Cabin crew rostering at KLM: optimization of reserves

Marco Bijvank[*] [†] *Jarek Byrka*[‡] *Peter van Heijster*[*] *Alexander Gnedin*[§]
Tomasz Olejniczak[¶] *Tomasz Świst*[¶] *Joanna Zyprych*[¶] *Rob Bisseling*[§]
Jeroen Mulder[‖] *Marc Paelinck*[‖] *Heidi de Ridder*[‖]

Abstract

In this paper, we will discuss the issue of rostering jobs of cabin crew attendants at KLM. Generated schedules get easily disrupted by events such as illness of an employee. Obviously, reserve people have to be kept 'on duty' to resolve such disruptions. A lot of reserve crew requires more employees, but too few results in so-called secondary disruptions, which are particularly inconvenient for both the crew members and the planners. In this research we will discuss several modifications of the reserve scheduling policy that have a potential to reduce the number of secondary disruptions, and therefore to improve the performance of the scheduling process.

Key words: airline crew rostering, reserve duties, soft flights

1 Introduction

KLM (Koninklijke Luchtvaart Maatschappij N.V. also known as KLM Royal Dutch Airlines) has more than 100 aircraft and over 8,000 cabin flight attendants. Every week, a new roster is received by the cabin crew, which shows their assignments for the next several weeks. There are about 6,000 flights assigned to crew members each week. Besides the flights other assignments are rostered as well, such as trainings, days off and reserve duties.

1.1 Rostering of Flights

The assignment of cabin crew to flights is a difficult problem in which many different aspects have to be taken into consideration. First, the cabin crew is divided into four ranks: Senior Pursers, Pursers, Business Class Flight Attendants and Economy Class Flight Attendants. The last two ranks are sometimes denoted by the stripes ('*band*' in Dutch) on their sleeves: two stripes for the Business Class Flight Attendants ('*2-bander*' or 2B) and one stripe for the Economy Class Flight Attendants

[*]Vrije Universiteit, Amsterdam
[†]corresponding author, `mbijvank@few.vu.nl`
[‡]CWI, Amsterdam
[§]Universiteit Utrecht
[¶]Adam Mickiewicz University, Poland
[‖]KLM
[0]We would like to thank the other participants of the KLM group during SWI 2007. Especially, Niels Oosterling, Allard Veldman and Erik van Werkhoven.

(‘*1-bander*’ or 1B). There are certain regulations on how many crew members of a particular rank have to be on a certain flight. In general, there are only Pursers and 1-banders for flights within Europe (*short-haul* flights), while all four ranks have to be scheduled on intercontinental (*long-haul*) flights.

Second, the crew has to be qualified to fly on a particular aircraft type. There are in total six different aircraft types, and each crew member is qualified to fly a maximum of three different types. Crew members of the same rank and with the same qualifications are grouped into divisions. Currently, the KLM cabin crew members are divided into 17 divisions.

According to legal regulations and the collective labour agreement of the KLM cabin crew, each flight duty should be followed by a minimum number of hours of time off. The length of this time off depends on the characteristics of the duty such as duration, time of day, time difference between origin and destination. Also, after several flight duties the flight attendant is entitled to a number of days of leave, which also depends on the characteristics of the previous assignments.

The combining of flights in such a manner as to optimize the balance between duty time and leave is a specialized process due to the complexity of the regulations. This is why flights are combined into predefined patterns prior to the assignment process. A combination of flight duties followed by the appropriate number of days of leave is called a *pairing*. The length of a pairing can vary from three days up to seventeen days. The process of constructing these pairings is called the *pairing process*. This is performed by an application called *Carmen Crew Pairing*.

Pairings are assigned to specific crew members several weeks ahead of the actual day of execution. This is done with consideration of the rank and aircraft qualifications of the specific crew member. But also the flight preferences and requests for leave on specific days are taken into account. The requests are evaluated according to certain priority rules so as to avoid conflicts (such as a number of crew members requesting the same assignment). These requests result in assignments of pairings to crew members prior to the rostering of the remaining pairings, which is performed by the application *Carmen Crew Rostering*. See Kohl and Karish [3] for rostering algorithms exploited in such a system. Carmen produces a roster for a period of two weeks, the *published period*. The resulting roster is rather fixed and cannot change much. The roster after these two weeks is also published, however, it can still change quite a lot.

Currently, about 80% of all the flight assignments are assigned as requests. This method has a serious drawback, because it does not allow for Carmen to optimize the crew rostering with respect to efficient allocation of resources. Therefore, a new strategy is currently being introduced called *Preferential Bidding System*. It consists in giving preferences rather than requests for specific flights. A preference could refer to a single flight or something more general, like the preference to start and finish early or the preference for flights to the Far East. The Carmen Crew Rostering system will then try to comply with a maximum number of preferences while also optimizing the efficiency of the rosters.

1.2 Disruptions

In general, a schedule is subject to changes. The flight assignments can become disrupted. In order to handle these disruptions, reserve duties are assigned to crew members in-between their regular activities. The reserves can take over flights that have become vacant. This process of resolving disruptions and adapting the schedule is an *online procedure*. That is, as soon as a disruption is reported, a reserve is assigned to the disruption. When a reserve is assigned to a disrupted flight, the entire pairing or its remainder is assigned to the reserve crew member.

There are several types of disruptions. *Internal disruptions* deal with disruptions of individual crew members, like illness. *External disruptions* affect more crew members, like short-term commercial changes, delays or problems on the day of operations. Next to these *primary disruptions*, there are also *secondary disruptions*. It often happens that the disrupted pairing consists of more days than the assigned reserve duty. In that case, the reassigned reserve crew member will no longer be able to perform the pairing which was originally assigned following his or her reserve block. As a consequence, this flight will become disrupted as well. This is what we call a secondary disruption. Such a disruption can cause another disruption, a *ternary disruption*. This *domino effect* will continue until a disrupted flight occurs outside the published period, or when it is assigned to a reserve who has enough days left to take over the remaining days of the disrupted pairing. This is illustrated with an example for three crew members in Figure 1. In the left figure, we see a possible roster produced by Carmen Crew Rostering. If on day 1 crew member 1 is disrupted, crew member 2 gets assigned to the disrupted flight. However, the pairing of crew member 2 starting on day 6 is also disrupted. Therefore, crew member 3 is assigned to this secondary disruption on day 6.

