

Finding a robust assignment of flights to gates at Amsterdam Airport Schiphol

G. Diepen · J.M. van den Akker · J.A. Hoogeveen · J.W. Smeltink

Published online: 18 October 2012
© Springer Science+Business Media New York 2012

Abstract In this paper we investigate the gate assignment problem as it appears at Amsterdam Airport Schiphol (AAS). Currently, the gate planners spend many hours on adjusting the automatically generated planning during the day of operation to make it proof against small deviations from the schedule. To alleviate this problem, we aim at finding a *robust* solution, given the planned arrivals and departures for the next day.

We present a completely new integer linear programming formulation that is based on so-called gate plans. Each gate plan consists of a subset of the flights that can be assigned to a single gate of the corresponding type; gates with identical characteristics are aggregated in gate types. The gate assignment problem then boils down to selecting the best subset of gate plans such that each flight belongs to one selected gate plan, and such that the number of selected gate plans for a

certain type of gate is equal to the number of gates of this type. In the first phase, we solve the LP-relaxation through column generation, and we describe specific features to find a very good solution to the ILP quickly. This solution is then handed to the planners at AAS in order to assign gate plans to physical gates. This consists of a number of relatively small problems that can be solved by hand and in which additional operational constraints can be incorporated. We also present the possibility of directly assigning flights to physical gates using the column generation formulation, where we then take into account other criteria as well.

Computational results with real-life data provided by AAS are promising and indicate that the algorithm is able to solve real-life instances within rather small running times.

Keywords Airport planning · Column generation · Integer linear programming

Supported by BSIK grant 03018 (BRICKS: Basic Research in Informatics for Creating the Knowledge Society). The research was performed while G. Diepen was at Utrecht University.

G. Diepen
Paragon Decision Technology, Schipholweg 1, 2034 LS Haarlem, The Netherlands
e-mail: Guido.Diepen@aimms.com

J.M. van den Akker (✉) · J.A. Hoogeveen
Department for Information and Computing Sciences, Utrecht University, P.O. Box 80089, 3508 TB Utrecht, The Netherlands
e-mail: J.M.vandenAkker@uu.nl

J.A. Hoogeveen
e-mail: J.A.Hoogeveen@uu.nl

J.W. Smeltink
National Aerospace Laboratory NLR, Anthony Fokkerweg 2, 1059 CM Amsterdam, The Netherlands
e-mail: smeltink@nlr.nl

1 Introduction

Between the time an aircraft lands at an airport and the time it departs again many things must happen. One of the most obvious things is that the passengers need to disembark the aircraft. Moreover, the aircraft needs to be refueled, new passengers need to board it, new supplies have to be put on board, and the aircraft has to get cleaned. These latter two actions are taken care of by the so-called *ground handler*. All of these actions take place while the aircraft is standing at a gate. We will refer to the arrival of an aircraft till the following departure of the same aircraft as a *flight*. Note that we have no freedom of coupling incoming and outgoing flights to aircraft; this is input.

Now the *gate assignment problem* is described as follows: assign a given set of flights to a (smaller) set of gates

while making sure that certain criteria are met. Examples of such criteria are:

- Gates can handle only flights operated by aircraft of certain sizes.
- Gates can handle only flights for certain origins/destinations (e.g. because of safety regulations).
- Gates can handle only flights assigned to certain ground handlers.
- Two adjacent gates cannot be assigned flights operated by big aircraft at the same time.
- Two adjacent gates should not be assigned flights with equal departure times.

Depending on the airport and its characteristics, many variants of the gate assignment problem have been researched and many solution methods have been suggested. A good overview of explored methods and models is given in Van Orden (2002). In this paper we consider the situation at Amsterdam Airport Schiphol (AAS). Depending on the time horizon of the planning we distinguish between the following three planning problems: *seasonal planning*, *daily planning*, and *tactical planning*. The seasonal planning problem is a capacity planning problem. In this problem the gate planners must decide to accept or decline new requests from airlines to have their aircraft visit AAS. The daily planning concerns the creation of a planning for the upcoming day on the basis of the available information about the flights of that day. Finally, the tactical planning concerns the resolving of conflicts that arise in the planned solution (i.e. the planning created the day before) due to disturbances. Observe that in Air Traffic Management tactical planning takes place within the last few hours before the actual operation of a flight, opposed to e.g. supply chain management where tactical planning is defined as the planning a few months before operation.

The problem we consider in this paper is daily planning. In contrast to the novel variant of *dial-a-flight*, as described in Espinoza et al. (2008a, 2008b), we assume that all flights have already been scheduled. Hence, for each flight we are given the arrival and departure time, the type of plane that carries out the flight, and the ground handler. Given these data, we are asked to find the best solution to the gate assignment problem. In the literature, there is no agreement on the criteria that constitute the objective function; this may also depend on the airport. A criterion occurring in many papers is passenger comfort, which is measured by minimizing the total walking distance (Bihl 1990; Haghani and Chen 1998; Xu and Bailey 2001). Other criteria are the number of flights that cannot be assigned to a gate, the total waiting time for the aircraft (i.e. the time an aircraft has to be held after landing before the gate is free), the total number of split flights (where the flight is towed from the arrival gate to the departure gate), the deviation of a given reference schedule

(which criterion is used to enforce for example that daily flights are assigned to the same gate each day), robustness (which is measured on basis of the buffer time between two consecutive flights on the same gate), and total preference score (each flight-gate combination obtains a score on basis of some preferences issued by the airport authorities). In many papers a linear combination of a number of the above criteria is optimized. For example, Yan and Huo (2001) consider the situation at Chiang Kai-Shek Airport. Here the optimization criteria are a linear combination of total waiting time for the aircraft and total walking distance of the passengers (where the walking distance of transfer passengers is estimated to avoid a quadratic model). Yan et al. (2002) present a simulation framework for testing the resulting solution for the daily planning problem in combination with a couple of heuristics for the tactical planning problem. Dorndorf et al. (2007) present a survey of the gate assignment problem in general and also discuss recent developments with regard to the multi-criteria objectives. In the follow-up paper (Dorndorf et al. 2008) evaluate a solution on basis of three criteria: number of split flights, total preference score, and robustness. The criteria are aggregated in a linear objective function, which is then solved through a heuristic based on the clique partitioning problem; by varying the relative weights of the criteria, an approximation of the Pareto optimal frontier is determined. In another follow-up paper, Dorndorf et al. (2012) consider two further criteria: the number of unassigned flights and the minimal deviation from a reference schedule. The authors show that the approach from Dorndorf et al. (2008) can be applied to find decent solutions for this problem as well. Drexl and Nikulin (2008) look for the Pareto optimal frontier, too. Here the criteria are the number of unassigned flights, total preference score, and walking distance; the latter criterion is quadratic when transfer passengers are included. Splitting of flights is not possible, but the authors have included this feature in Nikulin and Drexl (2010). Here the criteria are the number of split flights, the total preference score, and deviations from a given reference schedule. Again, this is tackled by using Pareto Simulated Annealing, where the possible uncertainty in the input data is treated by using fuzzy numbers.

