NASA/TM–2012-217762

Improvement to Airport Throughput Using Intelligent Arrival Scheduling and an Expanded Planning Horizon

Patricia C. Glaab
Langley Research Center, Hampton, Virginia

September 2012

NASA STI Program . . . in Profile

Since its founding, NASA has been dedicated to the advancement of aeronautics and space science. The NASA scientific and technical information (STI) program plays a key part in helping NASA maintain this important role.

The NASA STI program operates under the auspices of the Agency Chief Information Officer. It collects, organizes, provides for archiving, and disseminates NASA's STI. The NASA STI program provides access to the NASA Aeronautics and Space Database and its public interface, the NASA Technical Report Server, thus providing one of the largest collections of aeronautical and space science STI in the world. Results are published in both non-NASA channels and by NASA in the NASA STI Report Series, which includes the following report types:

- TECHNICAL PUBLICATION. Reports of completed research or a major significant phase of research that present the results of NASA Programs and include extensive data or theoretical analysis. Includes compilations of significant scientific and technical data and information deemed to be of continuing reference value. NASA counterpart of peer-reviewed formal professional papers, but having less stringent limitations on manuscript length and extent of graphic presentations.

- TECHNICAL MEMORANDUM. Scientific and technical findings that are preliminary or of specialized interest, e.g., quick release reports, working papers, and bibliographies that contain minimal annotation. Does not contain extensive analysis.

- CONTRACTOR REPORT. Scientific and technical findings by NASA-sponsored contractors and grantees.

- CONFERENCE PUBLICATION. Collected papers from scientific and technical conferences, symposia, seminars, or other meetings sponsored or co-sponsored by NASA.

- SPECIAL PUBLICATION. Scientific, technical, or historical information from NASA programs, projects, and missions, often concerned with subjects having substantial public interest.

- TECHNICAL TRANSLATION. English-language translations of foreign scientific and technical material pertinent to NASA's mission.

Specialized services also include organizing and publishing research results, distributing specialized research announcements and feeds, providing information desk and personal search support, and enabling data exchange services.

For more information about the NASA STI program, see the following:

- Access the NASA STI program home page at *http://www.sti.nasa.gov*

- E-mail your question to help@sti.nasa.gov

- Fax your question to the NASA STI Information Desk at 443-757-5803

- Phone the NASA STI Information Desk at 443-757-5802

- Write to:
  STI Information Desk
  NASA Center for AeroSpace Information
  7115 Standard Drive
  Hanover, MD 21076-1320

NASA/TM–2012-217762

Improvement to Airport Throughput Using Intelligent Arrival Scheduling and an Expanded Planning Horizon

Patricia C. Glaab
Langley Research Center, Hampton, Virginia

National Aeronautics and
Space Administration

Langley Research Center
Hampton, Virginia 23681-2199

September 2012

Available from:

NASA Center for AeroSpace Information
7115 Standard Drive
Hanover, MD 21076-1320
443-757-5802

Contents

List of Figures

Acronyms

ACC Arrival Capacity Calculator

ACES	Airspace Concepts Evaluation System
ATC	Air Traffic Control
ASAB	Aeronautics Systems Analysis Branch
BADA	Base of Aircraft Data
CD&R	Conflict Detection and Resolution
CDA	Continuous Descent Arrival
FDS	Flight Data Set
IAI	Intelligent Automation, Inc.
IDEAS	Intelligently Delayed and Expedited Arrivals Scheduler
M&S	Merging and Spacing
MF	Metering Fix
NAS	National Airspace System
OPD	Optimum Profile Descent
ROC	Radius of Control
RTA	Required Time of Arrival
STAR	Standard Terminal Arrival Route
Std	Standard
TFM	Traffic Flow Management
TOD	Top of Descent
TRACON	Terminal Radar Area Control
TSAM	Transportation Systems Analysis Model

Abstract

The first phase of this study investigated the amount of time a flight can be delayed or expedited within the constraints of the Terminal Airspace using only trajectory speed changes. The Arrival Capacity Calculator analysis tool was used to predict the time adjustment envelope using a generalized speed profile for standard descent arrivals and then for Continuous Descent Arrivals (CDA). These results ranged from 0.77 minutes to 5.38 minutes. STAR routes were configured for the Airspace Concepts Evaluation System (ACES) simulation, and a validation of the time adjustment envelopes found by the ACC analysis tool was conducted comparing the maximum predicted time adjustments to those seen in ACES.

The final phase investigated full runway-to-runway trajectories using the ACES simulation tool. The radial distance used by the arrival scheduler was incrementally increased from 50 nautical miles (nmi) to 150 nmi. The increased Planning Horizon radii allowed the arrival scheduler to arrange, path stretch, and speed-adjust flights to more fully load the arrival stream. The average throughput for the high volume portion of the day increased from 30 aircraft per runway for the 50 nmi radius to 40 aircraft per runway for the 150 nmi radius for a traffic set representative of a high volume day in projected 2018. The maximum useful radius for the arrival scheduler's Planning Horizon was found to be 130 nmi, which allowed more than 95% loading of the arrival stream.

THIS PAGE INTENTIONALLY LEFT BLANK

Introduction

This paper documents the simulation and analysis results for improvement to airport throughput gained with intelligent approach and landing scheduling, coordinated trajectory speed management, and an expanded arrival scheduling Planning Horizon for large scale arrival flow scenarios. The arrival scheduler's Planning Horizon is defined to be the radial distance from the destination airport at which the scheduler first considers the range of arrival options for any given flight and issues instructions for that flight to follow to the runway. The significance of throughput improvement obtained in this analysis is that it was achieved without any changes to current day wake spacing regulations, and without any arrival route lateral geometry or vertical profile modifications. All improvements are a result of opportunistic trajectory speed management to fully load the arrival stream, made possible by early scheduling to increase flexibility of leader/follower or RTA assignments to interim and final fixes.

This work was conducted as part of a larger milestone investigating the interoperability of advanced airspace management systems such as Merging and Spacing (M&S), departure and en route Traffic Flow Management (TFM), and Conflict Detection and Resolution (CD&R), particularly focused on the En Route to Terminal transition supporting the Airspace Systems Program. The study was divided into 3 phases to allow early work to commence while tools needed for later parts of the study continued development. This work detailed in this paper spanned Phases 1 and 2 of the experiment with results found in Phase 1 for trajectory analysis within the Terminal Airspace feeding into the design of the Phase 2 investigation of the size of the arrival scheduling Planning Horizon on the realizable airport throughput.

1.1 Experimental Phases

For high volume airspace, the ability to adjust aircraft arrival times to target fixes potentially allows an intelligent planning tool to more efficiently load arrival streams to increase throughput. If the pre-adjusted aircraft are travelling at their most fuel-efficient speeds, this obviously incurs fuel cost which must be weighed against the throughput improvement. The studies discussed here investigated this relationship between speed adjustment, recoverable time achievable by using these adjustments, and the impact of this scheduling flexibility when combined with an intelligent arrival scheduler planning tool on arrival throughput improvement.

Tasks conducted under Phase 1 focused on the Terminal Airspace, defined here to be the region of airspace spanning a 40 nautical mile (nmi) radius around the airport. Since the Standard Terminal Arrival Route (STAR) geometries vary through this airspace, Phase 1 targeted understanding the possibilities and limitations for managing Required Time of Arrival (RTA) deadlines exclusively through speed changes for aircraft flying STAR routes between the arrival fix and the runway threshold. Speed adjustments are particularly useful in the confined Terminal Airspace where arrival and departure corridors do not allow indiscriminant path stretching or holding patterns. In both the analysis cases and in the simulation runs, the speed adjustment was limited to +/-10% of the default trajectory crossing speed for each arrival route fix (including the threshold).

Phase 2 quantified the recoverable time to an arrival target within a 10% groundspeed adjustment for conditions representative of en route flight. The 10% adjustment used in this analysis was selected to be consistent with ACES simulation speed adjustments used later in the experiment. This analysis assumed a simple constant speed case to identify the range of recoverable times possible as a starting point to inform a TFM planning window. In this case, the arrival target would correspond to a Top of Descent (TOD) point since the analysis assumes a constant speed to the target location. Real-world arrival targets would likely also include a metering fix at the TRACON boundary for the arrival airport, which would have to account for changing speeds during descent.

Phase 3 of the experiment compiled knowledge gained in the initial two phases to test the results in a simulation environment that included aircraft performance considerations and runway-to-runway trajectories. These tests initially focused on individual flights to validate results seen in the analysis tests. Next, full days of arrival traffic were simulated in scenarios that varied the size of the Planning Horizon to correlate size of the Planning Horizon to the arrival scheduler's ability to efficiently schedule flights to load the arrival stream. The simulation results demonstrated a significant impact of the Planning Horizon size on the scheduler's ability to maximize airport throughput. A final validation phase for the Terminal Airspace under conditions observed for those simulation runs was conducted to verify the potential achievement of these higher throughput numbers under current wake spacing regulations and STAR route considerations.

1.2 Software Tools

Three software tools were used to conduct the study. The Arrival Capacity Calculator (ACC) was used to perform mathematical analysis on the configured Terminal Airspace for Phase 1. This tool uses parameterized models to mathematically determine the size of the throughput envelope under varying conditions. Results from the Phase 1 analysis defined the maximum recovery times possible as a function of arrival route length within the Terminal Airspace based on general trajectory speed profiles. These results show potential recovery time by distance for Standard Descent arrivals. The ACC tool was also used for the final validation of the ACES simulation runs in the rolled up Phases 1 and 2 testing.

The Airspace Concepts Evaluation System (ACES) simulation [1] was used to test scenarios for all phases using realistic trajectories based on aircraft performance capabilities and in consideration of physical flight paths through the airspace. The version of ACES used for this study was the January 2012 delivery of ACES with Merging and Spacing (M&S). This version of the simulation contains an intelligent scheduling tool developed by Intelligent Automation Inc. (IAI). It allows the Traffic Flow Management (TFM) component of ACES to operate concurrently with the M&S software.

The ACES Viewer tool [2] was used to inspect the simulation runs in post-processing. ACES Viewer is a java tool that reads MySQL files containing latitude, longitude, altitude, and time data points and presents the information as 4-D animations. This tool was used to verify the flown paths and to identify high volume flow characteristics.

1.3 The Intelligent Arrival Scheduler

An intelligent arrival scheduler was critical to obtaining the throughput improvements seen in the final phase of this study. The scheduler contained in the M&S software is the Intelligently Delayed and Expedited Arrivals Scheduler (IDEAS) by Intelligent Automation, Inc. (IAI).[3] The IDEAS scheduler was initially developed in 2009, but its methodologies have adapted to prioritize high efficiency scheduling techniques observed in the research studies for which the scheduler has been used. In this scheduling algorithm, aircraft are handled on a first-come-first-*scheduled* basis. Being scheduled earlier, however, does not necessarily imply landing earlier since later arrivals may fill earlier gaps in the stream, be assigned routes with less traffic, or travel at higher speeds. IDEAS prefers to pair an approaching aircraft to a leader when possible, and will then expedite the follower to minimize the trailing distance. Leader/follower pairs are not required to stay paired, and can diverge (to different runways, for example) to achieve an earlier landing time. If no previously scheduled flight is a viable leader for a new aircraft being handled, IDEAS schedules the flight on a waypoint-by-waypoint basis to the threshold using Required Time of Arrival (RTA) deadlines. The ACES M&S software also applies a non-static safety buffer that is approximately 5% of the wake spacing.

The leader/follower pairing technique improves throughput potential by consolidating flights early on the arrival route. If the flights were not paired, the traffic would consolidate itself anyway as it moved down the routes and faster followers gaining ground on slower leaders. When the traffic can be consolidated early into groupings of aircraft, otherwise wasted space grows to usable sizes for new slots for subsequently arriving or merging flights. The IDEAS scheduler is then able to identify gaps large enough to accommodate merging aircraft, and the usage of the available arrival stream space is improved.

1.4 Airport Selection

A set of arrival routes for Atlanta's Hartsfield-Jackson International Airport (Atlanta Airport, for short) was configured for use throughout this experiment. Atlanta Airport serves as a useful model for test scenarios since arrivals are completely de-conflicted with departures through the design of the STARs used for approach. Atlanta also uses dedicated arrival or departure runways and handles high volume very efficiently, providing real-world validation options for configured scenarios. All runways are parallel and independent, which removes the need to model runway interactions to simulate the airport. This low complexity of the Atlanta airport and airspace reduces the complexity required to modeling the system in simulations, which also reduced development efforts. Finally, the reduced simulation complexity simplifies post-analysis of results, making it an ideal starting point for the investigation of concepts and interactions.

