33RD CONGRESS
OF THE INTERNATIONAL COUNCIL
OF THE AERONAUTICAL SCIENCES
STOCKHOLM, SWEDEN, 4-9 SEPTEMBER, 2022

ICAS 2022 SWEDEN

AN AUCTION-BASED MECHANISM FOR A PRIVACY-PRESERVING MARKETPLACE FOR ATFM SLOTS

Christoph G. Schuetz[1], Sergio Ruiz[2], Eduard Gringinger[3], Christoph Fabianek[3] & Thomas Loruenser[4]

[1]Johannes Kepler University Linz, Linz, Austria
[2]EUROCONTROL Innovation Hub, Brétigny-sur-Orge, France
[3]Frequentis AG, Vienna, Austria
[4]AIT Austrian Institute of Technology, Vienna, Austria

Abstract

In case of reduced capacity of the European air traffic network, the Network Manager (NM) initiates a regulation, causing flight delay. The flights are issued air traffic flow management (ATFM) slots according to the principle “first-planned, first-served”. For airspace users, however, different flights have different priorities due to the individual cost structures of different flights. In this regard, the SlotMachine system will allow airspace users to submit preferences regarding the arrival or departure times of flights, which are then considered during a privacy-preserving optimization run that aims to find an optimal flight list while keeping the preferences a secret, even to the operator of the SlotMachine. In order to provide airspace users with an incentive to participate in an optimization run and submit truthful preferences, an appropritate market mechanism is required, which handles compensation for airspace users giving up favorable ATFM slots. In this paper, we present an auction-based market mechanism for the SlotMachine system with credits instead of real-world currency.

Keywords: Air Traffic Flow Management, Mechanism Design, Genetic Algorithm, Multi-Party Computation

1. Introduction

In case of reduced capacity of the European air traffic network, the Network Manager (NM) initiates a regulation, which typically causes flight delay; flights are issued air traffic flow management (ATFM) slots. In the event of a regulation, the conventional approach to re-planning the departure or arrival times of flights, i.e., to assigning ATFM slots, is to follow the principle of “first-planned, first-served”. For airspace users, however, different flights may have different priorities due to the individual cost structures of different flights. Those cost structures, however, are confidential information for the airspace users, which they are reluctant to share. Thus, in order to obtain an optimal flight list, airspace users must be provided the means to conduct optimization of the flight lists in a privacy-preserving manner.

The SlotMachine system, which is currently being developed in a collaborative effort between Frequentis AG, Johannes Kepler University Linz, AIT Austrian Institute of Technology, Swiss International Airlines, and EUROCONTROL, will allow each airspace user (AU) to submit preferences regarding the arrival or departure times of individual flights in encrypted form, which are then considered during a privacy-preserving optimization run that aims to find an optimal flight list. From a conceptual perspective, the flow of a SlotMachine optimization run is as follows (Fig. 1). Once the NM announces a regulation, the SlotMachine system retrieves the new first-planned, first-served (FPFS) flight list and any non-confidential data related to the flights affected by the regulation. The SlotMachine system

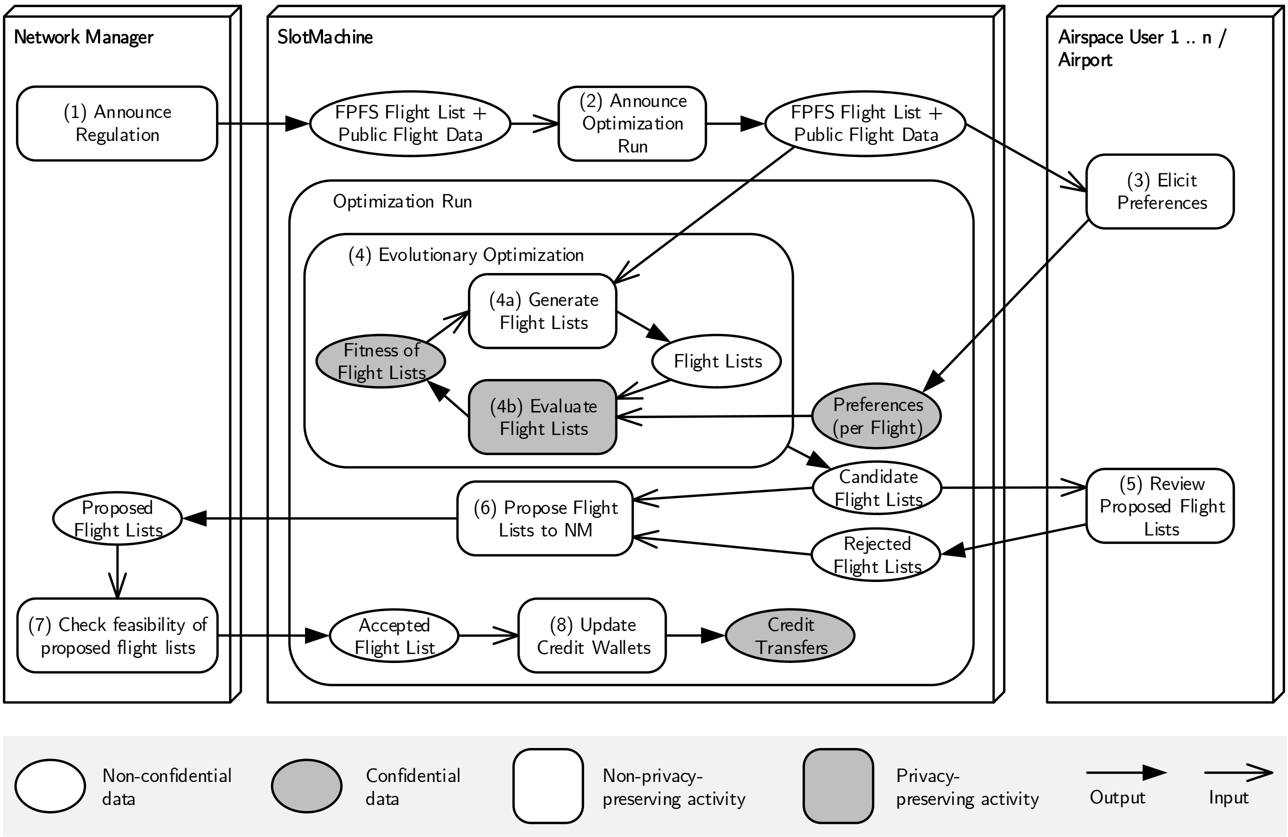

Figure 1 – Conceptual representation of the flow of a SlotMachine optimization run

