TRANSPORTATION SCIENCE
Vol. 40, No. 2, May 2006, pp. 147–164
ISSN 0041-1655 | EISSN 1526-5447 | 06 | 4002 | 0147

informs®

DOI 10.1287/trsc.1050.0141

An Airspace-Planning and Collaborative Decision-Making Model: Part II—Cost Model, Data Considerations, and Computations

Hanif D. Sherali
Grado Department of Industrial and Systems Engineering (0118), Virginia Polytechnic Institute and State University, Blacksburg, Virginia 24061, hanifs@vt.edu

Raymond W. Staats
Department of Operational Sciences, Air Force Institute of Technology, Wright Patterson Air Force Base, Ohio 45433, raymond.staats@afit.edu

Antonio A. Trani
Charles Edward Via, Jr. Department of Civil and Environmental Engineering (0105), Virginia Polytechnic Institute and State University, Blacksburg, Virginia 24061, vuela@vt.edu

In Part I of this paper, we presented a large-scale airspace-planning and collaborative decision-making (APCDM) model that is part of a Federal Aviation Administration (FAA)-sponsored effort to enhance the management of the National Airspace System (NAS). Given a set of flights that must be scheduled during some planning horizon, along with alternative surrogate trajectories for each flight, we developed a mixed-integer programming model to select a set of flight plans from among these alternatives, subject to flight safety, air-traffic control workload, and airline equity considerations. The present paper offers insights related to, and a detailed description of, implementing this APCDM model, including the development of a comprehensive cost model, a study for prescribing a set of appropriate parameter values for the overall model, and an investigation on incorporating a suitable set of valid inequalities in the model formulation. Computational results are presented based on several test cases derived from the Enhanced Traffic Management System (ETMS) data provided by the FAA. The results indicate that under plausible probabilistic trajectory error assumptions and with the incorporation of star subgraph convex hull-based valid inequalities, the model offers a viable tool that can be used by the FAA for both tactical and strategic applications.

Key words: air traffic management; national airspace; aircraft conflict resolution; sector workload measures; airline equity; mixed-integer programming; valid inequalities
History: Received: January 2004; revision received: February 2005; accepted: October 2005.

1. Introduction

The airline industry is a highly competitive business that has a significant impact on the overall U.S. economy. It is in the nation's continued interest to ensure safe and dependable use of the National Airspace System (NAS), even as increased demand strains air-traffic management capabilities. Furthermore, it is essential that we take steps to encourage new entrant airlines and develop services for smaller communities to ensure a healthy and competitive industry in the long term. Accordingly, the Federal Aviation Administration (FAA) is sponsoring an overall 10-year, $11.5B, effort to increase NAS capacity by 30%. Currently, air-traffic density is such that a single severe weather system can cause take-off and landing delays that cascade throughout the entire NAS. During 2000, the top 55 airports conducted more than 20.8 million such operations, with 425,000 of them subject to delays (see Crawley 2001).

The model developed in Part I of this paper deals with an effective management of air traffic, particularly under scenarios of disruptions due to the closure of sections of the airspace as prompted by severe weather or by Special Use Airspaces (SUA) designated during space launches. In this sense, it indirectly contributes toward the goal of enhancing airspace capacity. The essential operational framework that embodies the airspace-planning and collaborative decision-making (APCDM) model proposed in Part I of this paper is depicted in Figure 1 (see Sherali, Staats, and Trani 2003). Given a set of flights that must be scheduled during some planning horizon, the overall intent is to select a set of flight plans from among alternatives, called *surrogates*, subject to flight safety, air-traffic control workload, and airline equity constraints. The outer loop illustrated in Figure 1 is a possible augmentation to APCDM that

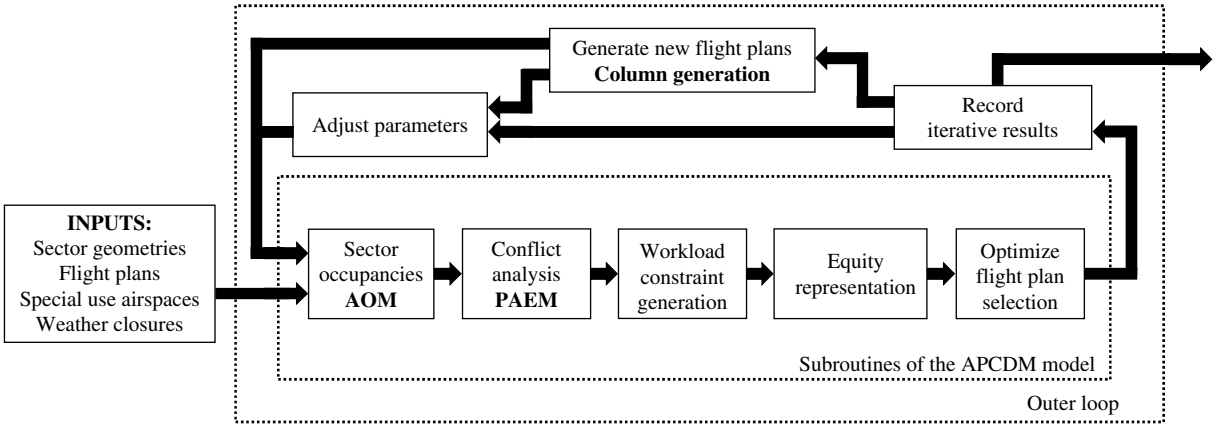

Figure 1 Conceptual Operational Framework for the APCDM Model

could be used to iteratively adjust parameters as necessary, and consider additional surrogate flight plans that might help achieve a greater degree of operational efficiency and equity. This is not presently implemented in the system, and our focus in the present paper is on the basic APCDM model.

The model's principal structure is comprised of several submodels, as depicted in Figure 1. The first of these, namely, the *Airspace Occupancy Model* (*AOM*), traces the trajectory of each flight plan as it traverses through the three-dimensional nonconvex sectors that partition the airspace, recording the duration of occupancy for each sector that it passes through.

The *Probabilistic Aircraft Encounter Model* (*PAEM*) conducts a conflict analysis between pairs of aircraft flight plans to ascertain whether any intruder aircraft can possibly penetrate the protective standard separation shell (assumed to be a rectangular box) enveloping some focal aircraft. This conflict analysis is performed assuming that the trajectories of both the focal and intruder aircraft are subject to either randomized or wind-induced probabilistic errors. Accordingly, a pair of threshold probabilities p_1 and $p_2 < p_1$ are specified, and the PAEM module determines the duration for which a penetration into the standard separation shell (called a Level-1 conflict) occurs with at least a probability of p_1, or a penetration into a shell of half this dimension (called a Level-2 conflict) occurs with at least a probability of p_2. In addition, if the intrusion penetrates a tight enveloping box having a half-dimension of 500 feet in the in-trail and cross-track dimensions, and 100 feet in the altitude dimension, then a *fatal conflict* is declared, and such pairs of flight plans are necessarily prohibited from coexisting. Otherwise, the conflict is declared to be *resolvable*, but subject to sector workload capacity restrictions.

The third module in Figure 1 addresses the formulation of the sector and conflict resolution constraints and penalties. The different resolvable conflict time intervals that are identified by PAEM are first augmented by a suitable prep-buffer duration to account for a reaction time that might be required to begin responding to the resolution of this conflict. By using the sector occupancy information available from the AOM module, this subroutine efficiently determines for each sector all maximally overlapping sets of resolvable conflicts, i.e., resolvable conflicts that overlap at some point in time, while not being a strict subset of some other such set of overlapping conflicts. Each maximally overlapping set of this type is represented by a graph, where the nodes represent the various flight plans involved and the edges represent the identified conflicts between pairs of flight plans. Using alternative mathematical formulations and valid inequalities, suitable constraints are imposed for each sector s to effectively ensure that the selected flight plans do not induce a simultaneous resolution of more than some $r_s \geq 1$ conflicts at any point in time, where r_s depends on the particular sector's capability to handle conflict resolutions. In addition, to further restrict the workload imposed on each sector s to be within its capacity, it is required that no more than some $\bar{n}_s$ flights should simultaneously occupy this sector, and therefore require monitoring. Furthermore, aside from ascribing a cost for the average workload in each sector s, the difference between the peak and average workload is penalized according to an increasing piecewise linear (convex) penalty function to deter extreme variations in workload.

The final module within this framework represents the collaborative decision-making (CDM) equity considerations. For each involved airline, $\alpha \in \{1, \ldots, \bar{\alpha}\}$, based on the ratio of its realized total fuel, delay, and cancellation costs in the overall prescribed solution (x, say) to that for its individually optimized flight plan selections, we define a *collaboration efficiency function* $E_\alpha(x)$ that takes on a value of 1 if this ratio equals unity, and a value of 0 if this ratio equals

some undesirable limiting value $D_{\max}$ (>1). Accordingly, ascribing a weight ω_α to airline α based on its level of operations (number of flights and passengers impacted), where $\sum_{\alpha=1}^{\bar{\alpha}} \omega_\alpha = 1$, we define a *collaboration equity function* $E_\alpha^e(x)$ for each airline α as the difference between its collaboration efficiency and the overall weighted mean (ω-mean) collaboration efficiency, i.e., $E_\alpha^e(x) = E_\alpha(x) - \sum_{\alpha=1}^{\bar{\alpha}} \omega_\alpha E_\alpha(x)$. The model then suitably penalizes both the *weighted* (ω-*mean*) *collaboration inefficiency* $\sum_{\alpha=1}^{\bar{\alpha}} \omega_\alpha[1 - E_\alpha(x)]$, as well as the *weighted* (ω-*mean*) *collaboration inequity* $x^e \equiv \sum_{\alpha=1}^{\bar{\alpha}} \omega_\alpha |E_\alpha^e(x)|$, the latter of which represents the total weighted absolute variations in the individual efficiencies from their mean value. In addition to this weighted mean absolute deviation term, we include another measure of dispersion related to the weighted spread of the efficiency values about their mean. More specifically, this term attempts to minimize the maximum *weighted inequity*, $E_{\max}^e = \max_{\alpha=1,\ldots,\bar{\alpha}}\{\omega_\alpha |E_\alpha^e(x)|\}$. (Some experiments concerning an alterative modeling strategy for minimizing the maximum "regret" $-\omega_\alpha E_\alpha^e(x)$, $\alpha = 1, \ldots, \bar{\alpha}$, are conducted in §5.)

Part I of this paper (see Sherali, Staats, and Trani 2003) focused on the aforementioned modeling issues and derivations. Part II addresses a detailed development of a cost model for representing the fuel, delay, and cancellation costs, along with an estimation of the various model parameters, and presents computational test results based on real Enhanced Traffic Management System (ETMS) data. These computational experiments aim to provide insights into the model, ascertain the algorithmic performance of various modeling constructs and valid inequalities, and study the sensitivity of the results to different pertinent model parameters.

The remainder of this paper is organized as follows. In §2, we review the notation and the APCDM model formulation, and identify the principal model parameters that need to be estimated. Some alternative conflict constraint representations and valid inequalities that require computational testing and comparison are also delineated. In §3, we propose a comprehensive flight plan cost model that includes, in addition to traditional fuel cost measures, a consideration of market-based factors such as hubbed operations, aircraft passenger load profiles, and arrival delay times. Next, in §4, we derive estimates for the various parameters used in the APCDM model. We present the results of our computational experiments for evaluating the different modeling constructs, and conduct some insightful sensitivity analyses, in §5. Finally, in §6, we provide a summary of our research and its findings, and discuss the proposed model's use in tactical as well as strategic contexts toward improving NAS operations.

