An Incentive Compatible, Efficient Market for Air Traffic Flow Management

Ruta Mehta, College of Computing, Georgia Institute of Technology.
Vijay V. Vazirani, College of Computing, Georgia Institute of Technology.

We present a market-based approach to the Air Traffic Flow Management (ATFM) problem. The goods in our market are delays and buyers are airline companies; the latter pay money to the FAA to buy away the desired amount of delay on a per flight basis. We give a notion of equilibrium for this market and an LP whose solution gives an equilibrium allocation of flights to landing slots as well as equilibrium prices for the landing slots. Via a reduction to matching, we show that this equilibrium can be computed combinatorially in strongly polynomial time. Moreover, there is a special set of equilibrium prices, which can be computed easily, that is identical to the VCG solution, and therefore the market is incentive compatible in dominant strategy.

1. INTRODUCTION

Air Traffic Flow Management (ATFM) is a challenging operations research problem whose importance keeps escalating with the unabated growth of the airline industry. In the presence of inclement weather, the problem becomes particularly serious and leads to huge monetary losses and delays[1], Yet, despite massive efforts on the part of the U.S. Federal Aviation Administration (FAA), airline companies, and even the academia, the problem remains largely unsolved.

In a nutshell, the reason for this is that any viable solution needs to satisfy several conflicting requirements, e.g., in addition to ensuring efficiency the solution also needs to be viewed as "fair" by all parties involved. Indeed, [Bertsimas et al. 2011] state that " ... While this work points at the possibility of dramatically reducing delay costs to the airline industry vis-a-vis current practice, the vast majority of these proposals remain unimplemented. The ostensible reason for this is fairness ... ." It also needs to be computationally efficient – even moderate sized airports today handle hundreds of flights per day, with the 30 busiest ones handling anywhere from 1000 to 3000 flights per day. The full problem involves scheduling flight-landings simultaneously for multiple airports over a large period of time, taking into consideration inter-airport constraints. Yet, according to [Barnhart et al. 2009], current research has mostly remained at the level of a single airport because of computational tractability reasons.

Building on a sequence of recent ideas that were steeped in sound economic theory, and drawing on ideas from game theory and the theory of algorithms, we present a solution that has a number of desirable properties. Our solution for allocating flights

[1]According to [Barnhart et al. 2009], the U.S. Congress Joint Economic Committee estimated that in 2007, the loss to the U.S. economy was $25.7 billion, due to 2.75 million hours of flight delays. In contrast, the total profit of U.S. airlines in that year was $5 billion.

Author's email addresses: rmehta, vazirani@cc.gatech.edu. Supported by NSF Grants CCF-0914732 and CCF-1216019.

DOI:http://dx.doi.org/10.1145/0000000.0000000

to landing slots at a single airport is based on the principle of a free market, which is known to be fair and a remarkably efficient method for allocating scarce resources among alternative uses (sometimes stated in the colorful language of the "invisible hand of the market" [Smith 1776]). At the same time, we prove that our solution can be captured through linear programming duality. Next, using results from matching theory, we show how to find equilibrium allocations and prices in strongly polynomial time. Moreover, using [Leonard 1983] it turns out that our solution is incentive compatible in dominant strategy, i.e., the players will not be able to game the final allocation to their advantage by misreporting their private information. Finally, we show how the robustness properties of this solution make it extendable to multiple airports.

We note that the ATFM problem involves several issues that are not of a game-theoretic or algorithmic nature, e.g., the relationship between long term access rights (slot ownership or leasing) and short term access rights on a given day of operations, e.g., see [Ball et al. 2005]. Our intention in this paper is not to address the myriad of such issues. Instead, we have attempted to identify a mathematically clean, core problem that is amenable to the powerful tools developed in the theories stated above, and whose solution could form the core around which a practical scheme can be built.