DAY1	DAY2	DAY3	DAY4	DAY5	DAY6	DAY7	DAY8	DAY9
FLT	FLT	FLT	FLT	FLT	OFF	OFF	FLT	FLT
RES	RES	RES	OFF	OFF	FLT	FLT	FLT	OFF
FLT	FLT	OFF	OFF	RES	RES	RES	RES	RES

DAY1	DAY2	DAY3	DAY4	DAY5	DAY6	DAY7	DAY8	DAY9
DIS	DIS	DIS	DIS	DIS	DIS	DIS	DIS	DIS
FLT	FLT	FLT	FLT	FLT	OFF	OFF	???	???
FLT	FLT	OFF	OFF	RES	FLT	FLT	FLT	OFF

Figure 1: The roster before and after the disruption of crew member 1.

Besides the concept of disruptions, crew members can also recover (after for instance illness). This means that they become available for duty again. After an internal disruption of a crew member his or her schedule is completely erased. As a consequence, a recovered crew member has no tasks left until the end of the published period and hence can be assigned new flights.

In the remainder of this paper, we refer to *flight blocks* when we talk about pairings. Note the difference between a *flight duty* (*reserve duty*) and a *flight block* (*reserve block*). The first definition does not include days off (or days of leave), while the latter does.

Currently, a reserve block consists of five consecutive days of reserve -by duty followed by two days off. Each day, the reserve duty is restricted by a start time and end time which can vary from one reserve to the other. There are six different start times over a day and the end time is always eight and a half hours later. This means that a disrupted flight may only be assigned to those reserves of which the flight starts within the reserve duty. The number of reserve blocks (*level of reserves*) that needs to be assigned to crew members each day is determined at the beginning of each season. There is, however, no model available with which these decisions are made. At the moment, this is done by employees of the Planning Department of KLM who mostly use their past experience and good judgement.

The goal of this research is to come up with a good reserve strategy for the cabin crew of KLM. A reserve strategy is defined by the number of reserve blocks that have to start each time unit and the configuration of each reserve block, i.e., the length of the reserve duty, the number of days off, and where to locate the days off. The quality of a specific reserve strategy can be determined by the following performance measures:

- The number of secondary disruptions: this is a measure of the uncertainty for crew members about the assignment following the reserve block.

- The number of unused reserve days: these days will make it necessary to roster extra crew, making the reserve strategy more expensive.

- The number of 'open' days. These are days between the end of an assigned disruption and the remainder of the roster for the assigned reserve. In order to be able to utilize these days, it will mostly be necessary to reassign a part of the existing assignment of the crew member. This is not desirable if one wants to maintain the remainder of the assignment intact as much as possible.

There is an intricate trade-off between these three measures. For example, open days can be avoided at the expense of creating secondary disruptions, whereas unused reserve days can be avoided by assigning duties to reserves at the earliest possibility, thus creating more open days.

The rostering process of all KLM crew is quite complex. Figure 2 shows the relationships between the different components that are discussed in this section. The construction of the actual schedule is performed by Carmen. In this paper, our goal is to define the input for Carmen about the reserve duties that have to be assigned to crew members. Section 2 discusses the assumptions and simplifications we made. In Section 3, we propose several methods to determine how many reserve blocks should start on a particular day, and how long these reserve blocks should be. In Section 4, we discuss the *soft flight approach*, which can be seen as an extra constraint for Carmen to prevent the domino effect to occur. The main idea of this approach is to have a pairing after a reserve block assigned to a reserve crew member on duty for the same length of time. This pairing is called a soft flight. This guarantees that no ternary disruption occurs. The possibilities for other configurations for the reserve block are considered in Section 5. After producing a schedule, we should also have a method to evaluate the performance of the schedule. We therefore introduce an algorithm to analyze how different schedules perform in Section 6. In the final section, we obtain some numerical results on the different approaches and compare them.

2 Assumptions and Simplifications

Airline Crew Rostering problems belong to the most difficult problems to solve since they are NP-hard (see Ní Éigeartaigh and Sinclair [1]). There are, however, principles and heuristic rules that result in solutions that are good enough in practice. In order to find such principles we have to make some assumptions since not all details can be modelled properly.

The first simplification we make is that we do not distinguish between cabin crew, i.e., we use no ranking and no specific qualifications for aircraft type of the crew members. Consequently, all crew members are interchangeable which simplifies the assignment of reserves to disruptions. This simplification is only justified when the reserve blocks are rostered to members of each of the 17 divisions (as explained in Section 1) proportional to the number of crew members of each division required to perform the scheduled flights.

For reasons of simplicity we assume that all crew members work full-time. An employee working part-time is entitled to have more days of leave. Another simplification we make has to deal with the time units. We round everything to days. So, a reserve duty is for 24 hours on a day instead of eight and a half hours.

We assume that a disruption can only occur on the first day of a flight block. As mentioned in the previous section, we can distinguish short-haul and long-haul flights. Most of the time, a long-haul flight duty consists of only one flight, where a pairing on a short-haul flight is likely to consist of multiple flights. Consequently, a flight duty for a short-haul can be disrupted on each day while a

