

AD-A265 323



93-12379



This document is disseminated under the sponsorship of the Department of Transportation in the interest of information exchange. The United States Government assumes no liability for its contents or use thereof.

1. Report No. ATC-192	2. Government Accession No. DOT/FAA/RD-92/28	3. Recipient's Catalog No.	
4. Title and Subtitle Evaluation of the Capacity and Delay Benefits of Terminal Air Traffic Control Automation		5. Report Date 14 April 1993	
		6. Performing Organization Code	
7. Author(s) Steven B. Boswell		8. Performing Organization Report No. ATC-192	
9. Performing Organization Name and Address Lincoln Laboratory, MIT P.O. Box 73 Lexington, MA 02173-9106		10. Work Unit No. (TRAIS)	
		11. Contract or Grant No. DTFA-01-88-Z-02009	
12. Sponsoring Agency Name and Address Department of Transportation Federal Aviation Administration Systems Research and Development Service Washington, DC 20591		13. Type of Report and Period Covered Project Report	
		14. Sponsoring Agency Code	
15. Supplementary Notes This report is based on studies performed at Lincoln Laboratory, a center for research operated by Massachusetts Institute of Technology under Air Force Contract F19628-90-C-0002.			
16. Abstract This report reviews the benefits that the CTAS component of the FAA Terminal Air Traffic Control Automation program (TATCA) offers to aviation users. In particular, the report evaluates the prospects that exist for increasing arrival capacity during Instrument Meteorological Conditions (IMC) by introducing CTAS functionality into current operations. The impact of anticipated capacity gains on air traffic delays is analyzed. Savings in delay are translated into dollar savings using FAA statistics on the fleet-weighted direct cost of delay to domestic air carriers. Also, the value of passenger time is considered. Economic impacts are estimated and reported on an annualized, nationwide basis. Adopting FAA projections of future traffic growth, estimates of delay and attendant cost savings to air carriers and their passengers are provided for fiscal years 1995-2015. Taking the nominal estimate of a 12% gain in IMC arrival capacity, a nationwide implementation of CTAS would be estimated to save an average of 412,000 hours of air carrier delay annually over this 21-year period, and 273 million gallons of fuel per year. With current fuel and labor costs, this amounts to average direct operating savings to air carriers of \$1.5 billion per year, and value to passengers of over \$3 billion per year, in constant 1988 dollars. There may be factors outside the scope of this study that restrict the implementation of CTAS to certain sites, or that limit the weather conditions in which CTAS is effective. Methods are discussed in the report for modifying benefits estimates in response to such considerations. However, since development and implementation costs of CTAS are estimated to be a small fraction of the benefits enumerated above, and since the delay savings recur annually, it is concluded that the development of ATC automation software such as CTAS is economically justifiable.			
17. Key Words air traffic control TATCA automation air travel delays CTAS airport capacity benefit/cost analysis FAST/TMA		18. Distribution Statement This document is available to the public through the National Technical Information Service, Springfield, VA 22161.	
19. Security Classif. (of this report) Unclassified	20. Security Classif. (of this page) Unclassified	21. No. of Pages 74	22. Price

EXECUTIVE SUMMARY

This report reviews the benefits that the CTAS component of the FAA Terminal Air Traffic Control Automation program (TATCA) offers to aviation users. In particular, the report evaluates the prospects that exist for increasing arrival capacity during Instrument Meteorological Conditions (IMC) by introducing CTAS functionality into current operations. The impact of anticipated capacity gains on air traffic delays is analyzed. Savings in delay are translated into dollar savings using FAA statistics on the fleet-weighted direct cost of delay to domestic air carriers. Also, the value of passenger time is considered. Economic impacts are estimated and reported on an annualized, nationwide basis.

Using data on historical trends in traffic volume and in air carrier delay, a quantitative and empirically supported description is developed of the relationship between capacity, demand, and traffic delay in the national airspace system. A simple formula is derived that predicts annual nationwide air carrier delays, in thousands of hours, if estimates are provided of annual traffic volume. A companion formula predicts the delay savings that would accrue from a given percentage increase in arrival capacity.

The report includes a detailed analysis of high-density IMC arrival traffic at Boston Logan airport. The factors that affect interarrival spacing are collected into two components, each reflecting a fundamentally different aspect of the task of approach control. One component, called final approach precision, is concerned with the positioning of aircraft for localizer intercept and subsequent fine tuning on the final approach course. The second component corresponds to flow management and preliminary positioning of traffic upstream of the final staging area. Relying upon FAA field experience with previous automation efforts, and upon analytical studies and laboratory simulations reported by major ATC research organizations, an assessment is made of the level of capacity enhancement that appears to be achievable with terminal area automation. It is estimated that average systemwide IMC capacity increases from CTAS will be between 8% and 16%, with 12% representing the best working estimate.

Adopting FAA projections of future traffic growth, estimates of delay and attendant cost savings to air carriers and their passengers are provided for fiscal years 1995-2015. Taking the nominal estimate of a 12% gain in arrival capacity, a nationwide implementation of CTAS would be estimated to save an average of 412,000 hours of air carrier delay annually over this 21 year period, and 273 million gallons of fuel per year. With current fuel and labor costs, this amounts to average direct operating savings to air carriers of \$1.5 billion per year, and value to passengers of over \$3 billion per year, in constant 1988 dollars. There may be factors outside the scope of this study that restrict the implementation of CTAS to certain sites, or that limit the weather conditions in which CTAS is effective. Methods are discussed in the report for modifying benefits estimates in response to such considerations. However, since development and implementation costs of CTAS are estimated to be a small fraction of the benefits enumerated above, and since the delay savings recur annually, it appears that the development of ATC automation software such as CTAS is economically justifiable.

CONTENTS

Executive Summary	iii
List of Illustrations	vii
List of Tables	ix
 1. INTRODUCTION	 1
2. DELAY REDUCTION	5
2.1 Introduction	5
2.2 Quadratic Model of NAS Delay History	5
2.3 Quantification of Delay Savings Expected from Capacity Increases	13
2.4 Economic Value of Capacity Gains and Resulting Delay Reductions	16
 3. CAPACITY ESTIMATION AND ENHANCEMENT	 21
3.1 Introduction	21
3.2 Baseline Performance of Terminal Air Traffic Control	24
3.2.1 Presentation of Field Data. Single Arrival Runway, IMC, at Boston Logan Airport	24
3.2.2 Models of Interarrival Spacing and Capacity	30
3.2.2.1 Three-Segment Model	31
3.2.2.2 Visual Separation Point	34
3.2.2.3 Final Approach vs. Upstream Sources of Variability	37
3.2.3 Summary of Baseline System Performance Measures	44
3.3 Performance Improvements	45
3.3.1 Improved Final Approach Precision and Flow Metering	45
3.3.2 Optimal Sequencing	50
3.3.3 Summary of Capacity Improvements	52
 4. SUMMARY AND CONCLUSIONS	 55
 Appendix 1. Notes on ATC Automation Studies	 57
Glossary	59
References	61

DTIC QUALITY INSPECTED 2

Accession For	
NTIS CRA&I	☒
DTIC TAB	☐
Unannounced	☐
Justification	
By	
Distribution /	
Availability Codes	
Dist	Avail and/or Special
A-1	

LIST OF ILLUSTRATIONS

Figure 1	Illustration of the "Ramp Model" for delay, as an approximation of departure delays issued by Central Flow to Boston-bound traffic.....	8
Figure 2	Plot of quadratic fits to SDRS total annual delay history	13
Figure 3	Operating curves relating air traffic delay to traffic volume. Comparison of current ATC system delays to reduced delays that would follow capacity increases achieved by CTAS.....	16
Figure 4	Illustration of final approach spacing precision when incoming aircraft positions are such that τ^0 is achievable.	22
Figure 5	("Closing" Aircraft Pairs) IAT's at Boston, 9/13/87, 10/11/87, 12/15/87, 5/18/88.	27
Figure 6	("Same" Aircraft Pairs) IAT's at Boston, 9/13/87, 10/11/87, 12/15/87, 5/18/88.	28
Figure 7	("Opening" Aircraft Pairs) IAT's at Boston, 9/13/87, 10/11/87, 12/15/87, 5/18/88.	28
Figure 8	Schematic overview of the three-segment model with linearly varying wind.....	32
Figure 9	Closure between successive aircraft on final approach.	35
Figure 10	Histograms of centered interarrival times, 4 episodes of IMC traffic at Boston.....	38
Figure 11	Fits of gamma mixtures to Boston interarrival time data.....	42

LIST OF TABLES

Table 1	SDRS Delay History	11
Table 2	SDRS Delay by Phase of Flight, Pre-strike vs. Post-strike	12
Table 3	Coefficients of Quadratic Fits to SDRS Total Annual Delay History	12
Table 4	Average Impacts of a Delay Hour to Air Carriers and Their Passengers	17
Table 5	Reference System Projections and CTAS Savings	18
Table 6	CTAS Delay Savings to Air Carriers and Passengers	19
Table 7	Sampled IMC Arrival Operations at Boston	24
Table 8	Aircraft Categories for Capacity Analysis, and Mix at Boston	26
Table 9	Descriptive Statistics of Interarrival Times, Boston IMC Data	26
Table 10	Airspeed Parameters, Evaluation of Three-segment Model	33
Table 11	Windfield Parameters, Evaluation of Three-segment Model	33
Table 12	Parameters of Two-component Gamma Mixtures, Fitted to IMC Interarrival Data at Boston	43
Table 13	Characteristics of the Baseline System	45
Table 14	Survey of ATC Automation Research	46
Table 15	Observed IMC Arrival Sequencing at Boston Logan	52
Table 16	Summary of Capacity Improvements	53

1. INTRODUCTION

Terminal Air Traffic Control Automation (TATCA) is a research and development program of the FAA intended to provide computer-aided sequencing, spacing, and management of air traffic flows in terminal areas. Recent investigations in the United States ([Credeur and Capron [6], Erzberger and Nedell [8], Simpson, *et al.* [26], etc.) and abroad (Schick, Schubert and Voelckers [25]) have indicated that automation, operating within the essential constraints of the current ATC environment, will be able to increase the number of operations that terminal facilities can sustain, particularly arrival operations, and particularly in instrument meteorological conditions (IMC). The purpose of this report is to quantify the benefits to airspace users of the CTAS component of TATCA, which is described below. Benefits are assessed by characterizing the effects of CTAS on system capacity and estimating likely levels of capacity increase, then translating capacity increases into user delay savings, and assigning dollar values to those delay savings.

Delay reductions generate attendant cost savings to airspace users, and the cost savings are aggregated here in three categories: fuel costs, other direct operating costs (ODOC) of operators, and the value of passenger time. The estimation of benefits is restricted to scheduled air carriers and their passengers, since they comprise the largest user group in controlled airspace, and since information on operating costs and delay distributions is more readily available for air carriers than for other user groups. Economic impacts are estimated and reported on an annualized, nationwide basis.

Increases in system capacity convey benefit to users in the form of reduced air traffic delays. It is also possible for capacity gains to yield benefit in the form of increased economic activity. In this report we make the assumption that capacity-generated benefits are exploited solely as reductions in delay. To the extent that the user community elects to trade potential delay reductions for increased traffic, one may consider that they do so because they perceive an additional benefit thereby, over and above the nature of the benefits described here.

Nature of the TATCA Program

As it is currently organized, the TATCA program includes two distinct development efforts, the Converging Runway Display Aid (CRDA), and the Center-TRACON Automation System (CTAS). The CRDA provides an on-screen display to help controllers judge the relative spacing between aircraft on separate ILS approaches. The CRDA enables ATC to extend the capacity advantages of converging runway operation to a broader range of weather conditions. Configuration managed CRDA software has been incorporated in the latest release (version A3.05) of ARTS IIIA software. The benefits of CRDA have been evaluated during field development and elsewhere (*e.g.*, Branscome [2]), and it will not be reviewed in this report.

The Center-TRACON Automation System, or CTAS, is an integrated set of capabilities for planning and controlling airside traffic movements. The CTAS concept originated primarily at NASA Ames Research Center, which, with the assistance of a System Development Team appointed by the FAA, has also developed prototype software to support field assessment. CTAS includes a scheduling function that generates a coherent set of timing objectives for all aircraft in its purview. CTAS also provides radar controllers

with visualization aids and clearance advisories to assist them in meeting the planned objectives, or in revising the objectives if deemed necessary. CTAS is comprised of three major components, the Traffic Management Advisor (TMA), Descent Advisor (DA), and Final Approach Spacing Tool (FAST).

The TMA is the executive component of CTAS. It creates a sequence for runway access and devises a movement schedule for arrival aircraft that, when necessary, provides for delay absorption. In devising the schedule, the TMA invokes the Descent Advisor, which models aircraft descent profiles in the prevailing wind field, and optionally provides recommended descent clearances to help aircraft meet the TMA schedule with minimum fuel consumption. FAST is a set of functions designed to improve the accuracy of spacing judgments made by controllers in positioning aircraft for final approach, and functions to help devise and implement tactics for making spacing adjustments. The CTAS software monitors surveillance reports, winds, and controller inputs, and updates its plans and displays continually to adapt to ongoing circumstances in the TRACON.

Briefly, CTAS has the ability to affect capacity in two ways. First, because it can compute an efficient schedule with good lead time, the TMA will be able to reduce the prevalence of gaps that otherwise develop in arrival traffic as it is organized into a final approach stream. Second, with the increased spacing precision of FAST, controllers will be able to deliver final spacings that are more consistent and more reliable, and thus they can deliver tighter spacings that, on average, are nearer the minimum required for safe separation.

Structure of the report:

Delay impacts attributable to anticipated capacity increases are discussed in Section 2. Section 2.2 examines the fundamental relationship between traffic demand, capacity, and user delays. Using historical delay trends, it develops a simple equation that can be used to forecast future air traffic delays as a function of traffic volume, assuming that the capabilities of air traffic control remain essentially the same as they are today. In Section 2.3 a companion expression is developed that enables one to modify delay forecasts, and therefore to estimate delay savings, given an assumed change in system capacity. The material in these subsections is independent of TATCA program specifics, and it can be used by anyone concerned with assessing the technical or economic impacts of a capacity initiative (or capacity shortfall) in the national air traffic system. In Section 2.4 the expressions for delay savings derived in Section 2.3 are evaluated with respect to CTAS and the TATCA program, and unit delay costs are applied to translate CTAS-derived delay savings into monetary terms.

Prospects for capacity enhancement by CTAS are discussed in detail in Section 3. It is concluded there that a TRACON-only implementation of CTAS offers an average capacity increase of about 12% in instrument conditions; also, low- and high-end estimates are given, amounting to 8% and 16%, respectively. Individuals willing to accept these estimates and uninterested in the technical rationale for the capacity increases will find it unnecessary to read Section 3. Section 3 will be relevant to those interested in evaluating the technical design or the feasibility of terminal automation.

Ground rules of the report

- Conventional equipage and procedures

It will be assumed that CTAS operates within the framework of existing ATC equipment and practices. Conventional aircraft equipage is assumed. Existing separation standards, procedures, airspace structure, nominal routings, and traffic management policies, are all assumed to remain as is, unless otherwise stated, in evaluating TATCA program benefits.

- TRACON-only implementation

For purposes of this report the CTAS system will be assumed to control aircraft only within the bounds of current TRACON airspace. The CTAS system itself imposes no such restriction. To the contrary, the origin of the CTAS concept and the design of the system place substantial priority on en route operations. The Descent Advisor (DA), for example, is designed to compute efficient descent profiles from cruise altitude and to advise controllers of a desirable top of descent point (TOD). The TMA and DA are also designed to help organize and space traffic from the TOD to the TRACON, as well as within the TRACON, all the while attempting to minimize fuel burn and workload and to maximize landing capacity. In future implementations the operational horizon of CTAS may indeed extend far enough to support selection of the preferred top of descent point, or to anticipate arrival delays and absorb them efficiently at cruise altitude. Near term implementations, however, must accommodate the divisions of airspace, infrastructure, and control responsibility that currently exist between en route and terminal areas. The exact form that near term or future CTAS implementations will take is still unsettled, and any decision regarding implementation issues is beyond the scope of this report. However, it is desired that the report be valid regardless of the implementation decisions that are made in the future. Therefore, whenever a distinction is necessary, the report considers the most restrictive CTAS implementation, namely one confined to operate within TRACON airspace, with the understanding that more complete realizations of the CTAS concept may offer additional benefits, beyond those stated in the report. Additional user benefits would include modest impacts on arrival capacity, and increased fuel efficiencies. A previous paper (Boswell, *et al.* [1]) briefly evaluated the role that terminal automation can play in supporting fuel efficient arrival operations, apart from the fuel savings obtained by capacity gain and delay reduction. It was concluded that, at current fuel prices, the potential economic benefits accruing from such flight efficiencies are about 25% as large as the benefits that are likely to derive from capacity gains. The TRACON-only assessments presented here may be considered to establish lower bounds for the future benefits of CTAS functionality.

- Balance of traffic and weather conditions in the national airspace system

The expectation is that CTAS, when fully implemented, will have charge of a substantial majority of ATC-supervised arrival and departure operations in the national airspace system (NAS). CTAS will be expected to remain in operation in all meteorological environments, in all runway configurations, and during configuration changes. For this reason, the analyses in this report attempt to weight their evaluations according to the time-averaged mix of weather, traffic density, and other operational conditions occurring in the

NAS today. Occasionally this weighting is performed in an explicit manner; more commonly it is performed by making use of annualized statistics on quantities such as traffic volume or delay distributions, where these statistics implicitly incorporate and average the impact of environmental influences that vary with time.

- Discounting of VMC capacity gains

The ability to coordinate traffic flows, as well as to alleviate controller workload, suggests that terminal automation can increase airside capacity even when visual approaches are permitted. However, the increase will be more modest in VMC than in IMC. Further, the peak throughput at some terminals may be constrained by factors external to TATCA, such as landside access, gate availability, or center volume restrictions. To avoid making claims that exceed such external limits, the report will disregard benefits that might accrue during visual meteorological conditions at the destination airport. Also, the capacity modelling done to evaluate IMC capacity enhancements will be based on data taken during IMC operations.

2. DELAY REDUCTION

2.1 Introduction.

In most transportation problems it is self-evident that user delays will decline if system capacity is increased. However, in a complex environment like the national airspace system (NAS), the exact size of delay reductions that would follow a capacity initiative is not obvious, nor, to some extent, is the intrinsic value of delay savings. In this section of the report a methodology is presented for relating capacity increases and delay savings on a nationwide basis, and for assigning economic value to the delay savings.

