

Quantifying Quality Operational Costs in a Multi-Agent System for Airline Operations Recovery

António J. M. Castro¹, Eugénio Oliveira²

Abstract – When recovering from operational problems, the Airline Operations Control Centre (AOCC) usually tries to minimize direct operational costs while satisfying all the required rules. In this paper we present the implementation of a Distributed Multi-Agent System (MAS) representing the existing real-life roles in an AOCC. This MAS includes software agents that cooperate through a distributed problem solving approach, to find the best solution for each problem. We propose a general approach to quantify quality operational costs, so that passengers' satisfaction can also be considered in the final decision. We present a real case study to introduce our approach to quantify the quality operational costs and solve several real unexpected crew problems. We show that our MAS with quality costs is able to reduce flight delays and increase passenger satisfaction without increasing significantly the direct operational costs. A comparison with two other methods is presented. **Copyright © 2007 Praise Worthy Prize S.r.l. - All rights reserved.**

Keywords: *Quality Costs, Disruption Management, Airline Operations, Agents, Multi-Agents*

I. Introduction

Operations control is one of the most important areas for an airline company. Through operations control mechanisms an 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 may arise during this stage [1]. Those problems can be related with crewmembers, aircrafts and passengers. The Airline Operations Control Centre (AOCC) includes teams of experts specialized in solving the above problems under the supervision of an operation control manager. Each team has a specific goal contributing to the common and general goal of having the airline operation running under as few problems as possible. The process of solving these kinds of problems is known as Disruption Management [2] or Operations Recovery.

To select the best solution to a specific problem, it is necessary to include the actual costs in the decision process. One can separate the costs in two categories: Direct Operational Costs (easily quantifiable costs) and Quality Operational Costs (less easily quantifiable costs). Direct operational costs are, for example, crew related costs (salaries, lodgement, extra-crew travel, etc.) and aircraft/flights cost (fuel, approach and route taxes, handling services, line maintenance, etc.). The quality operational costs that AOCC is interested in calculating are, usually, related with passengers satisfaction. Specifically, we want to include in the decision process the estimated cost of delaying or

cancelling a flight from the passenger point of view, that is, in terms of the importance that such a delay will have to the passenger. In this paper we propose a multi-agent system (MAS) to solve the airline operational problems, which include a generic model to quantify quality costs. This MAS is able to recover from operational problems taking into consideration the direct and quality operational costs in the decision process.

The rest of the paper is organized as follows. In section II we present some work of other authors. Section III introduces the MAS used to test our approach, including the decision mechanisms and operational costs. Section IV presents our model to quantify quality operational costs and in section V we show how we have applied this model to a real airline case. In section VI we present the scenarios and experiments performed to evaluate the system. The results and discussion is presented in section VII and in section VII we conclude.

II. Related Work

In this section we present a summary of the work of other researchers regarding operations recovery, divided in three main areas: aircraft, crew recovery and integrated recovery. We also list a brief summary of the application of agents and multi-agent systems in other domains.

II.1. Aircraft Recovery

Liu et al. [3] proposes 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”. As in other approaches, the authors do use the delay time in the objective functions but nothing is included regarding passenger quality costs.

Mei Yang Ph.D. thesis [4] investigates the use of advanced tabu search methodologies to solve the aircraft grounding problem and the reduced station capacity problem. The objective is to minimize the schedule recovery costs associated with flight schedule modifications and deviations from the original route. Mei uses cancellation and delay costs in the objective function. For the delay costs, Mei uses a value of \$20 if the delay is less than 15 minutes and \$20 each minute if the delay is greater or equal to 15 minutes. For flight cancellations it uses a combination of lost revenue, loss of passenger goodwill and other negative effects, specific and predefined for each flight. The main difference regarding our approach is that we allow the definition of several passenger profiles for each flight (Mei and others, do not consider profiles), each one with an associated cost formula, that reflects the delay costs from the passenger point of view.

Rosenberger et al. [5] formulates 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 process using SimAir [6], simulating 500 days of operations for three fleets ranging in size from 32 to 96 aircraft servicing 139-407 flights. Although the authors do try to minimize flight delays, nothing is included regarding the use of quality costs.

II.2. Crew Recovery

In Abdelgahny et al. [7] the flight crew recovery problem for an airline with a hub-and-spoke 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. The proposed model is an assignment model with side constraints. Due to the stepwise approach, the proposed solution is sub-optimal. Results are presented for a situation from a US airline with 18 different problems. This work omits the use of quality costs.

II.3. Integrated Recovery

Bratu et al. [8] 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. According to the authors, it is possible to include, although hard to estimate, estimations of delay costs to passengers and costs of future lost ticket sales. 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. For all scenarios are generated solutions with reductions in passenger delays and disruptions. The difference regarding our work is that we propose a generic model to calculate the delay cost to passengers, based on their specific profile and opinion (obtained through frequent surveys).

Kohl et al. [2] 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. Passenger delay costs are calculated regarding the delay at the destination and not at departure (we include both in our proposal) and takes into consideration the commercial value of the passenger based on the booked fare class and frequent flyer information. The main difference regarding our proposal is that we use the opinion of the passengers when calculating the importance of the delay.

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. Although the author takes into consideration the passenger, it does so regarding finding the best solution for the disrupted passengers. The difference regarding our approach is that we use the opinion of the passengers regarding the delay (expressed through a mathematical formula) to reach the best solution regarding delaying the flight. We do not approach the, also important, issue of finding the best itinerary for disrupted passengers.

Castro and Oliveira [10] present a Multi-Agent System (MAS) to solve airline operations problems, using specialized agents in each of the three usual dimensions of this problem: crew, aircraft and passengers. However, the authors ignore the impact that

a delay in the flight might have in the decision process and only use operational costs to make the best decision.

II.4. Other Application Domains

Agents and multi-agent systems have been applied both to other problems in air transportation domain and in other application domains. A brief and incomplete list of such applications follows. Tumer and Agogino [11] developed a multi-agent algorithm for traffic flow management. Wolfe et al. [12] uses agents to compare routing selection strategies in collaborative traffic flow management. For ATC Tower operations, Jonker et al. [13] have also proposed the use of multi-agent systems. As a last example, a multi-agent system for the integrated dynamic scheduling of steel production has been proposed by Ouelhadj [14].

