A Linear Altitude Rule For Safer and More Efficient Enroute Air Traffic

Russell A. Paielli

NASA Ames Research Center

Abstract: Current regulations designate cruising altitudes in discrete steps of 1000 ft (or 2000 ft above FL290, but those will eventually be reduced to 1000 ft also). By concentrating all enroute traffic into a few altitudes, this discrete altitude rule makes the enroute air traffic system less fail-safe and less fault-tolerant than it could be. The linear altitude rule proposed in this paper designates cruising altitudes as a linear function of heading or course. This rule spreads the traffic vertically and provides a default vertical separation that is proportional to path crossing angle. Monte Carlo simulation of enroute traffic with no air traffic control shows that the linear altitude rule greatly reduces both the collision rate and the mean relative speed of collisions. It also improves the efficiency of conflict resolution. It reduces the amount of altitude change required to resolve conflicts vertically, particularly for large-angle, high-speed conflicts, because those conflicts will have almost the required vertical separation by default. The inherent safety superiority of the linear altitude rule may also allow the horizontal separation standard to be reduced, which would increase airspace capacity and reduce the magnitude of speed or heading changes required to resolve conflicts horizontally.

Introduction

If power is lost to the control rods in a nuclear power plant, they fall into a position in which they stop the main fission reaction. Similarly, railroad crossing gates fall into the safe, down position if power is lost, and elevators are designed so that the tension in the supporting cable keeps spring-loaded emergency brakes *inactive*. These and many other fail-safe, fault-tolerant designs share a common design principle: safety is incorporated into the basic design and depends as little as possible on complicated monitoring and backup systems. Unfortunately, current enroute air traffic rules are not as fail-safe or as fault-tolerant as they could be. Without a foundation of fail-safe and fault-tolerant air traffic rules, safety can be maintained only at the expense of efficiency and capacity. Even if the current level of safety is adequate, therefore, safer air traffic rules

can facilitate improved efficiency and capacity, as will be shown.

Federal Aviation Regulations [1] designate level cruising altitudes in discrete steps of 1000 ft below FL290, or steps of 2000 ft above FL290. After RVSM (Reduced Vertical Separation Minimum [2, 3]) is in effect, the discrete steps will be 1000 ft for all altitudes up to FL410. The designated flight levels alternate between easterly and westerly traffic. Easterly traffic (traffic with a positive eastward component of velocity) is traffic heading from 0 deg through 179 deg; westerly traffic is traffic heading from 180 deg through 359 deg. Although this "hemispheric" or "semi-circular" discrete altitude rule automatically separates easterly traffic from westerly traffic, it does nothing to separate traffic in each of the two categories from itself. On the contrary, it concentrates all traffic in each category into a few discrete altitudes. This discrete altitude rule therefore makes the enroute air traffic system less fail-safe and less fault-tolerant than it could be, which will be corroborated by simulation later in this paper.

In other words, the fundamental problem with the current discrete altitude rule is that it concentrates all enroute traffic into a few horizontal planes, leaving most of the airspace unused (except as a buffer). In the event of a failure of the air traffic control system, the probability of a collision is therefore much higher than it would be if the traffic were more spread out vertically. Furthermore, the discrete rule is less tolerant of errors by air traffic controllers and pilots than a continuous altitude rule could be. The discrete altitude rule was instituted when air traffic was much less dense than it is today, and it was acceptable then. However, with traffic as dense as it is today, and considering the anticipated future growth, such wasteful use of airspace must eventually be reconsidered. Although the current enroute collision rate is extremely low, it is kept low by an inefficient system of static jet routes and a large horizontal separation standard, both of which could possibly be relaxed under a continuous altitude rule.

Collins [4] and Patlovany [5] each identified the problem with the current discrete altitude rule and proposed that cruising altitudes should be a linear function of heading. Their primary focus was on general aviation, but the basic principles are the same for all classes of aircraft and all altitudes. The linear altitude rule is more fail-safe and more fault-tolerant than the discrete rule because it spreads the traffic vertically and provides a default vertical

separation that is proportional to path crossing angle. Patlovany used a Monte Carlo simulation with no air traffic control to compare the collision rates and relative speeds for the discrete rule, random altitude placement, and the proposed linear rule. He found that simply letting aircraft fly at random altitudes reduces the collision rate substantially compared to the discrete rule. He also found that the linear rule further reduces the collision rate and also reduces the mean relative speed of collisions.

Patlovany's collision results are corroborated in this paper, and the broader implications of the linear altitude rule are also considered. Patlovany did not consider the major effects of the linear rule on air traffic control, but those effects are analyzed in this paper. For a given horizontal separation standard, the linear rule substantially increases the number of conflicts that need to be resolved, which tends to increase the workload of controllers and pilots. Fortunately, however, the inherent safety superiority of the linear altitude rule may allow the horizontal separation standard to be reduced, as will be discussed later. The resulting conflict rate would then be more comparable to what it is under the discrete altitude rule. More importantly, most of the conflicts will be far less dangerous than conflicts under the discrete rule, reducing the overall probability of collision. Furthermore, reducing the horizontal separation standard would increase airspace capacity and improve the efficiency of horizontal conflict resolution by reducing the magnitude of the speed or heading changes required.

The linear altitude rule also increases the opportunities for vertical conflict resolution by reducing the amount of altitude change required, particularly for large-angle, high-speed conflicts, because those conflicts will have almost the required vertical separation by default. Unlike the discrete rule, the linear rule allows most level conflicts to be resolved with less (often much less) than 500 ft of altitude change, making vertical resolution practical for level flight. Vertical resolution is desirable for several reasons, which will be discussed in more detail later. First, altitude maneuvers are much more efficient than speed or heading maneuvers. For level flight, vertical resolution is also much simpler than horizontal resolution, and it is also much less sensitive to trajectory prediction error. The efficiency and simplicity of vertical resolution could facilitate autonomous conflict resolution onboard aircraft, which has been proposed as part of Free Flight [6] and DAG-TM (Distributed Air/Ground Traffic Management [7]).

Because the linear altitude rule applies only to level flight, its benefits apply only to encounters between pairs of aircraft that are both flying level. Such level encounters are a majority of the encounters in some regions of airspace, but they are only a small minority in the terminal areas surrounding large airports. The linear altitude rule will therefore provide little or no direct benefit in terminal areas. Note, however, that non-level conflicts are usually simpler to resolve than level conflicts anyway. Most cruise/climb conflicts, for example, can be resolved

by simply leveling off the climbing aircraft below the altitude of the cruising aircraft, then letting it resume its climb after the cruising aircraft has passed. Cruise/descent conflicts are more complicated because the descending aircraft usually has an arrival time constraint, but the conflict can usually be resolved by simply starting the descent early.

The proposed linear altitude rule both expedites and benefits from the regulatory transition toward more direct or wind-optimal routing. The current system of static jet routes imposes structure and helps to maintain an orderly flow of traffic, but it also requires pilots to change course often and fly inefficient routes through a series of non-collinear waypoints. Pilots might be annoyed if they have to adjust altitude each time they adjust course, so the linear rule might not be practical under the current regime (although the altitude adjustments would be small and relatively easy to automate in an appropriately equipped aircraft). Fortunately, however, decision support systems are being developed to help controllers predict and resolve conflicts, thereby reducing the need for static jet routes. Eventually, aircraft will fly much longer straight segments than they now fly, and the linear altitude rule will be more practical. Conversely, the linear rule may further reduce the need for static jet routes by increasing tolerance to controller error.

Under the linear altitude rule, better altitude accuracy means decreased probability of collision, as will be shown later. Under the discrete rule, however, better altitude accuracy means *higher* probability of collision between aircraft at the same nominal altitude, because if they lose horizontal separation they are less likely to "accidentally" avoid each other vertically. TVE (Total Vertical Error) is defined by ICAO (International Civil Aviation Organization) as the difference between the actual and assigned pressure altitudes, accounting for both altitude measurement error and flight technical error. RMS (root mean square) TVE for RVSM-qualified aircraft is now approximately 50 ft [8] (because the mean is very small, RMS is essentially a synonym for standard deviation). Furthermore, when ADS-B (Automatic Dependent Surveillance–Broadcast [9]) and GPS/WAAS (Global Positioning System [10] with Wide Area Augmentation System [11, 12]) become operational, RMS altitude measurement error could be less than 10 ft. Hence, if GPS is ever used for altitude surveillance and control, RMS TVE could possibly be reduced to less than 25 ft. However, GPS determines geometric altitude rather than pressure altitude, so its use would have to be consistent among all aircraft. A combination of GPS/WAAS and baro-altimetry could eventually become very attractive.