The fundamental relation among capacity, traffic demand and delay is discussed in Section 2.2. It is suggested that, to a first approximation, individual episodes of congestion exhibit a simple pattern of linear rise and linear fall with respect to delay per aircraft. A consequence of this pattern is that overall delay in the National Airspace System (NAS) can be described as a quadratic function of traffic volume. Accordingly, the report describes the fitting of a quadratic function to annual statistics of the Standardized Delay Reporting System (SDRS), published by the FAA Office of Aviation Policy and Plans, 1976-1986. The resulting quadratic may be used in conjunction with traffic forecasts (*e.g.*, FAA [10-12]) to estimate air carrier delays in current and future years, assuming the continuation of current air traffic control technologies and procedures.

In Section 2.3 it is shown that adjustable parameters in the quadratic relationship between delay and demand can be adapted to assess the impact of capacity changes in the air traffic control system. A simple formula is derived that predicts nationwide annual delay savings, in thousands of hours, if estimates are provided of annual traffic volume and of the magnitude of proportional capacity increases. Adjustments are briefly discussed to modify the formula if capacity increases are restricted to certain meteorological conditions, or if partial rather than nationwide deployments of automation software are being considered.

Later in this report consideration is given to determining the magnitude of capacity increases that terminal automation can reasonably provide. Using estimates obtained from the capacity modeling, in Section 2.4 the delay savings that are anticipated from CTAS are calculated for the years 1995-2015. These years form the planning horizon used for some recent internal evaluations of TATCA. Then, using FAA figures on the direct costs to air carriers of operating delay, as well as the value of passenger time, the potential delay savings from terminal automation are converted into monetary terms.

2.2 Quadratic Model of NAS Delay History

From a national perspective, capacity, demand, and delay are dynamic, and they interact in a complex network of airways and control facilities. When delay arises, it is redistributed through the system by various mechanisms, including the national and regional traffic management programs of the FAA. The Standardized Delay Reporting System (SDRS), which was established by the FAA Office of Aviation Policy and Plans (APO) in 1976, has maintained an accounting of system delay broken down by phase of flight. Also, the APO has calculated fleet weighted delay costs separately for different phases of flight, and averaged again over all phases of flight following request for push

back from the gate (Geisinger [14]). Therefore, it is possible to estimate nationwide delay savings and assign appropriately weighted dollar values to the savings using only the SDRS annual total delay figures, which are available during the period 1976-1986. In the remainder of this section, we will present a model for use in summarizing how NAS delay accumulates, and how the level of delay in the system is likely to change as a function of system traffic levels.

The model development begins by considering delay dynamics around a single airport, and in doing so it takes a queuing point of view. The arrival runway, or the airport as a whole, is considered to be a service facility, serving one customer at a time. Aircraft wishing to land at the airport are the customers. An aircraft will be considered to enter the queuing system at the earliest moment when it would anticipate crossing the runway threshold, given its intended route and departure time, and considering winds and aircraft performance characteristics. Thus arrival into the queuing system is a conceptual quantity, not directly observable, and it is distinct from physical arrival at the airport runway. The difference between the conceptual arrival time and the observed physical arrival time is considered to be delay. Note that the delay need not occur in the vicinity of the destination airport. A large proportion of air traffic delays are taken on the ground at the origination airport, and less overt delays may occur because of route or speed restrictions en route, even when the root cause of the delay is capacity shortfall at the destination airport.

We will let $W^{(n)}$ denote the delay experienced by the n -th aircraft desiring service, following a first-come first-served sequence based upon the conceptual underlayed arrival time. Also boldface $W^{(n)}$ will be used to represent the total delay, summed over all arrivals from the first up to and including the n -th. Let $T^{(n)}$ represent the time elapsed between the n -th and $(n+1)$ -th (conceptual) arrivals in the queuing system. The service time for the n -th arrival, $S^{(n)}$, is the time elapsed between the n -th and $(n+1)$ -th arrivals' actual threshold crossings, that is to say the n -th (observable) interarrival interval. A prefix of "E" indicates the expected value for that quantity.

With the above notation a fundamental relation for any single server queue is

$$W^{(n+1)} = \max \{ W^{(n)} + S^{(n)} - T^{(n)}, 0 \}.$$

In a general setting, one must take proper account of occasions when there is no queue for service, and the customer experiences zero waiting time. However, if one restricts attention to an episode of congestion, during which a non-empty waiting line is maintained, then all aircraft arriving in that period experience a non-zero wait. Assume such an episode, and begin the count of aircraft just at the onset of congestion, so that $W^{(1)}$ is zero or negligible, but all subsequent $W^{(i)}$ are positive. Then, for any such episode,

$$W^{(n+1)} = \sum_{k=1}^n [W^{(k+1)} - W^{(k)}] = \sum_{k=1}^n [S^{(k)} - T^{(k)}].$$

The expected value of the waiting time of the $(n+1)$ -th aircraft is

$$EW^{(n+1)} = \sum_{k=1}^n E[S^{(k)} - T^{(k)}],$$

even if the expected service and interarrival rates are time- or state-dependent. The total waiting time accumulated by all arrivals through the $(n+1)$ -th arrival thus has expectation

$$EW^{(n+1)} = \sum_{k=1}^n (n+1-k) E[S^{(k)} - T^{(k)}] . \quad (2.2.1)$$

Any complete congestion episode must include intervals during which delay builds up (i.e., $S^{(k)} > T^{(k)}$), and intervals in which there is excess capacity to dissipate the backlog of demand (i.e., $S^{(k)} < T^{(k)}$). Perhaps the simplest complete congestion episode would consist of a single duration τ_1 during which delay increased linearly, followed by a period of constant rate delay decrease that lasts until the queue is again empty, or for a duration denoted as τ_2 . In fact, such a pattern occurs (in expectation) if, during the buildup phase the expected demand rate λ_1 remains constant, as does the average service rate μ_1 , and similarly, constant demand rates λ_2 and μ_2 apply during the recovery phase. Here λ_1 expresses the intentions of all scheduled and unscheduled flights wishing to land at the airport in the form of an average (desired) operation count per unit time, that is, $ET^{(k)} = 1/\lambda_1$. Also μ_1 measures the actual throughput rate ($ES^{(k)} = 1/\mu_1$), which is generally equivalent to the airport acceptance rate (AAR) used for traffic management. Similar statements hold for λ_2 and μ_2 . During the buildup phase, the expected waiting time experienced by a newly arriving aircraft increases linearly at a rate $(1/\mu_1 - 1/\lambda_1)$. During the recovery period the average delay decreases linearly at a rate $(1/\lambda_2 - 1/\mu_2)$. We shall refer to a delay pattern with the characteristics described in this paragraph as a "ramp pattern".

Though a ramp pattern holds true strictly only when one assumes two successive settings of otherwise constant airport demand and service rates, there are reasons to believe that it provides a reasonable approximation to most congestion events. First, runway configurations have a nominal AAR associated with them, and though actual throughput certainly varies with circumstances, deviations from the prevailing AAR are usually not extreme. Thus, over a short term, an airport will generally operate within a narrow band around one nominal throughput rate, or it will switch between two such rates via a configuration change. Second, at major, high volume U.S. airports, scheduled demand remains strong from early morning to early or mid evening. Any capacity shortfall that occurs during these hours will face a demand pattern that, integrated over an hour or so, is relatively uniform. Outside of these hours, major delay accumulations are infrequent.

Data depicted in Figure 1 illustrate the use of the ramp pattern as a summary of real sustained episodes of air traffic delay. The data correspond to four occasions of sustained IMC arrival operations during peak demand hours at Boston Logan airport, occurring on the dates indicated in the figure. On each occasion weather remained below circling minima and Boston approaches were confined to use of a single instrumented runway, leaving the airport with capacity well short of the scheduled demand. The TATCA program had observers on site with visual surveillance of the runway threshold. The observers collected detailed operational data which was reviewed for the purpose of capacity analysis (see Section 3, cf. Table 7). In addition the program collected worksheets from the FAA's Central Flow Computing Facility (CFCF), which, anticipating the capacity shortfall, issued mandatory departure delays to aircraft filing to fly to Boston. The issued departure delays are applied in 15-minute blocks and appear in the stair step histories in Figure 1. For example, on 9/13/87, aircraft with intentions to arrive at Boston between 5:45 pm and 6:00

pm EDT had their intended departure times delayed by 70 minutes, while those originally intending to arrive between 6:00 pm and 6:15 pm were delayed 80 minutes on the ground, and so on. Approximating ramp patterns are indicated by the inclined straight line segments. In all cases surveyed, a ramp pattern provides an acceptable summary of delay accumulation during the congestion episode.




Figure 1. Illustration of the "Ramp Model" for delay, as an approximation of departure delays issued by Central Flow to Boston-bound traffic.

Therefore, accepting the assumptions of the ramp pattern, let us consider for the moment only the buildup phase, during which demand exceeds capacity. Evaluation of (2.2.1) gives that the accumulated delay up through the $(n+1)$ -th aircraft has expectation

$$EW^{(n+1)} = \frac{n(n+1)}{2} \left(\frac{1}{\mu_1} - \frac{1}{\lambda_1} \right). \quad (2.2.2)$$

In air traffic it is more natural to measure or to predicate the duration of a congestion episode, rather than the number of aircraft involved in it. Even if the duration of the buildup phase is known, the actual number of aircraft that join the queue during this period will be a random variable, having mean $\lambda_1 \tau_1$, and an unspecified variance σ_1^2 . Indexing by duration, then, rather than by count of aircraft, the expected total delay developed over time τ_1 is

$$EW^{(\tau_1)} = \frac{\sigma_1^2 + (\lambda_1 \tau_1)^2 + \lambda_1 \tau_1}{2} \left(\frac{1}{\mu_1} - \frac{1}{\lambda_1} \right).$$

Flow regulation imposed by the ATC system, particularly when a terminal area is experiencing congestion, serves to restrict the variability of the interarrival time distribution, so that $\sigma_1^2 \ll \lambda_1 \tau_1$. Therefore, an adequate approximation for our purposes is that

$$\begin{aligned} EW(\tau_1) &= \frac{(\lambda_1 \tau_1)^2 + \lambda_1 \tau_1}{2} \left(\frac{1}{\mu_1} - \frac{1}{\lambda_1} \right) \\ &= 1/2 \left\{ \lambda_1^2 \frac{\tau_1^2}{\mu_1} + \lambda_1 \left(\frac{\tau_1}{\mu_1} - \tau_1^2 \right) - \tau_1 \right\}. \end{aligned} \quad (2.2.3)$$

In the recovery phase of the congestion episode there is a parallel result, except that the quantity $(1/\mu_1 - 1/\lambda_1)$ in (2.2.3) is replaced by the positive quantity $(1/\lambda_2 - 1/\mu_2)$ rather than the simple substitution $(1/\mu_2 - 1/\lambda_2)$. The expected total delay accumulated by aircraft involved in the complete congestion episode is thus approximately

$$EW(\tau_1 + \tau_2) = 1/2 \sum_{i=1}^2 (-1)^{i-1} \left\{ \lambda_i^2 \frac{\tau_i^2}{\mu_i} + \lambda_i \left(\frac{\tau_i}{\mu_i} - \tau_i^2 \right) - \tau_i \right\}. \quad (2.2.4)$$

To first order, delays occurring in the national airspace may be described as arising from a vast collection of congestion episodes, each generating a delay accumulation that is expressible in the form (2.2.4). Nationwide air traffic involves a large collection of such episodes, each requiring six quantities for full specification. It is clearly not feasible to specify all contributing congestion episodes along with their parameters λ_i , μ_i and τ_i . However, as will be seen below, expression (2.2.4) has embedded within it some simple and useful mathematical patterns. In particular, holding other quantities fixed, (2.2.4) is a sum of quadratics in λ_i . With a few additional assumptions it is possible to extend the quadratic pattern to annual and nationwide applicability, without knowing any of the parameters λ_i , μ_i or τ_i explicitly. Rather, the net effect of these parameters can be calculated indirectly by fitting to the historical record of traffic volume and air traffic delay.

To begin the process of extending to a national scale, suppose that a reference year is chosen to establish initial conditions, and that the nationwide traffic demand in that year, measured by number of domestic departures, is V_0 . In calculations below, 1982 is taken as the reference year, but the choice of year is immaterial. We will identify the reference year and today's air traffic control system with two parameters $(V_0, 1)$, whose meaning will become apparent below. Then, imagining all significant congestion events occurring in the NAS during this year, and summing all the corresponding individual delay contributions of form (2.2.4), the annual nationwide delay is expressible as:

$$EW(V_0, 1) = b_2 + b_1 + b_0, \quad (2.2.5)$$

where

$$b_2 = \frac{1}{2} \sum_i \lambda_i^2 \frac{\tau_i^2}{\mu_i} q_i ,$$

$$b_1 = \frac{1}{2} \sum_i \lambda_i \left(\frac{\tau_i}{\mu_i} - \tau_i^2 \right) q_i ,$$

and
$$b_0 = -\frac{1}{2} \sum_i \tau_i q_i ,$$

and the summation is now over all phases of all congestion episodes in the national airspace. In the above expressions, q_i is an indicator of sign for each term, that is, $q_i = \pm 1$, depending on whether delay is building up or dissipating during the term to which it is attached.

As stated above, it is impossible to know all the contributing parameter values, and thus we cannot expect to calculate the coefficients b_j directly. However, with additional assumptions it is possible to fit them to historical data. Let us assume³ that

- (i) the distribution of τ_i remains constant from year to year, that
- (ii) operational capacities (i.e., the collection of μ_i) have remained constant during the historical frame of reference, apart from a shift in traffic management policy originally occasioned by the 1981 controllers' strike, and that
- (iii) traffic increases have been distributed homogeneously; that is, if in the reference year national traffic volume is V_0 , and in the k -th outyear the traffic volume is $V_k = \alpha V_0$, then λ_i is assumed to have increased to $\alpha \lambda_i$ for all i .

With these conventions, expected delay in a year with traffic volume V will have the form

$$\begin{aligned} EW(V,1) &= \left(\frac{V}{V_0} \right)^2 b_2 + \left(\frac{V}{V_0} \right) b_1 + b_0 \\ &= V^2 c_2 + V c_1 + c_0 , \end{aligned} \tag{2.2.6}$$

³Each of these assumptions substitutes a simplifying premise for what in reality is a more diverse state of affairs. For example, assumption (iii) discounts the site-to-site variation in traffic growth that arises because of hubbing practices, the emergence of new economic centers (eg, Orlando), and other causes. The use of a single multiplicative growth factor at all sites understates traffic growth at some sites, and it will tend to understate delay growth at those sites. At other sites, traffic growth will be overstated, with a corresponding tendency to exaggerate delay growth. It is not necessary to accept assumption (iii) literally to make use of (2.2.6) and the development that follows. It is essential only that, to an accuracy sufficient for the purposes of the report, the overstatements and understatements comprised in (2.2.6) balance one another. Note that since delay is a convex function of demand level, any bias introduced by assumption (iii) and equation (2.2.6) will tend toward conservatism in representing delay growth (and in estimating delay-related benefits) in future years.

where $c_2 = b_2 / V_0^2$, $c_1 = b_1 / V_0$, and $c_0 = b_0$. In other words, with the assumptions listed above, the expression for annual nationwide delay, with its complicated summations for b_2 , b_1 , and b_0 , can be made to factor into a single parameter, V , which encapsulates all year-to-year change in expected delay, and constant coefficients, c_2 , c_1 , and c_0 , which do not change over time, or with traffic volume. Thus (2.2.6) has the form of a simple quadratic in V , and estimates may be obtained for the coefficients $\{c_j\}$ by fitting a quadratic to the delay trend data in Table 1. The fitted coefficients $\{c_j\}$ then implicitly incorporate the mix of operating circumstances (ie, the collection of individual terms comprising the summations in (2.2.5)) that is characteristic of the current airspace system. The data for 1986 were obtained from [14], while for previous years it was obtained from Figure D-7 in Butcher, *et al.* [3].

TABLE 1
SDRS Delay History

Year	Annual Departures (Millions)	Total Annual Delay (1000 Hrs)
1976	4.62	877
1977	4.79	929
1978	4.99	1011
1979	4.99	1047
1980	4.87	994
1981	4.62	965
1982	4.37	902
1983	4.75	1015
1984	4.90	1143
1985	5.18	1248
1986	5.27	1335

A least-squares fit of (2.2.6) to the data of Table 1 was conducted, subject to a couple of special considerations. One consideration is the increased use of FAA flow control procedures following the August 1981 controller's strike, and perhaps other less apparent changes in ATC service dynamics. As indicated in Table 2, obtained from Geisinger [14], the relative level of airborne delay decreased by almost 40% from 1976 to 1986, as flow control procedures shifted potential airborne delay to (certain) gate and taxi-out delays. This transfer of delay is considered to have safety and workload benefits, and it may reduce average unit delay costs, but, as is apparent in Figure 2 below, it occurs at the cost of an increase in total delay. In Figure 2, both the delay data and lines representing quadratic fits are plotted. Pre-strike data (1976-1980) are shown as 'x', data from 1983 onward are plotted with '+', and the years 1981 and 1982, which, because of their exceptional nature were excluded from the fitting process, are shown as open circles. For a comparable level of system traffic, total annual delays run about 10% higher after the strike. Therefore, while it is desirable to use as much of the available data as possible, provision must be made for the difference in pre- and post-strike levels.

TABLE 2
SDRS Delay by Phase of Flight, Pre-strike vs. Post-strike

Phase of Flight	Percentage of Total Delay		%Change, 1976-86
	1976	1986	
ATC Gate Hold	0.6	7.4	1,133
Taxi-Out	40.7	48.7	19.7
Airborne	39.0	23.7	-39.2
Taxi-In	19.7	20.2	2.5

A second consideration concerns the range of applicability of the modelling approach that leads to expression (2.2.6). The derivation of the expression hinges upon assumptions (*e.g.*, no changes in system capacity, 1983-1986) which we feel are close enough to reality to be useful. On the other hand, they are not exact. It is of little concern if the quadratic model becomes inaccurate at very low or extremely high traffic levels, but it should remain internally consistent in the current or anticipated operating range of the national airspace system. Given the very small number of data points in either the pre- or post-strike phase, we imposed a side condition specifying where the minimizer of both quadratics should occur, intending that the model should not be considered relevant near or below the corresponding traffic level, or, if the minimum is negative, near or below the greatest root of the quadratics. Minimizers at $V=0.54$, $V=0.66$ and $V=0.87$ were investigated, where traffic demand is expressed in millions of domestic departures. Note that the candidate minimizers represent 12.5%, 15%, and 20% of the 1982 traffic levels. The latter value, $V=0.87$, gave slightly smaller confidence intervals for the parameters, relative to their estimated values. Also, for the two smaller values of V , the minimum values were negative and the roots of the upper branch were higher than with $V=0.87$ (post-strike roots were at $V=1.66$ and $V=1.53$, respectively). Therefore, the quadratics with minimizer $V=0.87$ were chosen as both fitting the data slightly better and as being self-consistent over the widest range. The coefficients of the fit, giving annual delays in thousands of hours, are as given in Table 3:

TABLE 3
Coefficients of Quadratic Fits to SDRS Total Annual Delay History

	c_2	c_1	c_0
Pre-Strike	60.8	-106.4	56.0
Post-Strike	68.0	-119.0	56.0

Because of the side condition, values for c_1 are not separately estimable ($c_1 = -1.74 c_2$). Standard errors for c_2 were 5.1 pre-strike and 4.6 post-strike.

