

Towards a Market Mechanism for Airport Traffic Control

No Author Given

No Institute Given

Abstract. We present a multiagent decision mechanism for the airport traffic control domain. It enables airlines to jointly decide on proposals for plan conflict solutions. The mechanism uses weighted voting for maximizing global utility and Clarke Tax to discourage manipulation. We introduce accounts to ensure that all agents are treated fairly, to some extent. The mechanism allows an airport to determine the pay-off between optimality and fairness of schedules. Also, it compensates for agents that happen to be in practically unfavourable positions.

1 Introduction

Airports nowadays 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, like 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.

A current trend in air traffic control (ATC) automation is that of distributed planning. An example is the Free Flight program [1, ?], which enables aircrafts 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 [2, ?].

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.

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 arrive¹. 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 occupy 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. This last phase of planning just before execution is called *tactical planning*. In general the main aim for ATC 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.

An important criterion for ATC to observe in the tactical planning phase 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 burden². 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.

In general, it is hard for ATC to assess plan change costs for airlines and thus to involve airlines’ preferences. On the one hand ATC doesn’t know exactly the state of affairs of an airline, its schedule details and dependencies needed to correctly assess plan change costs. Also, it is hard for ATC to compare for instance the costs of changing gates to changing runways. On the other hand, airlines can not and are not willing to give all the information ATC would need. They can not because this would result in a communication overload for ATC. They also don’t want to give all their information because this can be disadvantageous to them, see for instance [4]. It would be beneficial if the private information of agents can be involved in the decision making process without revealing it.

From the multiagent point of view, ATP can be seen as a coordination problem between self-interested agents that have different preferences and private utility functions. These actors need to agree on an efficient and fair planning to

¹ The fair allocation of slots to airlines is a challenging problem on its own, see for instance [3].

be executed collectively, under continuously changing circumstances. The general challenge is to transform the current ATP situation of centralized planning, authority and responsibility to a decentralized one, utilizing distributed knowledge.

3 Coordination in multiagent systems

In the context of multiagent research, coordination is usually seen as a process between two or more autonomous agents that try to achieve their individual goals, but cannot do so without adapting their behaviour to each other. For instance when agents have plannings that involve shared resources or task interdependencies, they must coordinate their plannings. A framework developed for this kind of cooperation is *generalized partial global planning* (GPGP) by Lesser and others [5]. When some agents are specialized in certain tasks it can be useful to delegate tasks. This gives the problem of *task allocation*, for which the multiagent auction protocol *Contract Net* is developed [6].

In the above-mentioned systems, as in most multiagent systems, the coordination is based on the concept of *mutual benefit gain*; agents engage in a deal or contract when it is beneficial for all of them. In terms of Contract Net, agents engage in a contract if it is *individually rational* for them to do so. In terms of GPGP, an agent accepts a proposal if its marginally utility gain is greater than its marginally utility cost. The same concept is also called *principled negotiation*, for instance in [7] where it is applied to runway slot allocation.

Generally, if mutual benefit gain is used, agents will eventually arrive at a *Pareto optimal* point, i.e., an outcome that cannot be improved without an agent decreasing its utility. A negotiation protocol that arrives at a Pareto optimal solution is called *Pareto efficient*. There are usually many Pareto efficient solutions for a given problem, together forming the *Pareto front*. An important question is which one of these points we want our mechanism to reach. A possible solution is to maximize the product of the utilities of the different agents, generally known as the Nash point [8]. In case the Pareto front is convex, this corresponds to the intuition that a good solution is one that is Pareto optimal and gives agents approximately equal utilities. There are however many other ways to choose a best outcome. One could maximize a weighted sum of the utilities, maximize the minimum utility of the agents, minimize the differences in utilities, etc. Two important underlying motivations play a role here: *optimality* and *fairness*. Maximizing the sum of the utilities is a method ensuring optimality, while minimizing the differences in utility is aimed at fairness.

Using mutual benefit gain is not always desired in multiagent coordination. Take for instance the case where Anne and Ben have inherited a painting and want to decide how to distribute it between the two of them. There are three options: either Anne gets it, Ben gets it or they both get half of it. Obviously, half a painting is of little worth to them, but assume that they prefer half a painting above nothing. The three options can be denoted as tuples of the respective utilities: (10, 0), (0, 10) and (1, 1). All three options are Pareto efficient.

