Minimizing Airport Peaks Problem by Improving Airline Operations Performance through an Agent Based System

António J.M. Castro[1], António Mota[1], Luís Paulo Reis[1], Eugénio Oliveira[1]

[1] LIACC-NIAD&R, FEUP, Department of Informatics Engineering, University of Porto, 4200-465 Porto, Portugal
{ajmc, amota, lpreis, eco}@fe.up.pt

Abstract. Airports are important infra-structures for the air transportation business. One of the major operational constraints is the peak of passengers in specific periods of time. Airline companies take into consideration the airport capacity when building the airline schedule and, because of that, the execution of the airline operational plan can contribute to improve or avoid airport peak problems. The Airline Operations Control Center (AOCC) tries to solve unexpected problems that might occur during the airline operation. Problems related to aircrafts, crewmembers and passengers are common and the actions towards the solution of these problems are usually known as operations recovery. In this paper we propose a way of measuring the AOCC performance that takes into consideration the relation that exists between airline schedule and airport peaks. The implementation of a Distributed Multi-Agent System (MAS) representing the existing roles in an AOCC, is presented. We show that the MAS contributes to minimize airport peaks without increasing the operational costs of the airlines.

Keywords: Disruption management, multi-agent system, airport operation.

1 Introduction

Airports are a very important infrastructure for air transportation. They provide services for airlines and, also, for the passengers that fly on those airlines. Airport operators like to speak of the business of the airport in terms of throughput of passengers and cargo as represented by the annual number of passengers processed or the annual turnover of tons of air freight. This is entirely understandable because, most likely, the annual income is determined by these two parameters. However, from an operational point of view, it is the peak flows that determine the physical and operational costs involved in running a facility. Assigning staff and physical facilities are much more dependent on hourly and daily requirements than on annual throughput.

Airline companies, during their Airline Scheduling Process (especially during the Flight Schedule Generation phase) take into consideration the agreed schedule regarding departures/arrivals of the airports, especially important on hub airports (for more information regarding this process see [8]). Given the above, we believe that it is important for the airline company to operate according to the schedule, not only due to the fact that the airline schedule is the optimal one from the airline perspective but, also, because it takes into consideration the airport capacity and, therefore, the airport peaks. 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 [5]. Those problems are related to crewmembers, aircrafts and passengers. 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 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 [11] or Operations Recovery.

In this paper we propose a way of measuring the AOCC performance so that, in the decision process, the AOCC takes into account the relation that exists between airport peaks and the airline schedule. We present the architecture and specification of a multi-agent system that was developed for a real airline company, that uses our proposed measure (among other criteria) to solve operational problems.

The rest of the paper is organized as follows. In section 2 we present some related work. Section 3 explains the relation between airport peaks and airline schedule and proposes the AOCC performance criteria. Section 4 shows the architecture and specification of our MAS. In section 5 we present the scenario used to evaluate the system as well as the results of the evaluation. Finally, we discuss and conclude our work in section 6.

2 Related work

We divided the bibliography we have analyzed in three main areas: aircraft recovery, crew recovery and integrated recovery.

Aircraft Recovery: Liu et al. [12] 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 trying to minimize the delays for all aircrafts and flights. Although there are other differences regarding our approach, the main one is that we emphasize the role of the airport trying to minimize the difference between the real and schedule plan of the airline at each airport.

Crew Recovery: In Abdelgahny et al. [1] 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. According to the authors the tool is able to "solve for the most efficient crew recovery plan with least deviation from originally planned schedule". The major drawback is that it only includes one resource (crew) and does not consider the passenger dimension.

Integrated Recovery: Bratu et al. [4] 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. The objective is to minimize jointly airline operating costs and estimated passenger delay and disruption costs. According to the authors, "(…) decisions from our models can potentially reduce passenger arrival delays (…) without increasing operating costs". The main difference regarding our approach is that we emphasize the role of the airport trying to minimize the difference between the real and schedule plan of the airline at each airport. Castro and Oliveira [6] 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. The authors only use operational costs on the decision process ignoring if the AOCC is near the original schedule or not.

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 [14] developed a multi-agent algorithm for traffic flow management. Wolfe et al. [15] uses agents to compare routing selection strategies in collaborative traffic flow management. For ATC Tower operations, Jonker et al. [9] 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 [13].

