

Robust Airline Scheduling under Block Time Uncertainty

Milind Sohoni

Indian School of Business, Hyderabad, India, milind_sohoni@isb.edu

Yu-Ching Lee

University of Illinois at Urbana-Champaign, Urbana, IL, ylee77@uiuc.edu

Diego Klabjan

Northwestern University, Evanston, IL, d-klabjan@northwestern.edu

Airline schedule development continues to remain one of the most challenging planning activity for any airline. An airline schedule comprises of a list of flights and specifies the origin, destination, scheduled departure, and arrival time of each flight in the airline's network. A critical component of the schedule development activity is the choice of flight block-times, which depend on several factors. Many airlines decide schedule block-times based on fixed percentiles of block-time distributions built from historical data, however, such techniques have not resulted in significantly improved on-time performance of the schedule during operations. Thus, from a passenger's perspective, the service level guarantee of an airline's network continues to be low. We first define two service level metrics for an airline schedule. The first one is similar to the on-time performance measure of the U.S. Department of Transportation and we define it as the flight service level. The second metric, called the network service level, is geared towards completion of passenger itineraries. We then develop a stochastic integer programming formulation that optimally perturbs a given schedule to maximize expected profit while ensuring the two service levels. We also develop a variant of this model that maximizes service levels while achieving desired network profitability. To solve these models we develop an efficient algorithm that guarantees optimality. Through extensive computational experiments, using real-world data, we demonstrate that our models and algorithms are efficient and achieve the desired trade-off between service level and profitability.

Key words: robust scheduling, stochastic optimization, airline planning

1. Introduction

In a recent article (Associated Press 2007) The Associated Press reported that the U.S. airline industry's on-time performance (OTP) through the first eleven months of 2007, was the second

worst on record. According to the U.S. Department of Transportation, a flight is delayed if it arrives at its destination gate 15 minutes or more after its scheduled arrival time. Even in the previous year, i.e., 2006, statistics showed that there were 823,030 arrival delays out of a total of 3,805,313 commercial flights operated by all the major U.S. carriers (Bureau of Transportation Statistics 2009). Flight delays and cancelations have been attributed to several causes some of which include weather conditions, airport congestion, national air-space congestion, aircraft maintenance related issues, and more recently airline security related services. Consequently, such delays lower service reliability and adversely affect a commuter's travel experience.

While some of the causes of delays, such as weather conditions, are beyond the control of the airlines, previous research shows that some causes of delays are attributable to the network and schedule design decisions of an airline. For example, while an airline develops its hub-and-spoke network, it typically does not account for the congestion externality imposed on other carriers operating out of the same hub stations. In a recent paper, Mayer and Sinai (2003a) empirically demonstrate that the gains from *hubbing* activities offset the costs incurred by flight delays and congestions. In such cases, congestion pricing at certain capacity constrained airports, may be a solution to elevate the problem. In a companion paper, Mayer and Sinai (2003b) also hypothesize that wage cost minimization and aircraft utilization maximization result in airlines flying with very tight schedules. Such objectives are typical in most airline planning systems, which are designed to achieve cost efficient resource utilization. Schedule planning models do not address the following two important issues. First, they do not include passenger-centric service reliability measures in the schedule development process. Second, the schedules ignore block-time uncertainty (variance) and hence fail to capture robustness measures. In this paper we address these issues by developing schedule planning models that incorporate both, passenger centric metrics and block-time uncertainty, in the planning process.

Airline schedule development continues to remain one of the most challenging planning activity for any airline. An airline schedule comprises of a list of flights and specifies the origin, destination, scheduled departure, and arrival time of each flight in the airline's network. A critical component of

the schedule development activity is the choice of flight block-times. A flight block-time is defined as the total elapsed time between the time an aircraft pushes back from its departure gate and arrives at its destination gate. The block-time comprises of several components including taxi-out time, enroute time, and taxi-in time. Each of these components is subject to different causes of delay and the total block-time delay is the sum of all individual component delays. Since airline schedules must be published well in advance of the actual day of operation, block-times, for all the flights in the schedule, are typically decided using historical information of similar flights operated in the past. The Department of Transportation OTP metric is computed against these published flight block-times. Most airline operations are compared based on their OTP rankings and hence airlines perceive their OTP as an important operational measure of their schedule reliability. However, research indicates that airlines fail to adequately adjust block-times and typically do not incorporate uncertainty in their planned schedules. Since most planned resource costs, such as aircraft and crew utilization costs, depend on the cumulative hours in a schedule, airlines face a key trade-off decision between adjusting (increasing) flight block-times to improve schedule reliability and incurring additional planned costs. Using data made available by the Bureau of Transportation, Deshpande and Arikan (2009) argue that airlines systematically “under-schedule” flights, i.e., the amount of block-time allocated for a flight is less than the average block-time expected for the flight. Conversations with planners at a large U.S. carrier suggested that airlines do not judiciously allocate block-times to scheduled flights to balance costs versus operational benefits. Typically, planners use ad-hoc techniques to either lower or raise block-times across the entire flight network in the hope of increasing OTP. Results in Deshpande and Arikan (2009) also corroborate these findings and indicate that airlines do not maintain consistent service levels by adjusting their schedules based on the time of the day, origin airport congestion, and destination airport congestion.

Planning for uncertainty in the schedule building process becomes necessary not just to improve OTP rankings but also to improve passenger service levels. As stated earlier, the goal of this paper is to develop a robust optimization approach to schedule planning by specifically incorporating passenger centric goals and block-time uncertainty in the planning models. The key trade-off in

such a process is between higher service levels achieved through increasing (and better allocation of) flight block-times and higher planned costs (i.e., lower planned profits). In this paper we develop a model that re-times (perturbs) a proposed flight schedule by considering block-time probability distributions. First, we explicitly define notions of passenger and network service levels. Then we develop a model that maximizes the expected profit while guaranteeing minimum service levels. This model allows imposition of minimum service levels. Second, we develop a variant of this model that maximizes service levels while achieving required profitability. While the optimization models are complex, we also develop computational procedures, based on cut generation techniques, to efficiently solve these models. To this end, this paper also has a methodological contribution to the development of computationally efficient procedures. We provide extensive computational experiments, using real airline data from a large U.S. carrier, that validate our model and demonstrate potentially large operational gains for an airline. Overall network reliability is also improved.

The contributions of this paper are at several levels. First, to the best of our knowledge, this paper is an initial attempt at developing a comprehensive and holistic model that includes block-time uncertainty in developing robust schedules. Second, through chance constraints, we explicitly model block-time distributions allowing us to incorporate operational uncertainty in the schedule planning process. This makes the resultant schedule robust. We also incorporate network service levels, which probabilistically model passenger connections. Third, we propose a new cut generation algorithm to solve these stochastic binary integer programming models and establish its convergence. The analysis is non-trivial since the feasible region of the original problem is non-convex and first a linearization is required. Upon linearization, the resulting (modified) model is infinite dimensional with infinitely many constraints. Thus, our algorithmic procedure and optimal convergence result generalizes previously established convergence results for (1) semi-infinite linear programs with finitely many variables but infinitely many constraints, and (2) infinite dimensional problems with finitely many constraints and infinite number of variables.

Overall, this research is in line with the growing literature on linking operational variability (and hence costs) to planning models. For example, research in robust fleetings (Rosenberger et al.

2004); robust aircraft routing (Lan et al. 2003), robust crew scheduling (Shebalov and Klabjan 2006), and the robust approach to passengers rerouting in disruption management (Karow 2003) show this emerging trend. Another growing area is the development of simulation systems of airline operations, e.g., SimAir by Rosenberger et al. (2002) and MEANS by Bly et al. (2003). These systems play a crucial role in evaluating and comparing the performance of different schedules. This paper also contributes to several techniques developed in the airline schedule planning literature. In general, airlines, though plagued with low profitability margins, airspace and airport congestion, and high capital and operating costs are heavy users of mathematical optimization techniques (Dobson and Lederer 1993, Lohatepanont and Barnhart 2004, Barnhart et al. 2003). Barnhart and Cohn (2004) and Klabjan (2005) provide an extensive review of OR models used in airline schedule planning. There is other literature in the domain of stochastic scheduling that is also related to our work (see Portougal and Trietsch 2001). However, existing literature in stochastic scheduling ignores the need to achieve high customer service level.

The rest of the paper is organized as follows. First in § 2 we develop the two optimization models for schedule perturbation. Next, in § 3 we discuss issues related to the computational tractability of these models and develop the solution methodology and optimal algorithms. We provide extensive computational experiments in § 4. Finally, in § 5 we conclude the paper. Additionally, we provide a complete set of results of all the other computational experiments in an online appendix (see Sohoni et al. 2008).

2. Model Description

As discussed earlier, our goal is to develop a model to perturb the incumbent flight schedule to improve the service levels provided to the end consumers. Perturbing a flight schedule implies adjusting the scheduled departure times of flights¹ in the network within an allowable time window.

Soon after determining the flight schedule, the airlines determine capacity assignments (fleeting) and assign generic aircraft to routes. The latter, in the literature, is referred to as the aircraft routing

¹ Throughout this paper we use the terms “flight” and “leg” interchangeably.

problem. It is after these processes that we consider the issue of schedule re-timing (perturbation) to fine tune block-times and improve robustness of the schedule with respect to the service level metrics defined later. While perturbing the incumbent schedule, however, we must guarantee that the resulting schedule continues to remain feasible with respect to the *aircraft turns* built of the incumbent schedule. Every flight in the incumbent schedule is assigned to exactly one aircraft. An *aircraft turn* is essentially a pair of consecutive flights flown by the same aircraft. We assume that the set of turns associated with the incumbent schedule is known a priori.

A passenger travel plan, or *itinerary*, may comprise of multiple flight legs. Broadly defined, a *fare class* is the price an airline charges to book a passenger in a particular booking class. Airline seats are divided into several *booking classes*. Next, we define the important modeling notation and parameters.