2. Statement of the APCDM Model Formulation

In this section, we present the mathematical formulation of the APCDM model that is developed in Part I of this paper. To facilitate reading, consider the following list of model indices, variables, parameters, and coefficients.

(a) **Index Sets**

$s = 1, \ldots, S$: Sectors involved in the model analysis.

$\alpha = 1, \ldots, \bar{\alpha}$: Airlines involved in the model analysis.

$f = 1, \ldots, F$: Flights to be scheduled.

$p \in P_f$: Alternative flight plans or surrogates for flight f.

$P_{f0} = P_f \cup \{0\}$, with $p = 0$ representing the flight cancellation surrogate. (Note that canceling a flight is the prerogative of the particular airline and could be offered as an alternative in exchange for improving the status of other flights within the CDM framework. This surrogate ($p = 0$) is ascribed an inordinately high penalty if canceling this corresponding flight is not tendered as a possible option. Also, naturally, the surrogate $p = 0$ is not an option for airborne flights.)

(b) **Decision Variables**

(i) *Principal Decision Variables*:

x_{fp}: Binary variable, which equals one if flight plan $p \in P_{f0}$ is selected for flight f, and zero otherwise, for $f = 1, \ldots, F$. (These are sometimes denoted by x_P or x_Q (etc.) where each of P and Q (etc.) represent some actual flight plan combination (f, p) for $p \in P_f$, $f \in \{1, \ldots, F\}$.)

(ii) *Auxiliary Decision Variables*:

n_s: Peak occupancy level (number of flights) for sector s.

w_s: Average occupancy level (number of flights) for sector s.

$y_{sn} \in [0, 1]$, $n = 0, \ldots, \bar{n}_s$: Convex combination weights attached to the breakpoints of a piecewise-linear increasing convex penalty function, which represents the penalty ascribed to the difference between the peak and the average workload in sector s.

z_{PQ}: Binary variable, which equals one whenever conflicting flight plans P and Q are selected. (It is assumed in the model description that $z_{PQ} \equiv z_{QP}$.)

$E_\alpha(x)$: Collaboration efficiency for airline α.

$E_\alpha^e(x)$: Collaboration equity for airline α.

$x^e \equiv \sum_{\alpha=1}^{\bar{\alpha}} \omega_\alpha |E_\alpha^e(x)|$: ($\omega$-) Mean collaboration inequity.

ν_α: Variable used to represent the term $|E_\alpha^e(x)|$ for airline α.

$E_{\max}^e \equiv \max_{\alpha=1,\ldots,\bar{\alpha}}\{\omega_\alpha \nu_\alpha\} \equiv \max_{\alpha=1,\ldots,\bar{\alpha}}\{\omega_\alpha |E_\alpha^e(x)|\}$: Maximum weighted inequity. (As an alternative, which is discussed in Part I and is further exemplified in §5 of this paper, $E_{\max}^e$ can also be specified as a bounding constant imposed for each weighted inequity.)

(c) **Model Parameters and Coefficients**

ω_α: The weight factor ascribed to each airline α (as discussed above).

c_{fp}: The cost to execute flight plan $p \in P_{f0}$ for flight f.

γ_s: The airspace monitoring (penalty) cost factor for sector s, per unit average occupancy level workload.

ψ_{sn}: The (penalty) cost assessed when the peak monitoring workload in sector s exceeds the average workload by n.

φ_{PQ}: The (penalty) cost ascribed to resolve an en route conflict between flight plans P and Q.

μ^D: The (penalty) cost factor associated with the total weighted collaboration inefficiency attained.

μ^e: The (penalty) cost factor associated with the level of total weighted collaboration inequity attained.

$\mu^e_{\max}$: The (penalty) cost factor associated with the maximum weighted inequity attained.

H: The length of the time horizon under consideration (in minutes).

t^s_{fp}: The length of time (in minutes) that flight plan $p \in P_f$ of flight f occupies sector s.

Ω_s: The set of flight plans that occupy sector s during some subset of the time horizon.

$\bar{n}_s$: The maximum allowable peak monitoring workload (simultaneous flight occupancies) in sector s.

r_s: The maximum number of simultaneous conflict resolutions permitted to exist in sector s.

$c^*_f = \min\{c_{fp}: p \in P_f\}$, for each flight f.

$D_{\max}$: The maximum allowable ratio for any airline of the cost for the selected surrogates to that for the individually optimized surrogates.

A_α: The set of flights belonging to airline α.

W_f: Relative priority weight attached to flight $f \in A_\alpha$ by airline α, where $\sum_{f\in A_\alpha} W_f = 1$, $W_f \geq 0$, $\forall f \in A_\alpha$.

ν^e: The maximal limit imposed on x^e, the ω-mean collaboration inequity.

C_{si}, $\forall i = 1, \ldots, I_s$: The maximal overlapping sets of occupying flight plans for sector s.

M_{sk}, $\forall k = 1, \ldots, K_s$: The maximal overlapping sets of conflicting pairs of flight plans for sector s.

$G_{sk}(N_{sk}, M_{sk})$: The conflict subgraph of conflicting flight plans represented by M_{sk}, where the node set N_{sk} represents the respective flight plans (labeled as P, Q, R, etc.) that are involved within the edge set M_{sk}.

$J_{sk}(P)$: The set of flight plans Q that are adjacent to P within the graph G_{sk}.

FC: The set of fatally conflicting pairs of flight plans (P, Q).

A: The entire set of resolvable conflicting pairs of flight plans (P, Q).

APCDM

$$\text{Minimize} \sum_{f=1}^{F} \sum_{p\in P_{f0}} c_{fp}x_{fp} + \mu^D \sum_{\alpha=1}^{\bar{\alpha}} \omega_\alpha[1 - E_\alpha(x)] + \mu^e x^e + \mu^e_{\max} E^e_{\max} + \sum_{s=1}^{S} \gamma_s w_s + \sum_{s=1}^{S} \sum_{n=0}^{\bar{n}_s} \psi_{sn} y_{sn} + \sum_{(P,Q)\in A} \varphi_{PQ} z_{PQ} \quad (1a)$$

subject to:

$$\sum_{p\in P_{f0}} x_{fp} = 1, \quad \forall f = 1, \ldots, F \quad (1b)$$

$$\sum_{(f,p)\in C_{si}} x_{fp} \leq n_s, \quad \forall i = 1, \ldots, I_s,\ s = 1, \ldots, S \quad (1c)$$

$$w_s = \frac{1}{H} \sum_{(f,p)\in\Omega_s} t^s_{fp} x_{fp}, \quad \forall s = 1, \ldots, S \quad (1d)$$

$$n_s - w_s = \sum_{n=0}^{\bar{n}_s} n y_{sn}, \quad \forall s = 1, \ldots, S \quad (1e)$$

$$\sum_{n=0}^{\bar{n}_s} y_{sn} = 1, \quad \forall s = 1, \ldots, S \quad (1f)$$

$$x_P + x_Q \leq 1, \quad \forall (P, Q) \in FC \quad (1g)$$

$$\sum_{(P,Q)\in M_{sk}} z_{PQ} \leq r_s, \quad \forall k = 1, \ldots, K_s,\ s = 1, \ldots, S \quad (1h)$$

$$x_P + x_Q - z_{PQ} \leq 1, \quad \forall (P, Q) \in A \quad (1i)$$

$$\sum_{Q\in J_{sk}(P)} z_{PQ} \leq r_s x_P, \quad \forall P \in N_{sk}\colon |J_{sk}(P)| \geq r_s + 1, \quad \forall k = 1, \ldots, K_s,\ s = 1, \ldots, S \quad (1j)$$

$$E_\alpha(x) = \frac{D_{\max} \sum_{f\in A_\alpha} W_f c^*_f - \sum_{f\in A_\alpha} \sum_{p\in P_{f0}} W_f c_{fp} x_{fp}}{(D_{\max} - 1) \sum_{f\in A_\alpha} W_f c^*_f}, \quad \forall \alpha = 1, \ldots, \bar{\alpha} \quad (1k)$$

$$E^e_\alpha(x) = E_\alpha(x) - \left(\sum_{\alpha=1}^{\bar{\alpha}} \omega_\alpha E_\alpha(x) \right), \quad \forall \alpha = 1, \ldots, \bar{\alpha} \quad (1l)$$

$$\nu_\alpha \geq -E^e_\alpha(x), \quad \nu_\alpha \geq E^e_\alpha(x), \quad \forall \alpha = 1, \ldots, \bar{\alpha} \quad (1m)$$

$$x^e = \sum_{\alpha=1}^{\bar{\alpha}} \omega_\alpha \nu_\alpha \quad (1n)$$

$$z_{PQ} \geq 0, \quad \forall (P, Q) \in A,\ x_{fp} \text{ binary}, \quad \forall p \in P_{f0},\ \forall f = 1, \ldots, F,$$

$$y_{sn} \geq 0, \quad \forall n = 1, \ldots, \bar{n}_s,\ s = 1, \ldots, S,$$

$$E_\alpha(x) \geq 0, \quad \forall \alpha = 1, \ldots, \bar{\alpha},$$

$$n_s \leq \bar{n}_s, \quad \forall s = 1, \ldots, S,$$

$$x^e \leq \nu^e, \quad E^\alpha_{\max} \geq \omega_\alpha \nu_\alpha, \quad \forall \alpha = 1, \ldots, \bar{\alpha}. \quad (1o)$$

The first objective term is the summation of the fuel, delay, and cancellation costs for the selected

flight plans. The next three terms impose penalties on the attained levels of the ω-mean collaboration inefficiency, the ω-mean collaboration inequity, and the maximum weighted inequity, respectively. Note that although the cost factors $c_{fp}x_{fp}$ appear in both of the first two terms in (1a), the first term represents the total system cost, whereas the second term is dimensionless and includes weighting priority factors among airlines as well as weighting priorities for flights pertaining to each individual airline. Moreover, the second term is weighted with a factor μ^D, which provides the flexibility of running the model with different degrees of relative importance attached to this term. Later, in §4.4, we derive values for μ^D and μ^e that, in extreme cases of the ω-mean collaboration inefficiency and inequity, would yield a penalty of some μ_0% of $\sum_{f=1}^{F} c_f^*$ via the second and third terms. We also provide a sensitivity analysis with respect to μ_0 in §5.3.3 (see Figure 4). Additionally, note that the priority weights W_f, $f \in A_\alpha$ attached to the flights for an individual airline essentially affect its defined efficiency, which then governs the selection of flight plans via the three equity-based terms. Because the entire objective function (1a) represents a systems perspective, the influence of these equity terms is naturally offset by the remaining cost terms, depending on the value of μ_0 or that of the parameter vector $(\mu^D, \mu^e, \mu^e_{\max})$. The fifth term in the objective function (1a) ascribes a penalty to the average sector workloads, the sixth term penalizes the differential between the peak and the average sector workloads according to a piecewise-linear increasing convex function, and the final term penalizes the resolvable conflicts that will need to be addressed by the air-traffic controllers (ATC).

