Cockpit crew transition planning optimisation

by

A.A.J. Hooijen

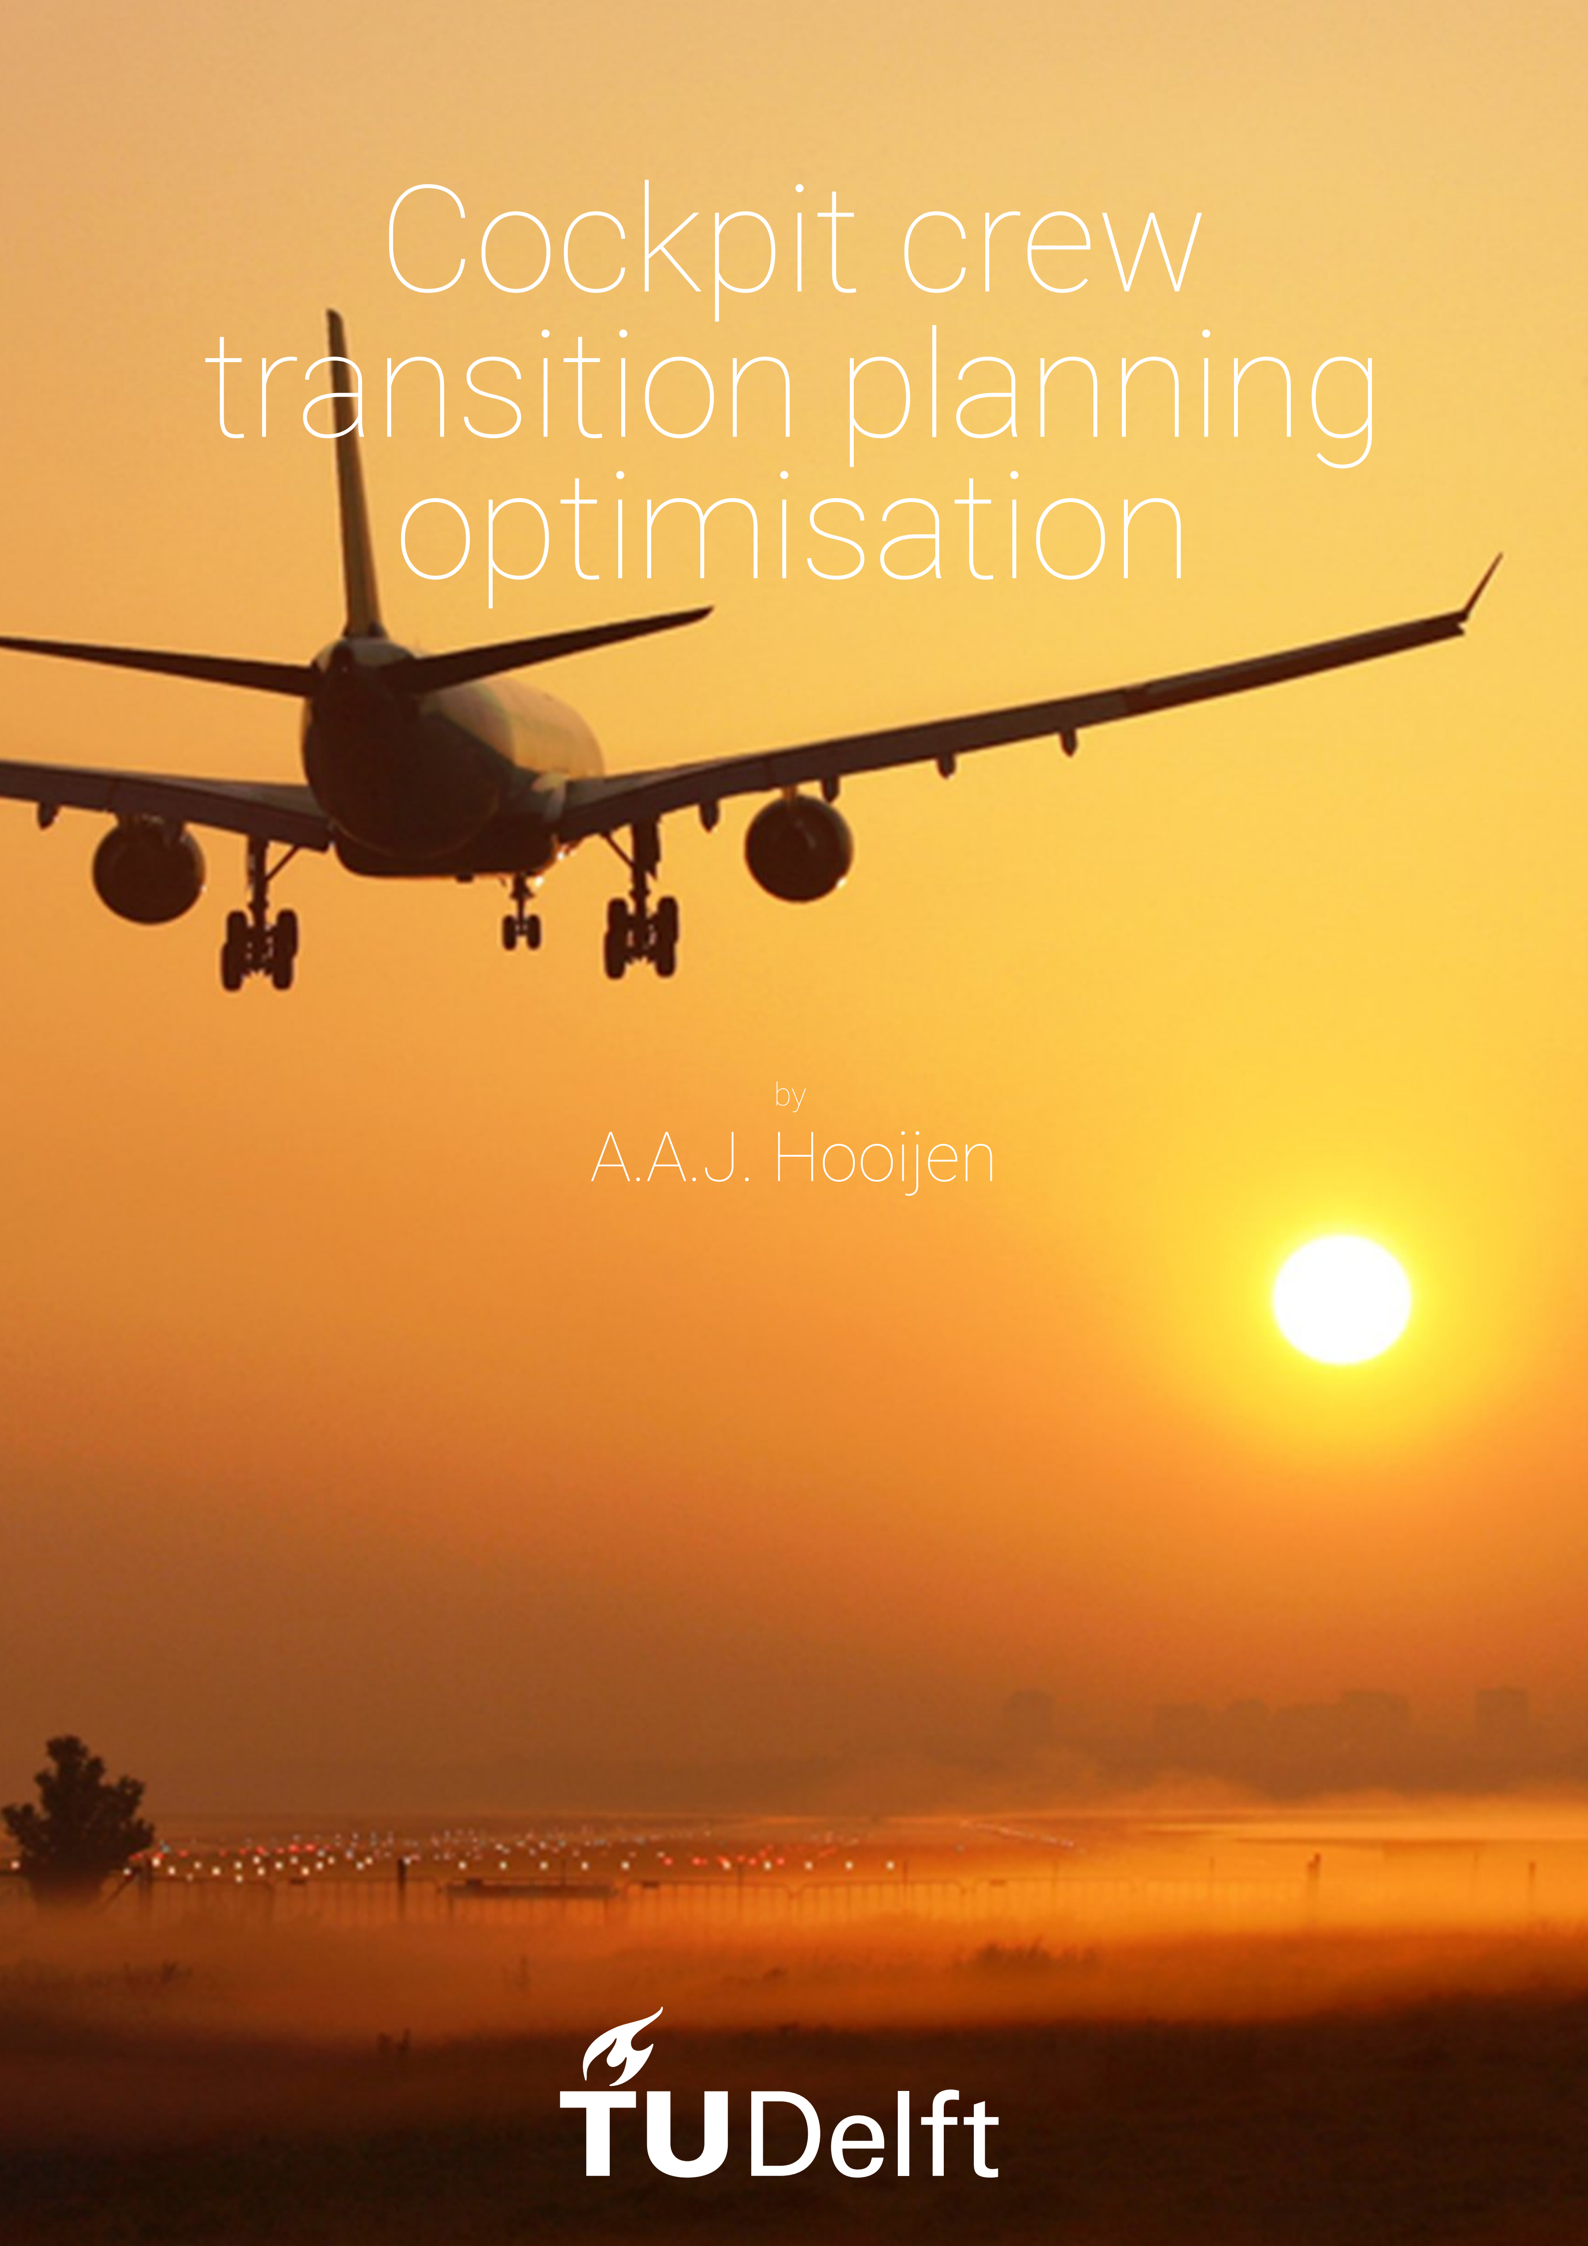

Cockpit crew transition planning optimisation

by

A.A.J. Hooijen

to obtain the degree of Master of Science
at the Delft University of Technology,
to be defended publicly on Friday May 3, 2019 at 14:30.

Student number:	4079418	
Project duration:	May 16, 2018 – May 3, 2019	
Thesis committee:	Prof. dr. R. Curran	TU Delft
	Dr. ir. B.F. Lopes dos Santos,	TU Delft
	Ir. drs. L. Scherp,	TU Delft

An electronic version of this thesis is available at http://repository.tudelft.nl/
Cover image source: https://pbs.twimg.com/media/Cwv5xnmWEAAYRbe.jpg

TU Delft Delft University of Technology

Preface

This thesis marks the final part of my Master's of Science degree in Aerospace Engineering. It presents my graduation research regarding the cockpit crew transition planning problem carried out in collaboration with KLM Royal Dutch Airlines. The opportunity to carry out this research in-house at the largest airline of The Netherlands has been an invaluable experience, for which I am very grateful.

Through this writing, I would like to express my acknowledgements to the people who have supported me throughout the project which started almost a year ago.

First of all, I would like to thank my supervisors during the project, Lennart Scherp and Bruno Santos, for their input in progress meetings and discussions, for setting up this project and for creating a welcoming and warm environment for doing the research. I would also like to thank my colleagues at KLM who have been both a supportive force as well as a welcoming distraction from time to time. Without this group, the experience would not have been the same.

Furthermore, A huge thanks to all people at KLM who have helped me reach my goals and have provided important input, support during the project. Especially, I would like to acknowledge Nico Scheeres for setting up the project and Danielle Damen for our frequent discussions on the progress and the research.

Finally, I would like to thank my family and friends who have been incredibly supportive of me throughout my academic career.

Toine Hooijen
Schiphol, April 2019

Executive Summary

Pilots are an important asset in the operations of airlines, as with insufficient pilots, an airline cannot operate the planned schedule. The problem of managing the pilots is done in the crew planning department. This department is in charge of hiring, promoting and scheduling pilots to flights. Within the field of airline cockpit crew planning, the long-term crew planning problem is defined to optimally plan the supply and demand for pilots as well as methods and strategies for closing the gap between these. The demand is determined by the number of pilots required at each crew position based on the demand for the flight schedule, contractual agreements for holidays and off-days and training requirements. The supply, on the other hand, is determined by the current pilot workforce and changes over time by factors such as retirements and illness. Finally, the inevitable gap between supply and demand can be closed by transitioning pilots to other crew positions and by recruiting new pilots.

These transitions are necessary as pilots are typically only certified for one crew position (i.e. a single function on a single aircraft type). Furthermore, airlines utilise a promotional hierarchy in which pilots can not move from every position to another but move through defined paths within the system (Ives, 1992). Recruits can only enter the system on some of the lowest ranked positions. Next to these systems that complicate the problem, the fact that the plan is made far in advance further convolutes the problem. This introduces a number of uncertainties for example in the flight schedule and fleet composition, pilot's preferences regarding crew positions and holidays, pilot's illness and other unavailabilities. Because of the large complexity, research often aims to solve a simplified or partial problem instead of aiming to solve and analyse the full problem.

Little research has been done to determine crew staffing and transitions at the strategic level specifically for airlines. It is, however, important to have more knowledge of the transition planning problem to be able to assign budgets with greater accuracy and be able to adjust to different scenarios. Because the transition planning problem is such a complex and complicated problem, large software packages are used which takes days to solve the problem and deliver a crew transition plan. Because of this large computation time, it is difficult to analyse the implications of different strategies, varying stochastic parameters or changes in the demand or supply. Such a scenario analysis tool can be incredibly useful in the decision-making process of airlines as it gives an airline more data-driven information regarding different scenarios and strategies.

The presented research focussed on developing a decision support tool that can be used to analyse the effect of different strategies and scenarios on the cockpit crew transition planning problem. A research framework has been set up from which the research question is to be answered. The chosen strategy was to first develop a heuristic planning model that is able to plan transitions with the goal to determine an optimal crew plan. This heuristic planning model uses a local search algorithm that aims to improve the crew plan by evaluating solutions in the neighbourhood of the current solution. This neighbourhood is defined as all solutions in which one transition is added to the current solution. This means that at each iteration, a transition is planned, all necessary parameters are updated and the balance is recalculated until a stopping criterion is met.

The objective function of the optimisation has been designed to accurately reflect the cost of having a higher (surplus) or lower (shortage) supply compared to the demand for different crew positions. In order to reflect the relative importance between different positions, the shortages and surplusses are multiplied with the salary cost of the position. Furthermore, the shortages are multiplied with the number of consecutive shortages in the balance for that position, as having a shortage for a longer period of time greatly increases the impact this has on the business operations.

As mentioned, the planning model is a local search algorithm which aims to plan transitions to crew positions in which shortages are present. In order to decrease the size of the neighbourhood search, several methods are used. A rule-based system aims to disregard transitions that do not comply with certain rules. These rules are based on parameters such as the balance of a given transition option or the available capacity. Furthermore, a tabu-search method has been implemented in which past actions are placed on a tabu-list for a number of iterations to prevent the model to get stuck in a local optimum.

With the heuristic planning model designed, the question arises which of the available transition options in the neighbourhood should be selected at each iteration in order to obtain the best possible solution. To do this, a tree search method is proposed. This method generates a tree of the available transition options per iteration and

explores the tree in order to find an optimal solution. Different configurations of the model can be created by changing a set of parameters which restrict the search space, amongst which are a naive selection algorithm and a greedy algorithm. Furthermore, a variation of the tree search method is developed which explores the tree similar to the shortest path algorithm developed by Dijkstra (1959). This method aims to further decrease the computation time of the model without compromising the solution quality.

In order to test the various configurations of the model, an experiment was designed in which each configuration was used to solve 10 different scenarios. These scenarios were created from data from a reference airline. By importing this data and solving the transition planning problem from a given start date for a period of 12 months, the performance of the different configurations in terms of solution quality, solution stability and computation time could be compared. In order to compare the results for different scenarios, the results have been scaled per scenario using min-max normalisation. This method scales the objective function values of the different configurations on a minimum value of 0 and a maximum value of 1.

From the results, several conclusions have been drawn. The shortest path method is able to determine solutions faster than the tree search method without compromising the solution quality. Furthermore, it is concluded that a combination of width and depth produces the best and most stable results, as opposed to configurations focussing on either width or depth. Finally, from the results of the naive selection and greedy configurations, it can be concluded that a minimum depth of 2 levels is required, as this greatly improves the stability of the solutions. The best performing configuration was the configuration utilising the shortest path method with a width of 3 and depth of 4. This configuration resulted in the second-best solution quality, the best solution stability and the fourth lowest computation time.

Following the discussion of the results, it was concluded that all presented configurations of the solution method are accurate enough to solve the cockpit crew transition planning problem. Depending on the application, different requirements with regards to computation time, solution quality and solution stability might be set. These requirements will also influence what configuration will be best suited for the application. The academic contributions of this research to the scientific body included a first definition of the cockpit crew transition planning problem, specifically designed solution methods for this problem and the wide applicability of the developed model to the airline industry as well as other industries. A more practical contribution of the developed model was that the model can be used by airline and other companies to quickly analyse different strategies, scenarios and the effect of variations in stochastic parameters in the system. This helps airlines in their decision-making process and ultimately improves business operations.

List of Illustrations

Figures

Tables

Contents

1 Introduction

Pilots are an important asset in the operations of airlines, as with insufficient pilots, an airline cannot operate the planned schedule. The problem of managing the pilots is done in the crew planning department. This department is in charge of hiring, promoting and assigning pilots to flights. Within the field of airline cockpit crew planning, the long-term crew planning problem is defined to optimally plan, up to five years in advance, the supply and demand for pilots as well as methods for closing the gap between these. The demand is determined by the number of pilots required at each crew position (which is a combination of the pilot's rank and the aircraft type), which is based on the demand for the flight schedule, contractual agreements for holidays and off-days, training requirements. The supply, on the other hand, is determined by the current pilot workforce and changes over time by factors such as retirements and illness. Finally, one of the methods for closing the inevitable gap between supply and demand is to assign transitions for pilots to other crew positions and to recruit new pilots.

These transitions are necessary as pilots are typically only certified for one crew position. Furthermore, airlines utilise a promotional hierarchy in which pilots enter the system at one of the lowest positions and get promoted to higher ranked positions troughout their career until they reach the highest crew positions and eventually retire (Ives, 1992). These rules complicate the problem as pilots cannot be hired for every position but instead have to be transitioned from other positions. Also, a crew plan is made far in advance, which further convolutes the problem as this introduces a number of uncertainties in, for example, the flight schedule and fleet composition, pilot's preferences regarding crew positions and holidays, pilot's illness and other unavailabilities. Because of the complexities, often a simplified or partial problem is solved in research instead of aiming to solve and analyse the full problem.

Because the transition planning problem is such a complex and complicated problem, software packages are used which takes days to solve the problem and deliver a crew plan. As crew costs are one of the largest expenses for an airline. This makes it even more important to be able to analyse different planning strategies, varying stochastic parameters or changes in the demand and supply. Such an analysis model can be used as a decision support tool and provide a more data-driven approach in the decision-making process in the airline cockpit crew planning process. No research has, however, been published that focuses on fast solution methods for the cockpit crew transition planning problem while such a tool could help airlines improve the business operations and decrease the cockpit crew costs.

The main topic of this report has been defined as the airline cockpit crew transition planning problem. The research aims the following research question:

How to model the transition planning of cockpit crew to provide insight into future staffing levels and transitions and analyse different planning scenarios, strategies and assumptions and their long-term effect?

The objective of this research is to create a decision support tool that can be used to give better insight into future staffing levels and transitions. This decision support tool should be able to analyse different planning strategies, scenarios and variations in input parameters such as the demand for cockpit crew.

In order to solve the cockpit crew transition planning problem, a selection algorithm has been developed that is based on a tree search method. This selection algorithm aims to select the optimal transition in a heuristic local search algorithm that is able to plan transitions and consequently calculate the implications to the supply and demand for pilots. The selection algorithm's configuration can be changed by varying a number of parameters that change the search space. By doing this, well-known methods such as the greedy algorithm or naive selection algorithm can be replicated. A variation of this tree search selection algorithm has also been developed which is based on Dijkstra's shortest path algorithm. The variation is developed as it is expected that this method is able to decrease the computation time without decreasing the solution quality. Using the selection algorithms, multiple model configurations are designed that are then tested in experiments created by actual supply and demand data from a reference airline.

This report is structured as follows: in Chapter 2, a research gap within the field of airline cockpit crew planning is defined, after which the current state-of-the-art regarding this research gap is presented in Chapter 3. Then, the design of the research, including a detailed presentation of the problem, is presented in Chapter 4. In order to answer the research question that has been designed, a heuristic planning model has been developed in Chapter 5 and the selection algorithm to be used in conjunction with the planning model is proposed in Chapter 6. Different configurations of the selection algorithm are tested in a set of experiments in Chapter 7. After the various selection algorithm configurations have been tested, the application of the model to a practical problem is presented in Chapter 8 and a sensitivity analysis on the model parameters is performed in Chapter 9. The report is then concluded with the conclusions drawn from the research, the contributions of the research and recommendations for further research in Chapter 10.

2
Research Gap

In this chapter, the domain of airline planning (Section 2.1) will be explained. Next, the cockpit crew planning is discussed more specifically (Section 2.2) and finally, within the domain of cockpit crew planning a research gap is defined (Section 2.3) that forms the basis of the literature survey in the following chapters.

2.1. Airline Planning

The airline planning department is involved in maximizing the airline's profitability by optimizing future aircraft and crew schedules. As this problem is characterized by numerous complexities like flight network, different aircraft types, crew labour agreements, restrictions on gate and airport usage, noise, maintenance requirements and many more, it is quite impossible to design a single optimization model to solve the problem.

For example, Clarke and Smith (2004) state that the airline planning process typically starts with fleet planning and is followed by four consecutive procedures; crew scheduling, aircraft maintenance routing, airport resource management and revenue management. Because of the complexities, the problem is generally separated into different subproblems that are solved sequentially (Barnhart, Belobaba, & Odoni, 2003):

- Schedule design: Design of the airline's network schedule by assessing the market and it's demand for different regions, countries and cities.
- Fleet assignment: Determining the aircraft size and type to optimally serve all flights.
- Aircraft maintenance routing: Routing aircraft in such a manner that maintenance requirements are met at minimum cost.
- Crew scheduling: Assigning cockpit and cabin crew to flight again at minimum cost.

To summarize, the airline planning process can be roughly divided in the development of a flight network, fleet plan, maintenance plan and crew schedule. Because of this separation into smaller subproblems that are solved independently of each other, suboptimal yet feasible plans are created. These plans are suboptimal as the solutions are constraint to the optimal result of the previous process instead of finding a global optimum over the entire process.

Out of these similarly defined subproblems, crew planning and scheduling is often regarded as one of the most, if not the most, important problems (Sohoni et al., 2004). Not only are the crew costs one of the largest operating expenses of an airline, it is also a highly complex problem because of strict labour agreements, large uncertainties in the planning process and shear problem size.

2.2. Crew Planning

The definitions of airline planning presented in the previous section all identify crew scheduling as part of the planning process. However, most do not recognize the steps taken before actually scheduling the crew as well as any measures that have to be taken after the schedule has been published, due to unforeseen circumstances and instead just limit the problem to the scheduling problem. As can be seen in Figure 2.1 (from Sohoni et al. (2004)), the crew scheduling process defined by Barnhart et al. (2003) only encompasses the final one to 1.5 months before the day of operations. The crew planning process, however, is a much more complicated and lengthy process that already starts approximately five years before the day of operations. The process before the crew scheduling consists of determining the number of crew required in the future, assign vacation, promotion and determine strategies to minimise the inevitbale gap between supply and demand. This process can be subdivided into long-range or

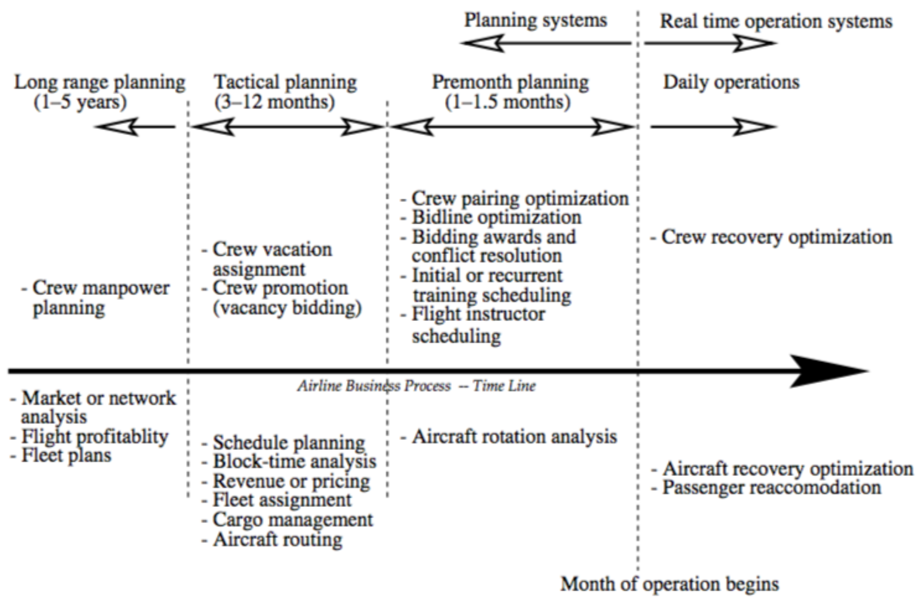

Figure 2.1: Graphic showing typical crew planning timeline (Sohoni et al., 2004).

strategic and tactical planning. After scheduling the crew, disruptions are always possible due to illness, delays and more. In order to minimise the impact this has on the schedule, recovery optimisation is necessary close to the day of operation.

For the scope of this project, the crew recovery problem is also added to the complete planning problem in contrast to the definition by Sohoni et al. (2004). To describe the entire process, working backwards in time from the day of operations, the crew planning problem consists of crew recovery on the day of operations, operational planning or scheduling for the final one to 1.5 months, tactical planning up to 1 year and strategic planning up to 5 years before the day of operations. These four parts are discussed in Subsection 2.2.1 - 2.2.4, respectively.

2.2.1. Crew Recovery

No matter how good an initial plan is, things never go perfectly according to plan. On the day of operations, airlines face disruptions to their flight schedule because of unexpected causes like severe weather conditions, aircraft unavailability and crew absence (Yu, Argüello, Song, McCowan, & White, 2003). If these disruptions are not managed properly, they can cause flight delays or even cancellations. It is, therefore, important to reassign crew and assign reserve crew in order to quickly cover open flights in case of regular crew's absence. Obviously, this recovery should be done in a cost-effective way while still honouring all government, contractual and collaborative labour agreements (Lettovský, Johnson, & Nemhauser, 2000).

Decisions to add, cancel, delay and divert flights create situations in which crews either arrive late for their next flight, or even not arrive at all, and cannot serve the flights in their schedule, calling for changes in the determined schedule. Somehow, crews have to be arranged to fill these open flights and if necessary, crew missing their flight have to be brought back to their home base or to the start location of their next flight. This can be partially done by scheduling a number of reserve crew in advance who are on standby for this kind of situations. The goal is then to cover as much of these open flights with reserve crew, thereby minimizing additional cost for cancelled and delayed flights, incremental crew compensation and passenger surcharges. Furthermore, creating solutions quickly limits the extent of disruptions and make that airlines can avoid additional delays and cancellations, improve on-time performance, reduce the number of passengers to reaccommodate, and preserve passenger goodwill (Yu, Argüello, et al., 2003). Sohoni, Johnson, and Bailey (2006) proposed an integrated planning and rostering approach by producing monthly schedules of on- and off-duty days at a tactical level. On the other hand, multiple researchers have investigated reserve crew pairing for both cabin and cockpit crew (Bayliss, 2016, Shebalov & Klabjan, 2006, Nissen & Haase, 2006).

2.2.2. Operational Planning

The operational crew planning phase at airlines is called crew scheduling in literature (Barnhart et al., 2003, Clarke & Smith, 2004, Medard & Sawhney, 2007). Within the crew scheduling problem, (Medard & Sawhney, 2007) define two subproblems. The first step is to create working patterns called pairings which are subsequently assigned to a specific crew in the second step. This method is usually used in European airlines. In contrast to this method, US airlines usually string a number of the pairings together to create bidlines and then let pilots make bids for certain bidlines.

The first step in this process is often called the crew-pairing problem (Clarke & Smith, 2004). It is often solved separately per aircraft type as these systems are independent of each other and consists of creating a set of pairings that cover a given set of flights at minimum cost. The pairings obviously have to conform to the prevailing labour and airline agreements. For the crew-pairing problem, Gershkoff (1989) has defined an optimization model that uses a set-partitioning framework in which the columns represent pairings that are constructed by the flights. The objective of the model is to minimize the cost of flying the published airline schedule.

After the pairings have been constructed, the problem to be solved in the second step is to assign crew to all available pairings in the crew-rostering problem. Rest periods, preassigned activities like leaves of absence, training and other tasks Clarke and Smith (2004) have to be taken into ccount in this step. This problem has already been formulated in 1989 by Gershkoff as an integer programming problem, solvable with any commercial optimization software. However, in these early efforts, the lack of computing power made it difficult to perform a global optimisation so most research focused on efficient heuristics to solve the problem. Later, global optimization models have been developed to solve the problem with a greater degree of precision (Fahle et al., 2002, Kohl & Karisch, 2004).

More recently, Cacchiani and Salazar-González (2017) proposed to merge the fleet assignment, aircraft routing and crew pairing problems into an integrated optimisation model. The objective of this model is to minimize a weighted sum of the number of aircraft routes, the number of crew pairings, and the waiting times of crews between consecutive flights. Additionally, it aims to maximize the robustness of the solution by minimizing the number of times a crew changes aircraft. Even though the model shows an improvement in the objective function of 5 to 10 percent, it also increases the computation time from a couple of minutes to hours. Shao, Sherali, and Haouari (2017) propose an integration of the same three planning parts and conclude that computational results obtained through the use of data from a major U.S. airline show the benefits of this integrated approach.

Jarrah and Diamond (1997) have developed a comprehensive approach for addressing the problem of generating bidlines consisting of a set of pairings for a major US airline. The approach uses a set partitioning model in which the objective is to maximise the covered credit time while simultaneously minimising the number of bidlines. The model reduced both the time and the amount of resources needed for generating bidlines. At the same time, because of the use of several parameters, it gave schedulers the possibility to perform what-if analyses and control over the quality of the bidlines.

2.2.3. Tactical Planning

As Verbeek (1991) defines it, the tactical manpower planning problem poses the question of when to recruit pilots and when to schedule transitions for pilots from one function to another. The goal is to minimize shortages and surplusses of pilots in all different positions as well as to minimize the transition training cost. Furthermore, in the tactical planning phase, the pilot's vacation allocation model aims to allocate annual leave to all pilots. For this allocation, the pilot's preferred vacation dates, planned training days and the minimized pilot surplusses or shortages have to be taken into account.

