

An Agent-based Electronic Market to Help Airlines to Recover from Delays

Luís Reis¹, Ana Paula Rocha² and Antonio J. M. Castro²

¹FEUP, University of Porto, Portugal

²LIACC, FEUP, DEI, University of Porto, Portugal

Keywords: Electronic Market, Multi-agent Systems, Negotiation, Case-Based Reasoning, Air Transport, Disruption Management, Flight Disruption, Irregular Operations, Recovery Process.

Abstract: The Airline Operations Control Center (AOCC) has the responsibility to ensure that flights meet their planned schedule or, if any problem arises, to find a viable solution that minimizes both the impact in the operational plan and its cost. The high cost of resources involved in this process (aircraft and crew members) leads to a lack of additional resources from the airline companies, implying a restricted solution space. Here, we propose an electronic market modeled as a multi-agent system where airline companies can negotiate and lease each other the required resources when solving a disruption problem, thus expanding their solution space. The proposed negotiation occurs in several rounds, where qualitative comments made by the buyer agent on proposals sent by the sellers enables these to learn how to calculate new proposals, using a case-based reasoning methodology.

1 INTRODUCTION

According to Kohl (Kohl et al., 2004), "research on the recovery operation to this date only deals with a single airline. Cooperation between airlines is not supported". Nowadays, each airline tries to solve the operations recovery problems with their own resources (Castro et al., 2014). If they have an open position for a specific type of crew in a flight, they try to find a suitable one from their own staff. The same happens with aircraft.

Sometimes, the airlines have to rent aircraft and crew members when needed (known in the industry as ACMI - Aircraft, Crew, Maintenance and Insurance), but through a direct contact with charter airlines. It is not a usual practice to use only crew members (without being part of the aircraft) from other companies.

The electronic market (EM) that we propose in this paper, is a permanently open virtual marketplace where registered airlines (represented by software agents) can meet each other to purchase services and has the possibility to be integrated with systems or tools for airline operations control, like the one we use in this paper. It has the following advantages:

- Airlines that participate in this EM will have more resources available to solve their problems.
- Airlines may take advantage of exceeding re-

sources in specific dates and times and sell services performed by these resources to other airlines.

- Can reduce costs and time for the airline that has a specific problem.

When compared with the work of Malucelli (Malucelli et al., 2006) we complete the work by proposing a negotiation algorithm for the EM.

Airlines have an organization called Airline Operations Control Center (AOCC) that has the responsibility to ensure that flights meet their planned schedule or, if any problem arises, to find a viable solution that minimizes both the impact in the operational plan and its cost. Research in the air transportation domain has shown that airline companies lose between 2% to 3% of their annual revenue as consequence of disruptions and, that, the impact caused by small disruptions in companies' profits can be reduced by at least 20%, through a better recovery process (Chen et al., 2010).

Currently, operations management is essentially a manual process, supported by tools, that among other functions include monitoring, event detection and problems resolution and, strongly depends on the tactical knowledge of the AOCC's members (Castro et al., 2014).

Every time an irregular event that has an impact on the scheduled plan is detected, the AOCC's team

has to plan carefully an alternative schedule, to ensure that it minimizes at most the disruption cost. A disruption can be viewed as composed by four dimensions (Castro et al., 2014): aircraft, crew member, passenger and flight. In disruption management, the AOCC commits to recover all the dimensions affected.

To test our proposed EM, we use MASDIMA (Multi-Agent System for Disruption Management) (Castro et al., 2014) that addresses the Aircraft, Crew and Passenger recovery problem using an approach that is able to recover all problem dimensions simultaneously.

In Multi-Agent Systems it is required for an agent to interact with other agents whom may not share common goals. This leads to the need to reach agreements (Wooldridge, 2009) through an automated negotiation process. The existent automated negotiation systems are composed of three major groups (Oliveira and Rocha, 2001):

- Auctions;
- Game Theory;
- *Negotiation*.

Although it only considers a single attribute, due to its simplicity and well predefined rules, auctions (Vulkan and Jennings, 2000) are a very popular negotiation mechanism.

Game Theory (Rosenschein and Zlotkin, 1994) is a mechanism that can only be applied to perfect information and rationality contexts.

