The Airline Container Loading Problem with Pickup and Delivery

V. Lurkin*, M. Schyns

University of Liege-HEC Management School-QuantOM – Belgium

Abstract

The present paper looks into the problem of optimizing the loading of a set of containers and pallets into cargo aircraft serving multiple airports. Due to the pickup and delivery operations occurring at intermediate airports, this problem is simultaneously a weight and balance problem and a sequencing problem. Our objective is to minimize fuel and handling operations costs. This problem is shown to be NP-hard. We resort to a mixed integer linear program. On the basis of a professional partner's real-world data, TNT Airways, we perform numerical experiments using a standard B&C library. This approach yields better solutions than traditional manual planning, which results in substantial cost savings.

Keywords: OR in airlines, Assignment problem, Fuel consumption, Weight and balance, Sequencing problem, ILP

1. Introduction

In the *Airline Container Loading Problem with Pickup and Delivery* (ACLPPD), a set of containers and pallets, known as *Unit Load Devices (ULD)*, must be loaded into a compartmentalized cargo aircraft. We take into account that pickup & delivery operations take place at different airports during the trip. The loading task is illustrated in Figure 1. We propose an exact solution approach relying on a mixed integer linear program to find the optimal assignment of the ULDs.

While air cargo represents but 10% of world trade volume, it is in excess of $6.4 trillion par annum, which roughly amounts to 35% of world trade value (IATA (2013a)). Air cargo transportation thus plays a highly significant economic role. Optimizing loading assignment on board is critical to airlines for several reasons. First and foremost, correct

*Corresponding author

Email addresses: vlurkin@ulg.ac.be (V. Lurkin), M.Schyns@ulg.ac.be (M. Schyns)

Figure 1: a ULD being loaded on board through the main deck side cargo door (left),through the nose door (middle), and through a lower deck side cargo door (right).

loading conditions safety. Inappropriate loading, on the other hand, can cause a lot of damage. The aircraft, the freight or even the crew could be at risk. The present research, therefore, models a wide set of constraints for operators to reckon with on a daily basis. This model applies to all aircraft and loads complying with international standards. Taking into account the same constraints as Limbourg et al. (2012), we adapt them to the case of a sequence of routes, called *legs* while considering the additional case of hazardous products and oversized ULDs. Second, optimal loading has a positive impact on aerodynamics, thus making for less fuel consumption, i.e. reduced cost and environmental impact. This issue is crucial for airlines, which are hit hard by rising oil prices and increasing pressure to reduce their carbon dioxide emissions. The present research analyzes fuel and handling operations in order to minimize costs. Handling these two first requirements is done through a proper distribution of the ULD weights inside the aircraft. This part is a weight and balance problem. The third reason why optimal loading is so important for airlines is that managing the operations on the ground is also a challenge, especially when the trip includes several legs with P&D operations. Reducing the number of handling operations saves time. Time saved means reduced labor costs per flight. It also allows shorter turnaround time, i.e. the time that elapses between the moment the plane lands on the runway and when it takes off again, and hence reduced airport fees. Time saved might also be used for other valuable operations. Optimizing the planning of the loading is also crucial and constitutes another reason why to consider this problem. Indeed, the loadmaster must build a planning within a very tight time window while this requires, by hand, a lot of time. With an interactive computerized efficient tool, he would be able to consider different alternatives and to select the best solution with respect to his experience and the real conditions faced on the ground.

In this context, the problem no longer consists merely, as in Limbourg et al. (2012), in positioning ULDs to reach a proper equilibrium but also in defining the sequence of unloading and loading operations at airports. As there is only one path between any ULD and the exit door, this path must be free to unload the ULD. The task is to minimize, at each airport, the number of ULDs in transit to be unloaded in order to have access to ULDs reaching their delivery point. The same question arises when pickup occurs. The problem is even more complex when, as it sometimes turns out, several doors can be used. The cost of these handling operations is the second element of the new objective function we propose. It is important to notice that we now face two conflicting objectives: optimizing the assignments on board for fuel and for the ground operations. Our contribution is to propose an exact approach to solve simultaneously both the weight and balance problem over a multi-leg trip and the sequencing problems associated to the pickups and deliveries. We resort to a mixed integer linear program where the objective is to minimize both costs.

Currently, this very complex problem (NP-hard) is still essentially solved manually on the basis of best practices. As load planners have very tight time windows to choose assignments, they mainly focus on finding a feasible and reasonable solution. As a rule, they do not incorporate P&D operations in the planning process. A common way to handle several legs is indeed to plan each leg independently. Accordingly, almost the entire cargo may be unloaded at intermediate airports and the ULDs that have not reached final destinations are reloaded afterwards, which is the worst possible scenario for ground operations. We show, on the basis of our first results on real data provided by industrial partners, that our approach allows significant savings.

The remainder of this paper is organized as follows. Section 2 outlines the problem and the assumptions involved. Related literature and contributions are presented in Section 3. Section 4 describes in more details the problem and provides the mathematical formulation of the model. Section 5 gives information on the theoretical complexity of the problem while Section 6 illustrates the performance of the approach through numerical results. Finally, some conclusions are drawn.

2. Problem summary and assumptions

The ACLPPD can be informally summarized as:

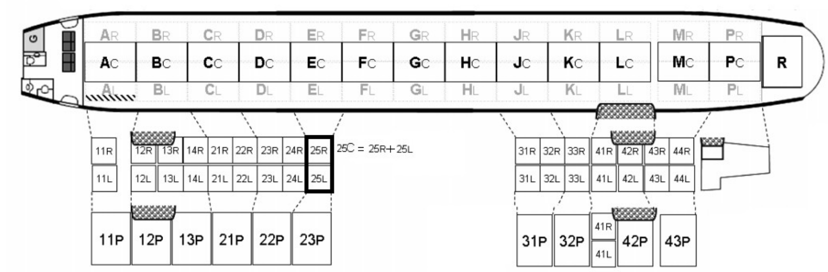

Figure 2: illustration of the structure of an aircraft

min	Fuel and loading operations costs on the entire trip (global optimization)
s.t.	Pickup & delivery sequences are feasible
	Customer demand is satisfied (each ULD is loaded)
	Each ULD fits in an aircraft position
	A position accepts only one ULD
	Some positions are overlapping and cannot be used simultaneously
	Longitudinal stability is within certified limits (ZFW,TOW,LW)
	Lateral stability is within certified limits
	Weight per position is below the certified limit
	Combined weight load limits are set
	Cumulative weight load limits are set
	Regulations for hazardous goods are fulfilled
	Oversized ULDs are managed

The decision variables are the location of each ULD inside the aircraft. The constraints are described in detail in Section 4.3. We make the following main assumptions. A cargo aircraft has to deliver goods to several airports. The flight plan is supposed to be known beforehand, which means that these airports as well as the order in which they will be visited are known. We also know all the containers and pallets (ULDs) to be delivered. For each ULD, we know its size, shape, weight, respective origin and destination. We follow international standards for the description of the ULDs. Full details on the coding standards can be found in the *IATA ULD Regulations (ULDR)* (IATA (2013b)). A cargo aircraft generally contains multiple decks with multiple configurations of positions for each. A position is simply a particular aircraft space accommodating exactly one ULD. The location of each position as well as all ULDs that fit into that position are also known. The location of the different doors is also given. A cargo aircraft has generally one side cargo door on the main deck and one for each of the three compartments of the lower deck. In addition, a nose door is sometimes available for the main deck. An example of cargo aircraft structure is illustrated on Figure 2. The focus of this research is on cargo transportation. While the central ideas remain the same and extensions of our approach could be considered, neither passenger transportation nor transportation of goods in the lower deck of passenger aircraft are covered in this paper.

