Towards an Autonomous and Intelligent Airline Operations Control

António J. M. Castro[1], Ana Paula Rocha[2] and Eugénio Oliveira[2]

Abstract— **Studies have estimated that irregular operations (flights affected by a disruption) can cost between 2% and 3% of the airline annual revenue and that a better recovery process could result in cost reductions of at least 20%. Even for small airlines this can represent millions of Euros. In this paper we propose a multi-agent system (MAS) whose members represent the roles, functionalities and competences existing in a typical Airline Operations Control Centre (AOCC), the airline entity responsible for managing the impact of irregular events on planned operations. This multiagent based system produces intelligent solutions in the sense that its outcomes are the result of an autonomous reaction and adaption to changes in the environment, solving partial problems simultaneously. We tested our MAS using real data from TAP Portuguese airline company and experimentally compared our system with solutions found by the human operators on TAP Portugal AOCC. A comparison was also made with a more traditional sequential approach that is the typical method followed by AOCCs when solving disruptions. Results from those comparisons show that it is possible to reduce costs and have a better integrated solution with the proposed system.**

I. INTRODUCTION

Airline companies face an important and difficult task in controlling their daily operation. Even for a small company like TAP Portugal[1] the size of the scheduling problem is considerable. TAP schedules 55 aircrafts of 5 different types to 76 cities in 34 different countries, covering 1850 weekly flights (approx. 8000 flights per month) and assigns 3132 crewmembers (composed of 821 pilots and 2311 flight attendants) to those flights (source: TAP 2010 Annual Report [1]). Just for comparison, consider an airline like American Airlines[2] with 510 aircrafts of 14 different types flying to 140 cities, with 2700 daily flights and 25000 crewmembers [2]. To operate such a system, airline companies use optimization techniques to be able to build their operational plan, maximizing the revenues and making an efficient use of resources (aircraft and crewmembers).

The optimal operational schedule that results from applying the optimization techniques has a strong probability of being affected, not only by large disruptions like the one that happened in April 2010 due to the Iceland Volcano Ashes but, specially, by smaller daily disruptions caused by bad weather, aircraft malfunctions and crew absenteeism, for example. These disruptions affected the original schedule plan, delaying the flights for example, and cause what is called an *Irregular Operation*. If nothing is done to manage the disruption, the delay can be propagated to other flights, making the problem more difficult to be solved. Studies have estimated that irregular operations can cost between 2% and 3% percent of the airline annual revenue [2] and that a better recovery process could result in cost reductions of at least 20% [3]. Consider the specific case of TAP Portugal, that according to the 2010 Annual Report [1], had an annual revenue of €1.986,3M. The irregular operations could have a cost between €39,7M to €59,5M. A better recovery process could mean a cost reduction between €7,94M to €11,9M. Based on this data we can claim that research on this domain is very important.

In this paper we present a multi-agent system (MAS) that represents the roles and functions that exist in a typical Airline Operations Control Centre (AOCC), the airline entity responsible for monitoring the execution of the operational plan and for managing the irregular operations. This MAS produces intelligent solutions in the sense that its outcomes are the result of an autonomous reaction and adaption to changes in the environment. The MAS has autonomic decision making and problem solving intelligent capabilities. It is able to interact with a human supervisor, showing him the solutions found, getting feedback from him and change its behaviour according to it. We used real data from TAP Portugal and experimentally compared our system with the human solutions found by the human operators on TAP Portugal AOCC as well as with a MAS with a sequential approach that is the typical method followed by AOCCs when solving disruptions. The results show two things: (1) when compared with a manual recovery process there is the possibility of reducing significantly the costs and delays and (2) when compared with a more automatic and sequential process (the common one used in AOCCs) there is the possibility to have better integrated solutions and cost reductions close to the value estimated by Irrang [3].

The rest of this paper is as follows: Section II introduces the Airline Operations Control problem, including the main type of events, the roles and functions of the AOCC, the common disruption management process and the main costs involved. Section III is the main section and presents our proposed Multi-Agent System. In Section IV we present the scenarios and experimentations and in Section V the related work. Finally we conclude in section VI.