N : The set of all flights (legs) in the airlines flight network,

B : the total available planned budget (depends on the total block-time across all flights),

O : the set of all passenger itineraries,

T : the complete set of aircraft turns,

F : the set of all fare classes,

α_i : the origin station of flight i ,

β_i : the destination station of flight i ,

m_{ij} : minimum passenger connection time between two flights i and j ,

t_{ij} : minimum turn-time between flights i and j on the aircraft rotation,

D_{of} : expected demand for itinerary o and fare class f ,

$[l_i, u_i]$: the allowable departure time-window for flight i ,

c_i : the per time unit cost incurred for flight i , which includes unit costs corresponding to crew pay and aircraft utilization,

- \mathcal{B}_{if} : booking limit for fare class f on flight i ,
- r_{of} : the average fare of itinerary o and fare class f ,
- d_i^s : the previously scheduled departure time of flight i in the incumbent schedule,
- e_i : the penalty for deviating from the preferred departure time of flight i , and
- δ : the Department of Transportation OTP measure for flight delay
(typically 15 minutes after the scheduled arrival time).

Next, we define the decision variables of the model:

- d_i : the published departure time of flight i ,
- a_i : the published arrival time of flight i ,
- X_{of} : demand of itinerary o and fare-class f satisfied, and
- z_{ij} : binary variable indicating if the passenger connection between flights i and j is feasible.

We define d and a to be the set of departure and arrival times respectively. The only random variables in the model are the block-times and are denoted by Y_{it} where t represents the departure time of flight i . We assume that these are continuous random variables. The relation between a flight's departure time, arrival time, and the corresponding block-time is as follows: $A_i = d_i + Y_{id_i}$, where A_i is the actual random arrival time of flight i . The probability density function of a flight's block-time is represented by $p_i(\cdot, t)$ since it might depend on the departure time t . The cumulative density function is assumed to have finite support $[\delta_l^i, \delta_u^i]$. To reduce the complexity of our computational experiments we assume the following.

ASSUMPTION 1. *The expected demand D_{of} for an itinerary does not vary significantly for reasonable deviations in departure time.*

Given that we disallow large perturbations of the departure time by controlling the time window $[l_i, u_i]$ for every flight i in the network, it is reasonable to assume the following:

ASSUMPTION 2. For each flight i we require that, $p_i(\cdot, t) = p_i(\cdot)$, i.e., the pdf of the block-time distribution does not depend on the departure time.

A flight j is said to follow-on flight i if passengers of flight i can connect to flight j . The set of all passenger connections for flight i depends on the arrival time of flight i and departure times of possible connecting flights. We define the connection set for flight i as follows.

DEFINITION 1. The connection set for flight i , with respect to the incumbent schedule, is defined as

$$C_i(d, a) = \{j \in N : d_j - a_i \geq m_{ij} \ \& \ \beta_i = \alpha_j\}. \quad (1)$$

Building on the definition of $C_i(d, a)$ we define a modified connection set \bar{C}_i , which denotes the largest set of possible connections for flight i under any departure and arrival time adjustment. For example, \bar{C}_i can be the set of all flights originating at station α_i , or we can further refine the set as

$$\begin{aligned} \bar{C}_i &= \{j \in N : \beta_j = \alpha_i \text{ and can connect to } i \text{ regardless of re-timing}\} \\ &= \{j \in N : \beta_j = \alpha_i, u_j - (l_i + \delta_i^i) \geq m_{ij}\}. \end{aligned} \quad (2)$$

The advantage of using set \bar{C}_i instead of the original connection set $C_i(d, a)$ is that, for any flight i the latter set is non-stationary, i.e., as the departure time of flight i changes, the flights in the set may change. Thus it depends on the decision variables. As we show later, this poses a modeling and optimization challenge since we cannot guarantee a convex feasible region.

We now define the Service Level, SL_i , of any flight $i \in N$.

DEFINITION 2. Service level SL_i is the probability that passengers from flight i can connect to any follow-on flight included in the set $C_i(d, a)$, i.e.,

$$SL_i = \Pr[A_i + m_{ij} \leq d_j \text{ for every } j \in C_i(d, a)]. \quad (3)$$

Observe that from definition 2 it follows that $SL_i = \Pr[Y_{id_i} \leq \min_{j \in C_i(d, a)} \{d_j - d_i - m_{ij}\}]$. The Network Service Level (NSL) is defined as follows.

DEFINITION 3. The NSL is defined as the minimum service level across all the flights in the airline's network, i.e.,

$$NSL = \min_i SL_i. \quad (4)$$

Finally, the Flight Service Level (FSL), also referred to as the OTP, is defined as follows.

DEFINITION 4. The FSL is the probability that a particular flight is not delayed based on the Department of Transportation acceptable arrival delay measure δ , i.e.,

$$FSL_i = \Pr[Y_{i,d_i} \leq a_i - d_i + \delta]. \quad (5)$$

Lastly, for notational convenience, we denote the fact that flight j follows flight i in an itinerary $o \in O$ by $i \rightarrow j$. Next, we describe the two optimization models.

2.1. Maximizing Operational Profits

We first consider the case when an airline must maintain a minimum FSL , γ_f , over all flights in the network and simultaneously guarantee a minimum NSL of γ_n . The profit maximizing model (PMM) reads:

$$\text{(PMM) } \max \sum_{o,f} r_{of} X_{of} - \sum_{i \in N} e_i |d_i - d_i^s| - \sum_{i \in N} c_i (a_i - d_i) \quad (6)$$

$$\Pr[Y_{id_i} \leq d_j - d_i - m_{ij}] \geq \gamma_n \quad i \in N, j \in C_i(d, a) \quad (7)$$

$$\Pr[Y_{i,d_i} \leq a_i - d_i + \delta] \geq \gamma_f \quad i \in N \quad (8)$$

$$\sum_{i \in N} c_i (a_i - d_i) \leq B \quad (9)$$

$$X_{of} \leq D_{of} \quad o \in O, f \in F \quad (10)$$

$$\sum_{o \in O, i \in o} X_{of} \leq \mathcal{B}_{if} \quad i \in N, f \in F \quad (11)$$

$$\sum_{o,f,j \in o, i \rightarrow j} X_{of} \leq \bar{K}_{ij} z_{ij} \quad i \in N, j \in \bar{C}_i \quad (12)$$

$$d_j - a_i \geq m_{ij} z_{ij} - K(1 - z_{ij}) \quad i \in N, j \in \bar{C}_i \quad (13)$$

$$d_j - a_i - t_{ij} \geq 0 \quad (i, j) \in T \quad (14)$$

$$l_i \leq d_i \leq u_i \quad i \in N \quad (15)$$

$$z_{ij} \in \{0, 1\}, d, a \text{ unrestricted.} \quad (16)$$

The first term in the objective function (6) corresponds to the net revenue due to satisfied itinerary demand, the second term is the net penalty due to deviation from preferred departure time (departure time specified in the incumbent schedule), and the third term represents the total operational cost. Constraint (7) ensures that the minimum NSL is at least as large as the desired value of γ_n . It is not difficult to observe that $NSL \geq \gamma_n$ if and only if constraint (7) is satisfied. Constraint (8) guarantees that the minimal FSL is at least γ_f . Constraint (9) restricts the total network operating cost incurred and constraint (10) restricts the fare-class itinerary demand to the maximum allowable. Since every flight i within an itinerary o can carry at most \mathcal{B}_{if} of a particular fare-class f , constraint (11) ensures that the booking limit constraint on each flight is satisfied. Constraints (12) and (13) ensure that we only account for those itineraries whose flight sequence is legal with respect to the minimum passenger connection time. Here $\bar{K}_{ij} = \sum_f \mathcal{B}_{if} + \sum_f \mathcal{B}_{jf}$. The constant K is the length of the time horizon, i.e., $K = \max_{i \in N} u_i - \min_{i \in N} l_i + \max_{i \in N} \delta_u^i$. Constraint (14) guarantees that the pre-determined aircraft turns are preserved and hence the aircraft routing solution always remains feasible. Finally, constraint (15) bounds the departure time adjustment for every flight and the constraint (16) restrict the choice of z_{ij} to be binary.

In § 3, we discuss issues regarding the computational tractability of the optimization model PMM . One peculiarity of PMM is immediately observable; the constraint set in (7) depends on the decision variables.

2.2. Maximizing Service Level

An alternate goal could be to maximize the service level across the entire flight network. However, the airline may only be willing to do so provided it maintains minimum operational profitability. In this case the optimization model differs from the PMM model described earlier, i.e., γ_f and γ_n are no longer parameters but are decision variables. Furthermore, the profit objective in PMM is now a constraint. We impose that the minimum operational profit must be at least U_o units. The service level maximizing model (SLMM) reads:

$$(\text{SLMM}) \max w_f \gamma_f + w_n \gamma_n \tag{17}$$

$$\Pr[Y_{id_i} \leq d_j - d_i - m_{ij}] - \gamma_n \geq 0 \quad i \in N, j \in C_i(d, a) \quad (18)$$

$$\Pr[Y_{i,d_i} \leq a_i - d_i + \delta] - \gamma_f \geq 0 \quad i \in N \quad (19)$$

$$\sum_{i \in N} c_i(a_i - d_i) \leq B \quad (20)$$

$$X_{of} \leq D_{of} \quad o \in O, f \in F \quad (21)$$

$$\sum_{o \in O, i \in o} X_{of} \leq B_{if} \quad i \in N, f \in F \quad (22)$$

$$\sum_{o, f, j \in o, i \rightarrow j} X_{of} \leq \bar{K}_{ij} z_{ij} \quad i \in N, j \in \bar{C}_i \quad (23)$$

$$d_j - a_i \geq m_{ij} z_{ij} - K(1 - z_{ij}) \quad i \in N, j \in \bar{C}_i \quad (24)$$

$$d_j - a_i - t_{ij} \geq 0 \quad (i, j) \in T \quad (25)$$

$$\sum_{o, f} r_{of} X_{of} - \sum_{i \in N} e_i |d_i - d_i^s| - \sum_{i \in N} c_i(a_i - d_i) \geq U_o \quad (26)$$

$$l_i \leq d_i \leq u_i \quad i \in N \quad (27)$$

$$z_{ij} \in \{0, 1\}, d, a \text{ unrestricted.} \quad (28)$$

The objective function (17) is a weighted sum of the minimal *NSL* and *FSL* quantities where w_f and w_n are the weights corresponding to the *FSL* and *NSL*, respectively. All the other constraints are similar to those described in *PMM*. The only additional constraint is (26) which ensures that any solution makes an expected operational profit of at least U_o .