3. Related literature and contributions

This problem is an Assignment Problem (AP) referred to in the literature as belonging to the family of Weight & Balance Problems. The scientific literature on aircraft cargo load planning is not extensive but still contains a number of papers. As in Limbourg et al. (2012), we classify these papers into three main categories. First, several papers consider how to optimize the loading of freight inside ULDs (Chan and Kumar (2006); Chan et al. (2006); Yan et al. (2008); Li et al. (2009); Wu (2010); Tang and Chang (2010); Tang (2011)) independently of the aircraft. This part essentially deals with Bin Packing Problems (BPP). A second important question is how to select the ULDs or items to be loaded in an aircraft or a fleet of aircraft; i.e. Knapsack Problems (KP). Papers on this subject cover military (Ng (1992); Heidelberg et al. (1998); Guèret et al. (2003); Kaluzny and Shaw (2009); Nance et al. (2011)) and commercial applications (Mongeau and Bès (2003); Fok and Chun (August 14-17, 2004); Tian et al. (2009); Verstichel et al. (2011)). Finally, some authors, like ourselves, optimize the location of ULDs in an aircraft. In this domain, the literature can be subdivided according to two approaches: Bin Packing or Assignment. In the BPP approaches (see e.g. Amiouny et al. (1992); Heidelberg et al. (1998); Mathur (1998); Guèret et al. (2003); Nance et al. (2011)), the authors attempt to fill the aircraft continuously by excluding empty spaces between the items while in the AP approaches (see e.g. Limbourg et al. (2012); Mongeau and Bès (2003); Verstichel et al. (2011)), they try to allocate ULDs into predefined standardized positions. These three families, however, are not exhaustive and some papers fall in between categories. This literature also varies on at least four other dimensions: the precise definition of the objective function, the nature of the shipments, the constraints taken into account, and the solution algorithm (exact methods or heuristics).

Related problems deal with the loading of containers in trucks, trains or ships. All these problems are combinatorial optimization problems which have to satisfy, several and sometimes similar loading constraints. Among these, the more distant is the loading of a truck since there are no predefined positions for the containers. The Truck Loading Problem

(TLP) is not an Assignment Problem but it fits under the more general Container Loading Problem (CLP) (Bortfeldt and Wäscher (2013)) which includes the Bin Packing Problems. Basically, in the CLP, a series of 3D boxes has to be stacked into a 3D container without overlapping. When the truck has to deliver goods at several destinations, the problem is sometimes referred to as the Multi-Drop Container Loading Problem (MDCLP). In this case, the minimization of the number of handling operations is also a major preoccupation. Among the most recent papers, we can cite the works of Pan et al. (2011) and Altarazi (2013).

The Train and Ship Loading Problems are closer to the definition of our problem. Indeed, for the loading of a ship, as stated e.g. in Dubrovsky et al. (2002); Ambrosino et al. (2010); Øvstebø et al. (2011), or of a train, as stated e.g. in Bostel and Dejax (1998); Corry and Kozan (2006); Bruns and Knust (2011); Ambrosino et al. (2011), available locations for containers are also predefined. In addition, as in our problem, the sequencing of loading/unloading operations also determines the ship or train turnaround time. For these specific carriers, Avriel et al. (1998); Imai et al. (2006); Li et al. (2008); Øvstebø et al. (2011) suggest keeping number of handling operations to the minimum. However, the structure of a ship or a train is significantly different from that of an aircraft. Consequently, the weight constraints and the loading operations are very specific, which makes the ACLPPD peculiar and basically different from these other two loading problems. Moreover, especially in the case of ships, the number of containers to be loaded leads to problems which are too big to be solved by exact methods. The literature mainly consists of heuristics and does not present exact mathematical models for real problems.

Some papers consider loading questions in conjunction with other important questions. In a recent survey, Pollaris et al. (2013) provide a list of research projects carried out on Vehicle Routing Problems (VRP) with loading constraints. These problems are extremely complex and few or limited exact methods are presented. All of them are based on bin packing models and not on assignment to predefined positions.

The studies of Mongeau and Bès (2003); Verstichel et al. (2011) and especially Limbourg et al. (2012), which we started from, relate most closely to our work. These three papers indeed deal with commercial cargo aircraft with predefined positions and standardized ULDs, use exact methods and consider the aircraft's center of gravity. We nonetheless significantly depart from these works. Our contributions are multiple. Our main contribution is to consider a trip consisting of several legs with pickups and deliveries occurring at intermediate airports. This compels us to model the handling operations on the ground and to adapt the basic constraints of the Weight & Balance Problem to the multi-leg trip. We analyze the impact of loading operations in terms of costs. The mathematical model that we propose simultaneously takes into account two conflicting monetary objectives. To our knowledge, this has never been done before. The integration of a larger set of realistic constraints is also a contribution. We also provide information about the complexity of the problem. Finally, we show measures of performance based on real data.

4. Problem description and mathematical formulation

4.1. Main parameters and variables

Our model is built on three main sets of parameters. The first one is the set $\mathbb{L}$ of legs, which are the different parts of a trip separating two successive airports. This model considers but two legs. However, generalization to more legs is straightforward. A trip composed of two legs is a common case for long range flights while considering more legs would essentially complicate notations. The second main set of parameters is the set $\mathbb{U}$ of ULDs. For each ULD, we know its type (IATA code), its weight w_i (supposed to be uniformly distributed), its airports of origin and destination. Knowing the origin and destination of each ULD, we can establish three distinct subsets of ULDs: $\mathbb{U}_1$ for ULDs loaded at the depot and unloaded at the first destination, $\mathbb{U}_2$ for those loaded at the first destination and unloaded at the second one, $\mathbb{U}_3$ for those loaded at the depot and unloaded at the second destination. For ease of notations, we also define $\mathbb{U}_k^L$ as the subset of ULDs present in the aircraft for the leg $k \in \mathbb{L}$. By definition, the intersection of $\mathbb{U}_1^L$ and $\mathbb{U}_2^L$ is $\mathbb{U}_3$. The last set of parameters is the set $\mathbb{P}$ defining the available positions, which are predefined spaces in the aircraft able to accommodate ULDs. Each position is located on a specific deck and can be situated on the left-hand side (L), on the right-hand side (R) or even cover the whole width (C). To identify all positions on either side of the center, we respectively use the subsets $\mathbb{P}_R$ and $\mathbb{P}_L$.

Our main variables are binary variables x_{ijk} defined as:

$$x_{ijk} = \begin{cases} 1 & \text{if ULD } i \text{ is in position } j \text{ during leg } k \\ 0 & \text{otherwise,} \end{cases} \qquad \forall i \in \mathbb{U}, \forall j \in \mathbb{P}, \forall k \in \mathbb{L}$$

4.2. Objective function

As stated in the previous section, the objective of the ACLPPD is to assign each ULD to one position while minimizing total costs, which is realised by minimizing fuel consumption on each leg and handling operations at intermediate airports. More exactly, we will minimize the overcosts.

