Landing slot allocation mechanisms in the air transportation system

Jackie Baek and Karthik Gopalakrishnan

June 9, 2017

1 Introduction

The number of aircrafts using the air transportation system is ever increasing. As the demand on the airspace and airport resources increase, there is no substantial increase in terms of the capacity. The rate at which new runways are being added or new airports being opened is not sufficient to cater to this increased demand, especially during peak operation hours or during low capacity periods.

Consider a single airport. During normal operations, flights get a clearance to land whenever they arrive at the airspace near an airport. However, during inclement weather, the Federal Aviation Administration (FAA), will control the allowed landing times of flights at an airport in accordance with the predicted arrival rate. For example, at Boston Logan Airport, the nominal arrival capacity is 60 landings per hour; but this rate may decrease to 40 or 30 in poor conditions.

A particular time when an aircraft is allowed to land is called a *slot*. During the period when visbility is low, the FAA will enforce a *slot-control*. Consider an example in Figure 1. The slots on the left are the flights in its original schedule that are affected by a slot control from 6:00 to 6:15 pm. These flights are then redistributed to the new slots on the right-hand side. The current practice by the FAA is to assign the slots in the same order as the scheduled flights. This is called the First Scheduled First Served (FSFS) allocation. This is the currently acceptable notion of fairness among the affected aircrafts. After this allocation, the airlines 'own' these new slots. So in our example, United Airlines is allotted slots 3 and 6, and now 'owns' these slots at 6:06 pm and 6:15 pm.

There are a few important observations we would like to make:

- O1 : Different flights have different sensitivities towards delays. An aircraft that has a lot of connecting passengers, or an aircraft that has a higher number of passengers may be more valuable to an airline than an aircraft that is only half-full and few connecting passengers. Further, different airlines may also have a different value per minute of delay based on their different operating practice.

- O2 : The delay sensitivity, or the value of every flight, is a private quantity. Airlines do not reveal the valuations to maintain competitiveness and to avoid disclosing their operation strategies.

- O2 : Every flight has a scheduled time of arrival, which is the time that is initially published on tickets. An important property that any slot allocation procedure must maintain is that the revised time slot for each flight must no earlier than its scheduled arrival time. This is important, because flight UA2 cannot be expected to take slot 2 which is at 6:03 pm as it

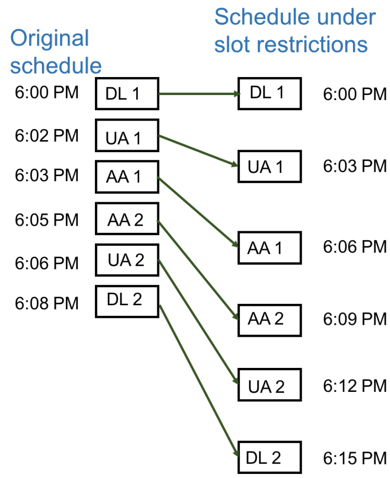

Figure 1: An example of a reduced capacity scenario from 6:00 pm to 6:15 pm where the minimum separation is prescribed as 3 minutes. Slot reassignment via a First Scheduled First Served (FSFS) policy.

would require it to take-off earlier than scheduled (and passengers would not have arrived at the airport by then).

Different private values for each aircraft introduces the opportunity for slot swaps. The focus of our project is to analyze mechanisms that facilitate this exchange. In particular, we are interested in exchange markets that meets the following requirements:

R1 : The market is a *secondary market* where airlines exchange slots after they are initially allotted by the FAA in the FSFS process. This is to enable compatibility with existing mechanisms. Nothing needs to change as per current procedures; simply, a new step is added in which airlines can trade the slots that they own using the secondary market.

R2 : For simplicity, we would like an exchange mechanism without any involvement of money.

R3 : We assume that airlines would have done any possible *intra airline swaps* to minimize the impact of the delay. In the example above, American Airlines has two flights, AA1 and AA2, that are scheduled to land at 6:06 and 6:09 pm. If the cost per minute of delay from schedule was \$1 for AA1 and \$2 for AA2, then American Airlines can swap the two slots internally. This would reduce the total cost for American from \$11 to \$8 . Note that this is a feasible swap, since the scheduled time for AA2 is before 6:06 pm (Observation *O3*). However, United (UA) or Delta (DL) cannot make such swaps since the swap would lead to a flight taking off before schedule. Hence, we only want to develop a market for *inter-airline* exchanges.