At present, no solution with optimization has been presented for the staffing problem, that is the decisions on who to transfer where. Most research instead assumes that manual decisions have been made for transfers and focuses on the transitioning and training parts in the scheme. It starts with the question when a pilot should transfer and not if or where he should do so. Verbeek (1991) presents a Decision Support System (DSS) for manpower planning in airlines. Gang Yu has, over the years ((Yu et al., 1998), (Yu, Pachon, & Thengvall, 2003), (Yu et al., 2004)), developed a DSS for Continental Airlines. In Yu et al. (1998) a heuristic solver for the integrated transition and training problem was presented. It starts with a manual solution to the staffing problem and tries to assign weeks and schedules for the pilots that need to take a course during the planning period. The heuristic works by assigning training courses in chronological order and for each week looking a number of weeks ahead to see if a course during that week

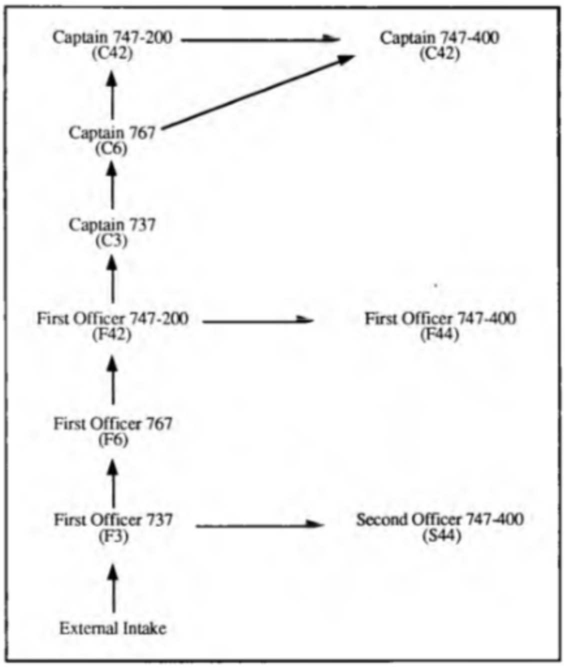

Figure 2.2: An example of the hierarchical transition system for airline pilots (Ives, 1992).

would cover a lack of supply in one of the upcoming weeks. The heuristic also uses a couple of different assigning rules but is at heart a greedy heuristic with some soft costs balancing the solution. Yu et al. continued to work on the problem and in (Yu, Pachon, & Thengvall, 2003) and (Yu et al., 2004) present an optimization-based solution to the transition and training problem. The solution algorithm splits the problem into two parts, the transition phase and the training phase, and does not try to integrate them as his earlier heuristic did. The transition algorithm works by first solving an LP-relaxation to find a lower bound of demands that is impossible to meet. This bound is used as input in the objective function of a MIP (Mixed Integer Program) that is solved by commercial software. In the second phase, the training problem is solved using the solution found in the transitioning phase (Thalén, 2010).

In a survey of the different kind of optimization problems within the field of airline operations concerning the tactical planning phase, Holm (2008) presents a number of different mathematical models and also tests an optimization of the transition, vacation and recurrent training scheduling problem. For this integrated problem approach, commercial software is used and a manual solution is used as a starting point of the optimization.

2.2.4. Strategic Planning

At present, transition planning is mostly only performed for one year or season (Holm, 2008), typically, this is also the period for which the pilot's preferences regarding transitions and vacation are known. This means that the only transition planning that is performed is within the tactical planning phase, however, outside this time window, it is also important to have a reasonably accurate prediction and planning of the pilot staffing levels and required transitions and thus training capacity. This is especially the case in airlines, where training of pilots is expensive and time-consuming and hierarchical systems are in place where pilots move up one step in the ladder each time (see for example the promotional hierarchy of Air New Zealand in Figure 2.2). This results in a slow moving system when a highly ranked pilot retires and no surplusses are available within the system as every transition from a lower ranked position creates new shortages in that position. No literature has been found regarding the practices of making an accurate long-term transition or manpower plan that takes uncertainties in supply and demand into account (Holm, 2008).

According to Holm (2008), strategic cockpit crew planning problem is a process of three elements:

1. Analyzing and predicting the cockpit crew staffing levels needed to achieve the objectives of the organization. This includes covering all flights, training requirements, additional functions and holidays.

2. Predictions of the future supply of staff in the organization through current staffing, retirement, planned and unplanned absence, etc.
3. Determining policies to minimize the inevitable differences between the former two by recruiting new pilots and transitioning pilots to different positions.

In contrast to the tactical planning phase, no information regarding the pilot's position bids is known during the strategic planning as pilots have to announce their preferences 6 to 12 months in advance. However, it would still be desirable to plan transitions alarger number of years in advance with some form of accuracy to provide more insight in future staffing levels and the airline's resilience to growth, policy changes, disappearance or emergence of a new aircraft type etc. Therefore, in the strategic planning phase, the problem at hand is first to determine the number of pilots required in the future based on long-term company- or network plans as well as predicted retirement, absence and more. Subsequentially, accurate estimations of staffing levels and required transitions should be made to give an overview of the staffing levels and potential bottlenecks at different crew positions.

2.3. Research Opportunity

Even though cockpit crew has been identified as one of the largest operating expenses for airlines (Holm, 2008) and a lot of effort has been made to optimize the process, almost all research in the past has focused mostly on scheduling and recovery problems. Meanwhile, planning problems have been largely ignored, eventhough solving this problem could lead to higher savings due to the large salary and transition costs for airline pilots. Little research has been done to determine crew staffing at strategic level specifically for airlines. It is important to have information about the crew planning problem in advance because of the slow-moving hierarchical system in place at most airlines. With the current system, accurate predictions are only known one season to one year in advance, while budgets, fleet and network plans are typically already defined up to 10 years in advance (Rosskopf, Lehner, & Gollnick, 2014). This leads to late decisions with respect to cockpit crew transitions and staffing levels which almost always results in higher costs. It is therefore important to have more knowledge of the long-term planning problem to be able to assign budgets with greater accuracy and to adjust to different scenarios.

The strategic crew planning problem consists of predicting the supply and demand for cockpit crew in the future and subsequently closing the gap between the two by recruiting new pilots and transitioning pilots to new positions. Ideally, this problem should be solved in an integrated manner, where supply, demand and transitions are dynamically adjusted to each other to generate a long-term solution under the uncertainties involved in the problem. This solution can be used for an analysis of the effect of different company strategies, new aircraft type and other managerial decisions.

In order to provide a clear definition of the goal of the research following the defined research opportunity, a research question can be determined. The research question has been defined as:

How to model the transition planning of cockpit crew to provide insight into future staffing levels and transitions and analyse different planning scenarios, strategies and assumptions and their long-term effect?

Subsequentially, this research question can be moulded into research objectives that define what should be attained in order to properly answer the research question. The objectives of this research are to:

[a] Create a decision support system by giving better insight into future staffing levels and transitions.
[b] Be able to study the effect of different scenarios and strategies on long-term cockpit crew planning.
[c] Perform an analysis on the different strategies and assumptions involved in strategic crew planning.
[d] Make recommendations on the assumptions and strategies for cockpit crew transition planning.

The next chapters in this report discusses the available literature with regards to the three different parts of manpower planning in airlines as defined by Holm (2008), manpower planning practices in other industries and some methodologies that have been developed for related problems, either in other industries or within different fields within the airline industry, that are thought to be applicable.

3

Literature Review

3.1. Current Practices

In this section, the available literature and current practices regarding the tactical and strategic cockpit crew planning process will be discussed. First, in Subsection 3.1.1, methods regarding the assessment and predictions of crew supply and demand will be discussed, followed by different transition strategies in Subsection 3.1.2. Finally, a synthesis of the current practices in cockpit crew planning will be presented in Subsection 3.1.3.

3.1.1. Demand and Supply

In an ideal world, the number of pilots to supply to the system would be equal to the demand generated by the flight schedule. However, this is seldom the case. First of all, the nett demand (dictated by the flight schedule) is increased by flexible demand for holidays, training requirements and additional functions. This demand is flexible as the precise allocation or schedule can be determined (to a certain extent) by the airline while the total over a given time window is predetermined by regulations and contracts. Secondly, the supply of pilots for flights is never the same as the number of pilots hired by the airline as pilots get sick, crew get disrupted, etcetera. Therefore the total number of pilots per seat to be hired to have enough pilots available has to be predicted in advance in order to be able to fly the scheduled flight plan.

Within the description of an integrated manpower management system, Yu, Pachon, and Thengvall (2003) discuss two important factors in the assessment of demand and supply of airline pilots. Demand forecasting is mentioned as the first critical step in manpower planning. According to Yu, Pachon, and Thengvall the need for cockpit and cabin crew is estimated based on the airline's fleet when planning more than a year in advance and based on the number of block hours in the published flight schedules when planning less than a year in advance. This estimation for the need of crew is divided by the expected utilization rate to account for training, vacation and other absences. As the proposed model is developed for Continental Airlines, it can be assumed this is the preferred method for this airline.

The next parts of Yu, Pachon, and Thengvall's model discuss absence and vacation management. For absence management, the module takes unplanned or unstructured absence requests from the crew as input in order to feed actual absence data to the system. With the lack of more information, it is assumed that this module estimates future absence based on the provided input, as, during the tactical or strategic planning phase, these unplanned absences are generally not known yet. For the vacation management, the module plans and proposes vacation periods for a crew to bid on. This allocation is based on the demand fluctuations of the airline since it is preferable to plan more vacation in periods with low nett demand. These two modules are subsequentially being followed by transition management, which will be discussed in Subsection 3.1.2.

Verbeek (1991) chooses to allocate the flexible demand within the year so as to minimise the cost of shortages and surplusses. According to Verbeek, this problem can easily be solved as a minimum-cost network-flow model. Similar to (Yu, Pachon, & Thengvall, 2003), this is done before transition management instead of in conjunction with it, which results in sub-optimal solutions since they are not independent. While at Continental Airlines, pilots' vacation is allocated by allowing pilots to bid for certain pre-proposed vacation blocks, it is also possible to allow pilots to register their preferred vacation dates and use these as input in the allocation of budget and planning of vacation

3.1.2. Transition Strategies

Verbeek (1991) describes the process of designing a decision support system (DSS) for crew planning at KLM. The goal of this system is to improve both the effectivity in terms of quality of the plans and a faster planning process and the efficiency by reducing the effort that is needed to create a plan. Even though the focus of this paper is on the design of a DSS in general, some important basics of the theory and mathematics behind the system are also given. Because of the large scale of the problem (19 seats, 120 months and 900 pilots), the mixed integer model could not be solved with any software and therefore a heuristic approach was used to come up with a feasible but sub-optimal solution. However, the paper does not go into detail about the (heuristic) methods used to solve the problem.

When a plan is generated using the system, several subproblems are solved to come up with an entire plan. After the allocation of flexible demand, the model is able to resolve shortages automatically by planning transitions from one seat to another. However, through these transitions, the shortages and surplusses change, and so does the optimal flexible demand allocation. Therefore, both problems are ideally solved in one formulation as a mixed integer problem. Finally, the allocation of instruction tasks per fleet can be modelled as an LP algorithm.

After a plan is generated by the model (Verbeek, 1991), the planner is able to evaluate the created plan manually using a number of metrics. The seat-year screen presents the shortages and surplusses for all seats and years in a matrix, while the transition matrix shows all transitions in one year grouped per 'from-to' pair. The experience development graph plots the ratio of experienced pilots in a seat over the total number of pilots versus time. For the flexible demand, a bar chart is shown with for every month a bar showing the amount of seasonal and yearly flexible demand allocated and a bar showing the balance between supply and demand. The training capacity graph plots the required and available simulator capacity over time per aircraft type. The mutation matrix meanwhile shows, for all seats, the values of gross supply, retirements, attrition, transitions out and transitions in per year and finally, the seat survey table details all parts of the supply and demand for twelve months for one seat.

As mentioned, the planning horizon for the model is 10 years, however, pilots could change their transition bids every season (6 months). Obviously, human behaviour is quite uncertain and this introduces some uncertainty in the pilots eligible for transition to a certain seat and thus also in the overall solution. To cope with this, Verbeek simulates the individual pilot's behaviour, however, no details about the structure of this simulation are given. So although the model is planning far into the strategic planning phase, no further information is given on how to handle the uncertainties associated with planning for this time frame. Verbeek also briefly mentions the supply and demand of pilots and while the flexible demand is allocated in the model, predictions of the supply are not taken into account for the described model.

A year later, Ives (1992) developed a linear programming model with the goal of finding the optimal promotion schedule for cockpit crew based on current staffing and changes and demand for up to two years in advance. As the hierarchical system used in this model is almost fully linear (except for the highest ranked position which has two preceding seats, see Figure 2.2), there is always only one position from which a pilot can be transitioned to a position with a shortages, which greatly simplifies the problem and makes the sequential solution of first planning transitions and then choosing which pilots to fill the transitions possible without deteriorating the solution. The solved problem incorporates a 79 week period with only two intermediate demand changes and does not take any other factors, like recurrent training, vacation and unplanned absences into account.

After its first definition of the manpower planning problem in (Yu, Pachon, & Thengvall, 2003), Yu et al. (2004) present the Crew ResourceSolver system for Continental Airlines in more detail. The model consists of modules to handle staffing, vacation, planning and training, as displayed in Figure 3.1.

The staffing module identifies shortages and surplusses for all seats based on the planned flight schedule, current staffing and already planned transitions and attritions. Based on the shortages and surplusses, pilots are transitioned to different seats based on seniority and their position bids. Since this module does not perform any optimization, no guarantee can be given that the awarded transitions are optimal and the module can only be used when pilot transition bids are known. Therefore, it is only applicable to the tactical planning phase. Next, the vacation module determines vacation periods on which pilots can bid. In contrast to the situation at KLM, pilots at Continental can bid on previously determined vacation blocks while pilots at KLM identify their preferred vacation dates. The vacation module now aims to avoid shortages by assigning more vacation blocks outside peak demand time windows.

Subsequently, the two optimization modules in the Crew ResourceSolver model provide both pilot-planning and

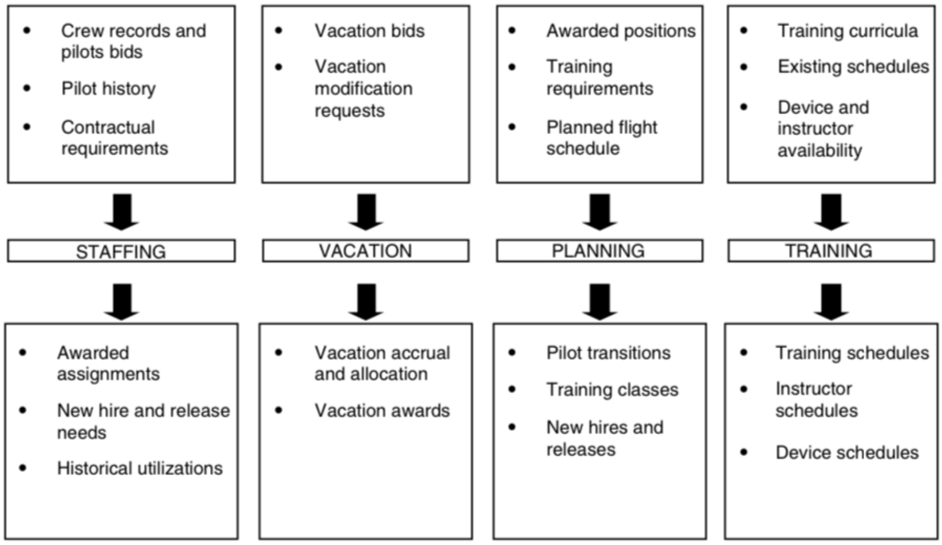

Figure 3.1: The four modules of the Crew ResourceSolver system with their in- and output (Yu et al., 2004).

training functions. The pilot-planning module is tasked with optimizing the transitions of pilots from one seat to another. This is done using a two-step approach in which the problem is first solved as a linear programming (LP) relaxation of the mixed-integer program (MIP). This relaxation is used to determine an estimate of the optimal shortages the airline can achieve. With this estimation, a cost associated with shortages is calculated which can subsequentially be used in the objective function of the MIP. The input of this model consists of current pilot positions, awarded transitions, the flight schedule, the airline costs and any existing constraints. The output is then comprised of the training and transition dates for all pilots with new positions, the timing and number of new hires, the number of training-class starts, and the training-class rosters. Secondly, the training module aims to optimise training and instructor schedules based on training curricula, existing plans and instructor availability using a MIP model. This model determines schedules for training resources (like simulators), students and classes.

Based on the pilot-transitioning model of Yu et al. (2004), Holm (2008) defines a model for the training and vacation allocation problem. Similar to the Crew ResourceSovler system, the model assumes the transitions are determined by the planner in an earlier stadium and the problem at hand is to plan the pilot's transition and vacation in an optimal way. A number of restrictions have to be considered by the model: the demand for pilots is covered at all times, the resources for training are constraint, pilots may have pre-assigned activities and all pilots with an awarded transition have to be trained within the planning horizon. Finally, seniority rules have to be enforced if they are used within the airline. In contrast to Yu et al. (2004), this model takes recurrent training courses into account and vacation is allocated concurrently with the courses as they have a similar effect; a loss of production due to pilot unavailability.

In contrast to the previous models, Thalén (2010) aims to solve the pilot transition allocation problem using a tabu-search heuristic. The aim of this model is to determine what pilots to transition, when, and to what seat, instead of having awarded transitions manually and then optimizing the planning of those transitions (Yu et al., 2004, Holm, 2008). Thalén identifies the possibility of finding superior solutions by planning a pilot different from the most senior pilot eligible for a certain position by either using pay-protection rules or making sure more senior pilots move to a different position in an earlier stadium. The heuristic solution algorithm makes use of two problem specific neighbourhoods are mentioned in several papers (Gendreau, 2003, Lü & Hao, 2008, Lü, Hao, & Glover, 2011), the best way to combine these two neighbourhoods is to sequentially shift between the two until the stopping criterion is met. The first neighboorhood used is that of adding or removing one course from the solution, which can be seen as equivalent to the add or drop neighbourhood regularly used in other problems. The second neighboorhood is moving a course one week forward or backwards. After the neighboorhood search, actions are stored in the tabu list for a predetermined number of iterations to prevent the solution to get stuck in a loop. However, two cases are allowed to breach the tabu rules. The first are moves that create a new best solution, while the second are a set of bad courses that are never allowed to be added if they do not improve the previous best solution.

The solution algorithm was found to yield similar results, but up to a factor 100 faster than a commercial MIP-solver. However, with a problem consisting of 1100 pilots, 13 positions and 3 bases, the heuristic method still took more than 7000 seconds to solve.

Recently, Altenstedt, Thalén, Sjögren, and Nilsson (2017) defined a model where it is assumed the predictions of supply and demand are given and accurate and the problem at hand is to award transitions to pilots in an optimal way. The problem is formulated as a MIP problem that, although too large to solve using commercially available software, can be used as a reference to test the heuristic algorithm used to solve the problem. For this heuristic algorithm, a very large-scale neighboorhood search is used, consisting of different types of neighbourhoods that aim to improve an initial heuristic solution. This initial solution is determined by relaxing the integrality first and subsequentially rounding off to produce a number of transitions and assign these to a crew.

Jarmar and Sörensson (2017) developed a model with the objective to assign cockpit crew promotions. The problem is described as a known stable matching problem, the hospitals/residents problem with ties and forbidden pairs. This is a problem where a set of residents and a set of hospitals have to be matched to each other. Both sets have preferences for the other and the hospitals have a given capacity. The problem is now to find a stable matching where residents are assigned to hospitals in such a way that the capacity conditions and preference lists are respected. When ties are allowed, the preference lists can contain ties in the order of preference. Also, forbidden pairs allow for the definition of pairs of residents and hospitals that might be preferred but are not allowed due to some rules (Manlove, O'Malley, Prosser, & Unsworth, 2007). This method is applied to the crew promotion assignment problem where pilots are residents and courses are hospitals. The courses have a capacity based on the gap between supply and demand. Ties are only allowed in the pilots' preference list, as no seniority ties exist in the system. A drawback of the proposed solution is that it is only able to solve for one point in time. If the method is repeated for multiple instances, no guarantee can be given that a global optimum is found as the solution is determined independently for all times. However, the solution method could provide to be useful as part of a larger integrated solution method as the results presented show to be promising.

3.1.3. Synthesis

The purpose of this section was to present the current practices, found in available literature, regarding tactical and strategic cockpit crew planning.

Not much research has been done regarding the assessment and prediction of supply and demand within the airline industry. As Yu, Pachon, and Thengvall (2003) mention, estimating the demand for pilots is usually done based on either the flight schedule (if available) or the projected fleet size and composition. Next, the estimated value is increased with a certain amount or percentage to account for the absence of pilots. Another part of the demand is for pilot's vacation and training. This demand can be allocated by the airline itself to a certain extent. Both Yu, Pachon, and Thengvall (2003) and Verbeek (1991) identify the possibility of allocating this demand based on fluctuations in the schedule demand using optimization techniques, however, no details about the used methods are given. It is therefore of important to study the possibilities regarding estimating schedule demand, predicting absence demand and optimizing flexible demand in the next sections in order to come up with useful methodologies for the integrated model proposed in the research question.

Most of the research in strategic and (predominantly) tactical cockpit crew planning has focused instead on transition strategies in order to close the gap between supply and demand. It has been found that in a number of cases (Verbeek, 1991, Yu et al., 2004, Thalén, 2010), this problem is simply too large to solve using optimization techniques and therefore, the problem has to be simplified greatly, split up into sub-problems or solved using heuristics. For example, Yu et al. (2004), choose to award transitions using a heuristic algorithm first and then use an optimisation model to plan the moment these transitions start. Also, available literature aims to solve the problem in a deterministic manner, thereby ignoring the stochastic elements involved in transition bids, pilot absence, etcetera. The most used method in literature is to make the problem smaller and thus easier to solve, therefore, it seems the focus of the propsed research project should also be geared towards splitting the complete problem into sub-problem in an effective and smart way and carefully integrating them after solving. Since the problem focuses on long-term planning and therefore does not have to be solved very often, the main focus should not be on computational time but on solution quality first.

3.2. Manpower Planning In Other Industries

The problem of manpower planning is1 not limited to the airline industry. Therefore, this section will discuss manpower planning literature originating from other industries. In Subsection 3.2.1 manpower planning practices for general applications are presented and some separate examples from other industries are given. Subsequently, in Subsection 3.2.2 workforce planning of nurses in hospitals is discussed and in Subsection 3.2.3, taskforce planning in armies is discussed. Finally, in Subsection 3.2.4, a synthesis of the reviewed literature is presented.

3.2.1. General Applications

In this first section, available literature is discussed with regards to manpower planning in industries other than the airline industry. The literature is subdivided into two sections, presenting prediction methods and transition management strategies, respectively.

Predictions For the first steps in manpower planning, Zhu, Sanil, Mardookhy, Sawhney, and Sun (2013) developed two statistical approaches for evaluation workforce requirements in a multi-purpose research organization. The proposed models aim to predict attrition for a large pool of employees with different skills as a resource with a certain demand using Logistical Regression and Time Series. This problem can be compared to the airline's case where different pilots seats are comparable to the skills and factors such as attrition, short-term leave and long-term absence have to be predicted.

The logistical regression method (Zhu et al., 2013) aims to find the probability of individual events (in this case the attrition of an employee) on the basis of a set of relevant predictors. In order to select these predictors, a stepwise selection method and the Akaike information criterion are used in the analysis of their significance. The minimum AIC value shows the relative goodness of fit (Akaike, 1974). The stepwise selection method is a semi-automated process in which variables are added or removed one after the other based on their t-statistics. Finally, the goodness of fit test designed by Hosmer and Lemesbow (1980) is used, p-value range from zero to one and models with a higher value are preferred. Both methods result in a list of relevant predictors, with an associated coefficient, that is used to determine the probability of attrition of an employee. Summing this for all employees with a certain skill in a certain time period yields the loss in the workforce which has to be accounted for.

In contrast to the logistic regression method, time series do not look at predictors of a certain event but solely on the numbers in the past. Some important time series models are the decomposition model, exponential smoothing, ARIMA and dynamic regression model. Using R^2 statistics, the models are assessed on their effectiveness. This statistic shows the amount of variability that is accounted for by using three sums of squares formulas; the total sum of squares, the explained sum of squared and the residual sum of squares. The results for the different time series models show that in this case, the decomposition model performs best with a R^2 value of 0.65.

Ho (2010) propose a single variable, first order grey model forecast for construction manpower one quarter ahead. Grey systems theory was developed by Deng (1982) and deals with problems with small samples or poor information. It searches for patterns based on this limited data. This is also one of the advantages of grey model forecasting: the limited amount of input data required, however, with that comes the inability to accurately capture seasonality in data. This can, however, be solved by deseasonalizing the data before applying the prediction method. The 'memory length' of the model is determined by the sample number, a larger sample number can be better used with very random series and a smaller sample number with smooth time series. The models with different sample number are tested on the minimum mean absolute percentage error criterion, from which it follows that optimal results are found, in this case (with a MAPE of 3.21%), with a sample size of 5.

Even though the only grey model tested in this paper is a GM(1,1) model, the author identifies the possibility that a number of more sophisticated grey models provide better results. Examples of these models are the remnant GM(1,1) model, GM(1,N) model, GM(2,1) model and Verhulst model. It is therefore suggested to develop a computer model that is able to forecast a time series with several different models and from the results choose the optimal for that specific case.

Transition Management For the process of transition planning for closing the gap between supply and demand, several authors have presented an overview of manpower planning models in the past. In his review of models, Price, Martel, and Lewis (1980) define two types. First of all, there are descriptive models that are used to describe how personnel moves through an organization and are usually defined by Markov models, fractional flow models and renewal models. Secondly, there are normative models that aim to determine a set of actions from all feasible solutions. This type of models can make use of, amongst others, linear programming or stochastic programming techniques.

Optimization models are seen as the optimal method for manpower planning in companies where costs are the prevailing criterion, where many conflicting objectives exist or with situations in which many complex constraints must be taken into account (Price et al., 1980). These models define an objective function which should be optimized by changing the decision variables while not violating the constraints (as imposed by labour agreements, logical constraints and more). The model can subsequently be solved by a chosen solution algorithm such as a heuristic method or commercial LP solver.

For the proposed research question, normative models are deemed to be more applicable. First of all, flow between different crew positions in an airline cannot really be defined by Markov or fractional flow models as they are largely determined by vacancies, crew bids and seniority rules instead of general flows between seats. Secondly, the problem at hand is not how personnel moves through the system. Instead, it focuses on the best actions the airline can take to meet demand in the future at minimum cost. For this reason, the following sections will focus more on normative models as opposed to descriptive models. Purkiss (1981) makes the same distinction between normative and exploratory models. It is mentioned that normative models can be more efficient in cases where clear objectives can be defined. Subsequently, these models are often made for specific situations and no generalised programs have been produced. Some of these models use dynamic programming or Markov programming techniques, but by far the most used models are linear programming models.

In another literature survey, Edwards (2007) identified important differences between the labour market in the U.S.A compared to Europe and the UK. As he mentions, the labour market in which a company is situated has a large impact on the way manpower planning is done. Due to more strict employment protection laws, it is more difficult to reduce staffing quickly in Europe when compared to the U.S.A. Also, the influence of trade unions is greater in Europe. For these reasons, his survey focuses on manpower planning for European organizations. From his survey, Edwards sees two main areas for improved planning models: first of all the trade-off between short- and long-term manpower needs and secondly the recruitment process.

In a review of operations research applications in military workforce planning, Wang (2005) identifies four different types of models: Markov chain models, computer simulation models, optimisation models and system dynamics models. For each of these models, the mathematics and concepts together with their advantages and disadvantages are presented. Of these four types of models, only optimisation models are ranked among the normative models while the rest are exploratory models and therefore, as mentioned before, not applicable to the identified research opportunity and the accompanied objective of this research.

These optimisation models as identified by Wang (2005), can be distributed into four different optimisation techniques. Linear programming aims to find the decision variables where the objective function and constraints are linear relations of these decision variables. With integer programming, a linear programming model is defined in which decision variables only take on integer values. Goal programming meanwhile aims to solve linear programming problems with multiple objectives. In this case, deviation variables are introduced which are defined by the deviation of the different objectives to their goal value. These deviations are subsequently minimised. Finally, dynamic programming is a method for solving more complex, multi-stage problems in which the output of a stage serves as the input for the next.

Bard, Morton, and Wang (2007) describe the process of workforce planning at USPS with stochastic demand. The two-stage solution method first determines the amount of full time and part time employees to hire given the stochastic demand and then, in the second stage, the actual demand becomes known and the schedule for all employees is determined and if necessary, overtime is assigned and casual workers are hired. Even though this problem focuses more on the operational planning phase, the fact that it incorporates stochastic demand is interesting for the proposed research.

Taking the inherent uncertainty within a problem into account often slows the model down. Therefore, it is important to think about the way the stochastic elements are incorporated into a model. A very easy way is to solve a deterministic problem in which the stochastic elements are replaced with the expected value, however, this greatly generalizes the problem and any non-linear effects are completely removed from the results (Birge & Louveaux, 2011). For the presented model, Bard et al. identified three methods for taking stochastic elements into account. The first method, the recourse problem (RP), minimizes the sum of the cost in the first stage problem and the expected (mean) cost of the second stage given the stochastic demand. The second method, the expected value problem (EV), instead computes the cost of the second stage using the mean demand and sums this with the cost of the first stage as the objective function. Finally, the third method, the wait-and-see value is obtained by taking the expected value of the full two-stage problem under stochastic demand. The results of the different methods show the RP method produces better results compared to the EV method. Also, it shows the RP method allocates a larger permanent workforce, but this is partly covered by thereby reducing the cost of overtime and casual workers.

Ng, Huang, and Ng (2008) defined a problem where workforce staffing levels have to be determined before employee attendance rate is known. This attendance rate is defined as the percentage of workers available for jobs. Subsequently, when attendance is known, workers have to be assigned to jobs. This problem can easily be translated to the cockpit crew planning problem where the same problems and uncertainties exist. Obviously, understaffing pilots might cause severe costs as flights could be forced to be cancelled if no pilots are available. With the uncertainty in attendance rate taken into account, the problem is to determine staffing levels using a trade-off between fulfilling the demand and the staffing cost. Ng et al. propose six approaches to make this trade-off. The first two approaches take either the average or the theoretical minimum attendance rate and therefore make the problem deterministic. These approaches are very convenient and serve as a baseline solution for the other approaches. The third approach takes the minimum attendance rate per worker type from a sample of attendance rates and uses these values to determine the staffing levels. The 'combined largest approach' solves the initial problem for different historical realizations of attendance rates and then takes the maximum of the staffing levels for each worker type as a decision variable. The fifth approach proposes a two-stage stochastic LP formulation which provides a good solution, however, it does so at a high computational cost. The final approach is based on robust optimization methodology using ellipsoid uncertainty sets (Ben-Tal & Nemirovski, 1999). In this approach, the set of attendance rate realizations is assumed to be confined in an N-dimensional ellipsoid. Once this ellipsoid is completely defined, a number of realizations are taken from the worst-case frontier and used in the two-stage stochastic model of the fifth approach. Computational studies showed this last approach appeared to be the best choice when modelling with the attendance rate uncertainty since it provides a high success rates with reasonable cost savings.

Zhu and Sherali (2009) propose a two-stage solution model for workforce planning that takes demand fluctuations and uncertainty into account. The solution method first aims to make personnel recruitment and allocation decisions based on the stochastic demand and then in the second stage determine centre-, shift- and month-changes for tasks in order to cover the demand for all individual units. This is done as recruiting and allocation decisions generally have to be made further in advance while shifts of tasks can be postponed to a later stage when more information about the actual demand is known.

The first step in the two-stage approach aims to determine, based on stochastic demand, the optimal number of recruited employees as well as their allocation in the company (as the triplet category k, planning unit u, month t). With these recruitment and allocation decisions made, the next stage is to determine, for all demand scenarios s, any required changes in location, shift and/or month for tasks in order to optimally use the determined workforce on all assignments. This second step is done with the Benders' decomposition method, in which additional constraints are imposed on the first stage if the second stage was found to be infeasible for any of the scenarios.

The results of the two-stage stochastic approach, when compared to a deterministic solution method, show that for all problem sizes, the two-stage method decreases the amount of shortage to zero and, especially for problems with a greater number of scenarios, decreases the number of changes and splits of demand over different centres.

Two heuristic solution algorithms have been proposed by Fozveh, Salehi, and Mogharehabed (2016) for a multi-skilled, multi-objective workforce planning problem. The objective of the problem is threefold: minimise the number of night-shift workers, minimise the total workforce cost and maximise the number of engaged employees.

The first solution heuristic is a bee colony optimization in which at first a set of feasible solutions is created. This initial population of solutions is then split into two groups based on their performance. A neighboorhood search is performed on the best scoring group while a random search is performed on the worst scoring group. This yields a new population of results. The described process is then repeated until a stopping criterion is met. The second solution algorithm is a differential evolution algorithm. This algorithm works by first creating a set of feasible solutions using a multi-start variable neighboorhood search after which, all solutions are altered by creating a mutant vector, applying a crossover operator combine the two vectors and apply a repair method two solve any infeasibility.

Subsequently, the best solutions are stored and the method is repeated until a predetermined number of iterations have been performed. The best solutions are in this case selected based on a Pareto optimisation of the different objectives.

Catalá, Moreno, Blanco, and Bandoni (2016) have formulated a multi-period mixed integer programming model for medium-term planning (which is defined as a one-year horizon, divided over 12 monthly periods) in the pome-fruit industry. A lexicographic approach is used in order to handle the multi-objective nature of the problem in which the planning of shifts of workers has to be determined in order to maximise the profit from growing, harvesting and selling fruit. This approach orders the relative importance of the different objectives, solves the most important first and then uses the results of this first optimisation as a constraint for the second objective, repeating the process until all objectives are evaluated and/or a single optimum is found (Sawik, 2007). In the model presented by Sawik, the most important objective is to minimise the dissatisfaction of clients, followed by maximising the total profit of the company. This approach could be used in cockpit crew planning problems with multiple objectives such as service level (defined by, for example, the gap between crew supply and demand), robustness, cost, number of transitions etcetera.

