

A Negotiation Based Approach to Airline Operations Recovery

António J. M. Castro, António Pereira, Leonardo Fraga, Ana Paula Rocha,
and Eugénio Oliveira

LIACC/NIADR, FEUP, DEI, University of Porto,
liacc@fe.up.pt,
<http://www.liacc.up.pt/>

Abstract. The Airline Operations Control Centre (AOCC) organization is responsible for monitoring and solving operational problems in day-to-day airline operations. It includes human expert teams specialized in solving problems related with aircrafts, crew members, and passengers, in a process called disruption management or operations recovery. We present a new and innovative negotiation-based approach to solve these problems, replacing traditional AOCC expert teams with intelligent agents in a cooperative multi-agent system (MAS). Human interaction focuses on supervision and critical decision actions, such as final approval of proposed solutions. The main research goal is to find the best solution for each problem in an integrated, dynamic and distributed way, by developing agents with their own objectives that work together to minimize the disruptions' effects in the operational plan. Our prototypes, implementing the described approach, led to experiments using real airline data, with problems and solutions validated by experts. Results are presented and discussed.

Keywords: disruption management, operations recovery, airline operations, multi-agent systems, intelligent agents, quality costs

1 Introduction

The Airline Operations Control Centre (AOCC) is one of the most important entities of an airline company. This organization is responsible to monitor the flight's schedule execution during day-to-day operations, checking if it follows what was planned. Unfortunately, problems arise during this phase. These problems are related with aircrafts (e.g., a malfunction), crew members (e.g., a missed sign-on of an employee) or passengers (e.g., a person who is going to miss a connection flight due to a delay). When any of these situations occur, it is necessary to find solutions for them and for that end, the AOCC is composed by teams of experts specialized in the mentioned problem under the supervision of an operation control manager. Each of these teams have a specific goal (e.g., ensuring that every flight has the appropriate crew) contributing to the overall goal of having the airline operation running with few problems as possible. This process is known as disruption management or operations recovery.

In this paper we propose a negotiation based approach to the Operations Recovery process. We believe a Multi-Agent System (MAS) can replace the AOCC in solving the previously mentioned airline operation problem. Our approach uses specialized agents, each one implementing different problem resolution techniques, to find the best solution to a specific problem related with crew, aircrafts or passengers using automated negotiation. All the agents have their own goals, resembling the human expert teams present in any AOCC and also working for a common objective, that is, to find the best integrated solution. By integrated we mean one that simultaneously considers the three dimensions of the problem, aircraft, crew and passenger, while solving it. With our system, we expect to obtain better results as the ones provided by a traditional AOCC.

The approach we propose is a generalization of the Q-Negotiation protocol [10], [11] applied to the Airline Operations Recovery domain and compared against the traditional way the AOCC solves disruptions.

The rest of the paper is structured in the following way: Section 2 presents some work of other authors regarding operations recovery. Section 3 introduces the problem. Section 4 presents our proposal of a MAS for airline operations recovery, including the architecture of the MAS, the evolution of the current system and the negotiation process between the intelligent agents. Section 5 describes the prototypes developed to prove our ideas and beliefs. Section 6 presents the scenarios we setup to evaluate our system, the data used and the results of that evaluation. Section 7 presents the conclusions and future work.

2 Related Work

A common way of analysing the existing work is to classify it according to the dimensions of the problem that is, aircraft, crew and passenger recovery. We add two additional classification types: Partial Integrated recovery (a process that is able to recover any two but not all of the problem dimensions) and Integrated recovery (a process that is able to recover all problem dimensions). Due to space limitations we will only give some examples of existing work. A more detailed and comprehensive state-of-the-art is available in [5].

Aircraft Recovery: In [7] the authors propose a multi-objective genetic algorithm to generate an efficient time-effective multi-fleet aircraft routing algorithm in response to disruption of flights. It uses a combination of a traditional genetic algorithm with a multi-objective optimization method, attempting to optimize objective functions involving flight connections, flight swaps, total flight delay time and ground turn-around times. According to the authors the proposed method has demonstrated the ability to solve the dynamic and complex problem of airline disruption management.

Crew Recovery: In [8] the Crew scheduling problems at the planning level are typically solved in two steps: first, creating working patterns, and then assigning these to individual crew. The first step is solved with a set covering model, and the second with a set-partitioning model. At the operational level, the (re) planning period is considerably smaller than during the strategic planning phase.

