Master Assignment

Capacity planning of police helicopters

How to improve and support the yearly capacity planning of police helicopters?

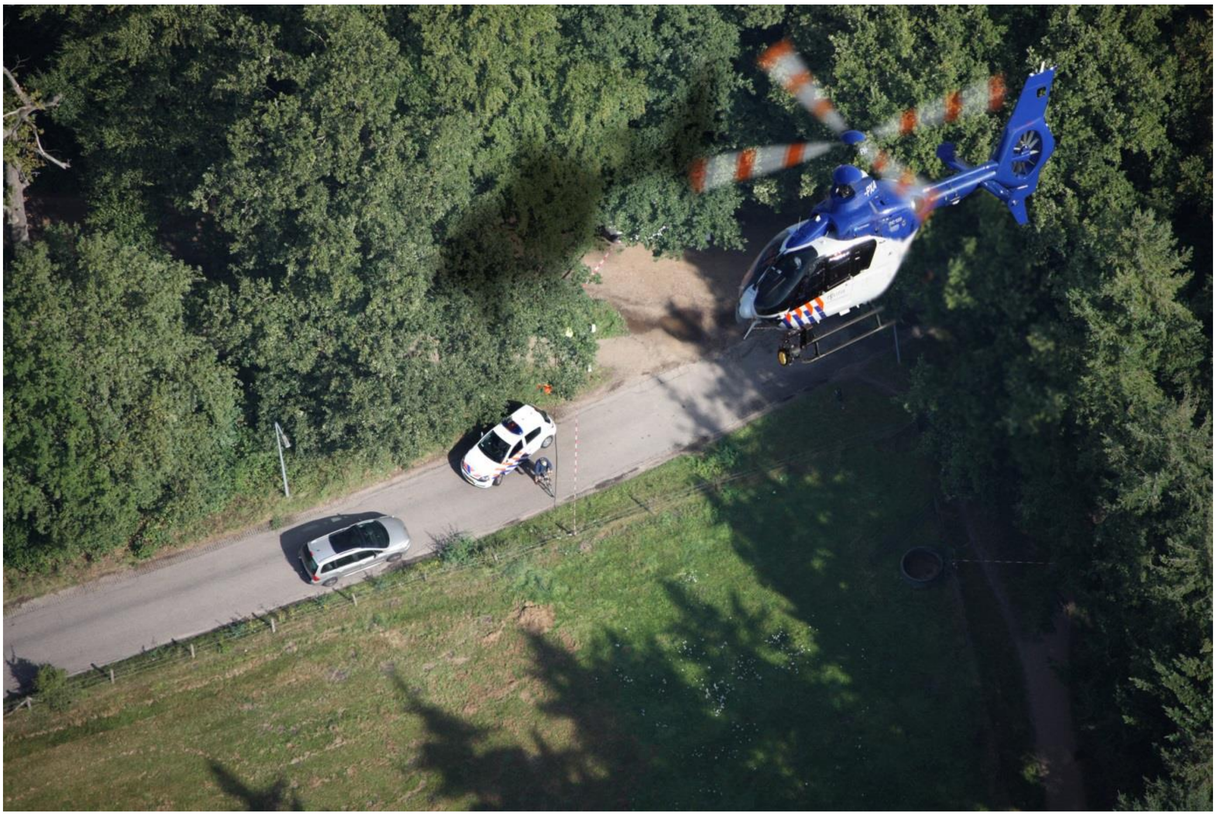

Author: Rob Vromans
Enschede, March 2014

POLITIE

Afdeling Luchtvaartpolitie
Politie | Landelijke Eenheid

Thermiekstraat 2 1117 BC Schiphol Oost The Netherlands	Correspondence address
Hoofdstraat 52 3972 LB Driebergen The Netherlands	Head office National Police
+31 (0)53 483 33 33 www.politie.nl	Telephone Internet

Document Title	Capacity planning of police helicopters: How to improve and support the yearly capacity planning of police helicopters? *Master thesis for the Master program Industrial Engineering and Management at the University of Twente.*
Date	07-03-2014
Author	Rob Vromans Robvromans@gmail.com +31 (0) 6 363 120 87
Graduation Committee University of Twente	Dr.ir. M.R.K. Mes School of Management and Governance Department Industrial Engineering and Business Information Systems (IEBIS) Dr.ir. J.M.J. Schutten School of Management and Governance Department Industrial Engineering and Business Information Systems (IEBIS)
National Police	Dr. ir. R. Rienks Team Leader
Dutch Police Air Support	E. van den Brink Project manager and pilot

UNIVERSITY OF TWENTE.

POLITIE

Preface

During my study Industrial Engineering and Management I worked on the optimization of supply chains and production processes. I never thought about police helicopters and that planning could improve the performance of a public service like the police force. Now I know better. Optimization of public services does not only save money, it also improves the quality of life for the people who benefit from it, in this case Dutch citizens. This component makes research on public resources interesting but also requires different viewpoints. For example, is the optimal distribution of the public resource a fair one?

This research focuses on the tactical, or medium-term, planning of surveillance flights by police helicopters. It has kept me busy and interested for ten months. Creating crime forecasts and directing police forces is an incredible field of research and I feel that I'm just discovering its potential. I wrote this report to share my current ideas about the tactical planning of police helicopters, and I hope to convince you of its future possibilities.

This report would not be here today without the help of all my supervisors. I am grateful for the amount of support I received. Most of all, I thank Edo van den Brink and Arjen Stobbe, my daily supervisors at the Aviation Police, for their continuous support and our interesting and enthusiastic discussions. Rick van Urk, my predecessor and advisor in this research, was always available for creative ideas, motivation, and feedback. Thank you Rick. I also thank Rutger Rienks, my supervisor of the National Police, for making the research possible. I thank Floris Korteweg and Daniel Bulten, analysts of the National Police, who provided me with data and ideas about helicopter effectiveness, and started interesting discussions.

I thank Martijn Mes and Marco Schutten, my university supervisors, for their valuable feedback and scientific insights, it certainly improved the quality of this thesis.

Finally, I thank my girlfriend Emma, my family, and friends for their continuous support and patience.

Management summary

Motivation
The Dutch Air Support and Aviation Police (LVP) currently has two disconnected decision support systems (DSS) that support strategic and operational decision making. The strategic DSS quantifies the long-term effect of different allocations of police helicopters over bases. The operational DSS supports short-term decisions such as the timing and route of the surveillance flights for the next day, given that the number of flights per day is determined at the tactical (medium-term) level. However, the distribution of helicopter capacity over the year is currently made without support. The LVP likes to bridge the gap between the strategic and operational decisions by using the operational crime forecast to support tactical decisions as well.

Research goals
The goal of this research is to integrate the strategic and operational decision support tools and create a prototype tool that combines crime data and police resources into a complete tactical plan. Furthermore, we want to create a simulation model to validate the current and prior research, and to provide insight for the LVP into the effect of different tactical decisions. Finally, we aim to implement the validated tactical planning methods into the procedures of the LVP.

Forecasting and routing
We extend the current forecasting tool of the LVP to a scope of one year and introduce seasonal and weekly crime patterns. To enable realistic computation times we propose a method that drastically increases the forecasting speed, without losing forecast accuracy. Furthermore, we propose a new heuristic to determine the best start time of a flight and evaluate the current route optimization method.

Tactical planning
We set up a tactical planning model to solve the planning problem under the assumption that there is at most one helicopter airborne at the same time. This model simultaneously schedules surveillance flights and helicopter crews. To relax this assumption without increasing the computation time, we propose a heuristic in which we sequentially schedule flights and crews. The heuristic determines at the tactical level the available resource capacity, and at the operational level the start times and routes of the surveillance flights.

Next to the tactical planning heuristic that optimizes performance, we propose a heuristic that considers the trade-off between performance and an equitable distribution of helicopter capacity over the Netherlands.

To determine the quality of the forecast methods, and the performance and fairness of tactical plans, we propose a set of performance measures for the LVP. For example, we propose to measure the average proximity of helicopters to criminal incidents to determine the quality of surveillance routes.

Results
Based on the simulation results, we conclude that the current operational DSS improves performance by approximately 20%. Furthermore, when the LVP is able to perform a variable number of flights per day and uses tactical planning, we find a performance improvement of at least 3.5%. Finally, we expect that the new estimator to determine start times and procedures to take incident priorities and night time effectiveness into account result in additional performance improvements.

Implementation

The simulation tool that we use to validate all proposed methods and extensions is useful for further analysis of the LVP processes and policies. However, it is not user-friendly enough and thus not ready for direct implementation. However, in cooperation with the LVP we made it possible to implement the resulting tactical plan and inform pilots about the optimal surveillance route during their flight.

Recommendations

We recommend the LVP to professionalize the current prototype decision support system. Furthermore, the system is currently based on an assumption that relates the arrival time of a helicopter to the probability that the helicopter successfully supports police officers on the ground. We suggest that the LVP starts measuring the actual relation between the arrival time and the success probability. This makes the system more reliable. Furthermore, we propose to share the insights from the simulation model and the forecast methods with other police departments. Geographical information systems (GIS) are very useful for this application. Finally, we propose to provide helicopter pilots with an existing rerouting tool to determine the optimal continuation of their surveillance route after they have handled an incident.

Future research

Based on our results and experiences, we find several new research directions. We propose to further investigate more advanced forecast methods and the simultaneous optimization of stand-by crew and flight scheduling. Furthermore, the model could be extended to include other helicopter types and police resources, like patrol cars and officers on foot. In this research we also shortly address the concept of equity. Since equity is crucial in the distribution of public resources, it requires more attention. Finally, since the added value of the police helicopters depends on the incidents that they are deployed to, we further recommend to investigate intelligent deployment strategies.

Table of contents

1 Introduction

Videos of criminals on the run, made out of police helicopters, are shown regularly on TV. Next to pursuing criminals, police helicopters also support police officers on the ground in their search for missing citizens or during the protection of the royal family. The Air Support department of the National Police (the LVP or 'Luchtvaartpolitie") has 8 helicopters to provide support with. Since flying these helicopters is expensive, an annual budget allows only a limited number of flight hours per year. Therefore, the deployment of helicopters for duties such as patrolling deserves attention. Furthermore, the added value of a helicopter at an incident decreases rapidly with the time it takes the helicopter to get to the incident. Therefore, the LVP developed, in earlier research, decision support systems (DSS) based on mathematical models to optimize the positioning of helicopters at airfields and to determine the routes helicopters should fly in anticipation of possible future incidents.

The aim of this research is to determine the effectiveness of the models that the LVP currently uses and to improve their performance. The added value of this research is an integrated decision support system that helps the police make operational and tactical decisions for the distribution of helicopter flight hours over the year.

This first chapter consists of the research plan: we introduce the Dutch Air Support organization and the current planning process (Section 1.1), followed by the motivation for this research (Section 1.2), the problem statement and research questions (Section 1.3), and the research plan (Section 1.4).

1.1 Organization

The introduction of the LVP starts with an outline of the organization of the aviation police (Section 1.1.1) a brief description of the current helicopter planning process (Section 1.1.2), and previously performed research by the LVP (Section 1.1.3). The introduction helps to understand the motivation as discussed in Section 1.2.

1.1.1 Organization

This research takes place at the Dutch Air Support & Aviation Police division of the Dutch National Police (NP). Currently, the police reorganizes its structure from 25 regional police divisions and one national division to an integrated national police with 10 regional units and one national unit. The national unit coordinates and supports regional departments' efforts. Figure 1.1 shows a summary of the planned organizational diagram of the national police. Appendix A contains the complete diagram.

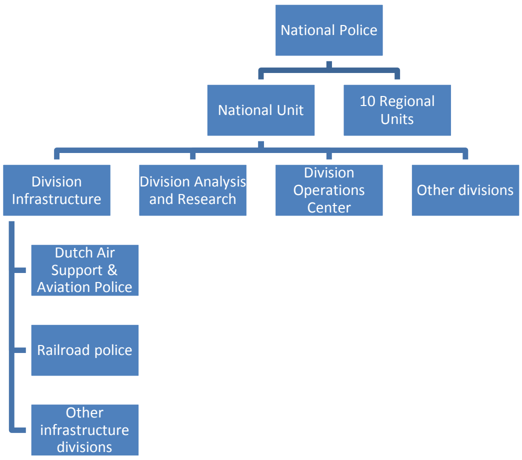

Figure 1.1: Summary of organizational chart of the new national Dutch police organization (Appendix A).

Within the National Police organization, the LVP performs two main tasks: it supports regional police departments in the air (Air Support) and it supervises aviation in the form of, e.g., airport supervision (Aviation Police). This research focusses on the air support task. The LVP consists of the following five air support functionalities:

- **The planning office** schedules flights and creates the monthly and yearly base plan.
- **Flight dispatch** is responsible for flight preparation, support, and completion.
- **The pilot** controls the helicopter during the flight.
- **The tactical flight officer (TFO)** controls the sensors of the helicopter.
- **Maintenance** is responsible to keep the air fleet available for use.

The aim of this research is to support the planning office in its planning activities, since the timing of helicopter flights has an influence on the performance of the flights.

The LVP has six Eurocopter helicopters (EC135) and two Agusta Westland helicopters (AW139), normally located at Schiphol. Although the LVP owns six Eurocopters, there are only five EC135s operational, since there is always one EC135 in maintenance. The EC135 has a cruising speed of 218 km/h, a top speed of 254 km/h, and has a maximum flight duration of two hours (at cruising speed). The AW139 has a cruising speed of 254 km/h, a top speed of 303 km/h, and is able to fly for four hours (at cruising speed). Figure 1.2 shows both helicopter types. The AW139 is better suited for police support in the eastern part of the Netherlands because of its speed and range, while the EC135 is less expensive and thus favourable in police work around Schiphol.

Figure 1.2: Photos of the Eurocopter EC135 (left) and the Agusta Westland AW139 (right).

1.1.2 Helicopter planning process.

The planning office is the part of the organization that aligns maintenance, personnel, and flights. The planning office determines the number of flights per day, the helicopters to use, and which crews need to be present at the base to fly the helicopters. The planning process has the yearly total flight hour budget as input, and results in a daily flight schedule with crews and helicopters assigned to flights, as output. This planning process is according to the hierarchical structure of Hans et al. (2007). This hierarchical structure consists of four levels:

1. Strategic planning: long term decisions such as the number and type of helicopters to use.
2. Tactical planning: medium-term decisions such as the selection of bases to station helicopters on, the scheduling of helicopter crews, and the planning of major maintenance.
3. Operational offline planning: short term decisions before the actual flight, such as the route that the helicopter is going to fly.
4. Operational online scheduling: short term decisions during the flight such as the decision to provide air support at an incident or not.

The decision levels indicate the degree of freedom in the decisions and the impact they have on other decisions. For example, on the tactical level the LVP determines the number of flights per day. This restricts the options for the operational offline schedule, since for that schedule the number of flights is known. We now simultaneously discuss the hierarchical planning structure and the LVP planning process to indicate to which hierarchical level the planning activities belong. Figure 1.3 shows the planning process and planning hierarchy.

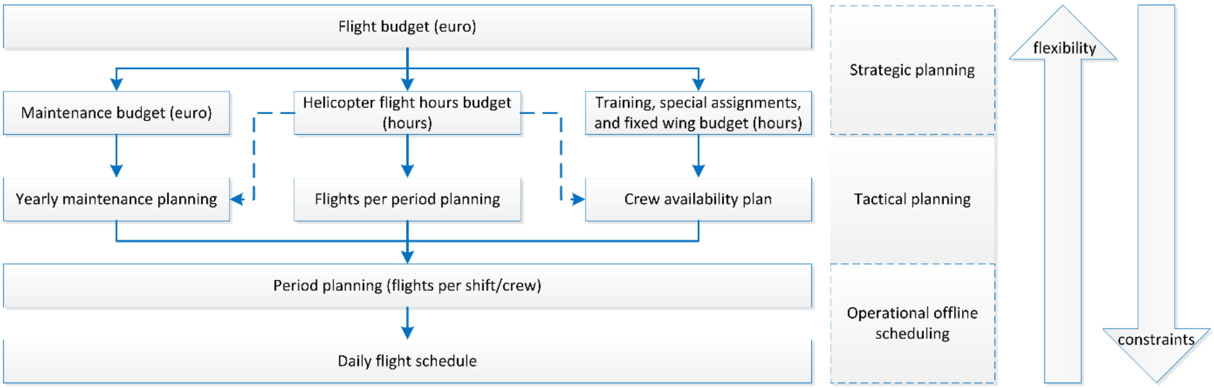

Figure 1.3: Planning process hierarchy.

For the LVP, the strategic planning process has two parts: the total budget that is available for flights and the decision on which bases to use. The LVP has a support contract with their helicopter suppliers to provide them parts for a fixed cost per flight hour. The LVP then has a fixed budget that enables a fixed number of flight hours per year (including maintenance). The budget is not completely available for surveillance flights; it includes also flight hours for training, airplane flights, and special assignments. For the location decision there are several options: the LVP has a hangar and office space on Schiphol and is able to hire the same facilities on Rotterdam and Volkel. Other Dutch airfields can be used to refuel and rest, but the LVP has no facilities there.

On the tactical planning level, the LVP makes three decisions. First, the maintenance department determines an annual maintenance plan, based on the total number of available operational flight hours. The planning office then determines the number of flights per period of 4 weeks, of which there are 13 per year, and finally determines how many crews should be available. The decisions on this planning level are constrained by the flight hour budget from the strategic level and constrain the scheduling on the operational level. The advantage of the tactical planning over the operational scheduling is the increased flexibility of crews and the maintenance department.

On the operational offline level, the number of flights hours per period is known. The planning office then makes a “period planning” in which they take the personnel constraints into account: crews can be unavailable due to scheduled training exercises (the training instructors are active pilots), simulator training trips, paid leave, and general meetings. Finally, the planning office assigns helicopters to flights while taking the maintenance activities into account.

Flight dispatch and helicopter crews make the operational online decisions. For example, flight dispatch cancels flights when the weather conditions are not good enough and helicopter crews decide when they start their flight. When helicopters are airborne, the Operations Centre of the National Police deploys them to incidents to provide support. The online decisions have an impact on the offline scheduling and tactical planning: the planning office can reschedule cancelled flights within the current period or move them to a different month.

1.1.3 Previous research

Previously, the LVP worked together with the University of Twente to develop tools to support decision making based on historical data. The research focused on different planning levels.

The research of Buiteveld (2011) focuses on the strategic planning level. Buiteveld presents a tool that supports the planning department in the decision which bases to put helicopters on, to cover as much incidents as possible. The tool has a positive effect on the performance.

The following research by Van Urk (2012) focuses on the operational planning level and proposes a tool that covers both offline and online scheduling. The tool is programmed in AIMMS (a modelling system for mathematical optimization) and makes a forecast for one day ahead, based on all historical data available, on which it bases the offline and online decisions. For offline scheduling, the tool helps planners to estimate the optimal start time of a flight within a crew schedule, and to determine the route of the flight. For online scheduling, the tool allows planners to optimally reroute a helicopter to an incident during its flight.

1.2 Motivation

The LVP now has one decision support tool to help allocate helicopters to bases and a second decision support tool to determine the optimal route to fly on a given time and from a given location. However, it is not clear how to connect the systems and there is no support on decisions such as how many helicopters to use or how many flights to perform on a given day. Therefore, the LVP now uses the model by Van Urk (2012) with four simple planning rules in general:

1. The flight budget is evenly distributed over all periods (flat planning).
2. Every day has three shifts of which the late shift is manned by two crews. The rest of the shifts are manned by one crew.
3. Every crew makes at least one flight during the shift.
4. The crews all fly from Schiphol.

The flat distribution of helicopter time over the days of the year does not seem appropriate when we compare it with Figure 1.4, which shows that the number of criminal incidents per day is relatively low in the summer when compared to the winter. This effect is commonly explained by the period between sunset and sunrise, which is longer in the winter period. More dark hours result in decreased sight and less possible witnesses on the street. Furthermore, Van Urk (2012) notes that besides the seasonal pattern between months, there is also a weekly pattern between days of the week.

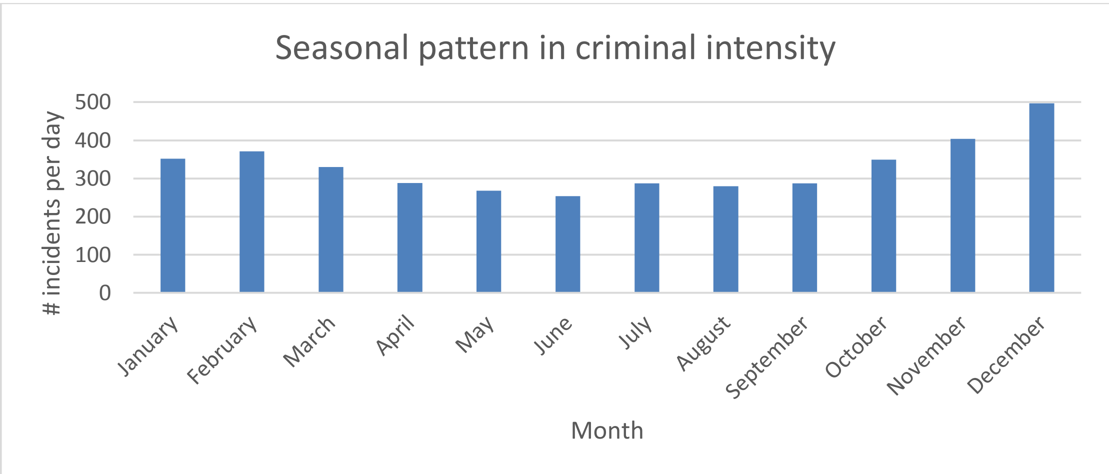

Figure 1.4: Number of incidents per day for every month in 2012.

The flat planning results in several problems and inefficiencies. First, it seems inefficient that the LVP flies the same number of hours during all periods, while there are less criminal incidents in summer time than in the winter period. By shifting more hours to the winter, the number of successful assists is likely to rise. Second, there are fewer pilots available during the summer period than during the winter period due to holidays, resulting in capacity problems.

The LVP uses three overlapping 9-hour shifts to provide 24/7 helicopter coverage from Schiphol. On every shift, all crews fly at least once. However, as Figure 1.5 shows, the intensity of incidents varies during the day. This figure raises the question whether it is efficient to fly during every shift and to have 24/7 availability of air support.

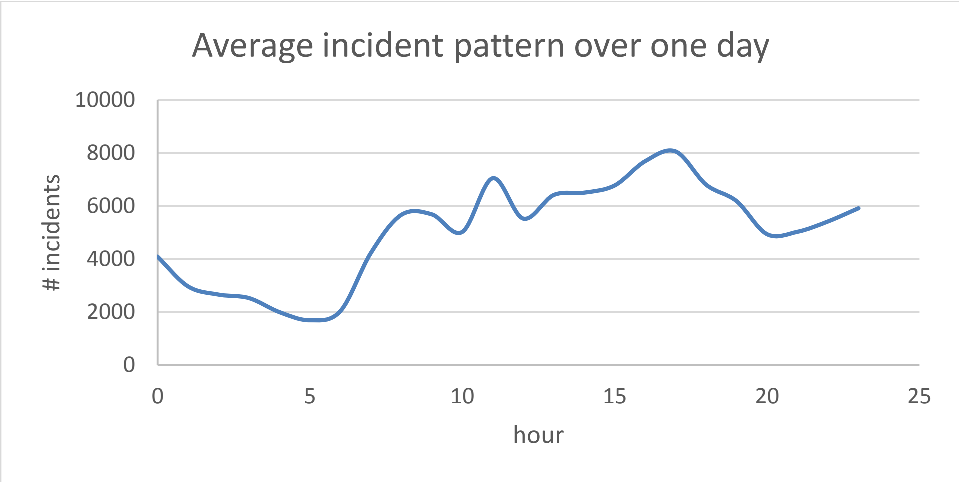

Figure 1.5: Number of incidents in 2012 per hour of the day.

Finally, there is also a discrepancy in incidents over the country. Figure 1.6 shows that crime is concentrated in the west and south-west of the Netherlands and there are fewer incidents per year in the rural areas in the east and north. The distribution of crime in Figure 1.6 is consistent with the theory by Sherman et al. (1989) of "hotspots", of which the main point is that crime clumps in relatively small places (that usually generate more than half of all criminal incidents) and is totally absent in others.

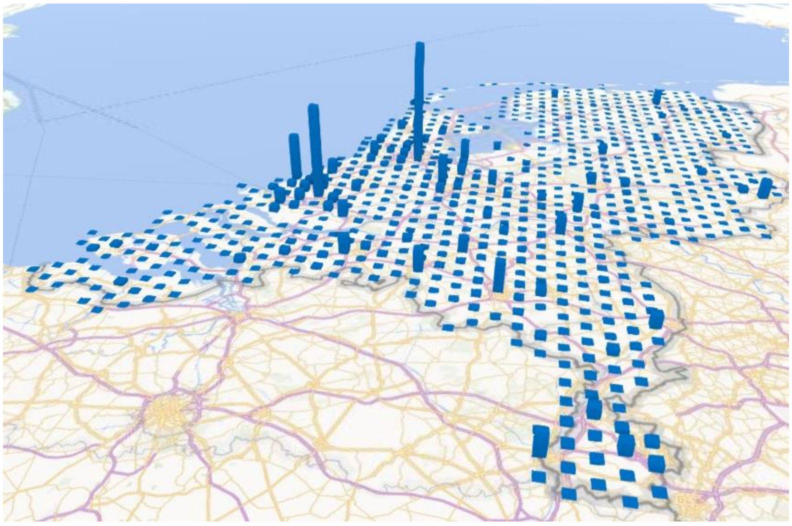

Figure 1.6: Geographical distribution of incidents in 2011 and 2012.

There are two clear hotspots, Amsterdam and Rotterdam. Optimal helicopter routing is not the only possibility to cope with this geographical distribution of incidents. The LVP has access to multiple airfields to land on and start from. Since the LVP has several helicopters, it is able to distribute its helicopters over multiple locations. The start and end locations of the surveillance helicopters thus deserve attention in the operational scheduling of flights. This shows the necessity of the link between the tactical helicopter positioning model and the operational routing model.

The LVP has an annual budget that limits the number of helicopter flight hours. Furthermore, the use of helicopters is also limited by the availability of helicopter crews with pilots and tactical flight officers (TFO's) and by the regular maintenance that helicopters require. With the limited budget, the LVP wants to maximize its performance. Therefore, The LVP would like to use an instrument that helps them to make tactical decisions, which are:

1. How to distribute helicopter flight capacity over the days of the next year?
2. How long should the flights be?
3. Where should the flights start and end?
4. Where should the LVP station the helicopters (per month)?
5. How to handle cancelled flights?
6. When should shifts start and how many crews should be available per shift?

1.2.1 Targets

The LVP has several targets. As part of the national police organization, it shares the common police goal to lower the overall crime rate. The LVP focusses on two types of High Impact Crimes (HIC): burglaries and (street) robberies. Since the time it takes to fly to an incident is a factor in the expected added value of the helicopter, the LVP aims to increase its proximity to these crimes. However, there is need for a balance between proximity and costs: since flying is expensive, the LVP aims to respond to as many incidents as possible from their stand-by situation. Surveillance flights should aim to cover the maximum number of incidents that cannot be covered from stand-by.

When we use the maximization of incidents covered as the goal of the tactical planning model, we focus on the places where most criminal incidents happen. This can lead to a situation where the entire helicopter capacity is divided over a small number of cities. However, fair allocation of public services, or "equity", is a critical and controversial factor when deciding how to allocate public resources (Stone, 2002). Currently, there is no definition of a fair or equitable distribution of helicopter capacity over the Netherlands. Therefore, the LVP wants to know how to translate equity to the tactical planning of helicopters, what the effect of optimization is on the equality of the helicopter support distribution, and how the LVP can influence the equality.

Furthermore, the planning department has a secondary objective. The LVP wants to minimize the number of flights that start in one shift and end in the next shift, since these flights require crews to fly in overtime.

1.2.2 Scope

Since there is limited time available for our research, we define a strict scope. This research adopts the forecasting and routing techniques as suggested by Van Urk (2012) and extends them to enable tactical planning. We use the existing AIMMS tool and adapt the software to make forecasts and plan multiple days ahead. We use the program to determine optimal flight times, durations, origin and destination, and routes for the surveillance flights. We also perform an analysis on the effect of the shift schedule on the expected overtime. Finally, we analyse the trade-off between performance and fairness of the distribution of helicopter support over the Netherlands.

1.2.3 Research goals

The first goal of this research is to validate earlier research results and the existing tools. The second goal of this research is to develop a prototype instrument that supports the LVP in its tactical decision making process based on historical data. The improved distribution of flights over months and days should lead to an increase of the number of incidents that the helicopter can reach in time. Furthermore, the instrument should provide insight into the effect of tactical decisions on the performance of the LVP. Finally, we implement the decision support system in the LVP organization.

1.3 Problem statement and research questions

The main problem of the LVP planning department is the lack of an integrated decision support system that determines the daily number of flights and their departure times on a tactical level. The main problem leads to the following problem statement:

What is the performance of the operational decision support tool and how can the LVP improve its performance by tactical planning?

To find a solution to the described problem, we answer the following research questions:

1. What is the current planning process at the LVP? (Chapter 2)
   Chapter 2 describes and analyses the current planning process and the tactical planning at the LVP in detail. Furthermore, it discusses the operational planning tool and draws conclusions on requirements for the final prototype tool of this research.
2. What literature is available related to the decisions of the LVP (Chapter 3)
   Chapter 3 positions this research in the existing literature and describes research fields that are relevant. By answering this question, we want to find techniques and models to improve the current system.
3. How can we extend the scope of the decision support systems and improve their performance? (Chapter 4)
   Chapter 4 discusses how we propose to adapt the current decision support system to provide input for the tactical planning.
4. What is the best approach to tactical planning for the LVP? (Chapter 5)
   Chapter 5 defines the tactical planning problem and explains several approaches to handle the problem. Furthermore, we discuss how we can incorporate fairness into the tactical planning method.
5. What are appropriate measures for the performance of the decision support systems of the LVP? (Chapter 6)
   Chapter 6 discusses measures for the quality of a forecast and the performance and fairness of a tactical plan.
6. What is the effect of different settings of the decision support system and what is the expected impact of the system on the LVP performance? (Chapter 7)
   Chapter 7 determines the added performance by the previous research. Furthermore, we determine the added value of tactical planning and the effect of several extensions on the tactical plan.
7. How to implement the tactical planning model in the LVP organization (Chapter 8)?
   Chapter 8 discusses the implementation process of the prototype tactical planning tool in the LVP organization.

1.4 Research plan

We start the research with interviews with planners and pilots to determine the current planning process. We then perform literature research to identify techniques and models to improve the current decision support system. Next, we extend the scope of the application and apply the techniques from literature to improve the system. We then use simulation to test and compare several tactical and operational planning policies for the LVP. Finally, we cooperate with the planning office and pilots to integrate the decision support system in the current planning process and to use the results in practice.

To be able to draw conclusions from the simulation results, the simulation model should represent reality well. Law (2007) notes that model programming is just part of the overall effort to design or analyse a complex system by simulation, and proposes steps that will “compose a typical, sound simulation study”. Figure 1.7 shows these numbered steps and this research organizes the simulation study accordingly.

Chapter 2 discusses the problem and study plan of step 1. For step 2, the LVP collects the data that is required as input for the simulation model that we propose. To construct a simulation model we then perform literature research in Chapter 3, discuss forecasting techniques in Chapter 4, create a tactical planning model in Chapter 5, and propose performance measures in Chapter 6.

Steps 3 to 6 are about verification and validation. Verification is the process to make sure that the computer program works well. Validation makes sure that the simulation model is an accurate representation of the system that we model. Chapter 7 discusses steps 7 to 10. We set up a simulation plan that defines the experiments we need to perform to underpin all choices in this research, and make the production runs. In step 9 we analyse the output data with the performance measures as discussed in Chapter 6. Finally, we present the results and conclude on the implications of the results for the LVP.

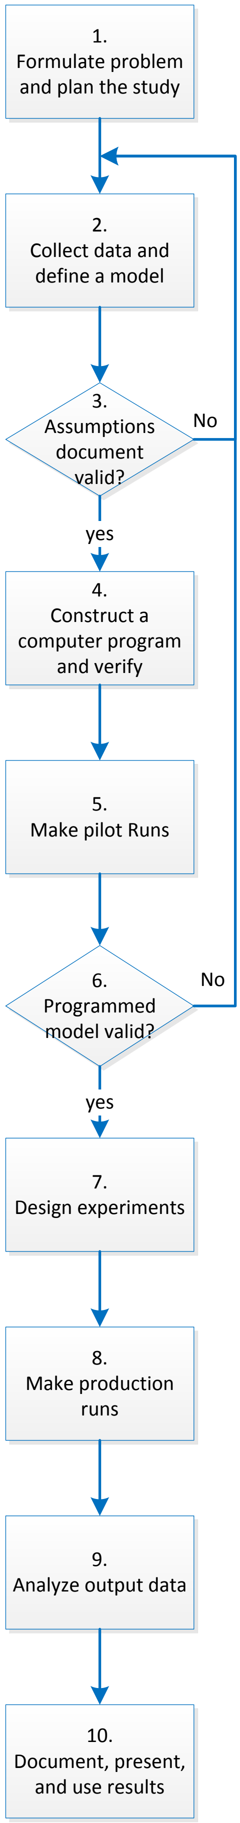

Figure 1.7: Steps in a simulation study (Law, 2007)

2 Current situation

This chapter describes the current planning process. The description of the planning process is necessary to understand the use and requirements of the decision support systems. As Chapter 1 explains, the tactical planning process consists of several steps and procedures. This section explains how the planning office performs these steps. Since the LVP budget is fixed, we consider the strategic planning of the LVP as fixed and thus start with an analysis of the tactical planning level in Section 2.1. Section 2.2 explains the operational scheduling process, Section 2.3 introduces the operational scheduling and positioning tool, and Section 2.4 discusses how the LVP measures its performance. Section 2.5 describes the desired planning situation by the LVP. Finally, Section 2.6 provides conclusions on the current situation.

2.1 Tactical planning

This section discusses the tactical planning process. First, we explain how the maintenance planning is set up and kept. Second, we discuss how the helicopter flight hour budget is used as input for the flights per period planning. Third, we discuss the crew scheduling characteristics.

2.1.1 Maintenance planning

The maintenance department makes an annual maintenance planning around October the year before. This planning consists of two types of maintenance: calendar based maintenance and flight activity based maintenance. Calendar based maintenance happens independent of the state of the helicopter and includes for example yearly maintenance: every year several parts are replaced. Flight activity based maintenance is performed after the helicopter has made a given number of flight hours after its last maintenance activity.

The maintenance department uses the flight hour budget as input and plans flight activity based maintenance as if it is calendar based maintenance: they determine a fixed date to perform the maintenance on. To make sure that the maintenance capacity is used efficiently, the planning office tries to make sure the helicopters are ready for maintenance in time. This means that when a helicopter is scheduled for its 100 hour flight time maintenance within one week, they make sure that the helicopter makes enough hours this week to reach the 100 hours, even when the helicopter has to fly on a Sunday morning, where the probability of a successful helicopter deployment is minimal. This means that the flight planning is currently maintenance driven, although maintenance in general is a support function.

2.1.2 Flights per period planning

Currently, the LVP has no system to support the planning office with the analysis of the optimal number of flights per period. Therefore, the planning office distributes the flight hours uniformly over the periods: every period gets the same number of flights. The advantage of this method is that personnel planning is very predictable and that, e.g., maintenance is also able to use a flat planning.

2.1.3 Crew planning

Helicopter personnel capacity is the main bottleneck in the tactical planning of the LVP: there is enough helicopter capacity to cover the budget and there are enough helicopters to cope with material breakdowns, but it is hard to find enough pilots to man all shifts.

The planning office uses manually updated excel sheets to determine the number of available pilots and TFOs per day. It then fills up the shifts along their priorities. Table 2.1 shows the current shift schedule. The crew composition changes between shifts and seasons: between sunset and sunrise, the crew should consist of 2 pilots. When the entire flight is during daylight, one pilot is enough. Furthermore, a crew always consists of a TFO to operate the on-board camera.

Since 2010, the LVP has agreed with the Ministry of Security and Justice to supply 24/7 coverage of emergencies from Schiphol. Therefore, the first priority of the planning office is to make sure that Schiphol has a crew available for emergency response 24/7. With 3 shifts, the planning office needs 3 crews per day to meet this requirement, or approximately 1,095 crews per year. We note that there is no crew stand-by on Schiphol, when the crew from Schiphol makes a surveillance flight. Since the LVP has personnel for approximately 1,500 crews, there are crews left. The LVP currently stations the extra crews on Schiphol to be able to simultaneously react to multiple incidents, but Rotterdam and Volkel are also possible stand-by locations.

The distribution of crews over the three LVP locations is based on the research of Buiteveld (2011). Buiteveld (2011) provides a static analysis of where incidents take place in the Netherlands and on which location a stand-by helicopter can reach most of these incidents. Buiteveld concludes, based on a limited dataset, that out of the available helicopter fields in the Netherlands, Rotterdam Airport is the best location to provide stand-by coverage from. The second location is a heliport near Amsterdam, third is Hilversum Airport, and fourth is Air Base Volkel. However, the LVP currently has its main facilities at Schiphol, and complementary facilities on Rotterdam and Volkel. Since the research was based on a limited dataset, the main base of the LVP is currently still Schiphol. Furthermore, since the LVP has no facilities on Hilversum Airport, Rotterdam and Volkel are the second and third choice to position stand-by helicopters on respectively.

Therefore, when there is personnel capacity left after scheduling 24/7 coverage on Schiphol, the planning office first schedules the shifts with the second priority: a crew that flies from Schiphol to Rotterdam, makes a surveillance flight from Rotterdam, and returns to Schiphol at the end of the shift. When there is still a crew left, then the 3rd priority is to schedule a crew on a flight to, around, and from Volkel. Table 1 gives an overview of the current shift schedule and the priorities.

Shift	Start time of shift	End of shift	Priority
Early	07:00	16:00	1
Late	15:00	23:00	1
2nd late crew to Rotterdam	17:00	02:00	2
3rd late crew to Volkel	17:00	02:00	3
Night	22:30	07:30	1

Table 2.1: Current shift schedule of the LVP.

Because most of the crews available to the LVP are mandatory scheduled on Schiphol to create 24/7 coverage, the LVP wants to optimize the allocation of the crews that are left, over all shifts and locations. There is one constraint: the LVP cannot schedule more crews than there are helicopters available during the shift.

When a crew is allocated to a shift on a different location than Schiphol, e.g., Rotterdam, there are two possibilities to get there. When there are only sporadic flights from Rotterdam, then the LVP keeps the helicopters at Schiphol. Crews then start at Schiphol and fly to Rotterdam at the start of their shift. From Rotterdam, they then make one or more surveillance flights. At the end of the shift, they fly back from Rotterdam to Schiphol. However, when flights from Rotterdam or Volkel become frequent, then the LVP is able to station a helicopter there. The helicopter can stay for weeks at Rotterdam or Volkel, before it is flown to Schiphol for maintenance. It is then possible to schedule crews at Rotterdam or Volkel that start and end on that location and can be stand-by during their entire shift.

2.2 Operational helicopter scheduling

This section explains how the planning office makes the operational helicopter schedule and how the planning office determines when, how long, and where to fly. Due to a lack of up-to-date data to provide input for the tool, and the lack of knowledge on the working of the tool, the tool is not yet used in a daily routine at the planning office. Therefore, the planning office is only able to tell crews how many flights they have to make during their shift and with which helicopter.

2.2.1 Crew scheduling

The planning office makes a final crew schedule at the start of a period and registers the flights with the designated crew and helicopter in its planning software. A crew can fly a maximum of 5 hours per shift and is allowed to rest for ¼th of the flight duration between flights. Furthermore, crews need the first hour of their shift to collect information and the last hour of their shift to enter their experiences during the flights into a log. These practical issues limit the number of flights per crew.

2.2.2 Flight duration

Currently, the planning office does not advice the pilots on the optimal duration of their flights. Since pilots want to catch as many criminals as possible every time they fly, pilots tend to fly as long as possible. However, this is sometimes inefficient and can lead to problems: when there is an emergency request just before the end of the flight and the helicopter has to land first to refuel, this will cost the LVP a successful assist. Furthermore, when the helicopter flies shorter flights in the low-season of criminal incidents, then there is more time left for the high-season.

2.2.3 Offline: timing and routing

Without the helicopter positioning tool, there is little guidance for crews on when to start their flights within the limits of their shift, and how to determine the flight route. Flight dispatch briefs the crews on weather conditions and some police intelligence. The crews that are allocated to Rotterdam or Volkel start their shift by flying to that location. The LVP has determined the general daily crime pattern and schedules routes on the expected peak times. When crews are at their start location, they start their surveillance flight at these peak times, and fly a route over the Netherlands guided by their personal knowledge on crime patterns and hotspots.

2.2.4 Online: rerouting

There are several reasons to change the offline made schedules during the day. Flights are sometimes cancelled due to weather conditions (fog, snow, hail) that limit the pilot's visibility or influence the safety of the helicopter. This happens more often in the winter period, when there is more criminal activity, than in the summer period. These flights are currently added later on in the period to make sure that maintenance budgets are kept.

Currently there are no in-flight route changes because pilots do not fly a predefined route. When the LVP starts flying optimal routes in 2014, in-flight route changes are the goal of the LVP surveillance flights: routes are interrupted by reactions to incidents. Since the LVP wants to be able to respond to as many incidents as possible, we could therefore say that the optimal route is the route that is interrupted as soon and as often as possible.

In-flight route changes are to be made by the central operations room in Driebergen that handles all emergency calls of the Netherlands. When a call requires police attention and includes a high impact crime, then the helicopter deployment protocol determines whether a helicopter is useful at the scene or not. Helicopters are useful when there is a description of the criminal or getaway vehicle and when the direction of travel is known. When a helicopter is useful at the crime scene then Operations contacts the LVP to check whether a helicopter is available (airborne or stand-by). When there is a helicopter available at a close enough distance to the incident, then the Operations room authorizes deployment.

When Operations deploys the helicopter to an emergency, the helicopter commonly flies straight to the incident location. However, in certain cases such as robberies near the Dutch borders, crews know from experience that the criminals will probably try to cross the border, and thus are able to anticipate by flying directly to the nearest border crossing.

Although the tool is not in current use, from now on we assume that the planning office does use the tool on a daily bases, and that pilots fly the routes at the times given by the tool. This enables us to compute the performance of the current situation and compare it with the performance after the improvements by this research and thus determine the added value of this research. The next section discusses how the tool currently works.

2.3 The helicopter positioning and routing tool

The helicopter positioning tool by Van Urk (2012) helps the planning office to determine the time to deploy a helicopter and where to send it. Section 2.3.1 discusses the first step: making a criminal intensity forecast. Section 2.3.2 explains how the tool estimates the best start time for a flight, and Section 2.3.3 describes how the program calculates the optimal route.

2.3.1 Forecasting

The forecasting module of the planning tool uses historical data on criminal incidents to create a forecast. This forecast shows the expected criminal intensity for every part of the Netherlands for one day in advance. Figure 2.1 is an example of the expected geographical distribution of incidents at a given time, where a dark red colour indicates a higher intensity than at the yellow parts of the Netherlands.

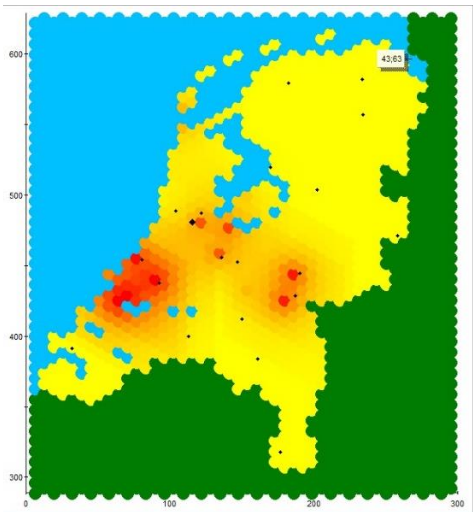

Figure 2.1: Visualization of geographical forecast of incidents at a given time.

As input for the forecast, the tool uses data on criminal incidents in the Netherlands. The tool currently requires four types of data per incident:

- Date of the incident: to determine the weekday and month of the incident.
- Time of the incident.
- Zip code of the incident: to determine the location of the incident.
- Priority of the incident: to distinguish between multiple types of incidents.

Van Urk (2012) notes that “a forecast based on fewer incidents has a higher forecast error”. Since the forecast is made for every 2 minute interval of the day ahead and there are thousands of locations in the Netherlands, Van Urk (2012) introduces several levels of data aggregation to produce a reliable forecast. Appendix C discusses the alterations to the tool we had to make to make the tool create the required output.

Geographical data aggregation: Hexagon grid

Van Urk (2012) proposes a grid of hexagons to combine incidents in multiple streets or districts. Figure 2.2 shows the hexagonal grid laid over the Netherlands. The forecast is thus made for multiple locations $l \in L$ and over two-minute time intervals $t \in T$.

Van Urk (2012) then summates all incidents per hexagon and discounts old data with a forget factor. The forget factor α determines the number of months between the incident and the forecasted time interval and discounts the data by the factor $(1-\alpha)^{months}$. The result of the aggregation per hexagon is a geographical pattern of incidents over time.

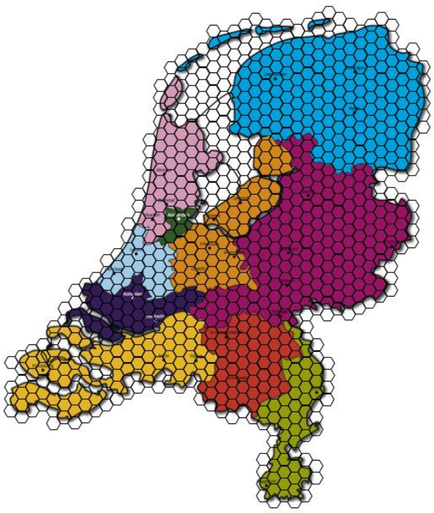

Figure 2.2: The hexagonal grid over the Netherlands by Van Urk (2012).

Yearly data aggregation: Weekday and month distribution factor

The forecasting model currently uses all available historical incident data to determine the forecast for 24 hours ahead. As said, for every hexagon it summates all historical incidents, multiplied by their priority. Van Urk (2012) shows that the timing of criminal peaks depends on the day of the week and the month of the year. Figure 2.3 shows the differences in distribution of incidents over the day for different months.