3 Airport Peaks, Airline Schedule and Operations Performance

According to Ashford et al. [2] there are four ways of describing variations in demand level with time:

1. Annual variation over time.
2. Monthly peaks within a particular year.
3. Daily peaks within a particular month or week.
4. Hourly peaks within a particular day.

The first one is very important from the viewpoint of planning and provision of facilities. For our work, we concentrate on monthly, daily and hourly peaks, because these are the ones that have more impact on day-to-day operations of the airports and airlines. The goal of airport operators is to spread demand more evenly over the operating day in order to decrease the costs associated with running the airport at peak times, avoiding, as much as possible, situations like the one presented in Figure 1.

On the other hand, airlines are looking to maximize fleet utilization and offer flights in more attractive slots. Additionally, airlines that operate in a hub-and-spoke, due to the characteristics of such an operation, want to minimize the total travel time and, for that, they need to rapidly connect the passengers that are arriving from long-haul flights to short-haul flights and vice-versa.

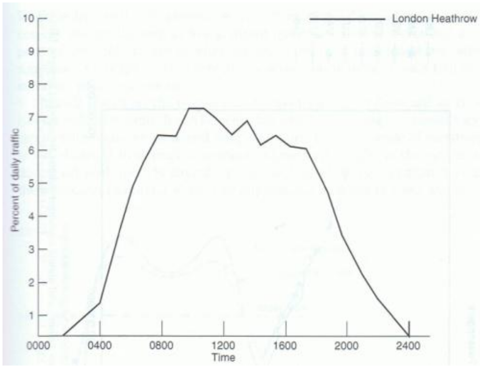

Fig. 1. Hourly variations of passenger traffic in a typical peak day (Source: BAA plc.)

As it is possible to see in Figure 2 this type of network makes airline companies to schedule *waves*, that is, a high number of aircrafts arriving or departing at the hub in a short time interval.

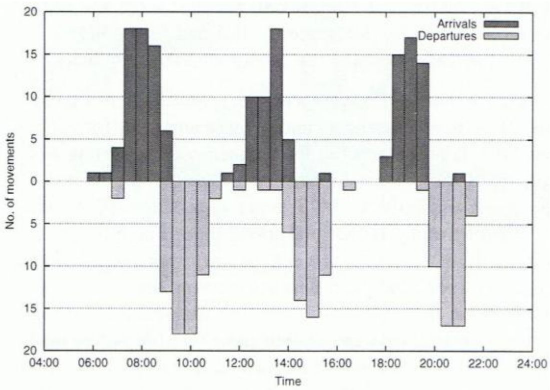

Fig. 2. Schedule structure of the Alitalia hub in Milan (MXP). (Source: [7])

Given the difference between the goal of the airport operator and the goal of the airline, there is, therefore, "a potential conflict between the airline satisfying its customer, the passenger, and the airport attempting to influence the demands of its customer, the airline" [2]. Because of that, it is important that airport operators and airlines cooperate regarding the flight schedule definition.

International Air Transport Association (IATA) has developed a general policy in scheduling so that, at some airports with official limitations, general government authorities carry out the coordination. From our own personal knowledge of the air transportation business as well as according to Ashford et al. [2], it is much more common the situation where the airlines establish themselves an agreed schedule through the mechanism of airport coordinator. The largest or national carrier of the airport, assumes the role of airport coordinator (TAP in Lisbon, Lufthansa in Frankfurt and BA in Heathrow, for example) and at semi-annual IATA scheduling conferences, they are able to set an agreed schedule for the airports they represent. As we stated before, the Airline Scheduling Process takes into consideration the agreed schedule. The goal of that process is to create an airline schedule that is optimal in regard to a given objective, usually operating profit [8], that is, minimum costs and maximum revenues. To operate according to the schedule is not an easy task. Airline companies face a lot of unexpected events during the operations of their flights [11]. However a good disruption management process should exist to minimize the impact of the unexpected events and, according to Yu [16] return to the original schedule as soon as possible[1]. For that, the AOCC should take decisions when solving disruptions that tend to the original schedule. We propose to measure the performance of AOCC's according to Equation 1.

$$\rho = \sum_{t=1}^{n} \sum_{a=1}^{|A|} \sum_{f=1}^{|F|} \left|\Delta dt_{\{f,a,t\}}\right| + \left|\Delta at_{\{f,a,t\}}\right| \qquad \textbf{(1)}$$