Within academia, research on this problem started with the pioneering work of Odoni [Odoni 1987] and it flourished with the extensive work of Bertsimas et. al. ; we refer the reader to [Barnhart et al. 2009; Bertsimas and Gupta 2010] for thorough literature overviews and references to important papers. These were centralized solutions in which the FAA decides a schedule that is efficient, e.g., it decides which flights most critically need to be served first in order to minimize cascading delays in the entire system.

A conceptual breakthrough came with the realization that *the airlines themselves are the best judge of how to achieve efficiency*[2], thus moving away from centralized solutions. This observation led to a solutions based on collaborative decision making (CDM) which is used in practice [Vossen and Ball 2006; Ball et al. 2006; Wambsganss 1996].

More recently, a market based approach was proposed by Castelli, Pesenti and Ranieri [Castelli et al. 2011]. This approach not only leads to efficiency but at the same time, it simply finesses away the sticky issue of fairness – whoever pays gets smaller delays, much the same way as whoever pays gets to fly comfortably in Business Class! [Castelli et al. 2011] also gave a tatonnement-based implementation of their market. Each iteration starts with FAA announcing prices for landing slots. Then, airlines pick their most preferred slots followed by FAA adjusting prices, to bring parity between supply and demand, for the next iteration. However, they are unable to show convergence of this process and instead propose running it a prespecified number of times, and in case of failure, resorting to FAA's usual solution. They also give an example for which incentive compatibility does not hold. Our market formulation is quite different, and achieves both efficient implementation and incentive compatibility.

1.1. Salient features of our solution

In Section 2 we give details of our basic market model for allocating a set of flights to landing slots for one airport. This set of flights is picked in such a way that their actual arrival times lie in a window of a couple of hours; the reason for the latter will be clarified in Section 4. The goods in our market are delays and buyers are airline

[2]e.g., they know best if a certain flight needs to be served first because it is carrying CEOs of important companies who have paid a premium in order to reach their destination on time or if delaying a certain flight by 30 minutes will not have dire consequences, however delaying it longer would propagate delays through their entire system and result in a huge loss.

companies; the latter pay money to the FAA to buy away the desired amount of delay on a per flight basis. Typically flights have a myriad interdependencies with other flights – because of the use of the same aircraft for subsequent flights, passengers connecting with other flights, crew connecting with other flights, etc. The airline companies, and not FAA, are keenly aware of these and are therefore in a better position to decide which flights to avoid delay for. The information provided by airline companies for each flight is the dollar value of delay as perceived by them.

For finding equilibrium allocations and prices in our market, we give a special LP in which parameters can be set according to the prevailing conditions at the airport and the delay costs declared by airline companies. We arrived at this LP as follows. Consider a traffic network in which users selfishly choose paths from their source to destination. One way of avoiding congestion is to impose tolls on roads. [Cole et al. 2003] showed the existence of such tolls for minimizing the total delay for the very special case of one source and one destination, using Kakutani's fixed point theorem. Clearly, their result was highly non-constructive. In a followup work, [Fleischer et al. 2004] gave a remarkable LP whose optimal solution yields such tolls for the problem of arbitrary sources and destinations and moreover, this results in a polynomial time algorithm. Their LP, which was meant for a multi-commodity flow setting, was the starting point of our work. One essential difference between the two settings is that whereas they sought a Nash equilibrium, we seek a market equilibrium; in particular, the latter requires the condition of market clearing.

We observe that the underlying matrix of our LP is totally unimodular and hence it admits an integral optimal solution. Such a solution yields an equilibrium schedule for the set of flights under consideration and the dual of this LP yields equilibrium price for each landing slot. Equilibrium entails that each flight is scheduled in such a way that the sum of the delay price and landing price is minimum possible. We further show that an equilibrium can be found via an algorithm for the minimum weight perfect b-matching problem and hence can be computed combinatorially in strongly polynomial time. In hindsight, our LP resembles the b-matching LP, but there are some differences.