Although minimizing passenger walking distance seems to be most frequently studied in previous work, this is not a natural objective in the case of AAS. Unlike some other international airports, the gate planning at Schiphol is performed by airport authorities and not by specific airlines who own part of the terminal. As a consequence the gate planners do not have complete passenger data and especially they do not have all transfer patterns. Although the planners consider passenger walking distances by putting flights with many passengers close to the terminal building, they are not in the position to perform a full-blown minimization of passenger walking distance.

The planning software currently in use at AAS is based on a greedy algorithm that assigns flights to gates on the basis of an optimal point score per flight, i.e. the total preference score criterion. This score includes different issues, such as preferences of airlines and ground handlers. An airport is a very dynamic environment and an actual day will hardly ever go completely as planned; flights arrive either earlier or later than planned due to all kinds of reason. Experiences from the gate planners at AAS reveal that currently each day a considerable amount of time (a number of hours) has to be spent on replanning the automatically generated schedule in order to make it more robust, because the current planning software does not consider robustness at all. Moreover, every change needed in the gate assignment during the actual day has an effect on a very broad range of parties: the ground handler, the security personnel, the passengers, etc. Therefore, creating a schedule that needs fewer changes during the tactical planning is of major importance for a variety of parties. For the above reasons we consider *robustness* as the objective function. We are concerned with creating a robust schedule in the sense that the schedule is able to cope with small perturbations (within approximately half an hour) in arrival or departure time without the need to replan big parts of the schedule.

The question is how to measure the robustness of a schedule, without knowing the actual instance, which only gets known during the day of operation. Since the arrival and departure times are subject to uncertainty, we could consider these times as stochastic variables. However, a stochastic optimization model has a much larger computational complexity than a deterministic model. Observe that the ability of the schedule to absorb small perturbations is determined by the probability that a conflict arises between two consecutively scheduled flights at the same gate. The idea of our approach is that, instead of a stochastic model, we use a deterministic model with a cost function that assigns larger cost to schedules in which the probability of conflicts is higher. The probability that a conflict arises between two consecutively scheduled flights at the same gate depends on the ‘reliability’ that the aircraft will stick to their assumed departure and arrival times and on the idle time between these two consecutive flights. The ‘reliability’ is related to the variance of the arrival and departure time. It cannot be influenced by the gate planners directly, and should be derived from historical data. Consequently, we focus on optimizing the choice of the pairs of flights that we put consecutively at a gate; we want to have as much idle time as possible between the consecutive flights, where the preference for larger idle time increases when the reliability of the involved flights becomes smaller. Since the total amount of idle time is constant, we look at the division of the idle times.

One approach for optimizing the individual idle times, and thus finding a robust schedule, is to minimize the variance of the idle time between successive flights at a gate.

This approach is used in Bolat (2000) where the problem is formulated as an Integer Linear Program (ILP) using binary decision variables that link flights to positions at a gate. Based on this research a similar approach was followed in Van Orden (2002) for modeling the gate assignment problem for AAS. Although the suggested model showed promising results, a major drawback of this model was that it was not possible to solve problems with more than 80 aircraft and 20 gates in a reasonable amount of time.

In this paper we present a new formulation for the gate assignment problem which enables us to compute a robust gate assignment schedule for a full day of traffic at AAS, i.e., about 600 flights, within a few minutes. We include all real constraints occurring at the airport that we have identified in discussions with the gate-planners of the airport. To attain robustness, we optimize the idle time between all consecutive flights at the gates by using an objective function that assigns high cost to pairs of consecutive flights leading to high probability of conflicts. Instead of minimizing the variance of the idle times, we develop an objective function based on the arctangent, and we make some adjustments to this function to take the ‘reliability’ of the flights into account. We formulate this problem as an ILP and use a completely different representation from the one used in Van Orden (2002). This different representation is derived from the one used in Freling et al. (2001) for the vehicle and crew scheduling. We present an algorithm based on column generation to find a very good approximation for the optimum in this model. Our experiments indicate that this algorithm is able to solve real-life instances to near-optimality within a few minutes. Because of the high quality of the approximation even a near-optimal solution will suffice and techniques as branch-and-price (Barnhart et al. 1998) to solve the problem to full optimality will not pay off.

The outline for the remainder of this paper is as follows: In Sect. 2 we give a detailed description of the gate assignment problem. In Sect. 3 we formulate the gate assignment problem as an ILP and present our algorithm to find an approximate solution. In Sect. 4 we will present the experimental results, and finally in Sect. 5 we draw some conclusions and indicate some future research topics.

2 Problem description

As mentioned in the previous section, our objective is to maximize the robustness of a solution to the gate assignment problem. To maximize the robustness we maximize the idle time between all pairs of consecutive flights at a gate, ensuring that each flight can arrive either a bit too early or a bit too late without the need for replanning the schedule. An example of a non-robust schedule can be found in Fig. 1, where we assume here that each plane can be assigned to




Fig. 1 Example of a non-robust schedule

both gates. This schedule does not have a lot of margin between flight 2 and flight 3. By modifying the schedule such that flight 3 is assigned to gate 2, while flight 4 gets assigned to gate 1 we introduce a lot more robustness in the solution.

To evaluate the robustness of a schedule more accurately, we take into account whether two consecutive flights assigned to the same gate, are operated by the same airline or share the same ground handler. Such a situation is *convenient*, since then the airline and or the ground handler will have an incentive to make the first flight leave on time. Moreover, some airlines or airports are known to be unreliable, meaning that if a flight of such an airline is due to depart at a certain time, then there is a great chance that it is delayed. In a convenient situation the amount of idle time is less critical, whereas in case of unreliability of the airline or departure airport the amount of idle time is very critical. We will model this by adjusting the cost associated to the amount of idle time between the two flights.

There are several hard constraints in the gate assignment problem. Obviously, each flight must be assigned to a gate, and two flights cannot be assigned to the same gate at the same time. But there are many more hard constraints, concerning the properties connected with the flights and the gates. The properties that are known for each flight are:

- The region of the origin of the flight.
- The region of the destination of the flight.
- The size category of the aircraft performing the flight.
- The ground handler for the flight.

The region can be either Schengen (which refers to the countries that signed the Schengen Agreement), European Union (EU), or Non-EU. Often the region of the origin and the region of the destination of a flight are the same, but there are many exceptions, including e.g. transit flights that do not have AAS as their final destination. With respect to the size category of the flight, there are eight categories at AAS: category 1 for the smallest aircraft up to category 8 for the biggest aircraft. Finally, the ground handlers are divided into two groups at AAS: KLM Ground Services and all other companies. Each group in principle works at its own part of the apron which means that the gates (including platform stands) are divided into gates that are served by KLM Ground Service and gates that are served by others.

