2019

Optimisation of patient blueprint schedules with downstream departments and uncertain care pathways

SJOERD MESMAN

Sint Maartenskliniek

UNIVERSITEIT TWENTE.

Supervisory committee

First supervisor:	dr. A.G. Leeftink
Second supervisor:	prof. dr. E.W. Hans
Organisation supervisors:	M.C. Carlier, MSc Healthcare logistics advisor
	R. Vromans, MSc Healthcare logistics advisor

Management summary

The Sint Maartenskliniek (SMK) is a hospital group in the Netherlands specialized in movement and posture, excelling in complex procedures.

Currently, the flow of patients across departments fluctuates during the day. These fluctuations result in additional waiting time for patients, uneven workload for staff, and staff overtime. The hospital wants to control patient flows in the orthopaedic care chain to reduce these issues.

Patients that have a consultation at the Orthopaedic Outpatient Department (OOD) often need to visit other departments as well. The current patient flow patterns show workload arrival variability at departments downstream from the OOD, with peaks during the day at the Radiology department and follow-up consultations at the OOD. The waiting time also shows an upward trend in the morning for all departments, with the Orthopaedic Outpatient department showing the strongest increase.

We want to control the workload at downstream departments caused by patients from the OOD. Data analysis shows the *probability*, *duration*, and *timing* of this downstream workload is different for each consultation type and unit, which combine into a demand profile for the consultation type. By spreading consultation types with high demand evenly throughout the day a more even spread of workload at downstream departments can be achieved.

An automated method of generating these schedules is needed. We define a mathematical (MILP) model that is able to generate blueprint session schedules based on the demand profile of consultation types. The model can generate multiple blueprint schedules in one session simultaneously, matching the combined outflow of patients as close to the desired pattern as possible. The desired pattern might be a flat line throughout opening hours of the department, but the norm can also be set to different values during the day for each department.

Figure M.1 shows the reduction of demand peaks and variability in downstream demand, achieved by changing the sequence and starting time of consultations in the schedule. The model achieves reduction of demand peaks to 40% of the original schedule, the sum of deviation from the desired pattern to 60% of the original schedule, and the reduction of demand variability to 50% of the original schedules.

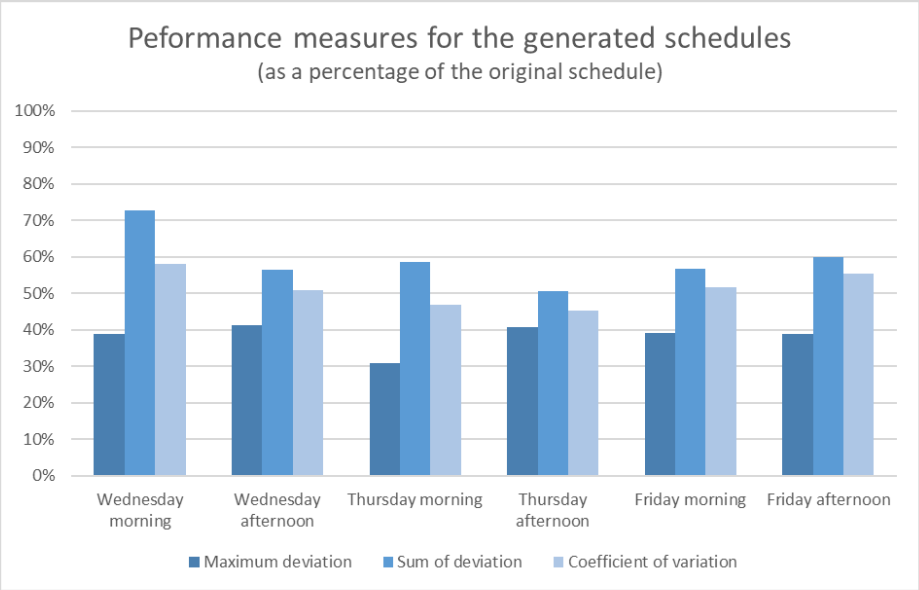

Figure M.1 – The achieved reduction of demand peaks and variability for 6 schedules (lower is better)

The experiments are performed on six blueprint session schedules that have been used in the hospital. The original schedule is entered into the model it calculates the expected workload pattern and the objective value. The model then schedules the same consultations, but it is free to adjust the sequence and starting time of the consultations. It attempts to match the desired pattern as close as possible.

Since the model incorporates a norm, its application becomes flexible. It is possible to achieve a very even spread of patient flow throughout the day, but other scenarios where less capacity is available at certain times of day the model can adjust the outpatient schedules to match.

The model shows promising results. Some further validation is needed to check if all constraints on schedules are present in the model. The main hurdle for complete implementation is the interface with schedule software HiX. At this time automatic import of the schedules is not possible, which means manual copying is needed. The recommendation is to generate certain standard blueprint schedules using the model and at the same time consult with HiX to hopefully allow automatic imports in the future. By using the new blueprint session schedules patient waiting time, staff workload fluctuations and perhaps staff overtime can be reduced.

Other articles have implemented methods that block complete parts of the day, or specifically book appointments for a patient. Our method for blueprint scheduling during the day can also be of interest for other researchers that seek a way to manage patient flow patterns during the day, with many adjustments possible in the model through the use of the norm.

Preface

I want to thank the Sint Maartenskliniek and Rhythm for the interesting research subject and all the support during the project. It was a pleasure to work with you and see integral capacity management in action in a hospital. It was a pleasure to be able to discuss the content and approach of the project with people of similar background during more in-depth meetings. I am especially grateful for the extensive collaboration and support by Mijke during our weekly meetings in Nijmegen. Your advice was really helpful throughout the entire project. I also want to thank Rob for the meetings and advice during the more challenging stages of the project.

I also want to thank the staff of the Sint Maartenskliniek that set aside time for interviews which were essential for this research.

Furthermore, I would like to thank my first supervisor Gréanne for the discussions at the university even during busy periods, and the solid advice, and my second supervisor Erwin for the final meetings.

I enjoyed the research of optimising patient blueprint schedules, and I hope you – the reader – will enjoy it as well.

Table of contents

Management summary i

Preface iii

1 Introduction 1

1.1 Organisation and departments involved 1

1.2 The orthopaedic care chain 1

1.3 Overview of planning hierarchy 2

1.4 Impact of the blueprint session schedule 2

1.5 Problem statement 3

1.6 Research goal 4

1.7 Research questions 4

2 The current situation 6

2.1 Current blueprint session schedule design 6

2.2 Key performance indicators 7

2.3 Current performance 9

2.4 Downstream demand characteristics of consultations 13

2.4.1 Transition probability 13

2.4.2 Demand duration 14

2.4.3 Travel time 15

2.4.4 Combining characteristics into the downstream demand profile 15

2.5 Conclusion 16

3 Literature review 17

3.1 Positioning our research question in literature 17

3.2 Multi-appointment planning 17

3.3 Conclusions 19

4 Model 20

4.1 Goal 20

4.2 Assumptions 20

4.3 Linear programming model definition 21

4.3.1 Sets and elements 21

4.3.2 Input parameters 21

4.3.3 Decision variables 22

4.3.4 Auxiliary decision variables 22

4.3.5 Constraints 22

4.3.6 Objective function 25

4.3.7 Model choices commentary 25

4.4 Model implementation in software 29
4.5 Model verification and validation 29
4.6 Model extension options 29
4.6.1 5-minute time slots 30
4.6.2 Downstream workload multiple time slots offset 30
4.6.3 Create multiple blueprint schedules simultaneously 31
4.6.4 Specify units, with different appointment characteristics per unit 32
4.6.5 Model DUO and TRIO sessions 33
4.6.6 Synchronised consultation starts across schedules 34
4.7 Conclusion 35
5 Experiments and results 36
5.1 Experiment design 36
5.2 Parameter settings 36
5.3 Algorithm running time 37
5.4 Experiments 38
5.5 Monte Carlo simulation 41
5.6 Conclusion 43
6 Implementation 44
6.1 Additional validation 44
6.2 User interface 44
6.3 Organisational support 45
6.4 Model in use 45
6.5 Conclusion 46
7 Discussion 47
7.1 Conclusions 47
7.2 Limitations 48
7.3 Future research 49
7.3.1 Stochastic version of the model 49
7.3.2 Using the model at the operational level 50
7.4 Recommendations 50
Bibliography 51
Appendix 53
A. Appendix data set columns 53
B. Appendix data analysis description 54
i. Location filtering 54
ii. Department filtering 54

iii. Unit classification 55
iv. Consultation type classification 56
v. Timestamp reliability 57
vi. Completed activities 58
vii. Identify source appointments 58
viii. Determine the linked appointments from each source appointment 59
ix. Downstream demand characteristics 59
x. Determining transition percentage before and after the consultation 60
xi. Travel time 61
xii. Downstream activity duration 62
xiii. Combining results into the final demand profiles 63
C. Appendix waiting time data quality check 64
D. Appendix literature review papers 65
E. Appendix appointment slot optimization papers 68
F. Appendix sliding window 69
G. Appendix base version of MILP model 70
H. Appendix sets and parameters for Thursday afternoon case 73
I. Appendix experiment results 75
i. Performance difference in minutes 75
ii. Complete overview of measures for the generated schedules 76
J. Appendix Monte Carlo graphs 76
i. Thursday afternoon – OOD 76
ii. Thursday afternoon – RAD 77
iii. Thursday afternoon – Plaster room 77
iv. Thursday afternoon – Screening 78
K. Appendix Monte Carlo script in R 78

1 Introduction

The Sint Maartenskliniek (SMK) is a hospital group in the Netherlands specialized in movement and posture. It is a specialised hospital performing complex procedures. The hospital aims to make its organisation and processes as excellent as the medical care.

Currently the flow of patients across departments is fluctuating during the day. These fluctuations result in additional waiting time for patients, uneven workload for staff, and staff overtime. The hospital wants to control patient flows in the orthopaedic care chain and to reduce these issues.

Section 1.1 introduces the hospital and relevant departments. In Section 1.2 the orthopaedic care chain is discussed. Section 1.3 describes the planning hierarchy, and Section 1.4 the blueprint session schedule, which is what this research will focus on. Section 1.5 discusses the problem statement, Section 1.6 the research goal, and Section 1.7 the research questions.

1.1 Organisation and departments involved

The SMK currently has five locations in the Netherlands, of which Nijmegen is the largest. The hospital in Nijmegen contains four medical specialties: orthopaedics, rheumatology, rehabilitation, and pain treatment. In 2017 over 44.000 unique patients were treated across all locations, and 1.100 FTE was employed ('Sint Maartenskliniek | cijfers jaarverslag 2017', n.d.).

The healthcare logistics department 'Logistiek Bedrijf' (LB) advises and supports departments in the SMK in decision making regarding logistics decisions. Their goal is to optimise the planning and control decision making for the hospital as a whole, accounting for impact of one department to other departments – i.e. over entire care chains. This research concerns the planning and control of one specific care chain, the orthopaedic care chain.

1.2 The orthopaedic care chain

The orthopaedic care chain consists of all steps a patient might take before, during, and after an orthopaedic intervention, such as a consultation at the outpatient clinic, diagnostics at Radiology, surgery, and inpatient care. An important step in this process is the outpatient consultation. New patients arrive at the outpatient clinic for a consultation, might get referred to other departments, for example to the 'Plaster and wound treatment' or 'Radiology' department on the same day on a walk-in basis. During their treatment patients return to the outpatient clinic multiple times for consultations. Figure 1.1 shows the possible patient flows for elective orthopaedic patients. The focus of this research is the orthopaedic outpatient department and the directly related departments.

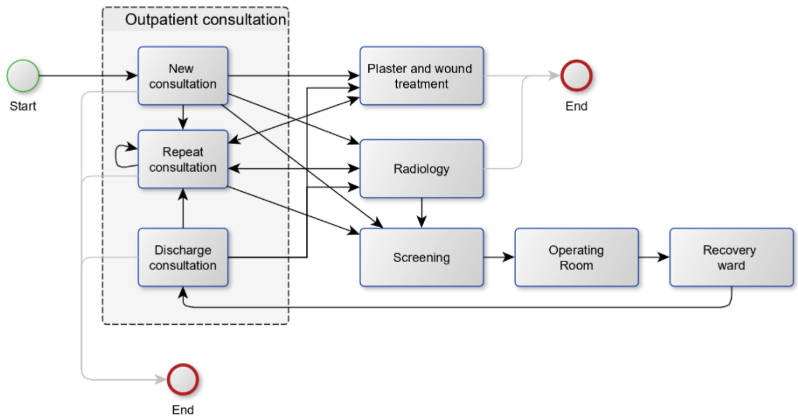

Figure 1.1 – Overview of patient flows for elective orthopaedic patients

1.3 Overview of planning hierarchy

The planning process for the orthopaedic outpatient department consists of the following steps:

Strategic

The **production outline for the entire year** is derived from the strategic goals of the hospital and the desired case mix. It describes for each type of activity (e.g. surgery, consultation) the target of how many will be performed in the entire year.

Tactical

Based on that, the **production plan per week** is derived. It specifies which activities take place in specific weeks.

Each day of the week is divided into a morning and afternoon session. The **staff schedule per session** is determined from the weekly plan. It specifies which doctors work in a session.

The **blueprint session schedule** determines in what timeslots the consultations start, and it specifies what type of consultation can be scheduled into each slot. For example, in the first two timeslots only new consultations can be scheduled, not repeat or discharge consultations. Per doctor currently different *blueprint session schedules* exists.

Operational

Every two weeks the *staff schedule per session* for the orthopaedic outpatient department is created for 12 weeks ahead. This specifies for each morning and afternoon if a doctor is scheduled in the operating room or for outpatient consultations. If a doctor is scheduled for outpatient consultations, **the appropriate blueprint session schedule is applied**, for example an afternoon session for unit 'Knee' for this doctor.

Once a *blueprint session schedule* is applied, patients can be booked in the slots in the way the schedule dictates. This means a new patient can be booked in the slots reserved for new patients, repeat consultations in the 'repeat' slots, and discharge consults in the 'discharge' slots.

The scope of this research is the design of the *blueprint session schedule*.

1.4 Impact of the blueprint session schedule

The sequence of consultation types dictated by the *blueprint session schedule* has an impact on other departments, because patients might visit these departments before and/or after their consultation at the orthopaedic outpatient department (see Figure 1.2 for an example).

Figure 1.2 – Example of demand for other departments before/after a consultation at the OOD

We expect that the *probability* a patient needs to visit a downstream department is different per consultation type (e.g. 'new' or 'repeat'). The downstream *activity duration* is expected to be different per consultation type. Furthermore, whether activities at other departments happen more often *before* or *after* the main consultation is expected to be different.

This means the timing and amount of workload at the downstream departments can likely be changed by changing the sequence of consultation types in the *blueprint session schedule*.

1.5 Problem statement

Currently several issues are experienced at the orthopaedic outpatient department and the related (downstream) departments in the orthopaedic care chain. Figure 1.3 displays the causal relationships between the experienced problems.

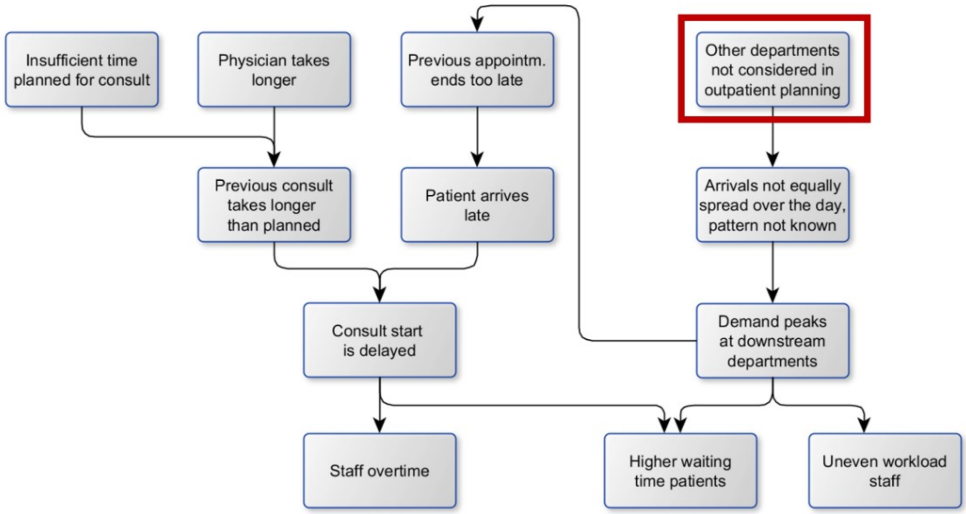

Figure 1.3 – Problem cluster indicating current problems and the causal relationships

Staff indicate consultations starts are often delayed, because the previous consultation takes longer, or the patient arrives late. These delays cause patients to wait for their appointment to start. The delays can remain or increase throughout the day, meaning staff can run into overtime in order to serve all the patients scheduled.

One of the reasons consultations take longer than planned is because insufficient time is scheduled for all these consultation types, or certain physicians takes longer than scheduled. Currently all consultation types have the same planned duration (15 minutes), however in interviews staff often noted that more time is needed for 'new' consultations. Another reason consultations take longer than planned is certain physicians typically take more time for their consultations as compared to their colleagues.

A reason the patient arrives late is when their previous appointment at another department ends later than planned. Other departments experience peaks in demand during certain times of the day. When these demand peaks occur and more patients arrive than can be served simultaneously, patient waiting time increases. The staff workload also increases during these times. If patients are spread more evenly over the day, patients would experience minimal waiting time, and staff experience minimal workload fluctuations.

Currently the outflow pattern of patients to other departments is not considered when scheduling patients at the orthopaedic outpatient department (the problem highlighted in red in Figure 1.3), the design of the blueprint session schedule does not take these same-day patient flows into account.

Core problem

We choose to focus on the right side of the problem cluster, the core problem that outflow of patients to other departments is not considered in the orthopaedic outpatient planning. By improving the *blueprint session schedule* the other problems listed in the problem cluster will likely decrease.

1.6 Research goal

The goal of this research is to investigate how the orthopaedic outpatient department patient planning can be improved, in order to achieve a more even spread of patient outflow to other departments during the day, reducing demand peaks and reducing problems related to the demand peaks.

We expect the design of the *blueprint session schedule* can improved, to evenly spread same-day patient outflow to other departments and thereby reduce delays, reduce patient waiting time, staff overtime and uneven workload. In order to have optimal *blueprint session schedules* for the many different situations, a method to automatically generate the schedules is desired.

An additional hope for Logistiek Bedrijf is to find an approach or a model that is generic, so it can also be used in other situations or hospitals.

The scope of this research is limited to the same-day patient flow effects of the *blueprint session schedule* on all directly related departments.

1.7 Research questions

The main research question is:

How can blueprint session schedules be generated for elective patients at the orthopaedic outpatient department, that match the expected pattern of same-day patient flows with the desired pattern?

The research question is answered with the following sub questions:

1. How does the blueprint session schedule currently work and perform?

How are the current blueprint session schedule designed, and how do they perform? Which measures are most appropriate to assess performance? Which factors (e.g. patient no-shows) are most relevant to this research? How does the Orthopaedic Outpatient Department relate to other departments? This is discussed in chapter 2.

2. What options for forming the blueprint session schedule exist?

What alternative methods for forming the blueprint session schedules exist? What is mentioned in other research on this subject? Which alternative method is most promising for the Sint Maartenskliniek? This is discussed in chapter 3.

3. How can the selected method for blueprint session schedule design be modelled and validated?

How can the most promising method be modelled? How do we check the model works as we expect it to and deliver valid results? What are the assumptions that the model is based on? What is the model capable of? This is discussed in chapter 4.

4. How does the performance of the model compare to the current performance?

How does the performance of the new method of creating blueprint session schedules compare to the current way of working? Does it improve the performance in a case study with several blueprint session schedules that have been used in practice? This is discussed in chapter 5.

5. How can this method be implemented at the Orthopaedic Outpatient Department?

How could the new method for creating blueprint session schedules be implemented at the Orthopaedic Outpatient Department? What steps need to be taken? This is discussed in chapter 6.

Finally chapter 7 will conclude with discussion on the strengths and the limitations of this research, and opportunities for future research.

2 The current situation

This chapter will discuss the research question "How does the blueprint session schedule currently work and perform?". Section 2.1 describes the current blueprint session schedule design. Section 2.2 discusses on what measures the blueprint session schedule will be evaluated. Section 2.3 discusses the current performance, and Section 0 discusses the downstream demand profiles of consultation types. Section 2.5 concludes this chapter.

2.1 Current blueprint session schedule design

The blueprint session schedule dictates for each slot what type of consultation can be booked into it, the sequence of these slots, and the appointment duration in minutes. Figure 2.1 is an example of a blueprint session schedule for the orthopaedic outpatient department.

Session	Unit	Doctor name	Sequence	Duration	Consultation type
Afternoon	Knee	Dr. Janssen	1	15	VCKN
Afternoon	Knee	Dr. Janssen	2	15	VCKN
Afternoon	Knee	Dr. Janssen	3	15	NKKN
Afternoon	Knee	Dr. Janssen	4	15	OCKN
Afternoon	Knee	Dr. Janssen	5	15	OCKN
Afternoon	Knee	Dr. Janssen	6	15	NKKN
Afternoon	Knee	Dr. Janssen	7	15	NKKN
Afternoon	Knee	Dr. Janssen	8	15	VCKN
Afternoon	Knee	Dr. Janssen	9	15	VCKN
Afternoon	Knee	Dr. Janssen	10	15	NKKN
Afternoon	Knee	Dr. Janssen	11	15	VCKN
Afternoon	Knee	Dr. Janssen	12	15	VCKN

Figure 2.1 – Example of a blueprint session schedule

In each slot one of the following four consultation types is specified:

- NK – "Nieuwe klacht": A consultation for a patient experiencing a new medical issue
- VC – "Vervolgconsult": A repeat consultation, the patient has been in the hospital before for this issue
- OC – "Ontslagconsult": A discharge consultation after surgery to check if the patient can be discharged from the hospital
- POP – "Preoperatief gesprek": A preoperative meeting between the patient and the orthopaedic doctor that will perform their surgery

Blueprint session schedules are currently different between units, and are different between doctors within the same unit. Figure 2.2 shows all blueprint session schedules belonging to just one doctor.

ORTH	KN MI 1	Knie Spreekuur dreamclinic middag UNO	13:00	16:15
ORTH	KN MI 2	Knie Spreekuur dreamclinic middag DUO	13:00	16:15
ORTH	KN MI 2+	Knie Spreekuur dreamclinic middag DUO met CDC	13:00	16:15
ORTH	KN MI 3	Knie Spreekuur dreamclinic middag TRIO	13:00	16:15
ORTH	KN MI 3+	Knie Spreekuur dreamclinic middag TRIO met CDC	13:00	16:15
ORTH	KN MO 1	Knie Spreekuur dreamclinic morgen UNO	08:30	12:20
ORTH	KN MO 2	Knie Spreekuur dreamclinic morgen DUO	08:30	12:20
ORTH	KN MO 2+	Knie Spreekuur dreamclinic morgen DUO met CDC	08:30	12:20
ORTH	KN MO 3	Knie Spreekuur dreamclinic morgen TRIO	08:30	12:05
ORTH	KN MO 3+	Knie Spreekuur dreamclinic morgen TRIO met CDC	08:30	12:20

Figure 2.2 – Blueprint session schedules belonging to one doctor (each row is one blueprint session schedule)

The morning shifts run approximately from 08:30 to 12:00, and the afternoon shifts typically run from 13:00 to 16:15. However, many different start times and end times currently exist between blueprint session schedules. Some start at 08:00, other schedules at 09:00, run until 11:45, 12:20 or 12:30. The afternoon blueprint session schedule can start at 12:45, 12:55 or 13:00, and last until 15:45, 16:00, 16:20 or 16:30.

A different blueprint session schedule is available per unit and UNO, DUO and TRIO versions exist. The UNO schedule is for a doctor without support, DUO for one additional employee (such as a doctor in training), and TRIO for a doctor with two support employees. The schedules with more support can see more patients, but some time for supervision is also needed. See Figure 2.3 for an example of a TRIO blueprint session schedule.

Session	Unit	Doctor name	Sequence	Simultaneous	Duration	Consultation type
Afternoon TRIO	Spine	Dr. Janssen	1	1	5	POSC
Afternoon TRIO	Spine	Dr. Janssen	2	3	15	NKWE
Afternoon TRIO	Spine	Dr. Janssen	3	1	5	-
Afternoon TRIO	Spine	Dr. Janssen	4	1	5	POSC
Afternoon TRIO	Spine	Dr. Janssen	5	2	15	NKWE
Afternoon TRIO	Spine	Dr. Janssen	6	1	5	-
Afternoon TRIO	Spine	Dr. Janssen	7	3	15	VCWE
Afternoon TRIO	Spine	Dr. Janssen	8	2	15	NKWE
Afternoon TRIO	Spine	Dr. Janssen	9	1	5	-
Afternoon TRIO	Spine	Dr. Janssen	10	3	15	VCWE
Afternoon TRIO	Spine	Dr. Janssen	11	1	5	POSC
Afternoon TRIO	Spine	Dr. Janssen	12	3	15	NKWE
Afternoon TRIO	Spine	Dr. Janssen	13	1	5	-
Afternoon TRIO	Spine	Dr. Janssen	14	3	15	NKWE
Afternoon TRIO	Spine	Dr. Janssen	15	3	15	VCWE
Afternoon TRIO	Spine	Dr. Janssen	16	3	15	OCWE
Afternoon TRIO	Spine	Dr. Janssen	17	2	15	OCWE

Figure 2.3 – A TRIO blueprint session schedule

The current blueprint session schedule designs do not consider the pattern of patient flow to other departments. The current blueprint session schedule templates do not change, when in some cases it might be beneficial to change the ratios of NK, VC and OC, or adapt to the combination of doctors in a session.

2.2 Key performance indicators

To evaluate the current blueprint session schedule and compare it with a different design, performance measures need to be chosen. We need to know what needs to be optimised. The research goal is to achieve an even spread of patient flow throughout the day, i.e. reducing the variability in minutes of workload that arrives at downstream departments. By reducing this variability the aim is to reduce the waiting time for patients, workload fluctuations of staff, and staff overtime.

For the measurement of workload, the number of patients or the minutes of workload they introduce can be measured. Since the time needed to help two different patients can be vastly different, the most appropriate measure is the minutes of workload they add to a department.

These research goals translate into the following key performance indicators:

Workload arrival variability

The blueprint session schedule is expected to influence how much workload is added to other departments, and at what time the patients arrive. This indicator measures how much the workload fluctuates during the day. Ideally the workload is evenly spread over the day. This also represents the workload fluctuations of staff.

Perspective: Staff
Definition: (minutes of workload arrived in interval – average workload per interval for this day)
Aggregate: Per time interval (e.g. 1 hour)

Average waiting time for an appointment

This measures the average waiting time for an appointment per time interval, to check trends of waiting time during the day. For the

Perspective: Patient
Definition: Start appointment – MAX(planned start time; actual arrival time of patient)
Aggregate: Per time interval (e.g. 1 hour), based on start time appointment

Overtime

This indicator measures how much overtime occurs per department. Ideally no overtime would occur.

Perspective: Staff
Definition: MAX(0; Actual end time of shift – closing time of department)
Aggregate: Per day

Additionally, because the initial idea is to control the workload arrival at other departments based on changing the sequence of consultation types in the blueprint session schedule, it is important to know if the expected workload is different per consultation types. If all consultation types have very similar outflow, changing the sequence of consultations will not impact the workload arrival pattern at downstream departments.

The three components that determine the workload pattern of a consultation type are:

- The **transition probability** to a downstream department per consultation type
- The **minutes of workload** one patient introduces at a downstream department (per consultation type)
- The **timing of the workload**, does the demand at downstream departments occur before or after the consultation at the outpatient department? Is there a different delay before the demand arrives for certain consultation types?

2.3 Current performance

The current performance of the blueprint session schedules is analysed with a data set containing two years of data on patient outpatient appointments. It contains the department of the activity, anonymised patient ids, the activity type (e.g. a 'repeat consultation'), and timestamps related to the patient arrival, activity start time, and activity end time. See Appendix A for the full list of columns.

Figure 2.4 gives an overview of the data set, and the number of appointments per department in Nijmegen. From the relevant departments the number of completed and cancelled appointments are counted to give an idea of the proportions of each department.