Constraint (1b) requires exactly one flight plan to be selected from among the set of available surrogates. The next four constraints are associated with sector occupancies, where (1c) limits the maximum simultaneous occupancy for each sector s, (1d) evaluates the average sector workloads, (1e) computes the peak-average workload differential, and (1f) requires the convex combination weights ascribed to the breakpoints of the peak-average workload differential penalty function to sum to one. All fatal conflicts are prohibited via (1g). The conflict constraint formulation, described in detail in Part I of this paper, is expressed via (1h) through (1j), where (1h) and (1i) jointly ensure that no more than r_s resolvable conflicts coexist within sector s at any point in time, and (1j) further tightens this representation via a set of star-subgraph convex hull-based valid inequalities. The airline collaboration efficiencies and collaboration equities are determined by (1k) and (1l), respectively. Constraints (1m) and (1n) provide a linear function for computing the ω-mean collaboration inequity. Finally, (1o) imposes the necessary logical and bounding conditions. As expounded in Part I of this paper, in case a model run indicates infeasibility, additional sensitivity analysis runs can be made by relaxing $E_\alpha(x) \geq 0$ and $x^e \leq \nu^e$ from these latter bounding constraints.

Remark 1. In Part I of this paper, we also developed two alternative conflict resolution constraint formulations in lieu of (1h)–(1j). The first of these, denoted C_2 for reference, was a generalization of the approach adopted by Sherali, Smith, and Trani (2002), and is obtained by simply omitting the set of valid inequalities (1j). The second representation, denoted C_3 for reference, was especially derived for the important special case of $r_s = 1$, $\forall s = 1, \ldots, S$, as treated in Sherali, Smith, and Trani (2002), and aside from substituting $r_s = 1$ in (1h), it replaces (1j) with an alternative set of valid inequalities as given by

$$x_P + x_Q + x_R \leq 2, \quad \forall (P, Q, R) \in T_{NC}. \tag{2}$$

Here, based on the conflict subgraphs G_{sk} defined above, T_{NC} is the set of distinct node triplets that do not admit a clique in any conflict subgraph, that is,

$$\begin{aligned} T_{NC} = \{&(P, Q, R),\ P < Q < R\text{: for some } (s, k),\\ &\text{a subgraph of } G_{sk} \text{ that is induced by the}\\ &\text{nodes } P, Q, \text{ and } R \text{ contains precisely two}\\ &\text{edges, but no such subgraph for any } (s, k)\\ &\text{contains three edges}\}. \end{aligned} \tag{3}$$

Sherali, Smith, and Trani (2002) examined another alternative representation of these conflict resolution constraints (1h)–(1j) that we refer to as C_1 in this paper, which does not include the quadratic variables z_{PQ}. Computational experiments reported in that paper, as well as by Sherali and Smith (2003) for a related generalized vertex packing problem, amply demonstrate that C_2 is a practically superior formulation of these constraints than is C_1, although there does not exist any theoretical dominance property between these two formulations. In contrast, we proved in Part I of this paper that C_3 yields a tighter representation of the conflict constraints than both C_1 and C_2 (the latter being obvious), and that the formulation given by (1h)–(1j), which we refer to as C_4, further tightens this representation. Hence, in our computational study, we shall focus on evaluating the relative merits of the alternative conflict resolution formulations C_2, C_3, and C_4.

3. Fuel, Delay, and Cancellation Cost Model

In our proposed cost model, the total cost of executing a flight plan p of flight f is taken as

$$c_{fp} = F_{fp} + D_{fp}, \quad \forall p \in P_f,\ f = 1, \ldots, F, \tag{4}$$

where F_{fp} and D_{fp} are, respectively, the associated fuel and delay costs as discussed in turn below.

3.1. Fuel Cost

Fuel consumption is one metric that is typically used to assess the cost of executing a flight plan. We utilized the base of aircraft data (BADA) operations performance model, developed by the Eurocontrol Experimental Center (1998), to calculate flight fuel costs for any proposed flight plan. The BADA database specifies a set of aircraft performance factors, airline procedure parameters, and performance statistics for 151 distinct aircraft types. The model includes ground movement costs (e.g., for taxiing and parking), as well as airborne costs. The fuel cost of executing a particular flight plan p of flight f, is hence given by a function $\mathsf{F}_{\text{fuel-cost}}(\bullet)$ as defined by this database:

$$F_{fp} = \mathsf{F}_{\text{fuel-cost}}\{\text{aircraft type, mass, flight envelope, aerodynamics, engine thrust, reduced power, fuel consumption, ground movement}\}. \quad (5)$$

3.2. Flight Arrival Delay Costs

Let τ_f^* be the originally scheduled arrival time of a given flight f, and let τ_{fp} be the arrival time pertaining to the particular flight plan p of the flight f. Then the arrival delay time can be expressed as

$$t_{fp}^d = \max\{0, \tau_{fp} - \tau_f^*\}. \quad (6)$$

Let us consider a flight whose arrival at the destination airport is delayed by some $t_{fp}^d \neq 0$. If the flight has only passengers who are presently traveling to their terminal destination, then the delay cost can be expressed simply as some function of t_{fp}^d. However, if there are passengers who must connect to other flights, then the delay cost must include a consideration of any impact of this delay to downstream flights. For example, if a passenger misses a connection due to a late arrival, then the planned outbound flight departs with an empty seat, and the airline must reschedule the passenger to depart on a subsequent flight (which perhaps has already been oversold). This phenomenon can cascade through the system, where a single delayed flight can impact a number of other flights as the affected passengers continue toward their respective terminal destinations.

Let us quantify the costs associated with delays on the basis of the flight's passenger profile (i.e., the proportion of passengers arriving at their terminal destinations). For example, a flight inbound to a regional airport would incur a lower per-passenger delay cost than a flight inbound to a major hub airport. One possible approach to quantify these costs might be to examine each flight and determine the actual proportion of passengers who are en route to their terminal destination. This method, however, would require a detailed data specification for a large number of flights that might not be available, or for that matter, might not be suitable in a strategic planning context. Alternatively, we may classify flights according to the respective destination airport's most common (or average) inbound passenger profile. While this approach only approximates flight profiles, it has the advantage that once each airport has been classified, any combination of flights or surrogates may be considered in the model without the need for meticulous data collection. Observe that such an approach includes an inherent estimation of downstream cost impacts, and permits us to consider each flight plan independently.

Accordingly, we assign to each airport a *connection delay cost factor*, based on the most common passenger profile for inbound flights. For simplicity, we group airports into one of three categories: low, medium, and high connection rate airports. These categories, respectively, contain small regional airports, small to medium-sized hub airports, and major hub and international airports (see Staats 2003 for a listing of the latter two groups of airports). The corresponding connection delay cost factors can be used to inflate the related consequences of delays. For example, we can take the low, medium, and high connection rate factors to be 1.0, 1.5 (or 1.25), and 2.0, respectively, based on a linear (or nonlinear) relationship, as desired. Note that an actual estimation of this factor is possible, but would require a detailed airport flight-connection analysis. In fact, Beatty et al. (1998) have conducted such an analysis for specific airports and have shown how one could compute a *delay multiplier* factor for each flight or, by aggregating data, for each airport, based on the time of day. In general, let d_f^c denote the multiplicative connection delay cost factor that is ascribed to flight f based on its terminal airport and its originally scheduled arrival time. In the absence of detailed flight-connection data, we take d_f^c to equal the particular connection delay cost factor ascribed to the category that contains the terminal airport.

Naturally, the per-minute delay cost is a function of the number of passengers affected. Rather than requiring extensive data on actual passenger counts for each flight, we use a *passenger load estimate*, l_f, for each flight f, that depends on the type of aircraft and the estimated load factor.

The delay cost for a surrogate p of flight f can now be written as

$$D_{fp} = \left(t_{fp}^d d_f^c\right)(l_f)(\delta), \quad (7)$$

where δ is the average delay cost per passenger-minute across all airlines and their respective flights. A typical value for δ is \$0.20, based on an average estimate provided by the Air Transport Association (see Chang et al. 2001).

3.3. Flight Plan Cancellation Cost

The cost model (4) pertains to the actively considered potential flight plans. We will also need to delineate a cost coefficient c_{f0} for the plan $p = 0$ of each flight f, which corresponds to the flight being cancelled, whenever this possibility is tendered as an option by the corresponding airline for consideration within the CDM framework. Because a cancelled flight is the least preferred surrogate, we should expect that $c_{f0} > \max_{p \in P_f}\{c_{fp}\}$. Let $t^d_{f0} > \max_{p \in P_f}\{t^d_{fp}\}$ be an estimated delay value ascribed to canceling flight f (which might depend on the average turnaround time for reestablishing a connection, the average passenger waiting time required to obtain an alternative connecting flight, or some upper bound on this value in the absence of reliable data). The cancellation cost coefficient should additionally include a consideration of the number of passengers affected, as well as the connecting flights that might be impacted as discussed previously. Accordingly, consistent with (7) and (4), we define the cost of a cancelled flight to be

$$c_{f0} = \max_{p \in P_f}\{F_{fp}\} + \left(t^d_{f0} d^c_f\right)(l_f)(\delta), \tag{8}$$

where the first term in (8) reflects a conservative estimate of the net fuel cost penalty attributed to the rescheduled passengers. Note that in case the cancellation of a particular flight f is not an option tendered by the corresponding airline, we assign c_{f0} a prohibitively high value.

Remark 2. Hansen, Gillen, and Djafarian-Tehrani (2001) have recommended that schedule disruption is a cost metric preferred to the average flight delay time that is traditionally used. Such schedule disruption costs (to augment the use of the proposed d^c_f-factors) can also be accommodated within the above approach, provided that the relevant data for estimating these costs exists.

Remark 3. (Interaction with Airlines) Observe that to ensure uniformity across airlines, we have suggested (in §3.1) the use of the BADA database for the computation of fuel costs for any flight plan, rather than have the individual airlines provide this information. This uniformity is important particularly from the point of view of the first term in the objective function (1a), and less so in regard to the equity terms in (1a) because of the scaling adopted in (1k) for defining the collaboration efficiency as a dimensionless quantity. Although different airlines might have different cost bases, the approach of using standardized cost factors is reasonable because the first term in (1a) reflects a system-based total cost entity. Moreover, by specifying desired relative values of W_f, $f \in A_\alpha$, each airline α has the opportunity to emphasize any subset of its flights in the CDM framework depending on the perceived relative cost bases for its own flights. This feature of using dimensionless scaled efficiency terms along with desired W_f-values would help mitigate any biases against airlines that might avoid major airports, and hence not incur heavy delay costs that are inflated by connective delay cost factors or delay multipliers. In particular, such airlines could de-emphasize flights that are not subject to significant detour and delay costs in the CDM consideration. It is also important to note that certain uniformity rules need to be established to govern the set of alternative surrogate flight plans that are submitted by the competing airlines. For example, an airline may choose to prescribe surrogates in such a way that the associated distribution of relative costs increases the likelihood that a preferred flight plan is selected by the optimization model. One approach to mitigate any bias that could be injected into the model by an individual airline is to ensure that the distribution of costs associated with the different surrogates is similar for each of the participating airlines. In other words, the airlines should be required to submit surrogate plans for each flight that are uniform (across all flights) with respect to their extents of delays (e.g., a surrogate for each 15–30 minutes of delay, as applicable depending on the disruption scenario, up to some limit). This way, airlines can focus on optimizing the individual flight plan trajectories, given the enforced departure delay and the existing environmental (wind, weather, imposed SUAs, etc.) situation, and the model would then equitably select a mix of viable flight plans. Also, note that the proposed equity framework is different from the presently implemented slot-based ration by schedule (RBS) and compression strategies (see Ball et al. 1998, 2000, and Vossen and Ball 2001), and requires further validation studies to promote acceptance in practice.