For Section 4.4, Determination of the Maximum Useful Planning Horizon Radius, routes to model arrivals into Dallas Fort Worth International Airport (DFW) landing south were configured. DFW uses a simple 4 corner post arrival configuration like Atlanta, and typically operates with dedicated arrival and departure runways. The DFW scenarios were configured to use one or two runways to provide a contrast to the 3-runway setup used for Atlanta.

2 Phase 1 – Terminal Airspace

2.1 Arrival Route Configuration

The arrival route pattern for this test was configured to match airport operations for current-day Atlanta Airport with westward flow. In this configuration, landing aircraft are headed almost due west as they touchdown on runways 26R, 27L, or 28 (Figure 1).

Figure 1: Configuration for Atlanta Airport with 3 Runways and Westward Flow

Traffic approaching from the northeast (NE) quadrant generally uses the FLCON arrival route. Traffic from the northwest (NW) quadrant uses the ERLIN route. Southeastern (SE) traffic uses the CANUK route, and southwestern (SW) traffic uses the HONIE route (Figure 2, left). The approaching traffic from the eastern half of the airport (combined NE and SE) has a shorter path since vectoring takes them from the bottom of their STAR route directly to the final runway approach line. These arrivals are referred to as using the "short side" of the airport since they avoid the extra flight path required to fly the downwind and base legs. Conversely, the "long side" traffic in this configuration is from the NW and SW. Extension of the final approach segment will have a greater impact on the long side traffic since each extra mile of final also requires a similar extension of the downwind segment.

In the published STAR definitions for these Atlanta routes, each STAR begins beyond the TRACON boundary with the fix of the same name (i.e., the ERLIN7 STAR begins at the ERLIN waypoint). Unlike class airspace or Mode C boundaries, TRACON boundaries are irregular in both shape and altitude to best accommodate arrival and departure traffic to airports within. For this experiment, a simple 40 mile radius circle was used to represent the Terminal Airspace. An experimental waypoint was created at each ordinal compass point along the boundary of this circle to serve as the experimental Metering Fix (MF) points. The MF is the location that the TFM uses to compute time deadlines for arrival to the terminal airspace. The actual STAR routes were configured, but terminated at these experimental MF points to simplify the setup and analysis (Figure 2, right).

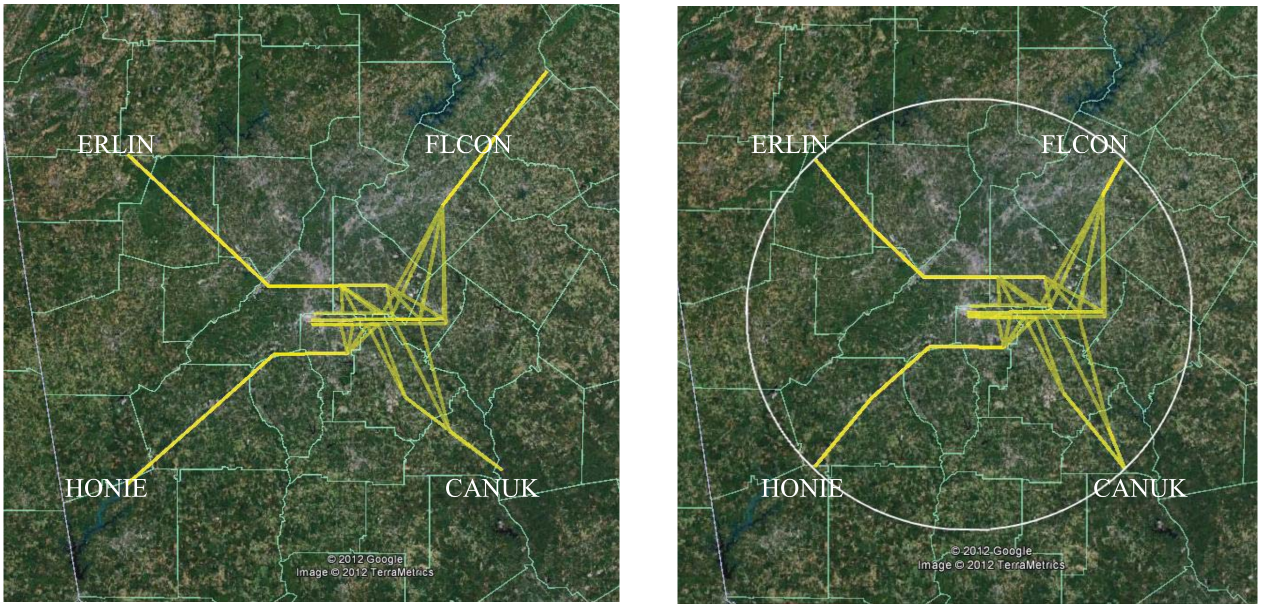

Figure 2 – Original (left) and experiment-limited (right) routes

The longest and shortest total path distances for each route were selected for the validation testing (using the ACES simulation), and are summarized in Table 1.

Route	Total Path Distance (nmi)
ERLIN_26R_5	46.43
ERLIN_27L_25	76.92
FLCON_26R_12	37.95
FLCON_27L_25	46.75
CANUK_28_10	37.93
CANUK_27L_20	43.46
HONIE_28_5	47.89
HONIE_28_12	56.14

Table 1 - Path Distance by Route

2.2 Mathematical Analysis of the Terminal Airspace

2.2.1 Standard Descent Arrival Route Analysis

2.2.1.1 Estimation of Time Recoverable Within TRACON

The first step of this phase performed a mathematical analysis using the ACC tool to determine the amount of delivery delay to the MF that can be tolerated and compensated for within the TRACON itself. This value is important because it provides a target delivery envelope to the MF at the TRACON boundary for an en route scheduler to measure against. The ability of a real-world airport to fully achieve these maximum time recovery values will depend heavily on the arrival route structure's ability to fully support necessary speed changes. For example, ATC might separate aircraft by altitude for a simple 4-cornerpost arrangement to allow some aircraft to expedite travel in the outer portions of the routes. Alternatively, parallel offsets might be used, or the activation of addition arrival routes within the TRACON with a final merge near the base leg.

To determine the amount of time recoverable, the aircraft speed profile from the ACC tool was used to compute the time to travel, evaluated at 1 nmi increments between the threshold and 100 nmi. The average velocity for each interval was used, and the velocity equation was solved for time. Time increments were summed to each distance point to obtain the travel time as a function of distance curve for the Standard Descent trajectory profile.

$$v = \Delta d / \Delta t$$
$$\Delta t = \Delta d / v$$
$$v = (v_{i+1} - v_i)/2$$
$$\Delta d = d_{i+1} - d_i$$
$$t = \sum_{i=1}^{n} (d_{i+1} - d_i)/(v_{i+1} - v_i)$$

Two additional time-to-go curves were then created for this speed profile using the default airspeeds +/-10%, which is the percentage typically used to limit the allowed simulation speed change. The resulting default, slowest, and fastest path times based on distance-to-go on the arrival route are presented in Figure 3 - Travel Times for a Standard Descent Profile Arrival Route.

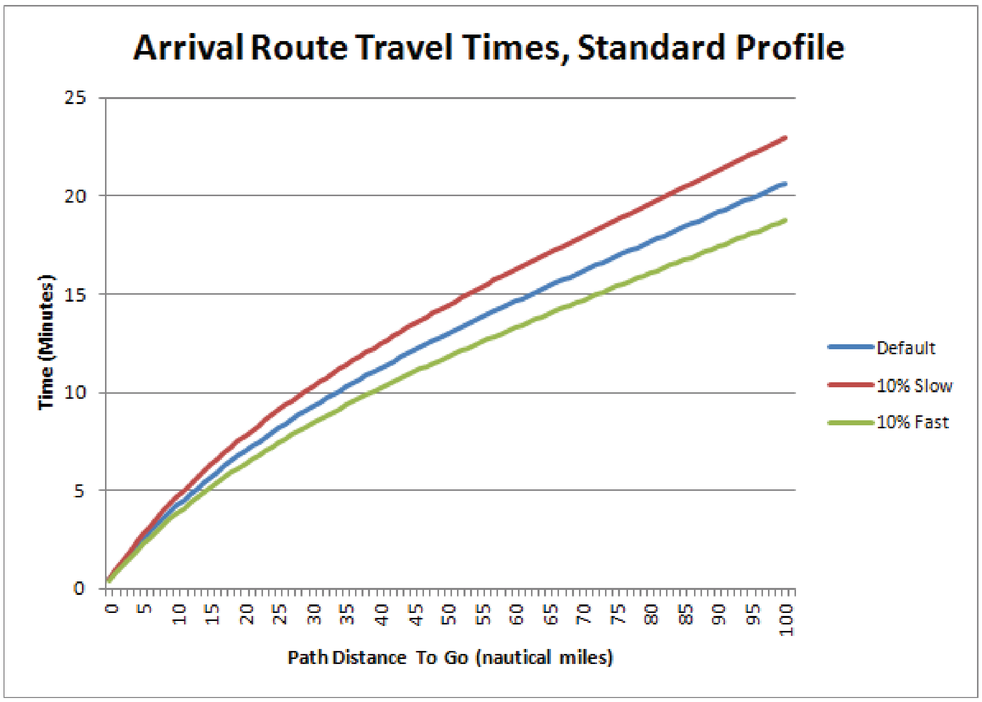

Figure 3 - Travel Times for a Standard Descent Profile Arrival Route

The maximum recoverable time by path distance-to-go can then be computed based on the differences in the times-to-go for the three trajectories (Figure 4).

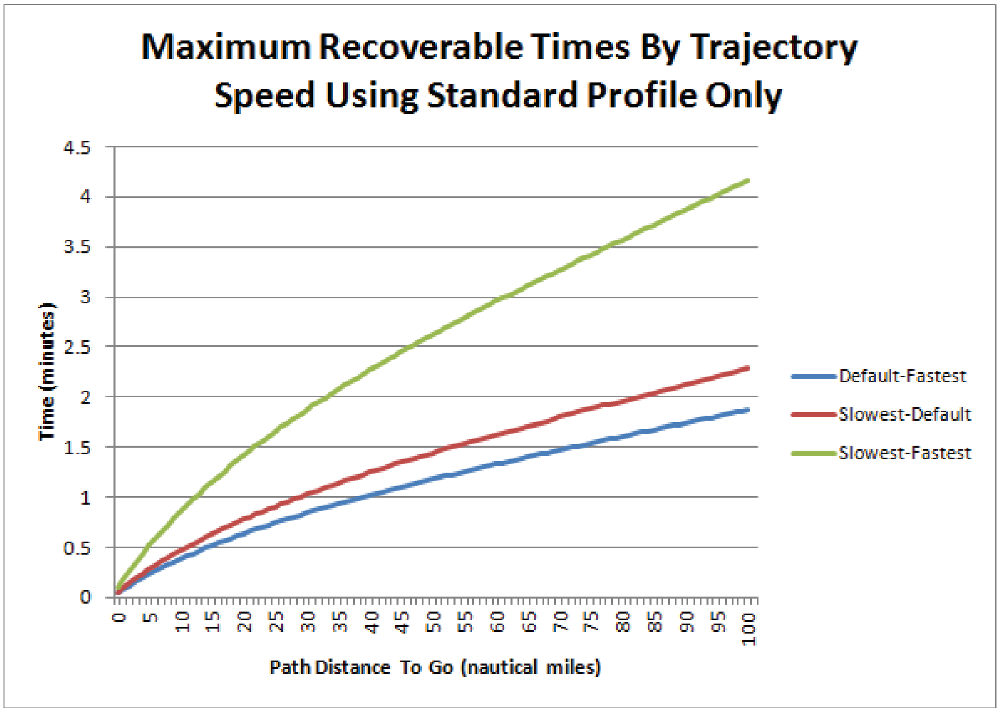

Figure 4 - Maximum Recoverable Times For Standard Descent Trajectories by Distance

The minimum and maximum path distances for each STAR arrival route in the experiment were computed previously in Table 1 - Path Distance by Route. Based on these minimum/maximum distances, the corresponding maximum recoverable times for the shortest and longest routes can be assigned for the default, slowest, and fastest profiles using the curves in Figure 4. The results presented in Table 2 are a summary of the recoverable time ranges achievable by switching from one profile type to another for each arrival route in the experiment. Though the results are shown by route to provide a scenario basis for the simulated runs for the next step, the recoverable times are not route-specific, but rather are profile-specific. However, presentation of the results by route is useful for visual correlation of the time savings achievable with typical Terminal Airspace paths.