A major regulatory change will occur when RVSM is implemented. RVSM will reduce the vertical separation standard (and the discrete steps between designated cruising altitudes) between FL290 and FL410 from 2000 ft to 1000 ft, so it will be 1000 ft all the way up to FL410. RVSM will theoretically double the enroute traffic capacity in the critical altitude range between FL290 and FL410, and it

will allow aircraft to fly closer to their optimal altitude for a given weight. It is also currently driving technical improvements in baro-altimeter accuracy and reliability. RVSM is already in effect over large oceanic regions, and it is planned to be operational in Europe in 2002. The United States is also likely to adopt RVSM by 2005. Because the linear altitude rule is unlikely to be implemented before that time, the remainder of this paper is predicated on the assumption that RVSM is in effect.

It is worth noting that the linear altitude rule could be a function of either heading or course. Course is the direction of flight relative to the ground, whereas heading is relative to the air mass, which is different in a crosswind. Aircraft at higher altitudes tend to have RNAV (area navigation) capability and can fly an arbitrary constant course, so the linear rule should probably be based on course at those altitudes. However, many general aviation aircraft at lower altitudes can measure heading but not course, so the linear rule should perhaps be based on heading at those altitudes. As the percentage of aircraft with RNAV capability increases, the linear rule could perhaps eventually be based on course at all (or nearly all) altitudes. The principles and results presented in this paper are the same regardless of whether course or heading is used, assuming consistency over large segments of altitude. In fact, the results change only marginally even if some aircraft use course and others use heading in the same area. The two terms will be used interchangebly in this paper.

This paper proposes a fundamental change in current air traffic rules and procedures, but the detailed operational implications are beyond the scope of the paper. The intention is only to present a theoretical analysis of the major effects and implications of the proposed change and to motivate further research and interest. The rest of the paper is organized as follows. The next section defines the proposed linear altitude rule and explains its advantages over the current discrete rule. The section after that explains how the linear rule can simplify and improve the efficiency of conflict resolution. Then a Monte Carlo simulation is presented to determine the effect of the linear rule on collision and conflict rates in the absence of air traffic control. Finally, a brief conclusion is presented.

Altitude Rules

The discrete and linear altitude rules for level flight are plotted in Figure 1, which shows designated cruising altitudes as a function of heading (or course) for a representative altitude range. For the discrete rule, the discontinuities at 0 deg (north) and 180 deg (south) are apparent. For the linear rule, the designated altitudes are a continuous, linear function of heading. Figure 2 is a simplified plot of the vertical distribution of traffic for the two rules. If altitude error is ignored for simplicity, the vertical distribution of traffic for the discrete rule is a series of delta functions or "spikes" at the designated altitudes,
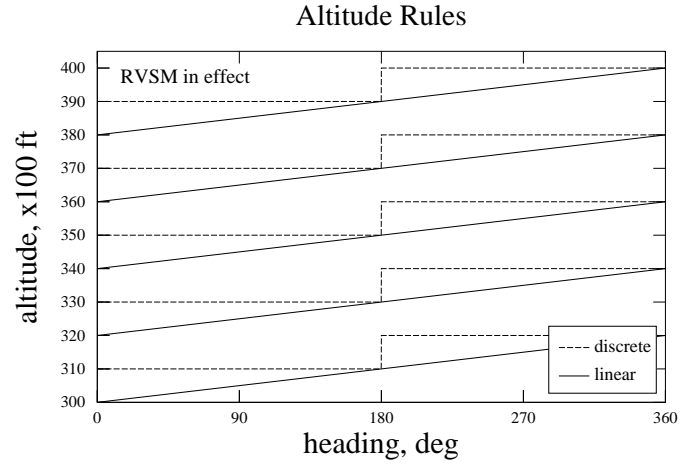

Figure 1: Designated altitude as a function of heading (or course) for level flight in a representative altitude range, for the discrete and linear altitude rules
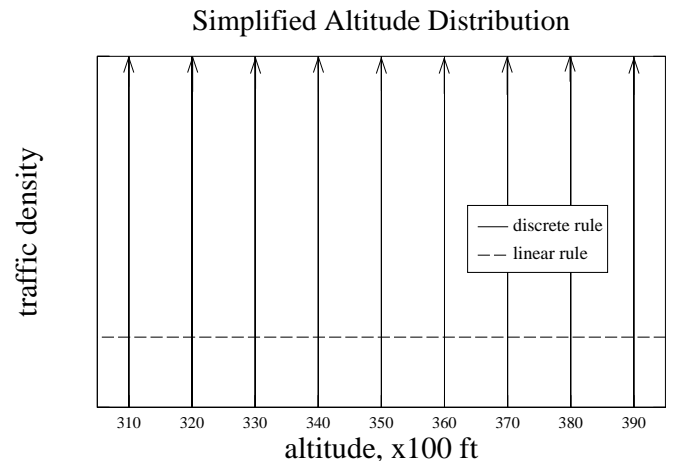

Figure 2: Simplified altitude distribution for enroute traffic in a representative altitude range, for the discrete and linear altitude rules

as shown. For Gaussian altitude error, the spikes become narrow Gaussian "bell curves." For a uniform heading distribution, the corresponding distribution for the linear rule is constant, as shown. For a non-uniform heading distribution, the altitude distribution is not constant, but it is still continuous.

For a given heading of ψ deg, the designated cruising altitudes A in units of 100 ft can be expressed for the discrete and linear altitude rules as

$$\text{discrete:} \quad A = 10\,\text{int}(\psi/180 + 2i + 1) \quad (1)$$

$$\text{linear:} \quad A = 20\,(\psi/360 + i) \quad (2)$$

where i is an altitude cycle integer, and the "int" function is integer truncation. The heading cycles for every altitude cycle of 2000 ft, and the heading/altitude slope for the linear rule is 18 deg per 100 ft. The formula for the linear rule could be programmed into Flight Management

Heading	+	+	Heading
351 - 008	0	10	171 - 188
009 - 026	1	11	189 - 206
027 - 044	2	12	207 - 224
045 - 062	3	13	225 - 242
063 - 080	4	14	243 - 260
081 - 098	5	15	261 - 278
099 - 116	6	16	279 - 296
117 - 134	7	17	297 - 314
135 - 152	8	18	315 - 332
153 - 170	9	19	333 - 350

Table 1: Altitude offset table: nearest integer multiple of 100 ft, modulo twenty, for the linear altitude rule
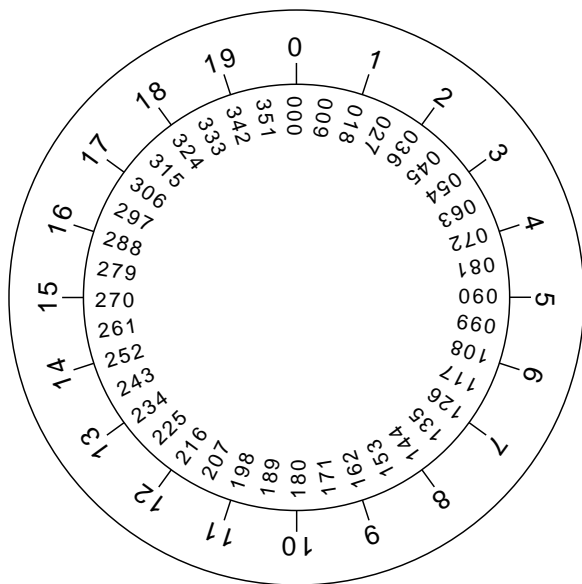

Figure 3: Altitude offset guide and overlay for a conventional heading indicator: outer ring shows designated cruising altitude in units of 100 ft, modulo twenty, for the linear altitude rule

Systems or other avionics equipment.

