?lanning Horizon Requirements for Automated Terminal Scheduling*

H. F. Vandevenne
J. W. Andrews
J. D. Welch

Lincoln Laboratory, Massachusetts Institute of Technology
Lexington, Massachusetts, USA

Abstract

This paper presents the results of an engineering analysis of the ability of an automated terminal scheduling process to achieve efficient use of runways. The motivation for the analysis is the need to understand possible architectures for an implementation of the proposed Terminal Air Traffic Automation (TATCA) system. The performance of TATCA is dependent upon metering precision and the controllability that TATCA can apply to aircraft entering the scheduling process. *Controllability* refers to the amount of time by which the flight time of an aircraft can be lengthened or shortened between the scheduling horizon and the chosen runway. The analysis concludes that when current en route metering mechanisms are used to deliver traffic to the terminal, the terminal scheduler needs a controllability window of 300 seconds or so in order to achieve full runway utilization. Because this amount of controllability is often achievable within the terminal area itself, a TATCA system can provide significant benefits prior to the implementation of further improvements in the en route metering process.

1. Introduction

By the end of the next decade, several significant new automation projects will have substantially altered the manner in which the air traffic control system in the United States regulates the flow of traffic into terminal areas. The principal impacts will come from the implementation of the Arrival Sequencing Program (ASP), En Route Spacing Program (ESP), Enhanced Traffic Management System (ETMS), Terminal Air Traffic Control Automation (TATCA), and the general metering functions provided by the Advanced Automation system (AAS). Clearly, whenever two or more automation systems are available to control the same flow of traffic, they must divide responsibilities in a compatible and efficient way. This paper addresses the relationship between terminal automation (TATCA) and the en route metering functions that deliver traffic to the terminal area.

The TATCA program is a major FAA RE&D project that focuses upon planning aids that will improve traffic flow in the airspace near major terminal areas. The objective of TATCA is to provide for safe and efficient traffic flow through computer-based planning functions that establish and coordinate arrival sequencing and scheduling. By design, the computer is concerned primarily with suggesting efficient sequencing and timing while, the controller is fully in charge of the execution of the plan. There is evidence that such efficient planning results in less delay and congestion, increased throughput and more economic flight profiles as well as a reduction of the workload devoted to coordination and spacing [1,2]. The design of TATCA draws upon research done over the last two decades in various parts of the world. We mention specifically TMA and FAST at NASA [3,5], COMPAS in Germany [4] and work in the UK [6].

A fundamental question that arises with the design of a terminal automation system concerns the extent to which such a system can achieve high runway throughput in the face of imperfect delivery of traffic from en route airspace. A detailed examination of the actual system suggests reasons why improvements within the terminal area itself may be important to full utilization of runway capacity. For example:

- At many airports, often tower-en route aircraft become known to the planning process at a point too late for effective en route metering. In such a case, the ability of the terminal to efficiently schedule the traffic is critical.

---

*This work was sponsored by the Federal Aviation Administration.

- The selection of the airport acceptance rate on which metering is based depends upon a detailed knowledge of the current state of the terminal. Among the dynamic events that affect terminal capacity are changes in wind, weight mix of arriving traffic, distribution of arrivals among fixes, missed approaches, skill of final vector controller, and proportion of departures to arrivals. Automation which can analyze the terminal situation and promote a responsive relationship to the metering system can ensure that the acceptance rate is not set too low (which wastes capacity) or too high (which overloads the terminal and ultimately results in a request for lowering the acceptance rate as a way of protecting against overload).

- The actual terminal throughput rate often changes too quickly for en route metering systems to respond. Throughput suffers from response delays, even when the achievable throughput is known. By fully employing the controllability that exists within the terminal area, terminal automation can recover this lost capacity.

- The achievable throughput rate depends upon the precision of the final spacing process, and this precision can only be addressed by improvements in terminal processes.

- Because of proximity to other congested terminals, some terminals will find it very difficult to extend detailed traffic planning into en route airspace. For such terminals, early progress toward automation assistance may require a system that can function within terminal airspace alone.

A key question concerns whether the controllability available to a terminal scheduler allows it to achieve an efficient schedule. In a simulation of Denver airspace Credeur [5] concluded that some runway throughput is lost when the runway scheduling process begins less than about 30 minutes prior to landing. If a 30 minute planning horizon is to be achieved, then the TATCA planning boundary would have to extend out of the current terminal boundaries to the top-of-descent point (or, alternatively, the en route metering process would have to be closely coupled with TATCA planning so that problems with the arrival stream begin to be corrected prior to terminal entry). For practical and operational reasons, it is desirable to implement early terminal automation within the current terminal boundaries without requiring significant modifications or enhancements to planned metering and flow control functions. Hence, a quantitative analysis is needed to determine whether the loss of efficiency produced by such a partitioning is significant.

