A Market Mechanism for Airport Traffic Control*

Geert Jonker, John-Jules Meyer, Frank Dignum
Institute of Information and Computing Sciences
Utrecht University

Abstract

A weighted-voting mechanism is presented in which agents are able to jointly decide on plan repair proposals in the airport traffic control domain. The mechanism uses Clarke Tax for incentive compatible bidding. It provides a slider value by which an airport can control the balance between optimality and fairness of plan repair solutions. Special care is taken to ensure that agents have no incentive to vote untruthfully and that no agent gets too much negotiation power.

1 Introduction

Nowadays airports are more and more faced with air traffic congestions as a result of increased capacity demands. Much effort has been put into the development of software tools to assist the air traffic controllers in their decision-making process. These tools typically try to optimize a part of the planning on an airport, typically the arrival and departure sequence and the gate assignment. Usually a strict hierarchy between planners exists to facilitate compliance to the several safety constraints. On the delay of an incoming aircraft, the arrival manager will typically replan its schedule first, to which the gate planner will adjust its schedule, after which the departure manager will adjust its planning.

Another, more progressive trend in air traffic control (ATC) automation is that of distributed planning. An example is the Free Flight program [4, 9], in which aircrafts are meant to plan their own path of flight while communicating with aircrafts around them to avoid collision. In the context of *collaborative decision-making*, a lot of work is done on information sharing between parties to increase quality of planning [5, 13].

This article focuses on distributed airport traffic planning (ATP), i.e., the planning of the arrival, gate and departure process. We will look at the most important aspects of this planning and present a coordination mechanism by which aircrafts can jointly decide on and enforce plan changes. In the remainder

*This research is supported by the Technology Foundation STW, applied science division of NWO and the technology programme of the Ministry of Economic Affairs. Project DIT5780: Distributed Model Based Diagnosis and Repair

of this article we use the terms 'aircraft' and 'airline' to represent the same party. When aircraft negotiate they in fact represent the interests of their airline; saying that two aircraft negotiate is in that respect equivalent to saying that their airlines negotiate.

2 Airport Traffic Planning

The planning of airport traffic starts months before it is executed. Based on the flight requests of airlines, provisional arrival and departure schedules are made. As time progresses and more information becomes available, these schedules become more and more detailed. On the day before execution the optimal gate assignment is determined and 'frozen', i.e., no more flight requests can be added. On the day of execution all flights are assigned *time slots*, 15 minute time periods in which they have to depart or arrive[1]. If a flight 'misses' its slot it has to request a new slot which is often not immediately available.

There are many reasons why things don't always go as planned. An aircraft might arrive at an airport later than planned, it might not be able to land on arrival because of congestions, a runway might be closed, etc. It might not be able to go on its gate on time because the previous aircraft hasn't left yet. The *turn-around process* of an aircraft, the time that it is at the gate and is cleaned, refuelled, boarded, etc., might take longer than planned. It might not be able to depart on time because of congestion on the runway. And so on. It is up to the air traffic controllers to deal with these disruptions as efficiently as possible. In general their main aim is to minimize the total amount of delay while complying with the safety constraints. The most important safety constraints are the separation constraints that indicate the minimal distance aircrafts should maintain in different situations. Other constraints follow from taxi distances, ground services (catering, refuelling, cleaning, etc.), transfer passengers, etc. Of course flights should be kept within their timeslots if possible.

When changes are made to the schedule, an important criterion for ATC to observe is *fairness*. In case global plan changes have to be made because of disrupting circumstances, the different airlines should each bear an equal share of the burden[2]. On a smaller scale, if at one occasion a flight from airline X has to be delayed in order to resolve a planning problem, the next time a flight has to be delayed it should be one from another airline than X. A factor that usually gets very little attention is the preferences of airlines themselves. It can very well be that an airline (or a group of airlines) prefers situation X over Y, while ATC has decided Y but wouldn't object to X. This might be because ATC doesn't have the time to research Y, or that it lacks information on the airlines' preferences.

[1] The fair allocation of slots to airlines is a challenging problem on its own, see for instance [12].

[2] Although fairness is important, in reality plannings shouldn't be *too* fair either. Every airport has a so called *hub owner*, the airline using the airport as its base. Traditionally, hub owners are slightly favoured by ATC in their planning. If an airport was to switch to a truly

From the multiagent point of view, ATP can be seen as a coordination problem between different actors that have different size, preferences, authority and responsibility. These actors need to agree on an efficient and fair planning to be executed collectively, under continuously changing circumstances. The challenge is to transform the current situation of centralized planning, authority and responsibility to a decentralized one.