4. Estimation of Objective Parameters

4.1. Airspace-Monitoring Penalty Costs (Parameters γ_s, $\forall s$)

For estimating the parameter γ_s in the objective function (1a), we consider costs that might be attributed to the average workload in terms of the related ATC sector crew expertise necessary to safely perform operations. The ATC infrastructure cost is taken to be a "sunk cost." Each sector within an ARTCC generally employs a three-person ATC crew of certified professional controllers (CPC). The crew consists of a flight data controller, a radar controller, and a radar associate/nonradar controller (see Nolan 1999). When ATC workload is minimal, relatively inexperienced crews (having minimal salaries) might be used, whereas more experienced crews might be necessary for increased workloads to safely perform more complex operations. When operating at sector capacity,

the most qualified crews (having maximal salaries) might likely be used to handle such highly stressful operations. The FAA publishes vacancy announcements for CPC positions. For example, the New York ARTCC advertisement dated September 2002 seeks to fill multiple CPC positions having annual salaries ranging from \$90,268 to \$126,375, with the median salary being \$108,322 (FAA 2002).

Suppose that each controller can monitor an average of five aircraft simultaneously throughout the horizon of H minutes. Then sector s would require an estimated average of $w_s/5$ controllers working H minutes. Assuming that each controller works 40 hours per week and 50 weeks per year, and a per-controller (median salary-based) annual cost of \$216,644, including a 100% employee overhead cost (hiring, training, medical benefits, retirement plans, etc.), the per-minute, per-controller cost would be $(\$216{,}644)/[(50)(40)(60)] = \$1.805$, leading to a total cost of $(\$1.805)H(w_s/5)$. Hence, the cost per unit of the average workload w_s for sector s can be taken as

$$\gamma_s = \frac{(\$1.805)H}{5} = \$0.361H, \quad \forall s = 1, \ldots, S. \tag{9}$$

4.2. Peak-Average Differential Airspace-Monitoring Penalty Functions (Parameters ψ_{sn}, $n = 0, \ldots, \bar{n}_s$, $\forall s$)

We use a similar rationale as in the foregoing subsection to generate a suitable penalty function to be associated with the peak-average differential in sector occupancy workloads. We assume a convex increasing rate penalty function ψ_{sn} as given by a quadratic function of n, say, $\psi_{sn} \equiv \frac{1}{2}a_s n^2$. Suppose further that we let the penalty ascribed to a peak workload exceeding the average by one unit to be $\gamma_s/5$, assuming that this has the same effect as increasing the average workload by 1/5. Hence, $\psi_{s1} = a_s/2 = \gamma_s/5$. This yields

$$\psi_{sn} = \left(\frac{\gamma_s}{5}\right)(n)^2, \quad \forall n = 0, \ldots, \bar{n}_s, \forall s = 1, \ldots, S, \tag{10}$$

where γ_s is given by (9).

4.3. Conflict Resolution Penalty Costs (Parameters φ_{PQ}, $\forall (P, Q) \in A$)

Conflict resolution activities increase the workload beyond that required for basic airspace-monitoring activities. Avoiding midair aircraft collisions is perhaps the most critical aspect of an air-traffic controller's duties, and intense conflict resolutions can be a highly stressful task. Ideally, only the most qualified and experienced controllers would perform these actions. Therefore, busy air route traffic control centers that experience a significant level of conflicts would need to recruit relatively highly skilled controllers. Using the upper range on the ATC salary range mentioned in §4.1 (including a 100% overhead), we recommend a commensurate additional conflict resolution per-minute penalty cost of

$$\frac{\$(126{,}375 - 108{,}322)(2)}{(50)(40)(60)} = \$0.301.$$

Accordingly, suppose that an ATC controller, on average, initiates conflict resolution actions 10 minutes prior to the projected aircraft separation violation. The controller issues new vectors to the affected flights, monitors compliance, and issues another set of vectors to the aircraft to return them to their respective nominal trajectories. Assuming symmetric conflict avoidance trajectories, from and back to the nominal flight plan trajectories, suppose that conflict resolution activities therefore persist for approximately 20 minutes for each conflicting pair of aircraft. Using the above per-minute cost factor, we let the conflict resolution penalty cost coefficient be given by

$$\varphi_{PQ} = I_{PQ}(\$0.301)(20) = (\$6.02)I_{PQ}, \tag{11}$$

where I_{PQ} is the *intensity* of the conflict (P, Q). In this study, we take $I_{PQ} = 1$ for Level-1 conflicts (see §1), and $I_{PQ} = 4$ for Level-2 conflicts, to reflect a quadratic intensity increase corresponding to the half-sized separation box (in the horizontal plane) that is violated in this instance.

Remark 4. We can modify (11) to account for additional conflict geometry information as generated by the PAEM. For example, conflict resolution actions may differ based on relative headings. An intruder moving along the cross-track axis of a focal aircraft's trajectory might avoid a potential conflict by either increasing or decreasing its air speed, rather than by altering its trajectory. Such conflict geometry implications could be explored to further refine (11).

4.4. Collaboration Inefficiency and Inequity Penalty Factors (Parameters μ^D, μ^e, and $\mu^e_{\max}$)

In this section, we ascribe values to the penalty parameters μ^D, μ^e, and $\mu^e_{\max}$ that are used as relative weights for the dimensionless equity terms so as to make the resulting sum commensurate with the other cost terms in the objective function (1a). The basic idea adopted is to examine scenarios under which the ω-mean collaboration inefficiency, $(\omega_\alpha[1 - E_\alpha(x)])$, or the ω-mean collaboration inequity (x^e) takes on its worst possible value, and assume that each of these scenarios yields a total penalty via the equity terms of some $100\mu_0\%$ of the best possible total flight plan cost, where $0 < \mu_0 < 1$ is specified in the sequel. More specifically, recall that $0 \le \sum_\alpha \omega_\alpha E_\alpha(x) \le 1$ and observe that $0 \le x^e \le 0.5$. Furthermore, we obtain $x^e = 0.5$ when, for example, given an even number of airline participants partitioned into two sets

A_1 and A_2 with $\sum_{\alpha\in A_1}\omega_\alpha = \sum_{\alpha\in A_2}\omega_\alpha = \frac{1}{2}$, we have $E_\alpha(x)=1$, $\forall\,\alpha\in A_1$, and $E_\alpha(x)=0$, $\forall\,\alpha\in A_2$, whence $\sum_\alpha \omega_\alpha[1-E_\alpha(x)]=\frac{1}{2}$ as well. Hence, the worst-case values for the ω-mean collaboration inefficiency and the ω-mean collaboration inequity occurs whenever the two-tuple $(\sum_{\alpha=1}^{\bar\alpha}\omega_\alpha[1-E_\alpha(x)], x^e)$ equals $(1,0)$ or $(\frac{1}{2},\frac{1}{2})$, respectively. In general, regardless of the data, we now examine these two worst-case scenario values and derive the related objective parameters so as to incur a certain penalty level whenever these values are achieved. Specifically, consider first the case where $E^e_{\max}$ is a specified constant. Suppose that we assign values to the factors μ^D and μ^e in (1a) in such a way that whenever the foregoing two-tuple equals $(1,0)$ or $(\frac{1}{2},\frac{1}{2})$, the related objective function expression, $\mu^D\sum_{\alpha=1}^{\bar\alpha}\omega_\alpha[1-E_\alpha(x)]+\mu^e x^e$, yields a penalty of, say, $100\mu_0\%$ of the total flight plan cost corresponding to each airline's individually optimized strategy. That is, we obtain a penalty of $\mu_0\sum_{f=1}^F c^*_f$ for either of the aforementioned worst-case situations. This yields

$$\mu^e=\mu^D=\mu_0\sum_{f=1}^F c^*_f. \tag{12a}$$

In addition, whenever the maximum weighted inequity $E^e_{\max}$ is taken as a variable in the model, we ascribe the corresponding penalty parameter $\mu^e_{\max}$ a value equal to $\xi\mu^e$, for some factor $\xi\ge 1$. Repeating the foregoing analysis with the additional penalty term $\mu^e_{\max}E^e_{\max}$ under the condition $\mu^e_{\max}=\xi\mu^e$, we get

$$\begin{aligned}&\mu^D=\mu_0\sum_{f=1}^F c^*_f,\\ &\mu^e=\frac{\bar\alpha}{(\xi+\bar\alpha)}\mu^D,\quad\text{and}\quad \mu^e_{\max}=\xi\mu^e.\end{aligned} \tag{12b}$$

In the next section, we study the sensitivity of flight plan selections at optimality with respect to varying the value of μ_0 from the baseline level of 0.1, as well as the effect on equity by varying the parameter $\xi\ge 1$.

5. Computational Experiments

In this section, we perform six sets of experiments to evaluate the effectiveness of the APCDM model and its submodules, and to provide insights into the sensitivity of the model to various pertinent parameters.

First, we examine the performance of our stochastic approach for generating aircraft trajectory realizations and conducting the overall conflict analysis. In addition to testing several probabilistic trajectory displacement distributions, we establish a set of baseline threshold probabilities for defining whether a significant conflict risk exists between any pair of flight plans. We also perform a sensitivity analysis for the three conflict threshold probabilities by varying them from their respective baseline values, and provide insightful comparisons against a deterministic conflict analysis approach. Following this analysis, four other sets of experiments are dedicated to testing the sensitivity of the APCDM model's solution to the CDM-related parameters, $D_{\max}$, $E^e_{\max}$ (when used as a constant), ξ (when $E^e_{\max}$ is taken as a variable), and μ_0 (see §2 and Equations (1a) and (12)). Finally, we explore the effectiveness of the alternative conflict constraint formulations, as defined in §2. The priority weights among airlines and for the flights within each airline are taken as $\omega_\alpha=|A_\alpha|/F$, $\forall\,\alpha=1,\ldots,\bar\alpha$, and $W_f=1/|A_\alpha|$, $\forall\,\alpha=1,\ldots,\bar\alpha$, throughout all problem instances described in the sequel. Each of the proposed flight plans in these experiments involve trajectories that traverse respective portions of the 88 sectors comprising the Miami-Jacksonville ARTCC. All reported computations have been performed on a Dell Inspiron 7500 laptop computer equipped with a 650 MHz Pentium III processor and 192 MB random access memory, and running the Microsoft Windows 2000 operating system. The various submodules were programmed using Microsoft Visual C++ Version 6, and the APCDM mixed-integer programs were solved using CPLEX 7.0.

5.1. Description of the Test Sets

The test cases for our computational experiments were constructed using real data based on the ETMS flight data information related to the Miami-Jacksonville ARTCCs, as obtained from the FAA. For each test set, described in Table 1 below, a subset of flights was taken from the database having intersecting flight trajectories that provide a representative sampling of the overall ARTCC traffic density. (The time horizon was selected to be sufficiently long to accommodate this subset of flights.) For each flight taken from this database, several alternative trajectories were generated between the respective origin-destination pair, assuming suitable diversions as prompted by the closure of a section of the airspace in the involved region. The schedules for the flight plan surrogates were then generated based on these alternative trajectories by suitably modifying the respective take-off times by 20-minute intervals. For example, to generate a particular data set, we might develop three alternative trajectories for each flight, and for each of these, we might adopt two alternative scheduled take-off times, thus yielding a total of six surrogates for each flight (a "cancellation surrogate" is additionally incorporated as well).