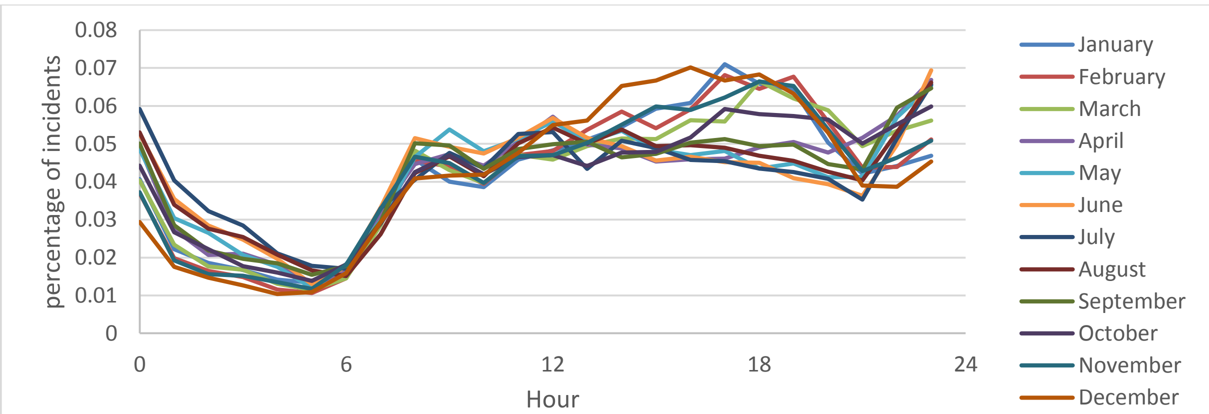

Figure 2.3: Average hourly incident distribution per month in 2010, 2011, and 2012.

Van Urk (2012) proposes to use weekday and month distribution factors to convert data from different months and weekdays to fit the distribution of incidents over the day of the target weekday and month. The month factor converts the distribution of incidents from another month than the target month to the distribution of incidents over time of the target month. The same works for the weekday factor. Figure 2.4 shows the method with which we convert a forecast of a Saturday to a forecast for a Monday. For every hour of the day, we determine the percentage of the incidents of that weekday or month that happens in that hour. We then multiply per hour-block all incidents by the conversion factor per hour h as in Equation 2.1. The conversion factor converts in this case the distribution of incidents over an average Saturday to the distribution of incidents over an average Monday.

$$conversion\ factor_h = \frac{target\ day\ distribution\ factor_h}{day\ distribution\ factor_h} \qquad \text{Equation 2.1}$$

In the fifth hour-block (5:00 AM to 5:59 AM) one percent of the incidents on Saturdays take place, while on Mondays 1.54 percent of the incidents takes place between 5 and 6 A.M. Therefore, we multiply the incidents that happened on Saturdays between 5 and 6 A.M. by the conversion factor 1.54. For hour-block 17 we multiply the incidents from Saturdays by the conversion factor 0.46.

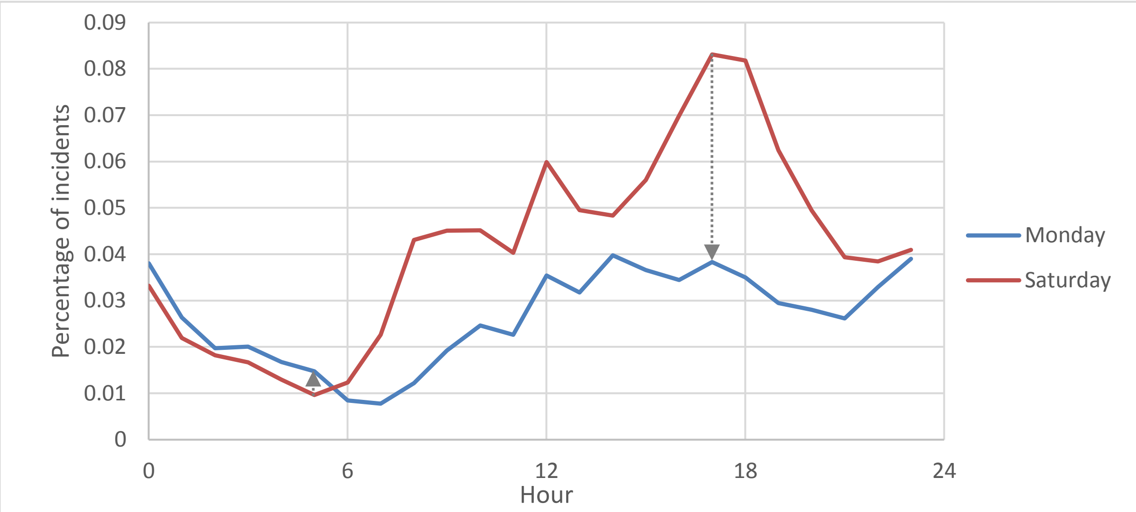

Figure 2.4: Example of conversion of forecast for a Saturday to a forecast for a Monday.

For the month factor, Van Urk (2012) uses the formulation as in Equation 2.2:

$$factor\ month_{m_t,m_i,h_i} = \frac{\left(\frac{N_{m_t,h_i}}{\Sigma_h(N_{m_t,h})}\right)}{\left(\frac{N_{m_i,h_i}}{\Sigma_h(N_{m_i,h})}\right)}, \qquad \text{Equation 2.2}$$

with the following notation:

- m_t Month for which the forecast is made (target month)
- m_i Month of the incident to be converted
- h_i Hour of the incident to be converted
- N_{m_t,h_i} Number of incidents in month m_t during hour h_i
- N_{m_i,h_i} Number of incidents in month m_i during hour h_i

Equation 2.2 has the same form as Equation 2.1, since the target day distribution factor is divided by the distribution factor of the day that is to be converted. The formula for the factor weekday is similar to the formula for the factor month, by replacing month with weekday in the description and notation above. Van Urk (2012) proposes to multiply every incident with its month and weekday factor. An incident from a Tuesday in June could thus be transformed to an incident on a Monday in January.

Spatial aggregation

To be able to create a forecast for all locations in the Netherlands, Van Urk (2012) assumes that the number of incidents in a location, or hexagon, have a predictive value for the number of incidents to happen in the surrounding hexagons. Therefore, Van Urk (2012) determines the values of neighbouring hexagons by adding to them a fraction of the value of the hexagon the incident is in. For the fractions, Van Urk (2012) proposes to use a quadratic function based on the distance to the surrounding hexagons. Equation 2.3 shows the formulation of the spatial aggregation fraction, where $distance_l$ is the number of steps from the hexagon where the actual incident is in and hexagon l.

$$spatial\ aggregation\ fraction = \frac{1}{(distance_l + 1)^2} \qquad \text{Equation 2.3}$$

The formula results in rings of hexagons with the same fraction around the centre hexagon. Van Urk (2012) limits the number of rings to five. Figure 2.5 shows an approximation of the proposed values.

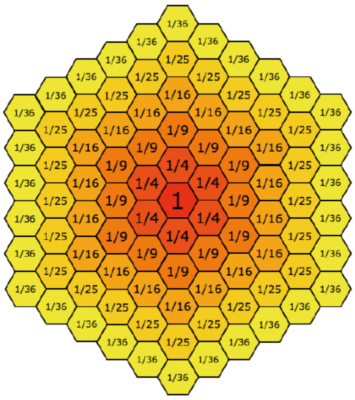

Figure 2.5: Approximate effect of an incident on its neighbouring area (Van Urk, 2012).

Temporal aggregation

Analogous to the assumption that incidents in a hexagon have a predictive value for incidents in surrounding hexagons, Van Urk (2012) assumes that incidents also have predictive value for the times around the time in which they actually happened. He proposes to represent all forecasted incidents in the previous spatial aggregation by normal distributions, with a likelihood of 95% that the incident will happen within thirty minutes before and after the forecasted incident time. Figure 2.6 shows the total forecast value of all locations in the Netherlands for one day of forecast, before and after the time aggregation step. We conclude that the temporal aggregation procedure results in a smoothened forecast over time.

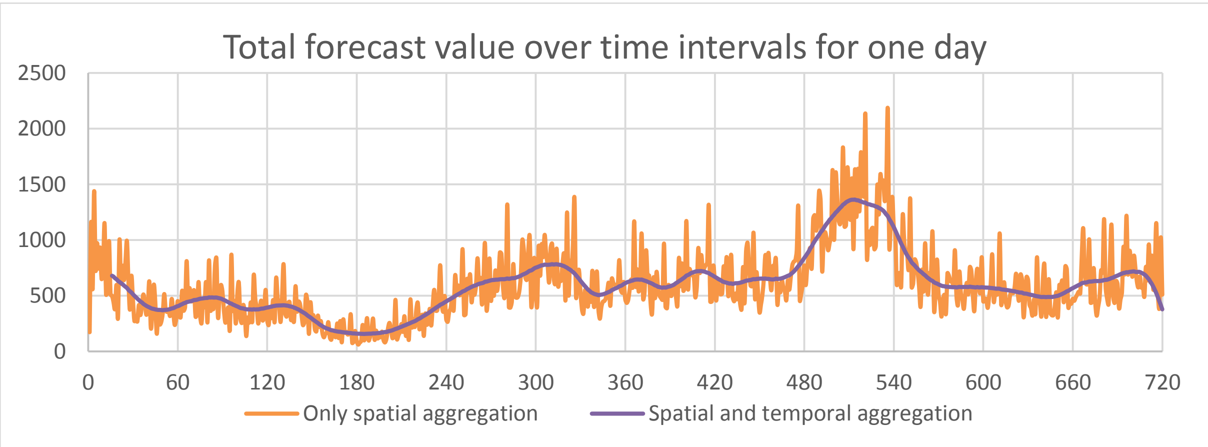

Figure 2.6: Total value of forecast over the Netherlands per time interval, with the new temporal aggregation factors.

Output

The output of the forecasting module is a forecast for one day ahead. The forecast consists of a unit-less *forecast value* $f_{l,t}$ for every hexagon $l \in L$ for every two minute time interval $t \in T$. The value is unit-less because of the aggregation steps but is determined only by the size of the dataset that we use as input. Since it is unit-less, the relative height of the value in comparison with other times and hexagons determines whether it is part of a hotspot or not.

The current limitation of the forecast procedure to 24 hours ahead has a drawback in the estimation of the route quality. For example, since the forecast stops at 23:59 it is never beneficial for a helicopter to start at 23:30 with a 1.5 hour flight since 1 hour of the flight will be outside the forecast and thus will yield no result. Therefore, the tool will not schedule helicopters during the last 1 to 1.5 hours of the day.

2.3.2 Routing

With a forecast and start time available, the tool determines the optimal surveillance flight route. The central concept in the optimization of surveillance routes is *coverage*. An incidents is covered by the LVP when the arrival time of the closest helicopter to that incident is short enough for the helicopter to provide support. When it would not matter how long police helicopters take to get to an incident, then there would be no use of route optimization or even flying. Helicopters could react from stand-by and be at every incident in time. However, to aid the ground forces in finding and following criminals, arrival time is critical. Buiteveld (2011) determines the relation between arrival time and the probability that the police helicopter can successfully support ground forces by interviewing LVP experts. She finds that helicopters are considered always successful when they are within 10 minutes at the location of the incident. When helicopters take between 10 and 15 minutes to arrive to the incident location, then the probability of success decreases. Helicopters that arrive after 15 minutes are never successful. In this case, we call 15 minutes the maximal coverage distance. Equation 2.4 shows the success function as formulated by Van Urk (2012), based on this description, where x is the arrival time in minutes.

$$f(x) = \begin{cases} 1 & 0 \leq x \leq 10 \\ 1.2 - 0.2 * 2^{\frac{x-10}{2}} & 10 \leq x \leq 15 \\ 0 & x > 15 \end{cases} \qquad \text{Equation 2.4}$$

We note that this function is continuous at $x = 10$ but discontinuous at $x = 15$ since $1.2 - 0.2 * 2^{\frac{15-10}{2}} \approx 0.069$. Figure 2.7 shows a graphical representation of the success function.

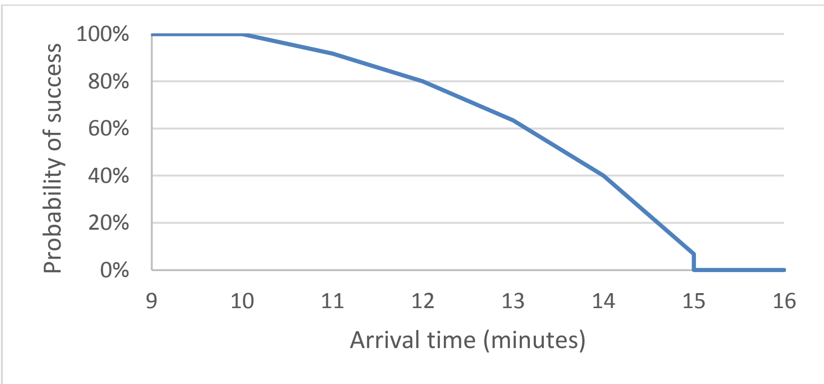

Figure 2.7: Graphical representation of the relation between arrival time and the estimated success function of a helicopter (Van Urk, 2012).

In this research, we define the coverage area of a helicopter as all locations that are covered to some extent by the helicopter. Since the helicopter can fly in every direction, the coverage area of a helicopter would have the shape of a circle. In the case of the hexagonal grid we determine that incidents are covered in all hexagons of which the centre can be reached within 15 minutes. We can thus determine the coverage $Coverage_{l,t}$ of a helicopter on a location $l \in L$ at a given time $t \in T$ by first determining which locations are (partly) covered by the helicopter. When the probability of successful support at a location by a helicopter is 80% then we assume that every incident at that time and location is covered for 80%. The coverage of that location by the helicopter is then the sum over all incidents, of the coverage of the incidents. For every location we thus multiply the number of incidents at the given time interval with the probability of a successful support at that location. The helicopter coverage is then the total of the coverage of all (partly) covered hexagons. Since a route is a combination of helicopter locations in time, the total route coverage of a route is the sum over the flight duration of the helicopter coverage.

The helicopter routing tool uses a Mixed Integer Linear Program (MILP) that maximizes the number of *forecasted* incidents in the hexagons that the helicopter can reach in time to provide support, during a surveillance flight. Equation 2.5 shows the objective function of the proposed MILP.

$$maximize \sum_{l,t} f_{l,t} * Coverage_{l,t} \qquad \text{Equation 2.5}$$

Since hotspots move over time and have a restricted lifespan, the result of the MILP is often a route that visits multiple hotspots after each other. Figure 2.8 shows an example of a helicopter that moves through time and space to visit forecasted hotspots. When the hotspots do not move during a flight, the optimal route can include time periods in which the helicopter visits the same location multiple times after each other. This implies that the helicopter should hover above a hexagon for multiple minutes. Since this creates noise problems for residents in that area and hovering is uncomfortable for pilots, Van Urk (2012) introduced a constraint that limits the number of times that a helicopter may visit a hexagon during one flight. We note that the maximum number of visits should always be more than one, when the helicopter has to start and land at the same location (e.g. Schiphol).

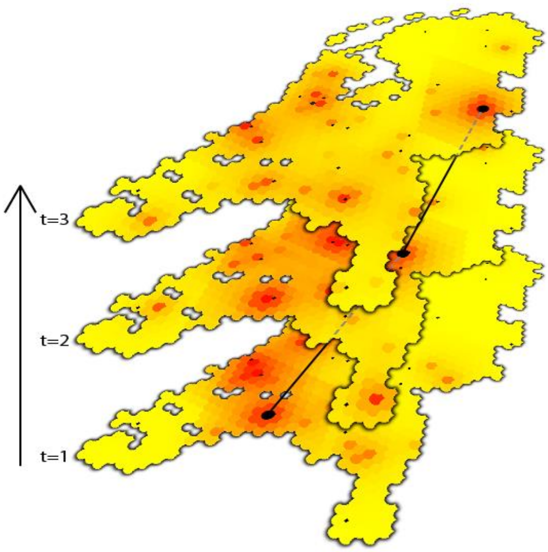

Figure 2.8: Graphical representation of a helicopter moving through time and space (Van Urk, 2012).

Before the tool determines the optimal route, it must first find the best start time, within the start and end of the shift in which the flight should take place. There are two ways to determine the optimal start time. First, the exact method: determine the optimal route for every possible two-minute time interval as start time and calculate the total route coverage of that route. Then we take the start time with the highest total forecast coverage as the optimal start time. It takes approximately 3 seconds to find the optimal route and calculate its forecast coverage. This results in a computation time of 36 minutes for the 720 time intervals per day.

Second, the tool enables the user to *estimate* the optimal start time "based on the assumption that it is more likely to have a successful assist when more accidents happen" (Van Urk, 2012). Van Urk (2012) proposes an estimator for the total coverage of a flight on a given start time. We call this estimator method the "flight value estimator". The flight value estimator takes the summation over all hexagons that the helicopter can reach in the time since it started and from which the helicopter can still reach the end location in the remaining time. After adjustments, this estimator now takes 1 second to compute the estimated flight value for approximately 400 time intervals and thus results in a computation time of 2 seconds per day or 12 minutes per year. The tool then determines the start time with the highest estimated flight value. Both techniques, the exact method and the flight value estimator, do not correct for the fact that the helicopter also covers incidents when it stays stand-by on the ground.

During the flight, the helicopter covers part of the forecasted incidents, as defined in the probability of success factor. Similar to the coverage of incidents we assume that when the probability of successful support at a certain location is 80%, 80% of the forecast value at that location is covered. For the next flight the flight value estimator should thus be updated to exclude the (percentage of the) incidents that is covered by the previous flight. Figure 2.9 shows the estimated value of flight for every time interval in a two day period. Flight one is scheduled at the highest peak and thus decreases the added value of future flights at the same point in time.

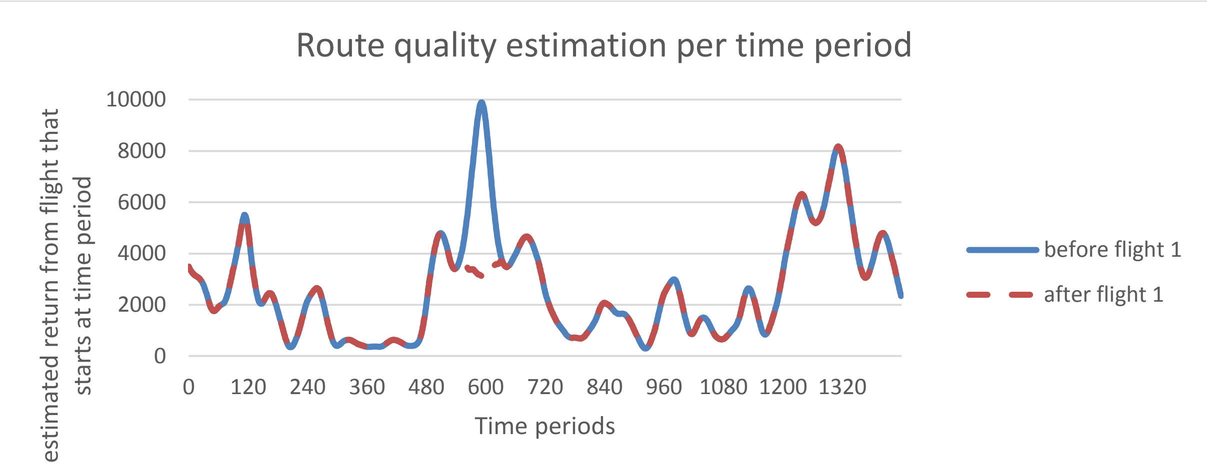

Figure 2.9: The route quality estimation is updated after every flight.

The number of time intervals that are affected depends on the duration of the flights. When the standard flight duration is 60 minutes, then a flight that starts 59 minutes before the start of the current flight will be influenced in its last minute. Every flight that starts after the start time of the current flight plus the current flight duration is not affected by the current flight.

2.3.3 Rescheduling

The tool also provides rescheduling support. When there is intelligence available about an incident, the user can input it into the tool. The tool then provides a new start time and route for the helicopter that makes sure the helicopter visits the intelligence location in time.

2.4 Performance management

The LVP currently measures performance by keeping track of the number of assists and arrests by the helicopters. When there is a helicopter available at a close enough distance to an incident where a criminal is caught in the act, then the Operations room authorizes deployment. This deployment is labelled as an "assist" of the helicopter. Furthermore, helicopter pilots never arrest criminals in person but are able to guide the police on the ground. Instead, in the case that a helicopter then enables the police on the ground to catch the criminal, the flight is registered as an "arrest". Figure 2.10 visualizes how a large number of emergency calls result in a smaller number of assists/arrest per year by the LVP. We conclude that we cannot assume that the helicopter has an added value to every HIC incident that happens within the area that is covered by the helicopter.

Figure 2.10: Visualization of incident filtering by the Central Operations Room.

Earlier experiments designed to determine the performance of the operational planning tool, ran for small time periods and assumed that the 100% coverage of an incident by a helicopter route equals an expected arrest in practice. When a helicopter covered the incident, it would fly to the incident, handle the incident, and fly back to its start location. However, as Figure 2.10 shows, not all incidents lead to an actual helicopter deployment. Furthermore, although a helicopter may arrive within 10 minutes at the crime scene, helicopters do not always provide crucial support at crime scenes and helicopter deployment thus does not always lead to an "arrest". Therefore, we conclude that conclusions on the proximity of helicopters cannot be directly translated to the number of expected arrests. Furthermore, we find that the performance of the operational helicopter scheduling tool has not been sufficiently validated on the performance measures of the LVP, by simulation or practice tests.

2.5 Desired situation

As noted by Van Urk (2012), the LVP wants to be able to "make a plan in such a way, that the expected number of arrests where a police helicopter makes a difference is maximized. A system should generate a plan for each helicopter showing when it has to be where. This system should operate with minimal required human intervention, however, it should allow for human input". We enlarge this goal by making the system ready to plan, schedule, and reschedule flights for up to a year ahead in time, taking into account crew scheduling and maintenance constraints.

Furthermore, we recognize the importance of validation and aim to provide the LVP with a simulation model that is able to test different settings of the decision support system and determine the added value of forecasting, routing, and tactical planning.

2.6 Conclusion

The planning office currently uses a uniform distribution of flight hours over the year, although the criminal intensity is not uniformly distributed over the year. Furthermore, the planning office has no tools to support its tactical decision making, while the current forecasting and routing tool needs refinement and validation. Since the output of a tactical decision support system is the optimal input for the operational planning tool, the planning office thus requires an integrated tactical and operational planning system.

3 Literature research

This chapter describes the relevant scientific literature on forecasting in Section 3.1. Section 3.2 discusses the calculation of forecast accuracy and Section 3.3 describes literature on routing. Section 3.4 presents related problems from literature and discusses tactical planning. Section 3.5 discusses the concept of fairness and how to measure it. Section 3.6 draws conclusions on the literature research for this research.

3.1 Crime forecasting

This research aims to create a medium-term forecast of crime. Gorr and Harries (2003) note that, while “conventional forecast methods are of little use in predicting the behaviour of individual serial criminals”, there are patterns in crime. Furthermore, they argue that crime forecasting became relevant when police began mapping crime using geographic information systems (GIS), and the *criminality of places* was established, based on routine activities and hot spots (Sherman et al., 1989). Hot spots are areas that have a relatively high level of criminal activity. There are two types of techniques to find hot spots. First, drawing hot spot circles, ellipses, or isopleths (lines through points on a map with the same crime density) (Block, 1995). Second, a fixed polygon grid as Van Urk (2012) uses.

Gorr and Harries (2003) state that forecast methods for fixed geographic units can easily draw on traditional time series methods of forecasting. Furthermore, they note that forecast methods for fixed boundary and ad hoc spatial cluster areas (i.e., hot spots) are complementary. After the use of fixed boundary forecasts to narrow down the search for problem areas, clustering methods are of use to diagnose and focus on specific targets. Gorr et al. (2003) provide evidence that geographic scale is the predominant factor influencing short-term crime forecast accuracy. They conclude that the average crime count in a forecast grid cell should be at least 25-35 per time period to achieve an acceptable forecast accuracy. Furthermore, Gorr et al. (2003) find that the Holt-Winters double exponential smoothing method with monthly seasonality is the most accurate time series forecast model for precinct-level crime series, in comparison with naïve forecasting methods and other types of exponential smoothing. Chen et al. (2008) compare the AutoRegressive Integrated Moving Average (ARIMA) model with Simple Exponential Smoothing (SES) and Hyperbolic Exponential Smoothing (HES) to make short-term forecasts for one city. They conclude that the use of ARIMA results in a higher forecast accuracy than the use of HES or SES. Felson and Poulsen (2003) find that crime varies per day and that the distribution of crime over the day varies per city and within cities.

Next to retrospective models, also prospective models are in development, although mainly for short-term forecasting. Caplan et al. (2011) propose Risk Terrain Modelling (RTM) to forecast shootings. The method uses risk terrain maps, based on a range of contextual information relevant to the opportunity structure of shootings. Liu and Brown (2003) model a transition probability of crime across a police jurisdiction. Cohen et al. (2004) specify a leading indicator model for a one-month-ahead forecast on the micro-level scale, that takes specific characteristics of areas or neighbouring areas (e.g., the number of shots fired in the month), to create future values for the expected crime level. Bowers et al. (2004) observe that a burglary in a household increases the probability of a burglary for houses within 400 meters of the household, especially on the same side of the street. This effect is apparent for up to two months. They use this observation to create prospective maps for burglaries. This section shows that there are several explanatory variables for criminal incident patterns in literature. However, we do not find an indication for which variables are most important.

To take multiple explanatory variables for crime like poverty and the unemployment rate into account while analysing aggregate crime rates, Osgood (2000) proposes Poisson regression. We can use Poisson regression when we assume that the number of crimes per time interval has a Poisson distribution. The Poisson distribution is a discrete probability distribution that gives the probability that a number of events occurs in a fixed time interval, given that these events occur with a known mean rate and the timing of the events is random and independent. Since crimes are discrete events and can be assumed to be independent, the Poisson distribution is suitable for investigating crime rates. Next, we should create a model that relates the expected crime rate to the explanatory variables, in which every variable gets its own weight. To determine the set of weights for the different explanatory variables that makes the model fit best to the historical data, we then use regression. The resulting weighted log-linear model can then be used to create the expected number of crimes in the next time interval, given that the future values for the explanatory variables are known.

3.2 Forecast accuracy

To determine the quality of a forecast we need to determine its accuracy. There are several forecast error measures available in literature. To determine whether we forecast the seasonal pattern of criminality well, the Mean Squared Error (MSE) is a common error measure. Equation 3.1 shows the formulation of the MSE, with the following notation:

- Y_t the actual number of incidents at time t
- F_t the forecasted number of incidents at time t

$$MSE = \frac{1}{n}\sum_{t=1}^{n}(Y_t - F_t)^2 \qquad \text{Equation 3.1}$$

Since we deal with intermittent occurrences of criminal incidents and since the forecast values are generally very small, the MSE does not provide clear results. Hyndman (2006) proposes a forecast error measure for intermittent occurrences: the Mean Absolute Scaled Error (MASE). The MASE compares the mean absolute error of the forecast, with the mean absolute error of a naïve forecast. The naïve forecast uses the actual value of the current period as the forecasted value of the next period. The division of the actual forecast performance by the naïve forecast performance scales the result. Equation 3.2 gives the formulation of the MASE:

$$MASE = \frac{\frac{\sum_{t=1}^{n}|Y_t - F_t|}{n}}{\frac{\sum_{i=2}^{n}|Y_i - Y_{i-1}|}{n-1}} \qquad \text{Equation 3.2}$$

To determine the forecast quality in two dimensions, time and space, we need a different forecast error measure. Willis (2002) discusses a measure called "spatial frequency analysis" to determine the spatial forecast error in geographical demand for electrical power. Spatial frequency analysis divides an area into regions and determines the frequency with which the forecast errors per area switch between negative and positive. Figure 3.1 shows two exemplary "error maps", with an indication on every grid area whether the forecast is over or underestimated. Both error maps have the same mean error, the same average absolute error and the same root mean square error. However, map A will probably lead to serious underestimation of required capacity in the eastern areas and overcapacity in the western areas, while map B will lead to no problems as power stations cover around 40 areas at the same time.

The discussed forecast error measures are designed to determine the accuracy of forecasts that determine the amount of, e.g., future demand. However, the forecasting procedures by Van Urk (2012) create a relative forecast in which the numerical value of the forecast at a given time and location is no indication of crime intensity. Only by comparing the forecast value with the forecast value at other times and locations we can determine where hotspots of crime intensity are. We did not find any specific forecast error measures for a relative forecast in literature.

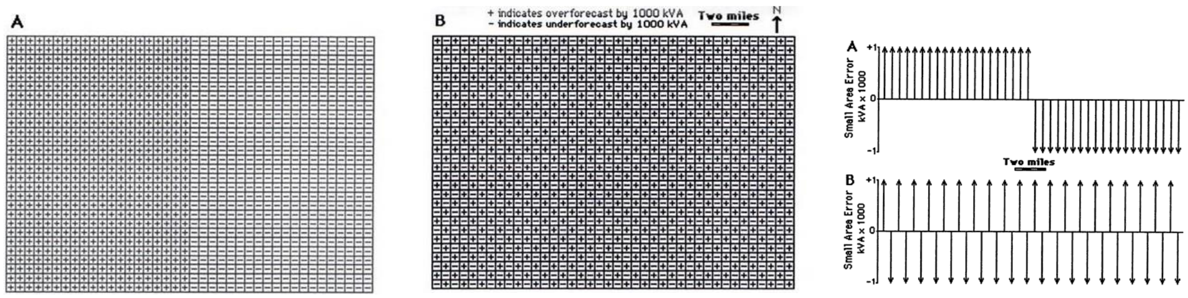

Figure 3.1 cross-sections through area error maps A and B reveal a difference in how fast the errors "oscillate" as a function of distance (Willis, 2002).

3.3 Emergency vehicle routing

Since the locations of hotspots of crime vary during the time of the day, it is not optimal to determine only fixed helicopter positions. By making surveillance flights with routes that follow the geographical pattern of hotspots, it is possible to decrease average arrival times to incidents. Van Urk (2012) found that the problem solved by the routing module in this research is a special case of both the Location Covering Problem (LCP) and the Dynamic Vehicle Routing Problem (DVRP).

The LCP is a combinatorial optimization problem in which the static location of resources is optimized to maximize the coverage of demand. There are several versions of the Location Covering Problem. Toregas et al. (1971) propose a model for the Location Set Covering Problem (LSCP) in which a set of locations is given where a resource may be positioned. The LSCP model then minimizes the number of locations required to cover all demand points. A demand location is covered by a resource when the distance between them is less than a predefined distance. Church & Revelle (1974) proposed the Maximal Covering Location Problem (MCLP). In the MCLP, a fixed number of resources and a set of possible locations is given. The objective function is then to maximize the number of demand points covered by the set of resources. Schilling et al. (1979) extend the MCLP to incorporate multiple resource types. Li et al. (2011) provide an overview of extensions for emergency response facility location models. The LSCP and MCLP are both static models. Owen & Daskin (1998) provide a review of probabilistic models that model emergency services as queues and determine availability of these services and waiting times. Kolesar & Walker (2012) developed a relocation model in order to take the effect of dispatches into account.

The DVRP is also a combinatorial optimization problem, where current and unknown future service requests are filled by multiple vehicles. Laporte (1992) notes that the Vehicle Routing Problem (VRP) is a special case of the DVRP, where all requests are known in advance. Eksioglu et al. (2009) provide an overview and taxonomy of the vehicle routing field.

Since we use the routing module as a fixed input in this research, we do not describe the literature on the LCP and DVRP in further detail. Van Urk (2012) concludes that the problem of positioning helicopters to maximize the number of arrests can be described as a DVRP with only delivery (of air support) with hard and soft constraints. However, Van Urk (2012) concludes that the problem is better described as a LCP due to the short time between incoming requests and the end of the hard time window. The focus of this research is to use the current routing functionality as proposed by Van Urk (2012) to create efficient long-term plans.

3.4 Tactical planning and related problems

Spatial and temporal forecasting, routing, and tactical planning methods are applicable to multiple problems besides police resource planning and scheduling. Davidson et al. (2009) mention that ocean current forecasting and routing is used to estimate the minimum cost ship route for freight ships and to determine the best route for icebreakers. Furthermore, Moreira-Matias et al. (2012) discuss the short-term prediction of taxi-passenger demand to real-time guide taxi drivers to areas where there is less congestion and more demand. Other examples include the temporal and geographic clustering of residential fires for targeted fire prevention by Wuschke et al. (2013), and the forecasting of medical demand to allocate medical facilities by Knight et al. (2012).

Next to police helicopter patrol routes, the forecasting and routing combination is also applicable to patrol routes for police cars and officers on foot. Chawathe (2007) proposes a patrol planning method to optimize coverage of important locations at minimum cost. However, he does not discuss where the importance of locations is based on. Reis et al. (2006) propose a genetic algorithm that determines optimal patrol routes given a set of simulated criminals that try to rob a number of targets. We did not find any patrol route optimization based on a forecast in literature.

Ambulance planning is related closely to police helicopter planning, as arrival time is the crucial factor for success and resources are expensive. Van Essen (2013) describes ambulance planning models on the strategic and tactical level. On the strategic level she discusses the problem of locating ambulance bases and on the tactical level she identifies the problem of the number of ambulances to allocate to a base to cover all demand. These problems are similar to the helicopter allocation problem of Buiteveld (2011), which we consider on the strategic level since the allocation to bases is permanent.

In Section 3.3 we discuss three types of ambulance planning models: deterministic, probabilistic, and dynamic models. Deterministic models determine the optimal static location of ambulances to maximize coverage. Probabilistic models model ambulances as servers and use queuing theory to model availability and waiting times. Dynamic models relocate ambulances after one or more ambulances are dispatched to emergencies. An example of a dynamic model is proposed by Reuter & Michalk (2012). We refer to Brotcorne et al. (2003) for a detailed discussion and classification of deterministic, probabilistic, and dynamic models for ambulance location and relocation. We identify several levels of accuracy with which the availability of ambulances is and can be modelled. For example, Van Essen (2013) uses a deterministic model with Euclidian distances between ambulances and patients. The probabilistic models is more accurate and takes into account that ambulances can be occupied, while the dynamic models allow the relocation of ambulances. To the best of our knowledge, there is no literature on the routing of emergency vehicles based on forecasts. We conclude that tactical planning problems can be solved with different degrees of detail.

3.5 Fairness

Fair allocation of public services, or “equity”, is a critical and controversial factor when deciding how to allocate public resources (Stone, 2002). During location optimization, it is often appropriate to combine equity and efficiency measures, as equity measures alone can produce undesirable results (Smith, Harper, & Potts, 2012). Figure 3.2 shows an example with customers C1 to C6 and Location L1 and L2. Location L1 at the approximate centre of the circular pattern of customers C1 to C6 equalizes the distance travelled by the customers and is therefore more equitable then L2, although L2 has a shorter distance to all customers.

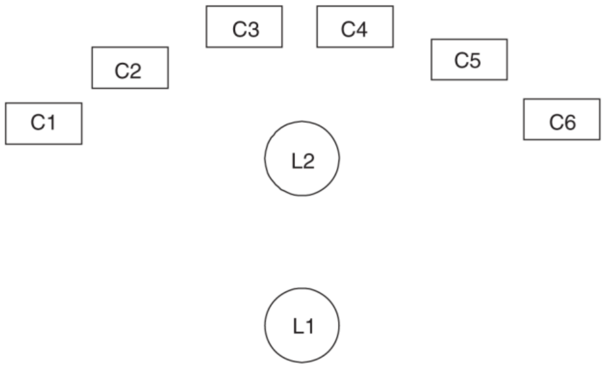

Customers C1–C6 are illustrated in a circular pattern. Location L1 equalises distance travelled, while L2 gives a shorter distance to all.

Figure 3.2: Example showing the possible inefficiency of equitable locations (Smith et al., 2012).

Since there are numerous different equity measures, Marsh and Schilling (1994) review the literature on equity measurement in Location Theory to create a common notation and framework to organize the measures. They use the following definitions:

- Equity: when each group receives its fair share of the effect of the facility siting decision.
- Effect: a change in the status quo of a value of interest caused by influences from the location and size of facilities being sited.
- Group: an aggregation of individuals on which effects are measured. A group can consist of one individual.

Marsh and Schilling (1994) then discuss 20 measures and organize them in a framework. It is important to point out that the measures actually measure inequity, which means that a bigger value indicates a less equal distribution. Equity would be maximized by minimizing the resulting value of the measure. Table 3.1 shows a summary of the framework with an example measure for every cell. They use the following notation:

- i, h: Index of group being compared.
- E_i: Effect on group i.
- $\bar{E}$: Arithmetic mean of the effect over all groups.
- A_i: An attribute of group i.

The measures are organized by the type of scaling and the reference distribution that the measures use, of which we now discuss the possible types. Scaling is used when the size or another attribute of group differs between groups. General scaling is the division of the effect on a group by the value of some attribute of that group. Normalization is a method of scaling by dividing the effects by the mean effect. Additionally, it is possible to divide the measure by the number of groups (N). Since the number of groups does not change during this research, this has no added benefits. The reference distribution can be described as the target solution. There are three types of reference distributions. When all groups should be treated equal, the quality of a solution is the difference between the effect on each group and the mean effect on all groups. When a group should receive an effect according to an attribute, the quality of the solution is the difference between the actual effect and the attributed effect on a group. Finally, it is also possible to compare all groups with a chosen group. Since choosing a group as the reference standard is not realistic, it is probably better to compare all groups to all other groups. This is called the peer reference distribution.

Scaling	Reference distribution: Peer	Mean	Attribute
None	$\lvert E_i - E_h \rvert^p$	$\lvert E_i - \bar{E} \rvert^p$	$\lvert E_i - A_i \rvert^p$
Normalized	$\frac{1}{N^2}\frac{\lvert E_i - E_h \rvert^p}{\bar{E}}$	$\frac{1}{N}\frac{\lvert E_i - \bar{E} \rvert^p}{\bar{E}}$	$\frac{1}{N}\left\lvert\frac{E_i}{\bar{E}} - \frac{A_i}{\bar{A}}\right\rvert^p$
General	$\left\lvert\frac{E_i}{A_i} - \frac{E_h}{A_h}\right\rvert^p$	$\left\lvert\frac{E_i}{A_i} - \frac{\bar{E}}{\bar{A}}\right\rvert^p$	$\left\lvert\frac{E_i - A_i}{A_i}\right\rvert^p$

Table 3.1: Framework for equity measures (Marsh & Schilling, 1994).

Implementation of the equity measures in location optimization problems is commonly done with multi-criteria mathematical programs. For example, for MCLP problems it is possible to incorporate the equity function into the objective function or in a constraint (Church and Revelle (1974), Chanta et al. (2011), Smith et al. (2012)).

3.6 Conclusion

This chapter discusses relevant literature on hotspot forecasting, forecast error measures, emergency vehicle routing problems, tactical planning, and fairness. In Section 3.1 we find that hot spot forecasting of crime in literature is mostly retrospective and prospective models are only used for short-term forecasts. Section 3.2 shows that literature contains forecast measures for temporal and spatial forecast, but the measures do not adequately measure the accuracy of the timing and location of hotspots in hotspot forecast. Section 3.3 describes literature regarding emergency service location, relocation, and routing, and we find that the routing model by Van Urk (2012) satisfies the requirements of this research. In Section 3.4 we discuss forecasting, routing, and tactical planning for other resources like ambulances, fire trucks, taxi's, and icebreakers, and conclude that tactical planning problems can be solved with different degrees of detail. Finally, Section 3.5 discusses the concept of equity and equity measures from literature.

4 Forecasting and routing

This chapter describes the improvements and extensions of the operational planning tool. Section 4.1 introduces a method to make the input data suitable for the tool, Section 4.2 discusses the adjustments to the forecasting technique to include weekly and seasonal patterns, and Section 4.3 explains the alterations to the flight value estimator to improve its accuracy. Section 4.4 identifies improvement options for the routing method and Section 4.5 provides the conclusion of this chapter.

4.1 Forecast data input

After the research of Van Urk (2012), the LVP increased the scope of the incident data to include not only robberies of commercial locations, but also burglaries and street robberies. Currently, a dataset is available of approximately 300,000 incidents from 1-7-2010 to 30-6-2013. Of these incidents we know the location (zip code), date and time, and type. We round up the incident times to the next two minute interval (e.g., 08:03:15 becomes 08:04, etc.), as the forecast tool works with two minutes intervals. Figure 4.1 shows the average distribution of incidents over the day, by adding the number of incidents in every time interval. We expect a gradual pattern of criminal intensity over the day but find a pattern with sudden peaks. We conclude that the distribution of incidents over the day does not seem natural. We expect that police officers, who report incidents to the police data system, round the time of incidents to hours, half hours, and quarter hours, as the peaks are at those times.

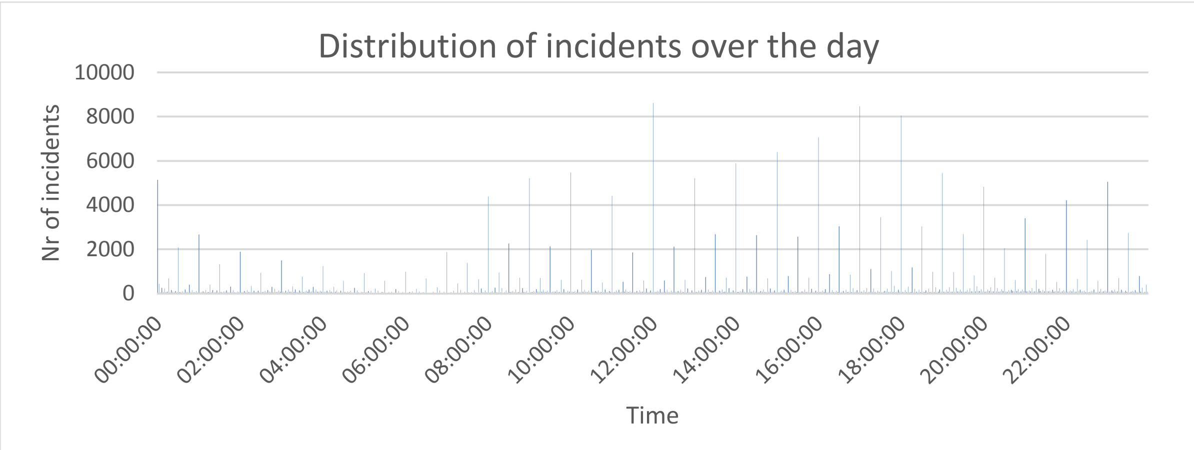

Figure 4.1: Distribution of incidents over time (initial data).

The incident pattern has an impact on the outcomes of this research. Based on this pattern, optimal flights would probably start just before full hours. The optimal flight duration would then be the shortest possible time it takes to get to a hotspot at the full hour, and fly back to a base. Since this would lead to inefficient routes in real-life, we want to restore a realistic pattern by correcting the timing of incidents for the rounding errors made earlier. Furthermore, we want to keep the geographical properties of the original data. We identify three methods to correct the timing of the incident data.

First, we can smooth the number of incidents per time interval with techniques such as exponential smoothing or the moving average. After we determine the number of incidents per time interval, we then have to determine the incidents to change the timing of, and how to adjust the timing. For example, when at 08:32 we only had 5 incidents but should have 30 according to the smoothing result, we should find 25 incidents from times with too many incidents to fill the shortage.

Second, we can aggregate the incidents into time buckets. For example, we can determine the number of incidents per hour and then distribute these incidents over the 30 time intervals in that hour. The redistribution of the incidents over the time intervals within the bucket can be done by, e.g., a uniform or normal distribution. As we do not want to create a bias, we recommend using a random number generator that generates time adjustments according to the chosen distribution.

Third, we can assume that incidents with specific times such as 08:39 are recorded correctly, and spread out only the peaks in the number of incidents. For every incident in a peak we then have to determine how much to adjust it. Again, we can use a random number generator with a distribution (e.g., normal or uniform) to spread a peak over the surrounding times.

Every described method has disadvantages. The resulting distribution of all methods depends mostly on the parameter setting, and we cannot choose between different parameter settings as we do not have a measure for the quality of the resulting distribution. Furthermore, the method that shuffles incidents around to fill up shortages after smoothing the distribution may be biased. Finally, the aggregation per time bucket creates unrealistic jumps in criminal activity at the cut-off point between two buckets.

Since the method to spread peaks is easy to understand and the underlying logic is clear to the LVP, we apply it. As we do not have any data to underpin the settings for a normal distribution, we use a uniform distribution. We let the bounds to the uniform distribution depend on the type of peak. The incidents that are booked on 8:00 could probably have happened between 7:30 and 8:30. Therefore, we use a random number generator with a uniform distribution to distribute the incidents over this time period. We use the same procedure for incidents on half hours (distribute uniformly over interval from a quarter hour before the incidents to a quarter hour after the incidents), and for incidents on quarter hours (spread over interval from 5 minutes before to 5 minutes after the incident time). Figure 4.2 shows a new distribution of incidents over the day after redistributing the incidents.

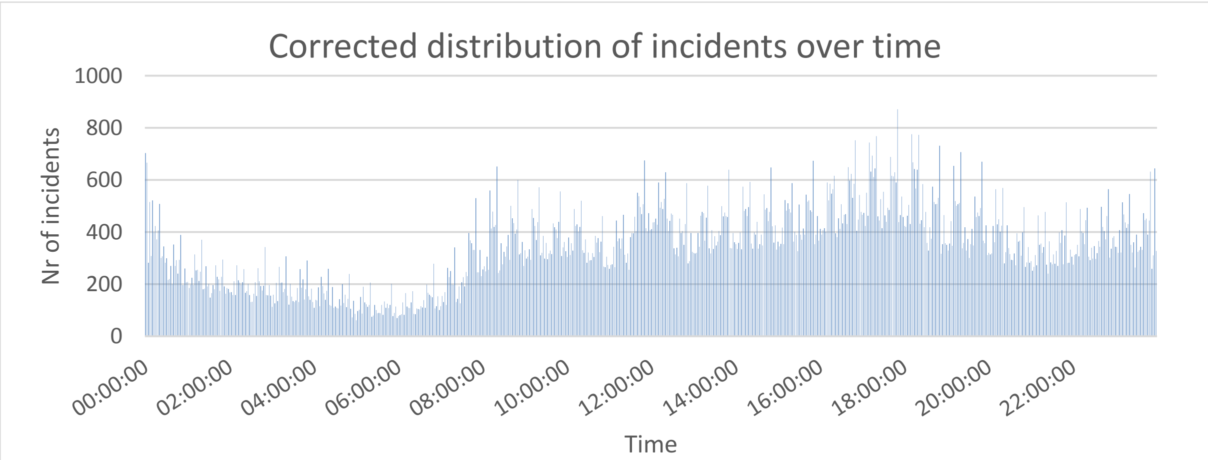

Figure 4.2: Distribution of incidents over time after redistribution of incidents.