Since the b-matching problem reduces to the maximum matching problem, our market is essentially a matching market. Leonard [Leonard 1983] showed that the set of equilibrium prices of a matching market which are simultaneously minimum for all goods corresponds precisely to VCG payments [Nisan 2007], thereby showing that the market is incentive compatible in dominant strategy. For our market, we give a simple procedure that converts arbitrary equilibrium prices to ones that are simultaneously minimum for all slots. Incentive compatibility with these prices follows. An issue worth mentioning is that the total revenue, or the total cost, of VCG-based incentive compatible mechanisms has been studied extensively, mostly with negative results [Archer and Tardos 2002; Karlin et al. 2005; Elkind et al. 2004; Conitzer and Sandholm 2006; Hartline and Roughgarden 2009]. In contrast, since the prices in our natural market model happened to be VCG prices, we have no overhead for making our mechanism incentive compatible.

The next question is how to address the scheduling of landing slots over longer periods at multiple airports, taking into consideration inter-airport constraints. Our contention is that any such solution obtained by taking into consideration static information will be inferior since over a long period, unexpected events happening even at a few places are likely to have profound effects at geographically distant airports. For this reason, in Section 4, we propose a dynamic solution by decomposing this entire problem into many small problems, each of which will be solved by the method proposed above. The key to this decomposition is the robustness of our solution for a single set of flights at one airport: we have not imposed any constraints on delay costs, not even monotonicity. Therefore, airline companies can utilize this flexibility to encode a wide variety of inter-airport constraints.

We note that this approach opens up the possibility of making diverse types of travelers happy through the following mechanism: the additional revenues generated by FAA via our market gives it the ability to subsidize landing fees for low budget airlines. As a result, both types of travelers can achieve an end that is most desirable to them, business travelers and casual/vacation travelers. The former, in inclement weather, will not be made to suffer delays that ruin their important meetings and latter will get to fly for a lower price (and perhaps sip coffee for an additional hour on the tarmac, in inclement weather, while thinking about their upcoming vacation).

2. THE MARKET MODEL

In this section we will consider the problem of scheduling landings at one airport only. Let A be the set of all flights, operated by various airlines, that land in this airport in a given period of time. We assume that the given period of time is partitioned into a set of landing time slots, in a manner that is most convenient for this airport; let S denote this set. Each slot s has a capacity $cap(s) \in \mathbf{Z}^+$ specifying the number of flights that can land in this time slot. As mentioned in [Ball et al. 2005] the arrival of each aircraft consumes approximately the same amount of airport capacity, therefore justifying the slot capacities as the number of flights while ignoring their types. We will assume that $cap(s)$ is adjusted according to the prevailing weather condition.

For $i \in A$, the airline of this flight decides the *landing window* for flight i, denoted by $W(i)$. This gives the set of time slots in which this flight should land as per prevailing conditions, e.g., if there are no delays, the earliest time slot in $W(i)$ will be the scheduled arrival time of flight i. For each slot $s \in W(i)$, the airline also decides its *delay cost*, denoted by $c_{is} \geq 0$. Thus, if time slot s is the scheduled arrival time of flight i, then $c_{is} = 0$[3] and in general c_{is} is the dollar value of the cost, as perceived by the airline, for delay resulting from landing in slot s.

A *landing schedule* is an assignment of flights to time slots, respecting capacity constraints. Each time slot will be assigned a *landing price* which is the amount charged by FAA from the airline company if its flight lands in this time slot. We will define the *total cost* incurred by a flight to be the sum of the price paid for landing and the cost of the delay.

We say that a given schedule and prices are *an equilibrium landing schedule and prices* if:

(1) W.r.t. these prices, each flight incurs a minimum total cost.
(2) The landing price of any time slot that is not filled to capacity is zero. This is a standard condition in economics; the price of a good that is not fully sold must be zero.

2.1. LP formulation

In this section, we will give an LP that yields an equilibrium schedule; its dual will yield equilibrium landing prices. Section 3 shows how they can be computed in strongly polynomial time.

For $s \in S$, x_{is} will be the indicator variable that indicates whether flight i is scheduled in time slot s; naturally, in the LP formulation, this variable will be allowed to take fractional values. The LP given below obtains a fractional scheduling that minimizes the total dollar value of the delays incurred by all flights, subject to capacity constraints of the time slots.