Results of the quadratic delay modelling are displayed in Figure 2, which displays annual air carrier delay as a function of annual traffic volume. Both the pre-strike and post-strike delay/demand curves are depicted. The post-strike curve gives expected annual delay in the current U.S. air traffic system as

$$EW(V,1) = 68 V^2 - 119 V + 56, \quad (2.2.7)$$

where again, V is millions of domestic annual departures, and delay is in thousands of hours. Expression (2.2.7) is proposed as a predictor of delay in the existing air traffic control system, and it will be taken as a baseline for the calculation of benefits that accrue because of CTAS capacity gains.




Figure 2. Plot of quadratic fits to SDRS total annual delay history.

Referring to Figure 2, it is worthwhile to note that a linear fit, or indeed many other fitting functions, would summarize the SDRS delay data well, and provide suitable near term predictors of air traffic delay. The quadratic is employed here because, as has been presented in this section, there is a generating explanation for delay accumulation that leads to the quadratic form. Also, as will be discussed in the next section, further examination of the quadratic coefficients leads to a formula for gauging the likely impact of changes in system capacity.

2.3 Quantification of Delay Savings Expected from Capacity Increases

In denoting capacity gains from CTAS we use the symbol β to indicate the supposition that the capacity of the existing baseline system is to be multiplied by a factor β , or equivalently, capacity is increased by a percentage $100*(\beta-1)$. The predicted annual delay at a given traffic level, under the assumption of the hypothesized increase in capacity, is then written $EW(V, \beta)$.

To investigate the effect of widespread capacity gains, we make a homogeneity assumption similar to that used for increases in traffic demand. Namely, if average IMC arrival capacity in the NAS is predicated to increase by a factor β , it is assumed that $\mu_i \rightarrow \beta\mu_i$ for all i in expression (2.2.4). With this assumption regarding capacity increases, expression (2.2.6) becomes

$$EW(V, \beta) = \beta^{-1} V^2 c_2 + V c_1(\beta) + c_0 , \quad (2.3.1)$$

where

$$c_1(\beta) = \frac{1}{2V_0} \sum_i \lambda_i \left(\frac{\tau_i}{\beta \mu_i} - \tau_i^2 \right) q_i . \quad (2.3.2)$$

Note that the homogeneity assumption carries an implication that capacity increases occur in all SDRS reporting sites, and in all delay-producing conditions. The implication is not appropriate in every case. However, the homogeneity assumption will be maintained until fundamental system delay/capacity relationships are established. Then, at the end of this section, we will estimate the prevalence of delay-producing conditions (*e.g.*, thunderstorms en route) that CTAS cannot address, and we will reduce estimates of delay savings accordingly.

Since β does not factor out of $c_1(\beta)$ immediately as it does in the quadratic term, some extra work is necessary to reduce $c_1(\beta)$ to a tractable form. A number of approaches to reducing $c_1(\beta)$ were considered. The approach adopted follows from the identity below, obtained by algebraic rearrangement of (2.3.2),

$$c_1(\beta) = \frac{1-\beta}{\beta} \frac{1}{V_0} \left(\sum_i \frac{\lambda_i}{\mu_i} \tau_i q_i \right) + c_1 . \quad (2.3.3)$$

The summation that appears in parentheses in (2.3.3) must be further simplified or approximated if the expression is to be useful quantitatively. It is helpful to recognize the summation as a weighted sum of terms $\tau_i q_i$, using positive-valued weights λ_i/μ_i , because the unweighted sum is related to the "current system" coefficient c_0 which appears in expression (2.2.6), and thus it already has an estimated numerical value:

$$\sum_i \tau_i q_i = -2 c_0 = -112 .$$

It is natural to attempt to express the weighted sum also as a multiple of c_0 , namely

$$\sum_i \frac{\lambda_i}{\mu_i} \tau_i q_i = k c_0 ,$$

for some value k . In such case (2.3.3) becomes

$$c_1(\beta) = \frac{1-\beta}{\beta} \frac{1}{V_0} k c_0 + c_1 ,$$

and (2.3.1) becomes

$$EW(V, \beta) = \beta^{-1} V^2 c_2 + V c_1 + \left(1 + k \frac{\beta-1}{\beta} \frac{V}{V_0} \right) c_0 . \quad (2.3.4)$$

It will be necessary to approximate k . In making an approximation, we note that forecast traffic volumes in the time period considered for this report lie between 6 million and 9 million departures per year, so that V has a value between 6 and 9. Thus, V^2 is at least 36 times greater than V/V_0 in all calculations to be conducted for this report. Also, for modest capacity increases, $(\beta-1)/\beta$ is an order of magnitude smaller than $1/\beta$. Therefore, k multiplies a quantity that is less than half a percent of the magnitude of the quadratic term in expression (2.3.4). Because k has such a limited influence on numerical calculation of the expected annual delay, we need only a rough estimate of its value.

Recall that τ_i denotes the duration of either the rising or falling phase of a congestion episode. It is always positive. Also, q_i is positive during the rising phase of congestion buildup, when, in addition, $\lambda_i > \mu_i$. Thus the contribution of positive terms $\tau_i q_i$ in the weighted sum is magnified by a weight greater than one, compared with their unit contribution to the unweighted sum. Similarly, during the falling phase, when $\tau_i q_i$ is negative, its contribution is reduced by the weight $\lambda_i/\mu_i < 1$. Therefore, the weighted sum must be less negative than the unweighted sum. Accordingly, $k > -2$. One could expect k to be roughly commensurate with the median value of the demand to capacity ratio, λ/μ , during episodes of congestion in the national airspace. We subjectively approximate k to have the value $k = -1.25$. While this approximation is admittedly rough, we note that, as indicated above, numerical calculations with 2.3.4 are relatively insensitive to the exact value supplied for k . With the assumption that $k = -1.25$, expected annual delay with a traffic volume V and proportional capacity increase β is predicted by the following expression:

$$EW(V, \beta) = \beta^{-1} V^2 c_2 + V c_1 + \left(1 + 1.25 \frac{\beta-1}{\beta} \frac{V}{V_0}\right) c_0. \quad (2.3.5)$$

By plugging in different values for β in (2.3.5) (e.g., $\beta=1$ to represent contemporary ATC and $\beta=1.12$ to represent a systematic 12% increase in arrival capacity) one may assess the impact of a capacity initiative on nationwide air carrier delays. A depiction of such impact calculations is given in Figure 3. The fitted post-strike delay curve from Section 2.2 is repeated in Figure 3, along with the delay curve predicted by (2.3.5) assuming a nominal increase of 12% in system capacity nationwide (dashed line), and also, as a form of sensitivity analysis, with capacity increases of 8% and 16% (dotted lines).

If V describes the traffic volume in a particular year, then a proportional increase in capacity of $(\beta-1)/\beta$ effected in that year is estimated to yield a total delay savings of

$$\Delta = EW(V, 1) - EW(V, \beta) = V \frac{\beta-1}{\beta} \left(V c_2 - 1.25 \frac{1}{V_0} c_0 \right). \quad (2.3.6)$$

If we ignore the c_0 term in (2.3.6), which as shown above is of secondary importance, then annual delay savings take the easily calculated form,

$$\Delta \approx 68 \frac{\beta-1}{\beta} V^2, \quad (2.3.7)$$

where V is the observed or forecast number of domestic departures in the target year.

As noted above, it must be recognized that some delay-producing conditions occur that CTAS is not in a position to address. Some are fundamentally independent of the terminal area. It is assumed here that CTAS will be able to address the root causes of 75% of all delay. By this we do not mean that it will shrink overall NAS delays by 75%. Rather, we mean that as CTAS increases capacity, 75% of delay events in the NAS will be somehow affected. Those delay events may be completely eliminated or only slightly reduced, depending upon the circumstances, but there will be a relationship between them and CTAS performance. A remaining 25% of NAS delays are assumed to result either directly or indirectly from causes, such as inherent center volume restrictions, that are outside the scope of CTAS. The number 75% was chosen on the basis of statistics assigning causes to delay events in the Air Traffic Operations Management System (ATOMS), which are reproduced in the 1990-91 Aviation System Capacity Plan [13]. In 1989 a total of 86% of delays exceeding 15 minutes were attributed by ATOMS to weather or to terminal volume. A remaining 8% was attributed to Center Volume, some of which is likely to have occurred as a secondary result of terminal area congestion. To allow for those weather events and volume restrictions which are beyond the influence of CTAS, we recommend that in calculating its delay impacts a 25% reduction should be applied to (2.3.6) and (2.3.7).




Figure 3. Operating curves relating air traffic delay to traffic volume. Comparison of current ATC system delays to reduced delays that would follow capacity increases achieved by CTAS.

2.4 Economic Value of Capacity Gains and Resulting Delay Reductions

The average systemwide costs of an hour of air carrier delay during 1986 were reported by Geisinger [14] as in Table 4. To assess the annual value of delay savings due

to a specified capacity increase, one may forecast the traffic volume in the year of interest, estimate the annual delay hours expected to be saved by (2.3.7), and multiply by the unit costs in Table 4. The Direct Operating Cost in the table includes fuel costs. Valued at \$0.75/gallon, fuel in 1986 accounted for approximately \$497 of the air carriers' DOC of an hour of delay. Other direct operating costs (ODOC) amounted to \$1057 in 1986 dollars. It is often desired in forecasts to treat fuel prices separately from other contributing costs.

TABLE 4
Average Impacts of a Delay Hour to Air Carriers and Their Passengers

Direct Operating Cost (DOC)	Fuel (Gal/hr)	Passengers/Flight	Passenger Time Value (\$/hr)	Total Delay Cost
\$1554	662	100	\$22.70	\$3824

A final step in evaluating CTAS delay benefits for any program plan is to account for the proportion of nationwide airport operations or emplanements that will be covered by CTAS implementations at each point in the program schedule. Rather than propose a program plan here, however, we will provide a benchmark in the form of forecast delay savings for the years 1995-2015, assuming complete nationwide implementation of CTAS during each of these years. The benchmark savings can then be downsized to fit any site-limited or phased CTAS implementation. Using FAA projections of future traffic growth [10, 11], and evaluating equation (2.2.6), we have forecast expected air carrier delays in the reference ATC system, that is, the contemporary ATC system without terminal automation. Also, projected delay savings from CTAS were calculated according to (2.3.6), and both quantities are given in Table 5. The delay savings are given for low and high end capacity estimates determined in Section 3.4, as well as a nominal estimate of 12%. Also, fuel use occasioned by the delay is included in the table, assuming the average fuel burn spent during an hour of delay, as reported by Geisinger [14]. Geisinger's fuel burn figure is appropriately weighted to reflect the fleet average fuel burn during each phase of flight (ie, gate holds, taxi-in, airborne, and taxi-out), and the proportion of delay experienced in each phase (*cf.*, Table 2).

Values which are needed to quantify the monetary benefits of delay savings, such as traffic volume, inflation, fuel prices, and passengers per flight, may be obtained from FAA forecasts (*e.g.*, references [10-12]). We have also obtained air carrier operating costs from Geisinger [14]. The estimated economic value of the savings estimated above, in constant 1988 dollars, is given in Table 6.

It is clear from Table 6 that substantial economic benefits derive from the capacity increases projected to be achievable by CTAS or similar terminal air traffic control automation. The rate of increase in benefits is greatest at the low end of capacity increases, which makes modest capacity increases more influential than might otherwise be thought. The benefits increase considerably with projected future increases in traffic volume, though they decline very slightly as a percentage of total expected delay costs, which also rise with increases in traffic volume.

TABLE 5
Reference System Projections and CTAS Savings

	Reference System		CTAS Savings					
			Low Estimate (8% Capacity Gain)		Nominal Estimate (12% Capacity Gain)		High Estimate (16% Capacity Gain)	
	Delay (K Hr)	Fuel (Mgal)	Delay (K Hr)	Fuel (Mgal)	Delay (K Hr)	Fuel (Mgal)	Delay (K Hr)	Fuel (Mgal)
FY								
1995	3081	2039	226	149	326	216	420	278
1996	3219	2131	234	155	339	224	436	289
1997	3361	2225	243	161	352	233	453	300
1998	3506	2321	253	167	365	242	470	311
1999	3604	2386	259	171	374	248	482	319
2000	3704	2452	265	175	383	254	494	327
2001	3805	2519	271	180	393	260	505	335
2002	3983	2637	283	187	409	271	526	348
2003	4130	2734	292	193	422	279	543	360
2004	4282	2835	301	200	436	289	561	372
2005	4440	2939	311	206	450	298	580	384
2006	4580	3032	320	212	463	306	596	394
2007	4724	3127	329	218	476	315	613	405
2008	4872	3225	338	224	489	324	630	417
2009	5025	3327	348	230	503	333	647	429
2010	5183	3431	357	237	517	342	666	441
2011	5322	3523	366	242	530	351	682	451
2012	5466	3619	375	248	542	359	698	462
2013	5613	3716	384	254	555	368	715	473
2014	5764	3816	393	260	569	377	732	485
2015	5919	3919	403	267	583	386	750	497
Ave/Yr	4456	2950	312	207	451	298	581	385

TABLE 6
CTAS Delay Savings to Air Carriers and Passengers
(Billions of 1988 Dollars)

	Low Estimate (8% Capacity Gain)			Nominal Estimate (12% Capacity Gain)			High Estimate (16% Capacity Gain)		
	Carrier	P'sgr	Total	Carrier	P'sgr	Total	Carrier	P'sgr	Total
FY									
1995	0.408	0.771	1.178	0.590	1.115	1.704	0.759	1.435	2.194
1996	0.448	0.859	1.307	0.648	1.243	1.890	0.834	1.600	2.433
1997	0.492	0.964	1.455	0.711	1.394	2.105	0.915	1.795	2.710
1998	0.539	1.074	1.613	0.780	1.553	2.333	1.004	1.999	3.003
1999	0.584	1.187	1.771	0.845	1.717	2.561	1.087	2.210	3.297
2000	0.632	1.318	1.950	0.914	1.906	2.820	1.177	2.454	3.631
2001	0.685	1.462	2.147	0.990	2.115	3.105	1.275	2.723	3.997
2002	0.750	1.597	2.347	1.084	2.310	3.394	1.396	2.974	4.370
2003	0.814	1.730	2.544	1.178	2.502	3.680	1.516	3.221	4.737
2004	0.884	1.874	2.758	1.279	2.711	3.990	1.646	3.490	5.136
2005	0.960	2.030	2.990	1.389	2.936	4.325	1.788	3.780	5.568
2006	1.038	2.189	3.228	1.502	3.167	4.668	1.933	4.076	6.010
2007	1.123	2.361	3.484	1.624	3.415	5.039	2.091	4.396	6.487
2008	1.214	2.546	3.760	1.756	3.683	5.439	2.261	4.741	7.002
2009	1.313	2.746	4.059	1.899	3.972	5.871	2.445	5.113	7.558
2010	1.420	2.961	4.381	2.054	4.283	6.337	2.644	5.514	8.158
2011	1.530	3.182	4.711	2.213	4.602	6.814	2.848	5.924	8.773
2012	1.648	3.418	5.066	2.384	4.944	7.328	3.069	6.365	9.434
2013	1.776	3.673	5.449	2.569	5.312	7.881	3.307	6.839	10.146
2014	1.914	3.946	5.860	2.768	5.708	8.476	3.563	7.348	10.911
2015	2.062	4.240	6.302	2.982	6.133	9.115	3.840	7.895	11.734
Ave/Yr	1.059	2.197	3.255	1.531	3.177	4.708	1.971	4.090	6.061

3. CAPACITY ESTIMATION AND ENHANCEMENT

3.1 Introduction

This section of the report is concerned with modeling airport arrival capacity and with assessing the nature and the amount of capacity enhancement that is achievable by terminal automation. The report confines its attention to capacity increases in instrument meteorological conditions (IMC), since these are the conditions in which arrival capacity is lowest, and since IMC capacities, being less than VMC capacities, are free of external constraints that might limit increases in VMC capacity. The report will also focus its capacity analysis on the case of a single runway dedicated to arrival operations. A capacity increase in this configuration, as in any other, depends, first, on having the ability to schedule operations for maximum runway utilization, adapting plans where necessary in response to changing conditions, and second, on having the ability to execute that schedule. When an arrival stream contends with departure aircraft or with other arrival streams for runway availability, the schedule for maximum utilization will supply a different set of objectives for the arrival stream than if it were operating independently. Nevertheless, once a set of objectives has been established, the airside factors that determine how well those objectives are met, namely spacing precision and the ability to sustain an uninterrupted flow of ready traffic, are applicable to any configuration. The percentage increases in capacity that are available to a dedicated single runway are approximately the same as those available in more complex configurations. Therefore, we concentrate on the single arrival runway, since it can be analyzed without attending to the collateral issues that arise in other configurations. Then we adopt the results obtained for it as an approximate standard for airport operations overall.

So long as visual contact among aircraft and between aircraft and the ground are precluded, arrival operations are subject to a set of radar separation rules that supply strict lower bounds on aircraft spacings. The radar minima are horizontal spatial distances that depend upon the weight classifications of the two aircraft being separated, and also depend upon the position of the lead and trail aircraft with respect to the runway threshold and the final approach fix (FAF) of the instrument landing system (ILS). Radar minima are given in Article 5-72 of FAA Order 7110.65, Air Traffic Control [9]. In most IMC circumstances, the control tower is able to see aircraft at some point on approach, prior to landing. Subsequently, as detailed in Article 7-10 of Order 7110.65, aircraft may be separated by less restrictive visual means. A convenient synopsis of FAA separation standards is contained in Mundra [22].

