NASA Public Access

Author manuscript

Proc Inst Mech Eng G J Aerosp Eng. Author manuscript; available in PMC 2020 February 04.

Published in final edited form as:

Proc Inst Mech Eng G J Aerosp Eng. 2016 February 8; 230(9): 1762–1779. doi:

10.1177/0954410016629499

Algorithms for Control of Arrival and Departure Traffic in Terminal Airspace

Heinz Erzberger,

NASA Ames Research Center, Moffett Field, CA, 94035

Tasos Nikoleris,

University of California Santa Cruz, Moffett Field, CA, 94035

Russell A. Paielli,

NASA Ames Research Center, Moffett Field, CA, 94035

Yung-Cheng Chu

AerospaceComputing, Inc., Moffett Field, CA, 94035

Abstract

This paper presents a design approach and basic algorithms for a future system that can perform aircraft conflict resolution, arrival scheduling and convective weather avoidance with a high level of autonomy in terminal area airspace. Such a system, located on the ground, is intended to solve autonomously the major problems currently handled manually by human controllers. It has the potential to accommodate higher traffic levels and a mix of conventional and unmanned aerial vehicles with reduced dependency on controllers. The main objective of this paper is to describe the fundamental trajectory and scheduling algorithms that provide the foundation for an autonomous system of the future. These algorithms generate trajectories that are free of conflicts with other traffic, avoid convective weather if present, and provide scheduled times for landing with specified in-trail spacings. The maneuvers the algorithms generate to resolve separation and spacing conflicts include speed, horizontal path, and altitude changes. Furthermore, a method for reassigning arrival aircraft to alternate runways in order to reduce delays is also included. The algorithms generate conflict free trajectories for terminal area traffic, comprised primarily of arrivals and departures to and from multiple airports. Examples of problems solved and performance statistics from a fast-time simulation using simulated traffic of arrivals and departures at the Dallas/Fort Worth International Airport and Dallas Love Field are described.

Keywords

automated air traffic management; algorithms for control of arrival traffic; autonomous system for air traffic control; separation assurance for air traffic

With the expected future growth of air traffic, higher levels of automation in air traffic management are likely to be needed to handle a denser and more diverse mix of air traffic, especially to accommodate the rapid growth of Unmanned Aerial System (UAS) traffic. Moreover, as the requirement for environmentally friendly aviation becomes increasingly important, higher levels of automation in air traffic control may enable all types of aircraft to fly more fuel efficient trajectories that are difficult for controllers to accommodate manually.

In response to anticipated need for handling the mix of conventional and UAS traffic efficiently, NASA has conducted research towards the development of a system for automated conflict resolution, arrival management, and weather avoidance that could be the basis for autonomous air traffic control. This system is referred to as the Advanced Airspace Concept [1] (AAC) and it has so far been developed for traffic in en-route airspace including descents to a meter fix. Recently, a research effort was initiated to extend AAC to terminal area airspace, referred to as Terminal AAC (TAAC). This paper aims to delineate the principal components of TAAC as well as its methods and algorithms for aircraft conflict resolution and management of arrival traffic.

The design objective of TAAC differs from previous work in that it aims to achieve a higher level of autonomy for controlling traffic in the terminal area. In earlier work, the Final Approach Spacing Tool [2,3] generated speed and turn advisories to assist controllers in sequencing and spacing traffic onto the final approach. In a recently developed system [4], which has been extensively evaluated in real time simulations, controllers are presented with speed advisories only to assist them in sequencing and spacing of arrivals. Since the premise of these designs is to assist rather than replace controllers, aircraft trajectories are generally not guaranteed to be conflict free against all traffic. This type of automation is referred to as a decision support tool. In using it, the controller retains responsibility for maintaining separation between all aircraft, including those for which automated advisories are presented to controllers on their monitors. When automated advisories are limited to solving only a subset of the problems occurring in managing traffic, controllers must maintain full situational awareness of all types of problems that can occur. That limits the ability of such tools to fully optimize traffic flows and trajectories. In contrast, the design described herein generates trajectories and clearances that obey the same separation requirements and operational limits controllers provide while also allowing aircraft greater freedom to fly preferred trajectories such as continuous descents and optimized climbouts.

Although the TAAC-generated trajectories could be adapted to provide classical decision support advisories to controllers, the more significant use of this approach is in a future air traffic control system that operates with a high degree of autonomy. In such a future system, the role of controllers would change from handling routine separation assurance and arrival sequencing tasks to handling exceptional traffic situations or accommodating special pilot requests. A key technology that will enable increased autonomy in air traffic management is the ground-air data link, which will allow ground-based automation systems to uplink clearances and trajectories directly to onboard systems without prior approval of a controller. Thus, the design proposed and studied in this paper represents a paradigm shift relative to NextGen in the U.S. and SESAR in Europe, which retain the traditional controller-centric operational concept of current systems.

Similar to AAC's architecture, TAAC incorporates two independent systems for conflict detection and resolution. One system is designed to handle conflicts and solve problems predicted to occur in the range of approximately 1–12 min. It is the mainstay for solving separation assurance and arrival management problems. This element can be considered the strategic problem solver in TAAC and is referred to as the Terminal AutoResolver (TAR), which is the main subject of this paper. The second separation-assurance system in TAAC

will focus exclusively on handling tactical conflicts, defined as those with times to loss of separation of less than approximately 1 min. This separate and independent system would come into play whenever TAR failed to detect or resolve a conflict for which the time to first loss had become less than approximately 1 min. The hierarchical structure for separation assurance in AAC and adopted for TAAC is described in [1]. This structure is intended to achieve autonomy with safety and thus enable reduced dependence on humans as primary traffic control and separation assurance agents. Finally, it should be understood that no algorithm can guarantee finding a conflict-free solution to every conceivable conflict scenario that can arise in practice. Therefore, the system must also provide preplanned safe-escape procedures, such as placing an aircraft in a holding pattern, to serve as a fallback solution for unresolved conflicts, although such conflicts should be extremely rare. Nevertheless, an autonomous system must be designed to find a safe solution, though not necessarily an efficient one, to every problem encountered. And when a handoff to a human controller becomes necessary, the system needs to provide sufficient lead time for humans to gain situational awareness of the traffic situation before they can safely take control. Safe-escape algorithms TAAC will be the subject of future research.

Terminal AutoResolver reuses several of the algorithmic functions originally developed for the enroute AutoResolver, the strategic problem solver of AAC. Its design and performance for resolving conflicts in enroute airspace are described in [1,5]. This paper focuses on the new functions that have been specifically designed for operations in terminal airspace and also updates the design of TAR first described in [6].

Section II of the paper provides an overview of TAR's process for conflict resolution and arrival management. Section III describes TAR's module for arrival management in greater detail, Section IV describes the principal maneuvers used for conflict resolution and arrival management while Section V describes an algorithm for rerouting arrival traffic around weather cells. The results of fast time simulation runs with traffic to and from Dallas-Fort Worth International Airport and Dallas Love Field are presented in Section VI. Finally, Section VII summarizes the key points of this paper and outlines future work to be pursued.

I. Overview of Terminal AutoResolver

The basic process for conflict resolution and arrival scheduling in TAR is depicted in a flow chart, Fig. 1. Aircraft-related information is an input, with data for type of aircraft, destination airport, aircraft's current state of motion and route intent, as well as weather and terminal airspace configuration information. Using these data as input, an external trajectory synthesizer/predictor module generates a predicted four-dimensional trajectory for each flight inside the Terminal Radar Approach Control (TRACON) airspace, with time being the fourth dimension. TAR operates cyclically on the aircraft track data that is updated at the normal radar sensor update cycle of 4–5 seconds. If ADS-B data becomes available the update rate could be increased to once per second, which would provide performance improvements.

TAR's conflict probe uses the predicted trajectories to check all aircraft in the TRACON airspace for future loss-of-separation conflicts. In addition to loss-of-separation conflicts,

each arrival aircraft is checked for future time-based spacing conflicts at the final approach fix (FAF) against every eligible aircraft in the arrival stream destined for the same runway. A spacing conflict arises when two aircraft crossing the FAF to the same runway are separated by less than a distance required by wake vortex rules, which depend on the sequence order and weight category of the two aircraft.