Maximizing the product of the utilities or any other decision function aimed at fairness would have the third option as outcome. Suppose that Anne and Ben settle for this option. The problem arises when Anne and Ben inherit another painting. If they use the same mechanism again, they would end up both with half a painting again. But now they are unhappy. Although the solution of the two negotiations together are fair, they are not Pareto efficient. Both of them would have preferred to get one painting each.

If we look at the two negotiations together it is easy to find a good solution. In this case the solution of one painting each is both optimal and perfectly fair. If we look only at the first negotiation however and we are not sure which other negotiations will ‘come up’, we need different motivations to choose to give the painting to one or the other. This choice can be justified only if we expect more negotiations to follow and we expect to maximize optimality by benefiting one agent now and the other next time. Thus, we have to relax the fairness constraint temporarily but remember the exact utilities of the enforced decisions, to arrive at more optimal and fair solutions in the long run. We will do this more formally in section 6.

4 Mechanism design

Mechanism design is concerned with the design of procedures in which several parties participate to reach a certain outcome, where the procedure has to meet a number of criteria. The previous section has already stressed two important criteria a negotiation mechanism should meet: optimality and fairness. In this section we will list these and a number of other important criteria. For more information on mechanism design we refer the reader to [9].

Optimality - The outcome of a mechanism should be optimal with respect to the utilities the agents ascribe to the outcome. The most common requirement is that the outcome is Pareto optimal. When the sum of the utilities is maximized, we speak of maximizing *social welfare*.

Individual rationality - It should be attractive for a user to participate in the process, i.e., it may only gain utility by participating.

Nonmanipulability - A mechanism should motivate agents to behave in the manner that leads to the desired outcome. Usually this means that an agents dominant strategy should be to be truthful. This is also called *stability* and *incentive compatibility*. We will go more deeply into this subject in section 5.

Fairness - Although usually not explicitly named as a criterion in mechanism design, the ATP case shows that fairness is of crucial importance in airline negotiations. An airline will simply not accept to be delayed repeatedly while others are not. A mechanism should thus to a certain extend provide for an even distribution of utility gain or loss among the agents.

5 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*. An well known method to reach an optimal decision 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 an agent might be tempted 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 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 even without the full force of its vote, he can submit a lower cost, pay less and still get his most preferred outcome. In the literature of economics this is known as the *free rider problem*.

A mechanism is called *nonmanipulable*, *stable* or *incentive compatible* when it yields the optimal social outcome when agents use their dominant strategy. In such protocols participants are best off when they are truthful. An example of a single-object nonmanipulable auction is the *Vickrey auction* [10], where every agent submits its bid for the single object on sale, and the object is 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 [10].

An example of a nonmanipulable decision protocol using cardinal preferences is the *Clarke Tax Mechanism* [11, ?]. 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 the 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 [11] and [12].

The Clarke Tax mechanism (CTm) can be used in subsequent negotiations, such as occur in our domain. An example of this is the nonmanipulable meeting scheduling system of Ephrati et al [13]. In this mechanism agents decide on the best time for a meeting, using credits to express their preferences over alternatives. Clarke Tax encourages the agents to be truthful. However, the article doesn't address the issue of fairness, which is important in our domain. Neither does it identify a phenomenon we will deal with later called *enriching*.

6 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* with *Clarke tax* as the main principle to determine the most preferred solution among the voters, and *accounts* to ensure fairness.

We will describe the problem more formally now. When a planning conflict occurs, ATC generates a set of plan repair proposals $Q = \{q_1, q_2, \dots, q_n\}$. We have k airlines, and each airline a has a utility function u_a describing a monetary valuation of a plan change. Such a utility will typically be a negative number, since plan changes often involve delays and gate or runway changes. As an example, $u_a(p, q) = -1000$ is to be read as "if the current planning is p , it would cost airline a 1000 euro to switch to planning q " or "airline a is indifferent between changing the current plan p to q or sticking with p and paying 1000 euro." It is of course possible that an airline has a positive utility for a plan change, for instance when an airline had a flight delayed against its will and the new proposal cancels this delay.