For operational convenience, the linear altitude rule can be rounded to the nearest 100 ft without significantly reducing its benefits, as will be corroborated later. Such rounding allows the designated altitudes to be specified as in Table 1, which lists the heading range for each integer multiple of 100 ft of altitude, modulo twenty. Alternatively, the guide shown in Figure 3 can be used either by itself or (with the center circle cut out) as an overlay on a conventional heading indicator. Although not shown, another circular figure is also available for use as an overlay on a conventional altimeter.

The linear altitude rule provides a default vertical separation that is proportional to path crossing angle, as shown in Figure 4. The default vertical separation is 1000 ft per 180 deg, or 100 ft per 18 deg, of path crossing angle. For the discrete altitude rule, on the other hand, the vertical separation is not even a function of path crossing angle; it is either zero or 1000 ft, depending only on whether the easterly components of velocity for each aircraft have the same sign or not. The linear altitude rule therefore has the advantage of an inherently fail-safe default vertical separation for all but small-angle, low-speed encounters. This default vertical separation is the fundamental key to the efficacy of the linear altitude rule in avoiding collisions.

The fundamental problem with the discrete altitude rule is that it concentrates all enroute traffic into a few horizontal planes. The traffic density at those designated altitudes is therefore much higher than necessary, and large path crossing angles occur at the same altitude. Figure 5 shows maximum path crossing angles as a function of heading. For the discrete rule, the maximum crossing angle at the same altitude is 90 deg for easterly and westerly headings, and for northerly and southerly headings it can approach 180 deg, a head-on encounter with the highest possible relative speed. The linear rule, on the other hand, spreads the traffic vertically and precludes large path crossing angles at the same altitude. If aircraft height and altitude error are ignored, the maximum crossing angle at the same altitude is theoretically zero, and the maximum crossing angle for all encounters within 250 ft vertically is only 45 deg, as shown in Figure 5.

The speed of one aircraft relative to another is obviously a strong function of the path crossing angle. For reference, Figure 6 shows relative speed ratio as a function of path crossing angle, with speed ratio as a parameter. The speed ratio is the ratio of the speed of a second aircraft to the speed of the subject aircraft, and the relative speed ratio is the ratio of the relative speed to the speed of the subject aircraft. The curves in Figure 6 are based on simple vector geometry and the law of cosines, and they do not depend on the altitude rule. These curves illustrate that large-angle encounters develop more rapidly than small-angle encounters, giving the controller and the pilots less time to respond. Figure 7 shows the dependence of the default vertical separation on relative speed ratio, for the linear altitude rule. Larger relative speeds correspond to larger vertical separations, as desired. The linear rule therefore gives controllers and pilots more time than the discrete rule to respond to critical conflicts.

The effects of aircraft size and altitude error have not yet been considered. Figures 4 through 7 are based on point-mass aircraft models capable of holding a given altitude exactly. As a point of reference, the height of large commercial aircraft (from the bottom of the fuselage to the top of the tail) is approximately 50 ft, and RVSM-qualified aircraft can hold altitude with an RMS TVE (Total Vertical Error) of approximately 50 ft. Of interest here is the probability of collision in level flight, given that horizontal separation has been lost. In that case, the probability of collision can be determined analytically as a simple
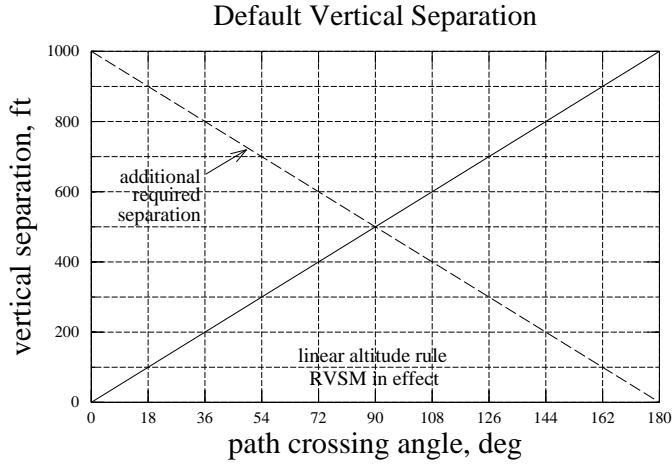
4

Default Vertical Separation

Figure 4: Default vertical separation as a function of path crossing angle for the linear altitude rule

Relative Speed Ratios
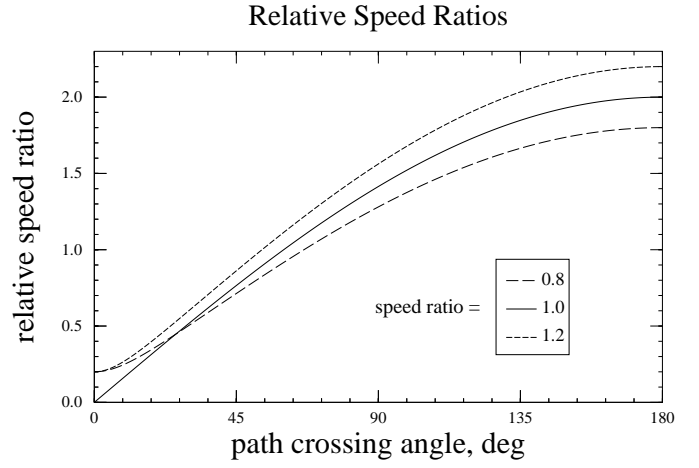
Figure 6: Relative speed ratio as a function of path crossing angle, with speed ratio as a parameter
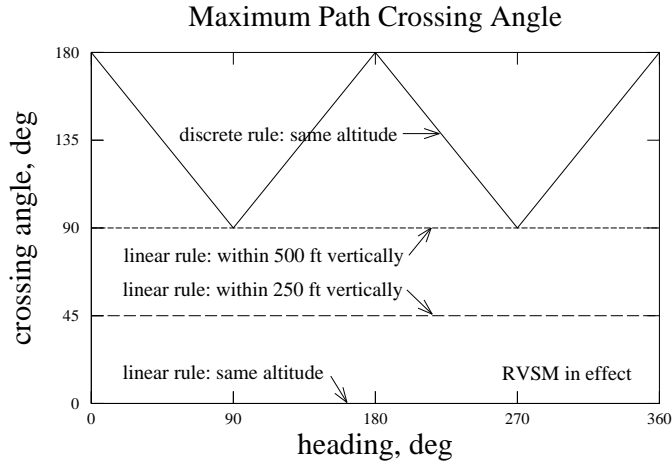

Maximum Path Crossing Angle

Figure 5: Maximum path crossing angles as a function of heading angle for the discrete and linear altitude rules
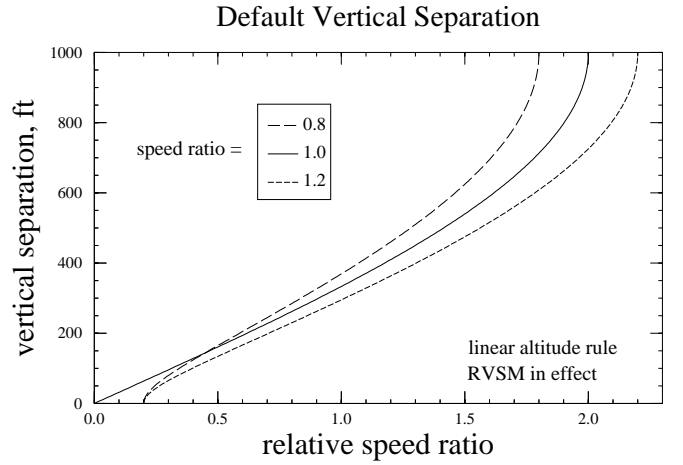

Default Vertical Separation

Figure 7: Default vertical separation as a function of relative speed ratio for the linear altitude rule, with speed ratio as a parameter

one-dimensional Gaussian cumulative difference. A vertical collision threshold of 100 ft will be used as a nominal, conservative value for the height of an aircraft. Collision probability is thus defined here as the probability that two aircraft centers of mass will come within 100 ft vertically, given zero horizontal separation.