Arrival Direction	Route Name	Approximate Path Distance Range (nmi)	Maximum Recoverable Times for Std Profile (minutes)		
			Default -> Fastest (shortest/longest)	Default -> Slowest (shortest/longest)	Slowest -> Fastest (shortest/longest)
NE	FLCON	38/47	0.99/1.13	-1.21/-1.39	2.19/2.52
NW	ERLIN	46/77	1.12/1.57	-1.37/-1.91	2.48/3.48
SE	CANUK	38/43	0.99/1.07	-1.21/-1.31	2.19/2.38
SW	HONIE	48/56	1.15/1.27	-1.40/-1.55	2.55/2.82
Max Range	(All)	38/77	0.99/1.57	-1.21/-1.91	2.19/3.48

Table 2 – Maximum Times Recoverable through Speed Change between Metering Fix and Threshold, Standard profiles

2.2.1.2 Results for Standard Descent Profile

Predictably, the routes with the longest path distance from MF to threshold provide the greatest opportunity for time recovery since there is more flight distance to work with. In the case of the transition from a default speed to a fastest speed, "Default -> Fastest" in Table 2 (which would correlate to the real-world situation of ATC expediting arrival traffic in the TRACON by issuing speed holds to descending aircraft to limit deceleration), the maximum recoverable times ranges from 0.99 to 1.57 minutes depending on route length. This means that an aircraft that is planning a default trajectory speed profile (probably a fuel efficient profile) and arriving up to approximately 1 minute late to the short side Metering Fix could still be expedited to the threshold to meet its original default touchdown RTA simply by using a 10% speed increase within the TRACON. For aircraft using the long side of the airport (downwind, base, and final pattern), up to 1.57 minutes can be made up in the TRACON with this type of speed adjustment.

The results shown for the profile transition from default to slowest, "Default -> Slowest" in Table 2, would correspond to a real-world scenario of ATC issuing speed reduction instructions. In this scenario, aircraft would have arrived to the MF early and needed to adjust their speed to avoid arriving early to the runway. In this case, an aircraft arriving up to 1.21 minutes early to the short side MF can be delayed using a 10% speed reduction to prevent early arrival to the threshold RTA. For routes using the long side, aircraft can be speed-adjusted to delay by up to 1.91 minutes.

The final column in Table 2 defines the maximum range for the routes investigated. This scenario would probably apply to an aircraft that was trying to delay using a slower speed to meet its original RTA, and was then offered an earlier slot to the runway. In high volume IFR conditions, aircraft can touch down to the runway with approximately 1.2 to 1.5 minutes separating them. This means that for the results shown in Table 2 (a range of 2.19 to 3.48 minutes of recovery time), aircraft could be accommodate a reassignment to slots 2 or 3 positions earlier at the runway to fill an otherwise wasted arrival slot.

Generally speaking, this demonstrates that the TRACON has considerable flexibility in accommodating delivery inaccuracy to the MF. This requires, of course, that the route structure within the TRACON has enough flexibility to allow the necessary speed changes to realize the full potential. Current day ATC already perform this type of delay or expedition of traffic to gain throughput during high volume operations, so the notion that traffic can be delayed or expedited to some degree is not novel. The numbers in Table 2 and the graphs in Figure 3 and Figure 4 simply define the range of possibility based on mathematical analysis.

2.2.2 Continuous Descent Arrival (CDA) Route Analysis

2.2.2.1 Estimation of Time Recoverable Within TRACON

The ACC analysis tool provides several generalized profiles, including a profile representing Continuous Descent Arrivals (CDA) created using the ACES simulation configured to fly this type of descent. The term Optimum Descent Profile (OPD) is more common in current literature, and is often used interchangeably with CDA, though there are subtle differences in the two terms. A CDA can be any profile that uses a continuously descending path between the TOD and the landing. With level-offs removed, these approaches tend to save fuel and reduce noise measured on the ground since microphones are crossed at higher altitudes and lower throttle settings are employed. An OPD, by contrast, is a CDA profile that is optimized to provide the lowest fuel burn possible between TOD and touchdown. These trajectories almost always use flight idle throttle settings which results in steep descents. Additionally, OPD paths are usually constructed to target a more efficient and direct route without being confined to the downwind/base/final pattern of current day routing. Unfortunately, the BADA performance database contained in the ACES simulation (upon which the generic profile for the ACC tool is based) does not provide a climb or descent rate for "optimal", so the term OPD cannot be applied in good conscience. Also, the path geometries for standard descent approaches were retained with the only change for this part of the study being the altitude and speed profile. So for this application, the term CDA is most appropriate.

An analysis was performed using the CDA speed profile in the ACC analysis tool. The CDA profile is based on ACES simulations using routes designed to force continual descent arrival trajectories. As with the Standard Descent profile, the times to travel were computed for the default and +/-10% speeds with recoverable times based on the difference between times required to traverse the arrival path for the different speed choices. The result is shown in Figure 5.

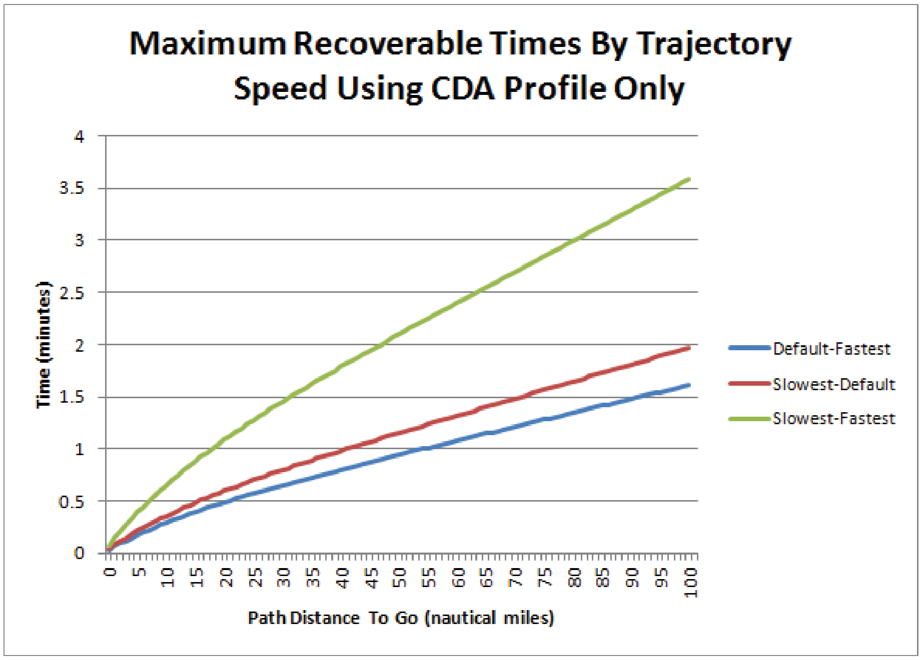

Figure 5 - Maximum Recoverable Times For CDA Descent Trajectories by Distance

There are several things to note about the CDA route configuration. Since the path distances use 2-dimensional projections only, the distance ranges for the CDA routes are equal to their Standard Descent counterparts because the same latitude and longitude crossing points were used. Since the profiles are based on groundspeed (groundspeed is altitude-independent), this does not produce error but might be a source of confusion. Also, for the CDA arrivals which require the aircraft to stay in cruise longer and to start their descent later, the Metering Fix (MF) crossing point would either have to be higher for the same fix location (as their standard approach counterparts), or would have to be closer to the airport to cross at the same altitude. For this analysis, the latitude and longitude locations of the MF points were held constant for the two profiles, effectively permitting the TRACON to reach any altitude necessary to accommodate these MF points on the 40-mile circle.

As with the Standard Descent analysis, the CDA versions of the maximum recoverable times by speed profile (Figure 5) were used to determine the maximum recovery ranges possible by transitioning to between profile speeds. Results are below in Table 3.

Arrival Direction	Route Name	Approximate Path Distance Range (nmi)	Maximum Recoverable Times for CDA Profile (minutes)		
			Default -> Fastest (shortest/longest)	Default -> Slowest (shortest/longest)	Slowest -> Fastest (shortest/longest)
NE	FLCON	38/47	0.77/0.90	-0.95/-1.10	1.72/2.00
NW	ERLIN	46/77	0.89/1.31	-1.08/-1.60	1.97/2.90
SE	CANUK	38/43	0.77/0.85	-0.95/-1.03	1.72/1.88
SW	HONIE	48/56	0.92/1.02	-1.12/-1.25	2.03/2.28
Max Range	(All)	38/77	0.77/1.31	-0.95/-1.60	1.72/2.90

Table 3- Times Recoverable Through Speed Change between Metering Fix and Threshold, CDA Profiles

2.2.2.2 Results for the CDA Profile

These results apply to aircraft already on a CDA route and simply switching to a new speed profile. This might cause a reduction in fuel efficiency, which would have to be weighed against the potential throughput improvement for real-world applications. The usage scenarios for speed profile transitioning are the same as presented in the *Results for Standard Descent Profile* section. Note that the times recoverable for aircraft using CDA profiles are smaller than for their Standard Descent profile counterparts. This is because the CDA speed profile results in higher airspeeds through the Terminal Airspace, shortening the arrival phase and reducing working time.

The maximum recovery time ranges for the default to fastest trajectory transition, “Default -> Fastest” in Table 3, are between 0.77 and 1.31 minutes depending on route length traversed. This means that an aircraft could either make up this much delivery error at the MF by the time it reaches the runway, or it could potentially transition to a runway RTA slot of up to 0.77 to 1.31 minutes earlier by simply increasing speed by 10% within the terminal airspace.

The maximum delay range afforded by switching from a default to the slowest trajectory (“Default -> Slowest” in Table 3), is from 0.95 to 1.60 minutes. This allows the aircraft arriving at the MF the option to delay within the TRACON with simple speed adjustments, or provides the option of moving the on-time MF arrival to a runway slot 0.95 to 1.60 minutes later (perhaps to accommodate other merging traffic) with speed reductions.

The maximum ranges for speed profile transitioning are shown by the slowest to fastest transition (“Slowest -> Fastest” in Table 3). As with the Standard Profile Descent cases, this would correspond to an aircraft planning a slow profile and then switching to a faster profile to use an earlier runway slot. Here, the maximum range possible is shown to be between 1.72 and 2.90 minutes.

2.2.3 Multi-Profile Arrival Route Analysis

2.2.3.1 Estimation of Time Recoverable Within TRACON

A final set of results was obtained by comparing the Standard Descent profile times with the CDA profile times. The curves for the profile changes are shown in Figure 6.

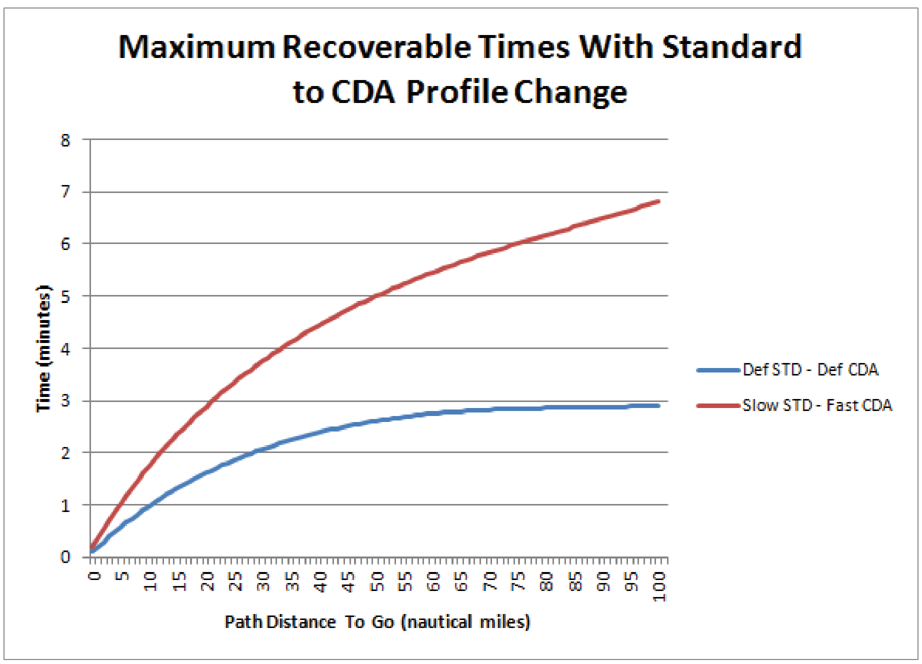

Figure 6 - Maximum Recoverable Times For Multiple Descent Trajectories by Distance

In this case, the final analysis was limited to two key scenarios. The first was the recovery possible by simply switching from a Standard Descent to a CDA profile using the default airspeed. The default airspeed represents the observed mean speed for flown and simulated trajectories, so this case preserves the most common speed choices. The second scenario was the recovery possible in switching from the slowest Standard profile to the fastest CDA profile. This case represents the largest possible recovery (probably at the cost of fuel efficiency). These results are shown in Table 4.