Radar minima are defined explicitly in terms of distance, since spatial position is the quantity determined by surveillance radars and displayed on the controller's scope. For analysis of a runway's arrival capacity, the relevant measurement of aircraft spacing is the time elapsed between successive crossings of the runway threshold, or the inter-arrival time (IAT). Spatial minima, combined with the integrated velocity of the lead and trail aircraft, translate into an implied minimum IAT. Consider an arbitrary pair of successive arrivals at a runway. If τ^0 is the implied minimum IAT between this pair and τ is the observed interval, then these two quantities differ by an amount b that varies from arrival pair to arrival pair. That is, as illustrated in Figure 4,

$$\tau = \tau^0 + b, \quad (3.1.1)$$

where b is a random variable taking on primarily positive values. The quantity b represents those sources of potentially recoverable capacity that a system such as CTAS has authority to address. Thus capacity analysis concerns itself with understanding the distribution of values that comprise b . Capacity enhancement is equivalent to concentrating that distribution closer to the origin.




Figure 4. Illustration of final approach spacing precision when incoming aircraft positions are such that τ_0 is achievable.

There are two fundamental sources of interarrival variability that combine to produce the overall distribution of b . One source is the intermittent occurrence of gaps which can form in the feed of aircraft being delivered to the final vector position. When such a gap forms, the final vector controller receives no aircraft that can feasibly achieve the minimum implied IAT behind a preceding aircraft in the landing sequence. In this case a positive, and sometimes a large value for b , is unavoidable. This source of variability, which in turn is a source of potentially recoverable capacity, will be referred to as the upstream feed.

Gaps in the upstream feed are to be expected in low demand situations, when there is simply too little traffic to keep the runways continually occupied. However, gaps also occur in heavy traffic situations, for two reasons. One is the inherent difficulty of coordinating traffic that arrives from different compass points and at different altitudes, and that mixes different performance capabilities. The process of aligning such traffic in preparation for final spacing is subject to miscalculation and errors of execution, just as is final approach spacing. Occasionally a gap forms in the alignment that cannot be closed up. A second reason for gaps in the upstream feed lies in the regulation of flow rates into the TRACON. For traffic flow to be stable, flow restrictions must maintain the average TRACON entry rate at a level slightly below the inherent instantaneous capacity of the airport. (This is a fundamental principle of any stable queuing system). To some extent, then, depending upon how accurately flow metering is matched to current runway capacity, flow regulation artificially creates a low demand traffic pattern, with attendant gaps in the arrival stream.

A second source of variability occurs after aircraft have been delivered to the final approach area. Controllers still face uncertainties about wind, piloting technique on final approach, and so forth, and the spacings that they attempt to establish are subject to a small random error in execution. This contribution to b will be referred to as final approach

precision. Separation must be assured in spite of such uncertainties. Therefore the spacing targets adopted by controllers may be said to incorporate an implicit buffer. The size of the buffer is considered to be proportional to the magnitude of the uncertainties. This magnitude may be measured by the standard deviation (σ) of the IAT for like aircraft pairs that have been fed to final approach with an alignment such that, by assumption, τ^0 is achievable.

This report takes an empirical approach to measuring the distribution of b , and to delineating the influence of factors that affect capacity. The empirical data base consists primarily of timed threshold crossings and airframe identifications, obtained during four occasions of single-runway IMC operation at Boston Logan International airport. A descriptive summary of the arrival data is given in Section 3.2.1.

The descriptive summary makes it clear that interarrival times can vary significantly from traffic sample to traffic sample. Variation occurs even between samples taken in weather conditions that, on the face of it, seem comparable from an ATC point of view. Some of the variation is random, but some of it reflects deliberate adaptation by controllers and pilots to differing wind fields, ceiling and visibility details, and runway surface conditions. The effect of these adaptations is to set τ^0 as appropriate for whatever are the prevailing conditions. Therefore, to analyze empirical interarrival spacings correctly, it is necessary to calibrate the interarrival times of each sample separately, that is to estimate from the traffic what values of τ^0 serve as implicit interarrival minima for the sample in question. This calibration is tantamount to sensing the inherent current capacity of a runway, and thus it is relevant not only to the review of field data, as conducted in this report, but also to the scheduling logic used in terminal automation, and to adaptive flow metering. Modelling efforts related to the calibration are discussed in Section 3.2.2.

In Section 3.2.2.1 a three-segment model is presented for the kinematics of aircraft on final approach. The model is easily understood and easily fit to observed traffic. The model has a very simple construction, but it is flexible enough to accommodate a wind field that varies with altitude, and to replicate the kinds of deceleration profiles and the rates of closure between aircraft that are observed in real traffic. As indicated in Section 3.2.2.1, this model, combined with radar separation rules, yields appropriate values of τ^0 for most aircraft pairings in most traffic samples. One exception in the traffic examined for this report appears to be explained by the anticipation and use of visual means of separation once aircraft have broken out of clouds or haze. A statistical procedure for detecting the relaxation of radar minima after such breakout, and for estimating the location of the breakout, is given in Section 3.2.2.2. Following the calibration efforts described in Sections 3.2.2.1 and 3.2.2.2, it is then possible to isolate the empirical distribution of b , which corresponds to flight-to-flight variability, and thus to excess spacing and recoverable capacity. In Section 3.2.2.3 a statistical procedure is suggested for differentiating the two components, final approach precision (1) and upstream feed (2), that together generate the distribution of b in a traffic sample. With the relative contribution of these two components identified, it is possible to establish a baseline against which capacity enhancements may be judged. In Section 3.2.3 the baseline is summarized. In Section 3.3 capacity increases obtainable with terminal automation are evaluated. The overall conclusion is that CTAS is expected to increase average IMC throughput by approximately 12% nationwide.

3.2 Baseline Performance of Terminal Air Traffic Control

3.2.1 Presentation of Field Data. Single Arrival Runway, IMC, at Boston Logan Airport

Table 7 below gives a sampling of IMC arrival operations during peak demand periods at Boston, with weather in each case below both vectoring and circling minima. The TATCA program had observers on site with visual surveillance of the runway threshold during each of the four dates listed. The observers were able to record runway threshold crossing times along with airline (if applicable) and aircraft type. For roughly the first hour of 12/15/87, there was also radar surveillance of the Boston TRACON provided by the Mode S Experimental Facility (MODSEF) radar beacon interrogator at MIT Lincoln Laboratory. A summary of the four data collections will be presented in this section.

TABLE 7
Sampled IMC Arrival Operations at Boston

Date	Time of Day (EDT)	Arrival Runway	Surface Wind (°/knots)	Ceiling	Visibility	Weather	Landings per Hour
9/13/87	16:00-20:00	4R	120/8	500-700	1/8 - 2	R+,F	34.8
10/11/87	16:00-21:00	4R	360/12	600-1100	2 - 6	R-,L-,F	35.3
12/15/87	16:00-20:00	15R,4R,15R	120/20	300-500	1/2 - 4	R-, IcePellets	28.7
5/18/88	16:00-19:00	4R	040/14	300-600	1/2 - 1 ¹ / ₂	L-F	38.1

Some of the observation periods included runway changes, and in such cases the runways are listed in order of use, but at all times a single ILS approach was in effect. With the exception of occasional side-steps on 10/11/87 from the 4R ILS to the close parallel runway 4L, arrivals were confined to a single active runway. The primary departure runway throughout was runway 9. This runway intersects with 4R, 15R, and 22L. However, the nature of the geometry and runway utilization is that for most practical purposes, the departure stream can be considered independent in analyzing arrival patterns. Therefore, the field data to be presented in this section can be said to be characteristic of a single dedicated arrival runway.

On each day the Central Flow Facility had delay programs in effect for Boston-bound aircraft. By adding a blanket delay to the scheduled departure times of flights bound for a congested facility, these programs force aircraft to absorb anticipated delay on the ground rather than in the air. Figure 1 gives a graphical summary of the delays issued by Central Flow on each observation date. The delays were substantial. On 9/13/87, the required departure delay was 35 minutes at the beginning of the observation period. The delay increased to 80 minutes during the first two hours, and declined to 40 minutes by end of the fifth hour. On 10/11/87 required departure delays ranged from 28 to 49 minutes. On 12/15/87 departure delays were 10 minutes when data collection began and climbed to over 80 minutes for the last half of the collection period. On 5/18/88 the departure delays

began at 19 minutes and climbed steadily to 80 minutes. Two points may be taken from these required departure delays. First is that throughout each observation period, even when the TRACON experienced a lull in traffic, there was a backlog of scheduled demand for runway service. The second is that delay impacts of the Boston weather were felt throughout the NAS, with the majority of experienced delay actually occurring far from Boston airspace, and primarily on the ground.

The observation periods, lasting from two to five hours, were long enough to experience some fluctuation in meteorological parameters. Ranges of values for ceiling and visibility are given in the corresponding columns in Table 7. Key characteristics of the prevailing weather pattern are given in the column headed "Weather". In this column, the symbols "R", "L" and "F" denote rain, drizzle and fog, respectively, while a suffix "+" or "-" denotes the qualifier "heavy" or "light". The surface wind given is a single value, representative of the sequence of wind measurements reported during the observation interval. Clearly some detail is lost in the abbreviation of meteorological state necessary to produce a summary table. Heavy rain on 9/13/87, for example, did not persist uniformly and without interruption for five solid hours. Nevertheless, heavy rain was reported during the majority of the observation period, and otherwise rain was reported at a light or moderate level. Therefore, "R+" gives a generally accurate summary of precipitation during the observation interval on 9/13/87. In other regards as well, circumstances during each of the intervals included in Table 7 were stable enough that the parameters listed in the table provide a useful characterization of the meteorological conditions pertaining during data collection.

It was not feasible to determine winds at altitude during the observation campaigns summarized in Table 7. One exception occurred on 9/13/87, when at approximately 18:53 EDT a pilot reported a 35 knot tailwind down to 500 feet, though surface measurements indicated almost a pure crosswind. The report was repeated to following pilots. It is fairly common at Boston Logan airport to conduct landings on runway 4R with a surface headwind but significant tailwinds or crosswinds in the air. Though no explicit indication of winds at altitude was recorded on 10/11/87 or 5/18/88, it will be seen in Section 3.2.2.1 that the assumption of a slight overriding tailwind improves the interpretation of measured landing time intervals in the context of radar separation requirements. Also of interest in data interpretation is the fact that on 12/15/87 the surface wind increased during the last hour of observation, to 27 knots, gusting to 36 knots. At 18:40 EDT on this date the tower began issuing LLWAS wind shear warnings to all flights, and around 19:30 EDT pilots were reporting loss of airspeed on final, both factors suggesting a stronger headwind on final than the measured surface wind.

In instrument conditions, interarrival times derive primarily from radar separation minima, and radar minima are established separately for each pairing of three weight classes. An aircraft's weight class is determined by its certificated maximum gross takeoff weight. The "Large" weight class includes both propeller driven and jet aircraft, and since these differ in performance characteristics and in treatment by air traffic control, it is advantageous to further subdivide the "Large" weight class into Large Props and Large Jets. The resulting categorization of aircraft, along with typical proportions of each in Boston IMC traffic, is as follows in Table 8:

TABLE 8
Aircraft Categories for Capacity Analysis, and Mix at Boston

Arrival Aircraft Category	Max. Gross Takeoff Weight (lbs)	Boston IMC Proportion
(S) Small	0-12,500	2%
(P) Large Prop	12,501-300,000	32%
(J) Large Jet	12,501-300,000	54%
(H) Heavy	over 300,000	12%

Statistics of interarrival times, calculated separately for each pair of arrival categories, and in the aggregate for each day and for all four days together, are given in Table 9. In calculating these statistics, subsets of the data were selected as necessary to restrict attention to a single runway with roughly the same operating environment. On 10/11/87, interarrival timings were excluded if either the lead or trail aircraft had side-stepped or circled to land on the nearby parallel runway, 4L, which is not ILS-equipped. There were twelve landings on 4L during the 5-hour observation period. On 12/15/87, statistics were calculated only for the latter of the two periods of use of runway 15R (the first period only involved a few aircraft, and ceiling and visibility changed substantially during the interim).

TABLE 9
Descriptive Statistics of Interarrival Times, Boston IMC Data

Seq	9/13/87			10/11/87			12/15/87			5/18/88			All Four Days		
	N	Mean	Std	N	Mean	Std	N	Mean	Std	N	Mean	Std	N	Mean	Std
SS	0	~	~	0	~	~	1	65.0	~	0	~	~	1	65.0	0.0
SP	0	~	~	2	86.0	12.7	1	125.0	~	0	~	~	3	99.0	12.7
SJ	2	88.5	26.2	6	97.0	28.5	2	129.0	28.5	1	93.0	~	11	100.9	28.2
SH	1	114.0	~	0	~	~	0	~	~	0	~	~	1	114.0	0.0
PS	0	~	~	0	~	~	1	233.0	~	0	~	~	1	233.0	0.0
PP	20	101.1	33.9	1	58.0	~	1	113.0	~	11	86.6	27.3	33	95.3	31.8
PJ	9	87.3	15.5	12	94.7	39.9	8	130.9	35.8	26	74.2	19.6	55	89.0	27.3
PH	1	97.0	34.2	1	114.0	~	3	94.0	16.7	7	71.9	17.6	12	83.0	17.4
JS	2	89.5	44.6	7	121.3	39.1	1	223.0	~	0	~	~	10	125.1	39.9
JP	20	90.1	27.8	7	100.7	13.0	7	133.4	34.7	24	103.7	29.1	58	102.2	28.0
JJ	41	96.6	28.7	70	93.2	25.5	24	130.2	40.4	22	103.1	48.0	157	101.1	32.7
JH	15	121.5	17.4	18	106.7	58.3	7	114.9	27.1	5	101.2	36.3	45	112.3	41.8
HS	1	78.0	14.1	0	~	~	1	148.0	~	1	159.0	~	3	128.3	0.0
HP	2	79.5	51.6	3	152.7	16.4	2	159.5	27.6	7	122.7	32.9	14	128.2	32.3
HJ	13	119.4	7.8	19	140.2	33.1	5	179.2	58.5	4	135.8	67.7	41	137.9	36.0
HH	3	129.7	38.7	6	101.7	20.7	0	~	~	2	115.0	18.4	11	111.7	26.2
All	130	101.1	26.9	152	103.9	33.5	64	133.9	38.5	110	95.7	33.7	456	105.3	32.5

Excepting those few arrival pairs that include a Small (type "S") aircraft, the interarrival time measurements underlying Table 9 are also displayed graphically in Figure 5 through Figure 7. In these figures, each asterisk represents a single measured interval between successive arrivals on the primary runway. The interarrival intervals are grouped by approach category pair, with the pair type indicated by a two-character string beneath the group (e.g., "PJ", "PH", "JH" in Figure 5). Figure 5 depicts so-called "closing" pairs, where the lead aircraft is in a lower performance category than the trail aircraft. Figure 6 depicts pairs having the same approach category in both positions, and Figure 7 depicts "opening" pairs, in which the lead aircraft is of a higher approach category than the trail aircraft.

Within each approach pairing, the measured intervals obtained during a given observation run are aligned in a vertical column. The columns, from left to right, pertain to the dates 9/13/87, 10/11/87, 12/15/87, and 5/18/88, respectively. On each column is superimposed a box, centered at the empirical mean for the corresponding date and approach pair, and extending 32 seconds on either side of the mean. The value of 32 seconds, as given in Table 9, is the calculated standard deviation in interarrival times, pooled over all dates and all approach categories, and it provides a measure of the level of dispersion in the combined data set.

In addition, centered on each column is a thickened horizontal line that depicts an estimate of the minimum interarrival time consistent with strict radar separation rules. The estimate is made using a three-segment model of aircraft motion on final approach, which will be described in Section 3.2.2.1.




Figure 5. ("Closing" Aircraft Pairs) IAT's at Boston, 9/13/87, 10/11/87, 12/15/87, 5/18/88. Boxes are one-sigma intervals about the mean of each arrival pair. Thick lines are approximate radar rule minima.




Figure 6. ("Same" Aircraft Pairs) IAT's at Boston, 9/13/87, 10/11/87, 12/15/87, 5/18/88. Boxes are one-sigma intervals about the mean of each arrival pair. Thick lines are approximate radar rule minima.




Figure 7. ("Opening" Aircraft Pairs) IAT's at Boston, 9/13/87, 10/11/87, 12/15/87, 5/18/88. Boxes are one-sigma intervals about the mean of each arrival pair. Thick lines are approximate radar rule minima.

Since radar separation rules are given in terms of distance, and the translation of these rules into time intervals depends on the velocity profiles of both the lead and trail aircraft, it is impossible to do more than estimate a likely minimum IAT. The estimated minimum IAT is useful as an indication of the spacing target used by controllers in a given traffic sample. However, since generic approach profiles are used in constructing the estimate, rather than separate profiles tailored for each individual approach, observed IATs that fall below the estimated minimum should not be construed as separation violations. Moreover, once aircraft have broken out visually, the tower at its discretion may apply visual separations (*cf.* Article 7-10, [9]), permitting closure below radar minima.

Reviewing Table 9 and Figure 5 through Figure 7, it is possible to make several observations regarding the interarrival spacings recorded there.

First, there are modest but systematic shifts in mean IAT from day to day. In large part, the shifts are a function of differing winds on approach. The shifts can be monitored in real time and incorporated in the automation logic. The CTAS software would thus provide the air traffic control system with timely feedback regarding the proper airport acceptance rate in any circumstance. CTAS would also help to construct and maintain arrival flows that are matched to the current arrival capacity.

Second, for aircraft pairings with enough data to form an identifiable pattern (in particular, any of the pairings involving a Large weight class turbojet, type "J") it is clear that the distribution of IATs is skewed, with a longer tail on the high end. Some of this skewness can be attributed to the concern for safety, which inclines controllers to correct more aggressively to avoid a small separation at the end of the approach than they do to tighten up a slack separation. However, the primary cause of the skewness appears to be the difficulty of maintaining an uninterrupted upstream feed, as discussed earlier in Section 3.1. In Section 3.2.2.3 methods will be presented for interpreting IAT distributional patterns in a way that clarifies the capacity implications of different automation functions (*e.g.* TMA scheduling and metering, FAST precision).

Third, referring specifically to Table 9, the IAT standard deviation on different days (see the row labelled "All") ranges from a low of 26.9 seconds on 9/13/87 and 10/11/87 to a high of 38.5 seconds on 12/15/87. Spacing precision could be expected to degrade in an exceptionally volatile or difficult operating environment, and such an effect may be present in the relatively greater dispersion observed on 12/15/87, a day with substantial wind shear and strong gusting surface winds. The differences between the standard deviations day-to-day do, in fact, test to be statistically significant with a probability of error of less than 0.01. (The test used is known as Bartlett's test for the homogeneity of variances [23], and the p-value of the test, conducted with 108, 138, 49 and 96 degrees of freedom in the four individual traffic samples, is 0.0085). However, as will be seen in Section 3.2.2.3, most of this difference in dispersion is associated with the upstream feed. The ability of controllers to form and execute spacing judgements in the final approach arena, while continuing to show some variability, is relatively consistent from day to day.