The process of arrival scheduling and resolution of predicted spacing and loss-of-separation conflicts is coordinated by TAR's scheduling module, referred to as Arrival Manager. The Arrival Manager's basic approach to problem solving is to first resolve a predicted spacing conflict and then a loss-of-separation conflict if needed. In this way, an orderly and efficiently spaced stream of arrivals is achieved, free of both types of conflicts. Arrival Manager checks whether an aircraft's remaining time to fly to the FAF has become less than a specified minimum time, referred to as the freeze time interval, in which case the aircraft is assigned a scheduled time of arrival (STA) at the FAF. The aircraft's STA is set equal to its unconstrained arrival time if this would not cause a spacing conflict with a previously scheduled aircraft or a loss of separation conflict with any other aircraft. Otherwise, Arrival Manager will employ its scheduling algorithm to assign an STA to the aircraft that is currently being scheduled. The landing sequence order is generally determined by the First-Come-First-Served (FCFS) rule, which is applied when an arrival crosses the freeze time horizon for the first time. The reference position for the rule is the FAF, and its sequence order is based on the nominal undelayed trajectory. However, deviations from FCFS are allowed if the aircraft currently being scheduled can be assigned a landing slot between previously scheduled aircraft without requiring such aircraft to be rescheduled. The Arrival Manager also has the capability to reassign aircraft to a different runway at airports with more than one active landing runway. Normally, aircraft entering the TRACON have a runway preassigned to them, but occasionally a reassignment becomes desirable in order to reduce delays as shall be explained later.

TAR uses a trial planning process as the basic algorithm for solving combined spacing and loss-of-separation conflicts as well as conflicts with convection cells. For each type of conflict, either meeting an aircraft's STA generated by the Arrival Manager or resolving a predicted separation or weather cell conflict, the Resolution Generator module generates appropriate maneuvers to solve the problem. These include speed control, horizontal maneuvers, and altitude changes. The maneuvers generally are computed from analytical formulas for a range of predetermined parameters. For example, an analytical formula determines the coordinates of a waypoint for a path stretch that approximately achieves a given delay, as specified by the scheduler. Each maneuver is a coarse approximation of a trajectory that has not been checked for conflicts, but contains sufficient information for generating a more detailed trajectory. The parameters for a candidate maneuver, referred to as a trial plan, are then sent to the external trajectory synthesizer module which uses the parameters to create a high-fidelity four-dimensional trajectory of the aircraft from its current position to the runway. The module integrates point-mass equations of motion together with performance models specific to each aircraft type as well as current atmospheric data. Next, the resulting trial plan trajectory is sent to the Conflict Probe to determine whether it has solved the original conflict without introducing new conflicts, referred to as secondary conflicts. The Conflict Probe checks the trial plan trajectory in

small time steps against all other trajectories for loss of separation. If the trial plan trajectory is found to be free of all problems, it is retained. If not, TAR loops back to the Arrival Manager/ Resolution Generator, which generates another trial maneuver, as indicated in Fig. 1. The Conflict Probe also provides failure diagnostics to help the Resolution Generator choose an improved maneuver. In complex traffic scenarios, more than a hundred iterations may be required to solve a problem, although most require fewer than 20 iterations. Instead of stopping the resolution process after finding the first successful resolution, TAR continues searching for additional resolutions using alternative maneuvers. When more than one successful resolution is found, the optimum resolution selector generally chooses the maneuver to be implemented as the one that gives the least delay, although other factors may also influence the choice. The selected resolution maneuver is then issued to the conflict aircraft.

The above process is repeated for all aircraft in the TRACON airspace, at a rate determined by the surveillance sensor update rate. Four-dimensional trajectories are recomputed at each update using the most recent sensor inputs. In the current version of TAR, it is assumed that the trajectories generated by the external trajectory synthesizer are flown with a high level of accuracy, essentially without errors. Future work will account for trajectory prediction and execution uncertainties.

The Terminal AutoResolver processes the list of detected conflicts in a time ordered sequence determined by time to loss of separation. The time ordered list is determined by checking all possible pairs of aircraft for conflicts and then sending each conflict pair to the Resolution Generator for resolution. Each resolution is constructed so as to avoid causing new (secondary) conflicts to be introduced. Thus, a new resolution trajectory becomes a constraint on the resolution of the next-in-sequence conflict. In other words, the resolution maneuvers for all previously resolved conflicts at any moment in time are taken into account when working on the next conflict in the time ordered list. This strategy guarantees that pairwise generated resolutions do not interfere with previous resolutions.

Software for the Terminal AutoResolver was coded in the Java programming language and comprises approximately 180,000 lines of code. This number is expected to increase substantially as additional functions are implemented.

II. Arrival Manager

TAR's Arrival Manager first attempts to first resolve a predicted spacing conflict and then a loss-of-separation conflict if necessary, as previously explained. An exception arises when the time until first loss of separation is below a certain threshold, typically one minute, in which case priority is given to resolving separation conflicts first, followed by resolution of spacing conflicts. If the nominal arrival trajectory produces a weather conflict, TAR first generates a new arrival trajectory that avoids the weather cell and then solves the spacing conflict if necessary, taking care not to penetrate the weather cell in the process. A later section describes this process in greater detail. Moreover, when an arrival aircraft is in loss-of-separation conflict with a departure or an over-flight that traverses the TRACON, preference is given to maneuvering the non-arrival aircraft to avoid disrupting the arrival

schedule. The following subsections describe TAR's logic for managing arrival traffic, which is encapsulated in the Arrival Manager module.

A. Eligibility for Scheduling

Trajectory Synthesizer provides the Arrival Manager with updated values of estimated times of arrival to the final approach fix (FAF) for all aircraft already in the TRACON and any new aircraft that have entered the TRACON since the previous update cycle. When an aircraft's remaining flying time to FAF first becomes less than a specified amount, referred to as freeze time interval, the aircraft becomes eligible to be assigned a scheduled time of arrival (STA) at the FAF. In case the aircraft's flying time to the FAF is already less than the freeze time interval by the time it is detected inside the TRACON, the aircraft is assigned an STA immediately upon entering the TRACON. To ensure stability of the arrival process, once an STA has been assigned to an aircraft it is generally not changed again unless an unplanned loss of separation is predicted or the aircraft fails to fly the assigned trajectory. This is an important characteristic of TAR's scheduling process and is similar to the so-called 'freeze horizon' used by the Traffic Management Advisor (TMA) [7], an operational system for managing arrivals installed in all Air Route Traffic Control Centers and large airports in the US.

For the case of the Dallas/Fort Worth TRACON, experience from simulation experiments suggests that suitable values for freeze time interval range between 9 to 12 minutes. Shorter values frequently resulted in more aggressive resolution maneuvers – such as large heading changes or large reductions in speed – due to smaller available time for aircraft to absorb delays. On the other hand long freeze horizons will be more susceptible to trajectory prediction errors which can result in additional maneuvers to avoid conflicts.

B. Spacing and Loss of Separation Conflicts

A spacing conflict arises when two aircraft are predicted to cross the FAF within less than a minimum time interval, which itself depends on the sequence order and weight category of the two aircraft. In practice, successive arrivals must maintain a minimum required separation distance throughout final approach, which starts at the FAF and ends at the runway threshold, as shown in Table 1. The values given in the table are determined by wake vortex separation rules and are generally larger than the 3 nmi minimums required for separation assurance. Arrival Manager, taking into account the speed on final approach of successive arrivals, converts the minimum required separation distances of Table 1 into minimum required time separations at the FAF, shown in Table 2. Whereas future versions of Arrival Manager will perform an aircraft-specific calculation, in the current version an average speed by aircraft weight category was used that yielded the minimum time separations listed in Table 2. It should be noted, though, that for cases when the trailing aircraft is significantly faster than the leading aircraft on final approach, the time separation of Table 2 is increased as necessary to ensure that minimum required separation distance is not violated anywhere on final approach.

A loss of separation (LOS) conflict occurs when two aircraft anywhere in TRACON airspace come within a certain horizontal and vertical distance. While many special

regulations exist in practice, the general rule of 3 nautical miles horizontal and 1000 ft vertical distance was used in this research. Future work will include additional rules for defining a LOS conflict applicable to special situations in TRACON airspace.