[3] All the results of this paper hold even if $c_{is} \neq 0$.

$$
\begin{array}{ll}
\text{minimize} & \sum_{i\in A, s\in S} c_{is}x_{is} \\
\text{subject to} & \forall i \in A: \ \sum_{s\in W(i)} x_{is} \geq 1 \\
& \forall s \in S: \ \sum_{i\in A, s\in W(i)} x_{is} \leq cap(s) \\
& \forall i \in A, \ s \in W(i): \ x_{is} \geq 0
\end{array} \tag{1}
$$

Let p_s denote the dual variable corresponding to the second set of inequalities. We will interpret p_s as the price of landing in time slot s. Thus if flight i lands in time slot s, the total cost incurred by it is $p_s + c_{is}$. Let t_i denote the dual variable corresponding to the first set of inequalities. In Lemma 2.1 we will prove that t_i is the total cost incurred by flight i w.r.t. the prices found by the dual; moreover, each flight incurs minimum total cost.

The dual LP is the following.

$$
\begin{array}{ll}
\text{maximize} & \sum_{i\in A} t_i \ - \ \sum_{s\in S} cap(s)\cdot p_s \\
\text{subject to} & \forall i \in A, \ \forall s \in W(i): \ t_i \leq p_s \ + \ c_{is} \\
& \forall i \in A: \ t_i \geq 0 \\
& \forall s \in S: \ p_s \geq 0
\end{array} \tag{2}
$$

LEMMA 2.1. *W.r.t. the prices found by the dual LP (2), each flight i incurs minimum total cost and it is given by t_i.*

PROOF. Applying complementary slackness conditions to the primal variables we get

$$\forall i \in A, \ \forall s \in W(i): \ \ x_{is} > 0 \ \ \Rightarrow \ \ t_i = p_s \ + \ c_{is}.$$

Moreover, for time slots $s \in S$ which are not used by flight i, i.e., for which $x_{is} = 0$, by the dual constraint, the total cost of using this slot can only be higher than t_i. The lemma follows. □

The second condition required for equilibrium is satisfied because of complementarity applied to the variables p_s:

$$\text{If} \sum_{i\in A, s\in W(i)} x_{is} < cap(s), \ \text{ then } \ p_s = 0.$$

At this point, we can provide an intuitive understanding of how the actual slot assigned to flight i by LP (1) is influenced by the delay costs declared for flight i and how LP (2) sets prices of slots. Assume that time slot s is the scheduled arrival time of flight i, i.e., $c_{is} = 0$ and s' is a later slot. Then by Lemma 2.1, slot s will be preferred to slot s' only if $p_s - p_{s'} \leq c_{is'}$. Thus $c_{is'}$ places an upper bound on the extra money that can be charged for buying away the delay incurred by landing in s instead of s'. Clearly, flight i will incur a smaller delay, at the cost of paying more, if its airline declares large delay costs for late landing. Furthermore, by standard LP theory, the dual variables, p_s, will adjust according to the demand of each time slot, i.e., a time slot s that is demanded by a large number of flights that have declared large delay costs will have a high price. In particular, if a slot is not allocated to capacity, its price will be zero as shown above.

It is easy to see that the matrix underlying LP (1) is totally unimodular. Therefore, it has an integral optimal solution. Further, minimization ensures that for every flight i at most one of the x_{is}s is one and the rest are zero. Hence we get:

THEOREM 2.2. *Solution of LP (1) and its dual (2) give an (optimal) equilibrium schedule and equilibrium prices.*

3. STRONGLY POLYNOMIAL IMPLEMENTATION

Since the matrix underlying LP (1) is totally unimodular, it has an integral optimal solution. In this section, we show that the problem of obtaining such a solution can be reduced to a minimum weight perfect b-matching problem[4], and hence can be found in strongly polynomial time, see [Schrijver 2003] Volume A. The equilibrium prices, i.e., solution of (2), can be obtained from the dual variables of the matching. Furthermore, we show that there exists equilibrium prices that induce VCG payments, and hence is incentive compatible in dominant strategy. Finally, we give a strongly polynomial time procedure to compute such prices.