$\underline{\textbf{\textit{where}}}$

$f \in F; F = \{flights\}, \quad a \in A; A = \{airports\}$

$t = time\ period\ (days), t \geq \alpha \leq \beta,$

$\alpha = start\ datetime\ of\ the\ AOCC$

$\beta = end\ datetime\ of\ the\ AOCC$

$\Delta dt_{\{f,a,t\}} : schedule/actual\ departure\ variation$

$\Delta at_{\{f,a,t\}} : schedule/actual\ arrival\ variation$

We might say that if ρ->0 (tends to) then the real operation of the airline is running near the original schedule contributing to improve the performance of the airport during peak times. At the same time, we might say that the AOCC is contributing to minimize the real operational costs. In the next section we present the multi-agent system (MAS) we have developed to help the AOCC. The MAS uses the performance criteria above (Equation 1) and, also, other criteria related with operational costs.

4 System Architecture and Specification

System overview: This section presents the architecture and specification of the multi-agent system (MAS) we have developed for the airline operations control centre (AOCC). Figure 3 shows one instance of the architecture of the system. There are seven types of agents:

---

[1] Assuming the original schedule as the optimal one.

- *Monitor*, which monitors the operation of the airline company.
- *EventType*, which defines the types of events that must be detected.
- *ResolutionManager*, which receives a problem and manages the resolution in cooperation with the specialist agents.
- *SimmAnneal* and *HillClimb*, specialist agents responsible for the resolution of a problem, using simulated annealing and hill climb algorithms, respectively.
- *Supervisor*, the agent that interacts with the human supervisor, showing the solutions proposed and requesting authorization to apply them.
- *ApplySolution*, the agent that is responsible to apply the solution in the environment.

Figure 3 shows also the existence of a data store, which has information about the airline company operations. The data store is accessed by the *Monitor*, *Specialist* and *ApplySolution* agent. The communication between the agents is done through the JADE system [3] and the data is passed between agents as serializable Java objects.

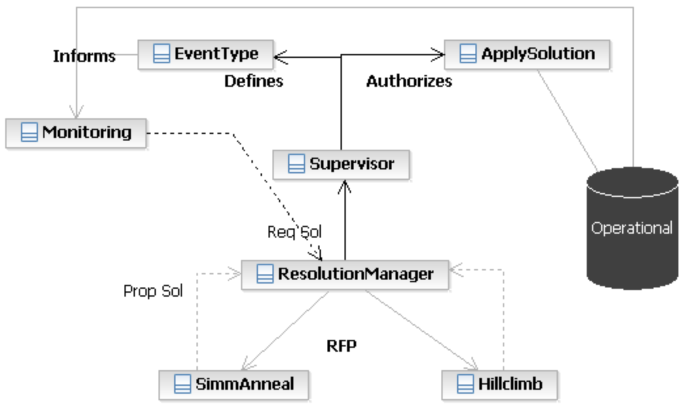

Fig. 3. Overall architecture of the Multi-Agent System

Airline Schedule and Actual Operational Plan Definition: In a simplified version of an airline schedule, we may say that it is composed of flights and the resources necessary to perform those flights (aircraft and crewmembers), in a specific period. The AOCC typically takes control of the airline operational plan some hours/days before the operation of a flight until some hours/days after. We can define the airline schedule plan S as:

$$S = \{s_{\{f,a,t\}}: s_{\{f,a,t\}} = (id, dp^s_{\{f,a,t\}}, ar^s_{\{f,a,t\}}, dt^s_{\{f,a,t\}}, at^s_{\{f,a,t\}}, ac^s_{\{f,a,t\}}, cm^s_{\{f,a,t\}}, to^s_{\{f,a,t\}})\} \quad \textbf{(2)}$$

$t \geq \alpha \leq \beta$

where

$id : flight\ identification$

$dp^s_{\{f,a,t\}} : sched.\ departure\ airport;\ \ dp^s_{\{f,a,t\}} \in A, dp^s_{\{f,a,t\}} = a$

$ar^s_{\{f,a,t\}} : sched.\ arrival\ airport;\ \ ar^s_{\{f,a,t\}} \in A, ar^s_{\{f,a,t\}} \neq a$