5.2. Computational Evaluation of the PAEM Submodule and the Determination of Threshold Probabilities

For the purpose of testing the PAEM submodule and determining an appropriate choice of threshold

Table 1 Test Data Sets

	Test Set 1	Test Set 2	Test Set 3	Test Set 4
Number of flights	30	80	100	120
Surrogates per flight	6	6	9	6
Centers (# of sectors)	ZMA/ZJX (29)	ZMA/ZJX (35)	ZMA/ZJX (16)	ZMA/ZJX (46)
Number of airlines	5	5	6	4
Airlines: flights per airline	AAL: 10 ABX/ACA/AMT: 5 CAA: 7 COA: 3 DAL: 5	AAL: 18 ABX/ACA/AMT: 5 CAA: 7 COA: 25 DAL: 25	AAL: 14 CAA: 18 COA: 25 DAL: 9 DEL: 22 UAL: 12	AAL: 40 DAL: 23 UAL: 22 USA: 35
Time horizon	1,200 min	1,200 min	200 min	1,800 min
Number of deterministic conflicts identified	Level 1: 69 Level 2: 15 Fatal: 0	Level 1: 559 Level 2: 235 Fatal: 5	Level 1: 3,999 Level 2: 5,498 Fatal: 1,285	Level 1: 1,211 Level 2: 704 Fatal: 55

probabilities, we used the first two data sets described in Table 1, and considered two types of probabilistic trajectory displacement regions that were developed in Part I of this paper: randomized and wind induced. The displacements for the latter were derived using two different (unitized) wind vectors, corresponding, respectively, to a south-southwesterly and a northerly wind, both measured at Palm Beach International (FL) Airport.

5.2.1. Evaluation of the PAEM Model Performance. We first examined the effort consumed by PAEM for a given problem as a function of the number of probabilistic trajectory realizations. Recall that for each piecewise-linear trajectory segment of the focal aircraft A's flight plan, a conflict analysis is performed for each intruder aircraft B for all pairs of trajectory realizations $(k_1, k_2) \in \{1, \ldots, n_A\} \times \{1, \ldots, n_B\}$ (where $n_A = n_B$ is the number of probabilistic realizations). Figure 2 plots the CPU seconds of computational time taken by PAEM to conduct the conflict analysis using eight randomized and wind-induced probabilistic trajectory displacements for each of Test Set 1 and Test Set 2.

From Figure 2, we see that there is an increasing quadratic relationship between the number of probabilistic realizations per flight and the CPU time. Clearly then, it is advantageous to limit the number of realizations, while maintaining a reasonable level of representability of the probabilistic distribution of trajectories, particularly when processing large data sets. Note that other factors that influence the PAEM effort across different problem instances include the number of flights (and the number of surrogates proposed for each flight), the nature of the respective flight plans (e.g., the number and duration of the piecewise-linear flight segments), and the extent of potential pairwise conflicts that are induced based on overlapping travel times and common assigned en route flight levels.

Next, we examined the conflicts identified by the PAEM using each of the randomized and wind-induced trajectory displacement instances. First, we noted that all the deterministic conflicts (based on filed nominal errorfree flight plans) were identified in the probabilistic analyses, as expected, because the deterministic trajectories are included within the set of probabilistic realizations. Second, we observed that these deterministic conflicts occurred in the probabilistic analyses with an average probability of occurrence ranging from 0.30 to 0.32. Moreover, for each problem, the deterministic and the probabilistic

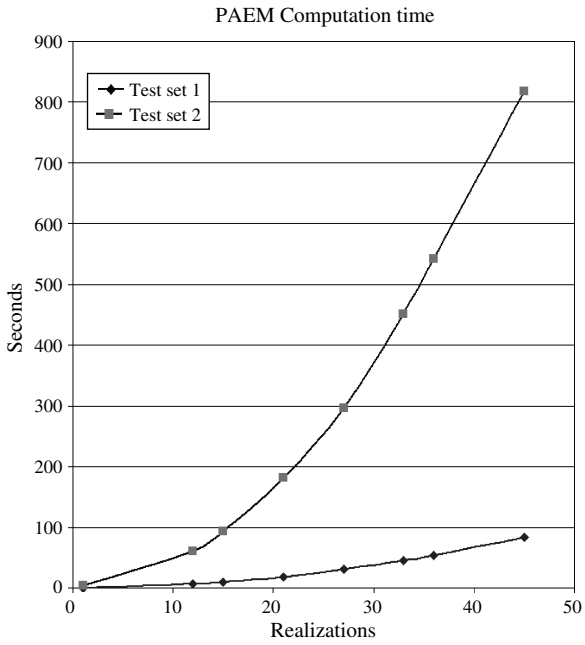

Figure 2 PAEM Processing Time vs. Number of Realizations

analyses exhibited somewhat similar, but not identical, pairs of conflicting aircraft.

We next compared the results using the three-dimensional trajectory displacements with the results using the counterpart two-dimensional trajectory displacements having the same number of realizations in the horizontal plane. Observe that for the three-dimensional displacements, the realizations for a given flight are geometrically separated from each other by approximately 267 feet in the vertical dimension. Given that the FAA imposes a vertical aircraft separation standard of at least 1,000 feet, we observed that these vertical displacement errors were relatively smaller than the enforced standard separation for the nominal flight plans, and hence, the three-dimensional consideration did not generate additional probabilistic conflicts over the two-dimensional analysis unless one or both of the focal or intruder aircraft were ascending or descending. Our primary interest for the APCDM model is in analyzing the alternative flight plans mainly in the en route airspace where aircraft are normally in level flight. Therefore, we recommend a reasonable trade-off between accuracy in representation and PAEM processing effort by eliminating the consideration of vertical trajectory displacements.

5.2.2. Determining Threshold Probabilities. Using the number of deterministic conflicts recorded as a benchmark, we examined the number of conflicts identified as the threshold probability parameter p_1 was varied. Table 2 presents the results for the number of Level-1 conflicts detected by PAEM as the threshold probability p_1 is varied, using two cases of randomized displacement realizations for each of the two test sets.

Note that the CTAS model of Erzberger et al. (1997) nominally identifies conflicts that exceed a (total) probability of 0.45. This threshold level might miss several potentially conflicting situations that yet have a significant probability of occurrence. To be more conservative, we selected $p_1 = 1/3$ as a reasonable threshold to identify potential Level-1 conflicts. Another important observation is that the conflicts identified during the deterministic analysis correspond to aircraft trajectories that actually have a very small probability of occurrence. Using a probabilistic analysis (with any satisfactory threshold probability), as opposed to a deterministic analysis, yields a significantly more robust prediction of potential aircraft conflicts, given that navigation errors do occur in practice.

Table 2 **Level-1 Conflicts vs. Threshold Probability p_1**

	Test Set 1		Test Set 2	
p_1	$n_A = 15$	$n_A = 21$	$n_A = 15$	$n_A = 21$
0.50	20	20	178	157
0.45	27	22	264	230
0.40	45	44	373	369
0.35	59	60	495	527
1/3	68	64	583	600
0.30	78	80	728	726
0.25	103	115	991	974

Table 3 **Level-1 Plus Level-2 Conflicts vs. Threshold Probability p_2**

	Test Set 1		Test Set 2	
p_2	$n_A = 15$	$n_A = 21$	$n_A = 15$	$n_A = 21$
1/3	68	64	600	583
0.30	68	64	600	583
0.25	68	64	600	584
2/9	68	64	612	590
0.20	84	66	638	604
1/6	88	88	660	647
0.15	93	88	703	657

Having selected a Level-1 threshold probability, we examined the total conflicts identified as a function of p_2, the threshold probability for identifying Level-2 conflicts. Table 3 presents the total number of Level-1 and Level-2 conflicts, which occur, respectively, with probabilities p_1 and p_2, as p_2 varies while holding $p_1 = 1/3$. Recall that the dimensions of the Level-2 separation box are half those of the Level-1 separation box, while the conflict probabilities are the product of two realization probabilities. Hence, taking the squares of the threshold probabilities to be in proportion to the respective cross-sectional box areas, we recommend $p_2 = p_1/2 = 1/6$. From Table 3, this appears to be a reasonable, conservative threshold for identifying Level-2 conflicts.

Note that throughout the range of threshold probabilities tested for these and other test cases, a similar number of conflict intervals were identified using either $n_A = 15$ or $n_A = 21$ realizations in the horizontal plane. Upon inspection of the respective sets of conflict intervals, we found that using $n_A = 21$ versus $n_A = 15$ realizations did not yield a substantially different set of conflict graphs. Therefore, we shall use the computationally less burdensome case of $n_A = 15$ in further analyses.

Turning our attention to determining a threshold probability for identifying fatal conflicts, observe that each pair of trajectory displacement realizations generated in the horizontal plane, for the several randomized displacement instances that we examined, were separated by a minimum geometric distance of more than 2,000 feet (occurring in the cross-track dimension). The inviolable airspace around each focal aircraft (which, when pierced by an intruder aircraft, defines a fatal conflict) extends 500 feet in both the

in-trail and cross-track axes. Given this fact, we expect that any fatal conflicts identified will have relatively small associated probabilities. This was the case with the two data sets examined for all of the randomized trajectory displacement distributions. (Note that there are no identical or closely parallel flight plans included in either data set.) The average fatal conflict probabilities for each distribution ranged from 0.011 to 0.043. The maximum fatal conflict probability encountered was 0.0623. On the other hand, the wind-induced probabilistic displacement realizations for the distributions tended to be relatively closer together geometrically by their natural propensity, having a minimum separation distance of approximately 350 feet. Using these distributions with the two data sets yielded similar average fatal conflict probabilities as before. However, as expected, the maximum fatal conflict probability encountered was significantly higher, being 0.1518. Consistent with the values of p_1 and p_2, we chose a fatal threshold probability $p_{\text{fatal}} = 1/18$. Note that this is greater than any of the averages of fatal conflict probabilities encountered in our experiments, but less than the maximum values of these fatal conflict probabilities.

To summarize, we have selected our baseline threshold probabilities to be

$$p_{\text{thresh}} = \{p_1, p_2, p_{\text{fatal}}\} = \left\{\frac{1}{3}, \frac{1}{6}, \frac{1}{18}\right\}. \tag{13}$$

Using these baseline threshold probabilities, Table 4 identifies the conflict information for the two test data sets. We shall next study the sensitivity of the APCDM solution to selecting more or less conservative thresholds (i.e., respectively, smaller or greater threshold probabilities), by simultaneously varying these probabilities (proportionately) from their nominal values given in (13).