A last word on safety. While efficiency and fairness are measures that can be satisfied until a certain extend, safety is much more stringent. Safety measures are physical and temporal separation rules that have to be satisfied at all times. Since they cannot be relaxed, we will not consider them in the optimization process.

3 Principled Negotiation

As a starting point we'll look at a fairly straight forward multiagent coordination technique called *principled negotiation* [7]. Suppose that a scheduling conflict has occurred which resolution needs the participation of several parties. In principled negotiation a central coordinating agent, in our case the air traffic controller, determines an initial solution. The involved parties then have the opportunity to propose alternative solutions. When the coordinating agent receives a proposal, it will send this proposal to all the agents involved in it. The agents then have the possibility to either accept or reject. If a proposal is accepted by all agents, this proposal replaces the old solution.

Wangermann and Stengel applied principled negotiation to runway slot allocation. They showed that it allowed airlines to reduce costs because they were given an opportunity to improve a given slot allocation according to their own preferences. At no point in the process would an agent see its situation deteriorate because it could always reject a disadvantageous proposal. Principled negotiation is also the underlying cooperation principle in the cooperation framework GPGP, where two agents use their combined utility gain and cost to arrive at a mutually beneficial solution[8, 14]. It is also used in the task allocation mechanism Contract Net, where two agents engage in a contract if it is beneficial to both of them [10].

The success of this algorithm largely depends on the character of the domain; if the initial solution enables a lot of small improvements to be made without putting any party at a disadvantage then this method is able to improve the initial solution. In the ATP case however, resolving planning conflicts typically requires parties to make unwanted changes to their schedules. Moreover, different solutions typically favour different parties. Therefore chances are high that, when given an initial solution, there will always be a party that favours this solution above all others. As a simple example, consider the conflict in which flights X and Y's schedules are partly overlapping at a gate. There are two possible solutions: either Y delays its gate entrance to avoid the overlap,

fair planning mechanism, the hub owner wouldn't agree. Another example where existing privileges are not to be ignored is slot allocation [12].

or X goes to another gate. At that other gate however flight Z has to delay its entrance in order to avoid a new overlap with X. The agent of flight Y would clearly prefer the second solution, while the agent of flight Z would prefer the first one. In this case principled negotiation would not be very successful. If the first solution is chosen as the initial one, it will be the final solution. If the second is chosen then that one will be the final solution. It could however very well be the case that the disadvantage of the first solution to agent Y is much less than the disadvantage of the second solution to agent Z. What we actually want is to find the solution that has the lowest combined disadvantage to all the agents involved.

Usually it is up to the air traffic controller to estimate which solution is the least disadvantageous to all agents. There are two reasons why he is not very well able to arrive at this optimal solution. First he is not aware of all the preferences that the involved parties have concerning plan changes, which he has to estimate. Secondly, he hasn't got enough time to generate and evaluate all the possible solutions to a problem. Air traffic controllers work under a lot of time pressure which forces them to use quite drastic rules of thumb to deal with conflicts. For instance, aircrafts that fail to arrive or depart at the scheduled times are often put 'on hold' until a space in the schedule appears in which they can then perform their actions.

In the next section we will look at voting, a multiagent decision technique better capable of involving agents' preferences.

4 The Clarke Tax Mechanism

A way to enable a group of agents to decide on which option to choose from a set of options is by *voting*. In its most simple form, all agents vote for their favourite solution and the solution with the highest number of votes wins. The reader will understand that this procedure will not always yield the optimal solution. It only takes into account the *ordinal preferences* of the agents, i.e., the order in which an agent ranks choices, without assigning relative weights[3]. To reach optimal outcomes we need to use the agents *cardinal ordering* over alternatives, i.e., the agents assign weights to their choices. The most straight forward procedure here is *sealed bidding*, in which every agent specifies an amount of money (positive or negative) for each alternative. The bids are added up and the option with the highest accumulated preference wins.