Consider the edge-weighted bipartite graph (A', S, E), with bipartition $A' = A \cup \{v\}$, where A is the set of flights and v is a special vertex, and S is the set of time slots. The set of edges E and weights are as follows: for $i \in A,\ s \in W(i)$, (i, s) is an edge with weight c_{is}, and for each $s \in S$, there are $cap(s)$ many (v, s) edges[5], each with unit weight (a multi-graph).

The matching requirements are: $b_i = 1$ for each $i \in A$, $b_s = cap(s)$ for each $s \in S$, and $b_v = \sum_{s\in S} cap(s) - |A|$ for v. Clearly, the last quantity is non-negative, or else LP (1) is infeasible. The following lemmas show that the equilibrium landing schedule and prices can be computed using minimum weight perfect b-matching of graph (A', S, E).

LEMMA 3.1. *Let $F^* \subset E$ be a perfect b-matching in (A', S, E) and x^* be a schedule where $x^*_{is} = 1$ if $(i, s) \in F^*$. F^* is a minimum weight perfect b-matching if and only if x^* is an optimal solution of LP (1).*

PROOF. To the contrary suppose x' and not x^* is the optimal solution of LP (1). Let $F' = \{(i, s) \in E \mid x'_{is} = 1\} \cup \{(cap(s) - \sum_{i;s\in W(i)} x'_{is})$ many $(v, s) \mid s \in S\}$ be the set of edges corresponding to schedule x'. Clearly, F' is a perfect b-matching. Note that the matching edges incident on v contribute cost b_v in any perfect b-matching. Since, x' and not x^* is an optimal solution of LP (1), we have,

$$\sum_{i\in A, s\in W(i)} c_{is}x'_{is} + b_v < \sum_{i\in A, s\in W(i)} c_{is}x^*_{is} + b_v \ \Rightarrow \ \sum_{(i,j)\in F'} c_{ij} < \sum_{(i,j)\in F^*} c_{ij}$$

Contradicting F^* being the minimum weight perfect matching. The reverse implication follows by similar argument in the reverse order. □

Using Lemma 3.1, next we show that the dual variables of the b-matching LP give an equilibrium price vector. In the b-matching LP there is an equality for each node to ensure its matching requirement. Let u_v, u_i and q_s be the dual variables corresponding to the equalities of nodes v, $i \in A$ and $s \in S$. Then the dual LP for minimum weight perfect b-matching in graph (A', S, E) is as follows.

$$\begin{array}{ll} \max: & \sum_{i\in A} u_i + \sum_{s\in S} cap(s)q_s + u_v b_v \\ \textbf{s.t.} & \forall i \in A,\ s \in W(i): u_i \le -q_s + c_{is} \\ & \forall s \in S: \qquad\qquad u_v \le -q_s + 1 \end{array} \tag{3}$$

There are no non-negativity constraints on the dual variables since the corresponding primal constraints are equality.

[4]The instance we construct can also be reduced to a minimum weight perfect matching problem with quadratic increase in number of nodes.
[5]This is not going to affect strong polynomiality, because we can assume that $cap(s) \le |A|, \forall s$ without loss of generality.

LEMMA 3.2. *There exists a dual solution (u^*, q^*) of (3) with $u^*_v = 1$, and given that, $-q^*$ yields a solution of LP (2).*

PROOF. If (u^*, q^*) is a dual solution then so is $v = (u^* + \delta, q^* - \delta)$ for any $\delta \in \mathbb{R}$. This is because, clearly v is feasible. Further, since $|A| + b_v = \sum_s cap(s)$ the value of objective function at v is same as that at (u^*, q^*).