3. Solution Methodology

In this section we discuss issues regarding computational complexity and tractability of the models discussed in § 2. More importantly, we exhibit two algorithms for solving *PMM* and *SLMM*. In the model *PMM* constraints (7) and (8) are non-linear. This makes the model difficult to solve computationally. Similarly, in model *SLMM* constraints (18) and (19) are non-linear. Additionally, objective function (6) and constraint (26) contain the absolute value function, however, it is straightforward to linearize these terms. A technical assumption regarding the block-time distribution allows us to simplify the model and reduce its computational complexity.

ASSUMPTION 3. *The block-time distributions are log-concave and stationary with respect to the departure time.*

Through extensive empirical studies using Bureau of Transportation Statistics data, Deshpande and Arikan (2009) estimate the best distribution fit for observed truncated block-times across several US airlines. Specifically, they use log-Normal and log-Laplace distributions. While the log-Normal distribution provides a reasonable fit, they show that the log-Laplace distribution is better. It is noteworthy that both of these cumulative distribution functions are log-concave (Bagnoli and Bergstrom 2005) and thus satisfy Assumption 3. The Laplace distribution is defined by two parameters: γ , a location parameter, and b , a scale parameter where the mean equals to γ and the variance is $2b^2$. The probability density function of the Laplace(γ, b) distribution is $f(x|\gamma, b) = \frac{1}{2b} \exp\left(-\frac{|x-\gamma|}{b}\right)$. Assumption 3 allows us to simplify the complicating chance constraints (8) and (19) into convex constraints. Given that we assume the block-time distribution is independent of the departure time we drop the departure time subscript, i.e., $Y_{id_i} = Y_i$. Constraints (7) and (8) are transformed as follows.

$$\log(\Pr[Y_i \leq d_j - d_i - m_{ij}]) \geq \log \gamma_n \quad i \in N, j \in C_i(d, a) \quad (29)$$

$$\log(\Pr[Y_i \leq a_i - d_i + \delta]) \geq \log \gamma_f \quad i \in N. \quad (30)$$

It is known that the feasible set of constraint (30) is convex due to log-concavity (see, e.g., Birge and Louveaux 1997). Unfortunately, constraints in (29) are not convex since their index depends on d and a . This fact poses a significant algorithmic and computational challenge. To devise an efficient solution strategy we first develop a linear approximation scheme to these constraints in § 3.1. The resulting mixed-integer model has an infinite number of variables and constraints. We then describe a cut generation algorithm that generates these linear constraints as needed and builds an optimal solution to the models.

3.1. Model Reformulation

Our goal in this section is to develop a linear formulation to the two models, *PMM* and *SLMM*. Recollect that the *NSL* constraints given by equation (29) are non-convex. To circumvent this issue, we construct a linear approximation for the *NSL* constraint over a stationary set of linear functions as follows. The added advantage of doing so is that the reformulation allows us to develop

an algorithm, similar to the Bender's cut generation algorithm (Birge and Louveaux 1997), to solve the model.

Recollect that the distributions have a finite support $\mathcal{K}_i = [\delta_l^i, \delta_u^i]$. Now, for every flight i , we define a function $g_i(x)$ as

$$g_i(x) = \log \Pr[Y_i \leq x], \quad x \in \mathcal{K}_i. \quad (31)$$

Since Y_i is log-concave, $g_i(x)$ is concave, see, e.g. Birge and Louveaux (1997). To build an outer linear approximation to equation (31), we consider a set of linear functions, U_{ik} , defined over interval \mathcal{K}_i . We show the form of these linear functions later. For now, using these linear functions we rewrite $g_i(x)$ as follows (this is a known fact in convex analysis):

$$g_i(x) = \min_{k \in \mathcal{K}_i} U_{ik}(x). \quad (32)$$

Using equation (32) we now reformulate the *NSL* constraint as

$$g_i(d_j - d_i - m_{ij}) \geq \log \gamma_n \quad i \in N, \quad j \in C_i(d, a). \quad (33)$$

The above equation can be rewritten as

$$z_{ij} g_i(d_j - d_i - m_{ij}) \geq \log \gamma_n \quad i \in N, \quad j \in \bar{C}_i, \quad (34)$$

where \bar{C}_i is defined by equation (2). Observe that $\log \gamma_n \leq 0$ and thus inequality (34) holds if $z_{ij} = 0$. If $z_{ij} = 1$, then $j \in C_i(d, a)$ and thus $g_i(d_j - d_i - m_{ij}) \geq \log \gamma_n$ must hold, which is guaranteed by constraint (34). Thus, constraint (29) is equivalent to

$$z_{ij} \min_{k \in \mathcal{K}_i} U_{ik}(d_j - d_i - m_{ij}) \geq \log \gamma_n \quad i \in N, \quad j \in \bar{C}_i. \quad (35)$$

It is noteworthy that in (35) if $d_j - d_i - m_{ij} \leq 0$, then $z_{ij} = 0$ and hence we need not worry about negative arguments, i.e., we restrict our attention to positive values only.

We now characterize the functions $U_{ik}(x)$. Given the probability density function $p_i(\cdot)$ for block-time Y_i , we can write these functions as

$$U_{ik}(x) = \frac{p_i(k)}{\int_0^k p_i(t) dt} (x - k) + \log \int_0^k p_i(t) dt. \quad (36)$$

To this end, notice that $U_{ik}(\delta_l^i - m_{ij}) < 0$ and $U_{ik}(\delta_l^i - m_{ij}) \leq U_{ik}(x)$ for all $x \geq \delta_l^i - m_{ij}$ (see Figure 1). This is the tangent of $g_i(x)$ at the point $k \in \mathcal{K}_i$. It is known that a concave function is the minimum of its tangents and thus equation (32) holds. We still need to linearize constraints (35).

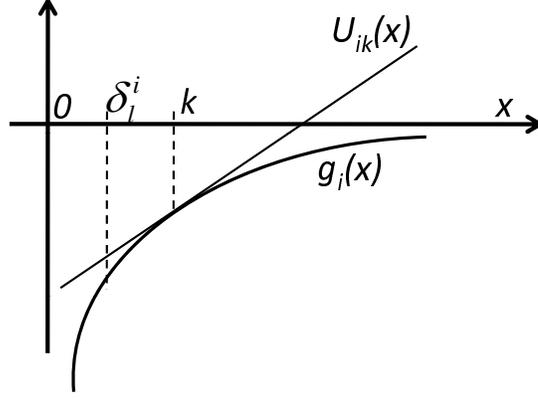


Figure 1 Linearization of constraints (35).

We now define additional continuous decision variables, s_{ijk} , for all $i \in N$, $j \in \bar{C}_i$, and $k \in \mathcal{K}_i$. Constraint (35) can then be replaced by the following set of linear constraints:

$$s_{ijk} \geq \log \gamma_n \quad i \in N, j \in \bar{C}_i, k \in \mathcal{K}_i \quad (37)$$

$$z_{ij} U_{ik}(\delta_l^i - m_{ij}) \leq s_{ijk} \leq 0 \quad i \in N, j \in \bar{C}_i, k \in \mathcal{K}_i \quad (38)$$

$$(1 - z_{ij}) U_{ik}(\delta_l^i - m_{ij}) + s_{ijk} \leq U_{ik}(d_j - d_i - m_{ij}) \quad i \in N, j \in \bar{C}_i, k \in \mathcal{K}_i. \quad (39)$$

If $z_{ij} = 0$, then (38) implies that $s_{ijk} = 0$ and thus (37) holds. In this case, (39) also holds since $U_{ik}(\delta_l^i - m_{ij}) \leq U_{ik}(d_j - d_i - m_{ij})$. On the other hand, if $z_{ij} = 1$, then we can assume that $s_{ijk} = \min\{0, U_{ik}(d_j - d_i - m_{ij})\}$ and thus (37) holds if and only if $U_{ik}(d_j - d_i - m_{ij}) \geq \log \gamma_n$.

Similarly, the *FSL* constraint given by equation (30) is equivalent to

$$\min_{k \in \mathcal{K}_i} U_{ik}(a_i - d_i + \delta) \geq \log \gamma_f \quad i \in N. \quad (40)$$

It is clear that (40) is equivalent to

$$U_{ik}(a_i - d_i + \delta) \geq \log \gamma_f \quad i \in N, k \in \mathcal{K}_i. \quad (41)$$

It is evident that the number of constraints in (37) - (39) and (41) is extremely large. Incorporating these constraints and variables a priori into the model is impossible. Hence, we must develop an iterative cut generation algorithm that generates relevant inequalities at each iteration as the solution progresses. A further complicating factor is the fact that we have an infinite number of s_{ijk} variables (uncountably many).

As discussed earlier, in addition to the above service level constraints, the term $\sum_{i \in N} e_i |d_i - d_i^s|$ is also a non-linear term in the objective function (6). However, this term can be linearized using standard techniques and hence we do not discuss this linearization technique in detail. Next, in § 3.2, we describe the cut generation algorithm for the profit maximizing model *PMM*.

3.2. The Cut Generation Algorithm for *PMM*

Based on the constraint linearization procedure described earlier, in this section we develop a constraint generation algorithm to solve our optimization model *PMM*.