A problem with this and many other kinds of voting is the fact that it might be attractive for an agent to vote *strategically*. This happens when they vote not in accordance to their own preferences in an attempt to manipulate the outcome. In the example above, where the agents submit their costs for each option so that the option with the lowest total costs will be chosen, an agent has an incentive to underbid. If it assumes that a certain outcome will be achieved

[3]Voting mechanisms that rely on ordinal preferences generally suffer from Arrow's *impossibility theorem*. This theorem tells us that any voting mechanism that meets a group of reasonable criteria and relies on ordinal preferences cannot be fair[1].

even without the full force of its vote, he can submit a lower cost, pay less and still get his most preferred outcome. In economics this is known as the *free rider problem*.

The area of game theory called *mechanism design* is devoted to interaction protocols that yield the optimal social outcome when agents use their dominant strategy. In such *nonmanipulable* or *incentive compatible* protocols participants are best of when they are truthful. An example of a single-object nonmanipulable voting procedure is the *Vickrey auction*[11], where every agents submits its bid for the single object on sale, and the object it sold to the highest bidder at the price of the second highest bid. Vickrey showed that a bidder's dominant strategy is to bid his true valuation[11].

An example of a nonmanipulable decision protocol using cardinal preferences is the *Clarke Tax Mechanism*[2, 3]. In this mechanism agents need to decide on which solution from a set of solutions to choose, and do this by giving their valuations to each of the alternatives. The agents however run a risk of having to pay a tax, which happens if their vote made a difference to the outcome. The tax is equal to the total value of the outcome minus the value of the outcome that would have happened if it hadn't voted, and not less than zero. This tax discourages an agent to overbid on an alternative it likes; if overbidding means changing the outcome, the amount of tax it has to pay might be higher than difference in valuation between this and its next best alternative. At the same time, the tax doesn't encourage them to underbid; if underbidding changes the outcome, the saved tax will never compensate for the loss of utility. Therefore, revealing true preferences is the optimal strategy in the Clarke Tax mechanism. For a formal proof see [2] and [6].

In general the Clarke Tax mechanism (CTm) has a number of drawbacks in practical implementations. It assumes for instance that the set of alternatives is fully determined in advance, that the participants are equally important in the decision making process, and that no coalitions are made. These and other drawbacks are dealt with in [3]. In this article we will focus on the question how the CTm can be applied in ATP and if it can help in achieving fairness.

5 A voting mechanism for ATP

The problem we are facing is the following: given a group of agents with different preferences and a set of proposals, which nonmanipulable decision mechanism allows them to find the most preferred proposal while fairness is maintained? In this section we will introduce such a decision mechanism. We use *weighted voting* as the main principle to determine the most preferred solution among the voters and *accounts* to ensure fairness.

We say that a coordination mechanism is *fair* if after the mechanism has been used many times, the parties involved each have experienced an equal utility gain or loss. In terms of ATP, the different airlines each have more or less suffered equal delays, gate changes and other plan changes[4]. If a certain

[4]For now we assume that all parties are of the same size, e.g., KLM is just as big as Easyjet.

situation forces ATC to put one airline at a disadvantage compared to another airline, the next time it would preferably be the other way around. Ideally, in the end each airline has suffered an equal disadvantage from the plan changes. A problem that arises is that it is not easy for ATC to assess and compare the effects of a plan change for the airlines involved. First it is hard to compare the utility loss of different events, for instance an arrival delay and a departure delay or a gate change. Secondly, an event may have different effects at different times; a delay could be very harmful to an airline at one moment and harmless at another moment. It would be most ideal if agents would be able to express the utility of a certain plan change in a numeric way. For this article we will assume that airlines are able to valuate a plan changes in terms of money. This means that they are able to say how much money a certain delay costs them, or in other words, they are indifferent between not suffering the delay and suffering the delay but being compensated for X credits.

We can now use the CTm as the voting mechanism by which airlines decide which plan change to enforce. The initial set of alternatives over which to vote is created by a central ATC-agent. This agent has full knowledge about the expected arrival, gate and departure times of the agents and the constraints between them. To solve a given plan conflict it creates possible plan repair schedules by shifting events forward or backward or putting them on a different resource, i.e., having an aircraft arrive on a different runway or putting it on a different gate. These proposals are sent to all the agents involved, who valuate them and sent the results back. Note that in this article a valuation of X stands for the costs of the proposal to the agents. The proposal with the highest value, i.e., the lowest accumulated costs, is chosen and enforced. This ensures optimality in terms of costs to the airlines involved.