$dt^s_{\{f,a,t\}} : sched.\ departure\ date\ time$

$at^s_{\{f,a,t\}} : sched.\ arrival\ date\ time$

$ac^s_{\{f,a,t\}} : sched.\ aircraft\ assigned\ flight;\ \ ac^s_{\{f,a,t\}} \in AC: AC = \{aircrafts\}$

$cm^s_{\{f,a,t\}} : sched.\ crew\ assigned\ flight;\ \ cm^s_{\{f,a,t\}} \in CM: CM = \{crewmembers\}$

$to^s_{\{f,a,t\}} : sched.\ total\ operational\ cost$

Similarly, the actual operational plan R can be defined as:

$$R = \{r_{\{f,a,t\}} : r_{\{f,a,t\}} = (id, dp^r_{\{f,a,t\}}, ar^r_{\{f,a,t\}}, dt^r_{\{f,a,t\}}, at^r_{\{f,a,t\}}, ac^r_{\{f,a,t\}}, cm^r_{\{f,a,t\}}, to^r_{\{f,a,t\}})\} \quad \textbf{(3)}$$
$$t \geq \alpha \leq \beta$$

In the case of Equation 3 the components reflect the actual data as opposed to the schedule data in S (the airline schedule plan).

Problem Specification: When agent *Monitoring* detects a problem that needs to be solved, a problem is raised and a solution is requested (through a *Fipa-Request*[2] protocol) from the *ResolutionManager* agent. Equation 4 represents the problem.

$$P = \{p_{\{f,a,t\}} : p_{\{f,a,t\}} = (id, dt^p_{\{f,a,t\}}, et^p_{\{f,a,t\}}, cm^p_{\{f,a,t\}}, dly^p_{\{f,a,t\}}, svr^p_{\{f,a,t\}}, tw^p_{\{f,a,t\}}, bd^p_{\{f,a,t\}}, csd^p_{\{f,a,t\}}, to^p_{\{f,a,t\}}, s_{\{f,a,t\}})\} \quad \textbf{(4)}$$
$$t \geq \alpha \leq \beta$$

$\underline{\textbf{\textit{where}}}$

id : problem id.

$dt^p_{\{f,a,t\}} = dt^s_{\{f,a,t\}}$

$et^p_{\{f,a,t\}} \in E : E = \{flight\ delay, crew\ delay, pax\ delay\}$

$dly^p_{\{f,a,t\}}$: minutes of delay

$to^p_{\{f,a,t\}} = to^s_{\{f,a,t\}}$

$svr^p_{\{f,a,t\}} \in SV : SV = \{warning, problem\}$

$tw^p_{\{f,a,t\}}$: time window for change, $\alpha < (tw^p_{\{f,a,t\}} - dt^p_{\{f,a,t\}});\ \beta > (tw^p_{\{f,a,t\}} + dt^p_{\{f,a,t\}})$

$$cm^p_{\{f,a,t\}} = \begin{cases} cm^s_{\{f,a,t\}}, & if\ et^p_{\{f,a,t\}} = \{crew\ delay\} \\ \emptyset, & otherwise \end{cases}$$

$$bd^p_{\{f,a,t\}} = \begin{cases} 0, & if\ p_{\{f,a,t\}}\ raised\ by\ Monitor\ agent \\ > 0, & after\ CFP\ by\ ResolutionManager\ agent \end{cases} : bid\ dealine\ in\ minutes.$$

$$csd^p_{\{f,a,t\}} = \begin{cases} 0, & if\ p_{\{f,a,t\}}\ raised\ by\ Monitor\ agent \\ > 0, & after\ accept\_proposal\ by\ ResolutionManager\ agent \end{cases} : candidate\ solution\ deadline\ in\ minutes.$$

For example: crewmember 231, delayed 10 minutes for flight TP438, departing from Lisbon would be represented as problem:

$p_{\{tp438,lis,0715\}}$=(001,09/07/01-0715,crew delay,231,10,problem,3,0,0,$s_{\{tp438,lis,0715\}}$).

At this stage, the *Monitoring* agent only adds to the *Problem*, information related with the operational plan time window that should be involved in the resolution process. In the example above, each specialist agent only considers flights between 04:15 and 10:15 of day 09/07/01 (09/07/01 07:15 ± 3 hours).