Arrival Direction	Route Name	Approximate Path Distance Range (nmi)	Maximum Recoverable Times for Std -> CDA Profile (minutes)	
			Std Default -> CDA Default (shortest/longest)	Std Slowest -> CDA Fastest (shortest/longest)
NE	FLCON	38/47	2.34/2.55	4.31/4.84
NW	ERLIN	46/77	2.53/2.85	4.79/6.07
SE	CANUK	38/43	2.34/2.47	4.31/4.62
SW	HONIE	48/56	2.57/2.70	4.89/5.28

Table 4 - Times Recoverable Through Trajectory Profile Change (Standard to CDA)

2.2.3.2 Results for the Standard Descent to CDA Profile

The results in Table 4 are significant because they indicate that by simply changing from a default Standard Descent profile to a default CDA profile, more than 2.5 minutes of en route flight delay can be recovered while traversing the arrival route from metering fix to threshold. Alternatively, this time benefit can be used to expedite an on-time aircraft to an earlier runway arrival slot. This number is limited to the time spent in the TRACON only. In practice, the profile change would be made before the aircraft reached its original TOD, so more delay recovery would be possible in the en route portion before the MF. In the case of the maximum range, the "STD Slowest -> CDA Fastest" column in Table 4, approximately 5 minutes of time can potentially be saved.

2.3 Phase 1 ACES Simulation Runs

For this set of tests, the objective was to configure ACES to mimic the scenarios used previously with the ACC analysis tool to verify that results for each tool were harmonious. The analysis tool specializes in estimating the limits of the throughput envelope, while the ACES tool is better suited to model nominal characteristics of flow. Ideally, when used together they provide a fuller picture of the expected flow outcome within the context of a known set of boundaries. In this set of tests, the ACES results would be expected to approach the theoretical throughput limits determined by the ACC tool, but yield results that seldom (if ever) exceed those limits.

2.3.1 Configuration Files

When using the M&S plugin with ACES, traffic into the airport will be managed by the arrival scheduler. Part of the arrival scheduler's job is to enforce regulated spacing between arrival aircraft, which it accomplishes with speed adjustments and path stretch maneuvering. Therefore, the traffic used for the ACES testing for these runs had to be delivered at intervals large enough to insure compliance with spacing regulation without intervention by the arrival scheduler. Otherwise, the intended maximum and minimum trajectory target speeds would be overruled by the arrival scheduler, which would skew the results. The traffic sets for this part of the analysis were carefully choreographed to space arrivals to such that the requested speed changes were only impeded by the performance limits of the simulated aircraft

The January 2012 version of M&S in ACES was used for this set of tests. ACES models simulated flights using traffic specified by the user in Flight Data Set (FDS) files. These FDS files were created by assembling individual flights from the 2006 Baseline Day FDS file selecting a range of aircraft types and paths from each of the four arrival quadrants, and staggering their departure times to insure a conflict-free arrival to the threshold. The 2006 Baseline Day file is one of the datasets available in the Common Scenarios database. This traffic set replicates actual logged flight tracks and cruise targets from September 26, 2006, which was a high volume and clear weather day in the National Airspace System (NAS).

Airports, runways, and arrival routes for the M&S component are specified by the user in XML formatted files. Each arrival route was configured with the same waypoints as were used in the first part of Phase 1 (Figure 2).

2.3.2 Simulation Scenarios and Data

The M&S software allows the aircraft to fly a default, fast, or slow trajectory speed target. These different targets are selectively employed to investigate optional paths for flights to interim waypoints during the trial planning process. These different speed options were employed to direct the flights to fly the required adjusted speed profiles for the simulation runs. The speed targets, however, must be vetted by the trajectory generator to be within performance limits for that aircraft as specified by the BADA data. Because of this, the full 10% speed adjustment requested was not always achieved. These aircraft had less of a speed increase, but still flew faster than their default counterpoints.

The results from 6 configurations -- shortest route with slow, default, and fast speeds and then longest route with slow, default, and fast speeds -- were used as the basis for the recovery time analysis. Once the flights were flown in simulation, the crossing time over the metering fixes were determined for each flight and compared to the runway touchdown time for that same flight to compute the time of travel. For each set, the mean and standard deviations were computed.

2.3.3 Standard Descent Arrival Route Analysis

The results achieved in switching between the three trajectory speed options for a Standard Descent profile are presented in Table 5 - *Simulation Results for Standard Descent Arrivals Using ACES with M&S*. Standard deviation for the results is provided for each run as an indicator for the amount of scatter in the travel times.

		Recovery Times for Simulated STD Profile (minutes)											
		Def -> Fast (shortest)		Def -> Fast (longest)		Def-> Slow (shortest)		Def-> Slow (longest)		Slow -> Fast (shortest)		Slow -> Fast (longest)	
Arrival Dir	Route Name	Mean	Std Dev	Mean	Std Dev	Mean	Std Dev	Mean	Std Dev	Mean	Std Dev	Mean	Std Dev
NE	FLCON	0.777	0.056	0.931	0.184	-0.889	0.085	-1.443	0.195	1.666	0.075	2.374	0.270
NW	ERLIN	0.895	0.051	1.433	0.212	-1.195	0.059	-1.890	0.185	2.090	0.013	3.331	0.406
SE	CANUK	0.639	0.126	0.816	0.050	-0.863	0.052	-1.038	0.132	1.599	0.080	1.854	0.107
SW	HONIE	0.922	0.106	0.995	0.341	-1.002	0.119	-1.589	0.240	2.412	0.129	3.000	0.081

Table 5 - Simulation Results for Standard Descent Arrivals Using ACES with M&S

Finally, the results from the simulated aircraft (from Table 5) are presented alongside the computed maximum values from the mathematical analysis (from Table 2) for comparison in Table 7.

Route Name	Default -> Fastest		Default -> Slowest		Slowest -> Fastest	
	Maximum (Analysis)	Measured in Simulation	Maximum (Analysis)	Measured in Simulation	Maximum (Analysis)	Measured in Simulation
FLCON	0.99/1.13	0.777/0.931	-1.21/-1.39	-0.889/-1.443	2.19/2.52	1.666/2.374
ERLIN	1.12/1.57	0.895/1.433	-1.37/-1.91	-1.195/-1.890	2.48/3.48	2.090/3.331
CANUK	0.99/1.07	0.639/0.816	-1.21/-1.31	-0.863/-1.038	2.19/2.38	1.599/1.854
HONIE	1.15/1.27	0.922/0.995	-1.40/-1.55	-1.002/-1.589	2.55/2.82	2.412/3.000

Table 6 - Comparison of Maximum Recovery Based on Math Analysis and Measured Recovery with Simulation

2.3.4 Results

In the ACC analysis-to-simulation comparison, all but 3 cases show the simulated traffic recovering less time than the prediction. Inspection of the results on a flight-by-flight basis reveals two reasons for this. The first is due to the use of more realistic trajectories for the simulated aircraft with rounded curves for aircraft using “fly-by” waypoints. With fly-by waypoints (as opposed to "fly-over" waypoints), the ACES KTG trajectory generator flies an inside curve for each corner with the resulting flown path being shorter than the point-to-point waypoint connection used in the analysis phase (Figure 7). The radius of the curve varies with the speed and performance of the aircraft flying the route. As a result, the simulated routes end up being several miles shorter, with less opportunity to recover time.

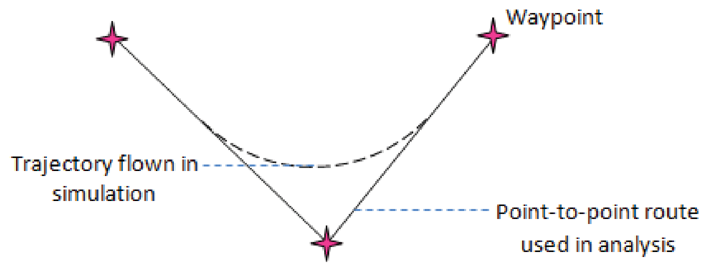

Figure 7 - Path Shortening With Fly-By Waypoint Trajectory

The second reason is that in many cases the simulated aircraft do not achieve the full 10% speed adjustment requested for the fastest and slowest trajectories. This is because the ACES simulation flights are constrained by the BADA performance limits. For these aircraft, the more restrictive performance limit is imposed for the speed adjustment. The result of both effects (shortened path and reduced speed adjustment) causes the analytical maximums to be slightly higher than simulated results.

3 Phase 2 – En Route Airspace

3.1 Mathematical Analysis of the En Route Airspace

A mathematical analysis was conducted to determine the time reduction possible by increasing an aircraft's groundspeed by 10%. Groundspeed is required here to relate speed changes to arrival time changes relative to a fixed location. In practice, the operating speeds of current day aircraft are limited by airspeed rather than groundspeed, so the corresponding airspeed (accounting for winds) to these groundspeed changes would have to be considered to assess any given aircraft's ability to achieve this change within its performance limitations. For aircraft experiencing extreme headwind or tailwind conditions, this 10% value might be either overly optimistic or overly constraining, respectively. However, the required groundspeed change is the necessary first step in the calculation.

The 10% parameter was selected to align the analysis results with the speed changes that will be applied by the M&S software in ACES as the next step. For the M&S software, a 10% speed adjustment limit is currently enforced to avoid adjusting aircraft speed outside the allowable range for aircraft below 10000 feet and within the terminal airspace. The FAA's Order 7110.65[4] requires that aircraft must be at or below 250 knots below 10000 feet, and ATC is restricted to adjust speed to no less than 210 knots for turbojet aircraft beyond 20 path miles from the threshold, and not less than 170 knots within 20 path miles of the threshold.[1] Though this regulation is not specifically applied in the software decision logic, restricting the speed adjustment to +/- 10% allows simulated flights to meet this guideline in most cases. For the simulated flights in the final phase of this experiment, the same

---

[1] Chapter 5 Radar, Section 7 Speed Adjustment.

M&S software will be allowed to trial plan the end of en route flight as the arrival scheduler's Planning Horizon control is extended to before TOD, and so will also be constrained to a 10% speed change.

For each data point in the distance and speed matrix, the original and reduced time to travel are computed for the 10% (0.1 in the equation below) speed increase and used to compute the reduction in time for that case:

$$timeToTravelMinutes_{Original} = \frac{distanceNM}{speedKts_{Original}}$$

$$timeToTravelMinutes_{Expedited} = \frac{distanceNM}{speedKts_{Expedited}}$$

$$speedKts_{Expedited} = speedKts_{Original} + 0.1 * speedKts_{Original} = 1.1 * speedKts_{Original}$$

$$timeReduction_{Minutes} = \left(\frac{distanceNM}{speedKts_{Original}} - \frac{distanceNM}{1.1 * speedKts_{Original}}\right) * hrsToMinConversion$$

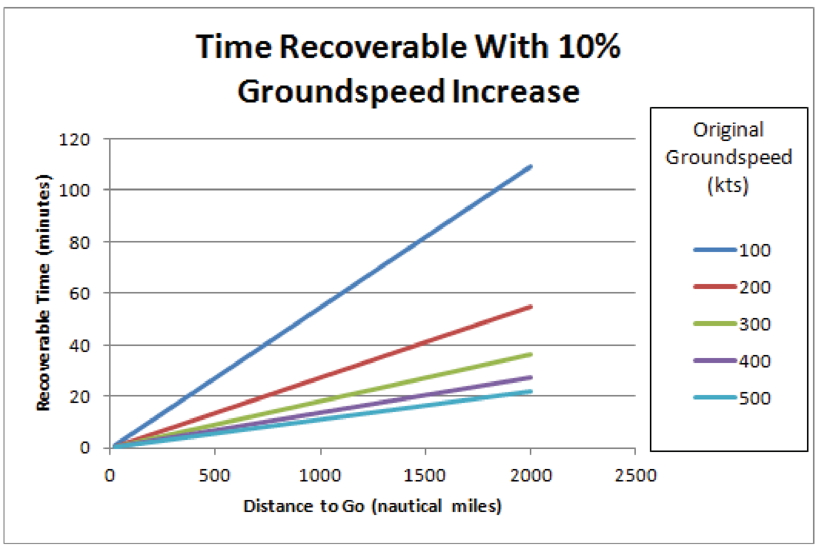

Figure 8: Recoverable time based on original groundspeed and distance traveled

This time recovery could either be used to expedite an aircraft to its final destination to fill an earlier arrival slot, or to recapture an RTA target after some in-flight delay. Note that recovery of even a small delay of 10 minutes or less requires a significant distance when aircraft are restricted to this 10% speed adjustment. For real-world aircraft, which target airspeed rather than groundspeed, the winds would certainly impact the achieved time recovery. This groundspeed presentation, though simplistic, allows easy identification of trends and magnitude. Simulation results, which are presented later in the document, are based on airspeed changes.