Negotiation is the generic name given to other techniques where agents must reach agreements on matters of mutual interest (Wooldridge, 2009). These techniques are more flexible than auctions and game theory in terms of preexistent protocols and rules, thus more suitable for open and dynamic environments (Oliveira and Rocha, 2001).

From the *Negotiation* group, stands out the multi-attribute negotiation, which is useful in the situation where the negotiation decision does not consider only one attribute but multiple attributes as it is in the case of AOCC. For instance, when buying any product, the buyer considers the price as an important attribute in its decision but the delivery time or the product quality may also be (usually are) factors to be considered in the decision of buying or not a certain product.

Giving different utility values to the different attributes under negotiation solves the problem of multi-attribute evaluation. The most common proposal evaluation formula is a linear combination of the attribute correspondent values, weighted by the respective utilities. Therefore, a multi-attribute negotiation is converted to a single-attribute one, to be made over the evaluation value (examples of this are

the work in (Oliveira et al., 1999), (Vulkan and Jennings, 2000), (Matos et al., 1998) and (Cardoso and Oliveira, 2000)). This is also the approach followed in our work, although we agree that, in some cases, it can be difficult to give an exact numeric value to an attribute utility. A solution which leads to a more intuitive situation, can be just to impose a preferential order over the domain values for the different attributes or on the attributes itself.

The multi-agent system based Electronic Market, presented in this paper, allows companies to negotiate among themselves the missing resources. The negotiation algorithm includes case-based reasoning to learn how to make a counter-proposal. According to Riesbeck and Schank (Riesbeck and Schank, 2013), "A case-based reasoner solves problems by using or adapting solutions to old problems.", i.e. case-based reasoning (CBR) focuses on the reuse of knowledge acquired from previous experiences in order to solve new problems. Like humans do, CBR is a problem solving paradigm that uses incremental and sustained learning since new experiences are retained each time a problem is solved making those available for future problems. A negotiation algorithm with a CBR approach has never been considered in the works mentioned.

The rest of this paper is as follows: section 2 is the main section and presents the proposed multi-agent system EM. In section 3 it is presented the scenarios but only one of the experiments done, as well as the results obtained. Finally, section 4 concludes the work presented.

2 AIRLINE ELECTRONIC MARKET SOLUTION

When a disrupted flight is detected, the AOCC's team should find an alternative trying to minimize both the delay and the disruption cost. The airline electronic market proposed here intends to help the airline company in this disruption management process, by allowing to find external resources, possibly less costly or available sooner than the company's own. In this market there are two types of entities:

- The buyer, that represents an injured airline company. This is the airline company that has a disrupted flight (an unexpected event causing a delay in the flight).
- The seller, that represents a service provider airline company

Being the object under negotiation, a *Need* is identified by the resource(s) needed: a list of crew mem-

bers and an aircraft fleet as well as relevant information related to the disrupted flight (scheduled departure time, trip time, delay, origin and destination), as depicted in equation (1).

$$Need = \langle Res, STD, TripD, Del, Orig, Dest \rangle \quad (1)$$

with $Res = \langle CrewList, Aircraft_Fleet \rangle$

where:

- Res is the resource(s) needed
- STD is the trip duration
- Del is the delay of the disrupted flight
- $Orig$ is the airport origin
- $Dest$ is the airport destination

Buyer and sellers will negotiate the resource that buyer identifies as its need. This resource can be a set of crew members, an aircraft or both. When the resource under negotiation is an aircraft, it is required that a crew to handle it should also be provided. For the negotiation to take place, buyers and sellers need to know each other. Sellers should register first, otherwise will be no one in the market to be asked for some resource(s). So, the first step is to have multiple sellers registered and wait for some buyer to register too. When a buyer registers in the market, it retrieves a list containing all registered sellers and starts a negotiation with them. The negotiation is a process where proposals are exchanged between buyer and sellers until an agreement is reached between the buyer and one of the sellers, and the negotiation ends successfully, or no agreement is reached and the negotiation fails.

2.1 An Adaptive Negotiation

The negotiation protocol proposed for the airline electronic market is based in the FIPA Iterated Contract Net and was chosen because it allows multi-round iterative bidding. This way, it is ensured that a wide space of solutions is subject to discussion and refinement, as it is the case of humans' negotiations. This protocol works with an initiator (buyer in this case) and multiple responders (sellers).