III. A MAS for Operations Recovery

It is important to point out that we arrived to the architecture of our multi-agent system, after performing an analysis and design using an agent-oriented software methodology [15]. The agent model and service model were the outputs of this process and the base for this architecture. A partial architecture of the MAS we built is presented in figure 1. The boxes represent agents and the narrow black dash lines represent requests/proposals made. The larger black lines represent the interaction between agents regarding negotiation and distributed problem-solving process. The narrow gray lines represent interaction within a hierarchy of agents and the normal black lines represent the interactions after a solution is found. It is important to clarify that Fig. 1 represents only one instance of the MAS. The MAS was developed using JADE [16] as a development platform and as the run-time environment that provides the basic services for agents to execute.

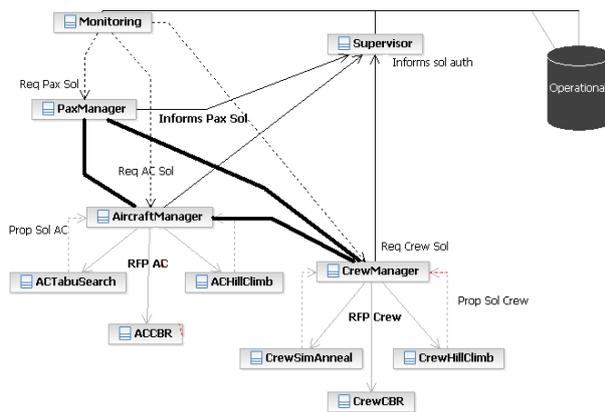


Fig. 1. MAS architecture

Considering the agent and multi-agent system characteristics as specified in [17] and [18], the following ones make us adopt this paradigm to the

AOCC problem:

Autonomy: MAS models problems in terms of autonomous interacting component-agents, which are a more natural way of representing task allocation, team planning, and user preferences, among others. In Fig. 1 the *PaxManager*, *AircraftManager* and *CrewManager* agents (among others) are agents that can choose to respond or not to the requests according to their own objectives.

Agents are a Natural Metaphor: The AOCC is naturally modelled as a society of agents cooperating with each other to solve such a complex problem.

Reactivity: Agents are able to perceive and react to the changes in their environment. The *Monitor* agent in Fig. 1 is an example of such an agent.

Distribution of resources: With a MAS we can distribute the computational resources and capabilities across a network of interconnected agents avoiding problems associated with centralized systems. Airline companies of some dimension have different operational bases. We use a MAS for each operational base, taking advantage of this important characteristic. Due to the social awareness characteristics of some of our agents (for example, *Monitoring* agent in Fig. 1) they are able to distribute their tasks among other agents with similar behaviour.

Modularity and Scalability: A MAS is extensible, scalable, robust, maintainable, flexible and promotes reuse. These characteristics are very important in systems of this dimension and complexity. Our MAS is able to scale in terms of supporting more operational bases as well as in supporting different algorithms to solve specific problems.

Concurrency/Parallelism: Agents are capable of reasoning and performing tasks in parallel. This provides flexibility and speeds up computation. The *CrewSimAnneal*, *CrewCBR* and *CrewHillClimb* agents in Fig. 1 are examples of concurrent agents. Additionally and according to [19] “if control and responsibilities are sufficiently shared among agents, the system can tolerate failures by one or more agents”. Our MAS can be totally or partially replicated in different computers. If one or more agents fail, the global objective is not affected.

Legacy Systems: The AOCC needs information that exists in obsolete but functional systems. We can wrap the legacy components in an agent layer, enabling them to interact with other software components.

In Fig. 1 each one of the agents *Monitoring*, *PaxManager*, *AircraftManager*, *CrewManager* and *Supervisor* has specific associated roles in the AOCC.

The *Monitoring* agent monitors the operational plan looking for events that may represent any of the usual three problem dimensions, that is, aircraft, crew and/or passenger problems. In case there are other instances of

this agent, they recognize and interact with each other, splitting the monitoring task. For example, if each instance corresponds to an operational base, each one will monitor the corresponding operational plan. This is one example of the social-awareness characteristics of our agents. The agent is autonomous in the sense that it will consider an event as a problem only if the event has certain characteristics.

The *PaxManager* agent has the responsibility to find solutions for passenger problems. The *AircraftManager* and *CrewManager* agents have the responsibility for finding solutions for aircraft and crew problems, respectively. These agents are *autonomous* in the sense that they can choose not to respond to the information received from the *Monitor* agent, i.e., if the problem is not related with their field of expertise or if they do not have local resources to solve that problem. These agents have similar *social-awareness* characteristics of the *Monitor* agent. These agents may decide to participate with their expertise in the integrated and distributed problem solving approach of the system.

The *AircraftManager* and *CrewManager* agents manage a team of specialized agents [10]. Each team should have several specialist agents, each one implementing a different problem solving algorithm, making them heterogeneous regarding this characteristic. The *ACTabuSearch* agent, *ACCB* agent and *ACHillClimb* agent implements algorithms dedicated to solve aircraft problems and present the candidate solutions they find to the *AircraftManager* agent. The *CrewSimAnneal* agent, *CrewHillClimb* agent and *CrewCBR* agent implements algorithms dedicated to solve crew problems and present the candidate solutions to the *CrewManager*.

The agent *Supervisor* is the only one that interacts with a human user of the AOCC. The *Supervisor* agent presents the solutions to the human supervisor, ranked according to the criteria in use by the airline (usually total operational cost), including details about the solution to help the human to decide.

All agents are able to act and observe the environment that is represented by the *Operational* database. This database includes information regarding the flight, aircraft and crew schedule as well as airport and company specific information.

III.1. Protocols and Decision Mechanisms

The protocols we use are the following FIPA¹ compliant ones:

Fipa-Request: This protocol allows one agent to request another to perform some action and the receiving agent to perform the action or reply, in some way, that it cannot perform it. Fipa-request is used in

interactions between the *Monitor*, *PaxManager*, *AircraftManager* and *CrewManager* agents.

Fipa-Query: This protocol allows one agent to request to perform some kind of action on another agent. It is used in the interactions that involve *PaxManager*, *AircraftManager*, *CrewManager* and *Supervisor* agent.

Fipa-Contract.net: In this protocol, one agent takes the role of manager which wishes to have some task performed by one or more other agents and further wishes to optimize a function that characterizes the task. We use a simplified version of this protocol in the interactions that entail the *AircraftManager* and its specialized agents, i.e., *ACTabuSearch*, *ACCB* and *ACHillClimb*; and *CrewManager* and its specialized agents, i.e., *CrewSimAnneal*, *CrewHillClimb* and *CrewCBR*.