3.2.2 Models of Interarrival Spacing and Capacity

The field data presented in Section 3.2.1 exhibits shifts from day to day in the level of IAT's for any given approach pair. In the terminology of Section 3.1, these are equivalent to shifts in τ^0 . As discussed below in Section 3.2.2.1, aircraft deceleration profiles and altitude-dependent winds can be incorporated into spacing calculations by modelling final approach in three segments, one of which includes a constant rate reduction in airspeed. The three-segment model maintains computational simplicity but also permits an accurate representation of the compression that occurs between aircraft as they slow to land. The results of such calculations, performed to estimate the minimum IAT consistent with strict radar separation rules, are depicted in the slightly thickened horizontal bars in Figure 5, Figure 6, and Figure 7. Each estimate is adjusted for prevailing winds on final approach. The estimates are made using deceleration profiles that, as supported by radar surveillance, are typical for commercial transports. They also assume operating speeds that are characteristic of each of the approach categories in Table 8. The same deceleration profiles and approach speeds are used on all four days. In other words, while the estimates of IAT target minima are adjusted for prevailing wind, there is no attempt in the estimate to model adaptations in piloting technique that might accompany different meteorological conditions.

The thickened bars generally match shifts in the mean IAT in Figures 5 through Figure 7, indicating that in most circumstances shifts in IAT level (hence capacity) may be explained, and tracked operationally, by using the three-segment model with appropriate adjustment for final approach winds. One exception is seen on 5/18/88 in the low interarrival times for arrival pairs that have a relatively low-performance aircraft in the lead position (especially, pair type "PJ" in Figure 5). A possible explanation for the interarrival times of these approach pairs involves use of visual separation near the runway and coordination between local and approach controllers to mitigate spacing constraints caused by the rapid closure between the two aircraft. The issue is discussed in Section 3.2.2.2, where a method is presented for estimating the point along the approach course where aircraft "go visual".

Application of the three-segment model to the data of Section 3.2, along with refinements such as the visual breakout point considered in Section 3.2.2.2, indicate that automation can adapt its spacing objectives as necessary to match prevailing meteorological conditions and controller intentions. For automation to improve arrival capacity, it remains necessary to increase the repeatability of delivered spacings, or, stated in a converse way, to reduce spacing variability (*e.g.* the scatter about the mean in columns of Figures 5-7). In Section 3.2.2.3 a methodology is described to quantify the dispersion in IAT's as observed in the Boston field data, with the intent of providing a baseline measurement of precision in manual air traffic control. As discussed in Section 3.1, it is helpful to decompose the distribution of IAT's into two constituents, namely the entering arrival flow, or upstream feed, which may be well or poorly matched to the real capacity of the runways, and final approach spacing precision. In reality, of course, every IAT results from a combination of both effects. However, it may be supposed that in most cases one of the factors was the primary determinant of the delivered spacing. For example, upstream feed is usually the primary determinant when the controller says of a spacing that "there was no way to close it up" or "the hole was there when I got it". Given observations at the runway threshold alone, one cannot classify every IAT with certainty as arising specifically from final

spacing judgments, or specifically from spacings in the incoming arrival stream. For this reason a statistical approach, built upon the notion of mixture distributions, will be presented in Section 3.2.2.3. Results of its application to the Boston field data are summarized in Table 12.

3.2.2.1 Three-Segment Model

If the position at time t of an aircraft on a straight line course is $\xi(t)$, its airspeed is $v(t)$, and $\omega(\xi(t))$ is the longitudinal component of wind at position $\xi(t)$, then aircraft position satisfies the equation

$$\xi'(t) = v(t) + \omega(\xi(t)) . \quad (3.2.2.1.1)$$

Under subcontract to the TATCA program, a study was authorized to investigate flexible mathematical models of the wind field and of aircraft motion, with the requirement that the models be simple enough and robust enough to support the dynamic scheduling function in a computerized automation support system.

Figure 8 gives a schematic overview of the methodology adopted by the authors of the study, Sorensen, *et al.* [28]. The instrument approach begins with level flight at the glide slope intercept altitude, followed by descent along the glide slope. While it is tempting to consider means for expressing phenomena such as wind shear, the difficulty of sensing fine structure in the wind field, particularly in real time, would likely preclude its incorporation in a scheduling function. In addition, expression 3.2.2.1.1 becomes difficult to solve unless the wind field has a particularly convenient form. Fortunately, the wind field may usually be approximated by a function that varies linearly with altitude. In this case, the longitudinal wind encountered by the aircraft on final approach is either constant ($\rho=0$) or, since the glide slope is inclined at a fixed pitch, the wind varies at a constant rate with changes in the ground position of the aircraft. Namely, for some value ρ ,

$$\omega(\xi(t)) = \omega_0 + \rho \xi(t) .$$

In this case, equation 3.2.2.1.1 is a first order linear differential equation, for which a formal solution exists, regardless of the equations of motion embodied in $v(t)$.

The simplest and most natural non-constant model for airspeed $v(t)$ is to assume a segment of deceleration at constant rate ζ , sandwiched between segments of constant initial airspeed, v_i , and constant final approach velocity, v_f . In Figure 8 the extent of the deceleration segment is indicated by the double-headed arrow in the "Airspeed" and "Geometry" boxes. Note that, in principle, the deceleration segment may take place entirely prior to, or entirely after glide slope intercept. Or, as drawn in Figure 8, it may overlap and include both level and descending flight. If the onset of deceleration is chosen as the origin for time and position, and if ω_0 is the longitudinal wind at this point, then while the aircraft is decelerating its trajectory (that is, the solution to Equation 3.2.2.1.1) is given by

$$\xi(t) = \frac{v_i + \omega_0}{\rho} (e^{\rho t} - 1) + \frac{\zeta}{\rho^2} (e^{\rho t} - 1 - \rho t) . \quad (3.2.2.1.2)$$

When the aircraft is not decelerating, calculation of its trajectory is straightforward.

For convenience we shall refer to the foregoing set of assumptions (linear wind, single interval of constant deceleration) as the "three-segment" model.

The three-segment model was applied by Sorensen, *et al.* [28] to traffic recorded during several arrival pushes at Memphis International Airport. Longitudinal wind on approach was obtained by interpolating between measured surface winds and National Weather Service gridded forecast winds at 3000 feet. Accurate radar surveillance was provided by the back-to-back Mode S antenna pair operating at Memphis in support of the Precision Runway Monitor program (PRM). Aircraft tracks were fit to the three-segment model using a nonlinear least squares algorithm with four free parameters, v_i , v_f , ζ and τ , the latter denoting the time elapsed between passage of a reference point on the approach course and onset of the deceleration segment. Sorensen *et al.* found that the three-segment model was quite adequate for fitting the approach tracks of the flights they examined. In one set of 28 DC9 and Boeing 727 approaches to parallels 36L and 36R, for example, the maximum divergence of any measured radar position from the three-segment fit was 240 feet, and 99% of the time the divergence was less than 150 feet, or under one second of travel time.




Figure 8. Schematic overview of the three-segment model with linearly varying wind

While there was considerable variability from flight to flight, deceleration as estimated by the fitting procedure typically lay at or inside the outer marker, between 3.8 nmi and 4.8 nmi from the runway threshold, though there were cases of both earlier and later decelerations. The outer marker at Memphis is 4.2 nmi from the 36L threshold and 4.7 nmi from 36R. The average rate of deceleration was 0.8 knots/second. Average

airspeed at 6 nmi out was 183 knots, and average final approach speed was estimated to be 138 knots. Since they are particular to the small traffic sample being fitted, exact values of the parameters are not important, beyond indicating that the results of the three-segment fits have a reasonable form. However, it is significant that decelerations delayed within the final approach fix and within 4 nmi of the runway seem acceptable in normal traffic flow. The traffic sample was obtained from dependent parallel ILS approaches, and though an arrival push was in progress, the necessity of maintaining a stagger meant that there was no pressure on either approach stream to tighten longitudinal spacings aggressively or to recover failing separations. In cases where a single arrival stream is under sustained pressure, it is reasonable to expect that late decelerations might be further encouraged and accepted in order to maximize runway throughput.

Calculating trajectories in accordance with the three-segment model, one may solve for the interval between successive arrivals which just satisfies whatever is the most binding separation requirement, and thus obtain a postulated minimum interarrival time. Assuming deceleration segments beginning at the Final Approach Fix (5 nmi from the runway threshold), deceleration rates of 0.8 knots/second, airspeeds as indicated in Table 10, and wind parameters as indicated in Table 11, minimum IAT's were calculated. These are pictured in the thickened horizontal line segments intersecting each data column in Figure 5 through Figure 7.

TABLE 10
Airspeed Parameters, Evaluation of Three-segment Model

Airspeed	Aircraft Approach Category			
	S	P	J	H
Initial	135	175	175	180
Final	100	125	135	140

TABLE 11
Windfield Parameters, Evaluation of Three-segment Model

Long. Wind	9/13/87	10/11/87	12/15/87	5/18/88
3000'	35	10	-20	6
Surface	-2	-9	-20	-13

The initial airspeed given in Table 10 is the airspeed assumed to obtain at the final approach gate, which for the Boston traffic was taken as 8 nmi from the runway threshold. Final airspeeds are the same as those used in producing Figure 5 to Figure 7. As previously noted, winds aloft were not measured. The values given in Table 11 were chosen to correspond to observers' field notes where applicable, and to match the patterns of recorded interarrival times. Note that positive wind values indicate a tailwind, negative a headwind.

Reviewing Figure 5, Figure 6, and Figure 7, it is clear that, overall, there is a strong correlation between minimum interarrival times estimated via the three-segment

model (shown as thickened horizontal bars), and actual intervals observed in the field. In other words, the three-segment model describes trajectories well enough to relate distance and time measures accurately, for all types of approaches.

With nominal airspeed and deceleration parameters as described above, the three-segment model does fail to account adequately for the spacing of closing aircraft on 5/18/88 (see pairs "PJ" and "PH" in Figure 5). A partial explanation for the discrepancy may be mismodeled winds; however, other aircraft pairs (*e.g.* "JP"), subject to the same wind field, are well characterized. Another possible factor, which is an aspect of situation-dependent piloting technique, involves the timing and aggressiveness of deceleration segments. The notion is that pilots are willing to maintain speed longer when assured of steady winds and visual acquisition of the airport well prior to decision height, as appears to have been the case on 5/18/88. Such a willingness, of course, leads to more adaptive, and on average smaller interarrival intervals. Conversely, in marginal conditions, decelerations are taken earlier, and, because piloting technique is less flexible, controllers have less ability to influence spacing deviations inside the final approach fix. Fine adaptations in piloting technique are not measurable on the basis of the Boston data used in this report, because radar surveillance is unavailable. However, as discussed briefly by Sorensen *et al.* [28] in their study of operations at Memphis, the three-segment model provides a vehicle whereby automation can detect situation-dependent adaptations in final approach conduct (by both controllers and pilots). The adaptations can then be incorporated in scheduling calculations. We remark here that further adaptations can be made on the basis of inputs from ATC supervisory personnel, and thus automation can function without crippling the ingenuity and flexibility that exists in manual control.

3.2.2.2 Visual Separation Point

This section will explore one likely explanation of the spacing between closing aircraft pairs on 5/18/88. The explanation involves another situation-dependent adaptation, suggested by Figure 9, which depicts the compression between a P-category lead aircraft and a J-category trail aircraft, as the lead aircraft moves from the approach gate (a range of 8 nmi) to the runway threshold. The vertical axis gives the distance by which the trail aircraft follows the lead, according to the three-segment model minimum IAT calculations, when the lead is at the distance from the runway threshold depicted on the abscissa. When the lead aircraft crosses the runway threshold, for example, that distance is 2.5 nmi. Points at which the lead and trail aircraft begin or stop decelerating are marked by vertical dashed lines and the letters "L" or "T". The trailing distance required by radar minima is indicated by the faint dotted line that jumps from a height of 3 nmi to 2.5 nmi when the lead aircraft is 2 nmi out, or equivalently, at the point when the trail aircraft is first able to cross the final approach fix (5 nmi out).

The lead aircraft completes its deceleration slightly more than 2 nmi from the threshold. Once the lead aircraft has decelerated, the trail aircraft has a 50 knot advantage in airspeed, and the rate of closure is such that the trail aircraft must be separated by 3.38 nmi at this point, in order to maintain 2.5 nmi all the way to the runway threshold. However, if visual contact occurs with sufficient reliability early enough in the approach, the tower may elect to apply visual separation rules near the runway threshold, and to permit closure below radar minima when it determines that such closure is safe. If this

closure is allowed near the threshold, then, as the following discussion indicates, the 3.38 nmi separation can be reduced substantially.

Suppose that at a distance r from the runway threshold the tower is able to apply visual separations. If this point is stable, and the local controller consents, then the approach controller can set up arrival spacings to enforce radar rules only to that point, after which naturally closing aircraft may be allowed to continue to close. Such a point, if it exists, will be referred to as the visual separation point or breakout point. In current operations a breakout point is not measured explicitly, but one may argue that it is calibrated implicitly in the operating balance worked out between local and approach controllers. The text that follows will discuss how the operational use of a breakout point can be detected (or discounted) by automation software, and how the location of the breakout point and its implications for capacity and arrival scheduling can be estimated.




Figure 9. Closure between successive aircraft on final approach. Lead aircraft is a Turboprop (type "P"), Trail is a Jet (type "J"), both in the "Large" weight class. The abscissa denotes range of the lead aircraft from the runway threshold, the ordinate shows range between the lead aircraft and trail aircraft. Vertical dashed lines indicate, in terms of the lead aircraft's position, points at which the lead ("L") and trail aircraft ("T") begin or end their deceleration segments. Radar minimum separation is shown by the faint dotted line.

Since 3 nmi separation is provided up to the final approach fix, which is typically about 5 nmi from the runway threshold, one would expect that closure beneath radar limits will take place only during the final 2 nmi of the lead aircraft's approach. It may be assumed, at this point, that the lead aircraft has already reduced to its final approach speed. Rather than modeling deceleration of the trail aircraft and altitude dependence of the wind field, via equation 3.2.2.1.2 for example, to simplify the present discussion we will approximate both lead and trial aircraft as having constant groundspeed, taking for each its

average groundspeed during the 2 nmi (or less) of closure inside the final approach fix. Let V_1 represent the average groundspeed of the lead aircraft during the period of closure, and V_2 represent the average groundspeed of the trail aircraft. Also let d be the minimum radar separation, expressed in nautical miles, when the lead aircraft crosses the threshold. Note that the values assumed by V_1 , V_2 , and d will depend upon the approach categories of the lead and trail aircraft (cf. Tables 8 and 10). Unless $V_1 > V_2$, which only occurs with a low performance aircraft in the trail position, the minimum interarrival time consistent with radar rules is approximately

$$\tau^0 = 3600 \frac{d}{V_2} . \quad (3.2.2.2.1)$$

However, if closure is allowed within the visual breakout range, r , the separation between the aircraft compresses by an amount

$$r \left(\frac{V_2}{V_1} - 1 \right)$$

by the time that the lead aircraft crosses the threshold. Thus the minimum IAT becomes

$$\begin{aligned} \tau^0 &= 3600 \cdot \left(\frac{d + r \left(1 - \frac{V_2}{V_1} \right)}{V_2} \right) \\ &= 3600 \cdot \left(\frac{d}{V_2} + r \frac{V_1 - V_2}{V_1 V_2} \right) . \end{aligned} \quad (3.2.2.2.2)$$

Because V_1 is essentially the final approach groundspeed of the lead aircraft, while V_2 is closer to the groundspeed of the trail aircraft at the outer marker, compression occurs even between aircraft of the same performance class. In principle the visual breakout point can be exploited for a capacity advantage with such pairs, including, for example, large jets (type "J") following other large jets. However, the capacity advantage is slight unless there is a pronounced difference between V_1 and V_2 . An inspection of the 5/18/88 spacing distributions in Figure 5, Figure 6, and Figure 7 reveals that it is specifically when propeller-driven aircraft are in the lead that the three-segment model, applied with strict radar separation limits, overstates interarrival times.

An implication is that, at least with the control team and the conditions applying on 5/18/88, controllers perceived a capacity advantage and exploited the visual breakout point when a propeller-driven aircraft occupied the lead position, but not otherwise. Accepting this premise, the following procedure can be used to solve for r on the basis of the observed pattern of interarrival times. Let $\{\eta_k, k=1,..n\}$ be the collection of measured interarrival times (excluding cases where a faster aircraft is in the lead). Form the k -th "response" as

$$y_k = \eta_k - 3600 \frac{d}{V_2} ,$$

where i and j indicate the lead and trail aircraft categories in the k -th arrival pair. Also form the $n \times 2$ matrix A , as follows:

$$a_{k1} = 1$$

$$a_{k2} = \begin{cases} 3600 \frac{V_1 - V_2}{V_1 V_2}, & \text{if a turboprop ("P") is in the lead} \\ 0, & \text{otherwise.} \end{cases}$$

Then calculate the least squares solution of the matrix equation

$$y = A \begin{pmatrix} b \\ r \end{pmatrix},$$

where r , again, is the distance from the runway threshold to the visual breakout point, and b corresponds to the spacing buffer discussed in Section 3.1.

This regression procedure was applied to the data of 5/18/88. In performing the fit, the single case of a Small aircraft in the lead was excluded, along with three clearly outlying interarrival times, each in excess of 175 seconds. There remained 65 interarrival measurements for estimating r and b . A point estimate of 1.6 nmi was obtained for r . The standard error of the estimate was fairly high, at 0.6 nmi. However, statistical tests confirm that the fit of the regression model to the data is very good, and they support a non-zero value for r with only a 0.01 probability of error. Note that while surface visibility dropped lower than 1.6 nmi (*cf.* Table 7), visibility reported by the tower, which is distinct from the visibility measurement reported in Table 7, was 2 nmi during most of the observation period. In other respects as well, the estimated value of 1.6 nmi makes sense operationally, as does the estimated buffer value ($b = 15.5$ sec, standard error of 3.9 sec).

In the case of Large weight class jet following a Large weight class turboprop (category "PJ"), the 1.6 nmi visual breakout point corresponds to a reduction in minimum IAT of 15-20 seconds, or more than 20%, compared to the strict radar rule minimum. Thus, the capacity implications of a situation-dependent adaptation like the visual breakout point can be significant. It may also be important for controller acceptance that automation software be able to accommodate adaptations such as this, and keep in step with controller intentions.

