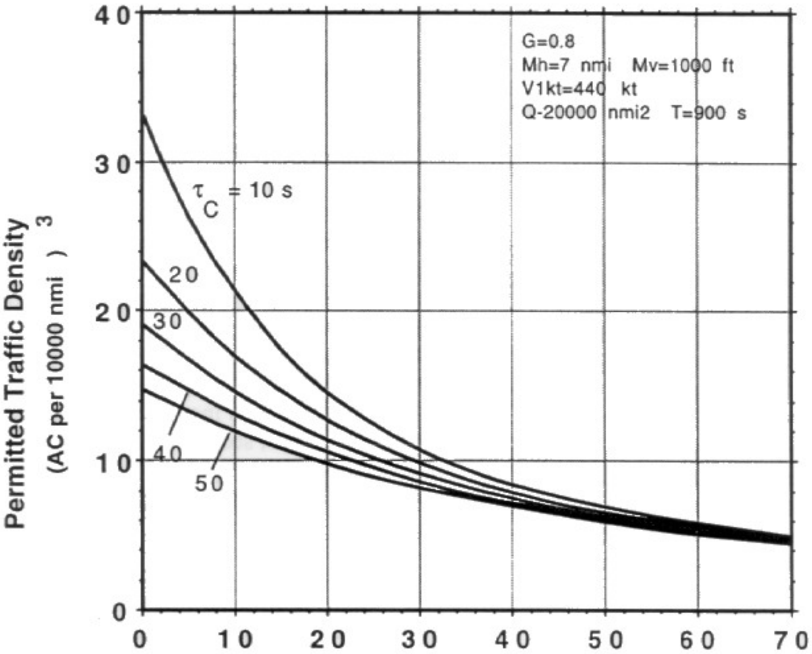

Routine Workload Time Parameter, ? τ_s (sec)

Figure 4.3. Permitted traffic density (G=0.8 limit) as a function of task worlkoad parameters

5.1 Queueing Model for Conflict Resolution

If we view the controller as the *server* in a single−server queue and view conflicts as *customers* to be serviced, then classical single−server queuing equations can be used to estimate the probability that simultaneous conflict resolution tasks will arise. Assume that the controller resolves conflicts sequentially and independently and that the average time required to resolve a conflict (to the point of devising and communicating instructions to aircraft) is τR. (Note that τR is somewhat less than the τS parameter used in the Schmidt model since τR includes only the time required to select and communicate the solution to the conflict whereas τS includes the time required to monitor the selected resolution and return aircraft to their desired trajectories.) In the calculations which follow, we will assume let the distribution of conflict resolution time be modeled as a simple exponential distribution,

$$P\left[\text{resolution time} < T\right] = -\exp\left[-T/\tau_R\right] \quad [5.0]$$

The traffic intensity of a queue is the fraction of time that the server is busy. In this case the traffic intensity is $\lambda C\ \tau R$. According to single−server queueing theory, the probability that there are n unresolved conflicts present at any given time is

$$Pn = (\lambda C\ \tau R)^n \quad (1 - \lambda C\ \tau R). \quad [5.1]$$

This is also the probability that a newly arrived conflict must wait for n previous conflicts to be resolved before receiving service. Thus, the probability that there is no delay in the initiation of conflict resolution is $p0 = 1 - \lambda C\ \tau R$. Figure 5.1 shows the probability of the conflict queue length according to this model. (Note: For other values of τR, the abscissa can be scaled by a factor $40/\tau R$.) It can be seen that as λC increases, the probability of two or more simultaneous conflicts increases inexorably. Given the large number of conflicts that must be resolved system−wide, even a small probability of conflict queueing can result in hundreds of multiple conflict instances per year of ATM system operation. Depending upon the system design and performance criteria, this may impose a more stringent limitation on traffic densities than the average workload criterion.

5.2 Multiple Server Queues

Adding additional servers is a standard technique for improving the performance of queueing systems. In the context of sector control, one could envision control processes in which conflicts were allocated dynamically between adjacent sectors, thus allowing the less busy controllers to assist a sector that was temporarily overloaded. In the case of conflict probes that look 15 or so minutes ahead, conflicts in one sector will often be detected in an upstream sector. There is the potential for assigning the conflict resolution responsibility from one sector to another, thus effectively bringing another server into the resolution process.

Insight into the value of adding additional servers can be inferred from the queueing performance of multi−server queues. Figure 5.2 shows the probability that an arriving customer will have to wait in the queue for systems with different numbers of servers. An M/M/c queue is assumed in which all servers have the same service rate.

Note that adding servers can have a significant effect upon the probability of wait. However, the effect of adding only one or two additional servers may not be dramatic when the traffic intensity is high. In an ATM system, it is doubtful that reallocation of responsibility for conflicts could involve more than one or two other sectors. Thus, while reallocation would be beneficial, it may not completely resolve the problem of multiple simultaneous conflicts.