To ensure fairness, the agents' contributions to plan changes are remembered. If a certain agent has already contributed a lot to repair plan conflicts and thus has incurred a lot of expenses, it should be spared in subsequent plan repair cases. To achieve this we give every agent an account with initially zero credits. Every time it participates in a repair event, the costs X of its participation is added to its account. Agents that have contributed a lot thereby get a positive balance, agents that have only benefited from plan changes will have a negative balance. Agents with a positive balance should then have an advantage in future decision procedures. The proposed way to achieve this is as follows: instead of choosing the proposal with the lowest total cost, we choose the proposal that leads to the lowest sum of the squared balances of the agents:

$$p = argmin_p \sum_{i=1}^{k} (b_i + v_{i,p})^2$$

where k is then number of agents, b_i is the balance of agent i, $v_{i,p}$ is the valuation of agent i to proposal p. This prevents an agent from acquiring a high balance, thus preventing it from becoming too much of a (relative) benefactor. In effect this principle gives agents that have contributed a lot already an advantage in future decision procedures; their share will be tried to be kept as small as possible in order to minimize the sum of the squared balances.

This mechanisms is not strictly optimal any more. If all the agents' balances are zero for instance, it favours a solution where everybody contributes a little bit above a solution where one agent contributes a lot, even if that last solution would have a lower total cost. In fact, this mechanism strikes a balance between optimality and fairness, where the power to which the balances are raised determines the exact balance between the two. If the power is set to one, the balances are simply added up, optimality is fully achieved and fairness is neglected. If the power is set very high, fairness is very well achieved while optimality is neglected. We therefore make this power variable, giving the airport a 'slider value' enabling it to regulate the balance between optimality and fairness by setting it not too high and not too low:

$$p = argmin_p \sum_{i=1}^{k} (b_i + v_{i,p})^s$$

where s is the slider value described above.

While this mechanism ensures optimality and fairness to a certain degree, there are a number of dangers to avoid. First, there is still an incentive for agents to lie, namely by submitting too high costs for *all* of the alternatives. If an agent adds a constant value to every valuation of the proposals, it is still playing a dominant strategy in terms of the CTm, but it is now gaining more credits on its account than if it wouldn't have added the constant value. Its dominant strategy now would be to make this constant infinitely high, thereby making the mechanism useless. The reason that this problem doesn't occur in the original CTm is the fact that here only relative valuations are taken into account. Whether agents add a constant to their valuations doesn't influence either the outcome of the vote or the advantage to the individual participant. The CTm doesn't reveal the *absolute* preferences of agents, but only the *relative* preferences.

We solve this problem by giving the agent an incentive to keep his prices low. One way to do this is too make sure that there are also proposals taken into consideration in which the agent doesn't participate. If the agent would then make his prices higher than they actually are, it risks the possibility that one of the proposals in which it doesn't participate is chosen. It then doesn't earn any credits on his account, which he doesn't like. A stronger condition is the following: every agent is required to vote on every alternative, but for every alternative an agent will find that at least one valuation is already fixed. This will typically be a valuation for a plan repair proposal in which he doesn't participate. For instance, there are three proposals to be voted over: A, B and C. Agent X is involved in both A and B, but C is a proposal in which the agent doesn't have to do anything. It must now submit its prices for A, B and C, but the ATC agent has already set the valuation of C on zero. It will now submit its *real* valuations of A and B to the ATC agent, since the original Clarke Tax principles apply: it doesn't want to overbid on A and B because it then runs the risk of paying more tax than the difference in utility with C.

6 Richness and Exploitation

We've given the agents accounts in order to achieve fairness, where a high balance gives an advantage in planning. This presents us with another problem however. Especially when the balance between optimality and fairness is set close to optimality, some agents might become *rich*. This can be because such an agent is often the one who is best able to help out in a problem situation, thus earning a lot of credits. Some agents might be in an economically better position than others, making it easier for them to earn credits. Richness has the disadvantage that it leads to locally suboptimal and unfair solutions. For instance, if A is a rich agent and (a b c) is a solution that costs a credits to A, b credits to B and c credits to C, then the solution (0 3 3) might be preferred above (1 1 1), while the latter is clearly locally more optimal. This is because the algorithm tries to spare rich agents and (0 3 3) is a solution that spares agent A. In other words, richness means power. One might say that if the slider is set close towards optimality, the effect described here will be very small. Although this is true, if the slider is set towards optimality, agents might also become richer than in a more fair mechanism, since richness is less suppressed. This enhances the effect again. Of course, having locally suboptimal solutions in order to achieve a globally optimal solution is not very alarming from a mathematical point of view. However, for a human air traffic controller it might be unacceptable to choose (0 3 3) over (1 1 1). Suppose that the first solution entails agents B and C to change gates, while the second solution only entails all three agents to be delayed for five minutes. A human air traffic controller might very well not allow agent A to 'buy out' B and C, no matter if it has contributed a lot in the past.