For each of the gates it is known which regions and which size categories can be served, and by which ground handler it is operated. When assigning flights to a certain gate, we need to satisfy the following three essential properties:

- The regions of origin and destination of the flight must match the possible regions of the gate.
- The size category of the flight must match the possible size categories of the gate.
- The ground handler of the flight must match the possible ground handlers of the gate.

One important issue of the gate assignment problem is that some of the flights stay at AAS for a longer period of time; for example, they arrive early in the morning and leave again late in the afternoon. If there are many such *long-stay* flights that stay at the gate, then the number of available gates quickly decreases. In case the capacity at the gates is not enough to handle all flights, the gate planners at AAS have the possibility of *splitting* the stay of long-stay flights into three different parts. According to operational rules used at AAS this proceeds as follows:

- *Arrival part.* After the aircraft lands, it will stay at the gate for 65 minutes, after which it is towed to some buffer stand.
- *Intermediate part.* During this part the aircraft resides on a buffer stand, where it does not use precious gate capacity.
- *Departure part.* The aircraft is taken from the buffer to the appropriate gate, 95 minutes before the aircraft will depart.

At AAS only flights that stay longer than three hours are considered for such a splitting. The advantages of splitting long-stay flights are three-fold. First, there is the obvious extra capacity that becomes available for the assignment of other aircraft. The second advantage concerns flights with different regions for the origin and destination: such flights normally would have to be assigned to a gate that is multi-purpose with respect to the region property, and these gates are quite scarce. When such a flight is split into parts, the separate parts do not have to be assigned to the same gate and thus can be assigned to two single-region gates. Third, the decoupling of the parts yields additional flexibility.

Currently the process of splitting the long-stay flights is done manually. First the gate planners try to solve the gate assignment problem without splitting any flights. If the available capacity is not sufficient to accommodate all flights, the gate planners determine which flights should be split and then solve the problem again. This step is repeated several times until sufficient capacity is available.

3 Solution approach

Recall that our main objective is robustness, which depends on the amount of idle time between two consecutive flights at the same gate. To value the idle time between two consecutive flights, Bolat (2000) considers minimizing the variance of the idle time. Here the idle times before the first




Fig. 2 Cost function to value idle time

flight and after the last flight on a gate are taken into account as well, which implies that the total idle time is a constant. As a result, minimizing the variance becomes equivalent to minimizing the sum of the squared idle times.

A disadvantage of minimizing the total squared idle time is that it looks counter-intuitive: large idle times are penalized heavily, whereas small idle times get low cost. As a consequence, it becomes hard to adjust the cost in case of a convenient pair of consecutive flights as discussed before. Therefore, we have defined a new cost function. An important criterion here was that it should penalize solutions with large conflict probability. Moreover, it should mimic the appreciation of a solution by a planner: the solution that is liked the most by the planner should be the one that gets the minimum score. Our cost function is based on the cost for idle time, with an adjustment for convenience. The function that we use to measure the cost of idle time should possess the following characteristics. First of all, it should penalize very small idle-times with very high cost and it should only mildly penalize rather large idle times. Second, the function should be steep in the beginning (for small idle times) and then flatten out, to reflect that for small idle times any improvement is very beneficial, whereas for already large idle times an extra increase is of minor importance. Third, it should be possible to combine this with a refinement reflecting the *convenience* of certain combinations and the *unreliability* of certain flights, as mentioned in the previous section. We found that a cost-function based on the arctangent fulfills the desired properties best. This function is defined as

$$c(t) = \begin{cases} 1000(\arctan(0.21(5 - t)) + \frac{\pi}{2}) & \text{if } t \geq 20 \\ \infty & \text{otherwise} \end{cases}$$

where t is the amount of idle time. The lower bound of 20 minutes follows the regulations at AAS. This function is depicted in Fig. 2. Initial experiments showed that it was not beneficial to make use of a *cut-off value* (i.e. a threshold after which any increase in idle time will not result in a decrease of the costs). Such a cut-off value resulted in longer running times for the ILP, which can be explained by the fact that the cut-off value introduces a lot of symmetry in the problem.

Recall from the previous section that there could be an advantage (or disadvantage) in assigning certain pairs of flights consecutively to the same gate. Since we want to mimic the appreciation that the planners have for a solution, we adjust the cost function for each pair of flights v and w , if necessary. Hereto, we introduce a convenience multiplier $conv(v, w)$ for flights v and w . To compute the cost of placing flight w immediately after flight v at the same gate, the cost corresponding to the idle time in between these flights will be multiplied by this multiplier. If putting flight w after flight v is desirable (e.g. flight v and w are operated by the same airline or handled by the same ground handler), then the convenience multiplier is given a value less than 1, thus decreasing the cost. On the other hand, inconvenient situations (e.g. when flight v is operated by an unreliable airline) can be penalized by giving the multiplier for these situations a value greater than 1, thus increasing the cost of such an assignment. In deliberation with the planners we define some criteria that we use in the computation, like for example “do the planes belong to the same airline and/or ground handler?”, “what is their reliability record?”, etc. On basis of this, we determine for each possible outcome a suitable value of the convenience multiplier. As a consequence, the convenience multiplier can be defined beforehand for all possible pairs of successive flights, while during run-time only one additional multiplication per pair of consecutive flights is needed.

Since instances of the complete gate assignment problem consist of around 600 flights and 120 gates, we would like to see if we can decompose the problem into a set of smaller subproblems based on one of the properties associated with each gate (i.e. the possible regions, sizes, and ground handlers). Unfortunately for each of those properties gates exist that are multi-purpose for that particular property. This means that a strict division into multiple independent subproblems based on any of the properties is not possible without discarding available capacity: a given multi-purpose gate can only be assigned to one of the subproblems. Hence, such a static division is not desired.

The standard approach for modeling the problem as an ILP uses variables that denote whether a flight is assigned to a certain gate, see Van Orden (2002). One major disadvantage of this kind of model is that many additional variables are needed to determine the order of the flights, which is necessary to compute the idle time between two consecutive flights and hence the cost of a solution. Moreover, the number of constraints grows rapidly when relations between multiple gates and flights are present in the problem. This rapid growth of both variables and constraints makes it impossible to solve a gate assignment problem of the size that occurs at AAS. Therefore, a new approach is needed.

Based on the approach used by Freling et al. (2001) for the single-depot vehicle scheduling problem, we split

the gate assignment problem into two phases. In the first phase we aggregate gates with the same properties (*identical gates*) into groups of gates and each such group we refer to as a *gate type*. A certain gate type is characterized by at least the three essential properties: origin and destination area, ground handler and size category, which have to be equal for two gates to belong to the same gate type. The model has the flexibility to make the gate type more specific; for example flights of the airline El Al have to be handled at specific gates. We now introduce a *gate plan* as a series of flights that are to be assigned to the same gate of a certain gate type. Gate plans enable us to check feasibility easily: all flights present in the gate plan must satisfy the properties of the gate type that the gate plan corresponds to and between each pair of consecutive flights there should be an interval of at least 20 minutes. From now on, we define a gate plan to be feasible if it satisfies these two constraints. The cost of a feasible gate plan is equal to the cost of all idle times between consecutive flights within the gate plan. With this representation solving the first phase comes down to finding a set of gate plans such that we have a gate plan for each physical gate, and such that each flight is present in exactly one of the gate plans.