Our system uses negotiation at two levels. The first level is the *Manager Agents* level, i.e., between *PaxManager*, *CrewManager* and *AircraftManager* agents. At this level the agents cooperate so that an integrated solution can be found. We define an integrated solution as one that considers the impact on the three dimensions of the problem, that is, aircraft, crew and passengers. Each manager agent looks for possible implications of a specific problem in their field of expertise and uses that information to help the other agents to fine-tune the parameters when looking for solutions. With this simple algorithm we are able to have a distributed problem solving approach to the problem. As of the writing of this paper, we do not have this level completely implemented.

The second level is the *Specialist Agents* level or *Team* level, i.e., between each manager agent and the specialist agents of the team. At this level we have used a simplified *fipa-contract.net* [20], [21].

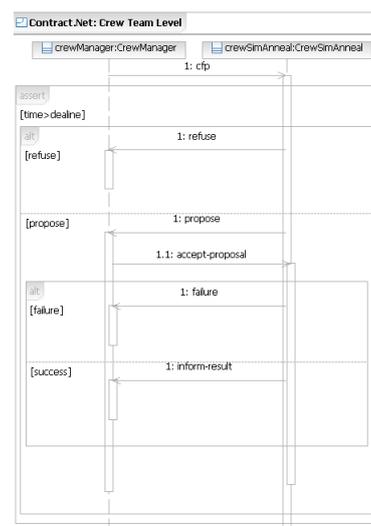


Fig. 2. Simplified contract.net protocol

Fig. 2 shows the simplified contract.net protocol

¹ <http://www.fipa.org>

applied to the *CrewManager* team (for simplicity only the interaction between *CrewManager* and one of the specialist agents is shown). After receiving a request from the *Monitoring* agent and case the *CrewManager* agent decides to reply, a Call for Proposal (cfp) is issued to initiate the negotiation process.

Please note that the content of the FIPA-ACL message is a serialized Java object, that contains the event description, as well as the deadline for receiving an answer (propose or refuse) and the deadline for receiving the candidate solution (i.e., the *CrewSimAnneal* agent needs to send a candidate solution before a specific period of time). The *CrewSimAnneal* agent may choose to answer refuse or propose. In our approach the *CrewSimAnneal* propose performative only means that it will look for a candidate solution according to the conditions of the cfp. The *CrewManager* agent will automatically answer back with an accept-proposal. It is here that we have simplified the contract.net protocol to speed-up the communication between our agents. In our case, the answer we get from specialist agents is a simple yes or no, because we want all available agents (i.e., that are not busy looking for candidate-solutions for other requests) to work in parallel to find candidate solutions. Because of that we do not need to choose between all the answers received. If there is a problem during the execution of the task, the *CrewSimAnneal* agent issues a failure performative stating the reasons for the failure, in the serialized Java object included in the message content. If the agent is able to perform the task with success, it will issue an inform-result performative that includes the serialized object with the candidate solution.

At the team level, the manager agent needs to select the best solution from the candidate solutions that were found by the specialist agents. Once the participant agent has completed the task (for example, agent *CrewHillClimb* in Figure 1), it sends a completion message to the initiator (agent *CrewManager* in Figure 1) in the form of an inform-result performative, with the details of the candidate solution including the *Total Operational Cost*. The manager agent sorts in descendant order all candidate solutions received by total operational cost and selects the first one. As of the writing of this paper, we use the *Total Operational Cost* as the only criteria for the selection. Other criteria, like *AOCC Global Performance*, are being tested but we do not have any results at this moment. Section III.2 details the criteria used at this level.

III.2. Operational Costs

The *Total Operational Cost* (tc) of a specific solution includes *Direct Operational Costs* (dc) and *Quality*

Operational Costs (qc) and is given by Equation 1. In this section we will detail the direct operational costs. The quality costs will be explained on section III.3.

$$tc = dc + \beta qc \quad \beta \in R, \beta \geq 0 \quad (1)$$

Coefficient β is used to define the weight of quality costs. *Direct Operational Costs* (dc) of a specific solution are costs that are easily quantifiable and are related with the operation of the flights, namely, *Crew Costs* (cc), *Flight Costs* (fc) and *Passenger Costs* (pc). It is given by Equation 2.

$$dc = cc + fc + pc \quad (2)$$

The *Crew Cost* (cc) (Equation 3) for a specific flight includes the salary costs of all crew members (*Salary*), additional work hours to be paid (*Hour*), additional per diem days to be paid (*Perdiem*), hotel costs (*Hotel*) and extra-crew travel costs (*Dhc*).

$$cc = \sum_{i=1}^{|F|} \sum_{j=1}^{|C|} (Salary_{\{i,j\}} + Hour_{\{i,j\}} + Perdiem_{\{i,j\}} + Hotel_{\{i,j\}} + Dhc_{\{i,j\}})$$

where
 $i \in F; F = \{\text{all flights in solution}\}$
 $j \in C; C = \{\text{all crewmembers in flight}\}$

(3)

The *Flight Cost* (fc) (Equation 4) for a specific flight includes the airport costs (*Airp*), i.e., charges applied by the airport operator like approaching and taxing; service costs (*Service*), i.e., flight dispatch, line maintenance, cleaning services and other costs; average maintenance costs for the type of aircraft that performs the flight (*Maint*); ATC en-route charges (*Atc*); and fuel consumption (*Fuel*), i.e., fuel to go from the origin to the destination (trip fuel) plus any additional extra fuel required.

$$fc = \sum_{i=1}^{|F|} (Airp_i + Service_i + Maint_i + Atc_i + Fuel_i)$$

where
 $i \in F; F = \{\text{all flights in solution}\}$

(4)

The *Passenger Cost* (pc) of the delayed passengers for a specific flight includes airport meals the airline has to support when a flight is delayed or cancelled (*Meals*), hotels costs (*PHotel*) and any compensation to the passengers according to regulations (*Comp*). The Passenger Cost of the delayed passengers for a specific solution is given by Equation 5.

$$pc = \sum_{i=1}^{|F|} \sum_{d=1}^{|D|} (Meals_{\{d,i\}} + PHotel_{\{d,i\}} + Comp_{\{d,i\}})$$

where
 $i \in F; F = \{\text{all flights in solution}\}$
 $d \in D; D = \{\text{all delayed passengers in flight}\}$

(5)

IV. Quality Operational Costs

The Airline Operations Control Centre (AOCC) has the mission of controlling the execution of the airline schedule and, when a disruption happens (aircraft malfunction, crewmember missing, etc.) find the best solution to the problem. It is generally accepted that, the best solution, is the one that does not delay the flight and has the minimum direct operational cost. Unfortunately, due to several reasons (see [22] for several examples), it is very rare to have candidate solutions that do not delay a flight and/or do not increase the operational cost. From the observations we have done in a real AOCC, most of the times, the team of specialists has to choose between candidate solutions that delay the flight and increase the direct operational costs. Reasonable, they choose the one that minimize these two values.