As the location of the center of gravity (CG) has an impact on fuel consumption, the CG, on each leg, should be located at its best position. Without going into too many technical considerations, the fuel consumption of an aircraft depends on the result of a continuous battle of forces; namely the weight, lift, thrust, and drag (see Figure 3). A slightly aft CG should reduce the drag and ultimately the thrust. Less thrust means less fuel consumption. The location of CG is restricted within a range of certified limits

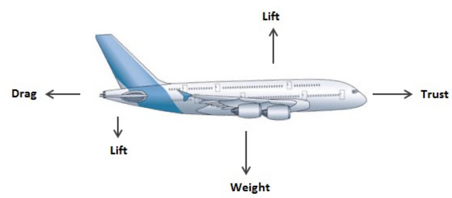

Figure 3: The forces acting on an aircraft in flight

defined by the aircraft's manufacturer. These limits are crucially important since a CG value outside of them can cause instability resulting in harmful effects. Within this range, some freedom is allowed. Airlines and pilots are aware that an aft CG usually saves fuel. Nonetheless, airlines typically define target location of the CG around the center of certified ranges. In our case, we define the best location close to the aft certified limit, less an additional security margin left to the operator's appreciation (see Section 4.3 for the longitudinal stability constraints and the computation of the optimal CG location). In mathematical terms, this gives:

$$\min \sum_{\forall k \in \mathbb{L}} \epsilon_k$$

Subject to:

$$CG_k - OCG_k - \epsilon_k \leq 0 \qquad \forall k \in \mathbb{L} \tag{1}$$

$$CG_k - OCG_k + \epsilon_k \geq 0 \qquad \forall k \in \mathbb{L} \tag{2}$$

where OCG_k is the requested optimal CG location, CG_k is the CG obtained after the assignment of the ULDs and ϵ_k is the resulting deviation from the target for leg k.

The second objective of the ACLPPD is to minimize the number of handling operations. All ULDs of $\mathbb{U}_1$ (resp. $\mathbb{U}_2$) must be loaded at the origin of leg 1 (resp. 2) and unloaded at the destination of the same leg. The number of ULD moves therefore is equal to the number of ULDs in these sets and cannot be decreased. ULDs of $\mathbb{U}_3$ doesn't a priori need to be unloaded at the intermediate airport (optimal case). However, all these ULDs are typically unloaded (worst case) in order to unload the ULDs belonging to $\mathbb{U}_1$ and load those belonging to $\mathbb{U}_2$. This comes from the fact that $\mathbb{U}_3$ ULDs could stand in the path of those that must be loaded/unloaded at the intermediate airport but also due to the fact that the load plan must be adapted to the new leg while still ensuring a feasible and correct location of the CG for the next leg. If the initial load plan takes into account the different legs, it should be possible to assign ULDs in $\mathbb{U}_3$ to positions that would not be in the entry/exit path of other ULDs and that would allow to reach a suitable CG for the second leg. We therefore focus on the re-handling operations at the intermediate airport.

Moreover, an aircraft compartment can have more than one door. In this case, the number of operations can be minimized by using the most accessible door. Without loss of generality, let's consider that each position can be reached from up to two doors: the first one in the direction of the nose (”*the nose door*”) and the second one in the direction of the tail (”*the tail door*”). There could be more than two doors for a same deck but only two of them are relevant for each position.

An example for the main deck of a Boeing 747 is illustrated in Figure 4. For the position j , the *nose door* is situated exactly at the nose and the *tail door* is the lateral door. For the position j', the *nose door* is the lateral door and there is no *tail door*. We

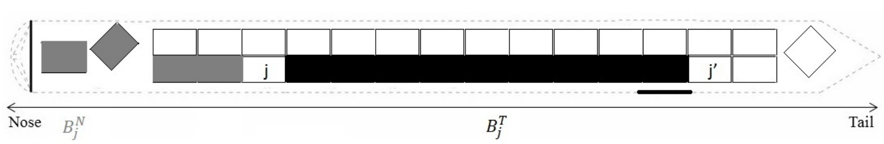

Figure 4: The different doors for the Boeing 747

can identify the set of positions in the entry/exit paths to/from the position j as follows (see Figure 4):

- $\mathbb{B}_j^N$ is the set of all positions situated between position j and the first door in the direction of the nose. $\mathbb{B}_j^N$ is an empty set if there is no nose door.

- $\mathbb{B}_j^T$ is the set of all positions situated between position j and the first door in the direction of the tail. $\mathbb{B}_j^T$ is an empty set if there is no tail door.

We also introduce the two following sets of binary variables to determine through which door a ULD is unloaded:

$$\alpha_j^N = \begin{cases} 1 & \text{if the ULD in position } j \text{ is unloaded through the } \textit{nose door}, \\ 0 & \text{otherwise .} \end{cases}$$

$$\alpha_j^T = \begin{cases} 1 & \text{if the ULD in position } j \text{ is unloaded through the } \textit{tail door}, \\ 0 & \text{otherwise .} \end{cases}$$

The values of α_j^N and α_j^T are determined by the following set of constraints:

$$\alpha_j^N + \alpha_j^T \leq 1 \qquad \forall j \in \mathbb{P} \tag{3}$$

$$\alpha_j^N + \alpha_j^T \geq x_{ij1} \qquad \forall j \in \mathbb{P}, \quad \forall i \in \mathbb{U}_1 \tag{4}$$

$$\alpha_j^N = 0 \qquad \forall j \in \mathbb{P} \ | \ \mathbb{B}_j^N = \emptyset \tag{5}$$

$$\alpha_j^T = 0 \qquad \forall j \in \mathbb{P} \ | \ \mathbb{B}_j^T = \emptyset \tag{6}$$

Constraint (3) ensures that when unloading a ULD in position j, one door only is used. Constraint (4) guarantees that the ULDs in $\mathbb{U}_1$ are unloaded at the intermediate airport. Constraints (5) and (6) stipulate that some doors may either not exist or not be used.

We still need to introduce two other sets of binary variables to count the number of ULDs belonging to $\mathbb{U}_3$ that are unloaded, then loaded again, at the intermediate airport. Since this airport is not their final destination, these operations ought to be spared.

$$n_j^N = \begin{cases} 1 & \text{if the ULD in position } j \text{ is unnecessarily unloaded through the nose door ,} \\ 0 & \text{otherwise .} \end{cases}$$

$$n_j^T = \begin{cases} 1 & \text{if the ULD in position } j \text{ is unnecessarily unloaded through the tail door ,} \\ 0 & \text{otherwise .} \end{cases}$$

The n_j are determined on the basis of the α_j values thanks to these constraints:

$$\sum_{j' \in \mathbb{B}_j^N} \alpha_{j'}^T \leq n_j^T |\mathbb{B}_j^N| + (1 - x_{ij1}) |\mathbb{B}_j^N| \qquad \forall j \in \mathbb{P}, \quad \forall i \in \mathbb{U}_3 \tag{7}$$

$$\sum_{j' \in \mathbb{B}_j^T} \alpha_{j'}^N \leq n_j^N |\mathbb{B}_j^T| + (1 - x_{ij1}) |\mathbb{B}_j^T| \qquad \forall j \in \mathbb{P}, \quad \forall i \in \mathbb{U}_3 \tag{8}$$

The sum in the left term of (7) is the number of ULDs that will go through position j to be unloaded by the tail door. The maximal value of this sum is equal to the size of $\mathbb{B}_j^N$. If this sum is strictly positive, the ULD in position j is blocking the exit path and must also be unloaded through the tail door. If this ULD belongs to $\mathbb{U}_3$, it is an "unnecessary" operation, x_{ij1} equals one and n_j^T can only take the value one. If it is not the case, the constraint is not binding and n_j^T is free but will take the value zero due to the objective function. The same reasoning applies to the other door with constraint (8).