From Figure 4.2 we conclude that the peaks have been decreased in size, and that the distribution of incidents over time is improved (more gradual) in relation to Figure 4.1. However, we do find a disadvantage of the uniform distribution with the current settings. Time intervals that are close enough to multiple peaks such as 08:18 (within 30 minutes of 08:00 peak, within 15 minutes from 08:30 peak, and within 5 minutes of 08:16 peak) become new peaks. Furthermore, time intervals at times such as 08:00 are underestimated because they only get their part of the spread, but do not have additional incidents that a time such as 08:02 already had. Therefore, we recommend for future research to determine the difference between the uniform distribution and, e.g., a normal distribution.

4.2 Forecasting

The goal of this section is to extend the current forecasting method to create the required data for tactical planning. As the tactical planning should distribute the flights over a one year period, we need to extend the scope of the forecasting method to one year. Section 4.2.1 extends the current forecasting method to a planning horizon of multiple months. Section 4.2.2 introduces a less complex alternative method and Section 4.2.3 proposes an adjustment to the geographical aggregation method to increase the forecast quality. Section 4.2.4 introduces methods to enable the LVP to adjust the forecast to include incident priorities and the difference in effectiveness of helicopters during the day and night.

4.2.1 Extend scope and increase speed of forecasting method

This section describes how we propose to extend the scope of the forecasting method from one day to multiple months and how we introduce weekly and seasonal patterns into the forecast. Furthermore, we discuss a technique to increase the forecast speed. Finally, we discuss the effect of the forget factor and propose an improvement for it.

Seasonal pattern

As discussed in Chapter 2, the forecasting method currently forecasts one day ahead. The total value of the one-day forecast is unit-less, and depends on the number of incidents in the input data set and the data aggregation settings. Since it is unit-less, only the height of the forecast at a given time and location in comparison with the forecast value at other times and hexagons is indicative.

We now want to extend the scope of this relative forecast to multiple months. The forecast output should not only forecast peaks in criminal intensity in place and time within one day, but also distinguish between different weekdays and months. For example, the forecast should show the relative difference in expected criminal intensity between a Tuesday morning in January and a Thursday night in September.

To extend the scope of the forecast we can create a new forecasting method that forecasts for a multiple day scope in one calculation, or create a multiple day forecast by combining several one-day forecasts. Since the current forecast method creates a daily pattern based on the month and weekday crime distribution characteristics, the multiple use of the method is straightforward. Furthermore, since forecasting is not the main goal of this research and we are limited to the available AIMMS software, we choose to apply the current method multiple times, instead of creating a new forecast method.

When we use the current forecasting method multiple times, the resulting forecast does not show a seasonal pattern. Since the total forecast value per day depends only on the number of incidents in the data set, we only see different daily patterns. Therefore, we note that the terms "weekday factor" and "month factor" are a bit deceiving, as the factors do not correct the forecast for the relative criminal intensity of the weekday and month that we forecast for. We propose to rename the current weekday and month factors the "weekday distribution factor" and "month distribution factor". The factors provide a conversion factor to transform data from a different weekday or month to the distribution of the target weekday/month.

The main influencing factors of criminal intensity are the day of the week and the number of non-daylight hours. We related the influence of non-daylight hours to months in Chapter 1. Figure 1.4 shows the monthly pattern and Figure 4.3 shows the average weekly pattern in 2012. From data analysis we found that this average weekly pattern holds for all months.

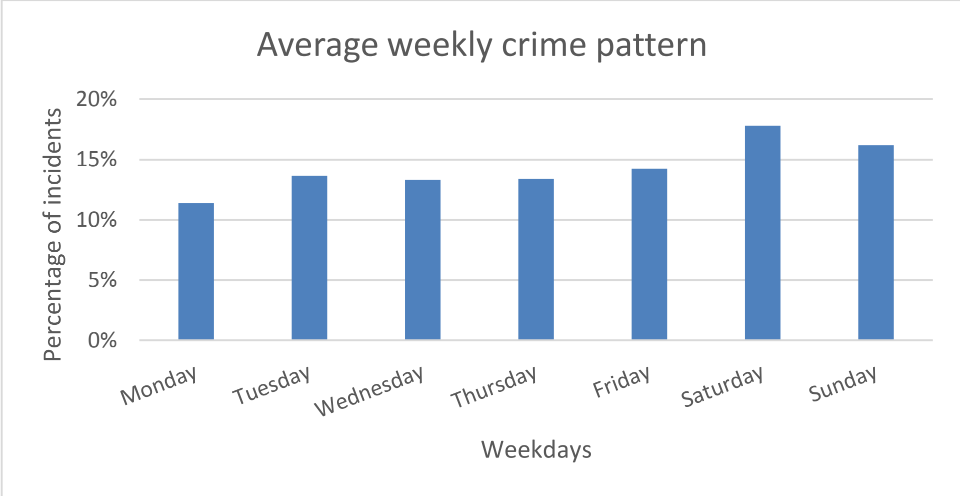

Figure 4.3: the average weekly crime pattern in 2012.

We propose to use a technique similar to the weekday and month distribution factors to convert data from days with a different total criminal intensity to the target day. We use the terms weekday and month factor as the relative criminal intensity of that weekday or month, related to the busiest weekday/month combination. Equation 4.1 shows the new formulation for the month factor with the same notation as Equation 2.2.

$$month\ conversion\ factor_{m_t,m_i} = \frac{\sum_h\left(N_{m_t,h}\right)}{\sum_h\left(N_{m_i,h}\right)}, \qquad \text{Equation 4.1}$$

Where $\sum_h\left(N_{m_t,h}\right)$ is the total number of incidents of the target month, and $\sum_h\left(N_{m_i,h}\right)$ is the total number of incidents in the month of the data we transform. We perform the same type of calculation for the weekday factor where we relate all weekdays to each other. Before we determine these factors, we note that we correct for the number of days per month and weekday, since we would otherwise constantly underestimate February and months with 30 days.

To convert data from months with higher criminal intensity to a forecast for a target month with a lower criminal intensity, we should alter the current weekday and month distribution factor by combining it with the month and weekday factor. Equation 4.2 shows the combined formulation:

$$month\ distribution\ factor_{m_t,m_i,h_i} = \frac{\left(\frac{N_{m_t,h_i}}{\sum_h\left(N_{m_t,h}\right)}\right)}{\left(\frac{N_{m_i,h_i}}{\sum_h\left(N_{m_i,h}\right)}\right)} * \frac{\sum_h\left(N_{m_t,h}\right)}{\sum_h\left(N_{m_i,h}\right)} \qquad \text{Equation 4.2}$$

We then combine the month and week distribution factor to create the day distribution factor for a certain target day as in Van Urk (2012).

$$day\ distribution\ factor_{m_t,w_t,h_i} = \\ month\ distribution\ factor_{m_t,h_i} * weekday\ distribution\ factor_{w_t,h_i} \qquad \text{Equation 4.3}$$

Figure 4.4 shows the result of the seasonal and weekly pattern extension. We scaled the total forecast value of the year by dividing it by the actual total number of incidents in that year, to compare both distributions of crime over time. The seasonal pattern does not take into account influencing factors like holidays, which shows clearly on the 31st of December. However, overall we see that the pattern is a combination of the week and month patterns and follows the actual seasonal pattern closely.

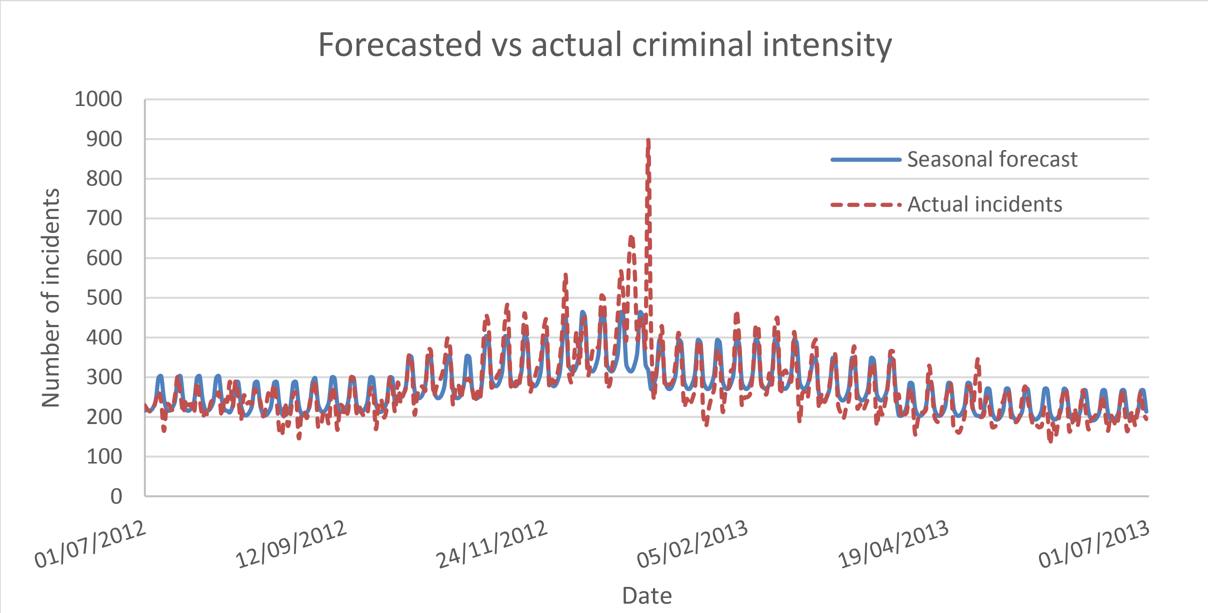

Figure 4.4: Seasonal pattern in both the forecast and the actual incident data of July 1st 2012 to July 1st 2013.

Fast forecasting for extended forecast scope

To extend the forecast to a year, we can use the original procedure, as described in Van Urk (2012) with the adjusted weekday and month factors, multiple times. For every day in the scope, we use all incidents in the dataset. We then use spatial aggregation and aggregation over time, which takes approximately half a minute and three minutes per day respectively. For a forecast of a year, this results in an estimated forecasting time of 20 hours and thus a smarter method is required.

Since we use all available data for every day of the forecast, the only difference between forecasts for different days is the applied month and weekday distribution factors. Therefore, we can convert a whole day of forecast into another day by these distribution factors. For an explanation of the procedure, we refer to Section 2.3.1.

We find that the average absolute difference between a forecast made with the faster method and with the original method is approximately 3% of the mean forecast value. Appendix D describes the causes for the inaccuracy. We conclude that the forecast error is minor and can be decreased by reducing the size of the blocks to half-hour blocks or even two-minute blocks. The fast forecasting method decreases the computation time by 97% to less than half an hour for a one-year forecast. Finally, we find that this technique is useful for forecasts with a scope of more than two days, since it requires the slower forecasting method to create the forecast for the first two days.

Forget factor

Van Urk (2012) uses a monthly forget factor to increase the speed of adaptation of the forecast to changes in the incident data. Every month, the forecasting method forgets a fraction α of the old data, which means that old data is multiplied by the factor $(1-\alpha)$. This forget factor calculation has two undesirable effects.

First, when we make a two-month forecast in January 2014, the historic data is one month older in comparison with February 2014 than with January 2014. For the February 2014 forecast we would thus forget the fraction α of the data and the expected criminal intensity would thus be a fraction α

lower. This leads to a significant underestimation of criminal intensity in the last months of a one-year forecast.

Second, a forget factor that is based only on the number of months between the date of the incident and the date of forecast does not distinguish between different types of months. For example, when we forecast for a day in January, one-year old data from a day in January is a better indicator than data from a day in June that is six months old, but is discounted by a higher factor.

Table 4.1 shows both effects for a forget factor of 0.01. The discount factors are higher for target months that are further in the future. Data from the same month as the target month is discounted by a higher factor than data that is less old, but from a completely different month. This first effect does not occur when we apply the fast forecasting method as described before, since we apply the forget factor once to create the first day of the forecast, that we then convert to create the rest of the forecast.

Discount factor	Target month January 2014	Target month February 2014
June 2013	$(1-0.01)^7 = 0.932$	$(1-0.01)^8 = 0.922$
January 2013	$(1-0.01)^{12} = 0.886$	$(1-0.01)^{13} = 0.877$

Table 4.1: example of current forget factor calculation

Therefore, we propose to replace the monthly forget factor by a yearly forget factor or to use a formulation to determine the relative quality of the data to the target forecast day (e.g., based on the month and weekday of the data).

4.2.2 New forecasting method without factors

As discussed in Section 4.2.1, Van Urk (2012) proposes to determine the daily crime distribution of a Monday in January by multiplying for every hour the average crime distribution factor for that hour in January by the corresponding factor for an average Monday. However, when we multiply the January factors with the Monday factors, we do not necessarily get the distribution of a Monday in January. Figure 4.5 shows that the result of the multiplication of the January and Monday factors over- and underestimate the actual distribution of incidents on a Monday in January.

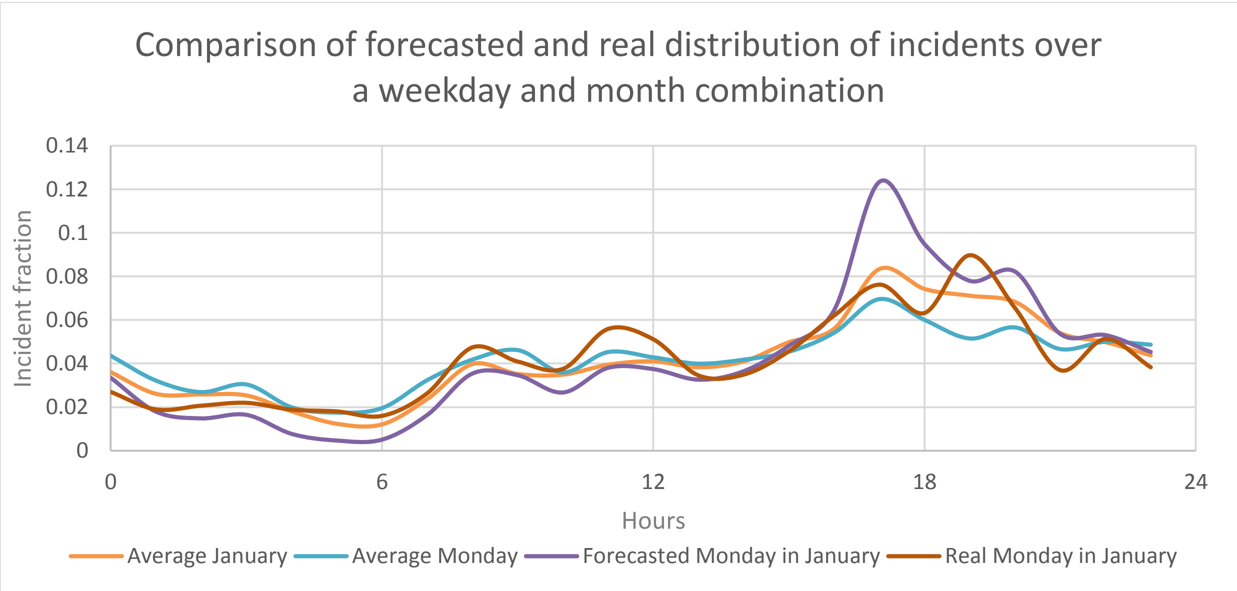

Figure 4.5: Distributions of an average day in January, an average Monday, the expected Monday in January, and the actual Monday in January.

We find that this over- and underestimation is due to the multiplication of the hourly weekday and month distribution factors. Figure 4.6 shows an example for the case where the average daily distribution is the same for days in January as for Mondays in general. The average distribution of crimes over a Monday in January should then be exactly the same as both distributions. However, we find that the multiplication leads to increased fluctuations, and thus overestimates peaks and underestimates relatively calm periods.

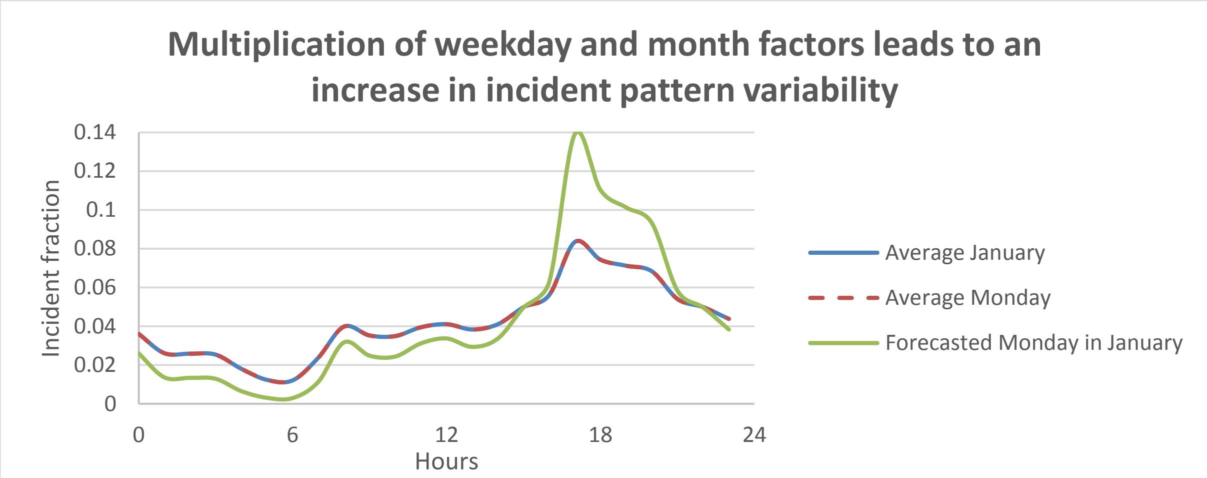

Figure 4.6: Expected distribution of a Monday in January, in case of an equal distribution of incidents over the average January and Monday.

There are several alternative methods possible to forecast the daily distribution of crime. Instead of multiplying the hourly weekday and month factors by each other, we can also take the average of both factors per hour. Alternatively, we can stop using weekday and month distribution factors. Instead, we can determine the distribution of criminal intensity over, or the entire forecast of, a Monday in January by using only data of historical Mondays in January. We call this the "weekday per month" forecasting method. This results in a forecast that is based on less data points, but has the added benefit of a geographical distribution of incidents that changes between weekdays and months. Since there is enough data available, and implementation at the LVP requires a simple procedure that can be automated in the current software tool, we choose to test the "weekday per month" factor method.

4.2.3 Spatial aggregation

Section 2.3.1 discusses the use of spatial aggregation to create a forecast for every location in the Netherlands, at any time. Spatial aggregation uses the assumption that when a historic incident happened in a certain hexagonal, then this data point is useful to predict the future crime level and pattern in that location, but also in surrounding locations. Figure 2.5 shows the factors that Van Urk (2012) proposed to spread incidents over rings of surrounding hexagons. Spatial aggregation was required because of the limited size of the dataset in earlier research. Figure 4.7 shows that the effect of the spatial aggregation is useful to create a forecast that covers most places in the Netherlands. However, it also shows that the approximation method fills the western area of the Netherlands with a high intensity forecast. When we compare this with the actual incident pattern, we see that the forecast looks less like the actual pattern than the initial geographical distribution.

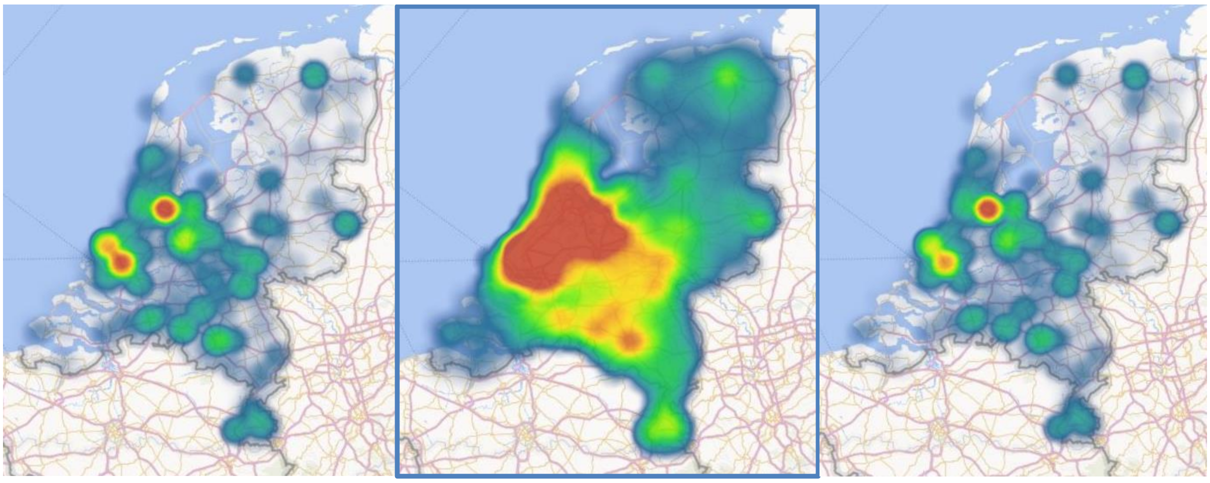

Figure 4.7: the forecasted geographical distribution for 2013 before spatial aggregation (left) vs. the forecasted geographical distribution after spatial aggregation (middle), and the actual distribution of forecasted incidents (right).

We conclude that the average distribution of crime over the Netherlands changes little between years and that spatial aggregation decreases the accuracy of the forecast. Because of the decrease in forecast accuracy, we propose to compare the accuracy and performance of a forecast with five rings around each incident's hexagon, with forecasts with one and no rings. The performance of a forecast is the performance of the set of routes that results from making a tactical planning with the forecast. We give exact definitions of forecast accuracy and performance in Section 6. We expect that less spatial aggregation increases the forecast accuracy and the forecast performance.

4.2.4 Forecast corrections for effectiveness and priority

This section discusses how we can take the effectiveness of helicopters into account in Section 4.2.4.1. Furthermore, Section 4.2.4.2 explains how to model different types of incidents.

4.2.4.1 Location population density and daylight

The LVP performed an analysis of the effectiveness of historical helicopter deployments. Table 4.2 shows the result of the LVP effectiveness and daylight analysis. The analysis shows that the effectiveness of helicopter deployment after sunset is higher than for deployment during daylight hours. Furthermore, the table shows (non-linear) differences between different levels of urbanization. Criminals tend to be harder to find in high density city areas and in low density rural areas, than in average density areas such as suburbs. In the case of high density areas this is probably because there are many buildings a criminal can enter, and because there are more civilians on the street. In low density rural areas, criminals can escape at higher speeds than in low density urban areas, and rural areas are commonly further away from the helicopter locations. Thus, incidents in the centre of Amsterdam are less interesting for helicopters than incidents in the suburbs. However, since the effectiveness of the helicopter per population density is based on a small dataset, the LVP concludes that there is only a significant difference between day time and night time effectiveness.

Population density	Effectiveness during daylight	Effectiveness after sunset
Non-urban municipality	5,2%	10,8%
Little urbanized municipality	5,8%	10,0%
Moderately urbanized municipality	1,4%	6,9%
Highly urbanized municipality	3,7%	11,2%
Very Highly urbanized municipality	3,9%	5,9%
Average effectiveness	4,1%	8,3%

Table 4.2: effectiveness of helicopter deployment per population density and light conditions.

We now want to adjust the forecasting and routing method to incorporate the effectiveness of helicopter flights. The timing of helicopter flights depends on the results of the flight value estimator, which depend on its input: the underlying forecast. Therefore, we have to adjust the input of the estimator to incorporate daytime and night time effectiveness. There are two ways to correct the input for extra information. We can correct for effectiveness during the calculation of the forecast, or use a filter to correct the forecast afterwards. For clarity, we choose to disconnect the forecasting method from the effectiveness corrections. We first create a forecast and then correct it with Equation 4.4, where e_t is the expected effectiveness at time $t \in T$.

$$corrected\ f_{l,t} = f_{l,t} * e_t \qquad \text{Equation 4.4}$$

Figure 4.8 shows the effect of taking into account the effectiveness on the temporal pattern of the forecast. We note that the transition from daytime effectiveness to night time effectiveness (around 22:00) and vice-versa creates a rather abrupt shift in the corrected forecast values.

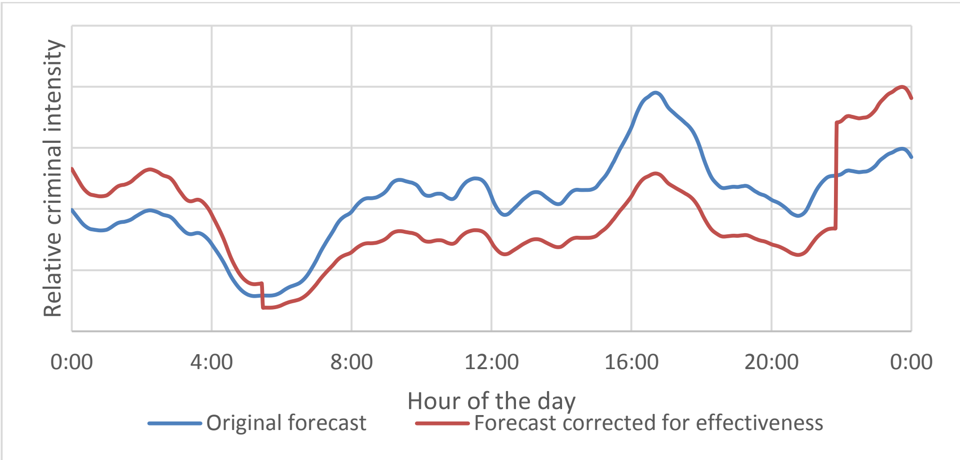

Figure 4.8: effect of the daylight effectiveness factor on the temporal forecast of 1-6-2012 on Schiphol.

4.2.4.2 Priorities

Next to the location and date/time of the incidents in the data set, we know of which type (robbery of commercial location, burglary, or street robbery) the incidents are. Since the police may find some types of incidents more important than others, the decision support system should enable the LVP to determine a priority value per incident. Furthermore, the LVP determined that the forecast should not be influenced more by types of crimes that happen more often. The LVP therefore uses the priority value to compensate for the difference in the number of incidents and proposes a priority value based on the relative difference between the number of incidents of different types of crimes. Table 4.3 shows the number of incidents and priority value per type of crime.

Type of crime	Number of incidents (2009-2012)	Priority value of Incidents
Burglary	335,595	1
Mugging	33,366	10.06
Robbery	10,592	31.68

Table 4.3: Number of incidents and priority value per type of crime (Korteweg, 2014).

To take priorities into account we have the same possibilities as for taking effectiveness into account. Priorities can influence timing and routes, and the forecast should thus be corrected for priorities. Since priorities depend on specific incidents instead of times and locations, we have to correct the forecast already in the forecasting process, before spatial and temporal aggregation. In Chapter 6 we investigate the effect of effectiveness and priorities on the results of the forecasting and routing method.

4.3 Optimal start time estimation

With the forecast available, we now want to determine the optimal start time. As discussed in Chapter 2, the tool by Van Urk (2012) uses an estimation method, because it takes too long to calculate a route for every time interval and choose the best one. The correct timing of flights is the most important step in creating a tactical planning and thus sets an upper bound to performance. Therefore, Section 4.3.1 evaluates the quality of the flight value estimator by Van Urk (2012). Section 4.3.2 then proposes a new flight value estimator.

4.3.1 Quality of the current flight value estimator

We determine the quality of the flight value estimator by comparing for 720 time intervals (one day) the difference between the expected and the actual coverage of the optimal route of that time interval. We determine the optimal flight value, or route quality, by calculating the optimal route at a start time and take the total forecast coverage by a helicopter on that route. Figure 4.9 shows the expected flight value according to the flight value estimator against the actual total coverage by the optimal route, for a 1.5 hour flight that starts and ends on Schiphol.

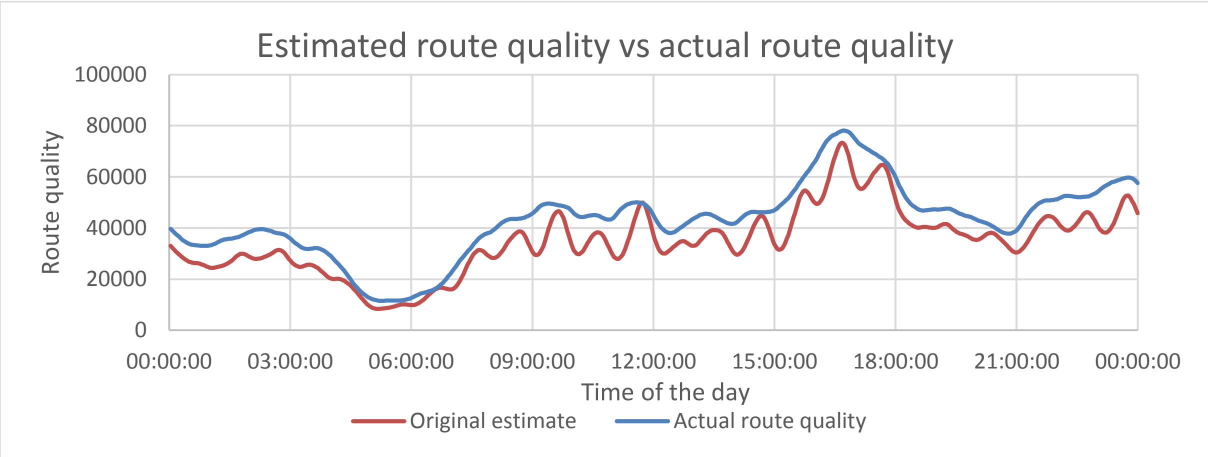

Figure 4.9: The estimated route quality does not behave like the actual route quality but correctly predicts the peak.

From Figure 4.9 we conclude that the maximum of the function of the flight value estimator is at approximately the same point in time as the actual maximum route quality. The flight value estimator thus seems able to find the optimal start time. However, when we compare the expected flight value of the best 1,460 peak times of an annual forecast with the actual forecast coverage of flights at these times, see Figure 4.10, we find that a higher expected forecast coverage not always results in a higher actual forecast coverage.

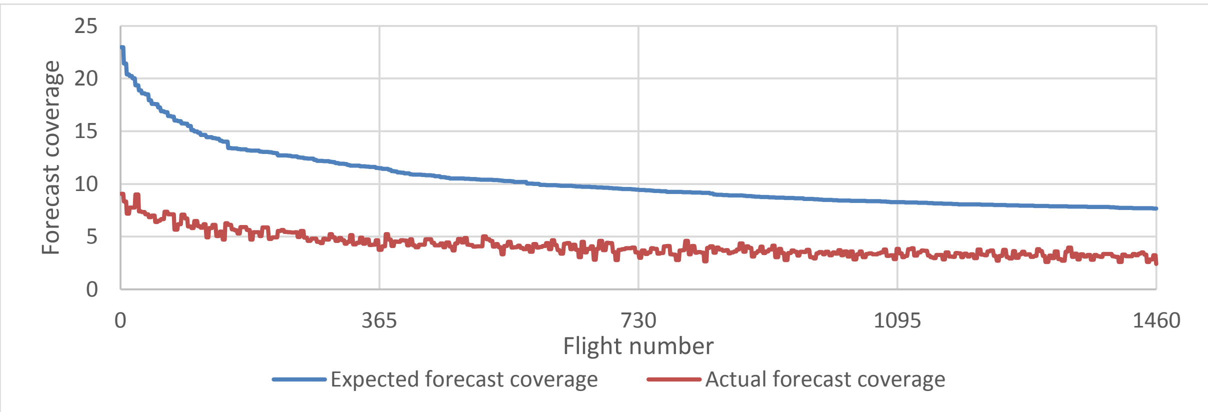

Figure 4.10: comparison of the expected forecast coverage and the actual forecast coverage for 1460 peak times.

Furthermore, since the estimator is independent of the coverage function of the helicopter and constraints to the possible route of a helicopter, it produces inaccurate results when we, e.g., want to compare route qualities from different start locations. Figure 4.11 shows that when the coverage function of the helicopter is set smaller, the actual coverage decreases but the estimated coverage stays the same. Thus, the estimated route quality is overestimated. Furthermore, analysis showed that this level of overestimation differs per location. This implies that we cannot use the original flight value estimator *directly* to find the optimal *start location.*

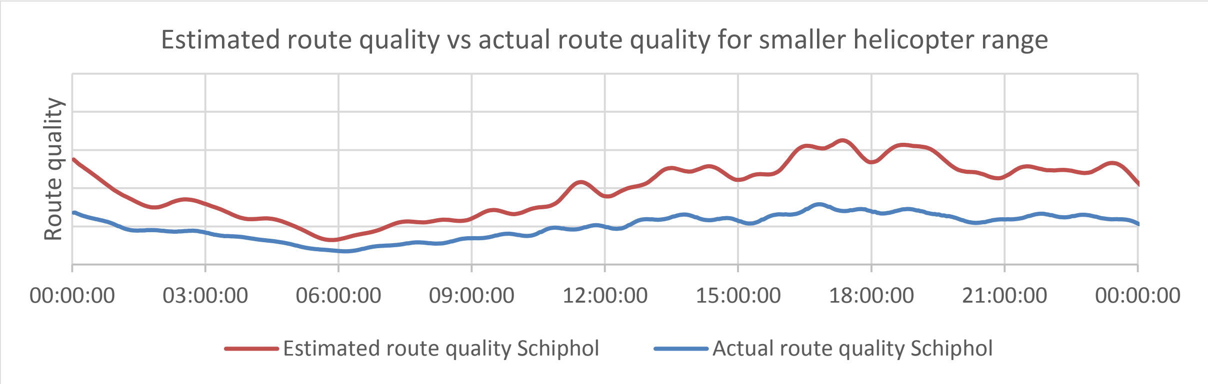

Figure 4.11: Comparison of the estimated and actual route quality of Schiphol for a helicopter with a smaller coverage area.

The heuristic currently determines the best start time by looking at the intensity of criminal incidents during the expected flight time. We cannot adjust it for the total coverage distance or success function, without determining where the coverage area is at a given time. To take all route constraints into account, we have to create an actual route. We want to use a more accurate estimator and thus create a new flight value estimator in Section 4.3.2.

4.3.2 New flight value estimator

We want to develop a new flight value estimator that has an increased accuracy, but is fast enough to be feasible to use in practice. The estimator should take into account the origin and destination of a flight, the size of the helicopter coverage area, and the maximum number of visits per location. To take into account the limited speed of a helicopter that restricts it to fly only to neighbouring hexagons in one time interval, we conclude that we need to construct an actual feasible route.

Since it takes too much time to create feasible routes with an exact method, we need a heuristic. The heuristic should create a feasible route between the given start and end location. We thus propose to use a heuristic that starts at the start location and stepwise determines the next best location to fly to, given the current location. We identify several restrictions to the route. The heuristic should make sure that the route finishes at the end location. The possible locations where the helicopter can fly to should thus be limited to make sure that the helicopter can always get to the destination in time. The heuristic should also count the visits of the route to different hexagons and be able to limit the possible next hexagons to the hexagons with less visits than the maximum number of visits. We propose to make the last constraint a soft constraint, since we rather have a feasible route with some hexagons visited more often than allowed, than an infeasible route. We identify two options to determine the helicopter location in the next time interval, given the current location:

1. The helicopter flies to the location that has the highest incident coverage in the next time interval, and is allowed given the restrictions.
2. The helicopter flies in the direction of the biggest hotspot in the Netherlands, when this is allowed by the route restrictions.

The first option is a greedy heuristic that looks only one time interval ahead. The heuristic first determines the locations $m \in L$ that the helicopter can go to from its current location, according to several criteria. First, the helicopter should be able to reach the locations in one time interval. Second, the locations should enable the helicopter to reach the destination in time, which means that the distance to the destination should be equal or less than the remaining time in the flight. Finally, locations are only allowed when the helicopter has not visited the location for the maximum number of times in earlier time intervals. For every allowed location, the greedy heuristic determines the forecast that the helicopter would cover on that location given the coverage function, by Equation 4.5. We use $f_{l,t}$ to denote the value of the forecast at a given location $l \in L$ and time $t \in T$. $Coverage_{l',t}$ is the coverage of a location $l' \in L$ at time $t \in T$ by next possible location $l \in L$.

$$forecast\ coverage_l = \sum_{l' \in L, t \in T} f_{l',t} * Coverage_{l',t} \qquad \text{Equation 4.5}$$

Next, a greedy approach is used which sends the helicopter to the location with the highest forecast coverage. The greedy heuristic commonly creates a route in which a helicopter flies straight to the centre of a hotspot, when part of the hotspot is within the coverage area of the helicopter. This works particularly well when we use several aggregation rings in the spatial aggregation step. However, this option does not know where to go when all its options have zero forecast coverage or when all possible locations to fly to are already visited the maximum number of times.

The second option determines the location of the biggest hotspot in the next time interval, and flies in the direction of that hotspot, regardless of the coverage of the location it flies to. Since the policy recognizes a hotspot only after it grows bigger than all other hotspots, at which point it is probably close to its peak time after which it will diminish again, we expect that the helicopter commonly arrives (too) late. The speed with which hot spots grow and shrink is based mainly on the settings of the temporal aggregation step.

We recommend to combine both options. The flight value estimator thus uses a greedy approach, but when all allowed neighbouring locations provide no forecast coverage, the helicopter flies in the direction of the biggest hotspot in the Netherlands.

Figure 4.12 shows that the new flight value estimator almost perfectly estimates the actual value of a 1.5 hour flight with a total coverage distance of 4 minutes at a given start time: the value of the estimator is commonly within 1% of the value of the optimal route.

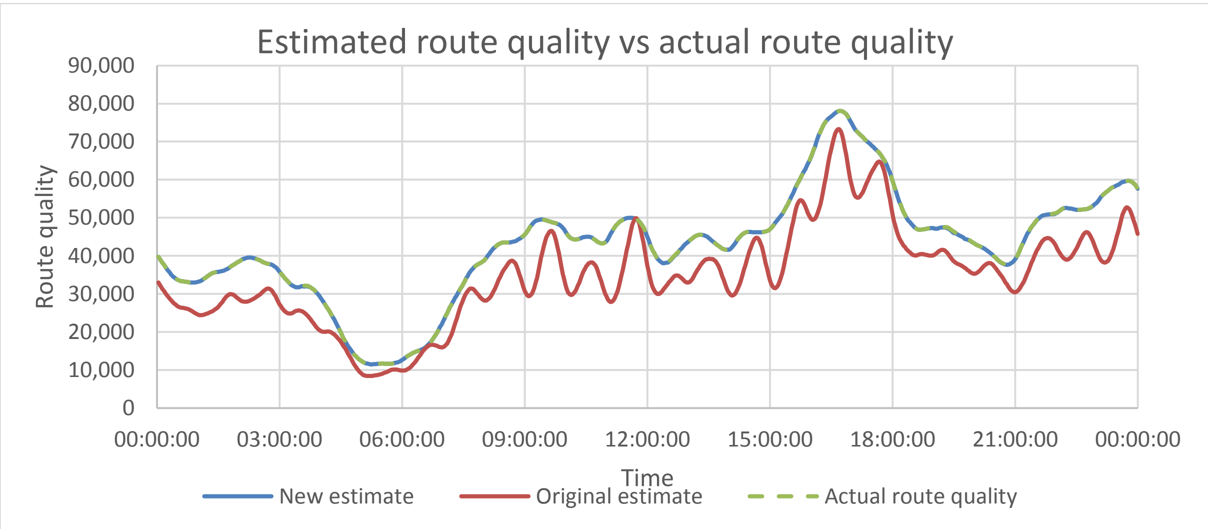

Figure 4.12: performance of the new and original flight value estimators.

We find that the new estimator is slower than the previous flight value estimator by Van Urk (2012): it takes half a second to determine the route quality of one time interval. Two estimations per second is fast enough for tactical planning as well as operational use, since it takes approximately 6 minutes and 40 hours to create a one-day and annual forecast respectively. Additionally, Appendix E describes a method to decrease the computation time. Since the quality of the solutions given by the flight value estimator approaches the optimal solution given by the MILP, it is in this case also possible to use the greedy heuristic to determine the actual flight routes. However, currently we are not sure whether this accuracy also holds for all forecasting methods.

4.4 Allocation of stand-by and airborne helicopters

This section discusses the current method to determine optimal routes and the effect of standing by. The performance of the routing method depends highly on the success function that is given as input, which thus requires evaluation. Section 4.4.1 discusses the current definition of the forecast area, Section 4.4.2 evaluates the current success function, and Section 4.4.3 proposes a success function for standing by.

4.4.1 Airborne helicopter coverage and scoring

The relation between the arrival time of a police helicopter to an incident and the probability of a successful intervention by that helicopter is unknown. Therefore, the area covered by a helicopter is uncertain. By increasing or decreasing the helicopter coverage function we can thus determine the number of incidents covered by the helicopter. Furthermore, the estimation of the probability of a successful assist has an impact on the routing model since it is the main influence on the behaviour of the helicopter. For example, the current function of the success factor by Buiteveld (2011) implies a 100% success factor for all incidents within 10 minutes flying of the helicopter location. Since the helicopter can fly in every direction, the coverage area of the helicopter has the shape of a circle. Figure 4.13 shows that a helicopter, e.g., hovering between Amsterdam and Rotterdam is thus theoretically able to perfectly cover all incidents in Rotterdam, Amsterdam, Den Haag, and Utrecht.

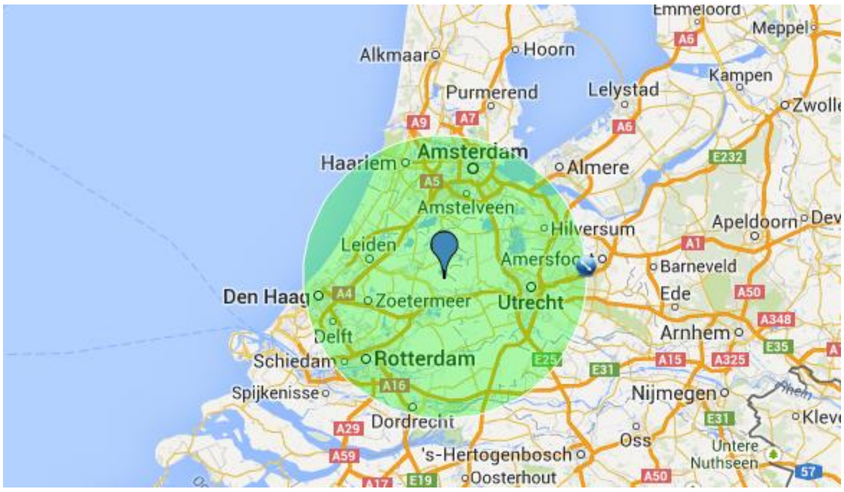

Figure 4.13: graphical representation of the 100% successful coverage area (arrival time ≤ 10 minutes) of a helicopter flying between Rotterdam and Amsterdam.

To cover incidents in multiple cities, the helicopter is best positioned between the cities. Therefore, the average distance to the incidents is larger than when the helicopter would be positioned above the biggest hotspot. The probability to catch criminals in the act thus decreases. The function for the probability of a successful assist thus influences the average proximity of the helicopter route to the criminal hotspots. A more pessimistic function would result in a smaller radius of the helicopter's coverage and would thus result in helicopters staying closer to hotspots.

4.4.2 Successful deployment function

During this research we found that the original function of the success factor of a helicopter deployment by Buiteveld (2011), which depends on the arrival time of the helicopter, has several flaws when we apply it to flying helicopters to determine optimal surveillance routes.

First, the success function ignores the preventive effect of the helicopter: criminals do not rob a bank with a helicopter hanging overhead. Since intended criminal activities right below a helicopter are delayed or relocated it is not always best to fly directly above a hotspot to maximize the incident coverage.

Second, helicopters are never 100% effective and cannot handle multiple incidents at the same time. The helicopter effectiveness is thus related to the criminal intensity on the ground (when there are 2 incidents at the same time, the maximum possible performance is 50%) and the performance of ground forces in cordoning off an area in which the suspect must be. However, the model currently assumes that all incidents within 10 minutes of the helicopter are completely covered.

Third, when the police officers on the ground have effectively cordoned off an area in which they know there is a criminal, then the arrival time of the helicopter is not decisive for the success of the helicopter: police helicopters have found criminals in closed down areas after hours of extensive search. However, the current function assumes that incidents more than 15 minutes flying from the closest helicopter positioning will never lead to an arrest.

Fourth, since it is independent from the time that the helicopter is airborne already, we ignore the probability that a helicopter is occupied by an earlier criminal incident (further explained in Appendix F), or that the maximum distance for which the helicopter can provide coverage depends on the fuel that is left in the helicopter.

To determine the effect of the success function on the behaviour of the system, we propose to test alternative definitions and compare the differences. As a function for comparison, we propose an exponential decay for the success function. Equation 4.6 shows an alternative function that is based on the underlying logic that when the LVP has to fly two times as long to get to an incident, then they are probably two times less successful.

$$f(x) = \min(1, \frac{1}{x}) \qquad \text{Equation 4.6}$$

We can implement a preventive effect into Equation 4.6 by adding a constraint. Equation 4.7 is an example of the assumption that of all incidents within 2 minutes of the helicopters, 90% is cancelled.

$$f(x) = 0.1 \quad x \leq 1 \qquad \text{Equation 4.7}$$

Figure 4.14 shows the probability of successful deployment to an incident to which the arrival time in minutes is on the x-axis, for the three identified functions for helicopter deployment success. In Chapter 7 we investigate the effect of these different success functions on the resulting set of routes.