IV.1. Perception of Quality Costs

Also from our observations, we found that some teams in the AOCC used some kind of rule of thumb or hidden knowledge that, in some cases, make them not choose the candidate solutions that minimize the delays and/or the direct operational costs. For example, suppose that they have disruptions for flight A and B with similar schedule departure time. To solve the problem, they have two candidate solutions: one is to delay flight A in 30 minutes and the other would delay flight B in 15 minutes. The direct operational costs for both candidate solutions are the same. Sometimes they would choose to delay flight A in 15 minutes and flight B in 30 minutes. We can state that flights with several business passengers, VIP's or for business destinations correspond to the profile of flight A in the above example. In our understanding this means that they are using some kind of quality costs when taking the decisions, although not quantified and based on personal experience. In our opinion, this knowledge represents an important part in the decision process and should be included on it.

IV.2. Quantifying Quality Costs

To be able to use this information in a reliable decision process we need to find a way of quantifying it. What we are interested to know is how the delay time and the importance of that delay to the passenger are related in a specific flight. It is reasonable to assume that, for all passengers in a flight, less delay is good and more is bad. However, when not delaying is not an opinion and the AOCC has to choose between different delays to different flights which one should they choose? We argue that the decision should take into consideration the passenger's profile(s) of the specific flight and not only the delay time and/or operational

cost. For quantifying the costs from the passenger point of view, we propose the following generic approach:

- 1) Define the existing passenger profile(s) in the flight.
- 2) Define a delay cost for each passenger in each profile.
- 3) Calculate the quality costs using the previous steps.

Most likely, every airline company will have a different method to define the passenger profile in a specific flight. Most of the airlines will just consider one or two profiles (for example, business and economy). To get the number of passengers that belong to these profiles is very easy. Airline companies can use the flight boarding information to calculate this number. In section V we present a real example of a company that used three profiles.

Most of the airline companies will choose to use a fixed delay cost value to each passenger of each profile. These numbers can reflect the perception of the costs from the point of view of the company or can result from a statistical analysis of the company information. In our opinion and that is one of the main contributions of our approach, we think that this cost should be calculated from the passenger point of view. This implies to use a formula to calculate the costs of each profile, that represents this relation. In section V we show how a real airline company used a passenger survey to obtain formulas to calculate the cost of each passenger profile.

Giving the above we believe that the quality costs should result from the relation between the number of passenger profiles in the flight and the delay cost for each passenger from their point of view, expressed by Equation 6.

$$qc = \alpha \sum_{i=1}^{|F|} \sum_{p=1}^{|PP|} (P_{\{p,i\}} * C_{\{p,i\}})$$

where
 $i \in F; F = \{\text{all flights in solution}\}$
 $p \in PP; PP = \{\text{flight passengers profiles}\}$
 $P = \text{number of passengers of profile } p$
 $C = \text{delay cost of each passenger on profile } p$
 $\alpha = \text{coefficient to convert to monetary costs}$

(6)

In our MAS we are prepared to get this information dynamically and for the specific flight(s) involved in the problem.

V. Airline Company Case Study

This section presents the use of the quality operational costs approach we proposed in section IV to the airline company that we are observing. We start by showing how we get the passenger's profiles, then how we get the formulas that express the cost for each passenger in the profile(s) and, finally, an example of

the quality operational costs for a specific flight.

V.1. Defining the Passenger Profiles

The final goal in this real example is to be able to have passenger profiles to every flight in the company, regarding the delay cost from the point of view of the passengers. To get this information, we have done a survey to several passengers on flights of the airline company. Besides asking in what class they were seated and the reason for flying in that specific flight, we asked them to evaluate from 1 to 10 (1 – not important, 10 very important) the following delay ranges (in minutes): less than 30, between 30 and 60, between 60 and 120, more than 120 and flight cancellation. From the results we found the passenger profiles in Table 1.

TABLE I
PASSENGER PROFILES

Profiles	Main characteristics
<i>Business</i>	Travel in first or business class; VIP's; Frequent Flyer members; Fly to business destinations; More expensive tickets.
<i>Pleasure</i>	Travel in economy class; Less expensive tickets; Fly to vacation destinations.
<i>Illness</i>	Stretcher on board; Medical doctor or nurse travelling with the passenger; Personal oxygen on board or other special needs.

For the profiles in Table I to be useful, we need to be able to get the information that characterizes each profile, from the airline company database. We found that we can get the number of passengers of each profile in a specific flight from the boarding database, using the information in Table II.

TABLE II
BOARDING INFORMATION

Profiles	Relevant fields for profiling
<i>Business</i>	#C/CL pax; #VIP's; #Freq. Flyer; #Pax according ticket price; Departure or arrival = business.
<i>Pleasure</i>	#Y/CL pax; #Pax according ticket price; Departure or arrival = vacation.
<i>Illness</i>	#Pax special needs; Stretcher on board=yes.

V.2. Defining the Passenger Cost Formulas

Besides being able to get the number and characterization of profiles from the survey data, we are also able to get the trend of each profile, regarding delay time/importance to the passenger. Plotting the data and the trend we got the graph in Fig. 3 (x – axis is the delay time and y – axis the importance).

If we apply these formulas as is, we would get quality costs for flights that do not delay. Because of that we wrote the formulas. The final formulas that express the

importance of the delay time for each passenger profile are presented in Table III.

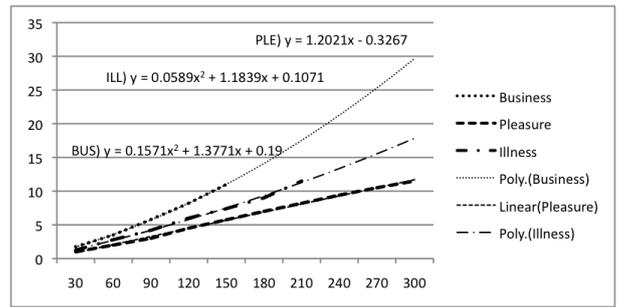


Fig. 3. Case study trend formulas for the profiles

It is important to point out that these formulas are valid only for this particular case and express the information we have from this specific survey data. Our goal is to update this information every year, using the annual company survey, and obtain different formulas according to flight destinations, flight schedules and/or geographical areas.