We want to minimize these unnecessary unloadings, which leads to the objective function for the ground operations:

$$\min \sum_{\forall j \in \mathbb{P}} (n_j^N + n_j^T)$$

Finally, both objectives have to be considered together and expressed in monetary terms. For the first element, the target location of the CG is typically expressed as a percentage of the aircraft Mean Aerodynamic Chord (MAC). The MAC is defined by the airfoils and is a mean distance between leading and trailing edges in the direction of the airflow. For the Airbus 330, a displacement of CG from its location of reference (28%) to a more aft CG (37%) could give rise to a 0.5% increase of air nautical miles per KG of fuel (Airbus Fuel Economy Material (2004)). This result can be converted in terms of fuel savings. The pilots and load planners consulted expect about a 2.5% decrease of fuel consumption for a 777 when the CG is displaced by a distance equivalent to 10% of the MAC to the aft. The exact relationship between CG location and fuel consumption depends on each type of aircraft. These examples provide rough approximations showing the amount of the potential savings. Knowing the exact relation between the CG location and the fuel savings for a specific aircraft and the absolute volume of fuel required for the trip with this optimal location of the CG, we can measure how many tons of fuel must be added for one percent forward shift of the CG. By multiplying this by the fuel price, we obtain the cost by unit deviation. ϵ_k measures the deviation. Set c_k^f as the monetary cost per unit deviation for fuel consumption, then $(\epsilon_k c_k^f)$ is the additional cost induced on leg k by an improper location of the CG (too forward). As the first goal is to push the CG to the aft, any positive value for the c_k^f would do the job. c_k^f can be interpreted as penalty coefficients. The main drawback of an approximative value is that it would lead to an approximative value of the total cost reduction, but it would have no impact on the optimal assignment (at least if this value remains a reasonable approximation with respect to the second objective: the cost minimization associated to the loading operations). If airlines and load planners prefer to select a sub-optimal location more inside the range of values certified by the aircraft manufacturer, the target CG does not correspond anymore to the upper certified limit and always pushing to the aft is not suitable. We decide to model it as a soft constraint. Any deviation of the CG to the aft beyond the target must also be penalized and at least by a coefficient c_k^f to give priority to this constraint. Therefore, the ϵ_k measure the absolute deviations with respect to the optimal CG. Alternatively, a hard constraint could have been easily implemented by defining a lower max threshold in constraint 26. When optimal solutions exist, as it is the case for most of our numerical experiments, both options are totally equivalent since the additional cost is null.

The second cost in the objective function is associated to the ground operations. The handlers assigned to this task need to operate quickly because the aircraft has a limited time window before the next departure. The total wages cost depends on the number of employees assigned to the task, which is directly proportional to the number of ULDs to be moved. Each company can measure it. The cost coefficient for handling one ULD is denoted as c^h.

We resort to the following multi-criteria objective function where all coefficients can be defined precisely knowing cost realities:

$$min \quad \sum_{\forall k \in \mathbb{L}} c_k^f \epsilon_k + c^h \sum_{\forall j \in \mathbb{P}} (n_j^N + n_j^T) \tag{9}$$

4.3. Constraints

Some P&D constraints are directly related to the objective function and have been presented in the previous section. Taking into account the sequences of loading and unloading leads to additional P&D constraints. We present first these additional constraints. Next, we introduce hazardous and oversized products. We conclude this section by adapting the realistic constraints presented in Limbourg et al. (2012) to the case of multiple legs.

4.3.1. Pickup & Delivery Constraints

Categories of ULDs. Due to their origin and destination, not all ULDs are present on each leg, which means that the corresponding x_{ijk} variables can be initialized to zero (and removed during the optimization). This is specified by constraint (10).

$$x_{ijk} = 0 \qquad \forall i \notin \mathbb{U}_k^L,\ \forall j \in \mathbb{P} \tag{10}$$

Correct unloading sequence. Constraints (11) and (12) prevent collisions of ULDs being unloaded through different doors. Constraint (11) states that the ULD in position j can only go out through the nose door ($\alpha_j^N = 1$), if there is no ULD between j and the nose door ($\mathbb{B}_j^N$) trying to go out in the opposite direction, i.e. through the tail door (the sum is null). Constraint (12) manages the move in the opposite direction.

$$\sum_{j' \in \mathbb{B}_j^N} \alpha_{j'}^T \leq (1 - \alpha_j^N)|\mathbb{B}_j^N| \qquad \forall j \in \mathbb{P} \tag{11}$$

$$\sum_{j' \in \mathbb{B}_j^T} \alpha_{j'}^N \leq (1 - \alpha_j^T)|\mathbb{B}_j^T| \qquad \forall j \in \mathbb{P} \tag{12}$$

No exchange of positions inside the aircraft. Any ULD assigned to a different position for the second leg, must first be unloaded. Constraints (13) and (14) ensure that each ULD not unloaded at the intermediate destination, i.e. when $(n_j^N + n_j^T)$ is null, keeps the same position for the second leg ($x_{ij1} = x_{ij2}$).

$$x_{ij1} - x_{ij2} \leq (n_j^N + n_j^T) \qquad \forall i \in \mathbb{U}_3 \tag{13}$$

$$x_{ij2} - x_{ij1} \leq (n_j^N + n_j^T) \qquad \forall i \in \mathbb{U}_3 \tag{14}$$

Feasible loading sequence. Constraints (15), (16), (17), (18) and (19) ensure a feasible loading sequence of the ULDs belonging to $\mathbb{U}_2$ at the intermediate airport. These ULDs can only reach positions for which the path is free. Constraint (15) checks if there exists a free path between the nose door and position j. At the time of loading, the path can only be blocked by ULDs of $\mathbb{U}_3$ that were not unloaded ($n_{j'}^T + n_{j'}^N = 0$). If there is no free path, the left term of (15) is strictly positive and the new binary variable β_j^N can only take the value one. Similarly, constraint (16) restricts β_j^T to one if the path through the tail door to position j is not free. If both paths are blocked, constraint (17) states that ULD i cannot be assigned to position j. When there is no nose (resp. tail) door associated to a position, then the path from this direction to the position j is automatically forbidden and consequently $\beta_j^N = 1$ (resp. $\beta_j^T = 1$). This is the purpose of constraints (18) and (19).

$$\sum_{i' \in \mathbb{U}_3} \sum_{j' \in \mathbb{B}_j^N} (x_{i'j'1} - n_{j'}^T - n_{j'}^N) \leq \beta_j^N |\mathbb{B}_j^N| \qquad \forall j \in \mathbb{P} \tag{15}$$

$$\sum_{i' \in \mathbb{U}_3} \sum_{j' \in \mathbb{B}_j^T} (x_{i'j'1} - n_{j'}^T - n_{j'}^N) \leq \beta_j^T |\mathbb{B}_j^T| \qquad \forall j \in \mathbb{P} \tag{16}$$

$$\beta_j^T + \beta_j^N - 1 \leq (1 - x_{ij2}) \qquad \forall j \in \mathbb{P}, \quad \forall i \in \mathbb{U}_2 \tag{17}$$

$$\beta_j^N = 1 \qquad \forall j \in \mathbb{P} \ \mid \ \mathbb{B}^N = \emptyset \tag{18}$$

$$\beta_j^T = 1 \qquad \forall j \in \mathbb{P} \ \mid \ \mathbb{B}^T = \emptyset \tag{19}$$

Inaccessibility of some doors. Due to the dimensions of ULDs, some doors might not be accessible. If a ULD i is assigned to position j on any leg k ($x_{ijk} = 1$) and cannot go through the nose (resp. tail) door, constraint (20) (resp. constraint (21)) ensures that this unloading direction is forbidden.