C. Scheduling Process

At each scheduling epoch, the Arrival Manager attempts to assign an STA to those aircraft that are eligible for scheduling and have not been assigned one. For each of those aircraft, the Arrival Manager computes its original estimated time of arrival (OETA) to the FAF. This OETA is the time the aircraft will cross the FAF if it follows its nominal flight plan inside the TRACON and no conflict resolution maneuver is assigned to it. Arrival Manager first checks if setting the aircraft's STA equal to its OETA causes any spacing or spatial conflicts with aircraft that have already been assigned a frozen STA. If no problems are detected the aircraft is scheduled to cross the final approach fix at its OETA.

As previously stated, the First-Come-First-Served rule based on the OETA at the FAF determines the landing sequence. This rule applies to all aircraft, including those that merge into a single stream at one or more route junction points before reaching the FAF. In-trail as well as crossing conflicts arising anywhere along the entire trajectory up to the FAF are resolved by the Resolution Generator using any of several route change maneuvers described in later sections. For these conflicts upstream of the FAF, the 3 nmi and 1000 ft. separation requirement is enforced. The more complex vortex separation rules given in Table 2 are used for in-trail spacing conflicts between the FAF and touchdown. This approach differs from sequencing rules used in existing decision support tools, which select a sequence at the intermediate merge points using logic that involves FCFS rules as well as controller preferences. This method of ordering was used in the Final Approach Spacing Tool [3] which also attempted to resolve conflicts between arrival aircraft on common route segments upstream of merge points. However, it was not designed to resolve all types of conflicts that can occur. Sequencing constraints at upstream merge points can produce additional delays and landing slot losses. Such losses are eliminated or substantially reduced by the methods used in TAR where controller preferences need not be considered.

If a spacing conflict arises, the Arrival Manager searches for gaps in the schedule of frozen aircraft where the new arrival can fit in. In the majority of the cases, a slot is added at the end of the scheduling queue behind all frozen aircraft. The additional flying time measured relative to OETA is defined as the scheduling delay. The Resolution Generator module of TAR (see next section), in collaboration with the Trajectory Synthesizer, is responsible for generating trajectories that are free of predicted loss of separation with other air traffic and deliver the new arrivals at the FAF with proper spacing from other frozen aircraft. To achieve that, Resolution Generator initially generates a trial maneuver that is estimated to absorb the delay needed for the new arrival, as computed by the scheduler in the Arrival Manager. However, if the four-dimensional trajectory corresponding to this trial maneuver does not meet the specified STA within specified tolerance bounds (typically ± 5 seconds) at the FAF, then Resolution Generator generates new trial maneuvers iteratively in small increments/decrements of time. After a trial trajectory has been found that achieves the specified STA within tolerance bounds and the Conflict Probe module determines that no loss of separation

conflicts have been introduced, the Arrival Manager designates that aircraft as frozen and scheduled. Thus, after this process has been completed for all new aircraft at the current scheduling epoch, the STA timeline shows only frozen aircraft. If, however, loss-of-separation is discovered in the conflict checking process, the Resolution Generator first tries an elliptic path stretch [5] to modify the trajectory spatially without changing the specified STA in an attempt to clear the loss of separation. If those attempts fail to clear the conflict, the Arrival Manager increases the STA in small increments of time. Each time-incremented STA is chosen not to conflict with STA's of frozen aircraft. Resolution Generator then generates the 4D trajectory for the incremented STA and checks that trajectory for loss of separation conflicts with frozen arrival aircraft as well as for short-range loss-of-separation conflicts with other traffic. This process continues until a conflict-free trajectory is found. Fig. 2 shows a simplified flow chart of the Arrival Manager focusing on the iterative interactions between the scheduler, elliptic path stretch algorithm and conflict resolution. It should be mentioned that if the fixed-time function of the elliptic path stretch algorithm does not resolve the conflict or cannot be used for a given conflict scenario, then it is necessary to increment the STA's. Thus, it is possible that the resolution of a loss of separation conflict remaining to be solved after a trajectory that meets the specified STA has been found will introduce additional delays. Simulation results described in Chapter VI indicate this situation occurs relatively infrequently.

The Arrival Manager then waits for real time to advance to the next scheduling epoch. When that time is reached, the Arrival Manager repeats the process for all new aircraft that have crossed the freeze horizon. The Arrival Manager performs this process for traffic flowing to any of the runways at an airport.

The procedures described are illustrated in Fig. 3, which shows both a time line plot and horizontal view of traffic converging on a final approach fix. A new arrival, A4, has crossed the freeze horizon and has become eligible for scheduling. Its OETA indicates it is in spacing conflict with the frozen aircraft A3. By computing the time difference between the A4's OETA and A3's STA, the Arrival Manager determines that a sufficient time gap is available to try decreasing the speed of A4 in order to resolve the spacing conflict between A3 and A4. The arrival manager then calculates the STA for A4 by adding the minimum required time separation – extracted from Table 2 - to the STA of A3. Next, the arrival manager asks the resolution generator to find a conflict-free trajectory for A4 to meet the specified STA. In this case, a reduced-speed descent profile that meets the required conditions is found. After the scheduling process for all new arrivals is completed, the arrival manager waits for time to advance to the next scheduling update cycle.

D. Change of Runway

Runway assignments are usually made before arrivals enter the terminal area airspace (TRACON) and are based on airline preferences, runway load balancing and other factors. Arrival Manager provides a function to change the original runway assignments at the time arrivals enter TRACON airspace if doing so reduces delays. Although changing runways adds to the flight crew's workload, it may be desirable in order to keep the arrival flow balanced on all active landing runways of an airport. Arrival Manager issues a change of

runway if landing on an alternate runway saves more than a specified minimum amount of delay, assumed here to be one minute, than landing on the preferred runway. The choice of one minute as the triggering criterion for initiating a runway change is somewhat arbitrary and can easily be changed in the software. However, because of the extra workload on pilots to reconfigure their Flight Management Systems for a change in runway, some minimum payoff in delay reduction before initiating such a change will most likely be required in practice.

In particular, Arrival Manager will attempt to find a conflict free trajectory for landing on any available alternate runway and assign a provisional scheduled time of arrival STA^{alt} . If the difference $STA^{pref} - STA^{alt}$ is larger than one minute, where STA^{pref} is the STA for landing on the aircraft's preferred runway, then a change of landing runway will be issued. In case there are more than two active runways for landing and all alternate runways yield more than one minute of delay saving, the alternate runway with greater delay saving is selected.

Finally, it should be noted that the Arrival Manager has been designed to operate either standalone or in conjunction with a separate scheduling and metering system such as TMA, which is widely used to control arrival traffic at US airports [7]. If the latter is the case, Arrival Manager then accepts STAs and runway assignments as inputs from the external metering system and only needs to generate conflict-free trajectories that satisfy those inputs.

III. Resolution Generator

The Resolution Generator is TAR's module that orchestrates the trial planning process for resolving predicted conflicts. Its role is to generate a set of alternative maneuvers for resolving a predicted conflict. The types of resolution maneuvers included here are patterned after those typically used by controllers in managing traffic in the terminal area. They have the advantage of supporting the transition from today's controller-centric operations to autonomous operations by making it easier for controllers to maintain situational awareness of the automated process. In general, higher priority is given to those maneuvers that create less delay and that deviate less from the nominal flight plan trajectory or, if delay is not a significant factor, to those maneuvers that follow rules controllers would typically use to resolve a similar type of conflict. Furthermore, it is important to note that all the maneuvers described here are compatible with trajectory specifications built into current generation of on-board Flight Management Systems (FMS's). This compatibility allows the resolution trajectories to be inserted into the FMS by the flight crew or to be uplinked directly into the FMS via a data link.