TABLE III
CASE STUDY FINAL QUALITY COST FORMULAS

Profile	Formula
<i>Business</i>	$y = 0.16*x^2 + 1.39*x$
<i>Pleasure</i>	$y = 1.20*x$
<i>Illness</i>	$y = 0.06*x^2 + 1.19*x$

V.3. Example

Using a real example from the scenario in section VI, let's calculate the quality operational costs for the following flight (assuming 10 as the coefficient to convert to monetary costs): Flight 103 will be delayed 30 minutes at departure. It has 20 passengers in the business profile (B), 65 in pleasure profile (P) and 1 in the illness profile (I). Applying the formulas in Table III, the cost of 30 minutes delay for each passenger in each profile is:

$$B_{\text{cost-103}} = 0.16*30^2 + 1.39*30 = 185.4$$

$$P_{\text{cost-103}} = 1.2*30 = 36$$

$$I_{\text{cost-103}} = 0.06*30^2 + 1.19*30 = 89.7$$

The quality operational cost for the flight 103 with a delay of 30 minutes is:

$$QC_{\text{cost-103}} = 10*(20*185.4 + 65*36 + 1*89.7) = 61377$$

VI. Scenarios and Experiments

To evaluate our approach we have setup a scenario that includes 3 operational bases (A, B and C). Each base includes their crewmembers each one with a

specific roster. The data used corresponds to a real airline operation of June 2006 of base A. We have simulated a situation where 15 crewmembers, with different ranks, did not report for duty in base A. A description of the information collected for each event is presented in Table IV.

TABLE IV
INFORMATION COLLECTED FOR EACH EVENT

Attribute	Description
<i>Event ID</i>	A number that represents the ID of the event. For tracking purposes only
<i>Duty Date Time</i>	The start date and time of the duty in UTC for which the crew did not report.
<i>Duty ID</i>	A string that represents the ID of the duty for which the crew did not report.
<i>Flt Dly</i>	Flight delay in minutes.
<i>C/Pax</i>	Number of passengers in business class.
<i>Y/Pax</i>	Number of passenger in economy class
<i>End Date Time</i>	The end date and time of the duty in UTC for which the crew did not report.
<i>Ready Date Time</i>	The date and time at which the crew member is ready for another duty after this one.
<i>Delay</i>	The delay of the crewmember. We have considered 10 minutes in our scenario.
<i>Credit Minutes</i>	The minutes of this duty that will count for payroll.
<i>Crew Group</i>	The crew group (Technical = 1; Cabin = 2) that the crewmember belongs to.
<i>Crew Rank</i>	CPT = Captain; OPT = First Officer; CCB = Chief Purser; CAB = Purser.
<i>Crew Number</i>	The employee number.
<i>Crew Name</i>	The employee name.
<i>Base ID</i>	The base where the event happened. We considered all events in base A.
<i>Open Positions</i>	The number of missing crews for this duty and rank. We used a fixed number of 1.

The events did not happen at the same day and each one corresponds to a crewmember that did not report for a specific duty in a specific day. Table V (at the end of the paper) shows the data for each of the events created. As you can see we have omitted the information regarding *Delay*, *Base ID* and *Open Positions* because we have used fixed values as indicated in Table IV. For example, the 10th event corresponds to the following situation: Crew Peter B, with number 32 and rank CPT (captain) belonging to the crew group 1 (technical crew), did not report for the duty with ID 1ZRH12X with briefing time (duty date time) at 15:25 in 15-06-2006. This flight did not delay on departure and has 5 passengers in business class and 115 in economy class. The event was created after a 10 minutes delay of the crewmember in reporting for duty and happened at base A. It is necessary to find another crewmember to be assigned to this duty. The duty ends at 09:30 on 17-06-2006 and the crewmember assigned to this duty will be ready for another one at 21:30 in 17-06-2006. The duty

will contribute with 1318 minutes (21h58) for the payroll. The new crewmember must belong to the same rank and group. After setting-up the scenario we found the solutions for each crew event using three methods.

VI.1. Experimental Methods

In the first method (*human*) we used one of the best users from the AOCC, with current tools available, to find the solutions. The user uses software that shows the roster of each crewmember 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. Each user has a specific way of trying to find the solutions. However, we have observed that, in general, they follow these steps:

1. Open the roster for a one month period, starting two days before the current day. For example, let's suppose that the current day is 7th of June of 2006, they open the roster from the 5th of June until the 4th of July.
2. Filter the roster by crew rank and base, where the base is equal to the base where the crew event happened and crew rank is equal to the crewmember rank that did not report for duty.
3. Order the information by month duty, in an ascendant order and by seniority in a descendant order.
4. Visually, they scroll down the information until they found a crewmember with an open space for the period of time that corresponds to the duty to be assigned. This period of time takes into consideration the start and end time of the duty and also the time required for resting (ready date time).
5. If they do not found a crewmember in the base specified, they try to find it in another base, filtering the information accordingly.
6. They assign the duty to the crewmember with less *credit hours*.

The data collected using this method is presented in Table VI. We point out that the data in columns marked with an asterisk where calculated manually, according to the equations presented in section III. The reason for this is that the information system that is available for the users does not include information related with any kind of costs.

In the second method (*agent-no-quality*) we have used our approach but with $\beta=0$ in Equation 1 (*Total Operational Cost*), i.e., although we calculate the *Quality Operational Cost* as indicated in Equation 6 we did not considered this value in resolution as well as in the decision process. The data collected is presented in Table VII.

In the third method (agent-quality) we have used our approach but with $\beta=1$ in Equation 1, i.e., considering the Quality Operational Cost in the resolution as well as in the decision process. The *Quality Operational Cost* was calculated using two passenger profiles (business and economy classes) and with $\alpha=0,1$. We used the first two formulas in Table III to calculate the delay cost of each passenger in business and economy profile, respectively. The data collected is presented in Table VIII.

VII. Results and Discussion

Table IX shows a comparison of the results obtained through the above methods. We point out that in method 1 (human) we did not calculate the quality costs, and in method 2 (agent-no-quality) we did calculate the quality costs but they were not used to find the best solution, although we use that value for comparison purposes. From the results obtained we can see that on average, method 3 (*agent-quality*) produced solutions that decreased flight delays in aprox. 36%. *Agent-quality* is, on average 3% slower than *agent-no-quality* in finding a solution and produces solutions that represent a decrease of 23,36% on the total operational costs, when compared with *agent-no-quality*.