The authors integrate both models to solve time critical crew recovery problems arising on the day of operations and describe how pairing construction and pairing assignment are done in a single step, and provide solution techniques based on simple tree search and more sophisticated column generation and shortest-path algorithms.

Partial Integrated Recovery: The authors in [1] introduce a large neighbourhood search heuristic for an airline recovery problem combining fleet assignment, aircraft routing and passenger assignment. Given an initial schedule, a list of disruptions, and a recovery period, the problem consists in constructing aircraft routes and passenger itineraries for the recovery period that allow the resumption of regular operations and minimize operating costs and impacts on passengers. The heuristic alternates between construction, repair and improvement phases, which iteratively destroy and repair parts of the solution. The whole process is iterated by including some randomness in the construction phase so as to diversify the search. This work won the 2009 ROADEF Challenge, a competition organized jointly by the French Operational Research and Decision Analysis Society and the Spanish firm Amadeus S.A.S.. In [2] the author presents two models that considers aircraft and crew recovery and through the objective function focuses on passenger recovery. They include delay costs that capture relevant hotel costs and ticket costs if passengers are recovered by other airlines. To test those models an AOCC simulator was developed, simulating domestic operations of a major US airline. For all scenarios solutions are generated with reductions in passenger delays and disruptions.

Integrated Recovery: In [9] the authors claim to be the first to present computational results on the fully integrated airline recovery problem. Given some time horizon, the recovery process seeks to repair the flight schedule, aircraft rotations, crew schedule, and passenger itineraries in a tractable manner and using a kind of backtracking algorithm. The authors present an optimization-based approach to solve the airline integrated recovery problem and test it using data from an actual U.S. carrier with a dense flight network. It is shown that in several instances an integrated solution is delivered in a reasonable runtime. Nevertheless and in our opinion, a sub-problem resolution order is naturally imposed by the algorithm making some sub-problems more important than others. In this case, the aircraft problem is more important than the crew problem and both more important than the passenger problem, that is, the output of one is the input of another. In the approach presented by us each sub-problem or dimension is taken care in parallel and simultaneously, making them equally important.

3 Airline Operations Recovery

To deal with the unexpected disruptions of their operational plan, airline companies have developed a set of operational control mechanisms to monitor the flights, and their resources. As mentioned, the AOCC is responsible for this, it reacts to disruptions by looking for solutions to overcome these problems with less impact possible to flights and their passengers. The majority of the disrup-

tions are difficult to predict (e.g., those caused by aircraft malfunctions or bad weather), and in nowadays are too common. Thus, our aim with this research is to have a system to handle disruptions that occur during the execution of an airline's operational plan, with close-to-real-time awareness of its environment, so that it can make informed, optimal, or near-optimal, and on-time decisions.

A typical Integrated Operational Control Centre is an hierarchical organization that depends on a supervisor (that takes the final decision) with a specific operation time-windows, i.e., that marks the responsibility boundaries of the AOCC. This operation time-window might be different from airline to airline. Besides the supervisor the other roles or support functions more common in an AOCC [6],[5] are the following:

Flight Dispatch: Prepares the flight plans and requests new flight slots to the ATC entities (e.g., FAA in North America and EUROCONTROL in Europe).

Aircraft Control: Manages the resource aircraft. It is the central coordination role in the operational control. In a disruptive situation, tries to minimize the delays by changing aircrafts and rerouting or joining flights, among other actions. Usually, uses some kind of computer system to monitor the operation that, in some cases, might include some decision supports tools. Much more common is the use of rules-of-thumb based on work experience.

Crew Control: Manages the resource crew. Monitors the crew check-in and check-out, updates and changes the crew roster according to the disruptions that might appear during the operation. Like the previous role, it uses some kind of system with or without decision support tools. The experience and the use, by human experts, of rules-of-thumb are still the most common decision tools. To use reserve crew and to exchange crew members from other flights, are among the possible actions used to solve crew problems.

Maintenance Services: Responsible for the unplanned maintenance services and for short-term maintenance scheduling. Changes on aircraft rotations may impact the short-term maintenance (maintenance cannot be done at all stations).

Passenger Services: Decisions taken on the AOCC will have an impact on the passengers. The responsibility of this role is to consider and minimize the impact of the decisions on passengers, trying to minimize the passenger trip time.