A. Speed Reduction

Speed reduction is the most desirable type of maneuver for pilots and air traffic controllers, since it does not modify the aircraft's planned horizontal route to the airport. It is therefore the preferred method for resolving spacing conflicts. Fig. 4 shows an aircraft's calibrated airspeed (CAS) as a function of time, both for the nominal case of its descent to the final approach fix and for the case of speed reduction. The aircraft's initial speed of 250 knots is

reduced to 220 knots. As a result, the aircraft crosses the final approach fix approximately 90 seconds later. Speed reductions can absorb up to about 2 minutes of delay inside typical TRACON airspace. If operational constraints dictate a smaller range of speeds than assumed here, less delay can be absorbed by speed reduction, and path stretch maneuvers will have to be used more frequently.

B. Hold Speed (maintain current CAS)

When an arrival aircraft is being scheduled to the FAF earlier than its OETA, increasing the aircraft's speed is usually not an acceptable maneuver since it is constrained by the speed limit of 250 knots CAS in the TRACON. Instead, TAR computes a maneuver that requires the aircraft to maintain its current speed for a longer time interval compared to its nominal descent procedure. In this way, the aircraft spends more time flying at a higher speed and it arrives at the FAF earlier (see Fig. 5 for an example). The extended hold of a higher speed must be terminated when the aircraft crosses a specified altitude and/or range to the final approach fix where deceleration to a reduced speed must begin. That limits the time reduction achievable by this method to less than one minute for most aircraft types.

This maneuver is defined by one parameter, namely the additional time interval of maintaining the aircraft's current speed. Maintaining a high speed for a longer time interval, however, might result in a higher descent rate to the FAF. The Trajectory Synthesizer module will reject those trial plans that yield a trajectory with a descent rates outside of the aircraft's performance envelope.

C. Extension of Base Leg ("Trombone")

For aircraft whose trajectory includes a downwind leg, extending this downwind leg prior to turning to base leg is a frequently used maneuver by air traffic controllers. Fig. 6 depicts a case where the base leg of the aircraft was extended by 3 nautical miles in order to resolve a spacing conflict. In TAR, such trombone resolutions are used if speed reductions alone are insufficient to solve the spacing conflict. Trombone solutions are always used in combinations with speed reductions as much possible. In the trial planning process, the turn to base is extended in increments of half a mile which yields adequate time control accuracy.

Unlike the other horizontal maneuvers of TAR that require the aircraft to deviate from its planned route immediately, the trombone maneuver begins only when the aircraft is about to turn to the base leg from its downwind leg. Therefore, it has the advantage of being available relatively close to the FAF to correct cumulative trajectory errors that produce conflicts or landing time errors. Such modifications can be done even if the aircraft is only a short time (30–60 seconds) from the revised turn-to-base waypoint.

The trombone maneuver has the capability of absorbing large amounts of delay as it can extend the downwind leg for an additional 5 or even 10 nautical miles. Moreover, it can do so in small increments, thus proving a powerful tool for efficiently resolving spacing conflicts.

D. Path Stretch

Two types of path-stretch maneuvers are used in TAR. Trial planner first attempts to find a resolution using a symmetric path stretch (see Fig. 7). It selects an auxiliary waypoint that lies on the perpendicular and midway to the vector that connects the aircraft's current position and the return waypoint. Symmetric path stretch generates a one-parameter family of maneuvers, defined by the turn angle from the aircraft's current heading. Symmetric path stretch maneuver is a powerful means to resolve spacing conflicts, where the primary goal is generating additional flying time for the trailing aircraft. As in trombone maneuvers, aircraft's current speed is first reduced as much possible, and it is this new speed profile that is used in generating a trial path stretch maneuver. The trial planner generates both left and right handed turns when maneuvering airspace is available in either direction. The alternative direction can be used to resolve a separation conflict if one of the directions generates a conflict. The choice of turn direction will also be influenced by the direction giving the greatest headwind, which has the benefit of reducing path deviations.

If no successful resolution is found using a symmetric path stretch maneuver, TAR will conduct a more extensive search for resolutions by employing the constant-delay elliptic path stretch algorithm. This type of path stretch is characterized by two parameters: a specified delay and the turn angle relative to current heading. A detailed description of the algorithm can be found in Ref. 5. Elliptic path stretches can sometimes resolve loss-of-separation conflicts while maintaining an aircraft's STA at a fixed value by trying different turn angles that yield the same time-to-fly until touchdown. It is the only maneuver available that can modify a trajectory spatially without changing its landing time. Path stretches of either type provide a versatile tool for resolving conflicts involving arrivals, departures and overflights.

E. Horizontal Vector Turn

The horizontal vector turn (hvt) algorithm was originally designed to resolve short-range conflicts where it is important to include the effects of turn rate limits in the generation of the resolution maneuver [8]. The algorithm has also been adapted for use in the TAR to resolve conflicts over longer time ranges while retaining its ability to resolve short-range conflicts. However, its primary use is for resolution of conflicts in the range of 1–3 minutes to LOS.

Given a maximum bank angle or turn rate, the hvt algorithm determines the heading change required to achieve a specified separation distance, assuming both conflict aircraft are flying at constant airspeed. The algorithm computes an explicit solution for a maneuver giving a specified minimum separation without trial planning and is the only algorithm in TAR with this property. However, trial planning must still be used in conjunction with the algorithm to generate a return path for the aircraft via an auxiliary waypoint and a return waypoint, as illustrated in Fig. 8. The hvt algorithm provides the trial planner/trajectory synthesizer with the coordinates of a point on the straight-line segment that immediately follows the heading change. It defines the point where minimum separation is reached, as computed by the hvt algorithm. Then, the auxiliary waypoint is located on the straight-line segment an incremental distance beyond this point. The incremental distance, taken to be the equivalent

of about 30s of flight time, ensures that a turnback maneuver starting at the auxiliary waypoint will not cause the original conflict to reappear. While the hvt generated trajectory segment resolves the original conflict up to the auxiliary waypoint, trial planning followed by a conflict check is still required to ensure that the entire trajectory is free of secondary conflicts. Should a secondary conflict be discovered, hvt provides a choice of alternative maneuvers which include a turn in the opposite direction of the initial solution as well as maneuvers by the other conflict aircraft in hard to resolve cases. In order to compensate for uncertainties in an aircraft's execution of the turn maneuver, the hvt algorithm could be used in a closed loop to update the resolution maneuver. In the current implementation a new hvt trajectory is generated if the conflict is found to reoccur after a certain time has passed.

F. Delayed Turn Back For Departures

Analogous to the trombone maneuver for arrival aircraft, the delayed turn back (dtb) maneuver is designed for resolving conflicts involving a climbing aircraft whose departure route inside the TRACON involves a large change of heading after take-off. Fig. 9 illustrates this case. Specifically, the maneuver is intended for an aircraft whose ultimate heading toward its destination requires the aircraft to make a large heading change (more than 90 degrees) from its initial heading following takeoff. For such an aircraft, the algorithm creates an auxiliary waypoint, located on the line segment that connects its current position with the next waypoint on the original route and an incremental distance beyond this point. The algorithm increases the delay in small increments to find the smallest delay that resolves the conflict. A second auxiliary waypoint, located at a 90 degrees angle relative to the takeoff heading is used as a turn back point where the aircraft can resume flight towards a downstream fix on its flight plan. The dtb algorithm adds to the set of available resolution options for departure aircraft. It is used as an alternative to path-stretch resolutions for situations where the departure route is similar to the one shown in Fig. 9.

G. Extension of Final Approach ("Fanning")

In certain situations, when an aircraft is on a heading to intercept the final approach path, an alternative to path stretching is changing the heading of the aircraft so that it intercepts the final approach path further upstream (see Fig. 10). In this way, the need for specifying an auxiliary waypoint is avoided and the entire maneuver can be communicated more easily to the pilot. The pilot is simply given a new heading and the instruction to intercept the final approach. Again, the aircraft's current speed is reduced to the lower limit of its speed envelope prior to trying a fanning maneuver.

A somewhat restricting factor for this maneuver is the interception angle with the final approach path. In order to ensure a smooth transition into the final approach, the algorithm requires that the interception angle not be greater than about 60 degrees. This limit can be changed if smaller values are required for operational acceptance. This constraint limits the amount of delay that can be absorbed by this maneuver to relatively small values. However, it is still a useful option for situations where speed reduction is not quite sufficient to absorb all delay, since it has the advantage of not requiring the specification of additional auxiliary waypoints.

