

# Solving Airline Operations Problems Using Specialized Agents in a Distributed Multi-Agent System

António J.M. Castro and Eugénio Oliveira

LIACC-NIADR&R, Faculty of Engineering, University of Porto  
R. Dr. Roberto Frias, 4200-465 Porto, Portugal  
{ajmc,eco}@fe.up.pt

**Abstract.** An airline schedule very rarely operates as planned. Problems related with aircrafts, crew members and passengers are common and the actions towards the solution of these problems are usually known as operations recovery. The Airline Operations Control Center (AOCC) tries to solve these problems with the minimum cost and satisfying all the required rules. In this paper we present the implementation of a Distributed Multi-Agent System (MAS) representing the existing roles in an AOCC. This MAS deals with several operational bases and for each type of operation problems it has several specialized software agents that implement different algorithms (heuristic, AI, OR, etc.), competing to find the best solution for each problem. We present a real case study taken from an AOCC where a crew recovery problem is solved. Computational results using a real airline schedule are presented, including a comparison with a solution for the same problem found by the human operators in the AOCC. We show that, even in simple problems and when comparing with solutions found by human operators, it is possible to find valid solutions, in less time and with a smaller cost.

**Keywords:** Distributed Multi-Agent Systems, Airline Operations Control, Operations Recovery, Disruption Management.

## 1 Introduction

One of the most important concerns in an airline company is the Operations Control. Through operations control mechanisms the airline company monitors all the flights checking if they follow the schedule that was previously defined by other areas of the company. Unfortunately, some problems arise during this phase [8]. Those problems are related with crew members (for example, a crew member that did not report for duty), aircrafts (for example, a malfunction or a delay due to bad weather) and passengers. When any of these problems appear it is necessary to find solutions for them. The Airline Operations Control Centre (AOCC) is composed by teams of people specialized in solving the above problems under the supervision of an operation control manager. Each team has a specific goal (for example, to guarantee that each flight has the necessary crew members) contributing to the common and general goal of having the airline operation running with few problems as possible. The process of solving these problems is known as Disruption Management [7] or Operations Recovery.

Based on the observations we have done on an AOCC of a real airline company we hypothesize that the objective of solving the operations recovery problems with the less cost as possible, will be much easier to achieve if we include information in the decision process related with various costs as well as if we take advantage of the fact that airlines usually have different operational bases with specific resources. Regarding crew recovery problems, we predict that if we take into account payroll information like *hour salary* and *perdiem value* of each crew rank, and costs related with hotels and *extra-crew travel* between the different operational bases, the solution will be less expensive. The same principle can be applied to aircraft recovery and passenger recovery if we use costs related with that domain. We also hypothesize that the use of different algorithms to solve the same problem (in crew and aircraft recovery) will contribute to the robustness of the system. We predict that using different algorithms (genetic algorithms, heuristic, etc.) in comparison with using always the same algorithm, to solve the same problem, will permit to always find the *best* solution (according to the criteria defined by the company) and to always find *a* solution, especially taking into account the fact that we might benefit from solutions presented by other operational bases.

In this paper we approach this problem so that it can be solved by a Multi-agent System (MAS) that represents the Operational Control Center of the airline company. We use specialized agents, each one implementing Artificial Intelligence algorithms, simple heuristic algorithms and/or Operations Research mathematical models, to find the best solution to a specific problem related with crew, aircrafts or passengers. We expect to obtain a considerable decrease in the costs of the solutions for the problems found when compared with the costs of the solutions found by the current method used in the airline we have observed. We also expect that the heterogeneity of the algorithms, specialized in different types of problems, will allow to find solutions especially for the non-trivial problems, contributing, in this way, to the robustness of the system.

The rest of the paper is structure in the following way. Section 2 presents some work of other authors regarding operations recovery. Section 3 presents our proposal of a MAS for airline operations recovery, including the architecture of the MAS, the algorithm used to choose the best solution and an example of the application of our MAS. Section 4 presents the scenario we setup to evaluate our system as well as the results of the evaluation. Section 5 presents the conclusion of our work.

## 2 Related Work

Traditionally, the Operations Recovery Problem has been solved through Operations Research (OR) techniques. The paper [2] gives an overview of OR applications in the air transport industry. The literature that exists related with this subject is usually divided according to the specific resource to be recovered. The most common division is by aircraft, crew and passengers. However, it is also possible to find papers related with more general approaches as well as related with integrated recovery approaches. We will present here the most recent published papers according to [6]. We divided the papers in four areas: general approaches, aircraft recovery, crew recovery and integrated recovery. For a more detailed explanation of the papers as well as for older papers related with each of these subjects, please consult [6].