then initiates an optimization run for the regulation [1]. The SlotMachine system may conduct multiple optimization runs per regulation, so the optimization run may only involve a subset of the flights affected by the regulation. The AUs (and possibly the airport) are notified about the regulation and receive a flight list for the time segment that the optimization run covers [2]. The AUs (and possibly the airport) then elicit preferences regarding the prioritization of the flights that are affected by the regulation. The preferences are confidential and, therefore, submitted to the SlotMachine system in encrypted form; the activities that process the preferences are privacy-preserving. Following the submission of the preferences by the AUs, the SlotMachine system conducts an evolutionary optimization run to find optimized candidate flight lists. The optimization run is an evolutionary process that consists of finding a population of flight lists and evaluating the flight lists. Evaluation of the flight lists in a population of flight lists is a privacy-preserving activity, employing (secure) multi-party computation (MPC) to compute fitness values for the proposed flight lists. After the evolutionary optimization run has finished, the best candidate flight lists are submitted to the AUs and the airport for review. The AUs and airport may reject individual solutions. The acceptable candidate flight lists are proposed to the NM. The NM checks the proposed flight lists' feasibility, chooses a flight list, and communicates the accepted flight list back to the SlotMachine system, which relays the accepted flight list back to AUs and airport before updating the credit wallets if a credit-based market mechanism is used. In this regard, the question is how to design this credit-based market mechanism.

In this paper, we present an auction-based market mechanism with credits instead of real-world currency, which serve to compensate AUs giving up favorable ATFM slots to other AUs that more urgently require an earlier slot. The proposed mechanism was briefly discussed in previous work [2], but this paper provides a thorough definition of the mechanism. AUs that would like to prioritize a flight issue bids for slots in terms of credits. For flights that still have margin AUs may specify how far the flight could be pushed back. The SlotMachine then finds a flight list that best fits the bids and margins submitted by the AUs. Based on those bids and margins, the SlotMachine computes the amount of credits that have to be exchanged between the different AUs.

The remainder of this paper is organized as follows. Section 2.provides background information and refer to related work. Section 3.describes the mechanism. Section 4.discusses privacy-preserving implementation of the SlotMachine system and the mechanism. Section 5.concludes the paper.

2. Background and Related Work

Each flight has its own individual, typically non-linear cost function in relation to the delay time (Fig. 2). Flights then have different delay targets in the event of a regulation. After a delay target is missed, the costs often bounce upwards considerably. However, pushing back a flight within the margins of the delay target is typically unproblematic.

In the event of a regulation flights could be reordered based on their non-linear cost functions. Figure 2 shows the cost functions of three flights (Flights A–C), with different delay targets, the costs considerably bouncing upwards in case a delay target is missed. The left side shows the original ATFM slots issued to the different flights as dashed lines. Hence, Flight C departs before Flight B, followed by Flight A. Using the cost functions as the basis, a flight prioritization would produce a different flight list, where Flight B departs before Flight A, followed by Flight C.

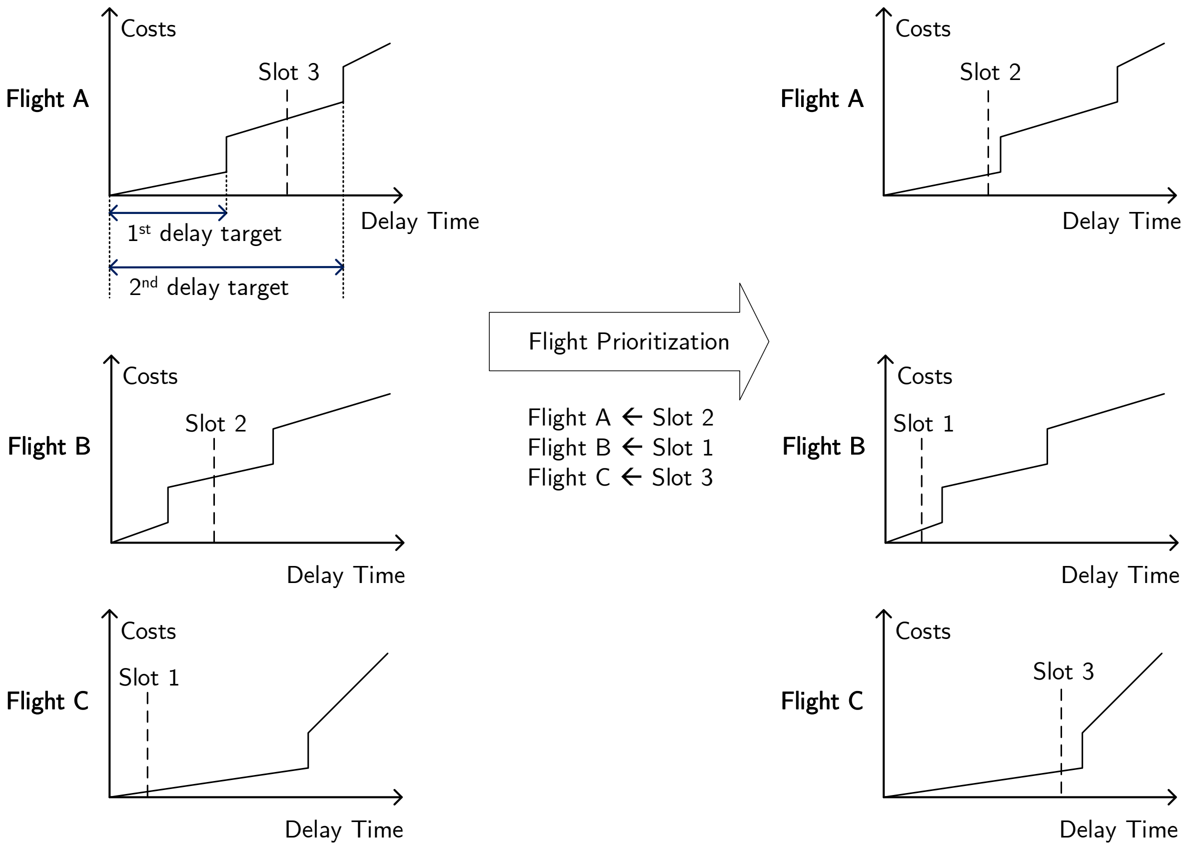

Figure 2 – Typical non-linear cost functions and flight prioritization based on cost functions