One trend, which at first may seem counter-intuitive, is that slower aircraft have a larger recoverable time than their faster counterparts over the same distance. This is counter-intuitive because a 10% speed adjustment represents a larger speed increase for the faster moving aircraft. However, the longer time to travel for the slower flights has a more significant influence on the outcome. This is consistent with trends results reported by Palopo, Nikoleris, Chatterji, and Almog in studies conducted in 2012 at Ames Research Center supporting the SAIE.SPA.3.03 milestone, System-Level Benefits Assessment of Combined Concepts. [6]

4 Phase 3 - Combined En Route and Terminal

4.1 Role of the Arrival Scheduler Planning Horizon

Since the arrival scheduler's capacity to efficiently assign RTAs is restricted by the scope of the time change that can be requested of the individual flights, the size of the Planning Horizon significantly impacts the arrival scheduler's range of options for trajectory adjustment. This means that the Planning Horizon radius is directly related to the realizable throughput of the subject airport for a scheduler than can speed or slow aircraft. As flexibility is gained in configuring the arrival stream, throughput is gained.

As the Planning Horizon is expanded, the throughput increases, but only to the limit of a fully loaded arrival stream. Therefore, there is a radius beyond which no more arrival scheduling efficiency can be gained. The arrival-scheduled airspace will necessarily be more restrictive, and over extending that space may place an unnecessary burden on en route flights. An analysis was conducted during this phase to determine the maximum useful radius for arrival scheduling.

The expanded Planning Horizon used in this phase represents a departure from current practice. Management of the Terminal Airspace and runway arrival sequencing is handled by Terminal Radar Area Control (TRACON) flow managers. Today's TRACONs vary in shape, and typically measure 40 to 60 nmi in radius. Changing the paradigm to use a Planning Horizon beyond a distance equal to the current day TRACON would obviously represent significant change. It would certainly also have to be accompanied by a set of enabling tools both for scheduling the flow (probably from the ground) and for following the instructions (probably from the flight deck). This analysis supposes the existence of these capabilities, and simply identifies the potential improvements afforded by those changes.

4.2 ACES Simulation Runs

4.2.1 Software Description

For this section of the experiment, the ACES simulation with integrated M&S and TFM was used to appropriately model the changing flight phases, the performance envelopes, and the physical aircraft trajectories. The M&S software runs as an ACES plugin, and allows the user to configure the size of the Planning Horizon (the "Radius of Control", or "ROC", in the M&S vernacular). The TFM software resides in ACES as a core component which can be activated at the user's selection.

This experiment required supplying the airport with an amount of traffic that was higher than capacity for each configuration. For these runs, the "No Growth, No Consolidation of Routes" version of the 2018 Baseline Day dataset was used for traffic. This version of the 2018 Baseline Day was created at NASA Langley using the Transportation Systems Analysis Model (TSAM) software tool [7], and is based on projected air traffic growth for the 2018 time frame, restricted to current route and air carrier fleet representations. Other versions of this high volume traffic set are available which modify the routes in response to this traffic using models that emulate carrier decision logic. However, this study required the higher volume of traffic with the current day route tracks. For Atlanta, this data set represents more than a 30% increase over current day volume, and is over capacity for current day Atlanta airspace configuration.

The 2018 dataset contains enough traffic to potentially keep the simulated model of Atlanta at full capacity for the entire period of interest in the simulation run. However, the TFM component of the simulation was needed to regulate the arrival rate. Without TFM operating, some aircraft would be delayed so long at the metering fix that they would reroute to a different airport. This would remove those aircraft from the supply pool, potentially resulting in under-capacity volume for later portions of the period of interest. The TFM component in ACES primarily delays aircraft using ground holds. During simulation initiation, TFM investigates the predicted arrival rates to the destination airport (in this case, Atlanta), and compares that rate to the airport's capacity. The capacity can be selected by the user, and was set to 50 aircraft per runway for the scenarios. For the 3-runway scenario, this fed 150 aircraft per hour into the terminal airspace. This traffic rate was over-capacity, but not to the degree that would cause aircraft to reroute to other airports. Instead, the over-capacity traffic was handled by the M&S component using in-flight delay.

The M&S software can impose in-flight delay only within the Planning Horizon boundary. It imposes delay in one of two ways, at its discretion. It is allowed to adjust an aircraft's speed by +/-10%, or within the limits of that aircraft's performance as dictated by the BADA performance database, whichever is more restrictive. Additionally, it can apply a path stretch maneuver between the Planning Horizon boundary and the first STAR waypoint. There is a maximum path stretch that can be applied, however, due to geometry and aircraft performance characteristics. This is because the trajectory generator flies a curve using the vertex of the path stretch triangle as a fly-by-waypoint. If that triangle becomes so

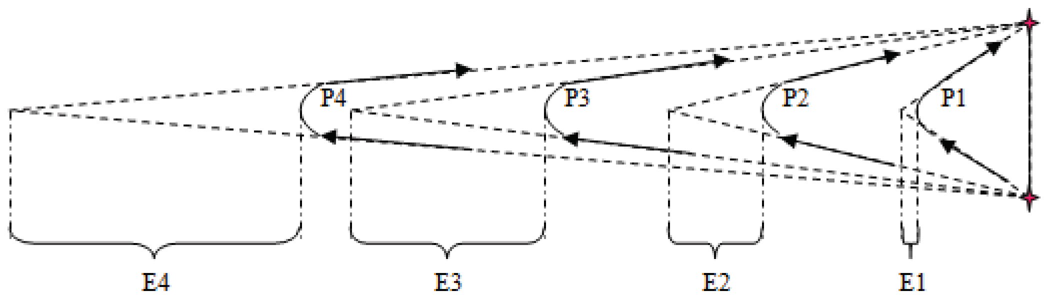

Figure 9: Increasing Path Error with Elongated Path Stretch

elongated that the distance between the flown path and the point-to-point path exceeds the trajectory generator's error tolerance (20 km), then the arrival route is no longer pursued for that flight (Figure 9). If all routes eventually fail in this way, then the flight is diverted to an alternate airport. For this study, the alternate airport was selected as the origin airport to facilitate the identification and removal of these flights from the result set. While this is not consistent with the way flights of this nature would be handled by real-world ATC, identifying and removing these flights was critical to accurate throughput measurement. This is analogous to pouring more water into a glass than the glass can hold, allowing the excess to run down the sides, and then measuring the amount of water that remained in the glass. The average amount of delay that caused this threshold to be breeched varied with Planning Horizon radius (since larger radii allowed larger base lengths for the path stretch triangle), but generally the threshold increased as the Planning Horizon radius increased.

An alternate maneuvering method to allow the traffic to delay indefinitely would obviously be a holding pattern. Earlier versions of the M&S software used holding patterns, but the granularity of the imposed delay was large (generally 3 to 5 minutes) with aircraft having to commit to a full lap of the maneuver. This resulted in lower throughput results because of the substantial decrease in flexibility to fine-tune the delay. The path stretch in combination with the speed adjustment allows examination of delay at 10 second increments, and so replaced the holding pattern method as the software matured. Though it is within the capability of the simulation, no CD&R was used for these runs except for in-trail spacing to meet required wake separation by the M&S software.

4.2.2 Arrival Routes

For the arrival routes, the longest and shortest routes used in the Terminal Airspace analysis for each STAR (see Table 1 - Path Distance by Route) were configured for the ACES runs for consistency with scenarios used in Phase 1 of this study.

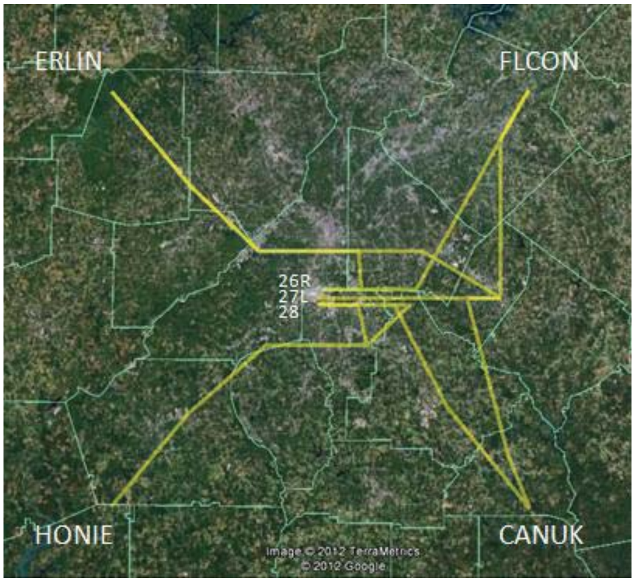

Figure 10: STAR Arrival Routes for Planning Horizon Radius Tests

4.2.3 Simulation Results

4.2.3.1 Processing of ACES Output Files

Landing time data was obtained from ACES log files, and sorted to order landings from earliest to latest. From this, the time offset relative to the first touchdown was computed for each landed aircraft. For each landing, all prior landings within 60 minutes were counted (using Excel) to generate the arrivals per hour at each data point. These results are plotted by simulation hour (relative to the first landing) for the 50 nmi, 100 nmi, and 150 nmi Planning Horizon cases in Figure 11.

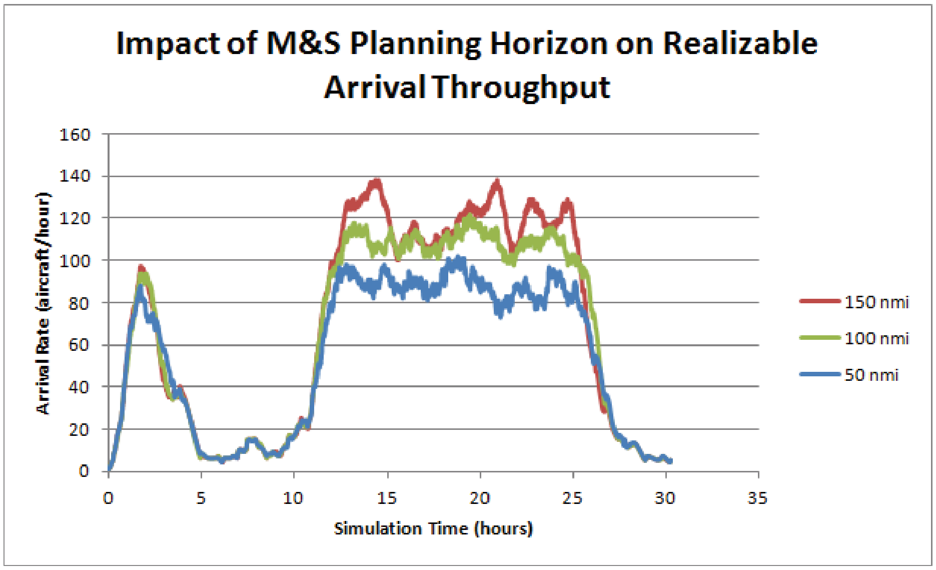

Figure 11: Throughput result with varying M&S Planning Horizon Radius

The throughput for the same set of routes and traffic increased as the radius of the Planning Horizon increased. The average throughput for the high volume period of the simulated day was computed for the 13 to 25 hour section of the run. The standard deviation was also computed to characterize the data range (Table 7).

Planning Horizon Radius (nmi)	Average Throughput At Peak Volume (aircraft/hour)	Highest Throughput (aircraft/hour)	Standard Deviation (aircraft/hour)
50	90	103	3.9
100	110	123	4.9
150	120	139	9.4

Table 7: Mean and Standard Deviation of Throughput for Peak Volume of Simulated Runs

The improvement in throughput seen in these runs is attributed to the increased flexibility afforded to the scheduler with respect to aircraft delay and expedition. As the Planning Horizon radius increases, there is an improved ability of the scheduler to pair followers with suitable leaders which results in significantly less wasted space in the arrival stream. There is, however, a radius of control beyond which little or no improvement is possible once the arrival routes become fully loaded. The size of that radius is investigated in the section "Determination of the Maximum Useful Planning Horizon Radius" later in this paper.

This efficiency is apparent when the traffic is visually inspected using the ACES Viewer tool for the peak volume portion of the runs. Figure 12 is a snapshot of one of the highest throughput traffic periods for the 50 nmi Planning Horizon radius. In this case, some aircraft are closely spaced using leader and follower, but the majority had to be instructed with RTA's to waypoints. Figure 13 is a snapshot of the same peak volume area of the simulated day, but with the 150 nmi Planning Horizon. In this case, almost all traffic could be accommodated using the leader and follower pairing selection resulting in much higher density of traffic along the arrival routes.