5.3. Sensitivity Analysis with Respect to the Conflict Risk Threshold Probabilities

The present analysis was performed using Test Sets 2 and 3. The sector capacity was nominally restricted with $\bar{n}_s = 15$, $\forall s = 1, \ldots, S$, for Test Set 2 and moderately constrained with $\bar{n}_s = 9$, $\forall s = 1, \ldots, S$, for Test Set 3. The Level-1 threshold probability was varied by up to ±12% of its nominal value of 1/3, with the other threshold probabilities for Level 2 and fatal conflicts being varied proportionately from their respective nominal values given in (13). The results revealed that the objective value, the ω-mean collaboration efficiency, and the ω-mean collaboration inequity were all insensitive to changes in the conflict risk threshold probabilities for both data sets (see Staats 2003 for detailed results). The reason for this relative insensitivity is that a reasonable range of variation in the threshold probabilities leaves the effective conflict graphs G_{sk} identified relatively unchanged, and the model is able to retain essentially the same solution, provided that it is able to accommodate the few additional conflicts within the sectors' handling capability. Even when the solutions were perturbed, we observed that the mix of available surrogates for these test instances was rich enough to provide a comparable trade-off between the two CDM penalty functions and the direct costs of the selected surrogates. Although additional tests might be useful to further establish these findings, our preliminary conclusion is that the model is not particularly sensitive to moderate changes in the conflict risk threshold probabilities.

Table 4 PAEM Computational Results

	Test Set 1			Test Set 2		
Number of flights		30			80	
Surrogates per flight		6			6	
Realizations (planar)	15	15	15	15	15	15
Wind direction	n/a	SSW	N	n/a	SSW	N
Solution time (seconds)	9.94	9.82	9.82	99.11	93.86	93.92
Probabilistic conflicts						
Level 1:	28	8	20	242	117	275
Level 2:	60	42	27	415	446	393
Fatal:	4	0	0	7	4	7

5.4. Sensitivity with Respect to Various CDM Parameters

In this section, we present, in turn, some sensitivity analyses with respect to the four CDM parameters, $D_{\max}$, $E^e_{\max}$, ξ (when $E^e_{\max}$ is taken as a variable—see Equation (12b)), and μ_0.

5.4.1. Sensitivity Analysis with Respect to the Parameter $D_{\max}$. Recall that the parameter $D_{\max}$ is the maximum allowable ratio for any airline of its cost pertaining to the set of surrogates selected through the CDM process to that of its individually optimized set of surrogates. In particular, defining the latter ratio as

$$d_\alpha(x) = \frac{\sum_{f \in A_\alpha} \sum_{p \in P_{f0}} W_f c_{fp} x_{fp}}{\sum_{f \in A_\alpha} W_f c^*_f}, \tag{14}$$

we see via (1k) that the efficiency $E_\alpha(x) = 1$ if $d_\alpha(x) = 1$, and $E_\alpha(x) = 0$ if $d_\alpha(x) = D_{\max}$. To provide some insights into the effect of $D_{\max}$, we examined the sensitivity of the objective function value to progressively decreasing values of $D_{\max}$, using four APCDM instances. We commenced with $D_{\max} = 1.5$ (an unreasonably high value because airlines would be unlikely to accept a 50% cost increase associated with participating in a group decision) and then gradually decreased this in discrete steps until $D_{\max} = 1.02$. Throughout this experiment, we held $\nu^e = E^e_{\max} = \infty$

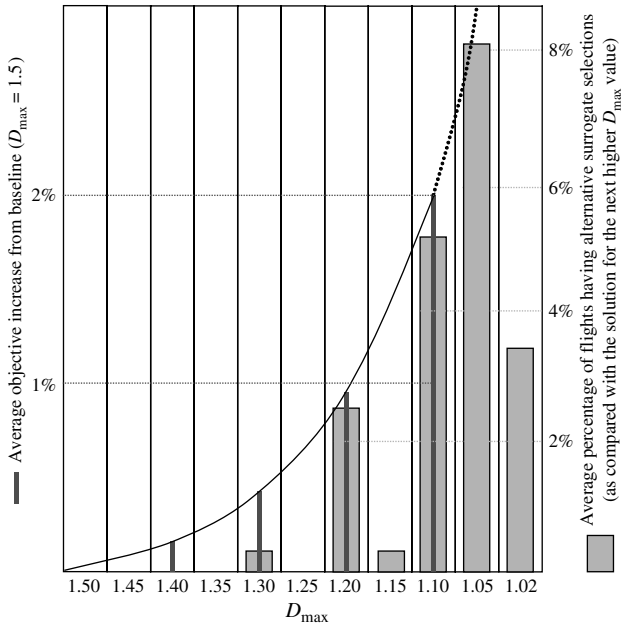

Figure 3 D_{max} **Analysis**

in (1o) to suppress the effects of the corresponding bounding restrictions. Figure 3 depicts the results obtained for this sensitivity analysis, using Test Sets 2 and 3.

Observe that the optimal objective function value increases nonlinearly as D_{max} decreases, but that the rate of this increase is quite small. Indeed, much of the increase can be directly attributed to the progressively steeper slope of the airline efficiency $E_\alpha(x)$ as a function of $d_\alpha(x)$, which occurs as D_{max} is decreased. Note that as the slope of this affine efficiency function is made steeper by virtue of decreasing D_{max}, any change in flight plan selections will produce a relatively greater impact on the airline collaboration efficiency. In fact, although relatively lower values of the ratio $d_\alpha(x)$ tend to be achieved as D_{max} decreases, several of these get proportionally closer to D_{max}, thereby reducing their efficiency values. Hence, due to the resultant dispersion of efficiency values, we noted a corresponding increase in the solution's ω-mean collaboration inequity as D_{max} was decreased. To gain further insights, we examined the impact D_{max} had on the selection of surrogates for the various flights. Specifically, we determined the number of flights where a different surrogate was selected at optimality from the one selected when D_{max} was at the next higher level. Referring to Figure 3, observe that for $D_{max} > 1.2$, this parameter had little impact on the surrogates selected. Below this value, the optimal selection of surrogates became significantly more sensitive to the value of D_{max}. By examining the APCDM model's output more closely, we noted a general trend toward an increase in the average and peak sector occupancies and in the number of resolvable conflicts as D_{max} decreased. This translated to an increase in the sector-monitoring and conflict resolution (penalty) costs, in addition to the aforementioned increase in the efficiency-based penalty costs for decreasing values of D_{max}.

Based on the foregoing analysis, we selected a value of $D_{max} = 1.20$ for implementation. This value also reflects a level where this parameter can be expected to influence the selection of surrogates at optimality.

5.4.2. Sensitivity Analysis with Respect to the Parameters E^e_{max} and ξ. Having selected $D_{max} = 1.20$, we first explored different values of E^e_{max}, considering this to be a specified constant in the model, and examined the resulting airline collaboration equities ($E^e_\alpha(x)$), the ω-mean collaboration inequity (x^e), and the ω-mean collaboration efficiency ($\sum_\alpha \omega_\alpha E_\alpha(x)$) that were achieved at optimality for each of four APCDM instances derived from Test Sets 2 and 3. These four instances (labeled CDM-2a, CDM-2b, and CDM-2c, based on Test Set 2, and CDM-3a based on Test Set 3) were generated by restricting selected sector capacities (as might be influenced by severe weather disruptions) to create scenarios in which the affected sectors could not accommodate all flight plans, as originally proposed. In doing so, these scenarios prompt a significant level of trade-off in selecting between the alternative flight plan surrogates due to the corresponding effects on individual airline equities and the ω-mean collaboration inequity. To begin with, we considered $E^e_{max} = \infty$, which we call the *unconstrained* case, where the individual airline inequities are not forcibly constrained to lie within any designated bound. The results obtained are displayed in Table 5.

Observe that the three instances that yielded relatively high ω-mean collaboration efficiencies (exceeding 94%) also resulted in relatively small ω-mean collaboration inequities. Hence, we focused on instance CDM-2c, for which this did not occur, and examined the sensitivity of its solution to reductions in E^e_{max}. Note that instance CDM-2c corresponds to a scenario in which certain critical airspace sectors are more tightly constrained than in the other three instances. Table 6 presents the values obtained for the ω-mean collaboration efficiency, the spread in the weighted

Table 5 **CDM Values at Optimality (Unconstrained Inequities)**

APCDM instance	$\sum_\alpha \omega_\alpha E_\alpha(x)$	$\max_\alpha\{\omega_\alpha \lvert E^e_\alpha(x)\rvert\}$	$\max_\alpha\{-\omega_\alpha E^e_\alpha(x)\}$	x^e
CDM-2a	0.9608	0.0034	0.0034	0.0070
CDM-2b	0.9474	0.0028	0.0027	0.0087
CDM-2c	0.5071	0.0234	0.0191	0.0823
CDM-3a	0.9928	0.0012	0.0012	0.0024

Table 6 Sensitivity Analysis for the Parameter $E^e_{\max}$ Used as a Constant

$E^e_{\max}$	$\sum_\alpha \omega_\alpha E_\alpha(x)$	$\max_\alpha\{\omega_\alpha \lvert E^e_\alpha(x)\rvert\}$	$\max_\alpha\{-\omega_\alpha E^e_\alpha(x)\}$	x^e	% objective increase with respect to the unconstrained case (%)
Unconstrained	0.5071	0.0234	0.0191	0.0823	0
$0.1/\bar{\alpha}$	0.5016	0.0189	0.0159	0.0612	~0
$0.09/\bar{\alpha}$	0.5000	0.0177	0.0154	0.0589	0.03
$0.08/\bar{\alpha}$	0.4975	0.0158	0.0146	0.0554	0.08
$0.07/\bar{\alpha}$	0.4951	0.0139	0.0139	0.0519	0.13
$0.06/\bar{\alpha}$	0.4919	0.0117	0.0096	0.0361	0.19
$0.05/\bar{\alpha}$	0.4898	0.0097	0.0089	0.0331	0.22
$0.04/\bar{\alpha}$	0.4862	0.0078	0.0078	0.0280	0.28
$0.03/\bar{\alpha}$	0.4825	0.0060	0.0024	0.0123	0.31

collaboration equity function values, the ω-mean collaboration inequity, and the changes in the objective value (relative to the unconstrained $E^e_{\max}$ case), as $E^e_{\max}$ was decreased for instance CDM-2c. Note that because of the type of restriction $\omega_\alpha|E^e_\alpha(x)| \le E^e_{\max}$ in (1o), we have explored values of $E^e_{\max}$ given by a constant times $(1/\bar{\alpha})$ to derive insights into the effect of problem-independent scaled values for this parameter.

Observe that as $E^e_{\max}$ decreases, the consequently enforced reduction in the weighted inequity values, and hence, in the ω-mean collaboration inequity (x^e), more than offsets the loss in the ω-mean collaboration efficiency. To comply with the more stringent equity requirement, in general, more costly surrogates were chosen for airlines that were more efficient in the unconstrained case, and vice versa for those that were relatively less efficient to balance the equity. However, the impact on the overall objective value was greatly mitigated by the reduced penalties accruing from the trade-off between the ω-mean collaboration efficiency and the ω-mean collaboration inequity. Table 7 displays the relative cost ratios $d_\alpha(x)$, $\forall\,\alpha$, as $E^e_{\max}$ varies.