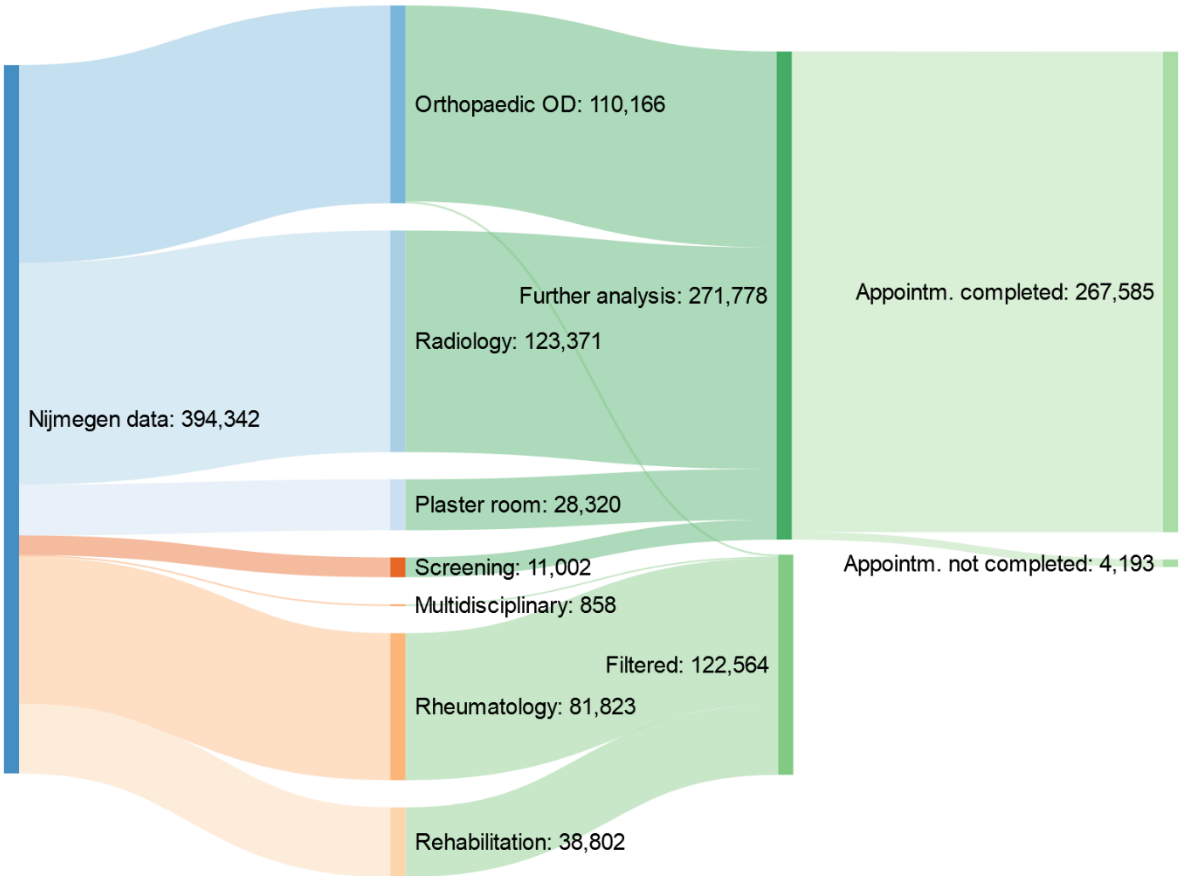

Figure 2.4 – Number of appointments per department over 2017 and 2018 (Sint Maartenskliniek)

Workload arrival variability

The patients that have an outpatient consultation at the Orthopaedic Outpatient Department (OOD) sometimes need to visit other departments on the same day as their outpatient consultation. We assume (based on interviews) that all visits to other departments by the same patient on the same day can be linked to the original consultation at the OOD.

Appendix B describes the data analysis steps in detail. Appendix sections B.vii and B.viii describe how appointments at other departments are linked to the main consultation at the OOD.

It is reasonable to assume the schedule of the OOD will influence the pattern of arriving patients at the downstream departments. To ensure we only analyse the arriving patients that originate from the OOD, we select only the activities that have been 'linked'.

For each individual day the following calculations are made:

1. **Sum of workload arrival per time interval**: Each patient that arrives at a department represents a certain amount of work, expressed in minutes of workload. At each time interval of 1 hour, the sum of workload from 'linked' activities that arrives is counted.
2. **Average workload arrival**: The overall average of workload arrivals of linked activities is calculated for that day:

$$\frac{total\ workload\ arrival \mid activity = linked}{number\ of\ time\ intervals} = average\ workload\ arrivals\ per\ time\ interval$$

3. **The difference from the average per time interval**: Given the average workload that arrives in any time interval, the difference between the average and the actual workload arrival can be calculated per time interval on that day:

$$sum\ of\ workload\ in\ interval - average\ workload = difference\ from\ average$$

Now the sum of differences is aggregated over all days. This method of calculation shows the pattern of patient arrivals throughout the day.

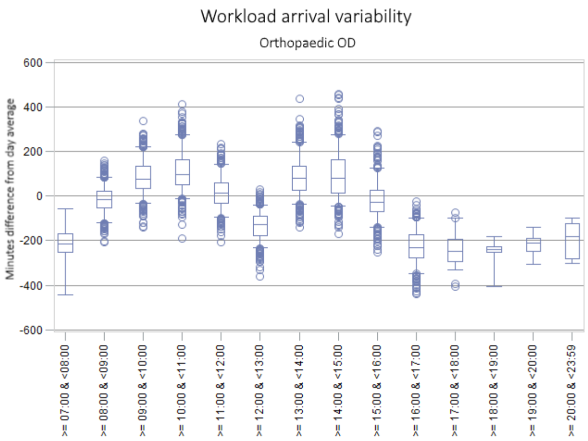

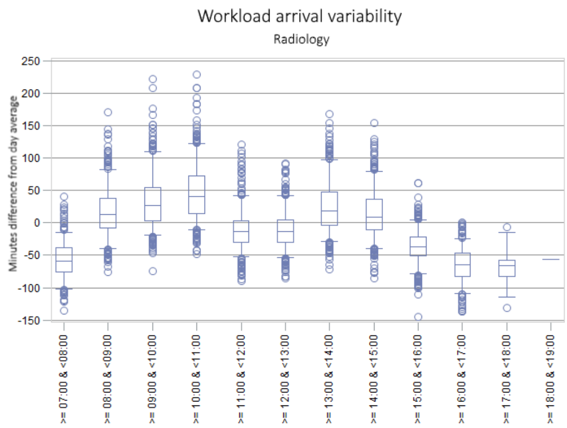

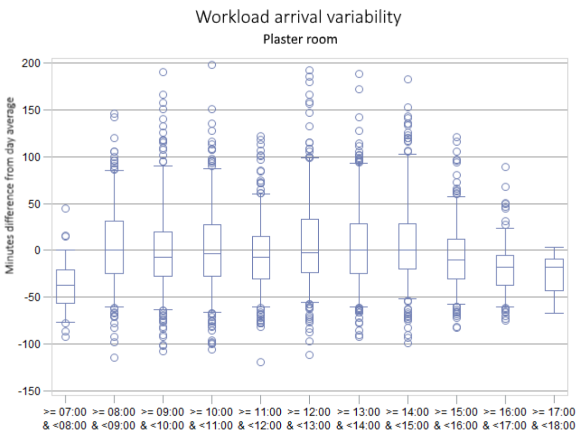

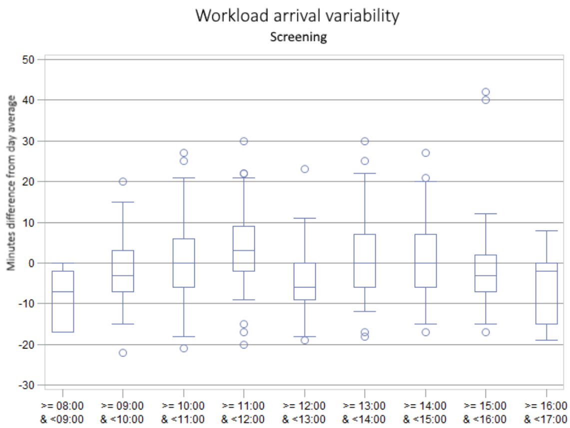

These graphs show workload arrival distributions for each department. It shows how much the arrivals (measured in minutes of workload) in that hour are different from the mean arrivals on the same day. The box plot displays the 5th, 25th, 50th, 75th, and 95th percentile. Especially for the Orthopaedic outpatient department and the Radiology a clear trend over the day is visible. For the Plaster room no standard trend is visible. For the preoperative screening the pattern seems to be stable over the day.

Overtime

The overtime is calculated by determining for each day the time the last linked patient left. If this time is before closing time of the department, the overtime on that day is 0 minutes. Otherwise the overtime of that day is the amount of time the patient left after regular opening hours of the department.

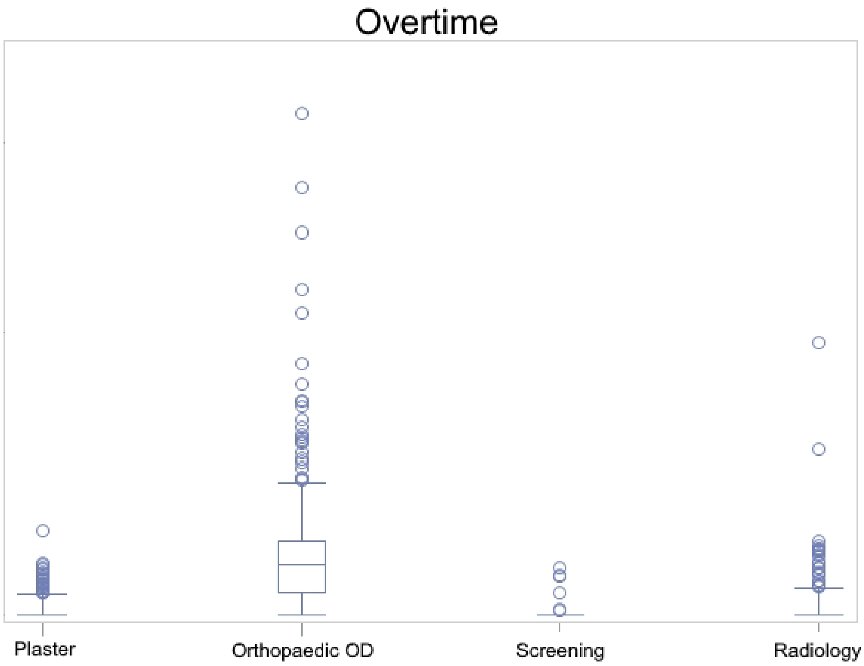

Figure 2.5 – Box plot of overtime per department (whiskers are 5th and 95th percentile)

Figure 2.5 shows the distribution of overtime for the four downstream departments. The results suggest this is quite an common occurrence. Especially at the orthopaedic outpatient department itself the overtime stands out. A caveat with this calculation is we assumed closing hours of 16:15, where in some cases the session schedules exceed 16:15. However 16:15 closing time was agreed in accordance with the department to be a suitable closing time for the analysis, and overtime for certain doctors was repeatedly mentioned in interviews. Based on the box plot a 25% of the days seem to run at significantly past 16:15 which will definitely be overtime. Overall linked patients arriving late seems to mostly affect the Orthopaedic Outpatient Department.

Waiting time per appointment

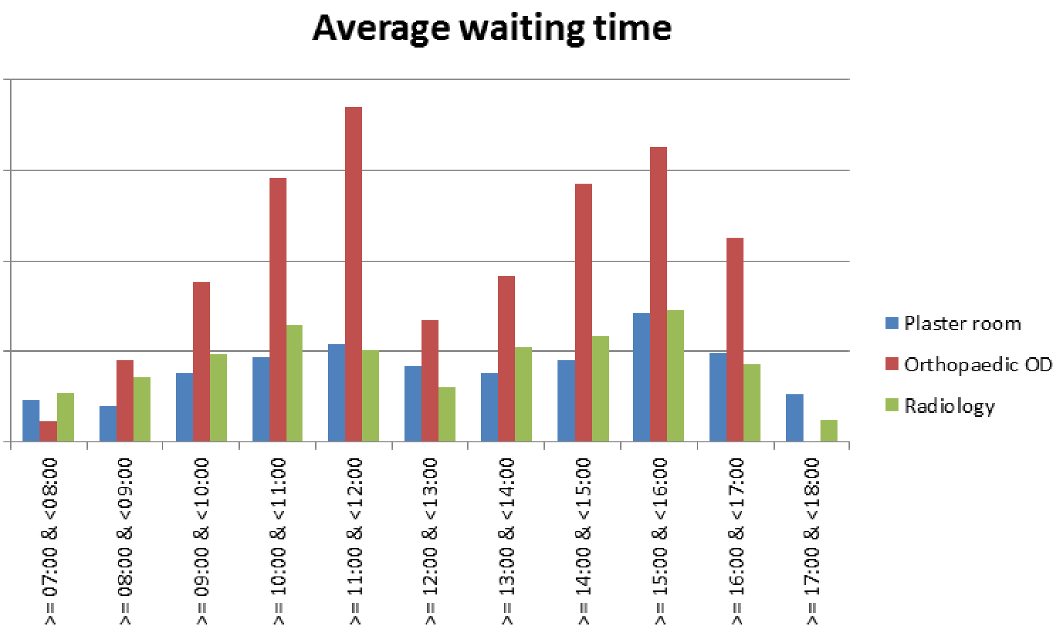

Figure 2.6 – Average waiting time per time interval for each department

Figure 2.6 shows the average waiting time per time interval, for each department except for the pre-operative screening department, they as waiting time is not recorded separately for the linked appointments in my data set. Most notable is the large increase in waiting time at the Orthopaedic OD during the day. The data has been cleaned so incorrect or contaminated data is not influencing these results, this is discussed further in Appendix C, but the data is cleaned. The same pattern is also visible in the median waiting time per hour, meaning the pattern cannot be explained by outliers.

A likely explanation is currently revisits to the OOD are not scheduled in, but rather the patient is called in when the doctor has downtime. This issue does not apply to the other departments, since far more is done on walk-in basis, thus they are not waiting to fit into a busy pre-booked schedule.

2.4 Downstream demand characteristics of consultations

The idea for the research project is to modify the pattern of patient flow that results from the blueprint session schedule, and the sequence of consultation types it defines. A core assumption that needs to be tested is if the different consultation types cause different patterns in demand at downstream departments.

As discussed in Section 2.2 the three components that will combine into the demand profile for a consultation type are:

- The transition probability
- The downstream demand duration (workload in minutes)
- The travel time or demand timing

The following sections will discuss the method of calculation, and finally combine the parts into the demand pattern per consultation type.

2.4.1 Transition probability

Source consultation

A patient might need to visit a different department before or after the consultation at the OOD. We identify for each patient the first consultation of that day at the OOD, this is considered the 'source consultation'. All other outpatient demand of that patient on the same day is assumed to be caused by the source consultation. See Appendix B.vii for a detailed description of identifying source consultations.

Link downstream demand to source consultations

We identify all other appointments and activities at the downstream departments (follow-up consultations at the OOD itself, visits to Radiology, Plaster room, or Preoperative screening), and we link them to the source consultation. The exact method for linking these activities is described in Appendix B.viii.

The intermediate results show that indeed the transition probabilities are different between consultation types and between units. The transition probabilities are different between the consultation types, as well as different between the units.

Determine whether demand occurred before or after the source consultation
Only a general transition probability is not specific enough, it is important to know if the downstream activity occurred before or after the source consultation.

Based on the linked appointments, we determine based on the timestamps if the activity at the downstream department occurred before or after the source consultation. If a patient gets a scan at Radiology from 10:00 until 10:10, and has their consultation at 10:30, it is clear the downstream demand at Radiology occurred before the source consultation. We label each transition as either ‘before’, ‘after’, or in case the timestamps are missing as ‘missing’.

Calculate transition probabilities (for each consultation type)
Now that for each source consultation the transitions (linked activities) are identified, and they are labelled as ‘before’, ‘after’, or ‘missing’ the transition probabilities for all consultation types can be calculated.

As an example, with 1000 source appointments, 130 transitions might be labelled as ‘before’, and 180 transitions might be labelled as ‘after’. Then the transition probabilities are 13% and 18% to before and after this source appointment.

However, sometimes the timestamps are missing. These linked appointments cannot be ignored, because the exact timing might be unknown, but it is certain the transition occurred. If the transitions with missing timestamps are ignored, the transition probability is underestimated.

We correct for the missing timestamps in the following way:

1. Count all transitions ‘before’, ‘after’, and with missing timestamps
2. Determine the proportion of ‘before’ and ‘after’
3. Add the transitions with missing timestamps to the ‘before’ and ‘after’ categories, while preserving the proportions calculated at step 2

After these steps we have successfully calculated the probability a patients will visit a downstream department before their consultation, and the probability they visit a downstream department afterwards.

Calculate transition probabilities (for each combination of consultation type and unit)
The previous calculations just considered the transition probabilities for each consultation type. However the model could be extended to simultaneously consider the consultation type and the unit. This involves more complex analysis, and since the unit is not always classified in the data, a smart way of dealing with those missing unit classifications is also required. Appendix B.x discusses how that is calculated.

2.4.2 Demand duration

The amount of time needed for a patient at a downstream department is also expected to be different per source consultation type and/or per unit. So if a patient comes in for a ‘New’ consultation, the amount of time needed at Radiology might be more than if a ‘repeat’ patient visits Radiology after their consultation.

We use the same basis of source consultations and linked appointments as described in the previous section, however instead of checking *if* a transition occurred, we check *how much time* the patient needed at the downstream department, given that a transition did occur. We filter out the unreliable timestamps in the way described in Appendix B.v.

First we perform an initial check to see if the downstream demand duration is indeed different per source consultation type and per unit. Again, like with the transition probabilities the demand duration is different between the different consultation types and units.

The differences between consultation types and units are present. We perform the more elaborate and reliable calculation, where the duration is calculated for each set of consultation type and unit, before and after separately. If too few observations are available a reasonable measure (such as the average of the unit) is used. The detailed description for the entire duration analysis is discussed in Appendix B.xii.

2.4.3 Travel time

When a patient completes a consultation and needs to visit a different department, the patient will not arrive their immediately, which means the workload the patient represents also does not arrive immediately. The most fundamental reason for this is simply travel time, the patient needs time to walk from one department to the next.

The travel time can be determined by analysing the timestamps of two appointments by the same patient on the same day. If the patient has their consultation at the OOD from 10:00 until 10:15, and arrives at the waiting room of the Plaster room at 10:20, the travel time was 5 minutes. In order to limit the effect of outliers in this analysis the median travel time is used. The calculation of travel time is discussed in detail in Appendix B.xi.

2.4.4 Combining characteristics into the downstream demand profile

The three characteristics *transition probabilities*, *demand durations*, and *travel time* are combined into the overall downstream demand profile of a consultation type and unit combination. This defines the amount of workload in minutes that can be expected, and the timing of that workload. Table 2.1 shows an fictitious example of a demand profile. Appendix B.xiii describes the calculations in detail.

	Expected workload - example demand profile															
	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
New																
OOD	0.00	0.00	0.00	3.22	3.22	1.29	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
RAD	0.00	0.00	0.56	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Plaster	0.00	1.31	1.31	1.31	1.05	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
PREO	0.00	0.00	0.00	4.51	4.51	4.51	4.51	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Repeat																
OOD	0.00	3.94	2.36	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
RAD	0.00	0.80	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Plaster	0.00	4.27	4.27	4.27	3.41	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
PREO	0.00	0.00	4.74	4.74	4.74	4.74	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Discharge																
OOD	0.37	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
RAD	0.00	0.00	0.00	0.00	0.56	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Plaster	0.00	0.00	0.00	1.90	1.90	1.90	1.90	1.90	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
PREO	0.00	0.00	0.00	0.14	0.14	0.14	0.14	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00

Table 2.1 – Demand profile for the expected workload after a source consultation

2.5 Conclusion

The data analysis shows distinct transition probabilities and resulting minutes of downstream workload between consultation types and units, which means changing the sequence can impact the workload arrival pattern at downstream departments.

For downstream workload arrival variability, the peaks at the Radiology department and follow-up consultations at the orthopaedic outpatient department show a trend during the day. For the Plaster room no strong trend during the day seems to be there, but the arrivals are still variable.

The waiting time shows an upward trend in the morning for all departments, but the Orthopaedic Outpatient department shows the strongest increase. Some of the solutions are not directly in scope of this project, such as adjusting the time per consultation or reserving time for patient revisits, but it is a good idea to accommodate these in a our proposed solution.

Considering in the new situation blueprint session schedules will be standardized per unit, it is important to look at how to achieve an even spread of workload arrivals over the day.

3 Literature review

This chapter will discuss existing research in this field. Section 3.1 breaks down the research question into key components. Section 3.2 will look at individual papers related to the decisions that need to be made for this research to find best practices. Section 3.3 will conclude with lessons learned from the literature.

3.1 Positioning our research question in literature

The goal of this research is to develop a model that can generate blueprint schedules, where the sequence of appointment types in the schedule is optimized to minimize the workload variability in downstream departments.

Hulshof, Kortbeek, Boucherie, Hans, & Bakker (2012) created a taxonomic framework to classify healthcare planning decisions. The design of a blueprint appointment schedule is classified as a **tactical** problem, in the **ambulatory care services** category.

In terms of multi-department or multi-disciplinary planning, A distinction between *flow shop*, *open shop*, and a hybrid *mixed shop* is made by Leeftink, Bikker, Vliegen, & Boucherie (2018). In the **flow shop**, related jobs (or multiple appointments for the same patient) have a fixed sequence, which applies to this research.

In order to recognise literature that applies to this problem, even for problem contexts that are different, we list the key problem characteristics:

- **Single server (main department)** – Appointments at the main department can only be fulfilled by a specific doctor
- **Downstream departments (multi-server) –** Multiple downstream departments are involved, where patients do not need appointments made beforehand. No single server is
- **Uncertain care pathway** – It is uncertain if a patient needs a follow up at downstream departments
- **Flow shop** – The sequence of multiple appointments for a single patient on the same day is assumed to be fixed
- **No walk-in patients** – All acute patients go to the acute outpatient clinic
- **Capacity is constant** – During one day (the time horizon of this schedule) the staff levels are constant

3.2 Multi-appointment planning

Developing the blueprint session schedule while minimizing the downstream workload variability can be classified as a multi-appointment planning problem, where the sequence of certain appointment types is changed to improve the situation in downstream departments.

In the review by Ahmadi-Javid, Jalali, & Klassen (2017) the research into appointment slot optimization is classified. These articles consider the problem of a single stage system, so no downstream departments are considered. The overwhelming majority minimize the patient waiting time, idle time and overtime using either a 1-stage or 2-stage stochastic program. See Appendix E for the complete overview.

We are interested in papers that handle multi-appointment planning, in our case over multiple departments, and with the objective to spread the workload. We use the literature review 'Multi-disciplinary planning in health care: a review' Leeftink, Bikker, et al. (2018) to find relevant papers and describe their approach.

Appointment slot optimization & multi-disciplinary scheduling

Dharmadhikari & Zhang (2013) reserve timeslots in different departments at different times of day, to increase the probability patients with appointments at multiple departments can be scheduled on one day. Different number of timeslots can be blocked to deal with different levels of multi-department demand. This approach with mutually exclusive blocks is not applicable in our case, since consultations need be scheduled all day, and it is unknown beforehand which patients need to visit downstream departments.

Liang, Turkcan, Ceyhan, & Stuart (2015) develop a mixed-integer linear programming (MILP) model that evenly distributes patients into timeslots to balance the workload throughout the day on two departments. The care pathway is assumed to be known, patient types are created for each pathway (e.g. type 1: only department 1, type 2: visits department 1 and 2, type 3: only visits department 2). The objective function minimizes the difference between the maximum and minimum workload. The output of the MILP model is input for a simulation model to evaluate operational decisions and improve on patient waiting time and clinic overtime.

Leeftink, Boucherie, et al. (2018) develop a 3-stage batch scheduling algorithm. The batch completion times are spread evenly over the day to reduce workload variability in the downstream department. The minimum interval between batch completion times is maximized to ensure the most spread of workload. All jobs need processing in all stages, which is different from our problem where the care pathway is uncertain.

Leeftink, Vliegen, & Hans (2019) develop a two-stage stochastic integer program for patients that have an appointment with a nurse practitioner, and afterwards have an uncertain care pathway to one the relevant clinicians (which is unknown beforehand). The downstream clinician is analogous to the downstream department in our research topic. In addition to the multi-disciplinary patients, regular patients that only need a single appointment at a known clinician (downstream department) also need to be scheduled. In the first stage the optimal blueprint schedule is determined, in the second stage for different scenarios with patient arrival realisations the overtime, waiting time and idle time is minimized.

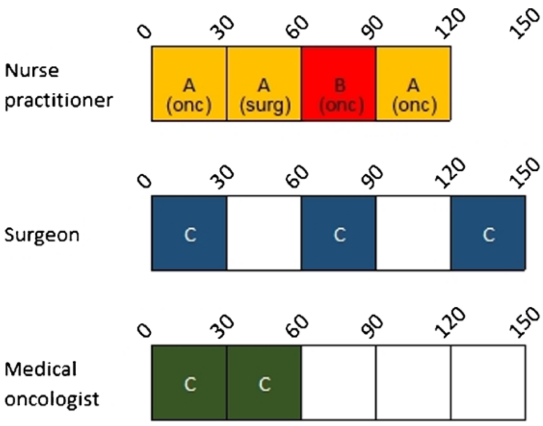

Figure 3.1 – Optimizing the patient sequence (nurse schedule) with downstream resources (Leeftink et al., 2019)

Their approach is to reserve a slot, where a patient might make use of with a certain probability. Our situation is different in the sense the downstream departments are more similar to multi-server queues, since no appointments are needed for the patients, and they can be helped by various staff members. Matching the expected workload to a norm, without booking appointments can be a suitable approach.

Surgery scheduling literature

In addition to the literature on appointment slot optimization or multi-disciplinary scheduling, the research into surgery schedules is closely related to our problem. Usually the sequence of surgery types need to be optimized, and a number of articles incorporate the downstream demand when optimizing the surgery type sequence. The time scale is different (usually one time slot equals one day), however the same approach can be useful in our blueprint schedule optimization problem.

Fügener, Hans, Kolisch, Kortbeek, & Vanberkel (2014) develop a master surgery schedule while considering the downstream demand, where the sequence of blocks is changed in order to minimize the downstream costs. The time scale is discrete, one time slot is one day. An exact multinomial distribution is derived to indicate the probability of x number of patients at a downstream resource at time t. Fixed costs are incurred by costs for beds, higher demand peaks mean more beds need to be available. In this way the model is incentivised to reduce demand peaks since this will lead to lower costs.

van Essen, Bosch, Hans, van Houdenhoven, & Hurink (2013) develop a model similar to Fügener et al. (2014), and evaluate two different approaches. The first approach is to use the exact multinomial distribution, and use a local search algorithm (Simulated Annealing) to approximate the global optimum solution. The second approach is to approximate the objective function by linearizing it in an MILP model. The MILP model returned better results solving the expected length of stay comprehensively instead of approaching the exact multinomial distribution with heuristics.

3.3 Conclusions

For appointment slot optimization problems where no downstream resources are considered, a 1-stage or 2-stage stochastic program is the most applied method to optimize the schedule. Typically a weighted sum of waiting time, idle time and overtime is minimized. These models are able to deal with stochastic service times, and alter the planned starting time and time between appointments in order to minimize the objective function.

For models that consider downstream departments various approaches exist. The article by Leeftink et al. (2019) models the problem with a two-stage stochastic programming approach to optimize the appointment sequence for the nurse (main department), considering an uncertain care pathway to clinicians (downstream departments). Our situation is different in the sense the downstream departments are more similar to multi-server queues, since no appointments are needed for the patients, and they can be helped by various staff members.

The literature for surgery schedule sequencing is similar with multiple downstream resources and stochastic length of stay (or appointment duration in our research). The article by Fügener et al. (2014) considers the exact multinomial distribution of length of stay, while van Essen et al. (2013) find that in their problem setting the expected length of stay outperformed the exact objective function with local search. This suggests that working with expected values rather than exact distributions might yield comparable or in some cases better results.

The research of Leeftink et al. (2019) and the linearized approach of van Essen et al. (2013) are quite similar to our approach, and both use an linear programming model. The possibility even exists to expand a MILP model into a Stochastic Linear Program to explicitly model uncertainty. This suggests using a MILP formulation of our problem to optimize the blueprint session schedules is a promising approach.

4 Model

In this chapter the mathematical model is defined. Section 4.1 describes the goal of the model. The assumptions we make are listed in Section 4.2. The model is defined in Section 4.3.