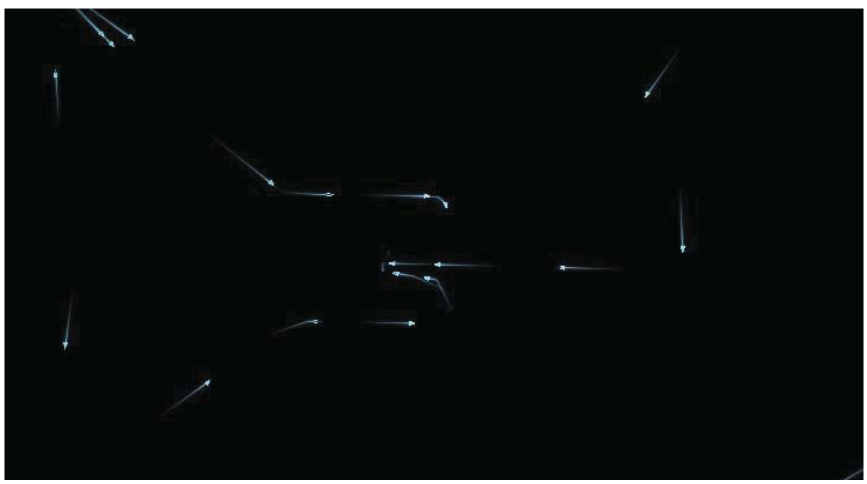

Figure 12: Snapshot of peak efficiency for 50 nmi Planning Horizon

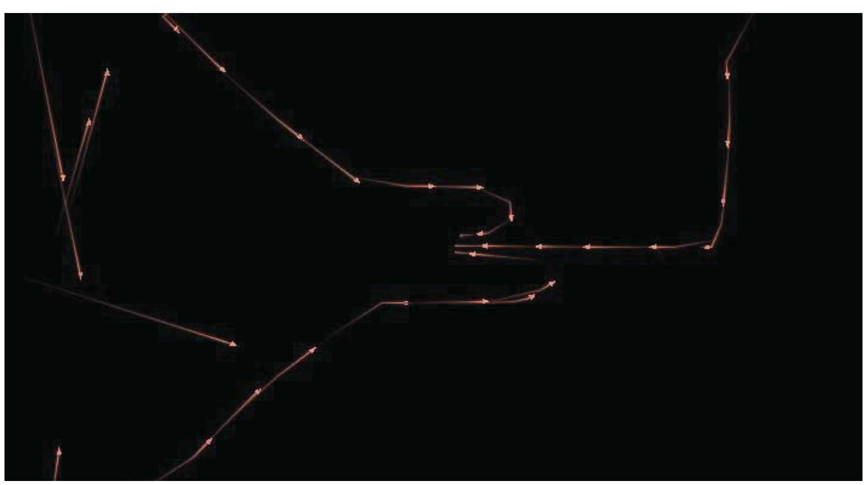

Figure 13: Snapshot of peak efficiency for 150 nmi Planning Horizon

4.3 Validation of Simulation Results

Current day Atlanta TRACON advertises an Airport Acceptance Rate to the Center of 126 aircraft per hour during Visual Flight Rules (VFR) operations, and 104 aircraft per hour during Instrument Flight Rules (IFR) operations.[2] Since the results seen in the ACES runs represents a change from current day advertised capacity, a validation of the results was performed to demonstrate that the higher throughput rates observed in the simulation are actually within the envelope of possibility for a Terminal Airspace under current day wake spacing constraints. Two scenarios are relevant to this study and were tested. First, the upper boundary of the Terminal Airspace capacity envelope was investigated to determine the throughput achievable when all available arrival slots are being used without waste. (Note that in these simulations and validations, safety buffer distance was NOT considered "waste", but was accommodated as part of the usable stream volume.) Though the simulation did not achieve 100% usage of the arrival stream, this number will identify an upper limit that the simulated runs might approach, but not exceed. Second, the analysis tool was used to approximate the percentage of use of the arrival stream for current day, based on the 50 nmi Planning Horizon case.

The ACC analysis tool was used for this final validation test. This tool was used for Phase 1 to quantify the potential reduction in travel time across the configured STAR routes, and allows the relevant parameters to be varied while restricting the traffic to regulated wake spacing and route geometry requirements. Since the spacing requirements vary based on aircraft type for a given leader/follower pair, the capacity results are sets of values dependent partially on the traffic mix. Additionally, the size selection of the safety buffer and nature of the airspeed profiles can vary in the research community. For these reasons, the validation used sweeps of relative parameters to demonstrate the mathematical range of arrival capacity for these varying conditions.

4.3.1 Parameter Selection for the Validation Test

4.3.1.1 Traffic Mix

At each waypoint along the configured routes, the ACC tool uses the regulated separation minima between aircraft as part of the calculation to determine waypoint capacity. Aircraft pairings that only require 3 miles in trail (e.g. *large* following *large*) will result in a higher number of aircraft crossing that fix over a given time than pairings requiring 5 miles in trail (e.g. *small* following *B757*).

The mix parameter attempts to assign a fair estimation of the effect of a non-optimal traffic mix on separation minima for the general stream of traffic being analyzed. The user can vary the parameter to test the impact of the mix on throughput. A stream of all large or all small aircraft would allow minimum spacing in trail and would correspond to a traffic mix of 1.0, the most optimal. An alternating stream of small and heavy aircraft would be the least optimal real-world situation, and would yield a mix of 0.625.

The traffic mix parameter for this test was swept from 0.6 to 1.0 to present the entire range of possibilities. Typical traffic streams fall in the 0.7 to 0.8 range, but since we are comparing against a simulation run that spans many hours, there will be variations in the observed throughput as subsets of the traffic are grouped more or less optimally. As a reference, the mix parameter for the entire traffic stream was computed for 3 cases using the methodology provided in the ACC Description Document – the 2018 Baseline Day, the 2006 Baseline Day, and a set of logged landings from current day Atlanta during the noon to 3:00 pm time period. These results are shown in Table 8.

[2] Provided by Mike Richardson, Support Manager of the Atlanta TRACON.

Data set	Normalized Traffic Mix Value
2018 Baseline Day	0.782
2006 Baseline Day	0.784
Current Day ATL	0.802

Table 8: Computed Mix values for 3 sets of traffic data

4.3.1.2 Safety Buffer

The sweep range for safety buffer size was selected to try to capture the range of values used in other studies in the industry, without judgment as to which is best. Some research favors a static buffer, some a percentage of the required wake spacing, and some a product of the speed variation, to name a few. For this analysis, the buffer was based on a percentage of the wake spacing with the percentage ranging from a very conservative 20% to an optimistic 2%.

4.3.1.3 Speed Profile

The ACC tool provides several optional profiles. The generic current day standard descent was selected for this analysis. Speed efficiency can be adjusted around this profile to capture delay or expedition of the traffic, but for these runs the standard profile was applied without modification.

4.3.1.4 Efficiency (of Occupancy)

The efficiency parameter represents the percentage of usable space in the arrival stream that is occupied. For this validation, efficiency was varied from 40% - 100% to capture the range of maximum throughput seen in the simulation cases, estimated with the help of ACES Viewer. For the simulation runs to be within the valid range, the achieved throughput should not exceed the maximum in the ACC analysis.

4.3.2 Results

Error! Reference source not found. shows the result for the 100% efficiency parameter set, the scenario that is approached by the 150 nmi Planning Horizon simulation case. The boxes shaded highlight the safety buffer and computed *average* mix parameter for the simulation runs using the 2018 Baseline Day. For the hour subsets used to compute the throughput, the mix varied around this average. The 150 nmi Planning Horizon case resulted in a throughput that peaked at 138 aircraft per hour for two brief intervals of the day, and averaged 120 aircraft per hour during high volume, and so the ACC results shown in Table 9 validate that the simulation results are within the capacity of the airspace. Table 9 also shows that more airspace capacity is available if improvement can be made in the mix (through more optimal pairings) or in the safety buffer (though precision speed and distance control, or with reduced wake spacing). In practice, of course, the throughput would still be limited by runway occupancy, which could become the more restrictive limit.

efficiency 1.0		mix				
		0.6	0.7	0.8	0.9	1
SafetyBuffer Percent	0.02	109.3514	120.1326	133.2721	149.6388	170.5882
	0.05	106.2271	116.7002	129.4643	145.3634	165.7143
	0.08	103.2764	113.4585	125.8681	141.3255	161.1111
	0.11	100.4851	110.3921	122.4662	137.5059	156.7568
	0.14	97.84076	107.487	119.2434	133.8873	152.6316
	0.17	95.33202	104.7309	116.1859	130.4543	148.7179
	0.2	92.94872	102.1127	113.2813	127.193	145

Table 9: Airport Capacity for 100% Usage Efficiency Based on Wake Spacing

The efficiency for the 50 nmi radius is represented by a range of solutions, since the conditions at any time point during the run vary in terms of mix and volume. The results for the 30, 40, 50, 60, and 70% efficiency of usage cases are shown in Table 14 – 15. The fields in the table that represent the normal mix (0.7 – 0.9) and safety buffer (0.05 – 0.11) applicable for this scenario are shaded. These shaded regions were inspected for values that most closely represent the average (90 aircraft per hour with a standard deviation of +/- 3.9) and the maximum (103 aircraft per hour) seen in the 50 nmi radius simulation run. The *average* throughput is most closely represented by the 30% efficiency case (Table 10). The *maximum* throughput is most closely represented by the 40% (Table 11) and 50% (Table 12) cases. The peak throughput achieved for the 50 nmi case is consistent with the arrival rate advertised by the Atlanta TRACON to the Center during IFR conditions, which is 104 aircraft per hour as introduced in section 4.3.1.1.

efficiency 0.3		mix					
		0.5	0.6	0.7	0.8	0.9	1
	0.02	80.96886	88.23529	96.93455	107.5368	120.743	137.6471
	0.05	78.65546	85.71429	94.16499	104.4643	117.2932	133.7143
	0.08	76.47059	83.33333	91.5493	101.5625	114.0351	130
	0.11	74.40382	81.08108	89.07499	98.81757	110.9531	126.4865
	0.14	72.44582	78.94737	86.73091	96.21711	108.0332	123.1579
	0.17	70.58824	76.92308	84.50704	93.75	105.2632	120
	0.2	68.82353	75	82.39437	91.40625	102.6316	117

Table 10 - Airport Capacity for 30% Usage Efficiency Based on Wake Spacing

efficiency 0.4		mix					
		0.5	0.6	0.7	0.8	0.9	1
SafetyBuffer Percent	0.02	83.73702	91.25189	100.2486	111.2132	124.871	142.3529
	0.05	81.34454	88.64469	97.38431	108.0357	121.3033	138.2857
	0.08	79.08497	86.18234	94.67919	105.0347	117.9337	134.4444
	0.11	76.94754	83.85308	92.12029	102.1959	114.7463	130.8108
	0.14	74.9226	81.64642	89.69607	99.50658	111.7267	127.3684
	0.17	73.00151	79.55293	87.39617	96.95513	108.8619	124.1026
	0.2	71.17647	77.5641	85.21127	94.53125	106.1404	121

Table 11 - Airport Capacity for 40% Usage Efficiency Based on Wake Spacing

efficiency 0.5		mix					
		0.5	0.6	0.7	0.8	0.9	1
SafetyBuffer Percent	0.02	86.50519	94.26848	103.5626	114.8897	128.999	147.0588
	0.05	84.03361	91.57509	100.6036	111.6071	125.3133	142.8571
	0.08	81.69935	89.03134	97.80908	108.5069	121.8324	138.8889
	0.11	79.49126	86.62509	95.16559	105.5743	118.5396	135.1351
	0.14	77.39938	84.34548	92.66123	102.7961	115.4201	131.5789
	0.17	75.41478	82.18277	90.2853	100.1603	112.4606	128.2051
	0.2	73.52941	80.12821	88.02817	97.65625	109.6491	125

Table 12: Airport Capacity for 50% Usage Efficiency Based on Wake Spacing

efficiency 0.6		mix					
		0.5	0.6	0.7	0.8	0.9	1
SafetyBuffer Percent	0.02	89.27336	97.28507	106.8766	118.5662	133.1269	151.7647
	0.05	86.72269	94.50549	103.8229	115.1786	129.3233	147.4286
	0.08	84.31373	91.88034	100.939	111.9792	125.731	143.3333
	0.11	82.03498	89.39709	98.21089	108.9527	122.3329	139.4595
	0.14	79.87616	87.04453	95.62639	106.0855	119.1136	135.7895
	0.17	77.82805	84.81262	93.17443	103.3654	116.0594	132.3077
	0.2	75.88235	82.69231	90.84507	100.7813	113.1579	129

Table 13: Airport Capacity for 60% Usage Efficiency Based on Wake Spacing

efficiency 0.7		mix					
		0.5	0.6	0.7	0.8	0.9	1
SafetyBuffer Percent	0.02	92.04152	100.3017	110.1906	122.2426	137.2549	156.4706
	0.05	89.41176	97.4359	107.0423	118.75	133.3333	152
	0.08	86.9281	94.72934	104.0689	115.4514	129.6296	147.7778
	0.11	84.5787	92.16909	101.2562	112.3311	126.1261	143.7838
	0.14	82.35294	89.74359	98.59155	109.375	122.807	140
	0.17	80.24133	87.44247	96.06356	106.5705	119.6581	136.4103
	0.2	78.23529	85.25641	93.66197	103.9063	116.6667	133

Table 14: Airport Capacity for 50, 60, and 70% Usage Efficiency Based on Wake Spacing

The 50 nmi Planning Horizon case only reached 103 aircraft per hour for one short interval and averaged 88 aircraft per hour over the high volume portion of the simulated day. This indicates that the scheduler, using a Planning Horizon more typical of current day TRACON sizes, was able to achieve only a 30 - 70% efficiency level, and only for a brief portion of the run. Again this is compared to the peak throughput achieved for the 50 nmi case and is consistent with the 104 aircraft per hour arrival rate advertised by the Atlanta TRACON to the Center during IFR conditions.