After we have obtained a solution for the first phase, for each gate type we have as many gate plans as there are gates of that gate type. We guarantee that all ‘single gate’ constraints concerning separation time, and concerning the correct gate type with respect to at least the size of the flight, origin–destination area and ground handler are met. In the second phase we undo the aggregation of identical gates into groups of gate types and we have to assign each gate plan to a physical gate. Here we consider ‘multiple gate’ constraints, such as that it is impossible to have simultaneous push-backs at opposite gates.

Besides the two step approach, we can also directly assign flights to gates. This is achieved by defining a separate type for each single physical gate. This allows us to include more details on the cost or benefits of assigning certain flights to certain gates, e.g. putting a flight with a large number of passengers far away from the main part of the terminal leads to an additional penalty.

We further want to remark here that our solution approach is independent from the cost function that we developed: it can be used for any cost function that is based on the idle times between consecutive flights. Suppose for example that we would know the exact probability distribution of the arrival and departure times, which allows us to compute the probability that putting flights i and j consecutively on a gate will lead to a conflict between these two flights. Then we can compute the cost of a gate plan as the expected number of conflicts for the flights involved. Hence, we can then use our approach to minimize the expected number of conflicts (where it might be wise to use a threshold probability

that must be respected for each pair of consecutive flights to avoid ending up with some gates in which all ‘difficult-to-schedule’ flights are accumulated). Moreover, we can easily work with reward functions for assigning flights to a given gate-type.

3.1 The ILP formulation

When we look at the number of possible gate plans, we can clearly see that the number of possible gate plans is enormous. Suppose for the moment that we do have the complete set of all possible gate plans. Then the gate assignment problem can be formulated as determining which of these gate plans we must select. Now for each gate plan i we introduce a decision variable x_i as follows:

$$x_i = \begin{cases} 1 & \text{if gate plan } i \text{ is selected} \\ 0 & \text{otherwise} \end{cases}$$

We define c_i as the cost of gate plan i which can easily be computed from the idle times between consecutive flights. Let V denote the number of flights, let A denote the number of gate types, let S_a denote the number of gates of type a , and let N denote the number of gate plans. Now the basic model for the gate assignment problem is as follows:

$$\text{Minimize } \sum_{i=1}^N c_i x_i$$

subject to

$$\sum_{i=1}^N g_{vi} x_i = 1 \quad \text{for } v = 1, \dots, V \quad (1)$$

$$\sum_{i=1}^N e_{ia} x_i = S_a \quad \text{for } a = 1, \dots, A \quad (2)$$

$$x_i \in \{0, 1\} \quad \text{for } i = 1, \dots, N \quad (3)$$

where

$$g_{vi} = \begin{cases} 1 & \text{if flight } v \text{ is in gate plan } i \\ 0 & \text{otherwise} \end{cases}$$

$$e_{ia} = \begin{cases} 1 & \text{if gate plan } i \text{ is of type } a \\ 0 & \text{otherwise} \end{cases}$$

We will extend this basic model to cover more possibilities. One important issue not addressed by the above model is the fact that it should be possible to promote placing a flight at a certain gate type. We will add constraints to model that certain flights are preferably assigned to certain gate types, e.g. because of the size of the waiting area or because of airlines having their ‘own’ gates where at least a certain percentage of their flights have to be assigned to. Such an additional constraint consists of a set of flights, a set of gate

types, and the minimum number of flights out of the given set that should be assigned to the given set of gate types.

Preferences can be modeled in the ILP by adding the following constraints to the model:

$$\sum_{i=1}^N \sum_{v=1}^V \sum_{a=1}^A p_{vak} e_{ia} g_{vi} x_i \geq P_k \quad \text{for } k = 1, \dots, K \quad (4)$$

where

$$p_{vak} = \begin{cases} 1 & \text{if flight } v \text{ has preference for gate type } a \\ & \text{in preference } k \\ 0 & \text{otherwise} \end{cases}$$

P_k denotes the minimum number of flights that have to be assigned to a given gate type according to preference k , e.g. an airline must have at least six flights at their ‘own’ gates. K denotes the total number of preferences.

We further extend the above model to deal with the case that there is not enough capacity to accommodate all flights. To solve this problem we add a penalty variable UAF_v (unassigned flight) for every flight v to constraint (1) in the following way:

$$\sum_{i=1}^N g_{vi} x_i + UAF_v = 1 \quad \text{for } v = 1, \dots, V \quad (5)$$

$$UAF_v \geq 0 \quad \text{for } v = 1, \dots, V$$

The extra variable UAF_v for flight v is added with a very high cost coefficient Q_v in the objective function. Now it is possible to have flights not being assigned to gates, but the cost of this option are so large it only occurs when it is infeasible to assign more flights to gates. The existence of these unassigned flights in the final solution is then reported to the planners. In contrast to our computer program, the planners are allowed to overrule some of the constraints, if necessary. Then, the planners can either make the solution feasible by manually planning the unassigned flights, or change the input according to their breaking of the rules, after which the program can be run again, starting from the final solution of the previous run. We can select the cost coefficients of the UAF_v variables to model the effort it takes to assign a flight manually; for example, in general it is easier to assign a flight operated by a small aircraft somewhere manually than a flight operated by a big aircraft. The values for these cost coefficients were determined via preliminary computations.

One thing not addressed in the model yet is the fact that flights with a long stay can be split into three parts: an Arrival part, an Intermediate part, and a Departure part. To model this possibility, for each long-stay flight v we create two *split flights* v_A and v_B , which refer to the arrival and departure part of flight v , respectively. The intermediate parts

of the flights are not modeled because the buffer stands for these intermediate parts are not part of the gate assignment problem.

Since these three flights v , v_A , and v_B concern the same flights, we must ensure we model their dependency. This can be achieved by splitting the original constraint (1) for flight v into two separate constraints

$$\sum_{i=1}^N (g_{vi} + g_{v_A,i}) x_i + UAF_{v_A} = 1 \quad \text{and}$$

$$\sum_{i=1}^N (g_{vi} + g_{v_B,i}) x_i + UAF_{v_B} = 1$$

Here UAF_{v_A} and UAF_{v_B} indicate the possibility of not assigning (a part of) flight v ; their cost coefficients each get a value half of the value of the original cost coefficient UAF_v .