Each airline gives its utility as a weighted vote for every proposal. Clarke tax is in effect to prevent manipulation. If we add up all the valuations of the airlines for a single proposal, we get the group utility of this proposal:

$$U_G(p, q) = \sum_{a=1}^k u_a(p, q)$$

We will need the average utility per agent later on:

$$\bar{u}(p, q) = \frac{U_G(p, q)}{k}$$

To measure the (group) fairness of a proposal, we could add up the amounts to which each agent's utility differs from the average utility of this proposal. This would give us an *unfairness* measure:

$$\mathbf{unf}_G(p, q) = \sum_{a=1}^k |u_a(p, q) - \bar{u}(p, q)|$$

However, we want to involve the agents' histories in the calculation of fairness. If an agent was previously burdened with a lower-than-average utility, it would be fair to compensate this with a higher-than-average utility the next time. Therefore we ascribe *accounts* to the agents on which the monetary utilities of the enforced decisions in the past are stored. The idea is that if an agent is involved in a planning conflict repair that has utility x to him, he earns the inverted utility, $-x$, on its account. So if an agent has an account of 20000 euro, it means that its contributions to conflict repairs in the past have cost

him 20000 euro. When calculating the unfairness of a proposal, we now involve these accounts to get an *contextual unfairness*, i.e., the fairness of a proposal in the context of previous decisions. Given the current balances of the agents $B = \{b_1, b_2, \dots, b_n\}$, the current planning p and a proposal q , the contextual unfairness is calculated as follows:

$$\text{unf}_C(B, p, q) = \sum_{a=1}^k |(b_a - u_a(p, q)) - \overline{u}_C(B, p, q)|$$

where

$$\overline{u}_C(B, p, q) = \frac{\sum b_a - u_a(p, q)}{k}$$

Thus, instead of comparing utilities of single decisions, we take for each agent the sum of its balance and the inverted utility for the proposal in question and compare those.

This gives us for each proposal in Q both an unfairness measure and a optimality measure. Figure 1 illustrates the difference between a fair but suboptimal solution and an optimal but not so fair solution. The bars denote for nine agents the balances plus the inverted utilities or *costs*. At the top the current balances of the agents are depicted; each agent has already contributed some costs to plan repairs in the past. There are two proposals to choose between. At the left the inverted utilities for the first proposal are added to the balances, resulting in the new balances if this proposal would be chosen. This shows a relatively evenly distribution of the new balances. On the right, the second proposal is depicted that is more optimal, i.e., has a lower group cost / higher utility, but leads to a less fair distribution of the balances. Thus, each proposal has its fairness and optimality value. For a whole set of proposals, we could depict these values in a diagram. Figure 2 shows such a diagram for an example proposal set. The point on the far right for instance corresponds to a proposal that has a global cost of only 1000 euro, but is not so fair since agents utilities differ on average 1600 euros from the average utility. The point on the far left top corresponds with a very fair but very suboptimal solution.




Fig. 1. Fair versus optimal.




Fig. 2. An example proposal set.

Out of these proposals one has to be chosen. It requires a weighing between optimality and fairness to decide on the winner. This weighing resembles the weighing of agents utilities we described in section 3. This time however fairness is explicitly weighed against optimality. This calls for other weighing functions - fairness and optimality are different concepts than individual utilities. Depending on the airport, the mechanism designer might choose to ensure a minimum level of fairness and then maximize optimality, or vice versa. It could also maximize a weighted sum of the two or choose the point closest to an ideal point. For the remainder of this article we will assume that the mechanism uses a weighted sum of optimality and fairness as its decision function.

7 Nonmanipulability

An important question is whether the introduction of accounts affects the non-manipulability of CTm. First note that CTm stimulates the agents to reveal their true preference relation over the proposals, but that this doesn't guarantee that they bid their *exact* valuations. An agent might just as well bid too high values for all of the proposals by adding a constant value c to its utilities, thus bidding $u_a(p, q) + c$ for every proposal q . CTm makes agents reveal their *relative* utilities, but not their *absolute* utilities. Adding a constant value to its bids doesn't influence the tax the agent might have to pay nor does it make one proposal more optimal compared to another. It can however influence the relative fairness of proposals in this way. Furthermore, if the agent uses a negative constant value, it will make him acquire a higher balance on its account than if it didn't.