The negotiation protocol is depicted on figure 1. The buyer initiates the negotiation, by sending to all sellers an *Invitation* message, (CFP - Call For Proposals), containing relevant information about the disrupted flight and resource needs (as indicated in (1)). After sending the *Invitation* message, the buyer gives a timeout for sellers to respond. If a seller did not respond after that timeout, it is removed from the negotiation.

When a seller receives the *Invitation* message, it processes the message verifying if it is able to provide the required resources or not. If yes, the seller replies with a *Proposal* message, containing the price and availability of its proposal (equation (2)).

$$Proposal = \langle \alpha, \rho \rangle \quad (2)$$

with $\alpha \in [0, delay_of_disrupted_flight]$

where:

- α is the proposed availability
- ρ is the proposed price

If the seller is not able to provide the resources required by the buyer, it replies with a *Refusal* message. The first round is now concluded and until the end of the negotiation, all rounds are processed the same way, explained as follows.

Buyer Role: Buyer receives one proposal from each interested seller, evaluates all proposals and selects the best one of the current round according to its utility (agents' utility is explained in section 2.2). If the best proposal of the current round is better than the best one found in previous rounds (if any), it is considered the new best proposal. If not, the best proposal remains unchanged. Buyer creates then a reply for each received proposal, issuing a qualitative feedback over the availability and price in it, by comparing these values with the ones in the best proposal. This reply or *Feedback* message is sent to all sellers that are currently in the negotiation (equation (3)). The qualitative comment included in the *Feedback* message ($QIEv$ in equation (3)) can assume one of the three options:

- OK: means there is no need to improve the attribute that received this feedback
- LOWER: means the attribute that received this feedback has a high value, should be reduced
- MUCH LOWER: means the attribute that received this feedback has a very high value, should be greatly reduced

$$Feedback = \langle QIEv_\alpha, QIEv_\rho \rangle \quad (3)$$

where:

- $QIEv_\alpha$ is the qualitative evaluation of the proposed availability
- $QIEv_\rho$ is the qualitative evaluation of the proposed price

Seller Role: When a seller receives the feedback for the proposal sent, it updates its experience history, by recording and reasoning the concerned feedback. Sellers will use its experience history (similar to what humans do) to formulate new proposals during the

current and future negotiations. If a seller does not have any more proposals to propose, it sends a *Refuse* message. This process is explained in more detail in section 2.3.

The negotiation is over when all sellers have sent a refusal message or a deadline is reached. In the last round of the negotiation, buyer sends an *Accept* message to the best proposal's owner, with the accepted proposal data and a *Reject* message to all others. The seller that received the accept message sends back to buyer a *Termination* message with all relevant data about the *Need*. Upon receiving the termination message, buyer unregisters himself from the market as the negotiation is over.

Note that messages exchanged during all the negotiation ensure that agents' (buyers or sellers) information is kept private. Agents never reveal their costs or utility. For instance, if sellers would know the buyer's disruption cost, their strategy would be to ask for a price slightly lower than that cost, making the market an unpractical alternative for buyer.

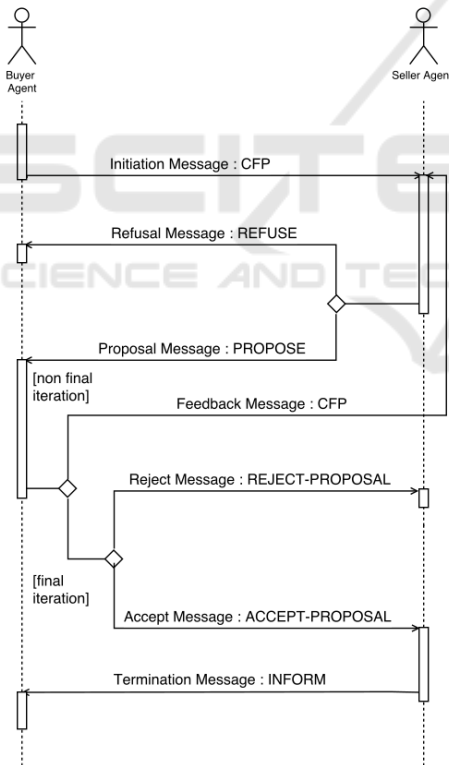


Figure 1: Negotiation Protocol.