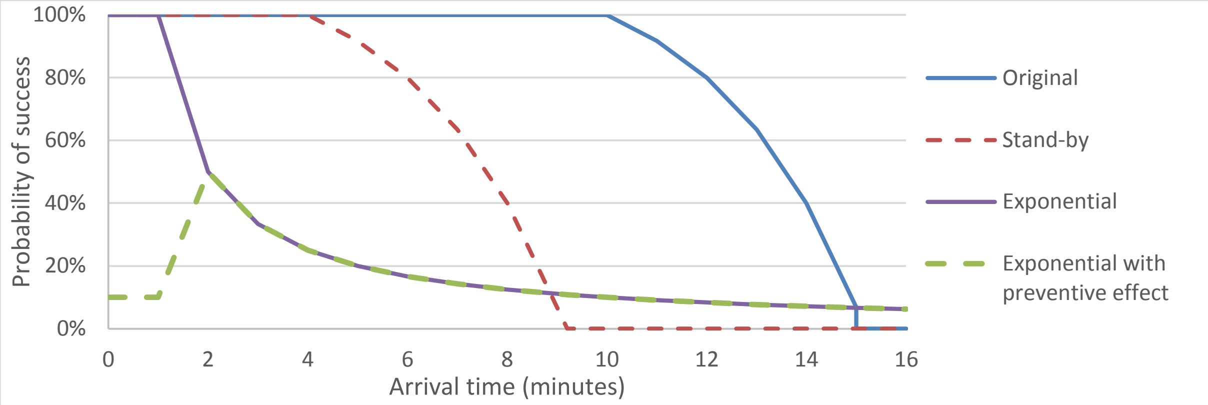

Figure 4.14: Graphical representation of original and proposed estimated success functions of a police helicopter.

4.4.3 Stand-by helicopter allocation

The original success function as discussed in Section 4.4.2 was proposed by Buiteveld (2011) to determine the value of stand-by locations and used by Van Urk (2012) to determine the optimal route of airborne helicopters. The success function defines that incidents in locations with a maximal flight distance of 15 minutes are covered by helicopters that are already airborne. We call 15 minutes the *maximal coverage distance* and 10 minutes the *total coverage distance*. Since starting from stand-by from, e.g., Schiphol takes the first 6 minutes, we create a stand-by coverage function by subtracting the first six minutes of the in-flight success function. The maximal coverage distance from stand-by is thus 9 minutes.

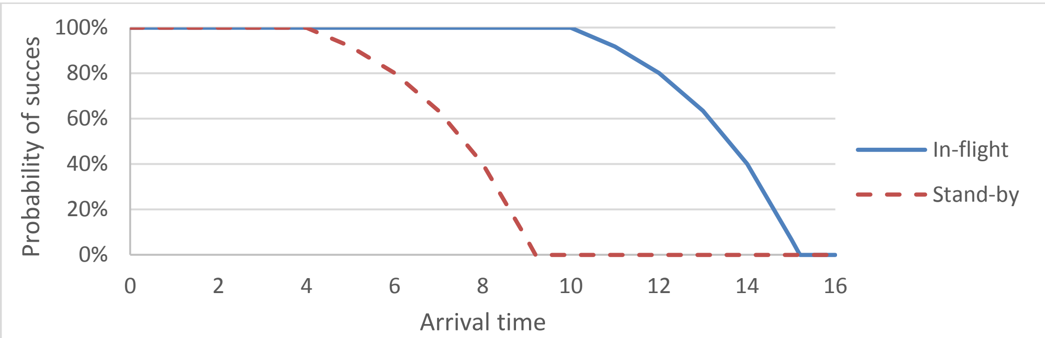

Figure 4.15: relation between the in-flight and stand-by success functions

When the *total coverage distance* is less than the start-up time from stand-by, then the total coverage distance of the stand-by location becomes negative. This means that there are no locations that are completely covered by a helicopter on stand-by. In Chapter 5 we use the stand-by success function to determine when and where to plan stand-by shifts.

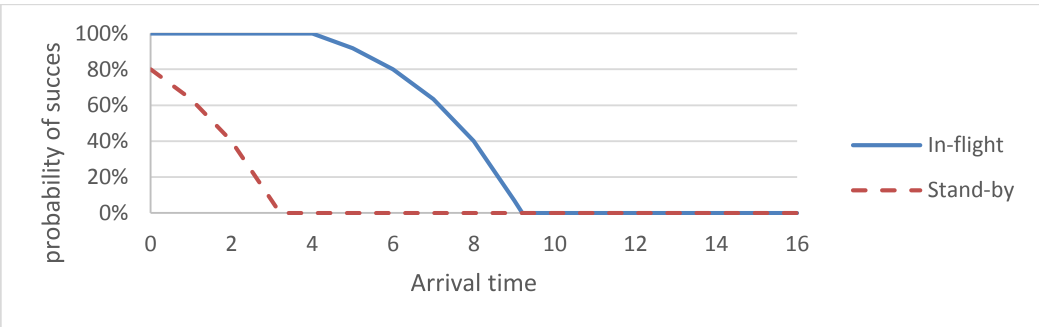

Figure 4.16: effect on total coverage distance when the start-up time from standby is longer than the total coverage distance.

Standing by versus flying surveillance flights

Knowing the coverage from stand-by and from being airborne, we can determine the added value of starting a flight at every point in time. For every time interval we can determine the expected added value of a flight that starts at that time interval, and we can determine the total coverage from stand-by during the flight duration at the start location.

The current flight value estimator determines the optimal start time of a flight based on the number of incidents in the neighbourhood of the start location. However, incidents close to the start location can also be covered by a helicopter on stand-by. Therefore, we propose to adjust the estimations for the possible stand-by coverage and determine the time interval with the maximum added value of being airborne over standing by. Equation 4.8 shows the correction formula, where $fv_{l,t}$ is the total coverage of a flight that starts at location $l \in L$ at time $t \in T$, fd is the flight duration, and $sv_{l,t}$ is the coverage value of standing-by at location $l \in L$ at time $t \in T$.

$$fv_{l,t} = fv_{l,t} - \sum_{t}^{t+fd} sv_{l,t} \qquad \text{Equation 4.8}$$

4.5 Conclusion

This chapter discusses the current forecasting and routing methods and proposes additions and improvements. We extend the scope of the forecasting model and include seasonal and weekly incident patterns. Additionally, we introduce methods to take into account time based helicopter effectiveness and different priorities for incident types. To improve the timing of flights we discuss a greedy heuristic as a flight value estimator. Finally, we discuss the impact of the success function on the helicopter routing behaviour and alternative success functions, to compare both in Chapter 7.

5 Tactical planning

This chapter describes the tactical planning problem of the LVP and provides methods to create a tactical plan. Figure 5.1 provides an overview of this chapter. Section 5.1 describes the tactical planning problem and provides notation. As discussed in Chapter 3, we can solve the tactical planning problem at different levels of accuracy. Section 5.2 discusses two approaches to tactical planning we consider: the rough-cut and integrated tactical planning approaches. In Section 5.2 we conclude that the integrated tactical planning and operational scheduling approach provides enough advantages over the rough-cut approach to justify the added complexity.

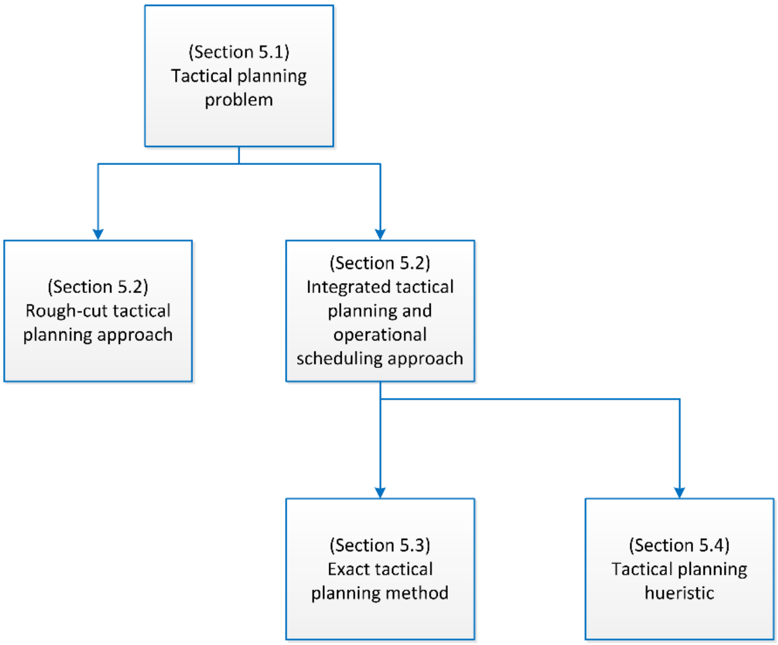

Figure 5.1: overview of the structure of Chapter 5.

Section 5.3 presents an exact method for the integrated approach, to solve the tactical planning model under the assumption that there is only one helicopter airborne at the same time. Section 5.4 proposes a tactical planning heuristic that enables the planning of multiple simultaneously airborne helicopters, but does not guarantee an optimal solution. Section 5.5 draws conclusions on the findings in this chapter.

5.1 Tactical planning of police helicopters

Tactical planning of police helicopters is the allocation of surveillance flights and stand-by shifts over multiple days. To create an appropriate tactical planning method, we first extend the notation and provide a formal description of the tactical planning problem. The LVP has a set of helicopters H, and each helicopter $h \in H$ is on a location $l \in L$ at every time interval $t \in T$. Every location $l \in L$ can handle multiple helicopters. The scope of a forecast is divided into days $d \in D$, and every day consists of multiple shifts $s \in S$. Every shift is manned by at least one crew $c \in C$ at one of the locations L. Shift starts at ts_s and end at te_s. A crew can fly for at most k hours; in the Netherlands this is 5 hours. The LVP uses the same fixed flight duration f_d for all flights and the number of flights thus depends on the surveillance flight hour budget fb and the fixed flight duration f_d . The number of crews is set by the crew budget cb. The LVP cannot make flights during a shift from a location, when there is no crew scheduled on that location for that shift.

The objective function of the tactical planning problem is the same as the objective function by Van Urk (2012) as given in Equation 2.5. However, coverage is now given by a combination of surveillance routes and crews on stand-by. There are some practical constraints that limit the maximization of the objective function. Currently the LVP commonly uses a single location to start and end surveillance flights (Schiphol), but in the future it could be able to use multiple locations (Schiphol, Rotterdam, and Volkel). We assume that the LVP makes only "single flights", where the origin and destination of flights are equal. When a crew would fly a multiple-flight route, it would for example fly from Schiphol to Volkel, and from Volkel back to Schiphol. We assume that the LVP currently flies a standard route that starts at Schiphol and flies over Amsterdam, Utrecht, Rotterdam, Den Haag, Leiden, and finishes at Schiphol again. The LVP is obliged to provide 24/7 stand-by coverage from Schiphol. Crews that are scheduled on a location during a shift are available for deployment from the start time of the shift. When there are two crews stand-by on a location at the same time, there is no added value for the second crew. When one of the crews makes a flight, then the other crew takes over the stand-by coverage.

This research considers two types of tactical planning. Variable tactical planning determines the distribution of resources over the scope of the planning problem, without constraints. Flat tactical planning creates a uniform distribution of resources over the scope. In the case of the LVP, we flatten the number of flights per day as much as possible.

Next to the incident data, the decision support system will need information regarding the number of crews to schedule, the shift schedule, the maintenance plan for the helicopters, and the bases that are available to land or fuel helicopters on. We assume that the helicopter bases in the model are always available to start and end flights at. For every helicopter, the LVP should be able to determine per day whether the helicopter is available or not.

5.2 Tactical planning approach

This section discusses the approach to tactical planning we take. We identify two possible approaches to tactical planning: the rough-cut tactical planning approach and the integrated operational and tactical planning approach. Section 5.2.1 discusses the rough-cut tactical planning approach and Section 5.2.2 describes the integrated tactical planning approach. Section 5.2.3 concludes on the choice between both approaches.

5.2.1 Rough-cut tactical planning approach

On the tactical planning level, there is a higher level of uncertainty than on the operational planning level. Therefore, it is common that a tactical planning is made on an aggregate level. Detailed scheduling is often avoided at this level, since it can lead to countless adjustments when detailed information becomes available.

An aggregated capacity planning requires an aggregated forecast of demand. Since we already have weekday and month factors on an hourly basis, we propose to use these factors to make an aggregate forecast per day of the scope. We can then determine, e.g., the 1300 hours in the year with the highest expected criminal intensity. This provides insight into the distribution of criminal intensity over the year and on which moments in the year the LVP should focus. Furthermore, it can provide insight into the effect of taking into account the difference between effectiveness in daylight hours and non-daylight hours. Figure 5.2 shows the number of flight hours per day for two examples of a tactical plan where we fly one hour with one helicopter on the 1300 most criminal hours of 2014, according to the aggregated forecast. It shows clear weekly and seasonal patterns, and is thus an indication of the added value of tactical planning. The figure contains two flight hour distributions, one without a correction for daytime and night time effectiveness as discussed in Section 4.2.4.1, and one with the correction. It shows that the correction for daytime effectiveness leads to more flights in July, although we expect flights to shift from the summer to the winter. Instead, we find that day-time flights in the winter are replaced by night-time flights in the summer.

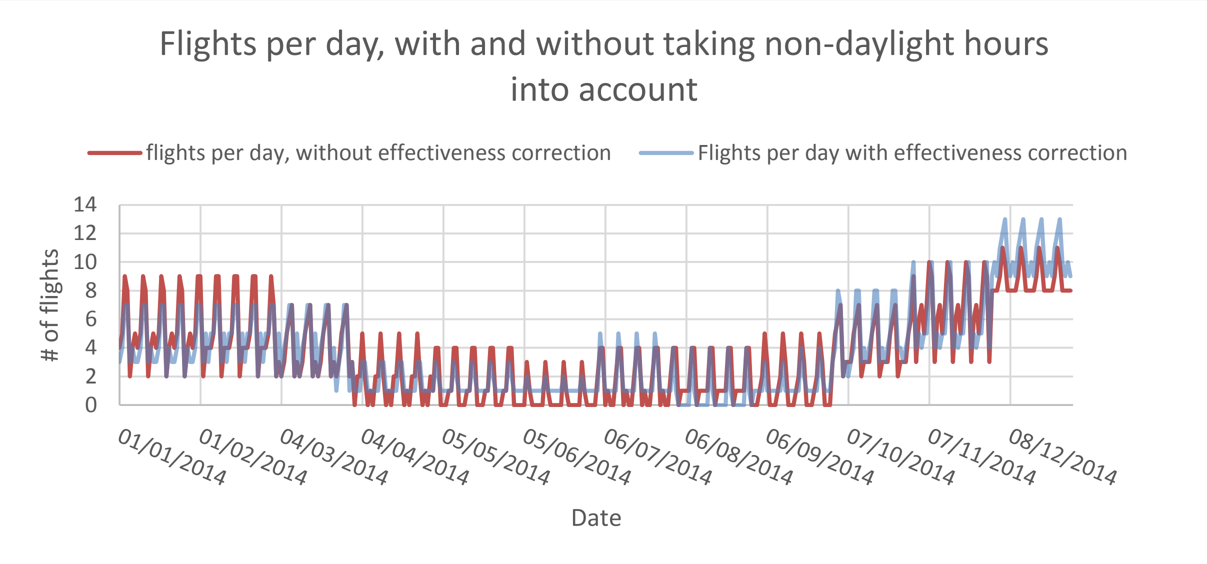

Figure 5.2: rough-cut capacity plan, based on the forecasted 1300 most criminal hours of 2014.

Although the rough-cut capacity planning gives valuable insight into the distribution of criminal intensity over days, weeks, and months, and provides feedback on how the LVP could improve the tactical planning, we identify the following disadvantages of this rough tactical planning approach:

- It does not provide an indication for how many airborne helicopters are useful at the same time during the hour.
- It does not provide feedback on the use of different locations.
- It does not distinguish between geographical patterns. A flight over one moderate hot spot in criminal intensity can be more beneficial than a flight in an hour where the total criminal intensity is higher but the criminal intensity is spread out and not all hot spots can be covered.

5.2.2 Integrated tactical planning and operational scheduling approach

Instead of making a tactical planning on an aggregate level, we can also make a more precise tactical planning that takes operational decisions into account, by extending the operational planning tool. Instead of making optimal schedules for one day forward, we can use the same methods to make an optimal planning for a whole year in advance. This is a more complex approach that will result in more rescheduling during the year due to factors such as the weather, but it enables the LVP to determine the optimal planning with the possibility of multiple simultaneously airborne helicopters and locations. Furthermore, the flight value estimators and routing module can distinguish between geographical patterns, where the rough-cut capacity approach cannot. Finally, it allows the LVP to take constraints from crew scheduling and maintenance scheduling into account.

5.2.3 Conclusion

We conclude that a rough-cut approach to tactical planning provides some insight into the added value of tactical planning for the LVP. However, based on the advantages of the integrated tactical planning approach, we propose to create a decision support system that creates a complete annual operational plan.

5.3 Exact tactical planning method

Van Urk (2012) proposed a Mixed Integer Linear Program to solve the helicopter routing problem, but found that this problem is too big to solve practical sized problems, based on a scope of one day and without taking shifts and crews into account. We find that, with some pre-processing to know the total coverage value of a flight and the value of standing by at a location at a given time in advance, we can determine the best set of flights with an Integer Linear Programming (ILP) formulation. We use the following notation:

- $A_{l,s,c}$ Availability of crew $c \in C$ at location $l \in L$ during shift $s \in S$
- $F_{l,t,c}$ Whether crew $c \in C$ starts a flight from location $l \in L$ at time $t \in T$
- $S_{l,t,c}$ Whether crew $c \in C$ is stand-by on location $l \in L$
- $sv_{l,t}$ The value of standing-by on location $l \in L$ at time $t \in T$
- $fv_{l,t}$ The value of starting a flight from location $l \in L$ at time $t \in T$
- fb Flight hour budget
- cb Crew budget
- fd Standard flight duration
- k maximum number of flight hours by a crew during one shift
- $tstart_s$ start time of shift $s \in S$
- $tend_s$ end time of shift $s \in S$

And we use the following objective function of the optimization problem:

$$maximize\ coverage = \sum_{l \in L, t \in T, c \in C} (fv_{l,t} * F_{l,t,c} + sv_{l,t} * S_{l,t,c})$$

Subject to the following constraints:

1. When a crew starts a flight from a location then they should be scheduled on that location during the corresponding shift. Every shift has a start and end time and when a flight starts between these times, the crew should be available for that shift.

$$A_{l,s,c} \geq F_{l,t,c} \quad \forall\, l \in L, s \in S, c \in C, \mathrm{tstart}_s \leq t \leq tend_s$$

2. When a crew is stand-by on a location then they should be scheduled on that location during the shift in which they are stand-by.

$$A_{l,s,c} \geq S_{l,t,c} \quad \forall\, l \in L, s \in S, c \in C, \mathrm{tstart}_s \leq t \leq tend_s$$

3. There is only one helicopter airborne at the same time.

$$\sum_{t'=t}^{t+fd-1} \sum_{l \in L, c \in C} F_{l,t',c} \leq 1 \qquad \forall\, t \in T$$

4. There is a limited crew budget.

$$\sum_{l \in L, s \in S, c \in C} A_{l,s,c} \leq cb$$

5. There is a limited flight budget.

$$\sum_{l \in L, t \in T, c \in C} F_{l,t,c} \leq fb$$

6. Every crew c can fly or be on stand-by on one location, not both at the same time.

$$\sum_{l \in L} (F_{l,t,c} + S_{l,t,c}) \leq 1 \qquad \forall\, t \in T, c \in C$$

7. Crews can only fly k hours per shift.

$$\sum_{l\in L,t\in T} F_{l,t,c} * fd \leq k \qquad \forall\, c \in C$$

8. There can only be one crew on stand-by at a location at a given time.

$$\sum_{c\in C} S_{l,t,c} \leq 1 \qquad \forall\, l \in L, t \in T$$

9. Crews are scheduled on at most one location and shift combination.

$$\sum_{l\in L,s\in S} A_{l,s,c} \leq 1 \qquad \forall\, c \in C$$

10. Constraints on the values of variables.

$$F_{l,t,c}, S_{l,t,c}, A_{l,s,c} \in \{0,1\} \qquad \forall\, l \in L, t \in T, s \in S, c \in C$$

This ILP formulation simultaneously determines when and where crews should be stand-by and from when and where flights should start. The ILP requires the value of flying and standing by at every location and time combination. We assume that the value of standing-by at a location does not depend on a flight that another crew can make at the same time from another location, or a second crew can make from the same location.

Furthermore, we assume that we know the added value of a flight in advance. This implies that there is only one helicopter airborne at the same time. When the optimal routes of multiple flights that overlap in time are determined independent, the coverage areas of the helicopters can overlap. The extent to which routes can overlap depends on the overlap in time of the flights and the start and end locations. We thus find that the value of $fv_{l,t}$ becomes uncertain when multiple helicopters can be airborne simultaneously. To determine the flight value $fv_{l,t}$ under the assumption of at most one airborne helicopter at any given time, we can use several techniques. We can determine for every location and time combination the optimal route and the forecast coverage of this flight. However, this takes too much computation time. Alternatively, we can use the faster flight value estimator as discussed in Section 2.3.2. When we create the forecast with month and weekday factors (see Section 2.3.1), we can use a third technique. The forecast technique with weekday and month factors uses all available data to forecast the first day of the forecast, and all data again to forecast for the second day of the forecast. The weekday and month distribution factors then correct the forecast to create the specific distribution over time of criminal intensity, but do not influence the geographical distribution. Since the geographical pattern per time interval is thus the same for different days, the optimal route is also the same for the same start time at different days. Therefore, when we find the optimal route for any given start time from every possible location for one day, we know the optimal route for the entire year. Currently, we consider 2-minute time intervals and 3 locations. We thus have to calculate for 720 possible start times and for 3 locations a total of 2160 routes. We then use these routes to determine the added value of making a flight or standing-by at every time interval during the scope. The calculation of the forecast coverage is faster than the calculation of the optimal route, and we thus provide exact flight values in a feasible computation time.

Finally, we assume that helicopters start and end flights at the same location and are stand-by during all times that they are not flying. We thus ignore refuelling times and consider crews available to respond to emergencies also when they have already flown their maximum number of hours for that shift. This kind of practical restrictions can be introduced into the ILP with extra constraints, which we do not discuss here.

We want to relax the assumption that there can be only one helicopter airborne at any given time. Since $fv_{l,t}$ then becomes uncertain, this exact formulation is not sufficient for the tactical planning problem. To create feasible solutions to this problem, Van Urk (2012) describes a sequential flight planning heuristic. In Section 5.4 we extend this heuristic.

5.4 Iterative heuristic method for tactical planning

We have a forecast method that is able to create a forecast with a scope of one year. Furthermore, we have a flight value estimator that is able to find the optimal start time of a helicopter flight, given a forecast. Finally, we have a routing module that determines the optimal route, given the forecast and a start time. We now use a heuristic to determine the best allocation of the budget over the scope of one year.

Figure 5.3 shows an overview of the heuristic we propose to create a tactical planning with the possibility of multiple simultaneously airborne helicopters. Before the heuristic starts, the LVP has to determine the start date and scope of the planning, the number of flights to be planned, the number of crews available, and the availability of helicopters due to maintenance. In Step 1 of the heuristic we create a forecast for the entire tactical planning scope. Since providing 24/7 coverage from stand-by is obligated, Step 2 makes sure that there is 24/7 coverage at Schiphol, by scheduling for every shift one crew at Schiphol. We cannot schedule the rest of the crews yet to the optimal stand-by shift/location combinations, since these crews can be necessary at other locations or shifts to simultaneously make multiple flights. Therefore, we first use an iterative process to determine the next best flight until the budget is fully used. For every cycle of the iterative process, we perform steps 3 to 9.

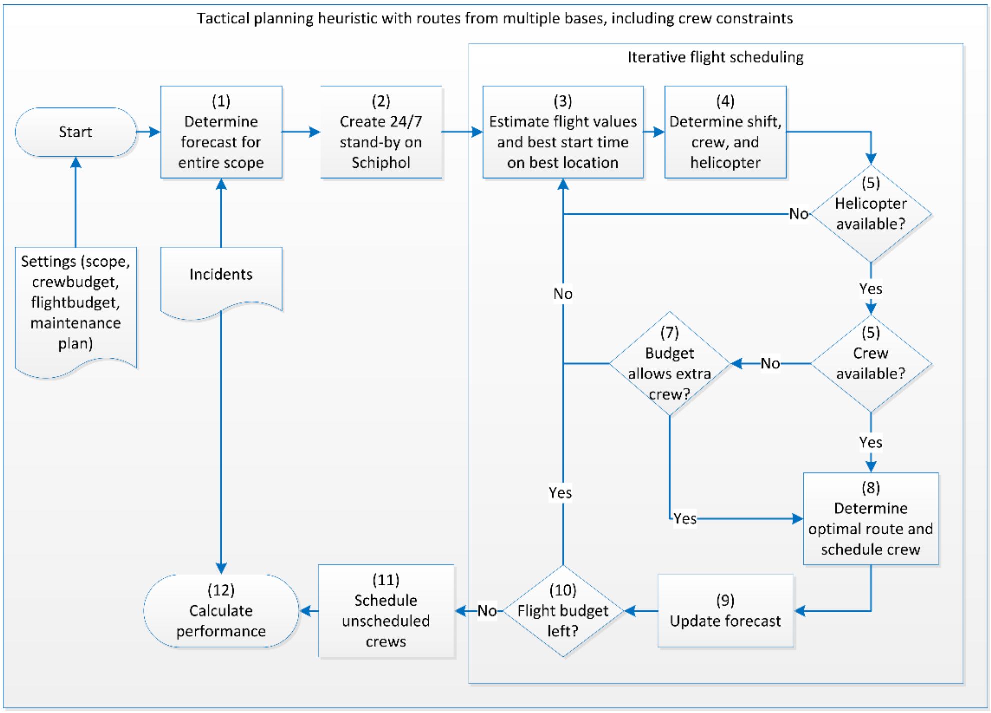

Figure 5.3: Graphical representation of the multi-location tactical planning heuristic with crew constraints.

To create a feasible tactical planning, we take into account crew and helicopter constraints while we determine in Step 3 the optimal start time and location of the next flight, and the corresponding optimal route. As Section 5.1 discusses, there are two approaches to tactical planning: flat and variable. When we create a variable tactical planning, there is no restriction on the number of flights per day. However, when we create a flat tactical planning, then the number of flights per day should be as equal as possible for all days. This is an interesting tactical planning restriction for the LVP, since it flattens the workload of, e.g., the maintenance department over the year. The maintenance department can thus handle the workload with a fixed number of employees. To create a flat tactical planning we use the following steps:

1. Before the iterative flight scheduling we determine the maximum number of flights per day. When we can distribute the number of flights equally over the number of days (e.g., 40 flights over 10 days), we use Equation 5.1 and skip step 3:

$$maximum\ flights\ per\ day = \frac{fb}{fd} \qquad \text{Equation 5.1}$$

   When we cannot distribute the number of flights equally over the days (e.g., 46 flights in 10 days or $mod(|F|, |D|) > 0$) we round up the maximum number of flights per day as in Equation 5.2:

$$maximum\ flights\ per\ day = \left\lceil \frac{fb}{fd} \right\rceil \qquad \text{Equation 5.2}$$

   When we round up the maximum flights per day, we also determine the number of days that can have the rounded up maximum number of flights ("rounded up days"). We use Equation 5.3 in which we distributing as many flights equally over the number of days and count the flights that are left (we have 6 flights left when we equally schedule 46 flights over 10 days):

$$rounded\ up\ days = mod(|F|, |D|) \qquad \text{Equation 5.3}$$

2. Before every flight we count the number of full days: days that have the maximum number of flights per day.
3. When the number of full days is equal to the number of rounded up days then we decrease the maximum flights per day by 1. We only once decrease the maximum number of flights by 1.
4. Days that are full are excluded from the possible start times.

To create a flat tactical planning of 46 flights over 10 days we thus find a maximum number of flights per day of 5 and 6 "rounded up days". When there are 6 "full days" with 5 flights, we decrease the maximum number of flights per day by 1 and fill the other 4 days to a maximum of 4 flights. When we instead create a variable planning, then the maximum number of flights per day is the number of flights that we can make in the entire tactical planning scope. We now know whether we restrict the number of flights per day and determine the corresponding set of flights.

As the optimal start time and route depend on the start location of the flight, we start by calculating the optimal start time out of the available set of start times for every location. Since the crew-helicopter combination that performs the flight cannot be stand-by on the start location during the flight, we decrease the expected route quality at a time-interval by the expected stand-by coverage during the flight, as in Equation 4.8. We then find the optimal start time not based on the moment that a helicopter flight covers the most forecast, but at the moment that the benefit of flying over standing by is highest.

Since the value of the estimated route quality is not accurate enough to determine the best location (see Section 4.3), we then calculate the optimal route for every location, given the location-specific best start time. Next, we determine the total forecast coverage of the route for the locations by Equation 5.4. We then choose the best start location.

$$total\ forecast\ coverage = \sum_{l \in L, t \in T} f_{l,t} * Coverage_{l,t} \qquad \text{Equation 5.4}$$

Step 4 determines in which shift the flight takes place and which crew and helicopter should make the flight. Step 5 checks for the best start location and start time when there is a helicopter and crew available at that location/time. When there is no crew scheduled at that location/time or when the available crews are at their maximum flight hours k, then we check whether there is budget left to schedule another crew (Step 6). When there is crew budget left, we know we are able to make a crew available. Next we check whether there is a helicopter $h \in H$ available. When there is no crew or helicopter, we cannot start a flight at the optimal start time and location and look for the next best start time and location combination.

When both a crew and a helicopter are available, Step 7 determines the optimal route and schedules the helicopter and crew. When the flight starts at Schiphol and the crew during the shift of the flight has time left to make the flight, we make the corresponding crew and helicopter unavailable for the entire flight duration. When the flight starts at Rotterdam and Volkel or when the crew on Schiphol is already at its maximal flight hours, we schedule a crew at that location for the entire length of the shift the flight starts in, and make the crew unavailable during the flight

In step 8 we update the forecast to take into account that the route covers part of the forecast, by decreasing every combination of time and location that is covered by a helicopter on the route, by the probability that the nearest helicopter is able to arrive in time and support the police officers on the ground. Equation 5.5 shows the formulation we use to update the forecast values.

$$f_{l,t} = f_{l,t} * (1 - coverage_{l,t}) \qquad \text{Equation 5.5}$$

We are now ready to create the next flight. In Step 9 we check whether there is flight budget left. When this is the case then we go back to step 3 and schedule the next flight. When the flight budget is used, we proceed to step 10. After we scheduled all routes, we determine in step 10 for every location all the times that there is at least one crew and a helicopter available, and determine the forecast covered by this stand-by availability. Finally, Step 11 is the performance measurement step. We determine the performance of the routes by the performance measures as discussed in Chapter 6.

5.4.1 Tactical planning output

The tactical planning heuristic results in a set of flights for which we know the start time, start and end location, and optimal route. Furthermore, the tactical planning determines per shift per location how many crews and helicopters are required and which crews and helicopters should make the flights.

5.4.2 Discussion

As discussed, the heuristic takes only the coverage of earlier routes of *flying* helicopters into account, while flying the designated route. This could result in a situation where a helicopter from Amsterdam flies to Rotterdam while there is a helicopter stand-by at the Rotterdam airport that covers the Rotterdam area already.

Furthermore, the heuristic does not take into account the impact on the stand-by performance while scheduling crews on shifts and locations. For example, when the heuristic chooses between scheduling a crew on Volkel to perform a flight with a forecast coverage of 5,000 and a stand-by coverage of 1,000, and scheduling a crew on Rotterdam to perform a flight with a forecast coverage of 4,800 and a stand-by coverage of 2,000, it schedules a crew on Volkel. This ensures that the tactical planning results in the best possible forecast coverage by the flights. Since a crew cannot be stand-by while it is flying, the final stand-by coverage of that crew is not known when the crew is scheduled on a location to make its first flight. Flights are scheduled sequentially and later flights by the same crew influence the time that the crew is stand-by and even the stand-by coverage of a possible second crew on the same location and shift.

As long as we are scheduling flights, we do not know for sure whether a crew that is available on a base, will stay available for stand-by (except when the crew is already scheduled for its maximum of 5 flights hours). Therefore, we cannot take the coverage of the stand-by crews into account while calculating the optimal flight routes.

5.5 Fairness

In this section we discuss how we propose to create a tactical planning that takes equity into account. For the sequential planning of simultaneous surveillance flights it is not possible to use a multi-criteria mathematical program as is common in literature. The first flight at a given time covers a part of the forecast of a number of areas and thus influences the route of later flights that (partially) overlap. Therefore, the focus and contribution of this section is the combination of equity objectives with a sequential routing model. As it is of no use to incorporate the equity objectives in the optimization of single routes, we adapt the sequential heuristic to introduce some level of equity over a set of flight routes.

We want to modify the iterative multi-route heuristic to be able to set the level of equity. As the calculation of routes depends on the forecast, we propose to adjust the forecast between flights to take into account the route of the previous flight. Furthermore, we want to visit locations at times with relative high criminal intensity. Therefore, we want to keep the temporal forecast pattern of the locations intact, while discounting the forecast of visited locations. Figure 5.4 shows how we propose to adjust the multi-route heuristic. After every flight we determine the coverage of every location by the last flight and update the forecast of the covered locations according to the total coverage of the flight (step 9). When all flights have been scheduled, we determine the level of inequity of the resulting route (step 13).

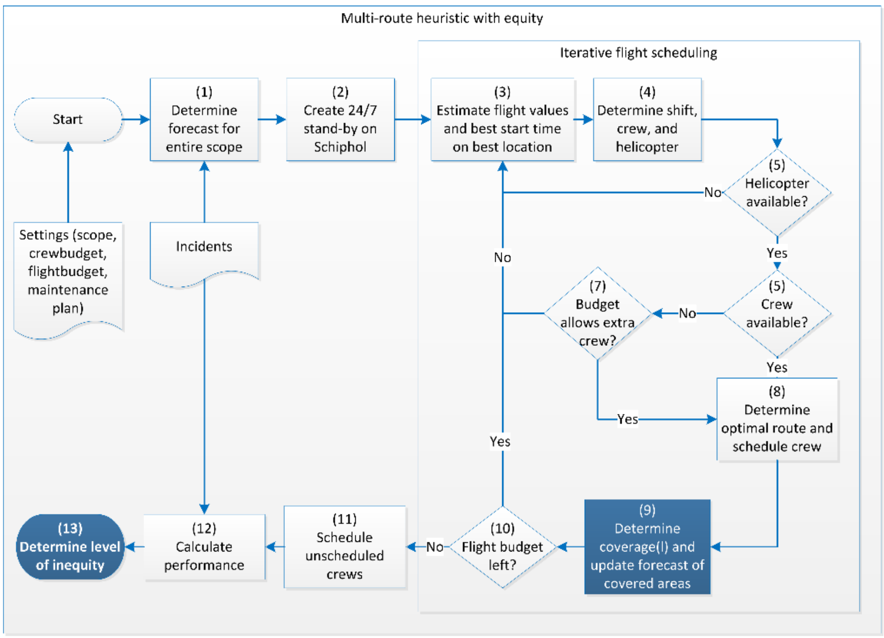

Figure 5.4: adjusted multi-route heuristic that incorporates a level of equity.

After every flight we know the optimal start time and route of that flight. We then use the success function, as discussed in Section 2.3.2, to determine to what extend the locations have been covered by the route. $flightcoverage_{l,t}$ is the coverage of location $l \in L$ by the last flight, during time interval $t \in T$. Based on $flightcoverage_{l,t}$, we can determine the coverage per location $coverage_l$, as defined in Equation 5.6.

$$coverage_l = \sum_{t \in T} (flightcoverage_{l,t}) \qquad \text{Equation 5.6}$$

To set the level of equity in step 9, we introduce the equity parameter q. Locations that have been covered should receive less attention in the next flight. Therefore, we decrease the forecast value of any location l that is covered by the last flight by the update factor $(1 - (q * coverage_l))$. Locations that have been covered more or better by the last flight are thus discounted more. For all locations that are covered to some extent, we decrease the value of the forecast for the entire year, according to Equation 5.7:

$$f_{l,t} = f_{l,t} * (1 - (q * coverage_l)) \qquad \text{Equation 5.7}$$

After updating the forecast values, we determine whether there is flight budget left. When there is enough flight budget for another flight, we go back to step 3 of the heuristic and determine the best start time of the next flight.

Since the update factor should not become negative, we note that the values that q can take are limited by the possible values of $coverage_l$. Since a flight can cover a location multiple times, influenced by the flight duration, the value of the equity parameter should take into account the length of the routes. Equation 5.8 gives the exact formulation of this constraint.

$$(q * coverage_l) < 1 \qquad \text{Equation 5.8}$$

5.6 Conclusion

This chapter discusses the police helicopter tactical planning problem. We introduce the required notation, explain the applicable constraints, and find that there are two approaches to the problem. First, we can use a rough-cut capacity planning approach that determines the expected criminal intensity for every hour in the scope of the planning and choose to fly at the best number of hours that are available for a given budget. Since this approach does not take into account the geographical distribution of crime, we choose to take a more complex approach, in which we combine operational scheduling and routing with the tactical planning problem. We provide an Integer Linear Programming (ILP) formulation of the problem under the assumption that there can be only one helicopter airborne at any given time. Since this is not a realistic assumption and we find that the problem cannot solved for realistic problem sizes when this assumption is relaxed, we propose a heuristic. The heuristic creates a complete tactical plan for helicopters and crews, and results in a set of operational surveillance routes for the entire tactical planning scope.

6 Performance measurements

This chapter discusses the performance measurements we use to evaluate the forecasting methods (Section 6.1), the flight value estimators (Section 6.2), and the tactical planning outputs (Section 6.3), based on the targets of the LVP. Section 6.4 explains how we propose to measure overtime and Section 6.5 translates the concept of equity to a measure of equity for a set of helicopter routes and crews scheduled on stand-by shifts. Section 6.6 draws conclusions on the findings in this chapters.

6.1 Forecast accuracy and performance

Since we make forecasts of periods for which we also have actual incident data, we can find the actual difference between forecast and reality, and determine the predictive value of the forecast.

Recall that an annual forecast consists of a *forecast value* for 808 locations L during 262,800 time intervals T. Since the forecast is a hotspot forecast (see Section 2.3.1), the forecast value is not the expected number of incidents at that time and location, but rather an indication for crime intensity. The distribution of expected crime intensity over the Netherlands during a time interval indicates where we expect the most criminal incidents. We use $f_{l,t}$ to denote the value of the forecast at a given location $l \in L$ and time $t \in T$.

The forecast error calculation methods, as discussed in the literature research in Chapter 0, measure the accuracy of a forecast technique in determining the actual number of incidents (priorities) at a given time and location. Recall that the forecasting methods in this research do not determine the expected number of incidents, but instead determine the expected distribution of crime over time and space. However, we can compare the distribution over time of several forecasts with the actual distribution of crime over time. To compare the forecast with the actual incident pattern we use the MSE (see Equation 3.1), since there are more than 150 incidents per day and thus no intermittent pattern. The mean squared error then determines the average squared difference between the total forecast value per day and the number of incidents per day. Since the total forecast value between forecasts with different levels of spatial aggregation differs, while the distribution of forecast value over time is equal, we first scale the forecasts. We scale the forecasts back to the total number of incidents in the scope of the forecast. We use $i_{l,t}$ to denote the number of incidents at a given location $l \in L$ and time $t \in T$. We determine the factor required to scale the total forecast value to the total number of incidents. We then multiply the entire forecast by the factor as in Equation 6.1, and thus keep the temporal and spatial distribution intact.

$$f_{l,t} = f_{l,t} * \frac{\sum_{l',t'} i_{l',t'}}{\sum_{l',t'} f_{l',t'}} \qquad \text{Equation 6.1}$$

We then apply Equation 3.1 to compare the forecasted criminal intensity with the actual criminal intensity, in which we compare the forecasted crime intensity with the actual crime intensity. We calculate the actual crime intensity with Equation 6.2 and use a comparable formula with $f_{l,t}$ for the forecasted crime intensity.

$$Y_t = \sum_{l \in L} i_{l,t} \qquad \text{Equation 6.2}$$

Since the MSE is sensitive for the total number of incidents, we note that we cannot compare the values for the MSE between datasets or for different scopes.

These measures are useful to determine the quality of a forecast that is aggregated to the day level. However, when we use them on the individual time interval and location level, these measures do not perform well. Figure 6.1 shows both a scaled forecast in the time dimension and a set of incidents. Since there are relatively few incidents in relation to the number of locations and time combinations, the discussed measures do not properly indicate the forecast quality. For example, the accuracy of a forecast that forecasts 0 incidents for every time and location combination is often better than any reasonable forecast.

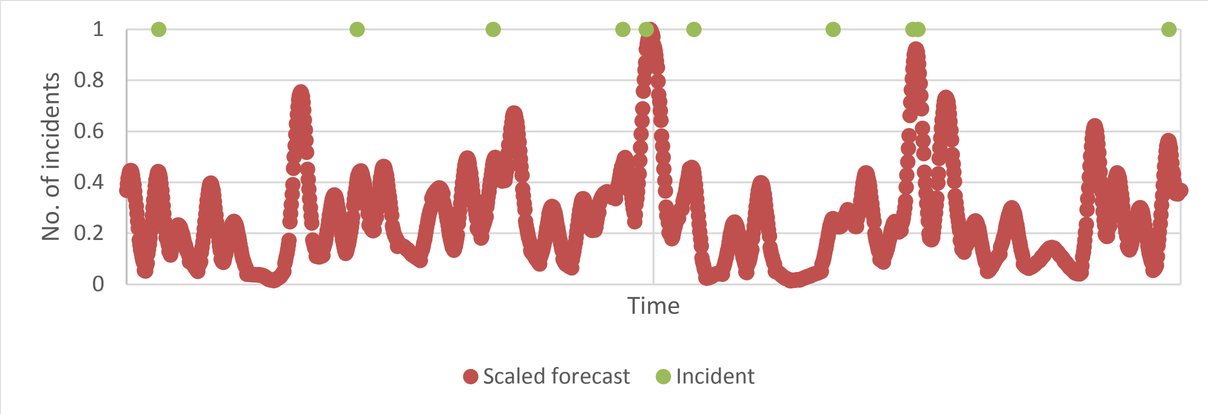

Figure 6.1: Comparison of two days of scaled forecast and the actual incidents for the hexagon over Schiphol

Furthermore, this research works with a relative forecast value: the actual value of the forecast is unit-less but when the value of the forecast is higher at time/location x than at time/location y, then we expect more incidents at time/location x than at time/location y. Therefore, we need a forecast error calculation method that determines whether incidents happen around peak moments and locations in the forecast.

Forecast timing and allocation estimation method

In this section we develop a new forecast error calculation method, which estimates whether the forecasted crime hotspots and peak times match the actual crime patterns. The assumption behind the accuracy measure is that how higher the forecast value is, the more probable is the forecast value to be part of a peak. Therefore, the higher the corresponding forecast value for an actual crime is, the higher is the probability that there is a peak in the forecast during the actual criminal incident.

We determine per historical incident within the forecast scope the value of the forecast at that time and location. The higher the forecast is at the moment and location of the incident, the higher the probability is that we would allocate capacity to that location and time. We already denoted the value of the forecast and the number of incidents at a given location $l \in L$ and time $t \in T$ as $f_{l,t}$ and $i_{l,t}$ respectively.

We now discuss several formulations that we can use to determine the forecast accuracy. The most basic formulation is given in Equation 6.3. We extend this formulation in Equation 6.4 and Equation 6.5.

$$Forecast\ accuracy = \frac{\sum_{l,t}(f_{l,t} * i_{l,t})}{N} \qquad \text{Equation 6.3}$$

Equation 6.3 sums the forecast value at times and locations where actual incidents have happened. When there are 2 incidents at the same time and location, then the equation counts the forecast value twice. We divide the sum of the relevant forecast values by the number of incidents to find the average forecast value per incident.

Figure 6.2 shows that this measure depends on the variability of the forecast (in this case influenced by the level of temporal aggregation). The figure compares the score of two incidents for two types of forecasts. Both forecast show a peak around time 20. The left incident scores 0.8 on both forecast. However, scaled forecast 1 has a narrow peak and forecast 2 has a wide peak. When we score the right incident on both forecasts, then it scores better on forecast 2 than on forecast 1. When we extend this concept then a totally flat forecast would score better than a forecast that only shows a peak at (almost) the right time.

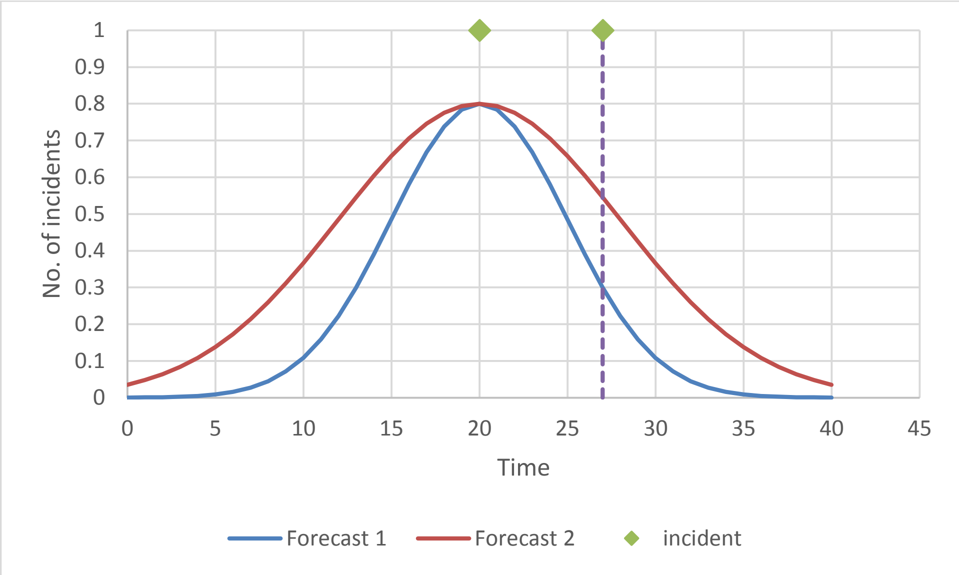

Figure 6.2: example of impact of temporal aggregation on forecast accuracy

The effect as in Figure 6.2 is also applicable to spatial aggregation. Therefore, we want to correct for the general shape of the forecast and thus propose to divide the total forecast value during incidents by the total forecast value over space and time. Equation 6.4 shows the corrected formulation.

$$Forecast\ accuracy = \frac{\sum_{l,t}(f_{l,t} * i_{l,t})}{\sum_{l,t} f_{l,t}} \qquad \text{Equation 6.4}$$

The denominator of the equation is the sum over all locations and times of the forecast value $f_{l,t}$. To take incident priorities into account we propose to adjust the equation by replacing the number of incidents by the total priority of the incidents $p_{l,t}$ at a given location $l \in L$ and time $t \in T$. Equation 6.5 shows the extended formulation for priorities.

$$Forecast\ accuracy = \frac{\sum_{l,t}(f_{l,t} * p_{l,t})}{\sum_{l,t} f_{l,t}} \qquad \text{Equation 6.5}$$

For incidents with a priority of 5, we thus count the forecast value 5 times. This implies that the value of the forecast accuracy not only measures whether incidents happen at times and locations that have a peak in the forecast, but also puts more emphasis on the forecast accuracy of incidents with a high priority, than on the forecast accuracy of incidents with a lower priority.

