Short-term Allocation of Time Windows to Flights through a Distributed Market-based Mechanism

Andrea RANIERI
Dipartimento di Elettrotecnica, Elettronica e Informatica, University of Trieste, Italy

Advisors
Lorenzo CASTELLI
Dipartimento di Elettrotecnica, Elettronica e Informatica, University of Trieste, Italy

Raffaele PESENTI
Dipartimento di Matematica Applicata, University Ca' Foscari of Venice, Italy

Abstract

The allocation of ground delays to flights is a tactical tool commonly employed to control the flow of air traffic and to ensure that available capacity of system resources is respected. Under the current European system these delays are assigned by the network manager according to a First-Planned-First-Served principle, without taking into account the individual cost of delay suffered by different flights. We formalize a decentralized, Individual Rational and Budget Balanced market-based mechanism for the collaborative assignment of delays among flights, based on the Lagrangian relaxation of the central assignment problem. It allows flights to pay for reducing their delays or to get compensations if they accept an increased delay with respect to the First-Planned-First-Served rule, in the general case of multiple capacity constrained resources and without requiring the disclosure of Airlines' private information. Some computational experience based on a real case instance is reported.

1 Introduction

The Air Traffic Flow Management (ATFM) is the service responsible to regulate flights in order to ensure that the available capacity of the system is efficiently used and never exceeded, to enable a safe, ordered and expeditious flow of traffic. In Europe this concept has recently evolved to the wider Air Traffic Flow and Capacity Management, to underline its role in managing the balance between demand and capacity by coordinating all actors involved. In Europe this service is provided by EUROCONTROL Central Flow Management Unit (CFMU), which acts as network manager by comparing the projected air traffic demand resulting from filed flight plans with the available capacity of system resources (i.e. airports and air traffic control sectors of airspace) as declared by the national Air Navigation Service Providers (ANSP) through the activation of ATFM regulations (or simply regulations). Each regulation specifies the affected geographical area, the maximal rate of flights that the area can accept and the period of activation. The maximal rate establishes the limit on the number of flights that can enter the specific area per period of time.

For each regulated resource, the number of slots available is defined by its maximal hourly rate of acceptance multiplied by the number of hours of activation. This set of slots constitutes the Slot Allocation List (SAL) for the given capacity-constrained resource. Each flight planned to cross the regulation is given a provisional slot based on its Estimated Time Over (ETO) the restricted location. Each slot has capacity of 1 flight, hence if the slot corresponding to the original ETO of the flight has not been already assigned to another flight, it is available and the flight will not be delayed. Otherwise the slot is assigned to the flight with the lowest ETO, according to a First-Planned-First-Served (FPFS) principle, and the flight with the highest ETO receives a later slot, corresponding to a new Calculated Time Over (CTO) which will be greater than the original ETO. In the case a flight is subject to several ATFM regulations, the highest delay, caused by the slot in the most penalizing regulation crossed, is forced also in the other ones.

This process ultimately results in the imposition of a Calculated Take-Off Time (CTOT) for flights affected, which is calculated backward from the CTO on the most penalizing regulation, which implies that aircraft must take off during the time range between CTOT - 5 minutes and CTOT + 10 minutes. This is the only effective ATFM slot that the flight must respect.

The ATFM slot allocation measure is based on the universally accepted principle that delays on the ground are safer and less costly than those in the air. Any forecast delay somewhere in the system is thus anticipated at the departure airport prior to the take-off and the traffic is controlled in a safe and simple manner.

1.1 Context of the study

The current ATFM procedures are going to be revised in the next future according to the new target concept developed by the Single European Sky ATM Research (SESAR) Programme, which identifies the technological steps and the modernization priorities necessary for implementing the new European Air Traffic Management (ATM) through a reorganization of systems, roles and procedures of the different air traffic stakeholders SESAR Joint Undertaking, 2007. One of the key concepts in SESAR is the notion of Business Trajectory, owned by the airspace user, which evolves through several phases and represents the best trade-off between the constraints imposed by infrastructural and environmental restrictions and the users internal business objectives. Business Trajectories will be expressed in 4 dimensions (latitude, longitude, flight-level and time).

A possible mechanism to formalize the Business Trajectory comes from the concept of Contract of Objectives (CoO) which has been developed by the Contract-based Air Transportation System (CATS) research project (www.cats-fp6.aero). The CoO is a formal and collaborative commitment of ATM actors, i.e., airspace users, airports and Air Navigation Service Providers (ANSPs), to the conduction of each flight. It establishes a sequence of spatial and temporal constraints which constitute milestones to be met during the flight execution. These 4D intervals are called Target Windows and are agreed upon all involved actors for specific transfer of responsibility areas (e.g. between en-route control sectors). They represent the commitment to deliver a particular aircraft inside such temporal and spatial intervals. In other words, the proposed CoO consists of a collection of Target Windows defined at each area where responsibility between actors is transferred and the Business Trajectory should then go through these different target windows.

When an unexpected imbalance between capacity and demand is detected on a short notice, SESAR states that airspace users (i.e., airlines in our context) will be offered the possibility to indicate to the Network Manager a priority order for flights affected by delays under the so called User Driven Prioritisation Process (UDPP).

Distributed decision making is also a key characteristic of the U.S. NextGen ATM Concept of Operations (CONOPS) for 2025 (Joint Planning and Development Office, 2007), which depicts a system where airspace users (represented by flight crew or the Flight Operations Center) have an increased level of decision making enabled by the access to a rich, timely and accurate information exchange that at the same time guarantees privacy and security. This implies that interacting stakeholders have the possibility to to better understand what are the prevailing constraints, the short-term and long-term effect of decisions and the interdependence among national, regional and local operations.

Since different airlines are in general competitors, market-based mechanisms seem natural ways to support dynamic, collaborative and effective decision making process in the allocation of capacity.

In this context, we propose a market-based negotiation scheme between flights and the Network Manager to allocate target windows in the context of European ATM system. The focus is on the 'critical' target windows which are those associated to the scarce air traffic resources, i.e., where the demand exceeds the available capacity. In fact, all the remaining target windows of each flight will be accordingly adjusted. For the sake of simplicity, in the remainder of the paper we only consider the temporal dimension of the critical target windows, thus in the following we simply refer to them as Time Windows (TWs).

Our work then relies on an extension of the current European ATFM system, under which just one departure ATFM slot may be allocated to a flight when the capacity does not meet the demand. Rather we assume that one TW is explicitly allocated to a flight for each congested resource crossed, in line with the 4-D trajectory management and the short-term capacity management principles proposed by SESAR. In particular, TWs are first allocated to flights following a First-Planned-First-Served rule at no cost, similarly to the current way of allocating ATFM slots. Then a market mechanism is initiated among flights coordinated by a Network Manager to negotiate available capacity, represented by a limited number of tradable TWs.

1.2 Literature review

Following the seminal work of Odoni, 1987, a number of researchers have focused their activity on the development of optimization models and algorithms for the assignment of ground delays as a short-term measure to regulate traffic flows. The problem of assigning ground delays to a set of flights in order to minimize an aggregated cost function, given airport capacity constraints, is known as Ground Holding Problem (GHP).