Furthermore, we have to include the split flights in the additional constraints (4). Since each split flight only represents half of the original flight, each gets a coefficient 0.5 by redefining p_{vak} as follows

$$p_{vak} = \begin{cases} 1 & \text{if flight } v \text{ has preference on gate type } a \\ & \text{in preference } k \\ 0.5 & \text{if the split version of flight } v \text{ has} \\ & \text{preference on gate type } a \text{ in preference } k \\ 0 & \text{otherwise} \end{cases}$$

The ILP formulation presented above models the problem correctly, but unfortunately the size of the problem is enormous, since the number of possible gate plans is enormous. To reduce the size of the problem, we only take a subset of the gate plans into account; when we solve the ILP problem for this subset, we find an approximate solution. In our selection of the gate plans, we try to identify gate plans that we expect to have some chance of getting selected in the optimum solution; we call these *presumably useful* gate plans. To identify these, we solve a relaxation of the problem; we assume here that gate plans that show up in the solution process of this relaxation will be reasonable candidates for the original ILP problem as well. The relaxation we consider is the LP-relaxation: thereto, we relax the integrality constraints (3) and solve the LP-relaxation through column generation. Eventually, we reinstate the integrality constraints and solve the resulting ILP formulation with the columns generated. As a side-effect, we can use the value of the LP-relaxation as a measure for the quality of the obtained ILP formulation.

3.2 Pricing problem

We start by generating a set of columns that constitute a feasible solution to the LP-relaxation. After we have solved the

LP-relaxation for a given set of columns, we find a dual multiplier π_v for the constraint (1) corresponding to flight v , a dual multiplier λ_a for the constraint (2) corresponding to type a , and a dual multiplier ψ_k for the constraint (4) corresponding to preference k . Therefore, the reduced cost for a gate plan i is equal to

$$c_i - \sum_{a=1}^A e_{ia} \lambda_a - \sum_{v=1}^V \left(g_{vi} \pi_v + \sum_{k=1}^K \sum_{a=1}^A g_{vi} e_{ia} p_{vak} \psi_k \right)$$

Note that for the original parts of a long-stay flight the above must be slightly changed. Since the original part of a long-stay flight is present in two constraints, we must subtract the two dual multipliers π_{v_A} and π_{v_B} instead of only π_v .

It is well-known from the theory of linear programming that we have solved the LP-relaxation to optimality if the reduced cost of each gate plan is greater than or equal to zero. To check this, we compute the minimum reduced cost over all feasible gate plans; this is called the *pricing problem*. We solve this problem by composing a network such that each feasible gate plan corresponds to a path in this network, and vice versa. Moreover, we choose the lengths of the arcs such that the length of a path equals the reduced cost of the corresponding gate plan. Hence, we can then solve the pricing problem by solving a shortest-path problem in this network.

We solve the pricing problem for each gate type separately. For each type of gate a we introduce a directed acyclic graph $G_a = (V_a, E_a)$. We add a vertex to this graph for every flight v that is allowed to be assigned to a gate of type a . Furthermore, we add vertices s and t , denoting the source and sink respectively. If two flights v and w are allowed on a gate of type a and the arrival time T_w^{arr} of flight w is greater than or equal to the departure time T_v^{dep} of flight v plus the minimum idle time T_v^{min} required after flight v (at AAS this is assumed to be 20 minutes independent from the flights v and w), then a directed edge from vertex v to w is added to the graph. Furthermore, a directed edge from the source vertex s to every vertex v is added, as well as a directed edge from every vertex v to the sink vertex t . Hence,

$$E_a = \{(v, w) \mid T_w^{\text{arr}} \geq T_v^{\text{dep}} + T_v^{\text{min}}\} \\ \cup \{(s, v), (v, t) \mid \text{for all } v\}$$

It can be easily seen that every path from s to t in G_a represents a feasible gate plan of type a and vice versa. What is left is to set the lengths of the arcs, such that the length of a path is equal to the reduced cost of the corresponding gate plan. The cost c_i of gate plan i is equal to the sum of the cost of the idle time intervals. Since including the arc (v, w) in the path implies that v and w are consecutive flights in the gate plan, we assign the cost of the idle time interval between v and w to the arc (v, w) . Looking at the remainder

of the reduced cost, we see that the additional contribution of including flight v in a gate plan of type a amounts to

$$-\pi_v - \sum_{k=1}^K p_{vak} \psi_k$$

which we add to the cost of the arc (v, w) in G_a . The additional arcs (s, v) and (s, t) get cost $-\lambda_a$, and the additional arcs (v, t) get cost equal to $-\pi_v - \sum_{k=1}^K p_{vak} \psi_k$. Since exactly one of the outgoing edges of vertex v will be used if v occurs in a path, the total cost of a path equals the reduced cost of the corresponding gate plan. Note that the cost $c(T_w^{\text{dep}} - T_v^{\text{arr}})$ of putting flight w immediately after flight v in the gate plan is constant; after each iteration, we only have to update the cost terms containing the dual multipliers.

For solving the pricing problem we need to find the gate plan with minimal reduced cost. Since each path in the graph corresponds to a possible, feasible gate plan and the path length corresponds to the reduced cost of the represented gate plan, finding the gate plan with minimal reduced cost comes down to finding the shortest path in the presented graph. Without loss of generality we assume that all flights are sorted by their arrival times. This assumption implies a topological order on the vertices in the graph, namely the order of the flight indices. Because we now have a directed acyclic graph with a topological order it is possible to find the shortest path in $\mathcal{O}(|V| + |E|)$ time (cf. Cormen et al. 2001).

When a gate plan with minimum reduced cost has been found, there are two possibilities:

- The new gate plan has *negative reduced cost*. This means that by adding this gate plan to the master problem, the objective value of the master problem might decrease and thus we add this gate plan to the master problem.
- The gate plan has *zero reduced cost* in which case there exists no gate plan with negative reduced cost.

For each of the gate types, we need to check whether a gate plan with negative reduced cost exists. If for none of the gate types a gate plan with negative reduced cost exists, then the master problem has been solved to optimality.

3.3 Solving the restricted ILP

After the master problem has been solved to optimality we only have a solution for the LP-relaxation. If this solution happens to be integral, then we have a solution to the original ILP formulation of the gate assignment problem, too. If the solution is fractional, then we have to convert it to an integer solution. One possibility for this is to make use of branch-and-price. But since this is computationally infeasible, and since a good approximation suffices, we go for an approximate solution. To this end, we use the ILP-solver

CPLEX to solve the ILP with the limited set of columns. In our preliminary computations, we only used the restricted set of columns generated by the column generation. It turned out that this took too much time and memory. Moreover, if solutions were found within reasonable running time, then the quality of these solutions was really bad. The large running time and memory consumption can be explained by the fact that the restricted set of columns generated for solving the LP is too restrictive for already covered flights. As the first-choice column was generated once, when solving a pricing problem, we decided to create a set of additional second-choice columns each time when solving the pricing problem.