*General Approaches:* In [7] the author reports on the experiences obtained during the research and development of project DESCARTES (a large scale project supported by EU) on airline disruption management. The current (almost manual) mode of dealing with recovery is presented. They also present the results of the first prototype of a multiple resource decision support system.

*Aircraft Recovery:* The most recent paper considering the case of aircraft recovery is dated from 2002 [12]. The proposed model addresses each aircraft type as a single problem. They formulate the problem as a Set Partitioning master problem and a route generating procedure. The goal is to minimize the cost of cancellation and retiming, and it is the responsibility of the controllers to define the parameters accordingly. It is included in the paper a testing using SimAir [13] simulating 500 days of operations for three fleets ranging in size from 32 to 96 aircraft servicing 139-407 flights.

*Crew Recovery:* In [1] the flight crew recovery problem for an airline with a hub-and-spoke (a system of air transportation in which local airports offers air transportation to a central airport where long-distance flights are available) network structure is addressed. The paper details and sub-divides the recovery problem into four categories: misplacement problems, rest problems, duty problems, and unassigned problems. Several means are used for recovery, including delaying, swapping, deadheading (extra-crew) and the use of stand by crew. Results are presented for a situation from a US airline with 18 problems.

*Integrated Recovery:* It is uncommon to find literature dedicated specifically to the passenger recovery problem. We believe the main reason for this is the fact that the passenger problems can be minimized if we solve the aircraft and crew problems. However, we would like to point out a recent paper [4] that, although presenting an integrated recovery approach, has a strong emphasis on reducing passenger arrival delays. This paper presents two models that considers aircraft and crew recovery and through the objective function focuses on passenger recovery. To test the models an AOCC simulator was developed, simulating domestic operations of a major US airline. It involves 302 aircrafts divided into 4 fleets, 74 airports and 3 hubs. Furthermore, 83869 passengers on 9925 different passengers' itineraries per day are used. Three different scenarios with different levels of disruption are presented. For all scenarios are generated solutions with reductions in passenger delays and disruptions.

Letovsky's Ph.D. thesis [9] is the first presentation of a truly integrated approach in the literature, although only parts of it are implemented. The thesis presents a linear mixed-integer mathematical problem that maximizes total profit to the airline while capturing availability of the three most important resources: aircraft, crew and passengers. The formulation has three parts corresponding to each of the resources, that is, crew assignment, aircraft routing and passenger flow. In a decomposition scheme these three parts are controlled by a master problem denominated the Schedule Recovery Model.

Finally, we would like to point out a tool called DART (Decision-Aided Rescheduling Tool) [11] that was developed to control the flight operations of IBERIA (the Spanish airline company). DART controls airline operations by gathering real time world-wide information about fleet and crew situation and providing decision support for handling incidents. It covers the daily execution of the ideal flight plan and is

responsible for tracking and solving any irregularities that might arise during its execution. The authors claim that DART has been able to solve some difficult problems, proposing, in some cases, better solutions than those proposed by the re-scheduling experts. The paper does not present any comparative results.

In addition to the above literature the conferences of the AGIFORS<sup>1</sup> organization often feature presentations related with operations recovery. The contributions from these conferences are, at best, available in the form of presentation slides. As such, we did not consider them here.

### 3 Airline Operations Recovery through an Multi-Agent System

#### 3.1 General Description

As stated before we approached this problem by developing a distributed MAS that represents the AOCC. A high-level graphical representation of the MAS architecture is presented in Fig. 1.