Figure 8 shows the collision probability as a function of path crossing angle for the linear altitude rule, and Figure 9 shows it as a function of relative speed ratio, with the RMS TVE as a parameter in each case. These figures show that the collision probability for RVSM-qualified aircraft (RMS TVE = 50 ft) is very small for path crossing angles larger than 60 deg and for relative speed ratios greater than one. If a realistic vertical collision threshold of 50 ft were used instead of the more conservative 100 ft, the collision probabilities would drop off even faster. To fully appreciate the significance of Figures 8 and 9, note that the corresponding curves for the discrete altitude rule would be large (the values at zero path crossing angle and

zero relative speed ratio) and constant all the way out to a path crossing angle of 180 deg and a relative speed ratio of two. For example, for RMS TVE = 50 ft, the collision probability is approximately 0.84 for all path crossing angles and all relative speed ratios.

Conflict Resolution

Under the discrete altitude rule, the vertical separation standard of 1000 ft is exactly equal to the spacing between designated cruising altitudes. The advantage of this arrangement is that easterly traffic (0 to 179 deg) is automatically separated from westerly traffic (180 to 359 deg) by a sufficient margin to avoid conflicts. Under the linear altitude rule, on the other hand, aircraft in the same altitude cycle do not automatically have sufficient vertical separation to avoid a conflict unless their path crossing
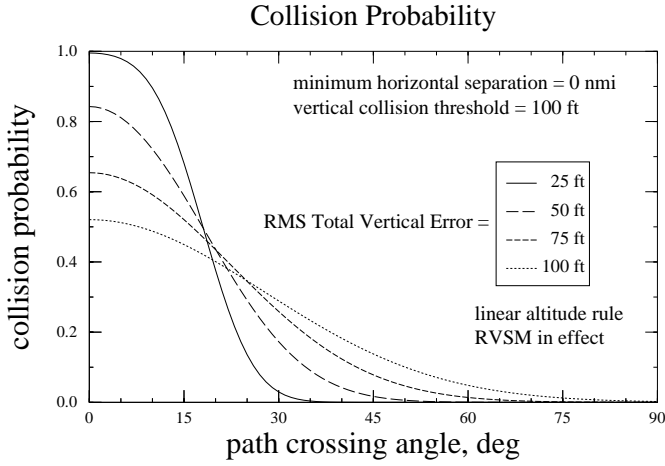
Collision Probability

Figure 8: Collision probability as a function of path crossing angle, given zero horizontal separation, with RMS TVE as a parameter, for the linear altitude rule
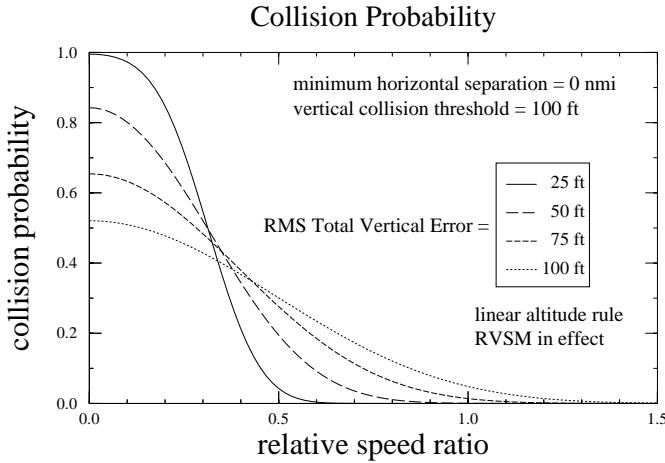
Collision Probability

Figure 9: Collision probability as a function of relative speed ratio, given zero horizontal separation, with RMS TVE as a parameter, for the linear altitude rule

angle is exactly 180 deg, so the conflict rate tends to be higher. Fortunately, however, the inherent safety superiority of the linear rule may allow the horizontal separation standard to be reduced, which would bring the conflict rate down closer to the level for the discrete rule, with the added benefits of increasing airspace capacity and improving the efficiency of horizontal resolution. Also, the default vertical separation of the linear rule increases the opportunities for vertical conflict resolution, which is simpler and more efficient than horizontal resolution. These effects are now discussed in more detail.

Horizontal Conflict Resolution

The primary purpose of separation standards is, of course, to render the probability of collision negligible. No physical harm is done if two aircraft come within, say, 4 nmi of

each other, but the enroute horizontal separation standard is 5 nmi as an added security buffer. However, the linear altitude rule provides each aircraft pair with a default vertical separation proportional to their path crossing angle, hence its inherent safety superiority may allow the separation security buffer to be reduced for better efficiency. Exactly how much it can be reduced depends on several factors. The horizontal separation standard should be a function of three basic factors: (1) aircraft surveillance accuracy and reliability, (2) the availability and accuracy of an automated conflict probe (a decision support system for conflict prediction and resolution), and (3) the minimum time that a controller or pilot will have to avoid a collision, starting from the threshold of conflict.

Current FAA (Federal Aviation Administration) rules [13] require a minimum horizontal separation of 5 nmi for aircraft that are more than 40 nmi from the tracking radar antenna, but they require *only 3 nmi* for aircraft within 40 nmi of the antenna. However, almost all enroute traffic is currently tracked by a "mosaic" multi-radar system in which only one of several possible radars is used to track each aircraft, but the controller is unable to determine which. Thus, according to Nolan [14], controllers using these systems must always assume that every aircraft is 40 nmi or farther from the radar antenna and therefore must separate every aircraft pair by a minimum of 5 nmi. But these concerns do not apply to ADS-B with GPS/WAAS, so the horizontal separation standard could possibly be reduced to 3 nmi when they become operational. Another possible solution might be to start using multi-sensor tracking algorithms to combine the tracking data from multiple radars rather than simply selecting a single radar.

When ADS-B and GPS/WAAS become operational, furthermore, position and velocity estimates will be one or two orders of magnitude more accurate and far more reliable than they are now with radar tracking, so the horizontal separation standard could perhaps be reduced even further to 2 nmi. Additionally, an automated conflict probe will help controllers detect and predict potential conflicts well in advance, further reducing the necessary separation security buffer. Other developing technologies such as CDTI (Cockpit Display of Traffic Information) may eventually help pilots to be safer and more comfortable in the presence of nearby traffic, particularly when the relative speed is low. The last line of defense against a collision is TCAS (Traffic alert and Collision Avoidance System), and it may have already reduced the necessary separation security buffer when it became operational in 1990. The possibility of wake vortex interactions may limit the reduction of the horizontal separation standard, incidentally, but that is beyond the scope of this paper.

The factors discussed above apply equally for both the discrete and linear altitude rules. However, the third factor, the minimum time that a controller or pilot will have to avoid a collision, starting from the threshold of conflict, greatly favors the linear rule. Figure 9 showed the collision

probability as a function of relative speed ratio, where collision probability was defined as the probability that two aircraft (centers of mass) will come within 100 ft vertically, given zero horizontal separation. Figure 9 showed that, for the linear altitude rule, the collision probability for RVSM-qualified aircraft is very small for relative speed ratios greater than one. As explained earlier, if the corresponding curves were plotted for the discrete rule, the collision probabilities would be large, constant values all the way out to a relative speed ratio of two. Thus, *under the discrete rule the collision probability is high for all speed ratios up to two, but under the linear rule it is very small for all speed ratios greater than one.* Recall also that Figure 9 is based on a vertical collision threshold of 100 ft, but if a more realistic threshold of 50 ft is used, the advantage of the linear rule over the discrete rule is even greater.