From the results (Table IX) we can see that our approach obtains valid solutions faster and with fewer direct operational costs when compared with the current method used in a real airline company (*human*). *Agent-no-quality* represents a decrease of aprox. 45,5% and *agent-quality* a decrease of aprox. 41%. *Agent-quality* has a higher direct operational cost than *agent-no-quality* because it uses the quality operational cost in the decision process. If we read this number without any other consideration, we have to say that the goal of having less direct operational costs was not achieved. An 8% increased on direct operational costs can represent a lot of money. However, we should read this number together with the flight delay figure. As we can see, although *agent-quality* has increased the direct operational costs (when compared with *agent-no-quality*) in 8% it was able to choose solutions that decrease, in average, 36% of the flight delays. This means that, when there are multiple solutions to the same problem, *agent-quality* is able to choose the one with less operational cost, less quality costs (hence, better passenger satisfaction) and, because of the relation between quality costs and flight delays, the solution that produces shorter flight delays.

From this conclusion, one can argue that if we just include the direct operational costs and the expected flight delay, minimizing both values, the same results could be achieved having all passengers happy. In general, this assumption might be true. However, when

we have to choose between two solutions with the same direct operational cost and delay time, which one should we choose? In our opinion, the answer depends on the profile of the passengers of each flight and on the importance they give to the delays (quality operational cost), and not only in minimizing the flight delays and direct operational cost. *Agent-quality* takes into consideration this important information when making decisions. This is the reason why we think that one of the main contributions of our work is the generic approach to quantify the passenger satisfaction regarding delaying a flight, from the passenger point of view, presented in section IV. It is fair to say that we cannot conclude that our MAS will always have this behaviour. For that we need to evaluate a higher number of scenarios, at different times of the year (we might have seasonal behaviours) and, then, find an average value.

Additionally, we found that the cooperation between different operational bases has increased with our approach, because we evaluate all the solutions found (including the ones from different operational bases where the event happened) and we select the one with less cost. In *human*, they choose the first one they find with less *credit hours*, usually from the same base where the event was triggered. This cooperation is also possible to be inferred from the costs by base. In Table IX is possible to see that the direct operational costs of base C using *human* represents only 7,58% of the costs of all bases, whilst in *agent-no-quality* and *agent-quality* it represents 88,77% and 51,73%, respectively. The same is possible to be inferred from the other bases (although with different figures). This means that our MAS uses more resources from other bases than the base where the problem happened (base A).

VIII. Conclusion and Future work

In this paper we describe our agent-based approach to the airline operations recovery problem, including the reasons that make us adopt an agent and multi-agent system (MAS) paradigm. We have detailed our MAS architecture, including: (i) agents and protocols as well as some agent characteristics like autonomy and social-awareness, and (ii) decision mechanisms, including the costs criteria and negotiation protocols used. One of the major contributions of our work is a way of quantifying quality costs that, we believe, represents better the passenger satisfaction and allows the airline company to include this important parameter when taking operational decisions. Using data from a real airline company, we tested our approach and discussed the results obtained by three different methods. We have shown that our approach is able to select solutions that contribute to a better passenger satisfaction and that

produce shorter flight delays when compared with methods that only minimize direct operational costs.

We are working on several improvements. Some of them are already implemented. However, we did not perform, yet, enough tests to have meaningful results. These are our current and future goals:

- Improve autonomy and learning characteristics of the *Monitor* agent, so that he is able to consider new events (or change existing ones) according to the experience he gets from monitoring the operation, without relying exclusively on the definition of events created by the human operator.
- Working on a protocol at the *Manager Agent* team level that allows a better coordination and improves the distributed problem solving characteristics of our approach. For example, including in each team, knowledge provided by other teams to improve the objective function of each specialist agent, with parameters of the other dimensions (aircraft, crew and passenger).
- Solving problems learning by example, applying Case-Based Reasoning (CBR).
- Increase robustness of future schedules by applying the knowledge gathered from learning by example.
- Study the behaviour and compare the results, of several problem solving algorithms, including the ones that implement heuristics to specific problems. The idea is to classify the algorithms according to their success rate in solving specific types of problems in this domain.

Acknowledgements

The first author is supported by FCT (Fundação para a Ciência e Tecnologia) under research grant SFRH/BD/44109/2008. The authors are grateful to TAP Portugal for allowing the use of real data from the company.