The square labeled BASE A shows the part of the MAS that is installed in each operational base of the airline company (e.g., NYC, LHR and LAX). Each operational base has its own resources that are represented in the environment, for example, *Crew Roster* and *Aircraft Roster* are databases of schedules for the crew members and aircrafts, respectively. Other resources represented are the *airport information system*

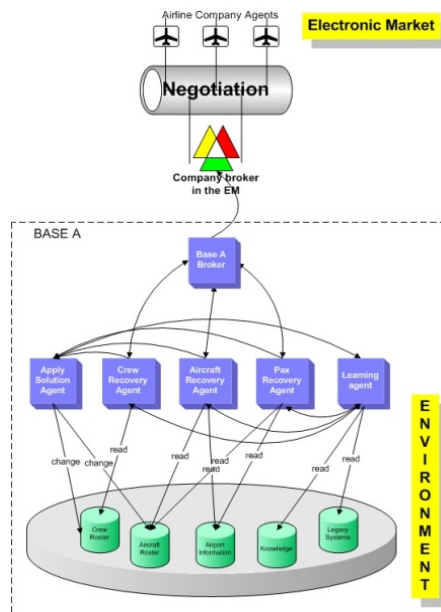


Fig. 1. MAS architecture

<sup>1</sup> AGIFORS, <http://www.agifors.org>, November 2007.

(to be able to get information regarding boarding gates and delays), *legacy systems* (to access information regarding costs, among others) and a *knowledge* database for the learning capabilities of the MAS (this characteristic of the MAS will not be explained in this paper). Each operational base has also software agents that represent roles in the AOCC. The *Crew Recovery Agent*, *Aircraft Recovery Agent* and *Pax Recovery Agent* are dedicated to solve crew, aircraft and passengers problems, respectively, and should be seen as sub-organizations inside the MAS. The *Apply Solution Agent* applies the solution found and authorized in the resources of the operational base. The *Learning Agent* is dedicated to the learning capabilities of the MAS. The MAS has also the possibility to interact with an electronic market of airline resources such as aircrafts and crew members, through the *Company Broker*. According to [7] “research on recovery operation to this date only deals with a single airline. Cooperation between airlines is not supported”. With this approach we try to foster the cooperation between airlines. More information about this electronic market can be found in [10].

The MAS was developed using JADE [3] as development platform and as the runtime environment that provides the basic services for agents to execute. The MAS was developed based on a previous analysis and design by Castro and Oliveira [5].

### 3.2 Sub-organization Architecture

As stated before, the Crew, Aircraft and Pax Recovery agents as presented in Fig. 1 should be seen as sub-organizations. These sub-organizations have their own architecture with their specialized agents. Fig. 2 shows the architecture for *Crew Recovery* in a UML diagram. The architecture for *Aircraft Recovery* and *Pax Recovery* are very similar.

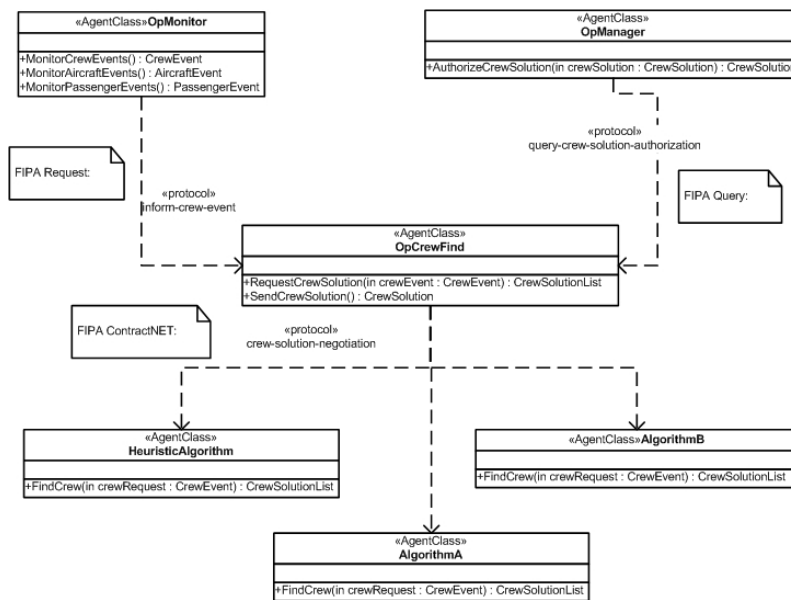


Fig. 2. Crew Recovery sub-organization architecture

CrewEvent	CrewSolutionList	CrewSolution
-dutyDateTime : string	-endDateTime : string	-dutyID : string
-delay : int	-startDateTime : string	-dutyDateTime : string
-dutyID : string	-baseID : string	-endDateTime : string
-endDateTime : string	-crewGrp : int	-readyDateTime : string
-readyDateTime : string	-rank : string	-baseID : string
-credMins : int	-crewNumber : string	-crewGrp : int
-crewGrp : int	-crewName : string	-rank : string
-rank : string	-seniority : int	-crewNumber : string
-baseID : string	-salaryCategory : string	-crewName : string
-crewNumber : string	-hourSalaryValue : decimal	-seniority : int
-crewName : string	-perdiemValue : decimal	-dutyPay : decimal
-openPositions : string	-dutyLimit : int	-perdiemPay : decimal
-eventID : string	-monthDuty : int	-cost : decimal
	-baseFactor : decimal	