A case can be made, therefore, that the horizontal separation standard under the linear rule can be reduced to approximately *half* of what it would be under the discrete rule, because in a worst-case (highest relative speed) potential collision, controllers and pilots will still have more time to avoid a collision, starting at the conflict threshold. Although further study is obviously required before any specific reduction can be considered, a reduction to 2 nmi seems reasonable (assuming that ADS-B with GPS/WAAS and an automated conflict probe are operational). Eventually, 1 nmi could perhaps even be considered. Note, however, that such a reduced separation standard would apply only to encounters in which both aircraft are in level flight and conforming to the linear altitude rule. Identifying those conforming aircraft might be error-prone for controllers without a conflict probe, but a conflict probe should be able to clearly distinguish them and apply the correct separation standard.

A reduction of the horizontal separation standard obviously increases airspace capacity by allowing more aircraft to coexist in a given region of airspace. It also improves horizontal resolution efficiency by reducing the magnitude of speed or heading changes necessary. Note that the cost of a horizontal resolution maneuver is not just proportional to the required separation. Krozel and Peters [15, 16] modeled the direct operating cost of horizontal resolution maneuvers in terms of (1) the additional fuel required due to the increased drag and flight path distance, and (2) the additional (non-fuel) operating costs due to the additional time required to execute the maneuver. They found that reducing the separation standard from 5 to 3 nmi reduces the overall cost of heading and speed maneuvers by a factor of three, and reducing the standard from 5 to 1 nmi reduces the cost by a factor of ten.

To understand why the cost of horizontal resolution is not just proportional to the required separation, consider the limiting case of a small heading change initiated far in advance of a potential conflict. Although the effects of trajectory prediction uncertainty are important [17], neglect them here for simplicity. Suppose the minimum predicted separation is r_0, and a heading change is initiated at a range R to increase the minimum separation to the required value of r_1, where r_1 is much less than R. Simple geometry with small angle approximations shows that the additional path distance is approximately $(r_1 - r_0)^2/R$, which is proportional to the *square* of the additional required separation, $r_1 - r_0$. For the case of an exact collision, $r_0 = 0$ and the additional path distance is proportional to the square of the required separation, r_1.

Vertical Conflict Resolution

Under the discrete altitude rule (with RVSM in effect), vertical conflict resolution for level flight always requires a minimum altitude change of 1000 ft (2000 ft above FL290 prior to RVSM). However, this minimum required altitude change puts the aircraft at a "wrong way" altitude facing oncoming traffic. Altitude changes of between 1000 and 2000 ft still create conflicts with the oncoming traffic, so an altitude change of 2000 ft is usually required to resolve conflicts vertically, unless traffic is light. Vertical resolution is therefore currently rarely used for level flight, except when traffic is light, such as late at night.

The linear altitude rule greatly increases the opportunities for vertical conflict resolution by reducing the amount of altitude change required, particularly for large-angle, high-speed conflicts, because those conflicts will have almost the required vertical separation by default. Under the linear rule, most conflicts, including all conflicts with path crossing angles greater than 90 deg, can be resolved with less (often much less) than 500 ft of altitude change. Note that aircraft will be allowed to temporarily deviate from the linear altitude rule during a vertical resolution maneuver. The linear rule therefore makes vertical resolution practical for level flight, which is important for several reasons.

According to Krozel and Peters [15, 16], altitude maneuvers are much more efficient than speed and heading maneuvers because (1) the separation requirement is much less vertically than horizontally, and (2) altitude maneuvers tend to be nearly conservative, converting kinetic to potential energy and vice versa, but consuming very little additional energy. For level flight, vertical resolution is also much simpler than horizontal resolution, both mathematically and operationally. Whereas horizontal resolution involves nontrivial vector computation or intuitive geometric visualization by controllers, vertical resolution for level flight involves only trivial scalar subtraction. Vertical resolution is also much less sensitive to trajectory prediction error in level flight because altitude is much more accurately predictable than horizontal position. Finally, the efficiency of altitude maneuvers is largely independent of the range at which the maneuver is initiated. Heading maneuvers can be very efficient too, but usually only if they are initiated at least several minutes before the conflict. Altitude maneuvers, on the other hand, are usually extremely efficient even if they are initiated only one or two

minutes before the conflict. The fact that altitude maneuvers can be of shorter duration, furthermore, means that they are less likely to create additional conflicts with other aircraft.

Given the efficiency and simplicity of vertical conflict resolution, the increased opportunity to use it under the linear altitude rule could greatly facilitate autonomous conflict resolution in the air, which is a key component of Free Flight and DAG-TM. Autonomous airborne conflict resolution reduces air traffic controller workload on the ground, of course. The DAG-TM [7] concept of "trajectory negotiation," in which the pilots negotiate a resolution between themselves, becomes almost trivial for vertical resolution. Altitude maneuvers are so efficient and so simple to calculate and execute that such negotiation is hardly necessary.

Monte Carlo Simulation

Simulation Methods

A Monte Carlo simulation was programmed in C++ to simulate the enroute air traffic in a region of the size typically controlled by an enroute control center. The Denver Center, for example, controls an area of approximately 280,000 square nmi and currently has a maximum of approximately 160 aircraft in level flight at FL290 and above during peak periods. The purpose of the simulation was to determine collision rates and conflict rates in the absence of air traffic control, hence no air traffic control was simulated. The initial positions and velocities of the simulated aircraft are based on a simple "random gas" model, with horizontal positions randomly distributed throughout a square region of airspace. All aircraft in the simulation *always* fly level at constant velocity, hence no integration with respect to time is necessary, and minimum separations can be determined analytically. The vertical positions were randomly distributed from FL300 to FL400, consistent with the altitude rules to be compared. The vertical distribution was determined by selecting the altitude cycle integer i in Equation 1 or 2 as a uniformly distributed random integer from 15 to 19 for each aircraft.

The simulation prompts the user for various parameters, such as the number of Monte Carlo runs, each of which starts with a re-shuffling of the aircraft positions and velocities. The results to be presented are based on 250,000 runs for each parameter set. The user is also prompted for the size of the region and the number of aircraft in level flight. For this paper, the region is a 500 by 500 nmi square with 200 aircraft, unless otherwise noted. TVE was modeled as a Gaussian zero-mean random variable with RMS magnitudes of 0, 25, 50, 75, and 100 ft, and is constant for each aircraft throughout each run. As explained earlier, these levels of altitude accuracy correspond approximately to GPS/WAAS (25 ft), RVSM-qualified baro-altimeter (50 ft), and low-cost baro-altimeter technologies. The user can select either the discrete altitude rule, the linear rule,

or uniformly random altitude placement. Overall, fifteen combinations of five RMS TVE magnitudes and three altitude rules were needed for the main results of this paper.

The simulated time horizon is significant because the traffic density begins to decrease after a certain period of time as traffic leaves the region under study. The traffic density was kept nominally constant for a given period of time by actually placing some aircraft outside the 500 by 500 nmi region, as follows. A larger, outer square region was determined by multiplying the mean aircraft speed by the time horizon and adding twice that distance to the (500 nmi) width of the inner region. Traffic was then randomly placed in the larger region such that the number of aircraft in the inner region was approximately 200. Collisions and conflicts were then counted only if they occurred in the inner region. A large time horizon is desirable, but the width of the outer region grows linearly with the time horizon, and the time required to run the simulation is roughly proportional to the fourth power of that width (the area and the number of aircraft are proportional to the square of width, and the number of aircraft pairs is proportional to the square of the number of aircraft). A simple analysis showed that a time window from -15 to 15 min is close to optimal.

Heading angles over the continental United States tend to have a higher density in the vicinity of 90 deg (east) and 270 deg (west) than 0 deg (north) and 180 deg (south). This non-uniformity is clearly illustrated in Figure 10, which shows the US national heading distribution in increments of 9 deg for FL290 and above, based on ETMS (Enhanced Traffic Management System) data for one full day. To make the simulation more realistic, the heading angle distribution was modeled as a combination of uniform and Gaussian random distributions to approximate the distribution shown in Figure 10. The ground-speed distribution was also modeled as Gaussian to approximate the ETMS data shown in Figure 11.