Use performance of planning policy for forecast accuracy estimation

Finally, we can also measure the quality of a forecast by the performance of a tactical plan (a set of surveillance flights) based on the forecast. In Section 6.3 we discuss how to determine the performance of a tactical plan.

6.2 Flight value estimator quality

To determine the quality of the flight value estimator, we determine the relation between the estimated forecast coverage of a flight at a certain time, and the actual forecast coverage of a flight at that time. Therefore, we calculate the correlation coefficient between set X with the estimated forecast coverage values, and set Y with the actual forecast coverage values. Equation 6.6 shows the formulation of the correlation coefficient.

$$\rho(X,Y) = \frac{cov(X,Y)}{\sigma(X)\sigma(Y)} \qquad \text{Equation 6.6}$$

We can use the same correlation coefficient calculation to determine the predictive effect of the total forecast coverage on the actual incident coverage of flights. Furthermore, we compare the performance of the resulting set of routes of the original and the new flight value estimator, as described in the next section.

6.3 Route and stand-by performance

This section discusses measures to determine the expected performance of the output of the tactical planning procedures. Section 6.3.1 discusses how we calculate the coverage of a set of helicopter routes and the stand-by schedule. Section 6.3.2 explains how we determine the expected proximity of helicopters to incidents. Section 6.3.3 evaluates how the LVP can combine both measures to create a fair performance measure.

6.3.1 Incident coverage

The output of the tactical planning is a set of routes with their respective start times and locations. To determine the quality of the set of routes we can compare them with the actual incident pattern of the period for which the routes are made. We note that we should use a different set of data to create a forecast and routes with, than the data set that we use to determine the quality of the routes with. Since the routes are based on the forecast, the timing of the routes depends on the flight value estimations, and the routes are made by the routing module, we determine the quality of all these elements at the same time. When we use the measures in Section 6.1 and 6.2 to find that the forecast methods and flight value estimators work correctly, and we find that the incident coverage does not perform well, then this is an indication that the tactical planning heuristic or route optimization method does not perform well. We identify multiple options to determine the performance of the set of helicopter routes.

1. We count all incidents that are covered to some extent by a helicopter on the route. Equation 6.7 shows the corresponding formula, where $i_{l,t}$ is the number of incidents at a location and time combination, and $Covered_{l,t}$ is a binary variable that indicates whether the location and time combination is to some extent covered by a route. Since a helicopter cannot cover multiple incidents at the same time, this is an overestimation of performance.

$$total\ incidents\ covered = \sum_{l \in L, t \in T} i_{l,t} * Covered_{l,t} \qquad \text{Equation 6.7}$$

2. Per incident that is covered, we determine the probability that the helicopter that is closest by is successful. Equation 6.8 shows the corresponding formulation, where $Coverage_{l,t}$ is the best coverage of a location $l \in L$ at time $t \in T$ by the set of routes, as discussed in Section 2.3.2. Since this measure uses the success function based on arrival time to determine the quality of a solution, it determines the average proximity of helicopters to incidents.

$$total\ incident\ coverage = \sum_{l \in L, t \in T} i_{l,t} * Coverage_{l,t} \qquad \text{Equation 6.8}$$

3. Since a helicopter cannot cover multiple incidents at the same time, both formulation above are overestimations. To create a real estimation we can simulate reality. A simulation model could model the reaction to incidents. For example, when there are one or more helicopters

on surveillance routes available to respond to an incident, the helicopter that is closest by and is not handling any other incidents, responds to the incident. The helicopter is then unavailable for other incidents for some time. After handling the incident, the helicopter determines the optimal route for the rest of the surveillance flight duration, from the location of the incident, and becomes available again. Since the handling duration is currently not sufficiently measured, we have no data to base this process on, which makes this measure not yet useful.

Since the LVP aims to increase its *proximity* to crimes and the second option measures just that, we propose to use the second measure when we can compare routes on only one measure. However, the combination of the first and second measure gives a more complete picture and is thus suggested. We find the third option interesting for future research.

For experiments with different priorities per incident type, or where the effectiveness per time is taken into account, we can adjust the second formulation to include priorities and/or effectiveness. Equation 6.9 shows the total coverage function where the number of incidents is replaces by the sum of the priorities of the incidents at a location and time.

$$total\ priority\ coverage = \sum_{l \in L, t \in T} p_{l,t} * Coverage_{l,t} \qquad \text{Equation 6.9}$$

Equation 6.10 extends the total priority coverage function by multiplying the expected coverage with the expected effectiveness of the time at which the incident happens e_t, as discussed in Section 4.2.4. We determine the expected effectiveness at time $t \in T$ by using the official sunrise and sunset times per date and the expected night time and day time effectiveness (see Table 4.2).

$$total\ coverage = \sum_{l \in L, t \in T} p_{l,t} * Coverage_{l,t} * e_t \qquad \text{Equation 6.10}$$

6.3.2 Average arrival time

The proximity of helicopters to incidents is one of the main aims of flying surveillance routes. Therefore, we propose to measure the expected arrival time to incidents. We assume that helicopters are only deployed to incidents that they (partly) cover. When an incident is more than 15 minutes from the closest helicopter, then, according to the current success function, the probability of successful support is zero. For every time interval we know for every location the number of incidents $i_{l,t}$. Furthermore, we know for every location at that time interval when there is a helicopter available: $L_{l,t,h}$. Finally we know the shortest distance (in minutes flight time) between any set of locations: $d_{l,m}$. We can thus determine the average arrival time for the entire set of covered incidents.

6.3.3 Combined performance measure

The total incident coverage measure and the average arrival time measure both depend on the success function. Recall that for the current success function the maximal coverage distance is 15 minutes: after 15 minutes helicopters are always too late to successfully provide support. In Section 4.4 we state that the results of the route optimization model depend on the maximal coverage distance: helicopters should choose positions closer to hotspots when we decrease the maximal coverage distance. When we determine the average distance for a smaller maximal coverage distance, it will improve since only incidents that are within the maximum distance are taken into account. However, with a smaller maximal distance the number of incidents covered and the total incident coverage will decrease.

To make a fair comparison between different maximal coverage distances or completely different success functions, we suggest to score the outcome of experiments with different coverage functions by Equation 6.11.

$$Score = \frac{total\ incident\ coverage}{average\ arrival\ time * maximal\ coverage\ distance} \qquad \text{Equation 6.11}$$

This scoring function rewards a helicopter position with more covered incidents, as well as a helicopter position that results in a short average arrival time. When we increase the maximal coverage distance, then the helicopter covers more incidents, but probably also takes longer to reach them. However, when we double the maximal coverage distance, then the total coverage area increases by a factor four. In return, the average arrival time will probably also just double. Therefore, we propose to divide the total incident coverage by the product of the average arrival time and the maximal coverage distance. When we now double the maximal coverage distance, the nominator will on average increase by a factor four and the denominator in Equation 6.11 will on average decrease by a factor 4.

6.4 Overtime

We want to determine the impact of shift schedules on the amount of overtime that crews make. After every flight, crews have to record their experiences during the flight. This takes on average approximately one hour. Therefore, crews work in overtime when a surveillance flight ends in the last hour of their shift or even after the end of the shift. A one-hour flight that starts one and a half hours before the end of the shift will thus probably result in a half hour of overtime while none of the flight time is in overtime. We propose two measures of overtime. Equation 6.12 shows the first measure that determines the percentage of total time that is worked in overtime, taking into account the registration process after flights. The expected minutes in overtime is thus the number of minutes of the hour after the last flight that happen after the end of the shift.

$$\%\ overtime = \frac{expected\ minutes\ in\ overtime}{total\ work\ minutes\ including\ overtime} \qquad \text{Equation 6.12}$$

The second measure does not take into account the registration hour since we assume that this time can be decreased when overtime proves to be dramatically increased by tactical planning. Therefore, Equation 6.13 calculates only the percentage of actual flight time in overtime. For every flight we determine the number of minutes that a flight finishes after the end time of the shift it started in. We then divide the total of all overtime minutes by the total minutes of flight to determine the percentage of flight time crews make in overtime.

$$\%\ overtime\ flight\ minutes = \frac{flight\ minutes\ in\ overtime}{total\ flight\ minutes} \qquad \text{Equation 6.13}$$

We can use these measures to determine the impact of shift schedules on overtime by comparing the expected overtime of tactical plans made with different shift schedules.

6.5 Fairness

In this section we translate the concept of equity to the police helicopter problem. The result of the tactical planning heuristic is a set of routes and stand-by shifts that defines the presence of a helicopter at a time and location. The set of routes can be summarized by the number of visits of a helicopter to a location. However, this is no indication of the timing of the helicopter and the impact that the helicopter has on the crime in a visited location. Therefore, we propose to summarize the set of routes as the forecast coverage of every location by the flying and stand-by helicopters.

To determine the equity of the tactical planning, we first translate the aspects of equity as discussed in Chapter 3 (groups, effect, and attribute) to the police helicopter problem. Recall that the LVP divided the Netherlands in 808 locations based on a hexagonal grid. We use the hexagons to define groups: every hexagon/location is a group. In the case of a flying helicopter, the effect of a helicopter route on a group is only distance-dependent: helicopters can fly in a straight line at a constant speed and are not slowed down by traffic jams. For stand-by helicopters this is the same, although the minimal arrival time to a group is affected by the response time from stand-by (see Section 4.4.3). When there are multiple helicopters available at the same time, the effect of the helicopters on location l is the coverage by the closest helicopter. The effect on a group/location at a given time $E_{l,t}$ is the largest coverage by the set of flying and stand-by helicopters. The total effect E_l on a group is the sum of the effect over all time intervals within the scope: $E_l = \sum_t E_{l,t}$. A_l is an attribute of group l against which a function of the effects can be compared or by which the effects can be weighted or scaled. In the LVP case, this can be the number of crimes that have happened in group l. For example, The LVP can determine that all locations should be covered an equal number of times (mean reference distribution) or according to their relative criminal intensity (attribute reference distribution).

We now determine how to calculate the equity of the tactical planning. Therefore, we need to select an equity measure. We assume that all groups have the right at the same percentage decrease in the overall crime level. When groups receive a higher or lower percentage decrease of crime level, this is considered unfair. We thus want to minimize the deviations from the mean. We translate these choices into the following measure of inequity:

- The deviations from the mean percentage of forecast covered per group.

Since we already scaled the measure by the number of HIC through the use of a percentage, we do not use additional scaling. Marsh and Schilling (1994) discuss multiple formulations for an equity measure without scaling and with a mean reference distribution: the absolute deviations, the variance in the deviations, the max deviation, and a combination of the log function with the variance of the deviations. Since we cannot simultaneously optimize for efficiency and equity, the calculation of the equity happens only after full completion of the heuristic and the formulation thus has no influence on the results. We thus choose the least complex calculation and sum the absolute deviations from the mean.

Since the LVP allocates its flights over a one-year scope, we calculate the equity measure over a scope of one year, without taking possible historical inequities into account. Furthermore, we find it impermissible that a location in the Netherlands is not covered at least once per year. Therefore, although the level of equity can change between tactical plans, there is one hard constraint for which we test solutions afterwards: every location/group should be covered at least once per year.

To determine the level of equity or inequity of a solution, we take into account the resulting set of routes and the stand-by availability. Since we can have multiple helicopters airborne at the same time and helicopter on stand-by also cover some locations, we determine per location and time combination which helicopter has the shortest arrival time to the location during that time interval. $coverage_{h,l,t}$ contains the coverage of location $l \in \mathrm{L}$ by helicopter $h \in H$ on time $t \in T$. Equation 6.14 is the definition of the effect on location $l \in \mathrm{L}$, in which we multiply the best coverage of the forecast values with the forecast values to determine the total forecast coverage. We divide the total forecast coverage on a location $l \in \mathrm{L}$ by the total forecast value of location $l \in \mathrm{L}$ to get the percentage of the forecast that is covered on location $l \in \mathrm{L}$.

$$effect_l = \frac{\sum_t \left(f_{l,t} * \max_h coverage_{h,l,t} \right)}{\sum_t (f_{l,t})} \qquad \text{Equation 6.14}$$

Equation 6.15 shows the corresponding definition of the level of inequity. We want to minimize the deviations from the mean percentage of forecast coverage and thus determine the mean effect $\overline{effect}$ (percentage of forecast covered per location) and the effect per location $effect_l$. The level of inequity is then the sum of the deviations of the individual effects from the mean effect.

$$inequity = \sum_l \left| effect_l - \overline{effect} \right| \qquad \text{Equation 6.15}$$

We propose to create several tactical plans with different degrees of inequity and performance, which we can measure by Equation 6.15 and Equation 6.8. The LVP can then determine which solutions has the best combination of inequity performance.

6.6 Conclusion

This chapter discusses performance measures to determine the accuracy of hotspot forecasts, the quality of the flight value estimator, and the performance of a set of routes that is the output of the tactical planning model. We find that measures like the Mean Squared Error are only applicable to aggregated hotspot forecasts and do not provide a useful measure for the accuracy of the forecasted geographic distribution of crime. Therefore, we develop an alternative measure that determines whether actual incidents correspond with the hotspot pattern of the proposed forecast methods. Furthermore, we conclude that the performance of surveillance routes is determined by their proximity to incidents and propose several measures. Finally, we discuss how overtime is measured by the LVP and how we determine the impact of shift schedules on overtime.

7 Simulation

This chapter describes our simulation study in which we use a simulation model to determine the quality of several proposed methods and the expected impact of tactical decisions on the performance of the LVP. To compare different methods, we perform multiple experiments with different settings, and compare the outcomes.

Section 7.1 describes the simulation plan, consisting of the goals, experiments, and parameter settings. Section 7.2 explains the verification procedure of the simulation model. Section 7.3 presents the results from simulation. Finally, in Section 7.4 we discuss the limitations of the simulation model and draw conclusions on the simulation results.

7.1 Simulation plan

This section describes the approach to simulation that we propose. We discuss how we use the available data to create input for the simulation and validate the results in Section 7.1.1 , describe the goals we want to reach by simulation (Section 7.1.2), create the necessary planning policies (Section 7.1.3) and set up a set of experiments to accomplish these goals (Section 7.1.4).

7.1.1 Simulation steps

The simulation process starts with forecasting. We determine a forecast value for every location in the hexagonal grid over the Netherlands over the scope of the experiment. For every experiment we use a scope of one year (365 days), to create a complete year plan.

We then apply a planning policy, which is a heuristic like the tactical planning heuristic as discussed in Chapter 5, in which we model a possible approach to tactical planning by the LVP. For example, in Section 7.1.4 we suggest a planning policy for flat tactical planning, and a planning policy for variable tactical planning.

To be able to validate the results of the forecasting procedures and the helicopter timing and routes that result from the planning policy, we want to compare actual incidents against the forecast and the planned routes. Therefore, we use a validation period: we exclude the last year of the available data. Figure 7.1 shows how we use part of the data to create a forecast, that we use during the simulated planning process, and use the rest of the data to determine the performance of the planning results.

Since the LVP has four years of incident data (from 1-1-2010 until 31-12-2013), it is possible to make a forecast and route planning one year ahead on 1-1-2013, based on three years of data. As Figure 7.1 shows, we can then compare the forecast and route planning over 2013 with the actual incident data of 2013 in the simulation module and determine the performance. We thus do not use the data of 2013 to create the forecast of 2013, to be able to objectively compare the forecast with the actual incidents in 2013. To determine whether the effects that we see are significant or due to random errors, we want to perform multiple measurements. Therefore, we also create simulations with one-year forecasts based on two years of data that start at 1-1-2012 and 1-7-2012.

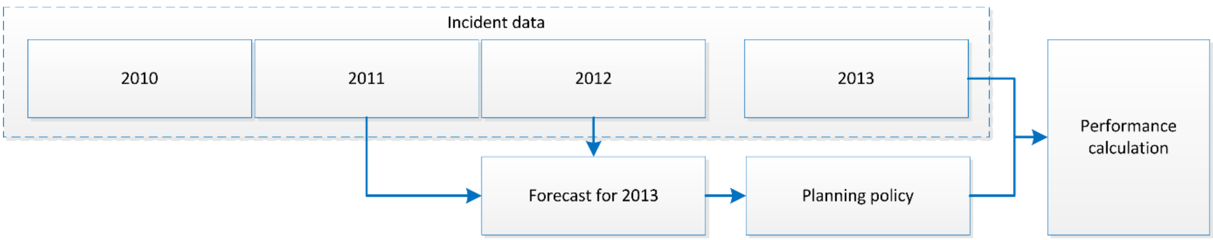

Figure 7.1: Visualization of simulation steps and validation period.

7.1.2 Simulation goals

We identify multiple goals to achieve with the simulation model. The first goal is:

1. To determine the effect of different forecast methods on the forecasting speed and the quality and performance of the resulting forecast. Furthermore, we vary the level of spatial and temporal aggregation, and the forget factor. The forecast methods, as discussed in Chapter 4, are:
    a. The extended daily forecast with monthly and weekly trends, for every day of the one-year scope.
    b. The fast forecasting method with month and weekday factors (Section 4.2.1).
    c. The forecast method without month and weekday factors, which uses only data of corresponding weekday and month combinations (Section 4.2.2).

Second, we determine the effect of different stand-by and tactical planning policies on the expected performance of the LVP (goals 2 to 10):

2. The effect of different stand-by tactics, as discussed in Section 7.1.3.
3. The added value of flying standard surveillance routes versus only standing by.
4. The added value of timing surveillance flights versus using a standard departure time.
5. The added value of the optimal route calculation versus the current standard route.
6. The added value of variable tactical planning versus flat tactical planning.
7. The cost of personnel and maintenance constraints (crew budget and helicopter availability).
8. The added value of using Rotterdam and Volkel as start and end locations of flights.
9. The effect of different total coverage distances for the current success function on the allocation of flights on bases and surveillance routes (Section 4.4.1).
10. The added value and performance of the new flight value estimator (Section 4.3).

Third, we determine the effect of adjustments to the current forecasting, routing, and tactical planning methods (goals 11 to 16):

11. The effect of different fixed flight durations.
12. The effect of a different type of helicopter coverage formulation (Section 4.4.2).
13. The impact of priorities on the forecast and route quality (Section 4.2.4).
14. The impact of taking day and night effectiveness into account (Section 4.2.4).
15. The impact of constraining surveillance routes to prevent hovering (Section 2.3.2).
16. The impact of several degrees of fairness in tactical planning on performance (Section 5.5) .

To achieve all goals we perform multiple experiments in which we make combinations of the forecasting, start time estimation, routing, and planning techniques. Since it takes between 8 and 20 hours to run one experiment, we cannot perform a full factorial analysis in which we perform experiments for every possible combination. Therefore, we propose a sequential approach with 3 steps. In this approach, we first determine the best forecasting method, which we then use to determine the best tactical planning policy, which we then use to determine the effect of, e.g., incident priorities. By using a sequential approach we cannot guarantee to find the best combination of techniques and we, e.g., ignore the effects of different forecasting techniques on the added value of tactical planning, but are still able to estimate the effect of different techniques on the tactical planning results.

Figure 7.2 shows the combinations of methods and policies we perform experiments with. For example, to determine the added value of the route optimization MILP by Van Urk (2012) we perform experiment 19 in which all flights follow a standard route as set by the LVP (planning policy

2), and experiment 20 in which we use the MILP to determine the optimal route for every flight (planning policy 3), and determine the difference in performance. We keep all other settings and methods equal to find only the effect of the routing method. Section 7.1.4 provides a complete overview of the experiments and their settings.

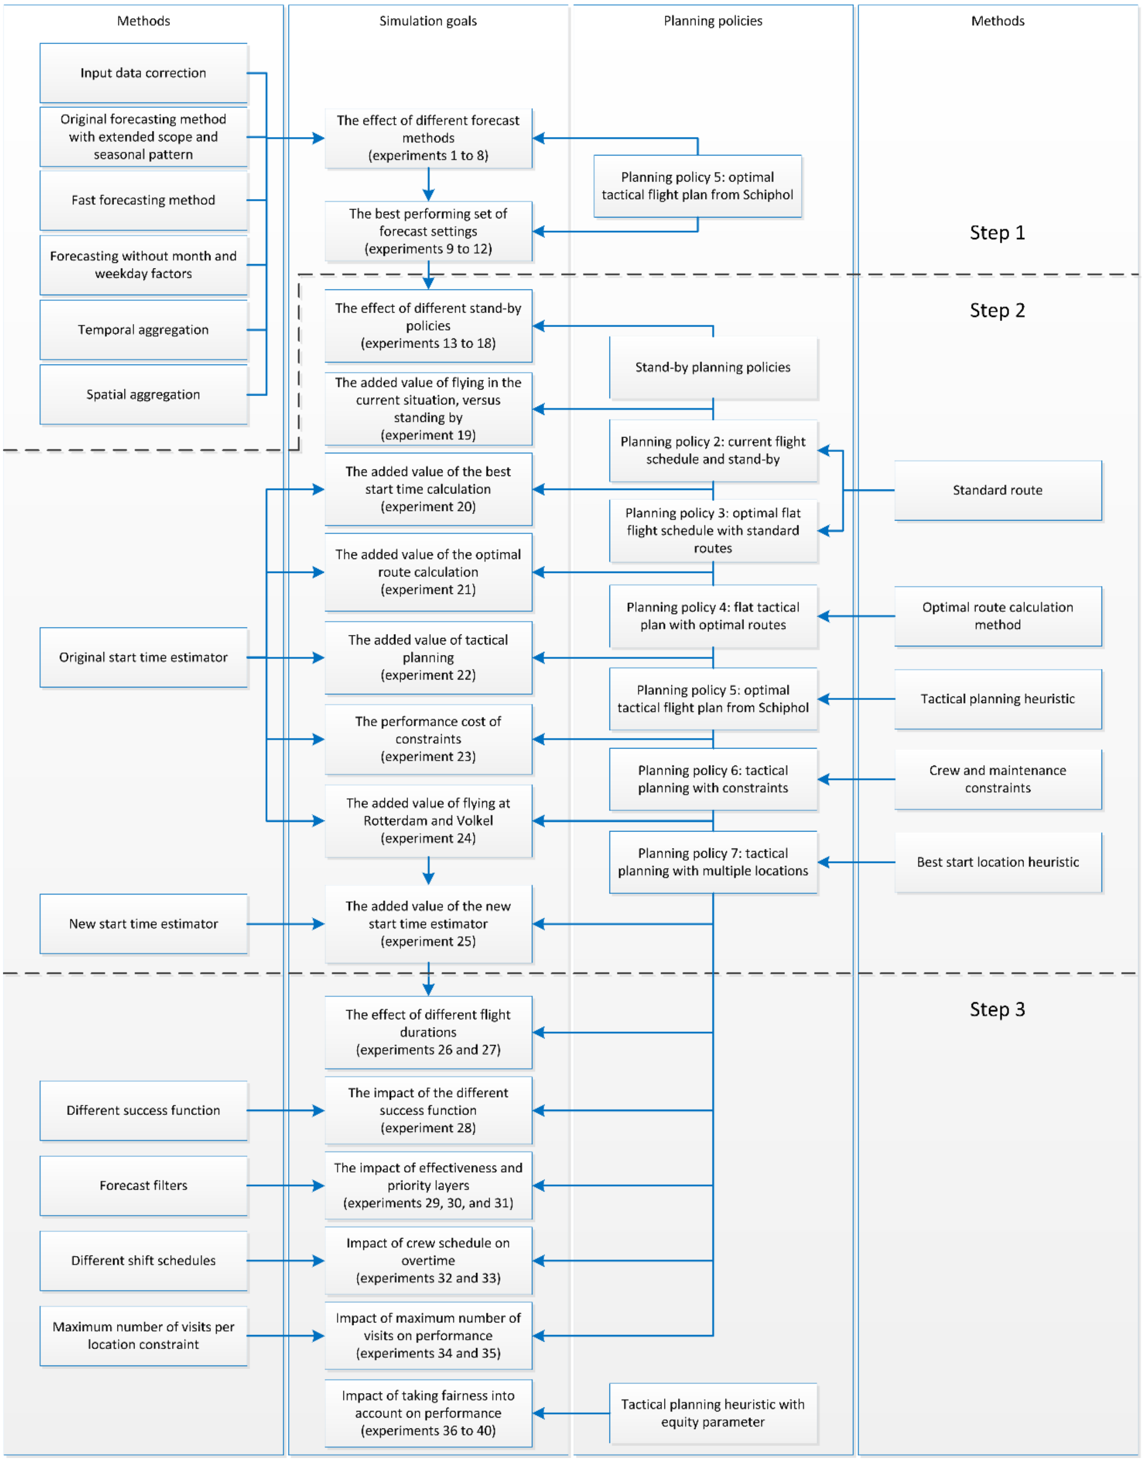

Figure 7.2: Overview of goals, planning policies, and methods.

For experiments 13 to 25 we compare the impact of different possible tactical planning strategies and thus create several stand-by policies in Section 7.1.3 and overall planning policies in Section 7.1.4. Stand-by policies are part of the overall planning policies. For every experiment we thus choose a combination of the stand-by and overall planning policies.

7.1.3 Stand-by policies

To determine the relative quality of the current stand-by policy, we introduce six stand-by policies. We use the stand-by policies to determine the added value of more stand-by time and standing by on other locations than Schiphol. Policy 1 is the basic policy in which the LVP provides 24/7 stand-by coverage from Schiphol. The difference between policy 1 and 2 shows the added value of stand-by shifts on other locations while the LVP already provides 24/7 stand-by coverage from Schiphol. The comparison of policy 2 with policy 3 shows the added value of enabling more flexibility in the location of the 24/7 stand-by coverage. Policy 4 then enables the LVP to plan without the obligation of 24/7 stand-by coverage. Policy 5 schedules stand-by time instead of shifts and shows the cost of the inflexibility due to shift blocks. Finally, policy 6 provides the maximal stand-by coverage the LVP can provide given the current three locations, and thus provides a reference point to determine the performance of the other stand-by policies.

1. The LVP is 24/7 stand-by on Schiphol only.
2. The LVP provides 24/7 coverage from Schiphol and provides additional optimal coverage from Rotterdam and Volkel as long as the crew budget allows it to.
3. The LVP optimally allocates stand-by crews over shifts and locations, with guaranteed 24/7 coverage from a location.
4. The LVP uses the optimal allocation of stand-by crews, without any constraints.
5. The LVP optimally allocates the available time for stand-by, without taking shift blocks into account. It is thus possible to be only 2 minutes stand-by on Volkel on a particular Sunday.
6. 24/7 coverage from Schiphol, Rotterdam, and Volkel. This is the best possible stand-by performance given the current set of locations.

For every stand-by policy we define a heuristic that determines the best possible allocation of stand-by time. In these heuristics, we do not schedule flights. We now discuss the heuristic of stand-by policy 2 as an example. To determine the added value of having a crew stand-by during a shift on a location, we first create a forecast. Since, for stand-by policy 2, there is always a crew available at Schiphol, we immediately schedule a crew on every shift on Schiphol. In step 3 we check whether there is budget left after step 2. When there is budget left, we determine in step 4 for the other locations (in this case Rotterdam and Volkel) for every shift the forecast coverage by a crew on stand-by for the entire shift $sv_{l,t}$. In step 5 we distribute the crews that are left in the budget over the shift/location combinations with the highest values through an iterative procedure: we determine the shift/location combination with the highest stand-by coverage and schedule a crew during that shift on that location. When there is still crew budget left, we do the same for the next crew. After having allocated all crews, we determine the performance of the stand-by crews by comparing the stand-by availability with the actual incident pattern in step 6. Figure 7.3 shows a graphical representation of this policy.

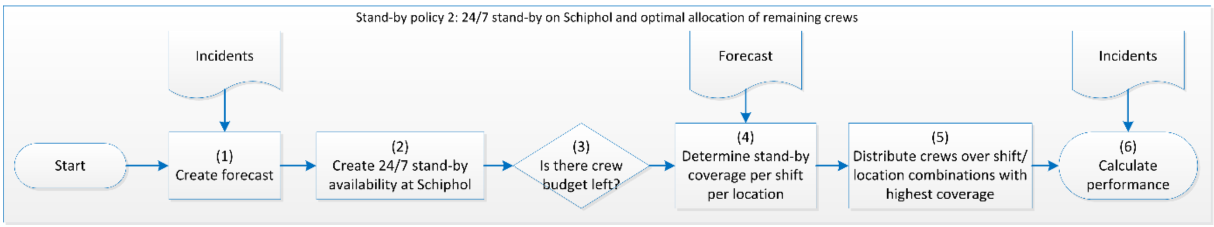

Figure 7.3: Graphical representation of the planning policy 1 heuristic.

For the tactical planning policies in the next section we commonly use stand-by policy 2, since this policy is most similar to the current practice at the LVP. Since the 24/7 coverage from Schiphol is required, we first schedule these crews. We then iteratively schedule flights, for which we need extra crews when they start at other locations than Schiphol, or when there is no crew available at Schiphol to perform the flight. After scheduling all flights, we determine the number of crews that is left in the budget and schedule them on shifts with the objective to cover as much forecast as possible that is not covered by the surveillance flights.

7.1.4 Planning policies

To compare different planning processes and methods, we define a set of *planning policies*. A planning policy consists of a set of tactical decisions such as the stand-by policy, what kinds of flights to make from which bases, and which practical constraints to take into account. Planning policies represent the possible interventions that the management of the LVP can make. Recall from Section 2.1.3 that a crew is in-flight or stand-by. A crew that makes a flight is thus not stand-by.

Planning Policy 1: 24/7 stand-by at Schiphol and optimal (variable) allocation of extra crews.
The first planning policy is the base policy in which we only take stand-by availability into account. With this policy we determine the effect on the performance when the LVP would not make surveillance flights. We use stand-by policy 2 to determine the allocation of the crew budget, in which all available crews are stand-by on the ground for emergency incidents and the LVP makes no surveillance flights. There is always a crew available at Schiphol and the rest of the crew budget is divided over the shifts on Rotterdam and Volkel with the highest forecast coverage.

Planning Policy 2: Current planning of ***single flights*** *from Schiphol with a* ***standard route***
The second planning policy is an estimation of the current practice within the LVP. The LVP flies standard routes that start and end at Schiphol at fixed moments of the day, irrespective of the weekday or month. Since the route and start times are fixed, we do not use the flight value estimator and the route optimization MILP. Again we use stand-by policy 2 to allocate the crews to shifts.

Planning Policy 3: Flat tactical planning of single flights from Schiphol with a ***standard route***
In the third planning policy, we add the timing of the surveillance flights based on a forecast. We create a tactical planning and distribute the available flight hours as equal as possible over the days in the planning horizon. Since the route is fixed and the calculation of the exact value of a known route at any time is fast, we determine the exact value of starting a flight at that time for every time interval in the scope of the policy. Since it can occur that flights overlap in this policy, we *sequentially* schedule the flights. The flights start at the optimal time and follow a standard one and a half hour route over from Schiphol via Utrecht, Rotterdam, The Hague, and Amsterdam, back to Schiphol.

Figure 7.4 visualizes the steps in the planning policy. In steps 1 and 2, we determine the maximum number of flights per day and create a forecast for one year ahead. Given stand-by policy 1 or 2, we then schedule one crew on Schiphol during every shift to create 24/7 stand-by coverage in step 3. When we use stand-by policy 3 or 4, we do not schedule crews yet. In steps 4 to 9 we sequentially schedule flights. For every flight we determine the best start time (step 4) and check whether the day of the optimal start time is not full yet (step 5). If there is room for a flight left in the flat tactical planning and when there is a crew and helicopter available (step 6), we schedule the flight in step 7, and update the forecast in step 8 with Equation 5.5. We make the crew unavailable for stand-by during its flight. if there is no crew or helicopter available, we exclude the start time from future analysis and find the next best start time. After scheduling all flights, we schedule the remaining crews for stand-by, according to the stand-by policy. Finally, we determine the performance of the set of routes and stand-by crews in step 11, by comparing the availability of helicopter support with actual incident patterns, as discussed in Section 6.3.

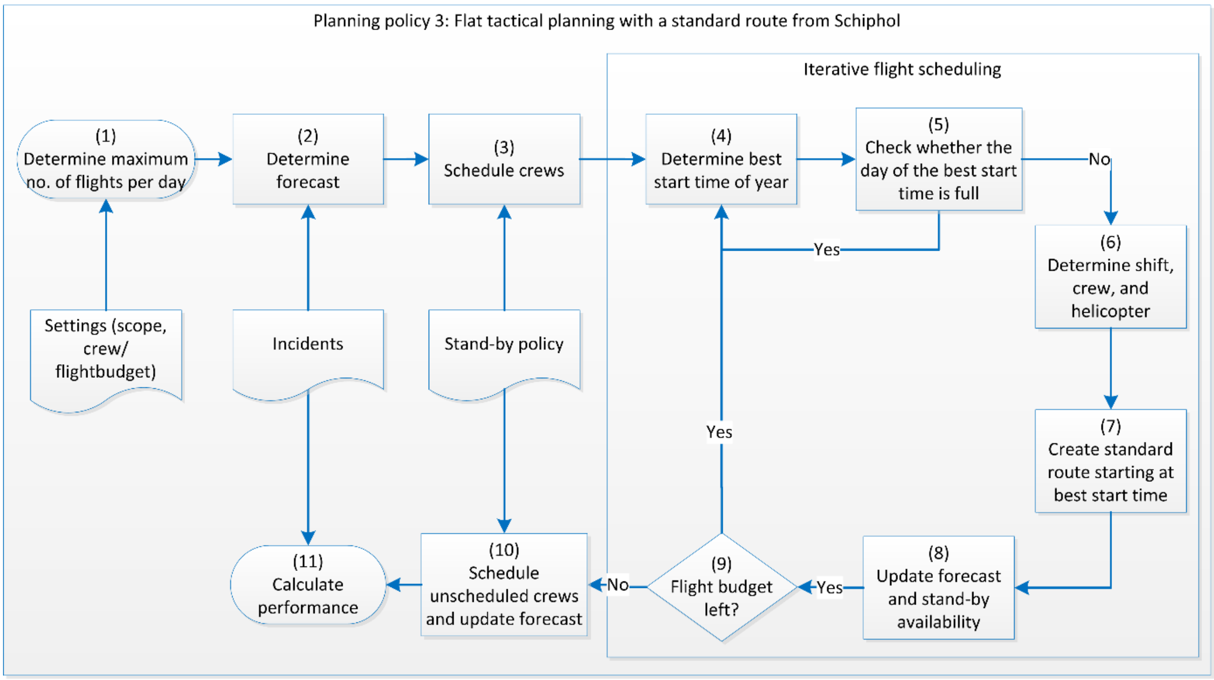

Figure 7.4: Graphical representation of the third planning policy heuristic.

*Planning Policy 4: Flat tactical planning of single flights with **optimal routes***

The fourth planning policy introduces the MILP by Van Urk (2012), as discussed in Section 2.3.2, to determine the optimal route. The heuristic is almost the same as the heuristic for planning policy 3, except that we replace the standard route in planning policy 3 for the optimal route as calculated by the routing module of the operational planning tool. Since we do not know the optimal route per time interval in advance, we use the flight time estimator, as discussed in Section 2.3.2, to estimate the added value of a flight at every time interval to find the optimal start time.

*Planning Policy 5: Variable tactical planning of single flights **without constraints***

The fifth policy introduces variable tactical planning. We relax the constraint on the maximum number of flights per day and the planning policy is free to variably allocate the flight budget over the year and calculates the best set of flights that is possible within the hour budget, without taking crew scheduling and maintenance into account. Since we do not take crew scheduling into account, we can increase the crew budget as long as it is necessary to create a better flight plan.

*Planning Policy 6: Variable tactical planning of single flights from Schiphol **with constraints***

Figure 7.5 shows the sixth policy, which is more realistic than the fifth policy and introduces crew and helicopter constraints. When there is no crew available, we check if the budget allows another crew. When there is no helicopter or crew available, and there is no crew budget left, then we look for the next best flight. This planning policy should result in a feasible flight schedule and the comparison of the expected performance with policy 4 shows the performance cost of the added constraints.

*Planning Policy 7: Variable tactical planning of single flights from **multiple bases***

This policy considers all known locations to start flights from and is discussed in Section 5.4. When we use the original flight value estimator, we cannot rely on the estimated flight value to determine the best location to start from. Therefore, we determine the best start time of the year for every location and calculate the corresponding coverage value of one single flight route (to and from the same location). We then compare the routes for the different start locations, schedule the route from the best location, and update the forecast. Figure 5.3 visualizes the seventh planning policy heuristic.

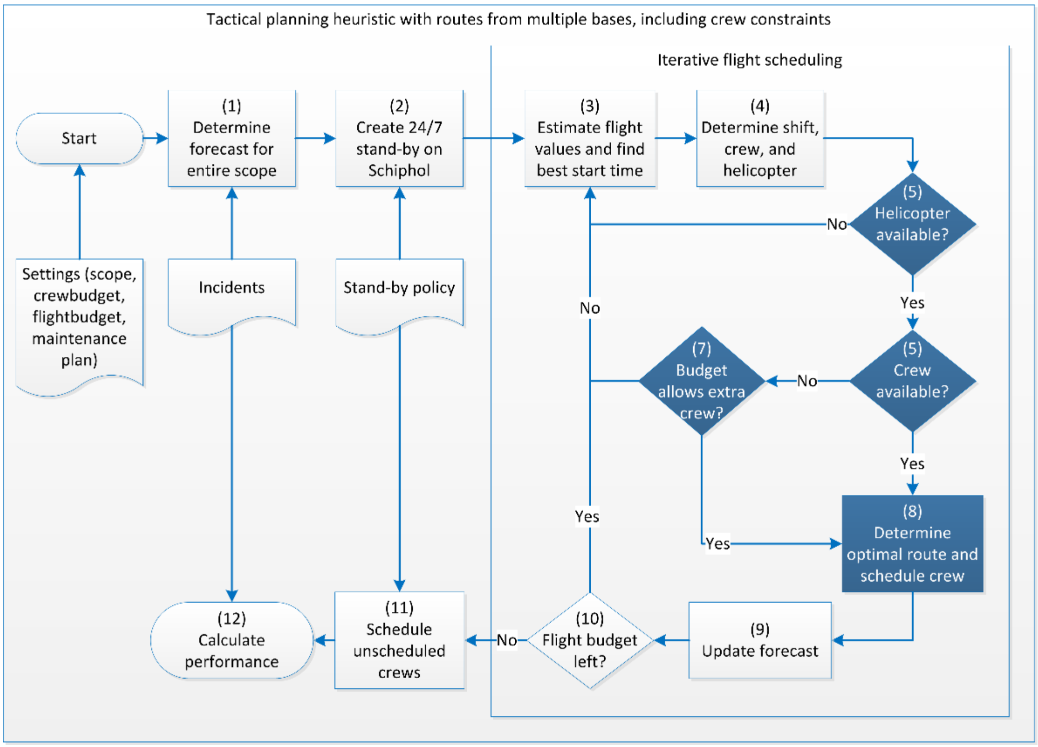

Figure 7.5: Graphical representation of the sixth planning policy heuristic.

7.1.5 Experimental setup

Since it takes between 8 and 20 hours to do one experiment, we cannot perform experiments for every combination of the forecasting methods, flight value estimators, planning policies, etc. Therefore, we propose a sequential approach. We follow the sequence of the goals and determine the corresponding required experiments. An important factor in the setup of experiments is that we do not change more than one settings between two experiments that we compare, to make sure that the difference between the results is due to the one change only.

Step 1: determine optimal forecast method (goal 1)

In Chapter 4, we identify several forecasting methods. First, the original forecasting method of Van Urk (2012) that uses all available historic data to create a forecast for every day in the scope. Second, the fast forecasting procedure transforms one day of forecast by the weekday and month factors to create forecasts for other weekday/month combinations. Finally, there is the forecasting method that only uses historical data from incidents that happened on a Monday in January to create a forecast for a Monday in January. We call it the “weekday per month” forecasting method.

For every forecasting method there are three possible settings. First, we can adjust the forget factor. Second, we can determine the level of spatial aggregation by choosing the number of hexagon rings that we spread the incidents over, see Figure 2.5. Finally, we can switch the aggregation in time on and off. To compare between different forecasting methods and levels of temporal and spatial aggregation, we use an initial constant forget factor of 0.03. We compare the results experiments 1 and 2 to determine the effect of fast forecasting. We then use different levels of spatial aggregation to determine the added value or cost of spatial aggregation in experiments 3 and 4. In experiment 5 we stop using temporal aggregation to determine its effect on the forecast. Experiments 6, 7, and 8 indicate the added value or cost of the weekday per month method, and we test again the effect of spatial and temporal aggregation. We perform all experiments for three time periods of a year, starting at 1-1-2012, 1-7-2012, and 1-1-2013, to determine whether the results are consistent for different datasets and monthly patterns.

Experiment	1	2	3	4	5	6	7	8
Fast Forecasting	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Temporal aggregation	Yes	Yes	Yes	Yes	No	No	Yes	Yes
Spatial aggregation	5 rings	5 rings	1 ring	0 rings	0 rings	0 rings	0 rings	1 ring
Data	All	All	All	All	All	Weekday /month	Weekday /month	Weekday /month

Table 7.1: Experimental setup to determine the effect of different forecast technique.

We use several measures to determine the impact of different forecasting methods on the resulting forecast. We record the time it takes to create the forecast (forecasting speed). We determine the seasonal and weekly distribution of the forecast by calculating the MSE for the aggregated forecast per day. We then determine the spatial and temporal accuracy by the measure as described in Section 6.1. To determine the impact on routes and the performance of routes that are based on the forecasts, we then apply planning policy 5 to the forecast, to see the effect of the forecast on the optimal route MILP. We give the experiments a budget of 100 flight hours to limit computation times and create one and a half hour flights. We set total coverage distance of the success function to 4 minutes to create routes close to hotspots, which generally results in extensive routes. Since the LVP currently provides 24/7 coverage from Schiphol, we use stand-by policy 2. Furthermore, the LVP currently has a budget for 1,500 crews and we use this budget in all experiments. Finally, we compare the tactical plan with the actual incident pattern and measure the incident coverage and the number of incidents covered when the LVP would have precisely followed the tactical plan in reality.

From the results of experiment 1 to 8 we choose the forecasting method that best fits the LVP requirements, and determine the influence of different forget factors for this method. We try several different values for the forget factor to determine the effect of the forget factor. We test two higher and two lower forget factors that should give an indication whether a higher forget factor leads to better results or not. We do not zoom in on the best forget factor, since finding the optimal forget factor is not the focus of this research. Table 7.2 shows the experiments and corresponding forget factors. As we already know the performance for the forget factor 0.03, we do not have to run a new experiment for that factor.

Experiment	9	10		11	12
Forget factor	0.005	0.01	0.03	0.1	0.5

Table 7.2: Experimental setup to determine the effect of different forget factors.

We use the same planning policy settings and performance measures as for experiments 1 to 8. Based on the computation time and performance measures, we choose the forecast method and forget factor that gives the best results in a feasible computation time.

Step 2: determine benefits of different planning policies (goals 2 to 10)

With the best forecast from Step 1, we now determine the cost or benefit of different (stand-by) planning policies that use this forecast as input. We start with measuring the effect of the stand-by policies as described in Section 7.1.3 by the experiments in Table 7.3. We assume that it takes a helicopter 6 minutes to depart from Schiphol and Rotterdam, and 4 minutes from Volkel. We perform every experiment for a total coverage distance of 4, 6, and 10 minutes and a crew budget of 1,500 shifts. As discussed in Section 4.4.1 the LVP currently assumes that a helicopter has a total coverage distance of 10 minutes, and that an airborne helicopter thus is able to cover Amsterdam, Rotterdam, Den Haag, and Utrecht at the same time. We compare this assumption with a total coverage distance of 6 minutes that enables an airborne helicopter to cover at most 2 of the big cities at the same time, and a total coverage distance of 4 minutes that restricts the coverage area of a helicopter to at most one of Netherlands biggest cities.

Experiment	13	14	15	16	17	18
Stand-by planning policy	1	2	3	4	5	6

Table 7.3: experimental setup to determine the effect of different stand-by policies.

We measure the performance of the various stand-by policies by the number of incidents covered and the incident coverage. Based on these results, we now know the performance of only standing-by and determine the effect of flying, according to different planning policies.

1. We determine the benefit of flying over standing-by, by comparing planning policies 1 and 2.
2. We determine the benefit of timing surveillance flights, by comparing policies 2 and 3.
3. We find the benefit of optimal routing over a standard route by comparing policies 3 and 4.
4. We identify the potential benefit of tactical planning by comparing policies 4 and 5.
5. We calculate the cost of crew scheduling constraints by comparing policies 5 and 6.
6. We estimate the benefit of using Rotterdam and Volkel as start locations of flights by comparing policies 6 and 7.

We already determined the performance policy 1, since it is equal to stand-by policy 2, and can thus start with tactical planning policy 2. Table 7.4 shows that we run six experiments, one for every planning policy.

Experiment	19	20	21	22	23	24
Tactical planning policy	2	3	4	5	6	7

Table 7.4: experimental setup for tactical planning policies.

To enable objective comparison of the policies, we use for every planning policy 5 helicopters, stand-by policy 2, a crew budget for 1,500 crews, 1:30 hour flights, a 2,190 flight hour budget (1,460 flights), and no limitations on the maximum number of visits per location. We determine the performance of the different planning policies through the same measures as we used for the stand-by policies: incident coverage and number of incidents covered.