The SlotMachine flight prioritization problem can be considered a permutation problem or an assignment problem [3]. If considered a permutation problem, the optimization goal is to find the optimal sequence of flights. Ultimately, the SlotMachine flight prioritization problem boils down to finding a bipartite matching between a set of flights and a set of ATFM slots, with the constraint that each flight must be matched to exactly one slot and each slot can be matched to at most one flight. Furthermore, the slot assigned to a flight must not have a time earlier than the flight's originally scheduled departure time (before the regulation).

The User-Driven Prioritization Process (UDPP) [4] allows AUs to maximize utility by swapping flights under the assumption that AUs are the ones that best know the optimal sequence of their flights. The current UDPP procedure allows AUs to swap the order of two flights only in case that the flights have the same “most-penalizing” regulation and that the flights are operated by the same airline or group of airlines [5].

Several market mechanisms with regards to ATFM slot swapping have been proposed with the goal that AUs may allocate the resources more efficiently, i.e., AUs that value a slot most should be allocated the slot. Rassenti et. al. [6] propose a mechanism based on combinatorial auctions with sealed bids to allocate airport time slots among competing airlines. The proposed mechanism would allow AUs to bid for combinations of landing or take-off slots at airports. Ball et. al. [7] propose market mechanisms as a solution to the disadvantages of slot allocation under the current system. Ranieri et. al. [8] propose a mechanism based on auctions for distributing delay among multiple flights. The proposed mechanism employs a centralized decision-making process to minimize the total delay. The following decentralized decision-making process maximizes each individual AU’s welfare. None of these mechanisms, however, satisfies all the requirements for a market mechanism for ATFM slot allocation in case of the SlotMachine system.

3. Mechanism Design

In the following, we first state the assumptions before describing the basic properties of the proposed market mechanism as well as the inputs. We then illustrate the mechanism using examples before discussing the mechanism.

3.1 Assumptions

We assume that it is not possible to trust completely the inputs submitted by an AU to the SlotMachine system, i.e., the truthfulness of the inputs cannot be guaranteed. Consequently, the SlotMachine system requires a strategy-proof market mechanism for proper functioning.

We further assume that exchange of money between AUs, even in the form of congestion and compensation fees following a market mechanism that avoids financial competion [2], is not allowed or not desired by the AUs. A system of credits, as a virtual currency with no monetary value, can be put in place to indicate the values that each slot has for each flight. The credits also serve as the basis for the redistribution of welfare gains achieved through flight prioritization, which is a common way to align the individual incentives with the system goals and a way to achieve fairness [9].

We follow the assumption of Ball et al. [10] that each AU’s decision potentially affects other AUs’ objectives, which may cause the other AUs to modify their decisions, meaning that an AU’s decisions depend on the other AUs’ decisions, making game theory a prime candidate to describe the decision-making processes. For additional background information we refer to Gibbons [11] for game theory in general as well as Evans and Kessides [12] for an application of game theory to a use case in the aviation industry.

The *Collaborative Decision Making* (CDM) paradigm according to Ball et al. [10] considers three “basic forces” that ensure cooperative behavior of the AUs. First, AUs show goodwill towards CDM, which may, however, not be enough to ensure cooperative behavior of AUs in the long run. Second, according to *Contestability Theory* (CT) [13], peer pressure may keep the AUs in line, postulating that AUs may show cooperative behaviour in order not to provoke countermeasures by its competitors. Third, the NM acts as a supervisor, enforcing certain rules, which may also ensure cooperative behavior to some extent.

Ball et al. [10] further suggest that in the short run, the introduction of monitoring of the AUs’ behavior and implementation of countermeasures on part of the authority (the NM) may foster cooperative behavior. In the long run, however, the CDM paradigm should better be achieved by ensuring that the AUs have economic incentives to release accurate information in a timely manner, i.e., cooperative behavior should ultimately benefit the AUs in terms of increased utility, rather than through establishing punitive actions. To this end, an individual AU’s interests must be aligned with the goals of the system as a whole.

Ball et al. [10] also mention the importance of fairness when introducing the incentives for cooperative behavior. In this regard, according to Ball et al. [10], an issue may be how to measure the objectives of the optimization, e.g., economic utility, costs, or safety, which may be difficult to capture. Slot swapping through auction-based procedures may provide appropriate incentives for AUs while ensuring fairness if the procedures are properly designed.

3.2 Properties and Inputs

In the following, we discuss the main properties of the proposed mechanism as well as the inputs that AUs are required to submit in order to participate in a SlotMachine optimization run.

Auction mechanism. The proposed market mechanism is based on combinatorial auctions. More specifically, the proposed mechanism is inspired by VGC auctions [14], which shares some useful properties with the proposed mechanism. However, since the proposed mechanism has been adapted to the specific needs and requirement of ATFM slot swapping, some of the limitations present in general VGC mechanism are not present in the proposed mechanism.

Use of credits. In the proposed mechanism money transfers are not allowed. In that case, a means other than money is needed to let the AUs provide their valuations for the slots and to exchange value with other AUs as a key part of the market mechanism. A system of credits, a non-monetary virtual currency, will be used to avoid the exchange of real money in the SlotMachine marketplace.