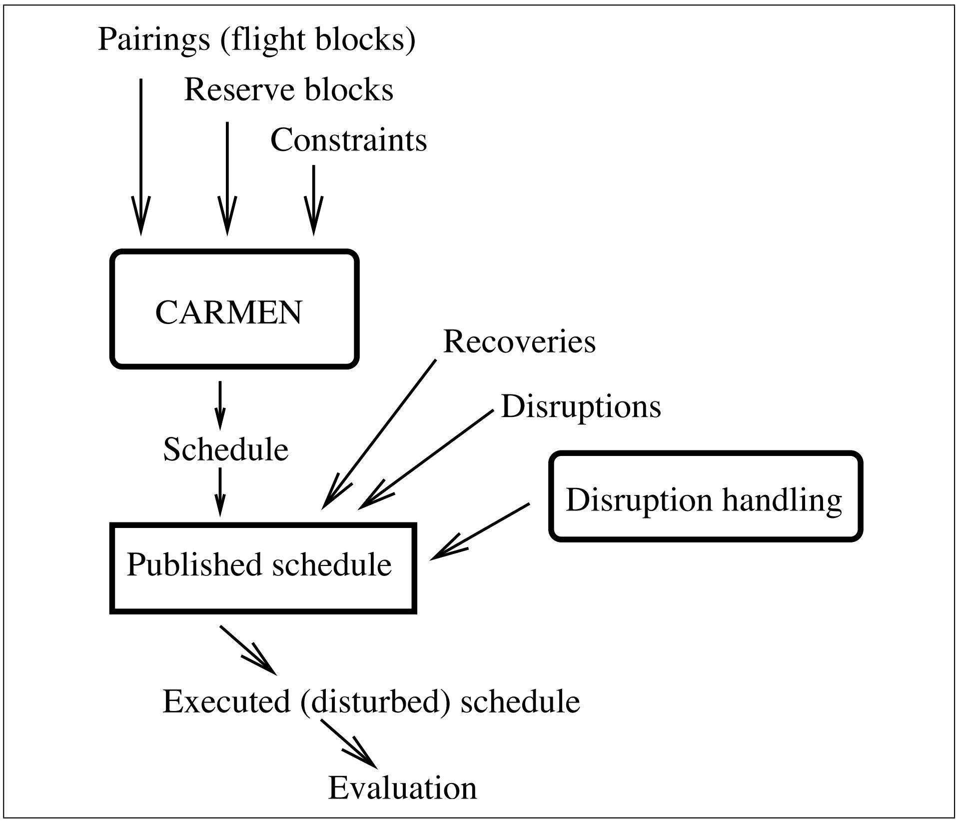

Figure 2: All aspects to be considered in the rostering problem.

flight duty for a long-haul can only be disrupted on the first day. Therefore, we make the assumption of considering long-haul flights only. Note that this assumption ignores the possibility of a crew member getting ill during his or her flight, or a malfunction of an aircraft at its foreign destination (flight blocks always start and end in Amsterdam). When we look at the length of short-haul flight blocks and the current reserve strategy, this assumption is not a problem. From historical data we know that about 29% of the flight blocks has a length of 6 days, 15% has length 8, 13% has length 7, 12% has length 11, 10% has length 10, 7% has length 9, and the remainder is small. The range of lengths is from 2 to 16 days.

Most short-haul flight blocks have duties of at most 5 days flying and 2 days off afterwards. This is exactly the same as the current configuration of the reserve blocks. Therefore, it is less likely that the current reserve strategy results in problems when these short-haul flights get disrupted. On the other hand, the current reserve strategy will most likely cause secondary disruptions when a long-haul flight is disrupted, since more than half of the long-haul flights exceed the length of the current reserve block. Therefore, the only way to deal with disruptions on long-haul flights without causing a secondary disruption is with recoveries. Since disruptions of short-haul flights are not a problem, it is reasonable to simply ignore short-haul flights completely in this research.

The final simplifications we make have to deal with the handling of internal disruptions, external disruptions, and recoveries. Internal disruptions result in less available crew members. Therefore, they require actual reserves. On the other hand, external disruptions will most likely result in changes of the reserve crew, since the crew becoming available due to an external disruption can partly be used as a reserve again. A disrupted crew member becomes available again for services after a predictable (e.g., external disruption) or unpredictable (e.g., internal disruption) time period. The unpredictable recoveries are modelled as the start of a reserve duty with infinite length, i.e., a length equal to the published period of two weeks.

The main idea of this paper is to model the crew members that get disrupted as *workforce out-flow*, while crew members returning to service after a disruption (i.e., recovering) are modelled as *workforce in-flow*. In the long-run, by constant number of employees, the total workforce in-flow and out-flow is approximately equal – a *workforce conservation principle*. However, at each time instant there is a mismatch of the flows, which must be resolved by reserves.

It seems only wise that a disrupted crew member should return to his or her original roster as soon as possible, since this was found to be optimal. Hence, keeping as close to the original schedule as possible means staying close to optimality. This approach is advocated in Kohl *et al.* [4, Section 3.2], where closeness to the original schedule stands as a principal objective of disruption management. In order to use this principle, we have to make sure that secondary disruptions are prevented as much as possible. Otherwise, the original schedule is mixed up even more. Therefore, the reserve blocks must cover the long-haul flights. This can be achieved by rostering longer reserve blocks as compared to the current reserve strategy. In the next section, we develop different techniques that exploit this concept.

3 Level of Reserves

The previous section made clear that secondary disruptions have to be avoided as much as possible. With the current situation nearly every disrupted long-haul flight will cause a secondary disruption since there are no indefinite recoveries. Therefore, longer reserve blocks are proposed. In this section, we develop three techniques that determine the number of reserve blocks that has to start on a particular day with a certain length. The first technique is based on the concept of constructing reserve blocks that are copies of the long flight blocks, see Section 3.2. The second technique copies the flight blocks proportional to their occurrence, see Section 3.3. In the third technique, we use a more statistical model to construct the reserve blocks, see Section 3.4. But let us introduce some notation first.

3.1 Notation

A given flight schedule is defined by the number of crew members starting their flight block of type j at day k (denoted by S_{jk}). A type j can involve the characteristics length, rank, and aircraft type. We write $j \succeq i$ if the characteristics of type j exceed (nonstrictly) the characteristics of type i (in the case of length or rank), or they are compatible (in the case of aircraft type). Plainly, $j \succeq i$ means that a reserve of type j can be used to serve a disrupted flight of type i. In our simplified model, where we ignore rank and aircraft type, we can identify the type of block with its length, so that $j \succeq i$ if and only if $j \geq i$.