H. Compound Horizontal Maneuvers

A combination of two horizontal maneuvers is available to resolve spacing conflicts that require a large amount of delay to be absorbed. When extension of final approach does not resolve the conflict, a symmetric or elliptic path stretch can be applied in conjunction with extending the final as illustrated in Fig. 11a. This type of resolution can absorb more delay compared to what a symmetric path stretch or extension of final can absorb separately.

Situations may arise where the return waypoint of a path-stretch maneuver lies on the final approach. Depending upon the position of the turn back point, it may happen that the aircraft must perform a turn of more than 90 degrees to align on final approach. Such a large heading change to capture the final approach is considered undesirable by pilots, causing the Resolution Generator to reject the trial plan. To resolve this issue, a special type of compound maneuver has been created that consists of a path-stretch and a constructed base-leg segment which provides a more acceptable transition to final approach; Fig. 11b provides an example. At the end of the path-stretch maneuver the aircraft transitions to a constructed base-leg segment, inserted only when the intercept angle with final approach exceeds 90 degrees.

I. Temporary Altitude during Climb

This maneuver is used to resolve conflicts involving an aircraft that is currently climbing toward its assigned cruise altitude. The time of first loss with the conflict aircraft may occur during the climb segment or after the aircraft has leveled out at its cruise altitude. It is illustrated in Fig. 12, wherein the dashed line represents the temporary altitude hold and return to cruise maneuver.

A temporary altitude maneuver for an aircraft in climb is defined by two parameters. The first specifies the flight level at which the maneuvering aircraft must level out (also referred to as the temporary altitude level). Only flight levels at least 1500 ft. above the current altitude of the climbing aircraft and 1000 ft. below the flight level of the altitude where first loss occurs are eligible for temporary altitude levels. All eligible altitude levels, starting with the highest one, are evaluated via trial planning iteration in an attempt to resolve the conflict. The second parameter that defines the temporary altitude maneuver is the time at which the maneuvering aircraft can resume its climb to capture the originally assigned flight level. The Resolution Generator estimates the earliest time to resume the climb by an analytical expression and then checks the validity of the estimate via the trial planning process. If the conflict remains unresolved, it will increment the resume-the-climb time up to a maximum value.

J. Resolutions generated for Departures at Brake Release Time

An important goal of TAR is to provide an environment that allows aircraft to fly more efficient trajectories in the terminal area. One approach to achieve this goal involves reducing certain procedural restrictions and/or airspace segregation between departures and arrivals. In today's manually operated system these restrictions often force aircraft to fly non-direct routes and non-optimum vertical profiles. Not surprisingly, an undesirable byproduct of removing these restrictions is an increase in the number of short range conflicts

between departures and arrivals, as fast time simulations discussed in Section VI have shown to occur. While these conflicts can be detected and resolved using one of the maneuver types described earlier, they are nevertheless undesirable and may be operationally unacceptable to pilots, especially if they occur within a short time after takeoff. The brake release resolution methods outlined in this section largely avoid short range conflicts during the initial climbout phase of flight. They differ from the separation assurance techniques previously described in that they are generated and issued while the aircraft is still on the ground, but in position on the runway to start the take-off roll.

It is generally recognized that the movement of traffic on the airport surface prior to takeoff is too unpredictable to be useful in conflict prediction between traffic on the ground and in the air. Variabilities in taxi times would generate an excessive number of false and missed alarms in the conflict prediction process to be operationally useful. Furthermore, a resolution process initiated while the aircraft is taxiing would pose complex implementation challenges.

To circumvent these prediction uncertainties, it is proposed to generate and issue resolutions only after an aircraft is in position on the runway, ready to release brakes for the take-off roll. When an aircraft is in position on the runway, gate hold and taxi time uncertainties have been eliminated and accurate trajectory synthesis, required for reliable conflict prediction with airborne flights, can now be performed. The known performance model of the aircraft, together with wind and temperature measurements at the airport, allows the trajectory synthesizer to compute the takeoff roll and climbout trajectory accurately starting at the moment of brake release. This trajectory together with the trajectories of flights in the air can then be used to detect conflicts. The detected conflicts are resolved by the methods described below.

The two types of resolutions that are proposed to be issued at brake release time are a delay in starting the take-off roll and a temporary altitude hold as illustrated in Fig 12. A third candidate for this type of resolution that was not implemented is a change in departure route, although it is an option that may be investigated in the future. Changes in departure routes at take off time are probably less desirable than the two types investigated here, because they would be more difficult for the flight crew to implement on short notice.

It is important to note that the conflict check and generation of the resolution maneuver for the detected conflicts are held in abeyance until the aircraft reports that it is ready to start the take-off roll. At that moment, TAR attempts to generate two resolutions, one for each of the two types of brake release resolutions described above. The resolution yielding the shortest delay is then immediately issued to the flight crew. For resolutions requiring a delay in brake release time, delays are computed and specified in multiples of 30 seconds, that being considered the smallest practical delay interval that the flight crew can execute accurately. TAR performs a series of trial plan iterations starting with the shortest delay interval and continuing in increments of 30 seconds, up to a specified maximum value (~3 minutes) until resolution of the detected conflicts with flights in the air is achieved. Similarly, TAR will attempt to generate a temporary altitude hold resolution by the method described in section I above. It should be mentioned that if only a single runway is used for both departures and

arrivals, the brake release procedure may have to be limited to the temporary altitude hold maneuver during periods of high traffic for interleaved arrivals and departures.

In generating the resolutions, TAR tries to achieve a conflict-free time interval of at least 6 minutes or until the departure flight leaves TRACON airspace, whichever occurs first. It should be noted that unlike airborne maneuvers that require changes in the trajectory, the delay in brake release time does not change the trajectory, only its starting time. On the other hand, the temporary altitude maneuver issued at brake release time does change the altitude profile of the original climbout trajectory. From an operational point of view it is preferable for the flight crew to receive the temporary altitude hold clearance before takeoff, thus giving time for the flight crew to enter the altitude hold restriction into the flight management system.

If both types of resolution are available, TAR generally chooses the one giving the least delay as the one to be issued. That selection criterion will often favor the temporary altitude resolution type, because a short period of level flight during the climbout will most likely introduce less delay than a 30 second brake release hold time. An alternative selection criterion is minimum fuel consumption, which could be used if fuel flow models for aircraft types are stored in TAR's aircraft performance data base. This criterion may favor short ground delays over temporary altitude resolutions since fuel flow rates with engines at idle on the ground are substantially less than fuel flow rates in flight.

It should be noted that the brake release delay method also required a new scheduling procedure to be added to TAR. This was necessary to account for backward delay propagation to aircraft lined up for departure on the same runway behind an aircraft that was issued a brake release delay.

Brake release control at the runway is a procedure already in use by controllers in today's operations to enforce a miles-in-trail program between aircraft when they transition into enroute airspace. That procedure is unrelated to its use here as a separation assurance/conflict resolution method for autonomous air traffic management in terminal airspace. However, a potentially important attribute of the two brake release resolution clearances is that they could be implemented as traditional controller decision support tools and used manually in today's operational system for air traffic management. These clearances could be delivered to the flight crew by conventional voice communications and would not require a data link. Pilot-in-the-loop simulations will have to be conducted to determine how much buffer to add to the separation requirements in order to avoid conflicts due to pilot errors in performing the brake release procedures.

V Combined Weather Cell Avoidance and Arrival Scheduling

At many airports in the U.S., thunderstorms, also referred to as convective weather cells, pose a significant flight hazard to airline traffic, especially during the summer months. When such cells are severe and pass through terminal area airspace, air traffic is often rerouted to avoid penetrating them. Because it is typical for cells to evolve quickly in areas affected by them, flight path changes with short lead times may be necessary to avoid them. While

departures can be held on the ground to avoid flying into them, arrivals operating in TRACON airspace have to be prepared to perform avoidance maneuvers with little advance notice. Depending on the density of arrival traffic, such flight path changes may require resequencing of the landing queue. This process poses a potentially challenging problem for the design of autonomous systems.