Resolution Manager: Agent *ResolutionManager* responds to the request from *Monitoring* agent, issuing a RFP to the specialist agents *SimmAnneal* and *HillClimb* and adding to the *Problem* the bid deadline and the deadline to receive candidate solutions (for example, $bd^p_{\{tp438,lis,0715\}}=1$ and $csd^p_{\{tp438,lis,0715\}}=5$). At this level a *Fipa-Contract.net* protocol is used to negotiate with the specialist agents. The specialist agents interested to respond to the RFP, should manifest that intention before the bid deadline. The specialist agents have a limited amount of time to find a candidate solution that is represented by $csd^p_{\{tp438,lis,0715\}}$ in $p_{\{tp438,lis,0715\}}$ (example above). The first step of the specialist agents is to obtain the flights that are in the time window of the problem, represented by $tw^p_{\{tp438,lis,0715\}}$. The set $FT_{\{p\}}$ represents these initial

---

[2] http://www.fipa.org

flights and will be the initial solution of the problem. The crewmembers and aircraft exchanges are made between flights of $FT_{\{p\}}$. Finally, when a candidate solution is found, the specialist agents send it to the *ResolutionManager* agent. Equation 5 defines a candidate solution.

$$PS = \{ps_{\{p\}}: ps_{\{p\}} = \left(id, fc^{ps}_{\{p\}}, ic^{ps}_{\{p\}}, \{ft^{ps}_{\{p\}}\}, \{cm^{ps}_{\{p\}}\}, r_{\{f,a,t\}}\right)\} \quad \textbf{(5)}$$

$\underline{\textbf{where}}$
$p \in PR$
id : problem solution identification
$ic^{ps}_{\{p\}}$: initial solution cost
$fc^{ps}_{\{p\}}$: final solution cost
$\{ft^{ps}_{\{p\}}\} \subset FT_{\{p\}}$
$\{cm^{ps}_{\{p\}}\} \subset CM$

A representation of a candidate solution for the example problem above could be:

$ps_{\{001\}}$=(1200,959,{flights},{crewmembers},$r_{\{tp438,lis,0715\}}$).

For each and all candidate solutions agent *ResolutionManager* calculates the AOCC Performance ρ, using Equation 1. The candidate solution with the minimum value of ρ will be the one that it is sent to the *Supervisor* for approval.

Solution generation and evaluation: The generation of a new solution, by the specialist agents *HillClimb* and *SimmAnneal*, is made by finding a successor that distances itself to the current solution by one unit, that is, the successor is obtained by one, and only one, of the following operations:

- Swap two aircrafts between flights that belong to the flights that are in the time window of the problem.
- Swap two crewmembers between flights that belong to the flights that are in the time window of the problem.
- Swap an aircraft that belongs to the flights that are in the time windows of the problem with an aircraft that that is not being used.
- Swap a crewmember of a flight that belongs to the flights that are in the time window of the problem with a crewmember that isn't on duty, but is on standby.

When choosing the first element (crewmember or aircraft) to swap, there are two possibilities:

- Choose randomly
- Choose an element that is delayed.

This choice is made based on the probability of choosing an element that is late, which was given a value of 0.9, so that the algorithms can proceed faster to good solutions (exchanges are highly penalized, so choosing an element that is not late probably won't reduce the cost, as a possible saving by choosing a less costly element probably won't compensate the penalization associated with the exchange). If the decision is to exchange an element that is delayed, the list of flights will be examined and the first delayed element is chosen. If the decision is to choose randomly, then a random flight is picked, and a crewmember or the aircraft is chosen, depending on the probability of choosing a crewmember, which was given a value of 0.85. When choosing the second element that is going to swap with the first, there are two possibilities:

- Swap between elements of flights.
- Swap between an element of a flight and an element that isn't on duty.

This choice is made based on the probability of choosing a swap between elements of flights, which was given a value of 0.5. The evaluation of the solution is done by an objective function that measures four types of costs:

- The costs with crewmembers. Those costs take into consideration the amount that has to be paid to the crewmember (depends on the duration of the flight), and the base of the crewmember (for instance, assign a crewmember from Oporto to a flight departing from Lisbon has an associated cost that would not be present if the crewmember's base was Lisbon).
- The costs with aircrafts. Those costs take into consideration the amount that has to be spent on the aircraft (depends on the duration of the flight), and the base where the flight actually is.
- The penalization for exchanging elements.
- The penalization for delayed elements. The cost associated with this aspect is the highest, because the goal is to have no delayed elements.