A remedy for this kind of manipulation is to fix some of the agents bids. If for a set of proposals, every agent finds at least one of his bids been set for him already, it will have to make its other bids relative to this one. Of course, if these fixed utilities are different than the real utilities of the agents, things are not fair. It is however possible to correctly fix the utilities for some of the proposals, namely the proposals in which an agent doesn't participate. It is realistic to say that a plan repair proposal in which an airline doesn't have to do anything has a utility of zero for this airline. So if we make sure that for every agent there is at least one proposal in the proposal set in which he doesn't participate, and the bids for these proposals are fixed to zero, we make it unattractive to overbid or underbid.

Nevertheless, because of the introduction of accounts the mechanism isn't strictly nonmanipulable anymore. In the original CTm the agent could never gain more than it would lose by lying. In our mechanism however, the extra credits earned by lying might in some cases be preferable above the extra tax the agent has to pay, especially when the amount of tax he has to pay is less than the amount he lied about. We will confine ourselves to giving some arguments against lying-for-credits in section ??.

8 Richness

In the mechanism described so far agents jointly vote over which proposal to enforce, are best of by being honest and can trust that they will not be very unfairly treated. An understandable worry that one might have is the question whether some agents aren't worse off than others still. In a credit based decision mechanism, an agent can be *rich*, i.e., it has a high balance on its account. Being rich means having a lot of negotiation power, since the mechanism favours agents with a lot of credits. Wouldn't it be reasonable to assume that some agents are in a *better economic position* than others, for instance because they are often able to help others, thus earning a lot of credits and becoming rich en mighty? Indeed these differences will surely exist. Some airlines might have their gates close together, making it easier for them to swap, other airlines might have flights from or to busy destinations, making it harder to change landing and take-off times. In other words, some agents will be in *practically preferable positions* compared to others. Fortunately, in the mechanism we described these practical differences don't lead to better economic positions. In fact, the mechanism smooths out practical advantages and disadvantages agents might have. Take for instance two airlines that can both solve the same problem. The first happens to be in such a situation that it is easy for him to solve it, for the second it is harder. The first is in this respect in a practically preferable position compared to the other and will bid a higher utility, the second agent will bid a lower utility. The first few times the problem occurs the first agent will be elected to solve the problem. When his expenditures are so high that it would no longer be fair that he is chosen again, the other agent is chosen. In this way agents in practically preferable positions are more often involved in plan repairs than others. Similarly, agents in practically unpreferable positions are compensated for this by having to contribute less to plan repairs.

There are other reasons however why richness can be undesirable. Note that richness can lead 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 and fair. 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.

From a global perspective, there is nothing wrong with a rich agent being favoured in decisions. This agent has already contributed a lot to the system, so it is only just if he is spared in future decisions. But from the perspective of an air traffic controller, this might not always be acceptable. It might very well be that in the example above a human air traffic controller would object against the fairer solution $(0\ 3\ 3)$ simply because it would in such a situation always choose $(1\ 1\ 1)$, no matter what happened in the past. Thus, *local fairness* or *fairness per decision* plays a role in human decisions. If we want our mechanism to have the same behaviour, we could do this by putting a maximum on the fairness per decision $U_G(p, q)$ or integrate it more subtly into the decision function.

There is however a more pressing reason why agents shouldn't become too rich. Privileges that are earned by helping others don't last forever. If for instance at a certain moment an aircraft from Air France is delayed and one from British Airways is not, and later that day it is exactly the other way around, both parties would agree. But if these two events are one month apart, British Airways would certainly not accept the first event as a justification for its delay. Apparently, earned privileges wear off. This means that in our mechanism, earned credits shouldn't last forever. The simplest way to achieve this is to impose a tax on the agents' accounts such that it reduces the differences in richness between agents. The tax function and frequency should be chosen such that it corresponds to reality; richness should be suppressed over time while fairness should still hold in short term. For instance, it could be chosen such that an earned amount of credits is halved after three days.

9 Exploitation

One pitfall that we haven't dealt with yet occurs in the following situation. Suppose we use the mechanism described as the decision making mechanism with taxes in effect. 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 changing the outcome and having to pay 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 loses 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 raise 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 200 euro, 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 5000 euro!”. 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.