$$x_{ijk} \leq (1 - \alpha_j^N) \quad \forall i \in \mathbb{U}_k^L \mid U_i \text{ doesn't pass through the nose door}, \quad \forall j \in \mathbb{P}, \quad \forall k \in \mathbb{L} \tag{20}$$

$$x_{ijk} \leq (1 - \alpha_j^T) \quad \forall i \in \mathbb{U}_k^L \mid U_i \text{ doesn't pass through the tail door}, \quad \forall j \in \mathbb{P}, \quad \forall k \in \mathbb{L} \tag{21}$$

Domino effect. Finally, in order to achieve a CG as close as possible to the target value, some ULDs could be assigned to a different position on the second leg. It will be the case each time the gain in terms of fuel is greater than the additional costs induced by the ground operations. This case, however, implies a domino effect: all ULDs in the path of those being unloaded to reach a better CG will have to be unloaded as well. This is ensured by constraints (22) and (23). With constraint (22), if the ULD in position j is unnecessarily unloaded through the tail door ($n_j^T = 1$), then all ULDs of $\mathbb{U}_3$ that were assigned to a position between j and the tail door ($\mathbb{B}_j^T$) over the first leg ($x_{ij'1} = 1$) must also be unloaded "unnecessarily" ($n_{j'}^T = 1$). Constraint (23) does the same for the nose door.

$$x_{ij'1} - 1 + n_j^T \leq n_{j'}^T \qquad \forall i \in \mathbb{U}_3, \quad \forall j \in \mathbb{P}, \quad \forall j' \in \mathbb{B}_j^T \tag{22}$$

$$x_{ij'1} - 1 + n_j^N \leq n_{j'}^N \qquad \forall i \in \mathbb{U}_3, \quad \forall j \in \mathbb{P}, \quad \forall j' \in \mathbb{B}_j^N \tag{23}$$

4.3.2. Other Advanced Constraints

Here are two new constraints linked to hazardous and oversized products.

Hazardous goods. Some loads may contain hazardous goods. Segregation requirements apply to ensure safety. Generally, dangerous goods can be classified into a limited number of categories. The required segregation distance (in inches) between the categories i and i' is known and denoted $s_{ii'}$. The effective longitudinal distance (in inches) between the positions j and j' is denoted $e_{jj'}$. This yields constraint (24).

$$x_{ijk} + x_{i'j'k} \leq 1 \qquad \forall i, i', j, j' \mid e_{jj'} \leq s_{ii'};\ \forall i, i' \in \mathbb{U}_k^L,\ \forall j, j' \in \mathbb{P},\ \forall k \in \mathbb{L} \tag{24}$$

Oversized ULDs. Some oversized ULDs do not fit into one position. In this case, several positions are exactly or partially combined in order to form larger ones. In order to introduce these specific ULDs into our model, we fictively divide each larger ULD into two smaller ones. However, we need to be sure that the model assigns the two parts of the larger ULD to two adjacent positions. If we denote by t_i the ULD linked to the ULD i and by $\mathbb{A}_j$ the set of positions adjacent to position j, constraint (25) ensures the right assignments of larger ULDs:

$$x_{ijk} \leq \sum_{\forall j' \in \mathbb{A}_j} x_{t_i j' k} \qquad \forall i \in \mathbb{U}_k^L,\ \forall j \in \mathbb{P},\ \forall k \in \mathbb{L} \tag{25}$$

4.3.3. Adapting basic constraints

Here are the constraints presented in Limbourg et al. (2012) but adapted to the case of multiple legs.

Stability constraints.

$$minCG_k \leq CG_k \leq maxCG_k \qquad \forall k \in \mathbb{L} \tag{26}$$

$$-\bar{D} \leq \sum_{i \in \mathbb{U}_k^L} w_i \Big(\sum_{j \in \mathbb{P}_R} x_{ijk} - \sum_{j \in \mathbb{P}_L} x_{ijk}\Big) \leq \bar{D} \qquad \forall k \in \mathbb{L} \tag{27}$$

Constraints (26) ensures the longitudinal stability by checking that the center of gravity of the aircraft on each leg is within the limits certified by the aircraft manufacturer. More precisely, the boundaries will depend on the weight of the aircraft under different scenarios: the weight at take off with fuel (Take Off Weight), the expected weight for landing (Landing Weight) and the total weight without fuel (Zero Fuel Weight). The three scenarios define three two-dimensional certified areas called feasibility envelopes. Using the expected consumption and the fuel curve relation between the three scenarios, one single lower bound $minCG$ and one upper bound $maxCG$ are determined. A security margin can be added here as mentioned in the previous section. Constraint (27) checks that the lateral balance is within reasonable limits. The difference between the weight allocated to either side of the fuselage centerline should not be too important (not exceed $\bar{D}$). An aircraft laterally unbalanced could force the pilot to adjust the aileron trim tab or to hold a constant aileron control pressure. Both measures will cause more drag, hence more fuel consumption and a lower efficiency.

Possible positions for ULDs.

$$\sum_{i \in \mathbb{U}_k^L} x_{ijk} \leq 1 \qquad \forall j \in \mathbb{P},\ \forall k \in \mathbb{L} \tag{28}$$

$$x_{ijk} = 0 \qquad \forall i \in \mathbb{U},\ \forall j \in \mathbb{P},\ \forall k \in \mathbb{L} \mid U_i \text{ does not fit in } P_j \tag{29}$$

$$x_{ijk} + x_{i'j'k} \leq 1 \qquad \forall i, i' \in \mathbb{U}_k^L,\ \forall j \in \mathbb{P},\ \forall j' \in \mathbb{E}_j,\ \forall k \in \mathbb{L} \tag{30}$$

Constraint (28) states that, for a given leg, a position can accommodate but one ULD. Constraint (29) makes certain that each loaded ULD physically fits in its position. If not, the corresponding assignment is forbidden and the only allowed value for the assignment variable x_{ijk} is zero. These variables are removed during the optimization. For a same aircraft, different configurations of positions are possible. Some positions are larger and are overlaying several smaller ones (see Figure 2). For each larger position j, the set E_j represents all the smaller positions covered by j. If the position j is used, then all positions in the set E_j must be discarded.

No selection among ULDs.

$$\sum_{j \in \mathbb{P}} x_{ijk} = 1 \qquad \forall i \in \mathbb{U}_k^L,\ \forall k \in \mathbb{L} \tag{31}$$

Constraint 31 ensures that all ULDs are accomodated on board.

Weight restrictions. :

$$w_i x_{ijk} \leq \bar{W}_j \qquad \forall i \in \mathbb{U}_k^L,\ \forall j \in \mathbb{P},\ \forall k \in \mathbb{L} \tag{32}$$

$$\sum_{i \in \mathbb{U}_k^L} \sum_{j \in \mathbb{P} | P_j \cap O_a^d \neq \emptyset} x_{ijk} o_{ija}^d \leq \bar{O}_a^d \qquad \forall d \in \mathbb{D},\ \forall a \in \mathbb{O}^d,\ \forall k \in \mathbb{L} \tag{33}$$

$$\sum_{i \in \mathbb{U}_k^L} \sum_{j \in \mathbb{P} | P_j \cap \bigcup_{c=1}^a F_c \neq \emptyset} \sum_{l=1}^{a} x_{ijk} f_{ijl} \leq \bar{F}_a \qquad \forall a \in \mathbb{F},\ \forall k \in \mathbb{L} \tag{34}$$

$$\sum_{i \in \mathbb{U}_k^L} \sum_{j \in \mathbb{P} | P_j \cap \bigcup_{c=1}^a T_c \neq \emptyset} \sum_{l=1}^{a} x_{ijk} t_{ijl} \leq \bar{T}_a \qquad \forall a \in \mathbb{F},\ \forall k \in \mathbb{L} \tag{35}$$