We want to determine the number of crew members starting their reserve block of type j at day k (denoted by T_{jk}). Therefore, we have to model the disruptions and recoveries. Internal disruptions are assumed to be independent over time, as well as independent over all crew members. Consequently, each crew member can have an internal disruption with some probability p_{int}, where p_{int} is based on the historical data. External disruptions are also assumed to be independent over time. For simplicity, external disruptions are analyzed for each flight attendant instead of per flight. Therefore, external disruptions are also independent over all crew members. Consequently, each crew member can get an external disruption with some probability p_{ext}, where p_{ext} is based on the historical data. We have to emphasize here that this is not always what happens in reality, but it is a useful simplification for an initial model.

The actual number of internal disruptions that requires a reserve of type j at day k is a random variable denoted by X_{jk}, where X_{jk} has a binomial (S_{jk}, p_{int}) distribution. In order to determine the number of recoveries used on a particular day, we assume independence over time. Consequently, we can define the number of recovered crew members on day k as a random variable Y_k with a probability distribution function $f_Y(y)$ with $y \in \{0, 1, 2, \ldots\}$, and expectation $\mathbb{E}[Y]$, both independent of k.

3.2 Mirror Longest Flights

As mentioned in Section 2, we would like to exploit the idea of preferring long reserve blocks over short ones for long-haul flights. In this section we use a maximum length for the reserve block configuration. Several options are presented in Table 1.

option	1	2	3	4	5	6
length reserve duty	5	6	7	8	9	10
number of days off	2	2	3	3	3	4
length reserve block	7	8	10	11	12	14

Table 1: Different options for the configuration of the (longest) reserve block.

The actual number of reserve blocks that has to start with this configuration depends on the flight schedule. When there are more long flights, the reserve blocks should be long as well. But it is meaningless to schedule more reserve blocks of a particular length than there are flight blocks with at least this length.

The idea of this approach is to copy long flight blocks and replace the flight duty with a reserve duty. This process of copying is referred to as producing a *mirror*. Therefore, we first determine the number of flight blocks that has to start on day k (i.e., $\sum_j T_{jk}$) based on a minimal flight reserve cover ratio α,

$$\frac{\sum_j T_{jk}}{\sum_j S_{jk}} \geq \alpha, \qquad \forall k.$$

Currently KLM uses a fixed α equal to 4%. This means that we would like to copy at least 4% of the longest flights scheduled for day k as reserves, where the maximum length of the reserve blocks is restricted to those of the fixed configurations (as in Table 1). This approach assumes flight block types to be characterized by their length only (as mentioned in Section 2). Since it is preferred that disruptions of the longest flights are solved by recoveries, we do not copy the first $\mathbb{E}[Y]$ longest flights. This is beneficial because the longest flights have a higher probability of causing a secondary disruption, and recoveries cannot have secondary disruptions since they have no schedule yet. Note that this strategy only incorporates averages.

The advantage of the *mirror longest flight approach* is that long reserve blocks are created. This approach, however, also has its disadvantages. Proportionally there are more long reserve blocks compared to the current situation. As a result, more reserve days are rostered as compared to the current situation. To resolve this, either the cover ratio of 4% reserves starting on a particular day has to decrease, or the cover ratio should be based on the number of reserves and flights scheduled at a particular day instead of only on the ones starting at day k, i.e.

$$\frac{\sum_{l=1}^{k}\sum_{j=k-l+1}^{\infty} T_{jl}}{\sum_{l=1}^{k}\sum_{j=k-l+1}^{\infty} S_{jl}} \geq \alpha \qquad \Rightarrow \qquad \sum_{j} T_{jk} \geq \alpha \sum_{l=1}^{k}\sum_{j=k-l+1}^{\infty} S_{jl} - \sum_{l=1}^{k-1}\sum_{j=k-l+1}^{\infty} T_{jl}.$$

This will be referred to as the *alternative cover ratio*.

3.3 Mirror Flights Proportional

The mirror longest flight approach prefers long reserve blocks. A disruption of shorter flight blocks will result in partial usage of such long reserve blocks. But also less reserves will start at a particular day when a crew member is rostered as a reserve for around 4% of the time. Another technique to solve the problem is to copy (or mirror) the flights in the same way as discussed in Section 3.2, but the number of reserve blocks with a particular length should be proportional to the number of flight blocks starting that day with the same length. Both options for the cover ratio can be used as mentioned in Section 3.2. We can deal with the recoveries in the same way as well.

3.4 Statistical Model

When all flight block types are unique and cannot exceed another type ($j \not\geq i$ for $i \neq j$), only reserves of type j can be used for a disruption of a flight block with characteristic type j. In such a situation, the probability of requiring more reserves than available should be low, e.g., less than 5%:

$$P(X_{jk} > T_{jk}) < 0.05 \qquad \Leftrightarrow \qquad P(X_{jk} \leq T_{jk}) \geq 0.95 \qquad \forall j, k. \tag{1}$$

In reality, however, characteristic types can exceed each other. In this paper, we assume flight and reserve block length to be the only characteristic, i.e., $i \geq j$ if and only if $i \geq j$. Consequently, disrupted flights of type j can only be resolved with reserves of type i if $i \geq j$ to exclude secondary disruptions. Furthermore, recoveries can be used as well. Therefore, Equation (1) can be reformulated as

$$P\left(X_{jk} \leq T_{jk} + Y_k + \sum_{i=j+1}^{\infty} (T_{ik} - X_{ik})\right) \geq 0.95, \quad \forall j, k. \tag{2}$$

Based on the central limit theorem (see Ross [5]), the binomial distribution for X_{jk} can be approximated by a normal distribution with mean $\mu_{jk} = S_{jk} p_{\text{int}}$ and variance $\sigma_{jk}^2 = S_{jk} p_{\text{int}}(1 - p_{\text{int}})$. For simplicity, the stochastic recovery variable Y_k is also assumed to be normally distributed, with parameters μ_{rec} and σ_{rec}^2.