3.1 Current Disruption Management Process

AOCC have a process to monitor the events and solve the problems, so that flight delays are minimized with the minimum impact on passenger and, preferably, with the minimum operational cost. The current disruption management process in use at most of the airlines is sequential with these five steps:

1. Operation Monitoring: At this step, flights are monitored to see if anything is not going according to the plan. The same happens in relation with crew members, passenger check-in and boarding, cargo and baggage loading, etc.
2. Take Action: If an expected event happens, e.g., a crew member is delayed or an aircraft malfunction, a quick assessment is performed to see if an action is required. If not, the monitoring process continues. If an action is necessary than we have a problem that needs to be solved.

3. Generate and Evaluate Solutions: Having collected all the information regarding the problem, the AOCC needs to find and evaluate the candidate solutions. Usually, a sequential approach is adopted when generating these solutions. First, the aircraft problem is solved. Then, the crew problem follows and finally, the one regarding the passengers. It is understandable that the AOCC adopts this approach. Without good software tools, it is difficult to take care of the problem, considering the three dimensions (aircraft, crew and passengers) simultaneously. Although there are several costs involved in this process, we found that the AOCC relies heavily on the experience of their controllers as well as on some rules-of-thumb (a kind of hidden knowledge from people's experience) that exist on the AOCC.
4. Take Decision: Having the candidate solutions a decision needs to be taken.
5. Apply Decision: After the decision the final solution needs to be applied in the environment, that is, the operational plan needs to be updated accordingly.

As mentioned in the Generate and Evaluate Solutions step, the process of finding a possible solution to an Airline Operations Control Problem (AOCP) involves looking into the problem from three perspectives (three dimensions): aircraft, crew and passenger. Thus, it is natural that the costs related to each of these must be weighted when choosing a solution. According to our observations these are the main costs associated to each dimension:

Aircraft/Flight Costs: aircraft takeoff and landing costs, airport parking costs, handling costs, average maintenance costs for the type of aircraft, ATC en-route charges and fuel consumption.

Crew Costs: the average or real salary costs of the crew members, additional work hours, per diem days to be paid, hotel costs and extra-crew travel costs.

Passenger Costs: passenger airport meals, hotel costs and compensations.

Finally, there is a less easily quantifiable cost that is also included: the cost of delaying or cancelling a flight from the passenger point of view. Most airlines use some kind of rule-of-thumb when they are evaluating the impact of the decisions on passengers. Others just assign a monetary cost to each minute of delay and evaluate the solutions taking into consideration this value. We use passenger profiles to calculate this important component according to [4].

4 A Negotiation Based Approach

If we think of each problem individually, that is, aircraft problem, crew problem, and passenger problem, the usual method consists of solving each problem one at a time. We propose something different, something more integrated, more concerned with the AOCP as a whole and not only the sum of its parts. We think that the problem may be solved in an integrated way and not sequentially as the traditional approach does.

We believe that to better solve the AOCP the several dimensions may be solved separately but the global view can not be overlooked and to achieve that, we propose a multi-agent negotiation system where entities responsible for

solving each part of the problem will debate to reach an understanding for a common and global goal but never forgetting their individual objectives.

Lets see an example to better understand the inter-dependencies that exist between each local (or partial) solution. Suppose that a specific flight is delay 60 minutes due to an aircraft malfunction. If nothing is done, this delay would case two crewmembers to break some duty hours limit and 20 passengers would miss their connection at the destination. A possible and better solution would be to exchange the aircraft of the flight with another one (this would reduced the delay in 40 minutes), find a crew that is qualified to the new aircraft (possible the same) and find new itineraries for only 2 passengers that are going to miss their connections. So, we could find the best local solution from the aircraft perspective (the one with less delay and costs) but we need to complement it with a crew (another local solution that needs to comply with some restrictions from the aircraft dimension) and with new passenger itineraries (another local solution that also needs to comply with some restrictions). These three local solutions together are compatible and the best solution from the aircraft dimension point of view.

However, we could apply the same approach from the crew and passenger dimension point of view and choose the best of the three complete solutions. Although each agent will strive to reach his goals, he will also be aware that the negotiations with the other agents are indeed to achieve a global objective. Thus, our system allows to take advantage of the best characteristics of both competitive and cooperative environments.

With our work, we set to find new and innovative ways of solving problems in the AOCC, not only fully Integrated but also in a Dynamic and Distributed way, using the Agent and MAS paradigm. By Dynamic and Distributed we mean:

Dynamic: in real-time and considering the changes in the environment (that is, the airline company operational plan).