3.2.2. Nurse Planning

The nurse planning problem is in many ways similar to the cockpit crew planning problem. Similar to pilots, nurses often have a specific qualification or specialisation which determines what tasks they can perform. Next to this, the nurse planning problem is solved in a similar way as cockpit crew planning. Punnakitikashem, Rosenberger, Behan, Baker, and Goss (2006) define four subsequent stages: nurse budgeting, nurse scheduling (which can be subdivided between staffing and rostering for long- and short-term, respectively), nurse rescheduling and nurse assignment. As for the defined research gap, the stage of interest is the nurse scheduling phase and more specifically nurse staffing.

In contrast to Punnakitikashem et al. (2006), the nurse planning problem is divided into three pases by Abernathy, Baloff, Hershey, and Wandel (1973). The first phase handles "policy decisions, including the operational procedures for service centres and for the staff-control process itself". The second phase determines "staff planning, including hiring, discharge, training and reallocation decisions", and finally, the third phase is for "short-term scheduling of available staff within the constraints determined by the two previous phases". Abernathy et al. identified the same problem as with cockpit crew planning research; most research has focused on short-term scheduling problems and improvements in solutions for this.

Several solution methods have been proposed by Abernathy et al.. The first method uses an iterative approach using a random loss function. This approach iteratively solves the staff-level problem and staff-allocation problem, it then updates the staffing cost function in the staff-level problem with the result from the allocation problem until the solution converges. The second solution method incorporates a chance constraint which states that an effective staffing level has to be achieved such that the probability of a shortage is smaller than a certain threshold. No recommendations are made in the paper with regards to the optimality of any of the two solution methods. However, as the first requires successively solving both the planning as well as the scheduling problem, it potentially becomes too large to be able to solve within a reasonable time.

Li et al. (2007) have developed an integrated staffing decision model that is divided into three stages: demand forecasting, planning and scheduling (as can be seen in Figure 3.2). The first stage is used for forecasting demand. In the nurse planning problem, the demand can be divided into two separate groups: known and unknown demand. Known demand entails planned appointments while unknown demand are walk-ins without appointment and emergencies. The output of this first stage are forecasts of the total demand on a weekly or monthly basis. The second stage then uses these forecasts to determine the staffing levels for each skill class for the entire planning time window. This stage takes as input the mentioned demand forecasts as well as planning requirements and regulations. The number of employees in each class, number of recruits and attritions serve as the output and are subsequently used in the third stage in which detailed schedules are constructed for all staff members generated by the planning model. This scheduling phase is constraint by several labour agreements.

The staff planning model (Li et al., 2007) will define the amount of staff to be hired using a multi-objective problem. The objectives of the staffing problem are defined with regards to the staffing cost, staff augmentation, staff task substitution, overtime and shortfall of professional development. The staffing problem is solved by converting the multi-objective linear programming model to a single-objective problem by using the Analytic Hierarchy Process

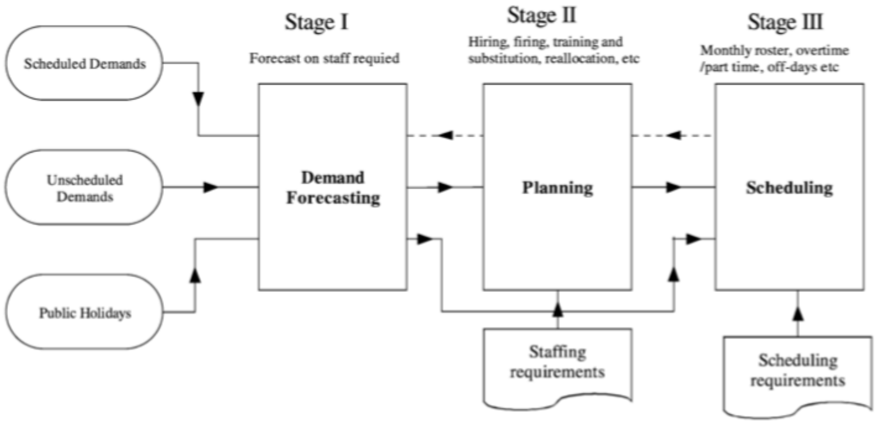

Figure 3.2: Block diagram of the three-stage integrated staffing model (Li et al., 2007).

(AHP) (Winston & Goldberg, 2004). In this method, the relative importance of all objectives compared to the others is put in a matrix and normalized, yielding a weight for each objective that can be used in the single-objective problem. Next, with the staff scheduling model defined as well, the full problem is solved in an iterative manner in which the planning model is solved first. From the results, it is analysed if the amount of overtime, number of recruits, number of task substitutions or the unachieved professional development are unacceptable, and if so, the weight of the relevant objective is increased. If all parameters are acceptable, the scheduling problem is solved for all time periods. If the amount of overtime or temporary staff is too large, the associated weight of these objectives are increased and the staffing problem is solved again, starting over with the algorithm. If everything is acceptable, the algorithm is finished.

The integrated staffing and scheduling model proposed by Maenhout and Vanhoucke (2013) aims to integrate two parts of the total problem using a Dantzig–Wolfe decomposition method in which the staffing problem is solved as the master problem and the scheduling of individual nurses as the subproblems. The staffing master problem then determines the required number and mix of nurses in each ward and recruitment plan for all wards. The paper shows the benefits of integrating the two phases in nurse planning. It is shown that by solving the problem for multiple departments simultaneously improves the schedule in terms of cost, job satisfaction and effectiveness. Also, the model can be used as a simulation tool in order to assess the effect of different staffing and scheduling strategies. From the presented results, it can be concluded that efficiency, effectiveness and job satisfaction can be greatly influenced by only minor differences in staffing and scheduling strategies. Because of the analogies with cockpit crew planning, it can be assumed that this is also the case for the identified research gap.

3.2.3. Army Planning

Similar to the problem with nurses, the army features several different positions and therefore decision makers face the problem of planning the strength of these different functions on a strategic level given a (regularly uncertain) demand.

Gass, Collins, Meinhardt, Lemon, and Gillette (1988) developed the Army Manpower Long-Range Planning System (MLRPS) for analysing the army's strength over a 20-year horizon. The MLRPS is a decision support tool providing the answers to two sub-problems. First of all, a projection of the strength of the army by grade, skill, years of service, and quality is made. This enables decision makers to evaluate the effect of different strategies over the 20-year time window and compare the projected strength to the desired one. Secondly, optimal transition probabilities are determined which enables decision makers to determine the desired strategies to ultimately reach desired workforce parameters defined by skill, grade, years of service and quality. The program utilises a Markov chain for analysing the flows between different states and can therefore be seen as an exploratory model. It furthermore assumes fixed transition probabilities and long-term behaviour.

In contrast to the exploratory model defined by Gass et al. (1988), Silverman, Steuer, and Whisman (1988) developed a multi-period, multi-criteria optimization model for the U.S Navy manpower planning problem. The model serves as a decision support system (DSS) and the goal is to come up with a recruitment and promotion strategy that best conforms to the seven defined goal trajectories. These seven trajectories are the salary cost, the strength of force of paygrade two and three, promotions to paygrade two and three and mean level-of-service for paygrade two and three. The objective function is then to minimize the maximum over-deviation in terms of salary and the minimum of the other six over deviations. An iterative approach is used called the augmented weighted Tchebycheff procedure. In this procedure, a sample of solutions is tested in an increasingly smaller subset of the deviation values. Based on the results of each iteration, the decision maker is able to alter the model and its goal trajectories based on any insights gained from the results, after which the next iteration is started. This method was chosen for several reasons. First of all, it does not require the user to specify any weights for the different objectives. Secondly, the Tchebycheff procedure does not produce large jumps in the solution following only minor changes in the weights, which provides the user with more control of the problem. Finally, the procedure provides multiple solutions, yielding the user the possibility to choose the best based on their expertise.

Škraba, Kljajić, Papler, Kofjač, and Obed (2011) have described the development and usage of a workforce planning model for the Slovenian Army. The model utilises a system approach, system dynamics and numerical optimization. As with all armed forces, the manpower system is hierarchical, similar to the cockpit crew system, however, in the Slovenian case, it is strictly hierarchical. This means higher rank member must be provided from the single next subordinate rank. The goal of the model is then to track the goal trajectories in the eight highest ranks in the system. Using a quadratic performance index for the deviations from the goal trajectories based on the three relevant parameters; recruitment, promotions and fluctuations, acceptable strategies were identified. Next, using Genetic algorithms and pattern search, these strategies were improved. From the two algorithms, it was found that the pattern search showed to be significantly more suitable to determine the optimal strategy. In later research, Škraba, Stanovov, Semenkin, and Kofjač (2016) describes the use of stochastic local search and genetic algorithms for manpower planning with the same application to the Slovenian Army. In this paper, a heuristic stochastic search algorithm is used to achieve the desired number of crew in each rank at each time by defining recruitments, promotions and fluctuations. This algorithm randomly chooses an action in each of the possible situations. For instance, if the number of crew in a certain rank is too low, possible actions are to decrease promotion from that rank to the next, increase promotion from the previous rank or decrease the outflow from that rank. This algorithm was able to find a feasible solution, however, the solution depended greatly on the initial values. Therefore, a genetic algorithm was proposed as an optimization technique where the optimised values were used as the initial values for promotion and fluctuation coefficients. The results of the model applied to the Slovenian Army show that for this specific situation, only a short modelling time is required to reach good solutions to the presented problem, however, this does require a proper modelling strategy because of the dynamic nature of the problem, creating a large search space.

3.2.4. Synthesis

As the literature regarding cockpit crew planning at airlines was quite limited and the problem has quite some similarities with manpower planning problem in other industries, the purpose of this section was to identify solution methods for the identified research gap in other industries.

Predictions have to be performed in all industries when planning for the future. The results of different prediction methods are often highly dependent on the specific application and system. Therefore, any methods found in this section should be compared and evaluated on the actual process in the cockpit crew problem. Identified methods include Logistical Regression and Time Series models for predicting employee attrition (Zhu et al., 2013) or Grey model forecasting for construction manpower (Ho, 2010). An important factor to account when using a specific forecasting technique is the ability of the method to account for seasonalities and trends and if the method is unable to factor this is, measures should be taken to alter the technique accordingly.

For transition management, two distinct types of models have been identified (Price et al., 1980, Purkiss, 1981, Wang, 2005): exploratory or descriptive models that describe how the system behaves and evolves and normative models that aim to determine actions to make the system evolve in a desired way. For the identified research gap, normative models are deemed more applicable. Some research has also tried to incorporate stochastic elements in the model in some way. Bard et al. (2007) does so for stochastic demand at USPS by using a two-stage approach that first determines the optimal amount of employees given the uncertain demand and then schedules the employees and if necessary assigns overtime and casual workers based on the actual demand. Again, a number of heuristic solution methods have been defined (Fozveh et al., 2016, Škraba et al., 2011), but another method often used is goal programming, in which the deviation from a number of goal values or trajectories is minimised (Škraba et al., 2016, Gass et al., 1988, Li et al., 2007). Because of the large computational time for optimisation models for the indentified research gap, a heauristic method will likely provide adequate results while still keeping the computational time acceptable.

3.3. Promising Methods from Related Problems

After moving the focus to manpower problems in other industries, a next step is to look at promising methods used in other, related, problems. First, fleet planning problems and the similarity with cockpit crew planning will be discussed in Subsection 3.3.1. Next, in Subsection 3.3.2, several useful prediction methods will be discussed en after that, some solution algorithms will be discussed in Subsection 3.3.3. Again, the section is finished with a short synthesis in Subsection 3.3.4.

3.3.1. Fleet Planning

Because of the similarities between fleet and crew planning, this section discusses methods used in fleet planning literature that could be used in the field of cockpit crew planning as well. Rosskopf et al. (2014) define the fleet planning problem for an individual airline as the problem to determine the optimal fleet composition, in terms of the amount of aircraft and type, for each time period within the planning horizon. Factors under consideration in this process include the usage of aircraft within the airline's operations, evolution of the fleet over time and fleet financing. One objective of fleet planning that is not present for the cockpit crew planning problem, however, are environmental goals. For the fleet and choice of aircraft, current and future environmental goals are an important driver. For this, Rosskopf et al. presented a methodology for balancing these environmental challenges with economic goals. It shows that reducing the fleet's NO_x emissions comes at a relatively high cost (around double the percentage of the emissions reduction).

Similar to the field of crew planning, a lot of advancements in modelling and optimisation approaches have arisen in the last years. Earlier, the problem was often solved with general linear programming methods such as mixed-integer programming (Marsten & Muller, 1980) or a multi-commodity flow problem (Listes & Dekker, 2005). Obviously, the fleet planning problem has to deal with some of the same uncertainties as for the crew planning problem, such as demand and availability uncertainty. To account for this, Listes and Dekker defined a scenario aggregation-based approach for dealing with the stochastic demand in the problem. In this approach, the problem is iteratively solved for different scenarios of the stochastic elements and the results are aggregated into an overall usable solution. Next, a decision solution is obtained by taking the weighted average of all scenario solutions. This solution is called the implementable solution, however, it may not be admissible (which means it is feasible for all scenarios). In search for an optimal solution that is both implementable and admissible, a sequence of converging estimates is made of the decision solution (Wets, 1989). Another important factor in the scenario aggregation-based method is the generation of scenarios. Careful generation ensures valid data as it is not possible to assess all possible scenarios. In this case, the descriptive sampling method is used (Saliby, 1990). Instead of picking stochastic realizations of a certain process using a random set and random sequence, the set of numbers is made deterministic while keeping the random sequence. The deterministic set is sampled from the known distribution at constantly spaced quantiles, ensuring more samples are drawn from high probability regions, but extreme values are also drawn.

Another solution strategy that focuses on uncertainty in the process is proposed by List et al. (2003) where a trade-off is the investment of a certain fleet and the risk associated with it. The uncertainty incorporated in the model includes uncertainty in future demand and vehicle productivity. Instead of focussing on the average performance of a solution, the model focusses on the probability that the total cost for a given solution exceeds a certain threshold. This probability is defined as the risk of the solution. By now assigning a weight κ to the amount of over-deviation of the threshold by a solution, the trade-off between risk and reward can be altered.

The approximate dynamic programming approach proposed by Requeno García (2017) is designed to determine adaptive fleet plans in order to better cope with demand uncertainty. This adaptive fleet plan defines, for each year in the planning horizon, an optimal strategy. However, for all years after the next, this plan is adaptive since it depends on the demand realization (most optimistic, most likely and most pessimistic) of the preceding year. In this way, a fleet policy decision tree is generated showing the optimal strategy based on the demand of previous years. The approach was tested on real-world data from Kenya Airways, which showed the applicability of the method, as well as the capability to estimate future operational profits. The approach's profitability was also larger than the current airline's fleet plan, showing the model's usefulness in coping with demand uncertainty using adaptive policies.

3.3.2. Predictions

Although prediction methods have already been discussed in the previous two sections regarding cockpit crew planning and general applications. some interesting approaches towards predictions will be given here to further develop the available knowledge regarding accurate predictions, especially regarding time series featuring trends and seasonality.

Tseng, Yu, and Tzeng (2001) developed a hybrid grey model used for forecasting time series. Usually, grey forecast models are not able to account for seasonality in its predictions. Tseng et al. instead propose a hybrid grey model with the ratio-to-moving-average deseasonalization method. This method calculates a seasonality index per period using the following approach:

1. Compute a k period moving average, where k is the period of seasonality.
2. Compute a seasonal index by dividing the actual value by the k period moving average.
3. Determine, for each period in k, the average seasonal index.
4. Divide individual values by their corresponding seasonal index to obtain the deseasonalized time series.

This method for deseasonalizing historic data can be used for any forecasting method. In this case, however, it was used together with a GM(1,1) model, which is defined by the first order differential equation:

$$\frac{dx^{(1)}(k)}{dk} + ax^{(1)}(k) = b, \tag{3.1}$$

where $x^{(1)}(k)$ is defined as the cumulative sum of the time series $x^{(0)}(k)$ and a and b the developing coefficient and the grey input, respectively. The solution can then be obtained using the least-squares method. In their computational results it was shown that the GM(1,1) model with deseasonalized data outperformed the in-sample forecast of the SARIMA-, GM(1,N)- and neural network model.

For the other special kind of time series, those with trends, Qi and Zhang (2008) have developed forecasting network using a neural network. In the paper, a Monte Carlo study is performed to determine what the optimal method is to forecast trend time series. The neural network used in this study is a standard three-layer feedforward network. Five different data generating processes where used and four modelling strategies were analysed: modelling with raw data, raw data with a time index, modelling with linearly detrended data and with differenced data. It was found that the differenced forecasting method is the most effective. Although neural networks are a powerful method for forecasting all sorts of time series, a large disadvantage is the lack of insight in the prediction process because of the network's hidden layers which gives the user less control of the solution as well as less insight in the process.

3.3.3. Solution Algorithms

As previous research has shown, the complete cockpit crew planning problem is often too large to solve using optimisation methods. Therefore, a lot of research focuses on heuristic methods to come up with a feasible, yet suboptimal, solution. In order to improve these heuristic solutions, several search algorithms exist that could potentially improve the solution significantly in a short time.

In his research with respect to multi-variable optimization problems, Abramson (2002) reviews a number of search heuristics and their applicability to his problem. The first heuristic, simulated annealing was originally devised by Metropolis, Rosenbluth, Rosenbluth, Teller, and Teller (1953). The method is based on the cooling of a substance to a desired temperature while keeping a certain thermodynamic equilibrium. This system was translated to optimization problems in which the hypothetical temperature is decreased (according to the user-defined cooling schedule) if the evaluated neighbour of the current solution is not better (smaller in a minimisation problem) than the current objective function. If it is better, the temperature is held constant. This process is continued until the minimum temperature is reached.

The next method, already briefly described in Subsection 3.1.2, is the tabu search. In this procedure (Abramson, 2002), it is tried to avoid local optima by accepting a worse solution if no better solutions are found amongst the neighbours and the solution is not found in the tabu list containing recently visited points. The efficiency and effectiveness of the tabu search rely on three parameters: aspiration, diversification, and intensification that define the strategies of overriding the tabu list with good points, global, and local searching, respectively.

Finally, Abramson discusses various evolutionary algorithms: evolution strategies, evolutionary programming and genetic algorithms. Only genetic algorithms are designed for discrete problems, while the others are designed for continuous problems. A genetic algorithm consists of four steps, named after biological processes:

- **Selection:** Procedure for the appropriate parents for reproduction.
- **Reproduction:** Of the selected parents in order to generate offspring
- **Mutation:** Introduction of random errors in the offspring.
- **Competition:** Evaluation of the performance of the offspring.

In contrast to the direct search methods presented above, pattern search methods evaluate the neighbours of a solution in a certain pattern. Torczon (1997) presented the method for solving non-linear, unconstrained problems. One of the pattern search methods, Generalised Pattern Search (GPS) starts by computing the objective function at an initial solution. It then determines the candidate points based on the defined pattern and picks either the first improving or best candidate point based on the user's preferences.

As with many computational problems, search heuristics can be improved using machine learning techniques. One example of such machine learning techniques often used for search algorithms with a finite set of actions is the so-called Learning Automata (LA). Thathachar and Sastry (2002) describe a number of different LA methods. They, however, all rely on the technique of improving the probability distribution over the action set based on the previously determined results of different actions. While learning, the algorithm increases the change of performing actions that are known to improve the objective function. Using techniques like LA could potentially decrease the time it takes to find a sufficiently optimal solution for the transition planning problem.

3.3.4. Synthesis

This section reviewed a number of methods and techniques found in various applications that are applicable to the cockpit crew planning problem. One problem that has quite some analogies with cockpit crew planning is (airline) fleet planning. In this field of research, several approaches to deal with uncertainties were found. The scenario aggregation-based approach (Listes & Dekker, 2005) is promising in determining an optimal solution given uncertain demand data and also discusses scenario generation methods when using a limited number of scenarios. One downside of their model, however, is the large number of computations and therefore computational time. Other research proposed to make a trade-off between risk (defined by the probability to go over budget) and reward (List et al., 2003) or an adaptive plan that presents an adaptive plan per year based on the demand of previous years (Requeno García, 2017).

An important part of accurate planning is the ability to make accurate predictions of the future. As literature with regards to predictions for crew or manpower planning has been found to be quite limited, some additional prediction methods have been reviewed in this chapter. For time series with seasonality, an approach by (Tseng et al., 2001) showed good results by deseasonalizing the time series using the ratio-to-moving-average method and forecasting that data with a GM(1,1) grey model. Qi and Zhang (2008) on the other hand proposed a method for forecasting trend time series using a neural network. Although the use of a neural network is very powerful and promising, it forms a sort of black box which does not give the user an idea of the source of the predictions.

Finally, a number of solution algorithms for the transition planning problem have been reviewed as current literature shows the problem is often too large to be solved using optimisation software. In these cases, a carefully designed heuristic algorithm could be used to speed up the solution and obtain sufficient solution quality. Examples of search heuristics that can be used in the crew planning problem are simulated annealing, tabu search, evolutionary algorithms and pattern search. The performance and convergence of these search algorithms heavily depend on the application and algorithm parameters and should, therefore, be tested on the actual problem to determine the optimal algorithm. With many heuristic algorithms, the solve time can be drastically decreased by designing machine learning techniques to aid the algorithm in searching in favourable areas. One of these machine learning techniques is Learning Automata which increase the probability of choosing actions that have shown to improve the objective function.

3.4. Literature Synthesis

The goal of this chapter was to present the available literature regarding the long-term cockpit crew planning problem. This problem was divided into three subproblems according to the definition by Holm (2008). The first subproblem aims to allocate flexible demand for training, vacation and more based on the nett demand distribution, thereby minimising shortages. The second subproblem is to make sufficiently accurate predictions for the absence demand. The third and largest subproblem then aims to close the still remaining gaps between supply and demand by transitioning pilots between different positions and by recruiting new pilots. In this section, the most important literature regarding these three subproblems is summarised. Next, in Chapter 4, this is used to define a project scope for the research.

The first subproblem in the cockpit crew planning process is the allocation of flexible demand for training, vacation and office days based on the gap between supply and demand for all seats and time periods. This allocation can be done using a rather straightforward linear programming model where the constraints define the contractual, governmental and labour agreements and the objective is to minimise the cockpit crew shortages.

The next sub-problem is tasked with making predictions for the absence demand, i.e. the amount or percentage pilots for each seat and time period that is absent due to unforeseen circumstances such as illness, pregnancy and injuries. Within the field of cockpit crew planning, not much literature was found regarding these predictions. Yu, Pachon, and Thengvall (2003) briefly mentions that a certain percentage is used to account for this absence demand, but provides no details about the predictions and their accuracy. Therefore, the scope of the search was broadened to manpower planning problems in different industries. Two types of models with various different applications have been identified by (Zhu et al., 2013): Logistical Regression in which relevant predictors are identified and used in the predictions and Time Series in which a number of techniques can be used to extrapolate the data. Grey model forecasting in a technique proposed by a number of authors (Ho, 2010; Tseng et al., 2001). An advantage of grey model forecasting is that it relies on a limited number of data points to make an accurate prediction. Depending on the data and prediction method, the data might have to be detrended or deseasonalized. This can be done using differenced data forecasting or ratio-to-moving-average methods, respectively. However, it should first be checked whether the data shows trends and seasonality. A significant number of prediction methods have been reviewed in this study, however, their usefulness and accuracy greatly depend on the specific application and data, therefore, these methods should be modelled and assessed on the actual data before choosing the most appropriate.

The final sub-problem is planning transitions for pilots between different seats as well as new recruits in order to close the gap between supply and demand. It was found that even without stochastic elements, this problem is too large to be solved using commercial optimisation software. Add to this the stochastic absence demand of pilots and the uncertain transitions bids and a problem is created that is too large to solve with any optimisation model. Available literature, therefore, focusses on only a small part of this problem and takes the rest of it as predetermined constant or aims to design efficient heuristic algorithms to obtain a solution of sufficient quality. Verbeek (1991) decided to use a heuristic approach but did not provide any details about the heuristics, Yu et al. (2004) uses a heuristic to award transitions and then optimises the plan the time of the transitions and two heuristic algorithms (bee-colony optimization and DE algorithm) are used by Fozveh et al. (2016).

As the problem at hand is large, it is also quite difficult to determine a single objective. The total cost is ultimately to be minimised but then factors such as crew satisfaction, transition capacity, overstaffing and understaffing surcharges are simply excluded. When working with a multi-objective optimisation problem, a careful combination of the different goals should be done in order to obtain optimal solutions. One method to combine the different goals into a single objective is goal programming (Wang, 2005), in which the deviation of goal trajectories with respect to the different objectives is minimised and Pareto optimisation (Fozveh et al., 2016) in which the optimal solutions are found on a frontier from which it is impossible to improve the objective in one field without aggravating the objective in other fields. On the other hand, Catalá et al. (2016) used a lexicographic approach in which the most important objective is solved first and then used as a constraint for the following objective and the Analytic Hierarchy Process (AHP) is used in (Li et al., 2007). This process identifies the relative weights of the different objective that can then be used to combine them into one objective function.

Using a simple heuristic algorithm, it is possible to come up with a feasible initial solution. This solution can then be improved by using efficient search algorithms combine with machine learning techniques in order to converge to an optimal solution. Examples of search algorithms that were found in the literature are simulated annealing, tabu search, evolutionary algorithms and pattern search. The performance and convergence of these search algorithms heavily depend on the application and algorithm parameters and should, therefore, be tested on the actual problem to determine the optimal algorithm. Machine learning techniques can then be used to further improve the search algorithm by steering the algorithm to search in favourable areas. One of these machine learning techniques is Learning Automata which increase the probability of choosing actions that have shown to improve the objective function in the past.

Finally, the stochastic elements of the problem have to be integrated into the solution approach, preferably without exponentially increasing the computational time. On two extremes of the amount of integration of these stochastic elements into a model, one can choose to simply use average or expected values of the stochastic parameters or choose to simulate the entire model for a number of realisations of the stochastic parameters. The major downside of the first method is that it loses any dependencies between different parameters, non-linear relations and extreme cases. The second method, however, increases the computational time too much to be practically applicable. An interesting method, stemming from the fleet planning problem is to construct a limited number of scenarios using the descriptive sampling method in order to create an accurate representation of the stochastic variables with only limited iterations (Saliby, 1990). From the results of the different scenarios, it is then important to analyse the data in an efficient and useful way, this has been done with the scenario aggregation method (Listes & Dekker, 2005), recourse problem approach (Bard et al., 2007) or a trade-off between risk and reward (List et al., 2003).

Following the definition of the research gap and available literature, the design of the research will be presented in Chapter 4. As the time for the project is limited, the focus of the research will be demarcated and the context of the problem will be given.

4
Research Design

The next chapter in this report introduces the design of the research as it has been carried out. In Section 4.1, the context of the research within the complete cockpit crew planning process is discussed. In Section 4.2, a simplified example of the cockpit crew transition planning problem is presented to better introduce the problem. Next, in Section 4.3, the structure and content of the models that have been used to solve the cockpit crew transition planning problem are presented and finally, in Section 4.4, the main assumptions made within the research are presented and discussed briefly.

4.1. Context

In this section, the context of the cockpit crew transition planning problem within the crew planning process is presented. The research will focus on this subproblem of the long-term cockpit crew planning problem as this is the largest and most complex problem of the three subproblems. As literature regarding this subproblem is limited and the problem is very complex, a local search method will be applied to solve the problem.

In Figure 4.1, a high-level overview of the full crew planning process is given. The input for the long-term crew planning problem is a flight schedule or fleet plan coming from the network department. Using this input, the necessary amount of pilots at each crew position (which is a combination of rank and aircraft type) and point in time can be determined using the relevant employment rules of the airline. Also, the supply of pilots has to be determined from the current workforce. In case the demand and supply of pilots at different positions do not match, measures have to be taken to resolve these shortages or surplusses. This can be done in various ways, but the most used method is to plan transitions for pilots from one function to another. The output of the crew planning problem is then a set of pilots that matches the planned demand as closely as possible, at minimal cost. These pilots can then be scheduled to the flights in the flight schedule and the flexible demand. This problem is solved in the crew scheduling problem. Afterwards, any disruptions to the schedule that arise close to or on the day of execution are resolved in the crew recovery process.

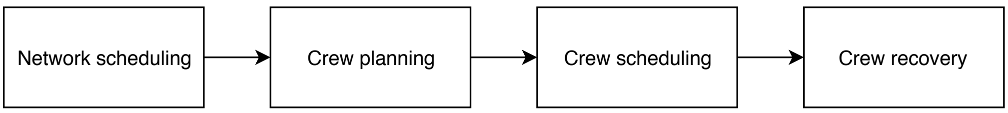

Figure 4.1: High-level overview of the airline crew planning process.

In his research into reserve crew pairings, Janssen (2018) defined the various parts in the crew planning process together with the problems in the different parts in more detail, as can be seen in Figure 4.2. The problems within the crew planning domain are defined to be: manpower sizing per crew position, transition planning, training planning and margins (or premises) for illness, operational disturbances and more.

For the transition planning problem, the input is a demand (both fixed demand and demand relative to the crew strength in a position) per position and time, and a current pilot workforce; the supply. These have been determined in the other three parts of the crew planning problem. From the input, the transition planning problem can determine the balance per crew position and date. This balance is created by subtracting the demand for pilots from the supply and shows any shortages and surplusses in crew. Now, the goal is to plan transitions and hire recruits in order to minimise both shortages and surplusses at minimal cost. These transitions minimise shortages by promoting a pilot from his current crew position to a new position.

As the transition planning problem is a complex problem, current solution methods using commercial software often take hours to even days to solve the problem, which makes it difficult to analyse the effects of different choices,

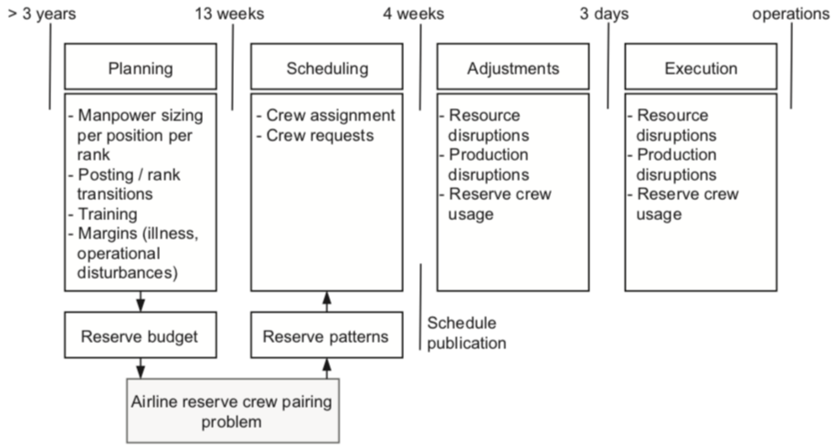

Figure 4.2: Detailed overview of the airline crew planning process (Janssen, 2018).

scenarios and strategies. The proposed research, therefore, aims to develop a model that is able to solve the transition planning problem in a limited time. This gives airlines the option to analyse the impact of different choices on the crew plan.

4.2. Simplified Problem

In order to create a better understanding of the transition planning problem, as small and simplified problem will be presented and solved in this section. In Figure 4.3, the promotional hierarchy for this simplified problem can be found. The simplified system features two aircraft types: one flying in Europe only (EUR) and one flying intercontinental (ICA). As can be seen, the problem consists of 5 different crew positions with for each position (except for captain ICA) two possible positions to transition to. There is one direct-entry position (second officer ICA), which means this crew position is increased by hiring recruit from outside the airline. For all possible transitions (depicted by arrows), the binding period is pictured next to the arrow. This value shows how long a pilot has to be in his current function to be able to be awarded the transition to the other position. Next to the binding period defined between two functions, binding from the date the pilot started at the airline are in place for various positions. For transitions to first officer (FO) ICA, the pilot has to be in service for 5 years, for transitions to captain (CP) EUR, this period is equal to 6 years, and for transitions to CP ICA even 9 years of service are required.

The first step in making a crew plan is to construct the current balance. This process starts with obtaining the nett demand (i.e. the crew demand for flying the scheduled flights) as can be seen in Figure 4.4. This nett demand is then increased with demand for a number of factors such as crew vacation, training requirements and illness to come up with the gross demand (Figure 4.5). The various factors contained in the gross demand are explained in more detail in Section 5.2.

On the other side of the balance is the supply of pilots to the available crew positions. This supply can be calculated from the number of pilots in a position multiplied with their full-time equivalence percentage (FTE). This results in a crew strength per position and time as can be seen in Figure 4.6. By now subtracting the gross demand from the supply of pilots, the balance per crew position and time is obtained, as shown in Figure 4.7.

As the current balance features a number of shortages in the various positions, measures have to be taken to resolve these. One of such methods is to award transitions to pilots from one position to another. However, these transitions are subject to a number of rules that complicate the problem greatly.