4.1 Goal

The goal of the model is to create a blueprint schedule that determines the optimal *sequence* and *starting time* of consultation types, where the *expected downstream workload* is as close to the *norm* per time slot as possible. Different consultation types can be specified, which have different expected workload at downstream departments.

For each downstream department a score is calculated based on how closely expected workload matches the norm. The final score of a blueprint schedule is the weighted sum of all downstream department scores.

Figure 4.1 shows a simplified example of a blueprint schedule. Figure 4.2 is a visual representation what the model tries to achieve at each downstream department.

Time slot	08:00-08:15	08:15-08:30	08:30-08:45	08:45-09:00
Consultation type	New	Repeat	New	Discharge

Figure 4.1 – Simplified example of a blueprint schedule

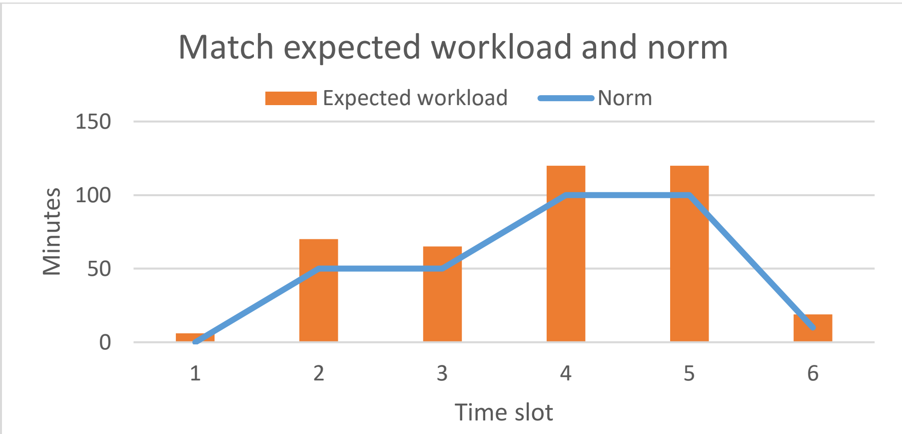

Figure 4.2 – Simplified representation of the goal of the model

4.2 Assumptions

The model has a scope of one day. The following assumptions are made in the model:

- Service times at the main department are assumed to be deterministic.
- The care pathway is uncertain. The transition probabilities to downstream departments are static, but are different for each appointment type.
- Patients and physicians are assumed to be punctual.
- One physician can only see one patient in a time slot (no overbooking)
- No walk-in patients in the OOD

4.3 Linear programming model definition

We model this optimization problem as a Mixed Integer Linear Programming (MILP) model. It consists of sets, parameters, variables, constraints and the objective function that are defined in this section.

4.3.1 Sets and elements

Set	Element	Description
C	c	Consultation types
D	d	Departments
S	s	Schedules
U	u	Units

Time sets:

Set	Element	Description
T = {1, …, numberOfSlots}	t, tt	Time slot
SW = {1, …, (numberOfSlots - *slidingWindowWidth* + 1)}	sw	Sliding time window. Used to aggregate workload of multiple time slots, similar to a moving average.
I	i	Time offsets, used to define a demand profile 1..i time slots before the start of the consultation, and 1..i time slots after the end of a consultation.

4.3.2 Input parameters

Parameter	Range	Description
$\boldsymbol{consultationDuration_{s,c,u}}$	$\in \mathbb{N} \setminus \{0\}$	The duration of a consultation of type *c* for unit *u* in schedule *s* in number of time slots
$\boldsymbol{consultationsToBeScheduled_{s,c,u}}$	$\in \mathbb{N}$	The number of consultation type *c* of unit *u* to be scheduled in schedule *s*
$\boldsymbol{norm_{d,t}}$	$[0, \infty)$	The capacity norm for total minutes of workload arriving at department *d*, per time slot t
$\boldsymbol{expectedWorkloadBefore_{c,u,d,i}}$	$[0, \infty)$	Expected workload (in minutes) at department *d, i* time slots *before* the start of consultation type *c* for unit *u*
$\boldsymbol{expectedWorkloadAfter_{c,u,d,i}}$	$[0, \infty)$	Expected workload (in minutes) at department *d, i* time slots *after* the end of consultation type *c* for unit *u*
firstSlot	$\{t_{min}, \ldots, t_{max}\}$	First timeslot a consultation can be booked into
lastSlot	$\{t_{min}, \ldots, t_{max}\}$	Last timeslot a consultation can be booked into
numberOfSlots	$\in \mathbb{N}$	The total number of timeslots
$\boldsymbol{departmentWeight_{d}}$	$[0, 1]$	Weight of each department for objective function
slidingWindowWidth	$\in \mathbb{N}$	Specifies the number of time slots one sliding window covers. The first

		window starts at time t=1, to t=1+slidingWindowWidth. The second window starts at time t=2, to t=2+slidingWindowWidth
bigM	150	Parameter to enable certain 'either-or' constraints to work

4.3.3 Decision variables

Variable	Range	Description
$X_{s,c,u,t}$	{0, 1} Binary	Start consultation type *c* of unit *u* at time slot *t* in schedule s (1=yes, 0=no)

4.3.4 Auxiliary decision variables

Variable	Range	Description
$Y_{s,c,u,t}$	{0, 1} Binary	Consultation type *c* of unit *u* is taking place at time *t* in schedule *s* (1=yes)
$Workload_{d,t}$	[0, ∞)	Workload for department *d*, at timeslot *t*
$WorkloadGrouped_{d,sw}$	[0, ∞)	Workload grouped, $Workload_{d,t}$ grouped into sliding time windows SW
$DeviationFromNorm_{d,t}$	$\in \mathbb{R}$	Absolute deviation from norm for department *d* in time window *sw*
$DeviationFromNormGrouped_{d,sw}$	$\in \mathbb{R}$	The absolute deviation from the norm for each time slot *t* is summed for sliding time window *sw*
$MaxDev_d$	$\in \mathbb{R}$	The maximum deviation from the norm for department *d* across all time slots

4.3.5 Constraints

1. Schedule the correct number of consultations for each schedule, consultation type and unit

$$\sum_t X_{s,c,u,t} = consultationsPerType_{s,c,u} \qquad \forall s,c,u$$

2. Block time slots for consultation duration

When a consult *c* is planned at time *t=5* in schedule *s*, $X_{s,c,u,5}$=1, with consultationDuration=2, then $Y_{s,c,u,5}$ and $Y_{s,c,u,6}$ need to be 1. When looking from the perspective of $Y_{s,c,u,t}$ we need to look back ($consultationDuration_{s,c,u}$-1) time slots.

$$Y_{s,c,u,t} = \sum_{tt=0}^{tt=consultationDuration_{s,c,u}-1} X_{s,c,u,t-tt} \qquad \forall s,c,u,t$$

To ensure the calculation $t - tt$ does not result in a value below 1 (out of bounds for the time horizon), the constraint is adjusted:

$$Y_{s,c,u,t} = \sum_{tt=0}^{tt=MIN(t-1;\, consultationDuration_{s,c,u}-1)} X_{s,c,u,t-tt} \qquad \forall s,c,u,t$$

3. Ensure total time blocked is correct

This constraint is necessary to ensure the model does not place a consultation close to the lastSlot, where not the entire consultation would fit.

$$\sum_t Y_{s,c,u,t} = consultationDuration_c * \sum_t X_{s,c,u,t} \qquad \forall s,c,u$$

4. Maximum one consultation per time slot within one schedule

$$\sum_c \sum_u Y_{s,c,u,t} \leq 1 \qquad \forall s,t$$

5. No consultations before first slot

$$\sum_{t=1}^{t=firstSlot\,-\,1} Y_{s,c,u,t} = 0 \qquad \forall s,c,u$$

6. No consultations after last slot

$$\sum_{t=lastSlot\,+\,1}^{T} Y_{s,c,u,t} = 0 \qquad \forall s,c,u$$

7. Maximum of two adjacent 'New' consultations in one schedule

We choose to allow no more than two 'New' consultations adjacent to each other. After two 'New' consultations only consultations of a different type can be scheduled (or that time is not booked at all, which also relieves the schedule in the desired way).

$$\sum_{tt=0}^{tt=consultationDuration_{New}*2+MIN(consultationDuration_{s,c})-1} \sum_u X_{s,New,u,t+tt} \leq 2 \qquad \forall s,t$$

To ensure the calculation $t + tt$ does not go above the latest time slot T, the constraint is adjusted:

$$\sum_{tt=0}^{tt=MIN(T-t;\, consultationDuration_{New}*2+MIN(consultationDuration_{s,c})-1)} \sum_u X_{s,New,u,t+tt} \leq 2 \qquad \forall s,t$$

8. **Set downstream workload**

$$Workload_{d,t} = \sum_{s}\sum_{c}\sum_{u}\sum_{i=1}^{I}\left(X_{s,c,u,t+i} * expectedWorkloadBefore_{c,u,d,i}\right)$$
$$+ \sum_{s}\sum_{c}\sum_{u}\sum_{i=1}^{I}\left(X_{s,c,u,t-consultationDuration_{s,c,u}-i+1}\right.$$
$$\left.* expectedWorkloadAfter_{c,u,d,i}\right) \quad \forall d,t$$

To ensure the calculation $t+i$ does not exceed maximum time slot T, and calculation $t - consultationDuration_{s,c,u} - i + 1$ does not reach values below 1, the constraint is adjusted. The set time offsets I needs to start from i=0.

$$Workload_{d,t} = \sum_{s}\sum_{c}\sum_{u}\sum_{i=0}^{i=MIN(T-t;\,I)}\left(X_{s,c,u,t+i} * expectedWorkloadBefore_{c,u,d,i}\right)$$

$$+ \sum_{s}\sum_{c}\sum_{u}\sum_{i=0}^{i=MAX(0;\,MIN(t-consultationDuration_{s,c,u};\,I))}\left(X_{s,c,u,t-consultationDuration_{s,c,u}-i+1}\right.$$
$$\left.* expectedWorkloadAfter_{c,u,d,i}\right) \quad \forall d,t$$

9. **Group workload into sliding time windows**

All workload of time slots t is aggregated to the workload per sliding window sw, which functions similar to a moving average.

$$WorkloadGrouped_{d,sw} = \sum_{t=sw}^{sw+slidingWindowWidth-1} Workload_{d,t} \qquad \forall\, sw, d$$

10. **Calculate positive and negative deviation from norm**

The following constraints linearize the absolute deviation from the norm variable.

$$PositiveDeviationFromNorm_{d,t} - NegativeDeviationFromNorm_{d,t} = Workload_{d,t} - norm_{d,t} \quad \forall d,t$$

11. **Ensure positive deviation and negative deviation are >0**

$$PositiveDeviationFromNorm_{d,t}, NegativeDeviationFromNorm_{d,t} \geq 0 \qquad \forall d,t$$

12. **Ensure only positive OR negative deviation takes a value >0 (constraint one)**

The variable $DeviationIsPositive_{d,t}$ is a binary variable. It is 0 if the deviation from the norm is negative, 1 if the deviation from the norm is positive. Using this variable in combination with bigM ensures only one of the deviation variables can take a value other that 0.

$$PositiveDeviationFromNorm_{d,t} \leq DeviationIsPositive_{d,t} * bigM \qquad \forall d,t$$

13. **Ensure only positive OR negative deviation takes a value >0 (constraint two)**

$$NegativeDeviationFromNorm_{d,t} \leq \left(1 - DeviationIsPositive_{d,t}\right) * bigM \quad \forall d,t$$

14. Calculate the absolute deviation from the norm

Now that the deviation has been split into its positive and negative parts, we can take the sum to get the absolute difference between workload and the norm (without losing linearity).

$$AbsoluteDeviationFromNorm_{d,t} = PositiveDeviationFromNorm_{d,t} + NegativeDeviationFromNorm_{d,t} \quad \forall d, t$$

15. Calculate DeviationFromNormGrouped per sliding time window

For each sliding time window, sum the deviation from the norm for individual time slots, for all time slots that belong to sliding window *sw*

$$DeviationFromNormGrouped_{d,sw} = \sum_{t=sw}^{sw+slidingWindowWidth-1} DeviationFromNorm_{d,t} \quad \forall d, sw$$

16. Calculate maximum deviation from norm

For each department, determine the maximum deviation from the norm over all the sliding windows

$$MaxDev_d \geq DeviationFromNormGrouped_{d,sw} \qquad \forall d, sw$$

4.3.6 Objective function

$$min \sum_d (MaxDev_d * departmentWeight_d)$$

4.3.7 Model choices commentary

Some model choices deserve further explanation, these are discussed in the following paragraphs.

First slot and last slot

closed	*closed*	*closed*	first slot			last slot	*closed*	*closed*	*closed*

Figure 4.3 – Example of schedule setup with first slot and last slot

Not all time slots in the model are used to schedule appointments. A certain number of time slots at the start and end of the schedule can be closed. The reason to include this in the model is that the expected workload at downstream departments can occur at an offset of multiple time slots before or after the consultation.

For example: the first 3 time slots are closed. If the downstream department (e.g., Radiology) is also closed at those times, the norm of workload for Radiology is set at 0. The model will then try to schedule a consultation type with *low* 'expected workload before' in the first open slot (slot 4 in this example), to keep expected workload at the closed time slots as close to 0 as possible.

Expected workload – one consultation

The expected workload indicates the expected time (in minutes) a department will spend to help a patient. The expected workload is calculated as follows: what is the probability a consultation of type *c* results in a patient visiting downstream department *d*? If a patient visits a downstream department, what is the average total number of minutes of work per patient? The expected workload is the product of the *transition probability* and the *average minutes of workload*.

Let's say the probability a patient coming for a 'New consultation' visits the Plaster room afterwards is 15%. When a patient visits the Plaster room, the average minutes of workload is 28 minutes. Expected workload in minutes at Plaster room for a new consultation: 0.15 * 28 = 4.2 minutes.

When the expected workload is determined, a choice is to be made on how to allocate this to time slots. We want the expected workload to span 28 minutes, and then be reduced to 15%, to give a good representation. See the second row of Table 4.1 – Example of placing expected workload (in minutes) into 5-minute time slots of what we want to do.

	t=1	t=2	t=3	t=4	t=5	t=6
1. Wrong	4.2	0	0	0	0	0
2. Right	5*0.15 = 0.75	5*0.15 = 0.75	5*0.15 = 0.75	5*0.15 = 0.75	5*0.15 = 0.75	3*0.15= 0.45

Table 4.1 – Example of placing expected workload (in minutes) into 5-minute time slots

The example above was about the expected workload to the Plaster room *after* the consultation. For each consultation type the expected workload *before* and *after* is calculated for *each* downstream department. Figure 4.4 shows an example of all expected workload for a single consultation type.

Figure 4.4 – Example of expected workload (in minutes) before and after for all downstream departments of one consultation

Expected workload – all workload at time *t* for department *d*
The calculation of workload from one consultation for one downstream department is discussed in the previous paragraphs, this section explains how the expected workload of multiple consultations and multiple schedules is aggregated. Figure 4.5 shows the example schedules. Doctor 1 starts a 'Repeat' consultation at time t=6, doctor 2 starts a 'Discharge' consultation at time t=1, and doctor 3 starts a 'New' consultation at time t=7.

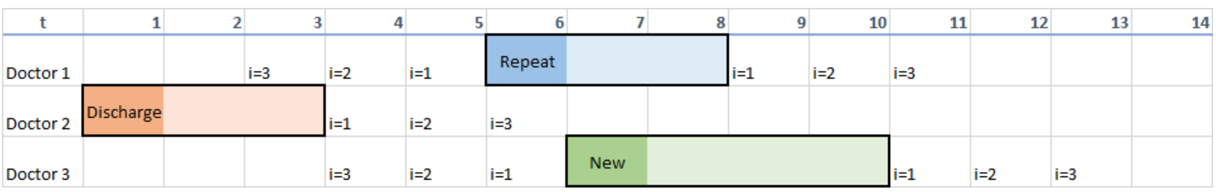

Figure 4.5 – Example of three schedules for workload aggregation explanation

A consultation may cause workload for downstream departments during a certain time *before* the consultation starts, and a certain time *after* the consultation has ended. The index *i* indicates the number of slots before (or after) the start (end) of the consultation. For this simplified example workload for downstream departments has a maximum offset of 3 time slots (i=3). For reference the real problem has a maximum offset i=16.

expectedWorkloadBefore	i=1	i=2	i=3
New	4.3 minutes	4.3 minutes	4.1 minutes
Repeat	3.9 minutes	3.9 minutes	1.2 minutes
Discharge	3.5 minutes	3.5 minutes	0.2 minutes

Table 4.2 – Mock figures of expected workload after for Radiology for each time offset i

expectedWorkloadAfter	i=1	i=2	i=3
New	3.8 minutes	3.8 minutes	3.2 minutes
Repeat	3.6 minutes	3.6 minutes	2.1 minutes
Discharge	1.7 minutes	1.7 minutes	1.1 minutes

Table 4.3 – Mock figures of expected workload after for Radiology for each time offset i

The tables above show some mock figures for expected workload at Radiology *before* (Table 4.2) and *after* (Table 4.3) the consultations. We will show two example calculations. The first example is simplified and only checks for workload *before* consultation starts. The second example calculation is complete and shows what the model does to calculate the total expected workload for a department per time slot.

Calculation example one: The expected workload for the Radiology department at time t=3 is calculated as follows.

+ i=1, does any consultation start one time slot later (t+i → 3+1 → t=4)? **No**
+ i=2, does any consultation start two time slots later (t+i → 3+2 → t=5)? **No**
+ i=3, does any consultation start three time slots later (t+i → 3+3 → t=6)? **Yes,** the repeat consultation of Doctor 1. The amount of workload expected i=3 time slots before its start is **1.2 minutes**

Thus, the total expected workload for the Radiology department at time t=2 is **1.2 minutes**. It works by checking if any consultations start at time $t + i$. Calculation example one is simplified since it only looked at workload *before* consultations. The second calculation example will be complete because it includes both the calculation of workload *before* and *after* consultations.

Calculation example two: the expected workload for the Radiology department at time t=4 is calculated as follows:

+ i=1
  - does any consultation start one time slot later (t+i → 4+1 → t=5)? **No**
  - does any consultation end one time slot before?
    - Consultations with duration 3: (t-consultationDuration-i+1 → 4 - 3 - 1 + 1 → t=1). **Yes,** the discharge consultation of Doctor 2 has a duration of 3 time slots, and starts at time t=1, and ends at t=3. The amount of workload expected i=1 time slot after its end is **1.7 minutes**.
    - Consultations with duration 4: (t-consultationDuration-i+1 → 4 - 4 - 1 + 1 → t=0). **No**

- + i=2
  - o does any consultation start two time slots later (t+i → 4+2 → t=6)? **Yes,** the repeat consultation of Doctor 1. The amount of workload expected i=2 time slots before its start is **3.9 minutes**
  - o does any consultation end two time slots before?
    - Consultations with duration 3: (t-consultationDuration-i+1 → 4 - 3 - 2 + 1 → t=0). **No**
    - Consultations with duration 4: (t-consultationDuration-i+1 → 4 - 4 - 2 + 1 → t=-1). **No**
- + i=3
  - o does any consultation start three time slots later (t+i → 4+3 → t=7)? **Yes,** the new consultation of Doctor 3. The amount of workload expected i=3 time slots before its start is **4.1 minutes**
  - o does any consultation end three time slots before?
    - Consultations with duration 3: (t-consultationDuration-i+1 → 4 - 3 - 3 + 1 → t=-1). **No**
    - Consultations with duration 4: (t-consultationDuration-i+1 → 4 - 4 - 3 + 1 → t=-2). **No**

The total expected workload at the Radiology department at time t=4 is **9.7 minutes** (1.7 + 3.9 + 4.1). The calculation works by checking if any workload from *before* a consultation needs to be added, and thus checks if any consultations start at time $t + i$. The model also checks if any workload from *after* a consultation needs to be added, and checks if any consultations end at time $t - consultationDuration - i + 1$.

This calculation example was for one department. For each department the calculation is performed separately (with the respective expected workload parameters for that department) to reach a total expected workload at each time slot *t* for each separate department *d*.

Sliding window

For the cases where matching the expected workload and the norm is not important at the time scale of one time slot (5 minutes in this case), the sliding windows give the option to aggregate deviation from the norm over multiple time slots. How many time slots are aggregated can be chosen by the user. If the smallest time scale is important, the sliding window width can be set to 1. See Appendix sliding window for a visualisation of how the aggregation spans multiple time slots.

4.4 Model implementation in software

The MILP model is implemented in the software package AIMMS.

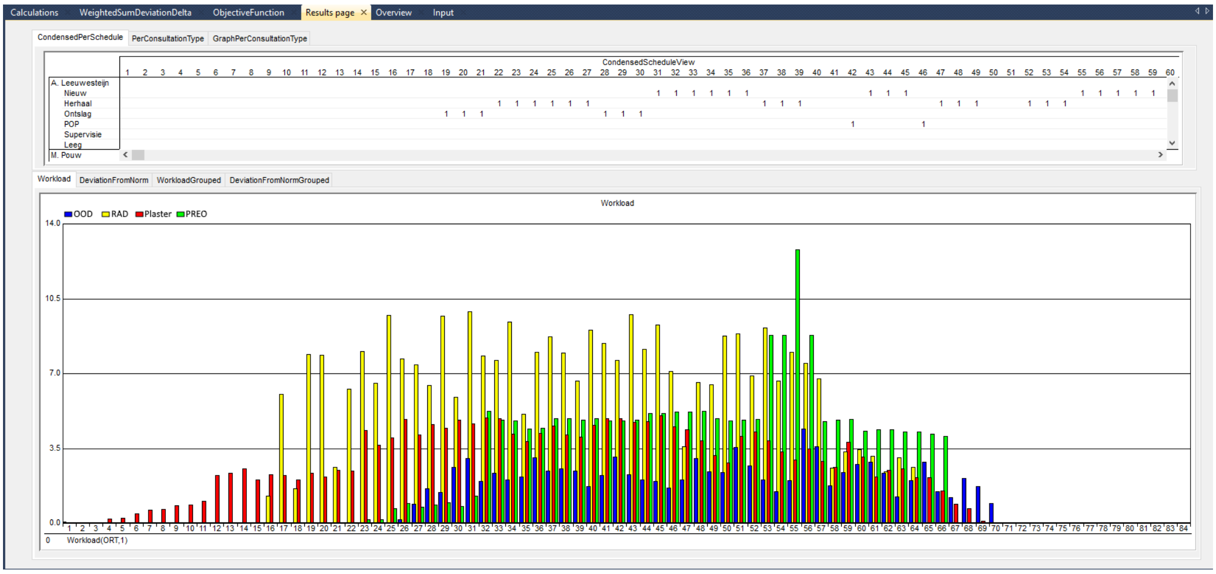

Figure 4.6 – Screenshot of the model implementation in AIMMS

See Figure 4.6 for a screenshot of the model implementation in the AIMMS software package and shows the result page. An overview of the generated schedule for each doctor is given at the top of the page, and the expected workload graphs at the bottom of the page. These visual representations give a clear image of expected downstream demand of the generated blueprint schedule.

4.5 Model verification and validation

Model verification and validation is performed. During model verification the model is checked to run like we expect it to. A few examples of items that are verified:

- ✓ The model schedules the correct number of consultations for each schedule
- ✓ The model blocks the correct number of time slots corresponding with the consultation durations
- ✓ The model adds expected workload *before* consultations correctly
- ✓ The model adds expected workload *after* consultations correctly
- ✓ The absolute deviation from the norm is calculated correctly

For model validation the output of the model is presented to hospital staff responsible for the session schedules to ask if the output is expected, to spot any irregularities and to ask if any requirements (constraints) are missing. Based on the feedback of hospital stakeholders and the problem owner in Logistiek Bedrijf the model is extended from the first version through a number of iterations to finally reach the model as described above in Section 4.3. In the following section the model extensions and how they are implemented in the model are discussed.

4.6 Model extension options

All extensions add to the base version of the model, see Appendix G for the base model definition. The new model definitions described in these sections are relative to the base model and the model extension options before it.

4.6.1 5-minute time slots

The first version of the model has one consultation per time slot. The limitation of this is all consultations need to be of identical duration. If all consultations are e.g. 15 minutes this works, but if one consultation type takes 10 minutes and another type takes 20 minutes it is impossible to use the first version of the model.

An interesting model extension is to work with smaller time slots, where a consultation can span multiple time slots. One slot can be chosen to represent 5 minutes. Consultation durations can then be defined as a multiple of that (e.g. 10 or 15 minutes).

Decision variable $X_{c,t}$ now indicates if a consultation of type *c* **starts** at time t. A new variable is added to the model: $Y_{c,t}$ which indicates if a consultation of type *c* is **taking place** at time t. The parameter $consultationDuration_c$ is added to specify per consultation type the number of time slots it takes.

Variable	Range	Description
$\boldsymbol{Y_{c,t}}$	{0, 1} Binary	Consultation type *c* is taking place at time *t* (1=yes)
Parameter	**Range**	**Description**
$\boldsymbol{consultationDuration_c}$	$\in \mathbb{N} \setminus \{0\}$	The duration of a consultation of type *c* in number of time slots

Constraint $\sum_c X_{c,t} \leq 1 \quad \forall t$ is changed to restrict no more than 1 consultation **taking place** at the same time: $\sum_c Y_{c,t} \leq 1 \quad \forall t$. Similarly, the constraints that ensure no consultations before first slot and after last slot switch $X_{c,t}$ for $Y_{c,t}$. For the constraint that sets the downstream workload the second term is adjusted to take consultation duration into account, the new formulation is:

$$Workload_{d,t} = \sum_c \left(X_{c,t+1} * expectedWorkloadBefore_{c,d}\right) + \sum_c \left(X_{c,t-consultationDuration_c+1} * expectedWorkloadAfter_{c,d}\right) \quad \forall d,t$$

Two new constraints are added. One constraint so the variable $Y_{c,t}$ is 1 when $X_{c,t}$ is 1, and for the correct number of time slots - the consultation duration - after $X_{c,t}$ as well. The other constraint to ensure if a consultation lasts three time slots, three time slots are blocked ($Y_{c,t}$ is 1), to avoid a consultation with duration three to start at the last slot, which would effectively schedule it as a consultation with duration one.

$$Y_{c,t} = \sum_{tt=0}^{tt=consultationDuration_c-1} X_{c,t-tt} \qquad \forall c,t$$

$$\sum_t Y_{c,t} = consultationDuration_c * \sum_t X_{c,t} \qquad \forall c$$

4.6.2 Downstream workload multiple time slots offset

In the base model the workload at downstream departments either arrives before the consultation at t-1, or afterwards at t+1. However, if workload for a patient takes longer than one time slot (in our case longer than 5 minutes) a better representation is to spread the workload over multiple time slots (as explained in Section 4.3.7). Additionally, some demand at downstream departments will not arrive right before or after the consultation, but after some travel time.

The parameters $expectedWorkloadBefore_{c,d}$ and $expectedWorkloadAfter_{c,d}$ are extended with index *i* that represents the number of time slots offset. The set I defines the time offsets {1..i}.

Set	Element	Description
I	i	Time offsets, used to define a demand profile 1..i time slots before the start of the consultation, and 1..i time slots after the end of a consultation.

Parameter	Range	Description
$\boldsymbol{expectedWorkloadBefore_{c,d,i}}$	[0, ∞)	Expected workload (in minutes) at department *d, i* time slots *before* the start of consultation type *c*
$\boldsymbol{expectedWorkloadAfter_{c,d,i}}$	[0, ∞)	Expected workload (in minutes) at department *d, i* time slots *after* the end of consultation type *c*

The workload constraint is also adjusted to call the correct expected workload at the correct time. See Section 4.3.7 for an explanation of the workload calculation.

$$Workload_{d,t} = \sum_{c}\sum_{i=1}^{I}(X_{c,t+i} * expectedWorkloadBefore_{c,d,i}) + \sum_{c}\sum_{i=1}^{I}(X_{c,t-consultationDuration_c-i+1} * expectedWorkloadAfter_{c,d,i}) \quad \forall d,t$$

4.6.3 Create multiple blueprint schedules simultaneously

In the base model only one blueprint schedule is created, however ideally the model can create multiple blueprint schedules simultaneously. If the same blueprint schedule is used multiple times the positive effect of the model for one schedule will likely be undone because the multiple identical schedules will amplify the low and high points of patient flow.