Distributed: considering three ways of distribution, that is, (a) *Functional* according to the different expertises involved in the domain, that is, resolution of aircrafts, crew members and passengers problems; (b) *Spatial* according to the partition of data/operational plan, e.g., operational plans may come from different geographically located bases and, (c) *Physical* when data and/or expertise is distributed by different computers in different geographical locations, e.g., according to the existing airline operational bases.

Our proposal is to replace some of the roles and the repetitive tasks done in the AOCC, with agents. With this new approach, the aircraft team will be replaced by a sub-organization of agents called Aircraft Manager and the crew team by a sub-organization called Crew Manager. Regarding the passenger services, we propose to use software agents (called Passenger Manager) with the task of finding the best solutions to the problems with passengers (usually a plan of alternative flights to each disrupted passenger) and keep the other tasks to be performed at the airports by human operators (e.g., the customer services). The human supervisor interacts with the software agents through an interface agent. In figure 1 it is possible to see the architecture of the MAS we have de-

veloped, including the above components and the Monitor agent that will take the responsibility for all the monitoring tasks of all the teams involved.

4.1 MAS Architecture and Negotiation Protocol

Figure 1 represents the agents we propose to use in our MAS, and every agent in it has its purpose, but the most important, and the core of our work, reside in the Monitor, Supervisor and Manager agents:

Monitor: this agent is responsible for finding events that may lead to problems, such as the malfunction of an aircraft, a crew member not signing on for duty, etc. He must then decide if the information provided by the event will cause a delay in any of the airlines scheduled flights.

Supervisor: this agent acts as a broker in the managers' negotiation and is also the bridge between the MAS and the human that has the final decision.

Managers: each manager agent is responsible for finding solutions for part of the problem, the part of their responsibility. Each has the personal objective of minimizing the costs related to its resource and, in conjunction with the other managers, reducing, or even removing, the delay from the problematic flight.

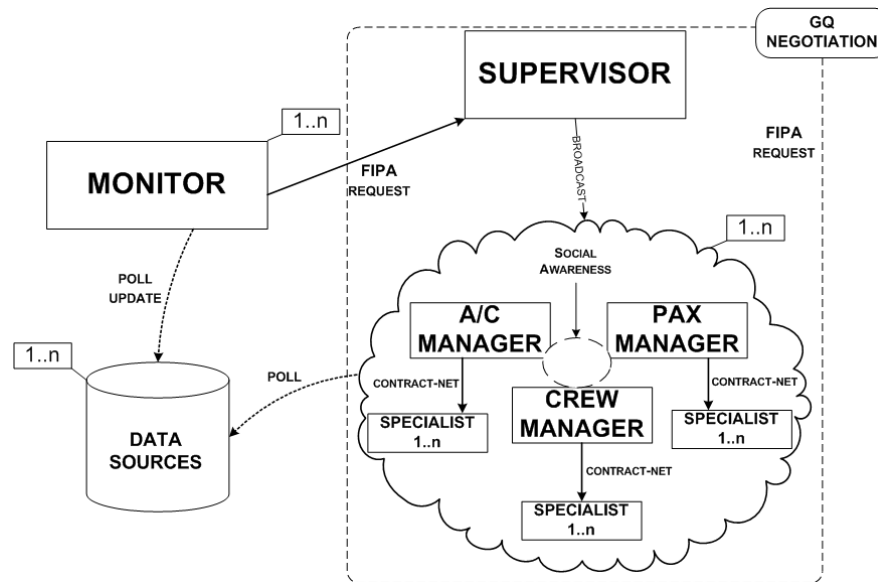


Fig. 1. MAS architecture

The general concept of this architecture is quite simple. The Monitor is constantly looking for events that may cause problems and when such happens, a fully characterized problem is sent to the Supervisor. The latter, will then act as a broker by forwarding the problem to the triad of Managers and controlling

the negotiation process happening between them. After the Managers attain a possible solution, they will inform the Supervisor of their findings. It is here that the human will be consulted by the Supervisor to approve or not the solution found.