Step 3: determine effect of the new flight value estimator and routing settings (goals 11 to 16)

We then use planning policy 7, since it is the most complete, in experiments to determine the impact on the performance of the improved flight value estimator (experiment 25) and different average flight durations (experiment 26 and 27). In experiment 28 we determine the impact on surveillance routes of the negative exponential success function with the preventive effect, followed by the effect of taking priorities and effectiveness into account in experiment 29, 30, and 31. Furthermore, we perform two experiments with different crew schedules to determine their impact on the expected overtime, two experiments with different levels for the maximum number of visits constraint, and five experiments to determine the trade-off between equity and efficiency. We thus perform a total of 40 experiments.

We measure the performance of the resulting set of routes and stand-by crews for all experiments by the number of incidents covered and the incident coverage. Furthermore, we determine the correlation between the expected performance of flights and the actual forecast coverage of flights (see Section 6.2) for experiments 24 and 25 to determine the effect of the new flight value estimator.

7.2 Verification

Verification is the step in the development of a simulation model that ensures that the model works correctly. We perform several checks to see whether the model works as it should. First, we have created several visual checks of forecasts and routes:

a. We compare every forecast with its input data and check whether it is spatially and temporally correct.
b. For every flight we check whether the route starts and ends at the right location and whether the route does not skip hexagons (since it can only fly to neighbouring hexagons in one time step).

Second, we have run the entire code step by step and ran several debugging procedures.

c. Every line of code in the simulation model has run for thousands of times without giving errors. Furthermore, the code contains parts that check whether the input and output are correct (for example, whether the tactical plan provides 24/7 coverage).

Finally, we have presented the outputs of the model to the LVP experts to verify whether the results match reality. The LVP experts concluded that the results match reality sufficiently.

7.3 Simulation results

This section discusses the results from simulation experiments. We first discuss the effect on the quality of different forecast methods and settings and choose the best method (Section 7.3.1). Second, we use the best forecast to test the effect of different planning policies (Section 7.3.2). Section 7.3.3 then discusses the effects of several extensions.

7.3.1 Step 1: the effect of forecast methods and settings

To compare the different forecast methods, we perform the experiments as defined in Table 7.1. Table 7.5 is a copy of Table 7.1 to enable quick comparison between experiments. We use the first two experiments to determine the impact of fast forecasting. Experiment 2 to 5 show the impact of different degrees of spatial and temporal forecasting. Experiments 6 and 7 are examples of the weekday per month forecast method.

Experiment	1	2	3	4	5	6	7	8
Fast forecasting	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Forget factor	0.03	0.03	0.03	0.03	0.03	0.03	0.03	0.03
Temporal aggregation	Yes	Yes	Yes	Yes	No	No	Yes	Yes
Spatial aggregation	5 rings	5 rings	1 ring	0 rings	0 rings	0 rings	0 rings	1 ring
Data	All	All	All	All	All	Weekday /month	Weekday /month	Weekday /month

Table 7.5: forecast method experiments.

Table 7.6 shows the results of the experiments with a scope of one year. To check whether the results hold for different datasets, we perform all experiments except the first three times, for the start dates 1-1-2012, 1-7-2012, and 1-1-2013. We use colours to indicate the relative quality of a measure compared to other forecasts. For example, we want to minimize the MSE and maximize the hotspot accuracy. We note that we can only compare results in the same row to determine the effect of forecasts. Since there is no difference between the results of experiments 1 and 2 for the start date 1-1-2013, we conclude that it does not cost accuracy or incident coverage to use the fast forecasting method (experiment 2), and we do not create forecasts with the slower method for the other start dates.

		Experiment	1	2	3	4	5	6	7	8
		Forecast speed (minutes)	1,800	30	30	24	15	15	20	30
		Route optimization time for 1,460 flights (hours)	6	6	6	10	30	30	8	6
Quality	MSE	1-1-2012		2,391	2,389	2,388	2,131	3,352	3,353	3,360
		1-7-2012		2,851	2,850	2,873	3,041	3,530	3,524	3,510
		1-1-2013	1,747	1,747	1,747	1,747	1,750	3,703	3,703	3,722
	Hotspot accuracy (%)	1-1-2012		0.208	0.376	0.716	0.750	0.773	0.732	0.387
		1-7-2012		0.206	0.384	0.732	0.707	0.747	0.731	0.385
		1-1-2013	0.196	0.196	0.358	0.694	0.737	0.749	0.696	0.363
Perfor-mance	Total incident coverage	1-1-2012		4,128	4,523	4,455	4,862	4,848	4,705	4,665
		1-7-2012		4,271	4,768	4,634	4,931	4,889	4,926	4,893
		1-1-2013	4,188	4,188	4,548	4,510	4,550	4,603	4,523	4,288
	flying coverage	1-1-2012		580	635	611	669	650	584	591
		1-7-2012		577	615	610	522	510	689	658
		1-1-2013	475	475	525	493	509	522	476	466
	Stand-by coverage	1-1-2012		3,548	3,888	3,844	4,194	4,198	4,120	4,074
		1-7-2012		3,694	4,153	4,024	4,410	4,378	4,238	4,235
		1-1-2013	3,713	3,713	4,022	4,017	4,041	4,081	4,048	3,822

Table 7.6: Results of the forecast method experiments.

Second, we see that the level of spatial aggregation does not the influence the MSE (compare experiments 2, 3, and 4), which is as expected since spatial aggregation multiplies the total forecast value at every time interval by the same value. However, less spatial aggregation does result in a higher forecast accuracy and a better incident coverage from stand-by as well as from surveillance routes. We see that the performance of a forecast with one ring of spatial aggregation (experiment 3) performs better than the forecast without spatial aggregation (experiment 4), and find that spatial aggregation makes it easier for the route optimization MILP to detect hotspots.

Third, from experiment 4 to 5 (with and without temporal aggregation) we see that the accuracy of the forecast improves, and that the incident coverage from flying the optimal routes increases. This indicates that forecasting without temporal aggregation increases the quality of the forecast.

Fourth, when comparing experiments 5 and 6, we find that the use of weekday/month forecasting method, instead of month and weekday factors, increases the hotspot accuracy and performance, but decreases the MSE. Figure 7.6 shows the forecast of experiment 6 and compares it with the actual incident pattern of the data from the validation period. We see that the forecast of experiment 6 overestimates criminal intensity of November and December, while underestimating January and February, thus leading to a higher MSE.

Fifth, experiments 7 and 8 show that the effect of temporal and spatial aggregation is the same for the forecast without weekday and month factors as for the forecast with these factors. We conclude that the forecasting method of experiment 6 provides the best results. However, it is not feasible to perform all following experiments with the forecasting method of experiment 6 due to its extensive computation time. Therefore, we perform the remaining experiments with the forecast of experiment 7.

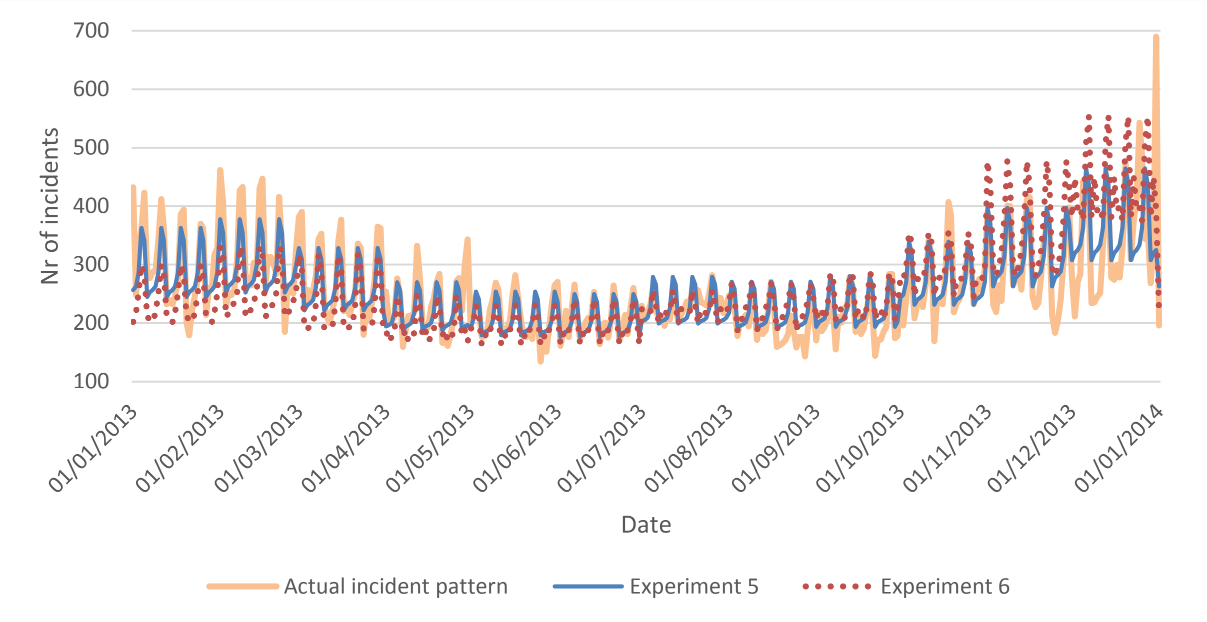

Figure 7.6: comparison of weekly and seasonal patterns in experimental forecasts and the actual pattern.

Overall, Table 7.6 shows quite consistent patterns between experiments for the different start dates. Only the flying incident coverage for start date 1-7-2012 deviates significantly, for which we did not find an explanation.

For the forecast method of experiment 7 we now determine the effect of the forget factor. Table 7.7 shows the resulting forecast accuracy and incident coverage for the experiments. We conclude that a forget factor of 0.03 results in the best performance, although we see no clear pattern or reason. From now on, we use the forecasting method of experiment 7, without spatial aggregation but with temporal aggregation and a forget factor of 0.03, for all following experiments.

	Experiment	9	10	11	12	13
	Forget factor	0.005	0.01	0.03	0.1	0.5
Quality	1-1-2012	0.730	0.730	0.732	0.746	0.883
	1-7-2012	0.732	0.732	0.731	0.721	0.633
	1-1-2013	0.694	0.694	0.696	0.711	0.824
Performance	1-1-2012	4,728	4,736	4,705	4,401	3,473
	1-7-2012	4,925	4,928	4,926	4,673	3,074
	1-1-2013	4,514	4,491	4,523	4,010	3,176

Table 7.7: results for experiments with forget factors.

Conclusions step 1

In this section we find that the fast forecasting method does not result in a decrease of performance. The best forecasting method is the method without weekday and month factors and without temporal and spatial aggregation. Furthermore, we find that the computation time of flights depends on the level of spatial and temporal aggregation. Since computation time is limited, we use the forecast without weekday and month factors but with temporal aggregation for the next experiments. For this forecasting method, the forget factor of 0.03 results in the best quality and performance.

7.3.2 Step 2: performance for the stand-by and planning policies.

We now test the different stand-by and planning policies and determine their effect on the performance of the LVP. We start with experiment 14 to 19, as discussed in Table 7.3.

Stand-by policies

Table 7.8 shows the resulting incident coverage per planning policy for different total coverage distance (see Section 2.3.2) and Figure 7.7 shows the performance per stand-by policy when we compare it with the maximum performance of experiment 18.

Incident coverage	Experiment	14	15	16	17	18	19
	Stand-by Policy	1	2	3	4	5	6
Total coverage distance (minutes)	4	1,500	4,105	5,720	5,711	6,297	7,677
	6	7,962	12,491	12,654	13,371	15,113	20,709
	10	17,042	25,836	25,944	27,865	33,712	47,657

Table 7.8: results for experiment with different stand-by policies.

When we compare experiments 15 and 14 we see that the 400 crews on top of the 1,100 crews on Schiphol increase the performance by at least 50%. Furthermore, we see that keeping the 24/7 stand-by coverage obligation but creating flexibility in the location from where this coverage is given (experiment 16), enables an expected performance improvement of 1% and 40% for a total coverage distance of 6 and 4 minutes respectively.

When the LVP is not obliged to provide 24/7 coverage (experiment 17), the performance increase is minimal. Figure 7.8 shows the forecast we can optimally cover with a number of crews, 100% being the forecast covered by having a crew on every location during every shift. It shows that with 1,500 crews, the LVP can cover approximately 80% of the maximum coverage for a total coverage distance of 4 minutes. We assume that stand-by policy 2 (experiment 15) is the current practice, and conclude that given the current tactical planning decisions, the LVP is able to perform between 50% and 60% of the maximum possible stand-by performance (experiment 19).

From the comparison of experiment 18 with experiment 17 we find that the scheduling of entire shifts (experiment 17) instead of the optimal set of times (experiment 18), costs at most 18% of performance. Figure 7.9 shows the distribution of crews (or stand-by time for policy 5) over the three locations for different policies and different total coverage distances. It shows that for different distances, the importance of, for example, Volkel changes. Furthermore, the ratio between crews on Schiphol and Rotterdam changes between policies and total coverage distances. Rotterdam is the best base when the total coverage distance is 4 minutes, probably since a helicopter at Rotterdam airport can then cover both Rotterdam and The Hague, while Schiphol only covers Amsterdam.

When we do not schedule full stand-by shifts (policy 5 versus policy 4), we see that Volkel becomes more attractive. It is thus interesting to be at Volkel at peak times, although full shifts are commonly less interesting than full shifts at Rotterdam. When the peak times at Volkel differ from the peak times at Rotterdam and Schiphol, combinations could be interesting. Future analysis of these peak times should reveal whether it is beneficial to stand by at multiple bases during one shift, and to use the surveillance flights to travel between the bases.

Overall, we find that the more advanced stand-by policies result in the biggest performance improvement for a total coverage distance of 4 minutes. Furthermore, Figure 7.8 shows that the expected forecast coverage per shift has a lower variation for higher coverage distances. This explains why the benefit of stand-by policy 6 depends on the coverage distance.

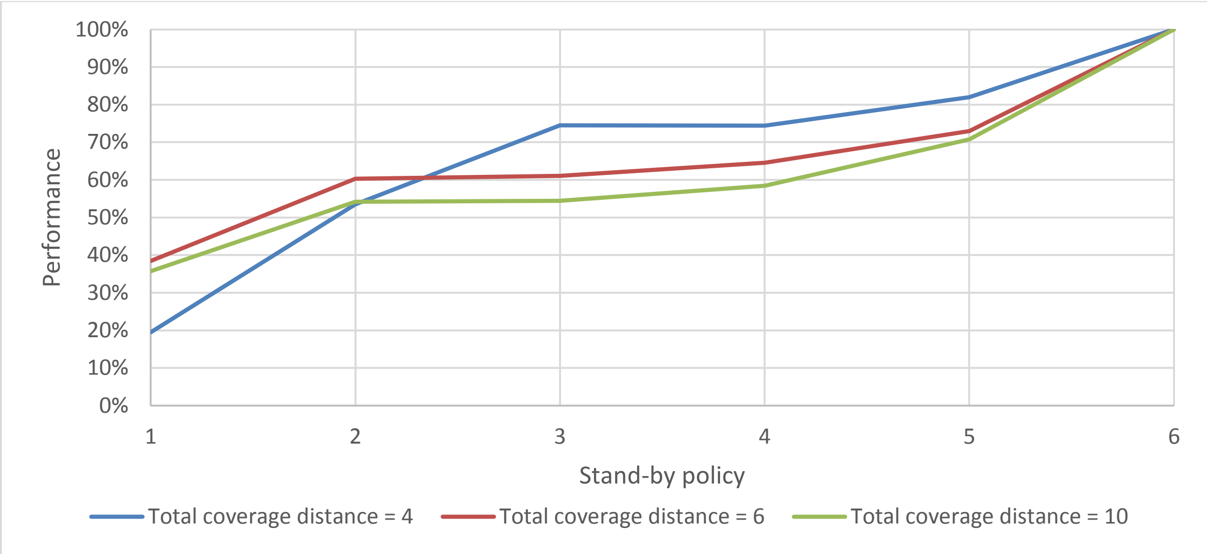

Figure 7.7: stand-by performance relative to the maximal performance, for several total coverage distances.

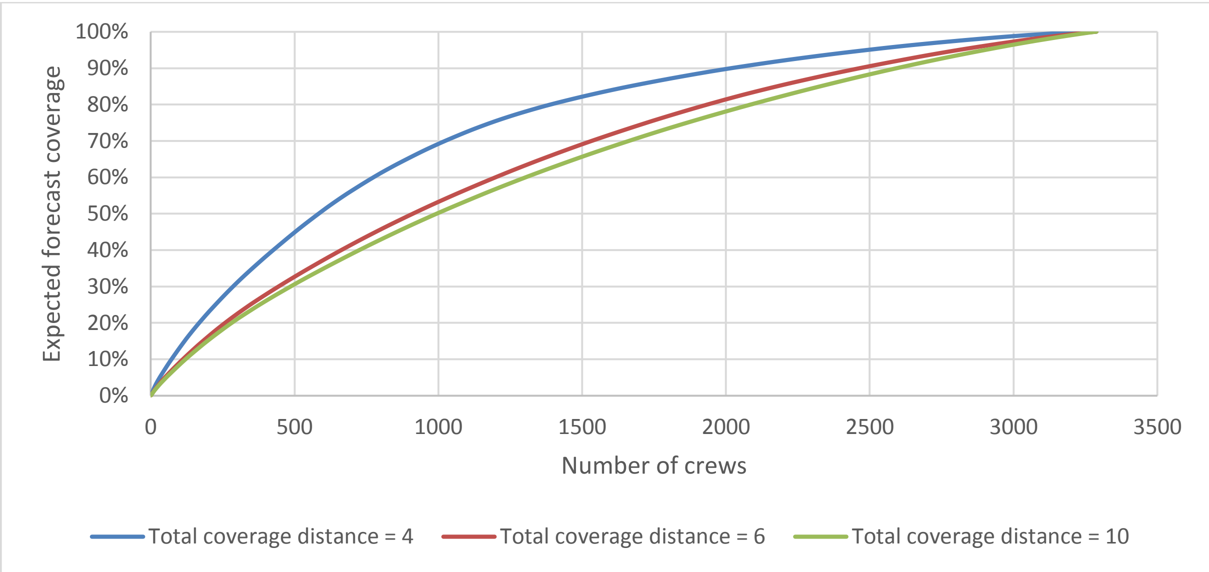

Figure 7.8: Expected forecast coverage per crew budget.

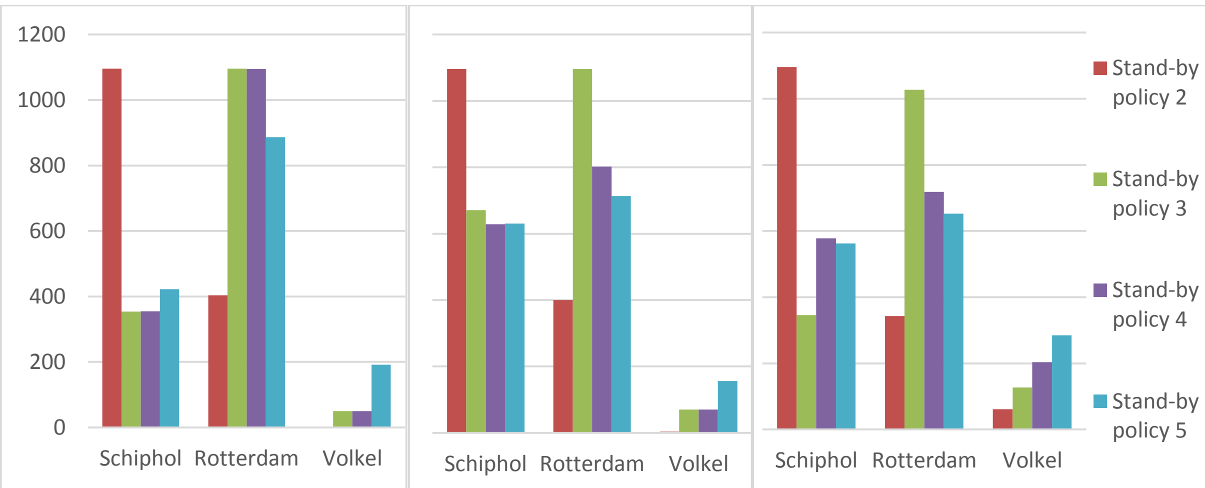

Figure 7.9: The number of crews at the stand-by locations for stand-by policies 2 to 6 and total coverage distance equal to 4 (left), 6 (middle), and 10 (right).

Planning policies

Recall that, since it is most similar to the actual stand-by planning, we use stand-by policy 2 as the default stand-by policy. We now discuss the results from experiments 19 to 24 with different planning policies, as explained in Table 7.4, and stand-by policy 2. Table 7.9 shows the total number of incidents covered, and incident coverage per planning policy, for different start dates (datasets) and total coverage distances. We did not perform experiment 23, since we found that the tactical planning with planning policy 5 is currently not restricted by the number of crews or helicopters.

Total incidents covered	Experiment	15	19	20	21	22	23	24
Total coverage distance (minutes)	Policy	1	2	3	4	5	6	7
4	01/01/2012	7,772	13,851	14,747	15,916	15,296		15,007
4	01/07/2012	8,153	14,125	15,106	16,481	15,067		16,269
4	01/01/2013	7,616	13,126	14,140	15,489	14,710		13,883
6	01/01/2013	19,039	22,841	23,665	26,973	24,300		25,524
10	01/01/2013	34,381	38,018	38,585	40,837	37,272		39,012

Total incident coverage	Experiment	15	19	20	21	22	23	24
Total coverage distance (minutes)	Planning policy	1	2	3	4	5	6	7
4	01/01/2012	4,260	8,995	9,737	10,460	10,461		10,201
4	01/07/2012	4,416	9,139	9,934	10,917	10,508		11,253
4	01/01/2013	4,105	8,406	9,300	10,186	8,249		9,598
6	01/01/2013	12,491	16,293	17,112	19,286	17,215		18,064
10	01/01/2013	25,836	29,911	30,778	33,177	30,065		31,614

Flying incident coverage	Experiment	15	19	20	21	22	23	24
Total coverage distance (minutes)	Planning policy	1	2	3	4	5	6	7
4	01/01/2012	0	5,569	6,393	7,094	8,511		8,291
4	01/07/2012	0	5,596	6,522	7,430	9,069		8,884
4	01/01/2013	0	5,080	6,140	6,959	7,804		7,699
6	01/01/2013	0	6,751	8,218	10,292	11,404		10,853
10	01/01/2013	0	10,396	12,830	15,790	15,972		15,960

Stand-by incident coverage	Experiment	15	19	20	21	22	23	24
Total coverage distance (minutes)	Planning policy	1	2	3	4	5	6	7
4	01/01/2012	4,260	3,426	3,344	3,366	1,950		1,910
4	01/07/2012	4,416	3,543	3,412	3,487	1,439		2,369
4	01/01/2013	4,105	3,326	3,160	3,227	1,120		1,899
6	01/01/2013	12,491	9,542	8,894	8,994	5,811		7,211
10	01/01/2013	25,836	19,515	17,948	17,387	14,093		15,654

Table 7.9: Total incidents covered, total incident coverage, flying incident coverage, and stand-by coverage per planning policy.

When we compare the incident coverage with the number of incidents covered per experiment, as Table 7.9 shows, we find that an increase in the number of covered incidents not necessarily leads to an increase of incident coverage. Since the incident coverage provides more information on the proximity of helicopters to criminal incidents than the number of incidents covered, we use the incident coverage for the rest of the analysis. Furthermore, we find that the trends in differences between planning policies are consistent for different planning periods, as defined by the start date, and for different total coverage distances. We now evaluate step by step the differences between consecutive planning policies.

First, we note that the benefit of flying surveillance flights is substantial for a total coverage distance of 4 minutes. 2,190 hours of standard surveillance flights (experiment 19, flying incident coverage) deliver more performance than standing by 24/7 (± 8,760 hours) on Schiphol (experiment 15, stand-by incident coverage). This effect is less for larger total coverage distances, which indicates that the standard routes perform less for larger total coverage distances, or that standing by profits more from a larger total coverage distance than flying. We note that there is no flying incident coverage for experiment 15 since this experiment includes only stand-by scheduling.

Second, we see that the timing of standard surveillance routes based on the forecast (policy 2, experiment 20) increases the added value of surveillance flights by approximately 20% in comparison to experiment 19 with policy 1. The best start time is based on the flight value estimator. Figure 7.10 shows the flight value per time interval as estimated by the original flight value estimator. We clearly recognize the weekly and monthly patterns as shown in the forecast pattern in Figure 7.6. Figure 7.11 shows that the timing of flights becomes spread out over the day in experiment 20, probably since the best start time changes between weekdays and months due to the different daily crime patterns.

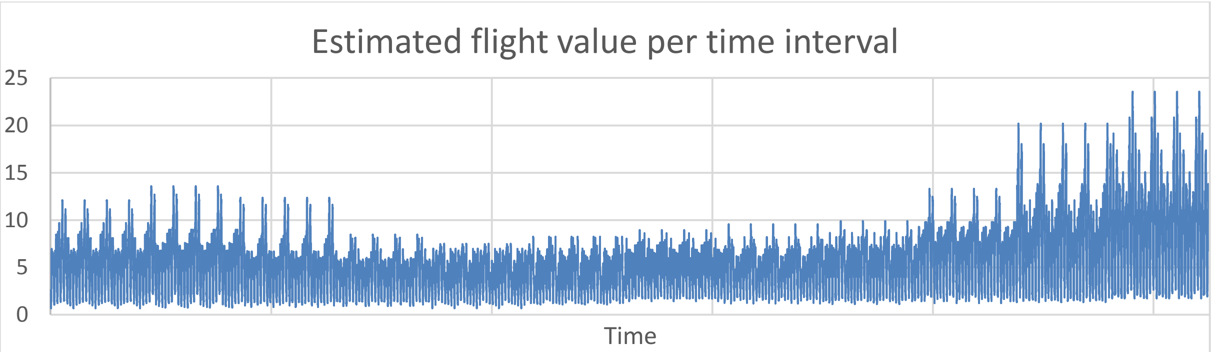

Figure 7.10: estimated added value of a surveillance flight from Schiphol to Schiphol per time interval for one year.

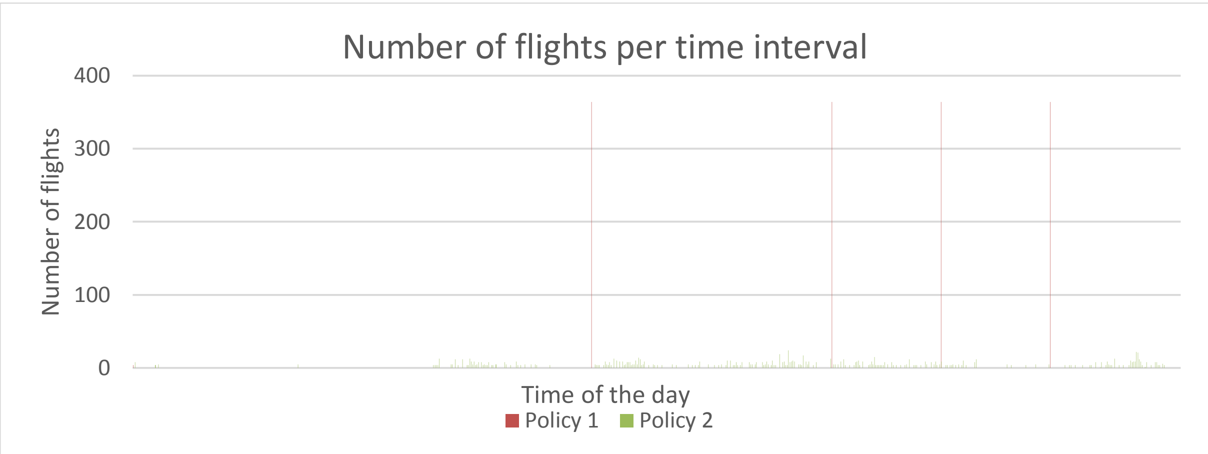

Figure 7.11: the number of flights that starts per time interval.

Third, we find that calculating the optimal route (experiment 21, policy 4) instead of using a standard route (experiment 20) increases performance for a total coverage distance of 4 minutes by approximately 13%. However, for a total coverage distance of 6 minutes this is 27% and for a total coverage distance of 10 minutes it becomes 24%. Figure 7.12 shows the difference in the routes for different policies, given a total coverage distance of 4 minutes. We find that the optimal routes focus on the centre of the “Randstad”, a group of cities in the West of the Netherlands. From within the Randstad, the helicopters can reach multiple large cities in time. However, the standard route passes directly over cities, which makes it cover in general one city at a time.

Figure 7.12: Accumulated routes for policy 1 and 2 (left), policy 4 (middle), and policy 7 (right).

Figure 7.13 shows the accumulated routes for different total coverage distances (in this case for policy 7). We find that for a total coverage distance of 4 minutes, the helicopter routes focus on areas where the helicopter can cover multiple hotspots. For example, these areas lie between Rotterdam and The Hague, and between Amsterdam and Utrecht. With a total coverage distance of 6 minutes helicopters are able to cover more than two cities at the same time and the routes focus on the area where one helicopter can cover Leiden, as well as Amsterdam and Utrecht. Furthermore, the blue shaded areas show that some of the routes start at Rotterdam, but most of the routes start at Schiphol. The total coverage distance of 10 minutes implies that one helicopter can cover the entire “Randstad”. We thus see a clear path from Schiphol to a location where the helicopter covers all major western cities at the same time.

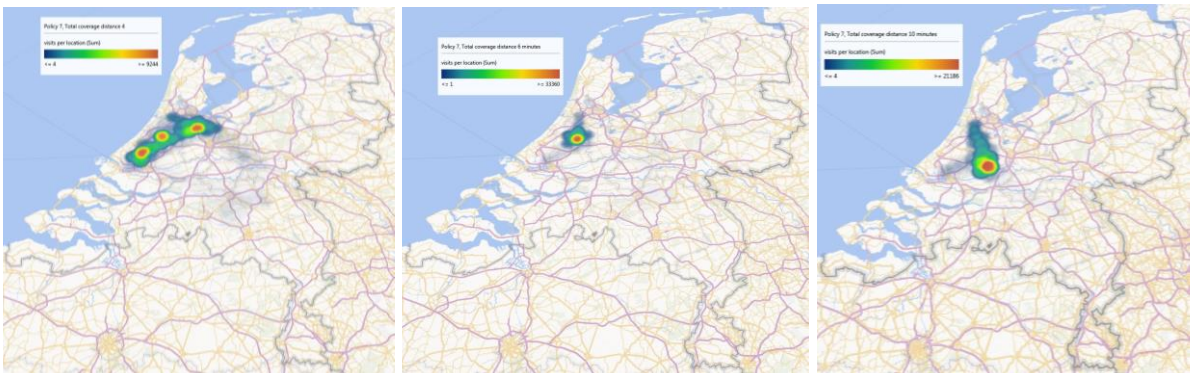

Figure 7.13: Accumulated routes for routes from policy 7 with total coverage distance of 4 minutes (left), 6 minutes (middle), and 10 minutes (right).

The total incident coverage thus has an impact on the positioning of the helicopters. Figure 7.14 shows the coverage per location for a given helicopter position. It shows that a helicopter in the centre of the “Randstad” with a total coverage distance of 4 minutes and the corresponding maximal coverage distance of 9 minutes, is able to (partially) cover multiple cities.

Figure 7.14: overview of the "Randstad" area with the coverage per hexagon (green), given a helicopter location (red) and a total coverage distance of 4 minutes and maximal coverage distance of 9 minutes.

To evaluate the effect of the different total coverage distances we compare the flying incident coverage of the resulting routes, with the average distance to the covered locations and calculate the corresponding score according to Equation 6.11. Table 7.10 shows the results for experiment 21 with policy 4.

Total coverage distance (minutes)	Flying incident coverage	Average arrival time (minutes)	Score
4	6,959	5.6	311
6	10,292	7.5	229
10	15,790	9.3	170

Table 7.10: the flying incident coverage, average distance to covered incidents and the combined score per total coverage distance.

Since one helicopter cannot cover simultaneous incidents in different cities and the LVP aims to increase its proximity, we find that the total coverage distances of 6 and 10 minutes are unrealistic. Furthermore, long total coverage distances result in an undesirable positioning of the helicopters, far away from every city in the "Randstad". We thus recommend the LVP to use a total coverage distance of 4 minutes for the current success function.

Experiments 19 to 21 consider a flat tactical planning. In experiment 22 we make this variable with planning policy 5. Table 7.9 shows that the incident coverage from surveillance flights increases, while the stand-by incident coverage decreases. This is due to the assumption that a second stand-by crew on a base has no added value during times that the first stand-by crew is not making surveillance flights. Variable tactical planning on Schiphol results in shifts with more flights than one crew can handle, or overlapping flights on times when there are multiple big hotspots. Then, multiple crews are available during the same shift, such that the incident coverage from flights increases, while the stand-by coverage decreases. Figure 7.15 shows that there are even shifts with more crews available on Schiphol than there are helicopters available. We find that in this case helicopters are used sequentially by multiple crews (due to the restriction on the number of flight hours per shift for crews).

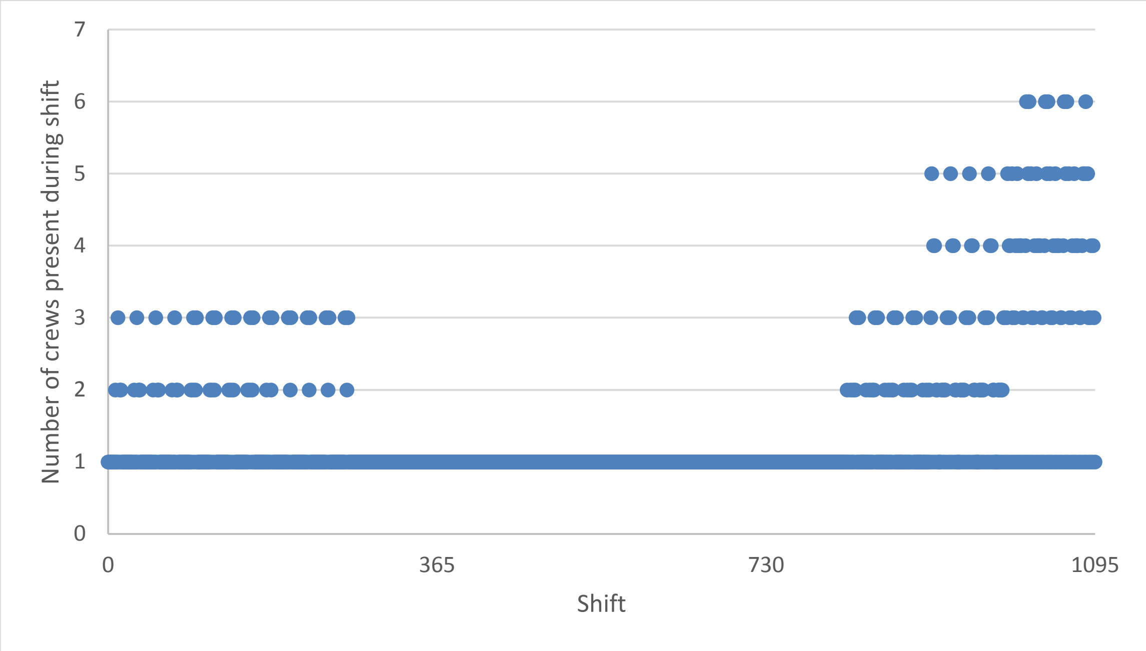

Figure 7.15: number of crews scheduled per shift in the variable tactical planning for 2013.

From Table 7.9 we thus conclude that variable tactical planning (experiment 22) results in an increase in the incident coverage during surveillance flights. However, due to the simultaneous availability of multiple crews on Schiphol, it leads to a decrease in the total incident coverage of approximately 4% for a total coverage distance of 4 minutes, 11% for 6 minutes, and 9% for 10 minutes. We now examine the effect of variable tactical planning on the number of flights per day and evaluate the quality of the estimator. For this analysis we use Experiment 22 with start date 1-1-2013 and a total coverage distance of 4 minutes. Figure 7.16 shows the number of flights per day for policy 4 (with the same start date and total coverage distance) and 5. We conclude that for variable tactical planning, the number of flights decreases significantly in the summer, and that these flights instead are made in the autumn and winter. Furthermore, we clearly see weekly patterns.

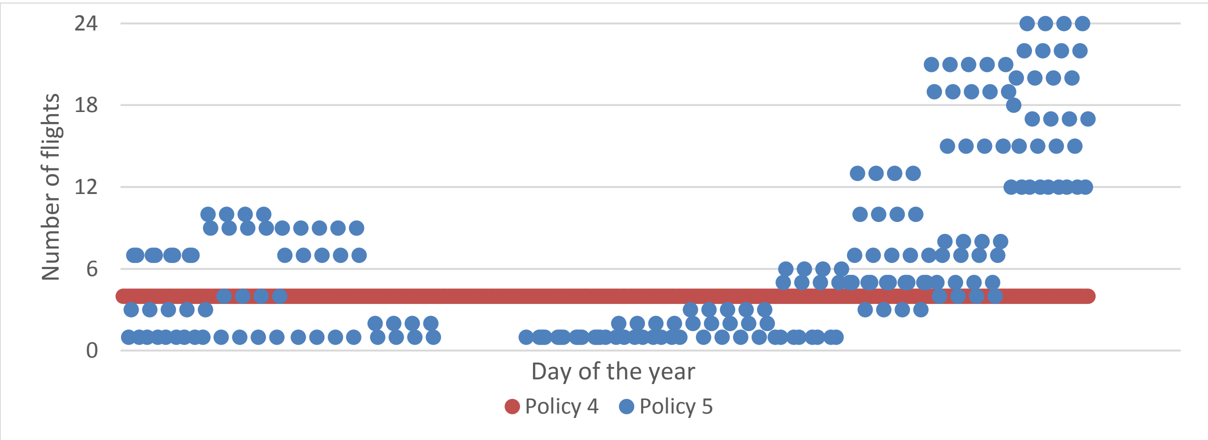

Figure 7.16: The number of flights per day for policy 4 (flat tactical planning) and 5 (variable tactical planning).

To determine the distribution of flights over the year and the timing during the day, we use the results of the flight value estimator. Figure 7.17 shows the expected flight value (forecast coverage), as determined by the estimator. Based on the left side of the graph we conclude that the estimator seems able to find times on which flights cover more than average. On the right side of the graph we conclude that the added value of timing flights decreases as the flight budget increases. After about half of the flights, the peak times have been covered by flights, and the difference between flights decreases. This is an explanation for the limited added value of variable tactical planning with the current LVP budget: when the peak times are covered, there is a limited benefit of comparing for example a flight on a winter morning, with a flight on a summer evening. We thus expect that the added value of variable tactical planning decreases when the flight budget increases.

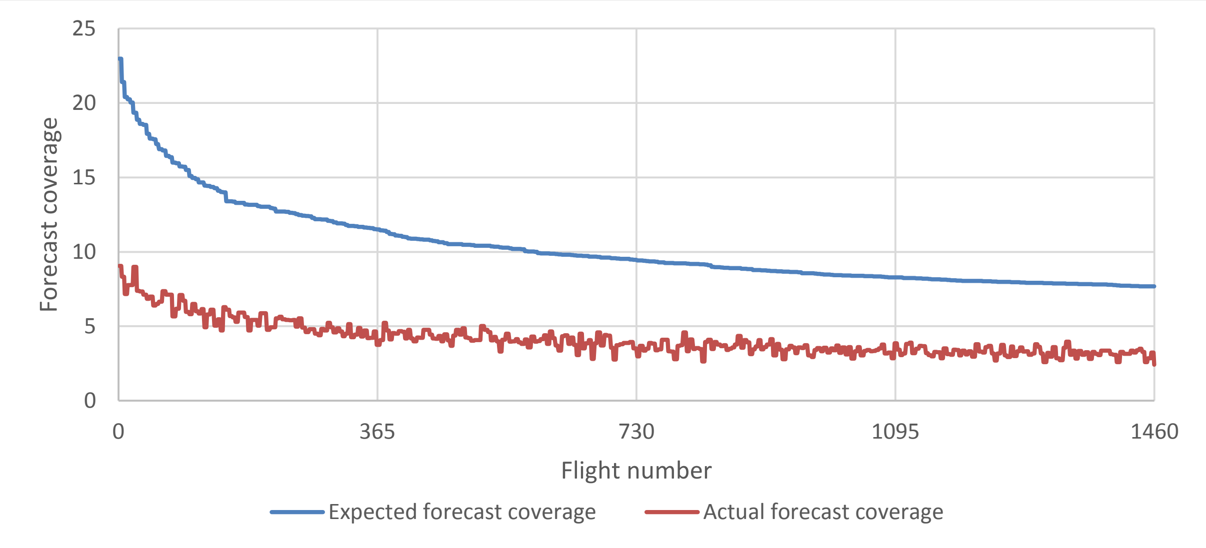

Figure 7.17: Expected and actual forecast coverage for the policy 5 flights.

Next to the quality of the estimator, we also want to determine the predictive value of the underlying forecast. Figure 7.18 shows per flight the actual forecast coverage and the incident coverage. We see no clear relation between both the forecast and incident coverage. With Equation 6.6 we find a correlation coefficient of 0.125. With a sample size of 1460 flights, we find that there is with 99.77% probability a correlation between the actual forecast coverage and the incident coverage.

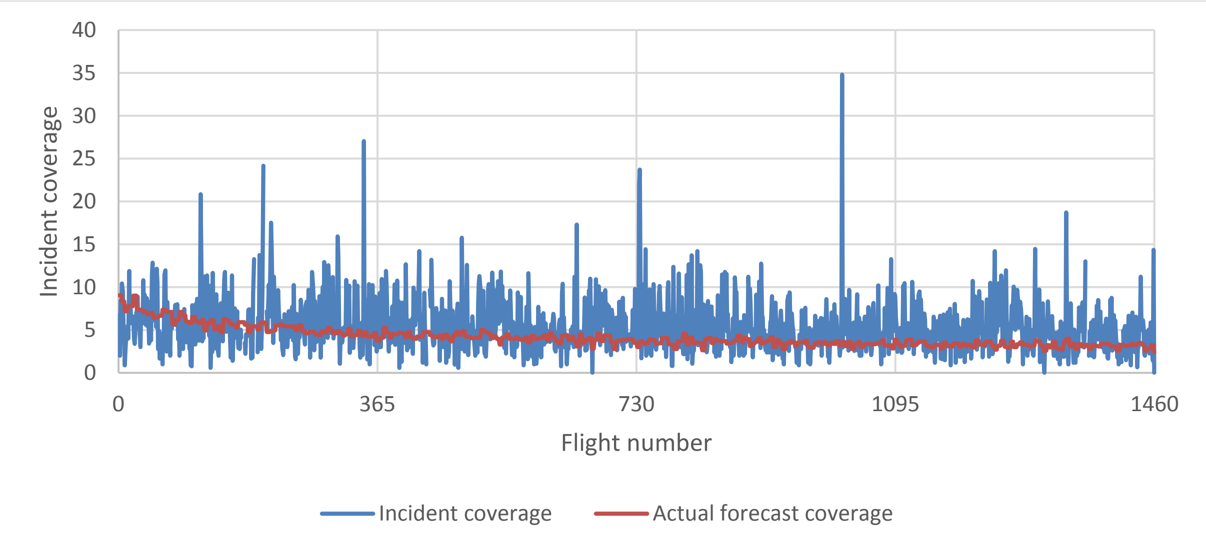

Figure 7.18: The actual forecast coverage and incident coverage for the policy 5 flights with original estimator.

In experiment 23 we aimed to introduce crew and helicopter constraints. However, we found that these factors do not have an impact on the performance, since the variable flight planning does not require more than 1,500 crews and 5 helicopters at the same time.

Finally, from experiment 24 we find that the use of multiple start locations does not increase performance. We note that the heuristic currently determines the optimal start location, based on the coverage of the flight only. We see that this works since the incident coverage from flights increases between experiments 23 and 24. However, when a crew and helicopter fly from Volkel, they are standing-by at Volkel for the entire flight, while otherwise they would have been standing-by at Rotterdam. Since standing by at Rotterdam is much more valuable (in the case of a total coverage distance of 4 minutes) than standing by at Volkel, flights from Volkel cost overall performance.

We performed two unscheduled extra experiments (22.A and 24.A) to test whether we can decrease this effect by first scheduling all stand-by crews on shifts and locations, according to stand-by policy 2, after which we schedule the flights with planning policy 6 and 7 (which do not use anymore crews than already scheduled). We used the forecast that starts at 1-1-2013 and a total coverage distance of 4 minutes. Table 7.11 shows the results of the extra experiments and the differences between the coverage values of flat tactical planning in experiment 21 and variable tactical planning in experiment 22.A, and using multiple start locations in experiment 24.A.

Experiment	21	22.A	Difference	24.A	Difference
Planning policy	4	6		7	
Total incident coverage	10,186	10,605	+4.11%	10,547	+3.54%
Flying incident coverage	6,959	7,349	+5.60%	7,993	+14.86%
Stand-by incident coverage	3,227	3,256	+0.90%	2,554	-20.86%

Table 7.11: effect on incident coverage of scheduling all crews before scheduling flights.

We conclude that the increase of performance by scheduling crews before scheduling flights is substantial and that the total performance of these experiments is at least 3.5% higher than the best possible performance of the flat tactical planning heuristic. We expect that the use of stand-by policy 4 will further increase both the performance from flying and standing by. The simultaneous scheduling of stand-by shifts and surveillance flights thus deserves attention.

When we compare the number of flights per day for experiments 21 and 22.A in Figure 7.19, we conclude that having at most one crew available on a location reduces the effect of variable planning. The maximum number of flights decreases from 24 to 8 and more flights are made in the summer.