We can rewrite Equation (2) as

$$P\left(\sum_{i=j}^{\infty} X_{ik} - Y_k \leq \sum_{i=j}^{\infty} T_{ik}\right) \geq 0.95, \quad \forall j, k. \tag{3}$$

The left-hand side of this inequality represents the number of required reserves of at least length j, while the right-hand side represents the number of available reserves of at least this length. The probability of a mismatch should be less than 5%. Since the disruptions and recoveries are independent of crew members, the number of required reserves ($\sum_i X_{ik} - Y_k$) also has a normal distribution with a mean equal to $\sum_{i=j}^{\infty} S_{ik} p_{\text{int}} - \mu_{\text{rec}}$ and a variance equal to $\sum_{i=j}^{\infty} S_{ik} p_{\text{int}}(1 - p_{\text{int}}) + \sigma^2_{\text{rec}}$. Therefore, Equation (3) equals

$$\frac{\sum_{i=j}^{\infty} T_{ik} - \left(\sum_{i=j}^{\infty} S_{ik} p_{\text{int}} - \mu_{\text{rec}}\right)}{\sqrt{\sum_{i=j}^{\infty} S_{ik} p_{\text{int}}(1 - p_{\text{int}}) + \sigma^2_{\text{rec}}}} \geq \Phi^{-1}(0.95), \tag{4}$$

where $\Phi(\cdot)$ is the cumulative distribution function of the standard normal distribution. This is equivalent to

$$T_{jk} \geq 1.645 \sqrt{\sum_{i=j}^{\infty} S_{ik} p_{\text{int}}(1 - p_{\text{int}}) + \sigma^2_{\text{rec}}} + \sum_{i=j}^{\infty} S_{ik} p_{\text{int}} - \mu_{\text{rec}} - \sum_{i=j+1}^{\infty} T_{ik}. \tag{5}$$

With this iterative procedure, we can determine T_{jk} by starting with the longest reserve block and then each time decreasing j and computing the next value of T_{jk}.

4 Soft flight approach

The previous section described several methods to determine the lengths and number of reserve blocks that start on a particular day. In this section, we will consider flight blocks following reserve blocks in a roster and give them a special status. We will call these special flight blocks *soft flight blocks* (SFBs). See also Figure 3. Note that this approach also deals with the input for Carmen, independent of the techniques discussed in the previous section. These SFBs have several special properties, and therefore, should be treated differently in the disruption management process. First of all, SFBs suffer the most from secondary disruptions. On the other hand, they account for just a few percent of all flight blocks (the number of SFBs is equal to the level of reserves, which is around 4% of the total flights at the moment). Hence they may potentially be treated with special care.

We propose to create dedicated reserve blocks to cover soft flight blocks. Although the idea here is similar as in the previous section, the method is quite different. First the level of reserves is determined, as in Section 3. Next they have to be assigned to crew members. In the SFB approach, we propose to start as many SFBs of a particular length as there are reserves starting that day with the same length. The latter is determined first. This principle is shown in Figure 3. Since reserves, and therefore, SFBs occur infrequently, most of these requests can easily be granted. If a dedicated reserve is not needed to prevent a secundary disruption, this will be known in advance and the reserve can then still serve as a regular reserve. This can be taken into account when determining the required level of reserves.

crew j	RES	RES	RES	RES	RES	OFF	OFF	SFL	SFL	SFL	SFL	SFL	OFF	OFF
crew j’								RES	RES	RES	RES	RES	OFF	OFF

Figure 3: Soft flight block for crew member j, with a dedicated reserve j'.

Ensuring the availability of reserve blocks having the appropriate length to cover a disrupted SFB should stop the domino effect, i.e., a secondary disruption will almost never cause a ternary disruption. Currently, the end of the published period is often used to stop this domino effect. Stopping the domino effect could help to extend the planning horizon in the future.

One could go a step further and exploit the regularity of used reserve patterns together with the soft flight approach. We could pre-assign a potential secondary disruption on a SFB of crew member j to the crew member j' whose reserve block follows the reserve block of j. We could inform crew member j' in advance that he or she is the backup of j. It would probably be advantageous for j' to have this information. Additionally, such an automatization of handling secondary disruptions could simplify the online disruption handling process. Note that this is true since each day the level of reserves is about the same and therefore after a reserve block another reserve block starts for a different crew member.

5 Configuration of reserve blocks

As mentioned in Section 1, we can also change the configuration of a reserve block, i.e., the distribution of the off days in the reserve block. Currently, a reserve block consists of five days of reserve duty followed by two days of leave. In this section we discuss the possibility of changing this configuration of reserve blocks. To be more specific, we consider the possibility of moving the days off from the end of the block to the middle. It is desirable for the crew to have some days off at the end of a reserve period, but not necessarily all of the days off.

When we consider the configuration of a reserve block, we have to make a choice which days of the reserve block are most appropriate to be the starting days for serving disrupted flights. If the length of the disrupted flight is at most the length of the reserve block, we would prefer to give this flight to the person whose remaining length of reserve block just covers the flight. Consequently, no secondary disruption is caused and no open days occur in the roster. But if the disrupted flight does not fit to any available reserve block (which is almost always the case with disruptions of long-haul flights), then the choice of the reserve depends on the flight he or she has scheduled just after the reserve block. This is one more reason for introducing SFBs and controlling their lengths.

Imagine a situation where no reserve block is available to cover a particular disruption. Consequently, a secondary disruption occurs. In such a case, one would prefer to assign this disruption to a person at the end of his or her reserve block, such that serving the disrupted flight block would finish at the same time as the SFB of the reserve. The open days are minimized with such a principle. In the SFB approach, the secondary disruption would be served by its dedicated reserve crew member and both reserves could return to their original schedules afterwards.

It makes little sense to speculate what would be the best configuration of a reserve block without extensive testing. Just to illustrate the idea, one could consider the situation illustrated in Figure 4. Assume we use the SFB approach. The first line shows a current configuration of 5 days reserve and 2 days off. When we use the first day to cover disruptions without causing open days, we can use it for disrupted flights of length 7 and 14. When this does not occur, the reserve can be used on the second day for disruptions of 6 or 13 days, and so on. Disrupted flight blocks of length 8 or 9 always result in open days. Such flight blocks are quite common for long-haul flights, as we know from the historical data mentioned earlier. Therefore, we propose to shift the two days off to day 4 and 5 of the reserve block, as shown on the second line of Figure 4. Note that on the other hand the reordering of off days may increase the level of reserves needed, since off days will occur more often during a reserve block.