Therefore given any solution of the dual, we can obtain one with $u^*_v = 1$ by an additive scaling. Replacing u_v with 1 and q_s with $-p_s$ in (3) gives $max\{\sum_i u_i - \sum_s cap(s)p_s + b_v \mid u_i \leq p_s + c_{is},\ p_s \geq 0\}$, which is exactly (2), and hence the lemma follows. □

Since a primal and a dual solution of a minimum weight perfect b-matching can be computed in strongly polynomial time [Schrijver 2003], the next theorem follows using Lemmas 3.1 and 3.2, and Theorem 2.2.

THEOREM 3.3. *There is a combinatorial, strongly polynomial algorithm for computing an equilibrium landing schedule and equilibrium prices.*

3.1. Incentive Compatible in Dominant Strategy

Since equilibrium prices of the market is in one-to-one correspondence with the solutions of the dual matching LP with $u_v = 1$ (Lemma 3.2), they and need not be unique, and in fact forms a convex set. In this section we show that one of them induces VCG payments, and therefore is incentive compatible in dominant strategy. Further, we will design a method to compute these prices in strongly polynomial time.

An instance of the perfect b-matching problem can be reduced to the perfect matching problem by duplicating node n, b_n times. Therefore, if we convert the costs c_{is} on edge (i, s) to payoffs $H - c_{is}$ for a big enough constant H, the market becomes an equivalent matching market (also known as assignment game) [Shapley and Shubik 1972] where the costs of producing goods, the slots in our case, are zero. It is not difficult to check that equilibrium allocations and prices of our original market and the transformed matching market exactly match.

For such a market, Leonard [Leonard 1983] showed that prices which are simultaneously minimum[6] for all goods exist and they correspond to VCG payments [Nisan 2007]. Clearly, such a price vector has to be unique. Next we give a procedure to compute the minimum equilibrium price vector, starting from any equilibrium price vector p^* and corresponding equilibrium schedule x^*.

The procedure is based on the following observation: Given an equilibrium prices p^* and corresponding schedule x^*, construct graph $G(x^*, p^*)$ where slots form the node set. Put a directed edge from slot s to slot s' if there is a flight i scheduled in s, and it is indifferent between s and s' in terms of total cost, i.e. $x^*_{is} = 1$ and $p^*_s + c_{is} = p^*_{s'} + c_{is'}$. An edge in graph $G(x^*, p^*)$ indicates that if the price of slot s' is decreased then i would prefer s' over s. Therefore, in order to maintain x^* as an equilibrium schedule the price of s also has to be decreased by the same amount.

LEMMA 3.4. *Prices p^{*m} gives the minimum equilibrium prices if and only if every node in $G(x^*, p^{*m})$ has a directed path from a zero priced node, where x^* is the corresponding equilibrium schedule.*

PROOF. Suppose slot s does not have a path from a zero priced node. Consider the set D of nodes which can reach s in $G^* = G(x^*, p^{*m})$; clearly, they have positive prices. Therefore, $\exists\epsilon > 0$ such that the prices of all the slots in D can be lowered by ϵ without violating the equilibrium condition (1), contradicting minimality of p^{*m}.

[6] Equilibrium prices p are minimum if for any other equilibrium prices p' we have $p_s <= p'_s,\ \forall s \in S$.

For the other direction, the intuition is that if every node is connected to a zero priced node in $G(x^*, p^{*m})$, then price of any slot can not be reduced without enforcing price of some other slot go negative, in order to get the corresponding equilibrium schedule. The formal proof is as follows:

To the contrary suppose every node is connected to a zero priced node in G^* and there are equilibrium prices $p' \leq p^{*m}$ such that for some $s \in S$, $p^{*m}_s > p'_s > 0$. Consider, one such s nearest to a zero-priced node in G^*. Since, $p'_s \geq 0$, we have $p^{*m}_s > 0$, and therefore s is filled to its capacity at prices p^{*m} (using equilibrium condition (2) of Section 2). Let x' be the equilibrium schedule corresponding to prices p'.