Fig. 3. Crew Events

The agent class *OpMonitor* is responsible for monitoring any crew events, for example, crew members that did not report for duty or duties with open positions, that is, without any crew member assigned to a specific role on board (e.g., captain or flight attendant). When an event is detected, the service *MonitorCrewEvents* will initiate the protocol *inform-crew-event* (FIPA Request) informing the *OpCrewFind* agent. The message will include the information necessary to characterize the event. This information is passed as a serializable object of the type *CrewEvent*. Fig. 3 shows the attributes of the *CrewEvent* class.

The *OpCrewFind* agent detects the message and will start a CFP (call for proposal) through the *crew-solution-negotiation* protocol (FIPA contractNET) requesting to the specialized agents *HeuristicAlgorithm*, *AlgorithmA* and *AlgorithmB* of any operational base of the airline company, a list of solutions for the problem. Each agent implements a different algorithm specific for this type of problem. When a solution is found a serializable object of the type *CrewSolutionList* is returned in the message as an answer to the CFP. Fig. 3 shows the attributes of the *CrewSolutionList* class.

The *OpCrewFind* agent collects all the proposals received and chooses the best one according to the algorithm in Table 1:

Table 1. Multi-criteria algorithm

---

```

foreach item in CrewSolution list
  totalDuty = monthDuty+credMins
  if (totalDuty-dutyLimit) > 0
    credDuty = totalDuty-dutyLimit
  else
    credDuty = 0
  end if
  perdiemDays = (endDateTime-dutyDateTime)
  perdiemPay = perdiemDays*perdiemValue
  dutyPay = credDuty*(hourSalaryValue/60)
  cost = (dutyPay+perdiemPay)*baseFactor
end foreach
order all items by cost desc
select first item on the list

```

---

**Table 2.** Computed values

totalDuty	Monthly duty minutes of the proposed crew after assigning the new duty
credDuty	Number of minutes to be paid case the crew exceeds the monthly duty limit
dutyPay	Cost of duty computed according to the hour salary of the crew
perdiemDays	Number of days of work for the specific duty
perdiemPay	Cost of duty computed according to the perdiem value of the crew
baseFactor	The cost associated with the base (extra-crew, hotels, etc.)
Cost	The sum of the cost of the perdiem plus duty multiplied by the base factor.

The algorithm in Table 1 is implemented in the service *SendCrewSolution* and produces a list ordered by the cost (a multi-criteria cost) that each solution represents. Table 2 explains each of the computed values in the algorithm in Table 1.

The first solution of the list in descendant order by cost corresponds to the less expensive one. The *SendCrewSolution* service initiates the protocol *query-crew-solution-authorization* (FIPA Query) querying the *OpManager* agent for authorization. The message includes the serializable object of the type *CrewSolution* as shown in Fig. 3.

### 3.3 Example

Consider the following situation: Airline Company A has two operational bases, one in London (LHR) and another in Paris (ORY), each with 150 crew members. On a specific day a crew member of the LHR base did not report for duty and it was necessary to find another crew member. In our MAS the *OpMonitor* agent of LHR base, would detect and characterize the event according to Table 3.

The agent starts the *inform-crew-event* protocol that includes the information from the *CrewEvent*, informing the *OpCrewFind* agent. This agent starts a CFP through the *crew-solution-negotiation* protocol requesting all the solutions from the specialized agents in both operational bases. The *OpCrewFind* agent receives the solutions as a *CrewSolutionList* from each agent according to Table 4 (this table does not show all the information that is included in the *CrewSolutionList* returned by the agents, like for example, the crew number and name).

The service *SendCrewSolution* of agent *OpCrewFind* computes the values indicated in Table 2 for each item of the *CrewSolutionList* and orders them, according to the algorithm indicated in Table 1. The result is indicated in Table 5.