RES	RES	RES	RES	RES	OFF	OFF	SFL	SFL	SFL	SFL	SFL	OFF	OFF
7|14	6|13	5|12	4|11	3|10									
RES	RES	RES	OFF	OFF	RES	RES	SFL	SFL	SFL	SFL	SFL	OFF	OFF
7|14	6|13	5|12			9	8							
							RES	RES	RES	RES	RES	OFF	OFF

Figure 4: Two reserve blocks with different internal structure. The numbers under the reserve blocks indicate the length of a disruption which they can fit ‘perfectly’, that is, without causing open days in the roster (in combination with their soft flight block). If we want to optimize the perfect fittings, we need to determine which disruptions have the highest probability to occur. In particular, if there are more disruptions of length 8 and 9 than with length 3, 4, 10 and 11 we prefer the reserve block pattern : 3 days active / 2 days off / 2 days active.

6 Comparison of scheduling techniques

In Section 3, we have developed different approaches to determine the number of reserve blocks that have to start at a certain day with a particular length. In order to compare the different approaches, we developed an analytical comparison procedure, which is the subject of this section.

For the comparison procedure we use the same assumptions as made in Section 2, i.e., no specific qualifications for air crew members (no ranks, no particular qualifications for air craft type), only long-haul flights, recoveries are reserves of infinite length (i.e., equal to the publishing period). Another assumption we make, such that reserves can be assigned to disruptions more easily, is complete knowledge of the availability of crew members being on reserve and of all disruptions on a particular day before any reserve gets assigned to a disruption. In practice, this means that all disruptions are reported before any of them is resolved by assigning reserves. Also all people who recover and become available again are reported before this assignment takes place.

For the assignment of reserves to disruptions we want to have a fixed disruption handling scheme. This scheme should result in using the reserve capacity in such a way that secondary disruptions and open days are prevented as much as possible. Therefore, when a disruption of length j occurs, we check recoveries first. Otherwise, we want to use a reserve block of length j. If there is no such reserve block, we gradually increase the length of the desired reserve block, each time by one day. When all options of reserve blocks of at least j days are checked and no reserve is available, a secondary disruption occurs. Now, use a reserve block of length $j-1$ and gradually decrease the length by one day until the disruption is resolved. This approach results in avoiding secondary disruptions as much as possible, and if they still occur making them as short as possible, thus minimizing the probability of ternary disruptions occurring.

Before we actually assign reserve blocks to disruptions, we make another simplification from reality. In reality, we distinguish between a crew member being allowed to be called as a reserve or having days off in a reserve block. In our comparison procedure, we do not want to keep track of this distinction. So, we only look at reserve blocks and flight blocks disregarding the days off.

Let us introduce some notations for the performance measures:

U_k = the expected number of unused reserves on day k

V_k = the expected number of secondary disruptions caused on day k

W_k = the expected number of disruptions that cannot be resolved with recoveries and reserves on day k

The first two definitions have already been mentioned in Section 1. We added the third performance measure, to have an understanding whether there is a lack of reserve crew. In the comparison technique we predict the number of available reserve blocks of a particular length based on the assignment scheme as described above and on the probabilities for recoveries and disruptions.

For each day, we would like to keep track of all flights that can get disrupted. Therefore we define the sequence set D_k as all possible disruptions on day k, where $k \in \{1, \dots, K\}$ and K being the publishing period (i.e., the planning horizon). A possible disruption $d \in \{1, \dots, |D_k|\}$ on day k is identified by the length of the flight $l_d^{(k)}$ and by its probability of requiring a reserve $p_d^{(k)}$. There are as many possible disruptions of length l on day k as the number of scheduled flight blocks with this length (i.e., S_{jk}). Consequently, $|D_k|$ equals $\sum_j S_{jk}$. Each possible flight block gets initially disrupted with probability $p_{\text{int}} + (1 - p_{\text{int}})p_{\text{ext}} = 1 - (1 - p_{\text{int}})(1 - p_{\text{ext}})$. We order the flights in D_k such that $l_d^{(k)} \geq l_{d+1}^{(k)}$ (i.e., in decreasing order of their length).

Set the initial probabilities for having i reserve blocks available of length j starting at day k (denoted by p_{ijk}) as

$$p_{ijk} = \begin{cases} 1, & \text{if } i = T_{jk} \qquad \forall j, k \\ 0, & \text{otherwise} \end{cases}$$

and for the availability of recoveries

$$p_{i,K+1,k} = f_Y(i) \quad \forall i, k.$$

Since disruptions occur, the probabilities of reserves being available (i.e, p_{ijk}) change over time. The analytical procedure to update these probabilities and to compute the performance measures is given in Appendix A.

7 Numerical Results

For the numerical results we used actual data of KLM. For simplicity, we assumed each day to be the same, i.e., $T_{jk} = T_{j,k+1}$ and $S_{jk} = S_{j,k+1}$. When we evaluated the data on long-haul flights, we observed the values as presented in Table 2 and Table 3, where $\mu_{\text{rec}} = 7.1$ and $\sigma_{\text{rec}}^2 = 8.353$.

j	S_{jk}	j	S_{jk}	j	S_{jk}	j	S_{jk}
1	0	5	8	9	27	13	13
2	8	6	108	10	38	14	3
3	0	7	49	11	46	15	5
4	0	8	55	12	12	16	2

Table 2: The number of flight blocks starting on a given day.

The probabilities for internal and external disruptions are $p_{\text{int}} = 0.06$ and $p_{\text{ext}} = 0.07$, respectively.

y	$f_Y(y)$	y	$f_Y(y)$	y	$f_Y(y)$
0	0	7	0.122905	14	0.011173
1	0.011173	8	0.139665	15	0.011173
2	0.033520	9	0.136872	16	0
3	0.064246	10	0.069832	17	0
4	0.092179	11	0.047486	18	0
5	0.103352	12	0.019553	19	0.002793
6	0.120112	13	0.013966		

Table 3: The probability distribution for the recoveries used to resolve disruptions.