We begin by ignoring the *NSL* and *FSL* constraints, i.e., constraints (35) and (40). Recollect that we replace the original constraints (7) and (8) with these new constraints. In addition the term $\sum_{i \in N} e_i |d_i - d_i^s|$ is linearized in the objective function. We refer to the resulting model, without these constraints, as the restricted profit maximizing model *R - PMM*. We initialize our algorithm with *R - PMM*. Let $h \geq 0$ denote an iteration step of the proposed algorithm. Further, let $Z^{h*} = \langle z_{ij}^{h*} \rangle$, $d^{h*} = \langle d_i^{h*} \rangle$, and $a^{h*} = \langle a_i^{h*} \rangle$ denote an optimal solution at the beginning of iteration h , i.e., after solving *R - PMM*. At every iteration let $\bar{\mathcal{S}}^{(n)}$ denote the set of new *NSL* constraints generated and let $\bar{\mathcal{S}}^{(f)}$ denote the set of additional *FSL* constraints generated. Let \mathcal{S} denote the set of combined *NSL* and *FSL* constraints added to the restricted problem *R - PMM*. Each time an *NSL* constraint is generated, the corresponding s variable is also introduced into *R - PMM*.

We list the steps of our constraint and variable generation algorithm in Algorithm 1. In Step 3 of the algorithm we gather all the current passenger connections. Since k_{ij} is the function argument in the right-hand side of (39), we need to consider tangents at this particular point (see Figure 1). Flight j_i is the index with the maximum violation in (37). Step 4.2, in Algorithm 1, introduces the new s variable and adds the corresponding constraints (37) - (39).

Algorithm 1 Algorithm for solving *PMM*

Step 1: Initialize $h = 1, S = \emptyset$ and let $R - PMM$ consist of objective function (6) and constraints (9)-(16).

Step 2: Optimize $R - PMM$ with constraints in \mathcal{S} . Let Z^{h*}, d^{h*}, a^{h*} be the corresponding optimal solution.

Step 3: Build updated connection sets, i.e., for each flight $i \in N$ collect

$$S_i = \{j \in N : z_{ij}^{h*} = 1\}.$$

Step 4: Check for the set of most violated *NSL* constraints. Set $\bar{\mathcal{S}}^{(n)} \leftarrow \emptyset$ and $k_{ij} = d_j^{h*} - d_i^{h*} - m_{ij}$. For each flight $i \in N$

1. Find

$$j_i = \arg \max_{j \in S_i} \{\log \gamma_n - U_{ik_{ij}}(k_{ij})\}.$$

2. If $\log \gamma_n - U_{ik_{ij}}(k_{ij}) > 0$, then define a variable $s_{i,j_i,k_{ij_i}}$, and generate constraints using (37) - (39) with $j = j_i, k = k_{ij_i}$. Add these constraints to $\bar{\mathcal{S}}^{(n)}$.

Step 5: Check for the set of violated *FSL* constraints. Set $\bar{\mathcal{S}}^{(f)} \leftarrow \emptyset$ and $\bar{k}_i = a_i^{h*} - d_i^{h*} + \delta$. For each flight $i \in N$

1. If $\log \gamma_f - U_{i,\bar{k}_i}(\bar{k}_i) > 0$, generate a constraint using (41) with $k = \bar{k}_i$, and add it to $\bar{\mathcal{S}}^{(f)}$.

Step 6: If $\bar{\mathcal{S}}^{(f)} \cup \bar{\mathcal{S}}^{(n)} = \emptyset$, **terminate**;

Step 7: Set $\mathcal{S} \leftarrow \mathcal{S} \cup \bar{\mathcal{S}}^{(n)} \cup \bar{\mathcal{S}}^{(f)}, h \leftarrow h + 1$, go to Step 2.

Next, in Theorem 1 we show that Algorithm 1 is guaranteed to converge to an optimal solution.

THEOREM 1. *There is a subsequence $\{h_q\}_q$ such that $d_i^* = \lim_{q \rightarrow \infty} d_i^{h_q^*}$, $a_i^* = \lim_{q \rightarrow \infty} a_i^{h_q^*}$ for every flight $i \in N$ is an optimal solution to *PMM*.*

PROOF. See the appendix, §6, for the proof. □

It is noteworthy that the proof of Theorem 1 also exhibits an optimal Z^* , s^* , and X^* . As a result, Algorithm 1 converges to an optimal solution for model *PMM*. Furthermore, it is worth emphasizing that the analysis is not trivial. As stated earlier the feasible region of the original problem is non-convex, but, the linearization procedure allows us to circumvent this issue. However, upon linearization, the modified model is infinite dimensional with infinitely many constraints. Algorithm 1 and Theorem 1 generalize the convergence results achieved with (1) semi-infinite linear

programs with finitely many variables but infinitely many constraints, and (2) infinite dimensional problems with finitely many constraints and infinite number of variables.

The computational time to convergence could still be an issue. While we cannot guarantee a bound on the computational time, in §4 we demonstrate that the algorithm converges within reasonable CPU time through extensive computational experiments using real airline data. Next, in § 3.3 we develop an approximate algorithm to solve *SLMM*.

3.3. Cut Generation Algorithm for *SLMM*

The algorithm to solve *PMM* can be modified to approximately solve the alternate model *SLMM*. Notice that constraints (37) - (39), and constraint (41) are also valid for *SLMM*; they replace constraints (18) and (19). To enable a complete linear transformation we define variables $\zeta_n = \log \gamma_n$ and $\zeta_f = \log \gamma_f$. Thus, constraints (37) - (39) and (41) transform as follows:

$$s_{ijk} \geq \zeta_n \quad i \in N, j \in \bar{C}_i, k \in \mathcal{K}_i \quad (42)$$

$$z_{ij} U_{ik} (\delta_l^i - m_{ij}) \leq s_{ijk} \leq 0 \quad i \in N, j \in \bar{C}_i, k \in \mathcal{K}_i \quad (43)$$

$$(1 - z_{ij}) U_{ik} (\delta_l^i - m_{ij}) + s_{ijk} \leq U_{ik} (d_j - d_i - m_{ij}) \quad i \in N, j \in \bar{C}_i, k \in \mathcal{K}_i \quad (44)$$

$$U_{ik} (a_i - d_i + \delta) \geq \zeta_f \quad i \in N, k \in \mathcal{K}_i. \quad (45)$$

We change the objective function of model *SLMM* using ζ_n and ζ_f , however, the objective function is now a non-linear function, i.e., $w_f \exp\{\zeta_f\} + w_n \exp\{\zeta_n\}$. Unfortunately, this is a maximization problem of a convex function and thus is not easily amendable to computational tractability. To simplify the computational procedure we use the first-order linear approximation of $\exp\{x\} = 1 + x$ which transforms the objective function to

$$\max w_f \zeta_f + w_n \zeta_n. \quad (46)$$

The new objective is an approximation of the original problem. Thus, any optimal solution to the transformed objective function may not result in an optimal solution to the original problem. However, the linear approximation allows us to solve for the service levels efficiently. We demonstrate this using several computational experiments in § 4.

In addition to the above service level constraints, the term $\sum_{i \in N} e_i |d_i - d_i^s|$, in constraint (26), is also non-linear. Just as in the case of *PMM*, this term can be linearized using standard techniques and hence we do not discuss it in detail here.

As with the solution methodology for *PMM*, to begin, we ignore the *NSL* and *FSL* constraints, i.e., constraints (42)-(45). We refer to the resulting model, without these constraints, as the restricted service level maximizing model *R-SLMM*. In addition, the objective function is replaced by (46). We initialize our algorithm with *R-SLMM*. Let $h \geq 0$ denote an iteration step of the proposed algorithm. Further, let $Z^{h*} = \langle z_{ij}^{h*} \rangle$, $d^{h*} = \langle d_i^{h*} \rangle$, and $a^{h*} = \langle a_i^{h*} \rangle$, ζ_f^{h*} , and ζ_n^{h*} denote the optimal solution at the beginning of iteration h . At every iteration let $\bar{\mathcal{S}}^{(n)}$ denote the set of new *NSL* constraints generated and $\bar{\mathcal{S}}^{(f)}$ denote the set of additional *FSL* constraints generated. Let \mathcal{S} denote the set of combined *NSL* and *FSL* constraints added to the restricted problem *R-SLMM*. We list the steps used to solve *SLMM* in Algorithm 2. The steps are similar to those in Algorithm 1.

Similar to Theorem 1 it is easy to verify that Algorithm 2 is guaranteed to converge to an optimal solution for the approximate model of *SLMM* with the objective function (46). We state this result as a corollary to Theorem 1 without proof.

COROLLARY 1. *Algorithm 2 converges to an optimal solution of the approximate model SLMM with the objective function (46).*

Next, we describe the computational experiments.

4. Computational Experiments

In this section we describe a series of computational experiments using real airline data. The goal of these experiments is twofold. The primary goal is to study the efficiency of the optimization models *PMM* and *SLMM* to solve the robust scheduling problem. A secondary goal is to study the trade-off an airline faces between higher passenger service levels (as defined by the *NSL* and *FSL*) and the possible degradation in profit using the models described earlier. To this end, we implemented the algorithms described in § 3.2 and § 3.3 to solve 5 airline network instances. We first describe

Algorithm 2 Algorithm for solving *SLMM*

Step 1: Initialize $h = 1, S = \emptyset$ and let $R - SLMM$ consist of objective function (46) and constraints (20)-(28).

Step 2: Optimize $R - SLMM$ with constraints in \mathcal{S} . Let $Z^{h*}, d^{h*}, a_i^{h*}, \zeta_f^{h*}$ and ζ_n^{h*} be the corresponding optimal solution.

Step 3: Build updated connection sets, i.e., for each flight $i \in N$ collect

$$S_i = \{j \in N : z_{ij}^{h*} = 1\}.$$

Step 4: Check for the set of most violated *NSL* constraints. Set $\bar{\mathcal{S}}^{(n)} \leftarrow \emptyset$ and $k_{ij} = d_j^{h*} - d_i^{h*} - m_{ij}$. For each flight $i \in N$

1. Find

$$j_i = \arg \max_{j \in S_i} \{\zeta_n^{h*} - U_{ik_{ij}}(k_{ij})\}.$$

2. If $\zeta_n^{h*} - U_{ik_{ij}}(k_{ij}) > 0$, then define a variable $s_{i,j_i,k_{ij_i}}$ and generate constraints using (42) - (44) with $j = j_i, k = k_{ij_i}$. Add them to $\bar{\mathcal{S}}^{(n)}$.