To create these second-choice columns we used the following procedure. We first solve the pricing problem, that is, we find the shortest path. We then take out the nodes in the shortest path one by one and solve the shortest path problems for each of the resulting graphs. These additional columns are added to a column pool. After the master LP problem has been solved to optimality, we determine the set of unique columns from the ones that are in the column pool. This set of unique columns is then added to the restricted ILP problem. After these unique columns were added to the ILP problem, the ILP solver was able to solve the problem in a matter of minutes and sometimes even seconds instead of running for hours or even days. A possible reason for this tremendous speed improvement may be that, since the number of unique gate plans that are in the column pool is quite large, it might be easier for the branch-and-bound subroutine of CPLEX to find a good lower or upper bound quickly, thus speeding up the process.

3.4 Assigning gate plans to gates

After solving the ILP, we have a set of gate plans with just as many gate plans of type a as there are physical gates of type a , such that the total schedule is robust against small variations caused by for example delays of flights. In the second phase of the problem we have to determine which gate plan is assigned to which physical gate.

In the first phase we have introduced constraints dealing with just one gate type. We did not consider relations between specific physical gates, like for example the constraint that two flights having the same departure time cannot be assigned to directly opposite gates due to the impossibility of a simultaneous push-back of both flights. We consider these types of constraint when we assign the gate plans to the physical gates in the second phase.

In Van Orden (2002) some additional constraints have been formulated that need to be addressed in the second phase. These are:

- Avoid putting two flights next to each other that have an overlapping wingspan.

- Avoid putting two flights with equal departure time next to each other because of conflicting push-back.
- Avoid putting two flights that have the same departure time on opposite gates because they cannot have a push-back at the same time.
- Minimize the walking distance for the passengers. This concerns both arriving or departing passengers and transfer passengers.

Although the first one of these constraints at first sight seems to be a good example of a second phase constraint, it turned out to be not important at all. After receiving detailed information of the gates at AAS from the gate planners it turned out that this constraint did not exist for any gate at AAS. But there does exist a strongly related constraint at AAS though, which decrees that it is possible to combine two gates of a small category to one gate of a bigger category. This constraint cannot be addressed in the first phase, since we do not know then which two gate plans will be next to each other in the final solution. In theory it is possible to take this constraint into consideration when solving the first phase by creating a new category of gates consisting of the two smaller gates. The pricing problem for this gate type would then consist of two paths that both contain the selected vertices corresponding to the aircraft of a bigger category. This problem is harder to solve. Therefore, and also since there are not so many of these possibilities, we have decided not to take this constraint into consideration.

After completing the second phase it will be known which gate plans and thus which flights will be assigned to the gates that can be combined. When there are big flights left that are still not assigned, the gate planner will be able to manually combine a set of these smaller gates into one gate of a bigger category and assign a bigger flight to this combined gate.

Although the gate planners at AAS do not have information regarding possible connecting flights of the passengers, one way they try to maximize passenger comfort is to minimize the maximum walking distance. This is achieved by putting flights with a large number of passengers at the best gates. Since we are assigning entire gate plans to gates now, we have less flexibility, but we can use the same principle. Gates that are closer to the beginning of the pier are considered to be better gates. So also in this case the number of passengers will be an important factor in the decision: when there are more gate plans to choose from, the one with the largest number of departing passengers will be assigned to the best gate.

The full assignment problem of the second phase can be decomposed into a number of smaller assignment problems, one for each type of gate. Most of these subproblems can be solved independently, but some are dependent, since they involve gates that have a neighboring or opposite gate of a

different type. The dependent subproblems need more attention. To determine the benefit of assigning a flight to a certain physical gate, the gate planners at AAS use a number of different rules.

Presumably, the best option is to present the gate planners with the results from the first phase and have them assign the gate plans to the physical gates. This can be done *manually* since the size of these problems is rather small; generally the maximum number of gates within one gate type is around eight. Only for remote stands this number is higher, but the flexibility for assigning gate plans to these remote stands is really high. Finally, sometimes it may turn out to be beneficial to swap two flights from two gate plans, resulting in a better solution with respect to the additional constraints, at the expense of a small deterioration of the robustness.

3.5 Directly assigning flights to gates

The two phases of first assigning flights to gate plans and then assigning the gate plans to gates can be integrated by modeling each single physical gate as a separate type.

We considered all the gates except for the remote stands as separate types. The reason we do not consider each remote stand as a separate type is that there does not exist any significant difference between these gates (e.g. they do not have waiting rooms and they all require a bus to transport passengers to and from the terminal building). Our computational experiments reveal that considering all gates as separate types is still computationally feasible. Furthermore, it decreases the amount of work in the second phase, because it limits the second phase to swapping some gate plans between physical gates and if necessary including unassigned flights and swapping flights between gate plans.

The advantage of directly assigning flights to gates is that we can include other criteria than just robustness to assess the solution. We can use here any criterion that can be computed when the combination flight-gate is known, like walking distance, total score (appreciation of a flight to be put on a gate) and deviation from a reference schedule. Given any composite objective function that we want to apply for this combination of criteria, we can compute the cost of the gate plans and solve the problem. The pricing problem for finding a gate plan for a given gate can be solved as before, since the flight-gate appreciation yields a constant reward that can be computed before. The problem can be solved for various composite objective functions (not necessarily linear) in an incremental fashion. Nonetheless, robustness remains most important.

The disadvantage is that there is no freedom left to deal with constraints in which a combination of gates is involved. For example, preventing simultaneous push backs will require additional constraints, since this requires that we forbid a combination of conflicting gate plans. To find

out which constraints should preferably be included in the model and which constraints can better be handled manually by the gate planners, is a matter of further research. Presumably, there will always be need for some kind of second phase especially if it is necessary to manually violate some of the constraints to accommodate all flights.

4 Computational experiments

We have implemented the model and its solution in C++ and performed extensive computational experiments. All of the tests were conducted on a Pentium 4, 2.8 GHz with 1 GB of RAM. For solving the LP problems and for solving the resulting ILP problem by branch-and-bound the Concert Technology interface of CPLEX 9.1.2 (ILOG 2005) was used.

AAS has provided us with six different sets of data, three of which contain the flights on busy high season days (HS), and three on low season days (LS). From each data set we derived two instances, one instance where we group gates with equal properties into a type and one instance where, except for the remote stands, each individual gate forms a separate type. The sizes of the different data sets are given in Table 1. For a data set X, instance X-GG denotes the instance with grouped gates and X-SG the instance with a gate type for each individual gate.