We extend the model to allow for multiple separate schedules to be generated by the model and it can optimize the combined outflow of all schedules simultaneously. The set of Schedules S is added, and the $consultationsToBeScheduled_{s,c}$ parameter now specifies the number of consultation type *c* to be scheduled in schedule *s*. Variables $X_{s,c,t}$ and $Y_{s,c,t}$ now specify the schedule *s*.

Set	Element	Description
S	s	Schedules

Parameter	Range	Description
$\boldsymbol{consultationsToBeScheduled_{s,c}}$	∈ ℕ	The number of consultation type *c* to be scheduled in schedule *s*

Variable	Range	Description
Y_{s,c,t}	{0, 1} Binary	Consultation type *c* is taking place at time *t* in schedule *s* (1=yes)
X_{s,c,t}	{0, 1} Binary	Start consult type *c* at time slot *t* in schedule s (1=yes, 0=no)

Many constraints are extended such that each schedule has the same restrictions as the single schedule in the base model. The maximum of one consultation taking place at one time per schedule: $\sum_c Y_{s,c,t} \leq 1 \quad \forall t, s$. The correct number of consultations need to be scheduled in each schedule: $\sum_t X_{s,c,t} = consultationsPerType_{c,s} \quad \forall c, s$. No consultations before firstSlot or after lastSlot: $\sum_{t=1}^{t=firstSlot\,-\,1} Y_{s,c,t} = 0 \quad \forall c, s$ and $\sum_{t=lastSlot\,+\,1}^{T} Y_{s,c,t} = 0 \quad \forall c, s$.

The calculation of workload now also sums over all schedules *s*:

$$
\begin{aligned}
Workload_{d,t} = & \sum_s \sum_c \sum_{i=1}^{I} \left(X_{s,c,t+i} * expectedWorkloadBefore_{c,d,i}\right) \\
& + \sum_s \sum_c \sum_{i=1}^{I} \left(X_{s,c,t-consultationDuration_c-i+1}\right. \\
& \left. * expectedWorkloadAfter_{c,d,i}\right) \qquad \forall d, t
\end{aligned}
$$

4.6.4 Specify units, with different appointment characteristics per unit

Another request is to add the unit specification (e.g. Knee or Spine) to the model. The transition probabilities for workload are different between units, and doctors will only see certain type of patients (e.g. knee or hip) anyways, so the information is available when the staff schedule is made (and before the blueprint session schedules are applied).

The set of units U is added to the model. Many parameters are extended with index *u* to specify the unit.

Set	Element	Description
U	u	Units

Parameter	Range	Description
$\boldsymbol{consultationsToBeScheduled_{s,c,u}}$	$\in \mathbb{N}$	The number of consultation type *c* of unit *u* to be scheduled in schedule *s*
$\boldsymbol{expectedWorkloadBefore_{c,u,d,i}}$	[0, ∞)	Expected workload (in minutes) at department *d, i* time slots *before* the start of consultation type *c* for unit *u*
$\boldsymbol{expectedWorkloadAfter_{c,u,d,i}}$	[0, ∞)	Expected workload (in minutes) at department *d, i* time slots *after* the end of consultation type *c* for unit *u*

Variable	Range	Description
***Y*$_{s,c,u,t}$**	{0, 1} Binary	Consultation type *c* of unit *u* is taking place at time *t* in schedule *s* (1=yes)
***X*$_{s,c,u,t}$**	{0, 1} Binary	Start consult type *c* of unit *u* at time slot *t* in schedule s (1=yes, 0=no)

Again, many constraints are extended with index *u* for unit. The correct number of consultations to be scheduled: $\sum_t X_{s,c,u,t} = consultationsPerType_{s,c,u} \quad \forall s,c,u$. Block time slots for consultations: $Y_{s,c,u,t} = \sum_{tt=0}^{tt=consultationDuration_{s,c,u}-1} X_{s,c,u,t-tt} \quad \forall s,c,u,t$. Ensure total time blocked is correct: $\sum_t Y_{s,c,u,t} = consultationDuration_c * \sum_t X_{s,c,u,t} \quad \forall s,c,u$. A maximum of one consultation per time slot in one schedule: $\sum_c \sum_u Y_{s,c,u,t} \leq 1 \quad \forall s,t$. No consultations before firstSlot or after lastSlot: $\sum_{t=1}^{t=firstSlot\,-\,1} Y_{s,c,u,t} = 0 \quad \forall s,c,u$ and $\sum_{t=lastSlot\,+\,1}^{T} Y_{s,c,u,t} = 0 \quad \forall s,c,u$.

The calculation of workload now includes the sum over all units U:

$$Workload_{d,t} = \sum_s \sum_c \sum_u \sum_{i=1}^{I} \left(X_{s,c,u,t+i} * expectedWorkloadBefore_{c,u,d,i}\right) + \sum_s \sum_c \sum_u \sum_{i=1}^{I} \left(X_{s,c,u,t-consultationDuration_{s,c,u}-i+1} * expectedWorkloadAfter_{c,u,d,i}\right) \quad \forall d,t$$

4.6.5 Model DUO and TRIO sessions

The base model considers UNO schedules (just a physician without supporting staff). However, DUO and TRIO sessions are common, but they are different from UNO sessions in the following ways:

- Support staff may need longer per consultation
- Supervision time needs to be scheduled, for the physician to give advice and training to the support staff

At the present we do not have information on the exact rules of current DUO/TRIO sessions (when exactly is support needed, how often), and from interviews an indication that the current way might not be the best suited either. This is likely not a straightforward extension. Two approaches are possible. We will first discuss the simpler extension, which is also the final version described in Section 4.3. The second more elaborate extension described in Section 4.6.6 is not used for this research, but is promising for further research and as part of a broader implementation.

The definition of consultation duration for support staff means we add index *s* to the $consultationDuration_c$ parameter so different durations can be set per schedule. Additionally, the feedback from stakeholders was to allow for different durations per unit *u*.

Parameter	Range	Description
***consultationDuration*$_{s,c,u}$**	$\in \mathbb{N} \setminus \{0\}$	The duration of a consultation of type *c* for unit *u* in schedule *s* in number of time slots

Constraints containing parameter $consultationDuration$ are updated to include the *s* and *u* indexes. For the addition of supervision time, a new consultation type ‘Supervision’ is added to set C.

4.6.6 Synchronised consultation starts across schedules

A likely requirement of DUO or TRIO session schedule is that supervision occurs at the same time for all linked schedules. So, if Doctor 1 has two support staff assigned, supervision time starts at the same time in all three schedules.

We model this by defining C_s, the subset of consultation types that need to have a synchronised start in all linked schedules. To define which schedules are linked, we define set $ScheduleGroups$ (SG), and parameter $scheduleIsInGroup_{s,sg}$ which is 1 when schedule s belongs to schedule group sg. Finally new decision variable $StartSynchronisedConsultAtTimeT_{sg,c,u,t}$ is added that is 1 when consultation type *c* and unit *u* starts at time *t* for all schedules in schedule group *sg*.

Set	Element	Description
SG	sg	Schedule groups, is used to link schedules that need synchronised starts of certain consultation types

Subset	Element	Description
C_s	C (consultation types)	Synchronised consultations, a subset of consultations that need to have the same starting time in all schedules *s* in a schedule group *sg*

Parameter	Range	Description
$scheduleIsInGroup_{s,sg}$	{0, 1} Binary	Indicates if schedule s belongs to schedule group *sg*

Variable	Range	Description
$StartSynchronisedConsultAtTimeT_{sg,c,u,t}$	{0, 1} Binary	Start (synchronised) consult type *c* of unit *u* at time slot *t* for all schedules in schedule group *sg*

Additional constraints are needed to make the synchronised consultation starts work.

Each schedule belongs to at most one schedule group:

$$\sum_{sg} scheduleIsInGroup_{s,sg} \leq 1 \qquad \forall s$$

If the decision is to *not* start a synchronised consultation, ensure no schedule in the group starts it:

$$\sum_{s \in sg} X_{s,c,u,t} \leq bigM * StartSynchronisedConsultAtTimeT_{sg,c,u,t} \qquad \forall sg, c \in C_s, u, t$$

When the decision is made to start the synchronised consultation, ensure all in the group start it:

$$\sum_{s \in sg} X_{s,c,u,t} \geq \sum_{s \in sg} scheduleIsInGroup_{s,sg} * StartSynchronisedConsultAtTimeT_{sg,c,u,t} \qquad \forall sg, c \in C_s, u, t$$

4.7 Conclusion

We have defined and implemented a MILP model that is able to generate blueprint schedules based on a given number of consultation types. It can generate multiple blueprint schedules simultaneously, where the number of consultations to be scheduled can be different for each schedule. Each schedule can contain consultation types from different units (e.g. Spine, or Knee), and the correct demand profile will be applied.

To accommodate different units and also doctors in training, the consultation duration can be set for each combination of schedule, consultation type, and unit.

The results of the model are checked to satisfy the requirements as derived from interviews with staff and an staff member experienced with the session schedules. More extensive validation was not possible due to insufficient spare time of schedulers and planners at during this stage of the research project. However, I expected no large changes during further validation. See Section 6.1 for more discussion about additional validation.

The model will determine the optimal sequence and consultation starting time to minimize the difference between expected downstream workload and the norm and create more smooth demand patterns at downstream departments. It achieves this by minimizing the maximum deviation from the norm for each department. In the next chapter the model will be tested with real world blueprint schedules that are currently in use in the hospital.

5 Experiments and results

In this chapter the model will be tested with blueprint schedules to check how much improvement can be expected by using this model over the current way of working. Section 5.1 discusses the setup of the experiments. Section 5.2 discusses the parameter settings used in the experiments. Section 5.3 discusses the algorithm running time, and Section 5.4 discusses the experiments. Section 5.5 describes the approach for Monte Carlo simulation used to asses the transitions in a stochastic manner.

5.1 Experiment design

For a sample of 6 sessions the used blueprint schedules are retrieved. These contain the schedules of multiple doctors in the same session.

	Number of doctors	*Total number of consultations*
Wednesday morning	6	99
Wednesday afternoon	6	80
Thursday morning	5	74
Thursday afternoon	8	111
Friday morning	6	88
Friday afternoon	5	61

Table 5.1 – Sample of blueprint schedules for used for the experiments

From each individual blueprint schedule the number of each consultation type is counted. The model then generates a blueprint schedule, given the same number of each consultation type for each doctor. The model outputs a weighted sum of the maximum absolute deviation from the norm, indicating how well the generated blueprint schedule matches expected workload to the desired norm.

This figure is compared with the score the original blueprint schedule gets from the model. The model in AIMMS allows for a schedule to be entered manually. By comparing the figure for the generated blueprint schedule and the original schedule, we get an impression of the type of improvement that can be achieved by using this model.

5.2 Parameter settings

Set	Values
Consultation types	New, repeat, discharge, POP, Supervision, Empty
Departments	OOD, RAD, Plaster, PREO
Schedules	*Depends on the case*
Units	Upper extremities, Hip, Knee, Spine, Foot

Parameter	Values
$\boldsymbol{consultationDuration_{s,c,u}}$	*Depends on the case*
$\boldsymbol{consultationsToBeScheduled_{s,c,u}}$	*Depends on the case*
$\boldsymbol{norm_{d,t}}$	The total workload for that department divided by the number of open slots. *Depends on the case*

expectedWorkloadBefore$_{c,u,d,i}$	Identical for all experiments. Maximum offset is 16 time slots.
expectedWorkloadAfter$_{c,u,d,i}$	Identical for all experiments. Maximum offset is 16 time slots.
firstSlot	19, which means 18 slots beforehand (that represent 90 minutes) are closed. A minimum of 16 slots is needed because of the maximum workload offset
lastSlot	60, which gives 3.5 hours of scheduled time. If the existing schedule in a case needs more slots, the last slot is increased until exactly that number.
numberOfSlots	84, since this allows for at least 16 closed slots before the first slot and after the last slot. In total this spans 7 hours.
departmentWeight$_d$	0.25 for all four departments
slidingWindowWidth	3, meaning workload is aggregated over a moving 15 minute window
bigM	150, since the expected workload in one time slot is not going to exceed 150 minutes of workload in the 5 minutes one time slot represents.

The sets and parameters indicated with *'depends on the case'* for the example case of Thursday afternoon are specified in Appendix H.

5.3 Algorithm running time

The experiments are performed on a laptop with the Intel i7-7600U CPU, with a clock speed of 3.05 GHz. Table 5.2 shows the running time for the largest problem instance (Thursday afternoon). The model seems to find a good solution quite fast. However since the problem is tactical in nature we choose a maximum running time of 8 hours which will allow a run overnight.

Creation method	Time limit	WeigthedDeviationFromNorm	Absolute improvement	% improvement	Improvement per second
Original		19.3			
Model (15s)	15	12.1	7.232	37.3%	0.02489
Model (30s)	30	12.1	7.232	37.3%	0.01245
Model (60s)	60	12.1	7.250	37.4%	0.00624
Model (2 min)	120	12.1	7.250	37.4%	0.00312
Model (5 min)	300	11.8	7.480	38.6%	0.00129
Model (10 min)	600	11.8	7.505	38.8%	0.00065
Model (30 min)	1800	11.8	7.565	39.1%	0.00022
Model (8 hours)	28800	11.7	7.622	39.4%	1.4E-05

Table 5.2 – Running time and improvement of objective function

Smaller problem instances need less time as show in Table 5.3, where most schedules reach their optimum before 8 hours.

Schedule	Runtime (sec)	Score manual	Score generated	Integrality gap	Iterations	Nodes
Wednesday morning	1762	8.94	3.48	0.00%	31239976	120271
Wednesday afternoon	28800	7.82	3.22	0.29%	207899699	317600
Thursday morning	19588	9.76	3.02	0.00%	164409378	218645
Thursday afternoon	28800	12.12	7.06	6.19%	144145462	859417
Friday morning	1351	7.41	2.89	0.00%	24124870	48289
Friday afternoon	476	7.17	2.78	0.00%	7977590	21395

Table 5.3 – Running time used for the experiments

5.4 Experiments

For each of the six sessions the experiments are performed. The original schedule is entered into AIMMS, and the maximum deviation from the norm is calculated. Then the model is used to generate a new schedule and the same calculation is used to get a score for the generated schedule. The runtime of the model is set to max. 8 hours, however some models reach their optimum earlier.

For the quality of the generated schedule we compare three performance measures against the original (manual) schedule:

- The maximum deviation from the norm
- The sum of deviations from the norm
- The coefficient of variation from the norm

The maximum deviation from the norm is calculated for each downstream department and is the biggest absolute difference between the desired amount of workload and the amount of workload the model expects.

The sum of deviations from the norm is the total absolute deviation difference between the norm and expected workload over all time slots.

The coefficient of variation (CV) is defined as the standard deviation divided by the mean of a data set: $\frac{\sigma}{\mu}$. It gives a good indication of how large the variability is relative to the mean, and allows comparison between departments.

These three measures are calculated per downstream department, and a weighted sum calculates the final measure, for now we have used a weight of 0.25 for each department. For all measures a lower score is better. The results of the generated schedules are discussed in the following paragraphs.

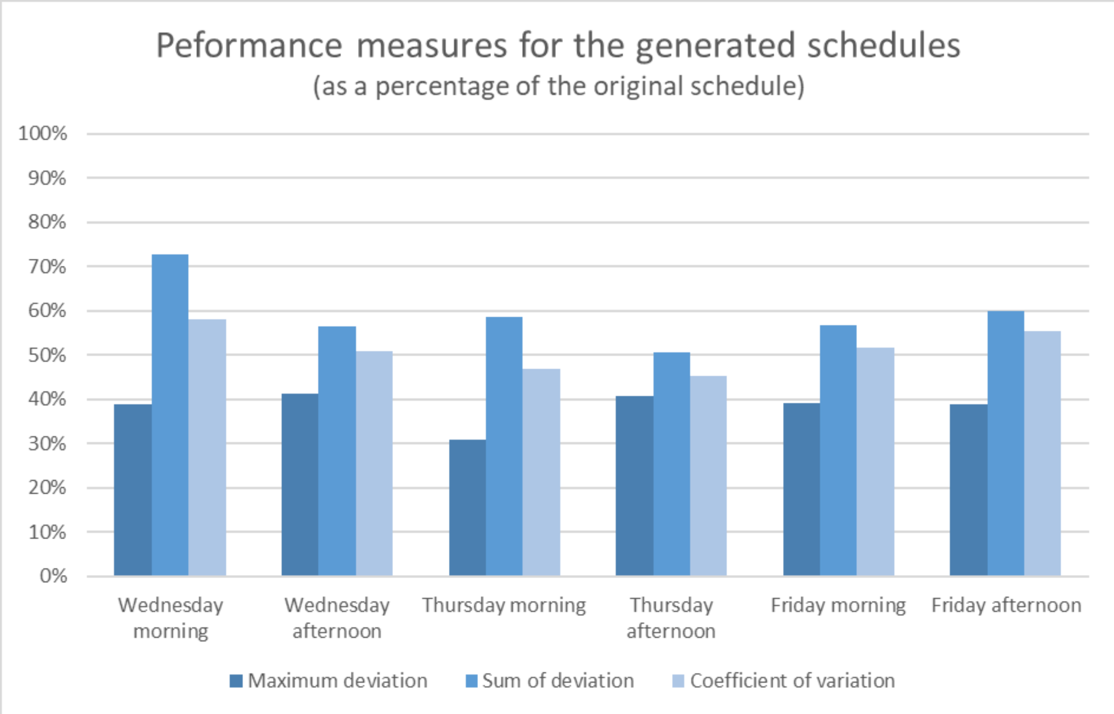

Figure 5.1 – The performance of the generated schedules in % of the original schedule

Figure 5.1 shows the solution quality on the three measures as compared to the original. The score for the original (manual) schedules is 100% for each of these measures. The results show a significant improvement in all three measures. The maximum deviation is reduced to approximately 40% of the original. The sum of deviation from the norm over all time slots is also reduced to approximately 60% of the original schedules. Finally, the coefficient of variation is significantly reduced to approximately 50% of the original schedules.

Appendix I.i contains the difference in performance in absolute difference in minutes. Appendix I.ii contains a complete overview of all percentage and absolute differences, split per department for a detailed overview.

For Thursday afternoon the workload graphs before and after are shown. Figure 5.2 shows the expected demand pattern, each colour is a different downstream department. This is made with a moving average (sliding window) of 3 time slots, and yet the fluctuations in expected workload are quite noticeable.

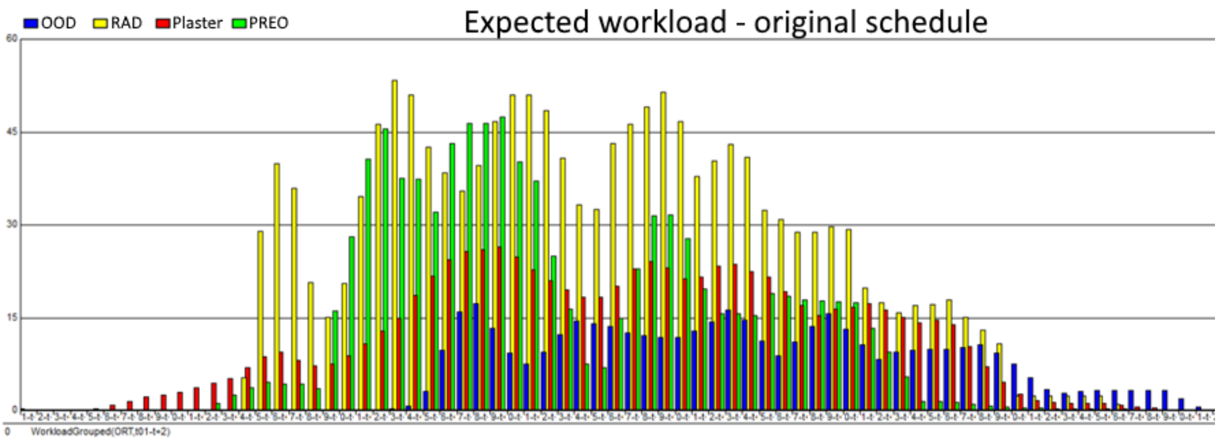

Figure 5.2 – Demand pattern for the original schedules on Thursday afternoon

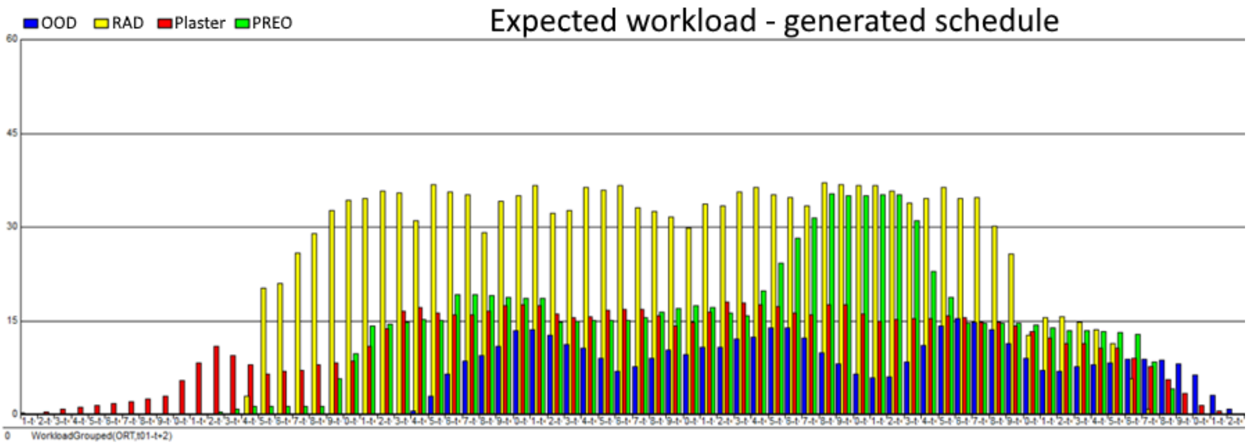

Figure 5.3 – Demand pattern for generated schedules on Thursday afternoon

Figure 5.3 shows the demand pattern of the generated schedules. It shows the maximum peaks of demand are reduced. In addition, the workload is spread much move evenly throughout the day. Finally, the green downstream department represents the preoperative screening department, where walk-in spots are available at the end of the afternoon. The norm for desired workload for the preoperative screening department is increased near the end of the afternoon, and the model behaves as expected and schedules the consultations with patients that are likely to go to preoperative screening closer to the available walk-in screening spots, reducing waiting time between departments for the patients. The changes in the sequence of consultation for all schedules are shown on the next page by Table 5.4 and Table 5.5.

	Doctor 1	Doctor 2	Doctor 3	Doctor 4	Doctor 5	Doctor 6	Doctor 7	Doctor 8
1	POP	POP	POP	New	New	New	New	New
2	New	New	New	Repeat	Repeat	Repeat	POP	POP
3	Repeat	Empty	Empty	Repeat	New	New	Repeat	Repeat
4	POP	POP	POP	POP	New	Repeat	New	New
5	Repeat	Repeat	New	New	Repeat	New	POP	Repeat
6	New	New	Empty	Discharge	New	New	Repeat	New
7	Discharge	Empty	Repeat	New	Repeat	Repeat	New	Repeat
8	Repeat	Repeat	New	POP	New	New	POP	New
9	New	POP	Empty	Repeat	Repeat	Discharge	Repeat	Repeat
10	POP	New	Repeat	New	Repeat	New	Repeat	Repeat
11	New	Empty	POP	Repeat	Repeat	New	Discharge	Repeat
12	Discharge	Repeat	New	POP	Repeat	Repeat	Discharge	
13	Repeat	New	Empty	Repeat			Repeat	
14	Repeat	Discharge	New	Discharge			Repeat	
15			Repeat				New	
16			Discharge				New	
17			Discharge				New	

Table 5.4 – Original schedule for Thursday afternoon

	Doctor 1	Doctor 2	Doctor 3	Doctor 4	Doctor 5	Doctor 6	Doctor 7	Doctor 8
1	Discharge	New	POP	Repeat	Repeat	Repeat	Repeat	Repeat
2	Discharge	Repeat	Empty	POP	New	New	New	New
3	Repeat	Repeat	Repeat	Repeat	Repeat	Repeat	New	POP
4	Repeat	Repeat	Repeat	Discharge	New	Repeat	Repeat	Repeat
5	New	Empty	Repeat	Repeat	New	New	POP	Repeat
6	Repeat	Discharge	New	New	Repeat	New	Repeat	Repeat
7	New	Empty	Discharge	POP	Repeat	New	POP	New
8	POP	POP	Discharge	Discharge	Repeat	New	Discharge	New
9	New	POP	New	New	New	New	Discharge	New
10	Repeat	New	New	POP	New	Repeat	POP	Repeat
11	POP	New	Empty	Repeat	Repeat	Discharge	Repeat	Repeat
12	Repeat	New	Empty	Repeat	Repeat	New	New	
13	POP	Empty	POP	New			New	
14	New	POP	Empty	New			Repeat	
15			POP				New	
16			New				New	
17			New				Repeat	

Table 5.5 – Generated schedule for Thursday afternoon

5.5 Monte Carlo simulation

The MILP model uses averages for the expected workload of the consultations. However, in reality the patient either transitions to the downstream departments or not, meaning either 0% or 100% of the workload for that specific transition is realised. We want to investigate whether the expected workload calculated by the MILP model gives a reasonable indication of the workload to expect, or if the interaction between multiple transitions produces unwanted patterns. If the actual workload distribution is heavily skewed to more or less workload than the MILP model predicts, then adjustments might need to be made to the MILP model, or it will at least aid with the appropriate interpretation of the expected workload results.

We approach this stochastic evaluation of the transitions through numerical experiments with the Monte Carlo method. The Monte Carlo method is based on repeated random sampling to obtain numerical results (Metropolis & Ulam, 1949). We develop a Monte Carlo simulation in the statistical software package R in order to test the behaviour of expected workload around the mean predicted by the MILP model.

In short, our Monte Carlo simulation takes the following input:

- Transition probabilities for each *consultation type* and *unit,* to each *department*, *before* or *after* the consultation
- The expected workload pattern for each *consultation type*, *unit,* for each *department* (given a 100% transition probability)
- Consultation durations for each *consultation type*, *unit,* and *schedule*
- The generated schedule from the MILP with the starting time for each consultation

In each Monte Carlo iteration the entire generated schedule from the MILP model is traversed, and for each consultation it is determined if any transitions to downstream departments occur.

For example, if a New consultation is scheduled, we determine if the patient will visit the Orthopaedic OD, Radiology, the Plaster room, and/or Screening before the consultation, and/or after the consultation. This means in total eight possible transitions are checked for each scheduled consultation.

We take eight independent draws from the binomial distribution to determine the realisation of each transition. The binomial distribution is appropriate since only two possible outcomes exist: either that transition occurs, or it does not. The binomial distribution has the following probability mass function (Hirsch, 1957):

$$P(X = k) = \binom{n}{k} p^k (1-p)^{n-k}$$

The formula describes the probability of getting exactly *k* successes in *n* trials where the probability of success in a single trial is equal to *p*. The number of trials (*n*) for each individual transition is 1. If the outcome is a transition occurs to that department, the appropriate workload is added to that department at the correct times. We perform 10,000 iterations in each Monte Carlo run, to reduce the stochastic uncertainty to a large extent, and the run time of the model is relatively short at approximately 20 minutes. The code can even be optimized further to utilize more than 1 CPU core for even quicker run times. Appendix J contains the complete R script used for the analysis.

From simulation output to workload graphs

The output of the simulations is analysed, and converted to graphs indicating the spread of workload for each department at each time *t*. See Figure 5.4 and Figure 5.5 for examples of the workload graphs that are calculated with the Monte Carlo simulation results. We show the spread of the workload through the following percentiles: 5^{th}, 25^{th}, 50^{th} (median line), 75^{th}, 95^{th}. In addition, we plot the deterministic average (as predicted by the MILP model) in red, and the average value from the Monte Carlo runs in yellow. Finally the dotted red line indicates the norm per time slot, what the model tries to match the workload to.

The meaning of the percentiles is as follows. For example in Figure 5.4 the 95^{th} percentile indicates at time t=33 is approximately 23 minutes of workload. This means that in 95% of the 10,000 iterations in the Monte Carlo simulation, the actual workload at time t=33 is 23 minutes or less. The 75^{th} percentile is approximately 18 minutes, so in 75% of the cases the workload is 18 minutes or less.