2.2 Agents' Utility

Agent's utility is a way of representing its preferences over a set of possible alternatives, in this case, a set of possible *Need*, which will constitute the negotiation

set. The higher the utility, more preferred is the correspondent *Need*.

Buyer Utility. Buyer uses the utility to evaluate the proposals sent by the different sellers in response to its *Invitation* or *Feedback* messages. Remember that a proposal tells the buyer what are the conditions under which a seller provides a *Need*, and contains two values: the proposed availability for the disrupted flight, and the rental price. In buyers perspective, the utility of a proposal must measure its availability and price, where it tries to minimize both. Its value ranges from 0 to 1, and its calculation follows equation (4).

$$\mu = \mu_{\alpha} \times \beta + \mu_p \times (1 - \beta) \quad (4)$$

with $\beta \in [0, 1]$ where:

- μ is the utility of the proposal
- μ_{α} is the utility of the availability parameter $[0, 1]$
- μ_p is the utility of the price parameter $[0, 1]$
- β is the importance factor of the availability parameter

As shown in equation (4), to the buyer, the utility of a proposal is composed of two parcels, the availability utility (μ_{α}) and the price utility (μ_p). The main purpose of the availability utility (μ_{α}) is to relate the availability proposed by the seller with the delay of the disrupted flight.

Seller Utility. In the case of sellers, the measure of how preferable a proposal is, is given exclusively by its price. This is materialized in the fact that sellers do not need to fulfill any disrupted schedule but it just has to compensate the leasing associated cost. From the seller point of view, a proposal is more profitable the higher is the price. The seller utility is given by equation (5), and its value ranges from 0 to 1.

$$\mu = \frac{\rho - \gamma \times \zeta}{(\sigma \times \zeta) - (\gamma \times \zeta)} \quad (5)$$

where:

- μ is the utility of the proposal
- ρ is the proposed price
- γ is the minimum price multiplier
- σ is the maximum price multiplier
- ζ is the leasing associated cost

The price multiplier is a static interval generated to ensure that seller is not impaired, what would happen if the price was only the leasing associated cost. This way, it is ensured that seller gets some profit even with a low utility deal.

2.3 Agents Learning through Case-based Reasoning

Sellers use CBR (Case-based Reasoning) to decide what to do upon receiving the buyer feedback over the proposal they have sent, consulting a record of previous experiences classified according to its usefulness. The object that represents an experience, along with its usefulness, is called *case* and is represented by a set of parameters, grouped into three types (*Features*, *Solution* and *Evaluation*), as shown in figure 2.

The parameters in *Features* identify the situation of the current case, regarding the feedback buyer gave, the number of sellers in the negotiation and the identification of the resource under negotiation (aircraft or crew). The parameters in *Solution* identify the actions performed (price changing, availability changing) in that specific situation. The parameter in *Evaluation* assigns an evaluation value to the case, that measures its usefulness.




Figure 2: Case Composition.

Find Similar Cases

CBR starts by retrieving similar cases to the one received. For this purpose, only parameters identified as *Features* in figure 2 are used to compare cases and to identify equal ones. Although all features are used, they do not have the same preponderance on the task, because the feedback over a proposal is more relevant to decide the action to be made than the others parameters. So, to each parameter in *Features* is given a weight. Similar cases are found through the euclidean distance between them, a distance of 0 means that the case is identical, a distance greater than 0 means a different case. To ensure that features' weight has relevance in the distance calculation, the weight was added to the well known euclidean distance formula. So, the distance between two cases is the weighted sum of their features distances, as presented in equation (6).

$$d(\kappa, \chi) = \sum_{\eta} \sqrt{(\kappa_{\eta} - \chi_{\eta})^2} \times \varepsilon_{\eta} \quad (6)$$

with $\eta \in \{\text{Features}\}$

where:

- κ is the case received
- χ is one of the case in the data set

- η is the current feature
- ε_{η} is the weight of feature η

Select a Case

After identify the set of similar cases, the seller has to select one of them to apply at the current situation, what it does using a softmax algorithm (Sutton and Barto, 1998). This algorithm applies a probability to each similar case retrieved from CBR, where a case's probability is greater the higher its evaluation value. The probability of a case being selected is given by equation (7).

$$P(k) = \frac{e^{Y_k}}{\sum_{i=1}^n e^{Y_i}} \quad (7)$$

where:

- Y_i is the evaluation of case i
- $P(k)$ is probability of the item k being selected
- n is the number of similar cases

After being assigned a probability to each case, a random value is generated and the first case with cumulative probability greater than that random value is selected. The new proposal to be sent by the seller is generated by applying the actions enumerated in the selected case.