During the column generation process we applied the dual simplex method for solving the LPs, which is default in CPLEX, as well as the primal simplex method. In our experiments this hardly made any difference and we report results on the dual simplex method only. To limit the size of the intermediate linear programs, we take out columns with large positive reduced cost, since these are unlikely to improve the current solution. After every given number of iterations we remove all columns from the model with reduced cost

Table 1 Sizes of the provided instances. LS is low season, HS is high season

| Instance | Flights | Gate types | Total gates |
|----------|---------|------------|-------------|
| HS-1-GG | 699 | 40 | 128 |
| HS-2-GG | 680 | 40 | 128 |
| HS-3-GG | 688 | 40 | 128 |
| LS-1-GG | 602 | 40 | 128 |
| LS-2-GG | 608 | 40 | 128 |
| LS-3-GG | 593 | 40 | 128 |
| HS-1-SG | 699 | 94 | 128 |
| HS-2-SG | 680 | 94 | 128 |
| HS-3-SG | 688 | 94 | 128 |
| LS-1-SG | 602 | 94 | 128 |
| LS-2-SG | 608 | 94 | 128 |
| LS-3-SG | 593 | 94 | 128 |

Table 2 Running time, number of iterations, and number of columns generated for various take-out frequency values

| Instance | Take out freq 40 | | | Take out freq 80 | | | No take out | | |
|----------|------------------|--------|-------|------------------|--------|-------|-------------|--------|-------|
| | sec. | iters. | cols. | sec. | iters. | cols. | sec. | iters. | cols. |
| HS-1-GG | 139 | 693 | 14157 | 139 | 634 | 12892 | 195 | 557 | 12236 |
| HS-2-GG | 126 | 625 | 13622 | 128 | 580 | 12627 | 165 | 582 | 11971 |
| HS-3-GG | 110 | 642 | 13629 | 118 | 595 | 12628 | 158 | 537 | 12287 |
| LS-1-GG | * | * | * | 61 | 685 | 9910 | 82 | 640 | 9198 |
| LS-2-GG | 62 | 681 | 10861 | 67 | 599 | 9944 | 89 | 562 | 9640 |
| LS-3-GG | 62 | 693 | 10881 | 67 | 651 | 10124 | 82 | 597 | 9783 |
| HS-1-SG | 331 | 765 | 33004 | 310 | 652 | 29976 | 447 | 572 | 29116 |
| HS-2-SG | 300 | 631 | 31628 | 257 | 582 | 28410 | 505 | 589 | 28958 |
| HS-3-SG | 257 | 641 | 31796 | 236 | 631 | 28163 | 385 | 573 | 28423 |
| LS-1-SG | 470 | 1238 | 27975 | 160 | 700 | 20895 | 169 | 641 | 19968 |
| LS-2-SG | 137 | 673 | 24041 | 146 | 625 | 21960 | 161 | 562 | 22169 |
| LS-3-SG | 127 | 668 | 23650 | 122 | 621 | 22119 | 137 | 590 | 22276 |

*Take out frequency 40 resulted in an unsolvable LP

above a certain threshold. This threshold is determined by taking the average reduced cost of the columns added in the previous iteration and multiply this average with -0.75 .

These removed columns are put in a special repository and are added again when we solve the ILP. We tested different frequencies of taking out columns: from every 40 iterations to every 110 iterations with step size 10, or no take out. In Table 2 the results are given for 40, 80, and no take out. We report on the time needed to solve the LP-relaxation, the number of iterations and the number of columns generated.

It is clear that the time needed for solving the LP is considerably bigger for the high-season data sets than for the low-season data sets. Taking out columns every 40 iterations is faster in many cases. However, in some cases the computation time is strongly enlarged, while in one case the problem could not even be solved at all. The reason for this is that CPLEX gets stuck in a cycle of generating new columns which need some time before they become useful and before they do become useful, they are taken out and CPLEX has to regenerate them again. Our experiments indicate that a *frequency of 80 is the best* although, there is not a large difference with neighboring values.

Recall that in order to solve the resulting ILP in a reasonable time, we need to add the columns from the column pool. Moreover, we add the columns that were taken out and put in a repository. To decrease the size of the branch-and-bound tree we *gamble* for a small integrality gap. In CPLEX we manually initialize an upper bound on the best integral solution at a value of 0.35 percent above the optimal value of the LP-relaxation, which implies that all nodes with a larger lower bound are pruned. This turns out to significantly decrease the solution time. Moreover, we experimented with different settings of CPLEX, such as more elaborate preprocessing and more aggressive cut generation; see Table 3.

Table 3 Running times for solving the ILP with and without enhancements, averaged over three different values for the take-out frequency

| Instance | ILP default (s) | ILP enhanced (s) |
|----------|-----------------|------------------|
| HS-1-GG | 7 | 6 |
| HS-2-GG | 18 | 9 |
| HS-3-GG | 9 | 8 |
| LS-1-GG | 118* | 52* |
| LS-2-GG | 101 | 24 |
| LS-3-GG | 58 | 24 |
| HS-1-SG | 21 | 19 |
| HS-2-SG | 48 | 24 |
| HS-3-SG | 26 | 15 |
| LS-1-SG | 168 | 120 |
| LS-2-SG | 73** | 126 |
| LS-3-SG | 94 | 113 |

*Take out freq 40 resulted in one unsolvable LP

**Default CPLEX was not able to solve three instances within 60 minutes

In Table 4 we show the running times for solving the ILP for the default settings of CPLEX and for the enhanced settings. For each instance we give the average running time over the different take-out strategies and the primal and dual simplex method for solving the LPs during column generation. In the enhanced settings, after 90 seconds, the solution process is aborted and more aggressive parameters settings are used for CPLEX, resulting in more time spent on preprocessing the problem which led to smaller running times for solving the restricted ILP. Moreover, using the enhanced settings we were able to solve more instances.

We will refer to the combination of using a take-out frequency of 80, dual simplex, and the above enhancements for the ILP as the *best variant*.

Table 4 LP and ILP for the best variant

| Instance | LP (s) | Convert (s) | ILP (s) | Total (s) | Iters. | Columns | Poolsize | Gap (%) |
|----------|--------|-------------|---------|-----------|--------|---------|----------|---------|
| HS-1-GG | 139 | 13 | 6 | 160 | 634 | 12892 | 85225 | 0.00 |
| HS-2-GG | 128 | 12 | 6 | 148 | 580 | 12627 | 82558 | 0.00 |
| HS-3-GG | 118 | 13 | 5 | 138 | 595 | 12628 | 84036 | 0.00 |
| LS-1-GG | 61 | 8 | 21 | 91 | 685 | 9910 | 58374 | 0.19 |
| LS-2-GG | 67 | 8 | 23 | 100 | 599 | 9944 | 59527 | 0.21 |
| LS-3-GG | 67 | 8 | 5 | 81 | 651 | 10124 | 59099 | 0.09 |
| HS-1-SG | 310 | 30 | 12 | 355 | 652 | 29976 | 185970 | 0.00 |
| HS-2-SG | 257 | 28 | 13 | 300 | 582 | 28410 | 173453 | 0.00 |
| HS-3-SG | 236 | 28 | 13 | 279 | 631 | 28163 | 172979 | 0.00 |
| LS-1-SG | 160 | 16 | 161 | 340 | 700 | 20895 | 113413 | 0.18 |
| LS-2-SG | 146 | 18 | 21 | 187 | 625 | 21960 | 122092 | 0.21 |
| LS-3-SG | 122 | 17 | 196 | 339 | 621 | 22119 | 118663 | 0.09 |

Finally, we report more details for the best variant. In Table 4, the running times, the number of iterations, number of generated columns, the number of columns in the column pool, together with the integrality gap are given. Here Convert denotes the time required to add all unique columns from the pool and repository to the model plus the time needed to convert the LP into an ILP. Again, a clear difference can be seen between the high season data sets and the low season data sets for both the number of generated columns and the number of columns in the column pool. The same holds for the groups of gates and single gates.