An important result of the Monte Carlo simulation is the collision rate with no air traffic control. The absolute collision rate is not as important as the ratio between the discrete and linear altitude rules. A nominal collision threshold is defined here as 300 ft of horizontal separation and 100 ft of vertical separation between centers of mass. For reference, the length and wingspan of a Boeing 777, one of the largest aircraft models in service, are approximately 200 ft, and the height (from the bottom of the fuselage to the top of the tail) is approximately 50 ft. The nominal collision threshold is therefore conservative and accounts for the fact that severe aerodynamic interactions might occur even if the aircraft do not actually contact each other. Another, less conservative collision threshold will also be used. A minimal collision threshold is defined here as 200 ft of horizontal separation and 50 ft of vertical separation between centers of mass, the minimal separation required to avoid actual contact between two large aircraft. These collision thresholds are shown in Table 2.
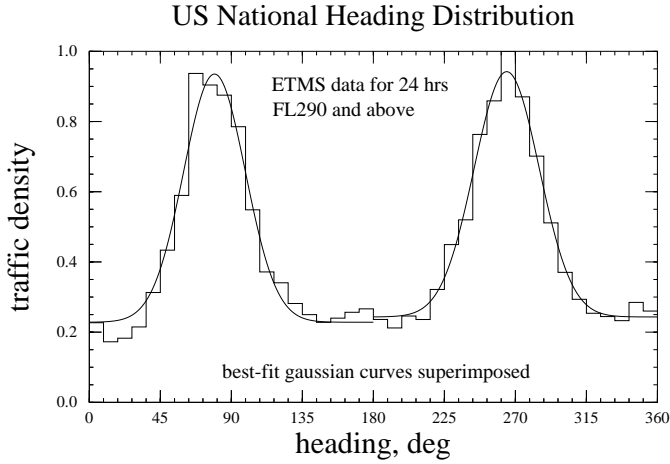
US National Heading Distribution

Figure 10: US national heading distribution in increments of 9 deg for FL290 and above

US National Groundspeed Distrubution

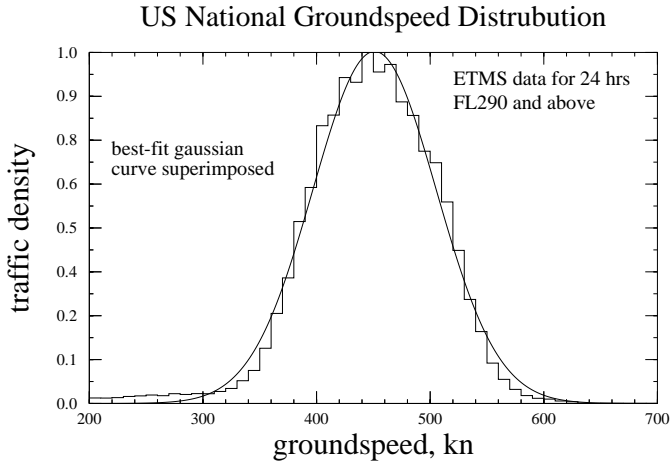Figure 11: US national groundspeed distribution in increments of 10 kn for FL290 and above

collision	separation	
threshold	horizontal	vertical
minimal	200 ft	50 ft
nominal	300 ft	100 ft

Table 2: Collision thresholds

Simulation Results

Collision Rates and Relative Speeds

As explained above, 250,000 Monte Carlo runs were executed for each altitude rule, with identical parameters and conditions for each run. The entire procedure was repeated for RMS TVEs of 0, 25, 50, 75, and 100 ft. The collision counts, based on the minimal collision threshold, are summarized in Table 3. Note that, for the discrete altitude rule, better altitude accuracy actually causes a higher colli-

altitude	RMS total vertical error			
rule	25 ft	50 ft	75 ft	100 ft
discrete	17190	10041	6323	5175
random	3425	3013	3531	3574
linear	508	721	1089	1082

Table 3: Collision counts, based on the minimal collision threshold, for 250,000 Monte Carlo runs with 200 aircraft in level flight from FL300 to FL400 in a 500 by 500 nmi region and a time window of 30 min

altitude	RMS total vertical error			
rule	25 ft	50 ft	75 ft	100 ft
discrete	1.0	1.0	1.0	1.0
random	5.0	3.3	1.8	1.4
linear	33.8	13.9	5.8	4.8

Table 4: Collision rate reduction factors relative to the discrete altitude rule (from Table 3), based on the minimal collision threshold

altitude	RMS total vertical error			
rule	25 ft	50 ft	75 ft	100 ft
discrete	1.0	1.0	1.0	1.0
random	2.8	2.9	2.2	1.7
linear	10.3	7.9	5.6	3.8

Table 5: Collision rate reduction factors relative to the discrete altitude rule, based on the nominal collision threshold

sion rate. This effect has been called the "navigation paradox," because pilots are "rewarded" with a higher danger of collision for following the rule more diligently. Under the linear rule, on the other hand, better altitude accuracy causes a lower collision rate. Collision rate reduction factors were calculated by dividing the collision counts for the discrete rule by the corresponding counts for the linear and random rules, and they are shown in Table 4. For the minimal collision threshold, the linear altitude rule with RVSM-qualified aircraft (RMS TVE = 50 ft) reduces the collision rate by a factor of approximately fourteen (13.9) compared to the discrete rule. For the nominal collision threshold, Table 5 shows that the corresponding collision rate reduction factor is approximately eight (7.9). The linear rule is therefore an order of magnitude more fail-safe than the discrete rule in this simulation.

Furthermore, the linear altitude rule does more than reduce the collision rate. It also reduces the relative speed of collisions. Table 6 shows the mean relative speed of collisions for the discrete, random, and linear altitude rules, where the nominal collision threshold is used, and Table 7 shows the reduction factors relative to the discrete rule. Although the random altitude rule reduces the collision rates compared to the discrete rule, Tables 6 and 7 show that it substantially *increases* the mean relative speed of

| altitude | RMS total vertical error | | | |
rule	25 ft	50 ft	75 ft	100 ft
discrete	462	458	446	450
random	675	663	686	692
linear	109	127	164	211

Table 6: Mean relative speed of collisions, kn, based on the nominal collision threshold

| altitude | RMS total vertical error | | | |
rule	25 ft	50 ft	75 ft	100 ft
discrete	1.0	1.0	1.0	1.0
random	0.7	0.7	0.6	0.7
linear	4.2	3.6	2.7	2.1

Table 7: Reduction factors for mean relative speeds of collisions, with the discrete altitude rule as the baseline (from Table 6), based on the nominal collision threshold
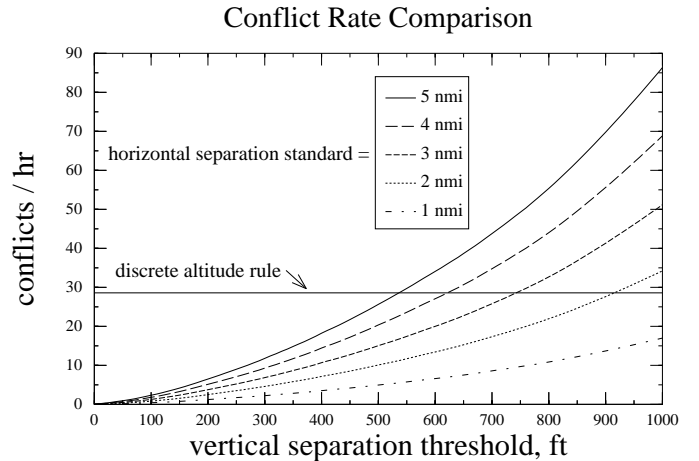

Figure 12: Conflict rate as a function of vertical separation threshold for the linear altitude rule, with horizontal separation standard as a parameter

collisions, giving controllers and pilots less time to respond. This result indicates that *it would be unwise to let airlines and/or pilots choose arbitrary cruising altitudes.* Tables 6 and 7 also show that the linear altitude rule substantially reduces the mean relative speeds of collisions. For RVSM-qualified aircraft (RMS TVE = 50 ft), the reduction factor is 3.6. Pilots and controllers therefore have much more time to notice an impending collision and maneuver to avoid it. In clear air, lower relative speeds also make visual avoidance much more feasible.