10 Related work

The mechanism described here resembles the nonmanipulable meeting scheduler of Ephrati et al. [13], in which the CTm is used in repeated decisions. The issues of richness and exploitation we have identified in this paper apply to this mechanisms as well. Therefore, our work adds to the already many insights into the workings and application of the CTm [14].

11 Conclusion and further research

Decentralization is a trend that is gaining ground in air traffic control systems nowadays. In this context we have looked at automating the tactical planning and plan repair phase of airport traffic planning. In current practise, this work is done by humans, who use stringent rules of thumb to comply to the strict safety constraints. Automating this process can lead to an improvement in the areas of optimality and fairness for the airlines.

We adopt a multiagent point of view, approaching ATP as a coordination problem between autonomous, competitive agents using private information whose plannings are highly dependent on each other. We introduced a weighted voting mechanism for distributed decision making with Clark Tax to discourage manipulation. To ensure a certain degree of fairness, we’ve introduced accounts on which an agent’s contributions are stored. These accounts ensure that disadvantages resulting from different planning disruptions are fairly distributed over airlines. We have shown that the mechanism designer can determine the pay-off between optimality and fairness. In reality, some agents will be in better practical positions than others. We have shown how the mechanism compensates for this.

We’ve identified a number of pitfalls in the mechanism 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 too much power, either local fairness or taxes can be enforced.

Also we would like to look more closer to the matter of nonmanipulability. Although the mechanism is not strictly nonmanipulable, we've given several arguments why lying is unattractive. Can we formally define a measure to which extent an agent can manipulate? What are the assumptions that make the mechanism manipulable? Could we find assumptions under which the *expected* utility of manipulation is negative?

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, to what extent richness will occur, the effect of taxes and voting rules to name a few.

References

1. Radio Technical Commission for Aeronautics: Final report of RTCA task force 3: Free flight implementation. Technical report, Washington DC (1995)
2. European Organisation for the Safety of Air Navigation: Airport CDM applications guide (2003)
3. Brough, W., Clarke, E., Tideman, N.: Airport congestion and noise: Interplay of allocation and distribution. *Transportation Research Record* **1450** (1995) 3–7
4. Chen, C., Ball, M.O., Hoffman, R., Vossen, T.: Collaborative decision making in air traffic management: Current and future research directions. Technical report (2000)
5. Lesser, V., Decker, K., Wagner, T., Carver, N., Garvey, A., Horling, B., Neiman, D., Podorozhny, R., NagendraPrasad, M., Raja, A., Vincent, R., Xuan, P., Zhang, X.: Evolution of the GPGP/TAEMS Domain-Independent Coordination Framework. *Autonomous Agents and Multi-Agent Systems* **9** (2004) 87–143
6. Smith, R.: The contract net protocol: High-level communication and control in a distributed problem solver. In: *IEEE Transaction on Computers*. Number 12 in C-29 (1980) 1104–1113
7. Wangermann, J., Stengel, R.: Optimization and coordination of multiagent systems using principled negotiation. *Journal of Guidance, Control and Dynamics* **22** (1999) 43–50
8. Raiffa, H.: *Lectures on Negotiation Analysis*. Program on Negotiation at Harvard Law School, Harvard Law School, Cambridge (1996)
9. Parkes, D.C.: *Iterative Combinatorial Auctions: Achieving Economic and Computational Efficiency*. PhD thesis, Department of Computer and Information Science, University of Pennsylvania (2001)
10. Vickrey, W.: Counterspeculation, Auctions and Competitive Sealed Tenders. *Journal of Finance* (1961) 8–37
11. Clarke, E.H.: Multipart pricing of public goods. *Public Choice* **18** (1971) 19–33
12. Groves, T.: Incentive in teams. *Econometrica* **41** (1973) 617–631
13. Ephrati, E., Zlotkin, G., Rosenschein, R.: A non-manipulable meeting scheduling system. In: *Proceedings of the 13th International Workshop on Distributed Artificial Intelligence*, Seattle, WA (1994)
14. Ephrati, E., Rosenschein, J.S.: Deriving consensus in multi-agent systems. *Journal of Artificial Intelligence* **87** (1996) 21–74