6. Conclusions

The analysis presented provides quantitative models that suggest that:

- The variability in sector loading (number of aircraft in a sector) increases as traffic flow structure is relaxed.
- Variability in sector loading can reduce sector productivity by forcing the average loading to be kept further below the maximum tolerable loading.
- The conflict rate experienced by the sector controller is often five or six times greater than the conflict rate experienced by individual aircraft.
- The sector controller workload intensity grows approximately as the square of traffic density whereas the traffic−related workload of the aircrew grows linearly.
- Simultaneous multiple conflicts will be a routine occurrence when traffic flow is not highly conditioned, and may impose a limitation on the tolerable traffic loading.

6.1 The Need for Productivity Increases

Traffic growth in a sectorized ATM system can be accommodated by a combination of increasing sector productivity and making sectors smaller (thus increasing the number of sectors in the system). Making sectors smaller has the drawback that it produces an increase in the controller work force, thus increasing the costs of ATM. A second and more fundamental difficulty arises from the fact that there are practical limits to the minimum size of air traffic control sectors. If sectors become too small, the sector controller does not have enough time (or airspace) to effect control actions, inter−sector coordination becomes more complex, and the transients associated with hand−off begin to degrade the reliability of the system. Thus, sector productivity must eventually increase in order to prevent the limitations of the sectorized system from becoming a barrier to traffic growth.

The traditional response to increasing traffic density is to institute more restrictive flow control and flow structure. Because of the costs and inefficiencies that result, the aviation community in the United States has made a commitment to avoid this response. A preferable solution is to increase sector productivity through automation support and decision−support tools. This can be important to the progress of free flight in the near term. However, as the analysis has shown, reasonable assumptions regarding the amount of workload relief provided by decision− support tools leads to the conclusion that only limited growth can be accommodated by this means. A transition to new ways of providing air traffic control services will be required in the foreseeable future.

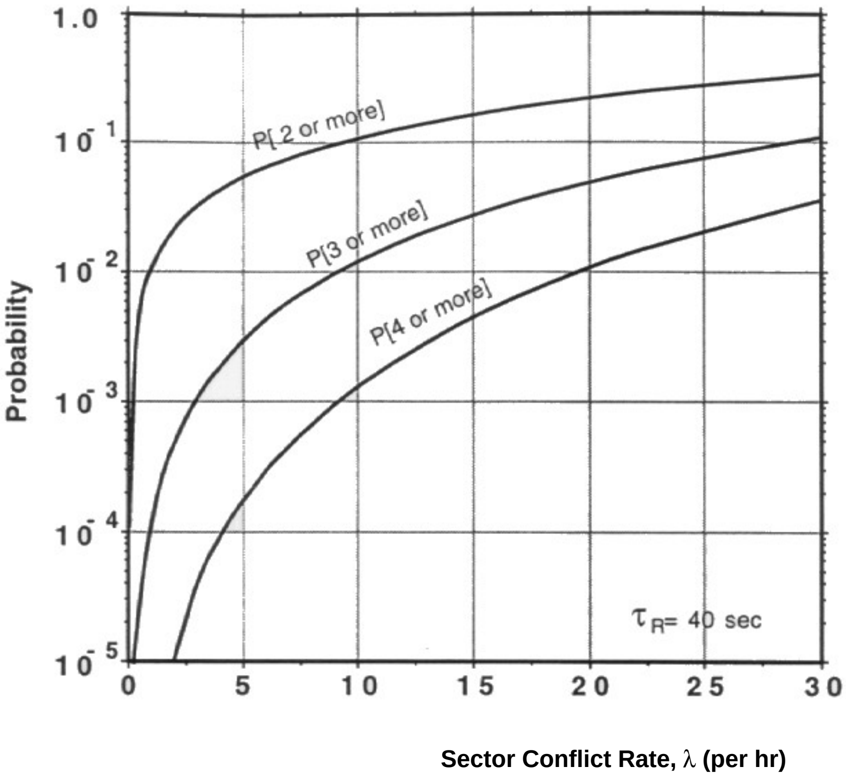

Figure 5.1. Probability for the number of conflicts simultaneously unresolved assuming that a conflict has just arisen