The approach taken here is an extension of the one described in Ref. [1] for combined weather avoidance and conflict resolution in en route airspace. Similar to the enroute case, it is assumed that weather radar sensors detect the location and severity of the cells, which are then processed by embedded convective weather analysis software that assigns risk levels for regions of airspace around the convection cells to be avoided by air traffic. These regions are input to algorithms that generate polygons enclosing them [10]. TAR uses these polygons as input to generate conflict-free avoidance trajectories. The polygons, which may be non-convex, include a sufficient number of vertices to ensure a close fit that covers the convective weather region.

The predicted trajectory of each flight crossing into TRACON airspace is input into a separate weather conflict probe that determines if it penetrates any weather polygon and, if it does, outputs conflict parameters such as time to penetration and length of conflict. Up to two waypoints may be used to construct a path around the conflict, although one-waypoint solutions are given preference in order to simplify the trajectory specification to the flight crew. A low limit on the number of waypoints may not be important in some applications such as UAV's or if uploaded via data link. The resolution algorithm uses linear algebra solutions from open software sources to locate two rays tangent to the polygon at two of its vertices. The first ray emanates from a point a short distance in front of the current position of the conflict aircraft and a second from the final approach fix in the reverse flight direction. The interception of the two rays specifies the location of the waypoint for the avoidance trajectory, as illustrated in Fig. 13. The large number of vertices used to form the polygon around the weather cell makes them indistinguishable in the figure. For exceptional combinations of initial aircraft positions and non-convex polygons, not generally occurring in TRACON airspace, the method generates two waypoint solutions as described in [1]. In all cases, two avoidance trajectories are computed for each weather conflict, one for a right turn and one for a left turn around the polygon. Then the trajectory yielding the least delay is selected if both turn directions are operationally acceptable.

Krozel et al [9] use dynamic programming to avoid weather cells in TRACON airspace, including limits on the number of turns. Their method could potentially be used as an alternative to the one described here if it is computationally fast enough to handle multiple aircraft in real time. However, the method does not include conflict resolution and scheduled time of arrival constraints as required for TAR. The algebraic method used here has several advantages over methods using optimal control algorithms such as dynamic programming. First, it is computationally fast, making it suitable for processing multiple aircraft in real time. Second, for the most frequently encountered weather conflicts in terminal area airspace, such as illustrated in Fig. 13, the tangent ray method yields minimum path length, given the constraints on the number of waypoints. Third, because the trajectories are structurally similar to path stretches used for conflict resolution and scheduled times of

arrival trajectories, they are also compatible with FMS capabilities. Furthermore, the method and algorithm have proven themselves in the Dynamic Weather Routes System [11] that has been jointly evaluated in operations by NASA and an airline for more than two years. The method has also been extended to avoid multiple polygons [12].

After the least-path-length avoidance trajectory has been determined, the next step is to send it to the Arrival Manager/Resolution Generator for resequencing and conflict checking via the trial planning process. If the arrival was previously frozen, the Arrival Manager first unfreezes the aircraft to make it again eligible for the scheduling. It replaces the original arrival trajectory penetrating the convection cell with two eligible weather cell avoidance trajectories (dashed lines), and then finds the earliest landing slot and corresponding STA relative to previously scheduled and frozen aircraft. That process may necessitate changing the landing order if the weather avoidance trajectory has introduced enough delay to cause it to lose its original place in the landing queue. Such is the case illustrated in Fig. 13 which shows that the lengthened path for C forces it to be resequenced behind D. Here it is assumed that C and D have similar speed profiles. As stated previously, Arrival Manager first attempts speed reductions along the avoidance path to meet the STA. If speed reduction is inadequate, it combines path lengthening with the lowest speed to achieve the STA. In this case, however, instead of using standard path stretch maneuvers, additional delay is created by moving the weather avoidance waypoint in small increments away from the weather polygon, as shown in the figure. This strategy avoids creating another waypoint for path stretching. Iteration via trial planning is used to accurately relocate the waypoint so as to achieve the specified STA. In the final step, the candidate trajectory is checked for conflicts using the Conflict Probe. In the circumstance that a conflict is detected, Arrival Manager/Resolution Generator adds more delay to the STA in small increments while staying clear of frozen STA's for other aircraft. Such delay is added by moving the weather avoidance waypoint in a direction away from the weather polygon. For each incremental STA the process described at the beginning of the paragraph is repeated. If near term conflicts are found, Resolution Generator inserts a separate path stretch maneuver and a new waypoint in front of the weather avoidance waypoint to avoid them.

VI Simulation Results

A fast-time simulation environment called the Airspace Concept Evaluation System (ACES) was employed to test Terminal AutoResolver's performance in spacing and deconflicting traffic. ACES is a gate-to-gate simulation of air traffic at airport, regional, and national levels, developed by NASA [13]. ACES generates flight trajectories using aircraft models obtained from the Base of Aircraft Data (BADA) [14] and traffic data consisting of departure times and actual flight plans obtained from Airline Situation Display to Industry (ASDI) files. These trajectories are then fed into TAR for conflict prediction and resolution of spacing and loss-of-separation problems.

In this study, air traffic to and from the Dallas/Fort Worth metropolitan area was simulated in ACES. Arrival and departure data for the two major airports of the area, the Dallas/Forth Worth International Airport (DFW) and the Dallas Love Field (DAL), were used as input to the simulation. Fig. 14 shows the Dallas/Fort Worth TRACON area, the runway systems of

DFW and DAL (not to scale) as well as the standard arrival (green) and departure (blue) routes to and from these airports. The flight plan data used in the analysis presented here include a 24-hour period that begins at 00:00 Coordinated Universal Time (UTC) on 4/25/2012. According to the methodology defined in Ref. [15], this traffic sample can be characterized as a high traffic volume and low delay day in the National Airspace System. The input scenario consisted of 900 arrivals and 915 departures at DFW and 229 arrivals and 245 departures at DAL totaling 2289 flights. This level represents an average day of traffic for DFW and DAL. Inbound and outbound flows are almost balanced, with the largest proportion of flights landing or taking off from DFW. Departures from the same runway were spaced a minimum of one minute apart. The simulation held departing flights at an airline's departure gate to achieve the one minute separation at the runway if departure times for flights were in conflict with each other. Delays resulting from such departure rescheduling at gates are not included in the delay analysis given below, since they are not under the control of TAR.

The south flow runway configuration was assumed for operations. For this flow configuration landings at DFW take place simultaneously and independently on runways 18R, 17C, 13R and 17L, and take-offs on runway 18L and 17R. At DAL aircraft were landing on runway 13L and taking off from runway 13R. Furthermore, in the simulation the TRACON area was treated as a single sector. Also, climb rate restriction on departures that are currently used to separate arrival and departure traffic were not enforced. The purpose of not enforcing these rules was to determine what kind of problems TAR would encounter without them. It should be understood that neither conflict nor delay rates given below can be easily compared with today's operations because of differences between the current manual operational environment and the relaxed constraints used in the simulation of a future automated environment.

During rush periods the chosen traffic sample produced delays which sometimes exceeded three minutes per aircraft. Moreover, it was assumed that TAR's Arrival Manager was operating with no external metering system to control the rate of arrivals from Center airspace into the TRACON airspace. This input scenario exposed TAR's Arrival Manager to free-flowing and uncoordinated traffic, which is useful in exploring its limits of performance. The uncoordinated flow sometimes resulted in flights being in LOS conflict at the moment they entered TRACON airspace. Since this would not be allowed to occur in practice, those flights found to be in conflict were delayed just enough to deconflict them prior to entering TRACON airspace. It should also be mentioned that convective weather was not simulated in the simulation runs that are presented here.

Table 3 provides traffic volume and conflict resolution statistics for the two airports in the study. At DFW, 900 arrivals were handled, of which 344 encountered spacing conflicts that were resolved. In only 15 of those cases, resolution of a spacing conflict created a LOS conflict with another aircraft that Arrival Manager treated as non-maneuverable, either because it was an arrival with a frozen STA or a departure that had already received another resolution maneuver. Arrival Manager resolved such additional LOS conflicts, which occurred upstream of the FAF, by adding further delay to the resolution maneuver. This additional delay would, of course, produce larger than minimum required separations along

the final approach path. The relatively small percentage of such cases supports the hypothesis that resolution of spacing conflicts also minimizes the occurrence of secondary conflicts, though it does not completely eliminate them. The result confirms the value of controller decision support tools that provide only time-of-arrival advisories without a method for detecting and resolving LOS conflicts. Controllers monitoring the traffic detect these occasional conflicts and resolve them by issuing additional clearances manually. However, the potential for such conflicts, even though infrequent, is not acceptable for an autonomous control system where controllers are not assumed to be in the loop to detect and resolve infrequent cases such as these.