Our results indicate that the integrality gap is very small. This implies that our method is able to find *practically optimal solutions* for real-life instances in about 10 minutes. Moreover, the results indicate it is feasible to consider single gates as a group. Although this does not make the second phase unnecessary, it makes it easier.

5 Conclusion and further research

We have investigated the gate assignment problem at AAS and developed a two phase solution approach. In the first phase we assign the flights to gate plans, while in the second phase we assign the gate plans to the actual gates. It turns out that the second phase boils down to a set of small problems that can be solved manually. For the first phase, we have presented a different way of formulating the gate assignment problem as an ILP. Our model includes all realistic constraints that are actually used for solving the gate assignment problem at AAS. Furthermore, we have described a method to find an approximate solution for this new ILP formulation. We have implemented it in C++, where we use CPLEX 9.1.2 for solving the (I)LPs.

Our implementation was tested with real life input data, provided by AAS. These instances could be solved in a mat-

ter of minutes to near-optimality and sometimes even to optimality.

The planning software currently in use at AAS is based on a greedy algorithm that assigns flights to gates on the basis of an optimal point score per flight. This score includes different issues, such as preferences of airlines and ground handlers, but it does not contain any robustness measure. Our advanced LP-based algorithm and the planning software of AAS using a greedy algorithm have comparable running times. During meetings we had with the gate planners at AAS, finding a robust schedule was identified as a major improvement in the support of their daily planning activities. Our algorithm focuses explicitly on robustness, and is expected to give the gate planners more time for solving larger conflicts that arise during the actual day.

A next step is to test the effect of optimizing our robustness objective in more detail and to analyze the quality of our solutions compared to the solutions presented by the current planning software. This needs an extensive simulation study, which is an important topic for further research. In Diepen et al. (2011) we presented a column generation algorithm for robust planning of platform busses at AAS, with a similar type of objective function. This paper also includes a simulation study which reveals that schedules from the column generation algorithm require fewer interventions from the planners than the schedules computed by a first-come-first-served algorithm as is used in current practice.

As a first step towards the integration of the first and second phase, we defined in our first phase model every individual gate as a separate type, except for the platform stands. Our experiments show that the problem then still can be solved within a few minutes. By increasing the number of types and defining the gate types more restrictively, we could include more constraints and preferences in the first phase. Further research possibilities are to model and implement more of the rules that the gate planners use to assign the flights to gates.

Finally, our algorithm allows parallelization, e.g. solving in parallel the pricing problems for different gate type and determining in parallel different columns for the column pool by omitting a single flight. This might be of interest when the algorithm is applied for replanning during the day of operation, because then speed is a crucial factor.

Acknowledgements We would like to thank Prof. Dr. Jan van Leeuwen for his helpful comments on an earlier version of the paper. Moreover, we thank two anonymous referees for their suggestions to improve the paper.

References

- Barnhart, C., Johnson, E. L., Nemhauser, G. L., Savelsbergh, M. W. P., & Vance, P. H. (1998). Branch-and-price: column generation for solving huge integer programs. *Operations Research*, *46*, 316–329.
- Bihl, R. A. (1990). A conceptual solution to the aircraft gate assignment problem using 0,1 linear programming. In *Proceedings of the 12th annual conference on computers and industrial engineering* (pp. 280–284). Elmsford: Pergamon Press
- Bolat, A. (2000). Procedures for providing robust gate assignments for arriving aircrafts. *European Journal of Operational Research*, *120*, 63–80.
- Cormen, T. H., Leiserson, C. E., Rivest, R. L., & Stein, C. (2001). *Introduction to algorithms* (2nd ed.). Cambridge/New York: MIT Press/McGraw Hill.
- Diepen, G., Pieters, B., van den Akker, J., & Hoogeveen, J. (2011, in press). Robust planning of airport platform buses. *Computers and Operations Research*. Available at: <http://www.sciencedirect.com/science/article/pii/S0305054811002267>
- Dorndorf, U., Drexl, A., Nikulin, Y., & Pesch, E. (2007). Flight gate scheduling: state-of-the-art and recent developments. *Omega*, *35*, 326–334.
- Dorndorf, U., Jaehn, F., & Pesch, E. (2008). Modelling robust flight-gate scheduling as a clique partitioning problem. *Transportation Science*, *42*, 292–301.
- Dorndorf, U., Jaehn, F., & Pesch, E. (2012). Flight gate scheduling with respect to a reference schedule. *Annals of Operations Research*, *194*, 177–187. doi:10.1007/s10479-010-0809-8.
- Drexl, A., & Nikulin, N. (2008). Multicriteria airport gate assignment and Pareto simulated annealing. *IIE Transactions*, *40*, 385–397.
- Espinoza, D., Garcia, R., Goycoolea, M., Nemhauser, G., & Savelsbergh, M. (2008a). Per-seat, on-demand air transportation part I: problem description and an integer multi-commodity flow model. *Transportation Science*, *42*, 263–278.
- Espinoza, D., Garcia, R., Goycoolea, M., Nemhauser, G., & Savelsbergh, M. (2008b). Per-seat, on-demand air transportation part II: parallel local search. *Transportation Science*, *42*, 279–291.
- Freling, R., Wagelmans, A. P. M., & Paixao, J. M. P. (2001). Models and algorithms for single-depot vehicle scheduling. *Transportation Science*, *35*(2), 165–180.
- Haghani, A., & Chen, M.-C. (1998). Optimizing gate assignments at airport terminals. *Transportation Research. Part A, Policy and Practice*, *32*, 437–454.
- ILOG (2005). ILOG CPLEX v9.1. <http://www.ilog.fr>.
- Nikulin, N., & Drexl, A. (2010). Theoretical aspects of multicriteria flight gate scheduling: deterministic and fuzzy models. *Journal of Scheduling*, *13*, 261–280.
- Van Orden, A. (2002). *Gate assignment: methods and models*. Master's thesis, Department of Mathematics, Utrecht University.
- Xu, J., & Bailey, T. G. (2001). The airport gate assignment problem: mathematical model and a tabu search algorithm. In *HICSS'01: proceedings of the 34th annual Hawaii international conference on system sciences (HICSS-34)* (Vol. 3).
- Yan, S., & Huo, C.-M. (2001). Optimization of multiple objective gate assignments. *Transportation Research. Part A, Policy and Practice*, *35*, 413–432.
- Yan, S., Shieh, C.-Y., & Chen, M. (2002). A simulation framework for evaluating airport gate assignments. *Transportation Research. Part A, Policy and Practice*, *36*(5), 885–898.