II. AIRLINE OPERATION CONTROL PROBLEM

As stated in the previous section, the outcome of the optimization process that airlines use to maximize the revenue is the *Airline Schedule*. The Airline Operations Control Problem (AOCP) starts when the planned is executed, i.e.,

---

[1]TAP Portugal and LIACC, Faculty of Engineering, University of Porto, Portugal `antonio.castro at fe.up.pt`

[2]LIACC, Department of Informatics Engineering, Faculty of Engineering, University of Porto, Portugal `arocha|eco at fe.up.pt`

[1]http://www.flytap.com

[2]http://www.aa.com

the flights are performed. If everything goes as planned the airline just needs to monitor the execution of the plan. Frequently, several unexpected events appear during this phase that can disrupt the plan. To monitor those events and solve the problems that may arise, it is necessary to define and follow a disruption management process. Airline companies have an entity called Airline Operations Control Centre (AOCC) that is responsible for the disruption management process. There are several types of AOCC organizations according to [4]. The adopted organization depends on multiple factors like airline size, airline network type and geographic distribution of the operation, as well as, tradition and/or company culture. Despite the different organization types, it is possible to identify the following characteristics and roles, that are common to all AOCCs [5],[4]:

Operation time-window: Defines the operation responsibility boundaries. Usually, ranges from 24 to 72 hours before to 12 to 24 hours after the day of operation.

Supervisor: responsible to take the final decisions.

Flight Dispatchers: Prepares the flight plans and requests new flight slots to the Air Traffic Control (ATC) entities.

Aircraft Manager: Manages the resource aircraft. It has the central coordination role in the operational control. In a disruptive situation, tries to minimize 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, may include some decision support tools. A much more common approach is the use of rules-of-thumb based on work experience.

Crew Manager: Manages the resource crew. Monitors the crew check-in and check-out, updates and changes the crew roster according to the disruptions that may 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 of rules-of-thumb are still the most common decision tools. To use reserve crew and exchange crewmembers 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 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. Part of this role is performed on the airports.

During the airline operation, several events may occur that will cause an irregular operation. From our observations in a real AOCC and from [6], we found the following typical events:

- Events that may cause a *Flight Arrival Delay*: Enroute air traffic delay, enroute weather, enroute aircraft malfunction, flight diversion and a flight departure delay.
- Events that may cause a *Flight Departure Delay*: Crew

delay, crew absenteeism, loading delay, passenger delay, air traffic control delay, aircraft malfunction, weather conditions and a previous flight arrival delay.

A delay either at departure or arrival, can cause another arrival or departure delay, respectively, and generate new problems on subsequent flights. This propagation characteristic makes the problem more difficult to be solved optimally in real time in a dynamic environment, like the one we have on the AOCC. After an event that causes a problem, a quick action should occur in order to minimize the propagation effect. AOCCs have a typical process to monitor the events and solve the problems. This process has five steps:

1) Operation Monitoring: The flights are monitored to check if anything is not going according to the plan. The same happens in relation with crewmembers, passenger check-in and boarding, cargo and baggage loading, etc.

2) Take Action: If an event arises, such as, a crewmember being late or an aircraft malfunction, a quick assessment is performed to check if an action is required. If not, the monitoring continues (step 1). If an action is necessary then there is a problem that needs to be solved.

3) Generate and Evaluate Solutions: Having all the information regarding the problem the AOCC needs to find and evaluate the candidate solutions. Usually, a sequential approach is adopted when generating solutions. First, the aircraft problem is solved. Then, using this partial solution, the crew problem is solved and, finally, using the previous solutions the passenger problem is also solved. It is understandable that the AOCC adopts this approach. Without good computer tools, it is difficult to handle 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 and in some rules-of-thumb that exist on the AOCC.

4) Take Decision: Having the candidate solutions a decision needs to be taken.

5) Apply Decision: After the decision was taken, the final solution needs to be applied in the environment, that is, the operational plan needs to be updated accordingly. An example of a solution to be applied is: exchange aircrafts between two flights, call a reserve crewmember to replace an absent one and send a *disrupted* passenger in another flight.

In the step *Generate and Evaluate Solutions* we should consider the main costs involved. According to our observations, these costs are the following:

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

Flight Costs: airport costs (approach and taxiing taxes), service costs (cleaning, handling and line maintenance services, etc.), average maintenance costs for the type of aircraft, ATC en-route charges and fuel consumption.

Passenger Costs: passenger airport meals, passenger hotel costs and passenger 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. For example, in the Westminster report [7] an average value of €36/minute of delay is proposed as *Passenger Opportunity Cost*. In a previous work we have proposed a different way of calculating this cost component [8], using passenger profiles.