When evaluating uncommon overload conditions and operations near system capacity, very large traffic samples are required in order to draw conclusions. For this reason, it is risky to attempt to answer some statistical performance questions using full scale human subject experiments for which typically only a few hours of traffic can be gathered. Although human subject testing is extremely valuable for system validation, a fast time simulation is still needed to explore statistical variation. This is particularly true when the system is being operated near capacity and throughput increases of 10 percent or less are to be verified. For these reasons, most of the figures presented in this study are based on the simulation of 24 hours of traffic data per curve.

Formulation of the Terminal Scheduling Problem

In order to reduce the problem to an analyzable form, this paper formulates the terminal scheduling problem as follows: An automated scheduling aid is available for the last portion of the flight (within some scheduling horizon). This aid can compute and execute a schedule that results in efficient runway utilization for the arrival stream delivered to it. The possible schedules that can be achieved are limited by the controllability of the landing times of arrivals. *TATCA controllability is the amount by which the nominal landing time of an arrival crossing the planning horizon can be shifted forward or backward under defined constraints on available path stretching and speed control.* The system that delivers arrivals to TATCA is subject to random fluctuations from the ideal arrival timings (i.e., aircraft do not always arrive regularly spaced at the exact rate at which they can be landed). Randomness will occasionally result in an aircraft being delivered so early or late that the TATCA scheduler cannot land it using the available controllability. *The average rate at which this occurs is called the controller intervention rate.*

The term "intervention" derives from the fact that the air traffic controller must, in some way, provide additional delay for aircraft that cannot be scheduled by the automation. The manner in which the TATCA system responds to such exceptions is still being refined. It is clear, however, that imposing such delay involves workload or less efficient flight profiles (as well as endangering the smooth traffic flow that TATCA seeks to maintain). Ideally, the intervention rate would be zero. However, in any practical system it is desirable to accept a non-zero intervention rate for the sake of increased throughput. In the analysis that follows, intervention rate limits between one and thirty percent are considered.

The purpose of this study is to determine the quantitative relationships between capacity, throughput, delay, and the controllability for a given intervention rate limit and to determine how those quantities depend on the level of randomness in the aircraft arrival stream.

2. Controllability in a Typical Terminal Area

The TATCA planner assumes that the arrival time for a particular aircraft can be altered in two primary ways: 1) by executing speed reductions earlier or later. 2) by flying a longer or shorter route selected from a predefined repertoire (this is referred to as path-stretching). In general, the controllers may have access to additional control techniques that are not available to the TATCA planner (such as holding within the TRACON, vectoring onto non-standard routes, flying S-turns, etc.). These additional techniques can be applied to assist scheduling, but cannot be specified for use by the automation.

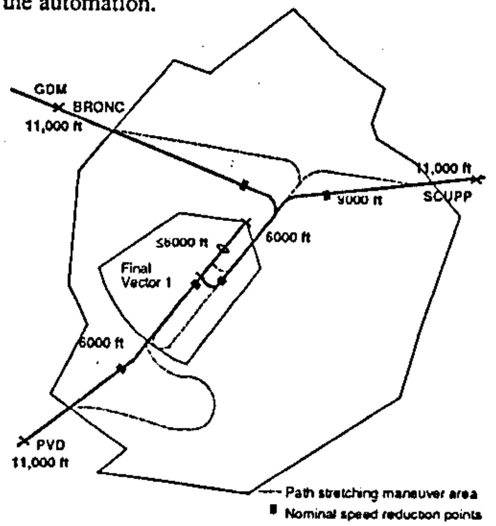

Figure 1. Boston TRACON high-altitude approach paths to runway 4R.

In order to provide some feel for the controllability that is available for actual approach routes, the Boston Terminal Radar Approach Control (TRACON) area will be examined (see Figure 1), for approaches to runway 4R.

Standard high-altitude turbojet traffic is fed to three fixes: traffic from the west holds at the Gardner (GDM) VOR and enters via the BRONC fix; traffic from the south holds at the Providence (PVD) VOR and enters via that fix, and trans-oceanic traffic, possibly augmented by overflow from PVD, holds at the SCUPP fix and enters via that fix. There are also several standard low-altitude traffic entry fixes and routes (not shown). Shown are the points on the nominal trajectory where the speed reductions from 250-210 and 210-170 knots are initiated under normal circumstances, and the areas where so-called path-stretching is allowed to occur.