These types of costs are taken into account in Equation 6:

$$tc = cmc + amc + exW * numE + dlW * numD \tag{6}$$

Where

$$cmc = \sum_{i=1}^{|CM|} (c_i * bcf)/numCm \tag{7}$$

$$\begin{aligned}
&where\\
&i \in CM; CM = \{all\ crewmembers\ in\ flight\}\\
&1 < bcf \leq 2 : base\ crew\ factor\\
&numCm : number\ of\ crew\ members\ in\ CM
\end{aligned}$$

$$amc = \sum_{j=1}^{|AC|} (ac_j * baf)/numAc \tag{8}$$

$$\begin{aligned}
&where\\
&j \in AC; AC = \{aircraft\ same\ fleet\}\\
&1 < baf \leq 2 : base\ aircraft\ factor\\
&numAc : number\ of\ aircrafts\ in\ AC
\end{aligned}$$

exW was given a value of 1000, and *dlW* a value of 20000.

Regarding the agent that implements a *Simmulated Annealing* algorithm [10], there is a probability that a new solution is selected even if the cost is not smaller than the previous one. Our agent has used the following values for calculating this probability: α=0.8, T=10 and T updated every N iterations (N=2).

5 Scenario and Experiments

Scenario: To evaluate our approach we have setup the same scenario used by the authors in [6] that include 3 operational bases (A, B and C). Each base, corresponding to a different airport, includes their crewmembers each one with a specific roster. Airport B is the Hub of the airline. In this small experiment it is included 15 flights, 36 crewmembers and 4 aircrafts. After setting-up the scenario we found the solutions for each crew event using our system (running only once). After that, the AOCC performance for each method was calculated according to Equation 1 and considering a one month period. As a final step, the solutions found by our system were presented to AOCC users to be validated regarding feasibility and correctness.

Results: Table 1 presents the results that compare our method (method B) with the one used by Castro and Oliveira [6]. From the results obtained we can see that method B increased 1.38 times the performance of the AOCC. As we stated in section 3, if the AOCC performance tends to zero it means that the airline is operating (in terms of flight departure and arrivals times) more according to the airline schedule and, because of the relation that exist between airline schedule and airport peaks (as we explained in section 3), it means that the airline contributes also to a better passenger flow at the airports. From Table 1 we see that performance of AOCC in our method (B) is closer to zero than previous method.

Regarding the performance in each airport, our approach improved the performance of the AOCC in airport A and B by 2 and 1.75 times, respectively. For airport C the performance is the same of previous method. Another important result is the total costs. Our method is 23% less expensive than the previous one.

Table 1. Comparison of the results

		Method A		**Method B**		**A/B**
	Flights	**ρ**		**ρ**		
Global		180		130		1.38
- Airport A	3	40		20		2.00
(0-13h)	1	30		20		
(13-20h)	2	10		0		
(20-24h)	0			0		
- Airport B	7	70		40		1.75
(0-13h)	2	10		10		
(13-20h)	4	50		30		
(20-24h)	1	10		0		
- Airport C	5	70		70		0.00
(0-13h)	3	20		15		
(13-20h)	2	50		55		
(20-24h)	0	0		0		
Total costs		11628		8912		-23%

6 Discussion and Conclusions

In this paper we proposed a way of evaluating the performance of the AOCC that takes into consideration the relation that exists between the airport peaks (airport capacity) and the airline schedule.

We have implemented a MAS that represents the roles in the AOCC and that solves the unexpected problems that usually happens on airline operations. Our MAS is able to take decisions taking into consideration the AOCC performance as well as the airline operational costs. Preliminary results show that it is possible to contribute to minimize the airports peaks without increasing the airline operational costs. However, due to the probabilistic nature of the simulated annealing algorithm and due to the fact that we have run our system only once to get the results, we cannot generalize the results presented here. Another conclusion that we are able to take from the results in Table 1, is that it would be important to collect the reason that caused the flight/crew delay, i.e., due to weather conditions, ATC and/or airport restrictions, aircraft malfunction, etc. This information would help to understand some of the results. For example, it could help to understand why the performance on airport C is the same as the one obtained by the previous method.