There are two ways by which the effects of richness can be suppressed. First we can introduce a notion of *local fairness* to play a role in the decision making. We say that a protocol is locally fair if its individual decisions are fair. An individual decision is fair if every participant has an equal share in the costs. Here we can use a optimality/fairness factor again, for individual proposals. The fairest proposal is:

$$p = argmin_p \sum_{i=1}^{k} {v_{i,p}}^t$$

where k is then number of agents, $v_{i,p}$ is the valuation of agent i to proposal p and t is a 'slider value' controlling local fairness. If we want to include local fairness in the mechanism we have used so far, we decide on a proposal as follows:

$$p = argmin_p \sum_{i=1}^{k} (b_i + {v_{i,p}}^t)^s$$

where k is then number of agents, b_i is the balance of agent i, $v_{i,p}$ is the valuation of agent i to proposal p, t is a slider value controlling local fairness and s is a slider value controlling global fairness. Although ensuring local fairness doesn't prevent agents from becoming rich, it does lessen the effect of it and it slows down enriching because costs tend to be more evenly distributed per enforced proposal.

Another way to lessen the effects of enriching is to prevent enriching itself. This can be done for instance by *progressive taxes*, where agents with a high balance are taxed - credits are taken off their account. In effect this means that privileges earned by helping others don't last forever. This has an intuitive motivation. If at a busy moment at a certain airport an aircraft from Air France is delayed to prevent congestion, then Air France would object if right after it it would be the same airline again that gets the next delay. However, if the previous time it happened was two days ago, Air France wouldn't object so much to the delay. There are two values that determine the effect of progressive tax: the tax function and the taxing frequency. The tax function should be chosen such that very rich agents become less rich, but fairness isn't affected too much. The frequency should should be chosen in such a way that earned privileges evaporate over a reasonable amount of time. For instance an agent that is X credits richer than the average agent should see this edge halved after two days.

One pitfall that we haven't dealt with yet occurs in the following situation. Suppose we use the CTm with progressive taxes as the decision making mechanism. Suppose that in a very tightly packed schedule a conflict occurs that can most easily be solved by one aircraft being delayed for a short time. The conflict can also be solved in other ways, but these involve many more actions and are much more expensive. The single agent that is involved in the easy solution knows all of this, or has at least a very strong suspicion that the simple solution is much cheaper than all the other solutions. For the simple proposal, it can now submit a price higher than its actual costs, thereby earning more credits. For the proposals in which it isn't involved its valuation is fixed by the ATC-agent at zero. This doesn't however keep him from overbidding; he can make a good estimation of how far he can overbid without running the risk of changing the outcome and thus paying a Clarke Tax. In this way he earns more credits than he is entitled to. We will call this principle *exploitation*, since the agent finds himself in an advantageous position and can exploit this.

Unfortunately, it is not easy to counteract exploitation. Just as in real life, people in key positions have more power than others. One drastic measure against exploitation would be the fixing of an initial solution by the ATC-agent. This initial solution is the solution the ATC-agent thinks is the most optimal. The airline agents can then negotiate over alternatives to this initial solution, using the CTm. All costs to the agents are now relative to the initial solution. The effect of exploitation is now greatly reduced, because if the conflict is such that it obviously needs the cooperation of a certain party to be solved, the initial solution will already entail this cooperation and the party involved looses its advantage. A clear deficit of this method is that it is not very fair. Agents are not compensated for the costs that are involved in executing the initial solution, not even if they decide on another solution. A workaround would be to let the ATC-agent estimate the costs of the initial solution to all the agents involved. Thus, when voting over alternatives, the valuations of the initial proposal are fixed for all agents. This will bring some fairness back into the mechanism, although a lot depends on the estimation capabilities of the ATC-agent of course. This method is also guaranteed to raised discussions. If the ATC-agent comes up with an initial solution where an aircraft X is delayed by 10 minutes and if it estimates its costs to be 10 credits, this aircraft might respond with: "No, no! There are transfer passengers on this plane who will miss their connection flight, the cost of this solution is 500 credits!". Thus, the estimation capabilities of the ATC-agent are crucial. It would be most interesting to see how this mechanism might be augmented with argumentation, by which agent could found their submitted costs. This is however out of the scope of this article.