Those constraints are related to the structural design of the aircraft. The first constraint (32) ensures that the weight exerted on each position does not exceed the maximum weight allowed by the position (denoted by $\bar{W}_j$). Constraint (33) gives the combined load limits. $\mathbb{D}$ denotes the set of decks augmented by an artificial deck corresponding to the entire aircraft. Let's imagine that the aircraft is cut in slices of one inch width. There is a specific weight limit for each slice and each deck d in $\mathbb{D}$. Under some uniform distribution assumptions of the weight inside the positions, the one inch slices can be replaced by broader ones. O_a^d denotes the a^{th} slice for deck d and $\bar{O}_a^d$ the maximal weight allowed for this area. Finally, o_{ija}^d represents the proportion of w_i falling in $\{P_j \cap O_a^d\}$. Constraints (34) and (35) are the cumulative load limits when all decks are considered simultaneously . With constraint (34), we are interested in the total weight from the nose to the end of each of the previously defined slices, and with constraint (35), in the total weight from the tail to the beginning of each slice. These consecutive forward and aft slices are denoted by F_a (forward) and T_a (aft). $\bar{F}_a$ (resp. $\bar{T}_a$) is the maximal cumulative weight allowed for the section starting at the nose (resp. the tail) and ending at F_a (resp. T_a). Finally, the variable f_{ija} (resp. t_{ija}) represents the proportion of w_i falling in $\{P_j \cap F_a\}$ (respectively $\{P_j \cap T_a\}$). See Limbourg et al. (2012) for details.

5. Complexity

Let us now give some insights into the complexity of the problem. The main component of the model is the first term of the objective function dealing with balanced loading and fuel consumption. It is shown below that the problem defined by this first part of the objective function is already NP-hard. Moreover, the other two main contributions of this paper, the introduction of several destinations with pickup and delivery and of several doors in the model, also significantly increase the complexity of the problem.

To show that the weight and balance part of the ACLPPD is NP-hard, we start from the Partition Problem. In this problem, the question is to decide whether it is possible to split a set of numbers $\{w_1, \ldots, w_n\}$ in two disjoint subsets so that the sum of elements in each subset is the same. Garey and Johnson (1979) define it more formally in their famous guide to NP-completeness. Let us now interpret the values $\{w_1, \ldots, w_n\}$ as the ULD weights. Let us also consider that the two partition subsets represent the aft and the forward areas of an aircraft. Assigning a value to a subset would therefore be interpreted as assigning the corresponding ULD to a position in this part of the aircraft. By setting each available position in the two parts of the aircraft at exactly the same distance of the ideal location of the center of gravity (CG), we get exactly the same objective for the two problems. For this specific ACLPPD, the CG after loading can only be at the requested location, i.e. that minimizing the fuel consumption, if and only if the total weight in each of the two areas is exactly the same, i.e. if the sums of values for the two subsets are equal. All the other constraints in the ACLPPD can be disregarded by setting appropriate and unbinding values of the constants. The partition problem is therefore a special case of the ACLPPD where all positions are at equidistance of the required CG location. Since the partition problem is NP-hard, it is also the case for the ACLPPD. Intuitively, the partition problem could represent a very robust aircraft (no weight limit) operated by volunteers (β set to zero), and in which the shipment can be stacked (the positions are at equidistance of the ideal CG location and overlaying list is empty).

The complexity of defining the assignments over two legs cannot be analyzed independently of the balanced loading problems. Considering two destinations implies solving simultaneously two (NP-hard) related instances of the balanced loading problem : one for each leg. The minimization of the number of operations (at the intermediate airport) is the link that makes these two problems dependent and that explains why it is far more complex to solve the problem with an intermediate airport than to independently solve two instances of the same size. Indeed, the inputs for the optimization over the second leg, i.e. mainly which positions are available, depend on the set of feasible solutions of the optimization problem restricted to only the first leg. For each feasible solution of the latter, there exists one instance to optimize over the second leg. This drastically increases the global domain space.

Finally, another originality of our work is that unloading can occur through several doors. For each ULD, when settled in a given position, the question therefore is as follows:

which door should it be unloaded through in order to minimize handling operations. This raises the question of the complexity associated with the introduction of the α_j^T and α_j^N in the model and, more operationally, what is the optimal sequence of unloadings? It is not the most complex part of the problem. A naïve algorithm with complexity $O(n^2)$, where n is the number of ULDs, already answers this question. Indeed, since the ULDs cannot collide once unloaded, there exists, for each deck and each lane, at least one position between two doors of the plane for which all the ULDs close to the tail door will leave by this door (or stay on board) and all the ULDs close to the nose door will go through the nose door (or stay on board). For each lane and deck with several doors, all the positions from the nose door to the tail door can be successively consider as a candidate for the partitioning. It globally takes n operations. Starting from this position, go back to the nose door and stop at the first ULD to be delivered at the intermediate airport (i.e. the first ULD that must be unloaded). From there, go on to the nose and count the number of ULDs to be delivered at the final airport, i.e. the sum of n_j^N. Do the same from the pivot position to the tail door to compute $\sum n_j^T$. This is again done in n operations. The best partitioning position, which defines all the α_j^T and α_j^N, is the one minimizing the sum of n_j^N and n_j^T and the complexity of the whole algorithm is $O(n^2)$.

6. Implementation and results

Our mathematical model has been tested on a set of real-world instances provided by TNT Airways, a wholly owned subsidiary of the TNT Express. Their main activity is to provide TNT Express with a air freight network connecting daily all TNT Express locations throughout the world and more specifically in Europe. TNT Express is one of the leading delivery integrators in Europe. The model was implemented in Java and relies on the IBM ILOG CPLEX 12 library (default parameters). Thanks to a graphical interface, it is possible to visualize, for each leg, the loading plans, the position of the CG and the different weight distributions. The tests were carried out on a personal computer (Windows 7, Intel Core i5-2450M, 2.50GHz, 8.00 GB of RAM).

6.1. The case of a Boeing 777

We first present the detailed results for a historical flight. This is an intercontinental one with a first leg of about 2740 nautical miles and a second leg of 3200 nautical miles. The aircraft is a Boeing 777 Freighter. The B777 is the successor of the famous B747, one of the most used freighters in the world. As illustrated in Figure 2, this Boeing 777 has four doors (one per compartment) and a rather large number, 114, of predefined positions. 40 predefined positions are on the main deck. This corresponds to 13 large positions overlapping 26 small ones (defined on the right and left sides of the aircraft) and one last central position at the rear. On the lower deck, 10 large positions (P) overlap 32 small ones (R and L). There is one more central position (C) for each couple of right and left positions.

The B777F was loaded at full capacity. The sets of ULDs are given in Table 1. For this shipment, we computed and set the optimal CG location for the first leg (resp. second leg) at 39.3% MAC (resp. 38.3% MAC) but we reduced it by 1% as a security margin. 48 tons of fuel were filled into the tanks for each leg. The cost c^h of an unnecessary operation is fixed at 40USD. Assuming an increase of 2% of fuel consumption for a 10% MAC shift of the CG location and a cost of one USD per fuel liter, we get an approximate fuel cost coefficient c_k^f of 100USD. The optimal solution found by our software is depicted in

Set	Origin	Destination	# ULDs
$\mathbb{U}_1$	Liège	DES1	19
$\mathbb{U}_3$	Liège	DES2	24 (1 large one)
$\mathbb{U}_2$	DES1	DES2	9
$\mathbb{U}_1^L$	Leg 1 (Liège-DES1)		43 (1 large one)
$\mathbb{U}_2^L$	Leg 2 (DES1-DES2)		33 (1 large one)

Table 1: Data for our main case

Figures 5 and 6, respectively for the first and second leg. All the positions are represented by boxes. When a ULD has been assigned to a position, the box is colored and the type as well as the weight of the ULD is indicated on the box. In Figure 5, the ULDs in light (resp. dark) gray belongs to $\mathbb{U}_1$ (resp. $\mathbb{U}_3$). This solution meets all requirements. The aircraft has a lateral weight imbalance well below the trehsolds: 235Kg for the first leg and 1527Kg for the second one. On the same figures, the two rectangles situated below each loading plan provide the level of combined weights. The lines situated above are the thresholds. The cumulative load limits are also complied with. For each leg, the three CGs (ZFW-TOW-LW) lie within their respective feasibility envelopes as depicted in Figure 7.