Intrinsic value of the credits. Credits, like any currency, can be linked to “assets” that give them an intrinsic value. Since the credits will only be valid in the SlotMachine system the design of the credit system should take into account the aim of the market mechanism, which is to facilitate the exchange of slots so that AUs have flexibility to manage the delay allocated to their flights and, in turn, the cost of such delay. The delay, which can be seen as a proxy of monetary costs and, therefore, has value for all the AUs, can be linked to the intrinsic value of the credits, which will be named delay credits (DC). A fixed exchange rate one credit being equal to one minute of delay could be set to facilitate the interpretability of the value of credits. Note that each AU may have a different economic valuation of one minute of delay, depending on the characteristics of the flight operated and on whether the flight has received a level of delay that is within acceptable operational margins. One natural and expected way for the AUs to operate in the market is accepting more delay in some flights when the delay is relatively cheap while spending credits in exchange of better slots (less delay) when the cost of delay is relatively high.

Initial endowment. All AUs have an initial endowment of a certain amount of credits, e.g., 100 DC. This endowment is not proportional to the number of flights operated by AUs since AUs with more flights may need more credits to protect their important flights but they also have more flights to generate those credits. AUs with fewer flights (e.g., one flight per day) may need proportionally more credits to prioritize their flights before they can generate credits on their own. Whether the equal initial endowment will produce the desired effects in practice, i.e., AUs with more flights are encouraged to offer slots while smaller AUs can participate in the marketplace from the beginning, will have to be further investigated in the future.

One-shot vs. iterative finding of auction solutions. A requirement for AUs is that the system should be easy to use since AUs want to keep the focus on their operations. The SlotMachine system should find solutions in a relatively short time (no more than 15 minutes), ideally a solution that improves, if feasible, the operational efficiency of AUs. For this reason, a one-shot mechanism seems more appropriate, instead of other alternate market mechanisms that find the prices that clear the market in an iterative way. This does not mean, however, that multiple optimization runs, starting from the accepted solutions from previous runs, could be periodically conducted. One-shot mechanism also does not refer to the technical solution with evolutionary optimization algorithms, which find auction solutions iteratively but without clearing by the AUs in between.

Concerning the inputs of the market mechanism: For each of the flights participating in an optimization run, the corresponding AUs will submit to the SlotMachine system (either manually or supported by automation) one of the following options.

- **Flexible flight**. A flexible flight is able to exchange its current slot for a new slot, which can be later or earlier than the current slot, plus a number of credits that is directly proportional to the delay variation accepted if the exchange is produced. If the exchange materializes the AU will receive one credit per each minute of delay variation accepted. The farther away the positions accepted from the baseline slot the higher the number of credits received in exchange if the exchange materializes. AUs then have an incentive for accepting more delay when it is relatively cheap for the AU compared to other flights of the same AU.
- **Priority flight**. A priority flight is bidding for a new (better) slot and offering credits to one or more AUs that would accept changing the positions of their flights to enable a flight to get the best position possible. AUs will bid for the positions within the margins time-not-after and time-not-before, i.e., the latest acceptable departure/arrival time and the earliest acceptable departure/arrival time, respectively. It is also expected, since the proposed market mechanism creates incentives, that the values offered for each of the slots will be consistent with the actual utility function of the flight. The average costs c of one minute of delay when the credits are earned, i.e., when a flight was offering its position because the delay cost was relatively low, may be different for each AU. As a consequence, each AU can offer a number of credits for each slot that is proportional to c, which is known (at least approximately) by the AU. For instance, let us consider that an AU was offering the position of FLight A and earned 60 credits. Let $c = 10$ euros be the cost of accepting an extra minute of delay for Flight A, thus the additional cost for Flight A was 600 euros. Now the same AU wants to bid for a better Slot S for Flight B, which will increase the utility of the AU by 500 euros if the exchange is realized. Since to reach this slot Flight B needs to reduce its delay by 20 minutes, the minimum bid made by the AU operating Flight B for that slot is 20 credits. However, the AU operating Flight B should bid 50 credits, i.e., $bid = 500 \div c$, for the wished slot to allow the system to find the most efficient slot allocation in case that several bidders are competing for the same slot. Let us consider that the slot S is occupied by Flight C of another AU, who is willing to exchange the slot with Flight B in exchange for credits. If there is no other bidder, then the price that the AU operating Flight B will pay is only 20 credits. If other AUs are bidding for the same slot, and in the case that the AU operating Flight B has made the highest bid, the price paid will be higher than 20 credits and normally less than 60 credits (never more than 60 credits). Therefore, in all the cases the AU will be better off in terms of utility. These rules will be explained later. Note that an AU is not obliged to enter the marketplace with all of their flights.
- **Priority flight with flexibility**. For some flights it may happen that the AU would like them to be allocated within their margins, to improve their utility, and therefore will be bidding for a better slot. However, it is sometimes not feasible to materialize a requested exchange. In such a case the AU could still be willing to accept credits in exchange of other later or earlier slots, which normally would be out of the margins, because the marginal cost of accepting even more delay is relatively low with respect to the baseline delay of that flight.

Given the above inputs, the SlotMachine system will find the combination of feasible exchanges of ATFM slots between flights and will reallocate the slots based on these feasible exchanges while optimizing the social utility.

3.3 Examples

In the following we present examples of feasible and unfeasible exchanges under the proposed auction-based market mechanism.

	S1	S2	S3	S4	S5	S6	S7
A	X	-1	-2	-3	-4	-5	-6
Z	60	55					X

Figure 3 – Bilateral exchange

In the first example (Fig. 3), there are two flights, Flight A and Flight Z, participating in the market in which seven slots, from S1 to S7, are valued. Flight A is offering its current slot S1 (the offered slot is identified by an "X") and Flight A is accepting up to six minutes of extra delay. For the simplicity's sake let us assume that each slot has a size of one minute. Thus, Flight A is willing to exchange S1 for any slot between S2 and S7, both included. On the other hand, Flight Z is bidding for slots S1 and S2, with a maximum valuation of 60 DC and 55 DC, respectively. The exchange is feasible because the AU operating Flight A is willing to take S7 and the AU operating Flight Z wishes to take S1, which are necessary conditions for any bilateral exchange. Since there is no other bidder, the maximum price to pay by the AU operating Flight Z for S1 is 6 DC, which is equal to the number of minutes of delay reduced for that flight. The 6 DC gathered by the system are given to the AU operating Flight A for accepting six minutes of extra delay. The credits awarded could be more than 1 DC per minute of accepted extra delay, but the awarded credits should in any case be proportional to the minutes of extra delay accepted because AUs should accept additional delay in order to be awarded credits.