Let $s' \rightarrow s$ in G^*. By choice of s we have that $p'_{s'} = p^{*m}_{s'}$. In that case, a flight, say i', allocated to s' at p^* will move to s at p'. Implying that $\sum_i x'_{is'} < \sum_i x^*_{is'} \leq cap(s')$. Hence $p'_{s'} = 0 \Rightarrow p^{*m}_{s'} = 0$ (using equilibrium condition (2)). Let $Z = \{s \mid p^{*m}_s = 0\}$. There are two cases at this point:

Case I - Flights in slot s at x^* remains in s at x', *i.e.*, $\{i \mid x^*_{is} = 1\} \subseteq \{i \mid x'_{is} = 1\}$: Since, $x'_{i's} = 1$ and $x^*_{i's} = 0$, implying $\sum_i x'_{is} > \sum_i x^*_{is} = cap(s)$, a contradiction.

Case II - Some flight i scheduled in s at x^*, reschedules at x', *i.e.*, $\exists i,\ x^*_{is} = 1, x'_{is} = 0$: Construct a graph H, where slots are nodes, and there is an edge from u to v if $\exists i,\ x^*_{iu} = 1, x'_{iv} = 1$, *i.e.,* flight i moved from u to v when prices are changed from p^{*m} to p', with weight being number of edges moved. Note that price of every node with an incoming edge should have decreased while going from p^{*m} to p'. Therefore, nodes of Z have no incoming edges. Further, nodes with incoming edges are filled to capacity at p^{*m} since their prices are non-zero. In that case, total out going weight of such a node should be at least total incoming weight in H.

If there is a cycle in H, then subtract weight of one from all its edges, and remove zero-weight edges. Repeat this until there are no cycles. Since, $s' \in Z$, it had no incoming edge, but had an edge to s. Therefore, there is a path in remaining H starting at s'. Consider the other end of this path. Clearly, it has to be filled beyond its capacity at x', a contradiction. □

Using the fact established by Lemma 3.4 next we design a procedure to compute the minimum equilibrium prices in Table 3.1, given any equilibrium prices p^* and corresponding schedule x^*.

Table I. Procedure for Computing Minimum Optimal Prices

```
MinimumPrices(x*, p*)
 1. Z ← Nodes reachable from zero-priced node in G(x*, p*).
 2. Pick a d ∈ S \ Z
 3. D ← {Nodes that can reach d in G(x*, p*)}, δ ← 0,
    and p*_s ← p*_s − δ, ∀s ∈ D
 4. Increase δ until one of the following happen
    - If price of a slot in D becomes zero, then go to 1.
    - If a new edge appears in G(x*, p*), then recompute Z.
      If d ∈ Z then go to 2 else go to 3.
 5. Output p* as the minimum prices.
```

LEMMA 3.5. *Given an equilibrium (x^*, p^*), MinimumPrices(x^*, p^*) outputs minimum prices in time $O(|A||S|^2)$.*

PROOF. Note that the size of Z and edges in $G(x^*, p^*)$ are increasing. Therefore, Step 3 is executed $O(|S|)$ many times in total. Step 4 may need $O(|A||S|)$ time to compute the threshold δ. Therefore the running time of the procedure MinimumPrices is $O(|A||S|^2)$. Let the output price vector be p^{*m}. The lemma follows from the fact that (x^*, p^{*m}) still satisfy both the equilibrium conditions, and every slot is reachable from a zero priced node in $G(x^*, p^{*m})$ (Lemma 3.4). □

Theorems 2.2 and 3.3, Lemma 3.5, together with [Leonard 1983] give:

THEOREM 3.6. *There exists an incentive compatible (in dominant strategy) market mechanism for scheduling a set of flight landings at a single airport; moreover, it is computable combinatorially in strongly polynomial time.*

4. DEALING WITH MULTIPLE AIRPORTS

As stated in Section 1.1, we will decompose the problem of scheduling landing slots over a period of a day at multiple airports into many small problems, each dealing with a set of flights whose arrival times lie in a window of a couple of hours – the window being chosen in such a way that all flights would already be in the air and their actual arrival times, assuming no further delays, would be known to the airline companies and to FAA. At this point, an airline company has much crucial information about all the other flights that have constraints with its flight. It is therefore in a good position to determine how much delay it needs to buy away for its flight and how much it is willing to pay, by setting c_{is}s accordingly. This information is used by FAA to arrive at a landing schedule. The process is repeated every couple of hours at each airport.