Currently KLM uses a 4% flight reserve cover ratio where a reserve block consists of 5 days of duty and 2 days off. So, in total 15 reserve blocks start on a daily basis (4% of $\sum_j S_{jk} = 374$). Since the length of a reserve block is 7 days, there are 105 crew members scheduled as reserve each day. We would like to keep this number more or less fixed when comparing the different approaches discussed in Section 3 (rounding real values to integers causes some small fluctuations). The cover ratio and the alternative cover ratio are the same since each day is identical. For the different configuration approaches we get T_{jk} as in Table 4. Since all days are the same, the performance measures are also the same for each day, as presented in Table 4. The current policy with 15 reserves starting a reserve block of length 7 each day is reflected in our simplified model by option 1 in the table. The statistical method without the restriction of the current 105 crew members being planned as reserves is given in the last column, as 'ideal statistical'. Notice that in our simplified model the number of needed reserves for the current policy is higher than the number available, by more than the accepted margin of 5% (about 5 disruptions), since $W_k = 18.64$, see column 2 in Table 4.

Table 5 presents a more detailed study of the requirements, only based on internal disruptions. As mentioned already, 374 flight blocks start each day, as well as 15 reserve blocks. Each day, there are on average around 7 recoveries. The expected number of internal disruptions of length j equals $p_{\text{int}} S_{jk}$ (column 3 in Table 5). On average, the 7 longest disruptions can be resolved with recovered crew members. The remaining disruptions have to be resolved with reserve crew (column 4 in Table 5). More than 17 reserves are required each day. So, this points to a lack of two reserves in our simplified model. This can also not be resolved with scheduling two disruptions in one reserve block. This is shown by multiplying the expected number of required reserves with their desired length, which represents the expected number of reserve days that have to be scheduled each day (column 5 in Table 5). This equals almost 120 reserves, compared to 105 in the current schedule.

8 Conclusions and Future Research

In this paper we developed several techniques to determine the number of reserve crew to be scheduled. The use of the statistical model is recommended since it incorporates the stochastical nature of disruptions and recoveries. This technique also uses a service level that can be interpreted quite easily, instead of the cover ratio. One topic for further research would be to investigate the effects of the SFB approach. This could be done by simulating the process of handling disruptions. Another option is to extend the comparison technique discussed in Section 6 to incorporate the configuration of the flight blocks and reserve blocks. It would also be interesting to look at more exact optimization methods like column generation and genetic algorithms (see for instance Guo [2] and Thiel [6]).

Our numerical results show that the statistical method halves the number of secondary disrup-

j	mirror longest flight opt 1	opt 2	opt 3	opt 4	opt 5	opt 6	mirror proportional	statistical	ideal statistical
1	0	0	0	0	0	0	0	0	0
2	0	0	0	0	0	0	0	0	0
3	0	0	0	0	0	0	0	0	0
4	0	0	0	0	0	0	0	0	0
5	0	0	0	0	0	0	0	0	1
6	0	0	0	0	0	0	4	0	8
7	15	0	0	0	0	0	2	0	4
8	0	13	0	0	0	0	2	2	4
9	0	0	0	0	0	0	1	2	2
10	0	0	11	0	0	0	1	3	3
11	0	0	0	10	0	0	2	3	3
12	0	0	0	0	9	0	0	1	1
13	0	0	0	0	0	5	1	0	0
14	0	0	0	0	0	3	0	0	0
15	0	0	0	0	0	0	0	0	0
16	0	0	0	0	0	0	0	0	0
$\sum_j T_{jk}$	15	13	11	10	9	8	13	11	26
$\sum_j j \cdot T_{jk}$	105	104	110	110	108	107	108	109	206
U_k	0.46	0.52	0.57	0.58	0.54	0.55	0.51	0.65	92.97
V_k	38.43	30.61	22.68	21.59	24.27	24.49	31.55	17.31	0.02
W_k	18.64	15.12	9.87	8.92	9.92	10.78	15.38	7.99	0

Table 4: The number of unused reserves U_k, secondary disruptions V_k, and unresolved disruptions W_k for the different approaches.

j	S_{jk}	disruptions	required reserves	scheduled reserve days
1	0	0	0	0
2	8	0.52	0.52	1.04
3	0	0	0	0
4	0	0	0	0
5	8	0.52	0.52	2.60
6	108	7.01	7.01	42.09
7	49	3.18	3.18	22.28
8	55	3.57	3.57	28.58
9	27	1.75	1.75	15.78
10	38	2.47	0.73	7.29
11	46	2.99	0	0
12	12	0.78	0	0
13	13	0.84	0	0
14	3	0.19	0	0
15	5	0.32	0	0
16	2	0.13	0	0
total	374	24.29	17.29	119.65

Table 5: The expected number of internal disruptions, required reserves, and scheduled reserve days.

tions compared to the current policy, comes close to the 5% accepted level of unresolved disruptions, while having the same low number of unused reserves. The current policy is already working quite reasonably in practice; the numerical results for our simplified model indicate that these can be improved by using the statistical model.

References

[1] Ní Éigeartaigh, T. and Sinclair, D. (2000) The Airline Crew Rostering Problem, School of Computer Applications, Dublin City University, *http://www.computing.dcu.ie/research/papers/2000/0300.ps* .

[2] Guo, Y. (2005) Decision Support Systems for Airline Crew Recovery, University of Paderborn, PhD thesis.

[3] Kohl, N. and Karish, S.E. (2004) Airline crew rostering: problem types, modeling and optimisation, *A*nn. Oper. Res. **127**, 223-257.

[4] Kohl, N., Larsen, A., Larsen, J., Ross, A. and Tiourine, S. (2004) Airline disruption management - perspectives, experiences and outlook, Carmen Research and Technology Report CRTR-0407, September 2004, *http://www.carmensystems.co.uk/research_development/research_reports.htm* .