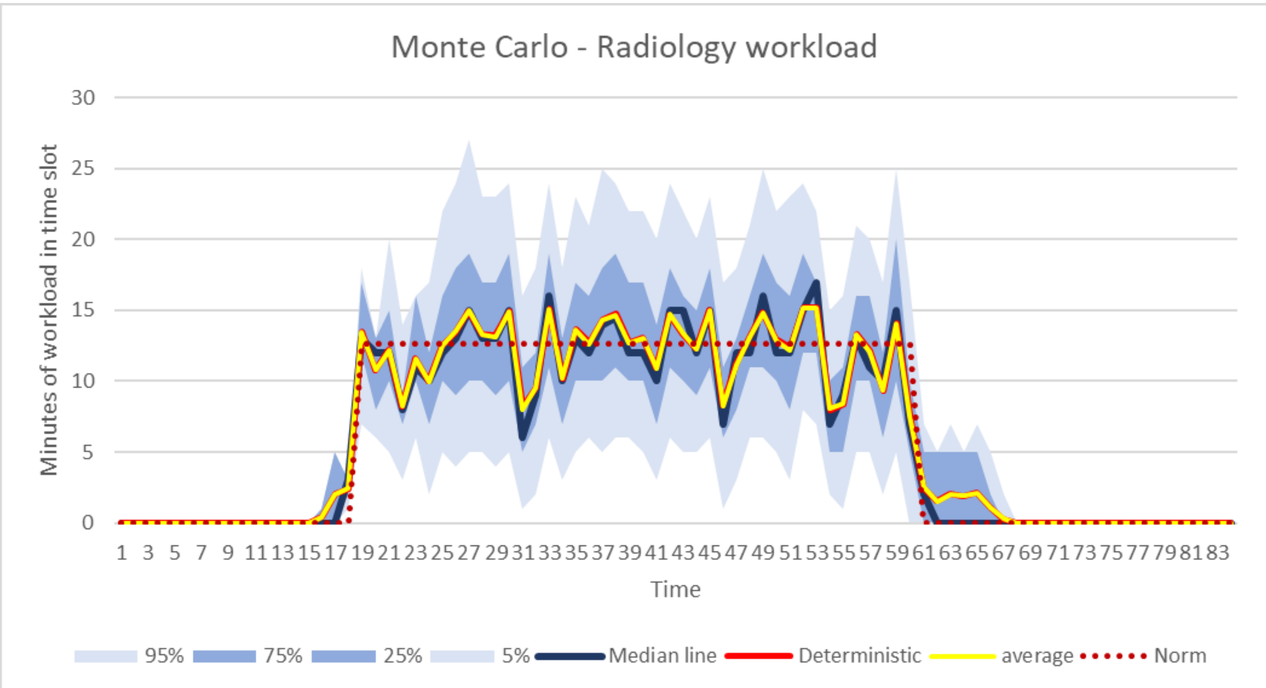

Figure 5.4 – A workload graph for the Radiology department, based on Monte Carlo simulation

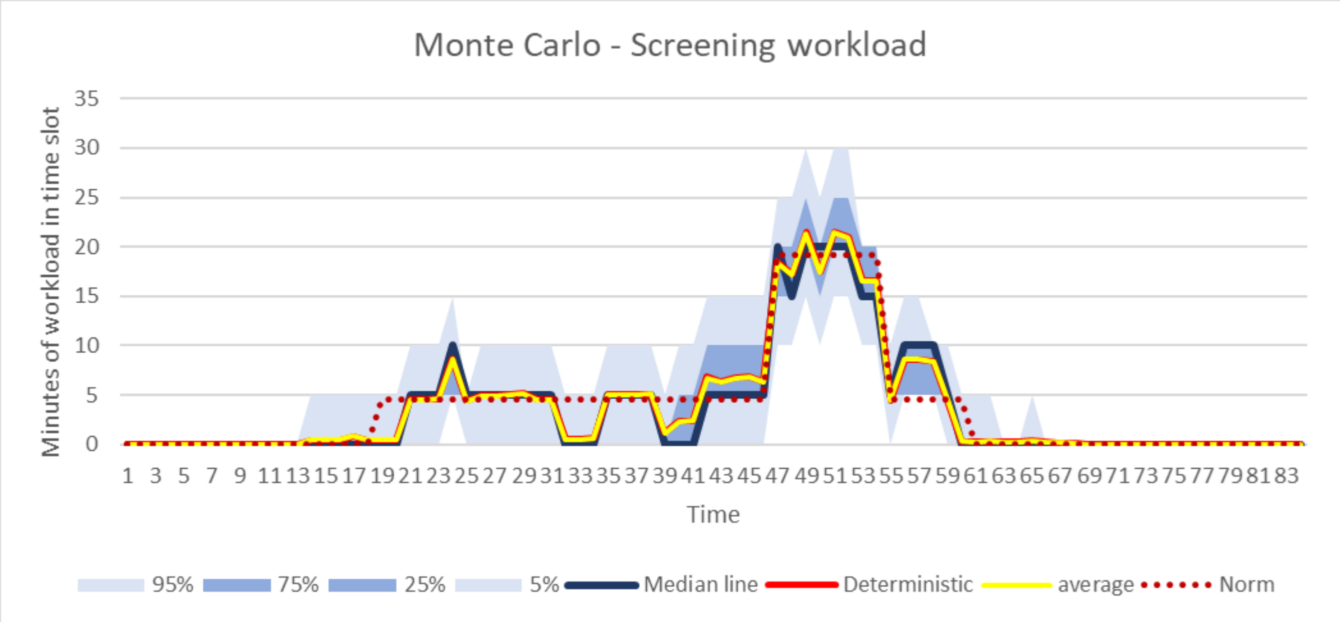

Figure 5.5 – A workload graph for the Screening department, based on Monte Carlo simulation

These two graphs show the average workload predicted by the MILP model (red line) is exactly matched by the average of workload from the Monte Carlo simulation (yellow line), which is correct. The mean and median are very close, and the percentiles are approximately the same distance above or below the median. This means the average workload used by the MILP model can give a reasonably good indication of the workload to expect per department. Some time slots do seem to have a wider spread than others, for example in Figure 5.4 the spread at time slot 19 is more narrow than time slot 27, but not at the level that we deem it a necessity to incorporate it into the model.

Appendix J contains the Monte Carlo graphs for the Thursday afternoon experiment of all departments.

5.6 Conclusion

The model improves on the original schedules significantly. The blueprint schedules generated by the model cause less demand fluctuation, reduce the height of expected demand peaks and can take timing into account, like with the preoperative screening patients to reduce waiting time for patients by simply switching the sequence of appointments. The Monte Carlo simulation is a useful tool to predict the spread of workload around the average, however the results from the MILP model still hold up, and give a good indication of the workload that can be expected.

In addition to this, the hospital is moving towards more standardized blueprint schedules. If identical schedules are used for multiple doctors, the demand patterns will be amplified, and will likely be worse than the current 'original' schedules, since they are currently quite different from each other.

6 Implementation

The model shows promising results with regards to reducing the variability of downstream demand, and Logistiek Bedrijf sees potential in the model. This chapter discusses what steps are needed for implementation of the model. Section 6.1 discusses the additional validation steps. Section 6.2 describes improvements in user interface. Section 6.3 discusses support in the organisation. Section 6.4 describes the first steps if the model is in use.

6.1 Additional validation

Validation with planners and schedulers

The model has been validated to a certain extent. The output of the model is checked with staff from Logistiek Bedrijf, and the schedules generated are generate reasonable schedules. However, the operational planners and schedulers did not have the time during this stage of the project to validate. Additional sanity checks need to be done with more hospital staff, in particular the planners and schedulers since they are aware of the requirements and wishes of doctors the schedules are created for.

Investigate additional doctor's requirements and wishes

Additionally, the doctors are likely to have more requirements and wishes that are not known until a schedule that violates an 'unwritten rule' is proposed. The valid requirements can be added as hard constraints to the MILP model, and wishes can be disregarded, or implemented as soft constraints with a certain penalty that depends on the extend the soft constraint is violated.

Implementation of model extension for DUO and TRIO schedules (synchronised consultations)

The model extension described in Section 4.6.6 is not yet implemented in the model. It allows for linking individual schedules, and defining which consultation types need to be scheduled at the same time for these schedules. By implementing this in the model the model should be ready for overall implementation, however additional checks with planners and schedulers is needed to check if all requirements for DUO and TRIO schedules are met.

6.2 User interface

Improve visual representation of generated schedule to match HiX

At this time, the AIMMS model produces clear graphs on the expected workload per department that results from a schedule, and it specifies for each schedule (e.g. doctor 1, doctor 2) the exact starting time of each consultation, and the consultation type and unit. The consultation duration is also specified.

Currently the generated schedule needs to be manually converted to a blueprint session schedule in scheduling software HiX. A visual representation that matches the user interface for input in HiX is recommended for easier use, and to make input error less likely. This can be achieved in AIMMS or by creating an automatic export to Excel.

Automatic loading of parameters from Excel

Currently the expected workload parameters, the norm, consultation duration etc are manually copied and pasted from many excel files. Merging parameters into one excel file, and configuring AIMMS to automatically load the parameters from this file will increase ease of use. Having all parameters in one file will also increase ease of use.

Input validation

Currently AIMMS will simply state the problem is not feasible if some parameters are set incorrectly. However, some validation on input should be performed to check the consultations actually fit in the schedule, if no consultations overlap in the manual schedule, if consultation durations are defined, etc.

6.3 Organisational support

Support from people in the organisation is critical for successful implementation. The schedulers, doctors, and managers need to be involved in the discussion to present the benefits of the new way of working by using the blueprint session schedules generated by the model. Depending on the hierarchical organisation of the hospital, not all staff need to agree before it can be implemented, but all stakeholders concerns should be considered seriously. Expected benefits of patient waiting time, and more stable workload are hopefully benefits they also want to strive for. The costs of implementation are increased complexity of scheduling, and a different order of consultations for doctors.

The daily operation will not become more complex for the doctors, since the booking of consultations is not performed by them. For the planners that schedule patients in specific slots the situation will be the same, only the sequence of consultations will be different, but this will be defined in HiX. The only step where additional operational complexity is introduced is for the scheduling staff that manages and applies the blueprint session schedules into HiX.

A problem that remains is how to import the generated blueprint session schedules into HiX. At this time no automated way is available, which means generated schedules need to be entered manually which is time consuming. The manual entry will be the main costs for implementation of the model. A suggestion is to have more standardised session schedules ready for the most common combinations. The schedules can be manually edited, and taking the generated schedules as a starting point will not reach the optimal value, it will still improve on the current situation significantly. Changes to the schedules can also be entered into the AIMMS model to evaluate if the changes will result in severe workload variability. If this is not the case, the manual adjustment in HiX can be carried out without completely importing an entirely new schedule into HiX.

The schedules can be made to still adhere to doctors requirements, while balancing the downstream demand simply by changing the sequence and starting time of consultations. It seems plausible the organisation can be convinced of this new way of working.

6.4 Model in use

Use in tactical planning

The model is suitable to use in the tactical level. Manual entry of certain standard blueprint schedules is recommended while automatic import into HiX is not possible.

Staff training

Scheduling staff will need to be trained to use this model. This should be relatively straightforward. The department Logistiek Bedrijf has a lot of experience with these type of models, and is able to understand the working of the model. They can train some scheduling staff on how to use the model, and check with them during the initial stages. The model itself should be relatively easy to use if the user interface improvements are completed. The expected workload

Initial manual checks on generated schedules

Initially the generated schedules will need to be checked by hand. If some problems arise, small adjustments can be made manually. If the number of adjustments is not too high, the performance is likely still better compared to the old blueprint session schedules. The adjusted version can be entered into the model again as a manual schedule, to check if the adjustments do not degrade the downstream demand pattern too much. Once the manual check is ok, the blueprint session schedule can be entered into HiX.

AIMMS license
The hospital already has a license for AIMMS, that is billed based on hours of use. This means the model could be immediately put to use, there is no need to implement the MILP model in different software.

6.5 Conclusion

In short, the most effort for implementation will be in the additional validation, and convincing the organisation on the new way of working. A rough guess for the time needed for additional validation is 4 weeks. Improvements in user interface should be possible in approximately 1-2 weeks. Convincing the organisation is likely possible since the required changes are limited for most staff. Implementation of the model in the hospital seems possible.

7 Discussion

7.1 Conclusions

The current patient flow patterns show downstream workload arrival variability, with peaks during the day at the Radiology department and follow-up consultations at the Orthopaedic Outpatient Department. The waiting time also shows an upward trend in the morning for all departments, but the Orthopaedic Outpatient department shows the strongest increase.

The data analysis shows distinct downstream demand patterns for the consultation types and units, which means the idea to change the sequence of consultations in the blueprint session schedule holds promise, and it could reduce the variability of patients flows across departments.

In the literature review different approaches are considered. Many studies consider only the outpatient department itself when optimizing the consultation schedules. They usually apply a 1-stage or 2-stage stochastic linear program and minimize a weighted sum of waiting time, idle time and overtime. For studies that consider downstream departments in their schedule various approaches exist. Most do not meet certain requirements for this project such as uncertain care pathways. Multi-appointment studies that are similar to our needs use linear programming to solve their scheduling problems, which confirms that creating a MILP model is a valid solution approach.

We define a MILP model that is able to generate blueprint schedules based on a given number of consultation types. It can generate multiple blueprint schedules simultaneously, matching the combined outflow of patients as close to the desired pattern as possible. The model is extended based on feedback, and it provides a relatively flexible and easy way to generate blueprint session schedules.

The results from using the model in experiments are promising. It shows the current blueprint session schedules maximum demand peaks can be reduced to roughly 40% of the original schedule, while smoothing out the patient outflow throughout the day. Adjusting the workload norm works, and the model reacts by scheduling certain consultation types at different times to approach the desired pattern of downstream workload as close as possible.

Implementation of the model will need more validation with planners and schedulers, and the requirements and wishes of doctors with regards to the blueprint session schedule should be inquired. However, additional constraints should be relatively straightforward to add to the model. To put the model into use only few employees need to be trained to work with the model. Once the blueprint session schedule is applied into HiX, the remainder of the processes stay the same for patient planners and doctors. The main usage costs will be the manual entry of schedules into HiX.

To answer the research question:

How can blueprint session schedules be generated for elective patients at the orthopaedic outpatient department, that match the expected pattern of same-day patient flows with the desired pattern?

The blueprint session schedules of the Orthopaedic Outpatient Department can be improved by generating them through the MILP model that is developed, and workload arrival variability will be reduced, the maximum peaks of demand can be reduced to approximately 40% of the original schedule.

Our approach of generating patient blueprint schedules for outpatient department is different from other studies. The consideration of downstream departments in outpatient scheduling is present in the literature, however the consideration of opening and closing hours into the model, and the ability to set different norms per time slot for each of the downstream departments is not discussed in the papers found in literature review papers. It can be used to manage patient flows during the day, and the model is flexible. If other departments are closed at certain times of day, or have reduced capacity available for walk-in patients, this model can spread the patients differently over the day to match supply (staff availability) and demand (expected patient workload).

7.2 Limitations

This research is has some limitations which are discussed in this section.

Deterministic consultation starting time and duration at the OOD itself
The scope of the research is limited to the patient outflow to other departments, so in order to limit the complexity we chose to model the part of the OOD itself as deterministic: consultation starting time and duration are assumed to be deterministic in our model. In reality, the delays that are currently occurring over the day will still occur when the new blueprint schedules are used. Consultation starts might be delayed, and consultations can take longer than planned. We still expect the patient flow pattern to achieve roughly similar or a slight degradation of performance as calculated in the experiments. Demand will likely shift more towards later time slots as the delay increases, but the same sequence of consultations will still be carried out.

The delays could be reduced by setting more appropriate *consultation durations*, and thus reduce the chance the planned duration will be exceeded, and reduce delays in the OOD in this way. The stochastic models that optimize the schedule of only the outpatient department typically schedule some slack (empty slots) to allow the doctor to catch back up to the schedule and reduce delays later in the day.

Unpunctual patients and all patients no-shows
Patients are assumed to be punctual in our model, and the no shows are not considered. If patient tardiness is high, and/or the no show rate is high, different schedule designs can be considered such as planning two patients per slot, and making the block twice the duration. The first patient that shows up gets server. This increases the waiting time for the patients that arrive second, but increases capacity utilization and decreases doctor idle time. The no show rate is approximately 1.5%, and patients punctuality was not an issue based on interviews, so we expect this to not be an likely issue.

Deterministic transitions
The model works with average transition probabilities. The model will consider a 100% probability on 30 minutes of workload identical to two times 50% probability of 30 minutes of workload. If a consultation type exists that is rare, but *if* a transition occurs a lot of workload is added this can lead to some sub optimal performance.

For example a consultation type with 5% probability of 2 hours of work. In a stochastic model the consultations would probably be scheduled to not overlap because of the chance the workload of both consultations is realised. The deterministic model uses the expected value of *probability * workload* and allows the two consultations to be scheduled at the same time.

This is mainly an issue when these transitions add *high workload*, have a *low transition probability* and are *rare* in the schedule, so scheduling them together can be avoided. In our case, the workload to the Plaster room has the longest consultation durations, but many consultations have somewhat comparable workload for the Plaster room, so the outflow is not as *rare* and there is not as much room to separate these consultation types. It will still be interesting to extend the model into a Stochastic Linear Program and evaluate transitions on a stochastic basis, however we expect the current deterministic to deliver quite good results, since not many *high workload*, *low transition probability* and mostly *rare* consultations are present.

Data quality – linking of appointments, timestamps and unit classification
The expected workload demand patterns are derived from the data set. However missing timestamps occur often (approximately 50% of the activities have reliable timestamps). The linking of appointments was also performed manually, since the dataset did not indicate which appointments were part of the same clinical issue for the patient. Additionally, the unit could not always be classified based on the activity code. The data analysis has corrected for this in a sensible way (by preserving the proportions) but higher accuracy of the downstream demand profile can be achieved with better data. This is only an issue for the calculation of the parameters however. If better data is available, the parameters simply need to be recalculated and entered into the model. The model itself will be unaffected and can continue to be used.

Similarly, if the demand characteristics change over time, the parameters can be updated and the model will be accurate again, no changes to the model itself will be needed.

7.3 Future research

Two aspects are of interest for further research, which are discussed in the following sections.

7.3.1 Stochastic version of the model

The most interesting area of further research will be to compare the current model with one where stochastic transitions are considered. The current model delivers relatively good results quite quickly, even within one minute. The run time seems to allow for extension into a Stochastic Linear Program.

Does the output change a lot compared to the current model? We expect that consultations with low transition probability, high workload, and that are relatively rare (so overlap can be avoided) to be scheduled farther apart. It will be interesting to see how much the schedules improves. This could be done by generating a blueprint schedule with the current model, and scoring it by the stochastic model. Then let the stochastic model generate the schedule with the same constraints to check what performance increase it predicts.

We expect it the improvement to be relatively small, especially for the current case. It will be interesting if the additional run time is worth it. A few components that can be evaluated as stochastic:

1. The transitions to downstream departments
2. The service times of the main department
3. The downstream demand durations
4. The effect of stochastic travel time to downstream departments

The Monte Carlo simulation model discussed in Section 5.5 takes the first component (stochastic transitions) into account, and it can be extended to evaluate the other components in a stochastic manner as well.

7.3.2 Using the model at the operational level

The initial implementation of the model would be to generate the blueprint session schedules that are applied approximately twelve weeks in advance. However in those twelve weeks many changes can occur. Maybe a different doctor from a different unit is scheduled, or the number of specific consultation types change. An example is when a doctor performs more surgeries then planned, more 'discharge' consultations will be needed.

When these changes occur, the changes will are made manually. A possible use of the model is to update the parameters, fix the consultations that have a patient booked, and let the model optimize the remaining consultations. When many consultations are already booked, those starting positions will be fixed, and the solution space for the new run of the model will be smaller, leading to short run times.

A trade-off needs to be made if the additional effort of running the model again is worth the improvements in patient flow, but it an application to consider. Perhaps for small adjustments it is not worthwhile, but if doctors are switched it is probably more worthwhile. A necessity is a good interface with HiX however, automatic import of the generated schedules might be necessary, or the time needed is likely to be too expensive.

7.4 Recommendations

We recommend Logistiek Bedrijf to implement the model and use it to generate the blueprint session schedules. Some additional validation is required, but no major issues are expected there. The implementation has potential since for most staff the day-to-day process would not change, only for the scheduling staff that creates and applies the blueprint session schedules. The main consideration is the amount of time needed to enter the blueprint schedules into HiX, my recommendation is to generate generic session schedules and enter those into HiX.

If the model is implemented, it is advisable to perform periodic checks if the parameters are still accurate, and update them if needed.

Additionally, we advise to look into data quality of *timestamps*, *unit classification* and the identification of *linked appointments* throughout the day. If the timestamps and identification of linked appointments is improved, the experience of a patient with multiple appointments can be investigated and possibly improved. The time a patient spends between appointments on a single is interesting, since it is not directly recorded. The waiting time only starts once the patient enters the waiting room for the next appointment, but he or she might be waiting elsewhere in the hospital.

The system for timestamp registration of the Preoperative screening department is not suited for the many steps the patient moves through. The current system is only suited for a single appointment. Increased quality of timestamps can reveal additional areas of improvement, and will give a more reliable indicator of patient and staff experience.

Implementation of these recommendations will contribute to operational excellence of the Sint Maartenskliniek.

Bibliography

Ahmadi-Javid, A., Jalali, Z., & Klassen, K. J. (2017). Outpatient appointment systems in healthcare: A review of optimization studies. *European Journal of Operational Research*, *258*(1), 3–34. https://doi.org/10.1016/j.ejor.2016.06.064

Dharmadhikari, N., & Zhang, D. J. (2013). *Simulation Optimization of Blocking Appointment Scheduling Policies for Multi-Clinic Appointments in Centralized Scheduling Systems*. *2*(11), 6.

Fügener, A., Hans, E. W., Kolisch, R., Kortbeek, N., & Vanberkel, P. T. (2014). Master surgery scheduling with consideration of multiple downstream units. *European Journal of Operational Research*, *239*(1), 227–236. https://doi.org/10.1016/j.ejor.2014.05.009

Gupta, D., & Denton, B. (2008). Appointment scheduling in health care: Challenges and opportunities. *IIE Transactions*, *40*(9), 800–819. https://doi.org/10.1080/07408170802165880

Hirsch, W. Z. (1957). *Introduction to Modern Statistics: With Applications to Business and Economics*. Retrieved from https://books.google.nl/books?id=KostAAAAIAAJ

Hulshof, P. J. H., Kortbeek, N., Boucherie, R. J., Hans, E. W., & Bakker, P. J. M. (2012). Taxonomic classification of planning decisions in health care: a structured review of the state of the art in OR/MS. *Health Systems*, *1*(2), 129–175. https://doi.org/10.1057/hs.2012.18

Leeftink, A. G., Bikker, I. A., Vliegen, I. M. H., & Boucherie, R. J. (2018). Multi-disciplinary planning in health care: a review. *Health Systems*, 1–24. https://doi.org/10.1080/20476965.2018.1436909

Leeftink, A. G., Boucherie, R. J., Hans, E. W., Verdaasdonk, M. A. M., Vliegen, I. M. H., & van Diest, P. J. (2018). Batch scheduling in the histopathology laboratory. *Flexible Services and Manufacturing Journal*, *30*(1–2), 171–197. https://doi.org/10.1007/s10696-016-9257-3

Leeftink, A. G., Vliegen, I. M. H., & Hans, E. W. (2019). Stochastic integer programming for multi-disciplinary outpatient clinic planning. *Health Care Management Science*, *22*(1), 53–67. https://doi.org/10.1007/s10729-017-9422-6

Liang, B., Turkcan, A., Ceyhan, M. E., & Stuart, K. (2015). Improvement of chemotherapy patient flow and scheduling in an outpatient oncology clinic. *International Journal of Production Research*, *53*(24), 7177–7190. https://doi.org/10.1080/00207543.2014.988891

Marynissen, J., & Demeulemeester, E. (2019). Literature review on multi-appointment scheduling problems in hospitals. *European Journal of Operational Research*, *272*(2), 407–419. https://doi.org/10.1016/j.ejor.2018.03.001

Metropolis, N., & Ulam, S. (1949). The Monte Carlo Method. *Journal of the American Statistical Association*, *44*(247), 335–341. https://doi.org/10.1080/01621459.1949.10483310

Sint Maartenskliniek | cijfers jaarverslag 2017. (n.d.). Retrieved 2 January 2019, from https://www.infographic.nl/ext/smk2018/

van Essen, J. T., Bosch, J. M., Hans, E. W., van Houdenhoven, M., & Hurink, J. L. (2013). Reducing the number of required beds by rearranging the OR-schedule. *OR Spectrum*. https://doi.org/10.1007/s00291-013-0323-x

Vanberkel, P. T., Boucherie, R. J., Hans, E. W., Hurink, J. L., & Litvak, N. (2010). *A Survey of Health Care Models that Encompass Multiple Departments*. 49.

Appendix

A. Appendix data set columns

The data set available for analysis of the outpatient appointments contains the following information for each appointment:

Column name	English translation	Description
AfspraakNr	Appointment id	The unique id for an activity or appointment (anonymised)
PatientNr	Patient id	The unique id for a patient (anonymised)
Locatie	Location	The 2 or 3 letter location code, it specifies which hospital, and the location within the hospital
LocatieOmschrijving	Location description	The location description in text
Specialisme	Specialism	The medical specialism, e.g. 'ORTH' for Orthopaedic
Agenda	Calendar	The calendar code for the main category
AgendaNaam	Calendar name	The calendar name for the main category
SubAgenda	Sub calendar	The code for the sub calendar. The sub calendar is where actual appointments can be booked into. Each sub calendar belongs to one main calendar type.
SubAgendaNaam	Sub calendar name	The name of the sub calendar. Usually the name of the doctor that the patient appointment is booked to.
AfspraakDatum	Appointment date	The date of the appointment or activity
AfspraakBeginTijd	Appointment start time	The scheduled starting time of an appointment
AfspraakDuur	Appointment duration	The planned duration of the appointment or activity
AfspraakCode	Appointment code	The 2-6 letter appointment code that specifies the exact activity type
AfspraakCodeOmschrijving	Appointment code description	The description of activity type in text
AANKOMST	Arrival	The arrival timestamp, to indicate when the patient arrives at waiting room of the department for this activity.
OPROEP	Call-in moment	The timestamp of the patient being called in for the start of the activity.
VERTREK	Departure	The timestamp of the end time of the activity.

AfspraakType	Appointment type	A single letter code to indicate if the activity is an appointment for a single patient, a group appointment.
Voldaan	Completed	Indicates if the appointment or activity has been completed or cancelled
RedenNietVoldaan	Reason not completed	The reason in text why the appointment or activity was cancelled

B. Appendix data analysis description

The source data are all activities (consultations, scans, screenings, etc.) of all locations of the Sint Maartenskliniek. Each activity is one row in the dataset.

The source data contains **524 086** data rows, from 02-01-2017 to 31-12-2018.

i. Location filtering

The hospital group has multiple locations throughout the Netherlands. We use the orthopaedic outpatient department at Nijmegen for the case study, so we isolate that data. Based on discussion with Logistiek Bedrijf the following highlighted location codes are selected.

A	B	BK	BO	BQ	C	CR	DO	GE	KL	LV	**N**	**ND**
	NF	**NG**	**NK**	NM	**NO**	**NP1**	**NQ**	**NV**	O	OG	OL	OM
	OW	P	RF	RH	T	VN	W	ZG				

This results in **394 342** data rows.

ii. Department filtering

The dataset contains multiple departments at location Nijmegen. We select only the OOD and departments a patient might visit on the same day as the OOD consultation.