	S1	S2	S3	S4	S5	S6	S7
A	X	-1	-2	-3	-4		
B			X	-1	-2	-3	-4
Z	60	55					X

Figure 4 – Multilateral exchange

In the second example (Fig. 4) the situation is similar to the first example but now the exchange between Flight A and Flight Z is not feasible, since for Flight A no more than 4 minutes of extra delay is acceptable. There is, however, another Flight B, which is placed in S3 and the operating AU is willing to accept up to four minutes of extra delay; Flight B could be placed into S7. The bilateral exchange between Flight B and Flight Z is also not feasible, since Flight Z is not bidding for S3. Fortunately, a multilateral exchange is feasible, e.g., Flight A can take S3 and Flight B can take S7, and Flight Z can take S1. Since there are no other bidders, the AU operating Flight Z will pay 6 DC, which will be distributed by giving 2 DC to Flight A and 4 DC to Flight B.

In case that multiple flights are bidding for the same slot, the system will find the most efficient allocation, i.e., the one that generates more credits, which is equivalent to more utility, and will calculate the negative externalities that the participants receiving a competed slot has generated to others. The participant receiving the slot will compensate the participants receiving a suboptimal slot in a proportional manner. To this end, the net optimized social utility, i.e., the utility after paying the slot prices, NU_s^{Opt}, will be compared against the ideal net utility, NU_s^{Ideal}, in which all the flights would be allocated their ideal slot. Note that in the presence of externalities this scenario would not be feasible, which is why it is considered ideal. The *equitable utility* will be found by calculating the externalities' degradation ratio as follows.

$$\nabla NU_s^{Ideal} = \frac{NU_s^{Opt} - NU_s^{Ideal}}{NU_s^{Ideal}} \tag{1}$$

The net equitable utility for each flight i is then found as follows.

$$NU_i^{Equ} = NU_i^{Ideal} \times (1 + \nabla NU_s^{Ideal}) \tag{2}$$

Finally, the cash flows CF for each flight i can be found as follows.

$$CF_i^{Equ} = NU_i^{Opt} - NU_i^{Equ} \tag{3}$$

	S1	S2	S3	S4	S5	S6	S7
A	X	-1	-2	-3	-4		
B			X	-1	-2	-3	-4
C		X	-1	-2	-3	-4	
Y	50	40				X	
Z	60	55					X

Figure 5 – Social optimum with fairness

Figure 5 shows an example with similar preferences for the Flights A, B, and Z with respect to the previous example. In this example, however, a new Flight C, which the AU is willing to accept additional delay for, and a new Flight Y, which the AU would like to move forwards, can be found. Flights Y and Z are competing for slot s1, which is the preferred slot for both of them. The Heuristic Optimizer component will find the combination that generates more utility, which in this case is to allocate s1 to Y and s2 to Z. Note that S1 is not allocated to the highest bid, but the solution still maximises the social utility. This will generate $50+55=105$ DC of optimized utility, instead of $60+40=100$ DC that would be produced if S1 was given to Flight Z and S2 to Flight Y. The net optimized utility of the optimal solution is $NU_s^{O}pt = (50-5)+(55-5) = 95$. The net ideal utility (both Flight Y and Flight Z hypothetically receiving S1) is $NU_s^{Ideal} = (50-5)+(60-6) = 99$. The degradation ratio is $\nabla NU_s^{Ideal} = (95-99) \div 99 = -0.04$. It follows that $NU_Y^{Equ} = 43.18$ and $NU_Z^{Equ} = 51.82$. Thus, the cash flows must be $CF_Y^{Equ} = 45-43.18 = 1.82$ and $CF_Z^{Equ} = 50-51.82 = -1.82$, which means that the AU operating Flight Y must pay 1.82 DC, which will be given to the AU operating Flight Z for compensating the negative externalities of Flight Y, which took the most preferred slot by Flight X. Note that if S1 were allocated to Flight Z and S2 to Flight Y the equitable utilities would be $U_Y^{Equ} = 40.91$ and $U_Z^{Equ} = 49.09$, which would be suboptimal for both flights. Therefore, both AUs have an interest in the solution that leads to the social optimum.

	S1	S2	S3	S4	S5	S6	S7
A	X	-1	-2	-3	-4		
B			X	-1	-2	-3	-4
C		X	-1	-2			
Y	50	40				X	
Z	60	55					X

Figure 6 – Social optimum and fairness when a bidder flight is left out of the margins

The example shown in Figure 6 is similar to the previous one, but now Flight C is only accepting two minutes of extra delay. A consequence for this scenario is that only one of the two bidders will be able to realize a feasible exchange. Following the same steps as before, the slot S1 will be allocated to Flight Z, and there will be a transfer of credits from Flight Z to Flight Y to compensate for the externalities. The degradation ratio is $(54-99) \div 99 = -0.4545$ and after the compensation cash flows the equitable utilities will be $U_Y^{Equ} = 24.55$ and $U_Z^{Equ} = 29.45$. Flights A, B, and C will be moved to later positions in FPFS order, i.e., A to S2, C to S3, and B to S4, while Y remains at S6. In the example from Figure 6, after the compensations, Flight Y may not receive the preferred slot but at least would be compensated fairly.

Note that the compensation mechanism makes the market more equitable and fairer, and aligns the individual interests towards achieving a social optimum. This mitigates the risks of potential cheating or system abuses. These risks can be further mitigated with privacy-preserving input and output verification processes brought by multi-party computation and blockchain technologies under research in SlotMachine.