As it is possible to see in figure 1 it is between the supervisor and managers and between managers that we use an automated negotiation protocol (GQN). This is a one-to-many, multi-issue and multi-round negotiation protocol suitable for cooperative environments with heterogeneous agents. It is an iterative process where the Supervisor assumes the role of organizer and initiator of the negotiation and issues a call-for-proposal (cfp) for a specific problem (a multi-dimensional one). The Managers assuming the role of respondents and participants on the negotiation will present proposals in response to the cfp. The proposals will be evaluated in each round and a round winner will be determined by the Supervisor. One Manager alone is not able to find a complete solution, since the problem has three dimensions and each Manager agent alone, only deals with one of them. So, in terms of minimal Manager agents requirements, we need to have at least one Manager for each dimension of the problem. In our application domain we will need three of them. These Managers will have to negotiate between them, to be able to complete their partial (or local) solutions, since they need to propose a complete solution (one that includes all the dimensions). This negotiation occurs at each round as well. The negotiation will terminate after a specific number of rounds or time.

Each Manager is a leader of a team of agents that are Specialists in solution finding, through the use of search algorithms and heuristics [3]. Every Manager is aware of its responsibilities and if, the problem received from the Supervisor, was caused by a resource of his responsibility, he will initiate the solution finding process through their team of Specialists. As mentioned before, these agents only find out part of the solution to the given AOCP. Thus, upon getting a solution, the Manager will make a request to the others in order to complete his findings. Managers receiving this request will search for their partial problem's solution. After this iterative negotiation between Managers, a complete solution is sent to the Supervisor. By complete solution, we mean a valid, compatible and doable implementation in all the dimensions the problem involves. The Supervisor and all the Managers have, at their disposal, an evaluation method they use to measure the utility that the solution just found out has to them, considering their design goals and the global objective. Since more than one solution can be found, this procedure helps the Supervisor agent to select the best one.

5 Approaches Developed for Simulation

An initial analysis on the method currently used in real AOCC for disruption management was done, to acquire the information needed for our work. We want to be able to simulate what happens in a real AOCC, to enable us to test, by comparison, the approaches subsequently developed. Thus, we developed a mechanism that imitates the traditional way AOCC workers deal with a

given problem to provide a feasible solution and, naturally, named it Traditional Sequential Approach (TSA).

The purpose of this approach was to provide us with a tool that could imitate, to a point, the traditional approach followed at TAP AOCC when dealing with a AOC. This way we could test it against other approaches, including our proposal, and have a valid mean to compare them and draw significant conclusions. Similar to the traditional way AOCC work, this mechanism of finding a feasible solution is sequential. Taking as example a problem caused by an aircraft malfunction. The Monitor will pass the information related with the problem to the Aircraft Manager (since it is responsible for resources of the aircraft dimension). After it gathers a solution, a partial one, he asks the Crew Manager to complete it according to some restrictions; this creates a kind of loop that ends when these two Managers agree on a solution (not a complete one, since the passenger dimension is still missing), it is sent to the Passenger Manager and it finds the missing part with both the restrictions from the Aircraft and Crew Managers. Only then the solution is complete and is communicated to the Supervisor for human approval. This sequential negotiation is repeated as many times as the human wants. For tests purposes we repeated this negotiation for ten rounds and, in the end, chose the solution whose evaluation (see section 5.2) renders a lower value.

By reviewing this approach, we found some steps in the problem solving method that could be easily improved. These enhancements led us to a new approach, which we entitled Integrated Sequential Approach (ISA), since it evolves the TSA in the sense of solving the three problem dimensions in a more integrated way. The aspect that mainly differentiates this from the TSA is the inclusion of the Passenger Manager in the loop negotiation present in the previous approach but only with the other two Managers. Although this new approach is still somewhat sequential, at least all the three agents responsible for the respectively dimension of the problem are consulted in the process of finding a solution. This way we believe this mechanism can produce better results than the TSA. The way the mechanism works is quite the same apart from the inclusion of the third Manager in the negotiation loop. In the tests conducted with this approach, the aforementioned negotiation was also run ten times and the best solution, according to the same formula, was chosen.

5.1 GQ Negotiation Approach

We want our MAS to have the capability to find the best solution it can, not from a single manager's perspective, but from a global three-dimensional perspective. This is partially accomplished by the ISA, since it looks for a solution keeping in mind not only the Aircraft and Crew resources but also the Passengers. The only issue here, and the one that stops the approach from fully looking for a solution with concern for the three dimensions of the problem, is that the type of problem determines the way the search starts. This may grant us the best overall solution but that is not certain, since the look for it is compromised from the starting point. Thus, the GQN was developed to allow us to test our

theory of involving all managers in the search process, as well as starting it from each agent's point of view and letting them communicate with each other in an effort to return a better and integrated solution for evaluation. As mentioned, this approach seems the most complex but also more complete, since it involves several communications between the managers before a solution is chosen, as can be observed in figure 2.