The first step, in this case, is to resolve the shortage at the captain function on the EUR fleet for September 2019. The airline is not able to choose which pilot will be awarded a transition to the EUR fleet at that moment directly. Instead, the most senior pilot with a valid bid will be awarded the transition automatically. In this case, the most senior pilot is currently a first officer on the ICA aircraft type. He is awarded a transition to CP EUR starting in August 2019 as in the first month after a transition, a pilot is being trained for the new function and is unable to contribute to the supply. After the first month, the pilot will contribute to the supply, but because this pilot also brings in extra demand for holidays, training and more, the demand will go up as well. The balance for CP EUR,

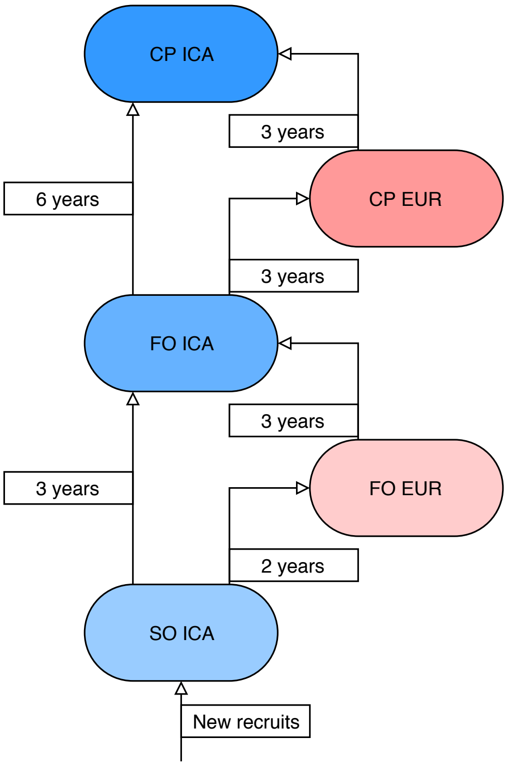

Figure 4.3: Schematic of the hierarchic promotion system for the simplified problem.

	Jul 2019	Aug 2019	Sep 2019	Oct 2019	Nov 2019	Dec 2019
CP ICA	421	432	432	431	421	421
CP EUR	395	394	404	402	392	392
FO EUR	412	414	424	423	415	415
FO ICA	419	422	422	421	411	414
SO ICA	325	333	332	334	348	346

Figure 4.4: Nett demand for the simplified example problem (in FTE).

	Jul 2019	Aug 2019	Sep 2019	Oct 2019	Nov 2019	Dec 2019
CP ICA	616	614	609	604	605	608
CP EUR	615	608	618	619	617	615
FO EUR	571	571	586	578	573	574
FO ICA	576	577	578	573	572	576
SO ICA	393	404	412	416	426	424

Figure 4.5: Gross demand for the simplified example problem (in FTE).

	Jul 2019	Aug 2019	Sep 2019	Oct 2019	Nov 2019	Dec 2019
CP ICA	617	616	616	613	609	604
CP EUR	623	619	614	614	611	611
FO EUR	584	578	574	574	574	574
FO ICA	587	586	585	584	584	584
SO ICA	405	405	405	405	405	405

Figure 4.6: Crew strength for the simplified example problem (in FTE).

therefore, rises 0.7 FTE and the balance for FO ICA decreases by 0.8 FTE as can be seen in Figure 4.8. This process completes the first iteration of the model, in which an iteration is defined as one step in the construction algorithm in which a transition or recruitment is planned. In this case, eight transitions (all from FO ICA, as this is the only position pilots can come from in the used system) are required to resolve the shortage of 4.1 FTE in September 2019. As can be seen in Figure 4.9, these transitions have also decreased the shortages on the CP ICA later in the planning window, but two shortages remain in the captain position on the EUR aircraft type.

With the first shortage resolved, a new option should be chosen. These options are shortages that have to be resolved. A rather large shortage can be found for FO EUR in September 2019. In this case, the only position a pilot can be transitioned from is the second officer ICA position. As can be seen in Table 4.1, the next iteration is indeed to plan a transition from SO ICA to FO EUR starting in August 2019. This, however, creates a shortage for second officers ICA in August 2019. Therefore, the next step is to hire a recruit to this position. In Figure 4.10, the balance after 21 transitions is shown. In this balance, the shortage in FO EUR in September 2019 has decreased by 4.5 FTE. This shows that quite a large number of iterations is going to be required to completely resolve this shortage and prevent larger shortages in the SO ICA position.

	Jul 2019	Aug 2019	Sep 2019	Oct 2019	Nov 2019	Dec 2019
CP ICA	1.1	2.5	7.2	8.6	3.9	-3.2
CP EUR	7.4	10.5	-4.1	-5.2	-6.1	-4.2
FO EUR	13.2	7.0	-11.7	-3.3	1.6	0.6
FO ICA	10.9	9.1	6.7	11.1	11.9	7.7
SO ICA	12.0	0.2	-7.1	-10.9	-21.4	-19.5

Figure 4.7: Initial balance of the simplified example problem (in FTE)..

	Jul 2019	Aug 2019	Sep 2019	Oct 2019	Nov 2019	Dec 2019
CP ICA	1.1	2.5	7.2	8.6	3.9	-3.2
CP EUR	7.4	10.5	-3.4	-4.5	-5.5	-3.5
FO EUR	13.2	7.0	-11.7	-3.3	1.6	0.6
FO ICA	10.9	8.3	5.9	10.3	11.1	7.0
SO ICA	12.0	0.2	-7.1	-10.9	-21.4	-19.5

Figure 4.8: Balance after the first iteration in the simplified example problem (in FTE).

	Jul 2019	Aug 2019	Sep 2019	Oct 2019	Nov 2019	Dec 2019
CP ICA	1.1	2.5	7.2	8.6	3.9	-3.2
CP EUR	7.4	10.6	0.5	-0.5	-1.6	0.4
FO EUR	13.2	7.0	-11.7	-3.3	1.6	0.6
FO ICA	10.9	4.0	1.6	5.9	6.8	2.7
SO ICA	12.0	0.2	-7.1	-10.9	-21.4	-19.5

Figure 4.9: Balance after 8 iterations in the simplified example problem (in FTE).

	Jul 2019	Aug 2019	Sep 2019	Oct 2019	Nov 2019	Dec 2019
CP ICA	1.1	2.5	7.2	8.6	3.9	-3.2
CP EUR	7.4	10.6	0.5	-0.5	-1.6	0.4
FO EUR	13.2	7.5	-7.1	1.3	6.2	5.2
FO ICA	10.9	4.0	1.6	5.9	6.8	2.7
SO ICA	11.4	-0.7	-8.0	-11.8	-22.3	-20.4

Figure 4.10: Balance after 20 iterations in the simplified example problem (in FTE).

Table 4.1: First 21 planned transitions in the simplified example problem (full list can be bound in Table A.1).

Iteration	Type	Employee	From	To	Date
0	Transition	2394	FO ICA	CP EUR	2019-08-01
1	Transition	7130	FO ICA	CP EUR	2019-08-01
2	Transition	9196	FO ICA	CP EUR	2019-08-01
3	Transition	8049	FO ICA	CP EUR	2019-08-01
4	Transition	8398	FO ICA	CP EUR	2019-08-01
5	Transition	6878	FO ICA	CP EUR	2019-08-01
6	Transition	0613	FO ICA	CP EUR	2019-08-01
7	Transition	0223	FO ICA	CP EUR	2019-08-01
8	Corrected	8294	SO ICA	FO EUR	2019-08-01
9	Recruit	R35	-	SO ICA	2019-07-01
10	Corrected	2231	SO ICA	FO EUR	2019-08-01
11	Recruit	R36	-	SO ICA	2019-07-01
12	Corrected	5571	SO ICA	FO EUR	2019-08-01
13	Corrected	2656	SO ICA	FO EUR	2019-08-01
14	Recruit	R37	-	SO ICA	2019-07-01
15	Corrected	1844	SO ICA	FO EUR	2019-08-01
16	Recruit	R38	-	SO ICA	2019-07-01
17	Corrected	3165	SO ICA	FO EUR	2019-08-01
18	Recruit	R39	-	SO ICA	2019-07-01
19	Corrected	2015	SO ICA	FO EUR	2019-08-01
20	Recruit	R40	-	SO ICA	2019-07-01

In Figure 4.11 the final balance after 105 iterations is shown. This simplified problem results in a feasible crew plan. Fesible, in this applicaiton, means a crew plan without any shortages. The fact that transitions cannot be chosen completely but are partially dictated by seniority rules makes it harder to choose the optimal transition. As transition capacity for certain positions gets filled, transitions have to be moved to earlier stages. This means the balance of other positions does not only get influenced in the future, but also in the past. These factors make the transition planning problem a complex and complicated process that is hard to model and solve. Even for this small problem, with 30 transition options at each iteration and 105 iterations to solve the problem, $30^{105} = 1.25 \cdot 10^{155}$ possible solutions can be created. Clearly, it is not achievable to evaluate all possible solutions and an efficient solution algorithm has to be designed to create a solution with sufficient solution quality within a reasonable time.

In the next section, the design framework of this research will be presented, which will show what the structure and content of the proposed model will look like.

	Jul 2019	Aug 2019	Sep 2019	Oct 2019	Nov 2019	Dec 2019
CP 777	0.5	7.0	9.6	1.7	1.0	0.0
CP 737	18.7	11.0	0.1	3.1	2.4	4.2
FO 777	12.2	0.4	1.4	1.2	14.2	17.0
FO 737	11.2	7.8	0.6	4.5	9.8	9.5
SO 777	20.1	2.8	0.4	0.1	0.3	0.5

Figure 4.11: Final balance of the simplified example problem (in FTE)..

4.3. Research Framework

The framework that is used to test and analyse the proposed model in order to answer the research question is shown in Figure 4.12. On the left side, the input of the research is shown. From the chosen planning window, the promotion plan and other CLA agreements and the transition bids of pilots, a scenario is constructed. Next to this, the current workforce is used to determine the supply of pilots. Finally, the schedule and additional demand together with the size of the different crew positions determine the gross demand for pilots per crew position and date. From the scenario, supply and demand, the initial crew plan can be obtained. This initial crew plan is basically a zero-solution with the supply, demand and balance during the planning window without any transitions planned.

In order to obtain an optimal crew plan, a number of methods are developed. A heuristic planning model (Chapter 5) is designed that is able to plan transitions and hire recruits to improve the crew plan. To do this, the planning model uses a selection algorithm that select a transition out of the available options. Also, constraints for the planning model are defined from CLA agreements as well as heuristic methods. Finally, an objective function is designed that aims to capture the quality of the crew plan in a single measure by summing the relative cost of shortages and surplusses. With this objective function, the crew plan can be evaluated and the optimal crew plan is assumed to be the crew plan with the lowest objective function value.

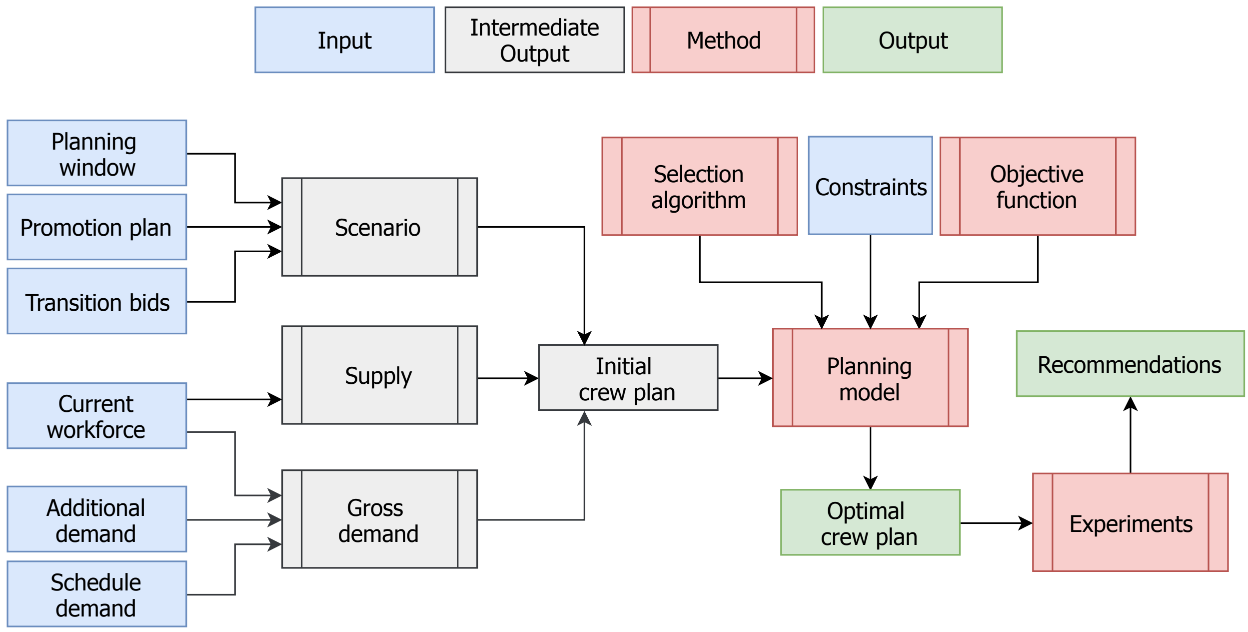

Figure 4.12: Flow diagram of the research framework for the cockpit crew transition planning problem.

The selection algorithm proposed in Chapter 6 is based on a tree-search method. In this algorithm, the different options in subsequent iterations are placed in a tree. A number of parameters are defined that limit the search space and thereby limit the computation time. By changing the algorithm's parameters and the way the tree is built, different configurations of the algorithm are developed. All of these configurations produce a different crew plan, dependent on the size of search space and build strategy as different choices are made in the process.

In order to test the performance of the various configurations, a number of experiments are performed from which the algorithms can be compared. The input for these experiments consists of the supply and demand for pilots from a major European airline for various start dates. In order to create a crew plan, all planned transitions from the start date to the end date (a year in advance) are deleted and the selection algorithm and planning model are used to develop a new crew plan. From these experiments, conclusions and recommendations can be made regarding the modelling op cockpit crew planning and the research question can be answered in the conclusion.

4.4. Research Assumptions

A number of assumptions have been made to clearly scope the research and its place within the cockpit crew planning and scheduling process. The major assumptions and their implications are presented in this section while any assumptions specific to a subproblem or model will be discussed in the relevant chapters.

- Transition planning problem is modelled after a European legacy carrier. In this reference airline, pilots have the option to place bids on certain positions and transitions are awarded in seniority order to pilots with a valid bid on that position.
- In cases when no bids are placed on a crew position, airlines usually assign a transition to the least senior pilot in a certain position, based on a number of parameters. In the reference airline, however, this rarely happend and the process of assigning the transition is quite complex. Therefore, in the model, no pilot is transitioned to that position if no bids are placed on that position.
- Pilots are only qualified to operate on one aircraft type and cockpit position (rank).
- The only methods to resolve shortages are to schedule transitions or hire recruits. In some cases, crew planning models are able to rearrange the distribution of flexible demand, however, the focus of this research is planning transitions only.
- The capacity for transitions is assumed to be only dependent on the maximum capacity as dictated by the simulator availability. The actual capacity might be decreased in some cases due to the limited availability of instructors.
- For each time period, the average supply and demand are used to determine the balance, therefore, it is assumed that any variances within these time periods can be resolved by scheduling strategies.
- The balance values are determined in FTE. Therefore, all shortage and surplus costs are determined as a function of the balance value only and it is assumed that the cost (shortage and surplus) is independent of the size of the position. This means that a balance value of -1 is evaluated the same for both a position group size of 10 FTE and 1000 FTE.

5 Planning Model

In this chapter, the heuristic planning model is presented. The first step is to define the objective function used in the model (Section 5.1) and input (Section 5.2). Next, the assumptions made in the development are discussed in Section 5.3 and the planning process is detailed in Section 5.4. The chapter is conluded with the model output in Section 5.5 and verification of results in Section 5.6.

5.1. Objective

The objective of the planning model is to create a feasible crew plan a number of years in advance by scheduling transitions for pilots to different crew positions and by recruiting new pilots. In this case, feasible means no shortages of supply given the demand for the different pilot positions. As this is not always possible given the constraints present within the problem (transition capacity for example), a trade-off has to be made between various parameters in order to obtain an optimal crew plan. An objective function has to be designed that accurately weighs all different aspects of a crew plan which can then be used to determine the optimal crew plan.

5.1.1. Objective Function

If the only focus of the model was to create a feasible crew plan, the objective function would be developed to minimize the total shortage in the system:

$$\min \sum_{p}^{P}\sum_{d}^{D}\left|B_{p,d}\right| \quad \forall d \in D,\ p \in P \quad \text{if } B_{p,d} < 0, \tag{5.1}$$

where D is the set of months in the planning window, P the available positions within the airline and $B_{p,d}$ the balance value at position p and date d. However, this does not accurately capture the entire objective. First of all, not all positions are equally important. This can be accounted for with the average salary of pilots in a position, or by a manually determined weight to account for more problematic positions (for example because of key positions in the hierarchical system). Secondly, the shortage cost is not the only factor in the problem. The goal of the model should not be to minimise shortages at all cost, which could result in large surplusses, a very high number of transitions and much more. Third of all, the absolute value of shortages as well as the total duration of a shortage in a position should be taken into account. For example, a shortage lasting for six months is harder to compensate than six isolated one month shortages. Also, when shortages start to become larger, the effort it takes to resolve this by other measures becomes increasingly harder.

As the amount of transitions is already limited by the available capacity, the number or cost of transitions is not incorporated in the objective function. The excess staff, however, should be penalised to a certain degree in order to steer the model to a solution with both minimal shortages and surplusses. Therefore, a surplus cost is added to the objective function. In addition, both the shortage and surplus are altered with an exponential factor. This exponential factor is used as the cost of having a shortage or surplus increases non-linearly with increasing absolute values. Furthermore, the different positions are weighted in order to capture the relative weight of the different crew positions. Finally, the cost of shortages is further increased by multiplying each shortage with the length of the consecutive shortage in that crew position. In mathematical form, the objective function now becomes:

$$\min \ \left(C_{\text{shortage}} + C_{\text{surplus}}\right) \tag{5.2}$$

$$C_{\text{shortage}} = \sum_{p}^{P} \left(w_p \cdot \sum_{d}^{D} l_n \cdot (-B_{p,d})^{\beta^-} \quad \text{if } B_{p,d} < 0 \quad \forall d \in D \right) \forall p \in P \tag{5.3}$$

$$C_{\text{surplus}} = \sum_{p}^{P} \left(w_p \cdot \sum_{d}^{D} (B_{p,d})^{\beta^+} \quad \text{if } B_{p,d} > 0 \quad \forall d \in D \right) \forall p \in P, \tag{5.4}$$

where w_p is the weight given to a position, l_n the length of consecutive shortages in a position, and β^- and β^+ the exponential factor for shortages and surplusses.

5.1.2. Parameter Value Selection

Within the scope of this research, some values are fixed while others can be analysed in a sensitivity analysis with the goal to find the best parameter value out of a range of values. First of all, the weight of a position is set equal to the average monthly salary of that position, thereby increasing the importance of those positions with the highest salary. This is done as these positions are more expensive and also feature on the top of the promotional hierarchy, this are the hardest to replace if no surplusses are available in the system. Furthermore, since the positive and negative counterparts in the equation aim to resemble the relative importance of shortages versus surplusses, one of those parameters can be set to 1 while the other is varied. In this case, β^+ is set to 1, thereby simplifying the shortage and surplus cost to:

$$C_{\text{shortage}} = \sum_{p}^{P} \left(w_p \cdot \sum_{d}^{D} l_n \cdot (-B_{p,d})^{\beta} \quad \text{if } B_{p,d} < 0 \quad \forall d \in D \right) \forall p \in P \tag{5.5}$$

$$C_{\text{surplus}} = \sum_{p}^{P} \left(w_p \cdot \sum_{d}^{D} B_{p,d} \quad \text{if } B_{p,d} > 0 \quad \forall d \in D \right) \forall p \in P, \tag{5.6}$$

where β^+ has been abbreviated to β. With this objective function definition, only one parameter has to be altered to change the objective function and thus the solutions generated by the model.

In the sensitivity analysis on the $beta$ parameter presented in Chapter 9, five different values (1, $\sqrt{2}$, 2, 3 and 4) are tested on four different scenarios. As minimising the shortages is more important than minimising the surplusses, values below 1 are not usefull and the minimum value chosen for $beta$ is equal to one. The maximum is chosen to be 4 as the relative cost of surplusses with respect to shortages with higher values for β will become negligible and the objective function almost only focuses on the shortages again. The results of the sensitivity analyse also do not show a trend that higher β values are better.

From the results of the sensitivity analysis, it was concluded that out of the five tested β values, a value of 2 produces the best results. This value is therefore used in the objective function defined for the planning model.

5.2. Model Input

The goal of the planning model is to develop an optimal crew plan as determined by the objective function. Various forms of input are necessary for this process. The input has been subdivided into four categories: supply, demand, labour agreement rules and others.

5.2.1. Supply

The supply for a crew plan consists of the pilots currently in service. In the supply, entries are created for every pilot-position pair, which means that pilots have multiple entries in the set if they are awarded a transition by the model. The pilots all have different characteristics that affect their salary, usage and more. These characteristics include the crew position, date of birth, the full-time equivalance (FTE) the pilot is working at, inservice and retirement dates and position bids. From the given pilot characteristics, the supply (in FTE) can be calculated per time-period and crew position. Also, the salary costs can be determined from the chracteristics of the individual pilots.

The supply is used in the model to calculate the balance per crew position and time-period. Furthermore, when transitions are planned, the supply for that period has to be changed according to the planned transition and the rules attached to that transition.

5.2.2. Demand

The demand for crew comes primarily from the flight schedule for which pilots are needed to operate the aircraft. Besides this nett demand, a number of other factors create demand on top of the nett demand and together comprise the gross or total demand. This additional demand originates from aspects such as pilot's vacation, training sessions required to keep one's certification, unavailability because of illness, transition periods and instructor duties.

After calculating the gross demand in FTE per time-period and crew position, the demand can be subtracted from the supply to determine the balance for each time-period and position. This balance then serves as the main input for every iteration in the planning model, as the balance dictates the next transition to be planned by the model.

5.2.3. Transition Rules

A number of rules with regards to transitions and recruitment have been agreed between the union and airline and written down in the collective labour agreement (CLA). When finding a pilot that is eligible for a transition to a certain crew position, the rules dictate whether and which pilot should be awarded the transition. These rules are thus used in the planning model to select the pilot for a chosen transition.

- **Seniority**: All transitions are awarded to pilots in strict seniority order. Every pilot has a seniority number, where a lower number equals a higher seniority, in other words, the pilot with the lowest seniority number should be awarded the transition.
- **Function Binding**: To prevent pilots from transitioning too often, a pilot is ineligible for a transition to a position for a certain period depending on the origin and destination position. In this period, the pilot can only be awarded a transition to a position if no pilot with a valid function bid without binding exists.
- **Employement Binding**: Similar to function binding, transitions to certain positions can only be awarded to pilots who have been employed by the airline. This is done to ensure pilots have sufficient experience for a certain crew position.
- **Retirement Binding**: In order to prevent pilots close to retirement from making an expensive transition, depending on the contract percentage (FTE), pilots within the last period before retirement are not eligible for a transition.
- **Disappearing Fleet**: When a certain aircraft type is disappearing or shrinking, surplusses in the relevant positions can be prevented by giving those pilots priority over more senior pilots in non-disappearing fleets.
- **Direct Entry**: In the hierarchical system used by an airline, only a number of positions (the lowest) can be filled by recruits while the other positions do not allow direct entry and therefore have to be filled by planning transitions for pilots from different positions.

5.2.4. Parameters

Some other input parameters that influence the process of awarding transitions to pilots and the effect these transitions have on the crew plan include transition bids, the capacity for transitions and the various characteristics different types of transitions have.

- **Transition bids**: In European airlines, pilots can make bids on the positions of their preference. Since for all the transitions, the most senior pilot with a valid transition bid (based on the various binding rules) is considered first, these bids can have a great influence on the origin of a pilot and therefore the resulting balance after the transition.
- **Transition capacity**: The capacity of transitions to a certain position is limited by both simulator capacity and instructor availability. These instructors often also operate as pilots in the same position, which further increases the demand when a transition is planned.
- **Transition characteristics**: The transitions that have to be planned have a number of characteristics that influence the supply and demand for a certain period of time. When a pilot starts a transition, training is first

done in a flight simulator, which means the pilot can not contribute to the supply of pilots in that crew position. After this period the pilot will be able to fly according to his new position with an instructor present on one of the other positions in the aircraft. Based on the type of transition, it will also take up a predetermined part of the given capacity (TQ), based on both the origin and destination position.

5.3. Model Assumptions

In this section, the assumptions made for the planning model are presented.

Simplifications
When modelling real-life systems, often, a number of assumptions have to be made to be able to translate that system to a mathematical model.

- The modelling time period (frequency) has been set to one month throughout the development, however, this can be changed for the entire period but also made variable with respect to time (for instance model the first year at one month frequency and the second year at a frequency of two months). A sensitivity analysis is presented in Chapter 9 to study the effect of different modelling frequencies.
- All transitions are scheduled to start on the first day of the month. By doing this, the number of decision variables decreases (from 28 to 31 per month to only one) and the solution quality decreases as not all transitions can be planned on the best date.

Pilots
Pilot behaviour influences the system in a large number of ways. A number of assumptions have to be made regarding this behaviour:

- Pilots do not change their transition bids. Usually, pilots are able to change their transition bids every six months or year, however, usually these changes are not large and with a planning window of 1 year, not applicable for the majority of the planning window.
- Planned transitions are assumed to always be carried out as planned. Transitions can be changed for a number of reasons. First of all, the date can be changed in a later stage due to circumstances. Secondly, the pilot can get sick and thereby postpone of completely cancel the transtion. Finally, in some airlines, pilots can decide to pass on a transition because they give priority to, for instance, a holiday. This creates disruptions to the crew plan since a new pilot has to be found for the transitions that may not come from the same position. However, these factors are hard to predict and are therefore not taken into account.
- Pilots do not change their FTE percentage. Pilots can decide to work more or fewer hours in a week if they desire. However, since the planning horizon used in this research is relatively small, the effect of these changes to the overall supply is not large and it is therefore not taken into account.

Recruits
When a new recruit is added to the airline, a number of assumptions regarding this recruit have to be made:

- Recruits start working at 1.0 FTE.
- The age of recruits will be set to 25.
- The retirement age of new recruits is set to be 58. However, this choice will not have any effect on the results as the retirement age will only influence transitions a couple of years before retirement.

5.4. Process

This section describes the process used in the planning model in order to create a crew plan. Various methods that have been applied in the model are presented and discussed.

5.4.1. Planning Model Overview

In Figure 5.1, a schematic overview of the planning model is presented in the form of a flow diagram. This diagram shows the process of creating an optimal crew plan as defined by the objective function from start to finish. In the next sections, some characteristics and methodologies of the planning model will be discussed in more detail. The blocks in this flow diagram can be described as follows:

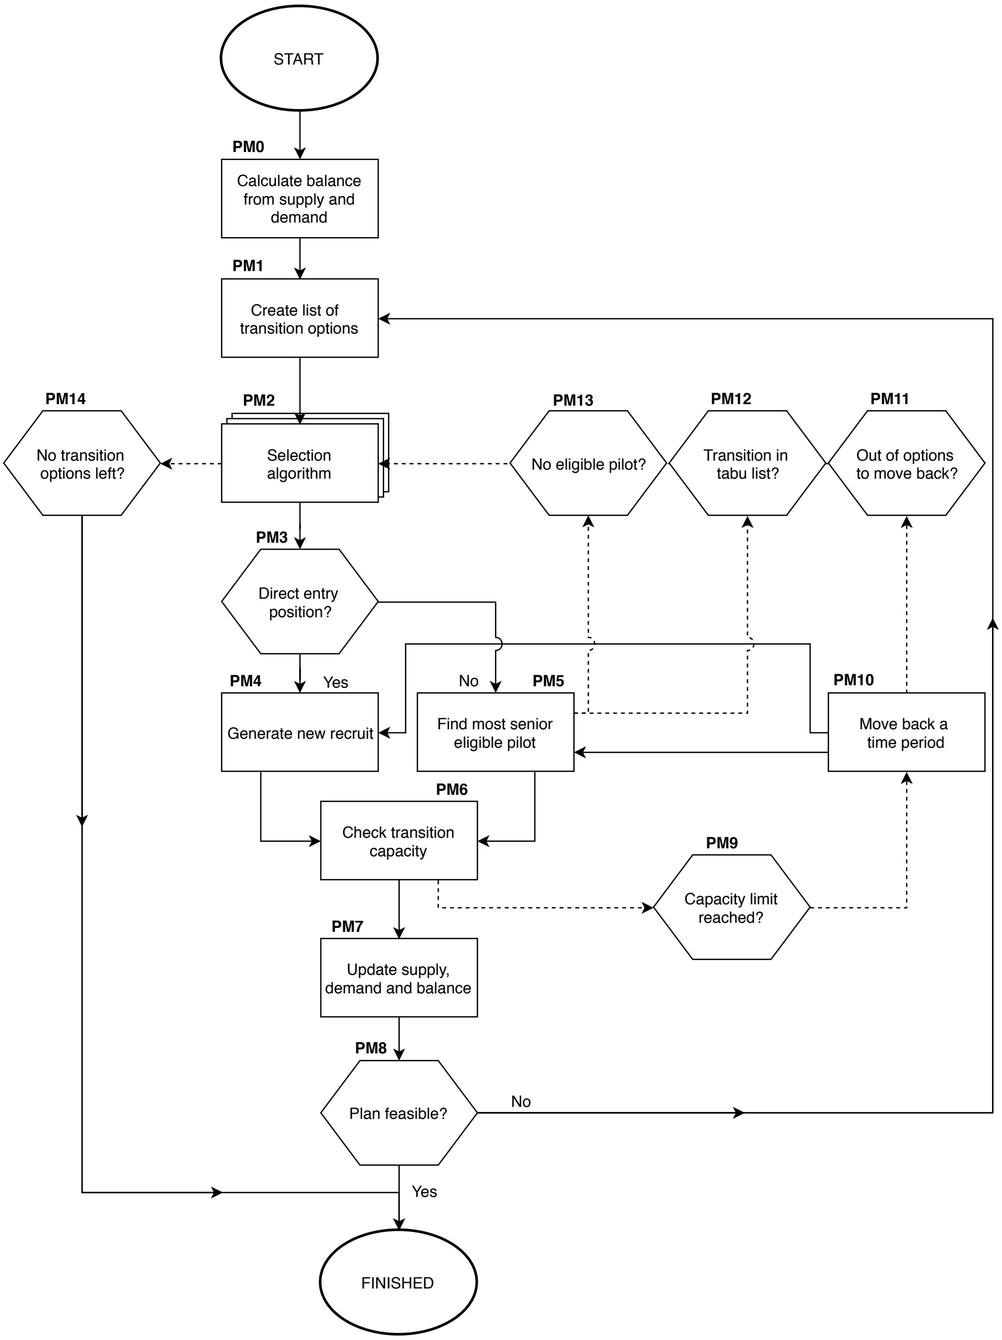

Figure 5.1: Flow diagram presenting the process of creating an optimal crew plan.

Calculating the balance (PM0)
From the demand for and supply of pilots, the initial balance can be determined per crew position and time period.

List transition options (PM1)
From the balance calculated in PM0, a list of possible and allowed transition options is created. These transition options are defined as a crew position a transition is planned to and a date. Whether transition options are allowed is determined by some of the rules that are presented in Subsection 5.4.3.

Selection algorithm (PM2)
In this block, the selection algorithm introduced in Section 4.3 is applied to the list of transition options determined in PM1. The chosen option is dependent on the configuration of the selection algorithm, which will be explained in more detail in Chapter 6.

Recruit or transition? (PM3)
In PM3, it is checked whether the crew position selected in PM2 is a direct entry position or not. If the position is a direct entry position, the planned transition is a recruitment of a new pilot, while in other cases a transition from a pilot in a different position has to be planned.

New recruit (PM4)
If the selected transition option is a transition to a direct entry position, it means a new pilot will be recruited and added into the system.

Plan transition (PM5)
If instead, the selected transition option is not to a direct entry position as determined in block PM2, a transition has to be planned for a pilot from his current position to the selected position. First, the most senior pilot with a valid bid who is eligible for that transition has to be determined based on the constraints as defined in Section 5.2.

Transition capacity (PM6)
For each transition and recruitment, the available transition capacity has to be checked. If no capacity restrictions exist, the transition or recruit is planned according to the assumptions presented in Section 5.3.

Updating supply, demand and balance (PM7)
If the transition or recruit that is planned, the supply, demand and balance can be updated. For the supply, the FTE percentage of the pilot awarded a transition should be subtracted from his old position and added to his new position from the start date of the transition onwards. The demand that is a function of the strength in a position should be recalculated for the new strength values and added to the constant demand. Furthermore, extra demand has to be determined for the time the pilot is in transition, as dictated by the length of the planned transition. When the supply and demand are updated, the balance can also be determined by subtracting the demand from the supply.

Exceptions (PM10, PM11, PM12, PM13)
A number of exceptions have to be dealt with when planning a transition or recruit. If in PM6 is was found that there is not sufficient capacity remaining to accommodate the transition, it is moved to an earlier date (in PM9 and PM10) as it is still desirable to resolve the shortage. This process is continued until a transition date is found where there is sufficient capacity, or until the model reaches the beginning of the planning window (which is checked in PM11). When planning a transition, it should again be checked which pilot is the most senior eligible pilot for that transition after the date has been changed as this pilot could also have changed. Furthermore, if no eligible pilots exist (PM13), or if the proposed transition is featured in the tabu list (PM12), the option is discarded and the next transition option is selected by the selection algorithm.