We first measure the quality of the solution by looking at the two terms of the objective function, i.e. the CG location and the number of rehandling operations. The solution found by our software is the best we could hope for. The requested CG location, minimizing fuel, is reached for each of the legs ($\epsilon_k \simeq 0$) and no additional rehandling operations are required

Figure 5: Optimal loading plan obtained by our model for the first leg

Figure 6: Optimal loading plan obtained by our model for the second leg

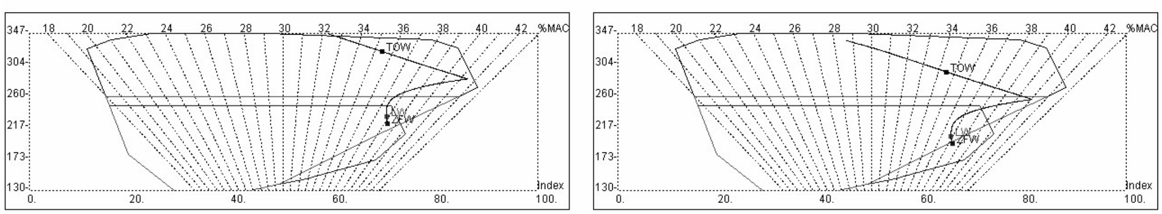

Figure 7: Feasibility envelopes as defined by the aircraft manufacturer

at the intermediate airport. This last fact clearly appears on Figures 5 and 6 since all the ULDs that must be unloaded at the intermediate airport are assigned positions close to the doors and the other ULDs keep the same positions over the two legs.

The second performance measure is the computation time. We limit the computation precision to one USD. The CPLEX optimizer stops as soon as it either finds the optimal solution or a feasible one with an objective value, i.e. the excess cost, of no more than one USD. We know that the problem is NP-Hard and that solving it could be time consuming. We indeed needed sixteen hours to get the optimal solution for this instance. However, this seems to be an extreme case and we usually obtain results in a few minutes (see the additional results in the next section). We believe that at least two reasons can explain why this specific instance remains difficult. First, the aircraft is loaded at full capacity on the first leg, which is rather seldom. Secondly, the number of positions to consider is very large due to the number of different types of ULDs that are involved. Just removing one ULD leads to an optimal solution achieved in only four minutes. It is also worth noting that the Branch and Cut process found quickly good feasible solutions but spent a lots of time to validate the optimal one. We therefore also provide in the Tables the best result obtained within a time limit of 10 minutes . Finally, another interesting insight is that optimizing independently over the two legs, as done in Limbourg et al. (2012), only takes a few seconds. Considering several legs with P&D operations significantly increases the complexity of the problem.

	Load Master	Our solution	With time limit
%MAC ZFW (leg1)	24.4	38.3	38.2
%MAC ZFW (leg2)	31.1	37.3	37.2
# unloadings at DES1	43	19	22
# loadings at DES1	33	9	12
# ULDs $\in \mathbb{U}_3$ unloaded at DES1	24	0	3
Computation Time	20 minutes	16 hours	10 minutes

Table 2: Loading into a Boeing 777 aircraft: main results

The last performance measure is to approximate the potential savings. For that, we compare our solution with the one obtained by the loadmaster working by hand. The loadmasters of the different companies we have met all follow similar procedures. They try, when time permits, to load the cargo in such a way that the observed CG lies close to a predefined value in the feasibility envelope, e.g. at 28% or 29% of the MAC range. The main results are summarized in Table 2. To save fuel, our approach pushes the leg CGs as much as possible to the aft, well beyond the values usually considered. Using our initial assumptions, it would imply a fuel saving of 2010USD. Moreover, no ULD in transit have to be unloaded at the first airport while, as observed in lots of companies, it is common to plan independently the two legs and to be required to unload all or a large part of the cargo. With respect to the worst case, this saves 24 unnecessary operations,i.e. about 960USD. The total savings for this single trip sums up to 2970USD. When we limit the computation time to only 10 minutes we still reach a very good feasible solution. The total savings in this case sum up to 2830USD. Let's imagine that this flight operates under the same conditions three times a week, every week and in both directions, the savings would amount up to 650 000USD per year. Moreover, the same optimization process may be applied to all the other aircraft of the fleet.

6.2. Additional Cases

The question now is to check whether the results for the specific case presented in the previous section are representative. To provide a partial answer, we solved hundreds of other real-world cases. We still consider the same aircraft but with different loads and different sets of ULDs. We provide the results for eight configurations in Table 3. All other simulations lead to similar results. Cases (A) and (B) imply Pickup and Delivery Operations. No delivery but only pickups occur with cases (C), (D) and (E) and the opposite with cases (F), (G) and (H). We set a computation time limit of 10 minutes. The optimal CG location is reached for all the instances, implying fuel savings. Case (H) is the only one stopped after 10 minutes of computation. The CG is at the optimal location but we observe 8 rehandlings of ULDs in transit. This remains better than the solution observed in practice for which 23 ULDs in $\mathbb{U}_3$ are moved at the first airport.

	P&D		Pickup			Delivery		
	A	B	C	D	E	F	G	H
$\lvert\mathbb{U}_1\rvert$	22	20	0	0	0	24	18	19
$\lvert\mathbb{U}_3\rvert$	15	21	23	22	14	14	23	23
$\lvert\mathbb{U}_2\rvert$	14	10	14	20	22	0	0	0
Status	O	O	O	O	O	O	O	F
ϵ %MAC	0	0	0	0	0	0	0	0
# unloadings	0	0	0	0	0	0	0	8
Comput. time	1'15"	6'58"	2'04"	5'15"	54"	3'26"	5'58"	10'

Table 3: Additional results for different cases of ULDs to load

We also randomly generated some hazardous goods among the ULDs loaded and we considered some cases in which a nose door was available. Since these experiments did not provide different results and insights, we have decided not to include them in this paper.

7. Conclusion

In this paper, we have analyzed the *Airline Container Loading Problem with Pickup and Delivery* (ACLPPD). This is a crucial problem encountered every day by airlines. We have considered trips made of several legs and at the end of which P&D operations may occur. We have proposed a new mixed integer linear model.

Our contributions are multiple. First, the model is based on international standards and is valid for most of the commercial operators. We have integrated, and adapted to the multi-leg context, a large set of the constraints they face. Most of the operators should be able to use this approach in real life and, if needed, to extend it, to any of their specifications. Second, we showed how to take care of the loading/unloading sequences when Pickup and Delivery arises. In this context, we again tried to keep close to reality by considering aircraft with several doors. Third, we showed that the weight and balance problem is NP-hard. Considering several legs and several doors make the problem even more complex to handle. Finally, another originality of our approach was to focus on the costs. We have analyzed two important costs directly linked to the loading of ULDs: the impact on fuel consumption and the cost of handling operations. We showed that locating the center of gravity closer to the aft should be done to decrease fuel consumption.