III. AOCC MULTI-AGENT SYSTEM

In our opinion the AOCC problem can benefit from the agent and multi-agent paradigm. According to [9] and [10], the following are some of the characteristics that make us adopt it to model this problem:

Autonomy: Although the aircraft manager, crew manager and passenger services, cooperate with each other to obtain a final global solution, they are also autonomous in the sense that they have their own goals when solving their part of the problem.

Natural Metaphor: The AOCC modelled as an organization of cooperating agents is a natural metaphor.

Reactivity: The operation monitoring is an example of a function that is reactive, i.e., perceives and reacts to changes in the operation plan.

Resource Distribution: Depending on the size of the airline, the AOCC can benefit from distributed computational resources.

Scalability and Modularity: Extensibility, robustness, maintainability, flexibility and scalability are some of the characteristics presented in MAS important for our scenario.

Figure 1 shows a simplified diagram of the MAS architecture we have developed to model the AOCC. The boxes represent software agents. The bottom ellipse represents the environment, i.e., the data sources available in the AOCC, e.g., operational plan, aircraft and crew roster, weather information and passenger booking system as well as data sources that support the system, e.g., agent history log, proposals and feedback, and human supervisor feedback amongst others.

The *GUI Agent* builds the user interface (see Figure 2) interacting with the environment (e.g., getting information about the operational plan), with the *Human Supervisor* (e.g., presenting the solution and getting feedback) and with the *Supervisor Agent* (e.g., getting the solutions found by the MAS and sending the feedback from the human supervisor).

The *Monitor Agent* is responsible for checking if there are events (e.g., aircraft malfunction, crew absenteeism, weather information, etc.) that could affect the execution of a flight and, in that case, perform the impact analysis to check if that represents a problem that needs to be solved. If that is the case, this agent will inform the supervisor agent and requests a solution to the problem. This agent is autonomous because it is able to consider an event as a problem only when specific conditions or characteristics are present.

The *Supervisor*, *Aircraft Manager*, *Crew Manager* and *Pax Manager* are agents that participate in the automated negoti-
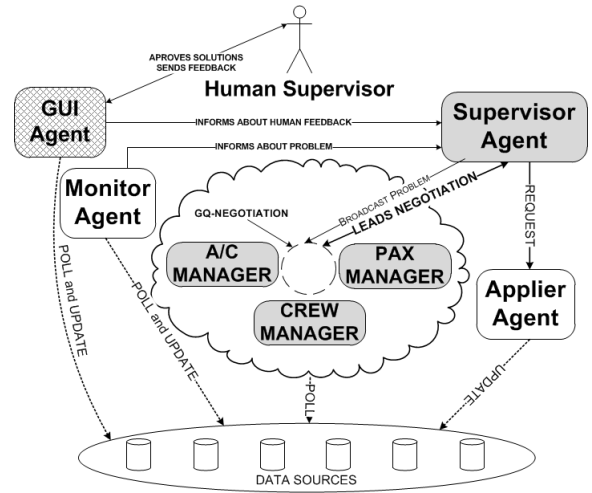

Fig. 1.   Multi-Agent System Architecture

ation using a protocol called GQN (Generic Q-Negotiation). We use negotiation as a decision mechanism, to obtain the best *integrated solution* to the problem presented by the monitor agent. By integrated we mean a solution that considers the three parts (or dimensions) of the problem: aircraft, crew and passenger, and the three manager agents (Aircraft Manager, Crew Manager and Passenger Manager) are responsible to find the best solution to each of these three dimensions.

The *Supervisor Agent* works as a mediator between the manager agents and initiates the negotiation through a CFP (Call-for-Proposal). This agent expects to receive a proposal, from each of the managers, that includes an integrated (candidate) solution to the problem. Each proposal will be evaluated according to Equation 1.