7 Conclusion and Further Research

In airport traffic planning, efficiency and fairness are the two most important factors besides safety. Last minute plan repair is nowadays done by humans using stringent rules of thumb that satisfy these principles to a small extend only. We've introduced a weighted voting mechanism with accounts to achieve both more efficient and fair plan repair solutions. We provided a slider value by which an airport can regulate the pay-off between optimality and global fairness. We've identified a number of pitfalls in the CTm and how to avoid them. In order to keep agents from overbidding, a central agent should fix for every agent at least one valuation of a proposal, preferably a proposal in which it doesn't participate. If an agent still has an incentive to overbid because of exploiting, the central agent should estimate the costs of an initial solution. To prevent agents from becoming rich and thereby gaining to much power, either local fairness or progressive taxes can be enforced. We've shown how local fairness can be achieved, again with a slider value to regulate the pay-off between local optimality and local fairness.

A number of questions about the mechanism still need to be answered: how are the alternatives generated and how many? Where does the collected Clarke Tax go? Should parties of different size should have different negotiation power? Further research could be done on the rules of the voting mechanism in terms of time: when does an agent need to submit its valuations and what happens if an agent is not on time? Also, if the conflict to solve is very urgent, can we make sure the voting procedure ends within a set time period? Another interesting line of research would be to see whether argumentation can be used to substantiate agents' cost submissions or to correct cost estimations. If so, the question whether agents' arguments can be verified becomes important, as well as what should be done if an argument turns out to be false.

We intend to test our mechanism in simulations of the ATP problem with real-world data. We expect this to give us an insight into a number of issues described above: the efficiency of the mechanism, the ideal pay-off between optimality and fairness, the effects of local fairness or progressive taxes and voting rules to name a few.

References

[1] Kenneth J. Arrow. *Social choice and individual values*. J. Wiley, New York, 1951. 2nd edition, 1963.

[2] Edward H. Clarke. Multipart pricing of public goods. *Public Choice*, 18:19–33, 1971.

[3] Eithan Ephrati and Jeffrey S. Rosenschein. Deriving consensus in multiagent systems. *Journal of Artificial Intelligence*, 87(1-2):21–74, November 1996.

[4] Radio Technical Commission for Aeronautics. Final report of RTCA task force 3: Free flight implementation. Technical report, Washington DC, October 1995.

[5] European Organisation for the Safety or Air Navigation. Airport CDM applications guide, 2003.

[6] Theodore Groves. Incentive in teams. *Econometrica*, 41:617–631, 1973.

[7] R. F. Stengel J. P. Wangermann. Optimization and coordination of multiagent systems using principled negotiation. *Journal of Guidance, Control and Dynamics*, 22(1):43–50, 1999.

[8] V. Lesser, K. Decker, T. Wagner, N. Carver, A. Garvey, B. Horling, D. Neiman, R. Podorozhny, M. NagendraPrasad, A. Raja, R. Vincent, P. Xuan, and X.Q. Zhang. Evolution of the GPGP/TAEMS Domain-Independent Coordination Framework. *Autonomous Agents and Multi-Agent Systems*, 9(1):87–143, July 2004.

[9] D.A. Daner R.L. Schultz and Y. Zhao. Free-flight concept. Technical Report AIAA-97-3677, AIAA, 2000.

[10] Tuomas W. Sandholm. *Negotiation among self-interested computationally limited agents*. PhD thesis, University of Massachusetts at Amherst, 1996.

[11] W. Vickrey. Counterspeculation, Auctions and Competitive Sealed Tenders. *Journal of Finance*, pages 8–37, 1961.

[12] Nicolaus Tideman Wayne Brough, Edward Clarke. Airport congestion and noise: Interplay of allocation and distribution. *Transportation Research Record*, 1450:3–7, 1995.

[13] Chien yu Chen, Michael O. Ball, Robert Hoffman, and Thomas Vossen. Collaborative decision making in air traffic management: Current and future research directions. Technical report, October 04 2000.

[14] XiaoQin Zhang, Rodion Podorozhny, and Victor Lesser. Cooperative, MultiStep Negotiation Over a Multi-Dimensional Utility Function. *Proceeding of the IASTED International Conference, Artificial Intelligence and Soft Computing (ASC 2000)*, pages 136–142, July 2000.