If no similar cases exist in the history set, the seller generates a new proposal by following the qualitative comments in the feedback received.

Update History

In order to prevent obsolete cases, every time an experience is reproduced, its evaluation is updated, where the latter the experience, the more important its evaluation is. The evaluation of an experience is updated as equation (8) shows.

$$Y = Y_{n-1} * (1 - I) + Y_n * I \quad (8)$$

where:

- Y is the updated evaluation value
- Y_{prev} is the evaluation value of an equal experience found in the history set
- Y_{curr} is the evaluation value for the current experience
- I is the weight given to the most recent experiment

If there is no previous experience equal to the current one in the history, the evaluation is simply: $Y = Y_{curr}$

Evaluate a Case

The evaluation of the current experience Υ_{curr} is calculated as the difference between the feedback over previous round proposal and the current round proposal, as presented in equation (9).

$$\Upsilon_{curr} = \Delta\rho_{feedback} + \Delta\alpha_{feedback} \quad (9)$$

where:

- $\Delta\rho_{feedback}$ is the price feedback variation;
- $\Delta\alpha_{feedback}$ is the availability feedback variation.

If the feedback variation is greater than 0, the evaluation is incremented by 0.5 for each attribute. This means that in the worst scenario, where feedback remains unchanged, evaluation is 0. If only one of the feedback values changed, evaluation is set to 0.5 and if both changed, best scenario, evaluation is set to 1.

3 EXPERIMENTS AND RESULTS

To validate our proposal, we have used data provided by a TAP Air Portugal expert in disruption management, regarding disruptions and solutions found for real problems. Each test reflected a disruption and assorted solution possibilities.

The data provided to test the electronic market is composed by 12 disruptions where each disruption contained a considerable amount of fields of which stand out the ID, delay, cost, disrupted resource, as well as the estimated departure time and number of passengers. The number of crew members of each category (captain, first office, senior cabin crew and flight attendant) was also included.

The following metrics will be used to measure the benefit of the solution's found with the electronic market:

- Buyer utility;
- Seller utility;
- Delay reduction;
- Price reduction.

To evaluate price and delay reductions and both agents utility variations, three different experiments were executed. The first experiment considered equal weights for the attributes price and availability. The second experiment valued the availability with a weight of 80% and the price with a weight of 20% in the utility calculation. The third and last experiment showed an inversion regarding the values of the second one, i.e. availability with a weight of 20% and the price with a weight of 80% in the utility calculation. In all experiments, disruption number 12 has no

results to present because does not exists in seller's data set any resource similar to the one required, so seller gives up the negotiation. Due to paper space limitations, we will only describe the first experiment in section 3.1. However, in section 4, the results presented consider the three experiments.

3.1 Experiment 1 (50/50 Experiment)

In this experiment, it was expected to find a similarity between cost and delay reductions, given that both are equally valued in the buyer's utility calculation. The chart in figure 3 presents the results obtained.



Figure 3: Experiment 1 Results.

The points below zero are solutions that had a price greater than the disruption cost, or in the case of agents' utilities, unfeasible solutions, like disruption number 6 that has a negative value for buyer's utility and for this reason it is not a viable solution. The fact that there is no delay reduction or the seller's utility being zero supports that statement, so this is a case that would never result in a leasing contract.

One interesting fact that can be extracted is that seller's utility and delay reduction lines are overlapped, which means that have the same values for every disruption. This is explained by both agents' utility formulas, detailed in section 2.2.