$$ev = \alpha_1 \left(w_1 \left(\frac{ad}{max(ad)} \right) + w_2 \left(\frac{ac}{max(ac)} \right) \right) \quad (1)$$

$$+\alpha_2 \left(w_3 \left(\frac{cd}{max(cd)} \right) + w_4 \left(\frac{cc}{max(cc)} \right) \right) +$$

$$\alpha_3 \left(w_5 \left(\frac{tt}{max(tt)} \right) + w_6 \left(\frac{pc}{max(pc)} \right) \right))$$

with

$$\sum_{i=1}^{3} \alpha_i = 1 \ and \ \sum_{j=1}^{6} w_i = 1$$

The w_j represents the importance of each attribute in the dimension and α_i the importance of each dimension in the problem. For example, α_1 represents the importance of the aircraft part, that includes the attribute *ad* (aircraft delay) and *ac* (aircraft cost). The attributes *cd*, *cc*, *tt* and *pc* represent crew delay, crew cost, passenger trip time and passenger cost, respectively. For each proposal received and evaluated the supervisor agent provides feedback that allows managers to improve their proposals in subsequent rounds. At the end of the negotiation, this agent presents the solution and waits for the approval and feedback from the human

supervisor. If the human supervisor does not approve the solution the supervisor agent uses its feedback to improve the CFP and, if necessary, to change some of the parameters of its evaluation function (e.g., changing the importance of a dimension and/or attribute), starting a new negotiation.

The *Aircraft Manager* agent represents the role that is responsible to find the aircraft solution to the complete problem. It uses a Simulated Annealing algorithm [11] to find solutions (a tuple $\langle aircraft, flight \rangle$) and a Q-Negotiation algorithm [12] to learn how to present new proposals considering the feedback received by the supervisor. This manager, like the other managers, does not have the knowledge to present a complete solution to the problem. For that, he needs to negotiate with the other managers, to be able to complete the proposal. It starts by getting a list of aircraft solutions that satisfies its goals, that is, that minimizes the flight delay and aircraft cost and, then, starting with its best solution, requests a compatible crew and passenger solution (to the crew and pax manager, respectively). If that is not possible, he relaxes by using the next best solution and repeats the process. This (and the other managers) agent, has the following negotiation roles: participant in the negotiation with the supervisor and, in the negotiation with the other managers, either as initiator or as participant (by responding to the requests of the other managers).

The *Crew Manager* agent represents the role that is responsible to find the crew solution (a tuple $\langle crewmember, activity \rangle$) to the complete problem. Like the aircraft agent, it uses a Simulated Annealing algorithm and a Q-Negotiation algorithm and it needs to negotiate with the other managers. Its goal is to minimize the crew costs and crew delay.

The *Pax Manager* agent represents the role that is responsible to find the passenger solution to the complete problem. It uses a Q-Negotiation algorithm to learn how to present proposals but, regarding finding solutions to the passengers (a new itinerary between the origin and destination airport) it uses the Dijkstra Shortest-path algorithm [13]. Its goal is to minimize the passenger trip time and the passenger costs.

Regarding the managers they are autonomous because they only respond to requests related with their area of knowledge and have adaptive characteristics due to the use of the Q-Negotiation algorithm, that encompasses a learning method (Q-Learning) in the generation of new solutions.

The *Applier Agent* is responsible to apply the solution in the operational plan when requested by the supervisor agent. It is important to point out that the solution found in the negotiation, i.e., the values for the attributes: aircraft delay, crew delay, passenger trip time, aircraft, crew and passenger costs, encodes an action plan to be performed on the operational plan. For example, exchange aircraft with tail *CS-TTA* from flight *TP101* with aircraft *CS-TNA* from flight *200*, and use reserve crewmember number *187336* in flight *TP101* and send disrupted passengers to Frankfurt in flight *TP300*.

Finally, the *Human Supervisor* has to approve or not the solutions found by the MAS. Besides the approval it also provides a feedback that quantifies the quality of the solution. If he approves the solution, the feedback is a quantitative classification (a number between 0 and 10) that expresses how pleased he is with the solution and the supervisor closes the problem by asking the *applier* agent to apply the solution in the operational plan. If he does not approve, then he provides quantitative feedback (very high, high, ok, very low and low) for each of the dimensions and for each of the attributes. With this information, the supervisor agent will improve the CFP and its preferences (by changing the weights of the evaluation formula (Equation 1)) and starts a new negotiation.
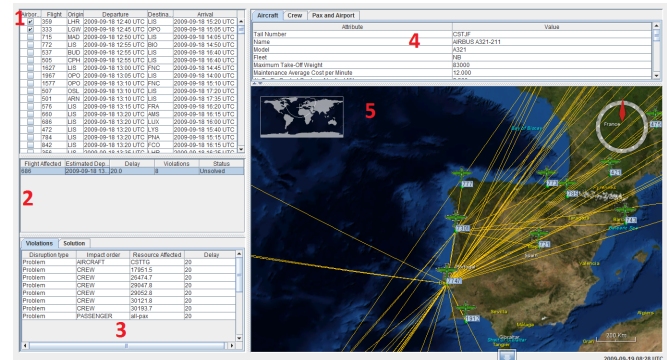