In its basic version the GHP assumes that only one airport in the system is subject to capacity constraints which are imposed only on arrival flights. This problem is referred to as the Single Airport Ground Holding problem (SAGHP) and has been formulated for different cases (see e.g. Andreatta and Romanin-Jacur, 1987; Terrab and Odoni, 1993; Andreatta and Romanin-Jacur, 1987; Richetta and Odoni, 1993; Hoffman and Ball, 2000)

In the case a network of airports is considered and propagation of delay can occur between successive flights, the problem of determining individual flight delays to minimize the total delay cost is referred to as the Multi Airport Ground Holding problem (MAGHP). Under this general case the problem is typically modeled through an integer program that increases computational burden with respect to the SAGHP (see e.g.Vranas et al., 1994a; Vranas et al., 1994b)

A further extension of the MAGHP is the one that also includes constraints on the capacities of en-route sectors of airspace and determines optimal speed adjustments of aircraft, besides their release time into the network. This is known as Air Traffic Flow Management problem (TFMP) and has been treated in Bertsimas and Stock Patterson, 1998; Bertsimas and Stock Patterson, 2000; Bertsimas et al., 2008; Alonso et al., 2000 and recently in Lulli and Odoni, 2007 for the European case.

The common characteristic of all these models is the the presence of a unique central decision maker, the ATFM authority, which is in charge of assigning individual delays to flights in order to minimize a global objective function, obtained by aggregating the direct operating costs caused to all regulated flights by ATFM restrictions. This approach is consistent with the current mode of operations of the European ATFM system, where the CFMU centrally calculates and impose ground holdings to flights according to a FPFS heuristic. This criteria for delay assignment does not consider the cost caused to users, which can be correctly estimated only by individual airlines depending on their internal priorities and business models.

This principle motivated the Federal Aviation Administration (FAA) to undertake in the '90s the Collaborative Decision Making (CDM) program to partially decentralize decision making regarding the assignment of ground delays to flights in the U.S. National Airspace System (Wambsganss, 2001). The Ground Delay Program represents today the CDM implementation in the U.S. system, which allows airlines to exert active control on their aircraft, by providing incentives to share up-to-date information and to cooperate in determining resource allocation. The first procedure used under this scheme to allocate ground delays is the Ration-by-schedule (RBS) algorithm, which produces an initial assignment of slots to flights following a FPFS criteria. Each Airline can then decide to modify this slot-to-flight assignment for its own flights, through cancellation and substitution procedures. In practice these processes transform the slot-to-flight first assignment obtained by RBS allocation into a slot-to-airline assignment, since Airlines are free to swap slots between pairs of flights or to cancel a less profitable flight and assign its slot to another flight, thus reducing its delay. After this intra-airline exchange, the Compression algorithm, which is carried out by the FAA, is run to maximize slot utilization by performing inter-airline slot exchanges, in order to ensure that no slot goes unused.

Vossen and Ball, 2006 demonstrate that the inter-airline slot exchange procedure implemented through the Compression algorithm can be interpreted as a mediated bartering, in which the FAA acts as a 'broker', matching offers proposed by the airlines. The idea behind the Compression algorithm is to reward airlines for slots they release, thus encouraging airlines to report cancellations.

1.3 Contribution of this work

The specific nature of the European ATFM problem, where several capacity constraints on airspace sectors and airports can be activated simultaneously, makes impossible a direct application of U.S. CDM procedures. In Europe two flights can swap their ATFM slots only if they are operated by the same Airline, they are subject to the same most penalizing regulation and only in the case CFMU confirms that this has no negative network effect on the system (EUROCONTROL Central Flow Management Unit, 2009).

This implies that more complex exchange procedures are required under the European system to provide users an incentive to participate while at the same time ensuring the respect of capacity and the fairness of the final allocation. We developed a practical market-based mechanism which responds exactly to these requirements.

Some papers have already proposed market mechanisms to assign limited resources in the context of air traffic. They are usually based on combinatorial auctions for the long-term strategic allocation of airport slots in the US (Rassenti et al., 1982; Le et al., 2004; Ball et al., 2005). However, an auction mechanism for air traffic resources is not in general appropriate, since some airlines should pay to get the same resources that they currently receive at no cost. To overcome this limitation, our mechanism allows participants to trade on the base of their individual cost of delay the TWs originally assigned by means of the FPFS rule, thus introducing a combinatorial exchange (Parkes et al., 2001).

A market mechanism is *Individual Rational* (IR) if it guarantees that every participant will have a payoff equal or higher than the payoff obtained without taking part to it. A mechanism is *Budget Balanced* (BB) if it does not require an external subsidization to run properly. In particular, it is strongly BB, if the payments collected by all actors sum up to zero (i.e. it just redistributes money among participants) or weakly BB if it can produce a monetary surplus that can be collected by the auctioneer.

Our market mechanism is IR and weakly BB. Every airline as well as the Network Manager (i.e. the auctioneer) are guaranteed to have a non-negative payoff from the participation. Furthermore, it is distributed meaning that flights do not need to communicate to the Network Manager their cost of the delay, which is a confidential information. Finally, our proposed mechanism falls within the Collaborative Decision Making (CDM) framework, as airlines have to actively interact with the Network Management.

This paper unfolds as follows. Section 2 illustrates the problem of allocating TWs to flights. Section 3 formalizes this problem as a combinatorial exchange, introduces TW prices and discusses potentialities and limitations of the centralized approach. Section 4 proposes a distributed market mechanism to determine the optimal TW exchange and the associated TW prices. An heuristic algorithm is presented along with some computational results based on a real case instance. Section 5 summarizes conclusions.

2 The central allocation problem

Let $\mathcal{F} = \{1, ..., F\}$ be a set of flights and $\mathcal{S} = \{1, ..., S\}$ a set of capacity constrained sectors and airports. Each flight $f \in \mathcal{F}$ is expected to cross a sequence of elements $S_f \subseteq \mathcal{S}$ according to its flight plan, hence it will need to be assigned a TW for each $s \in S_f$. A regulated resource $s \in \mathcal{S}$, with capacity limited to K_s entries per hour from st_time_s to end_time_s, has an associated TW Allocation List $L_s = [1, ..., N_s]$. Each TW $j = [I_j, U_j] \in L_s$ has capacity of one flight where:

$$\begin{aligned} N_s &= \left\lfloor \frac{end\_time_s - st\_time_s}{\frac{60}{K_s}} \right\rfloor \\ I_j &= \left\lfloor st\_time_s + (j-1) \cdot \frac{60}{K} \right\rfloor \text{ with } j \in \{2, ..., N_s\} \\ U_j &= I_{j+1} - 1 \text{ with } j \in \{1, .., N_s - 1\}, \end{aligned}$$

and $I_1 = st\_time_s$, $U_N = end\_time_s$. We assume that the Business Trajectory (i.e. the 4D extension of the current Flight Plan) also indicates an estimated time of entry into each element $s \in S_f$ traversed by flight f, i.e. E_f^s. Then f is allocated a list of TWs $q_f = [k_1, .., k_{|S_f|}]$, where k_i is the TW assigned on the i^{th} element of S_f and can not be earlier than E_f^i since flights cannot be anticipated, i.e. $E_f^i \leq U_{k_i}$ for all $k_i \in q_f$. Additionally whenever $|S_f| > 1$, we assume that the flying time between pairs of consecutive elements, (i, j) with $i, j \in S_f$ and $j = i + 1$, is fixed. An assignment q_f will cause a positive delay to flight f if and only if $E_f^i < I_{k_i}$ for some $k_i \in q_f$ and the amount of delay will be:

$$d_f^{q_f} = \begin{cases} \max\limits_{i \in S_f}(I_{k_i} - E_f^i) & \text{if } E_f^i \leq I_{k_i} \quad \forall i \in S_f \\ 0 & \text{otherwise} \end{cases} \tag{1}$$

Hence each assignment of TWs q_f to a flight f implies a nonnegative cost of delay $C(f, q_f) \geq 0$, which depends on several factors, e.g, the type of aircraft, the number of connecting passengers, the time of day, the amount of delay (Cook et al., 2004). Possibly, airlines only can correctly estimate the real costs of the delays of their flights, hence they have to be queried in a mechanism that seeks to minimize the cost of a TW allocation.