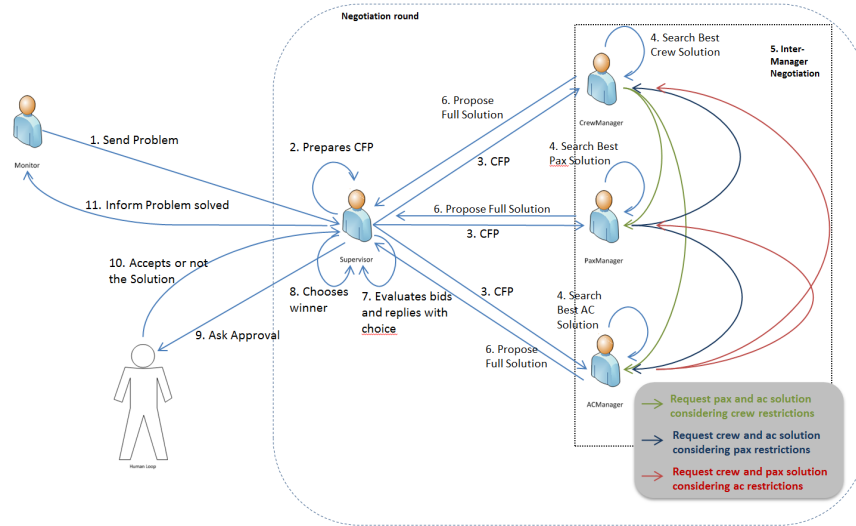


Fig. 2. GQ Negotiation.

In this approach the agents have identical behavior as in the two previous scenarios, but there are some differences worth to mention. The Monitor, after retrieving the event/problem, sends it to the Supervisor (step 1 in figure 2), instead of sending it to the Managers as happens in the ISA. Here, the Supervisor has a central role, acting as a broker of the negotiation taking place between the Managers. It receives the problem from the Monitor and forwards it to the Managers (steps 2 and 3). Upon receiving it, the Manager agents start the negotiation loop between them, much like in the ISA, but with a very important difference: all the Managers start the search process, asking each other to complement their partial solution (steps 4 and 5). In the ISA, only the Manager responsible for the resource affected started looking for a solution. When each of them has a complete solution, it sends it to the Supervisor (step 6), that chooses the best of the three, according to the same evaluation patterns/formula as the other approaches (step 7). Resembling what happens in the other two prototypes, and for tests purposes, the negotiation cycle — starting when the Supervisor broadcasts the problem and finishing when it decides a winner amongst the three solutions received — is repeated ten times and the best complete solution is selected at

the end of this cycle (step 8) and sent to the human user for approval (step 9). The human user has the possibility to accept or not the solution (step 10). If the solution is accepted, the problem is closed with the information from the supervisor to the monitor (step 11). For the tests we have done, all the solutions presented were accepted. In the current version of our protocol we did not implement, yet, the feedback from the human and the consequent supervisor strategy.

As aforementioned, tests were conducted with these three prototypes to see which had a better behaviour concerning the quality of the solutions retrieved.

5.2 Evaluation Formula

For all the approaches, a mathematical formula was developed to evaluate a solution. This evaluation function was prepared according to the guidelines and expertise of the current AOCC human operators at TAP Portugal.

$$Ev = \frac{\alpha * \left(\frac{da}{Da_{max}}\right) + \beta * \left(\frac{dc}{Dc_{max}}\right) + \gamma * \left(\frac{dtt}{Dtt_{max}}\right) + \frac{\delta}{3} * \left(\frac{ca}{Ca_{max}} + \frac{cc}{Cc_{max}} + \frac{cp}{Cp_{max}}\right)}{\alpha + \beta + \gamma + \delta}$$

Table 1 shows the description of the various variables present in the evaluation formula above and the values used (when it applies).

Table 1. Evaluation formula variables description and values.

Variable	Description	Value
da, dc, dtt	delay of aircraft, crew and trip time, respectively	not constant and solution dependent
ca, cc, cp	costs of aircraft, crew and passenger, respectively	not constant and solution dependent
Da_{max} , Dc_{max} , Dtt_{max}	maximum allowed for the aircraft, crew and trip time delay, respectively	100, 100, 100
Ca_{max} , Cc_{max} , Cp_{max}	maximum allowed for the aircraft, crew and passenger costs, respectively	3400, 2500, 8800
α , β , γ	importance given to the aircraft, crew and trip time delay, respectively	1, 0.33, 0.66
δ	importance given to all the costs	0.33

In our evaluation of a solution, we are giving a much larger importance to the delay in an aircraft than we are giving the delay in the crew or trip time. Also, the importance we are giving to the costs is much lesser than what we are giving the delays. All the maximum allowed variables got their values from the data that was used in the simulations. It is important to point out that the values of these parameters were also validated by the human experts of the AOCC. The

parameters that are in table 1 and do not have a defined value are the ones that each solution contributes with for its evaluation.