Step 5: Check for the set of violated *FSL* constraints. Set $\bar{\mathcal{S}}^{(f)} \leftarrow \emptyset$ and $\bar{k}_i = a_i^{h*} - d_i^{h*} + \delta$. For each flight $i \in N$,

1. If $\zeta_f - U_{i,\bar{k}_i}(\bar{k}_i) > 0$, generate a constraint using (45) with $k = \bar{k}_i$, and add it to $\bar{\mathcal{S}}^{(f)}$.

Step 6: If $\bar{\mathcal{S}}^{(f)} \cup \bar{\mathcal{S}}^{(n)} = \emptyset$, **terminate**;

Step 7: Set $\mathcal{S} \leftarrow \mathcal{S} \cup \bar{\mathcal{S}}^{(n)} \cup \bar{\mathcal{S}}^{(f)}, h \leftarrow h + 1$, go to Step 2.

the characteristics of these network instances in Table 1. Due to confidentiality issues we report only the underlying ranges. Instance 1 is the largest network covering all the fleets. Instances 3,

Instance	Flights	Stations	Itineraries	# Fleets
1	1500	85	50,000	5
2	450	75	30,000	2
3	850	80	45,000	3
4	1000	70	30,000	3
5	850	80	35,000	2

Table 1 Characteristics of network instances for the computational experiments.

4, and 5 are relatively medium sized networks and instance 2 is the smallest network. In Table 1, the largest network consists of 5 fleets with varying capacities. The set of itineraries consists of itineraries with up to 4 flights. All the networks are hub-and-spoke and the largest network has 5 hubs. Given that there are a large number of flights in each network, we do not report the block-time statistics for individual flights in these networks. To obtain block time distributions, we analyzed the realized block-times over two consecutive years. We concluded that arrivals are never earlier than 30 minutes before the scheduled arrival time based on the incumbent schedule, however flights can be significantly late. As a result, we assume that the block-times follow a truncated Normal distribution with a lower limit of 30 minutes before the scheduled arrival time and no upper limit. The truncated-Normal distribution satisfies Assumption 3 (Bagnoli and Bergstrom 2005). Obviously, our computational results depend on the form of the assumed distribution and we acknowledge the limitations of the results discussed in the section. However, as mentioned earlier, our main goal is to demonstrate that the solution methodology performs well on the real-world data. Additionally, our models allow planners to study important tradeoffs faced by an airline while increasing schedule reliability. Fine tuning the distributions would definitely provide more accurate results. The means of the block-time distributions vary from 36 minutes to 387 minutes and the variances range from 24.9 to 595.3.

All the problem instances were solved on an Intel Xeon 3.2 GHz dual core server running Redhat's 4.1 version of the Linux operating system. The cut generation algorithm, Algorithm 1, and its variant for the *SLMM* model, Algorithm 2, were developed using the g++ compiler, version 4.1. The mixed integer programming instances were solved using ILOG CPLEX version 10.1 and the models were developed using the ILOG Concert library, version 2.3.

In the accompanying online appendix (Sohoni et al. 2008) we list all the detailed computational results of all the instances described in Table 1. In this paper, however, to demonstrate the efficiency of our models and algorithm, we only summarize some of the performance metrics of *PMM*, for all the 5 instances, in Table 2. A priori, after adjusting for a few outliers, the *NSL* of the incumbent

schedule is 0.4 and the FSL is 0.6. Essentially, we ignored 10 flights with very low service levels to compute the NSL and 5 flights to compute the FSL .

For the first set of experiments the FSL was held constant at 0.8 and the NSL level was allowed to vary from 0.8 to 0.95 in increments of 0.1. We denote this set of experiments as Fixed- FSL . In the next set of experiments the NSL was held constant and the FSL was allowed to vary from 0.8 to 0.95 in increments of 0.1. We denote this set of experiments as Fixed- NSL . All these instances were solved to optimality. We report the maximum CPU time, maximum number of iterations, and maximum number of cuts generated using Algorithms 1 and 2 in Table 2. The results show

Metric	Network Instance				
	1	2	3	4	5
Fixed- FSL (max CPU secs)	3,842	167	952	838	801
Fixed- FSL (max iterations)	27	10	17	12	13
Fixed- FSL (max cuts added)	980	227	767	315	428
Fixed- NSL (max CPU secs)	3,921	187	921	873	792
Fixed- NSL (max iterations)	24	8	13	8	10
Fixed- NSL (max cuts added)	987	243	793	343	437

Table 2 Algorithm performance for model PMM .

that PMM performs reasonably well on all networks, especially, considering the fact that schedule development is performed several months prior to the day of operations and airlines do not mind spending additional computation time. Furthermore, the results also indicate that both algorithms converge within a few iterations.

For the remainder of the computational experiments we restrict our attention to instance 1 because it is the largest network. We discuss these experiments in §4.1 and §4.2. As mentioned earlier, results for all other instances can be found in Sohoni et al. (2008).

4.1. The PMM Model

The first set of experiments are for model PMM . Through several experiments we demonstrate the trade-off between higher service levels and planned profit. For all these experiments we restrict the flight departure times to be adjusted within 60 minutes of those specified in the incumbent schedule. Furthermore, the penalty for adjusting the departure time is assumed to be the same for all the flights $i \in N$ and is held at 1, i.e. $e_i = e_j = 1$ for all $i, j \in N$.

Effect on Profit. First, in Figure 2 we show how the profit (objective function) varies as the NSL is varied for different FSL levels. The profit is computed as a percentage of the planned profit of the incumbent schedule. In general the profits decrease as NSL increases. Since the NSL and FSL of the incumbent schedule are lower than those considered, these profits are also lower. We vary the NSL from 0.8 to 0.95. With lower FSL the decrease in profit is less pronounced as the NSL increases. To achieve extremely high NSL and FSL , substantial profit decrease must be tolerated. Next, in Figure 3, we show how the profit varies as the FSL is varied for different NSL

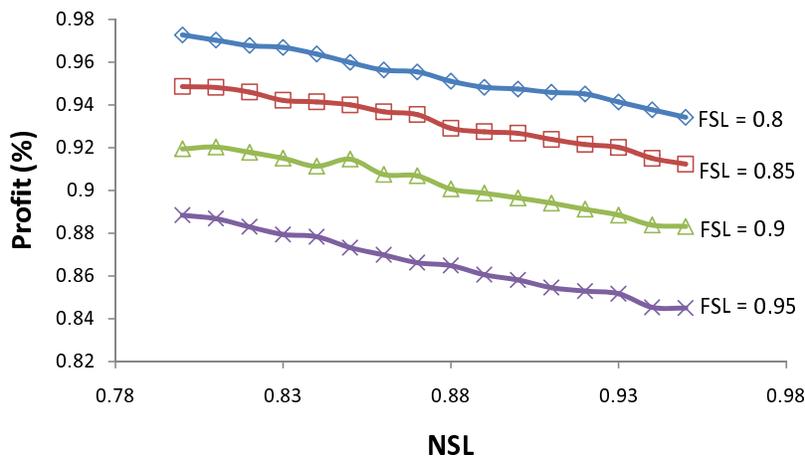


Figure 2 Effect of NSL on the profit level (percent of incumbent schedule profit).

levels. In this case too the profit levels decrease. Again, the profit is computed as a percentage of the planned profit of the incumbent schedule. Similar to the NSL experiments, we vary the FSL from 0.8 to 0.95. Finally, Figure 4 summarizes the reduction in profit, as a percentage of the planned profit of the incumbent schedule, as both the NSL and FSL vary from 0.8 to 0.95. It is noteworthy that at very high service levels the reduction in profit is 13%. However, the airline may be willing to consider a lesser degradation in profit to still achieve substantial improvement in FSL and NSL . We observe that the profit decreases almost linearly with respect to the FSL and NSL . This is confirmed also in Figures 2 and 3.

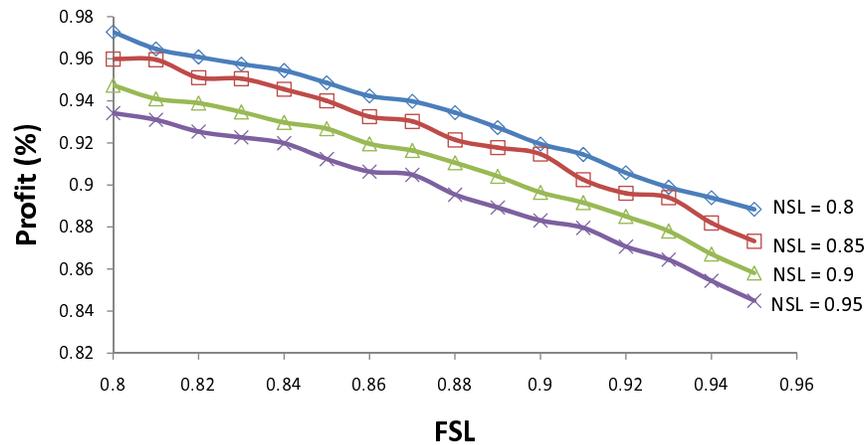


Figure 3 Effect of FSL on the profit level (percent of incumbent schedule profit).

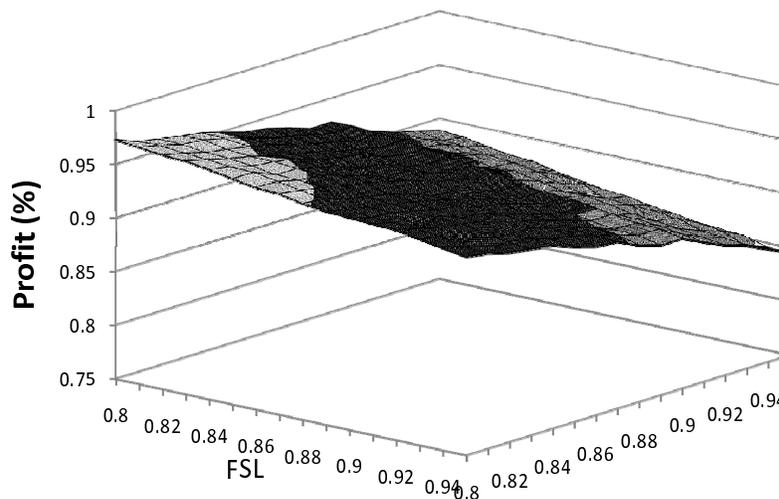


Figure 4 Effect of FSL and NSL on the profit level (percent of incumbent schedule profit).