The collision reduction factors cited above are for the US national heading distribution shown in Figure 10, but heading distributions vary both regionally and with altitude. In general, heading distributions at lower altitudes tend to be more uniformly distributed than at higher altitudes. Heading distributions above FL290 in Denver Center, for example, are very non-uniform, with far fewer aircraft heading north or south than are heading east or west. As a point of reference, collision rates for the discrete and linear altitude rules are exactly the same for the limiting case in which all traffic is heading either due east or due west. The collision reduction factor would therefore be less in Denver Center than it is nationally. It could be improved by making the altitude rule nonlinear (or piecewise linear), with a larger altitude/heading slope in the vicinity of east/west headings and a smaller slope in the vicinity of north/south headings. Such a nonlinear altitude rule may eventually warrant serious consideration for higher altitudes, but it would vary regionally and will not be considered in this paper.

Although the simulation only covered altitudes from FL300 to FL400, the basic principles are the same for all altitudes at which altitude rules apply. The only fundamental difference at lower altitudes is that altitude accuracy tends be worse and aircraft tend to be smaller. The former tends to diminish the collision reduction effect of the

linear rule, but the latter tends to enhance it. Although a given baro-altimeter tends to be more accurate at lower altitudes (due to larger pressure gradients), aircraft that cruise at lower altitudes tend to be smaller aircraft with cheaper equipment. The RMS total vertical error is unlikely to be worse than 100 ft at any altitude, however, so the collision reduction factors in Table 4 and 5 apply, in principle, at all altitudes. And because aircraft that cruise at lower altitudes are smaller, the *minimal* collision threshold of Table 2 can be regarded as the *nominal* threshold for lower altitudes.

The simulation also was run with 100 and 300 aircraft rather than 200. The collision reduction factors were not significantly different from those shown in Tables 4 and 5. The efficacy of the linear altitude rule at preventing collisions is therefore reasonably independent of traffic density. Finally, the collision rates for the linear altitude rule were not significantly affected when the designated altitudes were rounded to the nearest 100 ft. Thus, the use of Table 1 will yield nearly the same benefits as avionic equipment capable of computing the linear altitude rule exactly.

Conflict Rates

In addition to the collision rates discussed above, the other important results of the Monte Carlo simulation are the conflict rates, again with no air traffic control. Figure 12 is a plot of the conflict rate as a function of vertical separation threshold, with horizontal separation standard as a parameter. In other words, Figure 12 is a plot of what the conflict rate would be if the vertical separation standard were varied from 0 to 1000 ft, the significance of which will become clear later. In level flight, pressure altitude gets rounded to the nearest 100 ft onboard the aircraft, then on the ground it gets rounded again to the nominal designated altitude if it is within 200 ft of that altitude. The

10

vertical	horizontal separation standard				
threshold	1 nmi	2 nmi	3 nmi	4 nmi	5 nmi
1000 ft	0.59	1.20	1.79	2.41	3.02
900 ft	0.48	0.97	1.44	1.95	2.44
800 ft	0.38	0.77	1.14	1.54	1.94
700 ft	0.30	0.61	0.90	1.22	1.53
600 ft	0.23	0.47	0.70	0.95	1.19
500 ft	0.17	0.35	0.53	0.71	0.90

Table 8: Conflict rate factors for the linear altitude rule, with the discrete altitude rule as the baseline (28.6 conflicts/hr)

Figure 13: Vertical resolution opportunities: percentage of conflicts resolved as a function of vertical maneuver threshold

vertical separation standard is large enough to tolerate error within the rounding threshold. This altitude rounding means that the vertical separation standard is based on nominal rather than observed altitude. Thus, although altitude measurement and control error had a major effect on the collision rates discussed earlier, it has little or no effect on conflict rates. The effect of the altitude rounding was realized in the simulation by simply setting the RMS TVE to zero.

The baseline conflict rate is defined here as the conflict rate for the discrete altitude rule, which is 28.6 conflicts/hr, as shown in Figure 12. Table 8, which is based on the same data used for Figure 12, shows that for the current horizontal separation standard of 5 nmi and the vertical separation standard of 1000 ft, the linear rule yields a conflict rate approximately triple (3.02 times) the baseline rate. Recall, however, that the separation standard is currently 5 nmi only because aircraft are tracked with a "mosaic" multiradar system, but that may no longer be true when ADS-B with GPS/WAAS becomes operational, at which time the separation standard may be reduced to 3 nmi. The conflict rate would then be less than double (1.79 times) the baseline rate. The acceptability of increased conflict rates will have to be determined by further studies, but an automated conflict probe may help in that regard. As discussed earlier, the horizontal separation standard under the linear altitude rule could perhaps be reduced to 2 nmi. Table 8 shows that the conflict rate would then be twenty percent above the baseline rate.

Interestingly, the curves of Figure 12 are almost exactly parabolic, with the conflict rate proportional to the *square* of the vertical threshold. This parabolic form explains the paradoxical fact that the linear rule *increases* the conflict rate while it *decreases* the collision rate. The conflict rate corresponds to the standard vertical separation threshold of 1000 ft, whereas the collision rate corresponds to much lower threshol in the vicinity of 100 to 200 ft. The vertical separation standard obviously cannot be reduced to 200 ft, but the lower conflict rate at smaller vertical thresholds means that, as earlier results showed, *collisions* are much less likely under the linear rule. If two aircraft converge horizontally, 200 ft of vertical separation is certainly too close for comfort, but it is infinitely better
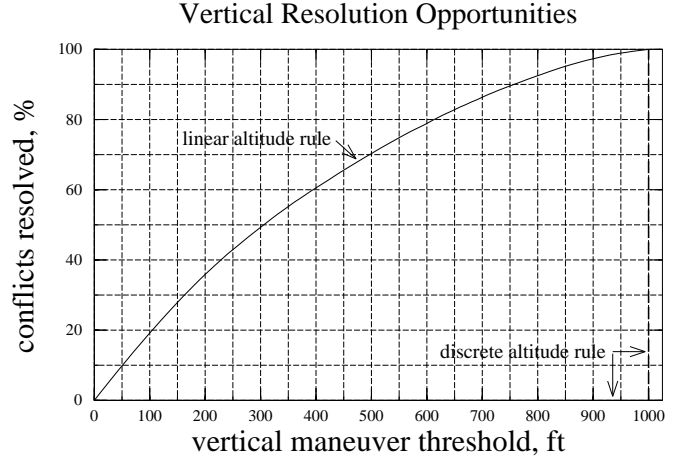
than 0 ft. In other words, the first 200 ft of vertical separation is worth infinitely more than the second or third 200 ft. And while the discrete rule delivers better than the linear rule on the fourth and fifth 200 ft of separation, it cannot compete with the linear rule on the critical first 200 ft.

Figure 13 shows the percentage of conflicts that can be resolved vertically as a function of the maximum size of the vertical maneuver. This plot, which applies for any horizontal separation standard, illustrates the improved vertical resolution efficiency under the linear altitude rule. For the discrete rule, no conflicts can be resolved vertically with an altitude change of less than 1000 ft. For the linear rule, however, Figure 13 shows that more than one third (35.9%) of conflicts can be resolved with an altitude change of 200 ft *or less*. These are the large-angle ($\geq$ 144 deg), high-speed conflicts that have nearly enough vertical separation by default to avoid a conflict. Furthermore, nearly one half (49.3%) of all conflicts can be resolved with an altitude change of 300 ft or less, and well over two thirds (70.3%) can be resolved with an altitude change of 500 ft or less. Horizontal resolution can also be used in all of these cases, and it may occasionally be preferable, but as explained earlier, vertical resolution is usually much simpler and much more efficient.