6 Simulation and results

The data used in all the simulations consisted of several files: one with data from 20 real events that caused problems for TAP’s AOCC to solve in September 2009; and three other, one for each dimension of the problem (aircraft, crew and passenger), containing groups of partial solutions, only concerning their dimension, for every problem present in the events’ file.

The events’ file contains all the necessary data to characterize a problem within our system. This file was used by the Monitor agent to collect information about occurring problems; it was our way of simulating the monitoring phase of the MAS.

The other three files contain information on candidate, possible and compatible solutions for the aforementioned problems. Although all solutions were created manually, they were validated by the AOCC experts. These files are accessed by the Managers, each accesses the file related to the resource of their responsibility, e.g., the Aircraft Manager reads the aircraft related solutions file but cannot attain the information enclosed in the crew and passenger related solutions’ file. This simulates the looking for solution stage that each managers does when prompted to. These files comprise information about costs, delays and evaluation, besides the required knowledge to characterize the solution.

The Manager uses the information enclosed in each of the solutions at his disposal, for the given problem, and pick the one he thinks is the best if he is starting the negotiation process or pick a random, and possible, one if he is answering a request from another agent, always taking into account the restrictions that are present (if any).

By using the same data in the different simulations and approaches, we are building a way to faithfully compare the results which allows us to draw conclusions with some level of assurance. As mentioned before, we defined the stop criteria for all the approaches as a finite number of negotiation loops. Taking into account the number of candidate solutions made available by each manager for each problem, we decided to run each loop ten times and each time keep the solution with the lower evaluation value. As random numbers were part of every simulation, we made three runs for each approach and then calculated an average of the values compared, in an effort to reduce the influence of randomness in the system.

Next, we will present some results. We are only showing part of our results due to size issues, so keep in mind that our conclusions regard all the simulations done (20, one for each event present in our data).

Figure 3 shows the evaluation value for each solution found in each round for one of the problems used in the tests, for the ISA and the GQN. An immediate conclusion is that the ISA only outputs one solution in each round of negotiations

while GQN approach outputs three. However, in this graph, only the solution that won the round — which had the lower evaluation value — is shown.

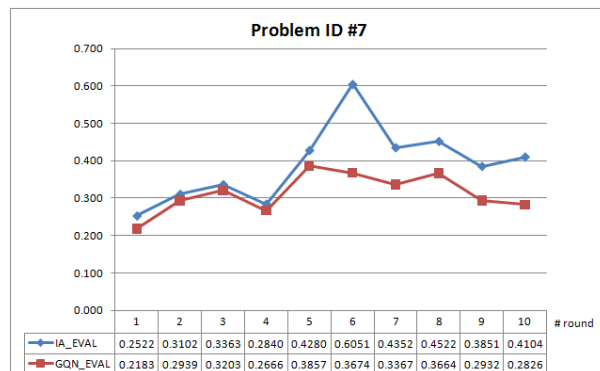


Fig. 3. Evaluation value for problem 7 for the ISA and GQN.

Three main conclusions can be taken from the results we gathered: (1) as the number of rounds increases, so does the evaluation value of the solutions from the ISA, meaning worst solutions; (2) in the majority of the rounds and in all problems, the GQN approach outputs better solutions; and (3) the solutions returned by the GQN approach have considerably lower evaluation values than the ones from the ISA.

The behaviour described on the first item can be due to the fact that the Manager starting the negotiation round, by each passing round, gets a worst solution than the round before. This is not so pronounced in the GQN approach because in each round not only one but every Manager starts a negotiation with the others. Despite the identical process of getting a worst solution than the round before being the same, a wider range of outcomes are possible, thus allowing a better choice to be available in each loop.

In figure 4 we have the choices made by the ISA and GQN approaches side-by-side for all the 20 problems. As observed, the GQN has, in half the situations, the best solution; in 30% of them, the two mechanisms return the same result and the ISA only has the best solution in 20% of the problems.