Figure 2 plots the amounts by which flight time can be decreased (negative values) or increased (positive values) over nominal, for the three approaches to runway 4R. For speed control, controllability increases steadily as the time horizon increases. When path stretching is available, controllability still benefits from increased time horizon, but it is more strongly dependent upon the availability of path options such as extension of the downwind path segment during final approach ("tromboning"). The PVD approach routinely allows only for a small dogleg path extension providing some 60 sec of delay capability, although occasionally a larger extension is used as shown on Figure 1. In case of more severe congestion the controller can divert a small number of aircraft to the SCUPP approach, thus imposing some 10 additional minutes of delay.

The dashed lines in Figure 2 correspond to variations in flight time observed for actual traffic at Atlanta and Denver [7]. These actual observations are generally consistent with the Boston calculations, although the greater controllability indicated by the second dashed line may indicate that the controllers at Atlanta employed more path stretching than was assumed for Boston.

Speed control alone provides roughly 120-200 seconds of controllability. When path stretching is added to speed control, roughly 400 seconds of controllability becomes available.

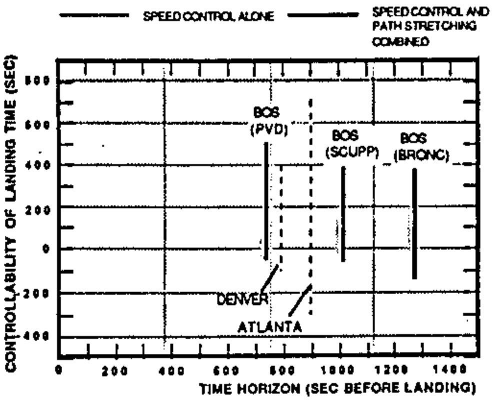

Figure 2. Ability to vary landing time starting at a specified time horizon (calculated for Boston and observed for Denver and Atlanta).

Some of the graphs that follow are simplified by assuming that controllability is equal for all aircraft entering the terminal area. In applying such curves to cases in which controllability differs for the different approach routes, it is worth noting that performance tends to be dominated by the route with the least controllability.

3. Scheduling Model

The prime function of the TATCA scheduling algorithm is to determine an efficient set of landing times (which also determines the landing sequence). Once a schedule is proposed, a Time-constrained Trajectory Generation (TTG) algorithm then attempts to generate a 4D trajectory to bring the aircraft from its present state to the runway at the scheduled time. In most cases, the aircraft will have to be expedited or delayed relative to the nominal time-to-fly.

If the TTG algorithm fails to find a trajectory that meets the schedule, then the aircraft becomes an exception and it contributes to the controller intervention rate, C, defined earlier.

3.1 First-come/First-served (FCFS) Scheduling

Because of a desire to impose delays fairly upon all classes of users, terminal air traffic control generally attempts to land aircraft in first-come-first-served (FCFS) order. *FCFS is defined here as scheduling landings in the order that the aircraft could have reached the runway if each one were flying the nominal path and profile, unconstrained by the presence of other aircraft.* In the analysis that follows, a strict use of FCFS order will be assumed.[1] The unconstrained landing times that are computed to determine the FCFS order will be referred to as the nominal landing times (NLT). When any degree of congestion arises, the scheduler must alter the NLT's in order to meet separation requirements and fill gaps in the arrival stream. The resulting modified landing times are called the scheduled landing times (SLT's). This scheduling process is illustrated in Figure 3.

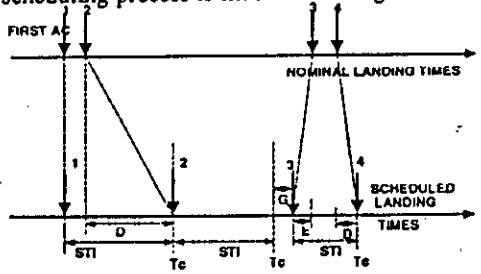

STI: Scheduled Time Interval (variable)
D: Scheduled Delay
Tc: Candidate Time Slot
E: Amount Expedited
G: Gap or Excess Time Separation Between Landings

Figure 3. FCFS Scheduling Process

The earliest acceptable landing time, T_c, is found by first adding the minimum required landing time separation (STI) to the scheduled landing time of the previous aircraft in the sequence. The value of STI is derived from the in-trail separation standards or from supervisory inputs to the automation. Note that the runway capacity (the theoretical upper limit on runway throughput) is $1/STI_{av}$, where STI_{av} represents the statistical average minimum landing separation. In order to maximize throughput under congested conditions, the scheduler attempts to land each aircraft as soon as possible after T_c. Whether the aircraft can actually reach the runway by that time depends on the controllability. If the earliest achievable landing time, T_e, is later than T_c, then the scheduler accepts an excess of separation of duration $T_e - T_c$ in front of the new arrival (this describes the case for aircraft 3 in Figure 3). This excess separation is called a gap and represents lost capacity. If $T_e < T_c$, then the scheduled landing time is set equal to T_c (as for aircraft 2 and 4) and no capacity is wasted.