Number of Departures Changed and Passenger Connections for the PMM Model.

In this set of experiments we vary the NSL and FSL between 0.8 and 0.95 and study the number of departure times and passenger connections affected. Figure 5 shows the effect of varying the NSL and Figure 6 shows the effect of varying FSL . In the former the FSL is fixed at 0.8 and in the latter set of experiments the NSL is fixed at 0.8. In both these figures the solid line with block markers represents the number of passenger connections achievable and the dotted line with diamond markers represents the number of flight departures affected. In Figure 5 the connections steadily decrease from 4,186 to 3,725 while the number of departures increases (as shown by the thin dark trend-line) from 171 to 273 as the NSL changes. Similarly, in Figure 6 the connections vary

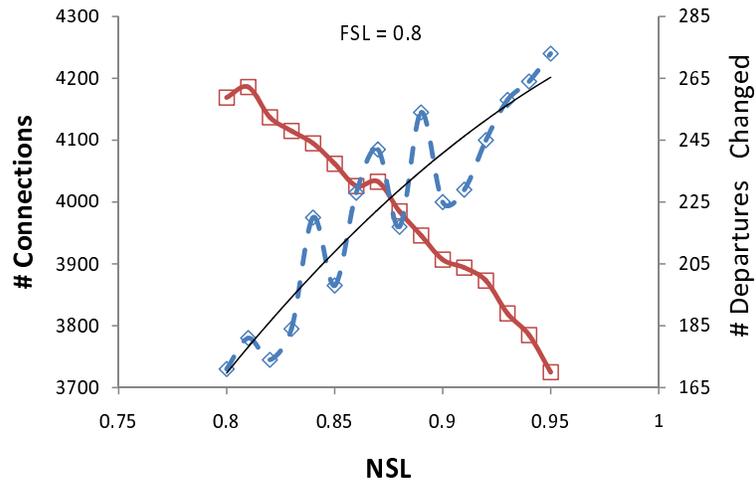


Figure 5 Effect of NSL on the departures changed and passenger connections.

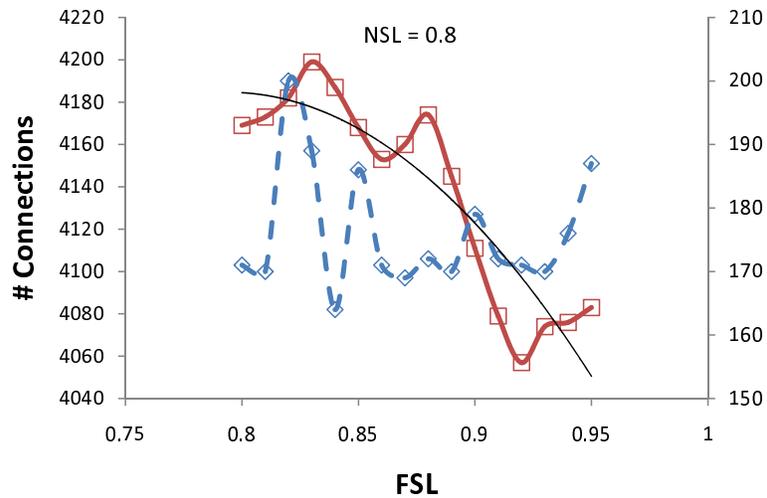


Figure 6 Effect of FSL on the departures changed and passenger connections.

between 4,197 to 4,055 and the total departures adjusted fluctuate between 164 and 200. There is not a clear trend in how the FSL affects the number of departures adjusted, however, as indicated by the thin dark trend-line the connections show a decreasing trend. Intuitively, to achieve high service levels, flexibility is required and thus more departure time changes are expected. Figure 5 confirms this, while it is not evident from Figure 6. The NSL captures passenger connections. If an airline operates a single flight, the NSL is 100%. We expect that as the NSL is increased, the number of passenger connections should decrease (clearly at the expense of diminishing profit).

This intuition is confirmed by both figures.

Effect of Deviation Penalty for the PMM Model. Here we vary the penalty from the departure time in the incumbent schedule ($e_i, i \in N$). In these experiments, however, we do not discriminate between flights, i.e., we assume $e_i = e_j$ for every $i, j \in N$. First, in Figure 7 we plot

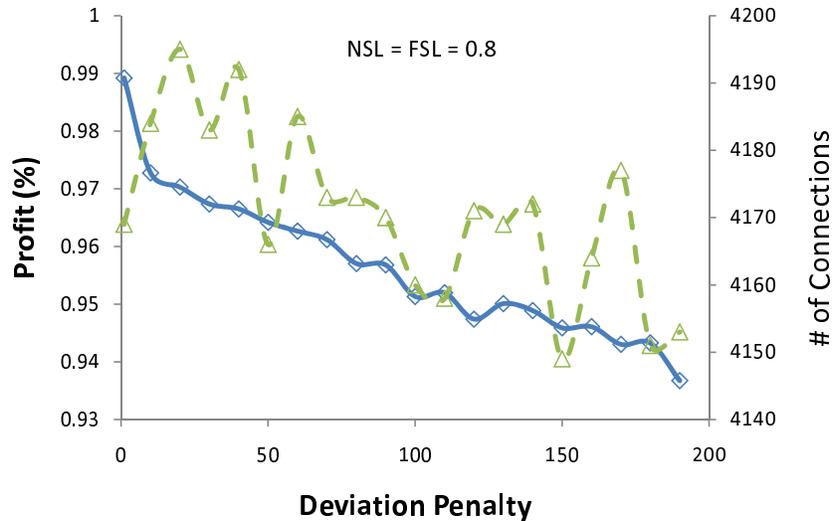


Figure 7 Effect of departure deviation penalty on profit and passenger connections.

the profit (as a percentage of the profit of the incumbent schedule) and the number of passenger connections changed as the deviation penalty is varied from 0 to 200. The thick-line represents the profit level and the dashed-line represents the connections affected. The NSL and FSL are held at 0.8 for all these experiments. The profit, as well as the connections, show a decreasing trend with respect to the deviation penalty. This is expected since higher deviation penalties imply more costly perturbations and thus the trade-off between schedule changes and profit sways towards schedule changes.

4.2. The SLMM Model.

Here, we report the experimental results for the $SLMM$ model. For these set of experiments we define the relative weight between the NSL and FSL as $\omega = \frac{w_n}{w_f}$. The first set of experiments for the $SLMM$ model study the effect of ω on the NSL and FSL achieved. The minimum profit level

(for constraint (26)) is restricted to 100% of the profit achieved with the incumbent schedule. Thus we do not allow any decrease in profitability. Figure 8 shows the results of these experiments. The

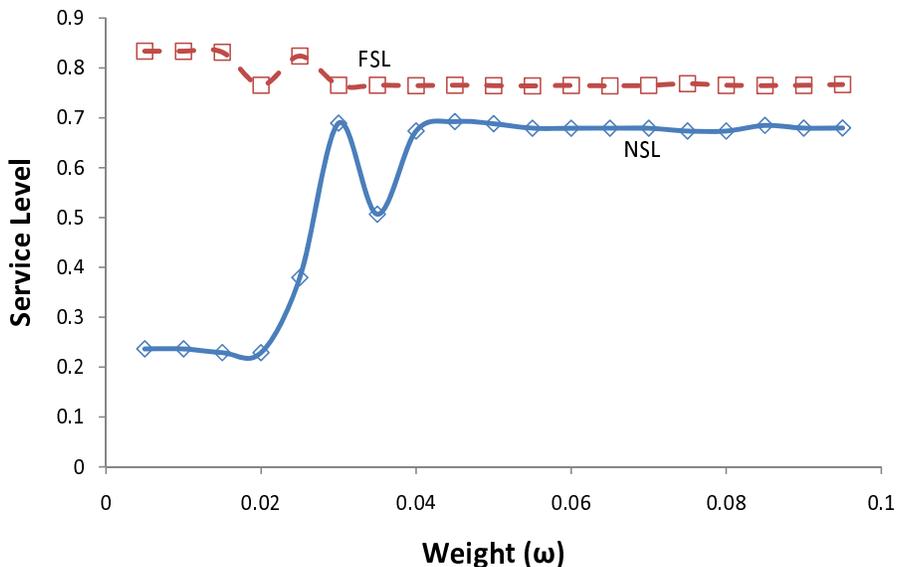


Figure 8 Effect of ω on NSL and FSL .

dotted-line represents the FSL and the solid-line represents the NSL . Since the objective function in Algorithm 2 is an approximation to the two service levels, given a solution, we have to compute the service levels from the schedule obtained. The experimental results, with the departure time window fixed at 30 minutes, show that FSL drops initially while the NSL increases as ω increases. Interestingly, though, both these service levels stabilize beyond a weight level and remain almost constant even for very large values of ω . Such a behavior is expected since ω captures the trade-off between the two service levels. From Figure 8 we can also observe the maximum possible service levels with the same profits.