References

- [1] Clausen, J., Larsen, A., Larsen, J., Disruption Management in the Airline Industry – Concepts, Models and Methods, *Technical Report, 2005-01*, Informatics and Mathematical Modelling, Technical University of Denmark, DTU, 2005.
- [2] Kohl, N., Larsen, A., Larsen, J., Ross, A., Tiourline, S., Airline Disruption Management – Perspectives, Experiences and Outlook. *Technical Report, CRTR-0407*, Carmen Research, 2004.
- [3] T. Liu, C. Jeng, Y. Chang, Disruption Management of an Inequality-Based Multi-Fleet Airline Schedule by a Multi-Objective Genetic Algorithm. *Transportation Planning and Technology*, Vol. 31, n.6, pp. 613-639, 2008.
- [4] M. Yang, *Using Advanced Tabu Search Techniques to Solve Airline Disruption Management Problems*, Ph.D. Thesis, The University of Texas at Austin, USA, 2007.
- [5] Rosenberger, J., Johnson, E., Nemhauser, G., Rerouting aircraft for airline recovery. *Technical Report, TLI-LEC 01-04*, Georgia Institute of Technology, 2001.
- [6] J. Rosenberger, A. Schaefer, D. Goldsmans, E. Johnson, A. Kleywegt, G. Nemhauser, A Stochastic Model of Airline Operations, *Transportation Science*, Vol. 36, n. 4, pp. 357-377, 2002.
- [7] A. Abdelgahny, G. Ekollu, R. Narisimhan, K. Abdelgahny, A Proactive Crew Recovery Decision Support Tool for Commercial Airlines during Irregular Operations, *Annals of Operations Research*, Vol. 127, pp. 224-232, 2004.
- [8] S. Bratu, C. Barnhart, Flight Operations Recovery: New Approaches Considering Passenger Recovery. *Journal of Scheduling*, Vol. 9, n. 3, pp. 279-298, 2006.
- [9] L. Lettovsky, *Airline Operations Recovery: An Optimization Approach*, Ph.D. Thesis, Georgia Institute of Technology, Atlanta, USA, 1997.
- [10] Castro, A.J.M., Oliveira, E., Using Specialized Agents in a Distributed MAS to Solve Airline Operations Problems: a Case Study, *Proceedings of IAT 2007 (Intelligent Agent Technology Conference)*, pp. 473-476, Silicon Valley, California, USA 2-5 November 2007, IEEE Computer Society, ISBN: 0-7695-3027-3.
- [11] Tumer, K., Agogino, A., Distributed Agent-Based Air Traffic Flow Management, *Proceedings of the 6th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2007)*, Honolulu, Hawaii, 2007.
- [12] Wolfe, S. R., Jarvis, P. A., Enomoto, F. Y., Sierhuis, M., Comparing Route Selection Strategies in Collaborative Traffic Flow Management, *Proceedings of IAT 2007 (Intelligent Agent Technology Conference)*, Silicon Valley, California, USA 2-5 November 2007, IEEE Computer Society, ISBN: 0-7695-3027-3
- [13] Jonker, G., Meyer, J.J., Dignum, F., Towards a Market Mechanism for Airport Traffic Control, *Proceedings of the 12th Portuguese Conference on Artificial Intelligence (EPIA 2005)*, Covilhã, Portugal, 2005.
- [14] D. Ouelhadj, *A Multi-Agent System for the Integrated Dynamic Scheduling of Steel Production*, Ph.D. Thesis, The University of Nottingham, School of Computer Science and Information Technology, England 2003.
- [15] A. Castro, E. Oliveira, The rationale behind the development of an airline operations control centre using Gaia-based methodology, *Int. J. Agent-Oriented Software Engineering*, Vol. 2, n. 3, pp. 350-377, 2008.
- [16] F. Bellifemine, G. Caire, T. Trucco, G. Rimassa, *JADE Programmer's Guide*. JADE 3.3 TILab S.p.A, 2004.
- [17] M. Wooldridge, When is an Agent-Based Solution Appropriate? In *Introduction to Multiagent Systems*, (West Sussex, England: John Wiley & Sons, Ltd., 2002, 225-226)
- [18] A. Elamy, Perspectives in Agents-Based Technology, In *AgentLink News 18*, August 2005, ISSN: 1465-3842.
- [19] P. Stone, M. Veloso, Multi-Agent Systems: a Survey from a Machine Learning Perspective, *Autonomous Robotics*, Vol. 8. N. 3, 2000.
- [20] Fipa, *FIPA Contract Net Interaction Protocol Specification*, <http://www.fipa.org/specs/fipa00029/SC00029H.html>, Accessed in 2009/07/18.
- [21] R. Smith, The Contract Net Protocol: High-Level Communication and Control in a Distributed Problem Solver, *IEEE Transactions on Computers*, Vol. C, n. 29, 1980.
- [22] N. Kohl, S. Karish, Airline Crew Rostering: Problem Types, Modeling and Optimization, *Annals of Operations Research*, Vol. 127, pp. 223-257, 2004.

TABLE V
EVENT DATA USED FOR TESTING

Duty Date Time	Duty ID	Flt Dly	C Pax	Y Pax	End Date Time	Ready Date Time	Cred Min	Crew Grp	Rank	Crew Nr	Crew Name
05-06 07:25	1ORY149S	0	7	123	05-06 13:35	06-06 01:35	370	2	CAB	80	John A
05-06 07:25	1ORY149S	10	11	114	05-06 13:35	06-06 01:35	370	2	CAB	45	Mary A
05-06 07:25	1ORY85P	0	10	112	05-06 13:35	06-06 01:35	370	1	CPT	35	Anthony
15-06 04:10	2LIS24X	30	0	90	16-06 16:15	17-06 04:15	1757	2	CAB	99	Paul M
15-06 04:10	3LIS25X	25	3	77	15-06 09:20	15-06 21:20	632	2	CAB	56	John B
15-06 12:50	2LHR63P	5	25	85	16-06 20:45	17-06 08:45	1549	1	CPT	57	Paul S
15-06 12:50	2LHR63P	0	20	95	16-06 20:45	17-06 08:45	1549	1	OPT	53	Mary S
15-06 14:15	1LHR31P	0	23	52	15-06 20:55	16-06 08:55	843	2	CCB	23	Sophie
15-06 15:25	2LHR19P	10	27	105	16-06 20:45	17-06 08:45	1341	2	CCB	34	Angel
15-06 15:25	1ZRH12X	0	5	115	17-06 09:30	17-06 21:30	1318	1	CPT	32	Peter B
25-06 05:20	1LIS16S	20	3	97	25-06 15:05	26-06 03:05	585	2	CAB	20	Paul G
25-06 05:20	1LIS16S	5	2	108	25-06 15:05	26-06 03:05	585	2	CAB	10	Alice
25-06 05:20	1LIS158T	0	4	92	25-06 15:05	26-06 03:05	585	2	CAB	15	Daniel
25-06 06:15	3LIS174S	0	1	129	27-06 16:15	28-06 04:15	1258	2	CAB	71	George
25-06 14:20	4LIS50A	0	2	83	28-06 19:40	29-06 07:40	219	1	OPT	65	Allan

TABLE VI
DATA COLLECTED (PARTIAL) AFTER USING METHOD 1 (HUMAN)

Duty ID	Base ID	Crew Grp	Rank	Hour Pay (*)	Perdiem Pay (*)	Quality Op. Cost	Op. Cost (*)
1ORY149S	A	2	CAB	0,00	72,00	0	72,00
1ORY149S	B	2	CAB	0,00	72,00	0	86,40
1ORY85P	A	1	CPT	942,90	106,00	0	1048,90
2LIS24X	A	2	CAB	939,00	144,00	0	1083,00
3LIS25X	B	2	CAB	0,00	72,00	0	86,40
2LHR63P	B	1	CPT	777,00	212,00	0	1186,80
2LHR63P	B	1	OPT	0,00	148,00	0	177,60
1LHR31P	A	2	CCB	687,65	72,00	0	759,65
2LHR19P	B	2	CCB	0,00	144,00	0	172,80
1ZRH12X	C	1	CPT	0,00	212,00	0	296,80
1LIS16S	A	2	CAB	0,00	72,00	0	72,00
1LIS16S	C	2	CAB	0,00	72,00	0	100,80
1LIS158T	B	2	CAB	0,00	72,00	0	86,40
3LIS174S	A	2	CAB	1051,60	216,00	0	1267,60
4LIS50A	A	1	OPT	246,40	296,00	0	542,40
Totals				4644,55	1982,00	0	7039,55