**Table 3.** Event characterization

Attribute	Value	Comment	Attribute	Value	Comment
dutyDateTime	05-10-2006 10:00		rank	FA	Crew rank
delay	10	Crew delayed 10 mins.	baseID	LHR	
dutyID	NBPNC-1LHR19		crewNumber	97	
endDateTime	05-10-2006 20:15	Duty end date/time	crewName	John	
readyDateTime	06-10-2006 09:15	Includes rest	openPositions	1	
credMins	615	Total work time.	eventID	1230	Internal
crewGrp	2	1=pilots; 2=flight att.			

**Table 4.** CrewSolutionList data

#	BaseID	Rank	Hr Salary	Perdiem	Duty Limit	Month Duty	Base Factor
1	LHR	FA	43	71	7800	7600	1
2	ORY	FA	30	71	7800	8120	1,3
3	LHR	FA	17	31	7800	8500	1
4	LHR	FA	14	31	7800	7950	1
5	ORY	FA	14	31	7800	5000	1,3

**Table 5.** Ordered CrewSolutionList data

#	BaseID	Total Duty	Crd Duty	Duty Pay	Perdiem Days	Perdiem Pay	Cost
5	ORY	5615	0	0	1	31	40
4	LHR	8565	765	178	1	31	209
1	LHR	8215	415	297	1	71	368
3	LHR	9115	1315	372	1	31	403
2	ORY	8735	935	467	1	71	699

**Table 6.** CrewSolution serializable object

Attribute	Value	Attribute	Value
dutyID	NBPNC-1LHR19	crewNumber	147
dutyDateTime	05-10-2006 10:00	crewName	Marie
endDateTime	05-10-2006 20:15	seniority	15
readyDateTime	06-10-2006 09:15	dutyPay	0
baseID	ORY	perdiemPay	31
crewGrp	2	cost	40
Rank	FA		

As it is possible to see, the solution with less cost is solution number 5. In this particular example, it is a crew member from a different operational base that is considered the best solution to substitute the one that did not report for duty. The *SendCrewSolution* service initiates the protocol *query-crew-solution-authorization* querying the *OpManager* for authorization. The message includes the serializable object *CrewSolution* with the complete information, as presented in Table 6.

## 4 Scenario and Experiments

### 4.1 Scenario

To evaluate our MAS we have setup a scenario that includes 3 operational bases (A, B and C). Each base includes their crew members each one with a specific roster. The data used corresponds to the real operation of June 2006 of base A. We have simulated a situation where 15 crew members, with different ranks, did not report for duty in base A. The events did not happen at the same day and each one corresponds to a crew member that did not report for a specific duty in a specific day.

After setting-up the scenario we found the solutions for each crew event using two methods. In the first method we used a real user from the AOCC, with the current tools available, to find the solutions. The user uses software that shows the roster of



each crew member in a Gantt chart for a specific period. The user can scroll down the information, filter according to the crew rank and base, and sort the information by name, month duty, etc. In the second method we have used the sub-organization Crew recovery of our MAS as indicated in Section 3.2.

## 4.2 Results

The data (partial) obtained using method 1 (user with current tools) is presented in Table 7 and the data obtained using method 2 in Table 8. We point out that the data in columns marked with an asterisk were calculated manually, according to the formulas in the algorithm presented in Table 1. The reason for this is that the information system that is available for the users does not include information related with any kind of payroll.

Table 9 presents the results that compare the two methods. From the results obtained we can see that in average, the second method took 25 seconds to find a solution and the first method took 101 seconds. Regarding the costs, the second method has a total cost of 3839.36 and the first method 7039.60. The second method

**Table 7.** Solutions obtained through method 1 (first 13 records only)

#	Base ID	Time (sec)	Rank	Duty Pay (*)	Perdiem Pay (*)	Cost(*)
1	A	90	CAB	0	72	72
2	B	115	CAB	0	72	86,4
3	A	75	CPT	942,9	106	1048,9
4	A	100	CAB	939	144	1083
5	B	120	CAB	0	72	86,4
6	B	100	CPT	777	212	1186,8
7	B	105	OPT	0	148	177,6
8	A	80	CCB	687,65	72	759,65
9	B	110	CCB	0	144	172,80
10	C	110	CPT	0	212	296,8
11	A	110	FA	0	72	72
12	C	120	FA	0	72	100,8
13	B	115	FA	0	72	86,4