The results obtained for this experiment show that with equal weights in the utility calculation, all but one disruption had its delay minimized by at least 68%. Relatively to the cost reduction the results are not as good as for availability, which is explained by the need to minimize a flight's delay. If the delay is largely reduced, as showed, then the buyer is willing to pay more than the disruption cost. Regarding this experiment results, the average values for each of the metrics used are presented in table 1.

Although the average delay reduction is good (74.91%), there is a cost increase (6.55%) instead of a cost reduction. Concerning to utilities, the aver-

Table 1: Experiment 1 - Average values.

Seller Utility	Buyer Utility	Delay Reduction (%)	Cost Reduction(%)
0.75	0.35	74.91	-6.55

age seller utility reveals that the electronic market is highly useful, at least for this experiment. Regarding the average buyer utility, it shows some improvement but in some cases at a great cost, which explains the considerable difference between seller and buyer utilities.

3.2 Results

The electronic market does not consider any costs unrelated to the disrupted resources. However, in order to choose the most cost-effective solution, passenger related costs must be considered after the market returns its solutions. For instance, the number of passengers that will miss a flight connection due to the delay carries an extra cost to the injured company (passenger cost) and will affect the passenger satisfaction, which also carries an extra cost to the company (passenger goodwill cost). These costs will be added to the aircraft and crew costs, being distributed as follows:

- Direct Costs: Aircraft cost plus crew cost plus passenger cost;
- Integrated Solution Costs: Passenger goodwill cost times passenger goodwill weight plus direct costs.

All these costs are considered by the human specialist (at the AOCC) when it must choose a solution to a disruption in its daily operation. This section intends to compare the solutions found by the electronic market to the ones chosen by a human specialist, by presenting the electronic market solutions to the human for him to analyze and validate. The passenger goodwill weight is 5, by default, according to the specialist.

The first step in the comparison between the solutions found by the electronic market and the ones chosen by a human specialist is to see how the three solutions (one of each experiment) obtained in the electronic market impacts in the flight delay and in the number of passengers missing the flight connections.

The second step is to see the disruption and each solution cost and its influence on the passenger and passenger goodwill costs.

The third step is to see the costs without considering the electronic market solutions: original direct costs and original integrated solution cost, i.e. the original integrated solution cost is the sum of aircraft, crew and passenger costs to which is added the re-

sult of the multiplication between passengers goodwill cost and its weight, as shown in equation (10).

$$IC_{orig} = c_a + c_{cr} + c_{pax} + (c_{paxgw} \times w_{gw}) \quad (10)$$

where:

- c_a is the aircraft cost;
- c_{cr} is the crew cost;
- c_{pax} is the passenger cost;
- c_{paxgw} is the passenger good will cost;
- w_{gw} is the weight of good will.

The new integrated costs represent the integrated costs of the electronic market solutions while the original integrated costs represent the company solution integrated costs. The final step is to see if there are savings provided by each one of the electronic market solutions because the specialist always chooses the solution with a higher value of integrated savings. After being introduced the methodology used to calculate the Integrated Savings, the results over all disruptions are presented in figure 4.

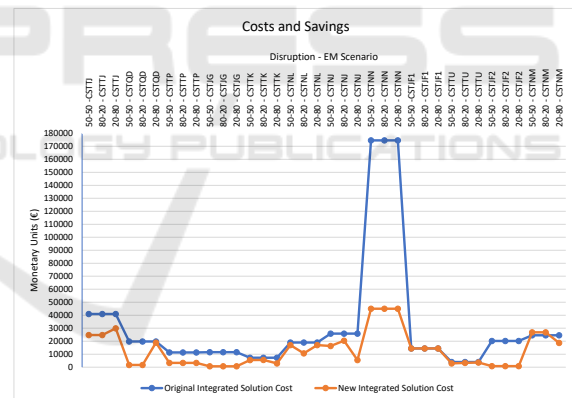


Figure 4: Costs and Savings.

As shown, the solutions obtained through the electronic market are more cost-effective than the company's solutions, except in the disruption which is identified by **CSTJF** that has no similar resources in the electronic market. When comparing the chosen solution from the electronic market with the disruptive solution, the electronic market solutions present an average delay reduction of 66.85% and an average cost reduction of 63.51%. Disregarding the **CSTJF** disruption, there is at least one solution obtained (considering the three experiments made) in the electronic market that is more cost-effective than the disruptive solution for each disruption, having a total of seven disruptions minimized in each experiment.