In order to generate more incentives for the market creators (the AUs offering slots) for accepting more delay, and therefore to increase the chances of finding feasible exchanges during the optimization process, an additional payment of credits would be given to them. In addition to the credits linked to their extra delay accepted, an additional number of credits consisting in 10% of the social utility generated by the bidders after paying the price of the slots would be allocated to the market creators. Applying this rule to the first example from Figure 3, Flight A would receive $(60-6)\times 10\% = 5.4$ DC in addition to the 6 DC paid by the operator of Flight Z for reducing the delay. In total Flight A would receive 11.4 DC. In the third example from Figure 5 the 10% bonus can be calculated from the equitable utilities, i.e., the operator of Flight Y should pay $43.18\times 10\% = 4.32$ DC and the operator of Flight Z should pay $51.82\times 10\% = 5.18$ DC, which makes a total of 9.5 DC to be distributed among the market creators A, B and C. The share for them should be proportional to the extra delay allocated to these flights in the optimal solution, i.e., one, four and four minutes of extra delay, respectively. Therefore, from the 9.5 DC that the operator of Flight A should receive $(1\div 9)\times 9.5 = 1.06$ DC whereas B and C should each receive $(4\div 9)\times 9.5 = 4.22$ DC. In total the operator of Flight A would receive 2.06 DC, the operator of Flight B would receive 8.22 DC, and the operator of Flight C would receive 8.22 DC. The bidders would have a utility surplus of 38.86 DC for Flight Y and 46.64 DC for Flight Z.

3.4 Discussion

The proposed market mechanism meets the following four desired requirements for market mechanisms.

- The mechanism satisfies *individual rationality* constraint because it is guaranteed that all the participants will be better off after the implementation of the mechanism, at least in the long run where all the AUs would be able to accumulate and spend credits over time.
- The mechanism is *budget-balanced* because the NM does not need to pay or subsidise to reach the equilibrium.
- The mechanism reaches *allocative efficiency* because the outcome after optimization will always be Pareto efficient, i.e., it is not possible to improve the welfare/utility of any of the involved actors without making the situation for one of them worse. Therefore, there must not be any mutually advantageous exchange between two agents that could improve the situation for both of them. However, in this mechanism the outcomes of the optimiser will normally be suboptimal with respect to the economic/monetary valuation of the slots. The reason is that there is no valuation in monetary units but in credits, whose value is linked to one delay minute. This implies that the slots will not be assigned to the ones that value them most in absolute terms but to the ones that value them the most in relative terms with respect to the cost structure and expected delay impact of their own flights over the time, i.e., each AU will accept extra delay in exchange of credits from flights that are in relatively cheap and will spend the credits in reducing the important delay of flights is to optimize the own utility. The slots will be allocated to the AUs that are willing to pay more credits for slots, but since the exchange rate between credits and euros would normally be different for each AU the final allocation might be sub-optimal with respect to the ideal case.
- The mechanism is *strategy-proof* because each AU can maximize their expected utility by reporting true valuations given that the other is expected to report honestly. The slots must be offered at a fixed price, which varies with regards to the extra delay accepted, but always at a fixed rate of 1 DC $= 1$ minute. This creates incentives for the “owner” to offer the slot when it is relatively cheap for them, and to accept as much delay as possible with the hope that a bilateral or multilateral trading will be feasible and can in consequence obtain as much credits as possible. On the other hand, the compensation mechanism ensures that reporting true valuations is the dominant strategy for the “buyer”.

Note that the Arrow impossibility theorem [15] seems not to apply to the proposed mechanism and, therefore, the four desired market properties are fulfilled. First, prices of offered slots are fixed and known. Thus, no one has market power to change the prices and there is no bilateral problem of asymmetric information since the price of the slot is known. Likewise, the AU "selling" (offering) slots maximizes their utility knowing that they will be able to use the credits obtained to buy other slots with the same constant price. Thus, the decision whether to offer a slot or not is based on the AU's own preferences on how to distribute the delay of their flights among their own flights and over the time, independently from the valuation of any potential "buyer".

4. Privacy-Preserving Implementation

The SlotMachine system separates the search for candidate flight lists from the evaluation of the flight lists. The evaluation is conducted by the Privacy Engine, which employs multiparty computation to evaluate the solutions. Figure 7 illustrates the basic principle of the SlotMachine system, particularly the interaction between Heuristic Optimizer and Privace Engine. The AUs employ their own business rules and user interfaces to capture and encrypt the preferences regarding the slots, which are submitted to the SlotMachine system in encrypted form. The SlotMachine system does not decrypt those preferences but computes the fitness of candidate flight lists using (secure) multi-party computation (MPC). The MPC nodes, which are hosted outside of the SlotMachine system, e.g., by different AUs, perform the privacy-preserving computations required for the Privacy Engine to be able to evaluate the candidate flight lists that a genetic algorithm finds in an iterative manner.