Table 1 is an example of a swap that we want to facilitate. In this example, United (UA) and American Airlines (AA) can trade slots which will to cost savings for both of the airlines. In Section 3 we discuss mechanisms that would achieve this.

2 Prior Literature

We first descibe two mechanisms that involve payments. In [CPR11], the authors propose two iterative methods to achieve the optimal social welfare: a Lagrangian relaxation approach and a

Airline	Flight	Original Schedule	FSFS Schedule	Delay Sensitivity	Initial Cost	Swap Schedule	Final Cost
AA	AA1	6:03	6:06	\$2/min	6+4=\$10	6:03	0+7=\$7
	AA2	6:05	6:09	\$1/min		6:12	
UA	UA1	6:02	6:03	\$5/min	5+60=\$65	6:06	20+30=\$50
	UA2	6:06	6:12	\$10/min		6:09	

Table 1: Example of a swap between airlines that decreases delay cost for both

Bertsekas auction method. These methods are individually rational, budget balanced and coalition-rational approximations to this optimal allocation; however, they are not incentive compatible, as acknowledged by the authors.

Another approach that involves side payments is an approximate Vickery based payment called the threshold rule [PKE01, Bal07]. This payment rule results in a mechanism that is individually rational and budget balances but only 'fairly optimal' and 'fairly incentive compatible'. Note that a Vickery auction can ensure IR, IC and maximum social welfare, but will not budget balanced (meaning that the FAA may lose money if it runs the secondary slot substitution market).

Our objective is to not involve payments, so that we may practically transition easily into any new mechanism we propose without challenges. [VB06] is the main motivation for our work. The authors propose a 2-for-2 slot exchange model in which airline submits any number of bids to the FAA of the following form: (f_u, s_u, f_d, s_d). This bid implies that the airline would like f_u to move up to at least slot s_u in exchange for flight f_d moving down to at most slot s_d. After receiving bids in this format, the FAA may use different objective functions to decide which trades to execute. For example, the FAA may want to maximize the number of trades that get executed, or maximize the upward movement of flights. An important fact that we observed, but was not mentioned in the paper was that the *mechanism is not incentive compatible*. This can be demonstrated using a simple example shown in Figure 2.

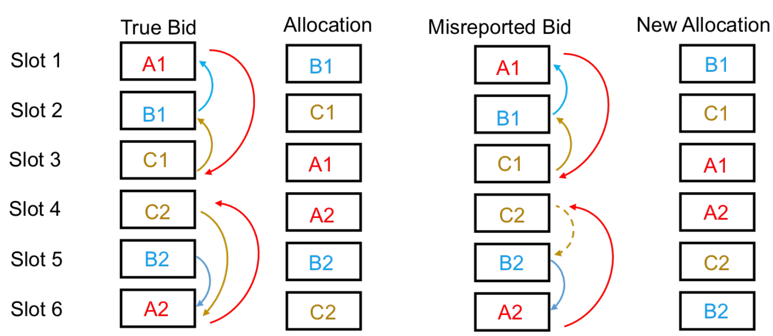

Figure 2: Example of misreporting a bid increasing payoff: There are three airlines with two flights each. The first column shows the initial schedule, and the solid lines represent the true bids for the 6 flights. The second column shows a possible allocation after trades are executed. In the third column, Airline C misreports, shown by the dotted line, and results are shown in the last column. In this scenario, Airline C strictly benefits from misreporting than telling the truth.

This example shows that the mechanism is not dominant-strategy incentive compatible - if Airline C knew the bids of all the other airlines, then it is better off misreporting. However, it is not clear whether this type of scenario can occur if Airline C did not know the bids of the other airlines. In our simulations, we show that if both the schedules and valuations of flights are randomized, then the airline is better off, in expectation, to misreport.