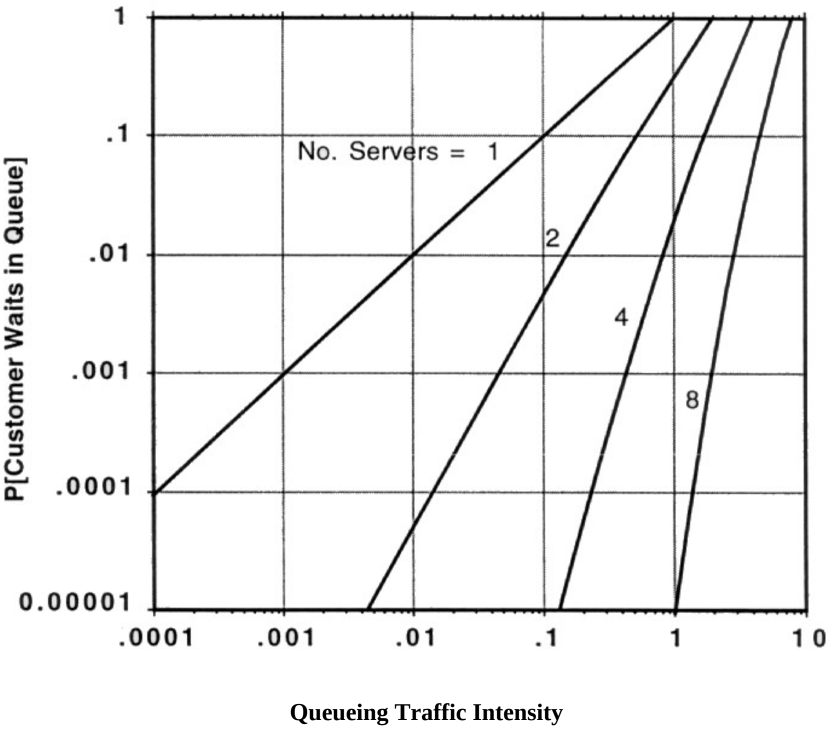

Figure 5.2. The probability that a customer must wait for service is decreased when multiple servers are available.

One avenue for exploration is the use of more dynamic allocation of controller resources through use of dynamic sector boundaries or dynamic distribution of workload among a wider team of human controllers. Such concepts are not well supported by the current ATM infrastructure and procedures, but may prove useful in the future.

Another potential solution to the dilemma is suggested by the finding that the conflict rate experienced by aircrews is tolerable at densities at which the sector controller workload cannot be supported. This means that much greater traffic densities could be tolerated if the sector controller can be relieved of routine conflict resolution tasks. This suggests a system in which automation analyzes trajectories and suggests resolution actions while the involved aircrews serve as the human monitors of the result. In effect, such an approach would escape the single−server queue limitations of sectorized control by introducing a multi−server queuing process. Because of the long lead times associated with the infrastructure improvements and the development of ATM automation, the necessary work to prepare for this breakthrough should begin now. With foresight and sufficient preparation for the future, our ATM systems should be able to provide safe and efficient service to airspace users for decades to come.

References

[1] RTCA Select Committee on Free Flight, Report of the RTCA Board of Directors' Select Committee on Free Flight, October 1995., RTCA, Inc., Washington , D.C.

[2] Andrews, J.W., and Hollister, W.M., "Impact of Intent Information on Alert Rates", Air Traffic Control Quarterly (to be published)

[3] Credeur, Leonard, and Capron, William R., "Simulation Evaluation of TIMER, a Time−Based, Terminal Air Traffic, Flow−Management Concept", NASA Technical Paper 2870, February 1989

[4] Erzberger, Heinz and Tobias, Leonard, "A Time−Based Concept for Terminal−Area Traffic Management", Advisory Group for Aerospace Research and Development, Conference Proceedings No. 410, 1989

[5] Vandevenne, H.F., Andrews, J.W, "Effects of Metering Precision and Terminal Controllability on Runway Throughput", Air Traffic Control Quarterly, Vol. 1(3) 277−297, 1993

[6] Andrews, J.W., "Further Analysis of the Ramp Model of Airport Delay", Memo 41A−1704, M.I.T. Lincoln Laboratory, 31 October 1995

[7] Politano, Arturo L., Flores, Michael, Winer, David E., and Czech, Horst C., "Toward Making the NAS More Productive: GPS−Enabled Direct Routing in the New England Region", Air Traffic Control Quarterly, Vol. 4, No. 2, pp. 81−107

[8] Schmidt, David K., "A Queuing Analysis of the Air Traffic Controller's Work Load", IEEE Transactions on Systems, Man, and Cybernetics, Vol. SMC−8, No. 6, June 1978, pp. 492−498

BIOGRAPHIES OF AUTHORS

John Andrews is Senior Staff in the Air Traffic Control Automation Group at M.I.T. Lincoln Laboratory. He was one of the senior system analysts during the initial development of the Traffic Alert and Collision Avoidance System (TCAS) and in that area has made contributions to aircraft altitude tracking techniques, analysis of computer logic, human subject flight testing, and the analysis of pilot visual acquisition performance. He has served as a consultant to the National Transportation Traffic Board in the investigation of mid−air collisions and has been involved in the technical management of the Terminal Air Traffic Control Automation (TATCA) Program at Lincoln Laboratory. He holds a B.S. in Physics from the Georgia Institute of Technology and a M.S. in aeronautical engineering from the Massachusetts Institute of Technology.

Dr. Welch is leader of Lincoln Laboratory's Air Traffic Automation Group responsible for assisting the FAA in the development of computer aides for air traffic controllers in all air traffic control domains. He was part of the team that led the development of the FAA's Mode S air traffic control radar beacon system. He led Lincoln Laboratory's program to develop surveillance techniques for air−to−air, beacon−based collision avoidance which led to the Traffic Alert and Collision Avoidance System (TCAS) surveillance design. He initiated the Terminal Air Traffic Control Automation Program at Lincoln Laboratory which led to the establishment of the FAA's current Center Terminal Automation System (CTAS) activity. He lead an effort to develop and test automation to improve safety of operations on the airport surface. He is leading Lincoln Laboratory efforts to help the FAA develop a new en route automation infrastructure. As a consequence, he is involved in studying the quantitative implications of free flight on future air traffic automation systems.