---

[1] In actual scheduling, there are additional constraints that can lead to the FCFS order being imperfectly realized. The analysis in this paper is not thought to be sensitive to moderate perturbations of the FCFS order.

However, the new arrival is delayed relative to the earliest time at which it could have landed in an non-congested situation.

3.2 Scheduling and Controllability Window

The extent to which the scheduling process alters the nominal landing times (NLT's) can be characterized by the probability distribution of the time shift $SLT_i - NLT_i$ where i is the index of a randomly selected aircraft. In order to simplify the presentation, it is helpful to note that it is the width of the controllability window alone that is important in determining the achievable throughput. That is, if aircraft can be expedited a time E and delayed a time D, then the same throughput is achieved by a system that cannot expedite at all but can delay by an amount W=E+D.

Figure 4 shows a sketch of such a distribution and indicates how the controller intervention rate, C, is determined by the probability that lies outside the controllability window, W. The actual value of C depends greatly on how close the arrival rate is to the capacity and the degree of randomness in the arrival pattern. We will study several such patterns.

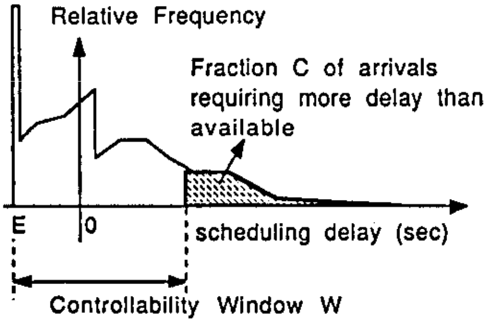

Figure 4. Distribution of delays imposed by the scheduling process.

4. Arrival Delivery Processes

The goal of this feasibility study is to determine the highest sustainable throughput rates corresponding to a set of established performance requirements and conditions. Three arrival delivery models will now be described.

Poisson Delivery Model: The final arrival stream is characterized by a specified average rate R such that the probability of an arrival in some small interval of time dt is Rdt. The arrivals are then generated by a classic Poisson process. Note that this is the process that might be expected as the result of a good flow control program which delivers aircraft to the airport vicinity at a specified rate (on the average), but where the actual arrival times vary greatly due to random effects (winds, delays at airport of origin, unanticipated ATC delays en route, etc.) Since the sum of Poisson processes is itself a Poisson process, the final stream can be viewed as resulting from the merging of N streams that are themselves Poisson.

Minimum Miles-in-trail Separation Model: For this model, the final arrival stream is the sum of N independent Poisson streams as before, but each stream is modified with a *minimum miles-in-trail separation requirement.* Such constraint clearly has a beneficial effect on reducing bunching in an individual arrival stream. This model is somewhat more realistic than the Poisson Delivery Model since consecutive arrivals over the same fix are required to have a minimum separation. In a sense each individual stream has been "scheduled" as reflected in the minimum separation requirement. If there were, for example only one entry fix, then the stream would not have to be modified, i.e., no flight time adjustments would be necessary. If, on the other hand, there are several fixes, the effect of the miles-in-trail procedure may well be greatly diminished in the combined stream. By how much will depend on the fix loadings, as will be demonstrated.

Metered Model: In this model, the delivery process is based on the specification of uniformly spaced landing slots (with spacing equal to STI). Each arrival is assigned to one of these slots, with some slots remaining unassigned if the arrival rate is below capacity. The metering system estimates the arrival time at the metering fix for that targeted landing slot. Aircraft are delivered to the fix with a time error that varies from aircraft to aircraft but can be modelled by a simple normal distribution, with a given bias and standard deviation. The error distribution may vary for each metering fix (since each flight path will be affected differently by wind estimate errors and other uncertainties).

The Metered Model approximates the effect that a conscientiously applied ERM-ASP Metering program might have on the NLT distributions.

5. Throughput Versus Controllability

A simulation can now be conducted to determine the maximum arrival rate, R, that can be tolerated without exceeding a specified exception rate, C. Results will be presented for all three delivery models.

5.1 Results for the Poisson Delivery Model

Figure 5 shows a delay histogram generated by simulating 1000 arrivals for an arrival rate of 35 per hour with runway capacity of 40 per hour, using a FCFS scheduler with no expedite capability. Observe the large pulse at delay zero representing the 14% of arrivals scheduled without any delay.