These departments are:

- Orthopaedic Outpatient Department (OOD / ORT)
- Radiology department (RAD)
- Plaster and Wound treatment (GIPS)
- Preoperative screening (PREO)

The dataset contains columns that indicate medical specialism, and the name of the calendar appointments are booked into. A combination of both is used to define the department of the activity. The code below is used to classify the department of each activity. The first condition that matches is applied.

```
CASE
  WHEN t1.AfspCode="BMDPHA"
    THEN "Multi-disciplinair"
  WHEN t1.Specialisme="GIPS" AND t1.AgendaNaam="GIPS"
    THEN "Gipskamer"
  WHEN t1.Specialisme="ORT" AND t1.AgendaNaam="PREO"
    THEN "Pre-operatief onderzoek"
  WHEN t1.Specialisme="RAD" AND t1.AgendaNaam="Radiologie Nijmegen"
    THEN "Radiologie"
  WHEN t1.Specialisme="ORT" AND (t1.AgendaNaam="NPPA" OR t1.AgendaNaam="ORTH")
    THEN "Orthopedie poli"
```

```
  WHEN t1.Specialisme="REU" AND (t1.AgendaNaam="NPPA REUM" OR t1.AgendaNaam="REUM")
    THEN "Reumatologie"
  WHEN t1.Specialisme="REV" AND (t1.AgendaNaam="NPPA REV" OR t1.AgendaNaam="REV")
    THEN "Revalidatie"
  WHEN t1.Specialisme="ORT" AND (t1.AgendaNaam="NPPA REV" OR t1.AgendaNaam="REUM" OR
t1.AgendaNaam="REV")
    THEN "Multi-disciplinair"
  WHEN t1.Specialisme="REU" AND t1.AgendaNaam="ORTH"
    THEN "Multi-disciplinair"
  WHEN t1.Specialisme="REV" AND (t1.AgendaNaam="ORTH" OR t1.AgendaNaam="REUM")
    THEN "Multi-disciplinair"
  ELSE "Onbekend"
END
```

A corner case are the emergency and urgent care activities (‘spoed’) at the outpatient department. These can be filtered based on the activity code. Since the urgent care outpatient clinic operates entirely separate from the regular outpatient clinic, we filter the urgent and emergency patients with the following regular expression:

```
CASE
  WHEN PRXMATCH('(SC[^AF]\w)', t1.AfspCode) > 0
  THEN "Spoed poli"
  ELSE "Niet spoed poli"
END
```

Meaning all activity codes that start with “SC” are labelled as urgent outpatient clinic (‘spoedpoli’), except for activity code “SCAF” which are preoperative screening appointments. This matches 1081 cases in the dataset.

After filtering other departments and the urgent outpatient clinic activities, **271 778 rows are left.**

iii. Unit classification

The activities for the OOD can be labelled with the unit (e.g. Knee). We label the OOD activities with the following regular expressions:

```
CASE
  WHEN t1.Afdeling="Orthopedie poli" AND (t1.AfspCode="POSC" OR t1.AfspCode="POSK" OR
t1.AfspCode="VCSCH") THEN "Not classified"
  WHEN t1.Afdeling="Orthopedie poli" AND PRXMATCH('(\w{2}BE)', t1.AfspCode)>0 THEN "Bovenste Extremiteit"
  WHEN t1.Afdeling="Orthopedie poli" AND PRXMATCH('(\w{2}HE)', t1.AfspCode)>0 THEN "Heup"
  WHEN t1.Afdeling="Orthopedie poli" AND PRXMATCH('(\w{2}KN)', t1.AfspCode)>0 THEN "Knie"
  WHEN t1.Afdeling="Orthopedie poli" AND PRXMATCH('(\w{2}SC)', t1.AfspCode)>0 THEN "Rug"
  WHEN t1.Afdeling="Orthopedie poli" AND PRXMATCH('(\w{2}VO)', t1.AfspCode)>0 THEN "Voet"
  WHEN t1.Afdeling="Orthopedie poli" AND PRXMATCH('(\w{2}WE)', t1.AfspCode)>0 THEN "Rug"
  WHEN t1.Afdeling="Orthopedie poli" AND PRXMATCH('(\w{2}WK)', t1.AfspCode)>0 THEN "Rug"
  WHEN t1.Afdeling="Orthopedie poli" AND PRXMATCH('(\w{2}CW)', t1.AfspCode)>0 THEN "Rug"
  WHEN t1.Afdeling="Orthopedie poli" AND PRXMATCH('(\w{2}VO)', t1.AfspCode)>0 THEN "Voet"
  WHEN t1.Afdeling="Orthopedie poli" AND PRXMATCH('(\w{2}HS)', t1.AfspCode)>0 THEN "Heup"
  WHEN t1.Afdeling="Orthopedie poli" AND PRXMATCH('(\w{2}KS)', t1.AfspCode)>0 THEN "Knie"
  WHEN t1.Afdeling="Orthopedie poli" AND PRXMATCH('(\w{2}VS)', t1.AfspCode)>0 THEN "Voet"
  ELSE "Not classified"
END
```

For some activity codes the unit cannot be determined (e.g. POSC -> preoperative screening). In these cases the label ‘Not classified’ is applied.

AfspCode Group	Unit	Percentage
1. ORT New	Unit 1	12,9%
1. ORT New	Unit 2	16,3%
1. ORT New	Unit 3	26,4%
1. ORT New	Not classified	3,0%
1. ORT New	Unit 5	24,7%
1. ORT New	Unit 6	16,7%
2. ORT Repeat	Unit 1	13,6%
2. ORT Repeat	Unit 2	11,5%
2. ORT Repeat	Unit 3	24,1%
2. ORT Repeat	Not classified	19,4%
2. ORT Repeat	Unit 5	16,6%
2. ORT Repeat	Unit 6	14,7%
3. ORT Discharge	Unit 1	26,2%
3. ORT Discharge	Unit 2	21,9%
3. ORT Discharge	Unit 3	36,5%
3. ORT Discharge	Not classified	0,1%
3. ORT Discharge	Unit 5	8,6%
3. ORT Discharge	Unit 6	6,8%
4. ORT POP	Not classified	100%

Figure 7.1 – Unit not classified

iv. Consultation type classification

The OOD consultations can be grouped into 4 main types that will be used in the blueprint session schedules:

- New
- Repeat
- Discharge
- Preoperative conversation with physician (POP)

Some other activities like a consultation by phone ('telefonisch consult') or taking a blood sample ('Prik') are distinguished. For the Radiology department a distinction is made between scans that are done on a walk-in basis, and scans that are only performed with pre-booked appointments. Especially the classification of level 1 CT scan is complicated and requires the following extensive regular expression: '(CT00(\d\d)(\w*)|CT1301(\w*)|CT20(\d\d)(\w*)|CT30(\d\d)(\w*)|CT40(\d\d)(\w*)|CT50(\d\d)(\w*)|CT700(\d)(\w*)|CT900(\d)(\w*)|CT9011(\w*)|CT910(\d)(\w*)|CT911(\d)(\w*)|CT9202(\w*))'

We achieve this with classification based on the activity code:

```
CASE
  WHEN t1.Afdeling="Orthopedie poli" AND UPPER(t1.AfspCode) LIKE UPPER('BP%') THEN '1. ORT New'
  WHEN t1.Afdeling="Orthopedie poli" AND UPPER(t1.AfspCode) LIKE UPPER('HNR') THEN '2. ORT Repeat'
  WHEN t1.Afdeling="Orthopedie poli" AND UPPER(t1.AfspCode) LIKE UPPER('NK%') THEN '1. ORT New'
  WHEN t1.Afdeling="Orthopedie poli" AND UPPER(t1.AfspCode) LIKE UPPER('OC%') THEN '3. ORT Discharge'
  WHEN t1.Afdeling="Orthopedie poli" AND UPPER(t1.AfspCode) LIKE UPPER('POS%') THEN '4. ORT POP'
  WHEN t1.Afdeling="Orthopedie poli" AND UPPER(t1.AfspCode) LIKE UPPER('SC%') THEN '1. ORT New'
  WHEN t1.Afdeling="Orthopedie poli" AND UPPER(t1.AfspCode) LIKE UPPER('SO%') THEN '1. ORT New'
  WHEN t1.Afdeling="Orthopedie poli" AND UPPER(t1.AfspCode) LIKE UPPER('VC%') THEN '2. ORT Repeat'
  WHEN t1.Afdeling="Orthopedie poli" AND UPPER(t1.AfspCode) LIKE UPPER('TC%') THEN 'ORT Telephone consult'
  WHEN t1.Afdeling="Orthopedie poli" AND UPPER(t1.AfspCode) LIKE UPPER('BELCON%') THEN 'ORT Telephone consult'
  WHEN t1.Afdeling="Orthopedie poli" AND UPPER(t1.AfspCode) LIKE UPPER('PRIK') THEN 'ORT PRIK'
  WHEN t1.Afdeling="Radiologie" AND UPPER(t1.AfspCode) LIKE UPPER('CR%') THEN 'RAD Bucky + CT niveau 1'
  WHEN t1.Afdeling="Radiologie" AND UPPER(t1.AfspCode) LIKE UPPER('MR%') THEN 'RAD Overig'
  WHEN t1.Afdeling="Radiologie" AND UPPER(t1.AfspCode) LIKE UPPER('US%') THEN 'RAD Overig'
  WHEN t1.Afdeling="Radiologie" AND UPPER(t1.AfspCode) LIKE UPPER('RF%') THEN 'RAD Overig'
  WHEN t1.Afdeling="Radiologie" AND UPPER(t1.AfspCode) LIKE UPPER('VE%') THEN 'RAD Overig'
  WHEN t1.Afdeling="Radiologie" AND UPPER(t1.AfspCode) LIKE UPPER('WO%') THEN 'RAD Overig'
  WHEN t1.Afdeling="Radiologie" AND
(PRXMATCH('(CT00(\d\d)(\w*)|CT1301(\w*)|CT20(\d\d)(\w*)|CT30(\d\d)(\w*)|CT40(\d\d)(\w*)|CT50(\d\d)(\w*)|CT700(\d)(\w*)|CT900(\d)(\w*)|CT9011(\w*)|CT910(\d)(\w*)|CT911(\d)(\w*)|CT9202(\w*))', t1.AfspCode) > 0)
```

```
  THEN 'RAD Bucky + CT niveau 1'
WHEN t1.Afdeling="Radiologie" AND UPPER(t1.AfspCode) LIKE UPPER('CT%') THEN 'RAD Overig'
WHEN t1.Afdeling="Gipskamer" THEN 'GIPS all'
WHEN t1.Afdeling="Pre-operatief onderzoek" THEN 'PREO all'
WHEN t1.Afdeling="Multi-disciplinair" THEN 'Multi disciplinary all'
ELSE "Not classified"
```

v. Timestamp reliability

Timestamp input for patient arrival in the waiting room, being called in by the physician, and the end time of the activity are all manually recorded. In interviews staff indicated timestamps are not always entered correctly. We check if timestamps are reliable. The following 3 timestamps are available:

1. Patient arrival (AANKOMST)
2. Patient called in (OPROEP)
3. Patient leaves consultation/activity (VERTREK)

The timestamps of an activity are only marked 'reliable' when:

- The time between a patient arriving (AANKOMST) and being called in (OPROEP) is positive or zero, i.e. a patient cannot be called in before he/she arrives.
- The time between a patient being called in (OPROEP) and leaving (VERTREK) and is positive, i.e. an activity needs to last at least 1 minute
- The time between a patient arriving (AANKOMST) and the patient leaving (VERTREK) is positive, i.e. the total time a patient is present at the department is at least 1 minute

All negative durations, and other occurrences with impossible sequences, or when timestamps are missing, in all these cases the timestamps of these activities are labelled 'not reliable'.

Afdeling	Oproep_min_aankomst	Vertrek_min_oproep	Vertrek_min_aankomst	Count
Gipskamer	Positive	Positive	Positive	13006
Gipskamer	Missing	Missing	Missing	8276
Gipskamer	Missing	Positive	Missing	3666
Gipskamer	Identical	Positive	Positive	1612
Gipskamer	Positive	Missing	Missing	1127
Gipskamer	Missing	Identical	Missing	316
Gipskamer	Identical	Missing	Missing	114
Gipskamer	Positive	Identical	Positive	67
Gipskamer	Identical	Identical	Identical	61
Gipskamer	Missing	Missing	Positive	56
Gipskamer	Missing	Missing	Identical	15
Gipskamer	Negative	Positive	Negative	2
Gipskamer	Negative	Missing	Missing	1
Gipskamer	Missing	Negative	Missing	1
Orthopedie poli	Positive	Positive	Positive	60138
Orthopedie poli	Missing	Missing	Missing	34603
Orthopedie poli	Missing	Missing	Positive	4734
Orthopedie poli	Positive	Missing	Missing	4356
Orthopedie poli	Identical	Positive	Positive	2334
Orthopedie poli	Positive	Identical	Positive	1181
Orthopedie poli	Missing	Positive	Missing	1110
Orthopedie poli	Missing	Missing	Identical	211
Orthopedie poli	Identical	Identical	Identical	153
Orthopedie poli	Identical	Missing	Missing	133
Orthopedie poli	Missing	Identical	Missing	130
Orthopedie poli	Negative	Positive	Negative	1
Orthopedie poli	Negative	Missing	Missing	1
Pre-operatief onderzoek	Missing	Missing	Missing	9925
Pre-operatief onderzoek	Positive	Positive	Positive	1038
Pre-operatief onderzoek	Missing	Positive	Missing	12
Pre-operatief onderzoek	Missing	Missing	Positive	10
Pre-operatief onderzoek	Positive	Missing	Missing	5
Pre-operatief onderzoek	Identical	Positive	Positive	5
Pre-operatief onderzoek	Positive	Identical	Positive	5
Pre-operatief onderzoek	Identical	Missing	Missing	1
Pre-operatief onderzoek	Missing	Identical	Missing	1
Radiologie	Positive	Positive	Positive	78030
Radiologie	Missing	Missing	Missing	24781
Radiologie	Positive	Identical	Positive	13661
Radiologie	Identical	Positive	Positive	5016
Radiologie	Identical	Identical	Identical	1046
Radiologie	Missing	Missing	Negative	353
Radiologie	Missing	Missing	Positive	304
Radiologie	Missing	Positive	Missing	72
Radiologie	Positive	Negative	Negative	57
Radiologie	Missing	Identical	Missing	18
Radiologie	Negative	Positive	Negative	9
Radiologie	Negative	Positive	Identical	9
Radiologie	Identical	Negative	Negative	7
Radiologie	Negative	Positive	Positive	3
Radiologie	Positive	Negative	Positive	3
Radiologie	Missing	Missing	Identical	2

Table 7.1 – Overview of timestamp quality per department

vi. Completed activities

For the analysis we do not need the cancelled appointments. Only 1.5% of activities were cancelled. We select only the completed activities which results in **267 585** data rows.

vii. Identify source appointments

We assume the first OOD consultation of the day for a patient is the consultation that ‘causes’ the other demand for that patient on the same day. We call this consultation the ‘source consultation’, because we assume in the model this consultation is the source for other same-day visits by the same patient.

We exclude a few consultation types as possible source activities:

- HNR – revisit at the OOD after a Radiology visit
- PRIK – taking a blood sample
- BELCON – consultation by phone
- TC – consultation by phone

With these four excluded, we find the consultation with the earliest scheduled starting time on each day per patient. We perform an INNER JOIN where the following three match:

- Date
- Patient id
- Consultation start time

Now the data is enriched again with all original columns (like department, activity type, etc.).

viii. Determine the linked appointments from each source appointment

For each source appointment, we check for each department if the same patient has any other appointments/activities on the same day.

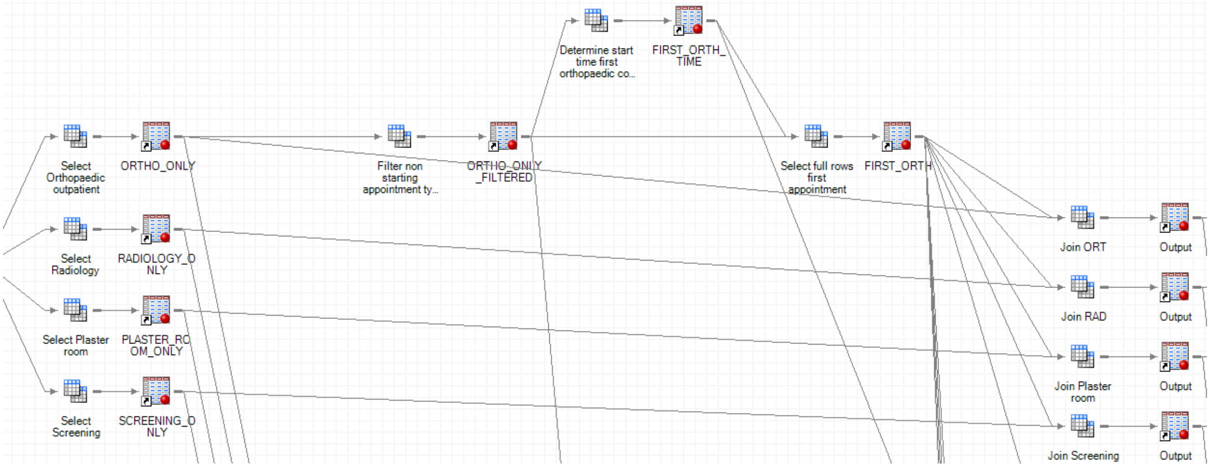

Figure 7.2 – Overview of the structure of source appointment and join data analysis

We perform a LEFT JOIN with the source appointments as the source (left) table where the following match:

- Date
- Patient id

If there is no match, that patient did not have any walk-in visits to other departments. If there is one match the information of that other consultation is added to the same row.

ix. Downstream demand characteristics

In order to determine the demand associated with a certain consultation type three components are required:

- The probability a patient visits another department on the same day as an outpatient consultation
- The minutes of workload per transition
- The travel time or delay between the consultation and the downstream demand

This section discusses the first part, the transition probabilities. Given that a patient comes in for an outpatient consultation, what is the probability he/she visits another department on the same day? And does this typically occur before or after the outpatient clinic consultation? Based on interviews we expect some source consultation types more often cause demand at other departments beforehand, other consultation types afterwards.

x. Determining transition percentage before and after the consultation

Based on the linked appointments described in Section viii, we determine if the secondary activity at the downstream department occurred before or after the source consultation.

- Before: when the 'patient called in' timestamp of the linked appointment is *before* the 'patient called in' of the source appointment
- After: when the 'patient called in' timestamp of the linked appointment is *after* the 'patient called in' of the source appointment

As an example, with 1000 source appointments, 130 might be labelled as 'before', and 180 might be labelled as 'after'. Then the transition probabilities are 13% and 18% to before and after this source appointment.

However, sometimes the timestamps are missing. These linked appointments cannot be ignored, because the exact timing might be unknown, but it is certain the transition occurred. If the transitions with missing timestamps are ignored, the transition probability is underestimated.

We correct for the missing timestamps in the following way:

4. Count all transitions 'before', 'after', and with missing timestamps
5. Determine the proportion of 'before' and 'after'
6. Add the transitions with missing timestamps to the 'before' and 'after' categories, while preserving the proportions calculated at step 2

			Original data			
Source appointments			*Before*	*After*	*Missing*	*Total transitions*
10000			2300	1200	1500	5000
			23%	12%	15%	50%
			Redistribute missing timestamp transitions			
			Before	*After*		
		Percentage of available transition timestamps	65.7%	34.3%		
		Add missing timestamps proportionally	985.7	514.3		
		New total	3285.7	1714.3		
		Check still same proportions	65.7%	34.3%		
			Calculate corrected transition percentage			
			Before	*After*		
			32.9%	17.1%		

Figure 7.3 – Example of correction for missing transition timestamps

Figure 7.3 shows an example of this correction method, and it arrives at the correct total transition probability of 50% (32.9% + 17.1%). If the correction was not applied, the total transition probability would have been underestimated at 35% (23% + 12%).

The source appointments can be all OOD consultations, which means the calculated transition probabilities are the overall transition probabilities. The interesting analysis for our research is the transition probabilities per source consultation type and source unit. Based on the analysis the following questions can be answered:

- What is the probability a patient visits Radiology *before* a 'New consultation' for unit 'Knee'?
- What is the probability a patient visits the Plaster room *after* a 'Repeat consultation' for unit 'Spine?
- Etc.

For each set of consultation type and unit the transition probability analysis is performed. However, yet another correction needs to be made for the source appointments where unit is not classified. They are a significant portion of the total number of source appointments, especially for the Plaster room (49.7%) and Preoperative screening (76.6%).

The correction works similar to the correction for missing timestamps mentioned above.

1. Select data from one department (e.g. Radiology)
2. Calculate the proportions of source units (e.g. 7% from Knee, 15% not classified)
3. Recalculate the proportions without unit 'not classified' (9% Knee, etc.).
4. Now all the data that belonged to unit 'not classified' is redistributed to the other units, while keeping the proportion of units the same.
5. Recalculate the transition probability of the other units with the added data from unit 'not classified'

Source unit	%	% without not classified
Unit 1	5,5%	11,0%
Unit 2	0,9%	1,8%
Unit 3	18,8%	37,3%
Not classified	49,7%	
Unit 4	2,8%	5,6%
Unit 5	22,3%	44,2%

Table 7.2 - The unit proportion calculation for the Plaster room department

Finally, the transition probabilities per set of source consultation type and source unit are calculated.

xi. Travel time

When a patient visits a different department before or after the source appointment, some delay is expected before the patient arrives at the other department, simply because the patient has to travel there. We need to incorporate this in the model for a more accurate prediction of workload arrival. In this analysis the median travel time is used to limit the effect of outliers and unusual situations.

The travel time is calculated as the time between the patient leaving department one and arriving at department two. For the departments Radiology, Plaster room, and Preoperative screening we only use a different travel time for *before* and *after* the source consultation, but make no distinction between source consultation types and source units. The detailed analysis was performed but figures were very similar (fluctuations of one or two minutes).

RAD before/after OOD appointment?	Median travel time
Before	5 minutes
After	4 minutes

Plaster before/after OOD appointment?	Median travel time
Before	20 minutes
After	7 minutes

Screening before/after OOD appointment?	Median travel time*
Before	5 minutes
After	5 minutes

*figures not reliable in data, but Screening counter is at same location as OOD, so 5 minutes for travel + arriving at the desk for filling in forms is assumed.

For the Orthopaedic Outpatient department we analyse different travel times per source consultation type and source unit combination, since those revisits somewhat often occur after a visit to a different department. This means the service time of the other departments is also included in the travel time to a follow up activity at the OOD.

For the analysis of the travel times to OOD, not all combinations of source consultation type and source unit have a big enough sample size for a reliable estimation. If 10 or less observations are seen we disregard the calculated travel time of the 'consultation type – unit' set, and overall median of that unit is used. If the overall travel time of that unit also does not include enough data points, the median of all observations that also happened *before* the source appointment is taken (if this appointment occurred *before,* if this appointment also occurred before).

xii. Downstream activity duration

Now that the transition probability and travel time are determined, the amount of demand needs to be determined. Given that a patient visits a downstream department, how many minutes of work does this add to the department?

For the duration the median is used in order to limit the effect of outliers in the data. The duration is specified on three levels: *before* or *after* the source consultation, the source consultation type, and the source unit.

xiii. Combining results into the final demand profiles

The three key items are analysed: the transition probabilities, travel time, and activity durations. For set of consultation type and unit a demand profile is calculated for all 4 downstream departments. Table 7.3 – Workload spreading (expected workload). Table 7.3 shows the first step: spreading the workload into multiple time slots. We define each time slot to span 5 minutes. If the activity duration is 13 minutes, first two time slots both contain the full 5 minutes of work, and the third time slot contains 3 minutes of work.

	Total work	1	2	3	4	5
New						
ORT	*12*	5	5	2	0	0
RAD	*3*	3	0	0	0	0
GIPS	*19*	5	5	5	4	0
PREO	*20*	5	5	5	5	0
Repeat						
ORT	*8*	5	3	0	0	0
RAD	*5*	5	0	0	0	0
GIPS	*19*	5	5	5	4	0
PREO	*20*	5	5	5	5	0

Table 7.3 – Workload spreading (expected workload fictious example)

The second step is correcting the workload with the transition probabilities. If a consultation has a 15% chance of occurring, the workload from step one is multiplied by 0.15 to get the expected workload per time slot. See Table 7.4 for the example.

	Transition %	1	2	3	4	5
New						
ORT	*64.47%*	3.22	3.22	1.29	0.00	0.00
RAD	*18.55%*	0.56	0.00	0.00	0.00	0.00
GIPS	*26.26%*	1.31	1.31	1.31	1.05	0.00
PREO	*90.29%*	4.51	4.51	4.51	4.51	0.00
Repeat						
ORT	*78.82%*	3.94	2.36	0.00	0.00	0.00
RAD	*15.95%*	0.80	0.00	0.00	0.00	0.00
GIPS	*85.32%*	4.27	4.27	4.27	3.41	0.00
PREO	*94.80%*	4.74	4.74	4.74	4.74	0.00

Table 7.4 – Correcting for transition probability (expected workload fictious example)

The third step is to correct the timing of downstream demand. If the travel time is 10 minutes, the demand should not start in time slot 1, but in time slot 3. See Table 7.5 for an example of this correction.

	Slots offset	1	2	3	4	5	6	7	8
New									
ORT	3	0.00	0.00	0.00	3.22	3.22	1.29	0.00	0.00
RAD	2	0.00	0.00	0.56	0.00	0.00	0.00	0.00	0.00
GIPS	1	0.00	1.31	1.31	1.31	1.05	0.00	0.00	0.00
PREO	3	0.00	0.00	0.00	4.51	4.51	4.51	4.51	0.00
Repeat									
ORT	1	0.00	3.94	2.36	0.00	0.00	0.00	0.00	0.00
RAD	1	0.00	0.80	0.00	0.00	0.00	0.00	0.00	0.00
GIPS	1	0.00	4.27	4.27	4.27	3.41	0.00	0.00	0.00
PREO	2	0.00	0.00	4.74	4.74	4.74	4.74	0.00	0.00

Table 7.5 – Incorporating travel time (expected workload fictious example)

After these three steps, a full expected workload pattern is calculated for each source consultation and unit set, to each downstream department. The demand patterns are different between consultation types. By changing the sequence of consultation types in the blueprint session schedule the variability of patient outflow could possibly be reduced.

C. Appendix waiting time data quality check

Analysis of Orthopaedic OD direct waiting time, to check if incorrect data is distorting the results. For each day (n=525) the minimum, average, and maximum direct waiting time is calculated. The minimum direct waiting time each day was 0 minutes (except for a few days with <10 observations, which we consider as outliers).

The results do not seem to reveal heavily contaminated or wrong data. We conclude these figures for waiting time are correct and paint an accurate image of direct waiting time.

D. Appendix literature review papers

Google Scholar is used with the following queries, and results from the first 3 pages (30 results) the papers with a promising title and abstract are selected. Papers older than 10 years are excluded.

Search query	Number of results	Review papers found
elective patient planning multiple departments	17900	a, f
healthcare appointment planning overview	58500	c, e, f
integrated appointment scheduling literature review	57200	e
multi appointment planning blueprint	40200	d, f
patient appointment planning literature review	17200	b, f

a. A Survey of Health Care Models that Encompass Multiple Departments (Vanberkel, Boucherie, Hans, Hurink, & Litvak, 2010)
b. Appointment scheduling in health care: Challenges and opportunities (Gupta & Denton, 2008)
c. Literature review on multi-appointment scheduling problems in hospitals (Marynissen & Demeulemeester, 2019)
d. Multi-disciplinary planning in health care: a review (Leeftink, Bikker, et al., 2018)
e. Outpatient appointment systems in healthcare: A review of optimization studies (Ahmadi-Javid et al., 2017)
f. Taxonomic classification of planning decisions in health care: a structured review of the state of the art in OR/MS (Hulshof et al., 2012)

Papers (a) and (b) are excluded since 3 recent review papers are available (c, d, e). Paper (f) is used to classify the research question in a standardised way.

Hulshof et al. (2012) classify the design of an appointment schedule blueprint as a tactical problem, and in the ambulatory care services. On the offline operational level, scheduling appointments can be distinguished between three types. S*ingle appointments* apply if only one appointment at a time is considered for patients. *Combination appointments* apply when multiple appointments for a patient are planned on the same day. *Appointment series* apply for example when a patient needs an appointment every month.