There were 101 predicted LOS conflicts resolved between arrival and departure traffic at DFW and 7 conflicts between DFW departures. All of these conflicts were resolved by one of the two brake release resolution procedures and therefore did not require maneuvers to be issued in the air. Both types of conflicts are rare in today's operations because procedural rules and airspace segregation between arrival and departure flows prevent them from occurring most of the time. However, such constraints generally impose some level of efficiency penalties on flight operations. Future simulation studies using TAR will quantify these penalties and determined the tradeoffs between reduced procedural constraints, conflict rates and traffic control complexity. Thus it will be possible to determine the fuel efficiency benefits obtained by the elimination of procedural constraints imposed in current operations. For completeness the table also includes the same set of statistics for DAL.

In some of the simulation runs a few cases of unresolved LOS conflicts were encountered. None of those cases were inherently due to TAR running out of resolution procedures. They were mostly caused by the trajectory generator failing to produce a trajectory in response to a trial plan request from the resolution generator. These failures can occur when an input state falls outside the envelope of an aircraft performance model. Improved performance models and bug fixes in the ACES simulation are expected to eliminate such failures. Other cases were due to input traffic scenarios not previously encountered and were eliminated with software changes. However, it must be recognized that in a large, complex, proof-of-concept software system that is still evolving, the process of bug fixes will continue for some time as new scenarios are encountered.

In Table 4, delay statistics are summarized. Total as well as average delay values are provided for both airports. Average delays per arrival flight are seen to be slightly higher for DFW compared to DAL as one would expect. While total delays for DFW are much higher than for DAL, the delays are roughly proportional to the number of arrival flights to each airport. This indicates that both airports operate at traffic levels that are commensurate with the number and capacities of runways available for landings, though DAL operates somewhat below capacity. The last row in Table 4 shows the additional delay that was required to resolve a LOS conflict after a spacing conflict had been cleared. Thus, the 15 arrivals at DFW that experienced a spacing conflict as well as a LOS conflict (see table 3) had to absorb, on average, an additional 30 seconds of delay to clear that LOS conflict. These types of LOS conflicts, though rare, can occur because of differences in speed profiles as well as at intersections of arrival trajectories that feed different runways.

In order to evaluate the delay reductions obtained from runway reassignment, a simulation was run in which runway assignment optimization was inhibited and instead all aircraft were nominally assigned to runways with the shortest distance to fly between arrival fixes at the TRACON boundary and final approach fixes to runways. However, some degree of balancing of runway loads for traffic from arrival gates with high traffic rates was still performed in order to limit excessive delays. Table 5 includes metrics that compare this scenario to the case where runway assignment optimization was allowed.

The first row in Table 5 indicates that Arrival Manager found 103 cases at DFW and 16 cases at DAL where delay savings for changing the landing runway was large enough (greater than one minute) to justify this action. The second and third rows show the delay reductions on the entire arrival traffic, not just on the reassigned aircraft.

The second row in Table 5 gives the dimensionless ratio of the total delay savings due to runway reassignment optimization relative to the total delay savings only from runway-reassigned flights. Similarly, the third row gives the delay saving per runway-reassigned aircraft. For each ratio we have subtracted out the individual delay savings from the reassigned aircraft. This procedure reveals explicitly the delay reduction amplification effect of the runway change. Thus, as shown in the second row for DFW, each minute of delay reduction for the reassigned aircraft generated 3.5 minutes of delay reduction for the non-reassigned aircraft. Equivalently, as shown in the third row, at DFW an aircraft that changed its landing runway generated an average of 2.8 minutes of additional delay savings for the aircraft that were not reassigned. The 4th and 5th rows list the total delays for the original and reassigned runway cases, respectively.

Thus, altering the landing runway of an arrival aircraft when the potential delay savings is estimated greater than one minute for that aircraft alone reduces the delay for all traffic significantly. This is a well known phenomenon in queueing theory when demand is at or above capacity of a system. The benefits of runway reassignment are somewhat less for DAL, since the traffic at that airport is below capacity most of the time and therefore offers fewer opportunities for delay reductions. In the current version of TAR, the algorithm for runway reassignment evaluates the delay benefits only one aircraft at a time, which takes place at the freeze horizon when the aircraft is scheduled. It does not take into account the delay effect on traffic that is known to follow behind it. Nevertheless, the delay reductions obtained from the current version of the algorithm for reassignments are substantial. Optimization techniques that take this traffic into account could further improve upon these results. As is the case for most optimization methods, the delay reductions obtained here depend strongly on the initial runway assignments used for comparing delays. Thus, the more optimum the initial assignment is the less there is to gain by optimizing the assignment.

Another important metric recorded in the simulation is the distribution of the types of maneuvers used to schedule and space the arrival traffic. To resolve the 406 spacing conflicts (both airports, as listed in Table 3), speed changes alone accounted for 198 resolutions while the remaining 208 required one of the several types of path change maneuvers in combination with speed reductions. While upstream metering of arrival traffic is likely to

reduce the frequency of path change maneuvers to lower levels than found in this simulation, it can never be completely eliminated. Thus, path change maneuvers, also called vectoring, will play an indispensable role in the arrival control process of an autonomous system just as they do in today's manual operations.

The two types of resolutions issued at brake release were also implemented and evaluated in the simulation. In comparison to simulation runs conducted without this function enabled, it achieved the objective of strongly reducing the number of short range conflicts (those with less than 5 minutes to first loss) that had to be resolved in flight between departures and between arrivals and departures shortly after takeoff. TAR attempted to generate both types of resolutions for each conflict detected at brake release time and succeeded in finding both types for most of the conflicts. Brake release delay in 30 second increments was limited to 3 minutes. The average hold time for brake release was found to be about one minute and average length of time for the temporary altitude hold was 2.5 minutes. A total of 125 brake release resolutions were issued of which 100 were temporary altitude holds and 25 brake release delays. As expected, TAR selected the temporary altitude hold type whenever available since it generated less delay. As a result, all brake release resolutions generated an average of only 14 seconds of delay per resolution. It is understood that controller and piloted simulations will have to be conducted to establish the operational acceptability and execution accuracy of these resolution types.

Finally, two examples of conflict resolutions are discussed. Fig. 15a shows a spacing conflict due to insufficient time separation at the final approach fix between leading flight AAL2069 (yellow path) and trailing flight DAL1299 (red path). The trial planning process resulted in a trombone type resolution which delayed the turn to the base leg of flight DAL1299 by 3 nautical miles (green path). It should be mentioned that path-stretch maneuvers beginning near the start of the trajectory also succeeded in solving this problem, but were considered less desirable than base extension for this specific situation.

In another example, shown in Fig. 15b, arrival flight COA1815 (pink path) is in a loss-of-separation conflict with departure flight EGF503 (blue path). To resolve the conflict, the departure flight was issued a delayed turn back maneuver (green path). In TAR's logic, departures are preferentially maneuvered to resolve conflicts with arrivals in order to preserve the scheduled time of arrival for landing aircraft. This is an example of the type of short range inflight conflict that was avoided in simulation runs in which the brake release resolution procedures were enabled.

VII Concluding Remarks

This paper describes the basic functions and algorithms of a system for conflict resolution and arrival scheduling of air traffic operating in terminal area airspace. The system, referred to as Terminal AutoResolver (TAR), employs a variety of maneuvers including speed control, path stretching, base leg extension, brake release control and other types to resolve conflicts, avoid weather cells and schedule arrivals to runways. The design attempts to solve the complex problem of integrating a diverse set of algorithms into software suitable for future real time applications. The key design approach used in the implementation involves

an iterative feedback process between a Resolution Generator containing simple analytical models and a Trajectory Synthesizer with more complete models of aircraft dynamics and atmospheric conditions. This approach enables TAR to generate four dimensional trajectories for resolving the complete range of conflicts and arrival traffic problems encountered at large airports. The trajectories, which are based on aircraft-specific performance and dynamics models, ensure that aircraft can fly them within acceptable limits of performance.