3.2.2.3 Final Approach vs. Upstream Sources of Variability

The previous two subsections dealt with modelling constructs designed to give an accurate and consistent estimation of the spacing aimpoints used by controllers as they adapt to prevailing circumstances. In other words, the models are designed to follow and give a meaningful physical explanation of the shifts in level depicted in Figures 5 through Figure 7. It will be important for CTAS to derive or be apprised of the minimum IATs that

are applicable in any given circumstance, so that it can schedule the runway without imposing unnecessary restrictions on capacity, compared to what manual control achieves. Beyond recognizing the need to target spacing objectives at the right level, however, it is not an objective of CTAS to influence minimum IATs. Rather, as discussed in Section 3.1, CTAS seeks to increase capacity by removing uncertainty or decreasing variability in delivered interarrival intervals, so that controllers may consistently set up spacings that are closer to the minimum, without increasing the likelihood of separation violations.

A feel for the nature of spacing variability in contemporary operations may be obtained from Figure 10, which contains histograms of so-called "centered" interarrival times for each of the data collections introduced in Section 3.2.1. Centered IATs are given on the horizontal axis in seconds. The vertical axis simply records frequency counts, and the scales vary from day to day because of the differing sizes of the traffic samples. On a given day, the centered IATs are obtained by first subtracting from each observed IAT the mean IAT for its associated landing pair type, and then collecting all these deviations from the mean into one data set. This makes it possible to look at all performance classes together in the same way, and analyze patterns for a data collection as a whole. Second, the overall mean IAT for that day is added back in, so that the spread in IATs can be expressed in correct proportion to the typical IAT value. The centered IATs thus have a distribution that mimics an underlying IAT distribution, though the centered values are no longer actual IATs. In particular, referring to Figure 10, there was no 20-second IAT observed on 9/13/87. There was, however, an IAT that fell 80 seconds below the mean for its group.




Figure 10. Histograms of centered interarrival times, 4 episodes of IMC traffic at Boston

On each of the four days pictured in Figure 10, the bulk of the data falls within a central peak of approximately 60 seconds in width. We may think of the bin with the highest count as representing a typical target IAT for the data set. The spread around that peak value gives an indication as to how precise final approach spacing is in practice. The overall distribution of centered IATs is slightly asymmetric, skewed to the right. On each day there is a subset of IATs that lie to the right of the main peak, say 150 seconds and above.

Given the locations and times at which they enter the final staging area, and projecting nominal travel times along conventional paths to the runway threshold, aircraft arrive for final spacing with a *de facto* arrival schedule and an implicit distribution of interarrival times. Depending on location and on circumstances, different flow management policies (e.g., free flow, Miles-In-Trail, metering) may govern the implicit IAT distribution. At Boston, when the airport acceptance rate drops below 40 arrivals per hour, as was the case in all the observation periods listed in Table 7, implicit IATs are determined primarily by the NAS En Route Metering function (ERM). The ERM software attempts to schedule traffic into the TRACON so as to maintain a uniform time interval between crossings of the outer marker. Approach control is an instance of a capacity-limited stochastic queuing system. Any such system must accept occasional idleness in order to protect against unacceptable and irremediable backlogs. In our setting, idleness can be equated with a temporary insufficiency in the incoming arrival stream, whereby aircraft fall outside a controllability window and lose the possibility of achieving a minimum spacing. Very high IATs usually reflect an insufficiency of this sort, preexisting in the implicit IAT stream, rather than gaps generated within the terminal area. Thus very high IATs present a relatively unperturbed image of the tail of the implicit IAT distribution, as it existed prior to the actions of approach control.

To understand spacing variability in current ATC, and to assess the impact that CTAS and its component functions (FAST, DA, TMA) are likely to have on spacing accuracy, it would be preferable to be able to observe implicit IATs and to distinguish the influence that these initial conditions have on final spacing patterns from the influence of feeder and final vector controllers in the terminal area. The reasons are several. We would like to insure that measurements of final approach spacing precision really address final spacing precision, and not have them inflated by variations that arise from distinct causes. Also, lack of control over implicit IATs at its planning horizon might limit the options available to CTAS and restrict its performance. If such limitations are significant, benefits assessments should be cognizant of them. On the other hand, the strategic planning capabilities that CTAS possesses should make it possible to precondition the implicit IAT stream and tailor it better to upcoming runway availability. Benefits assessments should also be cognizant of this potential advantage.

Unfortunately, implicit IATs (often called ETAs in CTAS terminology) are not observable in the data that was available for this report, which contained only threshold crossing times. And, as is suggested by the smooth transitions from peak to tail in Figure 10, air traffic control proceeds by a continuous process of modifying the current traffic arrangement, with the result that all contributing influences are melded in the final spacing pattern. Nevertheless, it is still desirable to recover an image of the implicit IAT distribution from the Boston threshold crossing data, and to characterize this distribution as

accurately as possible, at the same time recognizing that the characterization is a provisional one.

To do so, we will take a statistical approach, described below, by which a session's complete IAT distribution is modeled as a mixture distribution. The mixture will have a probability density function of the form,

$$f(x) = \pi f_1(x) + (1-\pi) f_2(x) , \quad (3.2.2.3.1)$$

with x representing interarrival time, the component f_2 representing the implicit IAT distribution, f_1 representing the distribution of IATs subject to distinct separate action of approach control, and π representing the probability of belonging to component f_1 . Either component may in principle include both large and small IATs, though relatively speaking one would expect a more diffuse and more skew distribution from the upstream component (f_2), and a more concentrated distribution from the final approach component (f_1).

A way to visualize the conceptual premise that underlies the mixture model is to imagine the presence of an observer with infallible judgment, and to present the observer with a forced choice about each arrival. The choice is to determine whether or not the actual IAT at the runway threshold is essentially the same as the implicit IAT computed at an earlier reference point (e.g., the TRACON boundary). Of course, approach controllers are equally responsible for all arrivals. However, the observer might indicate that in a certain proportion of the traffic, approach control performed some fine tuning, but did nothing of material significance to alter the arrival time as predicted by the implicit IAT. This situation may arise when, by following nominal approach profiles, a segment of traffic naturally falls into a particular arrival sequence with an acceptable spacing, or when the traffic contains embedded gaps that the final controller cannot remove.

Also, the observer would indicate that approach controllers did alter the remaining traffic, so that actual IATs differed materially from the prospective implicit IATs. The latter case would be expected when traffic is dense, but imperfectly synchronized, so that the controller must vector aircraft around to insure proper separation. There is, of course, no observer to make the force choice described here. Instead, the process of fitting the mixture model to the collected data will serve the function of "discovering" where in the threshold crossing data the implicit IAT distribution survives, and where it has been superseded by other control actions.

In constructing the mixture model, it is necessary to be able to capture both the spacing patterns that occur when upstream flow is well synchronized to runway availability, and patterns that occur when flow is insufficient or poorly synchronized. In turn, it is necessary to take as a mathematical representation of the implicit IAT distribution a family that includes both highly concentrated and highly dispersed distributions, and both symmetric and highly skewed distributions. The two-parameter gamma family of distributions provides a sufficiently flexible representation that is also well matched to the customary patterns in interarrival spacing. Thus, suppose that the i -th component of the mixture has a probability density function of the form,

$$f_i(x) = \frac{1}{\Gamma(\alpha_i)\beta_i^{\alpha_i}} x^{(\alpha_i-1)} e^{-x/\beta_i}, \quad (0 < x < \infty).$$

This is a gamma distribution with parameters α and β . The skewness coefficient of a gamma distribution has the value $2\alpha^{-1/2}$. The skewness coefficient is a measure of the degree of symmetry or asymmetry in a probability distribution. Gamma distributions can be skewed strongly to the right, if α is small. They can also be nearly symmetric and nearly normal, if α is large, and this suggests that they are also appropriate for the precision aspects of spacing on final approach.

The process of fitting the mixture 3.2.2.3.1 to a series of observed IATs can become rather involved. For one thing, the decision by a controller to accept an implicit IAT, or alternatively to intervene and modify it, would naturally depend upon the value of the implicit IAT. Therefore, a full development of the selective switching notion would make the probability of switching itself depend upon the implicit IAT, and the fitting procedure would include an explicit representation of some such value-dependent probability function. If such a function is to be employed, the most convenient approach to calibrating the fit would probably be maximum likelihood. The calculations required to support such a fitting procedure are likely to be complex. Among other complications, in the case of maximum likelihood it is known that the likelihood surface contains a large number of local singularities (Titterton, *et al.* [30]), each supplying a potential but generally spurious "answer" to numerical optimization routines.

Rather than undertake a level of effort that is out of balance with current needs, for each data set we will restrict the formulation to a single overall switching or mixture probability, applicable equally to all the implicit IATs in the data set. Also, we will make use of the convenient form assumed by the moment generating function of gamma variates, and employ a simple least squares criterion for fitting. The moment generating function (mgf) of a random variable X , assuming that the expectation exists, is given as

$$M_X(t) = E(e^{tX}).$$

In the case of a gamma distribution with parameters α and β , the mgf is,

$$M_X(t) = (1-\beta t)^{-\alpha}, \quad t < 1/\beta.$$

Therefore, the fit to a gamma mixture may be performed by calculating the empirical mgf,

$$\bar{M}(t) = (1/n) \sum_{i=1}^n e^{t x_i}, \quad (3.2.2.3.2)$$

where x_i is the i -th IAT, and then solving a nonlinear least squares problem to fit the five-parameter mixture mgf,

$$\hat{M}(t) = \pi (1-\beta_1 t)^{-\alpha_1} + (1-\pi) (1-\beta_2 t)^{-\alpha_2}$$

to the responses $\{\bar{M}(t_j)\}$, for a judiciously chosen set of values $\{t_j\}$. Optimal choice of the $\{t_j\}$ is an open question. However, since the variance of $\exp(tX)$ grows much more quickly with t than does its mean, indeed is undefined for $t \geq 1/2\beta$, the values in $\{t_j\}$ should be kept small. We used the set $\{t_j\} = \{-.025, -.020, -.015, -.010, -.005, .002, .004, .006, .008, .010\}$. A Newton-Raphson method with analytic gradient was used to compute nonlinear least square solutions.

Results of the fitting procedure are given in Figure 11 and Table 12. Figure 11 repeats the centered IAT histograms of Figure 10, along with dashed lines that depict the probability density functions (pdf) of the two gamma components, as well as the combined mixture density in the resulting fit (solid line). Table 12 gives parameters of the gamma mixture, followed by the mean and standard deviation of each component entity. By convention in Table 12, the more concentrated of the two mixture components is designated as f_1 . This is the component normally associated with final approach spacing precision. Therefore π represents the probability that controller-directed final spacing adjustments form the primary determinant of interarrival time, as opposed to spacings established *de facto* in the upstream feed.




Figure 11. Fits of gamma mixtures to Boston interarrival data

Note that only one curve is depicted in Figure 11 for the data of 12/15/87. Though it was seeded with a two component mixture, on this day's data the iterative nonlinear least squares fitting procedure converged to a solution in which the more concentrated of the two components vanished. The same result occurred regardless of the initial guess that was used. The vanishing final approach component is also evident in the zero value for π in Table 12. The interpretation we make of this is that, on 12/15/87, traffic was so light that the basic features of the interarrival pattern were set by TRACON entry times, and actions

of the final approach controller had little opportunity -- or incentive -- to affect the received IAT distribution. The traffic subjected to statistical analysis on 12/15/87 was that which occurred only after the final switch to runway 15R (cf. Table 3.2.1.1). Slightly less than two and one-half hours of traffic was timed during this period, during which 64 arrivals were recorded. This amounts to an average of 26 arrivals per hour, which is well below the single runway landing capacity at Boston, except in very unusual circumstances.

TABLE 12
Parameters of Two-component Gamma Mixtures, Fitted to IMC Interarrival Data at Boston.

	π	α_1	β_1	α_2	β_2	μ_1	σ_1	μ_2	σ_2
9/13/87	.72	25.6	4.0	4.9	19.5	102.3	20.2	95.4	43.1
10/11/87	.78	29.5	3.3	5.1	24.5	95.9	17.7	124.4	55.2
12/15/87	.00	...	...	15.1	8.9	...	...	135.1	34.8
5/18/88	.49	39.8	2.3	9.5	9.9	91.6	14.5	94.1	30.6

It should be recognized that the results displayed in Table 12 must be interpreted with caution. The main reason for caution is the provisional nature of the mixture model, and the restricted nature of the data used to develop it. Sampling variability is likely to be quite high in mixture problems of this nature, and no effort has been made within the context of this report to establish confidence intervals for the fitted parameters. A second reason is that convergence of the nonlinear least squares calculations was very slow on the data of 9/13/87 and 5/18/88, and it was necessary to take intermediate results, albeit after several thousand iterations. The intermediate results were chosen to give a best fit to the empirical cdf of the centered IAT distribution, since on both dates it appeared that during the course of its iterations, the nonlinear least squares algorithm passed through a local minimizer of that criterion.

However, at the same time that we acknowledge its tentative character, we would argue that the mixture model discussed in this section represents a reasonable effort to use the available data to obtain the best assessment possible of the relative roles played by metering and by final approach operations in determining airport capacity. Conclusions that we draw from Table 12 are, first, that in capacity-bound IMC operations two-thirds to three-fourths of the arrival traffic generally has its final IAT reformulated by the spacing skills of a final approach controller. The remainder of the IAT population follows from initial conditions established upstream. Some of the "upstream" aircraft will arrive already well spaced, but gaps in the arrivals will be prevalent enough that the upstream component represents a relative drain on capacity. Examining the modelled mean IATs for the upstream component (μ_2), which range from a low of 94.1 seconds on 5/18/88 to a high of 135.1 seconds on 12/15/87, it appears that the properties of the upstream component vary considerably from circumstance to circumstance. This is understandable, since the upstream component reflects the naturally variable success of traffic management policies that must adapt a long (in some regards a multi-hour) planning horizon to an uncertain and changeable terminal environment. The final spacing component appears to be more invariant. The mean IATs for this component lie in roughly a 10 second range. The standard deviations used to measure final approach precision (σ_1) range from the mid-teens to 20 or 21 seconds.

3.2.3 Summary of Baseline System Performance Measures

The modeling efforts in Section 3.2.2 were directed toward obtaining proper characterizations of the performance of today's manual air traffic control system, with a view to establishing a baseline for calibrating future capacity enhancements. In Table 9 the interarrival times from four extended periods of single-runway IMC traffic were summarized. The mean observed IAT, averaged over all this traffic, was 104 seconds. An estimate of the variance in IAT delivery, pooled over all four days and all aircraft pair types, yielded a composite standard deviation of 32 seconds. Subsequently, using a three-segment model of aircraft flight, estimates were made of intended IAT minima, conditional upon prevailing winds and typical aircraft deceleration profiles. The resulting IAT minima averaged 76.5 seconds over the four-day Boston data set. The minima were used to normalize interarrival time values for different weight classes, and to discriminate between required separation from separation that might be removable via automation.

In Section 3.2.2.3 the IAT distribution was discussed in terms of two generative factors, one being the inherent precision achievable in forming and executing spacing judgments on final approach, and the second being the feed of traffic delivered to the final staging area. Using techniques developed in Section 3.2.2.3, analyses of IMC traffic at Boston suggested that final approach precision is relatively consistent from traffic sample to traffic sample, that it can usually be described by a standard deviation of 15 to 21 seconds (we will take 20 seconds as a point estimate), and that the average spacing buffer for aircraft subject to final approach spacing precision is about the same magnitude as the IAT standard deviation, or 20 seconds. Also, a preponderance of about 70% of the arrival traffic was generally found to belong to the final approach group. As shown in Table 12, the upstream contribution to IAT patterns varied more from circumstance to circumstance than did the final approach contribution. The upstream contribution was also more diffuse and more difficult to estimate statistically. However, if one accepts a 20 second buffer and a 20 second standard deviation for excess spacing in the final approach group, and accepts that this group constitutes 70% of the arrival traffic, one may then estimate moments of the upstream group by choosing them to match the overall mean and standard deviation observed in the combined Boston data. Doing so, the average excess IAT in the upstream group is estimated to be 46 seconds, while the standard deviation is 48 seconds.

Thus, the reference system we take as a baseline for future capacity improvements appears as given in Table 13. Note the general equality between the size of the average separation buffer and the IAT standard deviation for the corresponding traffic group. This equality does not arise because of any theoretical premise in the modeling process. Indeed, according to the commonplace assumption of a Gaussian distribution for spacing perturbations, a buffer of only one standard deviation would produce an average of one missed approach in every 6.3 approaches, which certainly does not occur in practice. Nevertheless, the use of approximately a one-sigma spacing buffer has been observed in practice. The fact that it has been observed suggests that one or both of two premises hold. One is that local controllers, who frequently have access to visual separations, as well as timely knowledge of pilot intentions and likely runway occupancy times, are able to accept separations that, on a case by case basis, may drop below what we would calculate as a minimum by radar separation rules. Second is that some of the variability observed in IATs occurs intentionally, at the behest of the controller, as the controller adapts to slight changes in the environment, or perturbations in the flow of surrounding traffic. If the second premise is true, controllers' intrinsic precision is better than the 20 second figure we

have adopted. The value of 20 seconds reflects both the controller's ability to execute his intention, and variations in that intention from aircraft to aircraft, because of the presence of surface and departure traffic, as well as other imponderables.

TABLE 13
Characteristics of the Baseline System

Traffic Component	Proportion of Total Traffic	Minimum IAT (sec)	Average IAT (sec)	Excess IAT (sec)	Std Deviation (sec)
Overall	1	77	104	28	32
Final Approach	0.7	77	97	20	20
Upstream	0.3	77	123	46	48

Table 13 indicates that in contemporary air traffic control there is an average of 28 seconds of interarrival buffer, or excess spacing that automation can potentially reduce, compared to an average required interarrival interval of 76.5 seconds, rounded to the nearest integer in the table. Of the 28 seconds average excess, $0.7 \times 20 = 14$ seconds are attributable to excess in the final approach group. Approximately the same amount of removable excess, $0.3 \times 46 = 13.8 \approx 14$ seconds, is attributed to the upstream feed.

3.3 Performance Improvements

3.3.1 Improved Final Approach Precision and Flow Metering

Having determined a set of performance parameters for the current "manual" air traffic control system, it remains to estimate the impact that TATCA automation, particularly CTAS, will be able to exert upon system performance. Given the substantial cost of developing automation software that integrates with all the information flows, procedures and special considerations that arise in existing ATC systems, and given a reluctance to introduce developmental systems into critical operations, studies of the impacts of ATC automation have historically relied upon laboratory simulation and/or simple mathematical models. A tabular summary of a survey of such studies is given in Table 14, which will be discussed in the remainder of this section. One exception occurred with the computer-aided approach system (CAAS), which was installed at John F. Kennedy International Airport from December 1966 to April 1967, and utilized operationally "through the heaviest traffic periods at JFK, in all types of weather conditions" [21].