REFERENCES

ARCHER, A. AND TARDOS, E. 2002. Frugal path mechanisms. In *ACM-SIAM Annual Symposium on Discrete Algorithms*. 991–999.

BALL, M., BARNHART, C., NEMHAUSER, G., AND ODONI, A. 2006. Air transportation: Irregular operations and control. In *Handbook of Operations Research and Management Science: Transportation*, C. Barnhart and G. Laporte, Eds.

BALL, M. O., DONOHUE, G., AND HOFFMAN, K. 2005. Auctions for the safe, efficient and equitable allocation of airspace system resources. In *Combinatorial Auctions*, P. Cramton, Y. Shoham, and R. Steinberg, Eds. MIT Press, Cambridge, 507–538.

BARNHART, C., BERTSIMAS, D., CARAMANIS, C., AND FEARING, D. 2009. Equitable and efficient coordination in traffic flow management. Submitted to *Transportation Science*.

BERTSIMAS, D., FARIAS, V., AND TRICHAKIS, N. 2011. The price of fairness. *Operations Research 59(1)*.

BERTSIMAS, D. AND GUPTA, S. 2010. A proposal for network air traffic flow management incorporating fairness and airline collaboration. Submitted to *Operations Research*.

CASTELLI, E., PESENTI, R., AND RANIERI, A. 2011. The design of a market mechanism to allocate air traffic flow management slots. *Transportation Research Part C 19*, 931–943.

COLE, R., DODIS, Y., AND ROUGHGARDEN, T. 2003. Pricing network edges for heterogeneous selfish users. In *ACM Symposium on the Theory of Computing*. 521–530.

CONITZER, V. AND SANDHOLM, T. 2006. Failures of the VCG mechanism in combinatorial auctions and exchanges. In *AAMAS*. 521–528.

ELKIND, E., SAHAI, A., AND STEIGLITZ, K. 2004. Frugality in path auctions. In *ACM-SIAM Annual Symposium on Discrete Algorithms*. 701–709.

FLEISCHER, L., JAIN, K., AND MAHDIAN, M. 2004. Tolls for heterogeneous selfish users in multicommodity networks and generalized congestion games. In *ACM Symposium on the Theory of Computing*. 277–285.

HARTLINE, J. D. AND ROUGHGARDEN, T. 2009. Simple versus optimal mechanisms. In *EC*.

KARLIN, A. R., KEMPE, D., AND TAMIR, T. 2005. Beyond VCG: Frugality of truthful mechanisms. In *IEEE Annual Symposium on Foundations of Computer Science*. 615–624.

LEONARD, H. B. 1983. Elicitation of honest preferences for the assignment of individuals to positions. *Journal of Political Economy 91,* 3, 461–479.

NISAN, N. 2007. Introduction to mechanism design (for computer scientists). In *Algorithmic Game Theory*, N. Nisan, T. Roughgarden, E. Tardos, and V. V. Vazirani, Eds. Cambridge University Press, 209–241.

ODONI, A. 1987. The flow management problem in air traffic control. In *A. Odoni, G. Szego, eds.,* Flow Control of Congested Networks*, Springer Verlag, Berlin*.

SCHRIJVER, A. 2003. *Combinatorial Optimization*. Springer-Verlag.

SHAPLEY, L. S. AND SHUBIK, M. 1972. The assignment game i: The core. *International Game Theory 1,* 2, 111–30.

SMITH, A. 1776. *The Wealth of Nations*.

VOSSEN, T. AND BALL, M. 2006. Slot trading opportunities in collaborative ground delay programs. *Transportation Science 40*, 29–43.

WAMBSGANSS, M. 1996. Collaborative decision making through dynamic information transfer. *Air Traffic Control Quarterly 4*, 107–123.