**Table 8.** Solutions obtained through method 2 (first 13 records only)

#	Base ID	Time (sec)	Rank	Duty Pay	Perdiem Pay	Cost
1	A	20	CAB	0	72	72
2	B	31	CAB	0	72	86,4
3	B	18	CPT	0	106	127,2
4	C	27	CAB	563,4	62	875,6
5	B	32	CAB	0	72	86,4
6	C	26	CPT	0	212	296,8
7	A	25	OPT	0	144	144
8	B	15	CCB	229,17	72	361,4
9	B	29	CCB	0	144	172,8
10	C	23	CPT	0	212	296,8
11	A	27	FA	0	72	72
12	C	31	FA	0	72	100,8
13	B	32	FA	0	72	86,4

**Table 9.** Comparison of the results

	<i>Method 1</i>		<i>Method 2</i>		<i>Met1/Met2</i>
	Total	%	Total	%	%
Solution base:					
- From base (A)	7	47	3	20,0	-57,1
- From base B	6	40	7	47,0	16,7
- From base C	2	13	5	33,0	150,0
Time (avg sec)	101	100	25	24,8	-75,3
- Base A (avg)	88	21	24	24,0	-72,7
- Base B (avg)	110	27	24	24,0	-78,2
- Base C (avg)	115	28	26	26,0	-77,4
Total Costs:	7039,60	100	3839,36	54,5	-45,5
Costs by Base:					
- Base A	4845,55	92,4	288,00	11,2	-94,0
- Base B	1796,40	34,3	1275,80	49,8	-29,0
- Base C	397,60	7,6	2275,56	88,8	472,3
Costs by Rank					
- CPT	2532,5	36,0	720,8	18,8	-71,5
- FO	720,0	10,2	499,2	13,0	-30,7
- SFA	932,5	13,3	534,2	13,9	-42,7
- FA	2854,6	40,6	2085,1	54,3	-27,0

is, in average 4 times faster than the first method in finding a solution and produces solutions that represent a decrease of 45.5% on the costs.

The *Solution* part of table 9 shows the number of solutions found in each operational base. The *Time* part of the table shows how long it took to find the solutions in each base. The *Total Cost* shows the cost of the solutions by each of the crew ranks involved: CPT (Captains), FO (First Officers), SFA (Senior Flight Attendants) and FA (Flight Attendants). It also shows the costs by each of the operational bases.

## 5 Discussion and Conclusions

From the results obtained we can see that our MAS obtains valid solutions faster and with less costs when compared with the current method used in a real airline company. Regarding our first hypothesis we were expecting a considerable decrease in the costs of the solutions found by our MAS. From the results obtained (see Table 9) we can see that:

$$Cost(SolutionMAS(3839,36)) < Cost(SolutionMet1(7039,60))$$

It represents a decrease of 45.5% on the costs. Our hypothesis was accepted. Of course that we cannot infer that our MAS will always produce solutions that cost 45.5% less. It is not even possible to say that, in average, this decrease is valid. For that we need to evaluate much more situations, in different times of the year (we might have seasoned behaviors) and, then, find an average value. However, taking into consideration that our method includes information that is not available in the current method of the airline (for example, *hour salary*, *perdiem value*, *lodging* and

*extra-crew travel*), and that this information has an immense impact on the total cost, we can state that our method will never produce more expensive solutions.

From the results we can also obtain other interesting conclusions. These conclusions can be expressed by the following formalism:

1.  $Time(SolutionMAS(25s)) < Time(SolutionMet1(101s))$
2.  $Cooperation(SolutionMAS(BaseB(47\%))) >$   
 $Cooperation(SolutionMet1(BaseB(40\%)))$  and  
 $Cooperation(SolutionMAS(BaseC(33\%))) >$   
 $Cooperation(SolutionMet1(BaseC(13\%)))$

Regarding 1) our method was 75.3% faster than method one. The use of a computerized system to find and evaluate the solutions is the reason for our method to be faster than the present, almost manual, method used in the airline. Regarding 2) we can see that the cooperation between different operational bases has increased with our method, because we evaluate all the solutions found (including the ones from different operational bases where the event happened) and choose the one with less cost. In method one, they choose the first one they find. This cooperation is also possible to be inferred from the costs by base. In Table 9 is possible to see that the costs of base C had an increase of 472.32% while base A and base B decreased 94% and 29%, respectively. This means that our method used more resources from other bases than the base where the problem happened (base A).