An assignment q_f is feasible for flight f if and only if (i) it contains one TW for each $i \in S_f$ and each pair (i, j) of consecutive TWs is connected by the fixed flying time $E_f^j - E_f^i$, (ii) it assigns a nonnegative delay to f and (iii) the delay it assigns is bounded, i.e. $d_f^q < MaxDel_f$ where $MaxDel_f$ is a fixed parameter for each flight beyond which the flight prefers to be canceled.

Let us indicate with Q_f the set of all assignments that are feasible for flight f, then the optimal assignment of TWs to flights can be formulated as the following 0-1 IP program:

$$Z_{IP} = \min \sum_{f \in \mathcal{F}} \sum_{q \in Q_f} C(f, q) x(f, q) \tag{2a}$$

subject to

$$\sum_{f \in \mathcal{F}} \sum_{q \in Q_f : q \ni k} x(f, q) \leq 1 \quad \forall s \in \mathcal{S}, k \in L_s \tag{2b}$$

$$\sum_{q \in Q_f} x(f, q) = 1 \quad \forall f \in \mathcal{F} \tag{2c}$$

$$x(f, q) \in \{0, 1\} \quad \forall f \in \mathcal{F}, q \in Q_f. \tag{2d}$$

The objective is to find the minimal cost assignment such that each TW k is allocated at most to one flight (constraints (2b)) and each flight is assigned one bundle of TWs among its feasible ones (constraints (2c)). The binary decision variable $x(f, q)$ will be one if flight f receives the q^{th} bundle in its request set Q_f, zero otherwise. A feasible solution will exist if and only if there are enough TWs and requests, such that the assigned bundles are pairwise disjoint, i.e. they do not share any TW. To guarantee the existence of a feasible solution we assume that each regulated resource $s \in \mathcal{S}$ has an infinite capacity after the termination of its regulation and that each flight has a request $q_w \in Q_f$ that includes only TWs after the termination of each regulation traversed. The cost $C(f, q_w)$ associated to this bundle will be equal to either the cost of delay caused by such a bundle or to the cost of cancellation in the case this delay exceeds $MaxDel_f$.

Problem (2) can be reduced to the following standard Weighted Set Packing Problem:

$$Z_{SPP} = \max \sum_{f \in \mathcal{F}} \sum_{q \in Q'_f} (\overline{C} - C(f, q)) x(f, q) \tag{3a}$$

subject to

$$\sum_{f\in\mathcal{F}}\sum_{q\in Q'_f:q\ni k} x(f,q) \leq 1 \quad \forall s\in\mathcal{S}', k\in L_s \tag{3b}$$

$$x(f,q)\in\{0,1\} \quad \forall f\in\mathcal{F}, q\in Q'_f \tag{3c}$$

where $\overline{C} > |F|\cdot\max_{q\in\cup Q_f} C(f,q)$ is chosen such that the Lagrangian problem obtained from problem (2) by dualizing constraints (2c) gives the same assignment as the original problem. This prevent some flight $f\in\mathcal{F}$ to be excluded from the assignment since the cost induced by its exclusion is not compensated by any other possible assignment. The request sets Q'_f are obtained by appending one dummy TW k_d^f on an dummy resource d to each bundle q, such that each flight $f\in\mathcal{F}$ has one different dummy TW associated and this TW is included in all its requests $q\in Q'_f$. This prevent multiple assignments to the same flight. Then the new set of resources is $\mathcal{S}'=\mathcal{S}\cup d$. Problem (3) can be interpreted as a sealed-bid type of combinatorial auction, in which all flights communicate their costs $C(f,q)$ to the auctioneer which centrally solves problem (3) according to this information, to find an assignment $X^*=\{q_1^*,...,q_F^*\}$ that maximizes the social welfare. Successively the auctioneer computes a vector of nonnegative prices for the bundles $P=\{p(q_1^*),...,p(q_F^*)\}$ and charges each flight f the price $p(q_f^*)$ of the bundle assigned. As in standard auction theory, we assume that each flight has a quasi-linear utility $u(f,q_f^*) = -C(f,q_f^*) - p(q_f^*)$. Then the auction is not IR, since $u(f,q_f^*)\leq 0$ for all $f\in\mathcal{F}$.

To force IR we rather propose to calculate the standard $FPFS$ assignment $A=\{a_1,...,a_F\}$ as in the current system and to consider it as the initial endowment for each flight. This is a feasible solution for problem (3) but not necessarily optimal, then we propose to implement the exchange between A and X^* such that every flight f for which $a_f\neq q_f^*$ pays the price $p(q_f^*)$ for its optimal bundle but also receives the payment $p(a_f)$ for the released bundle a_f. Then each flight is a potential buyer and seller of TWs and its utility after the exchange will be $u(f,e^*) = [C(f,a_f) - C(f,q_f^*)] - [p(q_f^*) - p(a_f)] = -C(f,e_f^*) - p(e_f^*)$

3 Pricing the exchange

In this section, we formulate the optimal exchange problem as the following binary IP model:

$$Z_{IP-E} = \max\sum_{f\in\mathcal{F}}\sum_{q\in Q_f} V(f,q)x(f,q) \tag{4a}$$

subject to

$$\sum_{f \in \mathcal{F}} \sum_{q \in Q_f : q \ni k} x(f,q) \leq 1 \quad \forall s \in \mathcal{S}, k \in L_s \tag{4b}$$

$$\sum_{q \in Q_f} x(f,q) = 1 \quad \forall f \in \mathcal{F} \tag{4c}$$

$$x(f,q) \in \{0,1\} \quad \forall f \in \mathcal{F}, q \in Q_f \tag{4d}$$

where $V(f,q) = [C(f,a_f) - C(f,q)]$ is the value obtained by flight f by exchanging bundle a_f with bundle q and will be positive if the delay caused by bundle q is lower than delay caused by a_f, negative otherwise. We want to find prices such that for each flight f, $u(f,e^*) \geq 0$ in order to guarantee IR and $\sum_{f \in \mathcal{F}} p(e^*_f) \geq 0$ in order to guarantee (weak) BB.

We define Z_{LP-E} as the objective value of the linear relaxation of problem (4), whose dual problem is:

$$Z_{DLP-E} = \min \sum_{f \in \mathcal{F}} u_f + \sum_{s \in \mathcal{S}} \sum_{k \in L_s} p(k) \tag{5a}$$

subject to

$$u_f + \sum_{s \in \mathcal{S}} \sum_{k \in (L_s \cap q)} p(k) \geq V(f,q) \quad \forall f \in \mathcal{F}, q \in Q_f \tag{5b}$$

$$p(k) \geq 0 \quad \forall s \in \mathcal{S}, k \in L_s \tag{5c}$$

A competitive equilibrium (Bikhchandani and Mamer, 1997) is a situation in which there is a vector of prices p, usually referred to as market clearing or supporting prices and a feasible allocation $(b_1, ..., b_F)$ such that:

$$\begin{aligned} V(f,b_f) - \sum_{k \in b_f} p(k) &> V(f,q_f) - \sum_{k \in q_f} p(k) \quad \forall f \in \mathcal{F}, q_f \in Q_f \\ \sum_{s \in \mathcal{S}} \sum_{k \in L_s} p(k) &= \sum_{f \in \mathcal{F}} p(b_f) \end{aligned}$$

The formulation of problem (5) suggests the interpretation of dual variables as supporting prices $p(k)$ for TWs. Let us assume for the moment a linear structure of prices, i.e. for of each bundle of TWs $q \in Q_f$ its price will be $p(q) = \sum_{k \in q} p(k)$.