Fig. 2.   Multi-Agent System GUI

Figure 2 presents the GUI of our MAS. It is composed by five areas. *Area 1* shows the daily operational plan, i.e., the flights for a 24 hour period. It is possible to click on a flight and see more information about it on *area 4*, like aircraft and crew data and passenger and airport information (including costs). If a flight is delayed (meaning it has a problem) the flight row turns red and information about the problem will appear in *area 2*. Clicking on the problem on *area 2* will show the violations that affected the flight on *area 3*. In the example provided in Figure 2, the flight 686 has 8 violations and the problem is not solved yet. The violations are related with the aircraft with tail *CSTTG*, with six crewmembers and all passengers on that flight, all affected by a 20 minutes delay.

From this moment the system will try to find a solution using the negotiation approach explained earlier. When a solution is found it will appear on *Solution* tab in *area 3* and the human supervisor will have the opportunity to approve or not the solution and provide feedback.

Finally, in *area 5* it is possible to see the flights following the orthodromic route. We are using the NASA WorldWind API[3] and, as such, we have 3D visualization, zoom and weather information, amongst other useful features of this API.

IV. EXPERIMENTATION AND RESULTS

To evaluate our proposal we have defined the following three scenarios:

---

[3]http://worldwind.arc.nasa.gov/java/

TSA (Traditional Sequential Approach): We have used our MAS to implement the traditional approach used by the human operators at the AOCC when solving disruptions, i.e., after an event is detected, an analysis is performed to see if it affects a flight. In that case, the aircraft part of the problem is solved. Then, using the solution for the aircraft, the crew part is solved. After finding an aircraft and crew solution, the passenger part of the problem is solved. At the end, the supervisor evaluates and authorizes (or not) the solution. It is important to point out that this approach was developed based on interviews performed by us to the human operators of TAP Portugal AOCC. The problem strategy used by the managers were: Simulated Annealing [11] for the *Aircraft* and *Crew Manager* and Dijkstra Shortest-path algorithm [13] for the *Passenger Manager*.

GQNA (Generic Q-Negotiation Approach): We have used our MAS as described in Section III. The main difference regarding the *TSA* scenario is the use of negotiation as a decision mechanism to obtain the final solution.

TAP: We used the historic information from TAP Portugal related with the solutions found and applied by the human operators in real problems (the same ones we used on the other scenarios).

In all scenarios we used the same information from the September 2009 operational plan of TAP Portugal. We have randomly selected 20 events that affected 20 flights, 14 Narrow Body and 6 Wide Body aircrafts, 77 business and 1746 Economic passengers, a total expected delay of 496 minutes (average of 24,8 minutes per flight), total schedule aircraft costs of 33449 monetary units and 30579 schedule crew costs. The solution evaluation formula used (Equation

TABLE I
RESULTS FOR THE SCENARIOS TSA, GQNA AND TAP

Criteria	TSA	GQNA	TAP
A/C Delay & std	6 (5,53)	6 (5,46)	36 (24,28)
Pax Trip Time & std.	28 (20,57)	15 (5,64)	41 (40,02)
A/C Cost & std	1967 (497)	1814 (499)	2081 (521)
Crew Cost & std	1854 (436)	1825 (385)	2158 (504)
Pax Cost & std	2825 (2057)	1535 (564)	
Proposal Utility & std	0,787 (0,09)	0,852 (0,04)	

1) was the same for the *TSA* and *GQNA* scenarios ($\alpha_1 = 0,34; \alpha_{2,3} = 0,33; w_1 = 0,25; w_{2,3,4,5,6} = 0,15$). Due to the probabilistic characteristics of some of the problem resolution strategies, we have run the experimentation 100 times for each approach and, then, used the calculated average values, excluding the highest and lowest value of the sample. The average results for the three scenarios are presented in Table I. Due to lack of space we did not include the results for each problem solved. The first thing to point out is that the results obtained by the third scenario *TAP* are generally worse than the other two scenarios. A possible explanation for this might be that the *TSA* and *GQNA* are automated systems and, as such, capable of generating more candidate solutions and, from those, the best ones are chosen. The human operators base their resolution technique in *rules-of-thumb* and on training received from more experienced