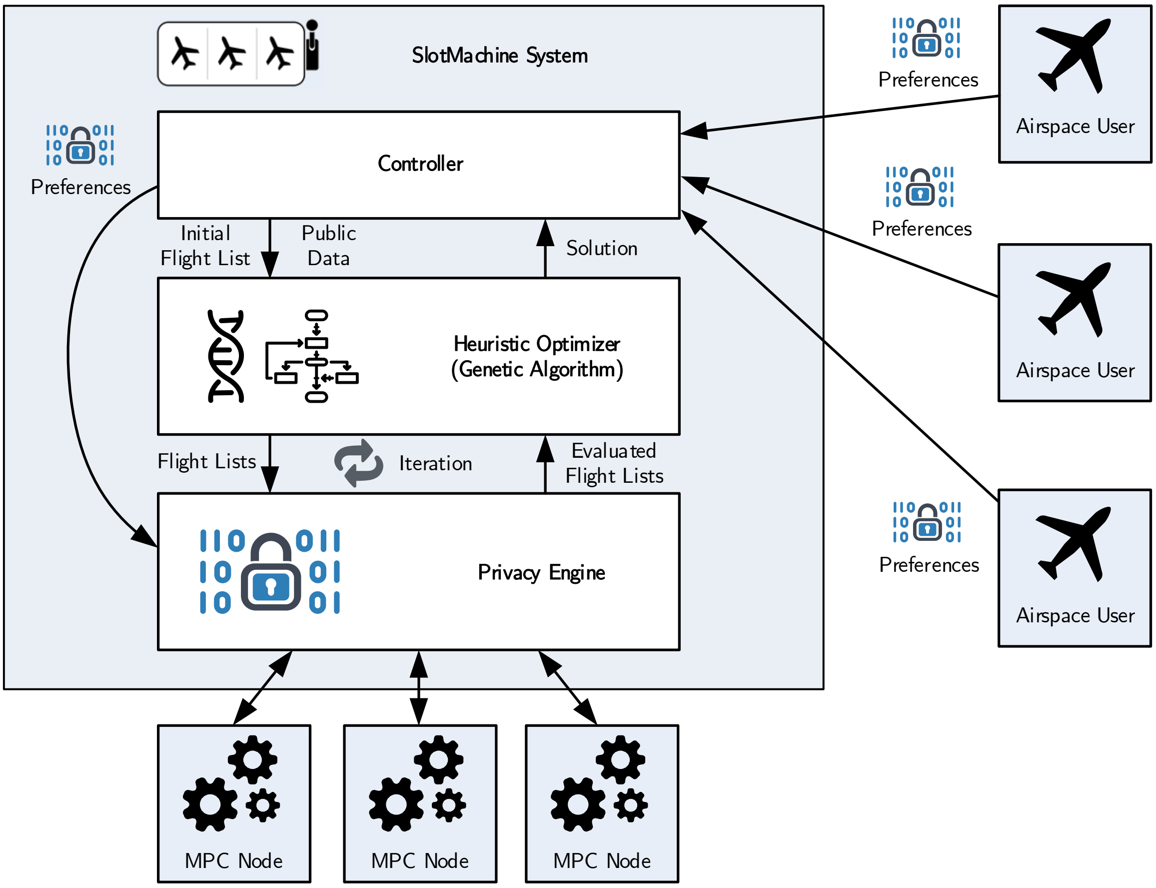

Figure 7 – Interaction between Heuristic Optimizer and Privacy Engine

The first approach that we have investigated in the course of the development of the SlotMachine system is to translate AUs' preferences, expressed in terms of margins (time-not-before, time-not-after, time-wished) and priority, into weight maps that express the utility of each slot for each flight. In order to allow for privacy-preserving optimization under the proposed market mechanism, the bids and offerings could be translated into weight maps. Alternatively, the evolutionary algorithm may conduct multi-objective optimization using the following criteria. The Privacy Engine would then rank the solutions according to the following criteria.

1. No margins must be violated. Any solution that violates a margin is always ranked lower than a solution that does not violate a margin.
2. Maximize the sum of the bids.
3. Minimize the number of slots that a flight is pushed back.

4.1 Input Data Quality

With air traffic on the rise again after a temporary slump during the years of the COVID-19 pandemic, more regulations will be put in place to handle limited capacity at airports and in the airspace. In those cases it will be in the interest of airlines to prioritize specific flights and optimize overall cost. There is already some experience from airlines to exchange departure times within an airline or within a group of airlines but for more flexibility and higher cost savings it is necessary to exchange slots between airlines.

An inhibitory effect for airlines sharing preferences regarding flight prioritization are concerns about disclosing internal information about the cost structure of an airline when requesting or offering an ATFM slot. In addition, not all airlines have data available about the exact cost structure of a flight and can easily participate in a marketplace to trade departure or arrival times when a regulation becomes active.

SlotMachine, in its basic form, uses weight maps as input for the Heuristic Optimizer, which describe the utility of a slot for a flight; there are different ways how to generate the weight maps. In the current stage of development, AUs provide their preferences in the form of margins, specifying the preferred time range for a flight as time-not-before and time-not-after as well as a preferred time-wished, and a priority of the flight. An alternative method could be to choose a preferred slot for a flight, or just selecting flights to participate in a flight list optimization run. It would also theoretically be possible to integrate SlotMachine with internal systems of an airline and compute a weight map using complex preferences for each flight.

For different airlines the available resources that the airline can afford to spend to participate in an optimization run vary. While some airlines have teams of data scientists optimizing their operations, others are not yet well positioned to provide inputs on short notice. An overall theme from airlines familiar with the SlotMachine concept is the request for fairness and transparency in the overall process.

- **Fairness**. Optimization models and calculation concepts must be known and be the same for all participating airlines. In case credits are used in the process, it needs to be established beforehand how many credits will be available in the beginning, how potentially planned deprecation works, as well as how issuing of new credits is planned.
- **Transparency**. Although input information is private to an airline, it should be possible to verify the correctness of an optimization process afterwards if all relevant airlines agree, i.e., disclosure of input data at a later point in time. In this regard, zero-knowledge proofs [16] could be employed.

4.2 Data Privacy

The basic idea behind SlotMachine is to build a trustworthy distributed platform which can be operated by stakeholders without the need for a central authority. It contains no single point of trust, especially with respect to the confidentiality of sensitive user input data, and protected data are only processed in encrypted form. Additionally, all crucial steps are made publicly verifiable by logging essential checkpoint data in a blockchain. To support verifiability for steps dealing with sensitive information, cryptographic zero-knowledge proof techniques are used to overcome the privacy–transparency trade-off (cf. [17]).

In order to ensure confidentiality of flight preferences from AUs, the inputs are first encoded for the use in and embedded secure multi-party computation (MPC) system and are sent to the respective MPC node with an additional layer of encryption only. Thus, data is protected from operators of MPC nodes by the very nature of the protocols used (secret sharing based MPC [18]) and even the SlotMachine operator has no access to them thanks to the additional encryption layer used (encryption based end-to-end secure channels from AU to each MPC node). These two mechanisms provide provable security guarantees for AU input data through the underlying cryptography concepts.

Nevertheless, a certain amount of information leakage, albeit small, in the setting of SlotMachine is inevitable for the system to work, e.g., a new flight sequence has to be published by the Network Manager. Therefore, the focus of our investigation is currently on studying the trade-off between the information that is needed for the platform to operate while still maintaining the confidentiality of sensitive information. In particular, we are minimizing the information revealed during the iterative process between Optimizer and Privacy Engine (PE), and researching optimization techniques which require only minimal information. This is necessary, because fully oblivious optimization is not feasible for SlotMachine in real-time (see [19]).