We also point out that these results, *per se*, do not mean that we are able to solve all the airport peak problems in a specific airport. The airport peak problems are the result of the passenger flow that is generated by several airline schedules that operate at the airport. It would be necessary that all airlines implemented a similar system to reach to such a conclusion. However, in the airports where an airline has a hub-and-spoke network, the majority of the passenger flow is generated by a single airline company. In those cases, our approach could contribute significantly to minimize the airport peak problems.

Finally, our MAS is an integrated system that automates much of the disruption management process, from the monitoring of the operation of the company in its several bases, to the detection of events and the resolution of the problems encountered. Additionally, our MAS is oriented to the future: its distributed nature and the fact that it is based on agents that are specialists in solving problems easily allows the insertion into the system of new agents that solve new kinds of problems that were identified in the meantime, or that resolve the current types of problems using different methods. It is thus a truly scalable solution, prepared to sustain the growth of the airline company. Although the goals have been achieved, it is important to consider a number of improvements that could be made on future developments, and that could enrich it. In terms of the algorithms used to solve the problems, other meta-heuristics can be implemented, as well as methods based in the area of operational research. The fact that this is a distributed system means that there is no theoretical limit to the number of agents that try to solve, at the same time, the same problem. It is also important to collect more data and run the system several times to get more conclusive and generic results.

Acknowledgments. 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. Abdelgahny, A., Ekollu, G., Narisimhan, R., and Abdelgahny, K.: A Proactive Crew Recovery Decision Support Tool for Commercial Airlines during Irregular Operations. Annals of Operations Research, 127:309-331 (2004).
2. Ashford, N., Stanton, H., and Moore, C.: Airport Operations. Second Edition, McGraw-Hill, ISBN: 978-0-07-003077-0 (1997).
3. Bellifemine, F., Caire, G., Trucco, T., and Rimassa, G.: JADE Programmer's Guide. JADE 3.3 TILab S.p.A. (2004).
4. Bratu, S., and Barhhart, C.: Flight Operations Recovery: New Approaches Considering Passenger Recovery. Journal of Scheduling 9(3):279-298 (2006).
5. Clausen, J., Larsen, A., and 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).
6. Castro, A.J.M., and 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.
7. Danesi, A.: Measuring airline hub timetable coordination and connectivity: definition of a new index and application to a sample of European hubs. European Transport 34, pp. 54-74 (2006).
8. Grosche, T.: Computational Intelligence in Integrated Airline Scheduling. Studies in Computational Intelligence 173, Springer-Verlag Berlin Heidelberg, ISBN: 978-3-540-89886-3 (2009).
9. Jonker, G., Meyer, J.J., and Dignum, F.: Towards a Market Mechanism for Airport Traffic Control, 12th Portuguese Conference on Artificial Intelligence (EPIA 2005), Covilhã, Portugal (2005).
10. Kirkpatrick, S., Gelatt, C.D., and Vecchi, M.P.: Optimization by Simulated Annealing. Science, New Series, Vol. 220, No. 4598, pp. 671-680 (May 13, 1983).
11. Kohl, N., Larsen, A., Larsen, J., Ross, A., and Tiourline, S.: Airline Disruption Management – Perspectives, Experiences and Outlook. Technical Report, CRTR-0407, Carmen Research (2004).
12. Liu, T., Jeng, C., and Chang, Y.: Disruption Management of an Inequality-Based Multi-Fleet Airline Schedule by a Multi-Objective Genetic Algorithm, Transportation Planning and Technology, Vol. 31, No. 6, pp. 613-639 (2008).
13. Ouelhadj, D.: A Multi-Agent System for the Integrated Dynamic Scheduling of Steel Production, Ph.D. dissertation, The University of Nottingham, School of Computer Science and Information Technology, England (2003).
14. Tumer, K., and Agogino, A.: Distributed Agent-Based Air Traffic Flow Management, 6th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2007), Honolulu, Hawaii (2007).
15. Wolfe, S. R., Jarvis, P. A., Enomoto, F. Y., and 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 (2007).
16. Yu, G., and Qi, X.: Disruption Management: Framework, Models and Applications. World Scientific Publishing Co. Pte. Ltd, ISBN: 981-256-017-3 (2004)