Figure 5 shows the results in terms of aircraft resource delay, collected from the experiments made with the three approaches developed and also from the solutions applied by TAP’s AOCC for these problems.

Although not perfect, the GQN approach seems to be by far the one that yields the best results.

7 Conclusions and Future Work

As shown in section 5, three approaches were developed, implemented and tested. One of these stood out by giving, for the majority of cases, more and better




Fig. 4. Evaluation values for all problems for ISA and GQN.




Fig. 5. Aircraft delays of all approaches and TAP's solution.

results, the GQ-Negotiation. This proves our beliefs when we developed our method since this approach encompasses all the three dimensions of the problem that needed solution: Aircraft, Crew and Passenger. It makes the Managers look for the solutions more actively despite it being of their responsibility or not, thus resulting in a more range of solutions found than the ISA.

Regarding the next developments in the project, we intend to add feedback (similar to the one used in [10] and [11]) to the negotiation process to make it more complete and more suitable for its purpose. We also intend to add, already developed by us, specialist agents that, regarding the problem, look for solutions in the environment using different algorithms and heuristics. This will take our system to a more complete level and will enable it to present solutions in real-time, as opposed to the current prototype used for the tests presented.

Acknowledgments. The first author is supported by FCT (Fundação para a Ciência e Tecnologia) under research grant SFRH/BD/44109/2008. The second and third authors are supported by the ResPlan project FEUP/SISCOG/2009. The authors are grateful to TAP Portugal for allowing the use of real data from the airline company. We are also grateful to SISCOG for support during the ResPlan project.

References

1. Serge Bisailon, Jean-François Cordeau, Gilbert Laporte, and Federico Pasin. A large neighbourhood search heuristic for the aircraft and passenger recovery problem. *4OR: A Quarterly Journal of Operations Research*, pages 1–19, 2010. 10.1007/s10288-010-0145-5.
2. Stephane Bratu and Cynthia Barnhart. Flight operations recovery: New approaches considering passenger recovery. *Journal of Scheduling*, 9(3):279–298, June 2006.
3. A.J.M. Castro and E. Oliveira. Using specialized agents in a distributed MAS to solve airline operations problems: a case study. In *IEEE/WIC/ACM International Conference on Intelligent Agent Technology, 2007. IAT'07*, pages 473–476, 2007.
4. António J. M. Castro and Eugénio Oliveira. Quantifying quality operational costs in a multi-agent system for airline operations recovery. *International review on computers and software*, 2009.
5. Antonio J. M. Castro and Eugenio Oliveira. A new concept for disruption management in airline operations control. *Proceedings of the Institution of Mechanical Engineers, Part G: Journal of Aerospace Engineering*, 3(3):269–290, March 2011.
6. N. Kohl, A. Larsen, J. Larsen, A. Ross, and S. Tiourine. Airline disruption management—Perspectives, experiences and outlook. *Journal of Air Transport Management*, 13(3):149–162, 2007.
7. Tung-Kuan Liu, Chi-Ruey Jeng, and Yu-Hern Chang. Disruption management of an inequality-based multi-fleet airline schedule by a multi-objective genetic algorithm. *Transportation Planning and Technology*, 31(6):613–639, 2008.
8. Claude P. Medard and Nidhi Sawhney. Airline crew scheduling from planning to operations. *European Journal of Operational Research*, 183(3):1013 – 1027, December 2007.
9. Jon D. Petersen, Gustaf Solveling, Ellis J. Johnson, Jonh-Paul Clarke, and Sergey Shebalov. An optimization approach to airline integrated recovery. Technical report, The Airline Group of the International Federation of Operational Research (AGIFORS), May 2010.
10. Ana Paula Rocha and Eugenio Oliveira. An electronic market architecture for the formation of virtual enterprises. In *PRO-VE '99: Proceedings of the IFIP TC5 WG5.3 / PRODNET Working Conference on Infrastructures for Virtual Enterprises*, pages 421–432, Deventer, The Netherlands, The Netherlands, 1999. Kluwer, B.V.
11. Ana Paula Rocha and Eugenio Oliveira. Electronic institutions as a framework for agents' negotiation and mutual commitment. In *EPIA '01: Proceedings of the 10 th Portuguese Conference on Artificial Intelligence on Progress in Artificial Intelligence, Knowledge Extraction, Multi-agent Systems, Logic Programming and Constraint Solving*, pages 232–245, London, UK, 2001. Springer-Verlag.