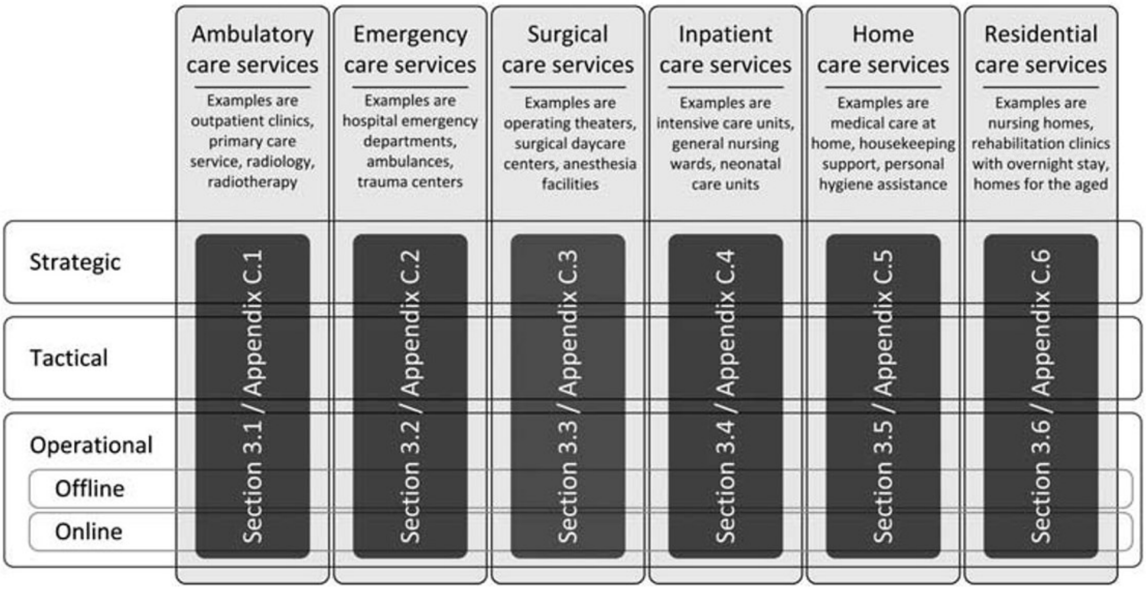

Figure 7.4 - The taxonomy for resource capacity planning and control decisions in health care (Hulshof et al., 2012)

The following decisions are named: *number of patients per consultation session, patient overbooking, length of the appointment interval, number of patients per appointment slot, sequence of appointments, queue discipline in the waiting room,* and *anticipation for unscheduled patients.*

Marynissen & Demeulemeester (2019) recognise that multi-appointment scheduling problems in hospitals cannot be seen separate from hybrid instances where some departments schedule appointments, and other departments do not schedule but take a queueing approach (patient flow). The scope of the paper is limited to cases where all related departments schedule appointments, which is different from the context of this research.

The main challenges they identify for outpatient departments are *uncertain service times* and *patient no-shows*. Approaches to the objective function are *single objective*, *multi-objective (with weights), and multi-objective (with different stages)*.

Leeftink, Bikker, et al. (2018) discuss designing a *blueprint schedule*. Objectives include *combining consultations, minimize waiting time, minimize access time,* and *minimize throughput time*. Robustness to different patient arrival realisations is important for the blueprint. Different design choices for the blueprint schedule are possible. The negative effects of patient tardiness can be reduced by scheduling multiple patients in a time slot.

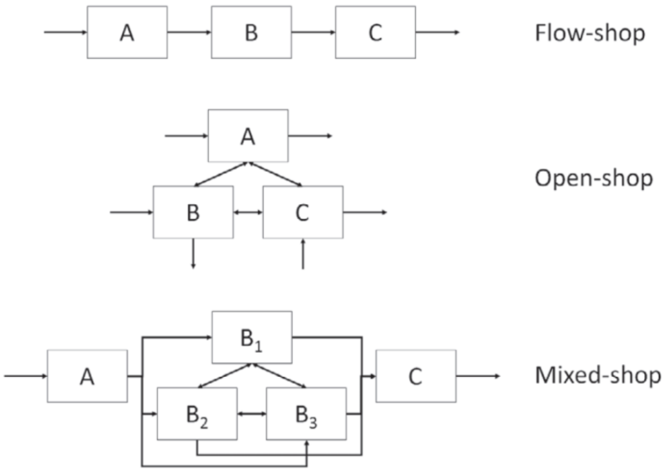

Figure 7.5 – Visualisation of flow-shop, open-shop, and mixed-shop systems (Leeftink, Bikker, et al., 2018)

The multi-appointment system can be characterised as a job-shop system. In a *flow-shop* a fixed sequence is followed. In an *open-shop* system the activities can be performed in any sequence. Finally a *mixed-shop* system is a combination of the two. This characterisation is useful to find papers dealing with the same type of scheduling problem. In this research a *mixed-shop* is appropriate.

Patient performance measures used in *mixed-shop* systems are *direct waiting time*, *length of stay*, and *levelled care load*. System performance measures are the *number of patients admitted, makespan, and completion times*. Performance measures not found in the literature but deemed important are *resource idle time, overtime,* and *utilisation*.

Variability exists in *patient arrivals, appointment durations, resource capacity,* and *care pathway*. Not taking variability into account may influence the robustness of the obtained solution, however adopting the stochastic approach often increased the model complexity and computation time. For all four variability sources mentioned here both deterministic approaches and stochastic approaches are available.

Finally, to increase the generality of the model, a wide range of parameter settings needs to be evaluated. A comparison of the performance of the proposed approach with the performance of relevant approaches can be made to increase scientific relevance.

Ahmadi-Javid et al. (2017) identify three major stakeholder groups: patients, system owners, and staff. The most common performance measures are *patient waiting time, server idle time, system overtime, number of patients seen,* and *number of rejected (or deferred) patients*. Indirect waiting time (time in between appointments for the same patient) is considered in a few recent papers. Rarely used performance measures include *number of patients exceeding waiting time targets, congestion, unfairness,* and *continuity of care*. The fairness measure can be formulated as the difference between the minimum and maximum waiting time. Most papers use a weighted sum to reach a single objective function. Non-linear, pareto and risk-averse objective functions are also possible.

Decisions related to this research problem are the *appointment intervals (slots),* where the optimum depends on service time distribution, interruption, physician lateness, patient lateness and no-shows. The *block size* (number of patients per appointment slot) can reduce the negative effects of patient no-shows. The *number of appointments per consultation session* can be adjusted to minimize patient waiting time and staff overtime. A decision can be made to assign a different priority per patient group, soft priorities can be set by applying different weights to waiting time per patient group. The *patient sequence* can be first-come first-serve, or ordered variance, or increasing no-show probability. The ordered variance is analysed in a paper and determined to be optimal if the service time distribution is positively skewed.

Important environmental factors for the optimal appointment system design are *patient unpunctuality, physician lateness, interruption, patient no-show, patient preference, service time distribution, patient heterogeneity, type of appointment required by patients*.

E. Appendix appointment slot optimization papers

The literature review "Outpatient appointment systems in healthcare: A review of optimization studies" (Ahmadi-Javid et al., 2017) classifies articles.

Title	Authors	Tactical & operational decisions	Single (S) or multiple (M) servers	Walk-ins allowed	Objective	Modeling approach	Solution method
	Anderson, Zheng, Yoon, & Khasawneh, 2015	T2 (appointment slots)	S	No	Min. costs of waiting time, idle time, and overtime	1-SSP (single stage stochastic programming)	S-SBO (Simulation based optimization)
	Berg et al., 2014	T2/O3/O6 (T2/O3: OBA, O6: OBA, RBA) (integrated)	S	No	Max. profit (revenue of patients seen – costs of waiting time, idle time, and overtime)	2-SSP (two stage stochastic programming)	AM (Analytical method) /BB (Branch & bound) /LD (L-shaped decomposition) /H (heuristic)
	Bikker, Kortbeek, van Os, & Boucherie, 2015	TO (physician's scheme design)	M	No	Min. weighted sum of lower bound of indirect waiting time, and difference between daily supply and demand	MILP (Mixed-integer linear programming)	SC (CPLEX)
	Chen & Robinson, 2014	T2/O3/O6 (T2/O3: OBA, O6: RBA) (T2/O3: integrated, (T2/O3)/O6: sequential)	S	No	Min. costs of waiting time, idle time, and overtime	1-SSP	BD (Benders decompositionn) /Heuristic (for O6)
	Denton & Gupta, 2003	T2/O3 (OBA) (integrated)	S	No	Min. costs of waiting time, idle time, and overtime	2-SSP	AM/LD
	Erdogan & Denton, 2013	T2/O3 (OBA) (integrated)	S	No	Min. costs of waiting time and overtime	a. 2-SSP b. M-SSP	a. – b. ND (nested decomposition)
	Erdogan et al., 2015	T2/O3/O6 (T2/O3: OBA, O6: OBA, RBA) (integrated)	S	Yes	Min. costs of direct and indirect waiting times, idle times, and overtime	2-SSP	AM/LD
	Hassin & Mendel, 2008	T2/O3 (OBA) (integrated)	S	No	Min. costs of waiting time and server availability	1-SSP	O
	Huang, Hancock, & Herrin, 2012	T2	S	No	Min. waiting time and idle time	1-SSP	S-SBO
	Klassen & Yoogalingam, 2009	T2/O3 (OBA) (integrated)	S	No	Min. costs of waiting time, idle time, and overtime	1-SSP	S-SBO
	Klassen & Yoogalingam,	T2/O3 (OBA)	S	No	Min. costs of waiting	1-SSP	S-SBO

2013	(integrated)			time and idle time		
Klassen & Yoogalingam, 2014	T2/O3 (OBA) (integrated)	S	No	Min. costs of waiting time and idle time	1-SSP	S-SBO
Kong et al., 2013	T2/O3 (OBA) (integrated)	S	No	Min. costs of waiting time and overtime	2-SSP/SP-O (DRO) Robust optimization /C-SDP Semi-difinitive programming	SO (numerical, accurate method, general purpose optimization software, other items)
Kuiper et al., 2015	T2/O3 (OBA) (integrated)	S	No	Min. costs of waiting time and idle time	1-SSP/O (QT)	SO
Kuiper & Mandjes, 2015	T2/O3 (OBA) (integrated)	M	No	Min. costs of waiting time and idle time	1-SSP/O (QT)	AM/SO
Luo et al., 2012	T2/T5/O3 (OBA) (integrated)	S	Yes (urgent)	Max. profit (revenue of patients seen – costs of waiting time and overtime)	1-SSP	AM/SO
Mak et al., 2014a	T2/O3/O6 (T2/O3: OBA, O6:OBA, RBA) (T2/O3: integrated, (T2/O3)/O6: sequential)	S	No	Min. costs of waiting time and overtime	SP-O (DRO)/C-SDP/SOCP/LP	AM
Robinson & Chen, 2003	T2/O3 (OBA, RBA) (integrated)	S	No	Min. costs of waiting time and idle time	2-SSP	AM/O
Tang et al., 2014	T2/O3 (RBA) (integrated)	S	No	Min. costs of waiting time, idle time, and overtime	1-SSP/O (QT)	AM/SO
Vink et al., 2015	T2/O3 (OBA) (integrated)	S	No	Min. costs of waiting time, idle time, and overtime	1-SSP	H (heuristic)

Table 7.6 – Overview of appointment slot optimization articles and solution methods in (Ahmadi-Javid et al., 2017)

F. Appendix sliding window

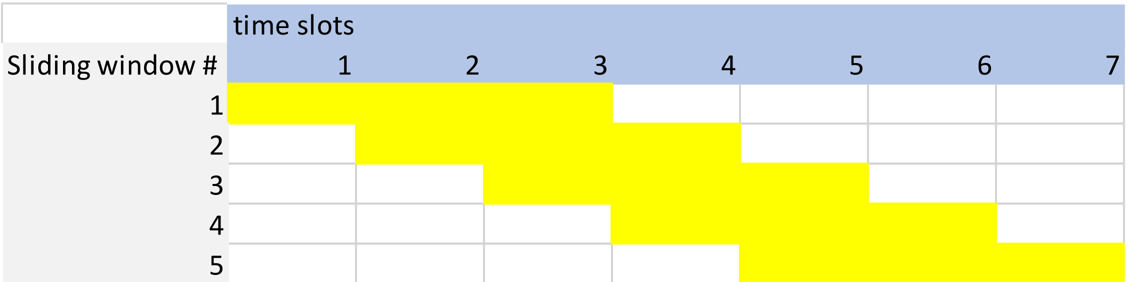

Figure 7.6 – Sliding window width 3 results in 5 sliding windows (with 7 time slots)

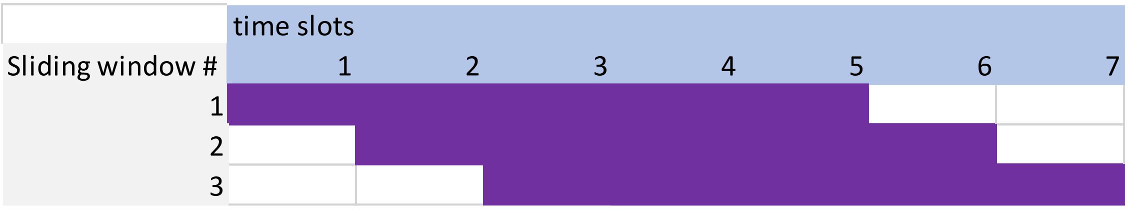

Figure 7.7 – Sliding window width 5 results in 3 sliding windows (with 7 time slots)

Figure 7.6 and Figure 7.7 give a quick visualisation of the sliding windows, and verify the formula for determining the number of sliding windows:

$$number\ of\ sliding\ windows = \ numberOfSlots - slidingWindowWidth + 1$$

E.g.

$$number\ of\ sliding\ windows = 7 - 5 + 1 = 3$$

G. Appendix base version of MILP model

The model described in Chapter 4 is the final version. The most basic version of the model is described in this appendix section, and all model extensions described in Section 4.6 are applied on top of this base model.

Set	Element	Description
C	c	Consultation types
D	d	Departments

Time sets:

Set	Element	Description
T = {1, …, numberOfSlots}	t	Time slot
SW = {1, …, (numberOfSlots - *slidingWindowWidth* + 1)}	sw	Sliding time window. Used to group workload into, similar to a moving average.

Parameter	Range	Description
$consultationsToBeScheduled_c$	$\in \mathbb{N}$	The number of consultations to be scheduled of type *c*
$norm_{d,t}$	$[0, \infty)$	The capacity norm for total minutes of workload arriving at department *d*, per time slot t
$expectedWorkloadBefore_{c,d}$	$[0, \infty)$	Expected workload (in minutes) at department *d* **before** the start of consultation type *c*
$expectedWorkloadAfter_{c,d}$	$[0, \infty)$	Expected workload (in minutes) at department *d* **after** the end of consultation type *c*
firstSlot	$\{t_{min}, \ldots, t_{max}\}$	First timeslot a consultation can be booked into
lastSlot	$\{t_{min}, \ldots, t_{max}\}$	Last timeslot a consultation can be booked into
numberOfSlots	$\in \mathbb{N}$	The total number of timeslots
$departmentWeight_d$	$[0, 1]$	Weight of each department for objective function
slidingWindowWidth	{1..numberOfSlots}	Specifies the number of time slots one sliding window covers. The first window starts at time t=1, to t=1+slidingWindowWidth. The second window starts at time t=2, to t=2+slidingWindowWidth

Variable	Range	Description
$Workload_{d,t}$	[0, ∞)	Workload for department *d*, at timeslot *t*
$WorkloadGrouped_{d,sw}$	[0, ∞)	Workload grouped, $Workload_{d,t}$ grouped into sliding time windows SW
$DeviationFromNorm_{d,t}$	$\in \mathbb{R}$	Absolute deviation from norm for department *d* in time window *sw*
$DeviationFromNormGrouped_{d,sw}$	$\in \mathbb{R}$	The absolute deviation from the norm for each time slot *t* is summed for sliding time window *sw*
$MaxDev_d$	$\in \mathbb{R}$	The maximum deviation from the norm for department *d* across the whole schedule

Decision variable	Range	Description
$X_{c,t}$	{0, 1} Binary	Schedule consultation type *c* at time slot *t* (1=yes, 0=no)

1. **Maximum one consultation per time slot**

$$\sum_{c} X_{c,t} \leq 1 \qquad \forall t$$

2. **Schedule the correct number of consultations for each type**

$$\sum_{t} X_{c,t} = consultationsToBeScheduled_c \qquad \forall c$$

3. **No consultations before first slot**

$$\sum_{t=1}^{t=firstSlot-1} X_{c,t} = 0 \qquad \forall c$$

4. **No consultations after last slot**

$$\sum_{t=lastSlot+1}^{T} X_{c,t} = 0 \qquad \forall c$$

5. **Maximum of two adjacent 'New' consultations**

$$\sum_{tt=t}^{tt=t+2} X_{New,tt} \leq 2 \qquad \forall t$$

6. Set downstream workload

$$Workload_{d,t} = \sum_c (X_{c,t+1} * expectedWorkloadBefore_{c,d}) + \sum_c (X_{c,t-1} * expectedWorkloadAfter_{c,d}) \quad \forall d, t$$

7. Group workload into sliding time windows

All workload of time slots *t* needs to be aggregated to the workload per sliding window *sw*.

$$WorkloadGrouped_{d,sw} = \sum_{t=sw}^{sw+slidingWindowWidth-1} Workload_{d,t} \qquad \forall\, sw, d$$

8. Calculate deviation from norm per time slot

For each time group and department, calculate how much the workload deviates from the norm

$$DeviationFromNorm_{d,t} = Workload_{d,t} - norm_{d,t} \qquad \forall d, t$$

9. Calculate positive and negative deviation from norm

As a step towards the absolute deviation from the norm without losing the linear property.

$$PositiveDeviationFromNorm_{d,t} - NegativeDeviationFromNorm_{d,t} = Workload_{d,t} - norm_{d,t} \quad \forall d, t$$

10. Ensure positive deviation and negative deviation are >0

$$PositiveDeviationFromNorm_{d,t}, NegativeDeviationFromNorm_{d,t} \geq 0 \qquad \forall d, t$$

11. Ensure only positive OR negative deviation takes a value >0 (constraint one)

The variable $DeviationIsPositive_{d,t}$ is a binary variable. It is 0 if the deviation from the norm is negative, 1 if the deviation from the norm is negative. Using this variable in combination with bigM ensures only one of the deviation variables can take a value other that 0.

$$PositiveDeviationFromNorm_{d,t} \leq DeviationIsPositive_{d,t} * bigM \qquad \forall d, t$$

12. Ensure only positive OR negative deviation takes a value >0 (constraint two)

$$NegativeDeviationFromNorm_{d,t} \leq (1 - DeviationIsPositive_{d,t}) * bigM \quad \forall d, t$$

13. Calculate the absolute deviation from the norm

Now that the deviation has been split into its positive and negative parts, we can take the sum to get the absolute difference between workload and the norm (without losing linearity).

$$AbsoluteDeviationFromNorm_{d,t} = PositiveDeviationFromNorm_{d,t} + NegativeDeviationFromNorm_{d,t} \quad \forall d, t$$

14. Calculate DeviationFromNormGrouped per sliding time window

For each sliding time window, sum the deviation from the norm for individual time slots, for all time slots that belong to sliding window *sw*

$$DeviationFromNormGrouped_{d,sw} = \sum_{t=sw}^{sw+slidingWindowWidth-1} DeviationFromNorm_{d,t} \quad \forall d, sw$$

15. Calculate maximum deviation from norm

For each department, determine the maximum deviation from the norm over all the sliding windows

$$MaxDev_d \geq DeviationFromNormGrouped_{d,sw} \qquad \forall d, sw$$

Objective function

$$min \sum_d (MaxDev_d * departmentWeight_d)$$

H. Appendix sets and parameters for Thursday afternoon case

For Thursday afternoon the set of schedules is set to the names of the 8 doctors. In this report we will simply name them by Doctor 1 through Doctor 8.

Set	Values
Schedules	{Doctor 1, Doctor 2, Doctor 3, Doctor 4, Doctor 5, Doctor 6, Doctor 7, Doctor 8}

		consultationsToBeScheduled				
		Unit 1	Unit 2	Unit 3	Unit 4	Unit 5
Doctor 1	New					4
	Repeat					5
	Discharge					2
	POP					3
	Supervision					
	Empty					
Doctor 2	New				5	
	Repeat				3	
	Discharge				2	
	POP				3	
	Supervision					
	Empty					
Doctor 3	New			4		
	Repeat			6		
	Discharge					
	POP			1		
	Supervision					
	Empty					
Doctor 4	New					4
	Repeat					5
	Discharge					2
	POP					3
	Supervision					
	Empty					
Doctor 5	New				5	
	Repeat				3	
	Discharge				2	
	POP				3	
	Supervision					
	Empty				4	
Doctor 6	New		5			
	Repeat		7			
	Discharge					
	POP					
	Supervision					
	Empty					
Doctor 7	New			7		
	Repeat			4		
	Discharge			1		
	POP					
	Supervision					
	Empty					
Doctor 8	New	6				
	Repeat	6				
	Discharge	2				
	POP	3				
	Supervision					
	Empty					

		consultationsToBeScheduled				
		Unit 1	Unit 2	Unit 3	Unit 4	Unit 5
Doctor 1	New					3
	Repeat					3
	Discharge					3
	POP					1
	Supervision					1
	Empty					1
Doctor 2	New				3	
	Repeat				3	
	Discharge				3	
	POP				1	
	Supervision				1	
	Empty				1	
Doctor 3	New			3		
	Repeat			3		
	Discharge			3		
	POP			3		
	Supervision			1		
	Empty			1		
Doctor 4	New					3
	Repeat					3
	Discharge					3
	POP					1
	Supervision					1
	Empty					1
Doctor 5	New				3	
	Repeat				3	
	Discharge				3	
	POP				1	
	Supervision				1	
	Empty				1	
Doctor 6	New		3			
	Repeat		3			
	Discharge		3			
	POP		1			
	Supervision		1			
	Empty		1			
Doctor 7	New			3		
	Repeat			3		
	Discharge			3		
	POP			1		
	Supervision			1		
	Empty			1		
Doctor 8	New	3				
	Repeat	3				
	Discharge	3				
	POP	1				
	Supervision	1				
	Empty	6				

I. Appendix experiment results

i. Performance difference in minutes

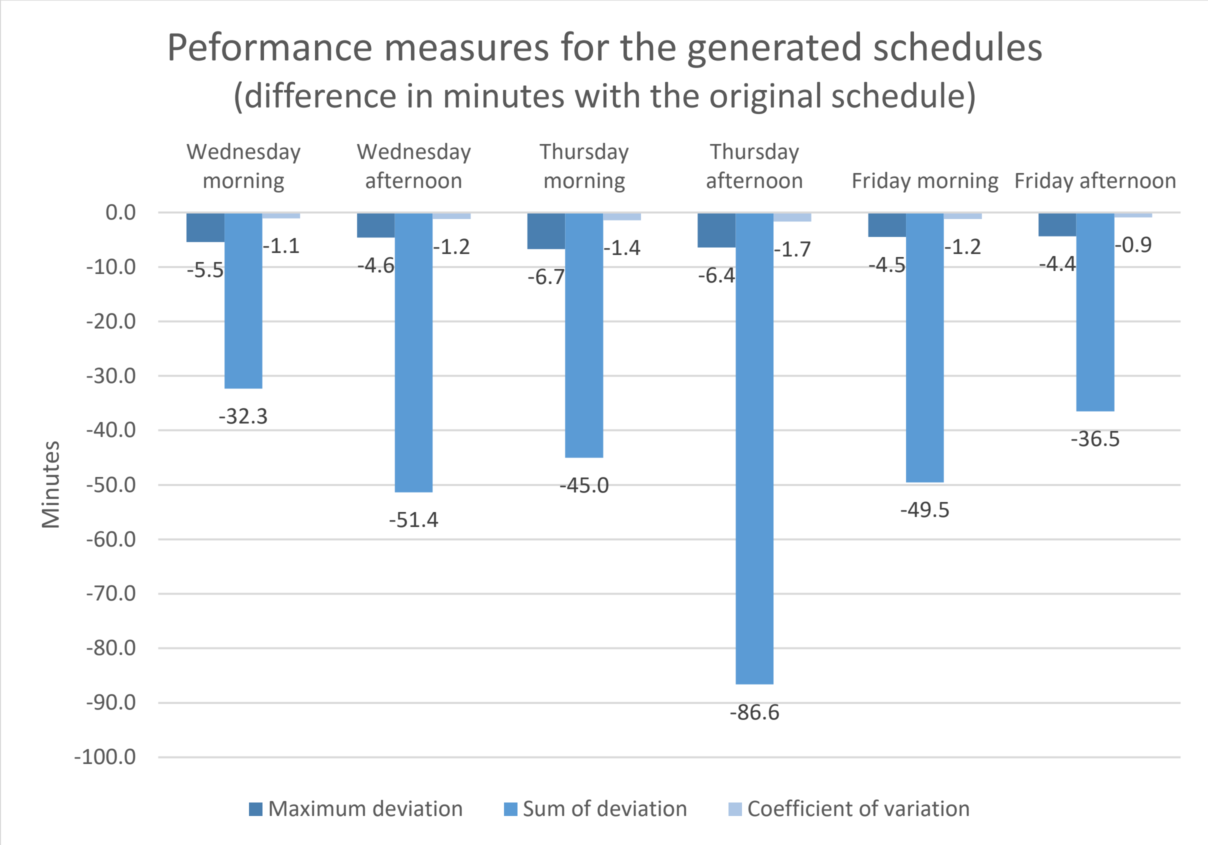

This graph shows the difference with the original schedule in minutes, with the weighted total.

ii. Complete overview of measures for the generated schedules

		Maximum peak	Sum absolute differences from norm	Variability (from n
		% of original	% of original	% of original
1. wednesdayMorning	OOD	101%	99%	120%
	RAD	23%	59%	41%
	Plaster	46%	80%	59%
	PREO	56%	74%	66%
	weighted total	39%	73%	58%
2. wednesdayAfternoon	OOD	126%	130%	103%
	RAD	25%	44%	33%
	Plaster	41%	78%	43%
	PREO	42%	33%	62%
	weighted total	41%	57%	51%
3. thursdayMorning	OOD	70%	107%	86%
	RAD	22%	58%	38%
	Plaster	49%	93%	48%
	PREO	30%	23%	46%
	weighted total	31%	59%	47%
4. thursdayAfternoon	OOD	109%	111%	134%
	RAD	21%	29%	27%
	Plaster	51%	78%	58%
	PREO	55%	47%	38%
	weighted total	41%	51%	45%
5. fridayMorning	OOD	90%	94%	94%
	RAD	21%	47%	26%
	Plaster	63%	98%	76%
	PREO	82%	39%	89%
	weighted total	39%	57%	52%
6. fridayAfternoon	OOD	102%	130%	92%
	RAD	21%	37%	34%
	Plaster	41%	77%	57%
	PREO	91%	92%	98%
	weighted total	39%	60%	55%
Average	OOD	100%	112%	105%
	RAD	22%	46%	33%
	Plaster	48%	84%	57%
	PREO	59%	51%	67%
	weighted total	38%	59%	51%

This table shows the complete overview of the measures for each department and the weighted total, in percentage of the original, for each of the six experiments.