The complementary slackness conditions for the primal-dual pair are:

$$x^*(f,q) > 0 \Rightarrow u^*_f + \sum_{s \in \mathcal{S}} \sum_{k \in (L_s \cap q)} p^*_k = V(f,q) \tag{6a}$$

$$u^*_f + \sum_{s \in \mathcal{S}} \sum_{k \in (L_s \cap q)} p(k) > V(f,q) \Rightarrow x^*(f,q) = 0 \tag{6b}$$

$$\sum_{f\in\mathcal{F}}\sum_{q\in Q_f:q\ni k} x(f,q) < 1 \quad \Rightarrow p(k) = 0 \tag{7a}$$

$$p(k) > 0 \Rightarrow \sum_{f\in\mathcal{F}}\sum_{q\in Q_f:q\ni k} x(f,q) = 1 \tag{7b}$$

Conditions (6) and (7) imply that a linear pricing rule based on individual prices $p(k)$ implements an individual rational mechanism for our TW-exchange problem, whenever the LP relaxation of problem (4) gives an integer solution. In fact each flight f will weakly prefer to sell the bundle of TWs $a_f \in Q_f$ received under FPFS allocation if it is not optimal for problem (2a) and to buy the bundle of TWs $q_f^* \in Q_f$ optimal for problem (4), because its utility increases:

$$C(f, a_f) + \sum_{k\in a_f} p(k) - C(f, q_f^*) - \sum_{k\in q_f^*} p(k) \geq 0 \tag{8}$$

Such a payment rule is (weakly) BB, since it can produce a monetary surplus for the auctioneer. In fact all TWs unassigned under FPFS which are assigned in the optimal allocation have a price $p(k) > 0$ according to (7b). This price has to be paid by the receiving flight but is not due to anyone since that TW was unallocated under FPFS. On the contrary, according to (7a) any TW that is not assigned in the optimal allocation has a price $p(k) = 0$, meaning that the flight giving up a TW received under FPFS which is not allocated by the market mechanism does not receive any compensation for it.

Unfortunately, supporting prices induced by conditions (6) and (7) exist only in the case that the integrality gap between problem (4) and its LP relaxation is null. This will always be verified in the case of gross-substitute valuation functions (Kelso and Crawford, 1982), i.e., if for every pair of price vectors $p' \geq p$ (component-wise comparison), we have that the optimal TW package demanded by a flight f at prices p' contains all the TWs in the optimal package demanded by f at prices p, whose price remained constant. The class of gross-substitutes is a largest set of valuation functions which contains unit-demand ones, i.e. when all flights demand bundles are composed by a single TW. In those situations problem (4) reduces to an assignment problem, thus implying that its linear relaxation gives integer feasible solutions, and dual variables define linear market-clearing prices. Whenever there are complementarities in valuations functions (i.e. in the case of multiple constrained resources), gross-substitutes property does not hold anymore.

An alternative approach to identify supporting prices is based on the Vickrey-Clarke-Groves (VCG) class of auction (Vickrey, 1961; Clarke, 1971; Groves, 1973). It is a central scheme in auction theory and mechanism design since it is the only general class of auction mechanisms which verifies the Incentive Compatibility property. A market mechanism is *Incentive Compatible* (IC) if the equilibrium strategy for participants is to report their preferences truthfully, i.e. if they can never increase their individual payoff by misrepresenting their cost functions and thus they have no incentive to untruthfully report these costs, independently from what is the information communicated by others.

Unfortunately when the VCG pricing rule is adopted in an exchange model, the resulting mechanism is IR and IC but not BB. In fact in some cases the Network Manager shall redistribute to participants a higher amount than the one collected from them, thus experiencing a loss. Moreover Myerson and Satterwhite, 1983 state that it is impossible to build an exchange mechanism that is at the same time IR, BB and IC . However several VCG-based payment rules that clear combinatorial exchanges while guaranteeing BB and IR and minimizing the distance from VCG prices have been proposed in Parkes et al., 2001. Experimental and theoretical analysis performed on several distance functions suggests that a "Threshold rule" has useful incentive properties and provides allocative efficiency higher than other rules, by removing easy opportunities for manipulation. In fact, the Threshold rule minimizes the maximal amount that a participant can increase its utility by misrepresenting its information.

4 A distributed market mechanism

The central allocation problem that determines the optimal exchange could result complicated to be adopted in practice, notably because (i) it requires the complete disclosure of airline private information regarding their cost and (ii) the computational burden for solving Problem (4) is entirely faced by the Network Manager which must solve one NP-hard problem for each allocation (plus other $|F|$ NP-hard problems in the case VCG prices are charged). In order to avoid this issues we propose in the following a distributed algorithm that exploits the decomposition properties of Problem (4). In fact, by dualizing constraints (2b) the corresponding Lagrangian formulation of problem (4) is:

$$ZLR_{LP-E}(\lambda) = \max \sum_{f \in \mathcal{F}} \sum_{q \in Q_f} V(f,q)x(f,q) + \\ + \sum_{s \in \mathcal{S}, k \in L_s} \lambda_k (1 - \sum_{f \in \mathcal{F}, q \in Q_f : q \ni k} x(f,q)) \tag{9a}$$

subject to

$$\sum_{q \in Q_f} x(f,q) = 1 \quad \forall f \in \mathcal{F} \tag{9b}$$

$$x(f,q) \geq 0 \quad \forall f \in \mathcal{F}, q \in Q_f \tag{9c}$$

Problem (9) is separable into F problems, one for each flight and can be solved locally by Aircraft Operators, according to problem (10), which is a linear problem and can thus be solved in polynomial time:

$$ZLR_{LP-E}(f,\lambda) = \max \sum_{q \in Q_f} V(f,q)x(f,q) + \\ + \sum_{s \in \mathcal{S}, k \in L_s} \lambda_k (1 - \sum_{q \in Q_f : q \ni k} x(f,q)) \tag{10a}$$

subject to

$$\sum_{q \in Q_f} x(f,q) = 1 \tag{10b}$$

$$x(f,q) \geq 0 \quad \forall q \in Q_f \tag{10c}$$

For each TW k at the iteration $t+1$, the price $\lambda(k)^{t+1}$ is centrally calculated according to the excess of demand at iteration t and then communicated to Aircraft Operators, which in turn modify the demand for k according to its price. The following algorithm can be employed to calculate prices

$$\lambda_k^{t+1} = \max(0, \lambda_k^t - S_r^t \cdot SG_s^t) \tag{11a}$$

$$SG_k^t = 1 - \sum_{f \in \mathcal{F}, q \in Q_f : q \ni k} x(f,q) \tag{11b}$$

where S_r^t is a positive stepsize chosen at iteration t and SG_k^t is a subgradient of $ZLR_{LP-E}(\lambda)$ at any λ for which x solves problem (9). Thus ideally the Network Manager seeks the prices λ that solve the following dual problem

$$ZLR_{LP-E} = \min_{\lambda} ZLR_{LP-E}(\lambda) \tag{12a}$$

subject to

$$\lambda \geq 0 \tag{12b}$$

Since $ZLR_{LP-E}(\lambda)$ is a convex, piecewise linear, non-differentiable function, this problem is typically solved through a subgradient algorithm. The resulting procedure iteratively alternates a central price-calculation phase (problem 11) with a local optimization one which finds the maximal value TW-exchange at current prices (problem 10). By appropriately choosing the stepsize S_r^t such that $S_r^t \to 0$ and $\sum_{i=1}^t S_r^t \to \infty$ for $t \to \infty$, the procedure converges to a solution which minimizes $ZLR_{LP-E}(\lambda)$ (Held et al., 1974). By duality theory it will follow that $Z_{IP-E} \leq Z_{LP-E} \leq ZLR_{LP-E}(\lambda)$, while $Z_{IP-E} = Z_{LP-E}$ in the case of null gap between the integer program and its linear relaxation (i.e. with gross-substitutability) and $Z_{LP-E} = ZLR_{LP-E} = \min_\lambda ZLR_{LP-E}(\lambda)$ when the subgradient algorithm converges to an optimal solution for problem (12).