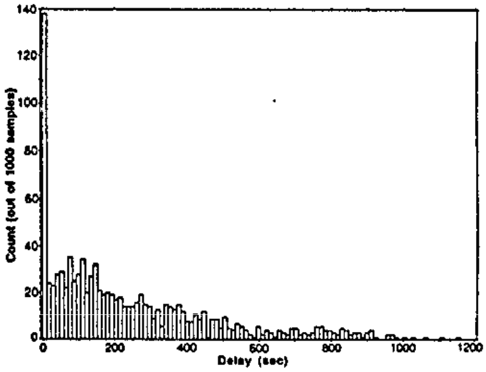

Figure 5. Typical Delay Histogram for Poisson Model.

Figure 6 shows the integral of the previous histogram (cumulative distribution). From this curve it is possible to read off the number of arrivals (on the ordinate) with scheduled delays less than a given controllability window (on the abscissa). Similarly, one can interpret the abscissa as the required controllability window size, W, at which the ordinate represents the number of arrivals scheduled without exceeding the controller intervention rate C. For example, for a maximum of C=10%, the window must exceed W=600 seconds.

The simulation was repeated for a range of arrival rates from 28 to 40 aircraft per hour (where 40 equals the capacity). The results plotted as Figure 7 for intervention rates of 1, 5, 10, 20 and 30% represent the primary output of this study.

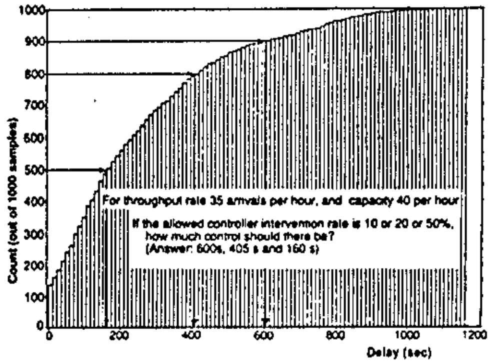

Figure 6. Cumulative Delay Histogram for Figure 5.

It should be noted that the gain in the achievable throughput rate is steepest for the lower values of W, with the curves becoming significantly flatter for W above 700 seconds or so. Above the figure are some typical values for W using different combinations of control mechanisms (speed control only, speed and path stretching, etc.).

For an intervention rate of C=5% and using speed controls only, the operational throughput rate is limited to less than 24 aircraft per hour (a 60 percent runway utilization rate).

If, in addition, path stretching is used for a total delaying capability of 600 seconds, the throughput rate can be increased to 34.2 aircraft per hour (.85 runway utilization rate) for the same 5% intervention rate.

If, in addition, top-of-descent control were made available (another 180 seconds) the throughput rate could be increased further to 35.0, but the gain is marginal (less than a .8 unit increase). The reason is that we are reaching the knee in the curve and enter in a regime of diminishing returns. That conclusion depends somewhat on the allowed intervention rate, however. On the 1% curve, the corresponding increase would have been somewhat more substantial: from a throughput rate of 32 to 34 aircraft per hour (a 2 unit increase).

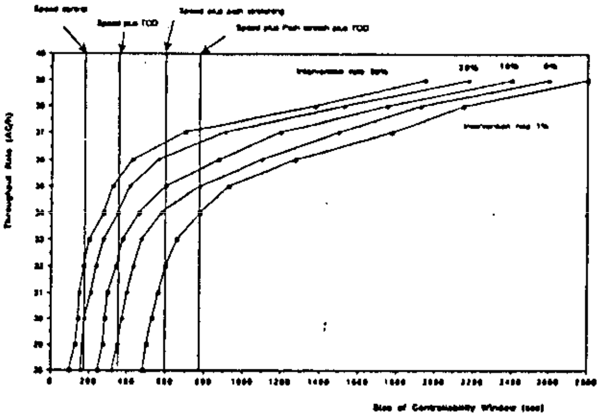

Figure 7. Throughput versus controllability for Poisson Model.

5.2 Results for the Minimum Miles-in-trail Model

This model represents an attempt to refine the Poisson Model by including the existence of minimum in-trail constraints (here taken to be 108 sec) in individual streams. At the same time new parameters, namely the fix loading and the number of arrival streams (or the number of entry fixes), are introduced. To simplify the discussion we assume three fixes and show only two cases: first when the fix loadings are equal, and next when they are 70%, 20% and 10%. The curves for Poisson Delivery Model will be used as reference.

Figure 8 shows that in the three stream and equal fix loading case the effect of minimum in-trail separations has no beneficial effect on throughput. This implies that even with only three feeder streams, the merged stream becomes quite random, differing little from the Poisson Model.