In SlotMachine we study different trade-offs between efficiency and privacy for optimizations and minimize the information that has to be revealed by the PE when optimizing a flight sequence. In our initial implementation the PE reveals the flight order together with the fitness of the best solution. To avoid publishing the fitness value of the best solution, another implementation just partitioned the flight order into quantiles together with the best fitness, i.e., it is not possible to link the fitness value to a single flight list. For configurations with a low number of flights, specific care must be taken to avoid having only a single flight list in the top quantile. This can be achieved through introducing dummy flights. Additionally, we are researching heuristic optimization methods which do not need concrete fitness values at all but can improve results from only learning relations between different solutions, i.e., only the order is revealed.

Initial tests have shown promising results that the genetic algorithms perform with the input of partitioning into quantiles as well as with the exactly ordered flight lists. Therefore, even a malicious platform operator cannot deduce the flight preferences of AUs when intercepting results between Optimizer and Privacy Engine.

5. Summary and Future Work

We presented a market mechanism based on combinatorial auctions for ATFM slot swapping using the SlotMachine system, which allows airspace users to submit preferences regarding flight prioritization in case of reduced capacity in the air traffic network. The proposed market mechanism ensures that airlines participating in the SlotMachine flight prioritization process are fairly compensated and have an incentive for participating. A privacy-preserving implementation ensures that the airlines' confidential inputs are kept private.

Future work will further investigate implementation aspects of the presented market mechanism, in particular whether optimization using weight maps is feasible or whether the Heuristic Optimizer should conduct a multi-objective optimization. Furthermore, future work will conduct simulation and experiments to more thoroughly investigate the properties of the presented market mechanism, ideally with participation of multiple airlines.

6. Contact Author Email Address

`mailto:schuetz@dke.uni-linz.ac.at`

7. Copyright Statement

Acknowledgment

This work was conducted as part of the SlotMachine project. This project received funding from the SESAR Joint Undertaking under grant agreement No 890456 under the European Union's Horizon 2020 research and innovation program. The views expressed in this paper are those of the authors.

References

[1] C. G. Schuetz, E. Gringinger, N. Pilon, and T. Lorünser, "A privacy-preserving marketplace for air traffic flow management slot configuration," in *2021 IEEE/AIAA 40th Digital Avionics Systems Conference (DASC)*, 2021, pp. 1–9.

[2] E. Gringinger, S. Ruiz, and C. G. Schuetz, "Business and economic concepts for a privacy-preserving marketplace for ATFM slots," in *2022 Integrated Communication, Navigation and Surveillance Conference (ICNS)*, 2022, pp. 1–11.

[3] T. Vossen and M. Ball, "Optimization and mediated bartering models for ground delay programs," *Naval Research Logistics*, vol. 53, no. 1, pp. 75–90, 2006.

[4] S. Ruiz, L. Guichard, and N. Pilon, "Optimal delay allocation under high flexibility conditions during demand-capacity imbalance: A theoretical approach to show the potential of the user driven prioritization process," in *SESAR Innovation Days, Belgrade, 2017*, 2017, pp. 1–8.

[5] S. Niarchakou and M. Sfyroeras, "ATFCM users manual, network manager, ed. 25," EUROCONTROL, Tech. Rep., 2021. [Online]. Available: https://www.eurocontrol.int/sites/default/files/2021-03/eurocontrol-atfcm-user-manual-25-29032021.pdf

[6] S. Rassenti, V. Smith, and R. Bulfin, "A Combinatorial Auction Mechanism for Airport Time Slot Allocation," *Bell Journal of Economics*, vol. 13, no. 2, pp. 402–417, 1982.

[7] M. Ball, R. H. Smith, G. Donohue, and K. Hoffman, "Auctions for the safe, efficient and equitable allocation of airspace system resources," in *Combinatorial Auctions*, 2003.

[8] A. Ranieri and L. Castelli, "A market mechanism to assign air traffic flow management slots," in *8th USA/Europe Air Traffic Management R & D Seminar*, 2009.

[9] T. Börgers, D. Krähmer, and R. Strausz, "An introduction to the theory of mechanism design," *OUP Catalogue*, 2015.

[10] M. O. Ball, C.-Y. Chen, R. Hoffman, and T. Vossen, "Collaborative decision making in air traffic management: Current and future research directions," in *New Concepts and Methods in Air Traffic Management*, L. Bianco, P. Dell'Olmo, and A. R. Odoni, Eds. Springer, 2001, pp. 17–30.

[11] R. Gibbons, *Game Theory for Applied Economists*. Princeton University Press, 1992. [Online]. Available: http://www.jstor.org/stable/j.ctvcmxrzd

[12] W. N. Evans and I. N. Kessides, "Living by the "Golden Rule": Multimarket Contact in the U. S. Airline Industry," *The Quarterly Journal of Economics*, vol. 109, no. 2, pp. 341–366, 05 1994. [Online]. Available: https://doi.org/10.2307/2118466

[13] E. E. Bailey and J. C. Panzar, "The contestability of airline markets during the transition to deregulation," *Law and Contemporary Problems*, vol. 44, no. 1, pp. 125–145, 1981. [Online]. Available: http://www.jstor.org/stable/1191388

[14] R. H. Coase, "The problem of social cost," *The Journal of Law & Economics*, vol. 3, pp. 1–44, 1960. [Online]. Available: http://www.jstor.org/stable/724810

[15] E. Maskin and A. Sen, *The Arrow impossibility theorem*. Columbia University Press, 2014.

[16] E. Ben-Sasson, A. Chiesa, E. Tromer, and M. Virza, “Succinct non-interactive zero knowledge for a von Neumann architecture,” in *23rd USENIX Security Symposium*, 2014, pp. 781–796.

[17] T. Lorünser, F. Wohner, and S. Krenn, “A privacy-preserving auction platform with public verifiability for smart manufacturing,” in *Proceedings of the 8th International Conference on Information Systems Security and Privacy*, INSTICC. SciTePress, 2022, pp. 637–647.

[18] R. Cramer, I. B. Damgård, and J. B. Nielsen, *Secure Multiparty Computation and Secret Sharing*. Cambridge University Press, 2015.

[19] T. Lorünser, F. Wohner, and S. Krenn, “A verifiable multiparty computation solver for the assignment problem and applications to air traffic management,” 2022. [Online]. Available: https://arxiv.org/abs/2205.03048