4.4 Determination of the Maximum Useful Planning Horizon Radius

As the Planning Horizon is expanded, the throughput increases, but only to the limit of a fully loaded arrival stream. Therefore, there is a radius beyond which no more arrival scheduling efficiency can be gained. The arrival-scheduled airspace will necessarily be more restrictive, and over extending that space may place an unnecessary burden on en route flights. This final analysis was conducted to determine the maximum useful radius for arrival scheduling. These runs were conducted using the ACES simulation with M&S and TFM.

The 3-runway configuration for the test used Atlanta Airport with an expanded 2018 baseline dataset to substantially overload the Terminal Airspace for all radii tested. This expanded dataset was created by replicating flights distributed throughout the dataset, which resulted in a 13.4% increase in traffic. This was done to insure that the Atlanta Terminal Airspace was fully saturated for each radius. Radial distances ranging from 50 nmi to 300 nmi were tested, with intermediate points judiciously selected to define the function. Because of this, the independent data points for the radii are not equidistant, but were arrived at methodically. This was done to minimize the number of runs necessary to define the relationship.

4.4.1 Results

Figure 14 presents the average throughput observed as a function of Planning Horizon radius. Figure 15 shows the corresponding throughput with simulation time for each test, for reference. For radii less than 130 nmi, efficiency increases almost linearly with increased Planning Horizon size. Beyond this radius, diminishing improvement is achieved up to the 150 nmi radius, after which no more improvement occurs.

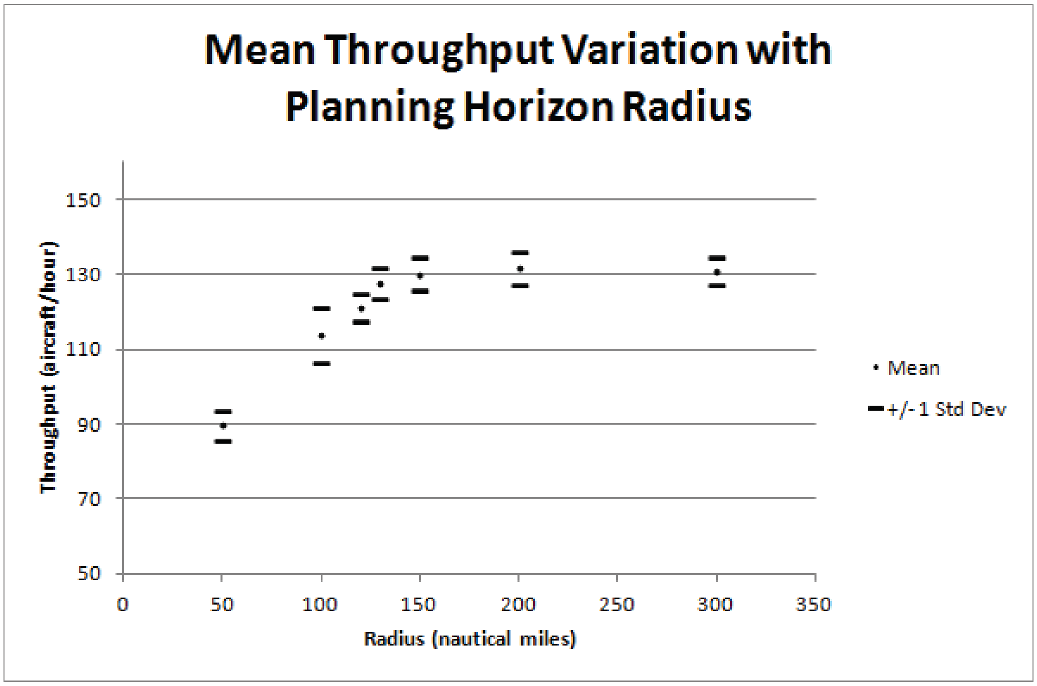

Figure 14: Variation of Throughput with Planning Horizon

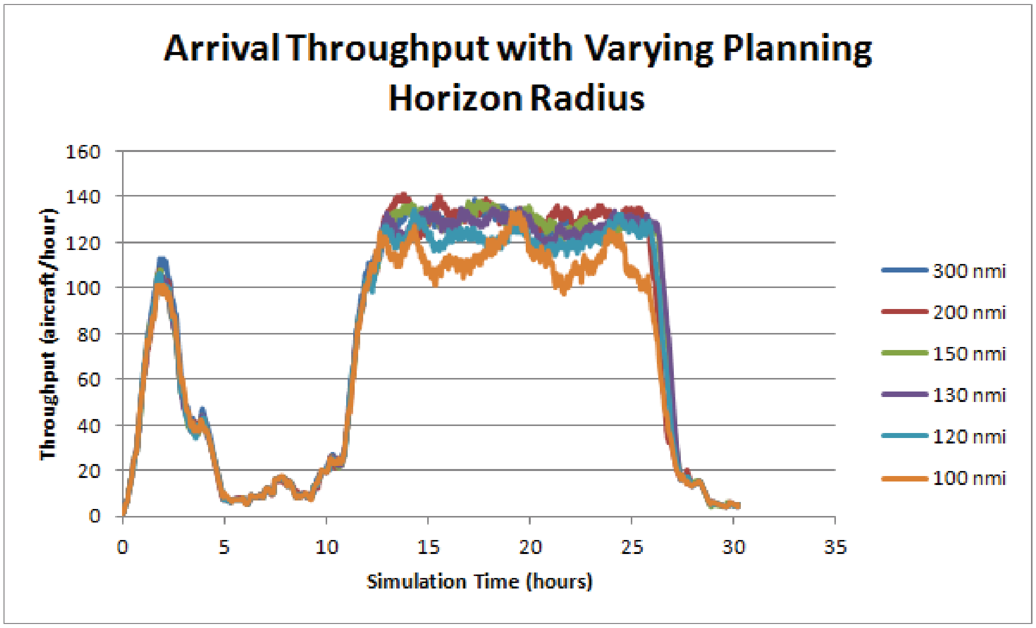

Figure 15: Arrival Throughput with Planning Horizon Size for Maximum Useful Size Test

The throughput efficiency can be computed for any single run by dividing the average throughput achieved during the high volume period by the maximum average throughput achieved for the high volume period across the entire set, and multiplying by 100 to convert the decimal to a percentage:

$$throughput\ efficiency = \frac{throughput}{throughput_{max}} * 100$$

The final column in Table 15 shows the computed throughput efficiency for the set of simulation runs using Atlanta with 3 runways operating.

Radius (nmi)	3 Runway Throughput (aircraft/hour)	Throughput Efficiency (%)
50	89.74	68.12
100	114.04	86.56
120	121.22	92.01
130	127.47	96.75
150	130.20	98.82
200	131.75	100.00
300	130.89	99.35

Table 15: Throughput Percentage of Maximum Achieved by Radius

The throughput efficiency demonstrates how fully the arrival scheduler was able to load the available arrival stream, regardless of that stream's size, compared to the highest throughput achieved for that set. Therefore, the 100% efficiency case does not necessarily represent capacity, but rather the closest that configuration was able to get to capacity for the test scenario. The relationship of the throughput efficiency to the radius is important because once the slope goes to zero, expansion of the Planning Horizon is no longer useful to throughput. This relationship is not dependent on the number of runways or whether the runways were supporting only arrivals or both arrivals and departures. This throughput efficiency is influenced by the amount of speed change allowed (10% in this case) and the arrival scheduler's ability to use this adjustment potential to efficiently arrange arriving traffic. Because of this, the same general results occur for a given arrival scheduler under a diverse set of airport configurations. To demonstrate this, a different traffic set (the 2006 Baseline Day traffic set) was used for ACES simulation runs into Dallas Fort Worth International Airport (DFW).

Figure 16: Dallas Fort Worth International Airport (DFW)

STAR routes matching current day were configured for the airport landing south with 2 arrival runways operating, 17C and 18L, and then again with 1 arrival runway operating, 17C (Figure 17). TFM was configured to deliver 50 aircraft/runway/hour (an over-capacity situation), and the simulation was run for radii of 50, 100, 120, 130, 150, and 200 nmi. As in the Atlanta over-capacity runs, the arrival scheduler diverted flights when the path stretch delay maneuver became too large.

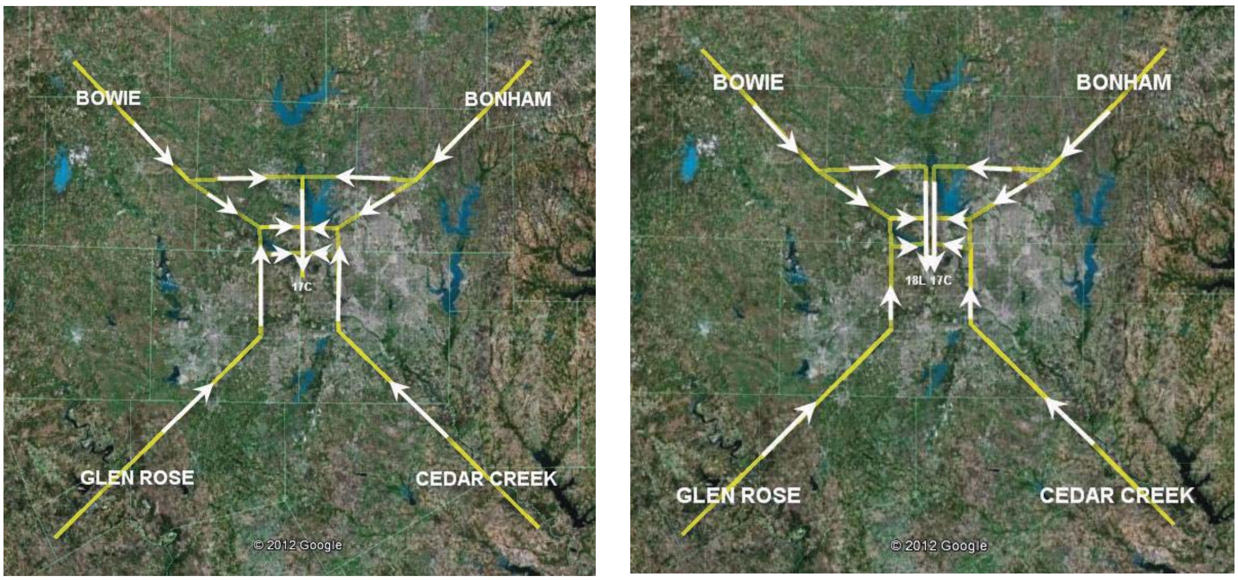

Figure 17: Arrival Routes for DFW with one runway (left) and two runways (right)

Figure 18 presents the results for throughput as a function of simulation time at various Planning Horizon sizes for the single runway and dual runway scenarios. As expected, the throughput for the single runway runs was significantly less than that of the double runway runs. However, the throughput efficiency was nearly the same. Figure 19 is a summary of the computed throughput efficiency for all the configurations (Atlanta with 3 runways, DFW with 2 runways, DFW with 1 runway) presented on a single plot. While the overall throughput varies with airport configuration and number of runways operating, the relationship of the throughput efficiency to Planning Horizon radius is very similar. In all cases, 95% efficiency is achieved between the 120 and 130 nmi radius circles.