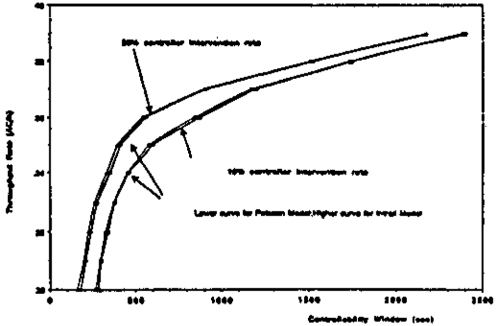

Figure 8. Throughput versus controllability for in-trail model with equal fix loadings.

Figure 9 shows a situation with significantly unbalanced fix loadings. As expected, some of the minimum in-trail spacing effect survives and greatly helps the TATCA scheduler, especially when the available controllability is small. For example for speed control only (180 s) and a 10% allowable intervention rate the throughput rate is 30 (versus 24 for Poisson). For a 600 s window it is 35.7 (up from 35), showing that the beneficial effect lessens when the knee in the curves is reached. The minimum miles-in-trail separation succeeds as a means of smoothing a Poisson arrival stream only if the bulk of the traffic arrives in one single stream or for several unbalanced streams.

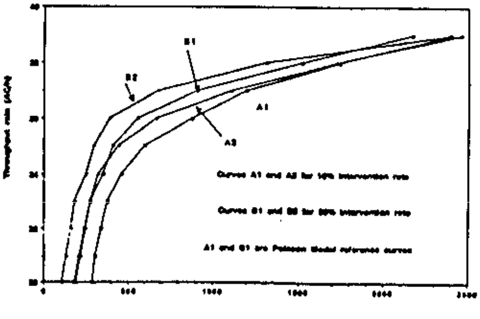

Figure 9. Throughput versus controllability for in-trail model with unequal fix loadings

5.3 Results for the Metered Model

Recall that for this model the TATCA scheduler is handed an arrival stream in which the arrival times have been adjusted across all entry fixes by an en route metering program in

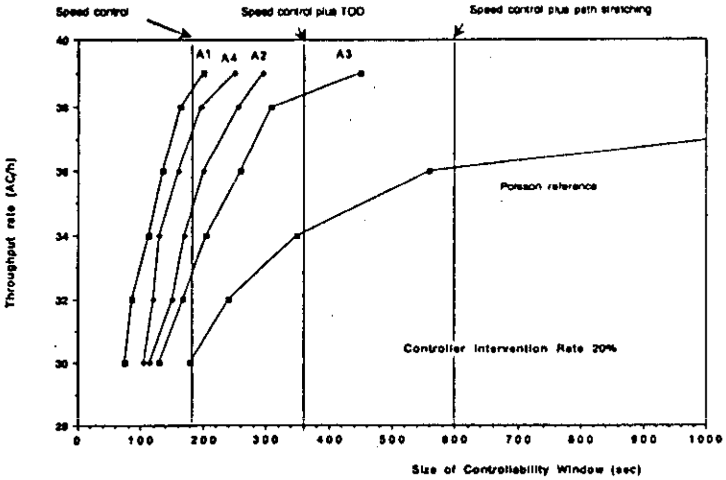

Figure 10. Throughput with Metering Model and three fixes with equal loading. Curve A1 is for error N(0,100) i.e. .normally distributed with 0 bias and standard deviation 100 sec. Curve A2 is for N(0,200); Curve A3 for N(0,300) and A4 for N(300,100). Average minimum landing time interval was 90 sec.

an attempt to achieve an ideal computed arrival time at the runway. The precision with which aircraft are actually delivered is diluted by normally distributed errors with specified mean b (for bias) and standard deviation S, denoted N(b,S)

Figure 10 shows that the magnitude of the metering precision error is hardly relevant if the amount of controllability available is of the same order as the spread of the delivery error. The main contribution of the metering program is that it has substantially reduced the potential for bunching in the arrival stream, thus greatly increasing the likelihood that the delay required by the scheduler will be within the controllability window. A relatively small controllability window consisting of speed control and some path stretching within the TRACON boundary allows TATCA to operate at very high throughput rates (39 out of a capacity of 40 or .975 runway utilization rate). Even when a particular stream is subjected to a large bias error (perhaps due to a wind error affecting one path) no large delays seem to ensue, as is shown by the curve labelled A4.

The Metering Model can be used to determine the effect of metering precision upon the achievable throughput. The system attempts to deliver aircraft at the metering fix at specific times and fails to do so because of errors.[2] Without errors, no controllability at all would be required. The system could then accommodate any arrival rate short of the actual capacity of the runway and the exception rate would be zero. The presence of metering errors leads to a requirement for a non-zero controllability window and a finite exception rate.