The study identifiers contained in the left-most column of Table 14 provide a shorthand for the technical approach or the experimental system forming the focus of each study. CTAS, described briefly in the introduction to this report, has been adopted as the primary operational concept of the TATCA program. It originated at NASA Ames Research Center, where it has been evaluated in simulation of the Denver airspace, and it is under continued development in preparation for field deployment. TIMER, like CTAS, is an integrated system for time-based terminal area sequencing and spacing, as well as extended area metering and descent management. TIMER has been developed and tested at NASA Langley Research Center. DTP is an automation system similar to CTAS and TIMER, with scheduling and display capabilities to assist controllers in sequencing and

spacing tasks. DTP was developed and tested at MIT Lincoln Laboratory. The entries under CDTI represent a series of investigations conducted to evaluate the spacing performance of pilots supplied with a cockpit display of traffic information. The performance was evaluated by students and faculty at MIT in several contexts, including

TABLE 14
Survey of ATC Automation Research

System/Study Identifier	Study Type	Sample Size		IAT s.d.		Capacity	Refs	Notes
		Man	Auto	Man	Auto	Gain		
CAAS	Live (NAFEC)	...	120	20.6	11.8	...		aop
	Operational	745	830	25.6	16.2	6-16%		gg
TIMER	FastSim	...	?	...	8-12	16-25%		af
	RealSim(C,P)	...	16	...	14.1	16%		bcdej
			19		10.1			
CTAS	RealSim(C,P)	104	155	25.2	16.9	11.9%		
DTP	RealSim(C)	386	385	10.2	7.1	...		hi
CDTI								
Connelly, MIT	RealSim(P)	...	9	...	11.1	...		kl
			42		2.4			
Kelly, NASA	RealSim(P)	...	100	...	8.1	...		kmn
MITRE	Sys Analysis	...	...	18	9	12-16%		

the presence or absence of ground based metering and spacing automation, different human-system interfaces, and different levels of on-board navigation and flight management. Related research has also been conducted at NASA Langley. Finally, a number of FAA-sponsored capacity analyses have been conducted by The MITRE Corporation, including Lebron [16].

The column "Study Type" in Table 14 distinguishes between fast time simulation, in which all human participants in the control process are emulated (usually perfunctorily) in software, and real time simulation, in which selected human participants perform their customary functions in as realistic a manner as possible. A suffix of 'C' in real time simulations indicates that there was a human controller in the loop. A suffix of 'P' indicates a human flight crew in a cockpit simulator.

Sample size, or the number of approaches in simulated or live trials, is given where available. Standard deviation of IAT should be self-evident. Both of these quantities are given separately for the current manual ('Man') control system, and with the provision of ATC automation ('Auto'). Estimated capacity gains are indicated when the conductors of the studies published them, or provided data which enabled them to be calculated. Indices in the list of references are given for relevant reports. Also, a collection of notes was assembled describing distinguishing characteristics of individual studies. The notes are reproduced in Appendix 1, and indices into the collection of notes are given in Table 14.

Laboratory simulations vary considerably in the fidelity and complexity of the simulation environment they emulate, in the operational scenarios they create, and in the way in which they measure and report performance statistics. Therefore, reporting differences are to be expected between even similarly intentioned simulations. Also, a fundamental conflict often surfaces in the design of simulations, as indeed in any experimental system. Namely, the ability of an experiment to provide evidence regarding a particular performance or system design issue is usually maximized when one controls factors that are distinct from the primary issue but potentially consequential, in order to preclude them from exerting a confounding influence on the specific comparison that is being sought. Such controls, however, may limit the straightforward transfer of simulation results to a full operational setting. For example, in the "DTP" experiments conducted by MIT Lincoln Laboratory, the primary objective was to evaluate the efficacy of a particular calculation and passive display strategy for improving final spacing at the outer marker. The experimenters wished to evaluate precision and workload effects of the spacing aid under sustained heavy traffic load, and at the same time to isolate these effects from the variable dynamics of high density traffic and from monitoring and flow management tasks that are interrelated but distinct elements of an approach controller's job. Therefore, the simulation was run with a high arrival rate, but with release times into final vector airspace timed perfectly or nearly perfectly so that aircraft would fall into proper sequence, already fundamentally well spaced. With this arrangement traffic flowed smoothly, yet a steady demand was maintained to challenge the fine tuning capabilities of the approach controller. The result was that the simulation was able to maximize the use of limited subject controller time to focus on its specific concerns. It was found that the passive display aid reduced inherent spacing variability by about 30%. The aid also appeared to have a beneficial effect on controller workload [24]. However, the absolute performance measures recorded during the simulation (IAT standard deviations of 10.2 seconds without the aid, 7.1 seconds with) are of uncertain extensibility to more erratic real world traffic flows. All the other simulations in Table 14 also rely upon some type of oversimplification, though in some cases the probable consequences of the simplification are more modest or more subtle. As far as we can tell from the descriptions provided, all of the simulations are limited in their emulation of the existence or the operational consequences of:

- dynamics and near term uncertainty regarding conditions at the airport, including
 - shifting and variable winds
 - changes in ceiling and visibility
 - surface traffic and departure demand
- the role of the local controller in reacting to the above dynamics, and in influencing aircraft motion inside the final approach fix
- RF congestion and distractions
- the kinematic performance and controller treatment of commuter and General Aviation (GA) traffic.

Also, except in cases where a piloted aircraft simulator is incorporated into the system simulation, and often in these cases as well, there is very limited emulation of real-world

- pilot response lags
- variations in piloting technique, particularly on final approach.

Taken together, the laboratory simulations provide the broadest and most current information available on ATC automation. However, because of the oversimplifications and omissions listed above, some of which are practical necessities, we regard the individual simulations only as guidelines, to be used to develop an informed estimate of automation impacts. Also, we accord substantial weight to the CAAS experience, because of its large sample size in a full operational setting, despite the twenty-five years that have elapsed since the experiment.

With minor exceptions, the simulations that control flow metering precisely to isolate precision and workload in final approach (CAAS Live, TIMER, DTP, CDTI MITRE) report automation-assisted IAT standard deviations of 8-12 seconds.

We will now present worst-case, best-case, and mid-range estimates of the improvements in IMC arrival capacity that may be anticipated from a mature implementation of CTAS in the TRACON.

Compounding worst-case assumptions, assume the least consensus improvement in final approach precision ($\sigma=12\text{sec}$), and no improvement in upstream delivery. Relying on the observation in Section 3.2.3 that mean buffer size (i.e., mean excess IAT above intended minima) is commensurate with IAT standard deviation, then a change in final approach precision from the baseline of $\sigma=20$ seconds to $\sigma=12$ seconds would reduce the IAT buffer by about 8 seconds. Referring to Table 13, an estimated 70% of IMC arrival traffic is typically subject to this improvement in final approach precision. Thus overall average IAT would be reduced by about $0.7 \times 8 = 5.6$ seconds, representing a capacity increase of about 6%. This is the same level observed in the field by the CAAS program, which in fact was designed originally to have companion functionality for en route release timing or metering, but lacked metering functionality in the operational implementation. Thus when the CAAS was turned on it did accept the same delivery of aircraft from the en route system as did the manual system.

Best-case assumptions would be, first, that final approach precision improves to $\sigma=8$ seconds. Also, suppose that, with automation, better synchronized metering and early recognition of potential spacing irregularities would serve to eliminate large gaps and so make the greatest possible proportion of aircraft subject to final approach precision. It must be recognized that air traffic control has to function in a physical environment that at times is volatile and uncertain. Because of this uncertainty, even if the ATC system synchronizes the upstream feed to match perfectly its best estimate of upcoming events on final approach, the actual timing of events may differ from what was anticipated. Also, upstream flow management involves the en route ATC network in ways that are beyond the scope of CTAS operation. Therefore, it is difficult to be specific about the extent to which CTAS will directly or indirectly influence what has been identified as the upstream component in contemporary air traffic control. Lacking an operational history for CTAS, we have to look to the studies cited in Table 14 for evidence as to its likely effect. The CAAS program and the TIMER studies, in particular, contain relevant information.

Despite the severe constraints that the CAAS program faced with respect to computing, communications and display capability, and without the support of an en route metering function, the CAAS program still exerted a positive influence on terminal area traffic flows. The number of aircraft in the final staging area was maintained at a much more consistent level. Also, controllers succeeded in narrowing delivery gaps when supplied with the "alpha" symbol, which signified a scheduled interval that was larger than a minimum interval, and thus an imbedded gap in the stream, compared to how they performed lacking display of that information. The CAAS final report [21] considers an "average delay for alpha aircraft", which corresponds conceptually to the average excess interarrival time for the upstream component in Table 13. In manual mode, the average delay was calculated by the CAAS staff to be 47 seconds (which is almost identical with the reference system estimate of 46 seconds in Table 13). With CAAS turned on, the average delay attributed to alpha aircraft was reduced to 38 seconds.

In the TIMER fast time studies (Credeur and Capron [6]) it was concluded that a modern automation system, given good adaptive scheduling, and operationally feasible amounts of flexibility in travel time along TRACON arrival routes, could deliver full final approach capacity with metering time inaccuracies (that is, standard deviations of coordination fix crossing time errors) of 45 seconds. A decline from ideal capacity of only 1% was indicated for inaccuracies as high as 60 seconds. En route flow management has changed since the era in which CAAS was operational. As discussed in Section 3.2.2.3, and noted in Table 13, an IAT standard deviation of 48 seconds is reported in this study for the subset of arrivals whose final spacing appears to be a direct function of their TRACON entry times, essentially unaltered by approach control. While this is not the same as measuring accuracy at a TRACON entry fix, it suggests that adequate structure exists in contemporary en route metering to keep the TRACON under sustained arrival pressure, assuming that airport arrival rates are properly set, without incurring uncontrolled workload.

The above studies suggest that automation has a potential for significantly reducing the presence of gaps in the upstream feed. For an optimistic best-case estimate we will suppose that CTAS is able to reduce the excess interval arising from the upstream feed by two-thirds, from 13.8 to 4.6 seconds. Coupled with the most optimistic estimate for improvement in final approach precision, $\sigma=8$ seconds, the remaining excess in interarrival interval is $0.7 \cdot 8 + 4.6 = 10.2$ seconds, which corresponds to a capacity increase of 20%.

For a mid-range estimate, which we view as the most likely capacity impact of a CTAS implementation, we assume that final approach precision is improved to an intermediate value of $\sigma=10$ seconds, and excess interarrival interval associated with the upstream component of traffic is reduced by one-half, from 13.8 to 6.9 seconds. The estimate of capacity gain resulting from such improvements is approximately 12%, which is consistent with results of the CTAS real time simulation, and in the middle range of most other study conclusions.

3.3.2 Optimal Sequencing

Wake vortex considerations have a significant impact on spacing and capacity when Small or Heavy aircraft are involved. When a Small weight class aircraft follows a Heavy aircraft, radar rules require the Small to trail the Heavy by 6 nmi when the Heavy crosses the threshold. Either the Small or the Heavy need trail another Small by only 2.5 nmi (or 3 nmi, depending on the airport and runway conditions), while the Heavy can trail another Heavy by 4 nmi. Therefore, a group of two Small, two Heavy aircraft and one Large aircraft may require a cumulative spacing of 17 nmi, if sequenced Heavy-Small-Heavy-Small-Large, while it will require only 14 nmi if sequenced Small-Small-Heavy-Heavy-Large. The above aircraft grouping occurs infrequently in practice, but it illustrates the point that simply by identifying and helping to organize preferred arrival sequences, without requiring any precision improvements, there is a possibility that automation software can create a capacity advantage. Numerous studies have sought to evaluate the practical benefits of sequencing algorithms. Among these we mention Thompson [29], Luenberger [18], Dear [7] and Venkatakrishnan [31].

It is clearly not possible to arrange a real stream of arrival aircraft in strict like-follows-like sequence. Routing constraints and limited controllability, among other factors, are such that both the design and the eventual capacity benefits of an effective practical sequencing algorithm remain open issues. In one geometry, for example, it may be impractical to change sequence by having one aircraft overtake another on a common approach segment. In another case overtakes may be routine. One common artifice for representing the limited ability that exists to adjust an aircraft's time to fly from the scheduling horizon to the runway threshold, is to limit each aircraft to a shift of no more than K slots from its ordinal position in a natural first-come-first-served (FCFS) sequence. The shifting constraint also tacitly incorporates an index of "fairness" in the computation of sequencing priorities, and it serves to control computational requirements as well. Limiting shifts to two positions, and using an aircraft mix of 15% Small with the remainder equally distributed between Large and Heavy, simulations by Leuenberger showed capacity increases of 4.5% to 5.5%, depending upon the extent to which overtakes were permitted. Simulations conducted independently by Thompson, who assumed 12% Small and 19% Heavy, with otherwise similar scheduling algorithms and operational assumptions, found a 6.6% capacity increase with the two-position constraint, compared to an idealized increase of over 17% with fully-optimized sequences. The increases were slightly reduced if overtakes along a common route were disallowed.

There are reasons to expect that the real percentage impact of sequence optimization on arrival capacity, averaged over the course of a year, for example, will be more modest than the figures reported above. First, the aircraft mix at major commercial airports includes fewer than 12% Small weight class aircraft, and this proportion typically decreases in poor weather. Potential benefits of sequencing decline along with the proportions of Small and Heavy aircraft. Second, the authors above assume as a baseline that no purposeful weight class sequencing is performed. Rather, they assume that the string of aircraft weight classes in the baseline system reflects random, mutually independent selections according to the marginal proportions in the arrival mix. The potential inefficiencies inherent in mixing aircraft with different performance characteristics are well known, and airports and ATC staff already address those inefficiencies in some ways that are not reflected in the baseline assumption of random sequencing.

In many airports, high and low performance aircraft can frequently be directed to separate landing runways, in which case additional wake vortex separation is seldom required.

Even when all components of the arrival mix are sharing an ILS approach to a single runway, data collected under the auspices of the TATCA program suggest that it is customary for smaller aircraft to mitigate their capacity impacts by decelerating late and maintaining highest possible approach speeds.

There is some indication, albeit inconclusive, that ATC already achieves a degree of sequence optimization. The counts of arrival sequence pairs observed during four periods of single-runway IMC operations at Boston are tabulated below in Table 15.⁴ Alongside these counts are expected counts, conditional upon the observed marginal category probabilities (obtained as the average of the lead and trail marginals) and the independence assumption. Evidence of preferential sequencing lies most clearly in the fact that turboprop aircraft are observed more frequently than expected behind other turboprops (36 observed BB sequences versus 25 expected), and slightly less frequently than expected behind Large jets or Heavies. The chi-square statistic, which is calculated as the sum of the ratios (observed-expected)²/expected across all the cells in the table, and which is widely used to test the assumption of independence in contingency tables (see Snedecor and Cochran [27]), has the value 12.04. A conventional test of independence proceeds by calculating the probability that a central chi-square distribution with, in this case, nine degrees of freedom, would exceed the value 12.04. That probability is slightly greater than 0.2, which would not usually be considered extreme enough to discount the independence (random sequencing) assumption. However, because the category assignments in Table 15 arise in a dependent manner (i.e., the trail aircraft in one pair becomes the lead aircraft in the next pair), conventional hypothesis testing with the chi-square statistic is not valid. Further analytical effort would be needed to define a satisfactory measurement of the extent to which potential benefits of optimal sequencing are already realized in current air

TABLE 15
Observed IMC Arrival Sequencing at Boston Logan

Observed Counts Arrival Pairs						Expected Counts (Assuming Independence)					
Lead	Trail					Lead	Trail				
	S	P	J	H	Total		S	P	J	H	Total
S	1	3	13	1	18	S	0.6	3.9	10.3	2.7	17.5
P	2	36	56	13	107	P	3.9	25.2	65.5	16.9	115.5
J	10	63	172	49	294	J	10.3	65.5	170.2	44.0	390.0
H	4	14	45	12	75	H	2.7	16.9	44.0	11.4	75.0
Total	17	116	286	75	494						

⁴ The counts in Table 3.3.2.1 are slightly larger than those in Table 3.2.1.3 because all of the traffic from 12/15/87 is included in Table 3.3.2.1 (restrictions to insure a uniform wind field which were imposed for Table 3.2.1.3 need not be imposed in discussing aircraft sequencing), and also because there were occasional arrivals for which aircraft type was noted, though the observer was unable to time the interarrival interval.

traffic control. Also, for nationwide applicability a broader data collection effort would be required, since the amount of resequencing that occurs in current practice undoubtedly varies with the route and airspace structures at each site.

For current purposes we will accept that random sequencing provides a reasonable approximation of current ATC practice, and thus an acceptable premise for a reference system for benefits analysis. However, reservations we have about that premise, in addition to other reservations cited above, incline us to be conservative in estimating the likely benefits of automation-assisted arrival sequencing.

Our summary conclusions about the benefits of optimized sequencing are as follows. To a certain extent, facilities already address the sequencing issue by segregating different categories of aircraft into different approach streams. The segregation is seldom complete, however, and it may disappear entirely when weather precludes multiple runway configurations. In the absence of such procedural sequencing, automation can provide a longer planning horizon than is feasible for human controllers, as well as the computational power to support a more exhaustive search through various timings and juxtapositions of aircraft types, in order to determine when opportunities for beneficial resequencing exist. When an opportunity does exist, the interarrival times for clusters of resequenced aircraft can be substantially reduced. The frequency with which such opportunities arise depends upon aircraft mix and local flow dynamics, and thus the benefits of optimized sequencing, as with most automation functions, can be expected to vary from time to time and from site to site. Our best estimate of the nationwide, time-averaged benefit of an effective sequencing function, is that it can increase single-runway arrival capacity by 2% - 4%, compared to the levels currently sustained in air traffic control.

3.3.3 Summary of Capacity Improvements

In Section 3.3.1 a nominal estimate of 12% was given for the likely increase in throughput of a single IMC arrival stream, as a result of scheduling and spacing improvement provided by CTAS. Compounding worst-case assumptions or best-case assumptions, capacity increases of 6% or 20%, respectively, were indicated. However, in producing meaningful low end and high end estimates, some degree of convolution (or averaging) should be appropriate, as opposed to compounding solely pessimistic or solely optimistic assumptions. Therefore, for purposes of evaluating CTAS impacts we will round both extremes to the middle slightly, and use 8% and 16% as low and high end expectations.