Regarding our second hypothesis we expected to increase the robustness of our system using heterogeneous algorithms to find solutions to the same problem, at the same time. We were not able to collect enough data to analyze the impact on robustness as the result of using different specialized agents. Preliminary results show that, most of the times, the MAS presents at least one solution even when the human operator cannot find one. Apparently this is the result of using different techniques to tackle the problem. However, the solution might have a cost that, when compared with other ways of solving the problem (for example, cancelling the flight), might be unacceptable. This tells us that our MAS need to have access to more information. For example and in the case of cancelling the flight, it would be important to have access to the cost of compensations due to passengers in these situations.

This paper has presented a distributed multi-agent system as a possible solution to solve airline operations recovery problems, including sub-organizations with specialized agents, dedicated to solve crew, aircraft and passenger recovery problems. We have detailed the architecture of our MAS regarding the sub-organization dedicated to solve crew recovery problems, including agents, services and protocols. We have introduced a multi-criteria algorithm for selecting the solution with less cost from those proposed as part of the negotiation process. A simple example was presented, following, step-by-step, our proposed method. A case study, taken from a real scenario in an airline company where we have tested our method was also presented and we discuss the results obtained. We have shown that our method produces faster and less expensive solutions when compared with the present and almost manual method, used in the airline company.

Further work is required in testing our method for large periods of time and in different times of the year (due to seasoned behaviors). We also need to test our MAS

with all the sub-organizations working at the same time (crew, aircraft and passenger) to see the impact that might exist in the results we have presented in this paper. Finally, we would like to apply and test the integration of the EM as presented in [10].

**Acknowledgements.** António Castro is grateful to TAP Portugal for the support during this research and for allowing the use of real data extracted from the information system used in TAP Portugal.

## References

1. Abdelgahny, A., Ekollu, G., Narisimhan, R., Abdelgahny, K.: A Proactive Crew Recovery Decision Support Tool for Commercial Airlines during Irregular Operations. *Annals of Operations Research* 127, 309–331 (2004)
2. Barnhart, C., Belobaba, P., Odoni, A.: Applications of Operations Research in the Air Transport Industry. *Transportation Science* 37, 368–391 (2003)
3. Bellifemine, F., Caire, G., Trucco, T., Rimassa, G.: JADE Programmer's Guide. JADE 3.3, TILab S.p.A (2004)
4. Bratu, S., Barnhart, C.: Flight Operations Recovery: New Approaches Considering Passenger Recovery. *Journal of Scheduling* 9(3), 279–298 (2006)
5. Castro, A., Oliveira, E.: A Multi-Agent System for Intelligent Monitoring of Airline Operations. In: *Proceedings of the Third European Workshop on Multi-Agent Systems*, Brussels, Belgium, pp. 91–102 (2005)
6. Clausen, J., Larsen, A., Larsen, J.: Disruption Management in the Airline Industry – Concepts, Models and Methods. Technical Report 2005-01, Informatics and Mathematical Modeling, Technical University of Denmark, DTU (2005)
7. Kohl, N., Larsen, A., Larsen, J., Ross, A., Tiourline, S.: Airline Disruption Management – Perspectives, Experiences and Outlook. Technical Report CRTR-0407, Carmen Research (2004)
8. Kohl, N., Karish, S.: Airline Crew Rostering: Problem Types, Modeling and Optimization. *Annals of Operations Research* 127, 223–257 (2004)
9. Lettovsky, L.: Airline Operations Recovery: An Optimization Approach, Ph.D. Thesis, Georgia Institute of Technology, Atlanta, USA (1997)
10. Malucelli, A., Castro, A., Oliveira, E.: Crew and Aircraft Recovery Through a Multi-Agent Electronic Market. In: *Proceeding of IADIS International Conference e-Commerce 2006*, Barcelona, Spain, pp. 51–58. IADIS Press (2006) ISBN: 972-8924-23-2
11. Martins, J., Morgado, E.: Managing Flight Operations. In: *Proceedings of ATTIS 1996*, London, England (1996)
12. Rosenberger, J., Johnson, E., Nemhauser, G.: Rerouting aircraft for airline recovery, Technical Report TLI-LEC 01-04, Georgia Institute of Technology (2001)
13. Rosenberger, J., Schaefer, A., Goldsmans, D., Johnson, E., Kleywegt, A., Nemhauser, G.: A Stochastic Model of Airline Operations. *Transportation Science* 36(4), 357–377 (2002)