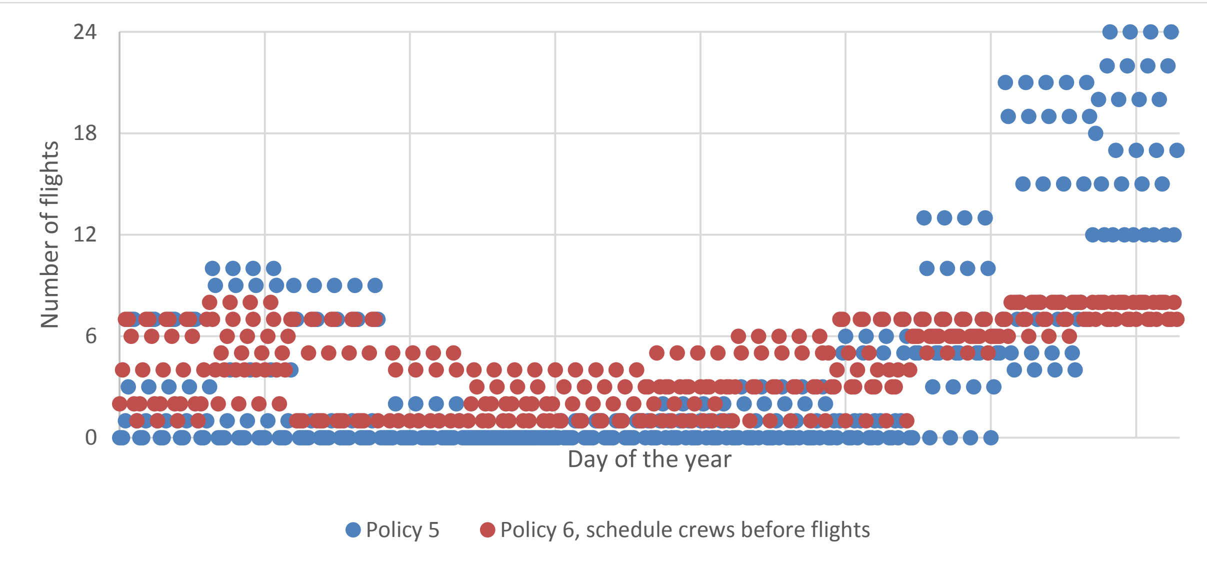

Figure 7.19: number of flights per day for policy 5 and policy 6 in which we schedule crews before flights.

New flight value estimator

In this section we compare the performance of the original flight value estimator with the new greedy flight value estimator. Due to a lack of computation time, Table 7.12 shows the flying incident coverage and correlation coefficient between the expected forecast value and the actual forecast value, for two experiments with a different forecast and total coverage distance. Figure 7.20 shows per flight the flying incident coverage and the forecast coverage, which we used to calculate the correlation coefficient with.

Estimator	Original	New
Total incident coverage	18,502	19,762 (+6.8%)
Correlation coefficient	0.134	0.54

Table 7.12: effect of the new estimator on the total incident coverage and correlation coefficient between the expected flight coverage and the actual incident coverage.

We conclude that the new estimator seems to result in an improvement of the total incident coverage as well as the correlation coefficient. Furthermore, the new method estimates the actual forecast coverage well. We thus recommend to use the new flight value estimator.

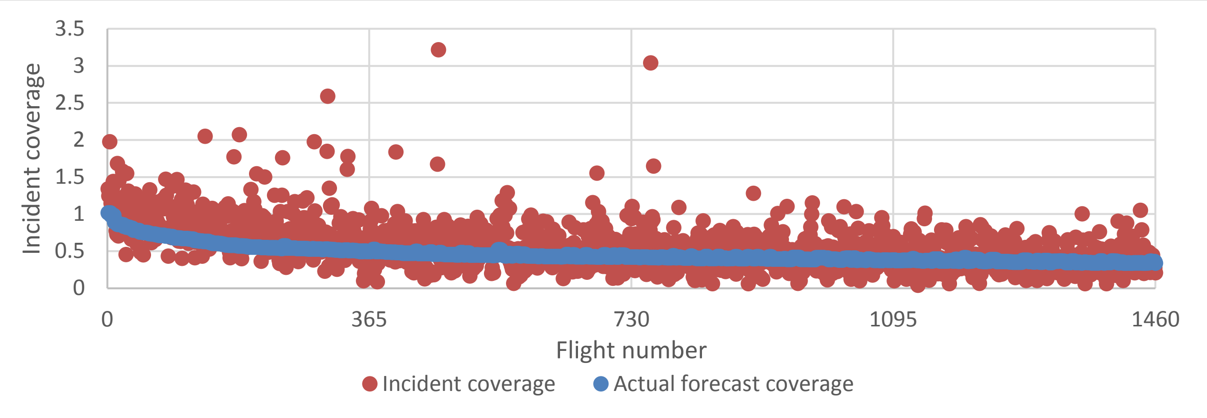

Figure 7.20: The actual forecast coverage and incident coverage for the policy 5 flights with new estimator.

Conclusions step 2

Overall, we find in step 2 that the added value of the previous research by Van Urk (2012) on the timing and route optimization of surveillance flights results in an expected performance improvement between 20%. Furthermore, we find that the total coverage distance is an important influencing factor on the surveillance routes and proximity of the helicopters to criminal incidents. We recommend to use a total coverage distance of 4 minutes from now on. With the tactical planning heuristics that we propose in Chapter 5 we find that variable tactical planning decreases the performance of the LVP. However, when we allocate all crews before we schedule flights, variable tactical planning of surveillance flights improves the performance by approximately 4% and the use of multiple locations increases performance by another %. We thus find that the simultaneous scheduling stand-by crews and flights deserves extra attention. Finally, the use of the new estimator improves the performance of the tactical planning heuristics substantially and is thus recommended.

7.3.3 Step 3: effect of extensions

In step 3 we examine the effects of the proposed extensions. First, we test the new estimator. We then experiment with different flight durations, a new success function, taking into account priorities and effectiveness, limits to the maximum number of visits to a location during a route, and the effect of shifts on overtime. For every experiment we use the forecast method from step 1, the corresponding forecast from 1-1-2013 to 31-12-2013, and a total coverage distance of 4 minutes.

Fixed flight duration

In this section we evaluate the impact of different flight duration settings on the total incident coverage. We compare the performance of 1,460 flights of 1.5 hours (from Experiment 24) with 2,190 flights of 1 hour and 1,095 flights of 2 hour (the maximum flight duration for the EC135). We compare both the total and flying incident coverage per experiment.

Experiment	24	26	27
Flight duration	01:30	01:00	02:00
Total incident coverage	9,598	9,583	9,592
Flying incident coverage	7,699	7,689	7,676
Stand-by incident coverage	1,899	1,894	1,916
Number of flying crews	674	759	645

Table 7.13: total incident coverage for different flight duration.

We conclude that a flight duration of 1:30 hours provides the best coverage. However, we also note that given this flight duration, not all locations in the Netherlands can be reached by the EC135. For 2-hour flights, the number of crews that makes at least one flight decreases, and the stand-by coverage increases, since less flights are made from Rotterdam and crews fly at least two hours.

Alternative success function

In this section we discuss experiments 28. In this experiment we replace the original success function by the negative exponential functions with the preventive effect (as discussed in Section 4.4.2). Figure 7.21 shows that this success function results in routes that focus on locations just outside major cities. On these locations the balance between the loss of coverage due to the preventive effect and the coverage of the hotspots is best.

Figure 7.21: Accumulated routes for the new success function.

Priorities and effectiveness

In this section we discuss the effect of incident priorities and the effectiveness of helicopters during day and night. Table 7.14 shows the results per experiment. The calculation of the performance measures that we use is discussed in Section 6.3.1.

Experiment	24	29	30	31
Priorities	No	Yes	No	Yes
Effectiveness	No	No	Yes	Yes
Flying Incident coverage	7,699	6,622	6,981	
Priority coverage	13,699	14,125 (+3.1%)	14,749	
Effectiveness coverage	484	462	536 (+10.7%)	

Table 7.14: the effect of priorities and effectiveness on the performance measures.

We find that the overall incident coverage decreases when we take priorities and effectiveness into account, while the priority and effectiveness coverage increase in comparison with experiment 24. Furthermore, we note that taking effectiveness into account also increases the priority coverage. Figure 7.22 shows the distribution of start times of surveillance flights over the day for policy 7 with and without effectiveness taken into account. It clearly shows that flights on average start later on the day, or just before sunrise, when we take the helicopter effectiveness into account. The LVP could use Figure 7.22 to determine a shift schedule that minimizes the number of flights that start in the last hour of the shift, to decrease overtime.

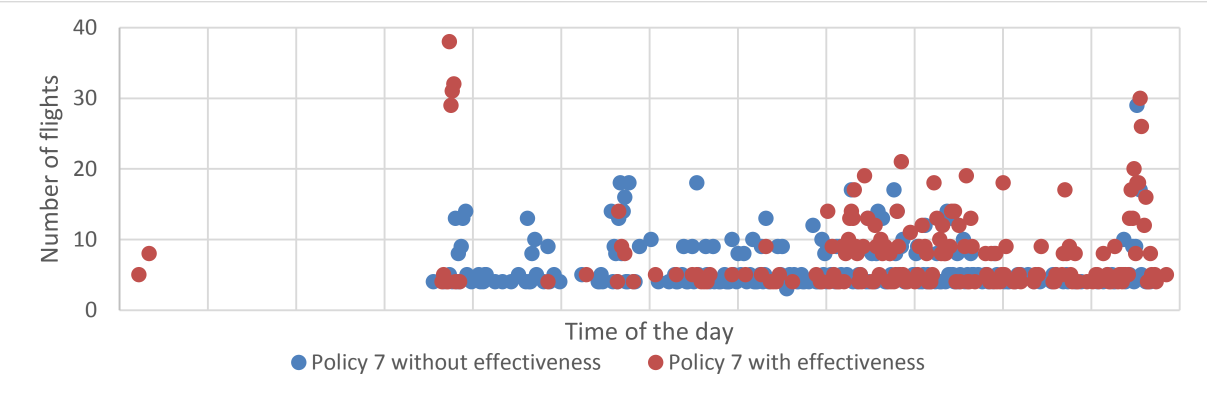

Figure 7.22: effect of taking effectiveness into account on the timing of surveillance flights.

Figure 7.23 shows the distribution of flights over days in the year. We find that taking effectiveness into account does not lead to more flights in the winter (in which there are more dark hours). On the contrary, it leads to more flights during summer nights and less flights during daylight in the winter. This is consistent with our conclusion in Section 5.2.1, based on the rough-cut tactical planning.

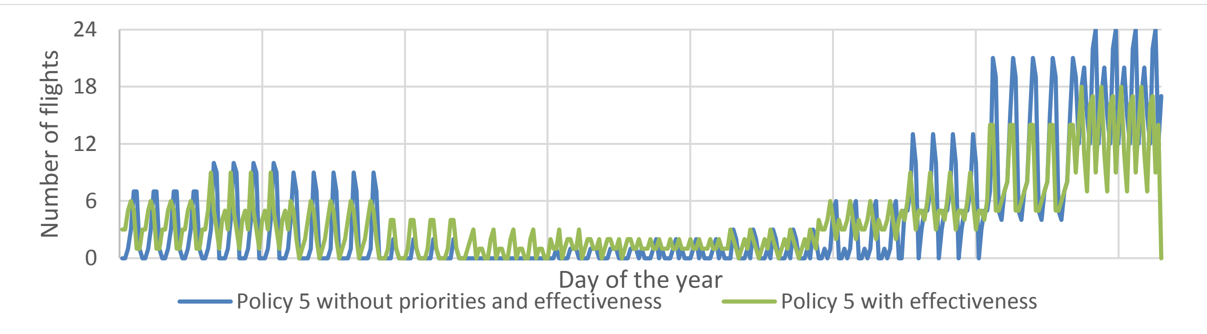

Figure 7.23: number of flights per day for policy 5 with and without priorities and effectiveness.

Overtime

This section describes the effect of different shift schedules on the percentage of work time in overtime. Table 7.15 shows the percentage in overtime for the current shift schedule, and two alternative shift schedules that start an hour earlier or later. We conclude that different shift schedules can decrease the percentage of work in overtime.

Experiment	32	24	33
start early shift	06:00	07:00	08:00
start late shift	14:00	15:00	16:00
start night shift	22:00	23:00	00:00
Percentage in overtime	7.1%	9.0%	12%

Table 7.15: effect of different shift schedules on percentage work time in overtime.

Maximum number of visits

In this section we examine the effect of constraining routes with a maximum number of visits per location during one flight. We determine the resulting flying incident coverage with the original success function and a total coverage distance of 4 minutes.

Experiment	24	34	35
Maximum number of visits	100	10	2
Flying incident coverage	7,699	7,519 (-2.3%)	7,454 (-3%)

Table 7.16: Flying incident coverage for different values of the maximum number of visits.

We find that the performance of the flights decreases significantly when we do not allow helicopters to hover above a location as long as is required for optimality. However, the resulting routes are also not attractive for helicopter pilots. Figure 7.24 shows exemplary flight routes for a maximum number of visits of 2, for the original success function (left) and the new success function (right).

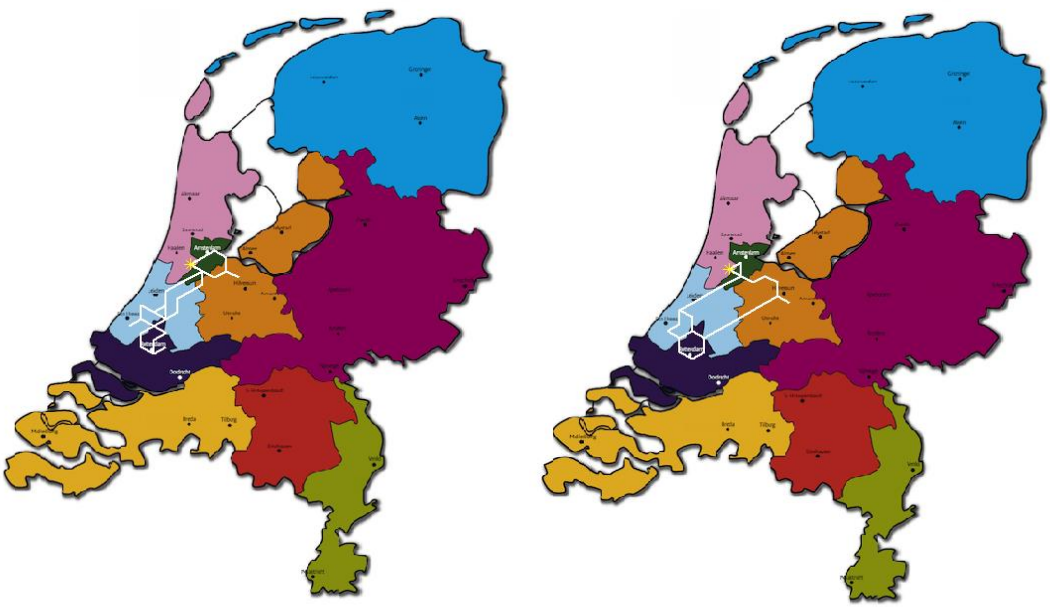

Figure 7.24: Exemplary routes for a maximum number of visits per location of two.

We conclude that the maximum number of visits constraints creates acceptable routes when we use it in conjunction with the new success function. With the original success function, it commonly results in an unclear mesh above the area between the major cities. We conclude that creating attractive surveillance routes requires more attention in future research.

Fairness

To investigate the trade-off between equity and efficiency, we perform five experiments with different values for the equity parameter q (as discussed in Section 5.5), stand-by policy 2, and the tactical planning heuristic that takes the equity parameter into account. We use the original success function with a total coverage distance of 4 minutes. Due to a lack of computation time, we create only 100 flights with a duration of 2 hours per flight, which may start on Schiphol, Rotterdam, and Volkel. As we currently assume that helicopters cannot refuel during their flight, a flight duration of two hours is required to be able to cover the North-East of the Netherlands. The goal of these experiments is to determine the trade-off between minimizing inequity and maximizing efficiency. Figure 7.25 demonstrates the trade-off effects of equity versus efficiency.

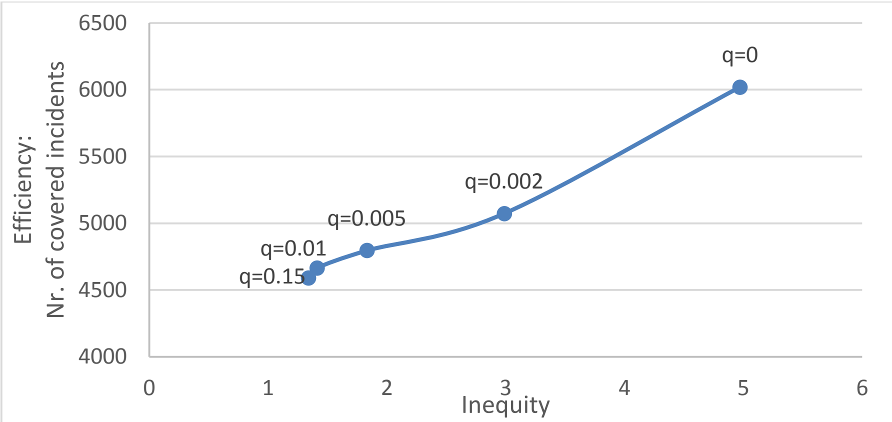

Figure 7.25: the equity/efficiency trade-off.

From Figure 7.25 we conclude that a larger value for the equity parameter leads to solutions that are less unequal. Furthermore, we see that an increase in equity requires a considerable decrease in efficiency. A solution that is half as unequal, is approximately 17% less efficient than the most efficient solution. To reduce the inequity by 80%, the incident coverage would reduce by approximately 23%. Table 7.17 shows that the decrease in efficiency is the case for both the set of routes and the timing of stand-by crews. We suspect that this is due to differences in timing of criminal hotspots in the big cities and the more rural areas of the Netherlands. Flights to the rural areas are thus scheduled at times that it is relatively quiet in the big cities, where the bases lie on which the crews are scheduled for the rest of their shift.

Experiment	36	37	38	39	40
Equity factor	0	0.002	0.005	0.01	0.015
Incident coverage from flight	1,903	1,036	871	760	695
Incident coverage from stand-by	4,116	4,035	3,926	3,909	3,894
Total incident coverage	6,019	5,071	4,794	4,669	4,589

Table 7.17: breakdown of total number of incidents covered into covered incidents from stand-by and from flight.

Figure 7.26 shows that the decrease in expected incidents covered by the surveillance flights corresponds with the decrease in the forecast covered by the surveillance flights as in Table 7.17. This indicates that there is a relation between the level of forecast covered and the actual incident coverage, and thus that the forecast is related to reality. Furthermore, we conclude that while the inequity and thus the deviations from the average percentage forecast covered decreases, the average forecast covered also decreases. A less unequal distribution thus leads to a decrease in average performance.

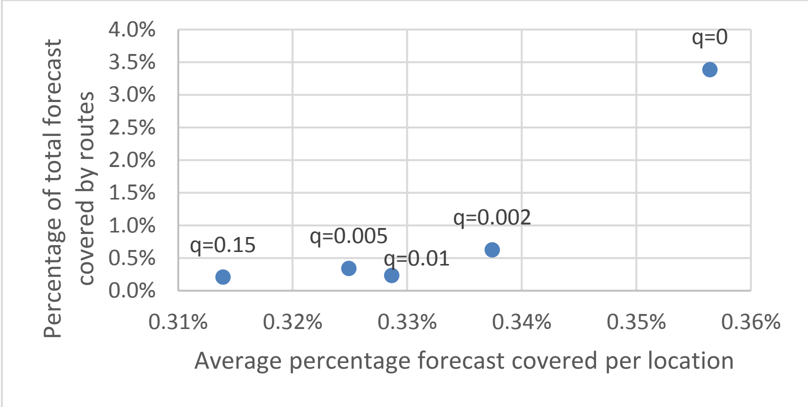

Figure 7.26: Comparison of average and total percentage of forecast covered for different values of the equity parameter.

We also check whether the solutions comply to the constraint that every location should be covered at least once every year. Figure 7.27 shows the percentage of forecast that is covered and the number of visits of helicopters during surveillance per location, for a set of routes created with q set to zero. We conclude that the constraint does not hold for the experiment with zero equity. Furthermore, we find that almost every flight flies directly to Alphen a/d Rijn, a city in the middle between Amsterdam, Rotterdam, and Utrecht. As described before, from here a helicopter covers all three big cities.

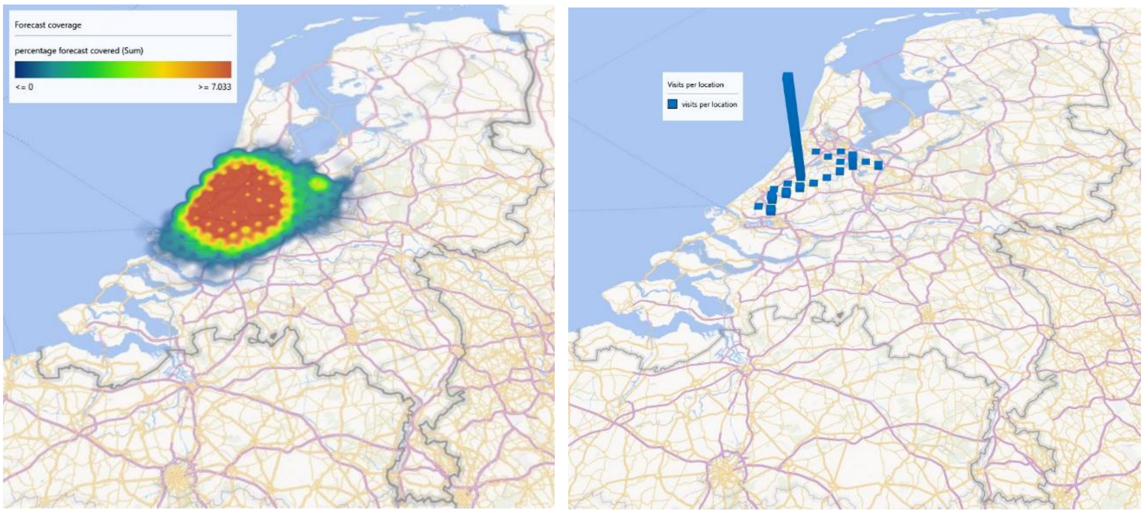

Figure 7.27: forecast coverage (left) and number of visits (right) per location, for equity parameter = 0.

Figure 7.28 shows the same effects, but then for an experiment with q set to 0.01. We compare Figure 7.27 with Figure 7.28 and see that the forecast coverage is more spread out in the experiment with a higher value for the equity parameter. When we exclude the north sea islands, which receive a minimal amount of coverage, we find that the coverage $coverage_w$ of the least covered location is 7, which means that the location is covered by one or more flights during at least 7 time intervals. Furthermore, we found that with the equity parameter set to 0.01, the minimum number of flights to cover all locations is 87. Figure 7.28 also clearly shows the effect of standing-by on the coverage of locations. The three areas with the darkest shade of red are all located around the helicopter bases at Schiphol, Rotterdam, and Volkel.

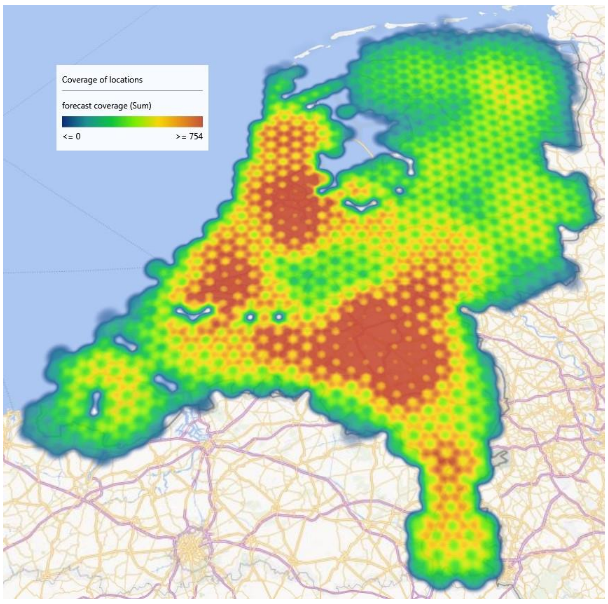

Figure 7.28: coverage of locations, equity parameter = 0.01.

Figure 7.29 shows per location the number of visits by a helicopter during the set of surveillance routes, through a heat map. Red hot areas are often visited. We conclude that Volkel is of strategic importance to cover the Eastern areas of the Netherlands. Since the LVP helicopters should not fly over water for safety reasons, helicopters from Schiphol that fly to the north eastern parts of the Netherlands have to cross the Ijsselmeer over the Afsluitdijk or fly around it on the southern side, which takes a lot of time.

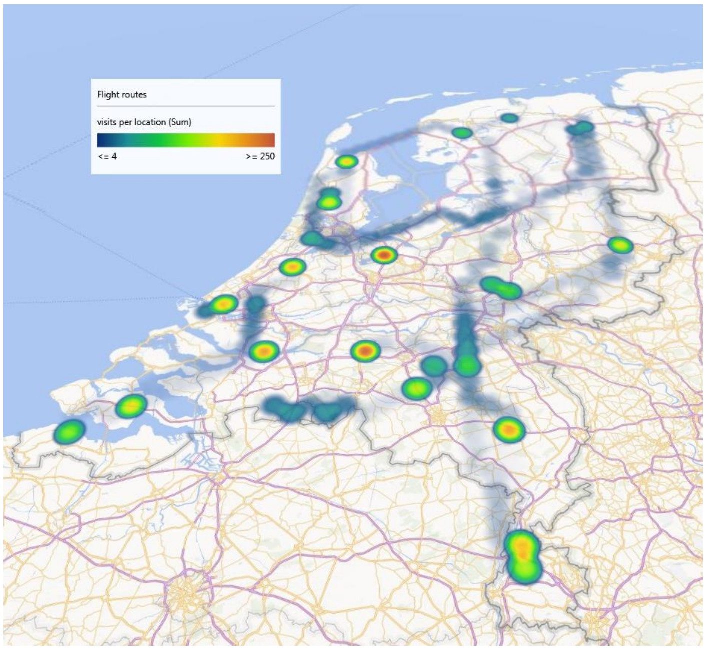

Figure 7.29: Location visits, equity parameter = 0.01

To determine the equity of the solution, we should relate the coverage of a location by the criminal intensity of the location. Therefore, we divide the covered forecast per location by the total forecast value per location as in Equation 6.14, to find the percentage of forecast covered by the helicopter capacity. Figure 7.30 shows the percentage of forecast covered per location. We see that the distribution of coverage is more equally distributed than in Figure 7.27.

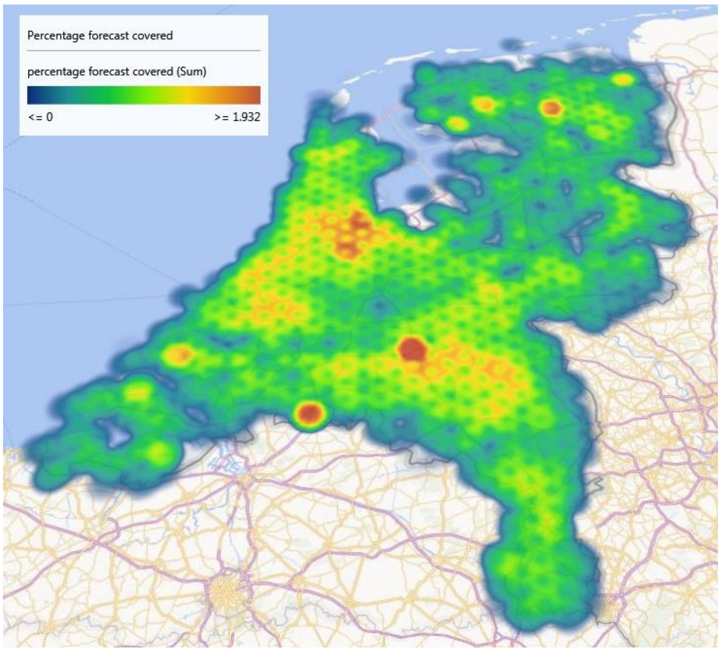

Figure 7.30: percentage of total forecast covered by solution, equity parameter = 0.01

We conclude that the use of the equity parameter q can create a more equal distribution of police helicopter capacity. The value of q determines the level of equity introduced and is thus useful for decision makers to investigate the impact of several levels of equity, and find the appropriate level.

Conclusions step 3

With the experiments in step 3 we show that 1.5 hour flights produce better results than 1 or 2 hour flights. Furthermore, we show the effect of an alternative success function. We introduce priorities into the forecast and adjust the forecast to take helicopter effectiveness into account and find that both adjustments have the desired effect. We quantify the effect of limiting the maximum number of visits to a location during one surveillance flights, determine the impact of different shift schedules on the expected overtime by helicopter crews, and find that the heuristic we propose to take fairness into account works as desired.

7.4 Conclusion and discussion of simulations

This section discusses the results of the simulation model that we use to determine the effect of different forecast methods and planning policies. Section 7.4.1 draws conclusions based on the results of Section 7.3. Section 7.4.2 discusses the limitations of the simulation model and the assumptions that were required to create the model.

7.4.1 Conclusion

This chapter discusses the goals we set and the simulation plan we follow to achieve them. We combine forecast methods, flight value estimators, and planning heuristics to determine the quality of the current way of working of the LVP and to determine in which ways the LVP can improve its performance. We perform the set of experiments in three steps for which we now give a short overview of the results.

In Step 1 we find that the forecast method without weekday and month factors results in the best performance. Furthermore, we find that the computation time of routes depends on the temporal and spatial aggregation steps in the forecast on which the routes are based. Due to a lack of computation time we use the second-best forecast for further experiments, for which a forget factor of 0.03 provides the best results.

In step 2 we simulate multiple stand-by policies and find that the current stand-by scheduling method of the LVP results in at most 60% of the possible performance (when there is always a crew stand-by on Schiphol, Rotterdam, and Volkel). When we compare the stand-by coverage between bases, we find that Rotterdam is on average a better stand-by location than Schiphol, and that Volkel is the worst base to provide stand-by from. Based on a total coverage distance of 4 minutes, making Rotterdam the prime stand-by location increases the performance by 40%. We thus recommend to provide 24/7 stand-by coverage while making the stand-by location variable between shifts.

We then compare the coverage from standing by with the coverage by surveillance flights. We find that currently one hour of flying surveillance flights creates on average as much incident coverage as 4 hours of standing by. We expect that timing and routing surveillance flights increases the performance by approximately 20% and find that variable tactical planning increases the performance by at least another 3.5%. Finally, we expect that the use of the new estimator will result in a performance increase of approximately 7%.

In step 3 we determine the effect of several adjustments. We find that a fixed flight duration of 1:30 hours results in the best performance, and that a negative exponential function with a preventive effect locates helicopters just outside the major Dutch cities. Taking incident priorities into account, the total incident coverage decreases but the total priority coverage increases by 3%. The same happens for effectiveness, where the total effectiveness coverage increases by more than 10%. Furthermore, we show the effect of different shift schedules on overtime and of the maximum number of visits constraint on performance and routes. Finally, we determine the trade-off between equity and performance and find that it costs about 20% of performance to reduce the inequity of the tactical plan by 50%.

7.4.2 Discussion

This section evaluates the proposed forecast methods, tactical planning policies, and simulation methods. We identify the following options for improvement:

- The effect of variable tactical planning strongly depends on the heuristic that we use. We experiment with two heuristics; in the first, crews are scheduled before we iteratively schedule flights, and in the second we schedule crews after we scheduled flights. We find that this sequential approach shows a negative or small positive added benefit for variable tactical planning. We thus consider the benefit of variable tactical planning that we find as a lower bound for the possible added value.
- The simulation results as discussed in this chapter are based on Dutch crime patterns. The simulation results are thus not representative for other geographical or temporal patterns of crime. We conclude that the added value of route optimization and variable tactical planning differ per country and region.
- We find that the total coverage distance determines where and when helicopters should be stand-by or fly surveillance routes. However, there is no source for the total coverage distance other than expert opinion.
- Based on comparison of the results of flat and variable tactical planning we find that the added value of variable tactical planning depends on the flight budget. With the current budget of 2,190 hours, the LVP is able to have multiple helicopters airborne during every major peak in crime. Extra flights would cover lower peaks and level out the variability in the number of flights per day. Based on this observation, we expect that a smaller budget on average results in a higher added value of variable tactical planning.

- The routing model does not take coverage of stand-by helicopters into account while routing. We first schedule the optimal flights and afterwards determine the moments on which a helicopter and a crew is still available for stand-by. Therefore, a crew can turn out to be stand-by on Rotterdam, while a part of the surveillance route of a crew from Schiphol goes over Rotterdam.
- We currently use the “incident time” as provided by the LVP to create temporal patterns. This data contains information about victim behaviour rather than about criminal behaviour, since it shows when criminals are spotted or when the act is discovered afterwards. Ratcliffe (2002) proposes to use the start and end time of criminal events, and spread incidents evenly over the time between these times to determine temporal patterns. The start time is then the last moment in time that the incident had not happened for sure, and the end time is the time at which it was over for sure. To us this seems more accurate than the current single time for incidents. A further benefit of using the start and end times is that we do not have to correct the input data and that we do not have to perform temporal aggregation.
- The simulation results currently do not consider cancellations and reschedules due to weather influences. We assume that all helicopter routes are actually flown. We thus overestimate the actual coverage of incidents in all scenarios.
- The planning policies do not take into account that a set of helicopter routes might lead to noise problems in case a location is visited too often. However, we expect that an increase in fairness also results in a decrease of noise problems.

8 Implementation

This chapter discusses the implications of the results on the tactical and operational planning procedures of the LVP. Section 8.1 discusses how the LVP can implement the current functionality into the tactical planning of police helicopters. Section 8.2 shows how we communicate the optimal surveillance routes to police pilots. Section 8.3 discusses how we implemented new performance measures. Finally, Section 8.4 draws conclusions on the implementation of tactical planning at the LVP.

8.1 Prototype tool

To perform the simulations of Chapter 7, we extended the current operational scheduling and routing tool by Van Urk (2012) to a complete simulation tool. Appendix B describes the functions of the simulation tool. As discussed in Chapter 1, one of the goals of this research is to create a prototype decision support tool to support tactical planning in the LVP organization. The current simulation tool is just that: it is a prototype that we use to test the effect of tactical planning and several (stand-by) planning policies. Based on the results of the simulations, the LVP can now determine which bases to use and which type of tactical planning to use.

The simulation tool is very useful for further analysis. The tool only requires criminal incident data and the organizational settings like the number of helicopters, shift hours, crew budget etc. All discussed algorithms, policies, and effects are implemented and the user can set the required combination of methods in one table. The tool then performs the entire simulation and with one click the user can export the resulting routes. The resulting tactical planning is shown on a separate page, but the tool is not yet able to synchronize the planning with the planning software as used by the planning office. However, the tool is not yet user-friendly or well documented enough for implementation. Furthermore, the underlying code is too complex for LVP employees without knowledge of programming. Therefore, we identify three options for the LVP to implement the current prototype tool into the organization:

1. The LVP can use the current version of the tool to test different scenarios and create more insight in the effect of tactical and strategic decisions. However, this requires the use of the tool by an experienced analyst with programming and IT skills.
2. The LVP can determine the forecast method, estimator, and tactical planning policy they want to use from now on, based on the results of this research, and fix them in the current tool. The tool then becomes less complex and probably more stable. However, we expect that the tool is still not stable enough to implement it in the current planning office, where there is no IT expertise.
3. Based on the results of this research, the LVP can also develop or outsource the development of a complete strategic, tactical, and operational scheduling tool. We would recommend the LVP to combine GIS functionality with the statistical functionality of programs like SPSS and R. Furthermore, we currently use a solver to determine the optimal route. However, with the current performance of the greedy estimator, we could also use the greedy routes.

Both the first and second option implicate the use of AIMMS as the platform for the decision support tool. However, there are several disadvantages of using AIMMS. First, it limits the possible forecast methods since it does not yet provide advanced statistical methods like Poisson regression. Second, the current graphical representation of the results by AIMMS is not able to meet the LVP requirements. Finally, the current problem with 6 helicopter and aggregated areas is small in comparison with the number of police cars in cities and the possible routes they can make. We expect that when the National Police expands the current program to optimize other police forces, the program will not be able to handle the problem size. Therefore, we recommend the LVP and the National Police to create a list of requirements for the final police force optimization software and determine which software package or supplier can match these requirements.

8.2 Flight route display

Since the tactical planning method creates a tactical plan at the operational level, with corresponding optimal routes, we can use these routes to direct police helicopter pilots to hotspot areas. As the police helicopter hardware cannot display pre-defined routes and the LVP cannot attach anything to a helicopter without certification, the LVP procured tablet computers that pilots can strap to their legs to display routes on.

The LVP chose the software package Skydemon[1] to display routes with. The program allows for planning routes in advance and provides information on the expected weather and relevant NOTAMs (short for Notice To Airmen). Furthermore, it constantly tracks the location of the device it runs on by GPS. In the prototype tool we created the possibility to export all routes as GPX files, which can be directly loaded by Skydemon. Figure 8.1 shows an artist's impression of how the LVP is currently able to display routes.

Figure 8.1: Artist's impression of in-flight display functionality of routes.

The LVP is currently implementing the hardware and software in the daily procedures of the pilots. We note that it is important that all pilots are taken into account during the implementation process. Guillemette et al. (2008) discuss the implementation of a RFID system to track security officers during

[1] www.skydemon.aero

their work. They note that it is important to have clear organizational directives, before implementing location tracking of employees. The superior's intentions regarding the use of the system appear to weigh heavily in the decision of employees to adopt the system.

We note that the routes that the tool currently creates often require the pilots to hover above one hexagon for an extended period of time. Since hovering helicopters create noise problems and hovering is very uncomfortable for helicopter crews, this is not desirable. The proposed solution by Van Urk (2012) to constrain the route optimization model to make helicopter routes that visit locations a limited number of time does not seem to work, since this leads to unclear routes with constant sharp turns, as shown in Figure C.4 in Appendix C. Furthermore, the routes consist of line between midpoints of hexagons, without taking the underlying geography into account. When pilots follow those lines consistently, the villages under these lines will experience more noise problems than villages outside these lines. Small adjustments by helicopter pilots to fly around villages to get to the next waypoint could thus be better.

Since we find that it is important to let pilots have some degree of autonomy and use their specific expertise in their work, we recommend the LVP to provide their pilots with the optimal route and enough information to let pilots determine how they want to make the route in practice. In addition to the optimal route we recommend to provide pilots with hot spot crime forecast maps and maps that show the pilots their expected added value at a certain time and location. The second map is a combination of the hot spot forecast and the coverage function of the helicopter. For example, the flight map could show that being at the optimal route at a given time results in the best possible performance for that time interval (100%), and that flying in a neighbouring hexagon results in an expected performance of 90%. The pilots can then choose between hovering at the best location and flying over a less optimal location. Figure 8.2 shows an exemplary hot spot forecast map (left) and the corresponding flight map (right) which indicates where the added value of a flying helicopter is maximal, based on the success function as illustrated in the middle. The success function shows the probability of success that a helicopter on the position of the red column has on the location of the green columns, where the height of the red column represents 100%.

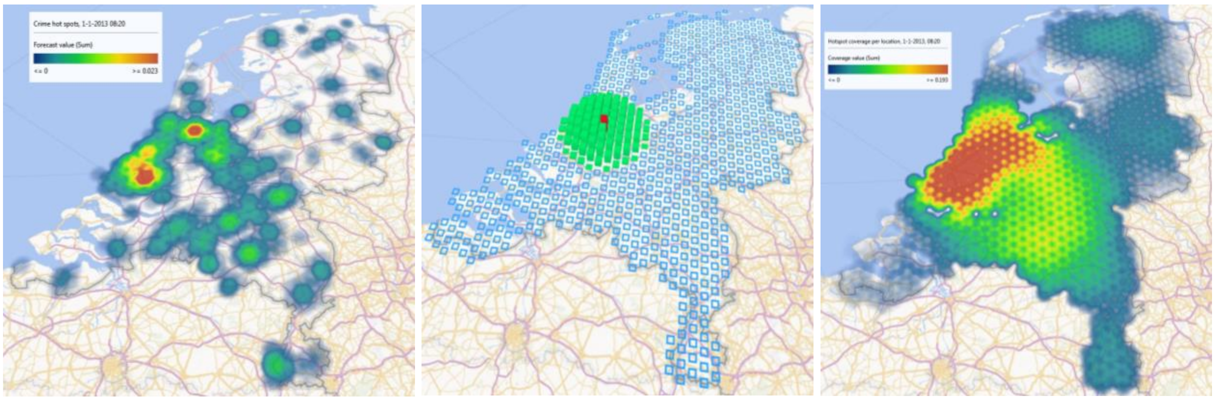

Figure 8.2: a hotspot forecast (left), coverage area of helicopter based on total coverage distance (middle), and resulting flight map (right).

In general, we find that is important to show decision makers and pilots the underlying reasons for tactical planning and helicopter routing. Therefore, we created functionality in the prototype tool to export forecasts, flight maps based on the forecasts and the success function, and routes to a format that can be used by the "Powermap" add-in for Microsoft Excel. The add-in is then able to make pictures such as Figure 8.2 and videos. Videos are especially useful to visualize forecasts, flight maps, and routes in time.

8.3 Actual performance measurement

To create visibility of real-time helicopter locations in the Central Operations Room in Driebergen, the LVP recently implemented hardware to track the police helicopters. We created an application that uses the position data that is recorded by this hardware to determine the performance of helicopters routes afterwards by comparing it with the actual incident data. Over a period of time the LVP can compare the actual proximity of police helicopters with the proximity of the planned routes at the planned times and determine the difference in performance. The LVP can use this information to build trust in the tactical planning methods amongst pilots.

8.4 Conclusion

This chapter discusses the prototype tool and how the LVP already implemented the results of the simulations. We find that the prototype tool is not yet ready for complete implementation and that the LVP and National Police have different options for implementation of the results of this research. Furthermore, we recommend and support the LVP to use the available hardware in the helicopters to measure its performance and provide feedback to its pilots.

9 Conclusion and recommendations

Section 9.1 contains the conclusions of this report. Section 9.2 discusses the limitations of this research and in Section 9.3 we provide recommendations for the LVP based on this research. In Section 9.4 we give suggestions for future research.

9.1 Conclusions

Before this research, the LVP had two disconnected decision support systems (DSS) that supported strategic and operational decision making. On the strategic level, the DSS showed the impact of locating stand-by crews on bases. The operational DSS determined optimal start times and routes of surveillance flights for one day ahead. The main problem of this research is the missing link in the support of the decision making process: the tactical planning, or the decision on when and where to allocate resources. This problem leads to the following problem statement:

What is the performance of the operational decision support tool and how can the LVP improve its performance by tactical planning?

To determine the long-term performance of the current DSS, we first extend the operational tool to a tactical tool. The tactical DSS requires input from a forecast method that takes into account seasonal and weekly trends, and is able to provide an annual forecast within a reasonable computation time. We thus extend the current forecast method and introduce crime patterns.

We propose a heuristic that is able to create a detailed tactical planning of surveillance flights for a year ahead in less than a day. The heuristic takes strategic and operational decisions into account, as well as spatial and temporal crime patterns on the operational level. Among others, the heuristic is able to use the flight value estimator and the route optimization model by Van Urk (2012) to determine the optimal timing and route of the surveillance flights.

We use simulation to determine the performance of the tactical planning with and without the methods by Van Urk (2012). We find that the expected performance improvement by the current operational decision support tools is approximately 20%. These results are limited by several constraints on the strategic and tactical level. For example, the LVP must provide 24/7 stand-by coverage from Schiphol, performs the same number of flights per day, and starts surveillance flights always from Schiphol. To determine the effect of loosening these constraints on the performance of the tactical planning, we create additional heuristics without these constraints.

By comparing the constrained with the unconstrained heuristics we find that being able to provide 24/7 stand-by coverage from other locations next to Schiphol can increase the incident coverage from stand-by by 40%. Furthermore, when the LVP allows a variable number of flights per day, we expect performance to increase by 6%. Finally, using multiple locations to start surveillance flights from enables another 12% of performance improvement.

During this research we found several other factors that influence tactical planning. The operational DSS tool by Van Urk (2012) was able to take incident priorities into account and restrict the number of times a route visits the same location. We find that taking priorities into account results in an increase in priority coverage by tactical planning of 3%. Constraining the route module leads to a decrease in performance and impractical flight routes. Furthermore, the LVP found that helicopters are more effective during the night. Taking effectiveness into account results in an increase of effectiveness coverage by more than 10%. Additionally, we find that shift schedules have an effect on the amount of overtime by helicopter crews, and based on the tactical planning we can thus optimize the shift schedules for overtime. Finally, since the optimal tactical planning is made to perform as efficiently as possible, the helicopter capacity is mainly allocated to the overall hotspots: big cities. As this can be considered unfair by other regions that also pay for the helicopter capacity, we propose a method to set the level of equity while creating a tactical planning. We find that we can reduce the resulting inequity by efficient planning by 50%, at a cost of 20% of performance.

The result of this research is a prototype decision support tool that integrates strategic, tactical, and operational decisions to create a complete one-year plan. The DSS is not yet user-friendly enough for direct implementation, but the results are already used by the LVP.

9.2 Limitations

This research has several limitations that we discuss in this section.

- The solutions that we propose in this research are limited by the available software at the LVP. Previous prototypes of decision support systems were developed in AIMMS and in this research we extend these techniques to enable tactical planning. While AIMMS is to the best of our knowledge not able to do, e.g., Poisson regression analysis, more advanced forecast methods were outside the scope of this research.
- The main assumption of this research is that helicopters are always in time for incidents that they can reach within 10 minutes. This research identifies that this is not accurate, and that the impact of the success function influences the tactical planning and operational routes significantly.
- We were not able to perform a full factorial analysis of all methods and settings due to the extensive computation time per experiment. Also, we use a limited sample size for our simulations since we base our performance measurement on historical data.
- We did not take operational constraints as detailed as refuelling times between flights and rest times for crews into account. Furthermore, we consider only Schiphol, Rotterdam, and Volkel as bases to start flights on and assume that flights start and end at the same location. However, it is also possible to combine multiple flights between locations into a round tour. Next to that, it would be possible to consider helicopters staying overnight on Rotterdam and Volkel.
- This research applies only to the EC135 helicopters of the LVP and creates flight routes that focus on the major Dutch cities when we do not take fairness into account. Since the AW139 helicopters are faster and have a longer range, they would be the ideal addition to the model to cover the more remote areas of the Netherlands with.

9.3 Recommendations

In this section we give recommendations that do not require further research. Our main recommendation is to professionalise the current prototype tool, since the current prototype tool is not user-friendly enough for implementation without support.

Second, we advise the LVP to use freely available GIS software to visualize the patterns in the available criminal incident data. We found that these visualizations have great educational value for police helicopter pilots.

Third, since helicopter routes and the allocation of crews to stand-by locations depend strongly on the relation between arrival time and the success of helicopter deployments, we propose to extend the current performance measurement by defining arrival time and measuring the arrival time to incidents.