However this exchange will be optimal for the original problem if and only if the gap between Z_{IP-E} and its linear relaxation is null, a condition which can only be guaranteed in the case of gross-substitute valuations, for example when all participants compete for TWs on a unique resource. Furthermore, even in the case of gross-substitutability, there is no guarantee of convergence in a finite number of steps. By stopping the procedure when a feasible exchange for the original problem (4) is demanded at current prices, the optimality of the solution will be verified if and only if $\sum_{s\in\mathcal{S},k\in L_s}\lambda_k(1-\sum_{f\in\mathcal{F},q\in Q_f:q\ni k}x(f,q)) = 0$ (Larsson et al., 1999). This condition is automatically verified in the case that all bundles are singletons (i.e. the single constrained resource problem) due to the complementary slackness condition, but this is no longer true in the general case. Hence if the procedure stops prematurely it will be $ZLR_{LP-E} > Z_{IP-E}$ and some lagrangian multipliers $\lambda_k \geq 0$ will be higher than minimal prices, but the correspondent exchange will still guarantee IR and weak BB. We propose in the following section an heuristic algorithm to implement the distributed exchange that exploits some of these properties to realize a practical mechanism for TW exchanges.

4.1 A heuristic approach

In order to implement a distributed market mechanism that achieves, in a reasonable amount of time, an exchange which is IR and BB we propose the following heuristic. A formula for S_r^t which has been proven effective in practice is:

$$S_r^t = \frac{\mu_t(ZLR_{LP-E}(\lambda^t) - Z^*_{IP-E})}{\sum_{s\in\mathcal{S},k\in L_s}(1-\sum_{f\in\mathcal{F},q\in Q_f:q\ni k}x^t(f,q))^2} \tag{13}$$

where $0 < \mu_t \leq 2$, x^t are the solutions to problem (10) at iteration t according to the vector of TW prices λ^t. Usually the scalar μ_t is taken at its higher values during first iterations and halved whenever $ZLR_{LP-E}(\lambda^k)$ has failed to decrease in a specified number of iterations (Fisher, 1985). In our case the Network Manager does not know the exchange values $V(f,q)$ and thus cannot calculate neither $ZLR_{LP-E}(\lambda^k)$ nor Z^*_{IP-E}. We then modify formula (13) in the following sense:

$$S_r^t = \frac{\mu_t(UB\_Z^* - ZLB_{IP-E})}{\sum_{s\in\mathcal{S},k\in L_s}(1-\sum_{f\in\mathcal{F},q\in Q_f:q\ni k}x^t(f,q))^2} \tag{14}$$

where UB_Z^* is an upper bound on the optimal value of the exchange which is held constant and ZLB_{IP-E} is a lower bound on the optimal value of the exchange for each instance, which is dynamically adjusted through the course of the distributed mechanism. At iteration t the Network Manager can in fact calculate for each bundle j demanded by flight f at current prices λ^t, a lower bound on the exact value $V(f,j)$ for the exchange:

$$LB(f,j) = \sum_{s\in\mathcal{S},k\in L_s:j\ni k}\lambda_k^t - \sum_{s\in\mathcal{S},k\in L_s:a_f\ni k}\lambda_k^t \tag{15}$$

For all $f\in\mathcal{F}$ and $q\in Q_f$, lower bound can be initialized to $LB(f,q)=0$ if $d_f^q \leq d_f^{a_f}$ and $LB(f,q) = -\infty$ if $d_f^q > d_f^{a_f}$, since we assume non-negative costs of delay. This implies that each flight would exchange its FPFS assigned bundle

a_f with q for a cost greater or equal to zero whenever q causes a shorter delay then a_f or for a negative cost (a taking) whenever q represents a longer delay then a_f. The amounts of delay can be easily computed by the Network Manager on the basis of the last Business Trajectory shared by the Aircraft Operator. At iteration t the Network Manager will calculate the $LB(f,j)$ value according to formula (15) and will store it in memory if it is higher than the previously calculated one. Also it is possible to update with this same value the $LB(f,b)$ for all the bundles $b \in Q_f$ such that $d_f^b < d_f^j$, since the value of an exchange is a decreasing function of delay and $V(f,b) > V(f,j)$. The Network Manager can then solve the following problem:

$$ZLB_{IP-E} = \max \sum_{f \in \mathcal{F}} \sum_{q \in Q_f} LB(f,q)x(f,q) \tag{16a}$$

subject to

$$\sum_{f \in \mathcal{F}} \sum_{q \in Q_f : q \ni k} x(f,q) \leq 1 \quad \forall s \in \mathcal{S}, k \in L_s \tag{16b}$$

$$\sum_{q \in Q_f} x(f,q) = 1 \quad \forall f \in \mathcal{F} \tag{16c}$$

$$x(f,q) \in \{0,1\} \quad \forall f \in \mathcal{F}, q \in Q_f \tag{16d}$$

Problem (16) is equivalent to problem (4) with $V(f,q) = LB(f,q)$ for all $f \in \mathcal{F}$ and $q \in Q_f$, then it will be $ZLB_{IP-E} \leq Z_{IP-E}$, the strict inequality holding whenever $LB(f,q) < V(f,q)$ for at least one request $q \in Q_f$ for some flight $f \in \mathcal{F}$. Hence all the integer (feasible) exchanges calculated by solving the linear relaxation of problem (16) with $V(f,q) = LB(f,q)$ and implemented at prices equal to the dual variables corresponding to constraints (16b), will guarantee IR and weak BB, whenever $ZLB_{IP-E} > 0$.

The solution obtained (exchanges and prices associated) is not however a competitive equilibrium because at the given prices there could be some flight better-off with another exchange than the one implemented and this implies that the solution must be somewhat forced by the Network Manager. Then after a pre-determined number of iterations or a given elapsed time, if an equilibrium cannot be found by simply alternating local optimization and price update (i.e. problems 12 and 11), the Network Manager can impose the best solution calculated so far by solving the linear relaxation of problem (16) at the dual prices, i.e. the integer solution which gives the highest positive-value according to LB, which is the last feasible solution obtained since LB are always updated by increasing them.

The complete procedure for TW exchanges can proceed as it follows. The first step is to create a partition of the grand coalition $\mathcal{F}$ into independent subsets $M_i \subseteq \mathcal{F}$, such that for every pair of different flights $f \in M_i$ and $g \in M_j$ with $i \neq j$ it will be $Q_f \cap Q_g = \emptyset$. Then each subset M_i constitutes an independent market, since all the tradable resources will be within the market itself. From each of these markets, smaller sub-coalitions (sub-markets) of predetermined size Z are formed and then processed, in order to increase the probability of obtaining integer solutions to the linear relaxation of problem (16), thus allowing a computation of exchange and prices in a polynomial time. In fact we have observed experimentally that by reducing the size of the sub-markets, at the same time the value of the optimal exchange reduces while the percentage of instances for which $Z_{IP-E} = Z_{LP-E}$ increases.