It is interesting to consider the speed changes that occur as a function of altitude change. Taking into account only a simple balance of kinetic and potential energy, the following results were calculated. If an aircraft is flying level at 500 kn and starts climbing, it loses only 2.3 kn of speed per 100 ft of altitude gain, and that rate is fairly constant for at least the first 500 ft. At 400 kn, the rate of speed loss is 2.9 kn per 100 ft of altitude gain, and at 300 kn the rate is approximately 3.9 kn per 100 ft. As discussed earlier, Krozel and Peters [15] found that even altitude maneuvers of 1000 ft are very efficient, but large altitude deviations expose the aircraft to large-angle con-

11

flicts. A reasonable policy, therefore, might be to consider vertical resolution for all conflicts that require 500 ft or less of altitude change. Those conflicts have path crossing angles larger than 90 deg and constitute over two thirds of all conflicts. The only compelling reason not to resolve such conflicts vertically is if the resolution creates a new conflict with a third aircraft. In that case the other aircraft can execute the altitude maneuver or, as a last resort, horizontal resolution can be used.

Referring back to Table 8, conflict rates are given for vertical separation thresholds other than the standard 1000 ft. Table 8 shows the rate of conflicts that require more than a given threshold magnitude of altitude change. For example, the row for a vertical separation threshold of 800 ft gives the conflict rate after discounting all conflicts that can be resolved with an altitude deviation of 200 ft or less. Table 8 shows that, even if the horizontal separation standard is kept at 5 nmi, the rate of conflicts that require more than 500 ft of altitude change is only 0.90 times the baseline conflict rate for the discrete rule. In other words, if vertical conflict resolution is limited to altitude deviations of 500 ft or less, the number of conflicts that must be resolved horizontally is 10% less than it is for the discrete rule. If the horizontal separation standard is reduced to 2 nmi, Table 8 shows that the rate of conflicts that will require more than 500 ft of altitude change is only 0.35 times the baseline conflict rate. The number of conflicts requiring complicated and inefficient horizontal resolution would therefore be cut by nearly two thirds.

To summarize these results, the linear altitude rule with an appropriate horizontal separation standard may yield a higher conflict rate than the discrete rule, but most of the conflicts will be less dangerous, and many of them will be simpler and less costly to resolve. The actual stress on controllers and pilots could therefore actually be less under the linear rule. Under the discrete rule, every potential conflict is a potential collision, but under the linear rule, the vast majority of potential conflicts will have enough vertical separation by default to avoid a collision, and many of them will be large-angle, high-speed conflicts with nearly enough vertical separation by default to also avoid a conflict. Under the discrete rule, every conflict makes the controller potentially responsible for lives, but under the linear rule, most conflicts simply require the controller to increase the vertical separation from a *fairly* safe magnitude to a *very* safe magnitude. Also, vertical resolution is much simpler and much more efficient than horizontal resolution for level flight, so its increased use under the linear rule could facilitate autonomous conflict resolution in the air.

Conclusion

The linear altitude rule proposed in this paper can make enroute air traffic safer and more efficient by spreading the traffic vertically and providing a default vertical separation that is proportional to path crossing angle. It can both expedite and benefit from the regulatory transition toward more direct routing, particularly in conjunction with an automated conflict probe and ADS-B with GPS/WAAS. Its inherent safety superiority can reduce the need for the current inefficient system of static jet routes. Monte Carlo simulation results for enroute traffic with no air traffic control show that the linear rule greatly reduces both collision rates and the mean relative speed of collisions.

The linear altitude rule can also improve the efficiency of conflict resolution. The inherent safety superiority of the linear rule may allow the horizontal separation standard to be reduced, thereby increasing airspace capacity and improving the efficiency of horizontal resolution by reducing the magnitude of heading or speed changes required. The linear rule also increases the opportunities for vertical conflict resolution, particularly for large-angle, high-speed conflicts, because those conflicts will have nearly the required vertical separation by default, and will require only small altitude changes. Because vertical resolution is much simpler and much more efficient than horizontal resolution, its increased use could facilitate autonomous conflict resolution onboard aircraft.

References

[1] Federal Aviation Regulations, Sec. 91.179, IFR Cruising Altitude or Flight Level. Available at "http://www.faa.gov/avr/AFS/FARS/far-91.txt".

[2] International Civil Aviation Organization: "Manual on the Implementation of a 300 m (1000 ft) Vertical Separation Minimum Between FL290 – FL410 Inclusive," Doc. 9574. Available from "http://www.icao.org".

[3] Eurocontrol: "Master Plan: European Reduced Vertical Separation Minimum Programme," V3.5, April 23, 1999. Available from "http://www.eur-rvsm.com". See also "http://www.faa.gov/ats/ato/rvsm1.htm".

[4] Collins, L.H.: "Automatic Altitude and Heading Separation", *Air Facts*, vol. 31, no. 2, Feb. 1968, pp 31-38.

[5] Patlovany, R.W.: "U.S. Aviation Regulations Increase Probability of Midair Collisions," *Risk Analysis*, vol. 17, no. 2, 1997, pp 237-248.

[6] Final Report of the RTCA Task Force 3, Free Flight Implementation. RTCA, Inc., Washington, DC, Oct. 26, 1995. See also "http://www.faa.gov/freeflight".

[7] National Aeronautics and Space Administration: Concept Definition for Distributed

Air/Ground Traffic Management (DAG-TM), version 1.0, Sept. 1999. Available from "http://www.asc.nasa.gov/aatt/dag.html".

[8] North Atlantic Reduced Separation Standards Implementation Group (RSSIG), Airspace Monitoring Sub-Group (AMSG): "RVSM Monitoring Results – Monitoring Pack," 9th Meeting, London, March 23-27, 1998.

[9] RTCA SC-186: "Minimum Aviation System Performance Standards for Automatic Dependent Surveillance Broadcast (ADS-B)," DO-242, Feb. 19 1998. See "http://www.rtca.org".

[10] Federal Aviation Administration: "GPS Transition Plan." See "http://www.faa.gov/asd".

[11] International Civil Aviation Organization: "U.S. GPS Wide Area Augmentation System (WAAS)," Fourth meeting of the Middle East Air Navigation Planning and Implementation Regional Group, Cairo, Dec. 1-5, 1997. See also "http://gps.faa.gov".

[12] Enge, P.K.; Van Dierendonck, A.J.: "Wide Area Augmentation System," chapter 4 of *Global Positioning System: Theory and Applications, Volume II*, AIAA, 1996, pp 117-142.

[13] Federal Aviation Administration: Air Traffic Control FAA Order 7110.65M, Sec. 5-5-3b.

[14] Nolan, M.S.: *Fundamentals of Air Traffic Control*, Second Edition, Wadsworth Publishing Co., 1994, pp 375-382.

[15] Krozel, J.; Peters, M.: "Conflict Detection and Resolution for Free Flight," *Air Traffic Control Quarterly*, vol. 5, no. 3, 1997, pp 181-212.

[16] Krozel, J.; Peters, M.: "Strategic Conflict Detection and Resolution for Free Flight," Proceedings of the 36th *Conference on Decision and Control*, San Diego, CA, Dec. 1997, pp 1822-1828.

[17] Paielli, R.A.; Erzberger, H: "Conflict Probability Estimation Generalized to Non-Level Flight," *Air Traffic Control Quarterly*, vol. 7, no. 3, Oct. 1999, pp 195-222.

Acronyms

ADS-B	Automatic Dependent Surveillance–Broadcast
CDTI	Cockpit Display of Traffic Information
DAG-TM	Distributed Air/Ground Traffic Management
ETMS	Enhanced Traffic Management System
FAA	Federal Aviation Administration
FAR	Federal Aviation Regulations
FL	Flight Level
GPS	Global Positioning System
ICAO	International Civil Aviation Organization
NASA	National Aeronautics and Space Admin.
RMS	Root Mean Square
RNAV	Area Navigation
RVSM	Reduced Vertical Separation Minimum
TCAS	Traffic Alert and Collision Avoidance System
TVE	Total Vertical Error
WAAS	GPS Wide Area Augmentation System

Russ Paielli received a BS degree in Mechanical Engineering from Oakland University in Michigan in 1982 and has worked at NASA Ames Research Center since then. He received an MS degree in Aeronautics and Astronautics from Stanford University in 1987. At NASA Ames he has done research in flight control theory and precision aircraft navigation and landing. He is currently working on advanced concepts in Air Traffic Management. He is a senior member of AIAA. His personal homepage on the web can be found at "RussP.org".