Finalising the model (PM8, PM14)
There are two ways the model can be terminated. When a feasible crew plan is reached in which no shortages exist, the algorithm is finalised in PM8 after recalculating the balance. In other cases, the algorithm reaches a point at which shortages still exist, but no transition options yield a valid transition (PM14). In that case, the algorithm will stop after trying to select the next option and finding out no options are left.

5.4.2. Local Search Algorithm

The planning model employs a local search technique to find the optimal crew plan. The model starts with an initial solution based on the current supply and demand of pilots. The model then moves from solution to solution by applying local changes, defined in the neighborhood of the solution. In this case, the neighborhood is defined as all solution with one additional transition or recruit compared to the current solution, as well as all solutions in which one transition is changed compared to the current solution. As this neighborhood is quite large, the search space is made smaller through a number of methods that are presented in the next subsections.

5.4.3. Rule-based

As the number of options at each iteration is very large (the amount of positions multiplied with the number of time periods in the planning window), a rule-based algorithm is applied to decrease the number of possibilities and thereby improve the computation time.

Recruit or transition
Different strategies have to be followed when scheduling a transition to a direct-entry position as opposed to another position. For direct-entry positions, a new recruit is generated and the transition can be planned. For other positions, however, the most senior pilot with a valid bid has to be found and not only the supply and demand of the destination position change but also those of the previous position of the pilot. This possibly creates more shortages in that position. This application of this rule can be found in Figure 5.1 in PM3.

Only shortages
As discussed in Section 5.1, the main objective of the planning model is to construct a crew plan without shortages. Therefore, the focus of the model is to resolve these shortages and ignore the surplusses. This, however, does not always mean that transitions will never be planned in time periods with a positive balance, as factors such as the transition time and capacity influence the transition date. This rule is applied to the planning model process in Figure 5.1 in PM1.

Transition date
Transitions generate additional demand in the first period after the transition date for training sessions. Because of this, transitions are not planned in the time period of the shortage if the balance is lower than a certain threshold (after sensitivity analysis of the threshold, a balance value of -0.1 for a time period of one month was chosen) as additional supply will be smaller than the extra demand. Instead, if the balance is lower than the threshold, the transition is planned one month earlier. When constructing the list of transition options in PM1 in Figure 5.1, this rule is immediately applied to the options.

Pilot selection
In order to find the most senior, eligible pilot for a transition to a certain position, a number of steps have to be taken. First of all, the most senior pilot, without binding of any kind with a current bid on the open position is found. Since pilots in the disappearing fleet can be given priority over other pilots, the most senior pilot that is free of binding and in a disappearing fleet is found and both transitions are planned and evaluated independently in order to choose the best transition. If no pilots with a current bid on the open position are free of binding, the most senior pilot who is bounded under the function binding rules is found and this pilot is awarded the transition. It is important to note that only pilots who are bounded to their current position are taken into account. Pilots who are bounded by their employment start date or retirement date are never allowed to be awarded a transition. This rule applies to block PM5 in Figure 5.1.

Transition capacity
Since transition capacity is limited because of a number of reasons (such as simulator capacity and instructor availability), the capacity for a certain transition should be checked prior to planning that transition. However, if no capacity is available, the shortage should not be to disregard as the goal of the planning model is to resolve all shortages. Instead, when no capacity is available to resolve a shortage, the transition will be moved a time period back (earlier in time). This increases the salary cost as the pilot is promoted earlier. The shortage will only be disregarded if the beginning of the planning window is reached and no possible transition date is found. This rule is applied in the planning model process in blocks PM6, PM9, PM10 and PM11.

Disappearing fleet
As mentioned before, pilots in the disappearing fleet can be granted priority over other pilots when awarding transitions. Large surplusses can occur in disappearing fleet since pilots often cannot be transitioned out of the crew position as fast as necessary. However, one thing to be careful of is that even though transitions from disappearing fleet can seem favourable at first, shortages in those aircraft types should be avoided as this would require additional transitions to that aircraft type. Therefore, the transition to a disappearing fleet is disregarded when it creates (additional) shortages in that position. When selecting the most senior eligible pilot in PM5 in Figure 5.1, this rule is applied.

5.4.4. Tabu Search

Tabu search is a local search metaheuristic previously applied to the cockpit crew planning problem by Thalén (2010). Because of the sneiority rules in the problem, it often happens that the models is trapped in a loop where pilots are awarded transitions between two crew positions continuously as every transition creates a shortage in the other position and vice versa. The core idea of the algorithm is, therefore, to block certain moves in order to steer the algorithm to new solutions and prevent getting stuck in a local optimum. Glover (1997) argues that intelligent choices are better than random choices, which is hard to argue. As Glover states: "Efficiency and quality can be greatly affected by using intelligent procedures for isolating effective candidate moves, rather than trying to evaluate every possible move in a current neighbourhood of alternatives." This is even more applicable when the neighboorhood of possible choices is large and difficult to examine, both of which are the case in the cockpit crew planning problem as many transitions can be scheduled at each iteration and the added benefit of a transition is hard to quantify directly.

A simple application of the tabu search principle is to store every previous solution in a tabu list and prohibit the algorithm from returning to these solutions. However, this method is inefficient as it only prohibits the exact solutions previously evaluated. Instead, a tabu list is created with in it the inverse of the moves performed in the past iterations. This is done as it prevents the model to plan transitions between two crew positions continuously, as in this case, the supply will remain the same in both positions while the demand increases because of the planned transitions. This is an undesired effect as it does not improve the crew plan and takes up capacity for transitions to other functions within the same aircraft type.

The tabu search thus prevents inverse transitions between specific crew positions and any date for a number of iterations. This technique is implemented in the planning model process in block PM12 in Figure 5.1. The length of the tabu list (i.e. for how many iterations previous transitions are blocked) is decided through a sensitvity analysis in Section 9.4 and is equal to 10.

5.5. Model Output

In this section, the various outputs of the planning model are presented and explained, these outputs are used by airline crew planners to evaluate the crew plan and if necessary, make changes to the system.

5.5.1. Key Performance Indicators

The first and most important KPI is the objective function as defined in Section 5.1. This objective function calculates the cost of shortages and surplusses in supply in the system. It has been defined in order to provide a single parameter to assess the quality of a given solution. To calculate the objective function, the balance of demand and supply have to be determined and the salary cost per crew position has to be known. Apart from the costs captured in the objective function, more cost factors are present in the system.

The largest part of the cockpit crew cost is the crew salary. This cost is based on the pilot's current position, working percentage, age and number of years in service. The surplus cost used in the objective function definition is actually part of the salary cost. For the objective function, it is not necessary to use the total salary cost as the largest part of this cost is simply needed for the airline to operate the planned flight schedule and only the surplusses should be minimised.

Part of the cost for transitions comes from additional demand for pilots and instructors for the transitions. This cost is, however, already captured in the salary cost. Another part can be calculated separately based on the cost for the different sessions in the planned transitions (this includes the cost for simulators, flight hours, etc.).

The final aspect of the crew plan cost is the shortage cost. This shortage cost is quantified relative to the cost of surplusses in the objective function.

The objective function is used in the model to assess the quality of the evaluated solutions. The other KPI's can be used, if necessary, to perform a deeper analysis into the quality of the solution.

5.5.2. Balance

An important output for an airline is the balance showing the shortages and surplusses at the different crew positions for different months, as can be seen in Figure 5.2. The model uses the objective function previously defined to assess the quality of the balance in a single parameter. However, in some cases the balance is needed to obtain more information about the proposed solution.

By looking at the balance, crew planners can assess any problems in the crew plan. In the provided example, high surplusses can be found in the first two months of the planning window. This might urge the crew planner to shift budgets between months. In later months, shortages can be found for a number of positions. If these shortages cannot be resolved with transitions and/or recruitments, other methods have to be applied to resolve the shortages as it is of vital importance to an airline to resolve all shortages. However, these additional measures that can be performed from an analysis of the balance are outside the scope of this research and, therefore, not incorporated into the model.

	Jul 2019	Aug 2019	Sep 2019	Oct 2019	Nov 2019	Dec 2019
CP ICA	1.1	2.5	7.2	8.6	3.9	-3.2
CP EUR	7.4	10.6	0.5	-0.5	-1.6	0.4
FO EUR	13.2	7.5	-7.1	1.3	6.2	5.2
FO ICA	10.9	4.0	1.6	5.9	6.8	2.7
SO ICA	11.4	-0.7	-8.0	-11.8	-22.3	-20.4

Figure 5.2: Example of the crew plan balance output.

5.5.3. Transitions

In terms of transitions, two important outputs are generated. The first is a list of all planned transitions. This list includes the employee id of the pilot that is awarded the transition, the position the pilot was previously on, the new position he is transitioned to and the start date of the transition, as can be seen in Table 5.1.

The second part is a table showing the used transition capacity per fleet and month (Figure 5.3). When comparing this to the maximum capacity as defined by the airline, possible bottlenecks can be identified. If for example the total transition capacity is used for an aircraft type for a long period of time, it might be profitable to increase the capacity for that aircraft type. It also serves as a verification tool to see if the capacity limitations are implemented correctly by checking if no values rise above the planned capacity.

	Dec 2017	Jan 2018	Feb 2018	Mar 2018	Apr 2018	May 2018	Jun 2018	Jul 2018	Aug 2018	Sep 2018	Oct 2018	Nov 2018	Dec 2018	Capacity
D	0	1	13	13	10	12	11	11	10	12	5	4	0	13
C	4	4	4	4	4	4	4	4	4	0	0	0	0	4
B	4	4	4	4	4	4	4	4	4	4	4	3	1	4
A	0	0	4	10	10	9	3	5	9	2	1	0	0	10
E	0	0	0	0	0	1	6	1	1	9	0	7	0	9

Figure 5.3: Utilised transition capacity per crew position and date combination (rounded to the nearest integer) as well as the maximum capacity.

5.6. Verification

In this section, the planning model is verified using the simplified example problem previously presented in Section 4.1. For a number of steps in the model, the calculations and choices of the model is evaluated manually to see that everything is done as it is supposed to be.

The gross demand and supply have already been presented in Section 4.1. From these two tables, the balance at each crew position and date can be calculated as:

$$B_{p,d} = S_{p,d} - D_{p,d}, \tag{5.7}$$

where $B_{p,d}$ is the balance value at position p and date d, $S_{p,d}$ the supply of pilots at the same position and date and

Table 5.1: Sample of planned transitions output.

Index	Employee	Seat from	Seat to	Date
23	5172	SO C	FO C	2018-01-01
28	2649	FO F	FO C	2018-01-01
36	5420	SO D	FO A	2018-03-01
46	5760	CP C	CP D	2018-03-01
73	2526	FO F	FO B	2018-03-01
88	5513	FO C	CP C	2018-03-01
95	5842	CP A	CP B	2018-04-01
105	5591	SO C	FO D	2018-02-01
114	5259	CP C	CP D	2018-02-01
116	2539	FO F	FO C	2018-01-01
120	5198	CP D	CP C	2018-01-01
189	7944	-	SO D	2018-05-01

Table 5.2: List of transition options used for verification of the planning model.

	Date	Position	Position weight
0	2019-09-01	CP EUR	24
1	2019-09-01	FO EUR	14
2	2019-09-01	SO ICA	11
3	2019-10-01	CP EUR	24
4	2019-10-01	FO EUR	14
5	2019-10-01	SO ICA	11
6	2019-11-01	CP EUR	24
7	2019-11-01	SO ICA	11
8	2019-12-01	CP ICA	29
9	2019-12-01	CP EUR	24
10	2019-12-01	SO ICA	11

$D_{p,d}$ the gross or total demand. Doing so yields the balance as shown in Figure 5.4, which is identical to the balance already shown in Section 4.1.

	Jul 2019	Aug 2019	Sep 2019	Oct 2019	Nov 2019	Dec 2019
CP ICA	1.1	2.5	7.2	8.6	3.9	-3.2
CP EUR	7.4	10.5	-4.1	-5.2	-6.1	-4.2
FO EUR	13.2	7.0	-11.7	-3.3	1.6	0.6
FO ICA	10.9	9.1	6.7	11.1	11.9	7.7
SO ICA	12.0	0.2	-7.1	-10.9	-21.4	-19.5

Figure 5.4: Calculated balance for the example problem.

From this balance, the list of transition options can be constructedbased on the rules defined in Section 5.4. For the first iteration, the tabu list is empty and all shortages appear in the options list, which is ordered according to the sorting method presented in Section 6.1. As can be seen in Table 5.2, the list is ordered by date first, and by the salary cost for the crew positions second.

The first transition option is selected and planned if possible. In this case, this means a pilot is awarded a transition to CP EUR in September 2019. However, since the balance value is lower than -0.1, one of the rules in the planning model dictates the transition should be planned a month earlier, in August 2019. The next step is to find the most senior pilot with a valid transition bid on this position. In Table 5.3, the ten most senior pilots with a function bid on CP EUR are shown. If the first pilot shown conforms to all binding rules, the transition can be awarded to this pilot. If not, the next pilot should be evaluated, until a pilot with a valid function bid is found.

Table 5.3: List of pilots with a bid on CP EUR, sorted by seniority (first 10 pilots shown).

Employee	Position	FTE	Seniority	Inservice	Position start	Retirement	Bids
0671	FO EUR	0.50	338	01/28/1987	01/26/2009	01/07/2023	CP ICA/EUR - FO ICA
2394	FO ICA	1.00	456	10/25/1993	03/28/2006	09/06/2024	CP ICA/EUR
8802	FO ICA	0.80	567	11/29/1985	01/23/2004	01/25/2022	CP ICA/EUR
7130	FO ICA	1.00	593	01/13/1995	11/09/2014	05/31/2025	CP ICA/EUR
7105	FO EUR	0.50	638	04/24/1995	12/12/2001	04/15/2030	CP ICA/EUR - FO ICA
9196	FO ICA	0.67	714	12/04/1995	11/10/2005	03/14/2036	CP ICA/EUR
8049	FO ICA	0.80	820	10/24/1994	10/19/2006	02/14/2027	CP ICA/EUR
8398	FO ICA	0.80	826	01/02/1995	01/02/2016	10/30/2029	CP ICA/EUR
6878	FO ICA	0.80	885	08/26/1996	07/07/2007	08/13/2029	CP ICA/EUR
0613	FO ICA	0.67	922	12/02/1996	04/06/2008	05/03/2024	CP ICA/EUR

The first pilot in the sorted list of pilots with a bid on CP EUR, presented in Table 5.3 is currently operating as a first officer at the European fleet, at 0.5 FTE. For the transition from FO EUR to CP EUR, a number of binding rules apply. As transitions from FO EUR to CP EUR are not allowed in the designed promotional hierarchy, this first pilot can therefore be disregarded. The second pilot in the list is currently employed as a first officer on the ICA fleet and working full time ($1.0FTE$). The transition from FO ICA to CP EUR is allowed, with a binding term of 3 years. Since the pilot has been at his current position since 2006, this poses no problem at all. The pilot should also be in service at the airline for 6 years to be eligible for a transition to CP EUR. As the pilot started at the airline in 1993, this also poses no complications. Finally, for pilots operating full time, to be eligible for a transition, their retirement date should be at least 2.5 years in the future. With a retirement date in 2024, the pilot also complies to this rule. The second pilot in the list is thus eligible for the transition to CP EUR in August 2019 and this transition can now be planned.

With the required transition found, the transition capacity should be checked. As this is the first transition to be planned in the simulation, no transition capacity is used yet and all values in Figure 5.5 are therefore equal to 0. As the selected transition from FO ICA to CP EUR requires 0.5 TQ, the updated transition capacity in Figure 5.6 shows a value of 0.5 for the EUR fleet in August 2019.

	Jun 2019	Jul 2019	Aug 2019	Sep 2019	Oct 2019	Nov 2019	Dec 2019	Capacity
ICA	0.0	0.0	0.0	0.0	0.0	0.0	0.0	10
EUR	0.0	0.0	0.0	0.0	0.0	0.0	0.0	13

Figure 5.5: Used transition capacity prior to the selected transition.

	Jun 2019	Jul 2019	Aug 2019	Sep 2019	Oct 2019	Nov 2019	Dec 2019	Capacity
ICA	0.0	0.0	0.0	0.0	0.0	0.0	0.0	10
EUR	0.0	0.0	0.5	0.0	0.0	0.0	0.0	13

Figure 5.6: Used transition capacity after the selected transition.

	Jul 2019	Aug 2019	Sep 2019	Oct 2019	Nov 2019	Dec 2019
CP ICA	617	616	616	613	609	604
CP EUR	623	620	615	615	612	612
FO EUR	584	578	574	574	574	574
FO ICA	587	585	584	583	583	583
SO ICA	405	405	405	405	405	405

Figure 5.7: Updated supply of pilots (in FTE).

	Jul 2019	Aug 2019	Sep 2019	Oct 2019	Nov 2019	Dec 2019
CP ICA	616	614	609	604	605	608
CP EUR	615	609	618	619	617	615
FO EUR	571	571	586	578	573	574
FO ICA	576	577	578	573	572	576
SO ICA	393	404	412	416	426	424

Figure 5.8: Updated gross demand for pilots (in FTE).

Following the transition, the supply and gross demand are recalculated, as can be seen in Figure 5.7 and Figure 5.8, respectively. With the new supply of pilots and gross demand calculated, the updated balance can be determined using Equation 5.7. This results in the balance in Figure 5.9. A comparions with the second balance presented in Section 4.1 (Figure 5.10) shows both balances are identical.

	Jul 2019	Aug 2019	Sep 2019	Oct 2019	Nov 2019	Dec 2019
CP ICA	1.1	2.5	7.2	8.6	3.9	-3.2
CP EUR	7.4	10.5	-3.4	-4.5	-5.5	-3.5
FO EUR	13.2	7.0	-11.7	-3.3	1.6	0.6
FO ICA	10.9	8.3	5.9	10.3	11.1	7.0
SO ICA	12.0	0.2	-7.1	-10.9	-21.4	-19.5

Figure 5.9: Manual verification of the balance after the first iteration in the example problem.

	Jul 2019	Aug 2019	Sep 2019	Oct 2019	Nov 2019	Dec 2019
CP ICA	1.1	2.5	7.2	8.6	3.9	-3.2
CP EUR	7.4	10.5	-3.4	-4.5	-5.5	-3.5
FO EUR	13.2	7.0	-11.7	-3.3	1.6	0.6
FO ICA	10.9	8.3	5.9	10.3	11.1	7.0
SO ICA	12.0	0.2	-7.1	-10.9	-21.4	-19.5

Figure 5.10: Balance after the first iteration in the simplified example problem (in FTE).

The presented process for manually verifying the model is repeated for a number of iterations. From this, it can be concluded the the planning model meets the requirements.

6
Selection Algorithm

In order to obtain a solution for the crew planning problem using the planning model presented in Chapter 5, a selection algorithm is required that chooses which transition to plan at each iteration in the construction algorithm, constraint to the rule-based tabu-search planning model. The goal of this selection algorithm is to find a solution with sufficient quality in a limited time. In this chapter, some possibilities of selection algorithms will be discussed after which the results of these methods are presented in Chapter 7. The investigated selection algorithms are based on a tree search method as already discussed in Chapter 4.

6.1. Algorithm Design

The local tree search is designed to be able to make a trade-off between an exhaustive search method, greedy and naive algorithms. The method creates a tree of options with a potential based on the results a number of iterations later. Furthermore, a search width and depth are defined as the number of options at each level that are taken into account and the number of iterations the tree evaluates, respectively. To bette explain the algorithm, Figure 6.1 shows an example of a local tree search with a width of 2 and depth of 3. Here, the current state or solution is shown on the top of the tree. Since the width is equal to two, 1 and 2 are the two options evaluated in the tree search. A depth of 3 means this process is repeated three times as shown in Figure 6.1. The choice to select either option 1 or 2 is determined by the the potential calculated using the objective function value of the lowest level in the tree under that option, in other words, the potential of 1 is determined using $111, 112, 121, 122$ and the potential of 2 is determined using of $211, 212, 221, 222$. If option 1 has the best potential, the tree is expanded below 1 with two more options for all bottom options: $111, 112, 121, 122$ and the part of the tree starting with a 2 is discarded. For calculating the potential, several strategies can be used, which will be discussed in the next subsection.

When setting the width and depth parameters to specific values, some well known methods can be replicated. An exhaustive search method is creating using a width = ∞ and depth = ∞), while a greedy algorithm is created using width = ∞ and depth = 1. Finally, a naive selection algorithm is created using width = 1 and depth = 1. By changing the parameters for the tree search (width and depth), a trade-off between the search space, solution quality and computing time can be made.

6.1.1. Determining Potential

For calculating the potential of different options in the search tree, two options can be employed. The first strategy is to use the average of the lowest iterations objective function values as the potential of the highest option. In the

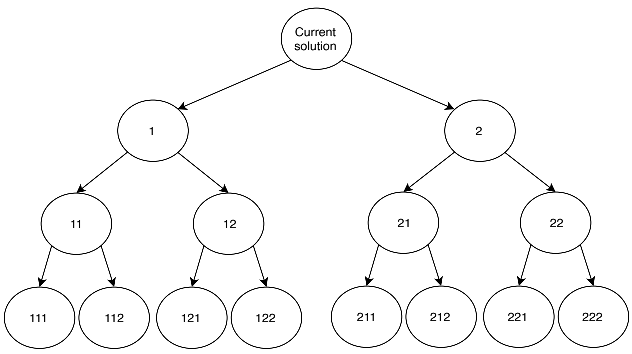

Figure 6.1: Visual example of the local tree search method with width = 2 and depth = 3.

example in Section 6.1 this means the potential of option 1 is determined as the average of the objective function values of 111, 112, 121 and 122. In mathematical form, this can be represented as:

$$P_x = \frac{1}{Y_{max} \cdot Z_{max}} \cdot \left(\sum_{y}^{Y} \sum_{z}^{Z} O_{x,y,z} \right) \tag{6.1}$$

A different strategy is to use the minimum of those same options. In this case, the potential can be determined as:

$$P_x = \min \left(O_{x,y,z} \ \ \forall\, y,z \in Y,Z \right) \tag{6.2}$$

Using the average to determine the potential is quite conservative. The other strategy, to use the minimum, is a more aggressive approach as it potentially reaches better solutions because of this strategy. It, however, also increases the chance to get stuck in a local optimum as the potential is only determined by one option.

In order to compare the two methods, a number of small tree search configurations have been used to solve different scenarios. The objective function is then normalized with respect to the naive selection algorithm's objective function. The resulting values can then be compared between the two strategies. This is done in Figure 6.2 which shows a box plot of the relative performance with respect to the naive selection algorithm of the two methods. From the plot, it can be seen that the minimum-approach is a more aggressive method, even though the mean of the two methods is almost identical, the spread of the minimum-approach is a lot larger (a spread of 0.13 for the mean-approach and a spread of 0.31 for the min-approach). As a stable solution is better suited for comparing different scenarios and strategies, the mean-approach is chosen as the best approach for this research.

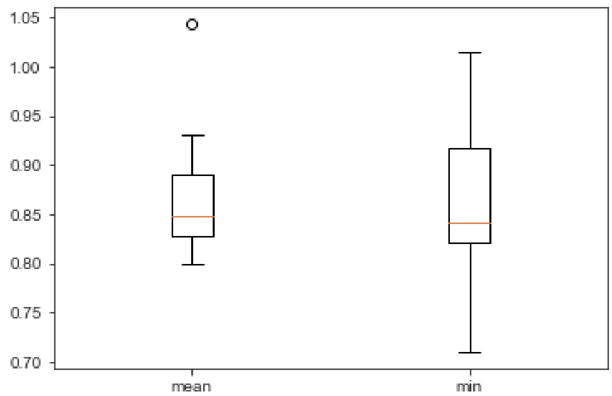

Figure 6.2: Box plot showing the spread of objective function values from 18 different crew plan scenarios (constructed from the reference airline's data) for the average and minimum tree search method.

6.1.2. Sorting Options

As the number of options evaluated in the tree search method can be changed with the width parameter, is has to be determined which of the available options is used in the method. This is done by sorting the available options based on rules which can be defined following expert knowledge of the process and behaviour of the system. Then, in the tree search method, the width parameter determines how many of the first available options are evaluated in the tree.

A number of sorting options have been evaluated by two different variations of the tree search method. By looking at the number of times a certain option is chosen, as can be seen in Figure 6.5, the sorting method where the first couple of options are chosen the most can be seen as the best sorting option.

There are a number of parameters that can easily be used to order the possible transition options. First of all, the date of the transition option can be used. Furthermore, the options can be sorted by the balance value (where higher shortages are solved first) and the minimum balance in that position (again, higher shortages are solved first). Finally, the options can be sorted by the same position salary cost as used in the objective function calculation. Two parameters are used in every algorithm with a primary and secondary sorting parameter. Furthermore,

Table 6.1: Performance of various naive selection algorithms.

Abbreviation	Primary	Secondary	Tree Search (10, 2)	Greedy (∞, 1)	Average
dpb	Date	Position min.	24.94%	19.42%	22.18%
db	Date	Balance	21.67%	19.58%	20.62%
ds	Date	Position weight	24.56%	24.92%	24.74%
pbd	Position min.	Date	21.67%	19.58%	20.62%
bd	Balance	Date	21.67%	19.58%	20.62%
sd	Position weight	Date	21.67%	19.58%	20.62%

the date should always be used in the algorithm as the first shortages are the most important to resolve and can potentially also resolve later shortages. This results in six possible sorting methods that are evaluated on the greedy selection algorithm (width = ∞, depth = 1) and a tree search solution method with a width of 10 and depth of 2.

In Table 6.1, the performance of the different sorting methods is presented. The first two values show the percentage of the iterations for which any of the first three options were chosen for the two chosen solution algorithms (a greedy algorithm and a tree search with a width of 10 and depth of 2). The third value is the average of the first two and is used to select the best naive selection algorithm. As all sorting options with the date of the shortage as the secondary parameter score identical results and also the same as the second sorting option, these last three are disregarded.

For the remaining three sorting options, detailed results for the two solution methods are shown in Figure 6.3 and Figure 6.4. In these figures, the cumulative distribution of option choices have been plotted. This means the plot show how many times that option or a higher ranked option are chosen in the solution method. The main difference for the tree search method is the performance of the $'db'$ option that performs worse compared to the other two. In the greedy solution, the option that uses the position weight performs significantly better than the other two. As this parameter is also used in the objective function equation, sorting the different shortages based on this parameter and thus solve the shortage with the highest salary cost first is the best option.

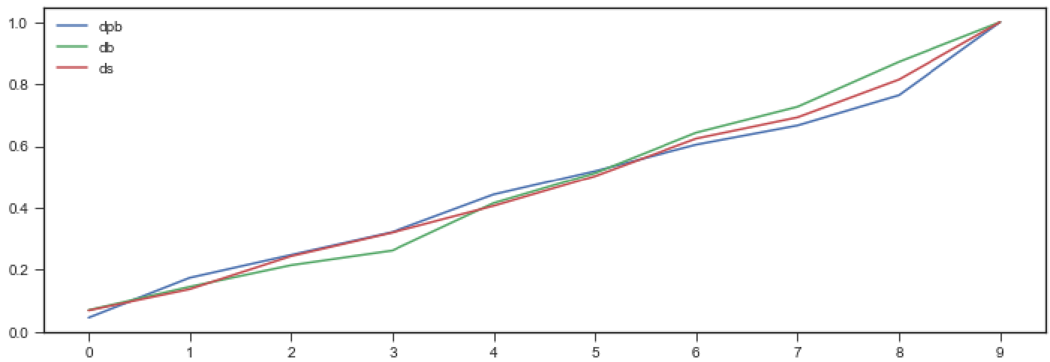

Figure 6.3: Cumulative distribution of option choices in the (10,2) tree search for three of the naive selection algorithms.

6.2. Naive Selection Algorithm

As mentioned before, the parameters of the tree search method can be set in such a way that the method replicates other algorithms. One of these in the naive selection algorithm which assumes that the best option can be determined using prior knowledge. By setting the width and depth of the tree search to 1, the algorithm will assume that the sorting done in the tree search method is accurate enough to blindly select the first option and plan the associated transition.

As in this case only one option is evaluated at each iteration and no potential has to be calculated using later iterations, it is to be expected that the computation time of the naive selection algorithm will be very low.

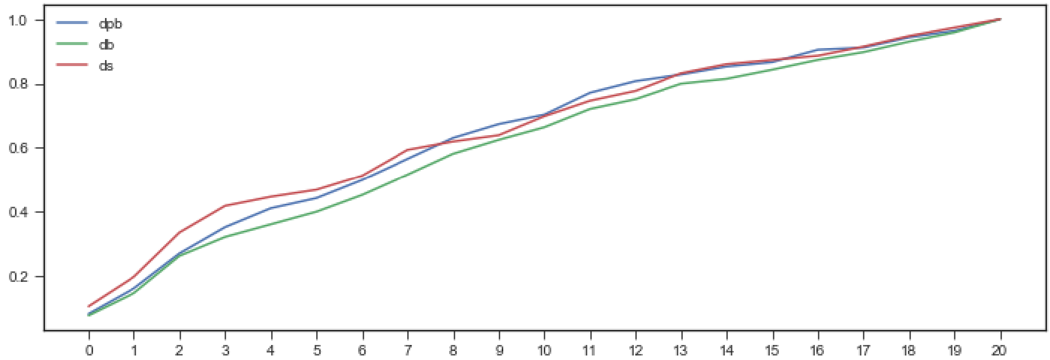

Figure 6.4: Cumulative distribution of option choices in the greedy solution for three of the naive selection algorithms.

6.3. Greedy Algorithm

The greedy algorithm is a heuristic that aims to make the local optimal decision at each iteration. In this specific case, it aims to find and plan the transition that has the greatest direct impact on the objective function value. As mentioned before, it can be replicated by a tree search method with a width of ∞ and a depth of 1. It only looks at the objective function value after one iteration instead of evaluating all options until the end (as an exhaustive search method does) and, therefore, produces solutions many times faster. However, a drawback of a greedy algorithm is that is has a high chance of finding a local optimum as opposed to the global optimum. In a lot of cases, this might not be the best strategy as an option that is the best local choice can end up creating bottlenecks, such as capacity restrictions, in later stages.

In Figure 6.5, the number of times a certain option was chosen in the greedy selection algorithm is presented, where the order is determined by the selected sorting method. Furthermore, Figure 6.6 shows the progression of the objective function value of the greedy algorithm versus the best random solution out of 25 independent runs for the random selection algorithm. It can be seen that the greedy solution makes significantly larger improvements to the objective function in the first stage of the solution process. This can be explained by the fact that the algorithm selects the option with the largest improvement as the optimal transition option. After about a quarter of the number of iterations in the greedy algorithm, the objective function already drops below the best random solution, after which only small improvements to the objective function are made for a while. The algorithm continues to search for a better solution but instead, the objective function only rises until eventually no options are left.

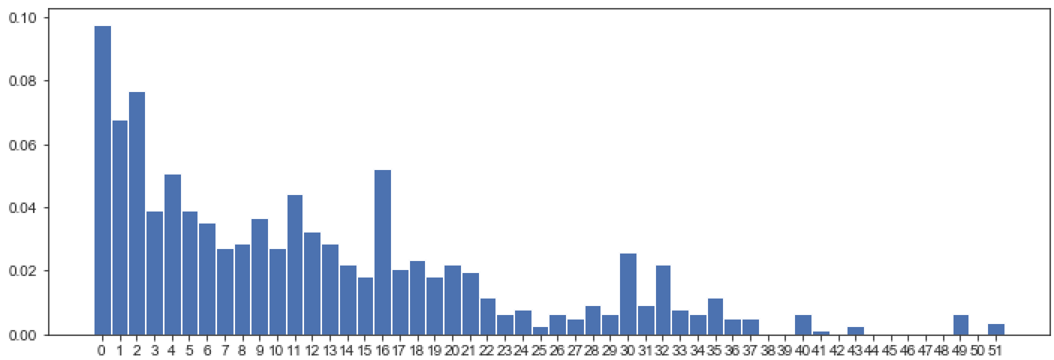

Figure 6.5: Histogram of the chosen options in the greedy algorithm (order based on the selected sorting method).

6.4. Shortest Path Algorithm

The local tree search provides a method in which a search space can be defined using two parameters and the best solution in this search space can be found. However, the design of the tree search makes it impossible to select a

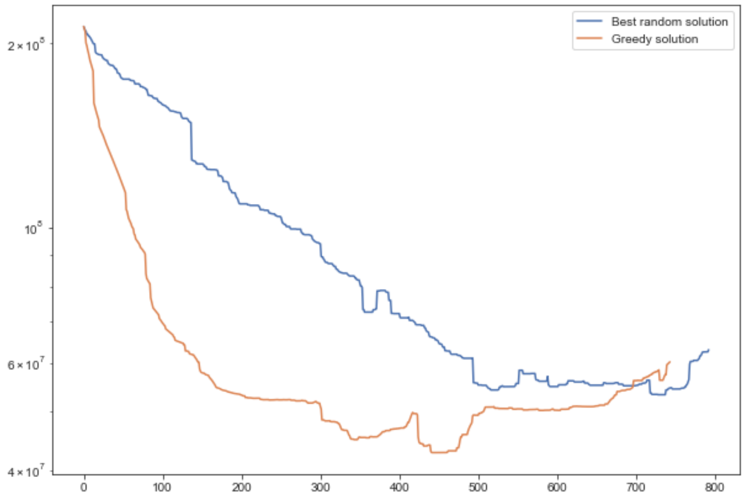

Figure 6.6: Objective function value progression for greedy selection algorithm and best random solution for a single scenario.

different path once a choice for a certain option is made and therefore potentially gets trapped in a local optimum if the depth of the configuration is not large enough. Also, the tree search method always evaluates all options in the tree. With an increasing tree size, this will greatly increase the required computation time eventhough sometimes it might be clear that a certain path is not improving the solution and does not have to be evaluated. A variation of the tree search selection method that aims to prevents this is the shortest path algorithm.