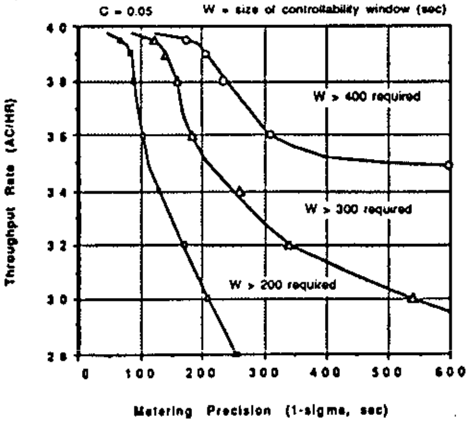

Figure 11. Effect of metering precision on throughput (for C=0.05).

---

[2] The same mathematical formulation can apply to errors caused by lack of coordination between the terminal requirements and the targeted times of the metering system.

Figure 11 shows how the achievable throughput is dependent upon the metering precision. In this figure the runway capacity is 40 aircraft/hour and the intervention rate, C, is allowed to be 5 percent. The steepness of the W<200 boundary implies that precise metering is of most value when the controllability is low. When the controllability is more than 400 seconds, the runway can be quite efficiently used even with very imprecise metering. As the metering error increases, the throughput deteriorates to the value that could be handled under the Poisson Model. For a TATCA system that employs both speed control and path stretching, the required metering precision appears to be within the range that is commonly achieved today where En Route Metering is in effect. The metering specifications called for in the System Level Specification for the Advanced Automation System requires precision of one minute and this is more than satisfactory for TATCA operation.

6. Summary of Principal Results

This analysis has established quantitative relationships between throughput rate, controllability which is a function of the planning horizon, controller intervention rate, randomness in the arrival pattern, fix loadings, scheduling discipline and capacity. A discrete time simulation model, incorporating only the salient features relevant to the throughput versus delay were included at this stage. Table 1 gives a short summary of the principal results.

If TATCA is handed an arrival stream whose average rate is reasonably controlled (e.g., by Flow Control), but where individual aircraft arrivals are fully independent from one another (Poisson) then the amount of controllability typically available within the TRACON boundary (180 s for speed control, and 600 s for speed and path stretching) will force TATCA to operate at lower throughput rates if only low controller intervention rates (5% or 10%) are allowed; or it can push up the throughput rate but then higher controller intervention rates are required. However, if TATCA were allowed to specify time-at-entry fix for a small subset of arrivals (in fact imposing that large delays be absorbed in the enroute control area, maybe in the form of a hold), it could operate at higher throughput rates.

The In-trail Model where aircraft entering the TRACON must have a minimum intrail separation seemed to perform very much as did the Poisson Model.

But in the Metering Model, where the arrival stream has virtually no bunched arrivals, TATCA needs little controllability to operate at throughput rates close to capacity, and this high performance seems quite insensitive to the precision with which arrivals are delivered to the metering fix.

7. Conclusions

From the analysis in this paper, several conclusions can be drawn. These conclusions are valid for the mathematical models employed, and their relevance to particular terminal areas should be validated on a case-by-case basis. The quantitative conclusions stated below assume that an intervention rate of C=0.05 is selected as a system operating point.

For a metered flow, the value of additional controllability decreases significantly beyond 200-300 seconds. For a Poisson model, the value of additional controllability decreases beyond about 600 seconds. This is seen in the "knee" in the curves of Figures 8, 9, and 10. One consequence of this is that when metering is in effect, a TATCA system that operates solely within the TRACON can obtain a significant portion of the total potential benefits.

The value of metering is confirmed by the throughput increases obtained when going from the Poisson arrival model to the metered model. With an unaided free-flow of arrivals as modeled by Poisson, the terminal area cannot handle a flow rate that is much greater than 75% of the actual runway capacity. The imposition of minimum in-trail separation restrictions, while helpful in preventing the terminal area from being overloaded, does not improve the throughput beyond that of the Poisson case. This is because the merging of streams with miles-in-trail restrictions often results in a combined stream that is essentially Poisson.

Table 1. Throughput versus controllability versus intervention rate for three arrival models.

Throughput (AC/h) Max 40				Degrees of Randomness in Arrival Process					
				Poisson	Minimum Miles-in-trail* Skewed Load**	Metered Input*			
						bias 0			bias 300 s (1)
						St.D 100	200	300	100
Controllability (Sec)	180	Intervention (%)	5	—	—	32.4	28.0	26.0	30.6
			10	<24	30.0	36.9	31.2	29.6	34.3
			20	30.0	33.0	38.6	31.5	29.6	37.5
	360	Intervention	5	29.5	33.6	39	35.1	32.9	39
			10	32.6	34.4	+	38.3	35.2	+
			20	34.2	35.6	+	39	36.8	+
	600	Intervention	5	34.2	34.9	+	+	38.3	+
			10	35.0	35.7	+	+	38.6	+
			20	36.1	36.7	+	+	39	+
	780	Intervention	5	35.1	35.6	+	+	+	+
			10	35.7	36.2	+	+	+	+
			20	36.7	37.2	+	+	+	+