Sub-markets are created according to their exchange potential. Flights in Market M_i are first ordered from the one with the lowest to the one with the highest assigned FPFS request. Then starting from the head of this ordered list one flight f is selected as well as its first potential seller g starting from the tail. A flight g is a potential seller for flight f if (i) they share at least one resource s ($S_f \cap S_g \neq \emptyset$) (ii) f prefers the TW k assigned to g on s than its currently assigned one j ($I_k < I_j$) (iii) TW k is feasible for f ($E^s_f \leq U_k$). If no potential seller exists the flight next to f is selected together with its first potential seller. Once a number of flights of the predetermined size Z has been selected, the first sub-market is created. Then a second sub-market which considers only flights which have not been previously included is built. This procedure may iteratively create up to $|M_i|/Z$ distinct sub-markets of size Z.

Relaying on the market M_1 of 425 flights as described in Section 4.2, Figure 1 illustrates the case in which 10 sub-markets of fixed size $|SM|$ are processed in order of their potential of exchange, in comparison with the situation in which they are formed by including flight at random. Each point corresponds to the average on 100 instances of different vectors of costs drawn from the same distribution.

By forming sub-markets with this procedure the exchanges attainable always give a higher global value than in the case trading occurs within random coalitions. The higher value exchanges are established during the first iterations, i.e. when the flights with higher trade potential, while after 5 iterations only a small number of residual exchanges occur.

The number of instances which give non integer solution to the linear relaxation of problem (4) increases with the size of the sub-market. When $|SM| = 40$ on average 0.15% of cases are non integer, 3.65% when $|SM| = 80$ and 4.35% when $|SM| = 120$. In these cases the dual variables are not supporting prices and one possible solution could be represented by the exchange optimal for the integer problem (4) and the prices calculated according to a VCG-based payment rule. Even if a competitive equilibrium with linear prices does not exist for such instances, our heuristic can still converge to a solution which guarantees IR and weak BB. Once the sub-markets have been formed according to criteria described before, the iterative mechanism can be applied to them.

If after a pre-determined number of iterations $MaxIter$ an equilibrium cannot be found, the last feasible solution calculated with LB is imposed by the Network Manager and the correspondent exchange is implemented at the dual prices. The diagram in Figure (2) represents the steps performed by the heuristic procedure.

4.2 Computational results

We simulated this procedure on a sample of traffic retrieved from real CFMU data relative to the two hours period from 09:00 AM to 11:00 AM on Friday August 15^{th}, 2008. There were a total of 60 capacity constrained resources and 482 regulated flights, that were clustered into 3 independent markets $(M_1, ..., M_3)$, with $|M_1| = 425$,$|M_2| = 34$,$|M_3| = 23$. Market M_1 included most of the flights interacting directly or indirectly in the exchange of TWs on 58 resources, flights in M_3 and only them were affected by a regulation on an upper en-route sector (LECMDOM) in the north of Spain that was limiting traffic flow rate to 43 entries/hour from 09:00 AM to 12:00 AM due to ATC capacity reasons. Flights in M_2 and only them competed for the assignment of TW on an upper en route sector (LFMMW2) located in the South of France that was closed from 10.15 AM to 11.30 AM due to ATC routing. In this such case, i.e. when capacity of a certain resource $z \in \mathcal{S}$ is null for a given time period $[st\_time_z; end\_time_z]$, we included just 1 TW in L_z with $I_1 = end\_time_z$ and infinite capacity in order to make problem feasible without rerouting (the same assumption is used in Terrab and Odoni, 1993).

For each flight $f \in \mathcal{F}$ we attached a vector cost of delay $CD_f \subset \mathbb{N}^3$, where each component represents the per-minute cost of delay according to the magnitude of the delay itself, which has been discretized into the three classes, while components $cd_f \in CD_f$ have been randomly drawn from the following uniform distributions in line with figures provided by Cook et al., 2004:

$$cd_f \sim \begin{cases} U(1;5) & \text{€/min} \quad \text{for } d_f \in [1;15) \\ U(15;25) & \text{€/min} \quad \text{for } d_f \in [15;45) \\ U(30;105) & \text{€/min} \quad \text{for } d_f \in [45; MaxDel_f] \end{cases}$$

Graph in Figure (3) shows the results obtained by applying the iterative Market Mechanism presented schematically in Figure (2). The highest value is achieved by letting all flights $f \in M_1$ participate in a unique repetition of the mechanism (i.e. $SM = M_1$), however the coordination tasks for the central authority may become slower and the percentage of non-integer solutions to the LP relaxation of problem (16) increases.

Instead by successively processing subsets of flights $SM \subset M_1$ according to their trade potential, the most valuable exchanges occur within the first five repetitions of the procedure. For the case with $|SM| = 40$, 90% of the final value is attained after 4 trades series, while with $|SM| = 80$ and $|SM| = 120$ just 2 repetition of the procedure are sufficient to achieve 95% of the final value. This is due to the highest number of exchanges which becomes possible by increasing the sizes of coalitions and it is a prove that remarkable cost savings are achievable also by employing a decentralized heuristic market mechanism.

We employed Xpress Mosel v.3.0.0 to code the heuristic procedure and Xpress Optimizer version 20.00.05 to solve all the linear problems. The result figures obtained regarding the cost savings achievable for Aircraft Operators by implementing exchanges are fairly high, especially if we consider that our traffic sample is relative to a 2 hours period. However they represent less than 20% of the total cost of delay originally caused by the FPFS allocation which equals for the traffic sample analyzed 740187 €, according to the algorithm we developed. Obviously it is not the same algorithm employed at CFMU for the allocation of ATFM slots, however this result is perfectly in line with the figures estimated by EUROCONTROL (EUROCONTROL Performance Review Commission, 2009).

5 Conclusions

We introduce the concept of Time Windows as a sequence of time periods through which flights execute their Business Trajectories. If not all airline requests can be accommodated, the Network Manager imposes TWs to flights following a FPFS rule, similarly to the current policy of allocating ATFM slots. Furthermore, we assume that airlines are interested in paying for delay reductions or receiving compensations for delay increases and we propose a market mechanism that allows airlines to trade their FPFS-allocated TWs. Our scheme is distributed as a centralized policy based on the airlines' costs cannot be implemented, unless the Network Manager knows the delay costs for each flight, which are private information internal to Aircraft Operators. Our mechanism let airlines decide autonomously for each flight whether it is preferable to keep the TWs obtained by the FPFS policy or to exchange it at the market price. In the case of multiple capacity constrained resources we show that an Individual Rational and weakly Budget Balanced market mechanism exists. We derive TW prices and their corresponding allocation by means of a primal heuristic using a distributed approach based on the Lagrangian relaxation. We apply our algorithm on a real instance considering approximately 500 flights and 60 sectors. Our results show that the market-based solution allows the participating airlines to decrease their overall ATFM delay-related costs with respect to the FPFS allocation by tens of thousands Euro per day.

6 Resume of the applicant

Andrea Ranieri received his M.Sc. in Control and Systems Engineering from the University of Trieste (Italy) in October 2005 and a Ph.D. in Information Engineering from the same University in March 2010. He developed part of his Ph.D. thesis at the EUROCONTROL Experimental Center in Bretigny-sur-Orge (France) and participated in several European projects (CATS, BLUEMED, ASSET) aimed at developing and validating new operational concepts for the next generation European ATM system.