Our approach was tested on real data and we conducted hundreds of experiments. It appeared that it was possible to find quickly optimal or near-optimal solutions and that our approach leads to substantial savings with respect to current typical practices.

Applicability TNT Airways has been and is still involved in this project. The Air France operations research department started this year a study to show if the findings of this research can be applied to their container loading process. Their first results show that the fuel consumption can be reduced by moving the center of gravity backward.

Acknowledgements This work was initiated as part of a research project with TNT Airways. Special thanks are due to E. Meyer, M. Clety, H. Marchal and JM. Urbani. We would also like to thank T. Kleyntssens for his preliminary work on some constraints. Virginie Lurkin is supported by the "Fonds National de la Recherche Scientifique" (FNRS). This work was also partially funded by the Interuniversity Attraction Poles Programme initiated by the Belgian Science Policy Office (grant P7/36). The paper, however, only expresses the authors' views.

References

Airbus Fuel Economy Material, 2004. Flight operations support & line assistance — getting to grips with fuel economy , `http://www.iata.org/whatwedo/Documents/fuel/airbus_fuel_economy_material.pdf`.

Altarazi, S., 2013. A new prioritizing-stacking heuristic algorithm for the inner-city truck loading problem. International Journal of Business and Management 8, 137–143.

Ambrosino, D., Anghinolfi, D., Paolucci, M., Sciomachen, A., 2010. An experiment comparison of different heuristics for the master bay plan problem. Lecture Notes in Computer Science 6049, 314–325.

Ambrosino, D., Bramardi, A., Pucciano, M., Sacone, S., Siri, S., 2011. Modeling and solving the train load planning problem in seaport container terminals, in: Proceedings of the 7th IEEE Conference on Automation Science and Engineering, pp. 208–213.

Amiouny, S.V., Bartholdi, J.J., Vande Vate, J.H., Zhang, J., 1992. Balanced loading. Operations Research 40, 238–246.

Avriel, M., Penn, M., Shpirer, N., Witteboon, S., 1998. Stowage planning for container ships to reduce the number of shifts. Annals of Operations Research 76, 55–71.

Bortfeldt, A., Wäscher, G., 2013. Constraints in container loading — a state-of-the-art review. European Journal of Operational Research 229, 1–20.

Bostel, N., Dejax, P., 1998. Models and algorithms for container allocation problems on trains in a rapid transshipment shunting yard. Transportation Science 32, 370–379.

Bruns, F., Knust, S., 2011. Optimized load planning of trains in intermodal transportation. OR Spectrum 34, 511–533.

Chan, F.T.S., Bhagwat, R., Kumar, N., Tiwari, M.K., Lam, P., 2006. Development of a decision support system for air-cargo pallets loading problem: A case study. Expert Systems with Applications 31, 472–485.

Chan, F.T.S., Kumar, N., 2006. A new heuristic embedded approach for multi-constraint air-cargo loading problem, in: Proceedings of the 4th International IEEE International Conference on Industrial Informatics, pp. 1165–1170.

Corry, P., Kozan, E., 2006. An assignment model for dynamic load planning of intermodal trains. Computers & Operations Research 33, 1–17.

Dubrovsky, O., Levitin, O.G., Penn, M., 2002. A genetic algorithm with a compact solution encoding for the container ship stowage problem. Journal of Heuristics 8, 585–599.

Fok, K., Chun, A., August 14-17, 2004. Optimizing air cargo load planning and analysis, in: Proceeding of the International Conference on Computing, Communications and Control Technologies, Austin, Texas, USA.

Garey, M., Johnson, D., 1979. Computers and Intractability; A Guide to the Theory of NP-Completeness.

Guèret, G., Jussien, N., Lhomme, O., Pavageau, C., Prins, C., 2003. Loading aircraft for military operations. Journal of the Operational Research Society 54, 458–465.

Heidelberg, K.R., Parnell, G.S., Ames, J.E., 1998. Automated air load planning. Naval Research Logistics 45, 751–768.

IATA, 2013a. Cargo e-chartbook. quarterly report, `http://www.iata.org/whatwedo/Documents/economics/eChartbook-Q3-2013.pdf`.

IATA, 2013b. ULD Regulations (ULDR), `http://www.iata.org/publications/Pages/uld-regulations.aspx`.

Imai, A., Sasaki, K., Nishimura, E., Papadimitriou, S., 2006. Multi-objective simultaneous stowage and load planning for a container ship with container rehandle in yard stacks. European Journal of Operational Research 171, 373–389.

Kaluzny, B.L., Shaw, D.R., 2009. Optimal aircraft load balancing. International Transactions in Operational Research 16, 767–787.

Li, F., Tian, C., Cao, R., Ding, W., 2008. An integer linear programming for container stowage problem, in: ICCS '08: proceedings of the 8th international conference on computational science, part I. Berlin, Heidelberg: Springer, pp. 853–862.

Li, Y., Tao, Y., Wang, F., 2009. A compromised large-scale neighborhood search heuristic for capacitated air cargo loading planning. European Journal of Operational Research 199, 553–560.

Limbourg, S., Schyns, M., Laporte, G., 2012. Automatic aircraft cargo load planning. Journal of the Operational Research Society 63, 1271–1283.

Mathur, K., 1998. An integer-programming-based heuristic for the balanced loading problem. Operations Research Letters 22, 19–25.

Mongeau, M., Bès, C., 2003. Optimization of aircraft container loading. IEEE Transactions on Aerospace and Electronic Systems 39, 140–150.

Nance, R.L., Roesener, A.G., Moore, J.T., 2011. An advanced tabu search for solving the mixed payload airlift loading problem. Journal of the Operational Research Society 62, 337–347.

Ng, K.Y.K., 1992. A multi-criteria optimization approach to aircraft loading. Operations Research 40, 1200–1205.

Øvstebø, B.O., Hvattum, L.M., Fagerholt, K., 2011. Optimization of stowage plans for RoRo ships. Computers & Operations Research 38, 1425–1434.

Pan, L., Chu, S., Han, G., Huang, J.Z., 2011. A heuristic algorithm for the inner-city multi-drop: Container loading problem. International Journal of Operations Research and Information Systems 2(3), 1–19.

Pollaris, H., Braekers, K., Caris, A., Janssens, G.K., 2013. The capacitated vehicle routing problem with loading constraints, in: Proceedings of HMS: The International Conference on Harbour, Maritime & Multimodal Logistics Modelling and Simulation.

Tang, C.H., 2011. A scenario decomposition-genetic algorithm method for solving stochastic air cargo container loading problems. Transportation Research Part E: Logistic and Transportation Review 47, 520–531.

Tang, C.H., Chang, H.W., 2010. Optimization of stochastic cargo container loading plans for air express delivery, in: IEEE Second International Conference on Computer and Network Technology, Bangkok. pp. 416–420.

Tian, C., Zhang, H., Li, F., 2009. Air cargo load planning, in: IEEE/INFORMS International Conference on Service Operations, Logistics and Informatics, pp. 454–459.

Verstichel, J., Vancroonenburg, W., Souffriau, W., Vanden Berghe, G., 2011. A mixed integer programming approach to the aircraft weight and balance problem. Procedia Social and Behavioral Sciences 20, 1051–1059.

Wu, Y., 2010. A dual-response forwarding approach for containerizing air cargoes under uncertainty, based on stochastic mixed 0-1 programming. European Journal of Operational Research 207, 152–164.

Yan, S., Shih, Y.L., Shiao, F.Y., 2008. Optimal cargo container loading plans under stochastic demands for air express carriers. Transportation Research Part E 44, 555–575.