In our project, we will restrict ourselves to a simple setting where each airline has only two flights that are slot controlled. With this assumption, we can set up a single dimensional environment where each airline just needs to bid one quantity, the number of slots it is willing to move an aircraft down for each upward movement in the other aircraft (which is basically the ratio of their values). Different objective functions- maximizing number of swaps, maximizing total number of upward movements and minimizing the total 'scaled delay cost' will be considered. In addition, we will compare the benefits that would be obtained on using an incentive compatible scheme- the Top Trading Cycle algorithm.

3 Mechanisms for slot allocation

We use the following notation. There are N airlines with two flights each. F is the set of all flights, S is the set of all slots that. For each flight $f \in F$, the scheduled time of arrival is SCH_f (the earliest possible time that flight f can land). t_i is the time assigned for slot s_i, $t_1 < t_2 < \ldots < t_{2N}$. Each flight $f \in F$ has a value v_f, which is the cost for delaying flight f for 1 unit time.

3.1 Social Optimal Allocation

The optimal social welfare is found by solving:

$$\min \sum_{i \in F j \in S} x_{ij} c_{ij} \tag{1}$$

$$\sum_{j \in S} x_{ij} = 1 \quad \forall i \in F \tag{2}$$

$$\sum_{i \in F} x_{ij} \leq 1 \quad \forall j \in S \tag{3}$$

$$c_{i,j} = \begin{cases} v_i(t_j - SCH_i) & \text{if } t_j - SCH_i \geq 0 \\ +\infty & \text{otherwise} \end{cases} \tag{4}$$

$$x_{ij} \in \{0, 1\} \tag{5}$$

x_{ij} is 1 when flight i is assigned to slot s_j and c_{ij} is the cost of assigning flight i to slot j. The constraints enforce that every flight has a slot allocated and no slot has more than one flight assigned to it. The c_{ij} constraints prevents the assignment flights to slots before its scheduled arrival time. Note that this allocation favors high-valued flights over low-valued flights. It is obvious that this allocation scheme is not incentive compatible, as reporting a higher value will always result in a better allocation for that flight.

3.2 Modified 2-for-2 Swaps

Every airline submits a bid (f_u, f_d, b), where f_u is the high valued flight (the one that wants to move up), f_d is the lower value flight and b is the reported ratio of the high value to the low value. The b ratio is interpreted as: "For every slot that f_u moves up, I am willing to move down flight f_d by at most b". This bidding mechanism protects the private information about the airlines valuation for each flight by asking them to only report the ratio. The ratio also results more equitable allocation to airlines with small aircrafts that may inherently have lower costs than an airline with larger aircrafts. $FSFS_f$ is the index of the slot assigned to the flight f under the FSFS policy. $EFS_f \coloneqq \underset{i}{argmin}\{t_i - SCH_f \geq 0\}$ is the index of the Earliest Feasible Slot for flight f.

$$U = \{f : f = f_u \text{ for some airline}\} \tag{6}$$
$$D = \{f : f = f_d \text{ for some airline}\} \tag{7}$$
$$\tag{8}$$

U and D are the set of flights that are high value and low value respectively. All airlines $a \in \{1, 2, \ldots, N\}$ submit the bid (f_u^a, f_d^a, b^a) and the central planner (FAA in this case) will decide which trades to execute.

$\forall f \in U$, $x_f \in \mathbb{N}$ is the number of slots that flight f is moved up by, and $\forall f \in D$, $y_f \in \mathbb{N}$ is the number of slots that flight f is moved down by. $s_f = \{1, 2, \ldots, 2N\}$ is the index of the slot that flight f is assigned after the exchange. We will now describe the optimization problem that the central planner follows.