6.4.1. Search Strategy

The shortest path algorithm is an vartiation of the tree search method based on Dijkstra's algorithm (Dijkstra, 1959) for finding the shortest path between two nodes. It picks the unvisited node with the lowest distance to the origin at every iteration, then calculating the distance through that node to all unvisited neighbours and finally updating the neighbour's distance if it is smaller than the current distance. The algorithm can be adjusted to the crew planning by using the objective function value as the distance from the origin and creating neighbours as the available transition options at an iteration.

An example of the algorithm is given in Figures 6.7 - 6.12. From the initial balance, a number of neighbours are created by planning the available transitions and calculating the resulting objective function value (Figure 6.7). The initial node is now visited and turns red. For the two created neighbours, the objective function value is determined and stored and they are marked as unvisited. In Figure 6.8, the unvisited node with the lowest objective function (in this case, option 2) is selected and expanded by planning new transition options, which now results in three unvisited nodes. In the next iteration (Figure 6.9), the best unvisited node (again option 2) is again expanded, creating 4 available options. This process is repeated until no unvisited nodes are available and the best solution is then the node with the lowest objective function value.

6.4.2. Search Restrictions

The computational time of the proposed algorithm will quickly increase when the algorithm starts looking for alternatives when a single path has been solved to completion and is unavailable for further exploration. Therefore, three search restrictions have been implemented to the problem to limit the computational time.

- **Width**
  Similar to the local tree search method, the width is defined as the number of options (or neighbours) to be evaluated from each node. In the visual example of Figures 6.7 - 6.12, the width is equal to 2, since for each node, two options are evaluated.
- **Height**
  Even with the width of the algorithm restricted, the algorithm is able to look back hundreds of iterations to

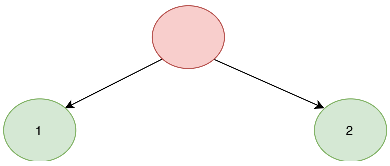

Figure 6.7: Visual representation of the local shortest path algorithm (step 1).

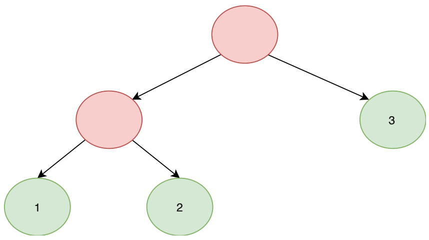

Figure 6.8: Visual representation of the local shortest path algorithm (step 2).

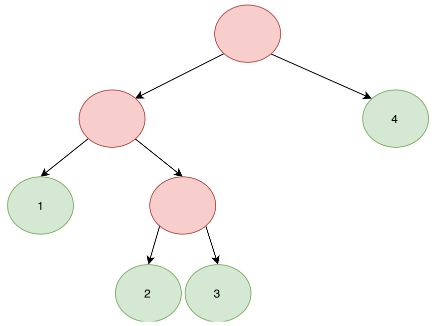

Figure 6.9: Visual representation of the local shortest path algorithm (step 3).

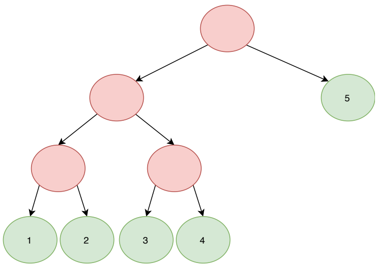

Figure 6.10: Visual representation of the local shortest path algorithm (step 4).

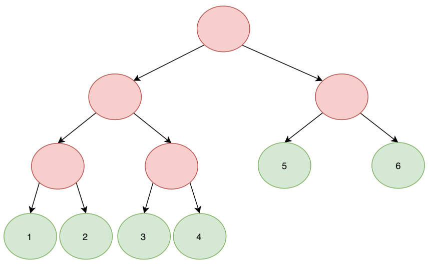

Figure 6.11: Visual representation of the local shortest path algorithm (step 5).

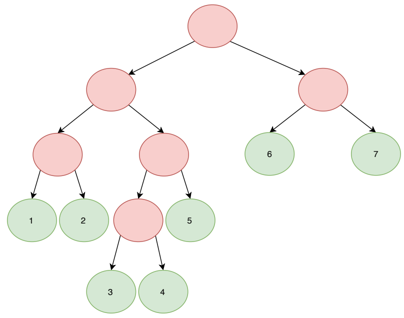

Figure 6.12: Visual representation of the local shortest path algorithm (step 6).

explore all possible transitions and find a better solution. In order to prevent this, the height parameter restricts the number of levels the algorithm is allowed to go back to search for unvisited nodes. In the algorithm, the level of a node is defined as the distance, or number of iterations, from the initial node or origin. This parameter is similar to the depth parameter defined for the tree search method, but because the orientation of the search is the other way around (the tree search method looks down to determine the potential of the option while the shortest path algorithm looks up to find the best, unvisited node), a different name is used. As an example for the restriction, in Figure 6.12, a height parameter of 1 would mean only options 1 to 5 are allowed, while options 6 and 7 are prohibited.

- **Relative distance**
  With the width and height defined, the search space of the algorithm is already a lot smaller. However, the amount of available options can still increase to large numbers which in turn greatly increases the computational time. Therefore, a final restriction is related to the 'family' of the available nodes. This means that the availability of nodes is also determined by nodes they originate from. With a relative distance of 2, only nodes that have the same origin two levels above the lowest level are allowed. In Figure 6.12, this means only options 3, 4 and 5 are allowed since options 1 and 2 do not share the same relative at the second level.

6.5. Selected Configurations

In order to test the performance of the proposed solution methods, different configurations of the defined parameters and search techniques will be tested in the experiments in Chapter 7. First of all, the two presented special cases of the tree search method (the greedy algorithm with a width of ∞ and a depth of 1) and the naive selection algorithm (with a width and depth of 1) are selected.

For the standard tree search method, which searches for the best transition option in a tree where the size is defined by the width and depth parameter, four configurations will be tested. First of all a baseline configuration (width = 2, depth = 2), secondly a configuration with a specific focus on width (width = 6, depth = 2), the third is a configuration with a specific focus on depth (width = 2, depth = 5) and finally a configuration focussing on both (width = 3, depth = 3). These configurations were selected as a variety of focus point for the tree search will give more information on the best strategy to use for solving the problem. Furthermore, the configurations where selected to have a computation time within the range set by the naive selection and greedy algorithm's computation time.

Similar to the configurations selected for the standard tree search method, a selection has been made for the shortest path algorithm by looking at the computational time and constructing configurations with different focus points. First of all, the relative distance will be fixed at the height parameter plus one. This means the options have to come from the relative node at the first unavailable level. The chosen configurations include a base configuration (width = 2, height = 2), a configuration with a specific focus on width (width = 8, height = 2), a configuration with a specific focus on height (width = 2, height = 6) and a configuration focussing on both (width = 3, height = 4). These configuration will be evaluated in the experiments in Chapter 7 alongside the tree search configuration, as well as the naive and greedy algorithms.

7

Experiments

This chapter presents the experiments performed in order to assess the performance of the different selection algorithms designed in Chapter 6. In Section 7.1 the experiment set up is presented, afterwards the results of the experiments are discussed in Section 7.2 and finally, the results are validated in Section 7.4.

7.1. Experiment Design

In this first section, the design of the experiments is discussed. This includes the method used to analyse and assess the performance of the various models, the scenarios constructed for the experiments and the models tested in the experiments.

The goal of the experiments is to test the performance of the ten selected selection algorithm configurations. To do this, the configurations are used to solve ten different scenarios and the parameters of the best found solution are stored for every simulation (100 simulations in total). From these solutions, the performance of the various configurations can be assessed and recommendations can be made regarding the best configuration for the cockpit crew transition planning model.

In order to compare the objective function values over the different scenarios and models and to be able to draw a conclusion from the performance of the different methods, the objective function is scaled using min-max normalisation. This means that the objective function is scaled to a value in the range from zero to one, where 1 is the highest objective function and 0 the lowest objective function for that scenario. This is done through the following equation:

$$x' = \frac{x - x_{min}}{x_{max} - x_{min}} \tag{7.1}$$

7.1.1. Scenarios

As discussed before, the development of the models is performed for a planning window of 1 year and a modelling frequency of 1 month. This means the scenarios will consist of 12 dates on which transitions can be planned. By selecting different starting dates and constructing datasets from the available data by disregarding transitions after the chosen date and calculating the supply and demand, several scenarios can be constructed.

Table 7.1: Summary of the different scenarios used in the experiments.

	Start date	Shortages	Lowest balance	Initial objective function
1	01/01/2018	66	-20.2	$1.448 \cdot 10^9$
2	01/04/2018	65	-37.7	$1.060 \cdot 10^9$
3	01/07/2018	60	-38.6	$2.390 \cdot 10^9$
4	01/10/2018	61	-33.3	$3.555 \cdot 10^9$
5	01/01/2019	61	-65.4	$3.762 \cdot 10^9$
6	01/07/2019	46	-66.8	$6.894 \cdot 10^9$
7	01/01/2020	40	-51.2	$9.812 \cdot 10^9$
8	01/07/2020	40	-44.9	$4.222 \cdot 10^9$
9	01/01/2021	46	-29.8	$3.022 \cdot 10^9$
10	01/07/2021	32	-38.4	$1.878 \cdot 10^9$

For every scenario, simulations have to be performed for all ten selection algorithm configurations. As some of these configurations result in high computation times (up to 1 hour), it is not possible to test the configurations on a very high number of scenarios. To limit to total time required for the experiments, ten scenarios have been constructed. In Table 7.1, a summary of the ten different scenarios is given. In this table, the amount of shortages are all combinations of date and position that have a balance below 0, out of a maximum of 12 (dates) $\cdot$15 (positions) = 180. Furthermore, the lowest balance value is given in FTE and the objective function value is calculated using the method as presented in Section 5.5.

7.1.2. Models

The selection algorithms that have been presented in Chapter 6 have to be tested in order to select the best method for the cockpit crew transition planning problem. The greedy and naive selection algorithms (as found in Table 7.2) present two of the extreme points of the tree search method with a width of ∞ and 1, respectively, and a depth of 1. For the tree search and shortest path methods, however, several different models have been developed by changing the two algorithm parameters. By doing this, the search space of the algorithms is changed which changes both the computational time and the solution quality.

Table 7.2: Summary of the different configurations tested in the presented experiments.

ID	Method	Width	Depth	Height
G	Tree search	∞	1	-
N	Tree search	1	1	-
T1	Tree search	2	2	-
T2	Tree search	6	2	-
T3	Tree search	3	3	-
T4	Tree search	2	5	-
P1	Shortest path	2	-	2
P2	Shortest path	10	-	2
P3	Shortest path	3	-	4
P4	Shortest path	2	-	7

For the tree search method, three models have been selected that focus on width (model $T2$ in Table 7.2), depth (model $T4$) and a combination of both (model $T3$) as well as a baseline model (model $T1$). For the shortest path method, the same strategy was used. A baseline model has been selected (model $P1$) as well as three models that focus on width (model $P2$), height (model $P4$) and a combination of both (model $P3$).

7.2. Results

In this section, the results of the presented experiments are presented and discussed. In Appendix B, the detailed results for all methods and scenarios are presented for further reference.

7.2.1. Tree Search

For the proposed tree search method, six standard configurations have been selected for the experiments. The first two are special cases of the tree search method in the form of the greedy algorithm (width = ∞, depth = 1) and the naive selection algroithm (width = depth = 1). The other four configurations where selected to represent a broad range of focus points for the tree search method while keeping the computation time between those of the naive selection and greedy algorithms.

Greedy algorithm
The first configuration that was tested is the greedy selection algorithm that follows a steepest descent approach and selects the transition option with the best local improvement of the objective function value at each iteration. In Figure 7.1, the progression of the best objective function value of the greedy algorithm for the ten different scenarios is plotted against the number of iterations. In the plot, the objective function value is normalised with respect to the initial objective function value. The objective function value drops rapidly in the early stages of the solution method. This can be explained by the usage of the steepest descent method in which the option with the highest improvement is selected. However, the progression quickly stagnates and eventually, the model is terminated since no valid options are left. From the end of the plots, it can be seen that for all scenarios, the best objective

function value found by the greedy selection algorithm lies between 0.8% and 8.8% of the initial objective function. As for the computation time of the greedy algorithm, the time required to solve the different scenarios ranges from around 30 minutes up to more than 1.5 hours, as can be seen in Figure 7.3. It can also be seen that the computation time for different scenarios has a lot of variation, which can be attributed to the number of shortages in the system and consequently the number of iterations required to solve the problem.

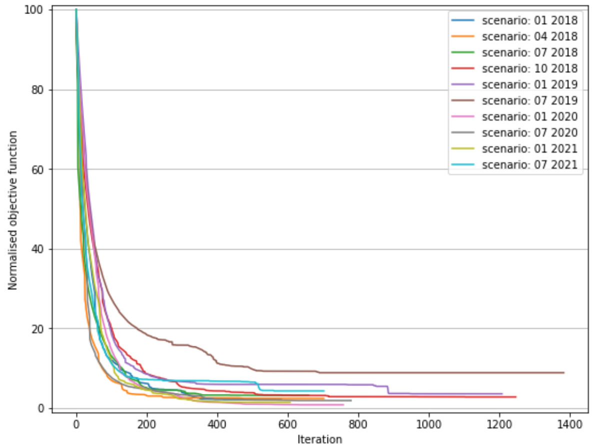

Figure 7.1: Objective function progression of the ten different scenarios for the greedy algorithm.

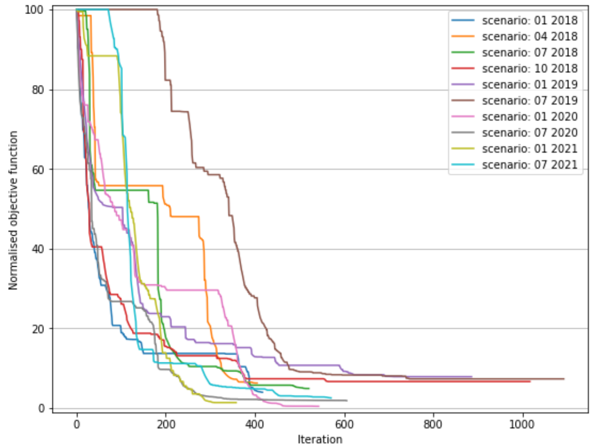

Figure 7.2: Objective function progression of the ten different scenarios for the naive selection algorithm.

Naive selection algorithm

The naive selection algorithm is created as a tree search method with a width and depth of 1. This means the first possible transition option is selected without evaluating the objective function value improvement of this option. The computation time for both the naive selection and greedy algorithms are plotted for the different scenarios in Figure 7.3 in order to compare the computation times of both configurations. It can be seen that the naive selection algorithm presents a large improvement in computation time: the computation time of the naive selection algorithm ranges from 2% to 6% of the greedy algorithm's computation time. For some scenarios, the computation time of the naive selection algorithm was below 1 minute, which provides the opportunity to perform many more simulations with this configuration. Looking at the objective function value progression of the naive selection algorithm in Figure 7.2, it can be seen that the progression is initially smaller than the greedy algorithm. However, the progression continues for a longer time and the end results do not differ much from the greedy results (ranging from 0.5% to 7.9% of the initial objective function value and an average of 4.3% versus 3.0% for the greedy configuration). Also, the number of iterations of both methods are in the same order of magnitude.

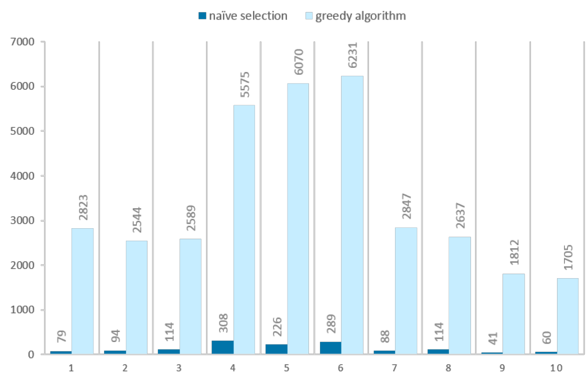

Figure 7.3: Comparison between computation time for the ten different scenarios for the naive selection and greedy algorithms.

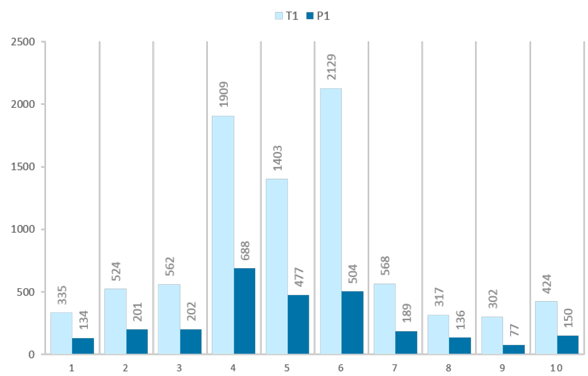

Figure 7.4: Comparison between computation time for the ten different scenarios for configurations T1 and P1.

Other configurations

The four remaining configurations represent different focus points and thus differnt search space sizes and shapes, created by varying the width and depth parameters of the algorithm. For model T3 (width = 3, depth = 3), the progression of the objective function is plotted in Figure 7.6. The plot shows a progression similar to that of the naive selection algorithm in Figure 7.2, which is to be expected as the search space (with 3 options per iteration) is not much larger than the naive selection's width when comparing it to the greedy algorithm's width. The search depth of the tree decreases the speed of the objective function decrease, as the model does not select the best option locally (and thus the steepest descent). Instead, it chooses the option that performs best on average three iterations later.

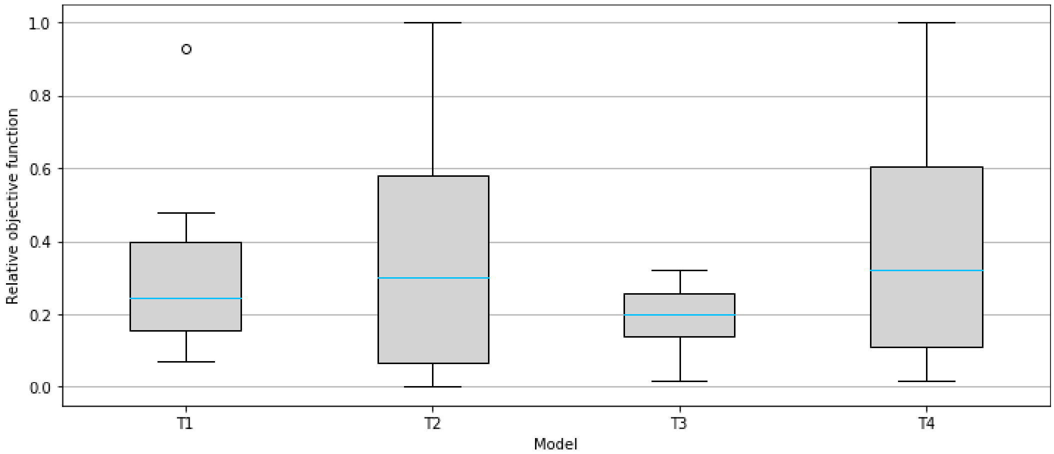

Figure 7.5: Box plot showing the spread of the objective function (normalized by the initial objective function) for the tree search models.

A comparison of the scaled objective functions for the four different tree search models is presented in Figure 7.5. It can be seen that the best performing model is the model focussing on a combination of width and depth (model T3) with a median scaled objective function of 0.2. Meanwhile, the worst performing model is model T4 (focussing on depth). It should be noted however that the median results of the four configurations do not differ much (ranging from 0.2 for the best configuration to 0.32 for the worst configuration). This means that all configurations are valid solution methods for the cockpit crew transition planning problem.

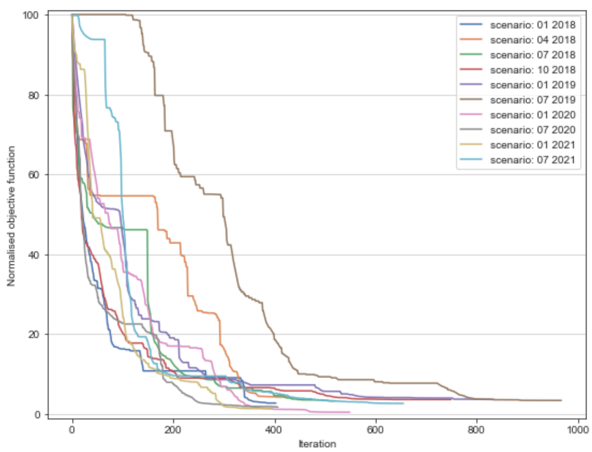

Figure 7.6: Objective function progression of the ten different scenarios for the third tree search configuration.

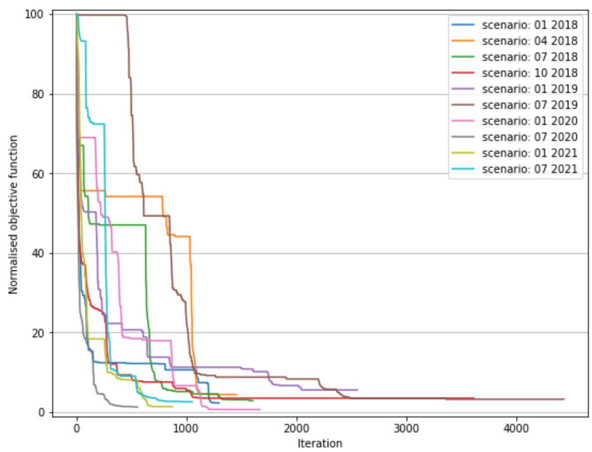

Figure 7.7: Objective function progression of the ten different scenarios for the third shortest path configuration.

7.2.2. Shortest Path Algorithm

The local shortest path algorithm only evaluates different paths in a restricted tree (similar to the tree search method) when the objective function increases instead of decreases. Therefore, the model finds improvements in the objective function value quickly when the chosen options decrease the objective function. However, when a local or global optimum is reached, the model starts to evaluate other paths in order to find better paths in the tree. The algorithm only continues when such an improving path is found, or the full tree is evaluated. This can also be seen in Figure 7.7 where the progression of the best objective function value of model P3 (width = 3, height = 4) is plotted. The progression stagnates a lot and compared to the tree search model, more iterations are needed. However, as each iteration requires a lot less options to evaluate (for model T3 and P3, 27 and 3 options are calculated and evaluated at each iteration, respectively) the shortest path algorithm is able to determine the best solution quicker than the tree search method. This can best be seen in the comparison of the computation time between models T1 and P1 in Figure 7.4.

The two plotted models utilise the same size tree. However, the tree search method always evaluates all options in the tree (with 4 options in the lowest level of the tree), while the shortest path algorithm selects the best one at each iteration without looking at the rest of the tree unless these options do not provide a decrease in objective function. This means, at every iteration, the tree search method evaluates four transition options, while the shortest path algorithm only evaluates two options.

Finally, in Figure 7.8, the scaled objective function values of the four different configuration of the shortest path method tested are presented. It can be seen that the median scaled objective function of the four models do not differ much and range from 0.18 to 0.26. In this case, the worst performing model is model P2, the configuration focussing on width. Furthermore, the best performing model is again the configuration focussing on a combination of width and height (model P3).

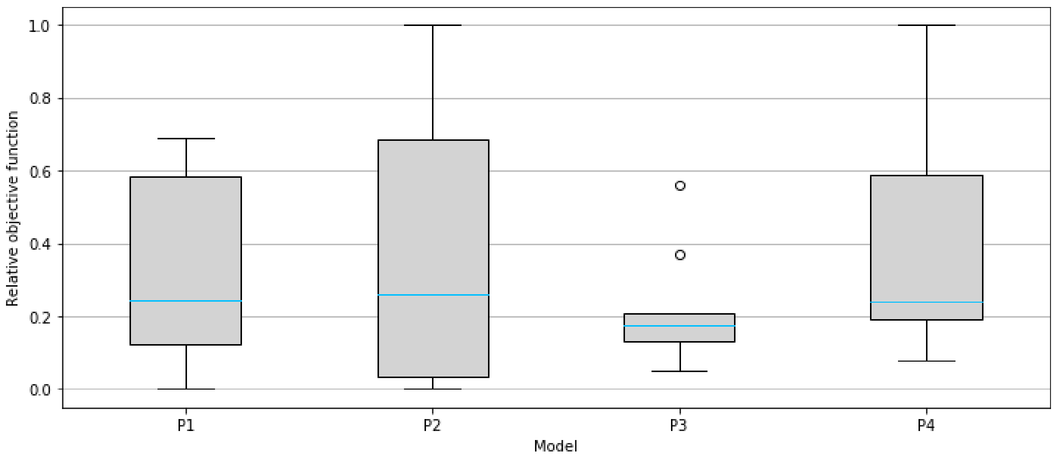

Figure 7.8: Box plot showing the spread of the objective function (normalized by the initial objective function) for the shortest path algorithm models.

7.3. Discussion of Results

In this section, the previously presented results are discussed. From the discussion with regards to the objective function and computation time of the models, a conclusion can be drawn regarding what model is best for the cockpit crew transition planning problem.

7.3.1. Objective function value

As all results are in the same range of values because of the min-max normalisation, the values can be easily compared. In Figure 7.9, the median of the scaled results for the different configurations is plotted against the average computation time. The error bars shown in the plot represent the $25th$ and $75th$ percentile of the results, which can be used as a measure of the variance of the results, or in other words, the solution stability of the configurations.

From the plot it can be concluded that the greedy algorithm obtains the best median solution quality. However, the results show a big variance, with an inter-quartile range (IQR) of 0.75. This can be explained by the fact that

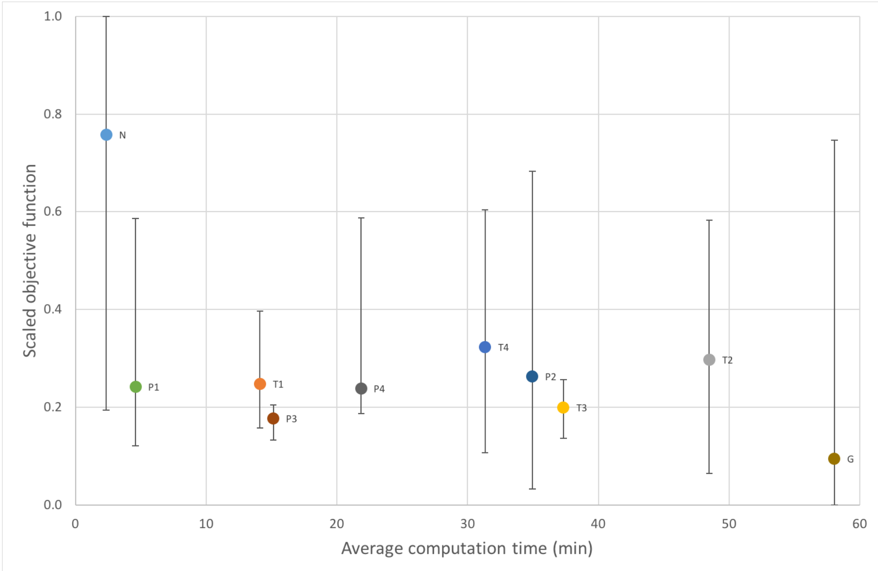

Figure 7.9: Scatter plot showing the average computation time versus median normalised objective function for the tested models.

the greedy algorithm only looks at the potential of options locally. This can result in bad choices being made even though the option has the best local performance. The next best performing models are T3 and P3, the tree search and shortest path models that focus on a combination of width and depth/height. These models also show a smaller variance when compared to all other models (an IQR of 0.12 and 0.07, respectively). When comparing the tree search and shortest path models focussing on width (T2 and P2), with the models focussing on depth/height (T2 and P4), no conclusion can be drawn regarding the benefit of width over depth or vice versa. It can be concluded that a combination of width and depth produces the best results. For both the tree search and shortest path algorithms, the configuration focusing on a balance between width and depth performs best in terms of solution quality and solution stability. The naive selection configuration shows the worst median solution quality (0.76) and a similar stability as the greedy algorithm with an IQR of 0.81.

7.3.2. Computation Time

The distribution of computation times for the different configurations is shown in Figure 7.10. The naive selection algorithm presents the shortest computation time while the greedy algorithm has the longest computation times (median value of 104 against 2730 seconds, respectively). The computation times for the other tree search and shortest path configurations between these two. Also, all shortest path configurations yield a smaller computation time when compared to their equivalent tree search models, eventhough most of these configurations have a larger search space compared to their tree search equivalent. Only configurations P2 and T2 have the same (maximum) search space, but the difference between computation time is also large (median values of 195 versus 543 seconds).

It can be seen that the naive selection has the most stable computation time of all configurations. Meanwhile, the greedy algorithm, as well as the shortest path configurations, show the largest variance in computation time. For the shortest path configurations, this can be explained by the fact that the amount of options to be evaluated is highly dependent on the scenario. For the tree search method, however, approximately the same amount of options are evaluated for all scenarios and the computation time is mainly determined by the number of iterations. For the greedy algorithm, the amount of options that is evaluated depends on the number of shortages left in the balance, which again means that the computation time is highly dependent on the scenario.

7.3.3. Comparison

In Figure 7.11 a radar chart is shown for four of the tested configurations: the naive selection algorithm (N, width = 1, depth = 1), greedy algorithm (G, width = ∞, depth = 1), configuration P1 (width = 2, height = 2) and configuration

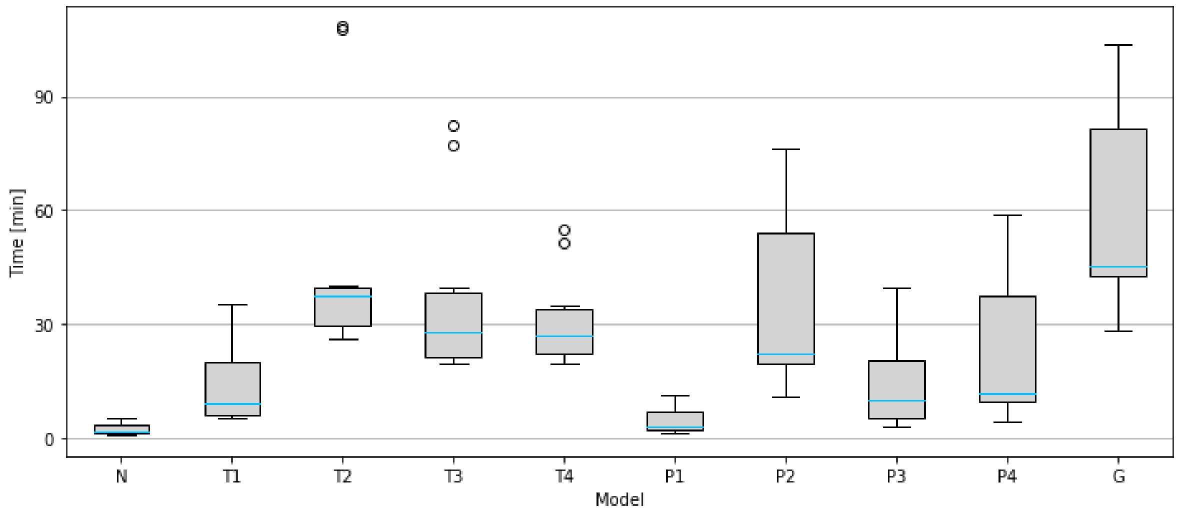

Figure 7.10: Box plot showing the spread of computation time for the tested models.

P3 (width = 3, height = 4). In the chart, the median scaled objective function value and inter-quartile range of the configuration's results are plotted together with the average computation time in hours on three different axes. From the shape and the size of the radar, the overall performance of a configuration can easily be read. From this chart and the previously presented results, the different configurations are compared and recommendations are made about which configuration should be used in different cases.

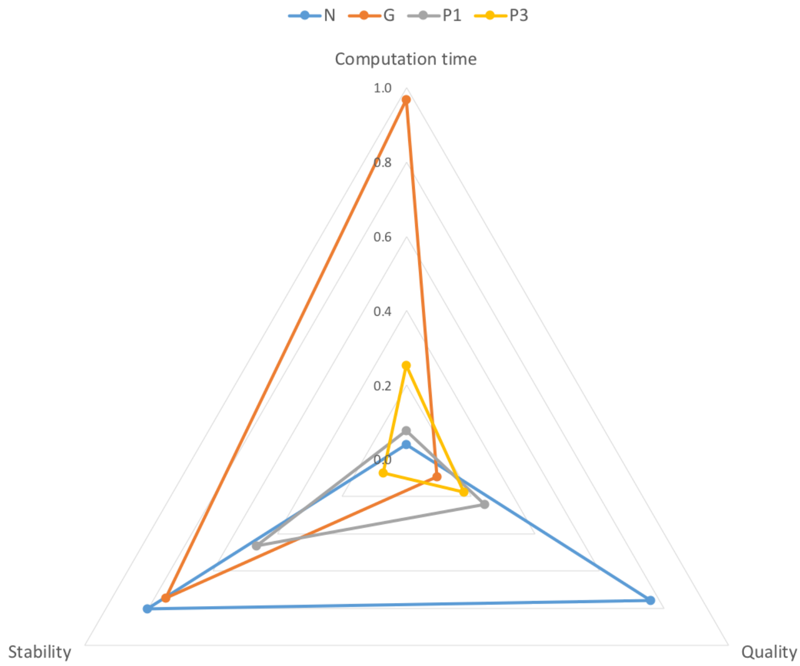

Figure 7.11: Radar plot showing the computation time (in hours), solution quality and solution stability of four of the tested configurations.