colleagues. As it is possible to see and regarding the delays, the *GQNA* is better with the exception of the aircraft delay where the *TSA* had the same performance. An explanation is that in *TSA* the aircraft dimension is more important since it is the first one to be solved and the resolutions of the other two dimensions are dependent on the resolution of the aircraft one. What is really interesting to see is that the *GQNA* result in the aircraft delay criterion equals the one in *TSA*, since in *GQNA* we are looking for the best integrated solution (the one that takes into consideration the three dimensions of the problem simultaneously). Regarding passenger trip time the *GQNA* is better than *TSA*. Additionally, the standard deviation in all criteria is better on *GQNA* meaning that the distribution is better behaved.

In general, these results mean that *GQNA* could decrease the aircraft delays and passenger trip time delays (on average) by approx. 83% and 63%, respectively, and the aircraft and crew costs by 27% and 15%, when compared with *TAP* scenario. Regarding the aircraft, crew and passenger costs and when compared with *TSA*, the *GQNA* could decrease by approx. 7,8%, 1,6% and 45,7%, respectively. If we compare the utility for the supervisor agent of the winning proposals we can see that, on average, the *GQNA* increases the utility in 8,25% meaning that, as expected, the *GQNA* is able to select the best integrated solutions.

V. RELATED WORK

We classify the literature we have reviewed in three categories: aircraft recovery, crew recovery and integrated recovery. By integrated we mean one that recovers at least two of the dimensions of the problem, e.g, aircraft and crew recovery or aircraft and passenger recovery. We are just going to point out the latest work we found about integrated recovery because it is more related with our proposal. For information about the other related work, please consult [14].

In Eggenberg et al. [15] a modelling framework is presented that allows the consideration of operational constraints within a Column Generation (CG) scheme and introduce the general concept of recovery network, generated for each individual unit of the problem, and show how unit-specific constraints are modelled using resources. The authors show this model applied to the Aircraft Recovery problem and give some insights on how to apply it to the Passenger Recovery problem. According to the authors, it is possible to apply the same model to the Crew Recovery problem although nothing is shown in the paper. However, even in the case of applying the model to the three dimensions, we still do not have an integrated and simultaneous resolution of the problem. The output of a sub-problem is the inputs of another.

In Bisaillon et al. [16] the authors 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. This

work was initiated in the context of the 2009 ROADEF Challenge, in which this work won the first prize. The main difference regarding our proposal is that theirs does not include the crew part. Finally, Petersen et al. [17] presents an optimization-based approach to solve the airline integrated recovery problem, including aircraft, crew and passengers. They tested their approach using data from an actual U.S. carrier with a dense flight network and shown that in several instances an integrated solution is delivered in a reasonable runtime. The used resolution algorithm applies a kind of backtracking. Nevertheless, 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 our approach, due to the use of a negotiation mechanism each sub-problem or dimension is taken care in parallel and simultaneously, making each equally important.

VI. CONCLUSIONS

In this paper we proposed a multi-agent system (MAS) to represent a typical Airline Operations Control Centre (AOCC). This MAS produces intelligent solutions in the sense that its outcomes are the result of an autonomous reaction and adaptation (through the Q-Negotiation algorithm) to changes in the operational plan. The proposed MAS achieves this through autonomic decision-making (after a negotiation process) as well as by integrating partial solutions provided by agents with intelligent problem solving capabilities to reach their ultimate goals. Our proposed solution includes the human-in-the-loop, enabling the supervisor to ask for approval for the solutions found and receiving feedback that allows to change its behaviour according to it.

We have setup three scenarios to test our approach: *TSA* (Traditional Sequential Approach); *GQNA* (GQN Approach); *TAP* (Human operators on the AOCC). We compared the results for the average values. The main conclusion is that the *GQNA* gets the best results in all criteria, except for the aircraft delay where the results were the same of *TSA* and *TAP* approach got the worst results. Concerning the achieved results, the *GQNA* leads to solutions that may decrease the aircraft delays by approximately 83% and aircraft and crew costs by approximately 27% and 15%, respectively, a combined average cost of 21%, when compared with human operator's scenario (*TAP*). When compared with a more automatic and sequential approach (*TSA*), the most common to be used in the AOCC, it has the potential to decrease aircraft, crew and passenger costs by approximately 7,8%, 1,6% and 45,7%, respectively, i.e., a combined average of approximately 18,4%. It is important to point out that the experimentation was **performed with data from one month operation only**. The airline operation domain **has seasonal behaviours**, so, we need to **perform further experimentations with one year data** to be able to get more realistic results. **Without these tests, it is difficult to generalize the results**. If these values were generic enough