Constraints that determine slot time:

$$s_f = FSFS_f - x_f \quad \forall f \in U \tag{9}$$
$$s_f = FSFS_f + y_f \quad \forall f \in D \tag{10}$$
$$s_f \geq EFS_f \quad \forall f \in F \tag{11}$$

Ensuring that the bid ratio constraint is maintained:

$$y_{f_d^a} \leq b^a x_{f_u^a} \quad \forall a \in \{1, \ldots, N\} \tag{12}$$

No two flights must be assigned the same slot, i.e. $s_{f_1} \neq s_{f_2}, \forall f_1 \neq f_2$. Equivalently, $|s_{f_1} - s_{f_2}| \geq \frac{1}{2}, \forall f_1 \neq f_2$. As linear constraints, for all $f_1 \neq f_2$ and a large M, we have

$$s_{f_1} - s_{f_2} \leq -\frac{1}{2} + My_{f_1 f_2} \tag{13}$$
$$s_{f_1} - s_{f_2} \geq \frac{1}{2} - M(1 - y_{f_1 f_2}) \tag{14}$$
$$y_{f_1 f_2} \in \{0, 1\} \tag{15}$$

A possible objective function could be to maximize the number of upward movements.

$$\max \sum_{f \in U} x_f \tag{16}$$

Another possible objective could be to maximize the number of trades, i.e Maximize $\sum_{f \in U} \mathbf{1}\{x_f > 0\}$. This can be written in the form of linear constraints as

$$\max_{f \in U} z_f \tag{17}$$

with additional constraints

$$z_f \leq Mx_f \quad \forall f \in U \tag{18}$$
$$z_f \in \{0, 1\} \quad \forall f \in U \tag{19}$$

Thus, the optimization for maximizing number of trades would be

Objective: (17)
Constraints: (9)-(15)

Objective for maximizing the total upward movement would be

Objective: (16)
Constraints: (9)-(15), (18)-(19)

3.3 Token based allocation

The idea here is that instead of maximizing the number of trades or the upward movement, we want to minimize the delay cost of the system. Since we want to maintain fairness across airlines, we give an airline one *token* for each flight they have in the system. Then, the airlines redistribute the tokens among their flights in accordance with the values for their different flights. Effectively, the airlines declare a value for each of their flights such that the mean value is 1.

Since we are assuming that each airline only has two flights, we can use the same bidding language for this mechanism as in Section 3.2. When an airline bids (f_u, f_d, b), the central planner can interpret the values of the flights as $v_{f_u} = \frac{2b}{1+b}$ and $v_{f_d} = \frac{2}{1+b}$. These values ensure that the ratio of the values of the flights is b, and the mean cost of the two flights if 1. Then, the optimization that minimizes the 'scaled delay cost' is

$$\min \sum_{f \in F} v_f s_f \tag{20}$$

subject to constraints (9)-(15) and $x_f, y_f \in \mathbb{Z}, \forall f \in F$.

3.4 Top Trading Cycles (TTC)

The last method we describe is the Top Trading Cycles (TTC) algorithm. This algorithm, developed in [SS74], is a method for trading objects among participants with different preferences without any payments. In the context of slot exchange, using this mechanism, airlines give away both of their slots together for two slots of another airline in return. This mechanism is more restrictive than the previous ones described because once two slots are initially allocated using FSFS, those two slots always belong to one airline.

For this to be implemented, each airline must submit a ranked preference of which airline's slots are the most desirable for them. The information of which slots belong to which airline must be public knowledge, but the valuations of the individual flights need not be shared. Given the list of preferences from every airline, the FAA uses the TTC algorithm to allocate the slots. The resulting allocation is such that no airline would need to trade with any airline that is ranked below its current assignment. The TTC algorithm is individually rational, and proven to be incentive compatible.

3.5 On the efficiency, incentive compatibility, and fairness of the mechanisms

In terms of social welfare, the Social Welfare Maximizing formulation in Section 3.1 is the theoretical maximum attainable. However, this need not a good benchmark for practical purposes. This formulation has a strong bias towards higher-valued flights, and therefore airlines always have incentive to report a high value.

Since FSFS is the currently practiced policy by the FAA, it serves as an ideal benchmark to compare the gains in social welfare using other methods.

Out of all the mechanisms described in the previous sections, only the TTC mechanism is incentive compatible. There is no benefit that any airline will get in lying about their preference about the slots it wants to obtain in a swap. The other mechanisms, which all report a cost ratio, fail in a similar manner as seen earlier in Figure 2. The strong incentive guarantees, along with the individual rationality and strong obfuscation of private information is the strength of the TTC algorithm in this problem's context.