We selected $E^e_{\max} = (0.07/\bar{\alpha})$ for implementation when $E^e_{\max}$ is taken as a constant, based on the results in Table 6 and the following observations. This bound impacts the objective function to a lesser extent than does $D_{\max}$. Note that for three of the four APCDM instances considered above, the unconstrained weighted airline inequities are below this threshold; hence, this imposition incurs no additional cost. Furthermore, suppose that the ω-mean collaboration efficiency achieved is midrange, i.e., $\sum_\alpha \omega_\alpha E_\alpha(x) = 0.5$, say. Then from Equations (1k) and (1l), we can assess that with $\omega_\alpha|E^e_\alpha(x)| \le E^e_{\max} = (0.07/\bar{\alpha})$ and $D_{\max} = 1.2$, and assuming that $\omega_\alpha = 1/\bar{\alpha}$, $\forall\,\alpha$, each airline α will have incurred an average collaboration cost $\sum_{f\in A_\alpha}\sum_{p\in P_f} W_f c_{fp} x_{fp}$ that lies between 8.6% to 11.4% greater than its respective individually optimized average cost, $\sum_{f\in A_\alpha} W_f c^*_f$ (i.e., from Equation (14), $d_\alpha(x) \in [1.086, 1.114]$).

Next, we considered the case when $E^e_{\max}$ is designated as a variable in the model (see (1a), (1o)). Using $\xi = \bar{\alpha}$ for our baseline value in (12b), we varied the level of ξ to ascertain the sensitivity of the obtained solution to this parameter. The results are shown in Table 8. Examining (12b) and the objective function (1a), observe that selecting $\xi = \bar{\alpha}$ would give roughly equal consideration to the efficiency and the inequity terms in the objective function, similar to the prescription (12a) when $E^e_{\max}$ is treated as a constant. Based on this observation and the reasonable compromise between equity and efficiency displayed in Table 8, we recommend using $\xi = \bar{\alpha}$.

In closing this subsection, we mention that we also experimented with the case when $E^e_{\max}$ is treated as a variable, but the objective term $\mu^e_{\max}E^e_{\max}$ reflects the situation where a decision maker seeks to minimize the maximum weighted "regret" $(-\omega_\alpha E^e_\alpha(x))$

Table 7 Relative Cost Ratios $d_\alpha(x)$ for Airlines α as $E^e_{\max}$ Varies

	$E^e_{\max}$							
Airline α	$0.1/\bar{\alpha}$	$0.09/\bar{\alpha}$	$0.08/\bar{\alpha}$	$0.07/\bar{\alpha}$	$0.06/\bar{\alpha}$	$0.05/\bar{\alpha}$	$0.04/\bar{\alpha}$	$0.03/\bar{\alpha}$
1	0.5858	0.5786	0.5676	0.5568	0.5200	0.5200	0.5200	0.4761
2	0.6883	0.6883	0.6883	0.6883	0.6799	0.6455	0.5883	0.5781
3	0.4566	0.4566	0.4566	0.4566	0.4566	0.4566	0.4566	0.4566
4	0.4671	0.4671	0.4671	0.4671	0.4612	0.4612	0.4612	0.4831
5	0.4506	0.4506	0.4506	0.4506	0.4747	0.4747	0.4747	0.4747

Table 8 Sensitivity Analysis for the Parameter ξ when $E^e_{\max}$ Is a Variable

ξ	$\bar{\alpha}-4$	$\bar{\alpha}-1$	$\bar{\alpha}$	$\bar{\alpha}+1$	$\bar{\alpha}+5$
$\sum_\alpha \omega_\alpha E_\alpha(x)$	0.5071	0.5071	0.5016	0.5016	0.5058
$\max_\alpha\{\omega_\alpha \lvert E_\alpha(x)\rvert\}$	0.0234	0.0234	0.0190	0.0190	0.0180
$\max_\alpha\{-\omega_\alpha E^e_\alpha(x)\}$	0.0191	0.0191	0.0159	0.0159	0.0172
$\max_\alpha \lvert E^e_\alpha(x)\rvert$	0.3746	0.3746	0.1867	0.1867	0.1825
x^e	0.0823	0.0823	0.0612	0.0612	0.0588

with respect to inequities; hence, $E^e_{\max}$ is given by (15) below in lieu of (1o):

$$E^e_{\max} \geq -\omega_\alpha E^e_\alpha(x), \quad \alpha = 1, \ldots, \bar{\alpha}. \tag{15}$$

In this context, we initially set $\xi = \bar{\alpha}$ in (12b) and then varied the level of ξ to determine the sensitivity of the obtained solution to this parameter, when (15) replaces (1o). As before, we focused on instance CDM-2c. The results turned out to be substantially insensitive to the value of ξ. For $\xi \in [\bar{\alpha}-4, \bar{\alpha}+4]$, we obtained identical optimal solutions, where the achieved ω-mean collaboration efficiency, the ω-mean collaboration inequity, the maximum weighted airline inequity, and the maximum regret were, respectively, 0.5071, 0.0823, 0.0234, and 0.0191.

In addition, the maximum (unweighted) inequity $\max_\alpha |E^e_\alpha(x)|$ for this solution was 0.3746. Observing the corresponding quantity ($\leq$0.1867) in Table 8 for $\xi \geq \bar{\alpha}$, we note that the maximum spread in the efficiency about the mean for the former solution is greater by a factor of two as compared with that reported in Table 8. Although the objective term $\mu_{\max} E^e_{\max}$ under either (15) or (1o) attempts to induce the same effect of reducing the spread of efficiencies about their mean, in general we observed a better control of this phenomenon when using (1o), and we therefore recommend this for implementation.

Observe that the results obtained in this last experiment might have been different, and more in line with the use of (1o), had we continued to lexicographically minimize the vector of regret values arranged in nonincreasing order. This would be in the spirit of achieving equity as delineated by Luss (1999) in general, and in the particular context of assigning arrival slots to airlines at an airport as studied by Vossen and Ball (2001). However, in the present context, such a lexicographic optimization would require a sequence of runs of the APCDM model based on constraining certain regret values and minimizing the maximum of the remaining values. Nonetheless, it might be worth investigating this for future research because this strategy considers the entire distribution of efficiency values in contrast with considering only the extremal values.

5.4.3. Sensitivity Analysis with Respect to the Parameter μ_0. Using the APCDM model instances cited in Table 5, we next examined the effect of varying the parameter μ_0 (see Equations (12a) and (12b)) on the ω-mean collaboration efficiency and the ω-mean collaboration inequity for the resulting solutions. Assuming that $E^e_{\max}$ is a specified constant as recommended in the foregoing section and noting that $\mu^e = \mu^D$ in (12a), we averaged the improvements exhibited for the two corresponding CDM measures as noted below, as μ_0 was increased from 0 to 0.3. For any value of $\mu_0 = \bar{\mu}_0$, define the *Average CDM Improvement* as

$$0.5\left[\left\{\sum_{\alpha=1}^{\bar{\alpha}} \omega_\alpha E_\alpha(x)\right\}_{\mu_0=\bar{\mu}_0} - \left\{\sum_{\alpha=1}^{\bar{\alpha}} \omega_\alpha E_\alpha(x)\right\}_{\mu_0=0}\right] + 0.5[\{x^e\}_{\mu_0=0} - \{x^e\}_{\mu_0=\bar{\mu}_0}], \tag{16}$$

where $\{\cdot\}_{\mu_0=\bar{\mu}_0}$ refers to the value of the corresponding quantity $\{\cdot\}$ at the solution obtained when $\mu_0 = \bar{\mu}_0$. Figure 4 depicts the results of this analysis, where the *Percentage CDM Improvement* is given by (16), expressed as a corresponding percentage of $0.5\{\sum_{\alpha=1}^{\bar{\alpha}} \omega_\alpha E_\alpha(x)\}_{\mu_0=0} + 0.5\{x^e\}_{\mu_0=0}$.

The intent of the CDM penalty terms in the objective function is to create a mathematical incentive to compose solutions that achieve an adequately high ω-mean collaboration efficiency and an adequately low ω-mean collaboration inequity. Figure 4 makes it evident that this goal is largely achieved at $\mu_0 = 0.1$. By increasing μ_0 further, we obtain only marginal gains in these measures, while significantly increasing the objective value (as much as 14% from the objective value when $\mu_0 = 0$), thereby substantially masking the direct costs that are incurred by the airlines. This justifies our selection of $\mu_0 = 0.1$. Based on the results in the previous subsection, for the case when $E^e_{\max}$ is a variable, we adopt a similar prescription under this strategy and recommend using (12b), along with $\mu_0 = 0.1$.

5.5. Comparison of the Alternative Conflict Constraint Formulations

In Part I of this paper we showed that the C_4 formulation yields a tighter representation of the conflict resolution constraints than the C_3 formulation,

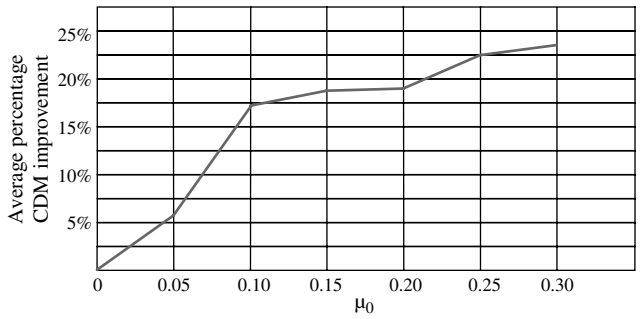

Figure 4 Sensitivity Analysis with Respect to μ_0

Table 9 C_3 **vs.** C_4 **Constraints Generated**

Label	Test set*	Edges in A	Maximum number of overlapping sets of conflicts in any sector	Maximum number of conflicts in any overlapping set	Number of C_3 constraints generated (beyond C_2)	Number of C_4 constraints generated (beyond C_2)
CFT-1, 2	1	44	3	9	59	25
CFT-3	1	283	18	20	453	207
CFT-4, 5	2	477	23	17	725	262
CFT-6	2	573	24	18	933	314
CFT-7	2	1,130	57	45	2,621	667
CFT-8	2	1,407	68	56	3,601	830
CFT-9	2	1,448	70	42	3,537	895
CFT-10	3	1,458	37	140	6,351	653
CFT-11	4	1,215	74	70	3,675	711
CFT-12	4	1,436	72	65	4,230	870

*All instances use randomized trajectory displacements, except CFT-7, which uses cylindrical SSW wind-induced displacements; and CFT-8 and CFT-12, which use N wind-induced displacements.

which, in turn, is tighter than the C_2 formulation. In the present analysis, we study the practical benefit of using these alternative formulations in terms of computational effort. Toward this end, for generating certain additional test cases based on the foregoing four primary test sets, we modified the conflict threshold probabilities from their respective nominal values to relatively small values, yielding instances having relatively denser conflict sets. The resulting test problem instances and their labels and specifications are presented in Table 9.

First, note that the C_4 formulation provides a more compact representation than the C_3 formulation in terms of the number of constraints generated, as evident from the last two columns in Table 9. Second, to better assess the usefulness of the three conflict constraint formulations, C_2, C_3, and C_4, we disabled the CPLEX program's automatic generation of various MIP cuts, and we examined the overall computational effort required using these three alternative formulations. The results of this experiment on 24 test problems are displayed in Table 10, where an additional set of 12 test cases (labeled by a suffix "a") have been generated based on the 12 instances of Table 9, by using more restricted sector capacities as shown.