and if we apply the smallest combined average cost decrease (18.4%) to the lowest estimated cost of irregular operations of TAP Portugal presented in Section I (€ 39,7M), we could say that TAP could have a cost reduction around € 7,3M. We believe that these results are encouraging and motivate us to improve our work.

REFERENCES

[1] TAP, "Tap portugal relatório anual 2010," July 2011.
[2] X. Chen, X. Chen, and X. Zhang, "Crew scheduling models in airline disruption management," in *2010 IEEE 17Th International Conference on Industrial Engineering and Engineering Management (IE&EM)*. Conference Publications, October 2010, pp. 1032–1037.
[3] M. E. Irrang, *The Handbook of Airline Economics*. Air Transport Association of America, 1996, ch. Airline Irregular Operations, pp. 349–365.
[4] A. J. M. Castro, "Centros de controlo operacional: Organizacao e ferramentas," ISEC - Instituto Superior de Educacao e Ciencias, October 2008, monograph for Post-graduation in Air Transport Operations.
[5] N. Kohl, A. Larsen, J. Larsen, A. Ross, and S. Tiourline, "Airline disruption management: Perspectives, experiences and outlook," Carmen Research, Technical Report CRTR-0407, 2004.
[6] N. Kohl and S. E. Karisch, "Airline crew rostering: Problem types, modeling, and optimization," *Annals OR*, vol. 127, no. 1-4, pp. 223–257, 2004.
[7] A. Cook and G. Tanner, "European airline delay cost reference values," Department of Transport Studies, University of Westminster, London, Technical Report, March 2011.
[8] A. J. M. Castro and E. Oliveira, "Quantifying quality operational costs in a multi-agent system for airline operations recovery," *International Review on Computers and Software (IRECOS)*, vol. 4, no. 4, pp. 504–516, July 2009.
[9] M. Wooldridge, *An Introduction to Multiagent Systems*, 2nd ed. West Sussex, England:: John Wiley & Sons, Ltd., May 2009, ch. When is an Agent-Based Solution Appropriate?, pp. 183–184.
[10] A. Elamy, "Perspectives in agents-based technology," AgentLink News 18, August 2005.
[11] S. Kirkpatrick, C. D. Gelatt, and M. P. Vecchi, "Optimization by Simulated Annealing," *Science*, vol. 220, no. 4598, pp. 671–680, 1983. [Online]. Available: http://www.sciencemag.org/content/220/4598/671.abstract
[12] A. P. Rocha and E. 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*. Deventer, The Netherlands, The Netherlands: Kluwer, B.V., 1999, pp. 421–432.
[13] E. W. Dijkstra, "A note on two problems in connexion with graphs," *Numerische Mathematik*, vol. 1, pp. 269–271, 1959. [Online]. Available: http://jmvidal.cse.sc.edu/library/dijkstra59a.pdf
[14] J. Clausen, A. Larsen, J. Larsen, and N. J. Rezanova, "Disruption management in the airline industry-concepts, models and methods," *Comput. Oper. Res.*, vol. 37, no. 5, pp. 809–821, 2010.
[15] N. Eggenberg, M. Salani, and M. Bierlaire, "Constraint-specific recovery network for solving airline recovery problems," *Comput. Oper. Res.*, vol. 37, no. 6, pp. 1014–1026, June 2010.
[16] S. Bisaillon, J.-F. Cordeau, G. Laporte, and F. Pasin, "A large neighbourhood search heuristic for the aircraft and passenger recovery problem," *4OR: A Quarterly Journal of Operations Research*, pp. 1–19, 2010, 10.1007/s10288-010-0145-5. [Online]. Available: http://dx.doi.org/10.1007/s10288-010-0145-5
[17] J. D. Petersen, G. Solveling, E. J. Johnson, J.-P. Clarke, and S. Shebalov, "An optimization approach to airline integrated recovery," The Airline Group of the International Federation of Operational Research (AGIFORS), Tech. Rep., May 2010.