Equality or fairness in the slot assignment has also been overlooked in the literature. One of the objectives that the FAA must consider when proposing a secondary market mechanism is to spread the benefits to all participating airlines. Maximizing the number of trades and the token mechanism try to achieve this objective in an indirect way.

In the next section, we run simulations on the methods described above and assess their efficiency, incentive compatibility, and fairness.

4 Simulation

4.1 Overall comparison

We compare the efficiency of all the methods in our first simulation assuming that everyone is truthful. Consider 10 airlines that own two flights each. Each of the 20 flights have a private value v_f drawn from a normal distribution with $\mu = 1$ and $\sigma = 0.4$. The initial schedule of the flights were 0.5 minutes apart, and the redistributed slots are 1 minute apart. In a FSFS allocation, the k^{th} flight gets a delay of $\frac{1}{2}k$ minutes and ends up with a delay cost of $\frac{1}{2}v_k k$. The schedules of the 10 airlines are picked randomly.

We run 1000 trials where we randomize over valuations and schedules. The average costs over these trials are shown in Table 2.

Mechanism	Total social cost	Improvement over FSFS
FSFS	95.6	-
Social Optimal	64.9	32.1%
Token mechanism	78.0	18.3%
Modified 2-2, max # trades	85.0	11%
Modified 2-2, max upward moves	80.7	15.5%
TTC	88.7	7.2%

Table 2: Mean social cost of different slot trading mechanisms (averaged over 1000 schedule and value realizations)

When bids are truthful, then the Token mechanism obtains about 18% improvement over the current practice of FSFS. The TTC algorithm, which is strategy-proof, gives a more modest 7% improvement only. This loss in social welfare is the price paid to achieve incentive compatibility. Among the 2-for-2 swap mechanisms, maximizing the number of upward movements leads to lower social cost compared to maximizing the number of trades.

4.2 Fairness and equitable distribution of benefits

Which airlines benefit the most from these mechanisms? In Figure 3, we assume everyone reports truthfully and plot the savings of the mechanisms (relative to FSFS) for one airline against the ratio of the airline's flights. The three methods in Figure 3a (2-for-2 with maximizing upward movement, Token and TTC mechanism) show an increasing trend, meaning that airlines with a higher bid get more benefits. This is not surprising. As all these mechanisms try to minimize the delay cost (maximize the savings), the best way to do it would be to reduce the delays for the most sensitive flights.

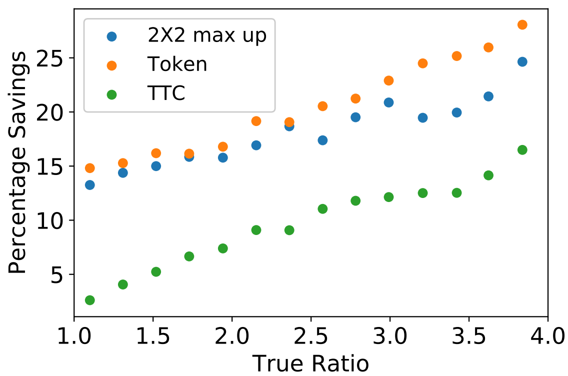

(a) Comparison of the cost savings for three methods

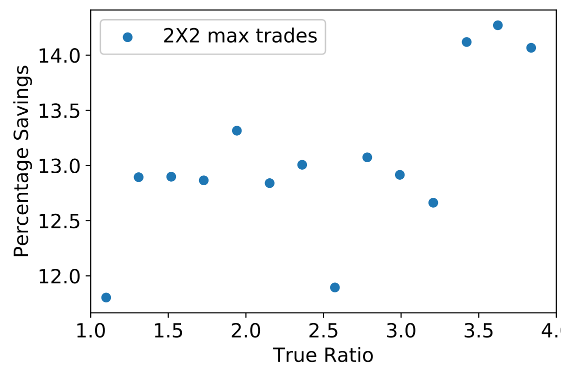

(b) Savings in cost for the Modified 2-for-2 swap that maximizes the number of trades

Figure 3: Percentage improvement over the FSFS allocation for airlines with different true bid ratios

It is interesting that the 2-for-2 mechanism with the objective of maximizing the number of trades does not show such a strong bias towards high bid airlines (Figure 3b). The nature of the objective function gets the most number of different airlines that get to trade slots. Thus, airlines with low bids, that would usually get less chance of getting a trade executed now have a higher chance, and consequently get greater savings in expectation. Note that all these plots are averages over 1000 realizations of the schedule and values.