The three formulations are indistinguishable in terms of computational effort for the first four instances, and for CFT-6 and CFT-1a, while the C_3 formulation exhibits a more intense effort for the remaining instances, and is of an untenable size for CFT-10, CFT-11, and CFT-12. The first four problems were such that few of the conflict resolution constraints produced any significant impact (given the allowable airspace traffic density), in that there were readily substitutable surrogates (albeit inferior in terms of cost) available to replace those that generated significant conflicts. For all of the instances, we obtained optimal solutions with zero flight cancellations, and with peak sector occupancy levels of $n_s = 5$ for the first nine cases, $n_s = 8$ for CFT-10, and $n_s = 12$ for CFT-11 and CFT-12. In the case of CFT-7 through CFT-12, observe that the greatly increased size of the constraint matrix for the C_3 formulation places a significant computational burden on the overall model.

Table 10 **Comparison of Computational Effort for Alternative Conflict Constraint Formulations**

Label(1)	$\bar{n}_s$	C_4 solution time (seconds)	Solution time ratio: C_2/C_4	Solution time ratio: C_3/C_4
CFT-1	15	0.771	1.013	1.065
CFT-2(4)	15	1.822	1.357	0.907
CFT-3	15	1.191	0.874	1.109
CFT-4	15	7.661	1.214	1.256
CFT-5(4)	15	20.088	0.486	1.080
CFT-6	15	9.053	1.246	1.086
CFT-7	15	30.253	0.949	1.468
CFT-8	15	61.338	0.747	1.362
CFT-9	15	29.752	0.907	1.854
CFT-10(4)	15	73.666	0.894	(2)
CFT-11(4)	15	274.504	1.147	(2)
CFT-12(4)	15	295.240	1.211	(2)
CFT-1a	3	0.600	0.950	0.900
CFT-2a(4)	3	9.105	0.925	1.056
CFT-3a	3	0.971	0.856	0.959
CFT-4a	3	66.856	0.797	0.774
CFT-5a(4)	3	69.420	0.947	1.073
CFT-6a	3	60.206	0.988	1.192
CFT-7a	3	53.056	1.338	1.575
CFT-8a	3	160.661	1.174	1.833
CFT-9a	3	230.690	0.848	1.743
CFT-10a(4)	7	772.506	0.937	(2)
CFT-11a(3)(4)	9	276.407	2.287	(2)
CFT-12a(3)(4)	9	207.888	1.450	(2)

(1) "a" label indicates an identical data set with sector capacity constrained as shown in the next column.

(2) Resulting constraint matrix size exceeded the computer's memory capacity.

(3) Solution obtained using an LP/IP gap tolerance of 5% (due to computer's memory limitations).

(4) The instances were examined using the prescribed parameter values as discussed in §§5.2–5.4.

The test cases having more complex conflicting situations, or relatively restricted values of sector capacities $\bar{n}_s$, prompted a more careful selection of available surrogates, thereby increasing the impact of surrogates made unavailable due to the conflict resolution constraints. The net effect on the LP relaxation was to fractionate the variables representing the selection of key surrogate alternatives. Consequently, the potential effectiveness of the tighter C_4 formulation became more evident for the instances CFT-11, CFT-12, CFT-7a, CFT-8a, CFT-10a, CFT-11a, and CFT-12a, where C_4 resulted in an average savings of 29% of the effort required by the C_2 formulation. On the other hand, for the instances CFT-4a, CFT-9a, and CFT-10a, C_2 resulted in a savings of 25%, 18%, and 7%, respectively, of the effort required by C_4. In these instances, the trade-off between the tightness afforded by C_4 in the vicinity of the optimum and its overall size gave C_2 an advantage because of the nature of the problems. Overall, we suggest using C_4 because its additional size does not significantly burden the model, while its relative tightness is advantageous when solving problems for which competitive solutions involve more complex conflicting situations.

Finally, we comment on the use of the special-purpose cuts as afforded by (1j) of C_4, as opposed to implementing CPLEX's default cut generation strategy in concert with C_2. It turned out to be far more beneficial to employ C_4, and thereby exploit the inherent special structures of the problem, rather than rely on general-purpose cuts. For example, on the most challenging instances CFT-10a, CFT-11a, and CFT-12a, where the formulation C_2, respectively, consumed 723.67, 623.14, and 301.44 CPU seconds as reported in Table 10, when the CPLEX default cut option was turned on, these times worsened to 728.83, 682.09, and 523.42 CPU seconds, respectively. Note that the CPLEX default cut option used in concert with C_4 also worsened the performance, consuming 794.63, 646.26, and over 1,200 CPU seconds, respectively, in contrast with the respective times of 772.51, 276.41, and 207.89 CPU seconds required by C_4. On the other hand, for the instances CFT-11 and CFT-12, the CPLEX default cut option enabled C_2 to solve these problems with about the same effort as that required by C_4, while two of the relatively simpler instances CFT-8a and CFT-9a were solved somewhat faster using this CPLEX option (62.16 and 139.78 CPU seconds, versus the 188.58 and 195.57 CPU seconds consumed by C_2). In general, because of the erratic performance of the general-purpose cuts and, in particular, their deleterious effort on relatively more complex problem instances, we recommend the sole use of special-purpose cuts in the context of the APCDM model.

6. Summary and Conclusions

In this paper, we have provided a detailed description of implementing the APCDM model, including the development of a comprehensive cost model and a study for prescribing a set of appropriate parameter values for the overall model. We have also presented results using the ETMS data, performed sensitivity analyses with respect to various pertinent model parameters to provide insights into the model and prompt reasonable choices of the different parameters, and studied the utility of the valid inequalities that were proposed to enhance the model representation. A summary of our findings is given below.

1. A probabilistic analysis (with any satisfactory threshold probability) of potential aircraft conflicts, rather than a deterministic analysis, provides a significantly more robust prediction of potential aircraft conflicts.

2. The APCDM model is relatively insensitive to moderate changes in the conflict risk threshold probabilities.

3. When the maximum allowable ratio for each airline of its cost pertaining to the set of surrogates selected through the CDM process to that of its individually optimized set of surrogates is set at a value of $D_{\max} = 1.20$, this parameter can be expected to substantially influence the selection of surrogates at optimality.

4. We observed a better effect of reducing the spread of efficiencies about their mean, in general, when using (1o) instead of (15).

5. The CDM penalty terms in the objective function, parameterized at a level of $\mu_0 = 0.1$, create a mathematical incentive to compose solutions that achieve a high ω-mean collaboration efficiency coupled with a low ω-mean collaboration inequity.

6. The potential effectiveness of the star-subgraph convex hull-representation-based valid inequalities (C_4 formulation) was evident in our experiments, and we recommend its use given the relative tightness they provide without imposing a significant computational burden.

7. The use of the foregoing specialized cuts in place of general-purpose cuts yields consistent and improved computational efficiency in the context of the APCDM model.

The APCDM model can be used in practice for tactical decision-making purposes as well as for strategic planning investigations. Potential tactical contexts include decisions pertaining to air-traffic control diversions and delays during spacecraft launch operations or during severe weather conditions; military theater operations (such as damage assessment, search and rescue, ground support, and counteroperations), and the generation of alternative flight plans (via the feedback outer loop in Figure 1). Strategic

applications include air-traffic control policy evaluations (e.g., revising aircraft separation standards, or analyzing the free-flight paradigm where aircraft are permitted to take wind-optimized routes); Homeland Defense contingency planning; spaceport location planning; fixed alternative or dynamic airspace resectorization strategies; military theater air campaign planning; and the construction of a priori plans to respond to various disruption scenarios (e.g., augmenting the FAA's *National Playbook* 2003). Furthermore, the model can be used to study the incorporation of the Small Aircraft Transportation Systems (SATS) into the NAS (see Trani et al. 2004). As work in this area progresses, we hope to report on results emanating from many such diverse studies conducted from this wide range of potential applications of the APCDM model.

Acknowledgments

This research has been supported by the National Science Foundation under Grant Number 0094462 and by a Federal Aviation Administration Grant under the National Center of Excellence for Aviation Operations Research (NEXTOR) program. Thanks are also due to Dr. Hojong Baik for his assistance in data acquisition and computations, and to two anonymous referees and an associate editor for their detailed comments, which significantly helped improve the perspective and organization of this paper. The views expressed in this article are those of the author(s) and do not reflect the official policy or position of the U.S. Air Force, Department of Defense, or the U.S. Government.

References

Ball, M., R. Hoffman, W. Hall, A. Muharremoglu. 1998. Collaborative decision making in air traffic management: A preliminary assessment. NEXTOR Technical Report RR-99-3, University of Maryland, College Park, MD.

Ball, M., R. Hoffman, D. Knorr, J. Wetherly, M. Wambganss. 2000. Assessing the benefits of collaborative decision making in air traffic management. *Proc. 3rd USA/Europe Air Traffic Management R&D Seminar*, Napoli, Italy, June 13–16.

Beatty, R., R. Hsu, L. Berry, J. Rome. 1998. Preliminary evaluation of flight delay propagation through an airline schedule. *Proc. 2nd USA/Europe Air Traffic Management R&D Seminar*, Orlando, FL, December 1–4.

Chang, K., K. Howard, R. Oiesen, L. Shisler, M. Tanio, M. Wambsganss. 2001. Enhancements to the FAA ground-delay program under collaborative decision making. *Interfaces* **31** 57–76.

Crawley, J. 2001. U.S. lays out 10-year plan to reduce flight delays. *Reuters News Service* (June 16).

Erzberger, H., R. Paielli, D. Isaacson, M. Eschow. 1997. Conflict detection and resolution in the presence of prediction error. *Proc. 1st USA/Europe Air Traffic Management R&D Seminar*, Napoli, Italy, June 17–20.

Eurocontrol Experimental Center. 1998. *Base of Aircraft Data* (*BADA*) *3.0 User Manual*. Eurocontrol Agency, Brussels, Belgium.

Federal Aviation Administration (FAA). 2002. Department of Transportation Federal Aviation Administration internal vacancy announcement: #AEA-AAT-02-209-66252. http://jobs.faa.gov.

Federal Aviation Administration. 2003. National severe weather playbook. http://www.fly.faa.gov/playbook.html.

Hansen, M., D. Gillen, R. Djafarian-Tehrani. 2001. Aviation infrastructure performance and airline cost: A statistical cost estimation approach. *Transportation Sci.* **37** 1–23.

Luss, H. 1999. On equitable resource allocation problems: A lexicographical minimax approach. *Oper. Res.* **47**(3) 461–478.

Nolan, M. 1999. *Fundamentals of Air Traffic Control*. Brooks/Cole, Pacific Grove, CA.

Sherali, H. D., J. C. Smith. 2003. A polyhedral study of the generalized vertex packing problem. *Math. Programming*. Forthcoming.

Sherali, H. D., J. C. Smith, A. A. Trani. 2002. An airspace planning model for selecting flight plans under workload, safety, and equity considerations. *Transportation Sci.* **36** 378–397.

Sherali, H. D., R. W. Staats, A. A. Trani. 2003. An airspace planning and collaborative decision making model: Part I—Probabilistic conflicts, workload, and equity considerations. *Transportation Sci.* **37** 434–456.

Staats, R. W. 2003. An airspace planning and collaborative decision making model under safety, workload, and equity considerations. Ph.D. dissertation, Virginia Polytechnic Institute and State University, Blacksburg, VA.

Trani, A. A., H. Baik, H. S. Swingle, S. Ashiabor. 2004. Integrated model to study the small aircraft transportation system. *Transportation Res. Record* **1850** 1–10.

Vossen, T., M. Ball. 2001. Optimization and mediated bartering models for ground delay programs. Robert H. Smith School of Business Working paper, University of Maryland, College Park, MD.