* In-trail separation 108 s
** Fix loading 70, 20, 10%

(1) Bias on one of 3 paths; balanced fix loading.
\+ More than 39 Arrivals/h

The use of speed control alone provides controllability windows on the order of 60 to 230 seconds, depending upon the approach path. Unfortunately, the achievable throughput tends to be constrained by the paths with the least controllability. Hence, a TATCA system that employed speed control alone within the terminal area would probably require an extremely precise metering system in order to avoid throughput loss. Such metering may be available in the AAS time frame, but is not available in the near-term environments.

The use of speed control and path stretching within the TRACON provides controllability windows on the order of 300-400 seconds. This results in a substantial easing of the requirements for metering precision. Metering errors of 200-300 seconds (one-sigma) then become acceptable (see Figure 10). Such a system could function well with existing metering systems, yielding throughputs up to 95% of the runway capacity.

In en route airspace, in-trail aircraft may be only 80-120 seconds apart. The tolerance of metering errors on the order of 300 seconds implies that the precise sequence of delivery and separation at the metering fix is not significant as long as arrival rate (as averaged over a few minutes) is well controlled.

The analysis indicates that efficient use of the controllability available within a scheduling horizon equal to the terminal area itself can lead to significant terminal throughput benefits, even with imperfect metering.

References

[1] Spencer, D., Welch, J., and Andrews, J., "An Experimental Examination of the Benefits of Improved Terminal Air Traffic Control Planning and Scheduling", The Lincoln Laboratory Journal, Vol.2, No. 3, Fall 1989, p. 527-536.

[2] Boswell, S., Andrews, J., and Welch, J. "Analysis of the Potential Benefits of Terminal Air Traffic Control Automation (TATCA)". Proceedings of the 1990 American Control Conference, San Diego, CA, 23-25 May 1990, pp. 535-542.

[3] Erzberger, H. and Nedell, W., "Design of Automated System for Management of Arrival Traffic", NASA Technical Memorandum 102201, Ames Research Center, June 1989.

[4] Voelckers, U. and Schenk, H.,"Operational Experience with the Computer Oriented Metering Planning and Advisory System (COMPAS) at Frankfurt, Germany", AIAA Guidance, Navigation and Control Conference, August 1989, Boston.

[5] Credeur, L. and Capron, W., "Simulation Evaluation of TIMER, a Time-based Terminal Air Traffic Flow-Management Concept", NASA Technical Paper 2870, Langley Research Center, February 1989.

[6] Magill, S., Tyler, A.C.F., Wilkinson, E.T., and Finch, W., "Computer-based Tools for Assisting Air Traffic Controllers with Arrivals Flow Management", RSRE Report No. 88001, Royal Signals and Radar Establishment, September 1988.

[7] McLaughlin, Michael P., "Transition/Terminal Air Traffic Control Automation - An Analysis of Two Planning Aids", MTR-88W87, The MITRE Corporation, June 1989.

[8] Sokkappa, Balraj G., "The Impact of Metering Methods on Airport Throughput", Journal of Air Traffic Control, Vol. 31, No. 3, July-Sept 1989.

Herman F. Vandevenne
Technical Staff
Lincoln Laboratory
Massachusetts Institute of Technology

Herman F. Vandevenne is a member of the technical staff of the Air Traffic Automation Group at MIT Lincoln Laboratory where he joined the Terminal Air Traffic Control Automation (TATCA) program in 1989. His prior work at Lincoln Laboratory was in the analysis and design of advanced tracking systems, digital communications, and signal intercept systems. He was a member of the original Mode S design team.

Dr. Vandevenne received the Master's degree in Electronics and the Master's degree in Computer Sciences from the University of Louvain in Belgium. He holds the M.S. and Ph.D. in Control and Decision Engineering from MIT. He was also a NATO Post-doctoral Fellow.

John W. Andrews
Assistant Group Leader
Lincoln Laboratory
Massachusetts Institute of Technology

John W. Andrews is Assistant Group Leader in the Air Traffic Automation Group at Lincoln Laboratory. He is responsible for technical management of the Terminal Air Traffic Control Automation (TATCA) Program. For almost 14 years, Mr. Andrews was involved in the development of the Traffic Alert and Collision Avoidance System (TCAS). He 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 also served as a consultant to the National Transportation Traffic Board in the investigation of mid-air collisions.

Mr. Andrews holds a B.S. in Physics from the Georgia Institute of Technology and a Master's degree in Aeronautical Engineering from the Massachusetts Institute of Technology.