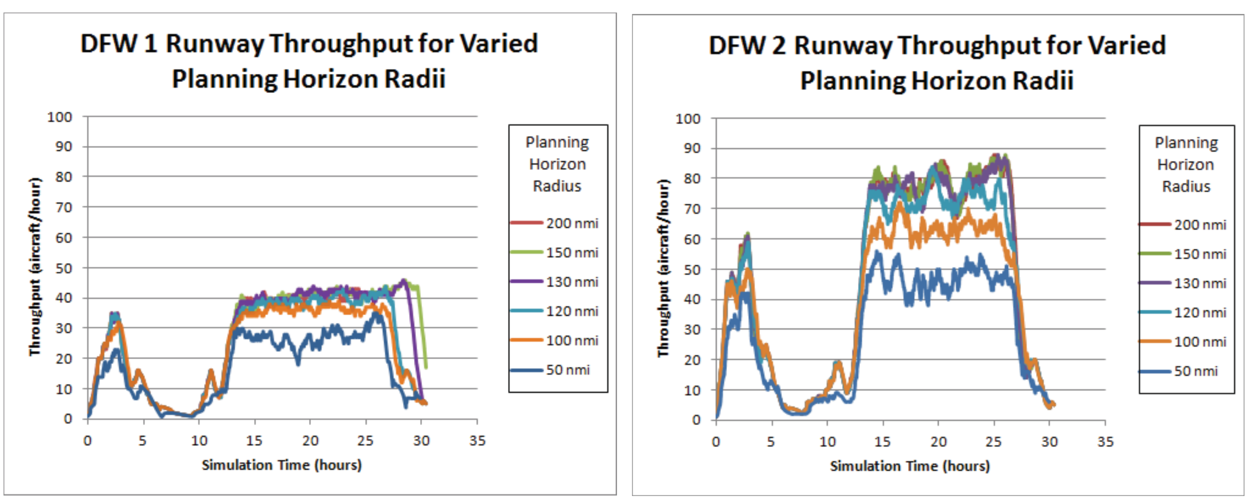

Figure 18: DFW Single and Double Runway Arrival Throughput for Varied Planning Horizon Radii

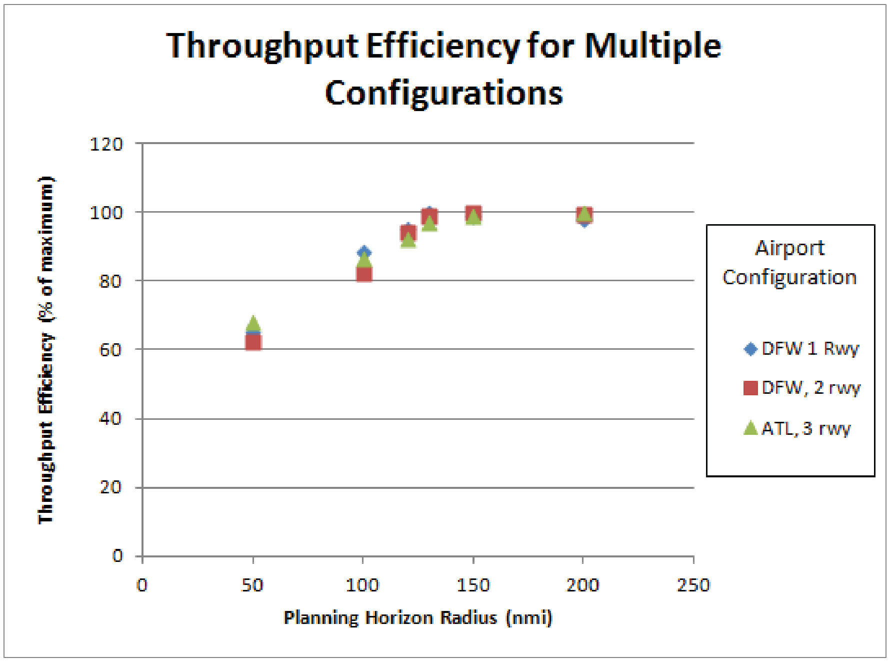

Figure 19: Throughput Efficiency Comparison for Varied Airport and Number of Runways

For shared arrival/departure runway operations, the same results are expected. Though the total arrival throughput would decrease with less of the simulated day supporting arrivals during shared runway operations, the throughput efficiency for any given arrival period would be comparable to the dedicated arrival runway runs as long as the specific blocks of time for arrivals could be isolated and conveyed to the scheduler in advance. Unfortunately, the interleaved departure/arrival scenario cannot be demonstrated here. While the simulation tool can model interleaved arrivals and departures, fixed assignment of arrival and departure block periods would be required for run comparison, but is not an option for this tool.

4.4.2 Results Discussion

Though the volume of traffic accommodated changed considerably depending on the number of runways operating, the throughput efficiency stayed the same. This was shown to be the case for three different combinations of routes and independent runways, and is predicted to also apply to runways supporting departures and arrivals. Furthermore, for the three configurations tested, achieved throughput efficiency increased almost linearly with increasing Planning Horizon radius up to a Planning Horizon radius of approximately 130 nmi. Beyond that size, diminishing improvement (if any) was obtained for the increased size of the radius. This is because the size of the Planning Horizon is large enough to allow the arrival scheduler to efficiently load the arrival stream to near capacity; no more benefit can be obtained.

While the throughput efficiency is *independent* of the route and runway configuration, it is very *dependent* on the efficiency of the arrival scheduler and speed adjustment imposed on it. The IDEAS scheduler used in this study operates very efficiently and can support delays, speed increases, following instructions, merging, and diverging at interim route fixes. Because of this, the scheduler was able to very effectively convert additional planning radius into throughput. In these tests, the maximum speed change was set to 10%. If larger speed changes were allowed and possible within the target aircraft's performance, the arrival scheduler would be able to fully load the arrival stream using a smaller Planning Horizon radius. Conversely, a larger radius would be required if the allowed speed changes were more restrictive. In a real world setting where a prevailing wind conditions would cause arriving aircraft from different compass points to see tail or head winds, a variety of speed adjustment ranges would ideally be managed by the arrival scheduler to increase the potential to expedite aircraft in a tail wind while increasing the potential to delay aircraft in a head wind, with each aircraft obeying their performance envelope for airspeed.

In practice, speed management by the arrival scheduler would potentially incur an increased fuel cost as aircraft were asked to vary from their default (and presumably more efficient) trajectories. The area inside the increased Planning Horizon would also be more restrictive. Therefore, for low volume periods or airspaces, a smaller radius might suffice and reduce the imprint of the more restrictive scheduling area. For projected traffic levels of 30 aircraft per hour per runway or less, for example, a 50 nmi Planning Horizon was sufficient for this scheduler using a 10% speed change maximum.

It is significant that observed throughput improvements required no change to current wake spacing regulations. This demonstrates that significant additional capacity exists in the current day Terminal Airspace if managed efficiently with enabling NextGen technologies. If on-going wake research is able to safely decrease wake spacing regulations in the future, even more throughput could be realized in addition to the results shown here.

5 Future Work

Investigations documented in this paper are a subset of a larger project milestone to study interactions that influence design and interoperation between M&S, TFM, and SA when these technologies have to work cooperatively. For this first analysis, the impact of the M&S arrival scheduler's Planning Horizon size was investigated to understand the role this plays in arrival throughput. This will inform future simulation studies using the arrival scheduler as one component in a cooperating system.

The next step will investigate some of the costs and benefits associated with the TFM's Planning Horizon size selection. Since TFM predicts airport loading, it could potentially inform a dynamically changing Planning Horizon for the arrival scheduler to maximize airport throughput when necessary with larger radii, yet minimize the arrival scheduler footprint on the airspace when traffic volume is expected to be low. Also, TFM monitors changing conditions for both the airport and for flights delayed unexpectedly due to weather or CD&R. A TFM system that can predict (or be informed of) a scheduled flight's late arrival to a metering fix early enough might turn otherwise wasted arrival slots into opportunities for other traffic to expedite. If done judiciously, throughput could be preserved with all arrivals occurring in the same window, but in a slightly different order. These investigations will be conducted using the ACES simulation (Version 8.1) with cooperative M&S and TFM, which will be available in June, 2012.

Another follow-on study is planned related to the combined use of standard decent and CDA arrival routes, and their impact on the arrival scheduler's Planning Horizon. In section 2.2.3, *Multi-Profile Arrival Route Analysis*, a significant increase in the delay recovery potential was predicted for aircraft originally planning a standard descent arrival, and switching to a CDA arrival route. This potentially increases an arrival scheduler's flexibility for ordering and merging flights in the TRACON if the arrival scheduler's Planning Horizon is crossed before flights reach their TOD. Since the CDA arrival profile keeps aircraft at cruise longer, any flight that has not yet crossed its standard descent TOD point is a potential candidate for a CDA profile. If the arrival scheduler can capitalize on this extra flexibility, the required size of its Planning Horizon might be smaller (than that of its standard-descent-only counterpart) for the same efficiencies. This hypothesis will be tested using the ACES simulation with M&S, using combinations of current day STARs and added CDA routes.

6 References

1. Airspace Concepts Evaluation System (ACES) Version 7.1 with M&S Plug-in, NASA Ames Research Center and Intelligent Automation, Inc., 2012.
2. ACES Viewer, Version 3, Ohio, Aerospace Computing, Inc., 2010.
3. Santos, Michel; Feinberg, Arthur; Zhang, Yingchuan; Teng, Yun; Chen, Jensan; Nigam, Nikhil; Smith, Jeremy, Scheduling the Use of Airborne Merging and Spacing Along Multiple Converging Routes to an Airport, AIAA Modeling and Simulation Technologies Conference, Toronto, Ontario, Aug. 2-5, 2010.
4. Order 7110.65U, Air Traffic Control, Washington, D.C., United States Department of Transportation and the Federal Aviation Administration, February 9, 2012.
5. Airliners.net, "Aircraft Technical Data and Specifications", http://www.airliners.net/aircraft-data/, accessed May 23, 2012.
6. Nikoleris, Tasos; Chatterji, Gano; Almog, Noam; and Palopo, Kee, Arrival Delay Absorption using Extended Metering with Speed Control, NASA Ames Research Center, April 30, 2012.
7. Viken, Jeffrey., et. al., "Utilizing Traveler Demand Modeling to Predict Future Commercial Flight Schedules in the NAS," 11th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference, September 2006, Portsmouth, Virginia.
8. Google Earth 4.3.7284.3916 (beta), Google, 2011.
9. AirNav, KATL Hartsfield-Jackson Atlanta International Airport, AirNav LLC, http://www.airnav.com/airport/KATL, accessed May 20, 2012.

REPORT DOCUMENTATION PAGE	Form Approved OMB No. 0704-0188

1. REPORT DATE *(DD-MM-YYYY)*
01- 09 - 2012

2. REPORT TYPE
Technical Memorandum

3. DATES COVERED *(From - To)*

4. TITLE AND SUBTITLE
Improvement to Airport Throughput Using Intelligent Arrival Scheduling and an Expanded Planning Horizon

5a. CONTRACT NUMBER

5b. GRANT NUMBER

5c. PROGRAM ELEMENT NUMBER

6. AUTHOR(S)
Glaab, Patricia C.

5d. PROJECT NUMBER

5e. TASK NUMBER

5f. WORK UNIT NUMBER
305295.02.07.07.02

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
NASA Langley Research Center
Hampton, VA 23681-2199

8. PERFORMING ORGANIZATION REPORT NUMBER
L-20174

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)
National Aeronautics and Space Administration
Washington, DC 20546-0001

10. SPONSOR/MONITOR'S ACRONYM(S)
NASA

11. SPONSOR/MONITOR'S REPORT NUMBER(S)
NASA/TM-2012-217762

12. DISTRIBUTION/AVAILABILITY STATEMENT
Unclassified - Unlimited
Subject Category 66
Availability: NASA CASI (443) 757-5802

13. SUPPLEMENTARY NOTES

14. ABSTRACT
The first phase of this study investigated the amount of time a flight can be delayed or expedited within the Terminal Airspace using only speed changes. The Arrival Capacity Calculator analysis tool was used to predict the time adjustment envelope for standard descent arrivals and then for CDA arrivals. Results ranged from 0.77 to 5.38 minutes. STAR routes were configured for the ACES simulation, and a validation of the ACC results was conducted comparing the maximum predicted time adjustments to those seen in ACES. The final phase investigated full runway-to-runway trajectories using ACES. The radial distance used by the arrival scheduler was incrementally increased from 50 to 150 nautical miles (nmi). The increased Planning Horizon radii allowed the arrival scheduler to arrange, path stretch, and speed-adjust flights to more fully load the arrival stream. The average throughput for the high volume portion of the day increased from 30 aircraft per runway for the 50 nmi radius to 40 aircraft per runway for the 150 nmi radius for a traffic set representative of high volume 2018. The recommended radius for the arrival scheduler's Planning Horizon was found to be 130 nmi, which allowed more than 95% loading of the arrival stream.

15. SUBJECT TERMS
Airports; Airspace; Arrivals; NextGen; Runways; Scheduling; Simulation

16. SECURITY CLASSIFICATION OF: a. REPORT	b. ABSTRACT	c. THIS PAGE	17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
					STI Help Desk (email: help@sti.nasa.gov)
					19b. TELEPHONE NUMBER *(Include area code)*
U	U	U	UU	37	(443) 757-5802