J. Appendix Monte Carlo graphs

i. Thursday afternoon – OOD

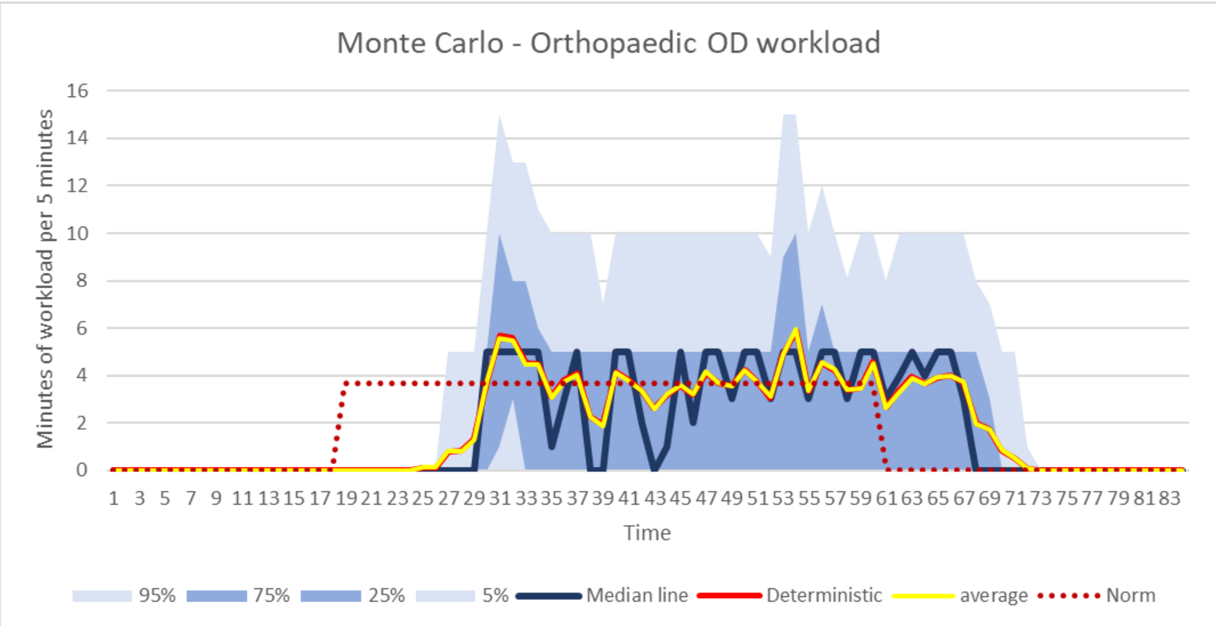

ii. Thursday afternoon – RAD

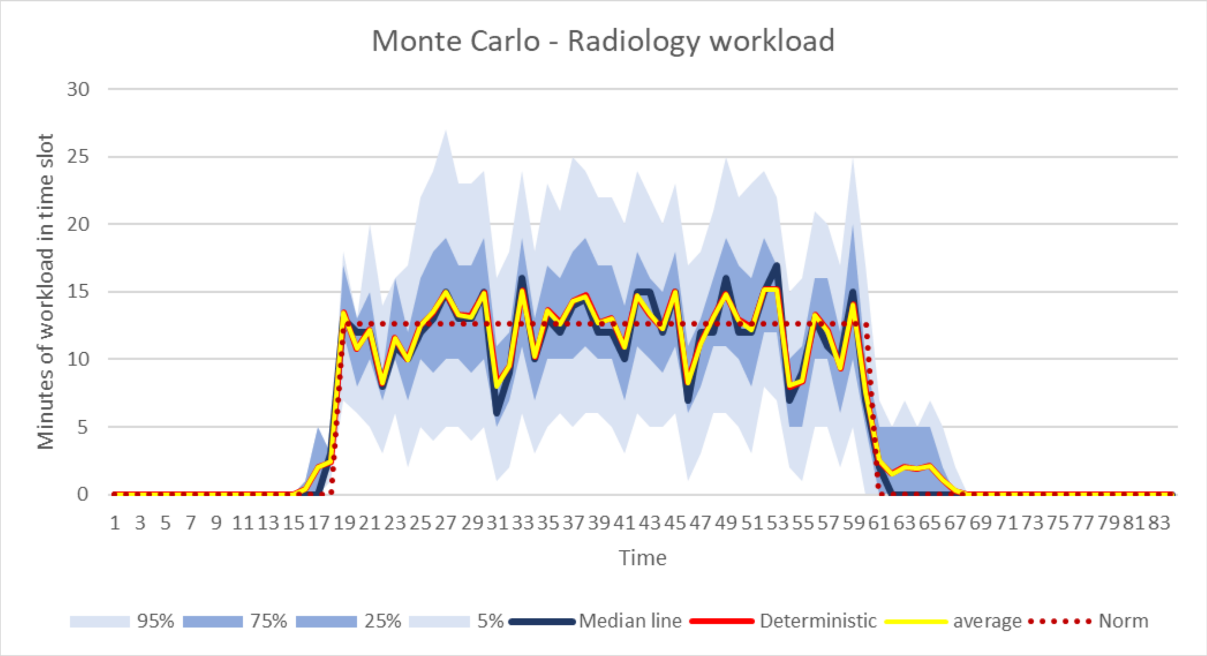

iii. Thursday afternoon – Plaster room

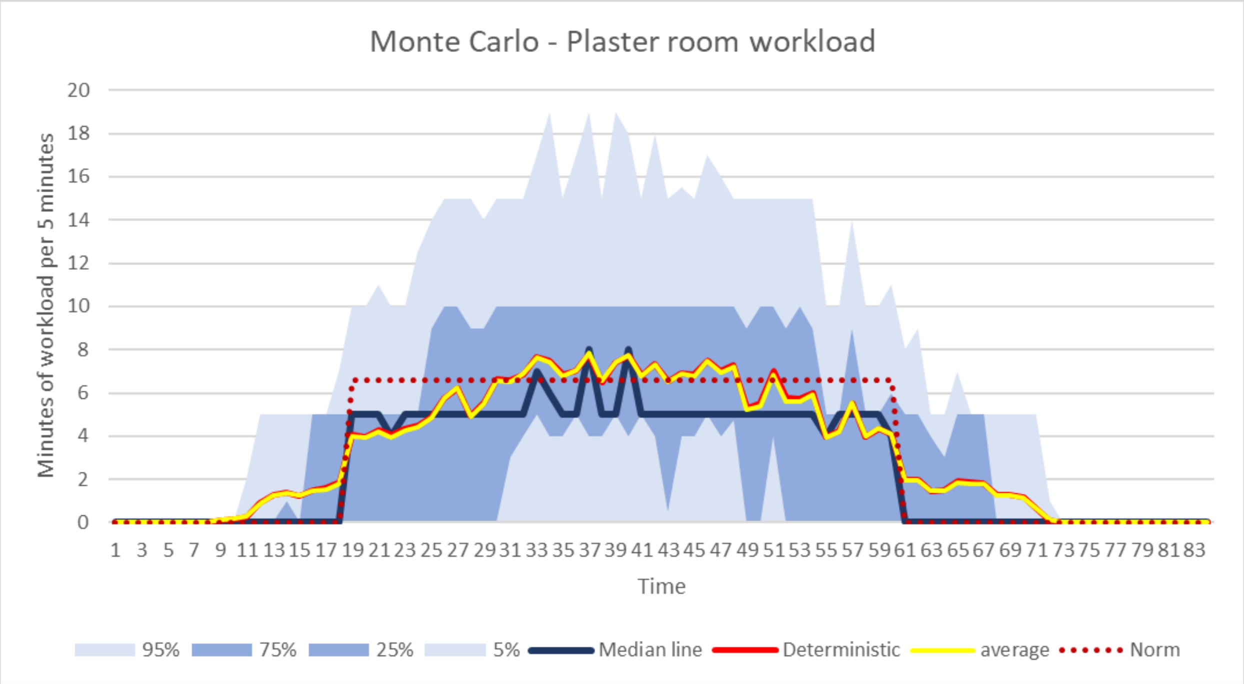

iv. Thursday afternoon – Screening

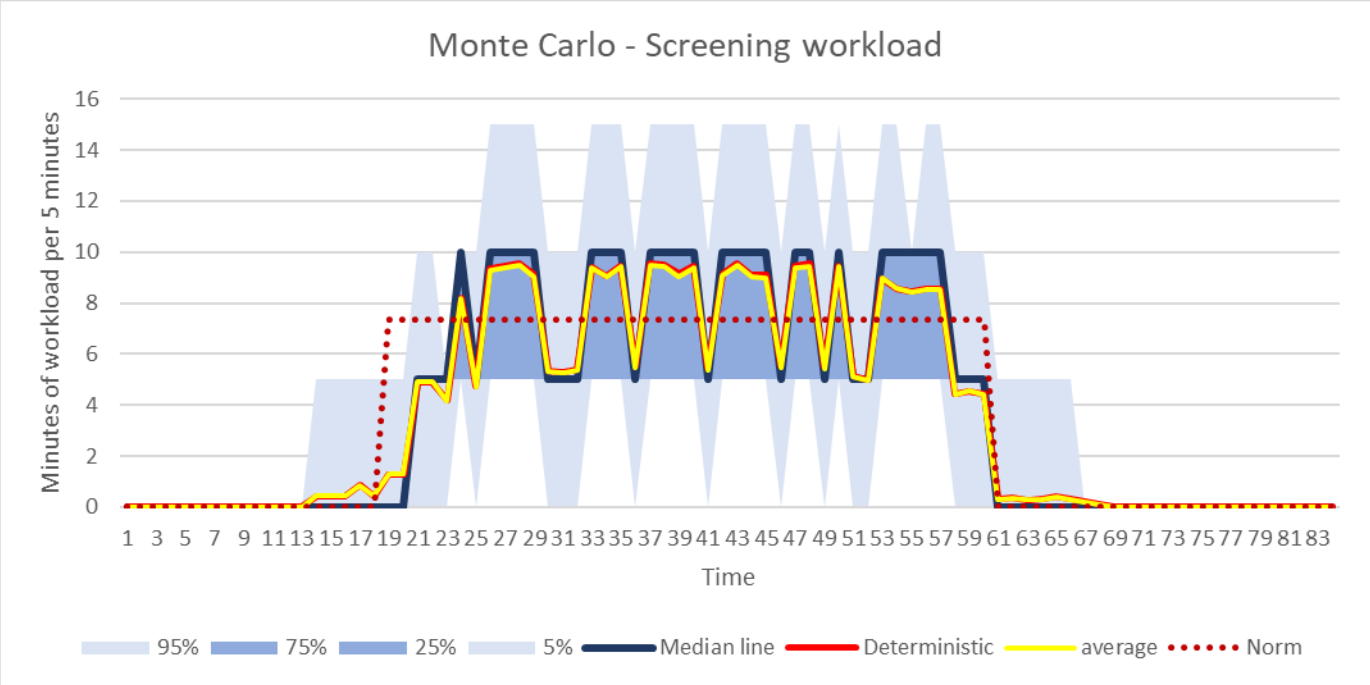

K. Appendix Monte Carlo script in R

```
setwd("G:/13. Monte Carlo simulatie")

# Libraries
library(tidyverse)
library(readxl)

# Clear workspace
rm(list=ls()); gc();

# --- Custom functions ---
transposeTibble <- function (tibbleInput, customColumnNames = FALSE, stripFirstColumn = FALSE) {
 firstColumn <- tibbleInput[[1]]

 if (class(customColumnNames) == "logical") {
  if (customColumnNames == FALSE) {
   columnNames <- firstColumn
  } else {
   stop("No column names provided")
  }
 } else {
  if (length(customColumnNames) != length(firstColumn)) {
   stop("Incorrect number of columnNames entered");
  }
  columnNames <- customColumnNames
 }

 dataFrame <- as.data.frame(tibbleInput)
```

```
  if (stripFirstColumn == TRUE) {
    dataFrameTransposed <- t(dataFrame[,-1])
  } else {
    dataFrameTransposed <- t(dataFrame)
  }

  colnames(dataFrameTransposed) <- columnNames
  outputTibble <- as_tibble(dataFrameTransposed)

  return(outputTibble)
}


# --- Prepare data input

# ---> FILE NAMES <---
consultationDurationsFileName <- "consultationDurations.xlsx"
generatedSchedulesFileName <- "generatedSchedules.xlsx"

# Specify case name (make sure to match the folder name on disk)
caseName <- "1. woensdagMorgen"
numberOfSchedules <- 6
numberOfTimeSlots <- 84

# Generic settings
departmentNames = c("OOD", "RAD", "Plaster", "PREO")
numberOfDepartments = length(departmentNames)
units = c("UpperExtr", "Hip", "Knee", "Spine", "Foot")
consultationTypes = c("New", "Repeat", "Discharge", "POP", "Supervision", "Empty") # Warning:
make sure order matches excel files
numberOfConsultationTypes <- length(consultationTypes)
maximumOffset <- 16

# Build a key value pair list, that can return the index of a given consultation type
# Purpose: To access the correct row in tibbles, since rows are not named
i = 1
consultationTypeIndex <- list()
for (consultationType in consultationTypes) {
  consultationTypeIndex[[consultationType]] <- i
  i = i+1
}
rm(consultationType, i)

# --- Create data templates:
# 1. expectedWorkloadBefore
# 2. expectedWorkloadAfter
# 3. transitionProbabilitiesBefore
# 4. transitionProbabilitiesAfter
# 5. consultationDurations
# 6. schedules
# 7. workload
```

```
# - 1&2: Create expectedWorkloadBefore & expectedWorkloadAfter objects
# Create tibbleTemplate
vectorTemplate <- rep(0, maximumOffset)
tibbleTemplateWorkload <- tibble("OOD"=vectorTemplate, "RAD"= vectorTemplate,
"Plaster"=vectorTemplate, "PREO"=vectorTemplate)
tibbleTemplateWorkload

rm(vectorTemplate)

# Construct object for expectedWorkloadBefore and expectedWorkloadAfter
# level 1: consultationType
#   level 2: unit
#     level 3: tibble with departments in columns, each row is one time offset i
expectedWorkloadBefore <- list()
expectedWorkloadAfter <- list()
for (consultationType in consultationTypes) {
  expectedWorkloadBefore[[consultationType]] <- list()
  expectedWorkloadAfter[[consultationType]] <- list()

  for (unit in units) {
    expectedWorkloadBefore[[consultationType]][[unit]] <- tibbleTemplateWorkload
    expectedWorkloadAfter[[consultationType]][[unit]] <- tibbleTemplateWorkload
  }
}
rm(consultationType, unit)

# - 3&4: Create transitionProbabilityBefore & transitionProbabilityAfter objects
vectorTemplate <- rep(0, length(consultationTypes))
tibbleTemplateTransitionProbabilities <- tibble("OOD"=vectorTemplate, "RAD"=vectorTemplate,
"Plaster"=vectorTemplate, "PREO"=vectorTemplate)

transitionProbabilitiesBefore <- list()
transitionProbabilitiesAfter <- list()

for (unit in units) {
  transitionProbabilitiesBefore[[unit]] <- tibbleTemplateTransitionProbabilities
  transitionProbabilitiesAfter[[unit]] <- tibbleTemplateTransitionProbabilities
}
rm(vectorTemplate)


# - 5: Create consultationDuration object
# No preparation needed (otherwise number of schedules would need to be hardcoded)
# structure will be a single tibble with rows=schedules, columns=consultationTypes

# - 6: Create schedules object
# level 1: scheduleName
#   level 2: unit
#     level 3: tibble with rows=time index, columns=consultationTypes
schedules <- list()
scheduleNames <- list()
```

```
vectorTemplate <- rep(0, numberOfTimeSlots)
tibbleTemplateSchedule <- tibble("New"=vectorTemplate, "Repeat"=vectorTemplate,
"Discharge"=vectorTemplate, "POP"=vectorTemplate, "Supervision"=vectorTemplate,
"Empty"=vectorTemplate)
tibbleTemplateSchedule

for (i in 1:numberOfSchedules) {
  scheduleName <- paste("Schedule_", i, sep = "")
  scheduleNames[i] <- scheduleName
  schedules[[scheduleName]] <- list()

  for (unit in units) {
    schedules[[scheduleName]][[unit]] <- tibbleTemplateSchedule
  }
}
rm(vectorTemplate, scheduleName, i, unit)



# --- Load data

# - 1. Load expectedWorkloadBefore
# Prepare for data import of excel file 1
excelFileName = "expectedWorkloadBefore all.xlsx"

# Define columns for each unit
startColumns <- list()
endColumns <- list()

# Unit UpperExtr
startColumns[[ units[[1]] ]] <- "C"
endColumns[[ units[[1]] ]] <- "R"
# Unit Hip
startColumns[[ units[[2]] ]] <- "S"
endColumns[[ units[[2]] ]] <- "AH"
# Unit Knee
startColumns[[ units[[3]] ]] <- "AI"
endColumns[[ units[[3]] ]] <- "AX"
# Unit Spine
startColumns[[ units[[4]] ]] <- "AY"
endColumns[[ units[[4]] ]] <- "BN"
# Unit Foot
startColumns[[ units[[5]] ]] <- "BO"
endColumns[[ units[[5]] ]] <- "CD"

startColumns
endColumns

# Define the first row with data
firstRow = 4
```

```
# Import first excel file with parameters
currentRow = firstRow
i = 1
for (consultationType in consultationTypes) {
  j = 1
  for (unit in units) {

    # Calculate current starting row
    currentRow = firstRow + (i-1)*numberOfDepartments
    rng = paste(startColumns[[unit]], currentRow, ":", endColumns[[unit]],
currentRow+numberOfDepartments-1, sep="")

    # Read correct range from excel file
    originalData <- read_excel(excelFileName, col_names = FALSE, range = rng)

    # transpose the data so each row=offset, each column=department
    transposedTibble <- transposeTibble(originalData, customColumnNames = departmentNames,
FALSE)

    # Enter data into prepared object 'expectedWorkloadBefore'
    expectedWorkloadBefore[[consultationType]][[unit]] <- transposedTibble

    # increment unit counter j
    j = j+1
  }
  # increment consultationType counter i
  i=i+1
}

# Clear unnecessary objects
rm(excelFileName, startColumns, endColumns, firstRow, currentRow, i, j, consultationType, unit, rng,
originalData, transposedTibble)


# - 2. Load expectedWorkloadAfter
# Prepare for data import of excel file 2
excelFileName = "expectedWorkloadAfter all.xlsx"

# Define columns for each unit
startColumns <- list()
endColumns <- list()

# Unit UpperExtr
startColumns[[ units[[1]] ]] <- "C"
endColumns[[ units[[1]] ]] <- "R"
# Unit Hip
startColumns[[ units[[2]] ]] <- "S"
endColumns[[ units[[2]] ]] <- "AH"
# Unit Knee
startColumns[[ units[[3]] ]] <- "AI"
endColumns[[ units[[3]] ]] <- "AX"
# Unit Spine
```

```
startColumns[[ units[[4]] ]] <- "AY"
endColumns[[ units[[4]] ]] <- "BN"
# Unit Foot
startColumns[[ units[[5]] ]] <- "BO"
endColumns[[ units[[5]] ]] <- "CD"

startColumns
endColumns

# Define the first row with data
firstRow = 4

# Import first excel file with parameters
currentRow = firstRow
i = 1
for (consultationType in consultationTypes) {
  j = 1
  for (unit in units) {

    # Calculate current starting row
    currentRow = firstRow + (i-1)*numberOfDepartments
    rng = paste(startColumns[[unit]], currentRow, ":", endColumns[[unit]],
currentRow+numberOfDepartments-1, sep="")

    # Read correct range from excel file
    originalData <- read_excel(excelFileName, col_names = FALSE, range = rng)

    # transpose the data so each row=offset, each column=department
    transposedTibble <- transposeTibble(originalData, customColumnNames = departmentNames,
FALSE)

    # Enter data into prepared object 'expectedWorkloadBefore'
    expectedWorkloadAfter[[consultationType]][[unit]] <- transposedTibble

    # increment unit counter j
    j = j+1
  }
  # increment consultationType counter i
  i=i+1
}

# Clear unnecessary objects
rm(excelFileName, startColumns, endColumns, firstRow, currentRow, i, j, consultationType, unit, rng,
originalData, transposedTibble)




# - 3. Load transitionProbabilitiesBefore
excelFileName <- "transitionProbabilitiesBefore all.xlsx"
for (unit in units) {
  transitionProbabilitiesBefore[[unit]] <- read_excel(excelFileName, sheet = unit)
```

```
}
rm(excelFileName)

# - 4. Load transitionProbabilitiesAfter
excelFileName <- "transitionProbabilitiesAfter all.xlsx"
for (unit in units) {
  transitionProbabilitiesAfter[[unit]] <- read_excel(excelFileName, sheet = unit)
}
rm(excelFileName)

# syntax example:
transitionProbabilitiesAfter[["Knee"]][["RAD"]][[consultationTypeIndex[["Discharge"]]]]

# - 5. Load consultationDurations

excelFilePath <- paste("Cases/", caseName, "/", consultationDurationsFileName, sep = "")
consultationDurations <- read_excel(excelFilePath)
rm(excelFilePath)

# - 6. Load schedules
excelFilePath <- paste("Cases/", caseName, "/", generatedSchedulesFileName, sep = "")

# Define columns for each unit
startColumns <- list()
endColumns <- list()

# Unit UpperExtr
startColumns[[ units[[1]] ]] <- "C"
endColumns[[ units[[1]] ]] <- "CH"
# Unit Hip
startColumns[[ units[[2]] ]] <- "CI"
endColumns[[ units[[2]] ]] <- "FN"
# Unit Knee
startColumns[[ units[[3]] ]] <- "FO"
endColumns[[ units[[3]] ]] <- "IT"
# Unit Spine
startColumns[[ units[[4]] ]] <- "IU"
endColumns[[ units[[4]] ]] <- "LZ"
# Unit Foot
startColumns[[ units[[5]] ]] <- "MA"
endColumns[[ units[[5]] ]] <- "PF"

# Define the first row with data
firstRow = 4

# Import first excel file with parameters
currentRow = firstRow
i = 1
for (scheduleName in scheduleNames) {
  j = 1

  # Determine first row of the range for this schedule
```

```
  currentRow <- firstRow + (i-1)*numberOfConsultationTypes

  #print(paste("CurrentRow: ", currentRow))

  for (unit in units) {

    # Determine the cell range in excel for this (schedule, unit) combination
    rng <- paste(startColumns[[unit]], currentRow, ":", endColumns[[unit]],
currentRow+numberOfConsultationTypes-1, sep = "")

    # Read schedule from excel file
    data <- read_excel(excelFilePath, col_names = FALSE, range = rng)

    # Transpose so rows=time, columns=consultationTypes
    transposedTibble <- transposeTibble(data, customColumnNames = consultationTypes, FALSE)

    # Store schedule in object
    schedules[[scheduleName]][[unit]] <- transposedTibble

    # increment unit counter j
    j = j+1
  }

  # increment schedule counter i
  i = i+1
}
rm(excelFilePath, startColumns, endColumns, firstRow, currentRow, i, j, scheduleName, unit, rng,
data, transposedTibble)




# ---------------------------------------------------------------------------------------------------------- #
#                                  BEGIN SIMULATION SECTION
# ---------------------------------------------------------------------------------------------------------- #
# Number of Monte Carlo iterations
numberMonteCarloIterations <- 10000

# Prepare Workload object
# rows=time, columns=departments, cell value is minutes of workload
vectorTemplate <- rep(0, numberOfTimeSlots)
tibbleTemplateWorkloadResults <- tibble("OOD"=vectorTemplate, "RAD"=vectorTemplate,
"Plaster"=vectorTemplate, "PREO"=vectorTemplate)
workload <- list()
for (i in 1:numberMonteCarloIterations) {
  workload[[i]] <- tibbleTemplateWorkloadResults
}
rm(vectorTemplate)


# --- Perform Monte Carlo simulation
for (MCiteration in 1:numberMonteCarloIterations) {
```

```
# --- Start of monte carlo iteration
# set a different seed for each iteration
set.seed(MCiteration)

if (MCiteration==1 | MCiteration%%50==0) {
  print(paste("Start of Monte Carlo iteration", MCiteration))
}

# start iterating over each schedule
s = 1
for (schedule in schedules) {

  # iterate over units
  for (unit in units) {

    # iterate over consultationTypes
    for (consultationType in consultationTypes) {

      # iterate over time slots
      for (t in 1:numberOfTimeSlots) {

        # Check if value is 1
        value <- schedule[[unit]][[consultationType]][t]
        if (value > 0) {
          #print(paste("In schedule", scheduleNames[s], "and unit", unit, "a consultation of type",
consultationType, "is started at time", t))

          for (department in departmentNames) {

            # --- Determine if a transition BEFORE will occur
            # Get the average probability
            probability <- transitionProbabilitiesBefore[[unit]][[department]][[
consultationTypeIndex[[consultationType]]  ]]
            #print(paste("Transition probability before for unit", unit, "to department", department, "for
consultation type", consultationType, "is:", probability))

            # Draw a single time from the binomial distribution
            transitionBeforeOccurs <- rbinom(n = 1, size = 1, prob = probability)
            #print(transitionBeforeOccurs)

            if (transitionBeforeOccurs) {
              #print("Transition before occurs!")

              # Add the workload for this department
              for (i in 1:maximumOffset) {

                timeIndex = t-i
                #print(paste("t:", t, "i:", i, "timeIndex:", timeIndex))

                # Check if not reaching out of bounds
                if (timeIndex >= 1) {
```

```
                # Add workload to the correct time  expectedWorkloadBefore$Repeat$UpperExtr$OOD
                workloadToAdd <- expectedWorkloadBefore[[consultationType]][[unit]][[department]][i]
                #print(paste("workloadToAdd:", workloadToAdd))

                #print(paste("workload beforehand:", workload[[department]][timeIndex]))
                workload[[MCiteration]][[department]][timeIndex] <-
workload[[MCiteration]][[department]][timeIndex] + workloadToAdd
                #print(paste("workload afterwards:", workload[[department]][timeIndex]))
              }
            }
          }



          # --- Determine if a transition AFTER will occur
          # Get the average probability
          probability <- transitionProbabilitiesAfter[[unit]][[department]][[
consultationTypeIndex[[consultationType]]  ]]
          #print(paste("Transition probability after for unit", unit, "to department", department, "for
consultation type", consultationType, "is:", probability))

          # Draw a single time from the binomial distribution
          transitionAfterOccurs <- rbinom(n = 1, size = 1, prob = probability)
          #print(transitionAfterOccurs)

          if (transitionAfterOccurs) {
            #print("Transition after occurs!")

            # get duration of the consultation
            consultationDuration <- consultationDurations[[consultationType]][s]
            #print(paste("Consultation duration for type", consultationType, "and schedule", s, "is:",
consultationDuration))

            # Add the workload for this department
            for (i in 1:maximumOffset) {

              timeIndex = t+consultationDuration+i-1
              #print(paste("t:", t, "i:", i, "timeIndex:", timeIndex))

              # Check if not reaching out of bounds
              if (timeIndex <= numberOfTimeSlots) {

                # Add workload to the correct time  expectedWorkloadBefore$Repeat$UpperExtr$OOD
                workloadToAdd <- expectedWorkloadAfter[[consultationType]][[unit]][[department]][i]
                #print(paste("workloadToAdd:", workloadToAdd))

                #print(paste("workload beforehand:", workload[[department]][timeIndex]))
                workload[[MCiteration]][[department]][timeIndex] <-
workload[[MCiteration]][[department]][timeIndex] + workloadToAdd
                #print(paste("workload afterwards:", workload[[department]][timeIndex]))
              }
```

```
            }
          }

        }
      }
    }
  }
}

  # increment schedule counter s
  s = s+1
 }
}


# ---------------------------------------------------------------------------------------------------------- #
#                                   END SIMULATION SECTION
# ---------------------------------------------------------------------------------------------------------- #



# --- Gather results, create tibble per department, rows=iterations, columns=timeslots
columnNames <- character(numberOfTimeSlots)
for (i in 1:length(columnNames)) {
  columnNames[i] <- paste("t=", i, sep = "")
  #columnNames[i] <- i
}
columnNames

# Construct OOD table
OOD <- tibble("Workload" = workload[[1]]$OOD)
OOD <- transposeTibble(OOD, customColumnNames = columnNames, FALSE)
for (i in 2:length(workload)) {
#for (i in 2:15) {
  tibbleFromTable <- tibble("Workload" = workload[[i]][["OOD"]])
  transposed <- transposeTibble(tibbleFromTable, customColumnNames = columnNames, FALSE)
  OOD <- bind_rows(OOD, transposed)
}

# Construct RAD table
RAD <- tibble("Workload" = workload[[1]]$RAD)
RAD <- transposeTibble(RAD, customColumnNames = columnNames, FALSE)
for (i in 2:length(workload)) {
  tibbleFromTable <- tibble("Workload" = workload[[i]][["RAD"]])
  transposed <- transposeTibble(tibbleFromTable, customColumnNames = columnNames, FALSE)
  RAD <- bind_rows(RAD, transposed)
}

# Construct Plaster table
Plaster <- tibble("Workload" = workload[[1]]$Plaster)
Plaster <- transposeTibble(Plaster, customColumnNames = columnNames, FALSE)
for (i in 2:length(workload)) {
```

```
  tibbleFromTable <- tibble("Workload" = workload[[i]][["Plaster"]])
  transposed <- transposeTibble(tibbleFromTable, customColumnNames = columnNames, FALSE)
  Plaster <- bind_rows(Plaster, transposed)
}

# Construct PREO table
PREO <- tibble("Workload" = workload[[1]]$PREO)
PREO <- transposeTibble(PREO, customColumnNames = columnNames, FALSE)
for (i in 2:length(workload)) {
  tibbleFromTable <- tibble("Workload" = workload[[i]][["PREO"]])
  transposed <- transposeTibble(tibbleFromTable, customColumnNames = columnNames, FALSE)
  PREO <- bind_rows(PREO, transposed)
}


# --- Write Monte Carlo results to excel files for each department
library(writexl)

departmentName <- "OOD"
outputPath <- paste("G:/13. Monte Carlo simulatie/Cases/", caseName, "/MC_results/",
departmentName,".xlsx", sep = "")
write_xlsx(OOD, outputPath)

departmentName <- "RAD"
outputPath <- paste("G:/13. Monte Carlo simulatie/Cases/", caseName, "/MC_results/",
departmentName,".xlsx", sep = "")
write_xlsx(RAD, outputPath)

departmentName <- "Plaster"
outputPath <- paste("G:/13. Monte Carlo simulatie/Cases/", caseName, "/MC_results/",
departmentName,".xlsx", sep = "")
write_xlsx(Plaster, outputPath)

departmentName <- "PREO"
outputPath <- paste("G:/13. Monte Carlo simulatie/Cases/", caseName, "/MC_results/",
departmentName,".xlsx", sep = "")
write_xlsx(PREO, outputPath)
```