The next set of experiments vary the departure time window between which the flight departure times are varied. Again, we study the effect on the service levels achievable while still assuring 100% of the original schedule's profit. Figure 9 plots both these curves. The dotted-line represents the FSL and the solid-line represents the NSL . As the departure time window is expanded, the NSL improves while the FSL marginally drops. This affirms that it is harder to increase the NSL

than the FSL . To increase the NSL , further flexibility has to be provided such as increased time windows. Figure 9 provides a clear trade-off with respect to the permissible schedule change and

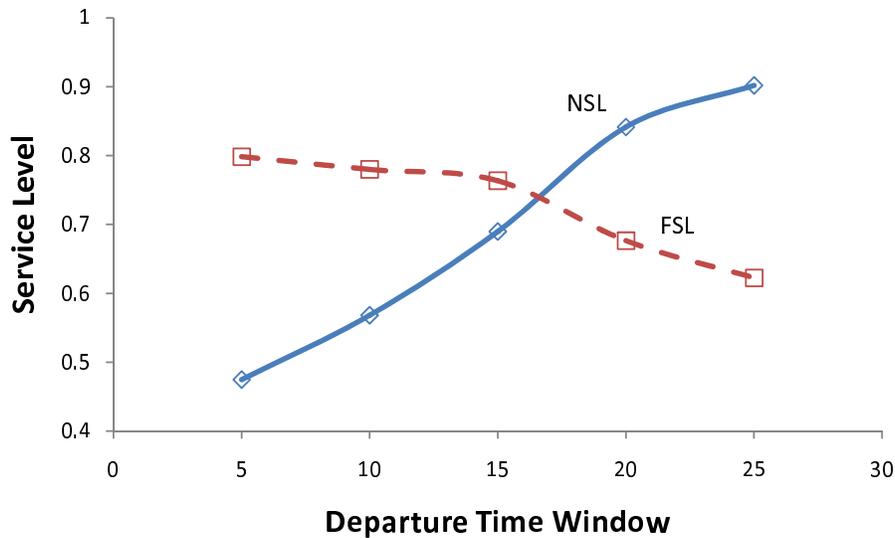


Figure 9 Effect of departure time window on NSL and FSL ($\omega = 0.7$).

the two service levels without reducing profitability. For example, if the departure time windows are 10 minutes (which may not be always acceptable, e.g., due to competition and implications on demand), then a NSL of about 58% and a FSL of about 79% is achievable at the same profit level. This is a substantial improvement over the original values of 40% and 60% for the NSL and FSL , respectively. For these experiments $\omega = 0.7$, which is a value of ω where the service levels are stable.

The final set of experiments study the trade-off between the service levels and the profit level. The profit is computed as a percentage of the profit for the incumbent schedule when $\omega = 0.7$ and the departure time window is set to 30 minutes. Figure 10 shows the variation in NSL , FSL , and the value $\omega NSL + FSL$ as the profit is reduced. At 100% profit, $NSL = 0.68$ and $FSL = 0.68$. However, as the profit percentage is reduced, NSL increases only slightly, i.e., up to 0.71 at 90% profit level, while FSL increases substantially to 0.91. Note that a different trade-off would be assessed if ω is changed. Thus, the airline may gain on service levels by slightly adjusting the profit.

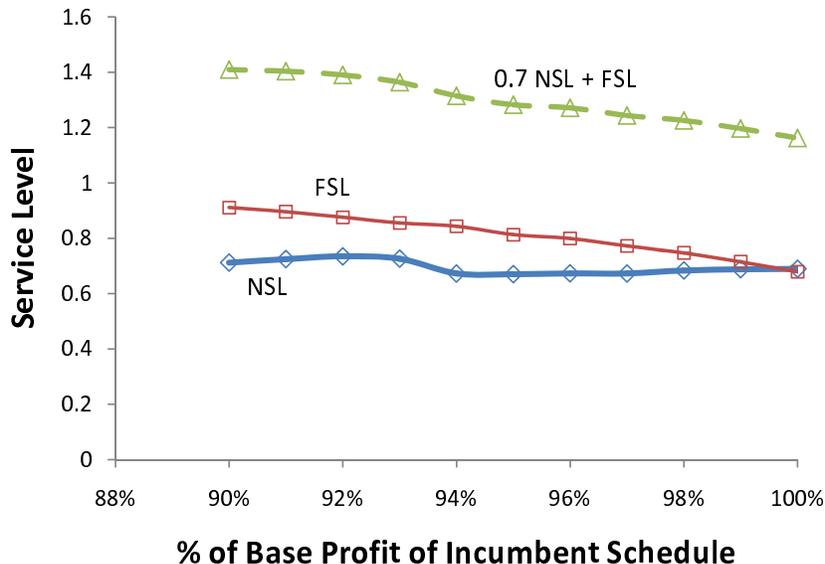


Figure 10 Effect of reduction in profit level on *NSL* and *FSL* ($\omega = 0.7$).

5. Discussion

In this paper we developed two models that incorporate uncertainty associated with block-times into the schedule development process. It is an initial attempt in developing a comprehensive and holistic model for incorporating block-time uncertainty in schedule planning. We explicitly model time distributions through chance constraints and hence the resulting schedule is robust with respect to the operational on-time performance measure. We also incorporate network service levels, which probabilistically model passenger connections. The new cut generation algorithm and linearization technique proposed are novel in the sense that the convergence result generalizes previously established results with (1) semi-infinite linear programs with finitely many variables but infinitely many constraints, and (2) infinite dimensional problems with finitely many constraints and infinite number of variables.

The benefits of our approach are two fold: (1) airlines could adjust the schedule to increase operational reliability, and (2) passengers could be guaranteed higher service levels. There are potentially other indirect benefits of adjusting the schedule by incorporating block-time uncertainty. For example, the schedule recovery cost due to a disruption during actual operations could be reduced because the planned block-times allow additional flexibility. However, we have not

specifically included such additional benefits in the models presented. Through extensive computational experiments we demonstrate the efficiency of our algorithms and models in trading off between profitability and service level guarantees. The algorithms perform well in achieving this trade-off and provide airline schedule planners the ability to decide on acceptable reduction in profitability to achieve desired passenger service levels.

There are several possible modifications and enhancements to the models described in this paper. First, the dependency of block-time distributions on the departure time can be included, if such information is readily available. This also allows for wider time windows to vary the departure times of scheduled flights and capture “time-of-the-day” effects related to block-time distributions. However, the resulting model is more complicated than those described by the *PMM* and *SLMM* because it requires the introduction of additional binary variables and several additional constraints to model the choice of the appropriate departure-time dependent block-time distribution. The key concept is to discretize each time window and assign a specific block-time distribution to each subinterval. Standard modeling techniques using piecewise linear functions capture these assignments.

Second, in the current model we assume that the block-times follow continuous log-concave distributions. It is possible that there may be a discrete jump in the actual block-time, i.e., the block-time distributions follow a discrete log-concave distribution. While our model can be considered as an approximation to the discrete case, incorporating discrete distributions may not guarantee convergence, unlike the case discussed in this paper with continuous distributions.

Third, to keep our analysis tractable and focus on the issue of schedule reliability, our model does not incorporate the trade-off between the local and through passengers. While constraints, (14) in the *PMM* model and (25) in the *SLMM* model, enforce the fact that the original passenger connections (itineraries) are feasible with any schedule perturbation, any solution to our model provides a lower bound to the potential revenue (and profit) achievable, if such a profitable trade-off (substitution of demand) were to be specifically included in these optimization models. Additionally, such data would also have to be captured. These constraints also enforce that the

aircraft routing solution to the incumbent solution continues to remain feasible under any schedule perturbations. In a more general setting it is possible to relax these constraints and allow larger perturbations of the schedule by embedding fleet assignment and aircraft routing constraints. In this case several additional passenger connections may also become feasible as the departure times are perturbed. Such a model would be an extension of our model and significantly harder to solve.

To address the issue of time dependent demand distributions for local and through passengers, one possible way is to construct multiple copies of the same flight, each with its own specific demand. This would necessitate the inclusion of additional constraints enforcing that exactly one of these copies is chosen as well as the connections of aircraft rotations and passenger flows remain feasible.

Fourth, in this paper, we focus on perturbing the schedule by including block-time uncertainty. However, as the block-times are varied, and the departure times are perturbed, the ground-times are automatically adjusted. It is possible, that airlines would want to trade-off between the allocation of ground-times and block-times to perturb the schedule. Currently, we do not explicitly model ground-time constraints because the form of the distributions for the ground-time variables is not known. Our solution methodology could be extended if these distributions are log-concave. It is possible that, in the current model, we could capture the change in the objective function due to increase or decrease in ground-times. This could be done by including some cost/profit associated with perturbing ground-times in the objective function. For example, if G_{ij} denotes the cost/profit associated with increasing the ground-time between flights i and j in an aircraft's rotation, we could include the terms $\sum_{(i,j) \in T} G_{ij}(d_j - a_j)$ in the objective function. This, of course, would not qualitatively change our main insights.

Finally, in the current setting we do not distinguish between various markets an airline serves (i.e., different portions of the network). It is possible to incorporate different service levels for different markets and use similar models, as described in this paper, to perturb the schedule and set suitable block-times. In the current setting we only guarantee a minimum service level for the entire network.

References

- Associated Press, The. 2007. Airlines' on-time marks in '07 near record low. URL <http://www.msnbc.msn.com/id/22490808/>.
- Bagnoli, M., T. Bergstrom. 2005. Log-concave probability and its applications. *Economic Theory* **26** 445–469.
- Barnhart, C., A. Cohn. 2004. Airline schedule planning: Accomplishments and opportunities. *Manufacturing & Service Operations Management* **6**(1) 3–2.
- Barnhart, C., A. Cohn, E. L. Johnson, D. Klabjan, G. Nemhauser, P. Vance. 2003. Airline crew scheduling. R. W. Hall, ed., *Handbook of Transportation Science*. Kluwer Scientific Publishers, Boston, Massachusetts, 493–521.
- Birge, J. R., F. Louveaux. 1997. *Introduction to Stochastic Programming*. Springer Verlag.
- Bly, E., J. Clarke, J. Loy, T. Melconian, F. Rabbani. 2003. MIT extensible air network simulation (MEANS). *Technical report*. Massachusetts Institute of Technology, Cambridge, MA.
- Bureau of Transportation Statistics, The. 2009. On-time performance: Flight delays at a glance. URL <http://www.transtats.bts.gov/HomeDrillChart.asp>.
- Deshpande, V., M. Arikan. 2009. The impact of airline flight schedules on flight delays: An empirical study. *Working paper*. Purdue University, West Lafayette, IN.
- Dobson, G., P. J. Lederer. 1993. Airline scheduling and routing in a hub-and-spoke system. *Transportation Science* **27** 281–297.
- Karow, M. 2003. Virtual hubs: An airline schedule recovery concept and model. *Master's thesis*. Massachusetts Institute of Technology, Cambridge, MA.
- Klabjan, D. 2005. Large-scale models in the airline industry. G. Desaulniers, J. Desrosiers, M.M. Solomon, eds., *Column Generation*. Springer Science + Business Media, New York, NY.
- Lan, S., J. Clarke, C. Barnhart. 2003. Robust aircraft routing. *Technical report*. Massachusetts Institute of Technology, Cambridge, MA.
- Lohatepanont, M., C. Barnhart. 2004. Airline schedule planning: Integrated models and algorithms for schedule design and fleet assignment. *Transportation Science* **38**(1) 19–32.
- Mayer, C., T. Sinai. 2003a. Network effects, congestion externalities, and air traffic delays: or why not all delays are evil. *The American Economic Review* **93** 1194–1214.