Best configuration

From the presented results, it can be concluded that the best configuration that has been tested is configuration P3. This configuration ranks fourth in terms of computation time (only behind the naive selection algorithm and the tree search and shortest path baseline models). Furthermore, the median result ranks second after the greedy algorithm and the results show the most stable performance, with an IQR of 0.07. This configuration presents the smallest total area in Figure 7.11, from which it can immediately be seen that this configuration has the best overall performance of the tested configurations.

Solution quality
In some cases, however, the solution quality of the results is more important than the computation time. In this case, the best configuration is the greedy algorithm. This model has the best median solution quality. However, the variance of the solutions is high (IQR = 0.75) as can be seen in Figure 7.11. Therefore, instead of using the greedy algorithm, it is recommendable to test a larger version of the tree search or shortest path algorithms to see if these produce better and more stable results. As it was found in the previous section that for both the tree search and shortest path algorithm, the best configurations were those focussing on a combination of width and depth, this should be one of the requirements. In other words, the width and depth or height of the configuration should be equal or at least almost equal.

Computation time
When the computation time of the configuration is the most important parameter, two configurations can be considered. First of all, the naive selection algorithm is the fastest configuration with an average computation time of 2.4 minutes. The solution quality and stability, however, are the worst of all tested configurations (median scaled objective function value of 0.76 and IQR of 0.81). Configuration P1, on the other hand, has a higher average computation time of 4.6 minutes, but this increase in computation time is accompanied by a large increase in solution quality and stability. The median scaled objective function value of this configuration is equal to 0.24 and the interquartile range (IQR) is equal to 0.47. Therefore, configuration P1 can be seen as the optimal configuration when a low computation time is required but the solution quality and stability are still important.

7.4. Validation

In this section, the validation process that is performed on the model and its results is discussed. This validation process aims to answer the question of whether the obtained results are sufficiently accurate to be used by an airline to analyse different scenarios. To do this, two different strategies have been used:

- Face validation with a number of people working on the cockpit crew transition planning problem on a day-to-day basis. This provides insight into the choices made by the developed model and might expose any factors that limit the validity. In these discussions, the assumptions made throughout the research as well as the implications of these assumptions to the results have been discussed. All relevant labour agreements have been implemented into the model as accurately as possible and are deemed valid within the scope of this research.

- Quantitative validation with respect to the number of planned transitions, balance values and objective function value. This comparison shows whether the obtained results fall within a reasonable range of possible values. In Figure 7.12, the balance of one position in the current scenario of the model application problem is comparde to the data in the system currently used. It can be seen that the final balance quite accurately follows the trends in the validation results. Especially in the first months, the difference is negligible. In later months, the results present higher balance values compared to the validation data. When looking at the number of transitions planned to one aircraft type in Figure 7.13, it can be seen that the model is able to plan more transitions compared to the validation data. To compensate for this, more transitions are planned in the validation data in the first months. Looking at the last 6 months, however, more transitions are planned in the validation data compared to the model. This can be explained by the fact the validation data plans further into the future and probably has to plan these shortages already for future shortages, while the model only evaluates the defined planning window.

From the results of the validation process, it can be concluded that the results of the model are valid within the project scope. The goal of the research is to create a decision support tool that helps airlines analyse different planning scenarios and strategies. To do this, a trade-off has to be made between computation time, solution quality and solution stability. From the presented results and validation of these results, it can be concluded that the quality of the solutions is accurate enough to perform these analyses. Furthermore, a high solution stability and low computation time can be reached by choosing the appropriate model configuration.

Although the results are deemed valid, the definition of the project within the research scope introduces some limitations to the validity of the model within the complete crew planning process:

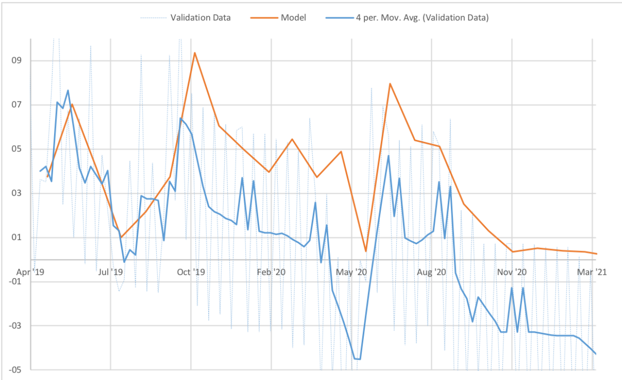

Figure 7.12: Validation data comparing the balance of the CP B position between the model's results and the actual data.

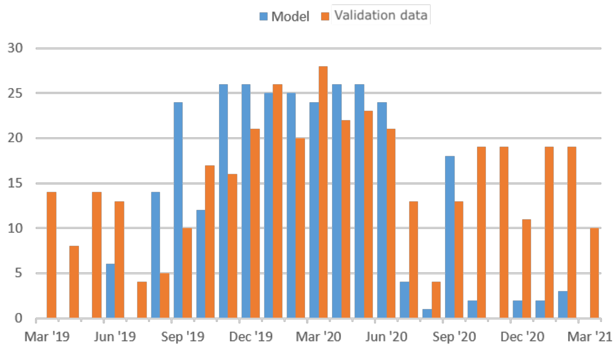

Figure 7.13: Validation data comparing the amount of transitions to aircraft type D between the model's results and the actual data.

- The scope of the research was defined such that only one method for resolving shortages was used. Alternatively, budgets can be shifted to change the demand and thus the balance. However, since this was not modelled, more transitions are planned in the model compared with reality. In order to create a more accurate model, it could be valuable to add these additional methods in future research into the cockpit crew transition planning problem.
- A number of assumptions have been made that limit the validity of the results. As seen in the quantitative validation, the model is able to plan more transitions in a number of months compared to the validation data. It is thought that this is due to the supply of pilots changing through planned transitions. For transitions, instructors are required that normally are available for the nett demand, but in case of planned transitions are required to train the pilots. Because of this, transitions are sometimes not planned by the airline as the decrease in supply is larger than the shortages. Since the model does not take this change in supply into account, the model does not find any objections to planning the transitions and is, therefore, able to plan more transitions compared to the validation data. Because of this, the balance as shown in Figure 7.12 is higher for the model compared to the validation data and more transitions are planned. Therefore, in future research, more accurate modelling of the usage of instructors could prove to further increase the accuracy of the model.

To summarise, the model provides results that are accurate enough to be used to analyse different planning scenarios and strategies. Furthermore, the model can be configured to provide solutions with a high stability in a low computation time. However, some recommendations can be made in order to further increase the accuracy of the results and thereby represent the actual data even better.

8
Model Application

In this chapter, an example of the application of the developed model will be presented. In Section 8.1 a the case that has been studied is presented. Then, in Section 8.2, the results of the current and alternative scenario simulations are presented and an analysis of the differences between the two scenarios is provided.

8.1. Planning Case

A case study has been developed together with the reference airline. In this case study, the effects of a number of demand changes for different aircraft types are analysed. These changes form the alternative scenario of this case study. The current scenario represents the system without the proposed changes.

- A decrease of flight for aircraft type A as can be seen in Figure 8.1.
- A shift in demand from aircraft type B to aircraft types C and D. This change is presented in Figure 8.3.
- A change in demand (both an increase and decreas in different periods) for aircraft type E, as shown in Figure 8.2.

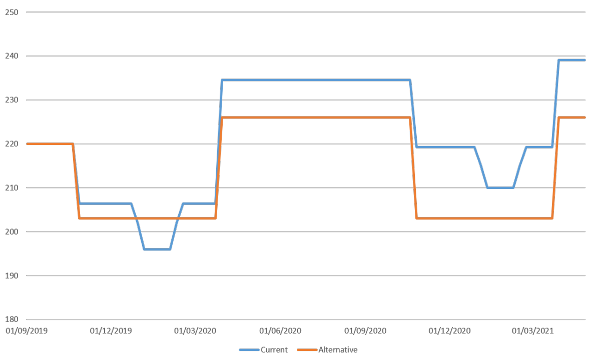

Figure 8.1: Visual representation of the change of demand for the captain position on aircraft type A.

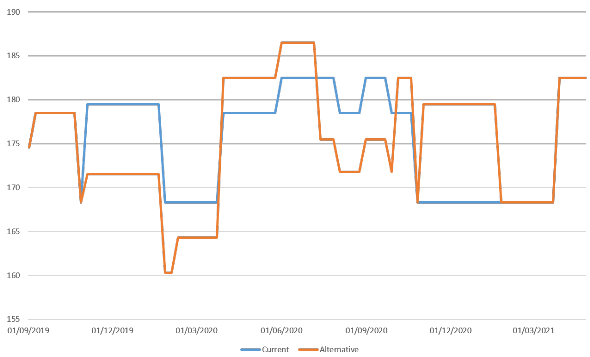

Figure 8.2: Visual representation of the change of demand for the captain position on aircraft type E.

A planning window of 24 months is used, with a modelling frequency of 1 month. The start date is chosen a number of months before the first change in order to allow the model to plan transitions in earlier months if necessary. The planning window of 2 years is chosen to be able to fully analyse all demand changes. For both the current and alternative scenarios, the problem is solved using configuration P3 with a width of 3 and height of 4 as this configuration was found the best performing configuration in the experiments of Chapter 7.

8.2. Results and Analysis

In this section, the results of the simulations for the two scenarios are presented. Afterwards, in the next section, the results are analysed and a conclusion from the results can be drawn.

8.2.1. Current Scenario

The current scenario presents the state of the supply and demand without the changes presented in Section 8.1. As can be seen in Table 8.1, the balance constructed from this supply and demand features 138 shortages with

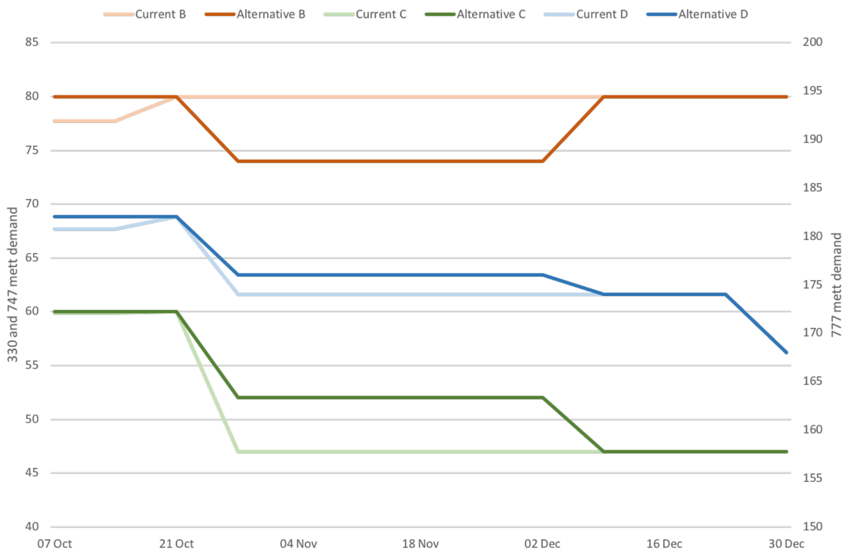

Figure 8.3: Visual representation of the change of demand for the captain positions on aircraft types B, C and D.

Table 8.1: Initial state of the current scenario.

Parameter	Value
Objective function	$6.369 \cdot 10^{10}$
Shortages	138
Minimum balance	-85.3

Table 8.2: Initial state of the alternative scenario.

Parameter	Value
Objective function	$7.120 \cdot 10^{10}$
Shortages	151
Minimum balance	-86.3

the largest shortage of -85.3. By solving this scenario with the developed model and comparing the results to the alternative scenario, the implications of the proposed demand changes can be analysed and recommendations regarding these changes can be made.

From the results of the solution for this scenario in Table 8.3, it can be seen that from the 138 shortages initially present, only 5 remain after solving the problem. This results in an objective function of $5.977 \cdot 10^7$ or 0.09% of the initial objective function.

When looking at the utilised transition capacity in Figure 8.4 and Table 8.5 it shows that in the first year in the planning window, most of the transition capacity is filled, while for the second year, the maximum capacity is utilised in only a few months.

	Mar	Apr	May	Jun	Jul	Aug	Sep	Oct	Nov	Dec	Jan	Feb	Mar	Capacity
A	0,0	1,0	1,0	1,0	1,0	0,9	0,7	0,4	1,0	1,0	1,0	1,0	1,0	1
B	0,8	1,0	0,1	1,0	1,0	1,0	1,0	1,0	1,0	1,0	1,0	1,0	1,0	1
C	1,0	1,0	1,0	1,0	0,5	0,8	0,4	0,9	0,9	0,9	1,0	1,0	0,0	1
D	0,0	0,8	1,0	1,0	1,0	1,0	1,0	1,0	1,0	1,0	1,0	1,0	1,0	1
E	1,0	1,0	1,0	1,0	1,0	1,0	1,0	1,0	1,0	0,2	0,0	0,1	0,0	1

	Apr	May	Jun	Jul	Aug	Sep	Oct	Nov	Dec	Jan	Feb	Mar	Capacity
A	1,0	0,5	0,0	0,0	0,4	0,0	0,0	0,0	0,0	0,0	0,0	0,0	1
B	0,8	0,0	0,0	0,0	0,0	0,1	1,0	0,1	0,1	0,1	0,0	0,0	1
C	0,0	0,0	0,0	0,0	0,0	0,0	0,0	0,0	0,0	0,0	0,0	0,0	1
D	1,0	1,0	1,0	0,2	0,1	0,8	0,1	0,0	0,1	0,1	0,1	0,0	1
E	0,0	0,1	0,2	0,0	0,0	0,0	0,0	0,0	0,0	0,0	0,0	0,0	1

Figure 8.4: Utilised transition capacity per aircraft type and month for the current scenario.

8.2.2. Alternative Scenario

Following the changes in demand for the alternative scenario, the balance for this scenario is also changed, resulting in a different initial objective function and balance. In Table 8.2, these changes can be seen. Because of the

Table 8.3: Solution parameters of the current scenario.

Parameter	Value
Objective function	$5.977 \cdot 10^7$
Shortages	5
Minimum balance	-4.6
Transitions	828
Recruits	273

Table 8.4: Solution parameters of the alternative scenario.

Parameter	Value
Objective function	$5.351 \cdot 10^7$
Shortages	5
Minimum balance	-4.7
Transitions	825
Recruits	294

Table 8.5: Total utilised transition capacity per aircraft type for the first and second year in the planning window for both scenarios.

	Current		Alternative	
Aircraft type	year 1	year 2	year 1	year 2
A	83%	23%	63%	32%
B	94%	28%	68%	55%
C	90%	0%	83%	1%
D	91%	46%	86%	54%
E	80%	2%	77%	23%

changed demand, the initial objective function has risen by 11.8% and the amount of shortages has risen from 138 to 151. The lowest balance value has remained almost constant; from -85.3 to -86.3.

Looking at the results found for the alternative scenario in Table 8.4, the final objective function of $5.351 \cdot 10^7$ is around 10% lower compared to the current scenario. The same amount of shortages is present in the final solution and the lowest balance value is almost identical. Also, more recruits are hired in the alternative scenario. The lower objective function can, therefore, be explained by the alternative scenario having lower shortages in the final solution. The addition of more pilots is not represented completely in the objective function as these added pilots will stay in the airline for a long time. This represents a large salary cost of which only a small part falls within the planning window and an even smaller part represents surplusses.

When looking at the utilised transition capacity per month in Figure 8.5 and the total utilised capacity in Table 8.5, it can be seen that less transition capacity is used in the first year of the planning window in the alternative scenario compared to the current scenario. For the second year, however, more capacity is used in the alternative scenario. Since the total number of transitions for both scenarios is almost equal, the changes in transition capacity utilisation can be attributed to a better distribution of the transitions over the two years.

	Mar	Apr	May	Jun	Jul	Aug	Sep	Oct	Nov	Dec	Jan	Feb	Mar	Capacity
A	0,5	1,0	1,0	1,0	1,0	0,4	0,0	0,0	0,0	0,8	1,0	1,0	1,0	1
B	1,0	1,0	0,6	1,0	1,0	0,0	0,0	0,0	0,4	1,0	1,0	1,0	1,0	1
C	1,0	1,0	1,0	0,9	0,4	0,8	0,9	0,1	0,4	1,0	1,0	1,0	0,0	1
D	0,5	1,0	0,8	0,9	0,4	0,6	1,0	1,0	1,0	1,0	1,0	1,0	1,0	1
E	0,0	0,3	1,0	1,0	1,0	1,0	1,0	1,0	1,0	0,1	0,7	0,9	1,0	1

	Apr	May	Jun	Jul	Aug	Sep	Oct	Nov	Dec	Jan	Feb	Mar	Capacity
A	1,0	0,5	0,6	0,1	0,7	0,0	0,0	0,0	0,0	0,0	0,0	0,0	1
B	1,0	0,2	0,1	0,0	0,9	1,0	1,0	1,0	0,0	0,1	0,0	0,0	1
C	0,0	0,1	0,0	0,0	0,0	0,0	0,0	0,0	0,0	0,0	0,0	0,0	1
D	1,0	1,0	1,0	0,9	0,3	1,0	0,0	0,0	0,1	0,1	0,1	0,0	1
E	0,6	1,0	0,1	0,0	0,0	0,0	0,0	0,0	0,0	0,0	0,0	0,0	1

Figure 8.5: Utilised transition capacity per aircraft type and month for the alternative scenario.

From the presented results, a number of conclusions regarding the demand changes can be made. Even though the number of recruits rises, this should not generate higher costs in the future as it is expected that the total demand will not decrease in the near future, in which case additional recruits are necessary to account for retirements. It appears that transitions can be spread out more in the alternative scenario, which is a benefit for the airline as capacity is available in a higher number of months in the first year of the planning window which can be used if necessary. The objective function value of the alternative scenario solution shows a 10% decrease compared to the current scenario. This shows the overall crew plan is improved with the changes in demand in this case study.

9 Sensitivity Analysis

This chapter presents the results of the sensitivity analysis that has been performed on some of the parameters used in the planning model in order to solve the transition planning problem. The parameters that have been tested in this analysis are the objective function parameter (Section 9.2), planning model frequency (Section 9.3) and tabu search list length (Section 9.4). Furthermore, a sensitivity analysis into the scalability of the model is presented in Section 9.1. The aim of this chapter is to test the change in the model's output for a change in the chosen parameters and to determine which of the selected parameter values yield the best performance.

9.1. Model Scalability

An important sensitivity analysis is the scalability of the model. This shows how the computation time and results of the model change when the model size is increased or decreased. In order to test this, an experiments has been designed in which the model is solved for a planning window of 6 months, 1 year and 2 years. This is done with the same start date and planning frequency. For these simulations, a number of parameters are stored. The amount of time a transition is evaluated and the number of iterations show how many times a transition has been planned and how many iterations have been performed. For configuration P1, 2 transitions are planned at every iteration, so the number of transitions should be roughly twice the number of iterations. Furthermore, the computation times for evaluating a transition, computation time per iteration and the total computation time are stored. All simulations are done using the selection algorithm configuration P1 with a width of 2 and height of 2, as presented in Chapter 6. From these times, it can be assessed what the increase in computation time is for increasing planning windows but also what causes this increase in computation time. In Table 9.1, the mentioned parameters for the three simulations with different planning windows are shown. These results are the averaged values based on three different scenarios per planning window size.

From the results, it can be seen that the amount of evaluated transitions and the number of iterations increases with a factor 3 if the planning window increases by a factor 2. However, the time required to evaluate a transition also increases, with a factor of around 1.7. Overall, this results in an increase in computational time with a factor of 5 for a simulation with a planning window double in size. This can be extrapolated to longer planning windows which would mean a planning window of 4 years has a computation time of almost 1.5 hours and the computation time for a planning window of 8 years is estimated at almost 7 hours.

Table 9.1: Averaged results for the model scalability sensitivity analysis.

Planning window	Amount of		Average time [s]		
	transitions	iterations	transition	iteration	total
6 months	474	239	0.078	0.155	37
1 year	1502	754	0.133	0.265	200
2 years	4512	2257	0.218	0.754	986

9.2. Objective Function

The objective function as designed in Chapter 5 utilises a parameter β that represents the relative importance of shortages versus surplusses. In order to choose the best value for this parameter out of the selected options, a sensitivity analysis on a number of different β values is performed.

In this analysis, a number of scenarios are solved using 5 different β values: 1, $\sqrt{2}$, 2, 3 and 4. For the value of β,

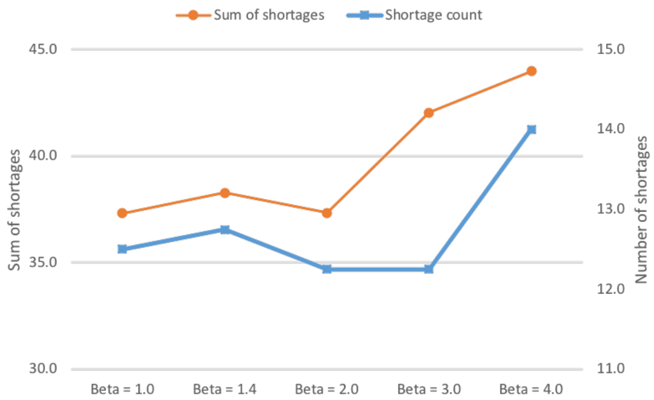

Figure 9.1: Plot showing the sum of shortages and the shortage count for the five models with different β values.

values below 1 are not helpfull as the goal of the parameter is to increase the relative importance of shortages, while values below one do exactly the opposite. Furthermore, value over 4 are assumed to bee too high as in that case the relative importance of surplusses becomes negligible. For all generated solutions, the balance is used to assess the quality of the solution by taking two parameters. The first is the sum of the shortages in the balance. Secondly, the number of shortages is used.

In Figure 9.1, the averaged results for the two parameters for the five different β values are shown. On the left y-axis, the sum of shortages is plotted, while on the right y-axis, the number of shortages is plotted. The average results are obtained from four different scenarios with the same settings as in the experiments in Chapter 7. It can be seen that the models with $\beta = 2$ and $\beta = 3$ score best on the number of shortages. The sum of the shortages of the model with $\beta = 3$ is, however, a lot higher. This means the model with $\beta = 3$ results in solutions with higher shortages than the model with $\beta = 2$. This model ($\beta = 2$) also scores best on average in terms of the sum of shortages (tied with the model with $\beta = 1$). Based on this analysis, it can be said that the best of the tested values for β is 2, which is why this value is used throughout the research.

9.3. Planning Model Frequency

In the development of the model, the supply and demand have been averaged over a period of a month and subsequently, transitions are planned on a monthly basis. Doing this decreases the accuracy of the simulation, but greatly decreases the number of decision variables as opposed to planning transitions on a daily basis.

The planning frequency of 1 month has been chosen since it is the highest frequency that allows a simplification of the changes in demand and supply through absence due to transitions. This is because the longest transitions take just under a month. In order to study the effect of a changing planning model frequency, a scenario with a planning window of 2 years has been solved with a frequency of 1, 2 and 3 months.

In Table 9.2, the initial parameters for the scenarios with three different planning frequencies are shown. It can be seen that because of the larger smoothing of the balance, the number of shortages, minimum balance value and objective function decrease with an increasing planning frequency.

The results for the 9 different simulations can be found in Table 9.3. From the results, it can be seen that the computation time decreases with an increasing frequency for all three different scenarios. In general, the amount of shortages of the optimal solution also decreases, which is to be expected as the total number of data points also decreases. The results in terms of the lowest balance value and amount of transitions do not show a obvious difference. Instead, from these varying results, it can be concluded that the solutions with a lower planning frequency are less accurate and stable because the effect of the assumptions becomes greater. From the results, it can be concluded that in terms of solution accuracy, it is best to develop the model with a planning frequency as high as possible, which in this case is equal to 1 month. However, the models with a lower frequency show a clear improvement in computation time, which could be useful for simulations in which a significantly longer planing window is necessary. In these cases, the planning frequency can be decreased for the periods further into the future. The loss in accuracy for these periods is not as important as the predictions further in the future are also less accurate and a

Table 9.2: Initial parameters for the planning model frequency sensitivity analysis.

Scenario	Frequency	Objective function	Shortages	Min. balance
S1	1 month	$7.268 \cdot 10^{10}$	144	−91.9
	2 months	$6.983 \cdot 10^{10}$	71	−85.4
	3 months	$7.403 \cdot 10^{10}$	48	−87.5
S2	1 month	$2.770 \cdot 10^{10}$	166	−72.8
	2 months	$2.818 \cdot 10^{10}$	82	−71.3
	3 months	$2.696 \cdot 10^{10}$	55	−65.8
S3	1 month	$1.448 \cdot 10^{9}$	66	−20.2
	2 months	$1.419 \cdot 10^{9}$	33	−18.2
	3 months	$1.288 \cdot 10^{9}$	21	−17.9

Table 9.3: Results for the planning model frequency sensitivity analysis.

Scenario	Frequency	Objective function	Time [s]	Shortages	Min. balance	Transitions
S1	1 month	$1.721 \cdot 10^{8}$	1711	26	−15.7	1363
	2 months	$4.064 \cdot 10^{8}$	806	16	−27.8	1251
	3 months	$7.448 \cdot 10^{8}$	698	22	−23.3	1276
S2	1 month	$5.211 \cdot 10^{8}$	435	41	−31.6	712
	2 months	$2.477 \cdot 10^{8}$	359	14	−20.2	795
	3 months	$3.110 \cdot 10^{8}$	317	13	−19.7	835
S3	1 month	$4.443 \cdot 10^{7}$	115	11	−8.4	291
	2 months	$4.392 \cdot 10^{7}$	59	8	−4.9	238
	3 months	$6.684 \cdot 10^{7}$	64	7	−5.7	247

very accurate model far into the future therefore does not yield more accurate results.

9.4. Tabu Search

The tabu search method used in the planning model blocks certain transitions from being planned for a number of iterations. This method is used to prevent the model from being stuck in local optima and keeping looping around this point. The length of the list of the tabu search method defines for how many iterations planned transitions are blocked.

To test the sensitivity of the model to different tabu list lengths, an experiment has been designed to test the model's results for a number of different scenarios for 6 different tabu list lengths: 0 (which means no tabu search method is utilised), 1, 3, 5, 10 and 20.

From the results in Figure 9.2, created using the selection algorithm configuration P1 with a width and height of 2, as the number of simulations required for this sensitivity analysis is quite high (18 simulations). It can be seen that the model with a tabu list length of 0 results in the worst objective function value for all scenarios. Furthermore, the model with a tabu list length of 20 performs worse than the other models with a length between 1 and 10. Out of those four models, the model with a tabu list length of 10 just outperforms the others. In the development of the planning model and in the experiments, a tabu list with a length of 10 iterations has, therefore, been used.

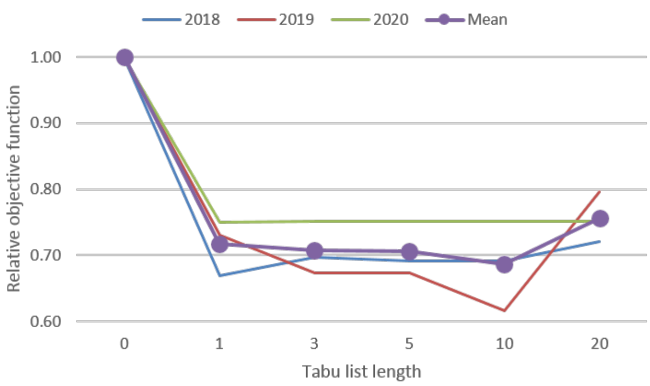

Figure 9.2: Plot showing the normalised objective function values for 3 different scenarios as well as the mean for the 6 different tabu list lengths.

10 Conclusions and Recommendations

This final chapter first presents the conclusions that can be drawn from the performed research in Section 10.1. Then, in Section 10.2, the contributions of the developed model will be discussed and finally, in Section 10.3, some recommendations for improvement or further research will be given.

10.1. Conclusions

The main topic of this report has been defined as the airline cockpit crew transition planning problem. In a literature study, a research gap and subsequently a research question has been defined as:

How to model the transition planning of cockpit crew to provide insight into future staffing levels and transitions and analyse different planning scenarios, strategies and assumptions and their long-term effect?

In a review of the available literature, it was found that there is not a lot of research into the domain of transition planning. Some commercial solvers are available that take days to solve the problem but no fast solution methods exists that accurately model the problem. It was also found that the problem is too complex and large to be solved using an optimisation method. Therefore, an efficient heuristic method is necessary to solve the problem. Solving the transition planning problem in limited time provides airlines with the opportunity to analyse different scenarios, variation in parameters such as illness or bidding behaviour and test different strategies. This gives the airline more insight into the process and helps airlines innovate and create better policies for cockpit crew planning.

In order to answer the posed research question, a research framework has been set up in which several methods can be compared. A heuristic planning model was developed that is able to create a crew plan using a local search method. This model is supported by a rule-based system and tabu-search method that decrease the number of available options in the model and is able to steer the model towards favourable options. This improves the speed at which solutions for the transition planning problem are found. In order to select the best transition at each iteration in the planning model, a tree search method is proposed. For this method, the extensiveness of the search can be changed by changing two parameters, the width and the depth which define the number of options evaluated for each solution and the number of levels in the tree. A variation of the tree search method has been developed which aims to decrease the computation time without compromising in solution quality. This variation is based on the Dijkstra algorithm for finding the shortest path between two nodes. For the selection algorithms, several different configurations have been developed that are tested in a number of experiments. The objective function of the model was designed to reflect the cost of having shortages in supply as well as the cost of having a surplus of pilots. The other crew costs for the airline, such as the cost for transitions or crew salaries have not been directly implemented into the objective function as they are already constraint to a maximum or incorporated into the shortages and surplusses indirectly.

Two extreme points of the tree search method have been identified. The first is a greedy selection algorithm in which the local performance of all options is compared and the best is chosen. This algorithm is created by defining a tree search model with a width equal to ∞ and a depth equal to 1. Secondly, a naive selection algorithm in which an option is selected based on prior knowledge of the options and the system. This algorithm is created by defining a tree search model with a width and depth equal to 1. For the tree search and shortest path methods, four more configurations have been selected as a baseline configuration, a configuration emphasizing width, one focussing on depth and a combination of both. The configurations have been selected in such a manner that the computation time falls within the computation times of the greedy and naive selection algorithms.

In an experiment based on data from a reference airline, the ten selected configurations of the selection algorithm have been tested on 10 different scenarios. These scenarios have been constructed by selecting a different start and end date for the planning window and constructing the supply, demand and balance for this time period. From the results, several conclusions can be drawn.

- The configuration with the highest median solution quality is the greedy algorithm (tree search with a width equal to ∞ and depth of 1). However, the computation time of this algorithm is the highest of all tested configurations (almost an hour on average) and the stability of the solution quality is the second worst of all configurations, only out-performing the naive selection configuration. The usefulness of this configuration is, therefore, limited and a minimum depth of 2 might be required in order to improve the solution stability.
- The best performing configuration was the shortest path algorithm focussing on a combination of width and height (width = 3, height = 4). With an average computation time of only 15 minutes, it had the fourth lowest computation time. The configuration produces results 4 times faster than the greedy algorithm. With a median scaled objective function of 18%, the model still outperforms all models apart from the greedy algorithm. Finally, the high of the configuration was the highest of all tested configurations (with an interquartile range (IQR) of 0.07 versus an IQR of the greedy algorithm of 0.75).
- Finally, the fastest configuration is the naive selection configuration, with an average computation time of 2.4 minutes. However, both the solution quality and solution stability of this configuration score worse than all other tested configurations. The second fasted configuration scores considerably better in terms of solution quality (a median scaled objective function value of 0.24 versus 0.76) and solution stability (an IQR of 0.47 versus 0.81). The computation time of this configuration is almost double the computation time of the naive selection algorithm (4.6 minutes), but still allows significantly more simulations in a limited time compared to the other configurations.

In general, the shortest path method is able to determine solutions faster than the tree search method without compromising the solution quality. Furthermore, it is concluded that a combination of width and depth produces the best and most stable results, as opposed to configurations focussing on either width or depth. Finally, from the results of the naive selection and greedy configurations, it can be concluded that a minimum depth of 2 levels is required, as this greatly improves the stability of the solutions. In summary, it can be concluded that all presented configurations are valid methods to solve the cockpit crew transition planning problem. Depending on the application, different requirements with regards to computation time, solution quality and solution stability might be set. These requirements will also influence what configuration will be best suited for the application. However, overall, model P3, a shortest path algorithm configuration with a width of 3 and height of 4 can be seen as the optimal configuration, scoring the second-best median objective function value, a limited computation time of 15 minutes on average and presenting the most stable results of all tested configurations.