4 CONCLUSIONS AND FUTURE WORK

In this paper, it is proposed an electronic market modeled as a multi-agent system to expand a company's solution space regarding disruptions management. This electronic market provides alternative solutions to companies affected by disruptions, using resources from other companies (and, as such, contributing to increase collaboration between airlines), which is achieved through automated negotiation, where agents negotiate the resource's availability and price for a disrupted flight. Human validation (at the AOCC) is also included to compare the solutions obtained through the EM with the ones obtained with the company's own resources. The Seller agent in the EM uses case-based reasoning to reuse or adapt previous experiences, to the current negotiation, which is also a contribution of our work.

Three different scenarios were tested to validate the concept, as described in section 3. As there were no available resources for only one disruption in the electronic market, the success rate is 91.7% considering the cost reduction parameter and 67.7% considering both cost and delay minimized.

Possible future directions to improve this work, could include firstly, different approaches in the whole process of identifying previous similar experiences (by the seller), like machine learning and q-learning in order to understand how the agent learning process influences the negotiation, either in terms of proposals' price and availability or in terms of utility for each agent. The methodology used (CBR) can also be improved by creating better evaluation scenarios and benefiting the accepted proposal (or the tree of the proposals that lead to the accepted one).

Secondly, the usage of heuristics to combine resources instead of doing all possible combinations, would be an interesting feature to include. The usage of clustering algorithms to classify resources (where the parameters would be availability and/or price) in order to have a better and more efficient resource combination is also something to explore.

Finally, it would be worthy to use trust models to evaluate the electronic market outcome when considering the relations established between agents and whether that trust measure would influence the agents' behaviour.

REFERENCES

Cardoso, H. L. and Oliveira, E. (2000). A platform for electronic commerce with adaptive agents. In *International*

tional Workshop on Agent-Mediated Electronic Commerce, pages 96–107. Springer.

Castro, A. J., Rocha, A. P., and Oliveira, E. (2014). *A New Approach for Disruption Management in Airline Operations Control*, volume 562. Springer-Verlag Berlin Heidelberg.

Chen, X., Chen, X., and Zhang, X. (2010). Crew scheduling models in airline disruption management. In *2010 IEEE 17th International Conference on Industrial Engineering and Engineering Management (IE&EM)*, pages 1032–1037. Conference Publications.

Kohl, N., Larsen, A., Larsen, J., Ross, A., and Tiourline, S. (2004). Airline disruption management - perspectives, experiences and outlook. Carmen Research and Technology Report CRTR-0407.

Malucelli, A., Castro, A. J. M., and Oliveira, E. (2006). Crew and aircraft recovery through a multi-agent electronic market. In Krishnamurthy, S. and Isaías, P., editors, *Proceeding of IADIS International Conference on e-Commerce 2006*, pages 51–58, Barcelona Spain. IADIS Press.

Matos, N., Sierra, C., and Jennings, N. R. (1998). Determining successful negotiation strategies: An evolutionary approach. In *Multi Agent Systems, 1998. Proceedings. International Conference on*, pages 182–189. IEEE.

Oliveira, E., Fonseca, J. M., and Steiger-Garçon, A. (1999). Multi-criteria negotiation in multi-agent systems. *CEEMAS'99*, page 190.

Oliveira, E. and Rocha, A. (2001). Agents advanced features for negotiation in electronic commerce and virtual organisations formation process. *Agent Mediated Electronic Commerce*, pages 78–97.

Riesbeck, C. K. and Schank, R. C. (2013). *Inside case-based reasoning*. Psychology Press.

Rosenschein, J. S. and Zlotkin, G. (1994). *Rules of encounter: designing conventions for automated negotiation among computers*. MIT press.

Sutton, R. S. and Barto, A. G. (1998). *Reinforcement learning: An introduction*, volume 1. MIT press Cambridge.

Vulkan, N. and Jennings, N. R. (2000). Efficient mechanisms for the supply of services in multi-agent environments. *Decision Support Systems*, 28(1):5–19.

Wooldridge, M. (2009). *An introduction to multiagent systems*, pages 15–23,105–111,129–148. John Wiley & Sons.