TABLE VII
DATA COLLECTED (PARTIAL) AFTER USING METHOD 2 (AGENT-NO-QUALITY)

Duty ID	Base ID	Crew Grp	Rank	Hour Pay	Perdiem Pay	Quality Op. Cost	Direct Op. Cost
1ORY149S	A	2	CAB	0,00	72,00	0	72,00
1ORY149S	B	2	CAB	0,00	72,00	501,31	86,40
1ORY85P	B	1	CPT	0,00	106,00	0	127,20
2LIS24X	C	2	CAB	563,40	62,00	1561,76	875,56
3LIS25X	B	2	CAB	0,00	72,00	1877,73	86,40
2LHR63P	C	1	CPT	0,00	212,00	658	296,80
2LHR63P	A	1	OPT	0,00	144,00	687,62	144,00
1LHR31P	B	2	CCB	229,17	72,00	0	361,40
2LHR19P	B	2	CCB	0,00	144,00	788,78	172,80
1ZRH12X	C	1	CPT	0,00	212,00	0	296,80
1LIS16S	A	2	CAB	0,00	72,00	961,95	72,00
1LIS16S	C	2	CAB	0,00	72,00	301,48	100,80
1LIS158T	B	2	CAB	0,00	72,00	0	86,40
3LIS174S	C	2	CAB	411,00	93,00	0	705,60
4LIS50A	B	1	OPT	0,00	296,00	449,84	355,20
Totals				1203,57	1773,00	7788,47	3839,36

TABLE VIII
DATA COLLECTED (PARTIAL) AFTER USING METHOD 3 (AGENT-QUALITY)

Duty ID	Base ID	Crew Grp	Rank	Hour Pay	Perdiem Pay	Quality Op. Cost	Direct Op. Cost
1ORY149S	A	2	CAB	0,00	72,00	0	72,00
1ORY149S	B	2	CAB	0,00	72,00	501,31	86,40
1ORY85P	B	1	CPT	0,00	106,00	0	127,20
2LIS24X	C	2	CAB	503,50	144,00	1060,92	906,50
3LIS25X	C	2	CAB	0,00	72,00	1420,78	100,80
2LHR63P	B	1	CPT	102,90	212,00	272,10	377,88
2LHR63P	B	1	OPT	37,22	144,00	0	217,46
1LHR31P	B	2	CCB	229,17	72,00	0	361,40
2LHR19P	B	2	CCB	0,00	144,00	788,78	172,80
1ZRH12X	C	1	CPT	0,00	212,00	0	296,80
1LIS16S	A	2	CAB	0,00	80,00	593,30	80,00
1LIS16S	C	2	CAB	0,00	80,00	144,34	112,00
1LIS158T	B	2	CAB	0,00	72,00	0	86,40
3LIS174S	C	2	CAB	411,00	93,00	0	705,60
4LIS50A	A	1	OPT	138,83	288,00	0	426,83
Totals				1422,62	1863,00	4781,53	4130,07

TABLE IX
SUMMARY OF THE RESULTS OBTAINED BY EACH METHOD

	Human	%	Agent-no-quality	%	Agent-quality	%
<i>Base of the solution:</i>						
- From crew event base (A)	7	47%	3	20%	3	20%
- From base B	6	40%	7	47%	7	47%
- From base C	2	13%	5	33%	5	33%
<i>Time to find solution (avr sec)</i>	101	100%	25	24,75%	26	25,74%
<i>Flight delays (avr min):</i>						
- Base A			11	100%	7	63,64%
- Base B			14	40%	7	30%
- Base C			9	26%	4	17%
- Base C			12	34%	12	52%
<i>Total direct operational costs</i>						
	7039,60	100%	3839,36	54,54%	4130,07	58,67%
- Base A	4845,55	92,42%	288,00	11,23%	578,83	14,02%
- Base B	1796,40	34,26%	1275,80	49,77%	1429,54	34,61%
- Base C	397,60	7,58%	2275,56	88,77%	2121,70	51,37%
<i>Total quality operational costs</i>						
			7788,47	100%	4781,53	61,39%
- Base A			1649,57	21,18%	593,30	12,41%
- Base B			3617,66	46,45%	1562,19	32,67%
- Base C			2521,24	32,37%	2626,04	54,92%
<i>Total operational costs</i>						
			11628,01	165%	8911,60	126,6%
- Base A			1937,57	16,66%	1172,13	13,15%
- Base B			4088,42	35,16%	2991,73	33,57%
- Base C			4796,80	41,25%	4747,74	53,28%

Authors' information

The photo must be 2.45 cm x 2.45 cm. The text (8 pt) wrapping style must be around the frame.

¹LIACC/NIAD&R, DEI/FEUP, University of Porto.

²LIACC/NIAD&R, DEI/FEUP, University of Porto.



António J. M. Castro was born in Portugal in 1965. He holds an MSc in Artificial Intelligence and Intelligent Systems from the Faculty of Engineering, University of Porto (2007), a BcS in Information Systems Engineering from the Polytechnic Engineering Institute in Porto (1997) and a Post-graduation in Air Transport Operations from ISEC in Lisbon (2008). António is a Software Engineer at TAP Portugal, the Portuguese airline company, where he manages software projects related with operations control and crew scheduling. He is currently a PhD student on Informatics Engineering at the University of Porto, Faculty of Engineering. His current research interests include multi-agent systems, agent-oriented software engineering and distributed systems.



Eugénio Oliveira was born in Portugal in 1948. He is a full Professor at the University of Porto in Portugal. He is coordinator of a Research Group on Distributed Artificial Intelligence and of the LIACC research Lab at the same University. He is also Director of the Doctoral Program in Informatics. He was first with Brown, Boveri C, in Switzerland and later, in the eighties he got his PhD in Knowledge Engineering and Logic Programming at the Universidade Nova in Lisbon. He was Guest Academic at IBM/IEC in La Hulpe, Belgium (84-85). He was awarded with the Gulbenkian Prize in 83. He has been involved in several European as well as national funded projects involving advanced features of intelligent agents for applications. Agent-based frameworks for B2B, multi-agent learning methods and agent-based team work strategies are among his main current topics of interest. He has published a large number of papers in international journals and conference proceedings.