10.2. Research Contributions

The contributions of the presented research are discussed in this section. The contributions state how the research adds value and knowledge to the scientific body of research already available in the field of cockpit crew planning as well as how the presented research and models can help the business of airlines and other companies.

Problem definition
Not much research has been published in the field of cockpit crew planning. Especially the transition planning problem has been untouched for years and no models have been developed that are able to quickly create a solution that can be used to analyse different scenarios. This research, therefore, fills a research gap in the field of cockpit crew planning that could aid airlines in their decision-making process.

Solution methods
The developed solution methods for the transition planning problem form a novel application of known methods to this problem. Specifically for the transition planning problem, no solution methods have been published in the past. This research, therefore, forms a first step on which future research can build.

Model applicability
As the model has been developed with generality in mind, it can be applied not only to the field of cockpit crew planning but to different problems and industries as well. The first problem that seems evident is the cabin crew planning problem. A lot of restrictions that are present in the cockpit crew transition planning problem are also applicable to cabin crew, which makes it easy to change the model to be able to plan cabin crew transitions. A distinction between cabin crew and cockpit crew, however, is that in almost all cases, the cabin crew is qualified to

operate on multiple aircraft types. Other industries for which the model can be used are nurse planning and army planning as has already been discussed in the literature review of this report. The manpower planning problem in these industries shows great resemblance to the cockpit crew transition planning problem.

Practical application
On the practical side of the research, the developed model provides airlines and other companies the opportunity to create crew plans in limited time. Through this, they are better able to analyse different strategies, variations of demand parameters and scenarios. This will provide more insight into the future and makes it possible to perform more analyses in the decision-making process.

10.3. Recommendations

Following the conclusions as well as the academic and practical contributions of the research, some recommendations regarding opportunities for further research should be made.

Assumptions
In order to increase the accuracy of the model's results, a more detailed analysis of the impact of the different assumptions made throughout the project can further increase the accuracy of the model. From the validation, it was already concluded that the change of supply due to the usage of instructors can influence the total amount of transitions planned by the system more than expected. When this is modelled in more detail, the results will more accurately match the validation results.

Methods
In the developed model, the only method to resolve shortages is to plan transitions. However, an airline can usually employ different methods depending on the exact state. An example of these methods is to change the planned vacation budgets if these budgets are not yet allocated to specific pilots yet. By doing this, shortages and surplusses can be decreased since the vacation budget in a period with a shortage can be moved to a period with a surplus. Implementing such a method asks for a decision tool that decides which method to employ at each point in time, creating a hyper-heuristic model for the cockpit crew planning problem.

Experiments
The experiments for the different models performed in this research are created from a single airline's system and data. This means that the different models and algorithms have only been tested on this system. Even though this system is comparable to the system used at a high number of other airlines, it is useful to test the models on different systems as this increase the confidence in the results and the conclusions that are drawn from these results. Also, testing the algorithms on systems from other industries might provide insight into the validity of the proposed models on a wider number of applications within the field of manpower planning. Furthermore, the experiments have only been done for planning windows up to 2 years to limit the total time needed for all experiments. As the cockpit crew transition planning problem is a long-term problem in which usually up to 3 to 5 years are planned in advance, experiments in which this planning window is larger could also benefit the overall results and conclusions that can be drawn from these. If necessary, the computation time for these experiments can be slightly decreased by decreasing the planning frequency for later periods and using smaller models.

References

Abernathy, W. J., Baloff, N., Hershey, J. C., & Wandel, S. (1973). A three-stage manpower planning and scheduling model—a service-sector example. *Operations Research*, *21*(3), 693–711.

Abramson, M. A. (2002). *Pattern search algorithms for mixed variable general constrained optimization problems* (Unpublished doctoral dissertation). Rice University Texas.

Akaike, H. (1974). A new look at the statistical model identification. *IEEE transactions on automatic control*, *19*(6), 716–723.

Altenstedt, F., Thalén, B., Sjögren, P., & Nilsson, S. (2017). Solving the airline manpower planning problem. In S. Albers, N. Megow, A. S. Schulz, & L. Stougie (Eds.), *Proceedings of the 13th workshop on models and algorithms for planning and scheduling problems* (p. 53 - 55).

Bard, J. F., Morton, D. P., & Wang, Y. M. (2007). Workforce planning at USPS mail processing and distribution centers using stochastic optimization. *Annals of Operations Research*, *155*(1), 51–78.

Barnhart, C., Belobaba, P., & Odoni, A. R. (2003). Applications of operations research in the air transport industry. *Transportation Science*, *37*(4), 368 - 391.

Bayliss, C. (2016). *Airline reserve crew scheduling under uncertainty* (PhD thesis). University of Nottingham.

Ben-Tal, A., & Nemirovski, A. (1999). Robust solutions of uncertain linear programs. *Operations research letters*, *25*(1), 1–13.

Birge, J. R., & Louveaux, F. (2011). *Introduction to stochastic programming*. Springer Science & Business Media.

Cacchiani, V., & Salazar-González, J.-J. (2017). Optimal solutions to a real-world integrated airline scheduling problem. *Transportation Science*, *51*(1), 250-268.

Catalá, L. P., Moreno, M. S., Blanco, A. M., & Bandoni, J. A. (2016). A bi-objective optimization model for tactical planning in the pome fruit industry supply chain. *Computers and Electronics in Agriculture*, *130*, 128–141.

Clarke, M., & Smith, B. (2004). Impact of operations research on the evolution of the airline industry. *Journal of Aircraft*, *41*(1), 62 - 72.

Deng, J.-L. (1982). Control problems of grey systems. *Systems & Control Letters*, *1*(5), 288 - 294.

Dijkstra, E. W. (1959). A note on two problems in connexion with graphs. *Numerische mathematik*, *1*(1), 269–271.

Edwards, J. (2007). A survey of manpower planning models and their application. *Journal of Operation Research Society*, *34*(1), 70 - 88.

Fahle, T., Junker, U., Karisch, S., Kohl, N., Sellmann, M., & Vaaben, B. (2002). Constraint programming based column generation for crew assignment. *Journal of Heuristics*, *8*(1), 59-81.

Fozveh, I., Salehi, H., & Mogharehabed, K. (2016). Presentation of multi-skill workforce scheduling model and solving the model using meta-heuristic algorithms. *Modern Applied Science*, *10*(2), 194.

Gass, S. I., Collins, R. W., Meinhardt, C. W., Lemon, D. M., & Gillette, M. D. (1988). Or practice—the army manpower long-range planning system. *Operations Research*, *36*(1), 5–17.

Gendreau, M. (2003). An introduction to tabu search. In F. Glover & G. A. Kochenberger (Eds.), *Handbook of metaheuristics* (pp. 37–54). Boston, MA: Springer US.

Gershkoff, I. (1989). Optimizing flight crew schedules. *Interfaces*, *19*(4), 29-43.

Glover, F. (1997). Tabu search and adaptive memory programming—advances, applications and challenges. In *Interfaces in computer science and operations research* (pp. 1–75). Springer.

Ho, P. H. (2010). Forecasting construction manpower demand by gray model. *Journal of Construction Engineering and Management*, *136*(12), 1299–1305.

Holm, Å. (2008). *Manpower planning in airlines : Modeling and optimization* (Master's Thesis). Linköping University, Department of Mathematics.

Hosmer, D. W., & Lemesbow, S. (1980). Goodness of fit tests for the multiple logistic regression model. *Communications in Statistics - Theory and Methods*, *9*(10), 1043-1069.

Ives, C. D. (1992). *Promotional planning for airline technical crew* . Department of Engineering Science, University of Auckland.

Janssen, R. (2018). *Airline reserve crew pairing optimisation* (Master's Thesis). Delft Univerisity of Technology.

Jarmar, J., & Sörensson, F. (2017). *A constraint programming approach to finding stable matchings within airline manpower planning* (Master's Thesis). University of Gothenburg, Chalmers University of Technology, Department of Computer Science & Engineering.

Jarrah, A. I. Z., & Diamond, J. T. (1997). The problem of generating crew bidlines. *Interfaces*, *27*(4), 49-64.

Kohl, N., & Karisch, S. (2004). Airline crew rostering: Problem types, modeling, and optimization. *Annals of Operations Research*, *127*(1-4), 223-257.

Lettovský, L., Johnson, E. L., & Nemhauser, G. L. (2000). Airline crew recovery. *Transportation Science*, *34*(4), 337 - 348.

Li, Y., Chen, J., & Cai, X. (2007). An integrated staff-sizing approach considering feasibility of scheduling decision. *Annals of Operations Research*, *155*(1), 361-390.

List, G. F., Wood, B., Nozick, L. K., Turnquist, M. A., Jones, D. A., Kjeldgaard, E. A., & Lawton, C. R. (2003). Robust optimization for fleet planning under uncertainty. *Transportation Research Part E: Logistics and Transportation Review*, *39*(3), 209 - 227.

Listes, O., & Dekker, R. (2005). A scenario aggregation–based approach for determining a robust airline fleet composition for dynamic capacity allocation. *Transportation Science*, *39*(3), 367–382.

Lü, Z., & Hao, J.-K. (2008). Solving the course timetabling problem with a hybrid heuristic algorithm. In D. Dochev, M. Pistore, & P. Traverso (Eds.), *Artificial intelligence: Methodology, systems, and applications* (pp. 262–273). Berlin, Heidelberg: Springer Berlin Heidelberg.

Lü, Z., Hao, J.-K., & Glover, F. (2011). Neighborhood analysis: a case study on curriculum-based course timetabling. *Journal of Heuristics*, *17*(2), 97–118.

Maenhout, B., & Vanhoucke, M. (2013). An integrated nurse staffing and scheduling analysis for longer-term nursing staff allocation problems. *Omega*, *41*(2), 485–499.

Manlove, D. F., O'Malley, G., Prosser, P., & Unsworth, C. (2007). A constraint programming approach to the hospitals / residents problem. In P. Van Hentenryck & L. Wolsey (Eds.), *Integration of ai and or techniques in constraint programming for combinatorial optimization problems* (pp. 155–170). Springer Berlin Heidelberg.

Marsten, R. E., & Muller, M. R. (1980). A mixed-integer programming approach to air cargo fleet planning. *Management Science*, *26*(11), 1096–1107.

Medard, C. P., & Sawhney, N. (2007). Airline crew scheduling from planning to operations. *European Journal of Operational Research*, *183*(3), 1013 - 1027.

Metropolis, N., Rosenbluth, A. W., Rosenbluth, M. N., Teller, A. H., & Teller, E. (1953). Equation of state calculations by fast computing machines. *The journal of chemical physics*, *21*(6), 1087–1092.

Ng, T., Huang, H., & Ng, J. (2008). Human resource planning with worker attendance uncertainty. In *Industrial engineering and engineering management, 2008. ieem 2008. ieee international conference on* (pp. 364–368).

Nissen, R., & Haase, K. (2006). Duty-period-based network model for crew rescheduling in european airlines. *Journal of Scheduling*, *9*(3), 255-278.

Price, W., Martel, A., & Lewis, K. (1980). A review of mathematical models in human resource planning. *Omega*, *8*(6), 639 - 645.

Punnakitikashem, P., Rosenberger, J. M., Behan, D. B., Baker, R. L., & Goss, K. (2006). An optimization-based prototype for nurse assignment. In *Proceedings of the 7th asian pacific industrial engineering and management systems conference* (pp. 17–20).

Purkiss, C. (1981). Corporate manpower planning: a review of models. *European Journal of Operational Research, 8*(4), 315 - 323.

Qi, M., & Zhang, G. P. (2008). Trend time–series modeling and forecasting with neural networks. *IEEE Transactions on neural networks, 19*(5), 808–816.

Requeno García, L. (2017). *Multi-period adaptive fleet planning problem with approximate dynamic programming* (Master's Thesis). Delft University of Technology.

Rosskopf, M., Lehner, S., & Gollnick, V. (2014). Economic–environmental trade-offs in long-term airline fleet planning. *Journal of Air Transport Management, 34*, 109 - 115.

Saliby, E. (1990). Descriptive sampling: a better approach to monte carlo simulation. *Journal of the Operational Research Society, 41*(12), 1133–1142.

Sawik, T. (2007). A lexicographic approach to bi-objective scheduling of single-period orders in make-to-order manufacturing. *European Journal of Operational Research, 180*(3), 1060 - 1075.

Shao, S., Sherali, H., & Haouari, M. (2017). A novel model and decomposition approach for the integrated airline fleet assignment, aircraft routing, and crew pairing problem. *Transportation Science, 51*(1), 233-249.

Shebalov, S., & Klabjan, D. (2006). Robust airline crew pairing: Move-up crews. *Transportation Science, 40*(3), 300-312.

Silverman, J., Steuer, R. E., & Whisman, A. W. (1988). A multi-period, multiple criteria optimization system for manpower planning. *European Journal of Operational Research, 34*(2), 160–170.

Škraba, A., Kljajić, M., Papler, P., Kofjač, D., & Obed, M. (2011). Determination of recruitment and transition strategies. *Kybernetes, 40*(9/10), 1503–1522.

Škraba, A., Stanovov, V., Semenkin, E., & Kofjač, D. (2016). Hybridization of stochastic local search and genetic algorithm for human resource planning management. *Organizacija, 49*(1), 42–54.

Sohoni, M. G., Johnson, E. L., & Bailey, T. G. (2004). Long-range reserve crew manpower planning. *Management Science, 50*(6), 724 - 739.

Sohoni, M. G., Johnson, E. L., & Bailey, T. G. (2006). Operational airline reserve crew planning. *Journal of Scheduling, 9*(3), 203–221.

Thalén, B. (2010). *Manpower planning for airline pilots — a tabu search approach* (Master's Thesis). Linköping University, Department of Mathematics.

Thathachar, M. A., & Sastry, P. S. (2002). Varieties of learning automata: an overview. *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics), 32*(6), 711–722.

Torczon, V. (1997). On the convergence of pattern search algorithms. *SIAM Journal on optimization, 7*(1), 1–25.

Tseng, F.-M., Yu, H.-C., & Tzeng, G.-H. (2001). Applied hybrid grey model to forecast seasonal time series. *Technological Forecasting and Social Change, 67*(2), 291 - 302.

Verbeek, P. J. (1991). Decision support systems — An application in strategic manpower planning of airline pilots. *European Journal of Operational Research, 55*(3), 368 - 381.

Wang, J. (2005). *A review of operations research applications in workforce planning and potential modeling of military training* (Tech. Rep.). Defence Science and Technology Organistion Salisbury (Australia) Systems Sciences Lab.

Wets, R. J. (1989). The aggregation principle in scenario analysis and stochastic optimization. In *Algorithms and model formulations in mathematical programming* (pp. 91–113). Springer.

Winston, W. L., & Goldberg, J. B. (2004). *Operations research: applications and algorithms* (Vol. 3). Thomson Brooks/Cole Belmont.

Yu, G., Argüello, M., Song, G., McCowan, S. M., & White, A. (2003). A new era for crew recovery at continental airlines. *Interfaces*, *33*(1), 5 - 22.

Yu, G., Dugan, S., & Argüello, M. (1998). Moving toward an integrated decision support system for manpower planning at continental airlines: Optimization of pilot training assignments. In G. Yu (Ed.), *Industrial applications of combinatorial optimization* (p. 1 - 24). Springer, Boston.

Yu, G., Pachon, J., & Thengvall, B. (2003). Optimization-based integrated manpower management for airlines. In T. A. Ciriani, G. Fasano, S. Gliozzi, & R. Tadei (Eds.), *Operations research in space and air* (p. 407 - 434). Springer, Boston.

Yu, G., Pachon, J., Thengvall, B., Chandler, D., & Wilson, A. (2004). Optimizing pilot planning and training for continental airlines. *Interfaces*, *34*(4), 253 - 264.

Zhu, X., Sanil, G., Mardookhy, M., Sawhney, R., & Sun, X. (2013). Statistical analysis for human capital forecasting in multipurpose research organization. In *IIE annual conference and expo 2013* (p. 3536-3542).

Zhu, X., & Sherali, H. (2009). Two-stage workforce planning under demand fluctuations and uncertainty. *Journal of the Operational Research Society*, *60*(1), 94-103.

A

Example Problem Data

This appendix displays the data of the simplified example problem as presented in Section 4.1.

Table A.1: Full list of planned transitions in simplified example problem (first part).

Iteration	Type	Employee	From	To	Date
0	Transition	2394	FO ICA	CP EUR	2019-08-01
1	Transition	7130	FO ICA	CP EUR	2019-08-01
2	Transition	9196	FO ICA	CP EUR	2019-08-01
3	Transition	8049	FO ICA	CP EUR	2019-08-01
4	Transition	8398	FO ICA	CP EUR	2019-08-01
5	Transition	6878	FO ICA	CP EUR	2019-08-01
6	Transition	0613	FO ICA	CP EUR	2019-08-01
7	Transition	0223	FO ICA	CP EUR	2019-08-01
8	Corrected	8294	SO ICA	FO EUR	2019-08-01
9	Recruit	R35	-	SO ICA	2019-07-01
10	Corrected	2231	SO ICA	FO EUR	2019-08-01
11	Recruit	R36	-	SO ICA	2019-07-01
12	Corrected	5571	SO ICA	FO EUR	2019-08-01
13	Corrected	2656	SO ICA	FO EUR	2019-08-01
14	Recruit	R37	-	SO ICA	2019-07-01
15	Corrected	1844	SO ICA	FO EUR	2019-08-01
16	Recruit	R38	-	SO ICA	2019-07-01
17	Corrected	3165	SO ICA	FO EUR	2019-08-01
18	Recruit	R39	-	SO ICA	2019-07-01
19	Corrected	2015	SO ICA	FO EUR	2019-08-01
20	Recruit	R40	-	SO ICA	2019-07-01

Table A.2: Full list of planned transitions in simplified example problem (second part).

Iteration	Type	Employee	From	To	Date
21	Corrected	1802	SO ICA	FO EUR	2019-08-01
22	Recruit	R41	-	SO ICA	2019-07-01
23	Corrected	2570	SO ICA	FO EUR	2019-08-01
24	Recruit	R42	-	SO ICA	2019-07-01
25	Corrected	3663	SO ICA	FO EUR	2019-08-01
26	Recruit	R43	-	SO ICA	2019-07-01
27	Corrected	2935	SO ICA	FO EUR	2019-08-01
28	Recruit	R44	-	SO ICA	2019-07-01
29	Transition	6538	SO ICA	FO EUR	2019-08-01
30	Recruit	R45	-	SO ICA	2019-07-01
31	Correction	8294	SO ICA	FO EUR	2019-07-01
32	Transition	5850	SO ICA	FO EUR	2019-08-01
33	Recruit	R46	-	SO ICA	2019-07-01
34	Correction	2231	SO ICA	FO EUR	2019-07-01
35	Correction	5571	SO ICA	FO EUR	2019-07-01
36	Transition	6869	SO ICA	FO EUR	2019-08-01
37	Recruit	R47	-	SO ICA	2019-07-01
38	Correction	2656	SO ICA	FO EUR	2019-07-01
39	Correction	1844	SO ICA	FO EUR	2019-07-01
40	Transition	8318	SO ICA	FO EUR	2019-08-01
41	Recruit	R48	-	SO ICA	2019-07-01
42	Correction	3165	SO ICA	FO EUR	2019-07-01
43	Correction	2015	SO ICA	FO EUR	2019-07-01
44	Transition	4599	SO ICA	FO EUR	2019-08-01
45	Recruit	R49	-	SO ICA	2019-07-01
46	Correction	1802	SO ICA	FO EUR	2019-07-01
47	Correction	2570	SO ICA	FO EUR	2019-07-01
48	Transition	8193	SO ICA	FO EUR	2019-08-01
49	Recruit	R50	-	SO ICA	2019-07-01
50	Correction	3663	SO ICA	FO EUR	2019-07-01
51	Correction	2935	SO ICA	FO EUR	2019-07-01
52	Transition	7804	SO ICA	FO EUR	2019-08-01
53	Recruit	R51	-	SO ICA	2019-07-01
54	Recruit	R52	-	SO ICA	2019-08-01
55	Recruit	R53	-	SO ICA	2019-08-01
56	Recruit	R54	-	SO ICA	2019-08-01
57	Recruit	R55	-	SO ICA	2019-08-01
58	Recruit	R56	-	SO ICA	2019-08-01
59	Recruit	R57	-	SO ICA	2019-08-01
60	Recruit	R58	-	SO ICA	2019-08-01

Table A.3: Full list of planned transitions in simplified example problem (third part).

Iteration	Type	Employee	From	To	Date
61	Recruit	R59	-	SO ICA	2019-08-01
62	Recruit	R60	-	SO ICA	2019-07-01
63	Transition	8078	FO ICA	CP EUR	2019-09-01
64	Recruit	R61	-	SO ICA	2019-09-01
65	Recruit	R62	-	SO ICA	2019-09-01
66	Recruit	R63	-	SO ICA	2019-09-01
67	Recruit	R64	-	SO ICA	2019-09-01
68	Recruit	R65	-	SO ICA	2019-08-01
69	Transition	9031	FO ICA	CP EUR	2019-10-01
70	Transition	0970	FO ICA	CP EUR	2019-10-01
71	Transition	1845	FO ICA	CP EUR	2019-10-01
72	Recruit	R66	-	SO ICA	2019-10-01
73	Recruit	R67	-	SO ICA	2019-10-01
74	Recruit	R68	-	SO ICA	2019-10-01
75	Recruit	R69	-	SO ICA	2019-10-01
76	Recruit	R70	-	SO ICA	2019-10-01
77	Recruit	R71	-	SO ICA	2019-10-01
78	Recruit	R72	-	SO ICA	2019-09-01
79	Recruit	R73	-	SO ICA	2019-10-01
80	Recruit	R74	-	SO ICA	2019-10-01
81	Recruit	R75	-	SO ICA	2019-10-01
82	Recruit	R76	-	SO ICA	2019-10-01
83	Recruit	R77	-	SO ICA	2019-10-01
84	Recruit	R78	-	SO ICA	2019-10-01
85	Recruit	R79	-	SO ICA	2019-10-01
86	Recruit	R80	-	SO ICA	2019-09-01
87	Transition	8869	CP EUR	CP ICA	2019-11-01
88	Transition	1121	CP EUR	CP ICA	2019-11-01
89	Transition	2139	FO ICA	CP EUR	2019-10-01
90	Transition	1217	CP EUR	CP ICA	2019-11-01
91	Corrected	1930	FO ICA	CP EUR	2019-11-01
92	Correction	1930	FO ICA	CP EUR	2019-10-01
93	Transition	1664	CP EUR	CP ICA	2019-11-01
94	Transition	0042	CP EUR	CP ICA	2019-11-01
95	Transition	2260	FO ICA	CP EUR	2019-10-01
96	Transition	0615	CP EUR	CP ICA	2019-11-01
97	Transition	4128	FO ICA	CP EUR	2019-10-01
98	Transition	3694	FO ICA	CP ICA	2019-12-01
99	Transition	7105	FO EUR	FO ICA	2019-11-01
100	Transition	9844	FO EUR	FO ICA	2019-11-01
101	Transition	8623	FO EUR	FO ICA	2019-11-01
102	Transition	6161	FO EUR	FO ICA	2019-11-01
103	Transition	1853	FO EUR	FO ICA	2019-11-01
104	Transition	3077	FO EUR	FO ICA	2019-11-01

B

Detailed Experiment Results

In this appendix, the detailed results for the twelve tested models in the ten different scenarios are presented. In the tables, the objective function and computation time columns speak for itself, where the computation time is given in seconds. The shortages column presents the amount of shortages in the optimal solution found in the model and the min. balance columns shows the lowest balance value in the same optimal solution. Finally, the transitions column shows the amount of transitions planned in the optimal solution. These last three parameters give some more insight in the solution found by the model. Also, model 0 in the tables represents the initial state of the scenario, therefore, the computation time and amount of transitions for these entries are equal to zero.

Table B.1: Detailed experiment results for the twelve models for the scenario with start date 01/01/2018.

Model	Objective function	Computation time	Shortages	Min. balance	Transitions
0	$1.448 \cdot 10^9$	-	66	-20.2	-
N	$5.751 \cdot 10^7$	79	14	-7,9	317
G	$2.910 \cdot 10^7$	2823	28	-4,8	255
T1	$4.287 \cdot 10^7$	335	12	-7,1	302
T2	$3.051 \cdot 10^7$	1674	12	-6,8	306
T3	$3.806 \cdot 10^7$	1191	13	-6,6	287
T4	$3.910 \cdot 10^7$	1207	16	-6,8	299
P1	$4.677 \cdot 10^7$	134	11	-8,2	298
P2	$3.954 \cdot 10^7$	860	23	-5,8	267
P3	$3.346 \cdot 10^7$	514	12	-5,9	295
P4	$4.843 \cdot 10^7$	690	11	-8,3	316

Table B.2: Detailed experiment results for the twelve models for the scenario with start date 01/04/2018.

Model	Objective function	Computation time	Shortages	Min. balance	Transitions
0	$1.060 \cdot 10^9$	-	65	-37.7	-
N	$6.595 \cdot 10^7$	94	21	-10,2	351
G	$2.522 \cdot 10^7$	2544	22	-4,5	254
T1	$4.011 \cdot 10^7$	524	16	-6,9	313
T2	$6.283 \cdot 10^7$	1703	23	-8,7	328
T3	$4.396 \cdot 10^7$	1324	19	-7,7	313
T4	$8.368 \cdot 10^7$	1270	26	-11,3	308
P1	$4.258 \cdot 10^7$	201	21	-6,9	281
P2	$7.128 \cdot 10^7$	1360	24	-10,2	286
P3	$4.703 \cdot 10^7$	554	18	-8,1	301
P4	$4.169 \cdot 10^7$	879	14	-6,9	273

Table B.3: Detailed experiment results for the twelve models for the scenario with start date 01/07/2018.

Model	Objective function	Computation time	Shortages	Min. balance	Transitions
0	$2.390 \cdot 10^9$	-	60	-38.6	-
N	$1.179 \cdot 10^8$	114	23	-12,4	412
G	$6.334 \cdot 10^7$	2589	36	-13,5	278
T1	$6.625 \cdot 10^7$	562	25	-6,4	401
T2	$7.337 \cdot 10^7$	1975	30	-9,5	356
T3	$7.701 \cdot 10^7$	1511	26	-8,3	366
T4	$9.846 \cdot 10^7$	1698	31	-10,2	372
P1	$6.241 \cdot 10^7$	202	22	-6,5	412
P2	$6.642 \cdot 10^7$	1548	23	-9,1	363
P3	$6.974 \cdot 10^7$	659	24	-9,1	401
P4	$7.324 \cdot 10^7$	738	24	-7,4	325

Table B.4: Detailed experiment results for the twelve models for the scenario with start date 01/10/2018.

Model	Objective function	Computation time	Shortages	Min. balance	Transitions
0	$3.555 \cdot 10^9$	-	61	-33.3	-
N	$2.377 \cdot 10^8$	308	43	-17,9	389
G	$9.658 \cdot 10^7$	5575	46	-9,7	431
T1	$2.280 \cdot 10^8$	1909	37	-13,1	417
T2	$1.527 \cdot 10^8$	2267	48	-12,9	423
T3	$1.270 \cdot 10^8$	2394	47	-10,1	437
T4	$1.629 \cdot 10^8$	2087	37	-16,3	466
P1	$1.900 \cdot 10^8$	688	40	-17,3	381
P2	$1.295 \cdot 10^8$	3811	45	-10,8	406
P3	$1.253 \cdot 10^8$	2260	38	-16,7	368
P4	$2.378 \cdot 10^8$	2701	44	-18,0	379

Table B.5: Detailed experiment results for the twelve models for the scenario with start date 01/01/2019.

Model	Objective function	Computation time	Shortages	Min. balance	Transitions
0	$3.762 \cdot 10^9$	-	61	-65.4	-
N	$2.966 \cdot 10^8$	226	39	-14,8	482
G	$1.334 \cdot 10^8$	6070	38	-25,9	397
T1	$1.767 \cdot 10^8$	1403	36	-13,5	475
T2	$9.933 \cdot 10^7$	6474	37	-12,6	433
T3	$1.357 \cdot 10^8$	4644	31	-11,6	521
T4	$1.203 \cdot 10^8$	3092	32	-13,4	520
P1	$2.348 \cdot 10^8$	477	37	-15,8	566
P2	$1.034 \cdot 10^8$	4434	38	-10,2	473
P3	$2.102 \cdot 10^8$	1433	43	-12,7	523
P4	$2.234 \cdot 10^8$	3547	32	-19,1	547

Table B.6: Detailed experiment results for the twelve models for the scenario with start date 01/07/2019.

Model	Objective function	Computation time	Shortages	Min. balance	Transitions
0	$6.894 \cdot 10^9$	-	46	-66.8	-
N	$5.054 \cdot 10^8$	289	34	-31,1	442
G	$6.055 \cdot 10^8$	6231	45	-20,9	356
T1	$2.723 \cdot 10^8$	2129	47	-21,8	571
T2	$2.151 \cdot 10^8$	6509	35	-24,1	534
T3	$2.376 \cdot 10^8$	4936	39	-19,0	502
T4	$2.140 \cdot 10^8$	3300	37	-19,9	539
P1	$2.261 \cdot 10^8$	504	44	-20,5	563
P2	$1.660 \cdot 10^8$	4573	35	-20,3	537
P3	$2.248 \cdot 10^8$	2389	41	-16,6	555
P4	$2.483 \cdot 10^8$	2705	43	-20,9	562

Table B.7: Detailed experiment results for the twelve models for the scenario with start date 01/01/2020.

Model	Objective function	Computation time	Shortages	Min. balance	Transitions
0	$9.812 \cdot 10^9$	-	40	-51.2	-
N	$4.557 \cdot 10^7$	88	13	-9,8	358
G	$8.304 \cdot 10^7$	2847	30	-10,0	261
T1	$5.553 \cdot 10^7$	568	23	-8,8	326
T2	$8.512 \cdot 10^7$	2261	21	-13,9	298
T3	$4.761 \cdot 10^7$	1855	18	-8,3	316
T4	$5.340 \cdot 10^7$	1879	23	-8,5	312
P1	$5.702 \cdot 10^7$	189	15	-10,7	328
P2	$1.438 \cdot 10^8$	1265	27	-17,4	265
P3	$6.544 \cdot 10^7$	647	22	-9,3	340
P4	$5.353 \cdot 10^7$	548	22	-10,8	322

Table B.8: Detailed experiment results for the twelve models for the scenario with start date 01/07/2020.

Model	Objective function	Computation time	Shortages	Min. balance	Transitions
0	$4.222 \cdot 10^9$	-	40	-44.9	-
N	$5.722 \cdot 10^7$	114	15	-5,7	315
G	$5.800 \cdot 10^7$	2637	11	-7,9	261
T1	$5.510 \cdot 10^7$	317	9	-4,9	300
T2	$5.708 \cdot 10^7$	2251	12	-6,5	285
T3	$5.412 \cdot 10^7$	1264	10	-5,1	270
T4	$5.880 \cdot 10^7$	1574	11	-4,5	284
P1	$5.561 \cdot 10^7$	136	11	-5,8	295
P2	$5.266 \cdot 10^7$	1161	8	-5,8	258
P3	$5.392 \cdot 10^7$	175	11	-4,6	276
P4	$5.550 \cdot 10^7$	264	10	-5,8	292

Table B.9: Detailed experiment results for the twelve models for the scenario with start date 01/01/2021.

Model	Objective function	Computation time	Shortages	Min. balance	Transitions
0	$3.022 \cdot 10^9$	-	46	-29.8	-
N	$4.261 \cdot 10^7$	41	13	-2,9	235
G	$4.258 \cdot 10^7$	1812	15	-2,9	231
T1	$4.345 \cdot 10^7$	302	12	-3,4	246
T2	$4.746 \cdot 10^7$	1571	16	-5,4	239
T3	$4.294 \cdot 10^7$	1227	9	-5,4	240
T4	$4.267 \cdot 10^7$	1176	12	-3,6	238
P1	$4.313 \cdot 10^7$	77	8	-4,2	240
P2	$4.691 \cdot 10^7$	665	15	-5,5	227
P3	$4.280 \cdot 10^7$	200	9	-4,3	242
P4	$4.310 \cdot 10^7$	355	11	-4,0	234

Table B.10: Detailed experiment results for the twelve models for the scenario with start date 01/07/2021.

Model	Objective function	Computation time	Shortages	Min. balance	Transitions
0	$1.878 \cdot 10^9$	-	32	-38.4	-
N	$4.820 \cdot 10^7$	60	4	-7,4	258
G	$7.200 \cdot 10^7$	1705	14	-7,3	197
T1	$5.124 \cdot 10^7$	424	5	-8,1	276
T2	$4.753 \cdot 10^7$	2405	8	-4,3	268
T3	$5.065 \cdot 10^7$	2054	8	-3,8	272
T4	$5.472 \cdot 10^7$	1522	6	-7,5	265
P1	$5.213 \cdot 10^7$	150	6	-8,1	288
P2	$5.474 \cdot 10^7$	1296	8	-8,0	237
P3	$4.904 \cdot 10^7$	250	6	-6,3	277
P4	$5.213 \cdot 10^7$	690	6	-8,1	288