4.3 How manipulable are the mechanisms?

In this series of experiments, we study the effects of non-truthful bidding. We know that TTC is the only mechanism that is incentive compatible. For the other mechanisms, reporting a smaller bid, which is a smaller willingness to move down a slot to gain another, could be beneficial to the airline. As long as a slot is going to be allocated, bidding smaller ratios would mean that the flight f_d was not moved down as much as it could have been in case of a true bid.

We run simulations assuming all all airlines are truthful except one. We have the same setup as Section 4.1 with 10 airlines that have two flights each, but we take the first airline and fix its true ratio to be 2. Then, this first airline misreports and bids a range of values from 0.5 to 3, while the other airlines bid their true value. The non-truthful airline will get a slot allocated based on its reported bid, and we compute its delay cost. This is repeated over multiple realizations of random schedules and costs for other airlines.

The results of the simulations are shown in Figure 4. In all of the figures, the x-axis is the reported bid of the first airline, and the y-axis is the percentage savings of cost (either to the misreporting airline or total social cost) in comparison to the FSFS allocation. Figure 4a and 4b show the savings for the misreporting airline with the Token and TTC mechanism respectively. The next two figures, Fig. 4c and 4d show the savings in the social cost (sum of all airlines) with respect to the FSFS allocation. We demonstrate the non-truthfulness of just the Token scheme for simplicity, but similar results hold for the other non-truthful mechanisms as well.