TAR's logic has been implemented in the Java programming environment and experiments have been run using the ACES simulation platform. The DFW and DAL airports were used as a case-study. Recorded flight plans of arrival and departure traffic for 24 hours of operations at these airports were an input to the simulation. Analysis of simulation results indicates that TAR is able to solve all types of conflicts encountered, including combinations of arrival spacing problems, separation conflicts and convective weather cell avoidance.

Future work will focus on such problems as handling missed approaches, giving priority to emergency aircraft, and generally accounting for errors in trajectory execution. Another step in the development of this concept will be to evaluate it in human-in-the-loop simulations wherein controllers will be assigned to handle primarily airport and airspace management problems, such as change of runway configuration for landings and take-offs, and responding to special requests by pilots. A data communications link will be simulated to send TAR-generated trajectories to the aircraft. Such a link is an essential requirement for a future autonomous air traffic management system using the TAR concept. Lastly, future work will also focus on achieving a high degree of autonomy by enabling TAR to operate under uncertainties and to compensate for system failures. For example, TAR must be able to generate corrective trajectories for aircraft that failed to properly execute maneuvers issued to it.

The long term objective of this research is to create and validate a system design that provides the basis for an air traffic control system where critical controller tasks, including separation assurance and arrival sequencing, can safely be delegated to an autonomous agent, allowing controllers to focus instead on handling the types of problems that require a human's unique problem solving skills.

References

1. Erzberger H, Lauderdale TA, and Chu Y-C "Automated conflict resolution, arrival management, and weather avoidance for air traffic management", Proceedings of the Institution of Mechanical Engineers, Part G: Journal of aerospace engineering, Vol. 226, No. 8, 2012.
2. Davis TJ, Krzeczowski KJ, and Bergh C, "The Final Approach Spacing Tool," IFAC Thirteenth Symposium on Automatic Control in Aerospace, Palo Alto, CA, Sep. 1994.
3. Robinson J III, E., Isaacson D. R, A concurrent sequencing and deconfliction algorithm for terminal area air traffic control, AIAA Guidance, Navigation, and Control Conference and Exhibit, 14–17 Aug. 2000, Denver.
4. Swenson H, N., Thipphavong J, Sadosky A, Chen L Design and evaluation of the terminal area precision scheduling and spacing system, Ninth USA/Europe Air Traffic Management Research and Development Seminar (ATM2011)

5. Erzberger H “Automated conflict resolution for air traffic control,” Proceedings of the 25th International Congress of the Aeronautical Sciences, Hamburg, Germany, Sep. 3–8, 2006.
6. Nikoleris T, Erzberger H, Paielli R, Chu Y-C “Automomous system for air traffic control in terminal airspace”, AIAA ATIO Conference, paper # 2014–2861; doi: 10.2514/6.2014-2861; June 2014, Atlanta
7. Swenson HN, Hoang T, Engelland S, Vincent D, Sanders T, Sanford B, and Heere K “Design and Operational Evaluation of the Traffic Management Advisor at the Fort Worth Air Route Traffic Control Center,” 1st USA/Europe Air Traffic Management Research and Development Seminar, Online Proceedings, 1997.
8. Erzberger H, Heere K Algorithm and operational concept for resolving short range conflicts. Proc. IMechE, Part G: J. Aerospace Engineering, 2010, 224(G): 225–243. DOI: 10.1243/09544100JAERO546..
9. Krozel J, Lee C, Mitchell S, Turn-constrained B route planning for avoiding hazardous weather, Air Traffic Control Quaterly, 2006, Vol. 14(2) 159–182.
10. Rubnich M, DeLaura R “An Algorithm to Identify Robust Convective Weather Avoidance Polygons in En Route Airspace”, 10th AIAA Aviation Technology, Integration, and Operations (ATIO) Conference, paper # AIAA 2010–9164; DOI: 10.2514/6.2010-9164; 13–15 Sep. 2010, Fort Worth, Texas
11. McNally D, Seth K, Gong C, Borger P Operational Evaluation of Dynamic Weather Routes at American Airlines, Tenth USA/Europe Air Traffic Management Research and Development Seminar (ATM2013), June 2013, Chicago.
12. Lauderdale T, A., Erzberger H Automated separation sssurance with weather and uncertainty, ENRI International Workshop on ATM/CNS, Tokyo, 2013.
13. George S, Satapathy G, Manikonda V, Palopo K, Meyn L, Lauderdale TA, Downs M, Refai M, and Dupee R, “Build 8 of the Airspace Concept Evaluation System,” AIAA Modeling and Simulation Technologies Conference, No. AIAA-2011–6373, Portland, OR, 2011.
14. “User Manual for the Base of Aircraft Data (BADA) Revision 3.9,” Eec note no. 11/08, Eurocontrol Experimental Centre, 11 2008.
15. Chatterji GB, and Musaffar B, “Characterization of Days Based on Analysis of National Airspace System Performance Metrics,” AIAA Guidance, Navigation and Control Conference and Exhibit, Online Proceedings, AIAA, Reston, VA, 2007.

Figure 1.
Functional diagram of Terminal AutoResolver

Figure 2.
Simplified flow chart of Arrival Manager

Figure 3.
Scheduling process for new arrival A4.

Figure 4.
Speed reduction profile

Figure 5.
Holding CAS for early arrival.

Figure 6.
Trombone.

Figure 7.
Path stretch

Figure 8.
Horizontal vector turn

Figure 9.
Delayed turn back maneuver.

Proc Inst Mech Eng G J Aerosp Eng. Author manuscript; available in PMC 2020 February 04.

Figure 11a.
Path stretch with extension of final approach

Proc Inst Mech Eng G J Aerosp Eng. Author manuscript; available in PMC 2020 February 04.

Figure 12.
Temporary altitude during climb

Figure 13.
Rerouting to avoid convection cell, followed by path stretch to resequence landing order.

Figure 14.
Diagram of the DFW TRACON. Runway lines not on scale.

Figure 15a.
Example of spacing problem resolution.

Figure 15b.
Example of loss-of-separation resolution.

Table 1.

Minimum required separation distance in nautical miles under Instrument Flight Rules (source: Ref. 4)

		Trailing aircraft			
		Heavy	B757	Large	Small
Leading aircraft	Heavy	4	5	5	6
	B757	4	4	4	5
	Large	3	3	3	4
	Small	2.5	2.5	2.5	3

Table 2.

Minimum required time separation in seconds used by Arrival Manager

		Trailing aircraft			
Leading aircraft		Heavy	B757	Large	Small
	Heavy	82	118	118	150
	B757	60	64	64	94
	Large	60	64	64	94
	Small	60	64	64	68

Table 3.

Traffic volume and conflict resolution statistics

	DFW	DAL
Total number of arrivals	900	249
Total number of departures	915	245
Total spacing conflicts resolved	344	62
Spacing conflicts that required an additional solution of a LOS conflict with a frozen aircraft	15	1
LOS conflicts resolved: Arrivals vs. Departures	101	14
LOS conflicts resolved: Departures vs. Departures	7	17

Table 4.

Delay statistics

	DFW	DAL
Total delay for departures (hh:mm)	00:14	00:12
Average delay per departure flight (sec)	1	3
Total delay for arrivals (hh:mm)	7:14	1:30
Average delay per arrival flight (sec)	28	21
Additional delay incurred for resolving a LOS conflict after the original arrival spacing conflict was solved (avg/flight) (sec)	30	45

Table 5.

Runway reassignment statistics

	DFW	DAL
Number of aircraft assigned to different runway	103	16
Change in delay for all arrivals except reassigned aircraft divided by change in average arrival delay of reassigned aircraft	3.5	0.4
Change in delay for all arrivals except reassigned aircraft divided by number of reassigned aircraft (min/flight)	2.8	0.4
Total delay, original runway assignment (hours)	13.5	1.9
Total delay, reassigned runway (hours)	7.2	1.5