His works have been presented at various international conferences including the XL Annual conference of the Italian Association of Operations Research (2009), the 8^{th} USA/Europe ATM R&D seminar (2009), the 4^{th} ODYSSEUS Workshop on Freight Transportation and Logistics (2009), the 26^{th} Congress of the International Council of Aeronautical Sciences (2008), the 3^{rd} International Conference on Research in Air Traffic (2008) and the 7^{th} USA/Europe ATM

R&D seminar (2007) where he received the best paper award for the track ''Finance, Deployment and Implementation issues".

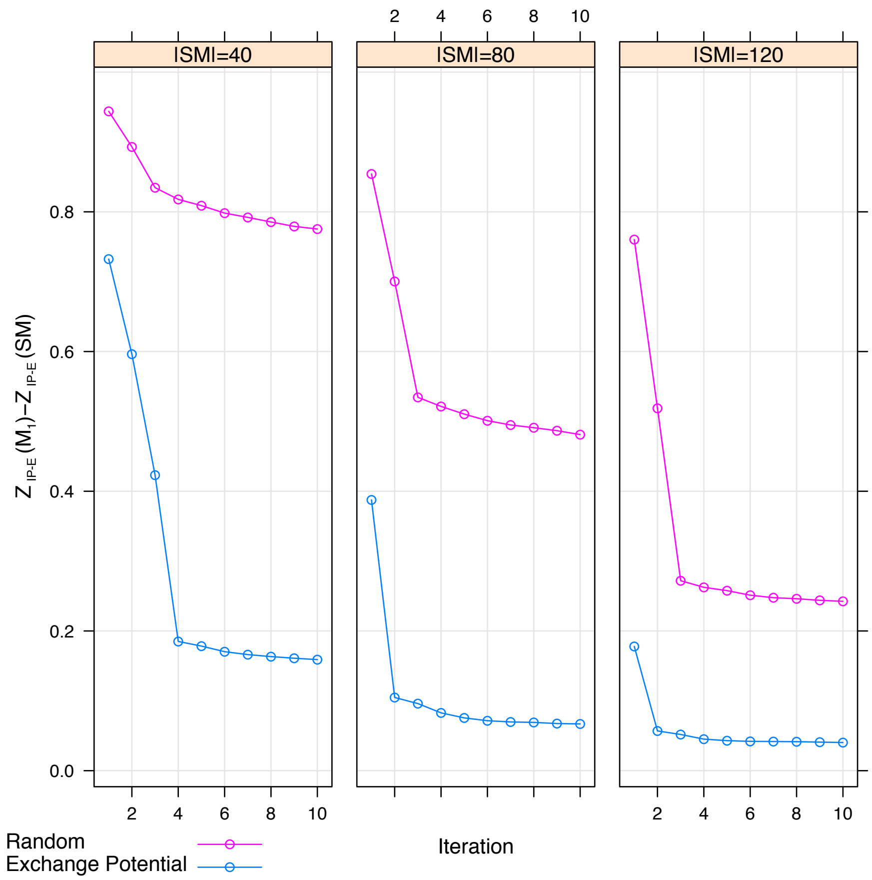

Figure 1: Gap between the optimal value of the exchange in the Main Market $M1$ and in Sub-markets SM.