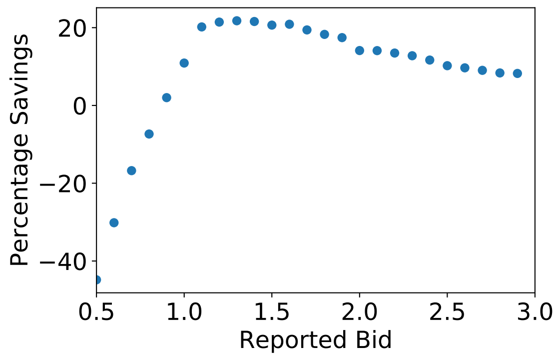

(a) Cost savings for the misreporting airline under the Token mechanism

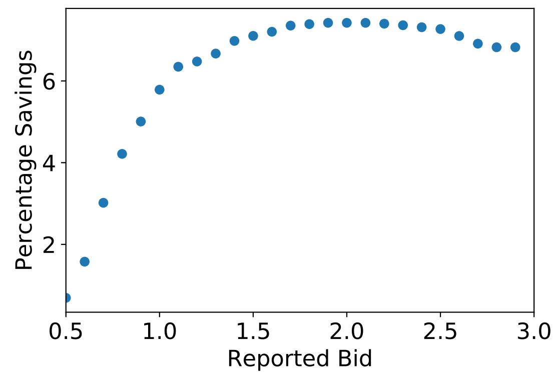

(b) Cost savings for the misreporting airline under the TTC mechanism

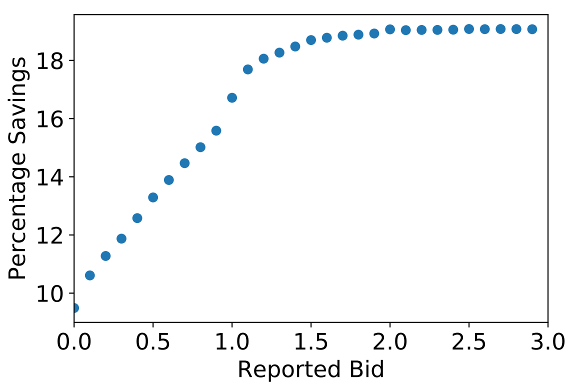

(c) Social cost for when an airline misreports under Token mechanism

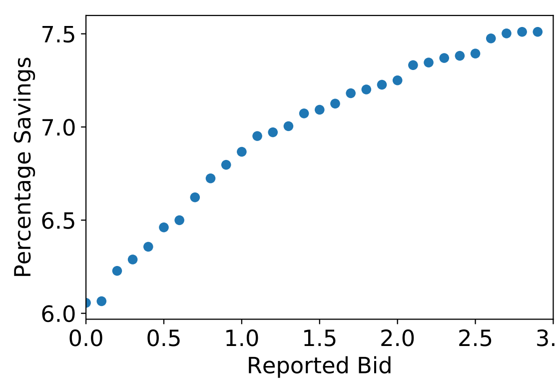

(d) Social cost for when an airline misreports under the TTC mechanism

Figure 4: Effect of misreporting bids: One airline with a true ratio of 2 reports different bids between 1 and 3

1. As predicted for TTC, an airline is always better off bidding truthfully. The peak of the savings, in Figure 4b is at 2, which is its true bid radio. The social cost decreases (increases savings) when the airline bids a higher value. When that happens it is sacrificing its profits and improving the system performance by providing more trading opportunities. On the other hand, when the airline under reports its bid, both the airline and the savings for the society decreases; its a lose-lose scenario for the airline and the system as a whole.

2. The Token mechanism clearly shows the benefits of under-reporting the bid. The optimal bid to minimize delay costs is 1.3, while the true value is 2. The social cost is monotonic in the bid of the non-truthful airline. Initially, when an airline starts under reporting, the airline gains (until it reaches the peak value) whereas the societal costs increase. However, under-reporting a lot increases the cost for everyone. If the airline bids higher, it has higher costs whereas the system costs decrease.

5 Conclusions and comments

We have studied the effectiveness of secondary slot markets that involve no payments with respect to efficiency, fairness, and truthfulness to airlines. Using the 2-for-2 slot exchange model introduced in [VB06] as a basis, we formalize three optimization models: maximizing number of trades, maximizing number of upward movements, and the token mechanism. Lastly, we also consider the Top Trading Cycles (TTC) mechanism, which is already known to be strategy-proof. We have showed that none of the mechanisms other than TTC are incentive compatible, but TTC has the smallest benefits in terms of efficiency. In terms of equity, maximizing the number of trades shows the least bias towards airlines with a high variance in costs of their flights.

In our work, we only consider the models in the specialized case when each airline only has two flights. There are many ways in which these models can be extended to more flights, like in the more general 2-for-2 slot exchange model described in [VB06]. However, the purpose of our simulations was to show that even in these very simplified scenarios, the mechanisms are not strategy-proof nor equitable. One direction for future work is to analyze how generalizing the models to more flights impacts the fairness and incentive compatibility to different airlines.

Another clear research direction is to either extend the TTC algorithm to the more general setting, or to design another mechanism that is incentive compatible. Our results suggest that system efficiency must be forgone in order to achieve incentive compatibility. It may also be worthwhile to assess, among the non-incentive compatible mechanisms, how much worse the social cost can be if everyone is non-truthful in some sort of Bayes Nash Equilibrium.

References

[Bal07] Hamsa Balakrishnan. Techniques for reallocating airport resources during adverse weather. In *Decision and Control, 2007 46th IEEE Conference on*, pages 2949–2956. IEEE, 2007.

[CPR11] Lorenzo Castelli, Raffaele Pesenti, and Andrea Ranieri. The design of a market mechanism to allocate air traffic flow management slots. *Transportation research part C: Emerging technologies*, 19(5):931–943, 2011.

[PKE01] David C Parkes, Jayant R Kalagnanam, and Marta Eso. Achieving budget-balance with vickrey-based payment schemes in exchanges. 2001.

[SS74] Lloyd Shapley and Herbert Scarf. On cores and indivisibility. *Journal of mathematical economics*, 1(1):23–37, 1974.

[VB06] Thomas WM Vossen and Michael O Ball. Slot trading opportunities in collaborative ground delay programs. *Transportation Science*, 40(1):29–43, 2006.