- Mayer, C., T. Sinai. 2003b. Why do airlines systematically schedule their flight to arrive late? *Working paper*. University of Pennsylvania, Philadelphia, PA.
- Portougal, V., D. Trietsch. 2001. Stochastic scheduling with optimal customer service. *Journal of the Operational Research Society* **52** 226–233.
- Rosenberger, J., E. Johnson, G. Nemhauser. 2004. A robust fleet assignment model with hub isolation and short cycles. *Transportation Science* **38** 357–368.
- Rosenberger, J. M., A. J. Schaefer, D. Goldsman, E. L. Johnson, A. J. Kleywegt, G. L. Nemhauser. 2002. A stochastic model of airline operations. *Transportation Science* **36**(4) 357–377.
- Shebalov, S., D. Klabjan. 2006. Robust airline crew pairing : Move-up crews. *Transportation Science* **40**(3) 300–312.
- Sohoni, M., Y-C. Lee, D. Klabjan. 2008. BlockTimeEstimationAndRobustAirlineScheduling_OnlineAppendix. URL <http://www.klabjan.dynresmanagement.com/articles/>.

6. Appendix

Proof of Theorem 1.

Recollect that h denotes the iteration index in Algorithm 1. Given that Z^{h*} only has a finite number of different values, there exists a subsequence such that $Z^{h_1*} = Z^{h_2*} = Z^{h_3*} = \dots = Z^*$. Furthermore, since for every flight $i \in N$, $d_i^{h_{q*}} \in [l_i, u_i]$, there exists a convergent subsequence in $\{d_i^{h_{q*}}\}_q$. From constraint (8) it follows that $d_i^{h_{q*}} \leq a_i^{h_{q*}}$ for every $i \in N$. Additionally, constraint (9) ensures that the subsequence $\{a_i^{h_{q*}}\}_q$ is upper bounded. Thus, we conclude that there exists a convergent subsequence in $\{a_i^{h_{q*}}\}_q$.

Let us denote the values of s_{ijk} in the optimal solution to $R - PMM$ by s^{h*} . We now assume that if $z_{ij}^{h*} = 1$, then $s_{ijk}^{h*} = \min\{0, U_{ik}(d_j^{h*} - d_i^{h*} - m_{ij})\}$ for every $k \in K_i$, $i \in N$, and $j \in \bar{C}_i$. If this is not the case, we can easily increase s_{ijk}^{h*} to satisfy this property without affecting feasibility. Furthermore, observe that for every itinerary o and fare-class f combination we have $0 \leq X_{of}^{h_{q*}} \leq D_{of}$. Therefore, there must be a convergent subsequence in $\{X_{of}^{h_{q*}}\}_q$. Here X^{h*} denotes the optimal itinerary fare-class demand values.

From the above set of arguments and due to a finite number of flight legs, there is a subsequence where all the departure and arrival times converge in addition to the itinerary fare-class demand

values. For ease of notation we denote this subsequence by $\{h_q\}_q$. Let d^* and a^* be the sets of departure and arrival times of flights, respectively, as defined in the statement of the theorem. Furthermore, we also define, for every $k \in K_i$ and $i \in N$, $j \in \bar{C}_i$

$$s_{ijk}^* = \begin{cases} 0 & z_{ij}^* = 0 \\ \min \{0, U_{ik}(d_j^* - d_i^* - m_{ij})\} & z_{ij}^* = 1. \end{cases} \quad (47)$$

It remains to be shown that d^* , a^* , Z^* , and s^* is an optimal solution to PMM , i.e., these values satisfy constraints (9)-(16), constraints (37)-(39), and (41). It is easy to verify that constraints (9)-(16) are satisfied because only a finite number of them exist. Hence, we first discuss constraints (37)-(39). On closer observation it is easy to note that constraints (38) and (39) hold by definition. Thus, we focus our attention on constraints (37).

To this end, let us fix a $i \in N$ and $j \in \bar{C}_i$. We first show that

$$s_{ijk_i}^{*h_{\bar{q}}} \geq \log \gamma_n \quad (48)$$

for every \bar{q} and $k_i^{h_{\bar{q}}} = k_{ij_i}$ in iteration $h_{\bar{q}}$.

First, let $z_{ij}^* = 1$ and $q \geq \bar{q} + 1$. We have $z_{ij}^{h_{q^*}} = 1$ for every q and thus $s_{ijk_i}^{h_{q^*}} \leq U_{ik_i}^{h_{q^*}}(d_j^{h_{q^*}} - d_i^{h_{q^*}} - m_{ij})$ and $\log \gamma_n \leq s_{ijk_i}^{h_{q^*}}$. We conclude $\log \gamma_n \leq U_{ik_i}^{h_{q^*}}(d_j^{h_{q^*}} - d_i^{h_{q^*}} - m_{ij})$. Since these constraints are not removed from $R - PMM$ in later iterations, we have $\log \gamma_n \leq U_{ik_i}^{h_{\bar{q}}}(d_j^{h_{\bar{q}}} - d_i^{h_{\bar{q}}} - m_{ij})$. Since U 's are continuous, by taking the limit as $q \rightarrow \infty$, we obtain

$$\log \gamma_n \leq U_{ik_i}^{h_{\bar{q}}}(d_j^* - d_i^* - m_{ij}).$$

Since $\log \gamma_n \leq 0$, we obtain

$$\log \gamma_n \leq \min \left\{ 0, U_{ik_i}^{h_{\bar{q}}}(d_j^* - d_i^* - m_{ij}) \right\} = s_{ijk_i}^{*h_{\bar{q}}}. \quad (49)$$

Thus, we have proved(48).

We now consider

$$\min_{k \in K_i} s_{ijk}^* = \min_{k \in K_i} \{0, U_{ik}(d_j^* - d_i^* - m_{ij})\} = g_i(d_j^* - d_i^* - m_{ij}).$$

Observe that we have

$$\begin{aligned} g_i(d_j^* - d_i^* - m_{ij}) &= g_i(d_j^{h_q^*} - d_i^{h_q^*} - m_{ij}) + g_i(d_j^* - d_i^* - m_{ij}) - g_i(d_j^{h_q^*} - d_i^{h_q^*} - m_{ij}) \\ &= U_{ik_i^{h_q}} \left(d_j^{h_q^*} - d_i^{h_q^*} - m_{ij} \right) + g_i(d_j^* - d_i^* - m_{ij}) - g_i(d_j^{h_q^*} - d_i^{h_q^*} - m_{ij}) \quad (50) \end{aligned}$$

$$\begin{aligned} &\geq U_{ik_i^{h_q}} \left(d_j^{h_q^*} - d_i^{h_q^*} - m_{ij} \right) - U_{ik_i^{h_q}} \left(d_j^* - d_i^* - m_{ij} \right) + \log \gamma_n \\ &+ g_i(d_j^* - d_i^* - m_{ij}) - g_i(d_j^{h_q^*} - d_i^{h_q^*} - m_{ij}) \quad (51) \end{aligned}$$

$$\begin{aligned} &= \frac{p_i(k_i^{h_q})}{\int_0^{k_i^{h_q}} p_i(t) dt} \left[d_j^{h_q^*} - d_j^* + d_i^* - d_i^{h_q^*} \right] + \log \gamma_n \\ &+ g_i(d_j^* - d_i^* - m_{ij}) - g_i(d_j^{h_q^*} - d_i^{h_q^*} - m_{ij}). \quad (52) \end{aligned}$$

In the above, equation (50) follows from the fact that $k_i^{h_q}$ maximizes the violation of constraint (37) (see Step 4.1 in Algorithm 1). Furthermore, (51) follows from (48) and (49). The last equality, i.e. equation (52), follows from the definition of U_{ik} .

Observe that the first term in equation (52) converges to 0 since $\frac{p_i(k)}{\int_0^k p_i(t) dt}$ is bounded for $k \in K_i$ for all $i \in N$. Similarly, the last two terms also converge to 0 since g_i is a continuous function. Thus, we conclude that

$$s_{ijk}^* \geq \log \gamma_n \quad \text{for every } k \in K_i, i \in N. \quad (53)$$

It is easy to verify that when $z_{ij}^* = 0$ constraint (37) holds. We conclude that (47) holds in general.

Using similar arguments it can be shown that constraint (41) also holds. Hence, Z^* , d^* , a^* , and s^* is a feasible solution to PMM .

It remains to show optimality. Let $V(Z^*, d^*, a^*, s^*)$ denote the objective value of the corresponding solution. Further, notice that in each iteration h , the optimal value V^{h^*} of $R - PMM$ is an upper bound on the global optimal value V^* , i.e., $V^* \leq V^{h^*}$. Thus, we must have

$$V(Z^*, d^*, a^*, s^*) \leq V^* \leq V^{h_q^*} = V(Z^{h_q^*}, d^{h_q^*}, a^{h_q^*}, s^{h_q^*}). \quad (54)$$

Since the objective function is continuous, by taking the limit, we obtain $\lim_{q \rightarrow \infty} V(Z^{h_q^*}, d^{h_q^*}, a^{h_q^*}, s^{h_q^*}) = V(Z^*, d^*, a^*, s^*)$. From (54) we obtain $V^* = V(Z^*, d^*, a^*, s^*)$. To arrive at this we must also have $\{X_{of}^{h_q^*}\}_q$ be a convergent subsequence, which was assumed earlier.

Hence, we have completed the proof. ■