[5] Ross, S.M., 2003, Introduction to Probability Models, Eighth Edition, Academic Press, San Diego.

[6] Thiel, M.P. (2005) Team-oriented Airline Crew Scheduling and Rostering: Problem Description, Solution Approaches, and Decision Support, Paderborn.

A Probabilities of having reserve capacity

1 `for` $k = 1$ `to` K
2 $V_k = 0$, $W_k = 0$
2 `for` $d = 1$ `to` $\sum_j S_{jk}$
3 $p\_r = p_d^{(k)}$
4 $l = l_d^{(k)}$
5 $\text{prob\_ext} = p_{\text{ext}} \cdot (1 - p\_r)/(1 - p_{\text{ext}})$
6 $j = K + 1$
7 $\text{prob\_not\_avail} = p_{0jk}$
8 `for` $i = 1$ `to` ∞
9 $p_{ijk} = \begin{cases} p_{i+1,j,k} \cdot p\_r + p_{ijk} & \text{if } i = 0 \\ p_{i+1,j,k} \cdot p\_r + p_{ijk} \cdot (1 - p\_r) & \text{otherwise} \end{cases}$
10 `next` i
11 $p\_r = \text{prob\_not\_avail} \cdot p\_r$
12 `for` $j = l$ `to` K
13 $\text{prob\_not\_avail?} = p_{0jk}$
14 `for` $i = 0$ `to` ∞
15 $p_{ijk} = \begin{cases} p_{i+1,j,k} \cdot p\_r + p_{ijk} & \text{if } i = 0 \\ p_{i+1,j,k} \cdot p\_r + p_{ijk} \cdot (1 - p\_r) & \text{otherwise} \end{cases}$
15 `next` i
16 `if` $(j > l)$
17 $\texttt{daysleft}_{j-l,k+l} \mathrel{+}= p\_r \cdot (1 - \text{prob\_not\_avail})$
18 `endif`
19 $p\_r = \text{prob\_not\_avail} \cdot p\_r$
20 `next` j
21 $V_k = V_k + p\_r$
22 `for` $j = l - 1$ `to` 1
23 $\text{prob\_not\_avail} = p_{0jk}$
24 `for` $i = 0$ `to` ∞
25 $p_{ijk} = \begin{cases} p_{i+1,j,k} \cdot p\_r + p_{ijk} & \text{if } i = 0 \\ p_{i+1,j,k} \cdot p\_r + p_{ijk} \cdot (1 - p\_r) & \text{otherwise} \end{cases}$
26 `next` i
27 $p_d^{(k+j)} \mathrel{+}= p\_r \cdot (1 - \text{prob\_not\_avail}) / \sum_i S_{i,k+j} \qquad \forall d \in D_{k+j}$
28 $p\_r = \text{prob\_not\_avail} \cdot p\_r$
29 `next` j
30 $W_k = W_k + p\_r$
31 `for` $i = 0$ `to` ∞
32 $p_{ilk} = \begin{cases} p_{ilk} \cdot (1 - \text{prob\_ext}) & \text{if } i = 0 \\ p_{ilk} \cdot (1 - \text{prob\_ext}) + p_{i-1,l,k} \cdot \text{prob\_ext} & \text{otherwise} \end{cases}$
33 `next` i
34 `next` d
35 $U_k = \sum_j \sum_i i p_{ijk}$
36 $p_{i,j,k+1} = \sum_{l=0}^{i} p_{l,j+1,k} \cdot p_{i-l,j,k+1} \qquad \forall i, j$
37 $p_j = \text{daysleft}_{j,k+1} / \sum_j T_{jk}$
38 $p_{i,j} = \binom{\sum_m T_{mk}}{i} p_j^i (1 - p_j)^{\sum_m T_{mk} - i}$
39 $p_{i,j,k+1} = \sum_{l=0}^{i} p_{lj} \cdot p_{i-l,j,k+1} \qquad \forall i, j$
40 `next` k

For every day k, we treat every scheduled flight as a possible disruption $d \in \{1, \ldots, \sum_j S_{jk}\}$. The impact of a disruption is represented by lines 3–33. Flight block d with length $l_d^{(k)}$ has an initial probability $p_d^{(k)}$ of getting disrupted. Let us introduce the following probabilities:

prob_not_avail	=	the probability that no reserve block of length j is available
prob_ext	=	the probability that there is an external disruption
p_r	=	the probability that a crew member on flight d is disrupted, and still requires a reserve crew member

To resolve this disruption d we first check whether a recovered crew member is available (i.e., a reserve block of length $K+1$). Therefore, the probabilities p_{ijk} are adjusted according to line 9. When there is no recovery available, a reserve is still required, and p_r gets updated (line 11). Next, we check whether a reserve block of at least $l_d^{(k)}$ days is available (lines 12–20). No secondary disruption will occur in these situations. The procedure to update p_{ijk} is the same as on line 9. If a reserve block with length j ($j > l_d^{(k)}$) is available, the remaining reserve days of this reserve block ($j - l_d^{(k)}$ days) become available after resolving flight d (at day $k + l_d^{(k)}$). This is represented on line 17. If there is still no reserve crew found (with probability p_r), a secondary disruption occurs. Therefore V_k increases. Again the same update procedure is used to update p_{ijk}. A secondary disruption occurs at day $k + j$ since the reserve crew member cannot perform its scheduled flight block after his/her reserve block. Since it can affect any of the flights, all flights at day $k + j$ obtain a higher chance of getting disrupted (line 27). If there is still no reserve available, an emergency crew member has to be called (i.e., W_k increases).

When the disruption is external, the crew member becomes available to resolve other disruptions (line 32). When all disruptions on a day are taken care off, we can calculate the expected number of remaining reserves U_k. These reserves can be used the next day (line 36). The probability of having reserve blocks of j days remaining and becoming available at day $k+1$ equals p_j (line 37). The actual number of reserve blocks becoming available with j days remaining has a binomial distribution (line 38). These have to be added to the reserve crew (line 39).