Fourth, we recommend the LVP to communicate the crime forecasts and helicopter routes to other police departments. The crime forecasts and decision heuristics create insight into the LVP decision process, and are valuable for other police departments to evaluate their own decision processes. Furthermore, when the timing and routes of surveillance routes are known to police districts, the districts expect the helicopter and can prepare for the helicopter capacity they receive by increasing the number of police officers on the street. This can increase the success rate of the helicopter deployments.

Fifth, we suggest the LVP to measure the effectiveness of police officers on the ground in different regions and provide feedback. The LVP could even sent the helicopter less often to regions with less than average effectiveness, after which the regions will hopefully request helicopter support only when there is an actual probability to find or catch a criminal, and regions will work harder to make the deployment of helicopters worthwhile.

Finally, we recommend the LVP to increase the intelligence available to pilots during the surveillance flight. For example, the LVP could provide the real-time rerouting functionality of the prototype tool of Van Urk (2012) for helicopters that have been deployed to incidents and become available again.

9.4 Future research

To improve the usefulness of tactical planning for police helicopters, we suggest to perform future research to simultaneously optimize the allocation of stand-by crews and timing of surveillance flights. To facilitate the implementation process of tactical planning of public resources, future research directions could extend and validate the equity introduction to the multi-route heuristic. Furthermore, the (equal) routing of police helicopters can be translated to other police forces, but also to ambulances and other mobile public services. Not only is it possible to apply the same principles to other police forces, it is even more interesting to let them work together. We thus propose to extend the forecasting and routing methods to include police officers on foot or in cars. Furthermore, we expect that more advanced forecasting methods can create more accurate forecasts. In this regard, we propose to determine the effect of holidays and other special dates on the criminal intensity at those dates.

Specific for the LVP, we recommend to extend the current methods to include a heterogeneous fleet. For the LVP this is useful, as they have two helicopter types with completely different specifications, which can handle different tasks. EC 135 helicopters are not able to perform a surveillance flight and handle incidents in the eastern parts of the Netherlands due to their limited range. Therefore, AW helicopters can be used to cover the areas that are too far away of the LVP helicopter bases. We expect that this is also the case for other air support departments and other emergency services. Unique for police resources is the preventive effect they seem to have. Future research is required to determine whether and how this effect works, and how the police can use this effect to maximize the performance of its resources.

As we did not test the ILP formulation as discussed in Section 5.3, we recommend the LVP to build and validate it. After validation, it is possible to extent the formulation with practical constraints like refuelling times and rest times for crews. Furthermore, it would be interesting to compare the results from rough-cut capacity planning with the results from the current integrated tactical planning heuristics. Also, we suggest to determine the effect of the flight and crew budget on the added value of variable tactical planning and the trade-off between the crew and flight budget. Furthermore, since flights are cancelled due to, e.g., weather influences, we recommend to support the planning office in making the decision when to schedule replacement flights.

During surveillance flights, commonly multiple incidents happen and the Central Operations Room has to decide whether to deploy a helicopter to an incident. Deployment to an incident far away from the current helicopter position results in a longer handling duration than deployment to incidents that are closer by. Other factors that the Central Operations Room can take into account are incident priorities and the expected effectiveness of the helicopter at the incident location. Since helicopters are sometimes deployed to low priority incidents, during which more important incidents happen, we see potential in the optimization of the deployment strategy. The deployment strategy for the helicopter could also be combined with deployment strategies for police officers on foot and in cars.

Bibliography

Block, C. R. (1995). STAC Hot Spot Areas: a statistical tool for law enforcement decisions.

Bowers, K. J. (2004). Prospective Hot-Spotting: The Future of Crime Mapping? *British Journal of Criminology*, *44*(5), 641–658. doi:10.1093/bjc/azh036

Brotcorne, L., Laporte, G., & Semet, F. (2003). Ambulance location and relocation models. *European Journal of Operational Research*, *147*(3), 451–463. doi:10.1016/S0377-2217(02)00364-8

Buiteveld, B. (2011). *Het oog in de lucht: het optimaliseren van de plaatsing van de politiehelikopters aan de hand van incidenten*.

Caplan, J. M., Kennedy, L. W., & Miller, J. (2011). Risk Terrain Modeling: Brokering Criminological Theory and GIS Methods for Crime Forecasting. *Justice Quarterly*, *28*(2), 360–381. doi:10.1080/07418825.2010.486037

Chanta, S., Mayorga, M. E., & McLay, L. A. (2011). Improving emergency service in rural areas: a bi-objective covering location model for EMS systems. *Annals of Operations Research*. doi:10.1007/s10479-011-0972-6

Chawathe, S. S. (2007). Organizing Hot-Spot Police Patrol Routes. In *2007 IEEE Intelligence and Security Informatics* (pp. 79–86). IEEE. doi:10.1109/ISI.2007.379538

Chen, P., Yuan, H., & Shu, X. (2008). Forecasting Crime Using the ARIMA Model. In *2008 Fifth International Conference on Fuzzy Systems and Knowledge Discovery* (pp. 627–630). IEEE. doi:10.1109/FSKD.2008.222

Church, R., & Revelle, C. (1974). The Maximal Covering Location Problem. *Papers of the Regional Science Association*, (32), 101–118.

Cohen, J., Gorr, W., & Olligschlaeger, A. Leading Indicators and Spatial Interactions: A Crime Forecasting Model for Proactive Police Deployment. , Heinz College Research (2004). Retrieved from http://repository.cmu.edu/heinzworks/200

Davidson, F., Allen, A., Brassington, G., Breivik, Ø., Daniel, P., Kamachi, M., … Sutton, M. (2009). Applications of Godae Ocean Current Forecasts To Search and Rescue and Ship Routing. *Oceanography*, *22*(3), 176–181. doi:10.5670/oceanog.2009.76

Eksioglu, B., Vural, A. V., & Reisman, A. (2009). The vehicle routing problem: A taxonomic review. *Computers & Industrial Engineering*, *57*(4), 1472–1483. doi:10.1016/j.cie.2009.05.009

Felson, M., & Poulsen, E. (2003). Simple indicators of crime by time of day, *19*, 595–601.

Gorr, W., & Harries, R. (2003). Introduction to crime forecasting, *19*, 551–555.

Gorr, W., Olligschlaeger, A., & Thompson, Y. (2003). Short-term forecasting of crime, *19*, 579–594.

Guillemette, M. G., Fontaine, I., & Caron, C. (2008). Hybrid RFID-GPS Real-Time Location System for Human Resources: Development, Impacts and Perspectives. In *Proceedings of the 41st Annual*

Hawaii International Conference on System Sciences (HICSS 2008) (pp. 406–406). IEEE. doi:10.1109/HICSS.2008.195

Hans, E. W., Herroelen, W., Leus, R., & Wullink, G. (2007). A hierarchical approach to multi-project planning under uncertainty. *Omega*, *35*(5), 563–577. doi:10.1016/j.omega.2005.10.004

Hyndman, R. (2006). Another look at forecast-accuracy metrics for intermittent demand. *Foresight: The International Journal of Applied …*, (4), 43–46. Retrieved from http://www.researchgate.net/publication/5055536_Another_Look_at_Forecast_Accuracy_Metrics_for_Intermittent_Demand/file/d912f50ff0c2ad9136.pdf

Knight, V. A., Harper, P. R., & Smith, L. (2012). Ambulance allocation for maximal survival with heterogeneous outcome measures. *Omega*, *40*(6), 918–926. doi:10.1016/j.omega.2012.02.003

Kolesar, P., & Walker, W. E. (2012). An Algorithm for the Dynamic Relocation of Fire Companies Author (s): Peter Kolesar and Warren E . Walker Reviewed work (s): Published by : INFORMS Stable URL : http://www.jstor.org/stable/169582 . An Algorithm for the Dynamic of Fire Companies Reloca, *22*(2), 249–274.

Korteweg, F. (2014). *Helicopterview*. Zoetermeer.

Laporte, G. (1992). The vehicle routing problem: An overview of exact and approximate algorithms. *European Journal of Operational Research*. doi:10.1016/0377-2217(92)90192-C

Law, A. M. (2007). *Simulation modeling and analysis*. *McGraw-Hill series in industrial engineering and management science*.

Li, X., Zhao, Z., Zhu, X., & Wyatt, T. (2011). Covering models and optimization techniques for emergency response facility location and planning: a review. *Mathematical Methods of Operations Research*, *74*(3), 281–310. doi:10.1007/s00186-011-0363-4

Liu, H., & Brown, D. E. (2003). C riminal incident prediction using a point-pattern-based density model, *19*, 603–622.

Marsh, M. T., & Schilling, D. A. (1994). Equity measurement in facility location analysis: A review and framework. *European Journal of Operational Research*, *74*(1), 1–17. Retrieved from http://www.sciencedirect.com/science/article/pii/0377221794902003

Moreira-matias, L., Gama, J., Ferreira, M., & Damas, L. (2012). A predictive model for passenger demand on a taxi network. In *2012 15th International IEEE Conference on Intelligent Transportation Systems* (pp. 1014–1019).

Osgood, D. W. (2000). Poisson-Based Regression Analysis of Aggregate Crime Rates. *Journal of Quantitative Criminology*, *16*, 21–43. doi:10.1023/a:1007521427059

Owen, S. H., & Daskin, M. S. (1998). Strategic facility location: A review. *European Journal of Operational Research*, *111*(3), 423–447. doi:10.1016/S0377-2217(98)00186-6

Ratcliffe, J. H. (2002). Aoristic Signatures and the Spatio-Temporal Analysis of High Volume Crime Patterns. *Journal of Quantitative Criminology*, *18*(1), 23–43. doi:10.1023/A:1013240828824

Reis, D., Melo, A., Coelho, A. L. V, & Furtado, V. (2006). GAPatrol : An Evolutionary Multiagent Approach for the Automatic Definition of Hotspots and Patrol Routes. *Society*, 118 – 127. doi:10.1007/11874850_16

Reuter, M., & Michalk, W. (2012). Towards the Dynamic Relocation of Ambulances in Germany: The Risk of Being Too Late. In *2012 Annual SRII Global Conference* (pp. 642–649). IEEE. doi:10.1109/SRII.2012.78

Schilling, D., Elzinga, D. J., Cohon, J., Church, R., & ReVelle, C. (1979). The Team/Fleet Models for Simultaneous Facility and Equipment Siting. *Transportation Science*. doi:10.1287/trsc.13.2.163

Sherman, L. W., Gartin, P. R., & Buerger, M. E. (1989). Hot spots of predatory crime: routine activities and the criminology of place. *Criminology*, *27*(1), 27–56. doi:10.1111/j.1745-9125.1989.tb00862.x

Smith, H. K., Harper, P. R., & Potts, C. N. (2012). Bicriteria efficiency/equity hierarchical location models for public service application. *Journal of the Operational Research Society*, *64*(4), 500–512. doi:10.1057/jors.2012.68

Stone, D. (2002). *Policy Paradox: The Art of Political Decision Making*. *New York W W Norton Company* (Vol. Revised, p. xiv, 394 p.). Retrieved from http://www.amazon.co.uk/Policy-Paradox-Political-Decision-Making/dp/0393976254

Toregas, C., Swain, R., Revelle, C., & Bergman, L. (1971). The location of emergency service facilities. *Operations Research*, *19*(6), 1363–1373.

Van Essen, J. T., Hurink, J. L., Nickel, S., & Reuter, M. (2013, October 1). Models for ambulance planning on the strategic and the tactical level. Beta Research School for Operations Management and Logistics. Retrieved from http://eprints.eemcs.utwente.nl/23756/

Van Urk, R. (2012). *Helicopter view, positioning helicopters where they make a difference*.

Willis, H. L. (2002). *Spatial Electric Load Forecasting* (Vol. 8, p. 760). CRC Press. Retrieved from http://books.google.com/books?id=_2HL-vs8HPgC&pgis=1

Wuschke, K., Clare, J., & Garis, L. (2013). Temporal and geographic clustering of residential structure fires: A theoretical platform for targeted fire prevention. *Fire Safety Journal*, *62*, 3–12. Retrieved from http://www.sciencedirect.com/science/article/pii/S0379711213001148

Overview of mathematical notation

Sets

- H Helicopters
- L Locations
- T Times
- D Days
- S Shifts
- C Crews
- F Flights
- N Incidents

Parameters

$f_{l,t}$	Forecast value at location $l \in L$ at time $t \in T$
$i_{l,t}$	The number of incidents at location $l \in L$ at time $t \in T$
h_d	Availability of helicopter $h \in H$ at day $d \in D$
$sv_{l,t}$	The value of standing-by on location $l \in L$ at time $t \in T$
$fv_{l,t}$	The value of starting a flight from location $l \in L$ at time $t \in T$
fb	Flight hour budget
cb	Crew budget
fd	Standard flight duration
q	Parameter that sets the level of equity
$d_{l,m}$	Distance between location $l \in L$ and location $m \in L$
$c_{l,m}$	Coverage of location $l \in L$ by location $m \in L$
$e_{l,t}$	Effectiveness of helicopter deployment at time $t \in T$
k	

Variables

C_s	Availability of crew $c \in C$ during shift $s \in S$
$F_{l,t}$	Whether a flight starts from location $l \in L$ at time $t \in T$
$S_{l,t}$	Whether there is a crew stand-by on location $l \in L$ c
$L_{l,t,h}$	1 when helicopter $h \in H$ is at location $l \in L$ at time $t \in T$, otherwise 0.
$Coverage_{l,t}$	Maximal coverage of location $l \in L$ at time $t \in T$ by the helicopters
$Covered_{l,t}$	1 when location $l \in L$ at time $t \in T$ is covered by a helicopter, otherwise 0.
E_l	Effect of a set of routes on location $l \in L$
A_l	Attribute of location $l \in L$

Appendices

Appendix A: complete organizational diagram National Police

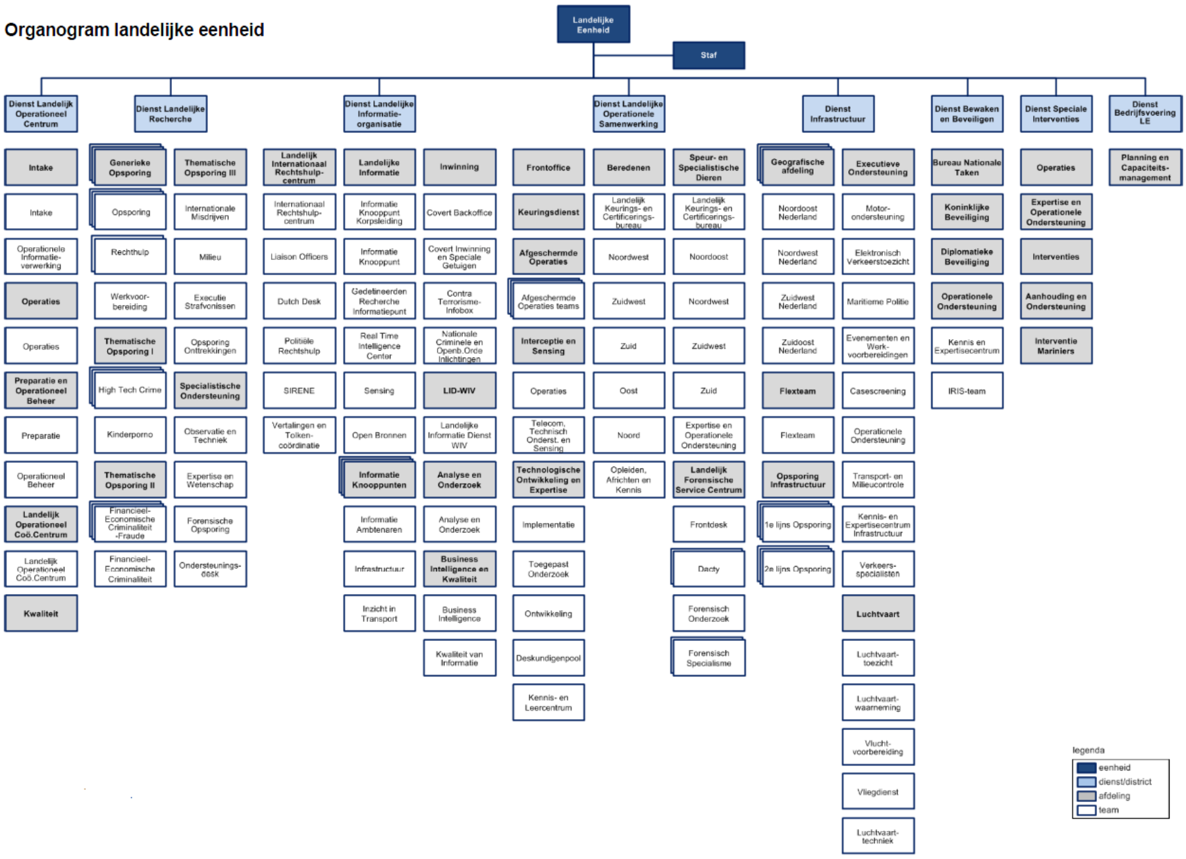

Figure A.0.1: Complete Organization Diagram of the National Police (inrichtingsplan-nationale-politie.pdf, http://www.rijksoverheid.nl/documenten-en-publicaties/kamerstukken/2012/12/07/inrichtingsplan-nationale-politie.html).

Appendix B: simulation tool functionalities

This appendix describes the current functionality of the simulation tool that we use in Chapter 7. Section B.1 discusses the user interface. Section B.2 explains which types of input the tool requires. Finally, Section B.3 discusses how we store and retrieve data for later use.

B.1 User Interface

This section describes the user interface of the simulation model. The goal of the user interface is to make it easy for users to perform multiple experiments in one run and to provide a clear overview of the differences in input and output per experiment. Section B.1.1 describes the possibilities for users to manipulate the settings per experiment. Section B.1.2 explains the workings of the main controls and Section B.1.3 discusses how the simulation model displays the results of the experiments.

B.1.1 Input

The main page of the simulation model is the input/output page where we determine the simulation settings and get a first glance at the results after the simulations are finished. Figure B.1 shows the user interface.

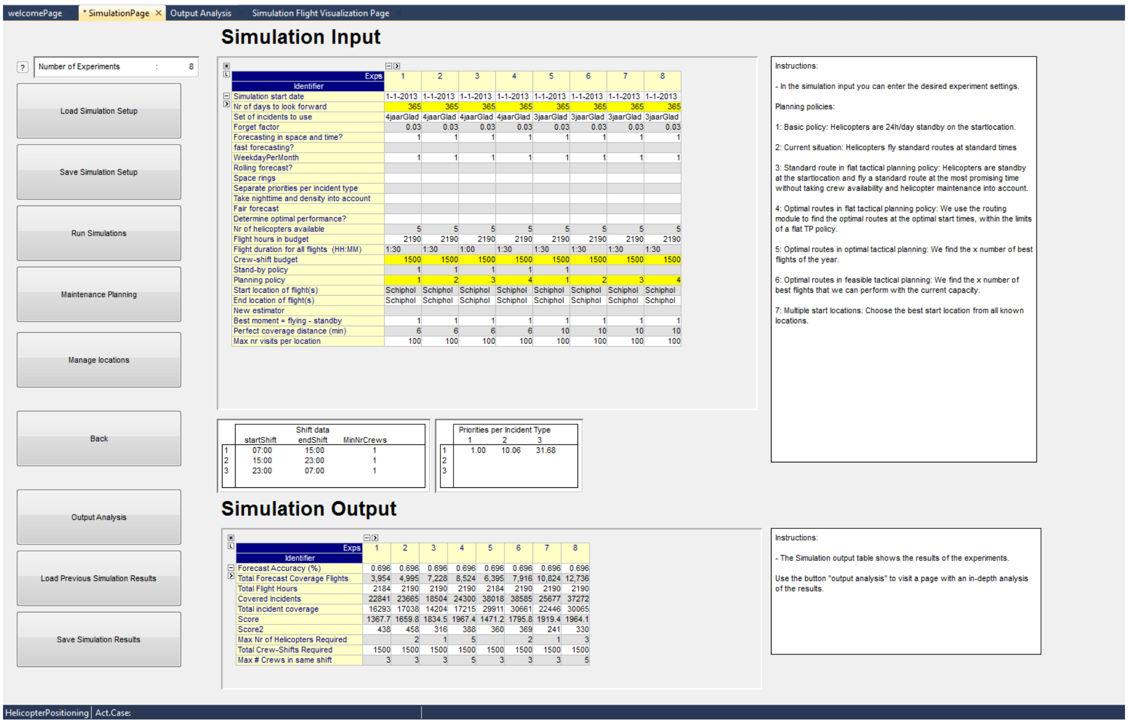

Figure B.1: User interface of the input page of the simulation model.

Per experiment (of which the number is set in the top left of the interface), we can determine the scope and type of the forecast, how we are going to plan and route (with stand-by and planning policies), and how much crew and flight capacity the LVP wants to use during the scope.

We start with the scope of the experiment. We enter the start date of the experiment and the number of days we want to look forward. The start date requires a date and the number of days to look forward should contain an integer value. When we would make a forecast for exactly one year in advance, the added value of flight routes in the last hour of the year would decrease since part of the route would take place in the next, not-forecasted, year. To prevent this, we always forecast for the scope of the experiment plus one day.

Most of the settings for the experiments concern the forecasting method. We base the forecast on historical data and since the model already contains several incidents sets, we choose which one to use. We then set the forget factor: This factor influences the impact of old data on forecasts. It should be a value between 0 and 1. Example: when the forget factor is 0.01 then data that is 12 months old will be taken into account with a factor 0,99^12≈0,887. Furthermore, we can switch the fast forecasting technique on and off, as well as the spatial and temporal aggregation and the number of rings to use in the spatial aggregation, to compare all combinations off settings during validation. Finally, we can set whether the model takes into account priorities, day and night time effectiveness, and fairness.

With the scope and forecasting method known, we determine the number of helicopters we would like to use for the experiment, the number of flight hours that are available in the budget, and the duration of the flights.

The next settings determine how and where we allocate the helicopter and crew capacity to the forecast: the stand-by policy, planning policy, and the start and end location of all flights. We already explained the different planning policies in sections 7.1.3 and 7.1.4. The effect of the start and end location settings on the model depends on the planning policy. For planning policy one the start location is the location of the 24/7 stand-by crew, for stand-by policies 1 to 3. For planning policy two and three the start and end locations do not matter, since we use a standard flight from and to Schiphol. For planning policy four to six, the start and end location determine the actual origin and destination of all flights. Planning policy seven does not require a start or end location, since this policy determines the optimal allocation of flights by itself.

Finally, we have two settings for the routing model. We can set the coverage of the helicopter by entering the total coverage distance, and we can set the maximum number of times a helicopter may visit a location during one flight. This last setting makes sure that routes do not include periods where helicopters hover above a hotspot. The value should be > 1 since otherwise the helicopter cannot visit its starting location again and thus cannot land at its starting location (commonly Schiphol).

B.1.2 Main Simulation Controls

The input page contains several buttons to control the simulation (data) with:

- Save simulation setup: This button allows users to save the current simulation setup.
- Load simulation setup: This button allows users to open an existing simulation setup, reducing time required to enter setups.
- Run simulations: This button activates the simulation process. The model will now sequentially run the experiments.
- Maintenance planning: this button opens a page with a maintenance planning Gantt chart, with which helicopters can be blocked for maintenance per entire day for the entire scope.
- Manage locations: this button opens a pop-up page on which we can enter new locations and determine the response time per location.
- Output analysis: This button opens a page where the simulation outputs are visualized.
- Save simulation results: This button saves simulation results (and the associated settings) to enable the user to look back to the data later without having to run the simulations again.
- Load previous simulation results: This button loads previous simulation results for analysis afterwards.

B.1.3 Outputs and analysis

When the simulations are finished, we use the button “Output Analysis” to go to a page on which we show the experiment results. Figure B.2 shows the user interface of the output analysis page.

The page is designed to analyse the information in two stages. First, the user is able to evaluate and compare the performance of different planning policies and other parameter settings. For this purpose there is a table with the outcomes per experiment and five bar charts with the most important data points. For example, the graphs show the incident coverage per experiment. With this last value we can compare experiments with different quantities of flight hours.

In the second step, the user is able to evaluate the output per experiment. The user can select the experiment to analyse and the bottom five charts are updated with the experiment specific data. The charts show the expected forecast coverage per flight, the actual incident coverage per flight, the number of crews and flights per location, and the number of flights per day.

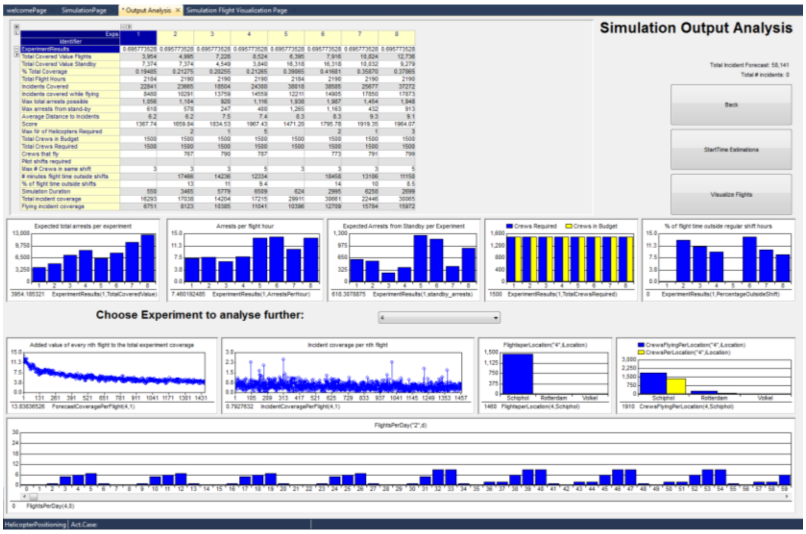

Figure B.2: User interface of output analysis page.

Finally, the individual routes and the annual tactical plan are available through use of the button “visualize flights”.

B.1.4 Visualizations of flights and flight planning

Figure B.3 shows the page at which we visualize flight routes and combine flight routes with the underlying forecast. This page also contains a Gantt chart functionality that shows the flight schedule over time, together with the shifts, the crews that are assigned to the flights, and the helicopter that is scheduled for the flight. Users can select for which experiment they want to see the Gantt chart and routes, and can choose which flight they want to see the route of. Users can automatically export the routes per experiment to files in the .kml or .gpx format, which can be used in Google Maps (functionality made by Van Urk (2012)) and Skydemon (as discussed in Section 8.2) respectively.

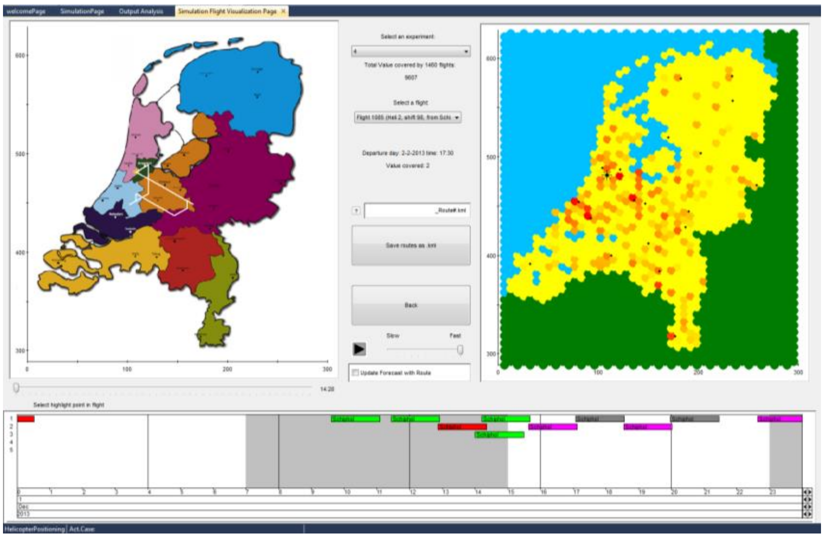

Figure B.3: User interface of the simulation visualization page.

B.2 Input requirements

The model requires two sets of inputs: crime data and capacity information. The tool has

We extended the LVP capacity settings by introducing shift and crew settings. We can enter the number of shifts and their respective start and end times and the model then creates the number of shifts that is required for the scope of the incident. Per shift we can then enter the minimum number of crews that should be present. After we have then set the number of crews that is available to fly during the scope of the experiment, the model plans the minimally required number of crews to fill all the shifts and leaves the rest open for later deployment. Before the model runs the simulation it checks whether there are enough crews to fill all shifts, and asks for more capacity when this is not the case.

B.4 Data manager

Since we do not want to make forecasts again for every experiment, we developed a module that saves new forecasts and checks for saved forecasts that are useful for experiments before it makes a new forecast. When the model contains a compatible forecast that has enough or more data than required, then the model takes the data that is required from the forecast and does not need any forecast calculations. When the model finds a partly compatible forecast than it takes the available data and only determines the extra forecast data that is required. This module saves forecasting time. For example, there is a forecast available that starts at 1-6-2013 and has a scope of one year. When the user then asks for a forecast that starts at 1-6-2013 and forecasts a half year ahead, then the module loads the yearly forecasts and cuts the last half year off.

Appendix C: AIMMS fixes

In this appendix we discuss two of the alterations we have made to the original AIMMS tool to solve earlier bugs. Section C.1 discusses the factors originally used for temporal aggregation, and Section C.2 explains how spatial aggregation led to undesirable results.

C.1 Temporal aggregation

Figure C.1 shows the factors Van Urk (2012) used to spread the forecast over time.

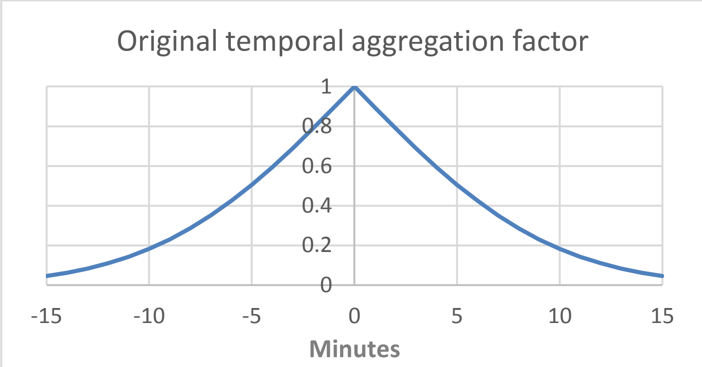

Figure C.1: Temporal aggregation factors in model from Van Urk (2012).

Figure C.2 shows the result of the temporal aggregation step on the total forecasted incident value of all time intervals in one day (720 intervals of 2 minutes).

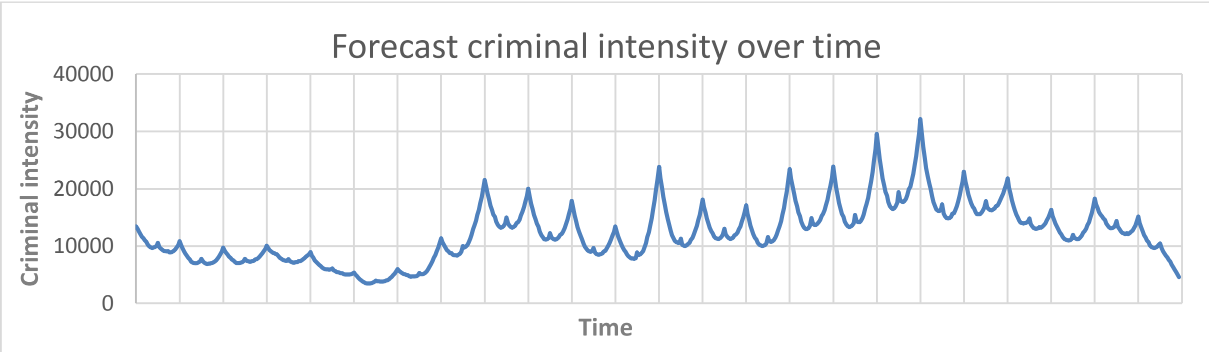

Figure C.2: Total value of forecast over the Netherlands per time interval over one day.

We propose to use the bell-shaped curve from a normal distribution with standard deviation 7.5 as described in Van Urk (2012). Figure C.3 shows both curves.

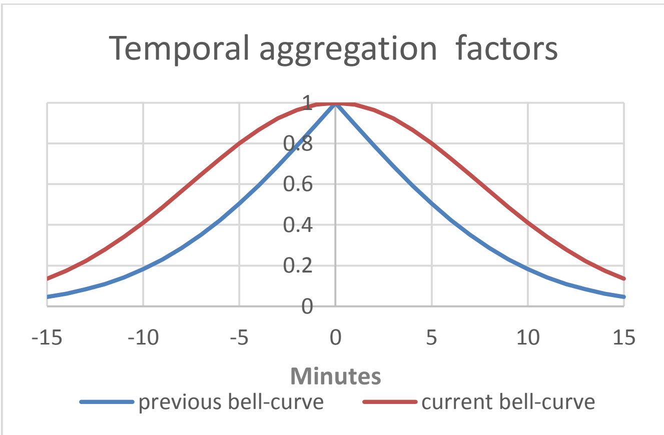

Figure C.3: Comparison of temporal aggregation factors from previous and current research.

C.2 Spatial aggregation outside the Netherlands

During the spatial aggregation step, incidents near the sea or the borders were spread out over the sea and outside the Netherlands. Figure C.4 shows that this has resulted in routes over water. We also noticed that the tool created routes across the Dutch borders. By constraining the program to forecast and route only for hexagons in the Netherlands we improve helicopter routes and decrease computation time.

Figure C.4: Helicopter routes over the sea due to unrestricted spatial aggregation.

Appendix D: fast forecasting procedure inaccuracy

There are two causes for inaccuracy in the fast forecasting method: the precision of the hour-blocks and the temporal aggregation effect between days. Figure D.1 shows that we create cut-off points at whole hours when we multiply the entire hour-block with a conversion factor. This commonly creates an inaccuracy that increases from the centre of the hour to the start and end of the hour.

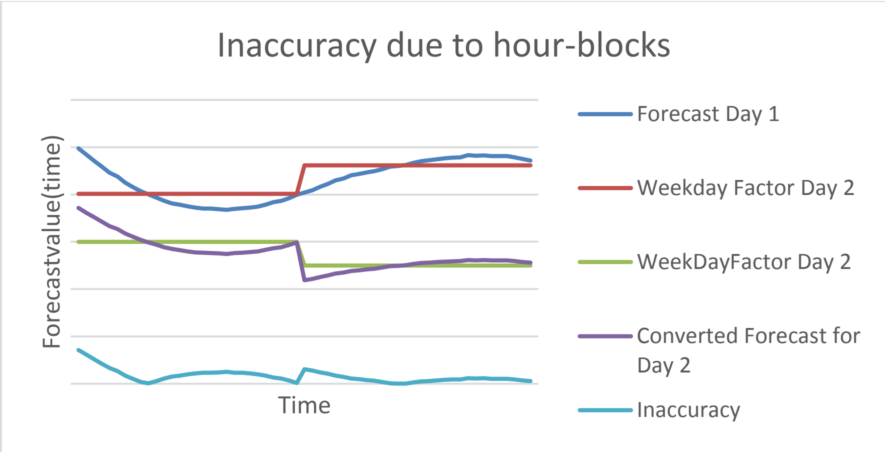

Figure D.1: Illustrative example of inaccuracy due to hour-blocks.

The second problem is the effect that days have on each other due to the temporal aggregation step. Since incidents in the last quarter hour of day 1 have an impact on the incidents in the first quarter of day 2, we observed inaccuracies. Figure D.2 shows the solution of this problem: we create a forecast for day 1, with its effect on day 2 and save it. We then create a forecast for day 2 with its effect on day 1 and 3 and save it. To create the complete forecast for day 1 and 2 we combine both forecasts. To create a forecast for day 1 to 3 we convert the forecast of day 2 to the distribution of day 3 and combine all forecasts. This means that we only perform the spatial and temporal aggregation steps for day 1 and 2 and create day 3 and further by converting day 2 and its time effects.

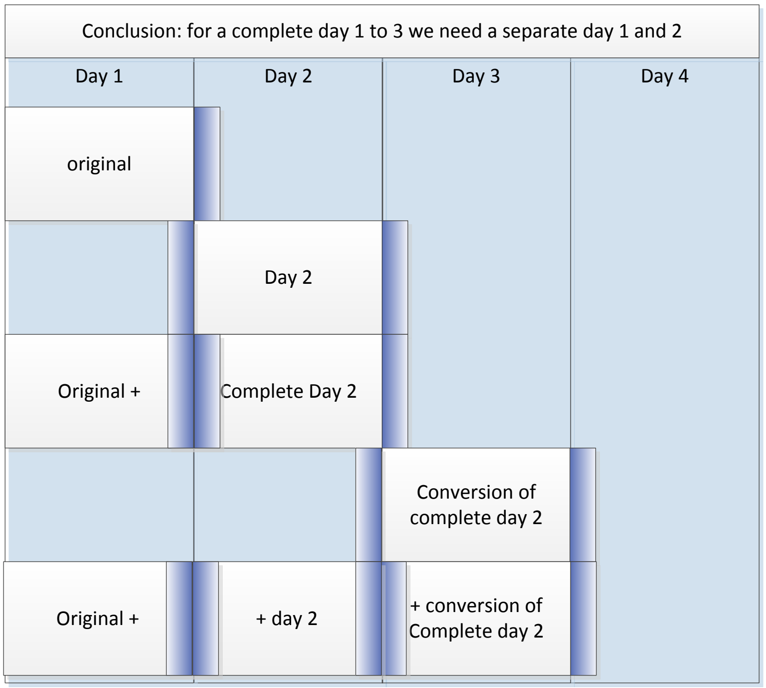

Figure D.2: Visualization of combination of forecast of day 1 and 2 to create a forecast for days 1 to 3.

Appendix E: fast estimation heuristic

We develop a simple heuristic that only calculates the most promising estimations at points that seem to be near a peak. We determine a step size n and calculate the estimation for every nth time interval. When the next estimation is lower than its predecessor, while its predecessor was bigger than the estimation before it, then we just missed a peak. We then start calculating estimates from the predecessor until we find the peak.

To determine the optimal step size we perform an experiment. We determine the number of calculations required per step size and the % of peaks that we miss for several step sizes. We use an estimation result for flights with a duration of 2 hours for one week of forecast (5040 time intervals). The step size should always be smaller than the flight duration, since the flight duration determines the width of the peaks (when the flight duration is 20 minutes, then all start times that are less than 20 minutes before a peak in the forecast are influenced by the peak). Since we estimate the minimum flight time as 40 minutes we take 20 time intervals (of 2 minutes) as the maximum step size. From Figure E.1 we conclude that the optimal step size is around 10 with an expected saving of 80% of calculation time. Smaller step sizes lead to more samples and bigger step sizes require more calculations for the exact peak and increase the risk that we miss a peak.

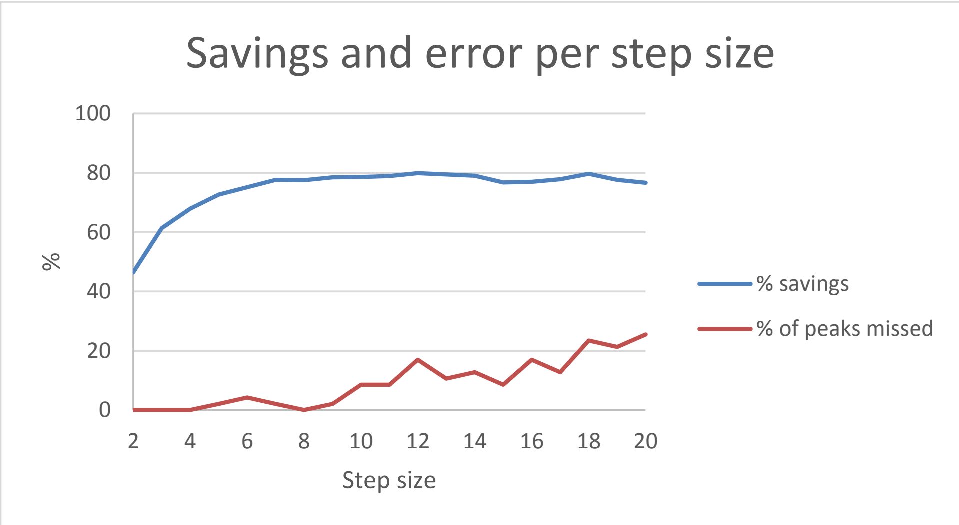

Figure E.1: Overview of the analysis on savings per step size for the most promising start size estimation system.

Figure E.2 shows the result of the heuristic in red against the full estimation graph in blue for a step size of 20. It shows that the heuristic misses some local maxima but covers all interesting peaks.

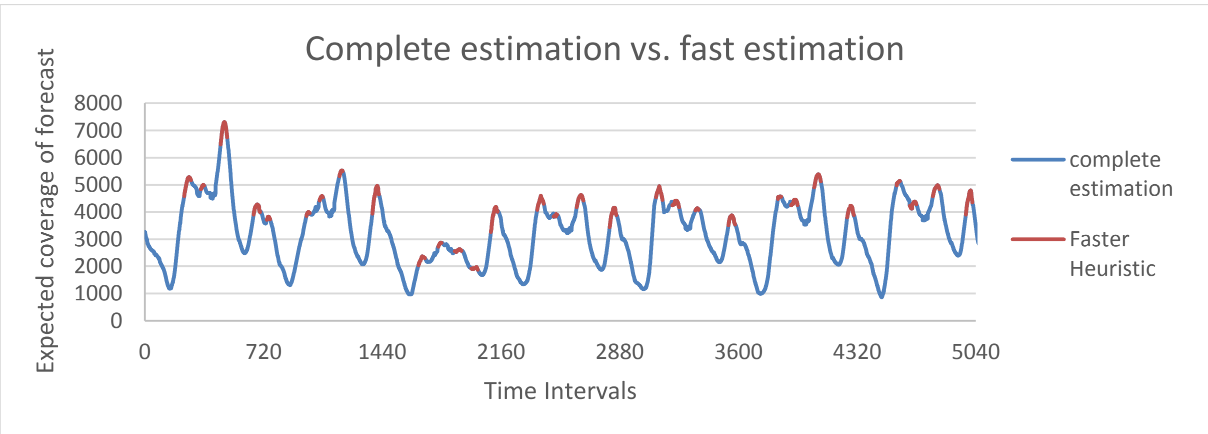

Figure E.2: Comparison of complete estimation of route quality and the fast heuristic.

Appendix F: coverage of helicopter degrades over the flight duration

In Section 4.4.2 we note that the success function may depend on the number of minutes. Since helicopter routes commonly cover multiple incidents, we assume that helicopters are deployed to an incident during every flight. Therefore, we assume that helicopters do not finish their route. When a helicopter is deployed to an incident it leaves its optimal route and goes the location of the incident. It then takes some time to handle the accident, after which the helicopter is available for deployment again. The helicopter then determines a new optimal route for the remaining flight duration from the new location, based on the flight routes of any other airborne helicopters. This optimal route will probably differ from the initial route, as areas visited in the initial route can be too far away for the remaining flight duration, or hot spots cannot be reached in time and thus are not interesting anymore.

We currently schedule flights sequentially and update the forecast after every flight. When two flights are scheduled around the same time, then the second flight is influenced by the first flight, since part of the forecast is already covered by the first flight. When the probability that the helicopter is on its optimal route decreases over time, then the end of the flight should have less influence on the next flight route than the start of the flight.

We assume that flights are not cancelled beforehand and thus know for sure that the helicopter flies the very first part of the route. Furthermore, we assume that the helicopter never finishes the flight. We then need a function for how the probability that the helicopter is on its initial route decreases, given the number of minutes that the helicopter is airborne. We identify two options:

1. Linear decrease: every time interval, the probability decreases by the same fraction.
2. Based on the underlying forecast: the probability that a helicopter is still on route decreases faster when the helicopter is above a hotspot, than when the helicopter flies between hotspots.

We conclude that it is most plausible and realistic to use the second option. Figure F.1 shows an example of the method we propose to use, for a helicopter flight that has a duration of 60 time intervals of 2 minutes (2 hours). We scale the total value of the forecast covered by the helicopter flight to the value of 1 and then decrease the probability of the helicopter availability per time interval by the covered forecast value of that time interval.

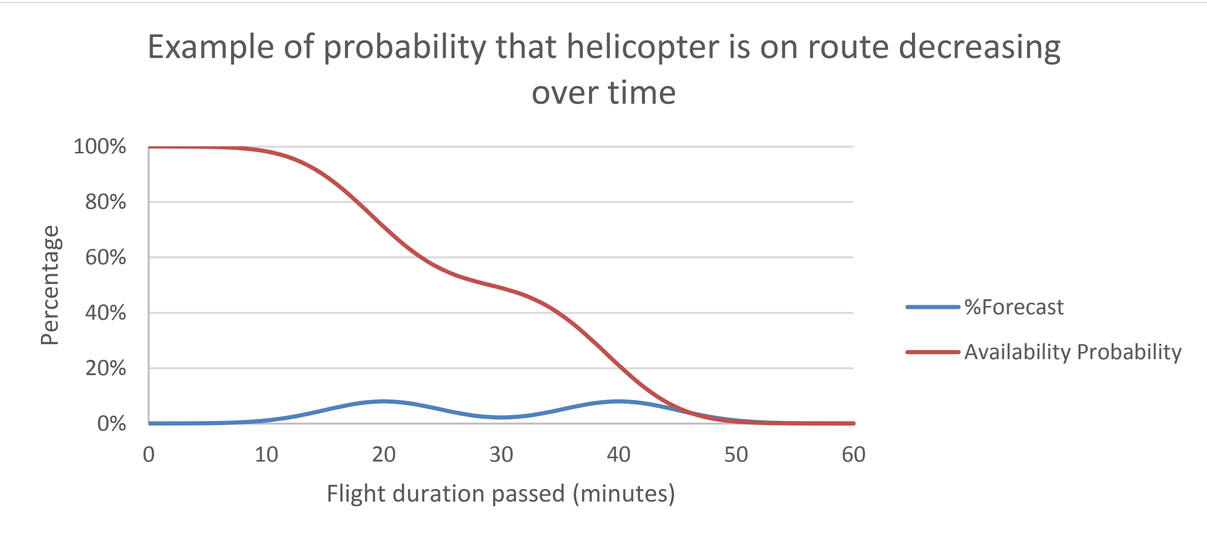

Figure F.1: Example of probability that helicopter is on route, decreasing over time