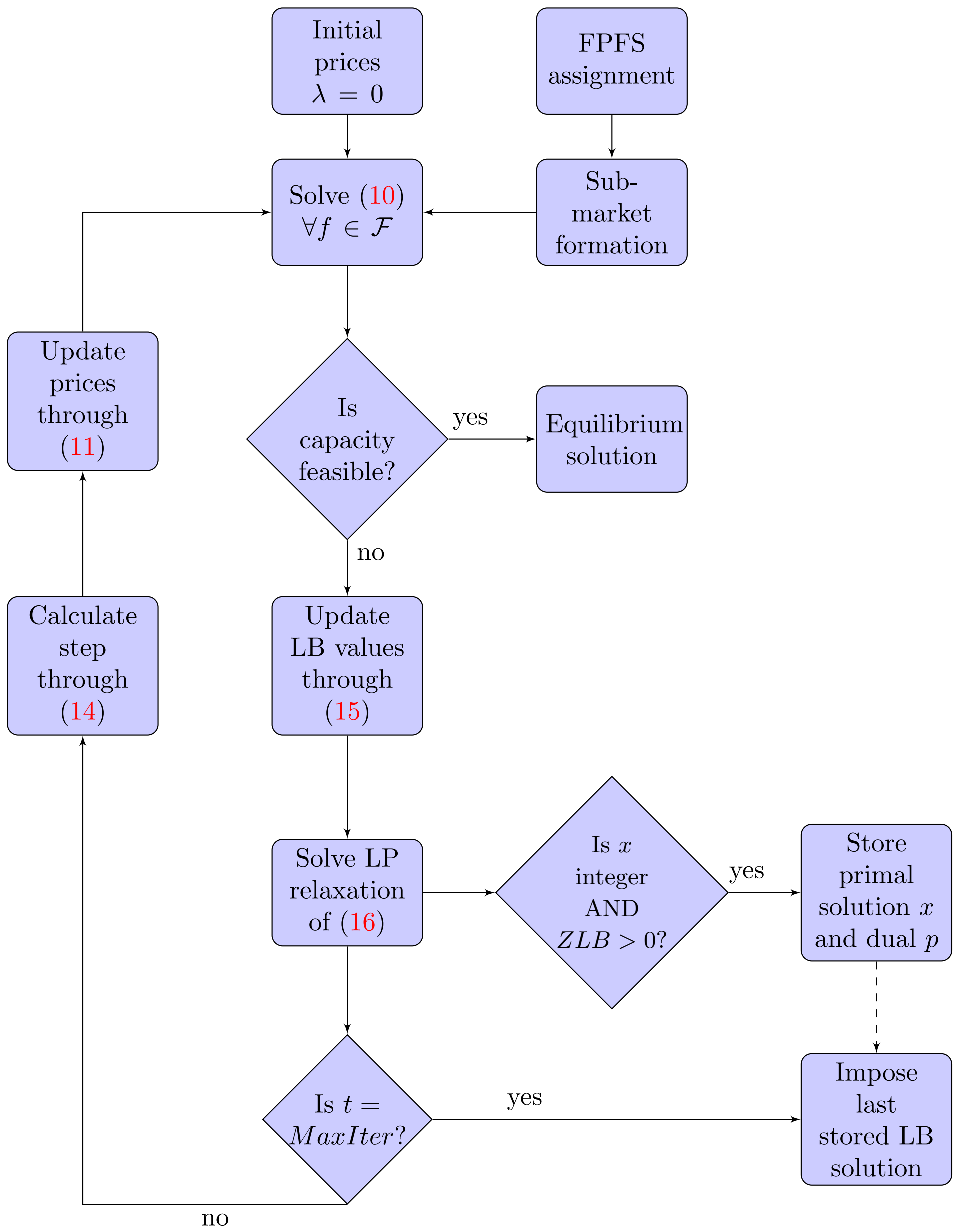

Figure 2: Schematics of the iterative Market Mechanism

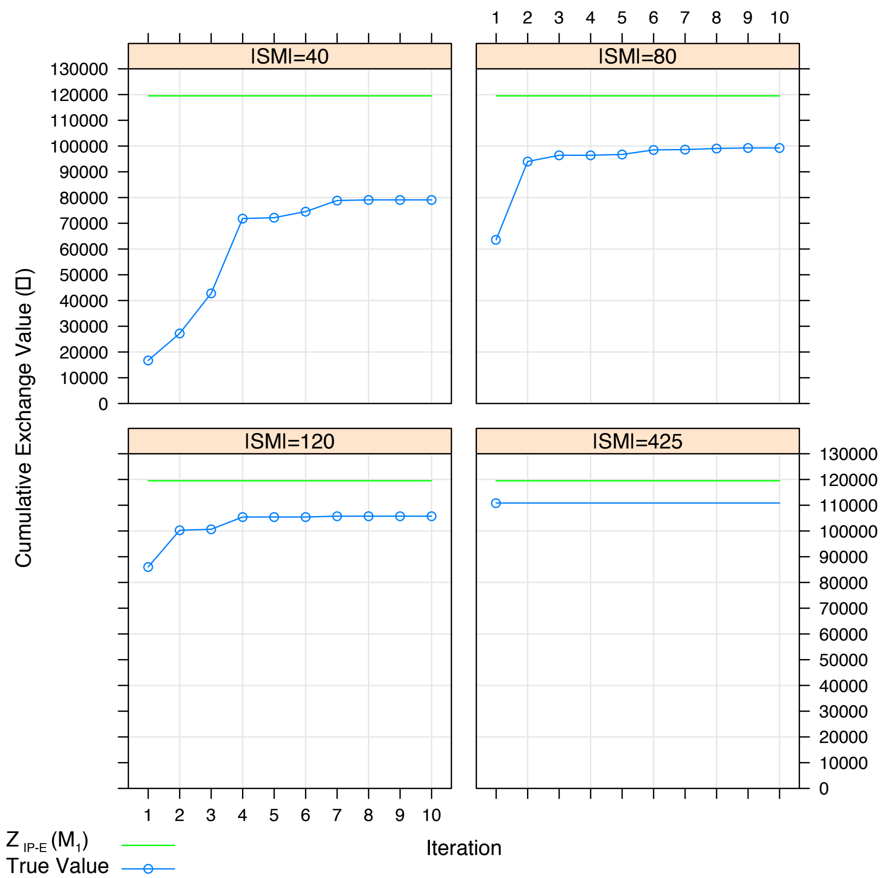

Figure 3: Value of the exchanges obtained through the iterative mechanism on several Sub-Markets

References

Alonso, A., Escudero, L. F., and Ortuño, M. T. (2000). A stochastic 0-1 program based approach for the air traffic flow management problem. *European Journal of Operational Research*, 120(1):47–62.

Andreatta, G. and Romanin-Jacur, G. (1987). Aircraft flow management under congestion. *Transportation Science*, 21(4):249–253.

Ball, M., Donohue, G., and Hoffman, K. (2005). Auctions for the safe, efficient and equitable allocation of airspace system resources. In Y., C. P., Shoham, and Steinberg, R., editors, *Combinatorial Auctions*, pages 507–538. MIT Press, Cambridge.

Bertsimas, D., Lulli, G., and Odoni, A. R. (2008). The air traffic flow management problem: An integer optimization approach. In Lodi, A., Panconesi, A., and Rinaldi, G., editors, *Lectures Notes in Computer Sciences*, volume 5035, pages 34–46. Springer-Verlag.

Bertsimas, D. and Stock Patterson, S. (1998). The air traffic flow management problem with enroute capacities. *Operations Research*, 46(3):406–422.

Bertsimas, D. and Stock Patterson, S. (2000). The traffic flow management rerouting problem in air traffic control: A dynamic network flow approach. *Transportation Science*, 34(3):239–255.

Bikhchandani, S. and Mamer, J. W. (1997). Competitive equilibrium in an exchange economy with indivisibilities. *Journal of Economic theory*, 74(2):385–413.

Clarke, E. H. (1971). Multipart pricing of public goods. *Public choice*, 11(1):17–33.

Cook, A., Tanner, G., and Anderson, S. (2004). *Evaluating the True Cost to Airlines of One Minute of Airborne or Ground Delay*. University of Westminster.

EUROCONTROL Central Flow Management Unit (2009). Basic cfmu handbook - general and cfmu systems v.13. Technical report, EUROCONTROL.

EUROCONTROL Performance Review Commission (2009). An assessment of air traffic management in europe during the calendar year 2008 - performance review report. Technical report, EUROCONTROL.

Fisher, M. L. (1985). An applications oriented guide to lagrangian relaxation. *Interfaces*, 15(2):10–21.

Groves, T. (1973). Incentives in teams. *Econometrica*, 41(4):617–631.

Held, M., Wolfe, P., and Crowder, H. P. (1974). Validation of subgradient optimization. *Mathematical Programming*, 6(1):62–88.

Hoffman, R. and Ball, M. O. (2000). A comparison of formulations for the single-airport ground-holding problem with banking constraints. *Operations Research*, 48(4):578–590.

Joint Planning and Development Office (2007). Concept of operations for the next generation air transportation system - version 2.0. Technical report, JPDO.

Kelso, A. S. and Crawford, V. P. (1982). Job matching, coalition formation, and gross substitutes. *Econometrica*, 50(6):1483–1504.

Larsson, T., Patriksson, M., and Strömberg, A.-B. (1999). Ergodic, primal convergence in dual subgradient schemes for convex programming. *Mathematical programming*, 86(2):283–312.

Le, L., Donohue, G., and Chen, C.-H. (2004). Auction-based slot allocation for traffic demand management at hartsfield atlanta international airport : A case study. *Journal of the Transportation Research Board*, 1(1888):50–58.

Lulli, G. and Odoni, A. (2007). The european air traffic flow management problem. *Transportation Science*, 41(4):431–443.

Myerson, R. and Satterwhite, M. (1983). Efficient mechanisms for bilateral trading. *Journal of Economic Theory*, 29(2):265–281.

Odoni, A. R. (1987). *Flow Control of Congested Networks*, chapter The flow management problem in air traffic control. Springer-Verlag, Berlin, Germany.

Parkes, D. C., Kalagnanam, J., and Eso, M. (2001). Achieving budget-balance with vickrey-based payment schemes in exchanges. In *Proceedings of the 17th International Joint Conference on Artificial Intelligence*, pages 1161–1168.

Rassenti, S., Smith, V., and Bulfin, R. (1982). A combinatorial auction mechanism for airport time slot allocation. *Bell Journal of Economics*, 13(2):402–417.

Richetta, O. and Odoni, A. R. (1993). Solving optimally the static ground-holding policy problem in air traffic control. *Transportation Science*, 27(3):228–238.

SESAR Joint Undertaking (2007). *SESAR Milestone Deliverable D3 - The ATM Target Concept*.

Terrab, M. and Odoni, A. R. (1993). Strategic flow management for air traffic control. *Operations Research*, 41(1):138–152.

Vickrey, W. (1961). Counterspeculation, auctions, and competitive sealed tenders. *The Journal of Finance*, 16(1):8–37.

Vossen, T. W. M. and Ball, M. O. (2006). Optimization and mediated bartering models for ground delay programs. *Naval Research Logistics*, 53(1):75–90.

Vranas, P. B., Bertsimas, D., and Odoni, A. R. (1994a). The multi-airport ground-holding problem in air traffic control. *Operations Research*, 42(2):249–261.

Vranas, P. B. M., Bertsimas, D., and Odoni, A. R. (1994b). Dynamic ground-holding policies for a network of airports. *Transportation Science*, 28(4):275–291.

Wambsganss, M. C. (2001). Collaborative decision making in air traffic management. In Bianco, L., Dell'Olmo, P., and Odoni, A. R., editors, *New Concepts and Methods in Air Traffic Management*, pages 1–16. Springer-Verlag.