In addition to the above improvements it was estimated in Section 3.3.2 that sequence optimization could provide an additional 2-4% increase in arrival capacity. Because sequence optimization is a natural concomitant of the scheduling function, which forms the basis for all other automation capabilities, we would expect some form of sequence optimization in mature automation products. However, as sequence optimization algorithms have received little scrutiny in field trials, and as they may encounter or

introduce unanticipated side effects, we prefer not to claim the potential sequencing improvements in forecasting program benefits. Rather, we leave them as a buffer, in effect, to help insure that the benefits we do claim are not exaggerated.

Therefore, the summary estimates of capacity improvements to be anticipated from CTAS are as given in Table 16.

TABLE 16
Summary of Capacity Improvements

Low End	Nominal	High End
8%	12%	16%

4. SUMMARY AND CONCLUSIONS

In this report we have evaluated the prospects that planning and advisory software have for assisting approach controllers in increasing IMC arrival capacity, and the prospects and value of reducing air travel delays. The report began by developing a forecast methodology that predicts nationwide air traffic delay levels as a function of traffic volume, and also projects the manner in which delay levels will respond to a broad change in system capacity. Using quantitative estimates of the size of capacity increases available from terminal automation, estimates of delay savings and attendant monetary savings for the years 1995-2015 were given in Table 5 and Table 6. These estimates were based upon FAA projections of future traffic growth, and assumed a nationwide implementation of the CTAS component of TATCA. As an example, in fiscal year 2000, the FAA forecasts approximately 8.25 million air carrier departures. Our forecast methodology predicts that, with the contemporary air traffic control system, air carriers will report 3.7 million hours of delay. A nominal 12% capacity gain is subsequently estimated for CTAS, and this capacity improvement is predicted to reduce the delays in fiscal year 2000 by 383 thousand hours, or by slightly more than 10%, representing a savings to air carriers and their passengers of almost \$2 billion in constant 1988 dollars. Since development and implementation costs of CTAS are estimated to be only about one-eighth of this amount, and since the delay savings recur annually, it appears that the implementation of ATC automation software such as CTAS will pay for itself many times over.

In order to evaluate the capacity gains achievable using terminal automation, the report began with a careful study of IMC arrival spacing in the current (so-called manual) air traffic control system, based as much as possible on data from high density final approach operations in the field. The many factors that influence arrival spacing were partitioned into two broad components that reflect interrelated but fundamentally different aspects of the task of approach control. One component, called final approach precision, is concerned with the positioning of aircraft for localizer intercept and subsequent fine tuning on the final approach course. The second component corresponds to flow management and preliminary sequencing and spacing upstream of the final staging area. Working conclusions about the role played by each component in current ATC were presented in Section 3.4.1. Subsequently, relying upon laboratory simulations and analyses reported by five major research organizations specializing in ATC, capacity increases achievable by an automation product such as CTAS were estimated, assuming conventional aircraft equipment and general conformance with current ATC procedures. Average systemwide IMC capacity increases were estimated to be in the range of 8-16%, with 12% chosen as the best working estimate.

Long-term forecasts of the behavior of complex systems, influenced by such quantities as inflation rates, fuel costs, or air traffic levels, are widely recognized to be uncertain. Similarly, the conclusions of this report regarding capacity increases and delay impacts are estimates, not known quantities, and it is difficult to place definitive confidence bands on the numerical estimates. However, there is no appreciable likelihood, in our opinion, that the overriding system relationships we have observed, or the apparent economic desirability of terminal air traffic control automation, are invalid.

While further effort does not appear necessary to establish whether the benefit/cost ratio of terminal automation exceeds a decision threshold, there are areas in which additional study might aid the design or clarify the benefits of CTAS.

If a small number of sites must be selected to receive CTAS, it may be necessary to consider the climatology and the configurations and capacities of candidate airports individually, as well as traffic volumes and growth rates between city pairs, in order to develop more implementation-specific benefit statements. It may be worthwhile to investigate the potential benefits of terminal automation in visual conditions. Also, there is value to flight plan and schedule reliability that is not fully captured using only direct operating costs, as has been done in this report. Padded block times represent a hidden cost to air carriers, as do flight cancellations and missed connections. Similar statements apply to travelling passengers as well. It should also be recognized that improvements in the air transport system have broad impacts on productivity and local and national economic growth, extending well beyond operating savings for air carriers [17]. The assessment of broader economic benefits has been considered beyond the scope of this report, and one would expect the benefits to be difficult to state with precision. However, such an assessment might be indicated if it becomes necessary to clarify the relevance of terminal automation to a wider cross section of the nation's population.

The collection of more detailed surveillance data, at different airports and in varying meteorological conditions, would give better confidence in the accuracy and range of applicability of capacity increases that are claimed for terminal automation. It should be noted, however, that the processing and proper analysis of such operational data tends to be time consuming and expensive, particularly when sites not previously studied are involved.

Both for benefits assessment and for refining the control mechanisms in the CTAS logic, it would be desirable to have a better quantitative understanding of the interdependencies and considerations that govern wide-area traffic management and influence the coordination (or lack of coordination) between en route traffic delivery and flow inside the TRACON. It may be worthwhile to analyze detailed surveillance data both inside and on the periphery of the TRACON, in order to supplant the mixture models of Section 3.2.2.3 with a more concrete description of flow dynamics, stated in more operational terms. Indications are that a useful monitoring capability can be developed to support effective metering and flow management strategies in real time. There is limited quantitative information available regarding control ability and piloting variation inside the final approach fix, especially in cases of marginal and changing weather. It may be desirable to obtain surveillance data inside the final approach fix and use it to review current practice, so as to determine what capabilities CTAS must accommodate near the runway. Finally, in future studies of the benefits of CTAS, it will be advantageous to assess the incremental benefit of supplying CTAS with improved forecasts of meteorological variables such as wind, ceiling, and visibility.

APPENDIX 1.

NOTES ON ATC AUTOMATION STUDIES

Following are notes made concerning specifics of published ATC automation studies, as cited in Table 14:

- a The two-tier results for computer-aided IAT standard deviation correspond to two different experimental circumstances, one in which an variable final approach airspeed is used, specific to each aircraft (the bottom value), and one in which a constant value, not aircraft-specific, is used for final approach speed in scheduling and spacing calculations.
- b The weather simulated was a ceiling of approximately 200', RVR of 0.5 nmi. Also, a wind with gradient by altitude was used, statistically tailored to common winds at the simulated site, which was Denver.
- c There were 16 runs conducted before a briefing in which pilots were informed about the TIMER logic and simulation objectives, 19 runs performed after the briefing.
- d The intent of the simulation was to determine delivery time precision (i.e., conformance to a scheduled runway threshold crossing time) achievable with conventionally equipped aircraft, lacking a 4D flight management system (FMS).
- e Delivery precision was also measured at the final approach fix. Standard deviations at the FAF were 9.1 seconds, before the TIMER briefing, and 5.4 seconds after (as compared with 9.7 seconds and 7.0 seconds, respectively, at the runway threshold).
- f This simulation included conscientious efforts to model the impact of system uncertainties, including piloting variations, aircraft weight and final approach speeds, communication and response lags.
- g In the course of the experiment there were 605 so-called "alpha" intervals, which, similarly to the unstream late-feed arrivals discussed in this report in Section 3.3.1, were judged by the computer scheduling function to be unable to achieve a desired minimum separation from the preceding aircraft. Of these, 290 occurred in manual operations, and 318 occurred during CAAS operation. Precision measures were calculated using only the remaining non-alpha arrivals, 455 of them manual and 512 computer-aided. Note that the CAAS final report variously refers at times to 458 and at other times to 455 non-alpha manual arrivals. We use 455 as being more consistent with other cross-tabulations (e.g., Table XV in Martin and Willett [21]).
- h Crossing intervals were measured at the final approach fix rather than the runway threshold.
- i An additional set of trials was conducted with the passive spacing aid ("slot markers") in force, and with slight perturbations in the release times from one of the three feeder sectors. The resulting standard deviation was 9.1 seconds.
- j Capacity gains are not specifically discussed in the report, but support of an IAT standard deviation of approximately 12 seconds implies support of approximately a 16% IMC arrival capacity gain, as discussed in the fast simulation report.

- k Each approach was run in isolation, aiming for a desired spacing behind a computer-generated aircraft, rather than in the context of an ongoing stream of merging traffic.
- l These simulations employed a Boeing 707 cockpit simulator with an air traffic situation display (ATSD), showing the position, flight id, and other selected ARTS information regarding nearby aircraft, and with instrument display of computer-calculated metering and spacing commands. The top figure for IAT standard deviation shows the dispersion observed with the ATSD overlaid on a moving RNAV map, but with a relatively crude, open-loop delivery of metering and spacing commands. The bottom figure for IAT standard deviation occurred with an improved HSI that permitted metering and spacing advisories to continually update to maintain a desired time of arrival at the next waypoint, and so enabled a form of closed loop control. The closed loop results were achieved in spite of simulated failures of the ARTS ground system at critical points during final approach. See Connelly [4,5].
- m Tests were conducted using the Terminal Configured Vehicle simulator (TCV) at NASA Langley Research Center, and a Cockpit Display of Traffic Information (CDTI). Each piloted approach followed a computer-generated aircraft of the same type as ownship. Two factorial experiments were performed, which examined the effect of different ground speed quantization levels and different lead ship approach speeds/deceleration profiles. Measurement was made of the self-spacing performance of pilots using CDTI, and the figure of 8.1 seconds is the IAT standard deviation averaged over all the combinations of variable test conditions. See Kelly [15].
- n In a separate study, piloted simulation tests with the TCV indicated that self-spaced approaches can function even in long queues of CDTI-equipped aircraft without inducing the oscillatory or accordion-like effect observed in automobile traffic.
- o There were 30 runs of 5 aircraft each, conducted at the airport at the National Aviation Facilities Experimental Center (NAFEC), now called the FAA Technical Center (FAATC) at Atlantic City, NJ. Thus there were 120 intervals in all.
- p The standard deviation for the manual system was obtained from a separate pilot study (Martin, D. [19,20]).
- q Capacity gains are discussed briefly in the final report [21]. It appears that mean interarrival times were slightly more than 5 seconds less with CAAS automation active than without, and on this basis a capacity increase of about 6% was observed in the field. Incorporating further exploitation of the increased IAT uniformity, the authors hypothesized potential capacity gains of up to 16%.

GLOSSARY

AAR	Airport Acceptance Rate. The number of arrivals per hour that the en route system will deliver to an airport's terminal airspace.
ARTCC	Air Route Traffic Control Center. The control complex for en route ATC.
ARTS	Automated Radar Terminal System. The computing system that supports terminal area ATC.
ATC	Air Traffic Control
ATOMS	Air Traffic Operations Management System. A system based on inputs from FAA operational staff, who catalog delays within their airspace. Only delays exceeding 15 minutes are recorded, but, unlike the SDRS, delay measurements are accompanied by a rough attribution of cause.
CAAS	Computer-Automated Approach System. A terminal automation program fielded at JFK International Airport from December 1966 to April 1967.
CDTI	Cockpit Display of Traffic Information
CFCF	The FAA Central Flow Computing Facility
CRDA	Converging Runway Display Aid
CTAS	Center-TRACON Automation System
DA	Descent Advisor. A component of CTAS.
DOC	Direct Operating Cost
DTP	Dynamic Time-based Planner. An automation system concept and demonstration program developed at Lincoln Laboratory, MIT
EDT	Eastern Daylight Time
ERM	En Route Metering. An operational software system that times aircraft transitions from en route to terminal airspace in order to insure a uniform entry rate matched to the AAR.
FAA	Federal Aviation Administration
FAF	Final Approach Fix
FAST	Final Approach Spacing Tool. A component of CTAS.
Host	The main computer, or the computing system, in an ARTCC.
HSI	Human-System Interface
IAT	Inter-arrival time. The time elapsed between successive runway threshold crossings. Also called the Landing Time Interval (LTI).

ILS	Instrument Landing System
IMC	Instrument Meteorological Conditions
LCL	Local Time
LLWAS	Low Level Wind shear Alerting System
LTI	Landing Time Interval. The time elapsed between successive runway threshold crossings. Also called the Inter-arrival time (LAT).
MGTOW	Certificated maximum gross takeoff weight of an aircraft.
MIT	Massachusetts Institute of Technology
NAS	National Airspace System.
PRM	Precision Runway Monitor
PVD	Plan View Display. The image displayed on a radar controller's scope, the geometric perspective implicit in the projection of this display, or the display equipment itself.
ROT	Runway Occupancy Time. The time elapsed between when an arrival aircraft crosses the runway threshold, and when it turns off the runway surface.
ODOC	Other (non-fuel) Direct Operating Costs
SDRS	Standardized Delay Reporting System. A system administered by the FAA with the participation of three major air carriers for measuring and obtaining statistics on air traffic delay. Delays are measured in one minute increments and by phase of flight.
TATCA	Terminal Air Traffic Control Automation
TIMER	Traffic Intelligence for the Management of Efficient Runway Scheduling. A terminal automation concept and demonstration program developed at NASA Langley.
TMA	Traffic Management Advisor. A component of CTAS.
TRACON	Terminal Radar Approach Control Facility. The radar controllers and control system around an airport, and, colloquially, the airspace they control.
TTMA	Terminal Traffic Management Advisor
VMC	Visual Meteorological Conditions

REFERENCES

1. 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, The Sheraton Harbor Island Hotel, San Diego, CA, 23-25 May, 1990, pp. 535-542.
2. Branscome, C. W. Benefit Analysis of IMC Converging Runway Operations at Boston's Logan International Airport. Mitre Corporation, MTR-87W125, August 1987. Conducted for the FAA Systems Engineering Service, Contract Number DTFA01-84-C-00001.
3. Butcher, A., Cheslow, M., Nelson, G., Reid, D., Russell, H., Sinnott, J., Walker, G. "The Advanced Automation System: A Benefit/Cost and Risk Analysis", Volume IV, The Mitre Corporation, MTR-87W00235-04, December 1987.
4. Connelly, Mark E. "Applications of the airborne traffic situation display in air traffic control," AGARD Conference Proceedings No. 188, Plans and Developments for Air Traffic Systems, ed. A. Benoît and D.R. Israel, Cambridge, MA 20-23 May 1975.
5. Connelly, Mark E. Simulation Studies of Airborne Traffic Situation Display Applications -- Final Report. Report ESL-R-751, Electronic Systems Laboratory, Massachusetts Institute of Technology, May 1977.
6. Credeur, L. and Capron, W.R., "Simulation Evaluation of TIMER, a Time-Based Terminal Air Traffic, Flow-Management Concept", NASA Technical Paper 2870, Langley Research Center, February 1989.
7. Dear, R.G. The Dynamic Scheduling of Aircraft in the Near Terminal Area. MIT Flight Transportation Laboratory Report R76-9, September, 1976.
8. Erzberger, H. and Nedell, W. "Design of Automated System for Management of Arrival Traffic", NASA Technical Memorandum 102201, Ames Research Center, June 1989.
9. Federal Aviation Administration. Air Traffic Control. Order 7110.65E, Washington, D.C.: Department of Transportation.
10. Federal Aviation Administration. FAA Aviation Forecasts, Fiscal Years 1990-2001. FAA-APO-90-1, March 1990.
11. Federal Aviation Administration. Terminal Area Forecasts, FY 1990-2005. FAA-APO-90-6, July 1990.
12. Federal Aviation Administration. FAA Long Range Aviation Projections, Fiscal Years 2001-2020. FAA-APO-90-9, August 1990.
13. Federal Aviation Administration. 1990-91 Aviation System Capacity Plan. DOT/FAA/SC-90-1, September 1990.
14. Geisinger, K. "Airline Delay: 1976-1986, Based upon the Standardized Delay Reporting System" FAA-APO-88-13, March 1989.

15. Kelly, J.R. "Effect of Lead-Aircraft Ground-Speed Quantization on Self-Spacing Performance Using a Cockpit Display of Traffic Information", NASA Technical Paper 2194, October, 1983.
16. Lebron, J. E. "Estimates of Potential Increases in Airport Capacity Through ATC System Improvements in the Airport and Terminal Areas", FAA-DL5-87-1, October 1987.
17. Lewis, D. "Primer on Transportation, Productivity and Economic Development", National Cooperative Highway Research Program Report 342, Transportation Research Board, National Research Council, September 1991.
18. Luenberger, R. A. "A Traveling-Salesman-Based Approach to Aircraft Scheduling in the Terminal Area", NASA TM-100062, March 1988.
19. Martin, D.A. "A Simulation Study of a Final Approach Spacing Concept Using a Manual Computer", Project No. 101-25X, Experimentation Division, Systems Research and Development Service, September 1963.
20. Martin, D.A. "Field Application of a Final Approach Spacing Concept Using a Manual Computer", Project No. 101-25X, Part B, Experimentation Division, Systems Research and Development Service, July 1964.
21. Martin, Donald A. and Willett, Francis M., "Development and Application of a Terminal Spacing System", Report No. NA-68-25, National Aviation Facilities Experimental Center, August 1968.
22. Mundra, Anand D. "A Description of Air Traffic Control in the Current Terminal Airspace Environment", Report No. MTR-88W00167, The Mitre Corporation, McLean, VA. March, 1989.
23. Neter, J., Wasserman, W., and Kutner, M. Applied Linear Statistical Models: Regression, Analysis of Variance, and Experimental Designs, 3rd Ed. Richard D. Irwin, Inc. 1990.
24. Picardi, M. C. "Test Report on a Candidate Human-System Interface for the Terminal Air Traffic Control Automation Final Approach Spacing Tool", Laboratory Technical Memorandum No. 41L-0380, Lincoln Laboratory, MIT, June 1991.
25. Schick, F., Schubert, M., Voelckers, U., "Simulation Test of the Planning System COMPAS", Deutsche Forschungs- und Versuchsanstalt für Luft und Raumfahrt, 1986.
26. Simpson, R., Odoni, A. and Salas-Roche, F., "Potential Impacts of Advanced Technologies on the ATC Capacity of High-Density Terminal Areas", NASA Contractor Report 4024, October 1986.
27. Snedecor, G. and Cochran, W. Statistical Methods, 7th Ed. The Iowa State Press, 1980.
28. Sorensen, J.A., Shen, M., and Hunter, C.G. Analysis of Final Approach Spacing Requirements, Part I. Report 91112-01, Seagull Technology, Inc. January 1991.

29. Thompson, S. D. "Aircraft Landing Sequence Optimization Under Constraints", Lincoln Laboratory Technical Memorandum, 41PM-TATCA-0001, August 1990.
30. Titterton, D.M., Smith, A.F.M., and Makov, U.E. Statistical Analysis of Finite Mixture Distributions. John Wiley & Sons, 1985.
31. Venkatakrishnan, C.S. "Analysis and Optimization of Terminal Area Air Traffic Control Operations", Ph.D. dissertation, Sloan School of Management, Massachusetts Institute of Technology, February, 1991.